"AN ADVANCED THYRISTOR VALVE FOR HVDC TRANSMISSION."

MARTIN EDWARD BRADLEY

Doctor of Philosophy

THE UNIVERSITY OF ASTON IN BIRMINGHAM

September 1988

THE UNIVERSITY OF ASTON IN BIRMINGHAM

"AN ADVANCED THYRISTOR VALVE FOR HVDC TRANSMISSION."

Martin Edward Bradley

Thesis submitted for the degree of Doctor of Philosophy, 1988.

SUMMARY.

The development of an advanced outdoor valve requires coordinated research in the areas of light-triggered self-protecting thyristors, light triggering systems, insulation, cooling and mechanical design aspects. This thesis addresses the first two areas primarily, with a conceptual discussion of the remainder.

Using the experience gained from evaluation of a prototype thyristor and computer modelling of turn-on behaviour, a light-triggered thyristor with immunity to damage from weak optical triggering and dv/dt triggering was designed, manufactured and evaluated. The optical turn-on process was investigated by measuring currents and voltages in the gate structure during turn-on, and this yielded insights not obtained through conventional measurement techniques. The mechanism by which the thyristor was immune to weak triggering damage is explained, and techniques for optimising the design of the gate structure are proposed.

The most significant achievement, however, was the first demonstration of the feasibility of self-protection against forward recovery failure damage. Using the Selective Failure Zone technique, thyristors were made to gate themselves into conduction under forward recovery failure conditions. Furthermore, this was achieved without the need for complex structures or high levels of irradiation. The performance of the devices was limited by the inrush capability of the Zones, but it is believed that this can be improved by conventional means.

A light triggering system was developed using semiconductor lasers, and this incorporated several improvements over prior art in terms of optical performance and flexibility. In addition, other features of an outdoor valve are discussed, including insulation and cooling systems as well as practical considerations such as maintenance philosophy.

The work was carried out at GEC Transmission and Distribution Projects Ltd, Stafford, and Marconi Electronic Devices Ltd, Lincoln, as part of a company-wide development program supported by the Department of Trade and Industry.

Key words: High Voltage Direct Current Transmission; Thyristor Valve; Light-Triggered Thyristor.

ACKNOWLEDGEMENTS.

Many people have taken direct and indirect parts in the work described in this thesis. I would like to express my appreciation to:

- My supervisors Dr Rod Latham, Dr Alan Bowden and Dr David van Rest, and my industrial supervisor Mr Mike Woodhouse, for their guidance, suggestions and encouragement. Also Mr Derek Lobley, who was responsible at TDPL management level for the project, and who was a source of valuable encouragement.

- The team at MEDL Lincoln, with whom it was a pleasure to work. In particular:

  David Crees and Paul Taylor, who were responsible for managing the thyristor development work, and who gave me full support and cooperation;

  Ashley Plumpton and Dave Courtney, who were responsible for the detailed design and much of the evaluation of the advanced thyristors;

  Lyn Harding and Mark Hancock, who carried out most of the manufacture and preparation of samples.

- Vic Humberstone and the team at PA Technology who carried out the development of the valve light triggering system.

- My wife, for her long-suffering nature.

The work was carried out under the auspices of the Interdisciplinary Higher Degree Scheme at the University of Aston in Birmingham.

This thesis is dedicated to our daughter Emma, without whose arrival it would have been completed sooner but with much less sense of achievement.

CONTENTS.

LIST OF TABLES.

LIST OF FIGURES.

CHAPTER ONE: INTRODUCTION TO THE PROJECT.

1.1 INTRODUCTION.

HVDC Transmission is the transmission of electrical power at high voltage by direct current (DC) rather than the more usual alternating current (AC). AC is normally used in power networks because it is naturally produced by rotating generating machines, and can be readily transformed from one voltage to another. This facilitates long distance transmission at high voltage (to reduce losses) followed by distribution to consumers at low voltage. Since all modern generators produce AC, the transmission of power by DC therefore requires rectification (AC to DC) at one end of the line and inversion (DC to AC) at the other. The converters required to carry out this rectification and inversion are expensive, but despite this there are situations where the use of DC transmission is attractive. These situations are explained in Section 1.2 below, and a brief introduction to converter technology is also given. The history of HVDC transmission is summarised in Section 1.3, and GEC's involvement in the field is set out in Section 1.4.

The conversion from AC to DC and back again requires high-voltage high-current switches operating at the same frequency as the AC system. Originally mercury-arc valves were used for this purpose, but today the switching function is performed by series strings of semiconductor thyristors. A short history of the development of thyristor valves is given in Section 1.5, and Section 1.6 then introduces the concept of the advanced thyristor valve, with which this thesis is concerned.

Section 1.7 describes the structure of the thesis, and the author's contribution to the work is outlined in Section 1.8.

1.2 APPLICATIONS OF HVDC TRANSMISSION.

The following paragraphs outline the situations where transmission by DC may be preferred to AC, for either economic or technical reasons.

a) Very Long Distance Transmission.

All modern power systems are based on 3-phase AC, and consequently an AC transmission line requires a minimum of three conductors. In contrast, a DC line only requires two for the same power level, and over several hundred miles the saving of this third conductor can more than offset the costs of the converter stations. If it is acceptable to use ground or sea as a return path, only one conductor is required, yielding even further savings. It is possible to calculate a "break-even" distance beyond which DC is cheaper than AC, but this can be misleading since the figure is very sensitive to parameters specific to each situation.

b) Cable Transmission.

Underground and undersea cables have a very high capacitance to ground per unit length, and for AC transmission this capacitance has to be charged and discharged twice every cycle. For a few tens of miles of cable, this charging current can reach the full rated capacity of the conductor, which effectively limits the maximum distance over which AC can be transmitted. Beyond this, DC must be used, which only requires charging the cable capacitance when the line is first energised.

Cables are also much more expensive per mile than overhead lines, and so the cost considerations of a) apply again, but with breakeven distances reduced to the order of tens of miles.

c) Interconnection of Asynchronous Power Systems.

If power is to be exchanged between two power systems that are asynchronous (either of different frequency or the same nominal frequency but unrelated phase), only DC can be used which by nature is asynchronous. If the two systems are geographically remote, a DC transmission line or cable may be employed, but sometimes the rectifier and inverter are located in the same building with AC connections to both sides. This is known as a back-to-back link.

d) Others.

There are other technical advantages of DC which may lead to its use in preference to AC in some situations. These stem largely from the fact that the power flow along a DC line is much more controllable than that of an AC line, and this may be beneficial or essential in some circumstances. This controllability can be used to stabilize a power system and in some cases increase the usable power capacity of an existing AC connection.

If a fault occurs at the end of an AC line, the line will usually contribute around 10 times its rated current into the fault, whereas a DC link will automatically sense the fault and stop transmission within one cycle. This means that extra power can be brought into a system using DC without increasing its fault-current level (and therefore the rating required of the system switchgear).

e) Conclusions.

Although there are some cases where DC is the only possible means of transmission, in most situations the justification for using HVDC is

based on a combination of technical and economic factors. Many of the most obvious applications for HVDC have already been carried out, and so the majority of potential new business is in the area of marginal applications, where it is either DC versus AC, or DC versus "do nothing". This has led to an increasing emphasis on the economics and technical performance of HVDC equipment, and the advanced valve described in this thesis will, if successful, improve converter technology in both of these areas.

1.3 THE HISTORY OF HVDC TRANSMISSION.

The very first power systems all operated with direct current, an example being Edison's Pearl Street network which was installed in 1882 [1.1]. However, the advent of transformers, polyphase circuits and induction motors rapidly led to the use of AC in preference to DC. In spite of this, the advantages outlined in Section 1.2 were still recognised, and at least 19 DC transmission schemes of up to 20MW were installed in Europe between 1880 and 1911. These schemes formed part of AC power systems, and they used series connection of DC generators and motors to achieve reasonable transmission voltages. At the receiving end the motors were used to drive AC generators. However, this technology could not really be extended to the levels of power transfer that became desirable, and the last system was dismantled in 1937.

Between the wars, some novel techniques for AC-DC conversion were proposed, but the first practical approach was based around a thyratron valve. A 5.25MW prototype system using this technology was built by General Electric (USA) in 1936. In 1940 the thyratrons were replaced with mercury-arc valves ("ignitrons"), but system operation was discontinued in 1945.

Some experiments using mercury-arc converters were carried out in Switzerland (by Brown-Boveri) and Germany (by Siemens and AEG) around the time of the Second World War. However, after the war development of mercury-arc valves was pursued only by the USSR, English Electric (now GEC) and ASEA of Sweden, of which ASEA was the most successful. HVDC transmission is particularly suited to the geography of Sweden, since the main hydroelectric sites are in the north and the principal load centres in the south. There are also attractive applications for

undersea cables.

In 1954 ASEA constructed a 20MW mercury-arc valve link between the Swedish mainland and the island of Gotland. This used a single conductor cable, with the return path being through the sea and earth. It can be regarded as the first commercial application of HVDC, since it was considered to be more economical than providing extra thermal generation on the island.

The next scheme to go into operation after Gotland was the 160MW link between England and France ("Cross Channel 1"), also built by ASEA. Following this a number of mercury-arc schemes were constructed by ASEA, English Electric and the USSR, as listed in Table 1.1.

The late 1950's saw the invention of the silicon controlled rectifier, or "thyristor". After several years of development, thyristor voltage and current ratings reached levels that were usable for HVDC, and it was these devices that provided the next milestone in the historical development of HVDC. Although initially the capital cost of thyristor valves was approximately equal to that of mercury-arc valves, the thyristor alternative offered lower maintenance and an elimination of some undesirable operational features of mercury-arc devices (such as "arc-back"). The Nelson River Bipole 1 scheme, completed between 1974 and 1977, was the last scheme to be built with mercury-arc valves.

In 1970 ASEA added a thyristor valve "extension" to the Gotland Link. In 1971, English Electric replaced a valve of the Cross-Channel link with an oil-immersed thyristor valve. The first scheme to be completely designed and built with thyristor valves was the 350MW Eel River back-to-back link, connecting the New Brunswick and Hydro-Quebec power

systems. This scheme, commissioned in 1972, marked the return of General Electric (USA) to the field of HVDC transmission after an absence of over 20 years. The good experience gained on Eel River persuaded electric utilities that thyristor valves were superior to mercury-arc valves, and no new orders for mercury-arc schemes were placed after this.

The next HVDC scheme to be commissioned after Eel River was the Cabora-Bassa link between Mozambique and South Africa. This scheme was built by a German-Swiss consortium of Brown-Boveri, AEG and Siemens, known as the HVDC Working Group. It is of particular historical interest because it was originally designed around mercury-arc valves, but in 1969 the decision was made that the Working Group would develop their own thyristor valves. In the absence of any prior experience the valves were designed as outdoor oil-immersed units. For later schemes, the Working Group changed to an indoor air-insulated valve design, as adopted by GE for the Eel River scheme. Cabora-Bassa was commissioned in stages between 1977 and 1979.

A summary of HVDC schemes in service is given in Tables 1.1 and 1.2.

Since the adoption of thyristors, the main developments in HVDC have been in the thyristors themselves and the converter control systems. The Eel River thyristors were rated at 2.6kV, 600A, whereas devices are now available at around 5kV, 3000A or higher, representing an almost ten-fold increase in switching power. This has had a major impact on the cost of valves. The valve control systems have been refined and extended in scope to provide features such as better protection, fast recovery of transmission after system faults, and modulation of the

real and reactive power flows in the link to enhance AC system stability. In Section 1.5, the development of thyristor valve technology is discussed in more detail.

1.4 THE ROLE OF GEC IN HVDC TRANSMISSION.

At the end of the 1960's GEC merged with English Electric, and the work of English Electric in the field of HVDC transmission now continues under the auspices of GEC Transmission and Distribution Projects Ltd (TDPL). English Electric had constructed the following HVDC transmission schemes:

- Italy-Sardinia, 200MW at 200kV;
- Kingsnorth, 640MW at ±266kV;
- Nelson River Bipole 1, 1620MW at ±450kV;
- Thyristor valve replacement in the first Cross-Channel link.

Since the merger, GEC has constructed the second Cross-Channel link, which at 2000MW and ±270kV is the largest submarine HVDC link in the world. The company is now constructing a back-to-back link in Alberta, Canada.

Whilst construction of the Cross-Channel link was underway, the company started a major development exercise to bring its HVDC technology up to and beyond the present state-of-the-art. This development exercise covered all aspects of HVDC technology, and the work described in this thesis forms part of the overall development program. Many of the incremental innovations developed by GEC are being incorporated into the back-to-back link currently under construction; amongst the novel features of this contract are the absence of a DC smoothing reactor and the use of single-circuit water-glycol cooling. The company presently employs approximately 500 people, and is also involved in the areas of static reactive power compensators, industrial rectifiers, traction equipment and AC system studies.

1.5 THE DEVELOPMENT OF THYRISTOR VALVE TECHNOLOGY.

This section briefly reviews the historical development of thyristor valves, with particular reference to the types of insulation and cooling systems adopted. This emphasis is adopted because the advanced valve concept to be introduced in Section 1.6 incorporates a change from the now conventional indoor, air-insulated approach to the use of outdoor valves. This has important implications for the insulation and cooling technologies applied in the valve. Whilst there have been developments in thyristors and valve electrical circuit design, these have been incremental in nature, and the basic electrical features of an HVDC valve have changed little. In particular, although there have been dramatic improvements in thyristor performance in terms of voltage, current and other ratings, today's devices are functionally identical to those employed in the Eel River valves in 1972.

Appendix One includes a literature review of the various valve insulation and cooling techniques that have been adopted in the past. The type of valve insulation system used is strongly influenced by the fact that conventional thyristors require an electrical gate pulse to trigger them into conduction. Since the thyristor may be at a potential of several hundred kV from ground, communicating a gate pulse to the thyristor presents difficulties. Some early valve designs employed high voltage pulse transformers to bridge this insulation gap, but these were fairly rapidly superseded (partly due to poor reliability) by local electronic circuits deriving their power supply from the voltage across the thyristor. Firing pulses were issued to the thyristor in response to signals received from ground level by a fibre-optic link.

Optical fibres are ideal for this high-voltage, high-interference environment.

In addition to providing normal trigger pulses, these local electronic circuits are also used to protectively fire the thyristor if it is in danger of being damaged (see Chapter 2). This makes the circuits fairly complex, with a correspondingly high component count. Although they have proved more reliable than the pulse transformer system, their component count still gives them a failure rate which is significant given the number of thyristors in a converter. In order to accommodate this (as well as the less frequent failure of other components), it is standard practice to include extra "redundant" thyristor levels into the valve, so that the valve can continue to operate with one or more thyristor levels failed to short-circuit. The converter is normally shut down annually to allow failed units to be replaced. This requirement for regular access has been a major constraint in the development of valve insulation systems.

In the early days of thyristor valve technology, two distinct design philosophies were developed. ASEA and GE adopted indoor designs based around air insulation and forced air cooling, whereas the HVDC Working Group and English Electric opted for oil-insulated designs, cooled by forced circulation of the oil. Toshiba and Hitachi in Japan also licensed this technology from the Working Group. However, practical experience with oil-insulated valves soon indicated that, because of the complexity and consequent unreliability of the electrical equipment described above, air-insulation was to be preferred because of its greater ease of access for maintenance. The use of indoor, air-insulated valves was therefore adopted by all manufacturers in the field.

The next major step in the development of valves was the introduction of water cooling. Although HVDC converters are very efficient overall, the losses in individual thyristors can be hundreds or thousands of watts, and this heat has to be removed efficiently to keep the silicon temperature down to an acceptable level. Water is intrinsically more effective than air as a cooling medium due to its higher thermal mass per unit volume, and it had already been used for mercury-arc valves. When applied to thyristor valves it not only allowed the construction of much more compact equipment, but also gave useful reductions in the electrical power consumption of the cooling plant itself. Water cooling of thyristors is now used by almost all manufacturers.

Because of the ease of access for maintenance and repairs, air insulation has been the standard technology for valves for the past decade, and it is likely to remain so for several years to come. However, during the early 1970's it was believed that there might be a need for valves that were more compact than could be achieved with air-insulated technology. In cities where the electricity demand was growing and expected to exceed the capacity of the existing AC network, the use of HVDC infeeds to strengthen the system would be preferable, for the reasons explained in Sections 1.2b) and d). Because of the high cost of land in such urban areas, these HVDC converters would need to be extremely compact.

To meet this anticipated requirement for compact converters, GE embarked on a programme to develop outdoor, SF_6-insulated, metal-clad HVDC equipment, including thyristor valves. This was intended to yield significant savings in land requirements over conventional air-insulated equipment. The valves were constructed using conventional electri-

cally-triggered thyristors with local electronic circuits, and the thyristors were cooled by circulating liquid Freon through enclosed heatsinks. Freon was used because the insulating properties of the SF_6 gas would not be degraded if there was a leak in the cooling circuit. The valve enclosures had a large removable panel to give access to the thyristor equipment for maintenance, and a spare valve was provided so that transmission could be maintained while a valve was being serviced. Valve maintenance was to take place in a special building provided on site.

Although a prototype installation rated at 100MW was completed in 1977, at least six years later it was still not commissioned, and so it must be presumed that unexpected difficulties were encountered with the equipment. Whilst this prototype station was under construction, some other manufacturers produced conceptual designs for schemes using metal-clad valves, but none pursued developments beyond that stage. At least one argued that if air-insulated valves were suitably designed, there was little saving in space to be achieved by using metal-clad technology, and that there was much more scope for space reduction in other areas of converter station equipment (most notably AC harmonic filters).

In spite of earlier predictions, the demand for compact urban converter stations did not materialise, and commercial installations continued to use air-insulated valves. Interest in compact valves died away until the early 1980's, when the Kansai Electric Power Co (a Japanese utility) and the Mitsubishi Electric Co started work on a joint development programme to build high pressure SF_6-insulated HVDC equipment. The GE valves had used SF_6 at only slightly above atmospheric

pressure in order to avoid strict US pressure vessel regulations; however, the Japanese design used SF_6 at several times atmospheric pressure, which allowed the valves to be more compact, and enabled the thyristors to be cooled by blowing the gas over open finned heatsinks. The cooling capacity of pressurised SF_6 is substantially greater than that of air, so that reasonably compact heatsinks could be used. No separate coolant circuit was therefore necessary.

As with the GE design, a spare valve was provided, but due to the increased need for cleanliness associated with high pressure gas insulation, valves requiring maintenance were to be returned to the factory. This equipment was developed because it was expected that Japan would move towards the use of large nuclear generation "parks" in remote areas, feeding urban loads via long HVDC lines with compact converter stations at the receiving end. Two prototype installations were completed in 1983, but as yet there has been no commercial application of this technology.

At the time of writing, the most common valve technology still embodies air insulation and water cooling. SF_6-insulated/SF_6-cooled valves and air-insulated/Freon-cooled valves have been offered commercially, but without success to date.

1.6 THE ADVANCED VALVE CONCEPT.

The finite failure rate of the present generation of thyristor valve electronics incurs important economic penalties for the utility. As mentioned earlier, it is necessary to provide regular maintenance intervals for the valves, which results in a loss of revenue to the utility due to non-operation of the equipment. Another consequence is that a benign environment must be provided to permit access for maintenance. Furthermore, in order to allow for component failures it is necessary to include extra redundant thyristor levels into a valve, increasing the capital cost and conduction losses. If a means can be found of eliminating the complex local electronics provided for thyristors in the present generation of equipment, this would considerably improve valve reliability, since the remaining valve components are inherently much more reliable. This is clear from the fact that, in general, the local electronics package is considered to have a failure rate five to six times higher than the rest of the thyristor-level components combined. If this source of unreliability was eliminated, lower levels of redundancy, longer intervals between maintenance and a more economic valve environment could be used.

There are several ways in which this desirable situation could be achieved:

- Integrate all of the functions presently provided by the local electronics onto the power thyristor slice, and use direct optical triggering. Such a device (as yet not in existence) is referred to as a light-triggered self-protecting thyristor;

- Re-design the triggering and protection philosophy so that those functions still needed at thyristor-level can be provided by simple "passive" components (ie, not requiring a power supply);
- Some combination of the above.

The advantages of eliminating thyristor-level electronics are recognised by all HVDC manufacturers, and in attempting to achieve this objective, all three of the above approaches have been tried. If a valve without complex local electronics could be developed, its reliability should be comparable with other major items of conventional equipment such as power transformers and circuit breakers. The valves could then be treated in the same way as these other items of equipment, ie, enclosed in steel tanks filled with either liquid or gaseous dielectric. This would protect the valves from contamination, and also enable them to be installed outdoors rather than in a special building. The main purpose of this encapsulation would not be to achieve compactness (although this benefit may also accrue), but to reduce the cost of civil works necessary for the converter station, and make the valves compatible with other items of metal-clad equipment. Such valves could be built and tested at the manufacturer's works, in order to reduce construction and commissioning times and minimise the risk of contamination or damage at site.

HVDC schemes today still only account for a tiny fraction of the total installed capacity of transmission equipment, and because of the special environmental facilities required by the valves, the converters are in a different class to conventional power system equipment. If the advanced valve described above became a reality, this would help to make HVDC converters a standard piece of equipment for electric

utilities. The purpose of the work described in this thesis is to advance the technologies necessary to achieve that goal.

1.7 THESIS STRUCTURE.

Chapter Two describes the present generation of valve technology based upon conventional electrically-triggered thyristors. Chapter Three then outlines the various technologies needed for the advanced valve concept to become a reality, and describes briefly some of the interdependencies between them. The need to coordinate work in the different technological areas of the advanced valve project is explained.

The development of a light-triggered self-protecting thyristor is the main technological breakthrough necessary for the advanced valve to become a reality, and therefore work was concentrated in this area. Chapters Four to Six describe this work as follows:

- Chapter Four: Development of the Thyristor, Stage 1. Description of the principles of light triggering and self-protection, the design philosophy of the first prototype thyristor, and the evaluation work carried out up to the point at which the author became involved.

- Chapter Five: Development of the Thyristor, Stage 2. Description of the evaluation work carried out by the author, or under his supervision, on the first prototype design.

- Chapter Six: Development of the Thyristor, Stage 3. Description of the design, manufacture and evaluation of a new set of test structures aimed at demonstrating the principles of light-triggering and self-protection (Note that "structures" in this context refers to special designs of silicon devices, some of which are not thyristors in the conventional sense of the term).

During the project, an optical system capable of firing a valve made up

of light-triggered thyristors was developed, and this is described in Chapter Seven.

Chapter Eight then reviews and summarises the work carried out in the project, and draws conclusions concerning the successfulness of each aspect of the work.

In addition to the main technical work covered in Chapters Four to Seven, consideration was also given to the insulation and cooling systems that would be used in an advanced outdoor valve. The results of these considerations are given in Appendix One. In Appendix Two, aspects of the mechanical design of an outdoor valve are discussed.

1.8 AUTHOR'S CONTRIBUTION.

Because of the scope of the project and the nature of the work involved, it was neither practical nor desirable for the author to work in isolation. On the contrary, the need to communicate and work with people from different technological specialities was essential for a successful outcome, and this is discussed further in Chapter Three. This section sets out the author's contribution to each of the areas covered in the thesis.

a) Thyristor Development.

It was originally conceived that the development of the light-triggered self-protecting thyristor would be carried out at Marconi Electronic Devices Ltd (MEDL) in Lincoln, with TDPL's involvement in the work consisting mainly of discussion of requirements and communication of results. Chapter Four describes the work carried out at MEDL under this arrangement. The author was involved in this discussion and communication as part of the development team at TDPL. However, in the course of the work, it became apparent that a closer level of technical liaison was required between MEDL and TDPL. The author was therefore seconded to MEDL for a period of four months, spending four days a week at MEDL and the remaining day at TDPL reporting on progress and maintaining links at Stafford.

It was during these four months at MEDL that much of the evaluation of the first prototype thyristor was carried out. In this period the author was primarily responsible for the evaluation of the light triggering and forward recovery protection functions, with assistance and safety cover from others who had little or no prior experience of

thyristor evaluation. Preparation of thyristor samples and investigation of built-in overvoltage protection was carried out by MEDL staff, with the author acting in an advisory role.

At the end of the four months, tne staff at MEDL had become sufficiently competent to carry out evaluation work with only general supervision and direction. The author therefore continued to visit MEDL one day a week to guide the ongoing work. Chapter Five describes the work carried out in this period as well as the original four months secondment at MEDL.

When the time came to design the new set of test structures (Chapter Six), the author was heavily involved in the conceptual design of these devices. The detailed design and calculations were carried out by MEDL staff. These devices were evaluated both at MEDL and the GEC Engineering Research Centre at Stafford, under the supervision of the author.

b) Light Triggering System Development.

Development of the valve light triggering system was subcontracted to a technology consultancy, PA Technology in Cambridge. At the start of the work, the author's involvement was simply attending the progress meetings with the consultants. However, by the end of the development (which lasted over a year) the author had main technical responsibility at TDPL for the work being carried out, and was involved in the strategic technical decisions as well as the development of the system philosophy. This was a deliberate policy on the part of TDPL management to ensure that the company gained experience as well as the prototype system. The work carried out is described in Chapter Seven.

c) Valve Insulation, Cooling and Mechanical Design.

The discussion of the insulation, cooling and mechanical design aspects of an outdoor valve given in Appendices One and Two was prepared by the author through study of the literature and consultation with GEC experts in the appropriate areas.

1.9 CONCLUSIONS.

This chapter has introduced the technology of HVDC transmission in terms of its applications and history, and covered the role of GEC in the field. The development of thyristor valves up to the present day has been described, and the concept of the advanced valve, with which this thesis is concerned, has been introduced. This has set the global context for the project. Chapters Two and Three go into more detail concerning the present state-of-the-art and the technologies required for the advanced valve, whilst Chapter Four opens the major technical content of the thesis.

The structure of the thesis and the author's contribution to the work have been outlined.

CHAPTER TWO: CONVENTIONAL THYRISTOR VALVE TECHNOLOGY.

2.1 INTRODUCTION.

Chapter 1 introduced the general principles of HVDC transmission and the concept of the advanced valve. In order for the issues and technical challenges involved in the advanced valve to be understood more clearly, particularly with respect to the light-triggered thyristor, this Chapter explains the conventional design of thyristor valves for HVDC transmission.

For an HVDC scheme, the operating voltage is selected to minimise the total cost of the scheme, including losses. For a long-distance transmission scheme, this usually results in a transmission voltage of several hundred kV, whereas for a back-to-back scheme, several tens of kV is more common. However, in both cases the operating voltage is well beyond the capability of any single semiconductor device, and therefore it is necessary to connect many thyristors in series to construct a suitable thyristor valve. These valves are then arranged in 12-pulse bridges (ie groups of 12 valves) in order to carry out the rectification and inversion; see Figure 2.1.1.

This Chapter discusses the operating stresses on the thyristor in an HVDC valve, some of which are inherent in the operation of a converter bridge, whereas others arise from the need to connect devices in series in order to achieve the desired operating voltage. More detail on this subject can be found in Reference [2.1].

2.2 THYRISTOR-LEVEL CIRCUIT DESIGN.

Associated with each thyristor in a valve is a set of components which, together with the thyristor, make up what is known as a thyristor level. The functions of these components are introduced under the following headings: The Off-State Thyristor; Thyristor Turn-On; Thyristor Turn-Off; Thyristor-Level Electronics.

a) The Off-State Thyristor.

If thyristors are connected in series as shown in Figure 2.1.1, the first problem that is evident is that of unequal sharing of voltage when the thyristors are not conducting. Each thyristor will have a different leakage current characteristic, and therefore since the same current must flow through each device, the voltage will not be shared equally between them; see Figure 2.2.1. It is therefore necessary to provide a grading circuit in order to smooth out the voltage distribution. If a resistor of suitable value is connected across each thyristor, then the current down the resistor chain will effectively dominate the variations in leakage current between thyristors, and an even voltage distribution will be imposed on the thyristor string. See Figure 2.2.2.

These resistors (known as "DC grading resistors") will provide grading for a direct voltage across the valve. However, if a rapidly changing voltage is applied across the thyristors, it will be distributed according to the stray capacitance in each level. Due to the physical design of the valve, it is likely that each level will have a different stray capacitance associated with it, and since the same current must flow through each stray capacitance, once again an uneven voltage

distribution will result. In order to enforce even distribution of transient voltages, it is therefore necessary to connect a capacitor across each thyristor level. (Alternatively, it may be connected across a group of thyristor levels of near-identical stray capacitance). The value of this capacitor is such that it dominates the effect of local variations in stray capacitance, and it is referred to as the "Fast Grading Capacitor"; see Figure 2.2.3.

b) Thyristor Turn-On.

If a thyristor is forward biased and is given a gate pulse, it will start to conduct. However, the whole thyristor does not turn on immediately, but rather it takes a finite time for the conducting area to spread from the gate region to occupy the full area of the device. During the early stages of turn-on, the area of silicon that is conducting is small, and so it is vulnerable to damage due to excessive current densities even though the total level of current may be relatively small.

In an HVDC converter, the thyristor valve is fed from the secondary winding of a transformer whose reactance is fairly high (typically 10 - 20%). Therefore the rate of rise of current from the transformer into the thyristor at turn-on is relatively slow, normally a few amps per microsecond. However, there are several sources of capacitance directly coupled to the valve: in addition to the fast grading capacitors described in the previous section, there is substantial stray capacitance in the transformer winding and, to a lesser extent, in the busbars and valve structure. When all this capacitance is referred to one thyristor level, its value can reach as much as 1uF; see Figure 2.2.4. The inrush current from this capacitance at turn-on would easily

be sufficient to destroy the thyristor.

It is therefore necessary to include inductance in the circuit to limit the rate of rise of current from the stray capacitance at turn-on. A conventional linear inductor cannot be used, since this would give the valve unacceptable operating characteristics. A saturating reactor is therefore used, which has a high inductance at the low currents present during the early stages of turn-on, but has low or zero incremental inductance at rated current. This is shown in Figure 2.2.5. (As with the fast grading capacitor, one reactor may be used to serve a number of thyristor levels).

Having put inductance in the valve, however, we have now created an oscillatory circuit between the stray capacitance across the valve and the inductance in the valve. In order to limit undesirable current oscillations, damping is therefore included in the circuit. This is achieved either by connecting a resistor across the inductor (possibly with a series diode to prevent inrush current bypassing the inductor), or by using an imperfect, "lossy" inductor which has a parasitic resistance. See Figure 2.2.6.

c) Thyristor Turn-Off.

At the end of the conducting period, the current through the thyristor is brought to zero by the action of the external circuit. However, the thyristor is not a perfect switch, and some reverse current will flow through it before it regains the ability to block reverse voltage. The form of this current flow is shown in Figure 2.2.7.

At the point when the thyristor starts to block reverse voltage (point T_1 in Figure 2.2.7) there is significant reverse current flowing

through the thyristor, and therefore energy is stored in the circuit inductance. At time T_1, this energy starts to transfer into the thyristor level capacitance, and would produce a large reverse voltage transient. To limit the magnitude of this transient, a damping circuit is connected across the thyristor consisting of a resistor in series with a capacitor. See Figure 2.2.8. During the recovery event, energy is dissipated in the damping resistor as the capacitor charges up, which reduces the magnitude of the voltage overshoot.

In addition to damping the recovery transient at turn-off, the damping circuit also performs several other functions, as follows:

- Since the circuit has a relatively low impedance at intermediate frequencies (10's to 1000's of Hz), it makes an important contribution to achieving even voltage distribution in the valve.

- Another consequence of its relatively low impedance at high frequencies is that when rapid rates of change of voltage are applied to the valve (for example due to a flashover), it forms a potential divider with the valve inductance such that only a small proportion of the dv/dt appears across the thyristor itself. The majority of the transient voltage is supported by the inductance, which allows time for the thyristor to be protectively triggered into conduction.

- Figure 2.2.7 showed the reverse recovery current flow in a thyristor. In practice, the integral of this reverse current (known as the thyristor "stored charge") will be different for each thyristor. Since the total current flow in each level must be the same, the difference current flows in the thyristor-level capacitance, producing a voltage offset on the level. The damping capacitor is sized to keep this offset to acceptable levels.

- When the thyristor turns on, the damping capacitor discharges into it, giving an initial step of current with a decay time constant usually of several tens of microseconds. This helps to prevent the thyristor current falling too low at any point during the oscillatory inrush transient. If the current does fall too low, the conducting area of the device may stop spreading, or even start to shrink.

d) Thyristor-Level Electronics.

For the valve to function, the control system must be able to send a gate pulse simultaneously to each thyristor in the valve at the required time. However, the thyristor level may be at a potential of up to 500kV or more to ground, and therefore it is impractical to have an electrical connection between the control system at ground level and the thyristor. Fibre-optic communication is therefore used to bridge this potential difference.

However, the use of fibre-optics means that there must be an active receiver at the thyristor level, as well as a means of supplying gate pulses to the thyristor. These functions are implemented by local electronic circuits, the power for these circuits being derived from the current flowing in one of the passive components in the thyristor level. See Figure 2.2.9.

By using this method to power the thyristor-level electronics, two limitations are incurred. Firstly, it normally takes several cycles of power frequency voltage to charge up the power supplies, and so the valves cannot be operated as soon as voltage is applied. If a rapid response is required from the converter, it may therefore be necessary to keep the valves energised continuously, which incurs losses in the

grading circuits. Secondly, in some situations it may be necessary to supply multiple gate pulses to the thyristor with little or no voltage appearing on the valve between them. In this case, the energy storage in the power supply must be sufficient to deliver these pulses without running down. This requires a large energy storage capacitor. One of the advantages of direct light triggering of thyristors is that the triggering energy is supplied from ground level, and therefore the valves are not subject to these limitations.

Given that a power supply is available at the thyristor-level, it is then convenient to provide various protective circuits local to the thyristor. Section 2.3 outlines the stresses that are experienced by a thyristor under normal and abnormal operating conditions, and describes the different protection strategies that are employed to prevent the thyristors being damaged.

2.3 STRESSES ON THE THYRISTOR.

Section 2.2 introduced the basic factors to be considered when designing an HVDC valve, some of which centred around the need to protect the thyristors from potentially damaging stresses. This section discusses in more detail the nature of these stresses, and the techniques used to restrict them to acceptable levels. Some of these stresses are normal operating phenomena which are repetitive, whereas others are transient conditions which may only occur a few times during the life of the equipment.

a) Protection Techniques.

Before considering how to protect the thyristor in any particular situation, it is useful to summarise the three general techniques which are used to protect the thyristors from damage. These are as follows:

- Passive circuit design;
- Thyristor design;
- Local electronic protection.

By careful design of the valve and the power circuit passive components in the converter, it is possible to reduce some of the stresses that may appear on the thyristors. In addition, the thyristor characteristics can be designed to give improved immunity to damage in some areas (although usually at the cost of some deterioration in other parameters). Finally, local electronic circuits can be provided to protect the thyristor, as follows. Many thyristor failure modes involve spurious turn-on of the thyristor, and damage can be avoided if this is pre-empted by a proper gate pulse being sent to the device. This requires the rapid detection of a potentially dangerous situation and

the sending of a gate pulse. Whilst this unscheduled turn-on of the thyristor may cause a temporary disturbance to converter operation, this is accepted in order to prevent damage to the device.

In many cases, all three of the above techniques are used in combination to protect the thyristor, and this requires coordination between the thyristor design, the valve design and the electronics design in order to achieve an optimised result. This need for coordination applies to the advanced valve design as well, and is discussed in this respect in Chapter Three.

b) Stresses on the Off-State Thyristor.

Overvoltage Stress.

In normal operation, the off-state thyristor experiences voltages which are typically 50 - 60% of its rating. There is therefore no danger of overvoltage damage to the thyristor under these circumstances. However, various situations can arise where the thyristor is stressed to the limit of its capability. The most important cases are as follows:

- Where one thyristor in a valve does not receive normal gate pulses;
- Where switching surges or lightning impulses are transferred from the AC system onto the valves via the converter transformer;
- Where a flashover occurs inside the valve hall.

In the first case, if all but one of the thyristors in a valve are fired, the remaining thyristor will potentially be exposed to the full valve firing voltage, which could be hundreds of kV. This would destroy it instantly, and so to prevent this, two techniques are used in combination. Firstly, the local electronics detects the overvoltage

before it reaches a destructive level, and sends a gate pulse to the thyristor. Secondly, the valve inductors and damping networks are designed to act together to limit the rate at which the voltage increases across the off-state level. This allows time for the electronics to fire the device and collapse the voltage on the level.

Switching surges and lightning impulses can appear across the valve in either polarity. In the reverse direction, the thyristor cannot be turned on to protect itself, and so a zinc-oxide surge arrester is connected across the valve to limit the magnitude of the voltage. The surge arrester is a highly non-linear resistor which carries minimal current at normal voltages, but has a very low incremental resistance at high voltages. By conducting hundreds or thousands of amps for a very short period, it limits the voltage transient across the valve. The level at which it limits the voltage determines the number of thyristors in the valve, and is known as the surge arrester protection level. Although the transient voltage may not be evenly shared within the valve (due for example to grading component tolerances), the more highly stressed thyristors are normally allowed to go into reverse avalanche. This prevents them supporting any further increase in voltage, whilst the voltage on the less highly stressed levels is still rising. The end result is near-perfect voltage distribution, which reduces the number of thyristors needed for a given surge arrester protection level. Thus the thyristor avalanche capability is used in conjunction with the passive converter components (in this case the surge arrester) to limit the stress on the device.

To keep the heat dissipated in the arrester during normal operation to acceptable levels, its characteristic must be chosen so that it

conducts very little current at normal voltages. This usually results in an arrester protective level which is about 1.8 times the normal peak working voltage.

If a switching surge or lightning impulse is applied in the forward direction, the thyristors cannot be allowed to go into forward avalanche because this would result in spurious turn-on, which would almost certainly be destructive. Therefore the whole valve is allowed to fire on its electronic overvoltage protection, before the transient reaches the surge arrester protective level. This firing will occur when the most highly stressed thyristor reaches its individual over-voltage firing level; when this level triggers, it causes a "domino" effect in the rest of the valve, otherwise known as "cascade firing". To prevent frequent occurrence of cascade firing, the utility normally specifies that it must not occur below a voltage of perhaps 1.5pu. The relationship between this level and the (say) 1.8pu surge arrester protection level determines the ratio between the overvoltage firing level in the forward direction and the thyristor avalanche level in the reverse direction. Allowing realistic grading errors for forward voltages, this means that the overvoltage firing level must not be below about 90% of the reverse avalanche level for an efficient valve design. This has important consequences when attempting to incorporate built-in overvoltage protection into the thyristor (Chapters Four and Five).

A flashover in the valve hall is much more severe than a switching surge or lightning impulse. The voltage across the valve terminals can reach the surge arrester protective level in less than one microsecond, and so the interaction between the valve inductors and damping capacitors becomes even more critical in holding off a forward voltage

until the electronics has had time to fire the thyristors. If the local overvoltage protection is implemented using Break-Over Diodes (BOD's), then the current drawn by these devices through the impedance of the valve inductors also helps to temporarily hold down the thyristor voltage.

Dv/dt Stress.

In addition to overvoltage, an off-state thyristor can also be damaged if a rapidly increasing forward voltage (or "dv/dt") is applied to it. This is because the dv/dt causes capacitive displacement currents to flow in the thyristor, which can cause spurious turn-on to occur (see Chapter Four). As with overvoltage, such turn-on is often destructive, and is known as "dv/dt failure".

High levels of dv/dt across the valve can occur as a result of the following events:

- firing of another valve in the converter from high voltage. The rapid voltage collapse on the valve being fired is coupled onto the off-state valves via stray capacitances, and in some cases will appear as a forward dv/dt;
- Lightning impulses transferred from the AC system onto the valves via the converter transformer;
- Flashovers in various locations, for example in the valve hall or inside the converter transformer.

As in the previous section, these events can appear as negative or positive impulses on the valves. Negative dv/dt's are not a threat to the thyristors, since they will not cause spurious turn-on.

Dv/dt turn-on in a thyristor is often a relatively slow process. This is because it creates weak gating of the device, resulting in a long turn-on delay time (perhaps several microseconds). Therefore if the positive impulse is of a large prospective amplitude, then the valve may fire on its overvoltage protection before damaging dv/dt turn-on could occur. However, some impulses will be fast enough to cause dv/dt turn-on without reaching the valve protective firing level. To prevent damaging spurious turn-on in these situations, it is necessary to electronically sense the dv/dt and issue a gate pulse to the device. This sensing is usually achieved by measuring the current flowing in the fast grading capacitor (Section 2.2a).

The thyristor has an inherent capability to withstand dv/dt up to a certain level without triggering, and this capability can be increased at the expense of other characteristics such as forward voltage drop. In addition, the valve circuit can be designed to reduce some of the dv/dt's which may appear across the thyristor. These techniques are used to ensure that the thyristors do not fire spuriously in response to dv/dt's experienced in normal operation. However, it would be very uneconomical to use these techniques to prevent dv/dt failure in all circumstances, and in any case there would still be some situations in which the capability of any practical thyristor would be exceeded. Therefore electronic protective firing of the thyristor is also used in conjunction with these other techniques.

In practice, if a large-amplitude positive impulse appears across the thyristor-level, the thyristor may be turned on either by dv/dt or overvoltage protection, depending on which mechanism operates first.

c) Stresses on the Thyristor During Turn-on.

The need to limit the rate of rise of current ("di/dt") into the thyristor at turn-on has already been explained in section 2.2b. To prevent excessive current density in the small area of the thyristor that is turned on initially, saturable inductance is included in the valve to limit the discharge of current from stray capacitances in the circuit. In addition to this use of passive inductance, the thyristor gate is made to turn on as large an area of silicon as is practical by the use of techniques such as "amplifying gates" (see Chapter Four). The design of the electronic gate unit is also an important part of the protection strategy, since if the gate pulse sent to the thyristor is too weak, then it may be destroyed even though the di/dt is within its theoretical rating. This is because a relatively strong gate pulse is needed to ensure that the whole of the thyristor gate is activated during turn-on, even though some proportion of the gate (with correspondingly reduced di/dt capability) may turn on in response to a weaker signal.

Failure due to excessive current density during turn-on is known as "di/dt failure".

d) Stresses on the Thyristor During Conduction.

The passage of normal load current presents no threat to the thyristors. However, large currents can flow in the valves if there is a flashover in the converter bridge. Under these circumstances, currents of around 10pu (possibly 30kA or more) may flow in the valves. Conventional circuit-breakers are used to clear these overcurrents within three or four cycles by isolating the converter from the AC system. On this time-scale, there is little risk of damage to the

thyristors from the overcurrents themselves, because once the thyristor is fully conducting, silicon temperatures of 300 to 400°C have to be reached before thermal runaway occurs. A far more significant risk exists from the voltage that may appear across the thyristor after it has been heated by a fault current, since the thyristor voltage rating falls off significantly between 100 and 200°C. It is the need to withstand voltage after a fault current that determines the level of current that can be tolerated. These current levels are controlled primarily by the leakage inductance of the converter transformer, which is an important design parameter. Voltage stresses after a fault current are considered further in the next section.

e) Stresses on the Thyristor at Turn-Off.

At turn-off ("recovery"), the thyristor is vulnerable to damage from excessive voltage and forward dv/dt, as in the off-state (section *b)*). However, its vulnerability to both these types of damage can be increased for the following reasons:

- Although current has stopped flowing through the device, there are initially high levels of free carriers in the silicon which have a relatively long thermal recombination time.
- If the thyristor has just carried a fault current, its junction temperature will be increased above normal steady-state levels. This causes a deterioration in most thyristor characteristics.

Recovery under Normal Conditions.

In normal conduction, the thyristor junctions are saturated with carriers. Immediately after current has ceased to flow in the device, there are still high levels of free carriers in the silicon, with the

result that the thyristor will turn on as soon as any forward voltage is applied to it. After a certain period (typically 300us), the free carrier levels have decayed to the point where the thyristor can support forward voltage. However, its ability to withstand forward dv/dt without spurious triggering remains substantially impaired, and this ability takes much longer to be recovered fully (possibly 2ms). In this period of recovery, there will be combinations of time and dv/dt at which the thyristor will only just turn on spuriously. This "weak spurious triggering" is the condition most likely to result in damage to the thyristor, since only a very small area of the silicon turns on. Failure under these conditions is known as "Forward Recovery Failure".

For a conventional thyristor, nothing can be done to reduce the likelihood of damage by this cause. There will always be a period between on-state and off-state during which the thyristor will be vulnerable. The duration of the period can be reduced by designing the thyristor to recover quickly, but this incurs significant penalties in device characteristics. Similarly, the valve circuit cannot be designed to eliminate the possibility of this type of damage. It is therefore necessary to program the thyristor's time-dependent recovery characteristic into the local protection electronics to a greater or lesser degree of accuracy. This protection system is known as Forward Recovery Protection (FRP).

The simplest version of this protection is to reduce the overvoltage firing level during the recovery period. This is based on the empirical evidence that, even if spurious turn-on does occur, it will not be damaging if the voltage at turn-on is less than a certain level (typically 500V). However, the coarse nature of this approach results

in the device being over-protected during much of the recovery period, so that unnecessary protective firing may result. This will disrupt operation of the converter. A more efficient technique is to program the thyristor's voltage and dv/dt capability during recovery into the local electronics, so that the protection settings are continually adjusted until true off-state is achieved. This minimises the incidence of protective firing.

At the instant of recovery, there is a reverse voltage overshoot as the circuit inductance goes into underdamped resonance with the valve damping circuits. For severe short-term operating conditions, the peak of this overshoot might be limited by the surge arrester. However, for recovery at normal conditions, the thyristor retains its full reverse voltage withstand capability at all times, and so there is no risk of damage.

Although it is usually conduction in the forward direction which creates free carriers in the silicon, giving rise to the need for forward recovery protection, conduction in the reverse avalanche mode will also generate carriers in the device. The thyristor therefore restarts its recovery process if it has gone into reverse avalanche, and the electronic protection should also be reset accordingly.

Recovery after a Fault Current.

At recovery after one loop of fault current (ie one power frequency cycle of asymmetric fault current), the thyristor junction temperature may be as much as 80°C hotter than it is under normal conditions. This not only extends the time for recovery by a significant amount, but even when true off-state is reached, the thyristor may have reduced forward voltage and dv/dt withstand capabilities due to its elevated

temperature. The valve circuit design, which is optimised for normal operation, will offer no protection under these conditions, and so the local electronics must reduce its overvoltage and dv/dt protection settings appropriately in order to protect the device. As the silicon cools, these settings can be restored to their normal levels.

From the thyristor's point of view, the best way to protect it after a fault would be to trigger it as soon as any forward voltage appeared across it. This would cause multiple loops of fault current to flow (presuming that the initial fault was still present), which would be cleared by the main circuit-breaker. However, this imposes arduous duty on the circuit-breaker and the converter transformer which has to carry the fault current, and so it is usually required that the valves do not fire protectively after a single loop of fault current. The above strategy is therefore unacceptable, and in effect, the valves must be used to suppress the fault before the circuit-breaker opens. If the valves are subjected to severely distorted voltage waveforms after the first loop of fault current, such that the thyristors are vulnerable to overvoltage or dv/dt failure, then it is accepted that the valves will fire protectively.

If the valve does refire after one fault loop, then the thyristors may be very hot after the second loop of fault current that will follow, possibly 200°C or more. At these temperatures they cannot support any significant forward voltage, and so the local electronics will trigger them if the valve voltage does go positive. However, they will still have to withstand the reverse recovery voltage transient at the end of the second fault loop. At the end of the third fault loop, the circuit-breaker contacts will normally be opening, so that the valves will be

isolated from the AC system and will not be subjected to any voltage.

2.4 CONCLUSIONS.

This chapter has provided a brief introduction to some of the principles of conventional HVDC valve design technology, followed by a consideration of the stresses on the thyristors in an HVDC valve. There are three different techniques available for preventing damage to the devices (ie passive circuit design, the design of the thyristor itself and the provision of local electronic protection circuits), and the ways in which these techniques are applied to keep each stress to acceptable levels have been described. In particular, the chapter has covered the functions which are presently implemented in the local electronic circuits, ie:

- The provision of adequate gate pulses to the device;
- Overvoltage protective firing;
- Dv/dt protective firing;
- Forward recovery protective firing.

It is these functions which must be provided either by the valve light triggering system (in the first case) or by the thyristor itself (for the others) if the local electronics is to be dispensed with in the advanced valve.

CHAPTER THREE: PROJECT STRUCTURE.

3.1 INTRODUCTION.

Chapter One has introduced the concept of an advanced thyristor valve for HVDC transmission, and Chapter Two explained conventional valve technology, some aspects of which will become obsolete in an advanced valve. For the proposed advanced valve to become a reality, five distinct areas of technology require development:

- The light-triggered self-protecting thyristor;
- The valve light triggering system;
- The valve insulation system;
- The valve cooling system;
- The valve mechanical arrangement.

All of these technologies interact with each other to varying degrees. Section 3.2 gives an introduction to the important issues in each of these technology areas, and the interactions between them are then described in Section 3.3. This illustrates the importance of effective coordination and direction of work in each technology area.

3.2 IMPORTANT ISSUES IN EACH OF THE TECHNOLOGY AREAS.

a) The Light-Triggered Self-Protecting Thyristor.

In conventional valves, an important part of the local electronics is the power supply which extracts energy from one of the thyristor-level grading circuits (cf Figure 2.2.9). This energy is used to supply the protection and triggering circuits. In order to eliminate the power supply, some form of direct optical triggering of a device is necessary. Two techniques can be used for this: the first involves connecting a small, light-triggered pilot thyristor between the anode and gate of a conventional electrically-triggered power thyristor, and the second involves making the main power thyristor optically sensitive. See Figure 3.2.1.

In the first technique, the pilot thyristor is fired optically from ground level. This then draws current from the voltage across the main thyristor, and uses it to trigger the electrical gate of that device. This approach is attractive in that if a poor yield is associated with an optically-sensitive structure, the cost implications of this can be reduced by producing many small chips from one silicon wafer. In addition, since the pilot thyristor is not conducting continually, its temperature remains low, and certain ratings (such as dV/dt capability) are more easily achieved at lower temperatures. In addition, a resistor is normally connected in series with the pilot device in order to limit the inrush current, so that turn-on duty in the light-triggered thyristor is reduced. However, the pilot device must still withstand the same voltage and dV/dt as the main thyristor, and if adequate yield can be achieved on the optical structure, it may be more economical to eliminate the extra pilot device, and integrate direct optical trigger-

ing onto the main power thyristor.

The technology for light-triggered pilot and main thyristors has existed for several years [3.1,3.2]. However, the difficulty with both these techniques lies in the need to provide protection for the main power thyristor during operation. Conventional valves use the local electronic circuits to detect the dangerous conditions listed in Chapter Two (ie forward overvoltage, excessive dv/dt and forward recovery failure) and trigger the thyristor into conduction to protect it. With direct optical triggering of the main thyristor, it is impractical to have two optical inputs, one for triggering from ground level and one for local protective triggering. Therefore any need for protective firing must be signalled to ground level so that an optical pulse can be transmitted back to the main thyristor. This process causes a delay in the protective action, which introduces uneconomic trade-offs in the valve design and renders the protection ineffective in some situations.

With the pilot thyristor technique, the electrical gate of the power thyristor is still accessible, so that the protective functions could be implemented locally and therefore rapidly. However, if these functions are provided conventionally using electronics, then the main advantage of direct optical triggering (ie, improved reliability through elimination of local electronics) has been compromised if not lost altogether.

Two solutions to this problem are acceptable for the outdoor valve. The first is to integrate all the protective functions onto the main thyristor, so that the device protects itself in all credible situations. This is the ideal solution, but requires important technical

advances in several areas to be practical. The other solution, using a pilot light-triggered device, is to provide protective functions locally using highly simplified circuits which avoid delays in protective action. Whilst overvoltage and excessive dv/dt can be sensed more or less easily using simple passive circuits, protecting against forward recovery failure is much more difficult without active measuring and logic circuitry. Indeed, there are no known practical methods for achieving simple forward recovery protection. This thesis therefore concentrates on the first approach (ie integrating all protective functions onto the main thyristor) but eventually concludes that, for the time being at least, a hybrid combination offers the best technical and economic performance.

b) The Valve Light Triggering System.

Whichever of the solutions presented in the previous section is adopted, an optical system is required to trigger the thyristors (pilot or main) into conduction. Light-triggered thyristors can be made to turn on in response to triggering energies which are orders of magnitude less than those required by their electrically-triggered counterparts. However, these optical energies are still at the limits of present optical technology, and therefore the design of the triggering system needs careful consideration.

The specification for the light-triggering system is set out in Chapter Seven, and centres on the following parameters:

- Optical performance (in terms of wavelength, pulse energy, power and power density);
- Reliability;
- Lifetime and maintenance;
- Cost, which should be compatible with the objective that the overall cost of an advanced valve should be no more than that of a comparable valve using conventional technology.

Although several light-triggering systems have been developed by other manufacturers (see Chapter Seven for details), none of these were considered to be suitable for the advanced valve under consideration, and so a novel system was developed during the project. This is described in Chapter Seven.

c) The Valve Insulation System.

If a reliable optically-triggered valve can be constructed at acceptable cost, it becomes desirable to encapsulate the valve. Consideration must therefore be given to the dielectric that would be used inside the valve tank. Various attributes of the dielectric are important, including:

- dielectric strength;
- permittivity;
- toxicity;
- environmental acceptability;
- flammability;
- availability;
- compatibility with other valve materials;
- cooling capacity;
- cost.

A major consideration which affects the choice of insulant is whether a separate cooling circuit is to be provided, or whether the insulant is also to act as a coolant. If a separate cooling medium is used, then the effect of leaks from the coolant into the insulant (and vice versa) may be of substantial importance.

Many insulants have been evaluated in laboratories, but only a relatively small number have achieved widespread use in electrical equipment (see Appendix One). The acceptability of the insulant in the market must therefore be considered, since if it is already employed in other equipment operated by the customer, it stands a much better chance of gaining acceptance than if it is an unusual insulant with little or no prior operating experience. Some undesirable attributes of an insulant (such as flammability, long term degradation) can be coped with, but at increased cost to the customer.

The subject of maintenance is also an important consideration in the choice of insulant. Although maintenance of the valve should be infrequent, it cannot be ignored altogether, and so the choice of insulant must be compatible with the need for access to the equipment. Possible contamination of the insulation system during maintenance must also be borne in mind.

Finally, the behaviour of many insulation systems is significantly different for direct voltage than for alternating voltage. This will be an important influence in the design of the insulation system for the valve, where both alternating and direct voltages are present.

d) The Valve Cooling System.

Adequate cooling of the thyristors and other components in the valve is a fundamental part of thyristor valve design. As explained in Chapter One, the most common technologies in service today are air cooling and water cooling. These are examples of the two different types of philosophy that can be adopted, ie, where the insulant is used as the coolant, and where the coolant is separate to the insulant. If the insulant is to act as the coolant, then in addition to the attributes outlined in the previous section, factors such as the specific heat capacity, viscosity and film heat transfer capability are also important. If a separate coolant is used, some of the considerations outlined in the previous section still apply, since the coolant will be flowing between points at different voltages, and therefore it must have reasonable properties as an insulant.

e) The Valve Mechanical Arrangement.

This area covers features such as the mechanical design of the valve and its tank, the number of valves per tank, and the critical areas of maintenance philosophy and the provision of redundancy. All these features affect both the initial cost of the scheme and the cost of keeping it operational over the full design life (usually at least 25 years). All manufacturers who have built outdoor valves have moved to using indoor air-insulated equipment (see Appendix Two), and the most important reason for this is the improved ease of access for maintenance. If an advanced encapsulated valve is to gain acceptance, not only must the general mechanical design of the valves be efficient, cost-effective and "elegant", but also the need for maintenance must be provided for in such a way that, should it prove necessary, it can be

carried out relatively easily and quickly. The provision of redundancy, for example in the form of spare valves, may have a substantial impact on the capital cost of the scheme, and therefore needs careful consideration.

3.3 INTERACTIONS BETWEEN THE TECHNOLOGY AREAS.

The interactions between the areas of the project discussed above are illustrated in Figure 3.3.1. There are two main "pools" of strong interaction: that between the thyristor and the light triggering system, and between the insulation, cooling, and mechanical design of the valve. It should also be recognised that there are weaker links between these two pools. Within each "pool", it is impossible to optimise one area of the technology in isolation from the others; therefore it is necessary to understand the issues and trade-offs involved in each area of the project in order to reach useful conclusions in any of them.

The wide range of interacting technologies shown in Figure 3.3.1 made it a very appropriate project to be carried out under the auspices of the Interdisciplinary Higher Degree (IHD) Scheme of Aston University; the purpose of this Scheme being to support research that extends beyond the boundaries of any single conventional field of investigation. As a result, projects of this type may not enter into the level of detail normally associated with a single-discipline PhD; however, it is recognised that achieving effective multi-disciplinary research can be an equally demanding and an equally valuable contribution to knowledge. The discussive, as opposed to focussed, nature of this thesis reflects its multi-disciplinary coverage, since all aspects of the advanced valve project are covered to a greater or lesser extent.

In the course of the work, it became evident that the project was a clear example of the need for an "IHD" approach. In particular, the effectiveness of the initial work on the light-triggered self-protecting thyristor (Stage 1 of the development) was impaired due to the

"isolation" of expertise in the areas of thyristor design and valve design. The author's background was in valve design, however, and so by becoming involved in the thyristor development, the overall efficiency of the research work was greatly improved. This improvement was a major contribution to the technical success of the development. Another important asect of the "IHD" dimension of the project was the need to consider the managerial function, an issue that is reviewed in more detail in Chapter Eight.

To highlight the need for a multi-disciplinary approach to the project, the following discussion describes some of the important interactions in the technology pools illustrated schematically in Figure 3.3.1.

a) The Thyristor and Light Triggering System.

The joint optimisation of the thyristor and light triggering system centres around the characteristics of the photo-sensitive gate of the thyristor. Altering the design of the thyristor gate affects both the thyristor capabilities and the design of the light triggering system, which in turn affects the overall cost and performance of the valve. The nature of these effects is explained briefly below.

The most critical parameter is the optical sensitivity of the thyristor, which has the biggest single impact on the cost of the light triggering system. Sensitivity can be improved by using a smaller optical gate, which increases the radiance required of the light sources, or by reducing the dv/dt capability of the gate. This then either requires costly dv/dt compensation on the device, or else it must be allowed for by changing the rest of the valve design, increasing costs elsewhere.

The vulnerability of the thyristor to damage by weak optical gating has an important effect on the optical overdrive required from the light triggering system, as well as the reliability of its control and monitoring circuits.

The combined characteristics of the thyristor and light triggering system must also be compatible with the design and performance requirements of the complete valve. The need to provide "dc gating" of the thyristor is one example of this; if dc gating could not be achieved, then the provision of a "pulse-on-demand" system would have a major impact on the rest of the valve design due to the need to provide voltage measuring on each level.

b) The Insulation, Cooling and Mechanical Design of the Valve.

Some of the interactions between these technology areas have already been mentioned, including the central issue of compatibility between the insulant and the coolant. A number of combinations of insulant and coolant can be envisaged where a leak from one into the other would have serious if not catastrophic effects, and therefore it is impossible to select either of these without consideration of the other. Also mentioned above was the need to allow for maintenance of the valve. If it is decided to have a redundant valve tank on site so that faulty units can be returned to the factory for repair, then this will allow the use of an insulation system requiring very high purity. However, if repairs are to be carried out on site, then a more contamination-tolerant system would be preferred. It would also be desirable to have an insulant that could be drained off quickly leaving the equipment ready to be worked on, rather than an insulant which

could leave either toxic substances or a viscous coating on the equipment. There is therefore an important interdependence between the choice of insulant and the maintenance philosophy for the valve.

There is also an interaction between the mechanical design, the maintenance philosophy and the cooling technique selected for the valve. If a liquid coolant is used in an enclosed circuit, then it is desirable that a thyristor can be changed without opening cooling circuit connections. This places a constraint on the mechanical design which would not exist if a gaseous or liquid spray cooling system was used.

3.4 CONCLUSIONS.

In order to illustrate the interdependence of the five technology areas involved in the outdoor valve, the interactions between them have been discussed briefly and some examples given. From the detailed study given in the rest of the thesis, the extent of some of these interactions will become clearer.

Section 1.8 in the first chapter described the author's involvement in coordinating work in the different areas of the project. Whilst the majority of the thesis is devoted to the technical aspects of the work carried out, it is clear that in this case technical research could have been of very limited value unless it was properly coordinated towards the overall project goals. A discussion of the successfulness of the project management is given in Chapter Eight.

CHAPTER FOUR: DEVELOPMENT OF THE THYRISTOR, STAGE 1.

4.1 INTRODUCTION.

All modern HVDC transmission schemes use silicon thyristors, which is because of their reliability, high power handling capacity and the low power required to control them. Thyristors can only conduct current in one direction; conduction will start if there is a forward voltage across the device and it is given a gate pulse. Once conducting, however, control of the device is lost, and the thyristor will remain in the "on-state" until the action of the external circuit causes the current to fall to zero. For a conventional thyristor, this gate pulse is an electrical signal applied to a terminal of the device, whereas for a light-triggered thyristor, it is an optical signal applied to a particular area of the silicon. A detailed discussion of thyristor operation can be found in reference [4.1].

The major draw-back with conventional thyristors is the complex electronic circuitry necessary to control and protect them (described in Chapter Two). A primary objective of the advanced valve project is to eliminate the need for this electronics by developing light-triggered thyristors with integral protection against destructive failure mechanisms. Chapters Four to Six therefore describe the work carried out to achieve this objective, as well as reviewing some of the theory involved.

The work on the thyristor took place in three stages. In the first part, development was carried out internally by MEDL, and monthly liaison meetings were held with TDPL to report on progress. During this period, which lasted from October 1984 until November 1985, the author

concentrated mainly on other areas of the advanced valve project. At the end of that time, it was felt that closer contact between MEDL and TDPL would be beneficial, and so the author spent four months working at MEDL, from December 1985 to March 1986. This started the second stage of the development, and after the initial period of full-time work, the author continued to visit MEDL one day each week to maintain contact. By September 1986, evaluation of the first prototype design was complete, which concluded the second stage of the development. The third stage was the design, manufacture and testing of a new set of test structures, which lasted from October 1986 to March 1987.

Chapters Four, Five and Six describe the work carried out in the three stages of the project. The sections in this chapter cover the following topics:

4.2 The theory of light triggering.

4.3 Thyristor failure mechanisms.

4.4 Literature review.

4.5 Summary of work carried out in Stage 1.

4.6 The first prototype design.

4.7 Evaluation carried out in Stage 1.

4.2 THE THEORY OF LIGHT TRIGGERING.

a) Introduction

An electrically-triggered thyristor is turned on by passing a current between its gate and cathode. This causes charge carriers to be injected into the base of the *npn* transistor in the structure, and when the charge reaches a sufficient level, turn-on occurs. The same effect can be achieved by illuminating the silicon with light of a suitable wavelength. This effect, known as optical or light triggering, is explained below. The implications of light triggering for dv/dt capability are then discussed, and finally a hybrid technique for light triggering (known as the 'slave thyristor' approach) is explained.

b) Optical Turn-on

Because silicon is photo-sensitive, photons of the correct wavelength incident upon it will cause the silicon atoms to be ionized, creating hole-electron pairs. If there is no electric field in the device, the hole-electron pairs will recombine thermally, and there will be no measurable effect on the silicon. However, if there is a field in the silicon, the hole-electron pairs will be separated, and current will flow through the device. This effect can be used to provide gate current for a thyristor.

In Figure 4.2.1, a cross-section through a forward-biased high-power thyristor is shown, with light being shone on the gate region. Because it is forward biased, a depletion layer exists around the middle junction of the device. When an atom is ionized inside this depletion layer, the electric field separates the hole and electron, which flow to the cathode and anode respectively. The hole current forward biases

the auxiliary thyristor cathode, which injects electrons into the p-base and from there into the depletion layer; this initiates regenerative turn-on in the auxiliary thyristor (Figure 4.2.2). The auxiliary current then forward biases the main emitter, causing this to turn on (Figure 4.2.3). The device has then been optically triggered into conduction. Although the optically-induced electron current to the anode does contribute to turn-on, its effect is small due to the poor efficiency of the pnp transistor.

A number of improvements on this basic design are possible. Firstly, when photons are absorbed outside the depletion layer, it is likely that the hole-electron pairs created will recombine thermally, since there is virtually no field to separate them. Consequently, they will not contribute to turn-on, and so it is desirable to inject photons as close as possible to the depletion layer. This can be achieved by etching down into the silicon, as shown in Figure 4.2.4. Since the depletion layer widens with applied voltage, the efficiency with which light is utilised will improve with increasing voltage; see Figure 4.2.5. This is fortuitous, since conventionally the strength of gate pulse required also increases with voltage. (Note that the sensitivity of the thyristor to gate current also improves with increasing voltage, since the transistor gains are voltage-sensitive).

A second improvement on the design of Figure 4.2.1 is to introduce another amplifying stage into the gate. The structure shown in Figure 4.2.1 is typical for electrically-triggered thyristors, where it is straightforward to supply a gate current of 1A or more to the centre of the device. However, the level of gate current that can be generated with practical light sources is much less than this, typically a few

tens of milliamps, so that a more sensitive structure is required. This can be achieved by introducing another thyristor into the gate structure, as shown in Figure 4.2.6. This extra "pilot" thyristor achieves increased sensitivity by virtue of two features:

- The emitter is diffused deeper into the thyristor, so that the sheet resistance underneath it is greater. This is shown in Figure 4.2.7. Thus for a given current density in the p-base, the voltage generated along the emitter is higher.
- The inside diameter of the emitter is very small, which also maximises the voltage generated in the p-base for a given level of optically-induced current. The relationship between the voltage generated along the p-base and the internal radius of the pilot thyristor is illustrated in Figure 4.2.8.

By combining the two techniques described above (ie etching into the silicon and adding a pilot thyristor), it is possible to make thyristors which will turn on in response to only a few nanojoules of optical energy. This is very attractive from the point of view of the control system driving the thyristor. However, if the gate of the thyristor is too sensitive, it becomes vulnerable to unintentional turn-on during rapidly rising forward voltages. This is discussed in the following section.

c) Dv/dt Capability

When a thyristor is not conducting (ie in the "off-state"), the depletion layer in the device causes it to behave as a parallel-plate capacitor, so that a changing voltage across the device will cause current to flow through it. Some of this current flows under the optical gate, and if it is of sufficient magnitude and the correct

polarity, the gate will turn on; see Figure 4.2.9. The design constraint for an HVDC thyristor is to withstand a specified level of dv/dt (typically 2-3 kV/us) without turning on, and this places a limit on the maximum optical sensitivity that can be achieved in the gate.

By comparing the horizontal voltage profiles generated along the p-base by photo-current (injected at a point) and dv/dt current (distributed), it can be seen that the ratio between these voltages is greatest when the radius of the well is small; see Figure 4.2.10. This ratio is a measure of the efficiency of the design in terms of trading off optical sensitivity against dv/dt capability.

The absolute level of voltage generated (from both sources of current) increases as the well is made deeper. To achieve maximum optical sensitivity, it is therefore necessary to make the well as deep as possible, and then increase the outer radius of the pilot thyristor emitter until the dv/dt constraint is encountered.

The practical limits to the depth of the optical well are the need to avoid punch-through from the depletion layer at maximum forward voltage, and the need to achieve repeatability between production batches of thyristors. As the optical well approaches the depletion layer, small changes in depth produce large variations in sheet resistivity (Figure 4.2.7), which will affect optical sensitivity and dv/dt capability. A compromise is therefore necessary between optimized performance and production yield.

If an acceptable compromise between optical sensitivity and dv/dt capability cannot be achieved, sensitivity can be improved by employing a technique known as dv/dt compensation. This alters the geometry of

the device such that the lateral voltage generated by dv/dt current in one part of the p-base offsets (or 'compensates') the voltage generated in a different area. In Figure 4.2.11, it can be seen that the positive potential generated at point "A" is applied to the n-emitter of the pilot thyristor to bias it off. The topology of this technique involves connections crossing over each other; this means that either a layer of insulator has to be included in the device processing, or that wires must be bonded to the device after processing. Since both of these options increase cost and complexity, compensation techniques are avoided where possible.

From Figures 4.2.8 and 4.2.10 it can be seen that for a given optical sensitivity, better dv/dt capability can be achieved if the light is delivered in a smaller radius (ie a smaller Ri can be used). This favours laser-based delivery systems, which can give very high power densities, rather than large-area sources such as LED's and flash-lamps. To deliver the same power, these latter sources require a larger delivery area, so that the dv/dt capability of the device is reduced (or conversely, a lower optical sensitivity is achieved for the same dv/dt rating).This is a major point of interaction between the design of the thyristor and the light triggering system, and it is discussed further in Chapter Seven.

d) The Slave Thyristor Technique.

So far, light triggering has only been considered as direct optical activation of the main power thyristor. However, as shown in Figure 4.2.12, an intermediate technique has been used whereby a small light-triggered thyristor is used to provide an electrical gate pulse to a conventional power thyristor. This has several advantages over direct

activation of the power thyristor:

- The light-triggered thyristor chip is small, so that the cost effects of poor yield are minimized;
- The slave device does not carry load current, and so remains at ambient temperature. This makes it easier to meet dv/dt withstand requirements;
- A resistor can be inserted between the light-triggered stage and the main device in order to limit inrush duty on the optical well;
- The electrical gate of the main device is still accessible, so that conventional local protection circuits can be used;
- One design of slave thyristor can be used with a wide range of conventional power thyristors;
- The mechanical design of the thyristor package is simpler. The absence of high-current pressure contacts means that the optical fibre can enter the package normal to the surface of the silicon, rather than having to enter from the side and then turn through 90° (see Chapter Seven).

Compared to direct optical activation, there will be the extra cost of having two discrete devices instead of one. However, both approaches still suffer the same drawback discussed earlier, ie that light triggering without self-protection offers little or no economic benefit compared to conventional techniques, since local electronic protection still has to be provided. Therefore, although the slave thyristor approach is technically easier to implement than direct optical triggering, neither has found widespread adoption for HVDC.

4.3 THYRISTOR FAILURE MECHANISMS.

a) Introduction.

Thyristors are vulnerable to damage from a number of causes, and since they are semiconductors, once this has occurred it is permanent and irreversible. It is therefore desirable to prevent damage occurring to thyristors in an HVDC converter for the following reasons:

- The cost of unscheduled repair time is very high (of the order of £100,000 per hour);
- The thyristors themselves are expensive (of the order of £1000 for a 100mm device);
- The pressure contacts to the thyristor packages must maintain low electrical and thermal resistance over many years, and it is therefore desirable to disturb them as little as possible.

When a thyristor fails, it normally becomes short-circuit, and this reduces the voltage withstand capability of the series string of devices comprising the HVDC valve. It is therefore normal practice to include a small number of extra 'redundant' thyristors in a valve (typically 3%) so that operation can continue between scheduled maintenance times in the event of a few devices failing. However, the inclusion of redundant thyristors incurs significant extra capital and running costs, and so the level of redundancy is kept to a minimum. There are therefore strong economic reasons for preventing thyristor failures, which go far beyond the replacement cost of the devices themselves.

Under some circumstances, it is possible for a thyristor to lose its voltage blocking capability in only one direction. This happens when

damage is limited to the surface of the device, so that only one blocking junction is affected. However, this type of failure is rare and can normally be prevented by good design practice, so that it is not considered again.

Normal thyristor failure (ie to a short-circuit condition) occurs when the local current density in some part of the device becomes sufficient to melt the silicon. This is usually brought about by 'weak' triggering in some area of the device; ie, part of the structure starts to turn on without the benefit of the normal strong gate pulse. Because the triggering is weak, only a small area of the device turns on initially, and the inrush current that flows into this small area from the circuit stray capacitance melts the silicon (see Chapter Two). Weak triggering can be caused by the following conditions:

- A fault in the gate drive circuit;
- Excessive forward voltage (overvoltage failure);
- Excessive rate of rise of forward voltage (dv/dt failure);
- Forward voltage transients applied during the recovery interval (forward recovery failure).

The reliability of the gate drive circuit is a separate problem to that of the thyristor and, for the light-triggered valve, it is considered in Chapter Seven. Excessive reverse voltage is not listed above because although it is possible for thyristors to be destroyed by excessive reverse voltage, they can be designed to operate safely in the reverse avalanche region of their characteristic. Limiting the energy dissipated in the device to safe levels during such operation is then a function of the external circuit design. The remaining causes of damage (ie forward overvoltage, dv/dt and forward recovery failure) are

discussed in the following sections. The events which can give rise to these failure mechanisms, as well as the conventional means for protecting against them, have already been considered in Chapter Two.

b) Forward Overvoltage Failure.

A thyristor will spontaneously trigger itself into conduction if the forward voltage across it becomes too great when it is in the off-state. This is because the current flow in some part of the device reaches the level necessary to cause turn-on, and this current can come from three sources:

- Leakage;
- Avalanche;
- Punch-through.

Leakage current arises from the thermal generation of carriers within the depletion layer, and is therefore a strong function of temperature. Avalanche occurs when a carrier inside the depletion layer gains sufficient energy between lattice collisions to ionize atoms in the lattice, leading to carrier multiplication. Since the mean free path between collisions falls with temperature (due to lattice vibration), the avalanche voltage increases slowly with temperature. Punch-through occurs when the depletion layer around the centre junction reaches one of the other junctions, causing a massive injection of charge. The width of the depletion layer is virtually independent of temperature, so that the punch-through voltage is effectively constant. Figure 4.3.1 shows these three limits to the forward voltage capability of the thyristor.

Forward overvoltage turn-on is inherently the most destructive of the

three failure mechanisms being discussed since, by definition, it always occurs from the maximum voltage rating of the device. Turn-on under dv/dt and forward recovery failure can occur from any voltage between zero and the voltage break-over level; however, at low voltages it may not be damaging.

In considering the forward voltage capability of the device, it is assumed that the edge of the device has been bevelled and passivated well enough to prevent this being a limiting factor. Otherwise, leakage currents flowing in this region will either trigger the device on, or cause sufficient heating to break down the passivation or induce thermal runaway in the silicon.

c) Dv/dt Failure.

As has been explained in Section 4.2, a changing voltage across the device will cause displacement current to flow (Figure 4.2.9), and this may be sufficient to turn the device on in some area. For a light-triggered device, the optical gate is made very sensitive, so that dv/dt turn-on would be expected to occur here first. For an electrically-triggered device with a less sensitive gate, turn-on could be anywhere on the device.

Dv/dt current flow is a strong function of temperature because the level of free charge that has to be swept out to form the depletion layer increases rapidly with temperature.

When a thyristor structure turns on, the delay time associated with that turn-on depends on the level of "gate drive" being applied relative to the threshold necessary to just cause turn-on. If a strong gate drive is used (ie the dv/dt is well above that necessary to cause

turn-on), then turn-on will be rapid, and the device voltage will collapse before the dv/dt ramp has had time to reach a high level. In this case, the turn-on may not be damaging, since the energy stored in the circuit capacitance (proportional to the square of the voltage) could be safely dissipated. However, if the dv/dt is only slightly above the capability of the device, then the turn-on is marginal and therefore slow. The ramp may reach damagingly high levels before the voltage collapses; see Figure 4.3.2.

d) Forward Recovery Failure.

Forward recovery failure is similar to dv/dt failure, except that in addition to the capacitive displacement current, there is also current flow due to the sweeping out of unrecombined carriers left over from a previous period of conduction. For low rates of dv/dt (eg ~10V/us), the point at which forward recovery failure occurs is defined as the turn-off time of the device ("tq").

When the thyristor is in conduction, the bases of the two transistors in the structure are saturated with charge. At recovery, there is substantial reverse current flow until the reverse blocking junction is formed. However, this current flow is not enough to remove all of the excess charge in the transistor bases, and once the depletion layer is formed, only leakage current can flow through the device. A significant quantity of the excess charge is therefore trapped and has to recombine thermally.

If a forward dv/dt ramp is applied to the device before the excess charge has recombined, a pulse of current flows as the ramp goes through zero volts and "sweeps out" the charge; see Figure 4.3.3. This

pulse continues until the forward blocking depletion layer is formed, when the voltage across the device starts to go positive. The current flow adds to the dv/dt current already tending to turn the device on, so that turn-on will occur at lower levels of dv/dt than for an "off-state" device. As with dv/dt turn-on, the worst situation is when just enough current flows to turn the device on, so that the turn-on delay time is long and the ramp can reach a high voltage before turn-on occurs.

The amount of excess charge remaining in the transistor bases is a function of time after current zero, and so the vulnerability of the device to forward recovery failure is also time-dependent. If a dv/dt ramp is applied soon after current zero, there is still sufficient charge in the bases for conduction to start immediately, so that no forward voltage appears across the device; such turn-on is completely safe. Once the charge decays below this level, some injection from the n-emitter is necessary to initiate turn-on, and there will therefore be a slight time delay before turn-on occurs. As the time after current zero increases, the turn-on time delay for a given dv/dt rate increases. Eventually the device will be able to withstand the dv/dt ramp, provided that it is less than the off-state capability of the device. Figure 4.3.4 shows the dv/dt withstand capability of the device as a function of time after current zero.

It can be shown by experiment that thyristors can safely withstand single-shot ungated turn-on (eg by the above mechanisms) up to a certain voltage level, which is circuit-dependent. In GEC's HVDC valve design, ungated turn-on up to around 1500V can be withstood at 90°C without damage. If an arbitrary upper limit to the delay time for marginal turn-on is assumed (perhaps 20us), then a dv/dt rate can be

calculated below which turn-on will never be dangerous (1500/20 = 75V/us in this case).

e) Self-Protection Against Failure.

Conventional valves using electrically-triggered thyristors provide electronic sensing and protection for each of the above types of failure (Chapter Two). Voltage, dv/dt and time after current zero are measured, and if circuit conditions could give rise to failure by any of the above mechanisms, the device is gated into conduction; this strong gate pulse pre-empts any weak triggering in the device. To eliminate the local electronics, then either the thyristor must carry out this protective triggering internally (ie be self-protecting), or else a simple passive circuit with a small parts count must provide the required protection. Either solution will fulfil the objective of the advanced valve.

An example of this simple external protection is the Break-Over Diode (BOD) presently used for forward overvoltage protection by several manufacturers. This is a simple assembly of avalanche thyristors which trigger themselves into conduction when the voltage across the thyristor reaches its protection level; see Figure 4.3.5. The current through the string is used to provide a strong gate pulse to the thyristor, turning it on and collapsing the volts before a dangerous level is reached.

Since the use of a BOD is consistent with the aims of the advanced valve (ie, it is simple and requires no external power supply), the choice between it and the use of internal VBO protection will be an economic one. A similar situation exists for dv/dt protection; if a

slave light-triggered thyristor is used which will also turn on safely when subjected to dv/dt's above the protective level, then this would provide adequate dv/dt protection. However, forward recovery protection cannot be provided by simple external components. During the course of the project, much thought was given as to how this could be achieved, but no satisfactory solution was found. The main obstacles to a solution were:

- The variation in the time at which recovery occurs for devices in a series string;
- The difficulty of detecting current zero in any one thyristor with only simple components;
- The difficulty of detecting whether ungated turn-on has occurred, so that the protection can be reset at the next current zero.

The only other solution (apart from true self-protection) would be to include so much inductance in the valve that no transient higher than the 'safe' level (eg 75V/us) could occur. However, this would require a substantial increase in the size, weight and cost of the valve reactors, which would render the valve completely uneconomic compared to conventional designs. It was therefore concluded that, whilst it might be economically preferable to use external components to provide VBO and dv/dt protection, forward recovery protection would have to be provided internal to the thyristor.

4.4 LITERATURE REVIEW.

The potential for light triggering of thyristors was recognised around the time that they were invented in the 1950's [4.2]. However, light triggering is only really economically attractive with high-power devices, and these did not become available for many years. In the early 1970's, light triggering was used to obtain ultra-fast switching of thyristors in high-energy discharge experiments [4.3]; however, this involved illuminating a large proportion of the device area with an extremely powerful laser pulse, and so was of little practical relevance to HVDC.

The first attempts at applying light triggering to HVDC used the slave thyristor approach [4.4]. However, for the reasons discussed in Section 4.2, there is little or no economic benefit in using this approach. This technique has only been applied to one HVDC scheme [4.5] (One of the two projects mentioned in this reference was subsequently cancelled).

Several manufacturers have developed high-power thyristors with direct optical triggering [4.6, 4.7, 4.8, 4.9, 4.10, 4.11], and three prototype valves have been built [4.12, 4.13, 4.14]. However, to overcome the lack of self-protecting features, these prototype valves have used at least 10% more thyristors in series than an equivalent conventional valve, thus making them economically unattractive. Valves using light-triggered thyristors without self-protection are therefore not competitive with conventional valves [4.11].

In order to overcome these limitations, attempts have been made to incorporate self-protecting features. The most popular feature has been

forward voltage break-over (VBO) protection [4.11, 4.15, 4.16, 4.17, 4.18]. Attempts to build in VBO protection have met with varying degrees of success. At lower voltages (around 1500V), reasonably good results have been obtained with the "laser-zapping" technique [4.17]. In one case, the development exercise has been reported as a failure [4.6], but others have reported successful results at 6kV [4.19]. However, none of the developments have reached a commercial status.

There has been little work on self-protecting features other than VBO protection. In one case [4.6], dv/dt protection was achieved apparently unintentionally. Dv/dt triggering was initiated by the optical gate, and the thyristor was able to withstand repetitive dv/dt triggering from its maximum forward voltage capability. Apart from this, however, MEDL is the only manufacturer to have reported work on either dv/dt or forward recovery protection [4.18].

4.5 SUMMARY OF WORK CARRIED OUT IN STAGE 1.

a) Introduction

MEDL have been working on light-triggered thyristors (LTT's) for many years. During the 1970's a small high-voltage LTT was developed for use as a slave device with conventional high-power thyristors (see Figure 4.2.12). However, due to the need for self-protection as well as light triggering for a light-triggered valve to be economically viable, this device was not applied commercially to HVDC.

In the early 1980's, work continued on developing a 56mm LTT incorporating VBO protection. The protection system used a large number of avalanche sites in the gate region surrounded by current limiting resistors. Avalanche in these regions was intended to gate the device into conduction. However, infra-red photographs taken during turn-on showed that only one or two sites were operating, so that most were redundant. This was probably due to small variations in the depth of the etched avalanche sites, and this experience led to the revised philosophy of having a single VBO site on the next design of LTT.

b) Work carried out in Stage 1

At the start of Stage 1, first samples of the 56mm LTT incorporating VBO protection were being produced. At that time, no work had been carried out on either dv/dt or forward recovery protection. After evaluation of the 56mm LTT's, MEDL produced a design for a light-triggered, fully self-protecting thyristor. Since this design formed the basis of much of the work carried out at MEDL during the project, it is explained in detail in section 4.6.

The design was first produced on 75mm silicon in March 1985, and a series of tests were carried out on these devices. The results of these tests are reported in section 4.7, but briefly were as follows. The optical sensitivity of the device was found to meet or exceed design targets; however, the dv/dt ratings of the samples was relatively poor, suggesting that either the optical well was too deep, or that the compensation mechanism was inadequate. The VBO function was also tested, but those devices that exhibited VBO operation all had a VBO level that was too low to be useful. The most encouraging feature of the initial tests was that the devices appeared to be immune to di/dt failure at turn-on. This characteristic, known as "controlled turn-on", is described in [4.15]. Turn-on transients of several hundred amps per microsecond could be withstood with only marginal triggering, a capability well beyond that of a conventional thyristor. This capability also has important implications for the design of the valve light triggering system; see Chapter Seven.

All of the batch of 75mm devices produced were used and destroyed in the above tests. In order to permit investigation of the forward recovery protection mechanism, two batches of 30mm devices were put into process. These samples, described at the end of Section 4.6, became available early in Stage 2 of the project.

4.6 THE FIRST PROTOTYPE DESIGN.

a) Overview.

The outline of the prototype is shown in Figure 4.6.1. Turn-on takes place through a two-stage amplifying gate; the first stage of this is the optical well in the centre of the device, which also acts as the site for VBO turn-on. The second stage is the auxiliary amplifying gate, which can be seen as three pairs of rectangular pads between the central optical well and the three "satellite" wells; these satellite wells are intended to perform the forward recovery protection. The auxiliary amplifying gates feed into the prominent "fingers" which extend into the main cathode to give a long turn-on edge.

Figure 4.6.2 gives an expanded view of part of the gate structure, and an equivalent circuit of the device is given in Figure 4.6.3. Its various features are described in the following sections.

b) Optical Triggering.

A cross-section through part of the optical well is shown in Figure 4.6.4, and will now be described region by region. Starting from the centre, region A is a secondary well etched in the centre of the main well. This secondary well serves three purposes:

- it minimizes the thickness of silicon that has to be penetrated by incident photons before they reach the depletion layer.
- it increases the initial turn-on area; since turn-on cannot occur where there is no n-emitter, turn-on should occur around the periphery of this well rather than at a point in the centre, thus improving the di/dt capability of the well.

- it triggers VBO action (See section *c) VBO Turn-on*).

Region B is the emitter of the pilot thyristor where turn-on initially takes place. This deep region of emitter is necessary because, to achieve good optical sensitivity, a high sheet resistivity is required under the emitter, typically 800 - 1000 ohms/square. The resistivity under the main cathode emitter is only 300 ohms/square, and so a deeper cathode diffusion is used for the pilot thyristor.

Regions A and B are both coated in silicon nitride (not shown), which acts as an anti-reflectant coating to improve the absorption of infra-red light.

Region C is diffused at the same time as the main n-emitter, and has a high emitter efficiency due to its thickness (60um compared to 5um for the deep emitter) and high doping concentration at the surface. After initial turn-on at the inner edge of region B, the conducting area spreads to region C. A ring of aluminium metallization is placed on top of region C to collect the current that flows in the optical well.

Current from the optical well flows down the three radial spokes "D", which can also be seen in Figures 4.6.1 and 4.6.2. Because of the silicon nitride insulating layer (Figure 4.6.4), current can only flow into the p-base in the region between the two auxiliary amplifying gates, as shown in Figure 4.6.5. Current flowing under these gates forward biases them and turns them on; this feeds current to the auxiliary cathode metallization and out to the main emitter, thus causing turn-on in the main body of the device.

In order to limit the current that flows into the gate structure during turn-on, two sets of control resistors are built into the gate. The

first of these, Rc1, is located between the radial spokes and the auxiliary amplifiers; the second, Rc2, is located between the auxiliary metallization and the main cathode. Both of these are shown in Figure 4.6.5. The resistors are created by etching away the highly-doped and therefore high-conductivity silicon at the surface of the device, so that the remaining p-base has a high sheet resistivity. (Figure 4.2.7 shows the variation of sheet resistivity with depth). The full equivalent circuit for optical turn-on is given in Figure 4.6.3.

c) VBO Turn-on.

VBO turn-on is intended to occur by avalanche at the centre of the optical well. The small secondary well is etched to a depth such that, when the forward voltage across the device reaches the required VBO level, the depletion layer in the p-base just reaches the bottom of the well. The discontinuity caused by the well distorts the electric field, thereby triggering avalanche. Alternatively, if the bottom of the well is damaged (eg by laser zapping), it cannot be depleted and will therefore act as a source of electrons; when the depletion layer comes within one carrier diffusion length, electrons will be injected into the depletion layer and turn-on will commence. This latter mechanism is more analogous to punch-through than avalanche.

Turn-on starts in Region B of Figure 4.6.4, and propagates as for optical triggering.

d) Forward Recovery Protection.

Forward recovery protection is activated by the three satellite wells shown in Figure 4.6.1. These wells are almost identical in construction to the optical well seen in Figure 4.6.4, the only difference being

that the metallization on region C is extended down onto region B. The proposed mechanism for providing forward recovery protection is described below.

After processing, the devices are selectively irradiated, which is achieved by placing thick metal shields over the three satellite wells and bombarding the device with electron irradiation; see Figure 4.6.6. The resulting radiation damage lowers the carrier lifetime in all areas of the device except those under the satellite wells. During normal conduction, charge diffuses under these wells from the main cathode, as shown in Figure 4.6.7. Once conduction ceases, this charge recombines more slowly than that in the main body of the device due to the longer carrier lifetime. The satellite wells are therefore the last areas of the device to recover the ability to block forward voltage. If a forward ramp is applied to the device whilst the main cathode is going through its period of vulnerability to forward recovery damage, the satellite wells will trigger the device into conduction. The ramp will sweep out the residual charge under the well, causing turn-on in a similar manner to optical triggering. The main cathode therefore does not see the potentially destructive ramp, whilst the satellite wells are protected from damage during turn-on by the control resistors. The control resistors are common to both the optical well and the satellite wells, since all the wells are connected by the radial spokes of metallization.

e) Dv/dt Compensation.

The metallization connecting the four wells is arranged such that when a forward dv/dt is applied to the device, the displacement current generated under the wells flows straight to the auxiliary metalliza-

tion. Figure 4.6.8 shows this current flow under the satellite well; the case for the optical well is similar. This current produces a potential drop under the n-emitter in the well, tending to turn it on. The compensation voltage to offset this is derived from Rc1, as shown in Figure 4.6.9. The displacement current flowing from the region between the auxiliary amplifiers creates a voltage between the spoke and the auxiliary metallization. This makes the spoke positive with respect to the auxiliary metallization, and since the spoke is connected to the n-emitter in the well, it tends to bias it off. The equivalent circuit is shown in Figure 4.6.10.

An unprotected thyristor in an HVDC valve can see dv/dt rates up to 10kV/us and beyond. It would incur serious penalties to make all areas of the thyristor able to withstand such dv/dt rates without turning on, and so it is necessary to ensure that safe turn-on occurs when the specified withstand level (usually 2-3kV/us) is exceeded. It is therefore important that the wells are not biased off for all dv/dt rates, but that when the specified level is exceeded, the wells turn on before any other area of the device; the control resistors will then protect them from damage. To achieve this requires accurate control both of the sheet resistance under the well, and the value of the control resistors. In this design, the control resistors have to be etched to within 3.5um of specification. If they are etched deeper than specified, the well will never turn on under dv/dt conditions; if they are shallower than specification, the well will turn on at a dv/dt rate below that required. This accuracy of control is difficult to achieve with conventional wet-etch techniques, and so this has important implications for the yield of the device in production.

f) 30mm Samples.

In order to permit more extensive evaluation of the forward recovery protection mechanism, a set of 30mm samples was manufactured which reproduced the relevant features of the mechanism. The design of these is shown in Figure 4.6.11, which illustrates how the gate structure contains two satellite wells linked by a single spoke, with one pair of auxiliary amplifier pads. There are no long gate fingers on the auxiliary metallization, since the devices were not intended for high di/dt testing. The design of the wells, amplifiers and control resistors is the same as for the main device described above.

4.7 EVALUATION CARRIED OUT IN STAGE 1.

This section summarizes the tests that were carried out on the first (75mm) samples of the prototype.

a) Optical Triggering.

Initially, optical sensitivity was measured with an infra-red LED, and tests were carried out at room temperature. A typical set of sensitivity curves is given in Figure 4.7.1, which display two features that were expected:

- Sensitivity improves with voltage (as explained in Section 4.2);
- As the gate pulse width is shortened its amplitude has to be increased, since turn-on of the pilot thyristor is "charge-controlled" for short gate pulses [4.20].

Later, an infra-red semiconductor laser was used for triggering the thyristors, with an optical pulse length of approximately 150ns. In Figure 4.7.2, the rapid improvement in sensitivity with voltage can be clearly seen. The reason why this improvement is so pronounced at low voltages is that the width of the depletion layer is proportional to the square root of the applied voltage [4.21], so that the rate of change of width with voltage is greatest at low voltages. This rapid widening of the depletion layer affects optical sensitivity via two separate mechanisms:

- The percentage of photons reaching the depletion layer increases;
- The efficiency of the npn transistor improves, since minority carriers injected from the n-emitter have less far to travel to reach the depletion layer.

The absolute levels of sensitivity measured were good, being substantially better than the target sensitivity of 40nJ at 60V.

b) Dv/dt Capability.

The dv/dt capability of the samples was not measured beyond 2kV/us or above room temperature. The target specification for these devices was 2-3kV/us at 110°C, and if this had been met, a considerably greater dv/dt capability would have been achieved at room temperature, since dv/dt current flow is a strongly temperature-dependent. Despite this, several of the devices failed to achieve 2kV/us at room temperature. A possible explanation for this is that the devices tested did not have the silicon nitride insulating layer shown in Figure 4.6.4, so that the dv/dt compensation would not be fully effective. However, even after taking this into account, the results were not promising.

The samples with the shallowest wells were found to have the highest dv/dt capability, which was expected since the p-base sheet resistance under the pilot thyristor emitter would be lowest.

c) Di/dt Capability.

Tests for di/dt capability were carried out using the simple LC discharge circuit shown in Figure 4.7.3. Destructive levels of di/dt were found to be 200 - 300A/us for threshold triggering, which is very good compared to conventional electrically-triggered thyristors. Device failures all occurred at the edge of the optical well, which was attributed to breakdown of the "parasitic" pn diode at the edge of the well, caused by the generation of excessive voltage across Rc1; see Figure 4.7.4.

d) VBO Turn-on.

Of the first eleven 75mm samples produced (with a variety of depths of secondary well), only three exhibited VBO turn-on; for the others, the forward withstand level of the device was reached before VBO triggering occurred. The results obtained are shown in Figure 4.7.5. The "best" device showed VBO turn-on at 3kV with a well depth of 48um; however, another device with a well 51um deep was tested up to 4.75kV without turn-on. This shows the unpredictability of the VBO mechanism, and thus the necessity for "interactive" etching of the VBO well (ie testing for VBO action, etching the device, testing again and so on).

The device which switched at 3kV survived single-shot turn-on in a circuit similar to that used for di/dt testing, with a stable VBO level. Six other devices were "zapped" with a ruby laser to try to induce VBO operation (refer to [4.17]); after each "zap", the device was tested for VBO switching. VBO levels of 2.5, 3.0 and 3.1kV were achieved on three devices, but for the remaining three devices, one laser "zap" took them from no VBO switching to short circuit. This again high-lighted the difficulty of controlling the VBO protection mechanism. There was also some doubt as to the stability of the VBO level when laser "zapping" had been used.

The design target for VBO protection was that switching should occur at 5kV $\pm$ 100V, and therefore substantial improvements were necessary in the performance of the VBO mechanism.

e) Conclusions.

By the end of Stage 1, promising results had therefore been obtained for optical sensitivity and turn-on capability, but dv/dt capability and VBO results were poor, and no work had been carried out on forward recovery protection. Chapter Five describes the work carried out in Stage 2 of the development.

CHAPTER FIVE: DEVELOPMENT OF THE THYRISTOR, STAGE 2.

5.1 INTRODUCTION.

This chapter describes the development work carried out on the thyristor between December 1985 and September 1986. In this period, extensive evaluation of the prototype thyristor design was undertaken, either by the author or under his supervision.

Testing of the prototype design was restricted by the availability of suitable representative devices. Thus, since 100mm samples were not available until March 1986, initial investigations were carried out on 30mm samples. The testing programme is described under the following headings:

- Optical sensitivity;
- Forward recovery protection;
- Charge extraction;
- Dv/dt capability;
- VBO protection;
- Turn-on capability.

The charge extraction experiments were related to the principle of forward recovery protection. The other headings are self-explanatory, and relate to the various features of a fully self-protecting light-triggered thyristor. A brief summary of the results obtained from this programme is given in Section 5.9.

5.2 EVALUATION OF OPTICAL SENSITIVITY.

a) 30mm Device Testing.

Optical sensitivity of the 30mm devices was measured as a function of applied voltage at room temperature in order to confirm the earlier results. For these tests and all subsequent tests involving optical triggering, one of the laboratory light firing systems was used (See Chapter Seven). Unless otherwise stated, the optical pulse width was set to 125ns.

A typical sensitivity characteristic for a 30mm device is given in Figure 5.2.1. Once again, this demonstrates the rapid improvement in sensitivity with voltage.

b) 100mm Device Testing.

The 30mm devices were known to have shallower optical wells (and therefore poorer sensitivity) than was proposed for the prototype thyristor. The 100mm samples produced later had the correct depth of optical well, and therefore better optical sensitivity. The tests on these 100mm devices were carried out at a single voltage (60V), since it was considered that changing the depth of the optical well would not alter the shape of the optical sensitivity characteristic, but only scale it. Figures obtained for the 100mm samples are given in Figure 5.2.2.

Effect of Metallization in the Well

Since the metallization in the satellite wells did not cover the central secondary well, it was possible to optically trigger these sites, and so the results for these wells are included in Figure 5.2.2.

The sensitivity of the optical wells was better than for the 30mm devices, as expected; however, it was observed that the sensitivity of a satellite well was generally around 20% better than the optical well on the same device. The only difference between the designs is that for the satellite wells, the metallization runs down into the well (Figure 4.6.6), and so it was concluded that a "significant" level of current must flow through the pn junction at the bottom of the well before triggering occurs. Since in the optical well this current has to flow through the lateral resistance of the thin pilot emitter, some of the bias across the junction will be offset. In the satellite well, the metallization is brought much closer to the centre of the well, so that the effect is reduced; Figure 5.2.3 compares the two well designs and their equivalent circuits. This conclusion (ie that significant current flows through the optical well before triggering occurs) was later corroborated by measurements made during turn-on; see section 5.7.

Comparison with Results from Other Workers.

Since the laser pulse was so short, the sensitivity in Figure 5.2.1 is quoted as an energy in nanojoules. However, most workers in this field have used LED's driven with long pulse widths (eg 100us) to trigger their light-sensitive devices, and so published figures for sensitivity are usually quoted as power in milliwatts (eg [5.1]). It is therefore difficult to compare sensitivities directly. If it is assumed that all the optical energy delivered to the device during its delay time contributes to turn-on, then it is necessary to know both the threshold power and the associated turn-on delay time in order to compute a threshold energy. The delay time near threshold triggering is normally long and very sensitive to small changes in the level of overdrive;

since it is not usually of practical interest, it is rarely quoted. If a delay time of 5us at threshold triggering is assumed for the data given in [5.1], then energies of around 25nJ result, making the prototype similar to state-of-the-art developments elsewhere. If a delay time of 20us is assumed, as could be the case, the prototype looks substantially better. In one of the light-triggered valves referred to in section 4.4 [4.14], a minimum operating energy of 20nJ was used.

Effect of Fibre Alignment.

The threshold optical sensitivity was also measured as a function of the alignment of the fibre over the well. It has been reported [5.2] that when the n-emitter is left in place in the centre of the well, sensitivity is improved for a shorted structure. This is said to occur because, in addition to the photo-current generated in the main depletion layer, current is also generated in the residual depletion layer between the p-base and the n-emitter. This current circulates through the shorted structure and contributes to forward-biasing the n-emitter. Although the depletion layer here is much narrower due to the junction being partially forward-biased, the photon flux will be substantially greater because it is close to the surface. The contribution to turn-on sensitivity is therefore claimed to be significant.

On the basis of this theory, one might expect the optical sensitivity to improve if the fibre was moved away from the secondary well to a position over region B (Figure 4.6.4). In practice, this was not found to be the case; the sensitivity deteriorated as the fibre was moved into region B, ultimately reaching a level about 40% worse than when it was over Region A. It was concluded that any benefit arising from the

circulating photo-electric current was more than offset by the increased attenuation in the silicon. This would be exacerbated by the high doping in the pilot emitter, which gives it a higher attenuation than intrinsic silicon. Further tests to investigate this theory are described in Chapter Six.

Effect of Optical Pulse Length.

At the outset, it was expected that optical sensitivity would be independent of pulse length for very short pulses. This was because pulses of less than one microsecond duration would be well within the "charge-controlled" triggering regime, so that a constant energy criterion would apply. Tests were carried out on a 30mm sample to confirm this, in which the sensitivity at 30V bias was measured for pulses of 100, 200 and 400ns duration. No variation was found within the bounds of experimental error. When the alignment of the fibre over the well was not disturbed, repeatability of sensitivity measurements was found to be better than $\pm$1%.

Effect of Repetitive Optical Pulsing.

Another feature of interest in the optical triggering of these devices is the optical integration time. The practical importance of this parameter is that it determines, for a train of short optical pulses to the gate, how close the pulses have to be to ensure that the thyristor will turn on immediately in response to any randomly-timed forward voltage ramp. This is because, under some conditions, it is necessary to ensure that the thyristor behaves as a diode, and the light triggering system aims to achieve this by giving the device a train of optical pulses (see Chapter Seven). The design was based around the

assumption that as long as the pulses were not separated by more than 20us, they would appear as continuous illumination to the gate. The following experiment was undertaken to test this.

For six 100mm samples and two 30mm samples, the optical sensitivity of the satellite wells was measured. The pulse output of the laser was then reduced to 20 - 30% below the threshold level, and the pulse generator set to repetitive operation. The frequency of the pulses was then increased from 10Hz until the device triggered. The results of this experiment are given in Figure 5.2.4. The frequencies necessary to cause sub-threshold triggering varied from 9kHz to 22kHz, and there was a reasonable correlation between the amount by which the pulse energy was reduced below threshold, and the frequency necessary to cause triggering. This experiment does not fully represent the conditions of interest described above, but it does give more confidence to the 20us figure assumed above, since the minimum pulse spacing required was 50us. Further consideration of these results is given in Section 5.8, "Theoretical Discussion of Optical Triggering".

c) Conclusions.

Results of the optical sensitivity tests were generally as expected. A good level of optical sensitivity was achieved, and the improvement in sensitivity with applied voltage was demonstrated. For short optical pulses, the triggering sensitivity was shown to be independent of optical pulse length, and the response of the gate to repetitive pulsing with sub-threshold pulses was characterised.

Only in two areas were the results unusual. Firstly, a layer of metallization on the pilot thyristor emitter was found to give a significant improvement in optical sensitivity, and an explanation for

this was presented. Secondly, illuminating the gate through a layer of n-emitter did not give an improvement in sensitivity (as reported by others), although this may have been because the optical well design was not optimised for this mode of operation.

5.3 EVALUATION OF FORWARD RECOVERY PROTECTION.

a) The Test Circuit

The circuit used for testing forward recovery protection is shown in Figure 5.3.1, and is based on one developed at the GEC Engineering Research Centre in Stafford. The components shown inside the dashed line represent the thyristor's electrical environment inside the valve (described in Chapter Two). L1 and L2 are non-linear inductors, R3 and C3 the damping circuit, and C4 the fast grading capacitor. R1 is the dc grading resistor. To the left of the dashed box is the main current pulse circuit; to the right is the impulse circuit.

Referring to Figure 5.3.1, circuit operation is as follows. A half-sinewave current pulse is passed through the device under test ("DUT") with a magnitude sufficient to fully spread conduction in the device. The pulse is triggered by firing TH1 and DUT, so that C1 discharges through L3 with a half-period of 4ms. At the end of the pulse, the device undergoes negative recovery, in which some of the excess charge in the transistor bases is removed by reverse current flow, and the rest decays by natural recombination. At this point, the state of the thyristor is the same as it would be after negative recovery in an HVDC valve.

At some predetermined time after current zero, the thyratron is fired. This discharges C2 into the thyristor-level circuit, and causes a forward dv/dt to appear across the thyristor. If the thyristor has had insufficient time to recover, the charge swept out by the ramp will trigger the device into conduction, and the ensuing inrush current may damage it (See Chapter Four). The rate of rise of the forward ramp can

be set to "high" or "low" depending on whether or not R2 is included in the circuit. If the resistor is in, the dv/dt is low (typically 10V/us), and the thyristor will not be damaged if it fails to recover. This is useful for determining the recovery time of the thyristor. Shorting out the resistor increases the ramp rate to around 100V/us, so that if the thyristor undergoes marginal forward recovery failure, it would be damaged unless it was self-protecting. Higher dv/dt rates could be achieved by reducing the value of L4.

b) Design of Irradiation Shields

The prototype thyristor design relied upon using selective electron irradiation to create regions inside the gate structure with longer recovery times than the main cathode (section 4.6). The dimensions of the irradiation shields to be used over the selected regions are discussed below.

Initially, three mask designs were used for electron irradiation. Mask #1 was designed to produce a long life-time region under the satellite wells, so that charge that diffused there during normal conduction would recombine more slowly than charge under the main cathode. See Figure 5.3.2. However, it was recognised that this might not be the ideal design for two reasons, as follows.

Firstly, there is no particular reason why the last area of the device to be in conduction should be next to a satellite well. As the current through the device falls, the conducting area will break up into filaments as those areas with a slightly higher Vf go out of conduction. Thus some areas of the device will start to recover sooner than others. Therefore, even if there is a longer life-time under the satellite well, the mechanism will be defeated if the last area to

conduct is some distance away from the well, so that the well starts its recovery period before the remote area. This effect will be more pronounced at low di/dt rates, when there will be a longer time delay between the onset of current filamentation and current zero. To compound this effect, the row of shorting dots along the inside edge of the main emitter (necessary for dv/dt performance) will raise the local Vf in that area, tending to force current away from the satellite well.

Secondly, because the mask only covered the well, the area between the well and the inside edge of the main cathode would be irradiated, as can be seen in Figure 5.3.2. Since any charge reaching the satellite well has to diffuse across this region, irradiating it will hinder this since the diffusion length will be reduced.

To overcome these drawbacks, Masks #2 and #3 extend onto the main cathode to create a region of low Vf round the satellite well; see Figures 5.3.3 and 5.3.4. This procedure should ensure that the last filament of current flow is next to a well, and that this is therefore the last area to start recovery. Mask #2 has uniform thickness, so that no irradiation reaches the main cathode near the well or the region between the main cathode and the well; this should facilitate maximum diffusion of charge to the well. However, it was possible that this would make the edge of the main cathode vulnerable to forward recovery damage, since it has the same carrier life-time as the well. To avoid this eventuality, Mask #3 was designed with a graded thickness, which should help in two ways: firstly, it ensures that the longest life-time region is under the well, and secondly, it creates a horizontal Vf gradient in the bulk of the device. This will encourage

lateral current flow under the well, increasing the amount of charge diffusion as shown in Figure 5.3.5.

The electron beam energy for these tests was 5.5MeV, and the devices were irradiated by placing them on a trolley which passed through the beam at a controlled speed.

c) 30mm Device Tests.

The first devices to be tested for forward recovery protection were the 30mm thyristors (see Figure 4.6.11). Devices were selected for testing on the basis of optical sensitivity and tq (recovery time). During processing, the depth of the optical well had been deliberately varied, so that there were wide variations in sensitivity between devices. However, although two wells on the same device had nominally the same depth, there was found to be a substantial variation in optical sensitivity between them. Typical variations were 20 - 30%, but in some cases one well was almost twice as sensitive as the other; see Figure 5.3.6. This was attributed to unevenness in the chemical etching used to produce both the main well and the secondary well, which was confirmed by measurement of the well depths using a "Talysurf" mechanical surface measurement device.

There was also found to be a very wide variation in tq between devices, with the recovery times varying from 56us to more than 500us under the same low dv/dt conditions (as shown in Figure 5.3.6). A spread as large as this is unlikely to be due simply to process variation, and so it was suggested that there might be inhomogeneities in the solder bond to the molybdenum backing disc. A region of poor contact in this bond would act as a reservoir of charge, since the charge extraction during reverse recovery would be reduced, and this would extend tq accor-

dingly. It would also defeat the mechanism of forward recovery protection by creating local regions with a recovery time possibly longer than that in the satellite well. Devices with very long recovery times were therefore excluded from the tests; otherwise, the wells which were most sensitive optically were used, since these would probably be most sensitive to forward recovery turn-on.

Interference Problems.

Initially, the device being tested was triggered electrically by means of an aluminium wire ultrasonically bonded to the thin spoke connecting the metallization on the two wells. However, it was found that spurious electrical signals appeared on this wire when the thyratron was fired to initiate the forward voltage ramp. This led to the device being turned on by the gate while the forward ramp was being applied, giving the appearance that the forward recovery protection mechanism was working. In some cases, the turn-on was very close to threshold, so that the delay times were very long. Under these long delay time conditions, the ramp would go up to its maximum value (eg 2kV) and stay there for perhaps 30us before the device triggered; in fact, a delay of over 100us was observed in one case. Although this did not demonstrate operation of the forward recovery protection, it at least proved that the device could survive marginal triggering from this voltage in the thyristor-level circuit.

The interference was traced to two sources: firstly, there was cross-talk in the electronic timing circuits, such that when a trigger pulse was sent to the thyratron, a brief spike appeared on the gate leads to the DUT, and secondly, interference was electro-magnetically coupled from the firing of the thyratron itself. In order to eliminate both of

these sources, it was decided to optically trigger the DUT using the arrangement shown in Figure 5.3.7. The laser drive circuit was fed from a pulse generator, and the input to the pulse generator was fed from the test circuit via an RC damping network to attenuate spurious signals. In order to demonstrate that no spurious laser pulses were being generated, the voltage on the energy storage capacitor in the laser drive circuit was monitored. (The design of the laser drive circuit is described in Chapter Seven). The 125ns electrical pulse to the laser was too short to see on the oscilloscope with the time-base used to monitor the forward recovery event; however, the storage capacitor had a recharge time of approximately 50us, and the dip in its voltage could be clearly seen when the laser fired.

In order to distinguish between random turn-on in the main cathode and protective turn-on through the gate under forward recovery failure, the voltage across the two control resistors was measured; see Figure 5.3.8. If turn-on was occurring via the gate, then current flowing through the control resistors would cause a voltage to appear there before the collapse in main cathode voltage. If turn-on was occurring in the main cathode, however, then the only voltage that would appear on the gate structure would be the normal p-base to n-emitter bias, ie of the order of 1V. This measurement required that connections be made to the gate structure, and therefore once again raised the possibility that interference might be coupled into the gate. However, it was felt that the oscilloscope input impedance was sufficiently high to prevent any significant currents flowing in the measurement loop due to electromagnetic pick-up.

Test Results.

First Set of Devices.

The first set of devices that were tested are listed in Table 5.3.1. These consisted of a control device with no irradiation, a device which was irradiated uniformly with 40kRad, and two devices with #1 masks and 40kRad and 80kRad doses respectively. The unirradiated control device was tested first, and whilst it was being electrically-triggered, the predicted phenomenon was observed during forward recovery failure, ie the voltage across the control resistors rose before the device switched. This meant that the device was being turned on by the gate during forward recovery failure, as shown in Figure 5.3.9 for a low dv/dt ramp. (Note that the zero current level before triggering is offset due to saturation of the current transformer caused by the main current pulse). This was initially taken to mean that the forward recovery protection mechanism was working without selective irradiation, which was at least partially plausible, because although there would be little charge left under the well after recovery compared to the main cathode, the well is much more sensitive to turn-on. However, once the presence of interference was discovered on the gate leads, and optical triggering was employed, these gate voltages no longer appeared during forward recovery failure. It was therefore concluded that the protection mechanism was not working. Despite this, prolonged attempts to damage the device by forward recovery failure were unsuccessful, and the device survived many ungated turn-on events from 1.7kV. This was the most severe recovery failure that could be induced, while the ramp voltage was set to its maximum of 2kV. Only when the device was raised from room temperature to 70°C did failure occur; in this case, a single ungated turn-on from 1.5kV destroyed the device. The failure site was

between shorting dots on the main cathode, which is consistent with normal forward recovery failure damage; see Figure 5.3.10.

The sensitivity of forward recovery damage to temperature is already known [5.3]. and is presumed to be due to two factors:

- the reduction in spreading velocity as temperature increases; at room temperature, the area in conduction increases rapidly after initial turn-on, so that the energy density and total energy dissipated in the turn-on event are reduced.
- at higher temperatures, the rise in temperature necessary to reach thermal runaway will be reduced.

After the failure of this control device, the other three devices were also tested under similar conditions at 70°C and yielded the same results, ie normal forward recovery failure damage between shorting dots on the main cathode. The protection mechanism was therefore not working with the masks and radiation doses used.

Second Set of Devices.

Following the unsuccessful results from the first batch of four devices, a second set of 15 test devices was prepared; see Table 5.3.2. These included mask types #2 and #3, with different radiation doses. There were also three devices which had a #1 mask on the main cathode and different doses; this was to try to determine the level of selective irradiation necessary to control the position of the last area to recover. Fast Neutron Detector ("FND") diodes were used to measure the radiation given to unmasked areas, and also to measure the dose penetrating the various masks. The change in Vf of these diodes

can be used to calculate the radiation dose, and the results of the FND diode measurements are given in Table 5.3.3.

Unfortunately, the tests on this larger sample of devices were also unsuccessful. None of the devices showed any voltage generated in the gate during forward recovery failure, and all the devices were destroyed by forward recovery damage when heated to 70°C. The devices with the #1 masks failed in the main cathode, as before; those with #2 and #3 masks usually failed on the inside edge of the main cathode in the shielded area (Figure 5.3.11), although some with the #3 mask and low radiation doses failed in a random location on the main cathode.

Of the samples which had a #1 mask on the main cathode, those which had a 11.5kRad dose failed on the main cathode, but not in the shielded location. A dose of 11.5kRad was therefore insufficient to control the position of the last area to conduct (the region of "vestigial current"). The one which received a 40kRad dose suffered voltage failure before marginal forward recovery failure could be induced, and so no result was obtained. The dose necessary to control the vestigial current is important because the normal radiation dose for lifetime control in HVDC devices is around 20kRad; thus if it is necessary to use levels higher than this for vestigial current control, a penalty in the Vf of the device would be incurred. Results from the main batch of devices showed that, where the mask was extended onto the main cathode and a 75kRad dose used, this did determine the position of the failure site. The dose necessary for controlling the vestigial current flow was therefore concluded to be between 11.5 and 75kRad. However, due to the uncertainties about alloying arising from the wide variations in tq, it was felt that the figure might be better for fully-floating devices due to their better contact uniformity.

d) 100mm Device Testing.

After completion of the above tests, some fully-rated 100mm devices became available, and so testing concentrated on these larger structures which had the correct sensitivity in the optical and satellite wells.

In view of the failures on the inside edge of the main cathode associated with masks #2 and #3 on the 30mm devices, a fourth mask design was proposed. The mask would have to come reasonably close to the main cathode so as to attract current flow by the reduced Vf, and avoid reducing the carrier diffusion length, but it must not come so close that forward recovery damage occurred on the main cathode edge. A compromise of half way between the auxiliary ring and the main cathode edge was chosen; see Figure 5.3.12. These #4 masks were produced from steel rod and were 8mm deep, steel being used since it is more easily machined than molybdenum. The mask thickness was increased to 8mm, compared to 3mm for the molybdenum masks #2 and #3, because of the reduced absorption of electron irradiation in steel.

The first 100mm device was given a 40kRad dose of radiation with two of the satellite wells shielded with #4 masks, and the test results are shown in Figures 5.3.13 to 5.3.16. Figure 5.3.13 shows the voltage generated across Rc1 and Rc2 during marginal forward recovery failure for conditions of low dv/dt, 350V ramp; the peak voltage across the control resistors is around 25V. Figure 5.3.14 shows failure under similar conditions on an expanded time-scale, and it can be seen that the rise in voltage across Rc1 and Rc2 precedes the main increase in anode current at turn-on. Figures 5.3.15 and 5.3.16 are records for

failures with a high dv/dt, 350V ramp; in these cases, the voltage across Rc1 and Rc2 reaches a peak of around 75V, and again the rise in this voltage clearly precedes the main increase in anode current and the collapse of anode-cathode voltage. The conclusion drawn from these results was that the protection mechanism was working.

The ramp voltage (and therefore dv/dt) was gradually increased, and the recovery time "tq" measured, as shown in Figure 5.3.17. When the ramp voltage was increased to its maximum level of 2.0kV, however, the device failed on first application of this ramp, even though the ramp was applied at the maximum delay of 1ms after current zero. Two failure sites were evident on the device; one on the inside edge of the main cathode, and the other in the bulk of the main cathode, which makes it difficult to define the cause of the failure. However, the device was known to have reduced voltage blocking capability before the tests started, with leakage starting to increase significantly at around 2kV, and so one site could be attributed to voltage breakdown. An alternative explanation is that, although the protection mechanism was working effectively at low ramp voltages, the inherent delay in its operation may have allowed turn-on to occur "simultaneously" on the main cathode. This could have created limited damage, making that area vulnerable to voltage failure.

In addition to forward recovery tests, the next two devices to be tested were also used for the "charge extraction" experiments described in section 5.4. The first of these was tested for low dv/dt forward recovery failure before irradiation; see Figure 5.3.18. From this figure, it is clear that there is negligible voltage across Rc1 and Rc2 during forward recovery failure, which reinforces the conclusion that the protection mechanism was working for the previous device,

where around 25V was seen across the control resistors under the same conditions. This new device was then irradiated with all the satellite wells shielded by #4 masks, but unfortunately it failed immediately when it was tested after irradiation, as described in section 5.4. The major failure sites were the scorched auxiliary metallization rings around two of the satellite wells. A further device was irradiated with two wells masked, but this also suffered the same fate. Possible causes of these failure and remedies for them are discussed in section 5.4.

To prevent further recurrences of this failure mode, subsequent devices had small breaks made in each end of the thin semi-circle of auxiliary metallization around each well, as shown in Figure 5.3.19.

The next device was prepared with 40kRad irradiation and all three satellite wells shielded with #4 masks. During forward recovery failure, no significant voltage (ie less than 1V) was seen across Rc1 and Rc2 under both low and high dv/dt conditions. However, during testing at room temperature, the device failed in the region of the optical well; see Figure 5.3.20. At the point where each radial spoke rises from the level of Rc1 to the ring of metallization on top of the optical well, the aluminium had melted. The voltage at which the device had switched during forward recovery failure was 300V.

It was already known that these points represented potential weak spots in the structure. This is because the aluminium metallization is evaporated onto the structure in a direction normal to the surface, and so where the tracks cross over "steps" in the surface, they can become quite thin. These weak points could then be destroyed by current flow resulting from triggering in either the optical well or one of the

satellite wells. The ring on the optical well acts as a distribution point, so that current from one satellite well can flow to all three sets of auxiliary amplifiers. If the failure was due to inrush following optical triggering for the main current loop, the device should have been damaged in the first few test shots; the fact that it happened after some time suggests that triggering from a satellite well was responsible. However, this is not consistent with the lack of voltage on Rc1 and Rc2 during forward recovery failure.

A second device was prepared as above, and to prevent a repeat of the failure mode described, aluminium wire was ultrasonically bonded over the steps round the optical well, as shown in Figure 5.3.21. This wire was 250um in diameter compared to the nominally 30um x 200um cross-section of the metallization. Under test, no voltage was observed across Rc1 and Rc2 during forward recovery failure, as before. However, the device again failed (for ungated turn-on from 450V) at one of the steps in the metallization; in this case it was where one of the spokes joined a satellite well, as shown in Figure 5.3.22. This failure site seemed to suggest strongly that turn-on was occurring at the satellite well, but this was still contradicted by the absence of voltage across Rc1 and Rc2.

A third thyristor was prepared with all the steps strengthened with aluminium wire bonds. This device also failed under similar circumstances to the previous two, with the failure site being underneath the aluminium bond over one of the steps at the optical well. The aluminium bond wire had to be removed to find the failure site. The fact that the wire was not melted suggested that something more than simple current density was causing the failure, and it may have been a recurrence of

the previous difficulties encountered with the parasitic diode breaking down (eg Figure 4.7.4). It had not been expected that much current would flow through the diode in this region, since the path to the auxiliary metallization was longer than at other points on the periphery of the well, and therefore would have a higher resistance. However, if the presence of the silicon nitride insulating layer reduced the breakdown voltage of the diode for some reason, then this path could attract significant current, resulting in destruction of the diode. For example, the etch to selectively remove the silicon nitride would go beyond the thickness of the nitride; this would result in the surface doping concentration of the diode under the nitride being higher than that elsewhere around the well, giving it a slightly lower breakdown voltage. Even so, this voltage should have been measurable at forward recovery failure, and no voltage was observed. Whatever the explanation, the fact that the devices were failing in the gate region rather than the main cathode did suggest that the protection mechanism was operating in some fashion.

Because of these unusual failures, testing was discontinued until a technique for measuring current flow through a well was developed; see section 5.7. This involved breaking the metallization at the step(s), and then reconnecting it through an external loop of wire so that the current flow in the wire could be measured. Figure 5.7.7 shows the arrangement used. It was therefore decided to apply this technique to see if the satellite wells were turning on at forward recovery failure, as suggested by some of the test results.

A device was prepared with all three wells shielded, and a 40kRad dose of radiation. The thin auxiliary metallization rings round the wells

were broken as before. All four wells were isolated from the gate structure by breaking the metallization over the steps, and one wire was bonded to each well and brought out. The three radial spokes were then connected together, brought out and joined to the four wells externally. In this way, the current through all the wells could be monitored, as shown in Figure 5.3.23.

During forward recovery failure, no current was observed to flow through the satellite wells on the device. This agreed with the fact that no voltage had been measured across Rc1 and Rc2 during forward recovery failure on the previous thyristors.

After testing under high dv/dt conditions, the device suffered partial failure. Its voltage withstand capability in the forward direction was reduced to 2kV, and the reverse blocking capability was lost. On the reverse (anode) side of the device, there were a number of small melt sites. A possible explanation of this is that the clamping pressure during the experiment was insufficient, so that only point contacts were formed. These sites were then melted by the main current loop, damaging the reverse blocking junction. However, no damage sites were observed on the cathode side.

A further device was prepared, and in order to give the protection mechanism as much chance as possible of working, two stages of radiation were used. Firstly, 10mm diameter shields were placed over each well, and a 40kRad dose given: this should have ensured that the last areas to be in conduction would be near the satellite wells. Secondly, the #4 masks were placed over the wells, and a dose of 20kRad given.

Even with this arrangement, no current was observed to flow through the

wells at forward recovery failure. After several high dv/dt recovery failures, the device was partially damaged in a manner similar to the previous one.

e) Conclusions.

Interpretation of the test results was hampered by the large number of unusual failures. However, the conclusion from these tests was that this mechanism for forward recovery protection, as embodied in the prototype design, was not successful. Although one device showed evidence of it working to a certain level, in the majority of cases it clearly did not operate. Although it is recognised that not all combinations of shield size and radiation dose were tried, and therefore the test devices may not have been fully optimized, a mechanism which would work for mass-produced devices operating in the field would be expected to produce more positive results under laboratory conditions. Although a modified version of this technique was retained in the next generation of test structures (Chapter Six), the majority of designs were therefore based on different protection philosophies.

5.4 CHARGE EXTRACTION TESTS.

a) Introduction.

The forward recovery protection mechanism in the prototype design relies on charge diffusing under the satellite wells during conduction, and then recombining more slowly there than in the main cathode. In order to attempt to verify this, experiments were carried out to measure the charge extracted from different areas of the device by a forward voltage ramp.

b) 30mm Device Tests.

The scheme used in the first attempt to measure charge extraction is shown in Figure 5.4.1. The metallization on top of one of the wells on a 30mm device was isolated from the central spoke by scratching away some of the metal with a scalpel. The forward ramp part of the test circuit was then modified so that the ramp was only applied to the isolated well; see Figure 5.4.2. A clip-on current probe would be used to measure the current drawn by the ramp. However, with this arrangement the voltage ramp drew charge from the whole device rather than just the well, and so the ramp was applied to the main cathode in addition to the well, as shown in Figure 5.4.3. This arrangement allowed extracted charge to be measured separately for the well and the main cathode. In these tests, the device was optically triggered from the other well.

Unfortunately, current sharing between the well and the main cathode resulted in approximately 1A flowing in the well during the main current pulse. This would render the results invalid, since the well would not normally conduct once turn-on is complete. Diodes were

therefore introduced into the circuit as shown in Figure 5.4.4. An external power diode was used to connect the DUT to the thyristor-level circuit in order to prevent current being drawn from the damping circuits during the forward ramp; this would allow charge extraction from the main cathode to be measured. A 1000R resistor was added in parallel with the DUT to discharge the snubber capacitance in the ramp circuit between shots.

Experimental results for a 30mm device are shown in Figures 5.4.5 to 5.4.7. This device had received a 40kRad dose with a #1 mask over the isolated well. Figure 5.4.5 shows the charge extracted from the main cathode, Figure 5.4.6 shows charge extracted from the masked well, and 5.4.7 shows charge extracted from the unmasked well. In these figures, each oscillogram is a superposition of several shots with different times of ramp application. The measurement of charge extracted can only be carried out after the recovery time tq has expired, since if the device fails to recover and turns on, the measurement will be meaningless. Consequently, the absolute levels of charge extracted are quite small.

Repeatability in the measurements was poor, with the amplitude of the current pulse varying by 20-30% for nominally identical conditions. This may be explained by the fact that the tail of the charge decay envelope is being measured, and therefore any small variations in device temperature caused by heating during the main current pulse would have a substantial impact on the residual charge under the well. Temperature determines carrier lifetime, which affects both the level of diffusion under the well and the subsequent decay rate.

In Figures 5.4.5 (the irradiated main cathode) and 5.4.6 (the

unirradiated well), the decay envelopes of the peaks of the current pulses both have a time constant of approximately 120us. The carrier lifetime in these two areas should be different due to the irradiation step, but given the level of experimental error in these measurements, one would not expect to be able to detect less than a 2:1 change in carrier life-time; it was not anticipated that the radiation doses used would produce a change this large. In the unmasked well (Figure 5.4.7), however, the current pulses are significantly smaller than those for the masked well; this result supports the basic principle of the forward recovery protection mechanism, ie that by selective irradiation, charge can be introduced under a satellite well at a level significantly higher than for a normal device.

c) 100mm Device Tests.

These experiments were repeated on a 100mm device. In addition to measuring charge extracted from the isolated well and the main cathode, however, the charge extracted from the auxiliary amplifier regions was also measured; see Figure 5.4.8. These currents were measured first of all for an unirradiated device, and are shown in Figures 5.4.9 to 5.4.11. It was intended subsequently to measure these currents after the device had received a 40kRad dose of radiation with all three satellite wells shielded with #4 masks.

Unfortunately, however, the device failed catastrophically the first time it was tested after irradiation, even though the test was only using low dv/dt and a 350V ramp. There were three failure areas on the device; the auxiliary metallization ring around two of the wells was scorched, and there was a short-circuit bridge from the auxiliary to the main cathode. Figure 5.4.12 shows the damage to one masked well,

and the auxiliary ring round the other damaged well was scorched in a similar manner.

The experiment was repeated on another 100mm device as before. The charge extraction prior to irradiation was measured, but this time only two of the wells were shielded, so that it would be possible to compare the charge extracted from a masked and unmasked well. However, after irradiation this device failed in the same way as the previous one, ie with the auxiliary metallization around the two shielded wells being scorched before any useful measurements could be made.

The reason for these two failures is not entirely clear. However, it seems almost certain that the damage was caused by the main current loop rather than the re-applied voltage ramp, since the ramp conditions were very gentle. One possibility is that since the shields covered part of the control resistor Rc2, this would give a lower value of resistance around the wells. A disproportionately large amount of turn-on current might therefore be attracted into these areas, possibly sufficient to melt the thin metallization (200um by 30um cross-section). This would be in addition to the fact that the curved geometry of the control resistance in these regions would already give it a slightly lower resistance per unit length of the edge than that in the linear regions. However, this does not explain why only two of the three shielded wells on the first device were damaged.

Since the existence of charge diffusion had already been satisfactorily demonstrated on the 30mm device, it was decided not to pursue these experiments in case any more samples were destroyed.

d) Implications for Future Designs.

The suspected problem of reduced control resistance in the shielded areas would need to be addressed if this type of design was pursued. The purpose of the auxiliary metallization ring is to apply the dv/dt compensation potential (generated across Rc1) uniformly around the satellite well. Since its function is to apply voltage rather than carry current, its performance would probably not be seriously affected if a small break was made in the metallization at each end of the thin ring. This would prevent the ring carrying substantial turn-on current, whilst the potential drop across this break under dv/dt conditions would be minimal. It was also questioned as to whether this ring was really necessary at all, or whether the main cathode metallization could be used as the equipotential point; see Figure 5.4.13. However, since it was decided not to use dv/dt compensation for the next generation of devices, the problem would not arise.

e) Conclusions.

Measurements on the 30mm devices confirmed that, by selective electron irradiation, it is possible to cause significantly higher levels of charge to diffuse into part of the gate structure than would be there normally. Whilst this does not prove that the proposed technique for forward recovery protection is viable, it at least demonstrates the underlying principle. No results were obtained from 100mm devices due to unexpected failures, which were related to the dv/dt compensation mechanism used on the prototype design. Means for avoiding these problems were discussed, but since it was not proposed to use dv/dt compensation on the next generation of devices, they were not pursued.

5.5 EVALUATION OF DV/DT CAPABILITY.

a) Introduction.

In view of the relatively poor dv/dt results obtained during testing in Stage One of the development (section 4.7), it was important to repeat dv/dt tests on structures with the correct depth of optical wells. No tests were carried out on 30mm devices, since because dv/dt capability is strongly affected by the horizontal geometry of the gate structure, these simplified structures would not have been representative.

b) 100mm Device Tests.

Initially, two devices were tested to 4kV/us at room temperature, with a ramp from 0 to 4.5kV. Both devices withstood this without triggering, which indicated that they were better than the previous 75mm samples in this respect. This was because the optical wells were of the correct depth, unlike the wells on the 75mm devices which were known to be too deep. However, when tested at 110°C both devices failed catastrophically at 2kV/us, one with a peak ramp voltage of 2kV and the other with 3kV. The failure site was on the inside edge of the main cathode, where the edge goes round a fairly sharp 300° corner; see Figure 5.5.1.

When the design of the shorting pattern on this corner is examined under a microscope, it is evident that the shorting dot on this corner has to carry displacement current from a much longer part of the auxiliary metallization than the dots along the straight edges of the cathode. This problem is not fundamental, and could be overcome by redesigning the horizontal geometry of the device in this area. In order to enable the dv/dt capability of the rest of the gate structure to be measured, all further tests were carried out with the auxiliary

metallization shorted to the main cathode. This would allow displacement current flowing in the gate region to bypass the shorting dots on the inside edge of the main cathode.

The next device to be tested at 110°C failed on the edge bevel at 1kV/us. Normally, the edge bevel of the thyristors is given a large radiation dose (around 1MRad) to "kill" leakage and displacement currents in this vulnerable area. However, this device had not received any edge irradiation, and so this result confirmed the necessity of this procedure for dv/dt capability.

These devices do have a peripheral short around their circumference, but beyond this there is a ring of main emitter n-type diffusion; see Figure 5.5.2. This ring is used as a getter layer to remove defects from the edge of the device, which historically has been found to improve performance; however, it does have the disadvantage of creating a thyristor structure in the most vulnerable part of the device. This n-type diffusion is already etched away in other parts of the thyristor, and so it would not cost anything to remove it in this area as well; only a simple change to the mask set would be required. This may be worthy of further investigation.

It has already been explained in section 4.2 that dv/dt compensation techniques incur extra cost and complexity in device manufacture. In addition, achieving safe dv/dt switching (ie dv/dt protection) with a compensated structure requires very accurate tolerances on control resistor values. It was therefore decided to test the optical wells with the compensation de-activated, to see if adequate performance was achieved without it.

The compensation was de-activated by connecting together the pilot and auxiliary metallizations, so that now the complete gate structure was shorted to the main cathode; see Figure 5.5.3. It can be seen that, from the point of view of displacement current, the well will appear as a shorted structure; current flowing underneath it will go directly to the auxiliary ring, which is connected to the emitter of the well. It was recognised that, by shorting out the gate structure, the device would probably be destroyed if it did switch under dv/dt, since the well would carry the full in-rush current from the ramp generator; however, this was the only way to test the uncompensated performance of these structures.

The devices were tested at three different temperatures. In each case, a dv/dt rate was set, and the peak voltage of the ramp increased in 1kV steps to 3kV. If the device did not switch, the dv/dt rate was increased by 500V/us, and the ramp voltage reset to 1kV. The results are given in Table 5.5.1. In each case, the device failed at one of the wells (as expected), resulting in a crack propagating across the device from the centre of the well concerned. The requirement for these devices is that they withstand 2 - 3kV/us at 110°C, and more than 100V/us at 150°C. It can be seen that the present well design exceeds these requirements without compensation. As a result of this, it was decided to use an uncompensated design for the next generation of devices (see Chapter 6).

Ideally, it would be desirable to test several devices at each temperature to give more confidence to these results.; however, there was only a limited number of 100mm devices available for the complete test program. It should be added that monitoring during processing of the samples showed a good uniformity within the production batch, and

so it was felt that these results were adequate for the present time. More extensive investigation was to be carried out on the next generation of test structures, and this is described in Chapter Six.

c) Test Conditions.

During dv/dt testing, there was some concern over the shape of the voltage ramp being applied to the thyristor. The capacitance of the device is highest at zero bias, and falls off rapidly as voltage is applied to it. This is because the capacitance of a parallel-plate system is inversely proportional to the plate spacing, and the width of the depletion layer increases rapidly when voltage is first applied to it. The device therefore presents a varying load to the dv/dt ramp generator, and since this has a finite internal impedance, the ramp rate will be a function of voltage; Figure 5.5.4 shows a typical result. Since turn-on is controlled in part by the peak current that flows in the device, this real situation will be less onerous than an ideal linear voltage ramp, which would generate a significantly higher peak current at the start of the ramp.

However, consideration of the actual service conditions for the device indicated that these tests would not give misleadingly optimistic results. The requirement for a thyristor in an HVDC valve is that it does not turn on in response to the worst dv/dt rates it will see under non-fault conditions, which occur when an adjacent valve is fired from its maximum voltage. A typical example of the thyristor voltage waveform resulting from this event is shown in Figure 5.5.5.

The most important feature of this transient is that it starts from a negative voltage. It can be shown that a dv/dt ramp starting from a

negative bias is less severe than one starting from zero volts [5.4]; by the time the ramp reaches zero volts, both the forward and reverse junctions are depleted, and they behave as two capacitors in series. This significantly reduces the total device capacitance at zero volts, and therefore also the peak current. The test condition of starting from zero volts is therefore more onerous than the actual in-service withstand conditions, and so provided the ramp is not too slow initially, the test will err on the safe side.

d) Conclusions.

A flaw in the design of the gate meant that the 100mm devices were limited to a dv/dt capability of 2kV/us at 110°C. Although better than the 75mm devices tested in Stage One, a higher capability would still be preferable. Fortunately this flaw could be overcome by shorting out part of the gate structure, and subsequent tests showed that the optical well design had adequate dv/dt capability over the whole temperature range of interest, without the need for dv/dt compensation. Dispensing with dv/dt compensation reduces device cost, and also makes it much easier to achieve dv/dt protection. Design of the next generation of test structures therefore proceeded without dv/dt compensation.

5.6 EVALUATION OF VBO PROTECTION.

a) Introduction.

Section 4.6 explained the mechanism for VBO protection employed in the prototype thyristor design. It was reported in section 4.7 that the initial tests for VBO protection on the prototype thyristor design were not promising. Not many devices exhibited VBO operation, and those that did had a VBO level that was too low to be useful. Attempts to use laser zapping to control the VBO level also gave poor results, and so work in Stage Two of the project aimed to improve on these results. However, before the experimental work is described, some of the practical difficulties associated with VBO protection are discussed, since these form the background to the experimental work.

There are two main difficulties in achieving VBO protection: the first is to be able to accurately set the level at which the protection operates, and the second is to ensure that the device is not damaged by the inrush currents that flow when the mechanism does operate. With regard to the first problem, the VBO level in the prototype design is effectively set by detecting the position of the depletion layer. The secondary well (Region A in Figure 4.6.4) is progressively etched until the depletion layer is encountered; this is detected by checking to see if VBO action occurs after each etch. Since the width of the depletion layer increases as the square root of the applied voltage, the detection accuracy required increases as the square of the VBO level. Thus although reasonably good results have been obtained by some workers at VBO levels of 2000V (see section 4.4), the dimensional tolerances required for a VBO level of 5000V, as needed for the prototype design, are over six times tighter. If higher device voltages

are contemplated for the future, the problem becomes even more difficult.

To some extent these difficulties are a result of conventional thyristor design techniques; the p-base is deliberately doped to minimize penetration of the depletion layer, so that by the time the anode-cathode voltage reaches 5000V, a tolerance of $\pm$100V on the VBO level (typical for conventional over-voltage protection systems) requires the well depth to be controlled to better than $\pm$0.1um. This places stringent requirements on the processing technology used to etch the VBO well.

These tight tolerances mean that the VBO level must be "trimmed" after the device has been manufactured. During manufacture, the depth of the junctions within the device cannot be predicted to sub-micron accuracies, so that the VBO mechanism must be trimmed interactively with voltage measurements when the device has been completed.

The second problem, ie that of preventing damage occurring when the mechanism operates, is more conventional in nature, although not necessarily easier to solve. The difficulty here is that both punch-through and avalanche are very fast turn-on mechanisms, which therefore allow little time for the turn-on area to spread. The inrush energy is always high, since by definition the mechanism operates at maximum device voltage. As a result, for this type of structure the VBO protection level often falls each time the mechanism operates, since damage is accumulating at the VBO turn-on site. The VBO level stabilizes when the turn-on energy at that voltage is insufficient to cause further damage.

Apart from the technical problems described above, there is the economic barrier which arises because a conventional technique for VBO protection (the series string of Break-Over Diodes) is relatively inexpensive, and is compatible with the advanced valve philosophy due to its simplicity. Since a series string of devices is used, it is easy to select these to achieve very good accuracy on the break-over level. The total cost of using this conventional technique with a thyristor which was otherwise light-triggered and self-protecting was calculated to be approximately £225 per device. The incremental cost of built-in VBO protection must therefore be less than this figure.

The built-in VBO level cannot be tested until the device has been bevelled and passivated to withstand full voltage, so that any yield loss is expensive. In addition, the manufacturing technique necessary to achieve the desired accuracy in the VBO level will almost certainly be expensive, so that even if integral VBO protection is technically feasible, making it economically feasible could be equally challenging.

As a result of both the technical and economic difficulties described above, work on VBO protection was suspended after four months, so that resources could be concentrated on the areas of self-protection where there was no conventional solution available. The following sections describe the work that was carried out in the first four months of Stage Two of the development.

b) Experimental Work.

Due to the limited number of prototype samples available, some of the VBO experiments were carried out on 75mm and 100mm electrically-triggered thyristors (referred to as "ETT's"). In these cases, the region inside the conventional amplifying gate was etched to create VBO

action, although it was recognised that the lack of control resistors would mean that VBO turn-on might well be destructive.

c) Effect of a Conducting Surface in the Bottom of the Etched Well.

The first ETT tested was given a 60um etch in the gate region, which resulted in a total thyristor leakage of 30mA at 3kV, room temperature. Thus since the device needed 50mA to trigger, VBO action did not occur at this level. However, when a drop of mercury was placed in the etched region, the device switched at 2.5kV.

This experiment confirmed the theory that VBO action would be best achieved with a conducting surface in the bottom of the secondary well (region A in Figure 4.6.4). If the bottom of the well is perfectly smooth, as is produced by chemical etching, it is possible that the depletion layer could move past the bottom of the well, and simply set up an electric field in the air (or nitrogen) in the well. Breakdown might then occur by arcing along the surface of the well, which would almost certainly result in damage and therefore a change in the VBO level. However, if the bottom of the well was covered by a conducting layer of finite thickness, this could not support any electric field, and would therefore distort the depletion layer once it reached the well. Under these circumstances, VBO action should eventually be triggered by avalanche in the bulk silicon, which would be non-destructive if the current was limited to a safe level. Estimated field patterns with and without a conducting layer are shown in Figure 5.6.1.

In a second experiment, an ETT was etched whilst the leakage current was actually being monitored. With the acid in the well, VBO action occurred at 3kV. However, when the device was washed and dried, VBO

action did not occur until 4.5kV was reached. This confirms the above conclusions, since the acid contains free ions and is therefore a conductor.

d) Plasma Etching Experiments.

With the proposed design of optical well, it would not be easy to introduce a conductive surface into the bottom of the well without shorting it to the pilot n-emitter layer in the bottom of the main well (region B in Figure 4.6.4). In addition, neither aluminium nor ion implantation could be used for creating the conductive layer, since both of these require high-temperature sinter/anneal processes which would damage the rubber passivation on the device. It was therefore decided to use plasma etching to produce VBO action in the well. Potentially, plasma etching has two main attractions: firstly, it leaves a rough surface, which might be sufficiently conductive to safely trigger VBO action without further treatment, and secondly, it is possible to achieve very accurate etch depths by this method.

An ETT was etched until VBO action occurred at 4.1kV. A leakage tester was used to measure the VBO level, so that the "in-rush" current at switching would only be a few hundred milliamps. However, in spite of this the VBO level was not stable; the second VBO switching occurred at 2.4kV, and finally the level settled down to 1.7kV. It was not possible to tell whether this damage was on the surface or in the bulk of the silicon, and so it could not be judged from this experiment whether or not the plasma etching had left a sufficiently conductive surface.

e) Gold Sputtering Experiments.

Gold can be sputtered onto a device without the need for subsequent

high-temperature processes, and so is the most likely candidate for producing a conducting layer in the bottom of a well. On the next acid-etched device, a thin layer of gold was therefore sputtered into the well. This device did show a stable VBO level under repetitive operation on a leakage tester, which suggested that the plasma etching technique tried previously had not left a sufficiently conductive surface. If the thyristor temperature was increased from room temperature to 110°C, the VBO level reduced by only 150V, which confirmed expectations that this mechanism for VBO protection would be largely insensitive to temperature.

On the next device prepared by etching and gold sputtering, the VBO level was 1650V, but when a thicker layer of gold was applied, it lost all forward blocking capability. When the gold was cleaned off, the device recovered its blocking capability. This result is very interesting, because it implies that the acid-etched well had gone very close to the middle junction of the device, so that with a thick layer of gold the device was virtually short circuit in the forward direction. However, with a thin layer of gold the device could withstand 1650V, suggesting that the depletion layer must have gone past the bottom of the well, but that the gold was sufficiently resistive to prevent triggering occurring until 1650V was reached. This raises the possibility of a novel type of VBO protection, which is discussed below.

This experiment was repeated on a second device, which was etched until, with a clean, dry well, the leakage characteristic started to turn up at 3200V. A thin layer of gold was then applied to the well, and VBO action occurred at 2640V. When a second gold layer was applied (half the thickness of the first one), the VBO level fell to 2520V.

This confirmed the conclusion that the VBO level is a function of the thickness (and therefore resistivity) of the gold layer.

f) Novel Techniques for VBO Protection.

VBO Triggering with an "Over-etched" Well.

Some interesting possibilities are opened up if it is possible to etch a well which penetrates the depletion layer without causing VBO operation. A VBO mechanism can be envisaged where the well is deliberately etched beyond the point where punch-through would normally occur. A resistive substance is then deposited in the well which would draw a controlled amount of leakage current from the depletion layer. This leakage current is then used to trigger a conventional gate; see Figure 5.6.2. This approach offers the following advantages over the VBO technique proposed for the prototype design:

- it removes the need for very fine control of the well depth.
- the VBO level can be set by trimming either the resistance of the substance coating the well, or the sensitivity of the gate.
- in the trimming exercise, the VBO level can be set low and then increased. This is easier than setting it too high and then reducing it, since if the VBO level is above the target figure, it cannot be measured without endangering the device.
- It might be possible to design the gate to be sensitive to leakage and displacement currents, so that the VBO level becomes a controlled function of temperature and dv/dt. This would be desirable for HVDC applications.

The resistive substance coating the well could be a thin layer of gold, a thicker layer of a more resistive metal such as Ni-Chrome (which can

be evaporated), or perhaps polysilicon.

If it is possible to have a well which penetrates the depletion layer, then it may be practical to use such a well for optical triggering. Since the photons would be injected directly into the depletion layer, this would give virtually 100% efficiency in their utilization, with a correspondingly substantial improvement in sensitivity over existing designs. Furthermore, this would be achieved without any deterioration in dv/dt capability.

Some surface passivation might be needed in the well to prevent arcing at high fields, but this would presumably depend on how far into the depletion layer the well penetrated. It might be possible to combine the optical triggering and VBO functions into one well, or to have two separate wells and etch them at the same time.

VBO Triggering by Short Wave-length Light.

An alternative method for VBO protection, based on the more conventional accurately etched well technique, was also considered.A well would be etched (probably by plasma-etching) until it reached the depletion layer at the required VBO level. However, instead of using a metallic coating in the bottom of the well to distort the depletion layer and trigger avalanche, short wavelength light (eg green) would be shone in the well. This would create hole-electron pairs close to the bottom of the well, which would prevent the silicon being depleted. For normal operating voltages, the hole-electron pairs would recombine thermally since there would be no depletion layer near the bottom of the well to separate them. However, once the depletion layer approached the well, the carriers would be separated by the field, and an optically-induced

current would flow. Current to trigger the device would then come either from this photo-electric current flow, or from avalanche caused by distortion of the depletion layer.

The attraction of this "green light" technique for HVDC applications is that the short wavelength light could be delivered down the same fibre as the normal infra-red triggering pulses. This would require precision etching of the optical well, but this would at least give the benefit of maximum optical sensitivity for normal triggering as well as VBO protection. In addition, if the optical fibre link to the thyristor was broken, the device would lose its VBO protection and be destroyed. This is actually an advantage for HVDC applications, because with conventional systems a broken fibre results in long-term repetitive operation of the VBO protection for that thyristor. This places severe stresses on the damping circuit for that level, and designing for this condition incurs extra cost. It also means that the valve cannot be fired if its terminal voltage is less than the VBO level of one thyristor, which for small valves can impose operating restrictions.

The main potential difficulty with the "green light" approach is that the VBO level would probably be sensitive to some degree on the intensity of the short wave-length light. Since there will be differences in attenuation in the optical paths to different thyristors, this means that there will be variations in the VBO level between devices. The extent of this effect could not be quantified without further investigations.

g) Conclusions.

Some very interesting results were obtained from the investigations into VBO protection, and some novel techniques proposed. However, the

technical difficulties with producing integral VBO protection, combined with the low cost and proven nature of the conventional BOD approach, meant that it was decided to concentrate development resources in those areas of the self-protecting device where feasibility was still not proven (in particular, forward recovery protection). If the other areas were shown to be successful, then work on VBO protection could be picked up again if it was felt to be economically attractive.

5.7 EVALUATION OF TURN-ON CAPABILITY.

a) Introduction.

The full turn-on test circuit was not available until some time into Stage Two of the development. Tests on the 30mm samples were therefore carried out in a simple RC circuit similar to that used in Stage One. The 100mm devices were later tested in the full circuit.

b) 30mm Device Tests.

The RC test circuit was intended to simulate the step in-rush current from the thyristor-level damping network. The rate of rise of this step current is limited only by the impedance of the device and any stray inductance. Although the proposed thyristor-level circuit includes a non-linear inductor between the device and the damping circuit, this inductor is lossy and therefore appears to have a "shunt" resistor connected across it, whose value is less than that of the damping resistor. This inductor therefore does not limit the rate of rise of the step in-rush current. Figure 5.7.1 shows the test circuit used; it was limited to a maximum operating voltage of 2kV, and so a small value of resistor was used to enable high step currents to be produced.

It can be seen from Figure 5.7.2 that the first device tested exhibited a remarkable di/dt capability. For 1.9kV turn-on, the initial di/dt is 550A/us, increasing to 850A/us above approximately 200A. It is presumed that the di/dt is initially controlled by the device impedance, and that the increase from 550A/us to 850A/us occurs when turn-on starts in the main cathode.

The effect of the control resistors can be seen in the "re-entrant"

form of the anode voltage collapse. As current starts to build up through the optical well, the anode voltage increases (region A in Figure 5.7.2). This can be explained by reference to the equivalent circuit in Figure 5.7.3 where, when turn-on starts, some of the capacitor voltage will appear across the stray inductance in order to generate the di/dt. However, as the current in the circuit increases exponentially towards its peak value (Vc/R) and the di/dt falls, this voltage transfers back to the resistance in the circuit, including the control resistance. This gives rise to the increase in anode voltage in region A. However, this situation does not persist, because the device impedance starts to fall again (Region B). Whether this fall is due to modulation of the control resistors or turn-on of the auxiliary stage cannot be determined; however, the start of turn-on in the main cathode can be clearly seen (Region C). The phenomenon of resistor modulation is discussed in the next section.

c) 100mm Device Tests.

Modulation of Control Resistors.

The values of the control resistors (Rc1 and Rc2) on the 100mm devices were measured. These resistors exhibited modulation, whereby the resistance falls significantly when turn-on occurs in the adjacent emitter [5.5]; Figure 5.7.4 shows a typical V:I characteristic. Figure 5.7.5 shows how electrons are injected from the emitter into the adjacent p-base, causing its bulk resistivity to fall dramatically. In this case, the slope resistance of Rc1 fell from 20R to only 1R above 250mA. This modulation would not affect the resistors if they were used for dv/dt compensation, since turn-on should not occur in the adjacent n-emitter, and the resistors would therefore stay in the high slope

resistance regime. However, modulation would seriously affect the protection offered by the resistors to the gate during turn-on.

Test Circuit.

The apparently excellent turn-on capability of the 30mm devices may have been due to the fact that they were not operating in a truly representative circuit. In particular, the time constant of the current decay was only 2us, and the transition from true off-state to true on-state is known to take longer than this. In order to test the 100mm samples in their proper electrical environment, the thyristor-level circuit from the forward recovery test equipment was modified and combined with a high-voltage supply originally designed for VBO testing; see Figure 5.7.6. The most significant modification to the circuit was the addition of 0.5uF stray capacitance across the "terminals" of the thyristor-level. An HVDC valve with 100 thyristor levels might typically have a stray capacitance of 5nF, and when referred to each thyristor this is equivalent to 0.5uF across each level (cf Figure 2.2.4). For completeness, the 400uH inductor and 6uF capacitor were used to provide follow-through current after the initial turn-on transient; however, the di/dt of the follow-through current in an HVDC valve is so low (eg 4A/us) as to impose negligible stress on the thyristors.

In order to measure the current flowing through the optical well, the well was isolated by breaking the contact with the three radial spokes, and then re-making the contact through an external bond passing through a clip-on current probe; see Figure 5.7.7.

Test Results.

For each test, the DUT was triggered at 50Hz with a high degree of optical overdrive, and the dc supply increased until the required turn-on voltage was achieved. With the oscilloscope on storage, the level of overdrive was then reduced until the device did not trigger. This gave an envelope of traces demonstrating the variation of delay time with overdrive, as shown in Figure 5.7.8. It can be seen that the delay times near marginal triggering become very long. For operation above 1kV, the triggering frequency of the circuit was reduced to keep dissipation in the current limiting resistors to an acceptable level.

The current measured through the external bond (Figure 5.7.8) has several interesting features. Firstly, the shape of the current pulse is almost independent of the level of overdrive used, with the only variation being in turn-on delay time. Secondly, the point at which the latching current is reached can be clearly seen. After the optical pulse is delivered, current through the well builds up at a rate determined by the degree of overdrive. When a certain threshold (approximately 500mA) is reached, the di/dt switches to a very high value. It is interesting to note that this current is greater than the measured threshold current for steady-state triggering of the auxiliary stage (typically 150mA).

Thirdly, as the turn-on voltage was increased to levels higher than that shown in Figure 5.7.8, the effect of the parasitic zener diode became apparent (cf Figure 4.7.4). The measured current developed a "plateau", above which the extra optical well current flowed through the internal zener diode, and was therefore not measured.

This device failed at the first turn-on from 1700V with 30x optical

overdrive. The failure site is shown in Figure 5.7.9, and the cause is clearly overheating of the parasitic zener diode. The high level of optical overdrive used suggests that the failure was unrelated to marginal triggering; the device had previously withstood repetitive marginal triggering from 1500V. Significantly, this failure site is the same as was observed during previous testing of the 75mm samples (section 4.7).

Testing with an External Zener Diode.

It was clear that the internal zener diode represented an inherent weakness in the prototype design, which would have prevented turn-on testing from above 1700V. To circumvent this, an external zener diode with a lower clamp voltage was connected in parallel with the internal one; see Figure 5.7.10. The internal diode had a clamp voltage of 14V, and so an 11V zener diode was used outside the device. It was recognised that this technique would introduce additional stress on the optical well for two reasons: firstly, it would divert current away from the auxiliary amplifiers, reducing their level of overdrive and therefore slowing down their turn-on, and secondly, Rc1 is bypassed, so that the slope resistance through the external zener would only be Rc2, which has a lower value (approximately 0.5R when modulated).

An example of the test results with the external zener diode is given in Figure 5.7.11. The turn-on voltage was increased in 250V steps to 2750V, at which voltage the device survived several threshold triggering events. However, when the voltage was increased with substantial optical overdrive being used, the optical well failed at turn-on from 2800V; the failure site was directly beneath the optical fibre. The peak current through the optical well was measured as 50A just prior to

failure, of which only 5A was flowing through the auxiliary thyristors.

To test turn-on from higher voltages, the surface of the junction forming the internal zener diode was plasma-etched. Doping concentrations are high at the surface, giving a low clamp voltage for the zener diode; etching down to a lower concentration region would increase the clamp voltage, forcing more current to flow through the auxiliaries and Rc1. The junction was etched to a depth of 32um (Figure 5.7.12), and the clamp voltage was found to increase from 14V to 18V. For these tests, no external zener diode was used, so as not to draw any current away from the auxiliaries.

The test results are shown in Figure 5.7.13, where the current plateau can be clearly seen. The device survived several threshold triggering events from 2750V, but as the voltage was increased towards 3000V with 30x optical overdrive, the device failed. The failure site was under the optical fibre, as before. From the shape of the current pulse flowing through the external bond, the peak current in the well was estimated to be similar to the previous failure threshold, ie 50A, after allowing for the current flowing in the internal zener diode.

One unexpected feature of the current pulses through the well was their duration, with it taking approximately 6us for the current to fall to 10% of its peak value. This probably arises from the low value of the control resistors after modulation; higher values of resistance in series with the optical well would tend to force the current to commutate out into the auxiliaries and main cathode. This was explored further with computer modelling, described in Chapter Six.

d) Conclusions.

From these test results, it was clear that any structures which included both control resistors and parasitic zener diodes would almost certainly be unsuccessful. For the control resistors to effectively limit the current flowing in the optical well, they would have to support (transiently) several hundreds or even thousands of volts, and no surface diode could support these voltages. The structure would therefore have to be concentric, with the transient voltage being dropped along the surface of the control resistors over a distance of several hundred microns. This was the final reason for abandoning dv/dt compensation, which requires non-concentric structures incorporating surface diodes. However, the encouraging part of these results was that the optical well design could withstand a 50A peak current pulse with only marginal triggering. Since a conventional single amplifying gate thyristor only requires an electrical gate current pulse of approximately 1A to survive worst-case turn-on, it was clear that the optical well had more than adequate capability to deliver a suitable gate pulse to a second amplifying stage. The only potential difficulty was the need to build in large enough control resistors to prevent the current capability of the well being exceeded.

5.8 THEORETICAL DISCUSSION OF OPTICAL TRIGGERING.

a) Introduction.

This section reviews some of the conclusions reached from the experimental work carried out so far on optical triggering. These conclusions then form an input to the design of the next generation of test structures described in Chapter Six.

b) Safe Optical Turn-on.

Reasons for Safe Optical Turn-on.

During testing of the early 75mm devices (Section 4.7), no conventional turn-on failures were observed. Normal di/dt failure on an electrically-triggered thyristor occurs at a site on the inside edge of the main cathode or an amplifying gate. However, all the failure sites on the 75mm prototype samples were located on one of the parasitic zener diodes in the gate structure, even when threshold optical triggering was used. Only when this type of failure was prevented on the 100mm devices by the use of an external zener diode was a more normal type of failure observed, ie melting of the inside edge of the optical well itself. However, even in these cases the failure threshold showed no dependence on the level of optical overdrive used, and seemed to be related only to turn-on voltage (and therefore inrush current). It is interesting to consider this in more detail.

Provided that the inrush is within the ultimate capability of the device, di/dt failure normally occurs when, as a result of a weak gate pulse, too small an area of the thyristor is turned on initially. Inhomogeneities in the device mean that gate current does not flow

evenly in the structure, but that a small area is turned on preferentially. Excessive current density during turn-on then melts this area. For optical triggering, one would therefore expect that a "weak" optical pulse would only turn on a small area of the optical gate, resulting in di/dt failure at high voltages. However, this was not observed, and the reason for this is attributed to the relationship between the latching current of the optical well and its horizontal geometry.

Under conditions of marginal triggering from quite low voltage (400V), approximately 500mA flowed through the well before the high di/dt part of turn-on occurred. The fact that current is flowing through the well, and not just under it, shows that the pilot emitter pn junction is already turned on, but the magnitude and duration of this current (over 30x the calculated threshold triggering current, for at least 3us) means that a large area of the well will be conducting. This would result in an excellent di/dt capability at turn-on, even for marginal triggering, as was observed.

The practical consequence of this is that for an optical well with a small inside diameter, the ultimate inrush capability of the well is always achieved whether marginal triggering or heavy overdrive is used. For a conventional thyristor, the inside diameter of the gate (typically 4mm) is large compared to the spreading that can be achieved during the turn-on delay time. If marginal triggering is used, one point on the inside edge will start to turn on, but even though some spreading will occur, by the time the latching current is reached this will still occupy only a relatively small proportion of the total gate edge. Consequently, the inrush capability of the gate will be much less than could be achieved if heavy overdrive was used to turn on the whole

gate edge. However, for an optical well with a small inside diameter (perhaps 200um), the spreading that occurs during the turn-on delay time may result in the whole inside edge of the gate conducting by the time latching occurs. The inrush capability would therefore be the same as if a substantial gate pulse had been used. As the optical drive is reduced, the area of the pilot thyristor turned on initially will shrink; however, the turn-on delay time will increase correspondingly, thereby allowing more time for spreading.

In addition to the above effect, turn-on in a small gate will be inherently more uniform anyway, since the effect of spatial inhomogeneities will be correspondingly reduced.

In conclusion, the fact that the thyristors had an inrush capability independent of the level of gate drive used was therefore not due to the optical triggering itself, but simply to the small dimensions of the gate that are permitted (and required) by optical triggering.

Implications for Control Resistor Design.

A further consequence of the above phenomenon for design philosophy is that control resistors are not needed to mitigate the effects of weak optical triggering. They are only required to ensure that the bulk energy absorption capability of the optical well is not exceeded. It has previously been suggested that calculation of the value of control resistors [5.6] should assume filamentary turn-on in accordance with the conventional understanding of weak triggering phenomena. On the basis of the discussion above, however, this may well result in unnecessarily large values of control resistance, with the consequent loss in cathode area.

Implications for Performance of Auxiliary Amplifying Gate.

It is interesting to consider that, since the latching current of the optical well is larger than the threshold triggering current for the auxiliary thyristor, the auxiliary amplifying gate will start its turn-on process before the optical well reaches its latching current. The auxiliary will therefore be "primed" before turn-on occurs in the optical well, thereby accelerating turn-on of the auxiliary amplifier and reducing stress on the optical well. This effect will probably be more pronounced with weak triggering, where the delay time is long and the auxiliary sees a current above its threshold level for some time.

As long as the optical well is more sensitive than the auxiliary amplifier, there is no danger of the auxiliary turning on with weak triggering before the optical well switches. However, it is possible that the performance of the device could be improved by deliberately giving the optical well a higher latching current, so that the degree of "priming" is increased. Alternatively, the same benefit could be achieved by making the auxiliary more sensitive, although the dv/dt capability of the auxiliary must still remain higher than that of the optical well.

Implications for Light Firing System Design.

If threshold triggering can be safely employed without fear of stressing local regions of the optical well, there is potential for considerable cost savings in the valve light firing system. Despite the many advances in optical systems, practical light triggering systems are still near the limits of technical feasibility, and so any relaxation in requirements can be translated into useful cost reduc-

tion. A further benefit is that monitoring and protection systems can be simplified; the light firing system described in Chapter Seven was designed on the assumption that the thyristors could be destroyed by weak optical triggering, and a significant level of complexity was incurred to minimise the possibility of this happening. However, if the results of a weak pulse are not catastrophic, it is acceptable to monitor the system and take action once weak pulses have been discovered, which is much easier to implement. This is discussed further in Chapter Seven.

Since it is no longer necessary to provide optical overdrive for high voltage turn-on, the optical drive used will probably be determined by the minimum voltage at which it is necessary to be able to trigger the thyristors. The drive necessary for threshold turn-on at 30V might well give at least 5x overdrive at peak voltage, which would ensure uniform turn-on delays within a series string of devices.

Implication for Metallization in the Well.

The experimental results and their discussion above corroborate the findings concerning metallization in the well in Section 5.2. In that section it was noted that the satellite wells, which had metallization extending onto region B of Figure 4.6.4, were 20 - 30% more sensitive to optical turn-on than the optical wells, where the metallization was restricted to region C. It was concluded that a "significant" level of current flows through the well before triggering occurs, and the extra lateral resistance in the pilot emitter in the optical well resulted in poorer sensitivity (Figure 5.2.3). This "significant" level of current has been shown to be the latching current for the well. For the next generation of optical devices, the metallization was therefore brought

down into the optical well.

One consequence of this phenomenon is that the control resistors necessary to protect the optical well during turn-on will adversely affect its sensitivity.

c) Continuous Gating.

As explained in Section 5.2, the light triggering system attempts to simulate continuous gating of the thyristors by sending a train of optical pulses. In the experiment described in that section, a train of sub-threshold pulses was sent to the device, and the frequency increased until the device triggered. The frequencies required were found to vary between 10 and 20kHz. This section attempts to relate this experiment to the actual conditions of interest, and the requirement is that when exposed to a train of pulses, the thyristor behaves as a diode. Should the voltage across it become positive, it will immediately start to conduct (although the delay time may be quite long). Prior to this, the device will be reverse biased, perhaps by approximately 10V.

The first important aspect of the above requirements is that [illegible]evice is reverse biased when the train of pulses is being sent. Th[illegible] that the central junction in the device will be for[illegible]ed, [illegible] will have a very narrow depletion layer. The [illegible] blocking junction will be on the remote side of the dev[illegible] result that the efficiency of conversion of photons into ca[illegible] be poor.

In contrast, the devices were forward biased d[illegible] with the result that the optical pulses would genera[illegible]ca[illegible]

down into the optical well.

One consequence of this phenomenon is that the control resistors necessary to protect the optical well during turn-on will adversely affect its sensitivity.

c) Continuous Gating.

As explained in Section 5.2, the light triggering system attempts to simulate continuous gating of the thyristors by sending a train of optical pulses. In the experiment described in that section, a train of sub-threshold pulses was sent to the device, and the frequency increased until the device triggered. The frequencies required were found to vary between 10 and 20kHz. This section attempts to relate this experiment to the actual conditions of interest, and the requirement is that when exposed to a train of pulses, the thyristor behaves as a diode. Should the voltage across it become positive, it will immediately start to conduct (although the delay time may be quite long). Prior to this, the device will be reverse biased, perhaps by approximately 10V.

The first important aspect of the above requirements is that the device is reverse biased when the train of pulses is being sent. This means that the central junction in the device will be forward biased, and so will have a very narrow depletion layer. The depleted reverse blocking junction will be on the remote side of the device, with the result that the efficiency of conversion of photons into carriers will be poor.

In contrast, the devices were forward biased during the experiment, with the result that the optical pulses would generate carriers within

the depletion layer, and these would constitute a leakage current which would then decay. If the next pulse arrived before the current from the previous pulse had decayed away, current would gradually build up until triggering occurred. The combination of pulse energy and frequency required to trigger the device is therefore an indication of the time taken for the photo-electric current pulse to decay; the smaller the optical pulse, the sooner the current in the device reaches negligible levels, and therefore the higher the frequency necessary to trigger it.

From the above discussion, it can be seen that the high-frequency pulsing technique will approximate to continuous gating when the device is forward biased. In the light triggering system described in Chapter Seven, each pulse in the pulse train would be of the same order as the low-voltage threshold sensitivity of the device. However, when the device is reverse biased, the reduction in photo-electric conversion efficiency may mean that when the voltage goes from negative to positive, the device may not trigger until the next optical pulse arrives. This may be 20us later (approximately 0.4 electrical degrees on a 50Hz system) which should be acceptable for the operating conditions concerned; if it is not, then the behaviour of the device under these conditions should be investigated in more detail.

5.9 SUMMARY OF RESULTS.

This section summarizes the important findings obtained in the evaluation of the prototype thyristor design in Stage Two.

a) Optical Triggering.

Various aspects of the optical triggering performance of the prototype design were investigated. As far as can be ascertained from published information, the optical sensitivity was at least as good as that obtained elsewhere, if not better. Other aspects of performance (such as delay time, effect of pulse length etc) were satisfactory and as expected.

b) Forward Recovery Protection.

Evidence of the correct operation of the forward recovery protection mechanism was only obtained on one device, and even in this case there was some doubt about its effectiveness. All other attempts to demonstrate operation of the protection were unsuccessful. It was therefore concluded that the technique for forward recovery protection, as embodied in the prototype design, was not practical. Other approaches would have to be tried in the next set of test structures.

c) Charge Extraction Tests.

Several problems were encountered in measuring the charge stored in the various areas of the device after recovery. However, enough results were obtained to demonstrate that it was possible, by selective electron irradiation, to significantly increase the amount of charge stored under a particular region of the gate. This supported the basic principle of the forward recovery protection technique, but clearly the

mechanism was not effective enough for the protection to work successfully.

d) Dv/dt Capability.

Apart from a minor flaw in the design of the prototype, the dv/dt capability was acceptable. In particular, the optical well design met or exceeded dv/dt design targets without the need for dv/dt compensation, which has significant benefits for device design.

e) VBO Protection.

Some interesting results were obtained from the investigations into VBO protection. However, other areas of work were more necessary for the success of the light-triggered self-protecting thyristor, and so these results were not pursued to a practical conclusion. For the time being, it was decided to rely upon a conventional BOD for overvoltage protection.

f) Turn-on Capability.

Investigations showed that the prototype gate design could not be damaged by weak triggering, unlike conventional thyristors. This was a very important result, since it allows considerable cost savings in the light triggering system. However, the gate design used to provide dv/dt compensation resulted in the devices having limited turn-on capability at high voltages. The control resistors also exhibited modulation, which reduced their effectiveness. Further work would therefore concentrate on designs with improved control resistors, and no dv/dt compensation. The reasons for the immunity to damage by weak triggering, and their implications, were considered. It is believed that the

phenomena discussed have not been previously reported in the literature.

g) Conclusions.

The evaluation of the prototype thyristor design yielded much useful information. Some aspects, such as optical turn-on and dv/dt capability, were successful. However, in the critical area of forward recovery protection, the tests were unsuccessful, so that considerable further work would be necessary to achieve a successful light-triggered self-protecting thyristor. Chapter Six describes the next program of development undertaken to achieve this objective.

CHAPTER SIX: DEVELOPMENT OF THE THYRISTOR, STAGE 3.

6.1 INTRODUCTION.

The evaluation of the prototype thyristor design revealed a number of areas where further work was required, most notably that of the forward recovery protection mechanism. A new set of structures was therefore designed, manufactured and tested with the objective of improving on those areas where the design was inadequate.

One major drawback with the prototype design was in the economics of its manufacture. The cost of producing a thyristor is approximately related to the number of photo-mask stages required in its manufacture, and whereas a conventional thyristor requires three mask stages, the prototype design required five. The increase in thyristor cost incurred by these extra two stages meant that even if the device worked successfully, it is doubtful whether it would be economically attractive. It was therefore decided that only those techniques needing four or less mask stages would be investigated.

For the investigation of forward recovery protection, the designs were based on small devices (30mm diameter) so that a large number could be produced. This would enable several different approaches to be tried. Any successful approaches could then be transferred to full-scale devices.

For work on optical triggering, some 30mm devices were produced to investigate "fine-tuning" of the optical well design, and a 56mm light-triggered thyristor was used to test turn-on in a representative HVDC circuit. The 56mm device incorporated the improvements in control

resistor design thought necessary to achieve the required turn-on performance. If successful, the design could be transferred without modification to 75mm and 100mm thyristors.

Two 100mm wafers were used to accommodate all the devices that needed investigation. The following structures were produced from each pair of wafers:

- Eight 30mm devices for investigating forward recovery protection techniques, including one control device;
- Two 30mm devices to give information for optimization of the optical well design;
- One 56mm light-triggered thyristor for investigating controlled turn-on.

The designs of the various devices are discussed in sections 6.2 and 6.3, and their testing is described in sections 6.4 and 6.5.

6.2 DESIGN OF THE FORWARD RECOVERY PROTECTION STRUCTURES.

a) Design Philosophies.

Two basic philosophies were proposed for providing forward recovery protection. The first was based around using gated turn-on to provide adequate protection against damage, and was therefore a derivative of the protection technique used in the prototype thyristor. The second approach allowed turn-on to occur randomly on the main cathode, but attempted to increase the area involved in turn-on to reduce the energy density. This second approach was based on proposals made by the GEC Engineering Research Centre at Stafford [6.1]. The principle of using gated turn-on to protect against forward recovery failure damage has already been explained with regard to the prototype design (sections 4.3 and 4.6). The second technique, that of increasing the area involved in turn-on, is explained below.

Turn-on in the main cathode during forward recovery failure occurs in much the same way as dv/dt turn-on (shown in Figure 4.2.9). When a voltage ramp is applied to the device, charge is swept out of the device through the cathode shorts, and so a potential drop is created under the n-emitter. If this potential becomes equal to the built-in voltage of the pn junction, turn-on will occur. From consideration of the cathode horizontal geometry for a hexagonal shorting pattern, it is clear that turn-on will take place in the centre of a triangle of shorts; see Figure 6.2.1. From the plot of potential drop between the p-base and n-emitter along the dashed line shown in Figure 6.2.1, it can be seen that under marginal conditions, turn-on tends to takes place in a small area. Furthermore, the current injected from the n-emitter is an exponential function of the bias voltage across the

junction (shown in the Figure), and this will exacerbate the tendency for turn-on to occur at a point.

Tests on the GEC HVDC valve design have shown that, for the type of thyristor being considered, ungated turn-on from up to approximately 1500V can be withstood without damage. Ideally, a self-protected device should withstand turn-on due to forward recovery failure from up to peak forward voltage, which is 5000V in this case. There is therefore a significant margin between the inherent capability of the device and the capability required of a self-protecting thyristor. However, if the area involved in ungated turn-on can be increased, then it is reasonable to suppose that the inherent capability of the device would be improved. Thus if the device can be made to survive ungated turn-on from 5000V, then "self-protection" will have been achieved through increased robustness. Even if the device can only survive ungated turn-on up to perhaps 3000V, then this may reduce the statistical probability of failure to a level where the thyristor can realistically be used without electronic protection.

b) Techniques for Achieving Increased Turn-on Area.

Two techniques were proposed to increase the area of the device involved in ungated turn-on. The first relies on using a high-conductivity region to "flatten out" the voltage profile in the region where turn-on occurs. If the conductivity of the p-base is increased in the centre of a group of shorts, then the voltage profile generated during a voltage ramp will be modified as shown in Figure 6.2.2. Although in the limiting case turn-on will still occur at a point, the area surrounding the point should be much closer to turn-on, so that initial spreading of the turn-on area will be faster. Consequently, the

device should be more resistant to damage.

The second technique proposed for increasing the turn-on area is the converse of the first, in that it relies on increasing the resistance in the turn-on zone. If the current in the n-emitter at turn-on is forced to flow horizontally through a fairly high resistance layer before it can reach the cathode contact, this will produce a lateral voltage gradient which will tend to increase the velocity of spreading; see Figure 6.2.3. This technique was originally put forward as an alternative to the amplifying gate, but was found to be less effective [6.2]. However, as with the first technique described above, it should help to increase the resistance to forward recovery damage.

To ensure that these special structures would always be close to the last area of the device to conduct, they should be distributed over the whole of the main cathode. The main cathode would therefore be made up of hexagons of shorting dots, with the special structures in the centre of each hexagon.

c) Implications for Other Protection Mechanisms.

If these techniques were shown to be successful in improving the resistance to forward recovery failure damage, then it is clear that they would also have the same effect for dv/dt failure damage, since the turn-on mechanism is similar. It is also possible that they might help prevent damage from forward voltage break-over on the main cathode, since here again this damage is caused by excessive energy density in the small turn-on zone. This would particularly be the case at high temperatures, where break-down is triggered by leakage currents rather than avalanche or punch-through.

d) Disadvantages of the Proposed Techniques.

The possible disadvantage of these techniques is the effect they might have on the forward voltage drop and off-state dv/dt withstand capability of the thyristor. Parts of the normal n-emitter are either lost or replaced by less efficient pn junctions, and the changes in the shorting pattern might reduce the dv/dt capability. To enable any deterioration in properties to be quantified, a control device with no special features was produced.

e) Design Details.

The following sections discuss the designs of the various forward recovery protection test structures.

"P+ Zone" Devices.

Three versions of the p+ zone design were produced. Starting from the normal hexagonal shorting pattern, the central shorting dot was removed from each hexagon and replaced by a region of thin n-emitter; see Figure 6.2.4. Since the n-emitter in this region was shallower than the normal diffusion, the p-base resistance underneath it would be lower (typically 25 ohms/square rather than 300 ohms/square).

The diameters of the p+ zones on the three versions were 1, 1.41 and 2mm respectively. This would enable not only the effectiveness of the zone to be evaluated, but also any implications for forward volt drop and dv/dt capability.

"P+ Grid" Devices.

An alternative to the p+ zone design was to have circles of normal n-

emitter around each shorting dot, buried in a "grid" of enhanced p-base conductivity; see Figure 6.2.5. This is the geometric inverse of the p+ zone design. The diameter of the normal n-emitter regions was set to 1mm, since this fitted with the existing hexagonal shorting pattern.

Lateral Field Devices.

As with the p+ zone devices, the special region for these structures was placed in the middle of the hexagons in the existing shorting pattern. Since the optical well on the prototype devices had been shown to have an excellent inrush capability under marginal triggering, this structure was reproduced in the special zones, as shown in Figure 6.2.6. The small secondary well was retained to encourage turn-on to occur around its circumference rather than at a point, and the metallization was brought part way down into the well to enable it to carry current during normal conduction. When turn-on occurred around the secondary well, the resistance of the thin emitter layer would set up a lateral field, encouraging rapid spreading of the turn-on area.

Selective Failure Zone Devices.

This design was a derivative of the prototype technique for forward recovery protection. However, in order to improve the likelihood of its working, the distance between the main emitter and the gate region was halved, to allow more charge to diffuse under the gate. It was believed that a single amplifying gate structure might provide an acceptable level of single-shot inrush capability if combined with an unmodulated control resistor. The use of only one amplifying gate would also allow the gate to be much closer to the main cathode, as shown in Figure 6.2.7.

The control resistor was designed to be nominally 25R, and a small anti-modulation ring was included next to the main cathode. An irradiation step with the selective failure zone shielded would be used, as for the prototype device. It was envisaged that for a 100mm device, there might be several of these small selective failure zones distributed around the main cathode.

Gated Turn-on Devices.

To avoid the need for selective irradiation, which increases the forward voltage drop of the main cathode, an alternative design was proposed which had a small amplifying gate at the centre of each hexagon in the shorting pattern; see Figure 6.2.8. This would ensure that one of these zones was always close to the last area to conduct, and its small size (2mm diameter) would mean that there would be a high level of charge diffused underneath it. The increased p-base resistivity under the thin emitter layer should make the gate more sensitive to turn-on than the main cathode under forward recovery failure conditions. Lastly, the lateral field generated in the thin emitter layer during turn-on would aid spreading as described above.

However, computer analysis of the voltage profiles generated by displacement current showed that, because of the horizontal geometry of the design, the amplifying gate could not be made more sensitive than the normal main emitter for any plausible design. Therefore, the emitter layer was omitted from the design, and the metallization alone was applied to collect displacement current and feed it uniformly to the inside edge of the surrounding main emitter; see Figure 6.2.9. By enforcing an equipotential region near to the turn-on edge, uniform turn-on should be encouraged.

6.3 DESIGN OF THE LIGHT TRIGGERING TEST STRUCTURES.

a) Introduction.

There were two main objectives in the field of light triggering: the first was to make a device that would withstand repetitive optical turn-on under the most severe conditions required (ie the 56mm device), and the second was to gain information that would help in the optimization of the optical well design (using the 30mm devices).

From the work carried out in Stage 2 of the development, it was considered that the inclusion of un-modulated control resistors would be vital for successful turn-on capability. However, it was not known what value of control resistance should be used, whether more than one amplifying gate was required, and if so, how the control resistance should be split between stages. Computer modelling of the turn-on process was therefore carried out to gain an appreciation of the behaviour of different gate designs. This computer modelling is described below, followed by an account of the design of the 56mm and 30mm devices.

b) Computer Modelling.

Finite element modelling of semiconductors is a very complex procedure, and so to obtain results in a realistic timescale, it was decided to represent the thyristor as a simple assembly of passive components, and analyse the interaction of this with the external circuit. The circuit analysis programme used to carry out the computer studies is known as VTECAP (Variable Topology Electrical Circuit Analysis Program). This proprietary program is described elsewhere [6.3], but its most useful feature in this context is that it allows switches and non-linear

elements to be controlled by a variety of parameters.

Modelling Technique.

Each thyristor in the device (ie the amplifying gate(s) and the main cathode) was modelled by a variable resistance in series with a switch. Turn-on was represented by closing the switch and then reducing the resistance in discrete steps, to simulate the collapse of impedance in a real device at turn-on. The rate of collapse of resistance was derived from voltage and current waveforms observed during turn-on of a conventional electrically-triggered thyristor. Although the thyristor observed had an amplifying gate, and therefore the waveforms were for two-stage turn-on rather than for a single structure, it was considered that the model would be a sufficiently good approximation for the purposes of this study. The resistance:time characteristic used in the model is given in Figure 6.3.1.

The basic circuit studied is shown in Figure 6.3.2. The non-linear inrush limiting inductors were represented with a piece-wise linear approximation, and the thyristor-level damping and grading circuits were included, in addition to stray inductance and capacitance.

The turn-on sequence was as follows. Firstly, the circuit capacitances were charged to the peak turn-on voltage (5300V in this case), and the switch for the optical well thyristor closed. The variable resistance for this stage then started to clock through its states on a time-controlled basis. When the current through the device reached the predetermined threshold level for the next stage, a counter was started which represented the delay time for that stage. When the delay expired, the switch for this stage was closed, and its variable

resistance started to clock through its states.

The delay time of all the thyristors was set to 0.55us, which was felt to be a reasonable upper limit for the delay times at high voltage and high temperature; if the delay times were shorter than this in practice, the duty on the optical well would be reduced. The optical well was considered to be the most vulnerable part of the system to turn-on failure, and therefore the studies were designed to produce "safe" design criteria for it. For this reason, there was no attempt to model the effect of overdrive on reducing the delay times of the subsequent stages. Note that although delay times for optical triggering can be shorter than for electrical triggering [6.4], the subsequent stages of the thyristor are electrically triggered by the current in the optical well, and so will have delay times similar to conventional devices. The computer studies were only concerned with the period after the optical well had turned on.

In addition to the normal effect of overdrive on reducing delay time, the "priming" effect described in section 5.8 might reduce delay time further. Since a relatively large current is already flowing through the structure before the optical well turns on properly (greater than the steady-state triggering threshold for the auxiliary), turn-on of the next stage of the device would be accelerated. However, this effect was ignored since it would also serve to reduce the duty on the optical well, and the purpose of the studies was to produce "safe" results.

As already explained, the optical well was thought to be the most vulnerable part of the device to turn-on failure. If the turn-on capability of the subsequent stages was found to be inadequate, then

this could be improved by conventional means. Therefore, only the results affecting the optical well are considered in detail below.

Studies Carried Out.

Two main sets of studies were performed. The first was to compare the performance of gate structures having one, two and three amplifying stages for a given total control resistance. The second was to investigate the effects of varying the split of resistance between stages. A third study to investigate the effects of control resistor modulation was also undertaken.

Study 1: Different Numbers of Amplifiers.

An important consideration in the use of control resistors is the area of the device they take up, since this is lost to the main cathode. The first set of studies were therefore designed to show how, for a given area loss, the best turn-on performance can be achieved. In addition, it was felt that the total control resistance in the gate structure would be the single most important factor in protecting the optical well, and so by keeping the total value constant, the benefits (or otherwise) of additional amplifying stages could be studied in isolation. Published work [6.5] indicates that there are diminishing returns in applying extra amplifying gates to a device, and so the purpose of the studies was to confirm this.

The maximum turn-on voltage predicted for the present generation of 100mm thyristors is approximately 5.3kV, whilst the peak current capability of the optical well was found in Stage Two to be approximately 50A. Thus, if it is assumed that in the worst case all of this turn-on voltage appears across the control resistors (and none across

the transient impedance of the device itself), then a total resistance of 105R is required to limit the optical well current to 50A.

Three thyristors were studied, whose designs are given in Figures 6.3.3 to 6.3.5. For the treble amplifying gate case, the resistances were split so that each stage had approximately one quarter of the control resistance of the previous stage. Although the threshold triggering sensitivities were set to plausible values, in practice the di/dt through the optical well was so high that the time delay between the current crossing the various thresholds was of the order of one nanosecond, and therefore turn-on of subsequent stages was effectively simultaneous [6.5].

The results of these studies are presented in Table 6.3.1 and Figures 6.3.6 to 6.3.10. For Case 1 with a single amplifying stage (the optical well itself), the energy dissipated in the well was 18.3mJ. Adding another amplifying stage only reduced this to 17.9mJ, and a third stage resulted in 17.6mJ. The reason that the effect of adding amplifying stages is so small is that the model predicts that virtually all the energy that is dissipated in the optical well appears during the very early stages of turn-on. Figure 6.3.6 shows the voltage and current waveforms obtained for Case 1, whilst Figure 6.3.7 shows the corresponding current and energy waveforms for the optical well. Since the subsequent stages do not turn on until 0.55us into the study, they cannot affect the energy dissipated in the optical well early on; rather, they can only speed up the commutation of current out of the well, when dissipation is relatively low.

In practice, the resistance:time model for the optical well should probably be scaled upwards in its later sections, since the area turned

on will be smaller than in the thyristor from which the model was derived. There would therefore be more dissipation in the tail of the current waveform, and consequently more opportunity for additional amplifying stages to reduce the dissipation in the optical well. Also, if the delay times of subsequent stages were shorter than those modelled, they would switch in at a time when the dissipation in the well was higher, and therefore they could have more effect. Despite this, these results do confirm the published results, ie that adding extra amplifying stages is subject to diminishing returns. It is interesting to note that for Case 3, the second stage never carries significant current (Figure 6.3.10).

Another significant feature of the waveforms in Figure 6.3.6 is that, in spite of stray inductance in the circuit, the current in the optical well approaches its peak value very quickly (approximately 300ns). The transient impedance of the device itself therefore cannot be relied on to limit the peak current in the well. This has two implications: firstly, in order to be safe, the value of the control resistors must be equal to the peak turn-on voltage divided by the peak acceptable current, and secondly, large transient voltages can be expected to appear across the control resistors. The possibility of surface breakdown must therefore be considered.

Study 2: Two Amplifiers: Varying Distribution of Resistance.

The second set of studies was designed to investigate the effects of varying the split in control resistance between amplifying stages. Using the second study above as a reference (ie two amplifying stages, 100R and 5R control resistors), a further study was carried out for a double amplifying gate with 65R and 40R control resistors. Figure

6.3.11 shows the circuit used, and Figure 6.3.12 the waveforms obtained, with a comparison of the results being given in Table 6.3.2. As above, the changes had little effect on the energy dissipated in the pilot thyristor, with the energy dissipated falling from 17.9 to 17.7mJ for the device with the 65R/40R control resistors. This fall in energy occurred because the increase in resistance in series with the auxiliary amplifier forced current to commutate out of the optical well more quickly, which is evident from a comparison of the optical well currents in Figures 6.3.8 and 6.3.12. As might be expected, the most significant effect of changing the split in control resistance was on the energy dissipated in the auxiliary thyristor, which fell from 80.2 to 11.9mJ. However, since the area of the auxiliary thyristor would typically be several times that of the pilot thyristor, a "constant energy density" criterion would dictate that the energy dissipated in the auxiliary thyristor should be greater than that in the pilot thyristor by the same factor. In fact, this is much more closely approximated to by the first case, where the auxiliary energy is approximately 4.5 times greater than the pilot energy.

Study 3: Effect of Control Resistor Modulation.

The final study was to model, for purposes of comparison, a thyristor with heavily modulated control resistors. Values of 2R and 1R were therefore used, since these approximate to the modulated values of control resistance in the 100mm samples. The circuit is shown in Figure 6.3.13, and Table 6.3.3 compares the results with those obtained from a device with unmodulated control resistors. It can be seen that both current and energy in the pilot thyristor are greatly increased.

Conclusions from Computer Studies.

The important conclusions to emerge from these studies are therefore as follows:

- The inductance in the circuit and the transient impedance of the pilot thyristor itself cannot be relied on to limit its dissipation. The control resistors themselves must therefore be adequate to keep the maximum current to acceptable levels.
- Because of the time delay between turn-on of the pilot thyristor and turn-on of subsequent amplifying stages, the energy dissipation in the pilot thyristor can only be marginally reduced by the addition of amplifying stages.
- Since the current in the pilot thyristor increases rapidly when turn-on occurs, all subsequent stages become heavily overdriven within a few tens of nanoseconds. Turn-on of the subsequent stages is therefore virtually simultaneous, so that it can be argued that there is little merit in having auxiliary amplifying stages. There is certainly negligible benefit in having more than one.
- With conventional electrically-triggered thyristors, the challenge in gate design is to make the device turn on safely with as low a level of gate current as possible. With light-triggered thyristors, the problem appears to be reversed, and the challenge is to limit the current delivered by the optical gate to safe levels.

c) Design of the 56mm Device.

In accordance with the conclusions reached in Chapter 5, the 56mm device was designed with a shorted optical well (ie without dv/dt compensation). Using data from dv/dt withstand tests, the well was designed to have a dv/dt capability of 3kV/us at 110°C. Also, with

guidance from the results of the computer studies (although these were not complete when the design was fixed), a total control resistance of 100R was provided, with nominal values of 80R in series with the optical well and 20R in series with an amplifying gate.

In order to utilise the test technique developed during earlier turn-on testing, both the optical well emitter and the auxiliary emitter were isolated, so that an external bond was required to short them to the p-base. This would enable the current in each stage to be measured. Anti-modulation rings were provided for each control resistor, and for further investigations, these would allow the auxiliary thyristor to be by-passed and one or both of the control resistors to be shorted out. The anti-modulation ring for the second control resistor was a novel design which combined the anti-modulation function with the distribution of current to the edge of the main cathode. This enabled the control resistor to be annular, thereby using much less device area than if it had run parallel with the inside edge of the main cathode. The thyristor design is shown in Figures 6.3.14 and 6.3.15.

In order to prevent dv/dt failure at the corners of the inside edge of the main cathode, the corners were rounded, and much smaller shorting dots were used along the edge. This meant that the ratio of short diameter to corner radius was improved, so that the discontinuity of the horizontal geometry at the corners was reduced; see Figure 6.3.16.

d) Design of the 30mm Devices.

In order to investigate the sensitivity and inrush capability of different designs of optical well, two 30mm devices were manufactured containing a total of six different well designs. The well designs were

based on that used in the prototype device, shown in Figure 4.6.4. In all cases, the outside diameter of the normal n-emitter (Region C) was fixed.

Design #1: Different Secondary Well Diameters.

In the first set of three wells, only the diameter of the secondary well (Region A) was varied, as shown in Figure 6.3.17. This parameter would be expected to affect optical sensitivity, inrush capability and, to a limited extent, dv/dt capability. The first case with no secondary well (ie zero diameter) would correspond to the structure previously discussed, which makes use of the photocurrent generated in the top pn junction to help turn-on of the device (refer to Section 5.2). If it was shown that light transmitted through the thin n-emitter is overall less efficient than that transmitted through a secondary well, it follows that there would be a trade-off between peak optical sensitivity (achieved with the secondary well being the same diameter as the optical fibre) and the sensitivity that can practically be achieved with realistic alignment accuracies; see Figure 6.3.19. By varying the length of the turn-on line, the diameter of the secondary well might have an influence on the inrush capability of the gate.

Design #2: Different Primary Well Diameters.

For the second set of three wells, the outside radius of the primary well (Region B) was varied in order to investigate the trade-off between optical sensitivity and dv/dt capability; see Figure 6.3.18. The secondary well diameter (Region A) was constant. Since one design was the same as a well on the #1 device, the metallization was brought down into this well to check the effect on sensitivity and inrush capability. The wells were designed with theoretical dv/dt capabilities

of 3, 4.5 and 5.8 kV/us.

The outline of one of the complete 30mm devices is shown in Figure 6.3.20. There is no main cathode on the device, and the wells are not shorted so that all the wells are isolated initially. The well being investigated can then be "activated" by connecting an external wire bond between the emitter of the well and the metallization on the surrounding p-base. Between the central island on the device (containing the three wells) and the peripheral ring of metallization is a control resistor; the purpose of this was to try to ensure that dv/dt tests would not be destructive.

6.4 TESTING OF THE FORWARD RECOVERY PROTECTION STRUCTURES.

a) Introduction.

Testing of the forward recovery protection structures took place both at MEDL in Lincoln and the GEC Engineering Research Centre (ERC) in Stafford. Because of the relatively large number of different designs to be tested, it was decided to subject two samples of each design to a standard "screening" test to see which, if any, showed promising signs. More detailed evaluation could then be carried out as appropriate.

The test circuits used at Lincoln and Stafford were the same as that shown in Figure 5.3.1. At MEDL, optical triggering of devices was used, whereas at ERC electrical triggering was employed, with a timed relay being used to open-circuit the gate after the main current pulse had started. The standard test was a 350V/us (MEDL) or 500V/us (ERC) ramp, 2kV peak, 90°C. This ramp was initially applied soon after current zero, and then the time of application put back until the marginal forward recovery failure region was encountered. The delay was then gradually increased (eg in 5us steps), and the voltage at turn-on measured. The test was stopped either when the device recovered completely, or was damaged.

b) Test Results.

Typical forward recovery failure waveforms are presented in Figure 6.4.1. It soon became apparent that none of the devices showed any improvement in resistance to forward recovery failure damage compared to a conventional unprotected thyristor; the ungated turn-on voltages at which damage occurred are shown in Table 6.4.1, and these all lie within the range that would be expected from conventional devices.

However, the Selective Failure Zone samples showed clear signs that the device was being gated into conduction during forward recovery failure, and that the devices failed because the inrush capability of the gate was insufficient. This technique therefore held out the prospect that, with further development, it could form the basis of a commercial technique for forward recovery protection.

The failure mechanisms for the various designs are discussed below.

"P+ zone" Devices.

For each version of the p+ zone design, damage occurred on the edge of one of the zones, as shown in Figure 6.4.2. This was not the location that was expected; Figure 6.2.2 shows that turn-on was predicted to occur in the centre of the p+ zone. Since the failure did in fact occur under part of the normal emitter, the structure would therefore be expected to have the same resistance to forward recovery failure damage as a normal device. In effect, the "robustness" of the p+ zone has not been tested; all that has happened is that the location of the failure site has been moved to another part of the normal emitter.

Two reasons were put forward to explain this. The first was that because the doping profile was different in the p+ zone, the built-in potential of the pn junction in this region would be different. The thin emitter layer is less efficient as an emitter than the normal cathode, and so it requires a higher bias to cause carrier injection and turn-on. If the built-in potential was significantly higher than that in the normal emitter, then turn-on might never occur in the p+ zone; see Figure 6.4.3. This was recognised at the design stage, but because of the difficulty in predicting the built-in potential of

a diffused junction (particularly an unusual one such as that formed by the thin emitter layer), it was considered that experimental results would be required to prove whether or not the mechanism would work.

The second possible reason why the p+ zone did not turn on is related to the vertical geometry of the device. The pn junction formed by the thin emitter layer is approximately 60um further away from the forward blocking junction than the normal emitter (Figure 6.2.4). Thus even if the centre of the p+ zone is emitting carriers more strongly than the edge of the normal emitter, it is possible that the carriers from the normal emitter might reach the forward blocking junction first, causing turn-on in that area. Turn-on therefore might always occur first under the normal emitter, preventing the p+ zone from fulfilling its protective function. The p+ zone could only operate when the ramp conditions were insufficient to cause turn-on of the normal emitter, in which case the protection would not be needed anyway.

Whatever the reason was for the failure of the p+ zones to turn on, it was clear that the principle behind that technique for protection (ie increasing the area involved in turn-on) had not been tested. However, any other implementation of this principle would incur more cost in manufacture of the device, and so these tests at least showed that the simplest method of implementation (via the thin emitter layer) was unsuccessful.

Attempts were made to force turn-on to occur in the p+ zone by irradiating all of the device except one zone. However, turn-on still occurred under the normal emitter, even for very high radiation doses.

"P+ grid" Devices.

These devices failed on the edge of one of the regions of normal n-emitter, in a similar manner to the p+ zone devices; see Figure 6.4.4. The reasons advanced for failure are the same as for the p+ zone devices.

Lateral Field Devices.

The lateral field devices failed in the region of thin emitter layer between the secondary well and the metallization in the primary well; see Figure 6.4.5. If the lateral field mechanism was working in these cases, it was clearly insufficient to provide any measurable improvement in the resistance to forward recovery failure damage.

Selective Failure Zone Devices.

Four devices were tested, the first of which had received a radiation dose of 80kRad, with a mask that extended to half way between the edge of the gate and the inside edge of the main cathode (similar to the #4 mask in Figure 5.3.12). In this case, a layer of silicon rubber was applied to the control resistor to prevent surface breakdown. During forward recovery failure testing, a large voltage was observed between the gate and the main cathode, as shown in Figure 6.4.6. In the final shot which destroyed the device, approximately 530V appeared between gate and cathode for turn-on from 1400V.

Subsequent examination of the device showed a failure track across the control resistor, as shown in Figure 6.4.7. The voltage measured would have given rise to a field of approximately 1.7kV/mm, and so the failure could have been due to breakdown of the resistor surface. Alternatively, it is possible that the resistor could have gone into

thermal runaway, since its volume was small, and a temperature rise of perhaps 20 to 30°C could have been sufficient to cause runaway.

In order to prevent similar failures, the control resistor was shorted out on subsequent devices, although this would still leave a region of resistor which could be modulated by turn-on of the main cathode edge. Three more devices were tested, with doses of 80, 40 and 20kRad, and in all three cases significant voltages were generated between gate and cathode during forward recovery failure; see Table 6.4.2. This is particularly encouraging in the case of the device with only 20kRad dose, since this suggests that it might be possible to achieve forward recovery protection without any increase in radiation dose (and therefore forward volt drop) over that already used for lifetime control.

All three devices failed on the inside edge of the main cathode, but it was difficult to tell whether this was caused by breakdown of the remaining resistor in the gate, or inrush failure on the cathode edge. However, the fact that a positive voltage was produced in the gate still suggests that the gate was turning the device on, rather than that the device was suffering normal forward recovery failure on the inside edge of the main cathode. This was confirmed by tests on a 56mm device described later in this section.

Gated Turn-on Devices.

The gated turn-on devices suffered damage on the inside edge of the main cathode at one of the special sites, as shown in Figure 6.4.8. Even though the metallization in the well may have helped to enforce uniform potential along this edge, the effect was clearly insufficient

to prevent weak turn-on occurring at a point location, with a correspondingly limited energy absorption capability.

56mm Device.

One 56mm device (#13) was subjected to forward recovery testing, having received an 80kRad radiation dose with a mask over the optical well. The purpose of this test was to see whether the small amount of current that would flow through the optical well during normal conduction (due to current sharing with the main cathode) might cause the well to trigger during a forward recovery failure event. Tests were carried out at 90°C, as before.

In practice, the well did not trigger, and the device was damaged on the main cathode between shorting dots. However, two interesting features were noted during the test, the first of which was that it was very difficult to induce marginal forward recovery failure. For previous tests, the "vulnerable window" during which marginal failure occurred was several tens of microseconds long, and so it was straightforward to induce marginal failure events of increasing severity by advancing the point of ramp application through the vulnerable window. However, for the 56mm device the vulnerable window appeared to be only one or two microseconds wide, thus making it difficult to find and explore. The first shot that was successfully placed within the window was very severe, and damaged the device; see Figure 6.4. 9. This oscillogram shows the last two shots applied to the device, for the first of which the device recovered, but for the second (approximately 0.2us earlier) the device switched from 2350V.

The second interesting feature of these tests was the capacitive coupling that could be observed between the optical well and the main

cathode. The voltage between the pilot thyristor (ie the optical well) and the main cathode is also shown in Figure 6.4.9, and it can be seen that when the ramp is applied, there is a positive voltage spike between pilot and cathode. This arises because the optical well is coupled to the main cathode through the control resistors (in this case, a total of 414R at 22°C), and so the voltage developed across the junction capacitance of the well will lag behind that developed across the main cathode capacitance. This should help the dv/dt performance of the optical well, since the source impedance of the ramp generator has effectively been increased, so that the front end of a dv/dt ramp will be smoothed out and the peak current drawn from under the well reduced. This effect will not normally be so pronounced, since in practice the control resistors would be of a lower value than was the case with this sample.

The effect just described supports the previous conclusion that the Selective Failure Zone devices were operating correctly. Where the damage site for these devices was on the edge of the main cathode, it was difficult to tell whether this was normal forward recovery damage or an inrush failure caused by insufficient di/dt capability of the gate structure. However, if turn-on had occurred first on the inside edge of the main cathode (ie normal forward recovery damage), then there should have been a negative voltage spike between the gate and cathode similar to that in Figure 6.4.9, whereas in fact a positive voltage was observed. This reinforces the conclusion drawn previously, ie that the devices were being gated into conduction during forward recovery failure, as intended.

c) Conclusions.

None of the test structures showed an increased ability to withstand forward recovery failure damage compared to conventional devices. For the p+ zone and grid structures, turn-on failed to occur in the required region, possibly due to the diffusion profiles or vertical geometries used. In other structures, the protection strategy failed to produce any measurable effect on the resistance to damage. However, the Selective Failure Zone devices showed correct operation with radiation doses as low as 20kRad. Unfortunately, it appeared that the design of the gate on these devices was unable to cope with the single-shot turn-on conditions required. Having established that the principle works, development of the gate structure is now needed to improve its inrush capability.

6.5 TESTING OF THE LIGHT TRIGGERING STRUCTURES.

a) Introduction.

Repetitive turn-on testing of the 56mm devices was carried out at the Engineering Research Centre in Stafford on a high-power test circuit capable of producing 50Hz turn-on at voltages up to 5kV. Single-shot turn-on testing of the 56mm devices and evaluation of the 30mm structures was carried out at MEDL. The laser light sources used for all the tests were the laboratory light firing systems described in Chapter Seven.

The 56mm devices were tested both for optical and dv/dt turn-on capability. A useful spread of results was obtained because the samples had a wide range of control resistor values; this had been achieved by varying the depths of the two chemical etch stages used in production.

b) 56mm Devices: Repetitive Optical Turn-on Testing.

The turn-on test circuit was similar to that shown in Figure 5.7.6, except that a resonant energy recovery circuit was used to reduce losses in the circuit. Device temperature was measured with a thermo-couple in contact with the remote side of the molybdenum anode washer.

A summary of test results on six devices is given in Table 6.5.1. Two failure modes were encountered: the first was surface breakdown of the control resistors, and the second, inrush failure on the inside edge of the main cathode. The #0 device used for setting up the test circuit suffered accidental mechanical damage during modifications to the circuit, and one device failed due to voltage breakdown.

The laser drive circuit was supplied from a variable voltage source, so that the optical drive could be varied in a controlled fashion during the tests. The usable control range was from 80V to 200V supply voltage, and the laser output increased by 30 times over this range.

At the start of the tests on each device, the thyristor was repetitively triggered from 500V anode-cathode voltage, and with the laser supply set at 80V, the optical attenuation in the circuit was increased until threshold triggering occurred. This level of optical drive was then used as the "threshold" level (1x overdrive) for the rest of the tests. Even though it was known that the true threshold level would fall with increasing voltage, it was decided for two reasons not to try and find the threshold for each voltage level: firstly, it was not known whether the devices would survive threshold triggering from high voltage, and so this could have been a hazardous procedure, and secondly, any practical light triggering system would operate with constant optical drive, so that the tests would resemble real-life conditions.

A substantial volume of information on turn-on behaviour, sensitivities and delay times was collected during these tests. Tests on each device are summarized below, and the most interesting results are discussed.

Test 1: Device= #0 *Rc1=* 76R *Rc2=* 18R *Total=* 94R

Waveforms for 1.5kV turn-on at room temperature for this thyristor are given in Figure 6.5.1. Two sets of traces are shown, for 1x and 30x overdrive, and the effect of overdrive on delay time and the rate of build-up of anode current before turn-on can be clearly seen.

Test 2: Device= #43 *Rc1=* 47R *Rc2=* 10R *Total=* 57R

A set of waveforms at 30°C showing anode-cathode and pilot-cathode voltage is given in Figure 6.5.2. The pilot-cathode voltage was measured between the metallization on the optical well (the "pilot" thyristor) and the main cathode; it therefore represents the voltage dropped across the control resistors. Two features are of interest in this oscillogram:

- The effect of initial anode voltage on the speed of the turn-on process can be clearly seen.
- The peak pilot voltage reached is approximately 53% of the anode voltage before turn-on, showing that the control resistors are playing an important role in limiting current during turn-on.

Figure 6.5.3 shows the form of the anode current when the optical drive is just below and just above threshold for 500V turn-on. The positive feed-back mechanism inherent in thyristor triggering can be seen in both the build-up of current (successful triggering) and the decay of current (failure to trigger).

This thyristor failed as the turn-on voltage was being raised from 3000V to 3500V. Tracking had taken place across both control resistors, as shown in Figure 6.5.4. It is assumed that the tracking across Rc2 occurred as a result of the breakdown of the more highly stressed Rc1. Examination of the failed device showed a large number of incomplete tracks emanating from the inner edge of Rc1, which were reminiscent of the incomplete failure tracks sometimes seen in solid insulators. These are shown in Figure 6.5.5. The existence of these tracks suggests two things: firstly, that the breakdown process is gradual, and secondly (as a consequence), that the failure was due to surface breakdown and

not thermal runaway. Surface breakdown is a stress-relieving mechanism which can be "self-healing" up to a point; one would therefore expect that some breakdowns could occur without being catastrophic. In contrast, thermal runaway is a self-destructive mechanism, and the first time it occurred would be catastrophic. Subsequent examination of the #0 device showed similar surface tracking, and since this thyristor had not failed on test, it supports the above conclusions.

The possibility of this type of failure was predicted as a result of the computer studies (section 6.3), which showed that a large proportion of the anode voltage would appear across the control resistors. To prevent surface breakdown on future devices, the control resistors were therefore covered with a layer of silicone rubber, which successfully eliminated the problem.

Test 3: Device= #11 *Rc1=* 485R *Rc2=* 40R *Total=* 525R

This was a very high resistance device, and as a result it was possible to observe distinct turn-on of all three thyristor stages in the anode current waveform when triggered from low voltage; see Figure 6.5.6. Also as a consequence of its high resistance, the pilot-cathode voltage reached 90% of the initial anode voltage during turn-on. Waveforms for 3kV turn-on are shown in Figure 6.5.7

The thyristor failed as the turn-on voltage was being increased from 3kV to 3.5kV, with the failure site being on the inside edge of the main cathode near the root of one of the gate fingers. This is normal for di/dt failure, and shows that the performance of the auxiliary amplifying gate is reduced to an unacceptable extent by including a 40R series resistor (Rc2).

Test 4: Device= #46 *Rc1=* 110R *Rc2=* 25R *Total=* 135R

This device survived repetitive turn-on for ten seconds from 5kV at 30°C, and 4.7kV at 105°C; ten seconds is the longest period that the thyristor would be required to turn on from maximum voltage. Testing was not continued to 5kV at the higher temperature because part of the thermal insulation around the device started to break down electrically.

Figure 6.5.8 shows turn-on from 5kV at 30°C. In this figure, the change in device resistance during turn-on can be clearly seen, since initially the anode current rises at 115A/us, but later shifts to a 200A/us characteristic. Since there are only two slopes evident from the oscillogram, it is presumed that two of the three stages in the device turn on virtually simultaneously. Measurements made during single-shot turn-on tests (reported below) show that, contrary to predictions from the computer modelling, the auxiliary thyristor turns on very soon after the pilot thyristor. The transition that can be seen in Figure 6.5.8 is therefore concluded to be the transfer of current from the auxiliary thyristor to the main cathode; the transfer of current from the pilot to the auxiliary occurs too soon to be distinguishable.

The transition between the two characteristics in the oscillogram is interesting. When Rc2 is "shorted out" by turn-on of the main cathode (point "A" in the oscillogram), the voltage that was dropped across it transfers to the circuit inductance, causing a rapid increase in di/dt. The current then rises until it reaches the characteristic it would have followed had there been no control resistance in the thyristor. A similar effect can be seen in the anode current waveforms derived from

the computer model (eg Figure 6.3.8), although in this case the major transition occurs when current transfers from the pilot thyristor to the auxiliary and main cathodes.

It is also clear from Figure 6.5.8 that the difference in delay time between threshold and strong triggering is much smaller than at low voltages. This is useful, since the need for coherent turn-on in a series string of thyristors (ie minimum variation in delay times) increases as the voltage is raised.

Figure 6.5.9 shows 1kV turn-on waveforms for 20°C and 95°C. The effect of temperature on the speed of turn-on is evident; also, the peak gate voltage has increased due to the rise in the value of the control resistor with temperature.

Figure 6.5.10 shows the turn-on delay time at 95-105°C as a function of anode voltage and optical overdrive. The uncertainty about device temperature arises from self-heating during the tests; the device started the tests at 95°C, but its temperature was increased by the inrush current, and for 4.5kV turn-on it had risen by 10°C during the test. The curves in Figure 6.5.10 are corrected to allow for the reduction in threshold triggering level with voltage, and the shape of the characteristics is generally as expected.

The effect of thyristor heating caused by the inrush current was evident during the high-temperature high-voltage tests; for turn-on from 4.5kV at 105°C, the turn-on delay time increased visibly during the test, and the audible noise from the test circuit also indicated that this was happening. There may be a danger that thermal runaway could occur; as the gate heats up, the delay time increases, causing

the gate dissipation to increase, and therefore further rises in temperature. However, the devices being tested were isothermally heated to the required temperature; the effect of gate heating would probably be much less pronounced for an actively cooled device, as would occur in service. The gate structure will normally be cooler than the rest of the device, since it does not carry load current.

In order to see if the device would withstand true threshold triggering, the supply voltage to the laser was gradually reduced while the circuit was running, until the device stopped triggering. As threshold was approached, the delay times became very long and erratic. It was possible to slowly reduce the laser supply voltage until the device stopped triggering, and then slowly increase it until triggering restarted. This remarkable procedure was carried out for turn-on voltages up to 3kV at 50Hz without damage. However, when the device stopped triggering, the circuit left a substantial dc voltage across it, and so the test was not carried out beyond 3kV for fear of damaging the device. However, even at 3kV this represented a significant achievement.

Test 5: Device= #45 *Rc1=* 58R *Rc2=* 16R *Total=* 74R

This device survived turn-on from 4.5kV at 30°C for ten seconds. However, when the voltage was increased to 5kV, the device failed after five seconds due to voltage breakdown at the outside edge of the main cathode. The test devices had not been screened for repetitive voltage withstand capability up to this level. Figure 6.5.11 presents the variation of peak gate voltage with anode firing voltage and optical overdrive for this thyristor, and it can be seen that the peak gate voltage is highly linear as a function of anode voltage. It was

interesting to note that, for all the devices, the peak gate voltage rose slightly as the optical overdrive was increased. This could be due to two effects: firstly, a strongly-gated optical well will collapse its impedance faster than a weakly gated one, and so more voltage will be transferred to the control resistors, and secondly, the current that flows in the well before turn-on occurs will discharge the circuit capacitance slightly, and therefore the longer the delay time, the lower the circuit voltage at turn-on.

Test 6: Device= #16 *Rc1=* 162R * *Rc2=* 38R * *Total=* 200R

This device withstood turn-on from 5kV at 30°C and 4kV at 90°C for ten seconds. It failed after three seconds with turn-on from 4.5kV at 90°C. The failure site was in a similar position to that for Test 3, ie the inside edge of the main cathode, and indicated normal di/dt failure due to an excessively large value of Rc2. It is interesting to note that, for this device, the control resistors withstood a total of 3.7kV, whereas on device #43 (Test 2), they broke down at less than 1.9kV without the silicone rubber coating.

(*The values of Rc1 and Rc2 were not measured directly, and have therefore been estimated from the total measured value of control resistance).

Summary and Conclusions.

The repetitive turn-on tests were generally successful, and constituted a substantial improvement over the results obtained on the prototype thyristor. Three devices withstood the most onerous turn-on conditions at 30°C with low levels of optical overdrive, and although two of these did not withstand the same conditions at elevated temperature, this was

due to an excessively large control resistor in series with the auxiliary amplifying gate. There were no failures in the optical well, and with the correct values of control resistor, there is good reason to suppose that this design will withstand the most severe turn-on conditions required with very low levels of optical drive. Passivation of the control resistors will be required to prevent surface breakdown.

c) 56mm Devices: Single-shot Optical Turn-on Testing.

The purpose of the single-shot turn-on tests was to measure the current flowing in the different parts of the gate structure. Relatively large external wire bonds were therefore used (Figure 6.3.15) so that current probes could be attached to them. Whilst this would introduce some extra inductance into the circuit, it was considered that the voltage dropped across this would be small compared to that across the control resistors and other stray inductances in the circuit.

Three devices were tested with both pilot and auxiliary amplifying gates activated; another three were then tested with the auxiliary by-passed. In all the results discussed, a high level of optical overdrive was used, since experiments showed that (as for the prototype thyristor) the level of optical overdrive had only a limited influence on turn-on characteristics other than the delay time.

For the first two tests, the circuit was limited by passive component ratings to a maximum voltage of 3kV.

Test 1: Device= #49 *Rc1=* 438R *Rc2=* 101R *Total=* 539R (22°C)

Rc1= 690R *Rc2=* 150R *Total=* 840R (90°C)

Voltage and current waveforms for 3kV turn-on at room temperature are given in Figure 6.5.12, and there are several interesting features in

these waveforms. Firstly, it is clear that the auxiliary amplifier turns on very quickly. The auxiliary current is the difference between the two gate current waveforms I_P and I_{P+AX}, and this difference appears less than 0.2us after the pilot turns on. In contrast, the main cathode does not turn on until 1.6us, and this fast turn-on of the auxiliary could be due to the "priming" effect discussed in section 5.8, whereby current flow below the latching level of the optical well initiates turn-on in the auxiliary thyristor. It will render invalid the computer modelling carried out for devices with two or three amplifying gates, since in the model the amplifying gate was assumed to have a turn-on delay time of 0.55us. In particular, the addition of an auxiliary amplifying gate will have significantly more benefit in reducing the pilot energy dissipation than the model predicted.

A second feature of the waveforms in Figure 6.5.12 is that, if there had been no direct measurement of gate current waveforms, it might have been assumed that only the pilot thyristor was conducting up to the time the anode voltage started to collapse (1.6us); turn-on of the auxiliary and main cathodes could have been assumed to be simultaneous. In fact, the pilot starts to commutate off after only 0.4us, and for the rest of the time up to 1.6us, it is the auxiliary control resistor Rc2 which is limiting the anode current, not the pilot resistor Rc1. This information is a valuable input to the gate design process.

The third interesting feature of the waveforms is that at the moment the anode volts start to collapse, approximately 20A is flowing in the anode, whereas only 8A is flowing through the pilot and auxiliary thyristors combined. Whilst the subtraction of these two waveforms is likely to be prone to some inaccuracy, it does suggest that a current

of the order of 10A is flowing through the depletion layer in the (off-state) main cathode. This would lead to a high instantaneous power dissipation (about 27kW) in the turn-on region. If the gate currents had not been measured, it would not have been possible to discriminate between dissipation in the main cathode turn-on zone and dissipation in the control resistors (about 22kW).

There is a large "hump" in the anode-cathode voltage after turn-on, the peak value being around 1350V at 5us. This feature arises because conduction has not had time to spread far into the main cathode, and so the small area that is conducting has a high forward voltage drop. As a result, the instantaneous dissipation in this region is approximately 400kW. Since it is related to the local current density in the turn-on region, the magnitude of the voltage "hump" can be used as an indicator of impending di/dt failure. This device would be expected to be prone to di/dt damage, since in the repetitive turn-on tests, a device with an Rc2 of 40R (#11) was destroyed by turn-on from 3kV at 30°C. The fact that this device survived 3kV turn-on with an Rc2 of 101R illustrates the difference in severity between single-shot and repetitive turn-on tests.

A final observation is that the effect of the large control resistors on the decay of pilot current can be clearly seen. The pilot is almost completely turned off after only 2us, which will keep its energy dissipation low. The device was not tested to destruction, and was re-used for Test 6.

Test 2: Device= #25 *Rc1=* 64R *Rc2=* 16R *Total=* 80R (22°C)

Rc1= 96R *Rc2=* 27R *Total=* 123R (90°C)

Figures 6.5.13 and 6.5.14 show turn-on for this thyristor at room

temperature from 1kV and 3kV respectively, and Figure 6.5.15 shows 3kV turn-on on an expanded time-base. Because the control resistors are much lower than for the previous device, the anode voltage droops more before turn-on occurs in the main cathode. This will reduce energy dissipation in the main cathode turn-on region, but increase dissipation in the gate due to the correspondingly higher currents. The peak gate current (I_{PEAK}) has increased from 8A for the previous device to 27.5A for 3kV turn-on, which will increase the size of the turn-on region in the main cathode. As a result, the "hump" voltage has fallen from 1350V to around 700V.

It is also of interest to note that the peaks of the gate current waveforms are noticeably sharper for 3kV turn-on than for 1kV turn-on. This is because the collapse of impedance in the thyristor stages is faster, and so the transition between the different impedance states is more rapid.

For 1kV turn-on, the peak currents in the pilot and auxiliary thyristors are 6A and 11A respectively; however, at 3kV the figures are 19A and 13A, so that the auxiliary carries less peak current than the pilot. This is because the main cathode switches before the auxiliary current has had time to build up any higher. It suggests that Rc1 should be increased at the expense of Rc2, so that the peak current duties are more in line with the cathode areas of the two stages.

Decay of the gate current is slower than for the previous device, as would be expected from the lower values of control resistors. The pilot current decays to virtually zero after approximately 5us compared to 2us previously.

Test 3: Device= #4 *Rc1=* shorted *Rc2=* shorted

The purpose of this test was to investigate the capability of the gate structure with no control resistors. The thyristor failed as the turn-on voltage was being increased above 1500V; there was no obvious damage site, but the device was cracked. This was presumed to have occurred due to stress set up by local heating during turn-on.

Waveforms for 1500V turn-on are shown in Figure 6.5.16. The gate currents are very high, as would be expected with no control resistors: the pilot current reaches a peak of 60A, and is still 20A at 8us after turn-on. Turn-on in the main cathode occurs approximately 0.5us after the pilot, and so the current that flows in the pilot later than this (resulting in significant energy dissipation) does not contribute to safe turn-on. The design is therefore inefficient, in that energy loss is incurred in the pilot thyristor without any resultant benefit for the rest of the gate structure.

Test 4: Device= #35 *Auxiliary amplifier by-passed Rc2=* 60R (90°C)

For this and the last two devices tested, the auxiliary amplifying gate was by-passed, so that current flowing in the optical well was only limited by Rc2; see Figure 6.5.17. The purpose of this arrangement was to test whether an auxiliary amplifier was necessary.

Figure 6.5.18 shows voltage and current waveforms for turn-on from 3250V at 90°C. Since the main cathode does not turn on until approximately 1us after the pilot, the circuit reaches a "steady-state" condition in the intervening period, such that the pilot current rises to a peak and then levels off. The waveform shapes are similar to those obtained from the computer studies (eg Figure 6.3.6), after

allowing for the fact that the main cathode turn-on delay time at 3250V is longer than at 5300V. In the computer study, the pilot reached its peak current after 0.4us, whereas the measured pilot current reaches its peak after approximately 0.3us.

In this test, the optical well melted after three shots from 3500V. Unfortunately, however, there are no other results with which this can be compared directly to see whether, for the same total control resistance and test conditions, a device with an auxiliary amplifying gate is more robust than than one where the pilot feeds the main cathode directly. The nearest comparison is Test 5 of the repetitive turn-on tests, where device #45 (total control resistance = 74R) withstood repetitive turn-on from 5kV at 30°C. However, the combination of a higher control resistance and lower test temperature mean that this alone does not prove conclusively that a double amplifying gate is better than a single one. Nonetheless, some conclusions are drawn at the end of this section.

Test 5: Device= #15 *Auxiliary amplifier by-passed Rc2=* 108R (90°C)

Waveforms for 2kV turn-on at 90°C for this device are given in Figure 6.5.19. In this case, it can be seen that the pilot current trace did not have such a well-defined "plateau" at its peak, but continued to increase slowly during the delay before the main cathode switched. The reason for this is not clear; for example, it may have been a measurement phenomenon rather than a real effect. For each new device, the measuring system had to be dismantled and rebuilt around the device, and so there could be differences between one device and the next in terms of inductive pick-up or earthing effects.

This device failed as the turn-on voltage was being increased above

2kV. The failure site was between one of the gate fingers and the main cathode, and is believed to be due to uneven etching of the metallization; in fact, the finger had bulged out towards the main cathode at this point, which would lead to local current concentration and overstressing of the p-base. The level at which failure occurred is therefore not meaningful.

Test 6: Device= #49 *Auxiliary amplifier by-passed Rc2=* 150R (90°C)

Waveforms for 3750V turn-on at 90°C are given in Figure 6.5.20. It is interesting to note that the main cathode did not suffer di/dt failure even though it was only receiving 9A drive. The large value of control resistor is evident in the ratio between the anode current di/dt before and after the main cathode switches; beforehand the anode current increases at 27A/us, whereas after the control resistor is "shorted out" by the main cathode, it increases at 63A/us.

The device failed at turn-on from 4kV, as a result of the control resistor tracking across. It was not possible to tell whether this was electrical surface breakdown or thermal runaway; however, thermal runaway is quite likely, since for Rc2 to be as high as 150R, the control resistor must be etched fairly deep, which will make it more prone to thermal runaway.

Summary and Conclusions.

Direct measurement of currents in the gate structure yielded insights into the turn-on process which would not be evident from conventional measurements of anode voltage and current. The auxiliary amplifying gate was found to turn on faster than expected, thus relieving stress on the optical well; this could be due to the current "priming" effect

discussed in section 5.8. For the levels of total control resistance likely to be used in the future (ie similar to device #25, Test 2), the split between Rc1 and Rc2 should be changed so that Rc1 constitutes a larger proportion of the total; this will keep the current duties on the pilot and auxiliary stages more in line with their cathode areas.

From these tests alone, it could not be proved conclusively that a double amplifying gate would be necessary for the light-triggered thyristor. However, when combined with results from the repetitive turn-on tests, it is apparent that two amplifying stages are required. When the auxiliary was by-passed in Test 4 above, the control resistance of 60R was insufficient to prevent di/dt failure of the optical well. If only one amplifying stage (the optical well) was used, more than 60R control resistance would therefore be required to protect it. However, we know from the repetitive turn-on tests that the gate stage feeding the main cathode must have a series resistance of less than about 40R, otherwise di/dt failure will occur at the inside edge of the main cathode (Repetitive Tests 3 & 6). Since these two requirements are incompatible for a single amplifying gate design, it follows that two amplifying stages will be required.

Figure 6.5.21 shows the peak pilot current as a function of anode voltage for the different tests. The final conclusion that can be drawn from these results is that the total control resistance should probably be more than 60R, as was the case in Test 4. Although the final device will have two amplifying stages (unlike Test 4), so that the optical well will see reduced stress, the device must survive repetitive turn-on from 5kV rather than single-shot turn-on from 3.5kV. The move to 5kV and repetitive firing will probably more than offset the reduction in

stress resulting from the presence of an auxiliary amplifying gate.

d) 56mm Devices: Dv/dt Turn-on Testing.

The purpose of these tests was to ascertain whether or not the 56mm devices were dv/dt protected, ie whether they could withstand ungated turn-on due to excessive rate of rise of voltage. Since they had already been shown to withstand marginal optical triggering from high voltage (with suitable values of control resistors), it was expected that they would also withstand marginal dv/dt turn-on.

The test circuit, which is shown in Figure 6.5.22, had previously been found to produce representative waveshapes across the thyristor and its associated reactor. The damping circuit resistance is included, but the damping capacitor is not, since it will be effectively short-circuit on the time-scales of interest. The 0.25uF capacitor is charged from a 0-20kV supply, and then the spark gap is triggered to collapse its voltage very rapidly. This produces a fast-rising forward voltage across the thyristor, whose dv/dt rate was progressively increased until the thyristor fired. Two devices were tested at 110°C, and the results are discussed below.

Test 1: Device= #21 *Rc1*= 232R *Rc2*= 52R *Total*= 284R (22°C)

Figure 6.5.23 shows a dv/dt turn-on event for this device. The dv/dt at which turn-on occurred was approximately 3kV/us, averaged over the first microsecond of the ramp. Turn-on was marginal, since the anode-cathode voltage reaches its peak value before turn-on occurs.

The device failed after three turn-on events, with the failure site being at the inner edge of the main cathode. This was similar to the failure on device #11 (repetitive inrush Test 3), as could be expected.

since Rc2 on that device was 40R whereas in this case it was 52R at room temperature. In addition, the control resistor value will be much higher at 110°C, further weakening the gate pulse to the main cathode edge.

Test 2: Device= #46 *Rc1=* 110R *Rc2=* 25R *Total=* 135R (22°C)

This device had already withstood repetitive turn-on from 4.7kV at 105°C. For the dv/dt test, the pilot-anode voltage was measured to check that the pilot was turning on first, and the results for two dv/dt rates are given in Figures 6.5.24 and 6.5.25.

In Figure 6.5.24, turn-on is seen to be marginal, since the pilot voltage collapses approximately 0.6us before the main cathode voltage, and the turn-on of the auxiliary and main cathodes can be clearly seen in the shape of the waveforms. In Figure 6.5.25, turn-on is stronger, and it is not possible to distinguish between turn-on of the auxiliary and main cathodes. This device survived over 20 dv/dt turn-on events, which was taken as adequate proof that the thyristor was dv/dt protected.

Conclusions.

When combined with previous evidence about the ability to survive turn-on with threshold triggering, it was concluded that the design will be immune to damage by dv/dt turn-on, provided the correct values of control resistors are used.

e) 56mm Devices: Optical Sensitivity Measurements.

The optical sensitivities of the 56mm devices were measured at room

temperature with 60V bias, and the results are given in Table 6.5.2. Average sensitivity was 22nJ, with the spread being from 8nJ to 30nJ. However, it is interesting to consider the relationship between the optical sensitivity and the control resistor values: both the secondary well in the optical gate and the control resistors are etched to the same depth, and so the value of control resistance can be taken as an indication of the depth of the optical well. As the etch goes deeper, the optical efficiency improves and the control resistor values increase.

Optical sensitivity versus total control resistance is plotted in Figure 6.5.26, from which it can be seen that as the resistance value increases, the threshold triggering level first of all falls rapidly, reaches a minimum, and then starts to increase slowly. It appears that beyond a certain etch depth, the increase in resistance in series with the optical well more than compensates for its improving efficiency, so that the performance of the gate as a whole deteriorates. The high value of control resistance limits the optically-induced current flow, so that a larger optical pulse is required to trigger the device. This agrees with the conclusions reached in section 5.2, that resistance in series with the optical well reduces the gate sensitivity.

If the dimensions and depth of the control resistors are designed such that the "nominal" device lies at the lowest point of the characteristic in Figure 6.5.26, then two benefits will accrue. Firstly, optimum optical sensitivity will be achieved, and secondly, the production spread in sensitivities will be minimised. Etch depths above and below the nominal will both tend to increase the threshold triggering level, so that the production spread will be approximately halved compared to a design lying some way away from the minimum point. The 56mm devices

had been produced with a fairly wide range of etch depths, and in practice it should be possible to achieve substantially closer "grouping" of properties than that shown in Figure 6.5.26. It should be noted that the depth of the secondary well in the optical gate has no effect on dv/dt capability, and so the above optimization in the design of the device can be carried out independently of dv/dt rating.

A series of tests were carried out to provide general information on optical sensitivity. The sensitivity of device #46 was measured as a function of voltage and temperature up to 120V, and this is presented in Figure 6.5.27. Figure 6.5.28 then shows the relative variation of optical sensitivity with voltage up to 3000V at 20°C, where 100% represents the threshold triggering level at 500V. Finally, Figure 6.5.29 shows the optical sensitivity of the #0 56mm device as a function of temperature up to 150°C.

Conclusions.

All of the devices showed good optical sensitivity, being better than the design target of 40nJ sensitivity at 60V. Variations of sensitivity with voltage and temperature were as expected. The optical sensitivity was found to be affected by the values of the built-in control resistors, and this phenomenon can be used to find a design optimized for maximum sensitivity and minimum production variations.

f) 56mm Devices: Dv/dt Withstand Capability.

The results of dv/dt tests performed on the 56mm devices are given in Table 6.5.3, and it can be seen that all but one (#16) of the seven devices tested met the target specification of 2-3 kV/us at 110°C. Device #36 was a "limit case" device, with a capability of 2.0kV/us at

110°C and 180V/us at 150°C.

Comparative tests on devices #46 and #49 showed that when the optical well was isolated, the dv/dt capability of the device increased. This confirmed that the optical well was the most sensitive part of the device to dv/dt turn-on, as is necessary for the dv/dt protection to work.

It is significant that when the control resistor Rc1 on device #49 was shorted out, there was no measurable effect on the dv/dt withstand capability. Since this device had very high values of control resistors (Rc1 = 441R), any practical device with values lower than this would definitely see no benefit in dv/dt performance compared to a structure with no control resistors. This is probably because the impedance of the depletion layer is large compared to the control resistor values, so that the effect on the amplitude of displacement current flow is small.

Conclusions.

All but one of the devices met or exceeded the required dv/dt withstand capability, and the relative sensitivities of the different parts of the gate structure were shown to be as designed. Thus, given reasonable production control of sheet resistivities in the structure, the test results indicate that it should be possible to achieve a good yield on dv/dt withstand capability with the present design.

g) 30mm Devices: Optical Sensitivity Measurements.

Figure 6.5.30 and 6.5.31 show respectively the optical sensitivity as a function of voltage at 22°C and 90°C for the 30mm devices from

slice #43.

The wells labelled 1,2 and 3 had different diameters of secondary well (Figure 6.3.17) in order to investigate the trade-off between optimum sensitivity, obtained with a small secondary well, and ease of fibre alignment, requiring a larger well. Well 1 had no secondary well, to check the benefit of leaving the n-emitter in place; see section 5.2 . It can be seen that at both room temperature and 90°C, Well 1 is significantly less sensitive than the other two cases, indicating that the benefit of leaving the n-emitter in place is more than offset in this design by the increased optical attenuation that results.

At room temperature, Well 2 is more sensitive than Well 3, which was expected since its secondary well is smaller. The sensitivities shown are effectively for perfect fibre alignment. However, it is interesting to see that the difference between the two has almost disappeared at 90°C. The reason for this is not immediately apparent; it may be caused by a temperature-dependent change in the relative importance of the thin emitter layer and the main emitter parts of the optical well structure, so that the significance of the secondary well diameter diminishes as the temperature increases. However, since the design of the valve light triggering system is likely to be determined by the low-temperature, low-voltage sensitivity of the thyristors, it will still be worthwhile keeping the secondary well diameter as small as realistically possible.

Wells 2, 5 and 6 have a constant secondary well diameter, and different primary well diameters; see Figures 6.3.17 and 6.3.18. Theoretically, Well 2 should be the most sensitive of these structures, since it has the largest diameter primary well, and therefore the highest resistance

under the thin emitter layer. However, it can be seen from Figures 6.5.30 and 6.5.31 that the expected ranking of sensitivity is reversed; the smallest primary well is the most sensitive, with Well 2 being the least sensitive of the group. This anomaly occurred as a result of the chemical etching process used to produce the wells; for small features such as the secondary well, the heat produced by the chemical reaction cannot escape as easily as for a large, plane surface, and so there is a local temperature rise in the etchant. This accelerates the etching process, so that the depth of the etch becomes proportional to the feature size. Thus the primary well in Well 6 was deeper than that in Well 5, which was in turn deeper than that in Well 2. The depth of the primary well determines the sheet resistivity under the thin emitter layer (and therefore the optical sensitivity), and the difference in etch depths has more than offset the effect of the well diameters on sensitivity.

Wells 2 and 4 have the same dimensions, but Well 4 has metallization that runs down into the primary well, whereas for Well 2 the metallization stops at the edge of the well. This experiment was to confirm the improvement in optical sensitivity that results from having metallization in the well. It can be seen from the results that the effect is even more pronounced than for the prototype device; in that case, the improvement was of the order of 30%, whereas in this case the well with metallization is approximately 60% more sensitive than that without. The improvement is greater in this case because the wells are larger than those used on the prototype design, and so the lateral resistance of the thin emitter layer will be that much greater. The improvement is maintained with temperature.

Figure 6.5.32 shows the results for the devices on slice #46 at 22°C.

The results for wells 2,4,5 and 6 are closely grouped, and Figure 6.5.33 shows these on an expanded vertical scale. All the wells on this device are significantly more sensitive than those on slice #43; since Well 1 is more sensitive (10nJ at 60V rather than 49nJ), then this must be due to a deeper primary well etch, because Well 1 has no secondary well. The control resistors for the 56mm device from slice #46 were not unusually high, which indicates that the deep primary etch has been followed by a relatively shallow secondary etch. This affects some of the "rankings" for this device. Well 2 is more sensitive than Well 1 as before, but Well 3 is much the same as Well 1: the improvement in sensitivity gained by etching the large secondary well is offset by the loss of voltage generated under the thin emitter layer in the centre of the well.

The results for Wells 2,4,5 and 6 are all closely grouped together. This should not be taken as too significant; to achieve sensitivities as high as those indicated, the wells must be close to the forward blocking junction and therefore at very high sheet resistivities. Any minor variations in etch depth will therefore produce changes in sensitivity far outweighing those arising from the differences in horizontal geometries between wells.

Also shown in Figure 6.5.32 is the sensitivity of the 56mm device from slice #46. The optical well on the 56mm device has the same dimensions as Wells 2 and 4; the metallization runs into the well, but not as far as that on Well 4. It can be seen that the 56mm device is significantly less sensitive than both Wells 2 and 4, and this is taken to be due to the effect of the control resistors limiting the optically-induced current flow, and therefore reducing overall optical sensitivity.

Conclusions.

Optical sensitivity measurements on the 30mm devices confirm the basic design principles that were followed in the 56mm device optical gate, ie that a secondary well should be used, that it should be no larger than necessary, and that metallization should run down into the well. Some unexpected features were encountered; the effect of the secondary well diameter diminishes as the temperature increases, and the increased chemical etching rate for small features means that a small primary well may in practice be more sensitive than a larger one of the same nominal depth.

The beneficial effect of increased temperature on sensitivity with a large secondary well probably cannot be utilised, since normally HVDC valves will be required to start up from "cold" (typically 10°C). If the effect of locally accelerated etching of small features cannot either be measured or predicted, then it may be best that the primary well diameter is kept large enough to avoid this phenomenon. This will maintain uniformity of optical sensitivity and dv/dt capability.

h) 30mm Devices: Dv/dt Withstand Measurements.

Dv/dt withstand was measured for wells from three slices, using a 3kV voltage ramp in all cases.

Figure 6.5.34 shows the results for Wells 2 and 5 from slice #13. Since Well 5 has a smaller primary well than Well 2, it should have a higher dv/dt rating; the dimensions were calculated to give it a rating 50% higher. Although the measured values for Well 5 were generally higher than those for Well 2, it was only by a small margin, and this was

attributed to the accelerated local etching of the primary well discussed above.

Figure 6.5.35 shows the results for Wells 3 and 4 from slice #5. These results are the opposite way round to that expected; Well 4 should be more sensitive than Well 3, since it has a smaller secondary well and metallization in the primary well. However, whilst these features are important for optical sensitivity, they will have less impact on dv/dt capability. This is because the dv/dt rating is not very sensitive to the design of the centre of the well; most of the voltage build-up under dv/dt current flow occurs at the outer edge of the well, as shown in Figure 4.2.10. It follows that Wells 3 and 4 would have similar theoretical dv/dt ratings, and that the measured difference between them is simply normal production scatter.

Figure 6.5.36 shows the results for Wells 1 and 6 from slice #16. These results are as expected, ie Well 6 has a high dv/dt rating due to its small primary well diameter, and Well 1 has a low dv/dt rating due to its large primary well and no secondary well. Well 6 will have undergone accelerated etching in the primary well; however, since its diameter is small anyway, the consequent reduction in dv/dt capability would be minimal. As above, most of the voltage build-up under dv/dt current flow will occur under the normal emitter part of the well, away from the small primary well. The high optical sensitivity and control resistor values for the 56mm device from slice #16 suggest that it had a deep primary etch; This would explain the large difference between the dv/dt rating of Well 1 (mainly determined by the primary etch) and the rating of Well 6 (mainly determined by the depth of the normal n-emitter).

For any given slice, the wells on the 30mm devices appear to have a lower dv/dt rating than the 56mm device. For example, the 56mm device #13 had a dv/dt capability of 3.3kV/us at 110°C, whereas Well 2 from the same slice, with nominally the same dimensions, had a withstand of only 1.7kV/us at the same temperature. This is considered to be a measurement effect rather than a real phenomenon; the 56mm devices have over four times the cathode area of the 30mm devices, and so will present a much larger dynamic load to the dv/dt ramp generator. (The diameter figures quoted are for the silicon before it is cut out and bevelled). This means that the small devices will see a significantly "stiffer" voltage ramp, with a consequent reduction in apparent dv/dt rating. This highlights the importance of the way in which parameters such as dv/dt rating are specified. Section d) above described how the 56mm devices were tested for dv/dt turn-on in a representative circuit under realistic conditions, which gives confidence in the performance of the devices in service.

Conclusions.

The dv/dt measurements indicated that the diameter of the secondary well and the presence or absence of metallization in the well have minimal influence on dv/dt withstand capability. Up to a certain point, the accelerated etching effect will reduce the dv/dt rating of a well by increasing the sheet resistivity under the thin emitter layer; however, for a fixed outer diameter of normal n-emitter, as the primary well becomes smaller its influence on the dv/dt rating of the structure diminishes anyway, and so the accelerated etching effect becomes unimportant.

One conclusion that can be drawn from these measurements is that there

is a need for both practical experimentation as well as theoretical calculation in designing an optical gate for a particular dv/dt rating. The combination of different effects (such as accelerated etching), and the variation of parameters at different rates with temperature, mean that it is difficult to predict both the absolute value and the spread of dv/dt capabilities that will be achieved with a particular gate design.

6.6 CONCLUSIONS.

A substantial number of experimental devices were produced and evaluated, generating a large volume of information which can be used for further development of a light-triggered self-protecting thyristor.

Whilst none of the new structures showed any improvement in resistance to damage during forward recovery failure, the Selective Failure Zone devices gave clear evidence of correct operation with only low levels of selective irradiation. The failures that did occur on these devices were due to inadequate in-rush capability in the gate, and there is therefore the prospect that normal design techniques could improve the performance of this type of structure. This is a very important result, since it illustrates for the first time that built-in forward recovery protection may indeed be feasible. It will be recalled from Chapter Three that the difficulty of developing this type of protection was identified as the most significant barrier to the achievement of a light-triggered self-protecting thyristor, and therefore the advanced outdoor valve.

The 56mm optically-triggered thyristor showed excellent characteristics, and with the correct values of control resistors, the devices were immune to damage by dv/dt turn-on and weak gate pulses. Optical sensitivities were also better than the target value. By measuring current flows in different parts of the gate structure during turn-on, valuable insights were gained into the turn-on process.

The 30mm optical test structures yielded useful information on the optical sensitivity and dv/dt capability of various designs of optical well. This confirmed the assumptions made in the design of the 56mm

device, and will prove of value in any future exercises aimed at optimising the design of the light-triggered thyristor.

CHAPTER SEVEN: DEVELOPMENT OF THE VALVE LIGHT TRIGGERING SYSTEM.

7.1 INTRODUCTION.

Chapters Four, Five and Six described the work carried out towards the development of a light-triggered self-protecting thyristor. In order to operate, the thyristor needs a system which will deliver optical energy of the correct magnitude, waveshape and wavelength to its photo-sensitive gate through an optical fibre. This system is referred to as the valve light triggering system, since one such system would control all the thyristors in one HVDC valve.

Whilst great improvements have been made in the sensitivity of light-triggered thyristors since they were first developed, the level of optical energy required to trigger them is still orders of magnitude greater than that used in normal optical fibre data transmission systems. At the time the valve light triggering system was developed, the optical energy levels required placed it at the limits of available light source technology.

The development of the light triggering system was subcontracted by GEC to a technology consultancy, and the author was responsible for the immediate supervision of the work carried out. This chapter describes the development of the valve light triggering system under the following headings:

- Specification for the light triggering system;
- Review of existing technology;
- Review of light sources;
- Selection of the light source;
- Basic system design;
- Design of the thyristor package optics;
- Design of the optical mixer;
- Design of the laser drive circuit;
- Design of the control system.

The optics inside the thyristor package are an important part of the overall optical system, and so these are also described in this chapter.

7.2 SPECIFICATION FOR THE LIGHT TRIGGERING SYSTEM.

The purpose of the light triggering system is to make the thyristor behave as a diode during the time that the valve is required to conduct by the control system. Normally, this is achieved by giving the thyristor (either conventional or light-triggered) a single gate pulse at the appropriate time. As the thyristor turns on, the current rapidly rises above the latching level, and the thyristor stays in conduction until the circuit reduces the current to zero by natural commutation; see Figure 7.2.1. Only one pulse is therefore required per power frequency cycle.

The characteristics of this "normal" gate pulse are important. Conventional thyristors can suffer di/dt failure if the gate pulse is too weak, and so it is normal practice to make every gate pulse strong enough for the device to withstand turn-on from the highest possible voltage. The amplitude of the gate pulse is therefore usually set to be ten times the threshold triggering level. The pulse should also have a fast rise time, because it must reach peak amplitude before the thyristor starts to turn on. Once turn-on has started, current or optical energy delivered to the gate makes no contribution to the turn-on process. When the light triggering system was developed, it was believed that the light-triggered thyristors would be damaged by a weak gate pulse in the same way as conventional thyristors, and therefore the system was designed to give 10x gate overdrive and prevent (as far as possible) a weak gate pulse being sent. The subsequent demonstration of controlled turn-on (and therefore immunity to damage from weak gate pulses) means that the system developed could be simplified, and this is discussed further in Chapter Eight.

In addition to the normal "one pulse per cycle" operation described above, there are conditions when more frequent gating is required, such as during start-up, operation at low current and shut-down. Under these conditions, the current in the valve can fall to zero before the end of the required conduction interval, as shown in Figure 7.2.2. The thyristors will then start to recover their ability to block forward voltage, and so if the voltage across the valve becomes positive, a new gate pulse will be needed if they are to conduct again. Two strategies can be used to achieve the required "diode-like" behaviour under these conditions, which are illustrated in Figure 7.2.2. The simplest is to continually gate the thyristors during the required conduction interval, which is referred to as "dc gating". However, this requires substantial power from the gate drive (whether optical or electrical), and so the more common technique is therefore to measure the voltage across the thyristor, and send a new gate pulse if it becomes positive during the conduction interval. This is referred to as "pulse-on-demand". Its drawback is that it requires a measurement of voltage at each thyristor level, which would be difficult to achieve in a valve that had no active electronics. The dc gating option, in contrast, requires no feedback from the valve.

In the initial feasibility study carried out for the system, several techniques for passive sensing of the voltage across the thyristor were considered, but none of them seemed promising in terms of technical or economic feasibility. The dc gating approach was therefore adopted, but to increase the laser life-time, dc gating would only be used when the valve operating conditions were such that discontinuous current might occur.

Conventional HVDC valves achieve very high levels of reliability and availability, and a light-triggered valve must have performance at least as good in this respect. The light triggering system should therefore have high reliability and availability: reliability is usually achieved by the inclusion of redundant equipment, whereas good availability can be achieved by making it possible to replace failed redundant equipment whilst continuing to operate the system.

System failure in this context is defined as sending an optical pulse of less than the specified magnitude to a thyristor, since conventionally this could result in the destruction of the thyristor. In the limiting case where no gate pulse is sent at all, the thyristor will fail to trigger; although it will not be damaged, the converter performance will deteriorate. Maintenance will therefore be required to restore normal operating capability.

The required service life for HVDC equipment, including the light triggering system, is usually at least 25 years. The specification for the valve light triggering system can therefore be summarised as follows:

- Function : to make the light-triggered thyristors behave as diodes during the period required by the valve control system.

- Optical requirements: $\geq$ 400nJ delivered to the thyristor photosite within 1us for normal triggering.

  $\geq$ 5mW continuous power delivered to the thyristor photosite to maintain continuous conduction.

- Reliability : 1x10E-8 failures per hour per thyristor.

- Lifetime : 25 year system life.

- Cost : £200 per thyristor for a system designed for a valve with 100 thyristors.

7.3 REVIEW OF EXISTING TECHNOLOGY.

A number of HVDC manufacturers have developed light triggering systems for use with light-triggered thyristors. Due to the technical and commercial nature of this field, the pattern has been for manufacturers to develop both the thyristor and the light triggering system to be used with it. Brief descriptions of the various systems developed are given in the following sections.

a) General Electric/EPRI.

General Electric (USA) have been developing light-triggered thyristor valves for HVDC applications in conjunction with the Electric Power Research Institute (EPRI) [7.1]. A number of light sources were investigated for use in the light triggering system, including Xenon flash lamps, CW YAG lasers, semiconductor lasers and light emitting diodes. Initially, semiconductor lasers were favoured, but later development was based on the use of a caesium vapour lamp. Large slab waveguides are used to collect the light from each lamp before it is mixed and then fed to the fibres going to the thyristors; see Figure 7.3.1.

GE/EPRI have been working on the light-triggered thyristor for this system since the mid 1970's. In 1977, EPRI took out a patent on a design for a light-triggered thyristor package [7.2]; this patent is very general in nature, and the commercial implications for any light-triggered thyristor design would need careful consideration.

b) Toshiba.

An outline of the Toshiba system is given in Figure 7.3.2. A redundant

pair of high power light-emitting diodes feed into a bifurcated optical fibre bundle [7.3, 7.4]. This fibre bundle has a diameter of 2mm, and since the photo sensitive area of the thyristor is smaller than this, the light guide inside the thyristor (known as the "hockey-stick") is tapered. No attempt is made to monitor the output of the light-emitting diodes, but instead they are replaced every 2 years. As with all the fibre bundle systems described in this section, there are no in-line connectors because of the high losses that they incur. These losses (typically around 3dB) arise from the random alignment of individual fibres in the two fibre bundle ends.

c) Hitachi.

The Hitachi system is based on a single light-emitting diode feeding into a large-diameter, single-core optical fibre [7.5, 7.6]. The diameter of the fibre was chosen to match the area of the photo-sensitive gate on the thyristor, and therefore tapering of the hockey-stick was not necessary. There is no attempt to measure the output of the light-emitting diode, nor to provide redundancy.

d) Mitsubishi.

An early Mitsubishi system used a light source that consisted of 20 light-emitting diodes each feeding into a single strand of a large-diameter fibre bundle [7.7,7.8]. This novel arrangement meant that the system could be designed to be tolerant to failure of individual light emitting diodes. As with the Toshiba system, the hockey-stick had to be tapered to focus the light down onto the photo-sensitive area of the thyristor. However, more recently they adopted a system very similar to Hitachi, i.e. a single LED feeding into a large diameter (500um) single fibre [7.9]. The LED is a double-heterostructure device with three

light-emitting points.

e) ASEA.

ASEA have stated that they are working on a light-triggered thyristor [7.10] and are known to have developed a high power light-emitting diode, but otherwise no details of their system are available.

f) Siemens/Brown Boveri.

In the late 1970's Brown Boveri were working on an auxiliary light-triggered thyristor, activated by an LED and fibre bundle [7.11, 7.12]. However, few details of the light triggering system were published, and no recent papers exist on this development. A publication by Siemens in 1985 [7.13] describes an auxiliary thyristor which bears some resemblances to the Brown Boveri device, and appears to use the same LED/fibre bundle combination. The Siemens paper discusses direct light triggering of the main thyristor, and suggests that flash-lamps or double-heterostructure semiconductor lasers would be necessary for a system utilising such thyristors.

g) Westinghouse.

Westinghouse built a prototype 8-level light-triggered thyristor valve for a Static VAR Compensation System [7.14]. This used single-heterostructure laser diodes feeding into fibre bundles which were multi-furcated to provide redundancy; see Figure 7.3.3. However, problems were encountered with multiple laser diode failures [7.15], and present work is based around the use of LEDs [7.16].

h) Others.

Both Fuji [7.17] and AEG-Telefunken [7.18] are known to have worked on developing light-triggered thyristors fed by LEDs or laser diodes, but neither have published details of a light triggering system. Fuji is not known to be currently active in HVDC.

7.4 REVIEW OF LIGHT SOURCES.

a) Basic Principles.

The light source is a critical component in the valve light triggering system, in terms of its technical feasibility, cost and reliability. The work carried out by others did not indicate one type of light source as being ideal, and in any case the specification for the light source is dependent on the thyristor design and other aspects of the system configuration. It was therefore necessary to review the various types of light source available in order to select the one optimised for this application.

In order to make the selection easier, a screening parameter was defined, based on the radiance required by the system. Radiance is a fundamental characteristic of a light source, which cannot be increased by the use of lenses. Radiance is defined as:

$$R = P / (A \times S)$$

where: R = radiance

P = optical power

A = area of optical flux

S = solid angle of convergence/divergence of optical flux.

Lenses can be used to alter the balance between A and S, but the product of the two remains constant. Reflections, losses and incorrect matching of optical components will all reduce the radiance in the system.

The radiance required to trigger the thyristor can be readily calculated, since the area of the photosite and the maximum angle at which

incident light will be absorbed rather than reflected are known. The optical power required for triggering is also defined, so that the minimum radiance required can be calculated from the above formula. A separate radiance figure for the dc gating condition can be derived in a similar manner.

When classified by radiance, light sources generally fall into two categories: laser and non-laser sources. Laser sources typically have small emitting areas and well collimated beams; this results in a very small product A x S, and thus very high radiance. Non-laser sources usually have a much higher divergence, which gives them a low radiance. The source area varies widely within this category.

For normal triggering, optical pulses of only a microsecond or less are needed, and some types of light source are naturally suited to this mode of operation. However, sources only suitable for continuous operation ("CW" or Continuous Wavepower) could be controlled by a shutter of either mechanical or electro-optical design, the latter having no moving parts. Such an arrangement would also provide the continuous illumination necessary to hold the thyristors in conduction, and so both types of light source (pulsed and CW) were evaluated.

b) Light Sources Considered.

The following types of light source were included in the evaluation:

- Non-laser sources:
  - Arc: Xenon arc (CW/pulsed)
    Mercury arc (CW)
    Caesium arc (pulsed)
  - Metal Halide: Tungsten Halogen (CW)
    HPI (Hydrogen/Phosphorus/Iodine) (CW)
  - Semiconductor: Light Emitting Diodes (CW)
- Laser sources:
  - Gas: Copper vapour (pulsed)
    Gold vapour (pulsed)
    Krypton (pulsed)
    Helium Neon (CW)
    Argon Ion (CW)
    Helium Cadmium (CW)
  - Solid-state: Nd:YAG (Neodymium: Yttrium Aluminium Garnet) (CW/pulsed)
  - Semiconductor: Laser Diodes (CW/pulsed)

A number of other sources, particularly gas lasers, were omitted since their wavelength fell outside the 800 - 1000 nm range required by the thyristor. The following sections discuss each type of light source.

Arc Sources.

Xenon and Mercury arc lamps offer high levels of radiance. However, their lifetime is limited by erosion of the electrodes and blackening of the tubes. For continuously rated lamps, the lifetime is about 1000 hours, which is unacceptable. Xenon pulsed sources have a life of

around 10 million pulses, but at 50 Hz operation this only yields 55 hours operating life. Caesium vapour lamps (as used by GE/EPRI) do not suffer the same degradation mechanisms as the first two types, and consequently they can have a very long life. However, suitable versions of these lamps are not commercially available, and so special development would be needed for this application.

Metal Halide Sources.

Although powerful, these sources have a low radiance due to their large emitting area, and they were inadequate for this application.

Light Emitting Diodes.

High power light-emitting diodes have been specially developed by a number of manufacturers for use in conjunction with light-triggered thyristors. They offer reasonable radiance and life; however, as with caesium arc lamps, they are not commercially available, and so it would be necessary to develop one specifically for this application.

Gas Lasers.

Gas lasers are available which can give extremely high pulsed and continuous power outputs. Optically, this makes them ideal for this application. However, they are generally expensive and have short lifetimes. Some of the commonest industrial lasers (eg CO_2) are of unsuitable wavelength, and the only credible contender in terms of life and cost was found to be the copper vapour laser. Present systems require that the copper vapour is replaced approximately every 300 hours, but systems are being developed which can continuously replace the vapour. However, such systems were not available at the time of the

project, and cost would still be a drawback.

Solid State Lasers.

In operation, Nd:YAG lasers absorb incoherent light and emit coherent light, and so they are dependent on some other source for their optical energy. Flash lamps are commonly used as light sources, which suffer the drawbacks described in the section on Arc sources. Solid-state lasers are being developed which can use LEDs or semiconductor lasers for light sources, but these were not yet available. When being driven by a continuous light source, high pulse powers can be achieved by Q-switching the laser.

Semiconductor Lasers.

There are two main types of semiconductor lasers available: pulsed and continuous (CW). Pulsed lasers generally use a single-heterostructure construction developed in the early 1970's, and they can be purchased from a number of suppliers. Their performance is characterised by high pulse powers but low duty cycles, typically 0.1% or less. CW laser diodes are of a more complex construction, and the technology is newer and advancing quickly. There are substantial consumer markets for these devices, mainly in the field of read/write optical systems, and this is providing impetus for their development. Because of high development costs, state-of-the-art devices are usually very expensive, typically several thousand dollars each.

7.5 SELECTION OF THE LIGHT SOURCE.

Due to the budget and timescale available for the project, GEC did not want to develop a light source specially for this application unless absolutely necessary. Sources that were commercially available in a suitable form were therefore much preferred. This placed an important extra constraint on the light sources that could be used, in addition to the technical and economic requirements already discussed.

In addition, it was considered desirable that the chosen light source would be one that was being developed elsewhere for other applications; thus TDPL would benefit from improvements in cost and performance arising from developments in a wider market.

For all the above light sources, information was obtained from manufacturers on products that were commercially available, and on developments that were likely to reach the market place in the near future. This information was used to assess the different types of light source. Where a light source was technically satisfactory but uneconomic, it could still be considered if there was the prospect of sufficient price reduction in the future. On this basis, the following judgements were reached:

Non-laser sources:

- Arc: Xenon and mercury arc lamps would require such frequent replacement as to be uneconomic and inconvenient. A caesium arc lamp would require special development.
- Metal Halide: Inadequate radiance.

- Semiconductor: LED's with sufficient power for triggering duty would require special development. Commercial devices are potential candidates for dc gating duty.

Laser sources:

- Gas: Generally too expensive and with insufficient lifetime. Copper vapour lasers with continuous gas replenishment could be suitable in the future, but are not near enough to be used as the basis of a prototype system.
- Solid-state: Limited by the light sources used. LED sources might make these attractive in the future.
- Semiconductor: Pulsed lasers relatively cheap and available; attractive for triggering duty. CW lasers expensive, but potentially suitable for dc gating duty.

Pulsed laser diodes were found to have several advantages for the light triggering system:

- They give high pulse energies with maximum pulse lengths of 150 - 200ns.
- The wavelength of a standard GaAs pulsed laser is 905 nm, corresponding to nearly peak sensitivity for a silicon light-triggered thyristor.
- The emitting facet for a high-power device is typically 200 x 2um, which can be coupled fairly efficiently to a convenient size of optical fibre.
- GaAs lasers were developed for military applications, which means that second sources are available and there is reasonable certainty of continuity of supply.

The combination of all these factors meant that pulsed semiconductor lasers were the preferred light source to be used for triggering the thyristors.

The choice of light source for the dc gating duty was less clear. Both LEDs and CW semiconductor lasers are naturally suited to this mode of operation. However, LEDs have relatively low radiance, requiring large diameter optical fibres or fibre bundles and a tapered hockey-stick, and CW lasers are expensive. Pulsed lasers could also be used for dc gating by rapidly pulsing them at short intervals, and it was considered by MEDL that if the pulses were not separated by more than 20 us, this would appear as continuous illumination to the thyristor. As long as the average power was equal to 5 mW, the thyristor would therefore behave as a diode. The choice between these three alternatives could not be made without a more detailed consideration of possible system designs, which is given in the next section.

7.6 BASIC SYSTEM DESIGN.

A brief consideration of the reliability performance of a valve and its light triggering system reveals an important constraint on the design of the system. HVDC valves normally incorporate a certain amount of redundancy in the number of thyristor levels used; this means that the valve can continue to operate even with a small number (typically 3%) of short-circuit thyristor levels. If each thyristor is driven by one light triggering system, then there is automatically the same level of redundancy in the light triggering system as for the valve as a whole. However, if one light triggering system serves a significant number of thyristors, eg 10% of the total, then failure of one system will require the whole valve to be taken out of service for maintenance. There is therefore a strong incentive that the light triggering system should serve either one thyristor or the whole valve; if failure of one system requires the valve to be shut down for maintenance, then there should be as few of these systems as possible.

In view of the above, four possible system designs were considered. The first three used one system per thyristor, whereas the fourth used one per valve. The designs are shown in Figures 7.6.1 to 7.6.4.

System 1 uses pulsed laser diodes to trigger the thyristor and LEDs for dc gating; see Figure 7.6.1. Two of each type of source are provided for redundancy, and their outputs are combined in a fibre bundle. A photo-diode measures back-scatter, in order to monitor degradation of the light sources. Due to the area of the fibre bundle required, the thyristor hockey-stick has to be tapered.

System 2 uses a similar arrangement of light sources to System 1, but

instead of a fibre bundle, a single fibre is taken from each source up to the thyristor hockey-stick; the hockey-stick then acts as an integrating rod. See Figure 7.6.2.

System 3 is the same as System 2, except that CW laser diodes are used to provide dc gating instead of LEDs; see Figure 7.6.3. The CW laser diodes can be coupled to much smaller fibres than are necessary for LEDs.

In System 4, the outputs of a number of pulsed laser diodes are combined in an optical mixer, from which individual outputs are taken up to each thyristor in the valve. A spare output on the optical mixer is used to monitor the light outputs of the laser diodes. A simplified version of System 4 is shown in Figure 7.6.4; in practice, there would be at least four times as many thyristors as laser diodes in the system.

For each of the above systems, the cost, light source replacement interval and other factors affecting system performance were calculated.

System 1 was rejected on the grounds of cost and the fact that system performance requirements could not adequately be met with commercially available LEDs, which would need annual replacement. Fibre bundles are not desirable in a high voltage environment because of the possibility of moisture wicking between the fibres; in this case it would also not have been possible to provide an in-line connector because of the high losses that this would incur.

System 2 was cheaper than System 1 but still suffered the same problems of insufficient performance and frequent LED replacement.

System 3 was effectively a further improved version of System 2. However, despite being able to meet system performance and lifetime requirements, the cost per thyristor was unacceptable due primarily to the high cost of CW laser diodes.

System 4 had a number of attractions, although as originally conceived it still exceeded the cost target. The single main advantage of the system was flexibility, since it allowed light sources to be configured as desired at the input of the mixer. By firing laser diodes in sequence, pseudo-"DC" gating could be provided to hold the thyristors in conduction. The system originally proposed used single pulsed laser diodes on the input, but with this arrangement the laser life was predicted to be very short due to the overdriving required to meet system performance requirements. Subsequent designs were therefore based around fibre-coupled arrays of pulsed lasers, which make better use of the high radiance of the laser sources. Each array contains ten lasers, and each laser in the array is coupled to a rectangular fibre with a 10:1 aspect ratio. This avoids much of the loss of radiance that occurs when the laser stripe (with an approximate aspect ratio of 100:1) is coupled into a circular fibre. These rectangular fibres are then brought together to form a square source, which couples more efficiently into a circular fibre; see Figure 7 6.5.

Another major advantage of the design of System 4 is that it allows the level of redundancy to be tailored to the system requirements. With one system per thyristor, the minimum level of light source redundancy is 100%, ie two sources either of which is capable of triggering the thyristor. This level of redundancy, which is excessive in the context of the whole system, applies whatever advancements may be made in laser

technology. Since the light source cost is likely to be a major component of total system cost, the economic attractiveness of this approach suffers accordingly. With System 4, each redundant laser is shared between all the thyristors, and so the minimum level of redundancy is equal to the reciprocal of the number of laser diodes required for system operation. Adequate reliability can therefore be achieved at significantly reduced cost.

In a similar way, System 4 can take advantage of higher power devices as they become available, by using fewer lasers to supply the same number of thyristors. "One system per thyristor" designs do not have this advantage.

Finally, by sequential firing of the lasers, one output of the optical mixer can be used to monitor the powers of all the lasers in the system. This minimises the cost of the monitoring and the loss of radiance caused by the "tap-off" for the monitor. Systems 1 to 3 need one monitor per thyristor (if monitoring is to be used), so that both these factors deteriorate.

The advantages of System 4 can therefore be summarised as follows:

- minimum cost;
- flexibility in the provision of redundancy;
- flexibility to use higher power devices as they become available;
- the provision of dc gating by commercially available pulsed laser diodes;
- one photo-diode can monitor all light sources.

The remaining sections of this chapter consider the implementation of System 4 in more detail.

7.7 DESIGN OF THE THYRISTOR PACKAGE OPTICS.

The optical path inside the thyristor package is an important part of the light triggering system as a whole, and its design is discussed below.

The high levels of current and heat flow in a typical HVDC thyristor require high pressure, large-area contacts and high-efficiency heat-sinks. Conventional designs use water-cooled heatsinks clamped against both faces of the thyristor package. This makes it mechanically impractical to adopt the simplest form of optical coupling to the thyristor, ie to have an optical fibre entering the thyristor package normal to the surface of the silicon slice. The only realistic solution is to have a fibre entering the side of the package and then turning through 90° to meet the photosite; see Figure 7.7.1. This fibre configuration is referred to as a "hockeystick".

When a fibre is bent, the numerical aperture (NA) of light travelling inside it increases; the smaller the bend radius, the larger the increase in NA. The thyristor package used imposed a bend radius of approximately 8mm, and to avoid optical loss, the fibre needed to have a NA of at least 0.7. This can be met easily by an unclad fibre, but such fibres incur optical loss whenever there is a discontinuity at the fibre surface. Losses would therefore result from a hermetic seal to the fibre, a mechanical locator used to position the fibre over the photosite, and any minute scratches on the surface caused during handling. The use of a clad fibre is therefore greatly preferred.

Most commercial clad fibres have a NA of 0.3 or less, but a fibre with a NA of 0.7 was specially produced for this application. This

considerably eased problems in the three areas above, ie handling, sealing and locating.

The problems of assembling the thyristor package need careful consideration. The fibre itself is fragile, being less than 0.5mm diameter, and it must be located over the photosite to an accuracy of approximately ± 10um. Once the fibre has been bonded to the photosite and the thyristor package, any movement of the thyristor slice relative to the package (particularly along the fibre axis) will set up large stresses in the fibre, potentially shortening its life. The proposed package design is shown in Figure 7.7.2, and its features are discussed below.

An extended feedthrough tube is used to bring the fibre into the package, providing the hermetic seal and mechanical protection for the fibre. Indium solders are available which will solder directly from the tube to the glass. A solder preform is slid onto the fibre and into the feedthrough tube, and a furnace cycle then melts the solder and forms a hermetic seal. Alternatively, the fibre could be gold or copper coated (by evaporation) and a more conventional solder used. The feedthrough tube extends into a slot in the copper pole-piece; when the heavy pole-piece is being positioned, the fibre cannot be seen, and is therefore vulnerable to damage. The tube provides rotational alignment for the pole-piece, and prevents it from crushing the fibre.

Alignment of the fibre over the photosite is achieved by means of a precision brass ferrule, shown in Figure 7.7.3. The ferrule has a cross section which will sit in the optical well, and alignment is achieved by placing the ferrule over the well and feeling when it locates. The fibre is bonded into the ferrule and polished, and then a UV-cured epoxy is used to bond the ferrule to the thyristor slice.

Before assembly, the fibre is moulded into a silicon rubber coating; see Figure 7.7.4. This gives it additional mechanical strength and protects it from gross mechanical scratches along much of its length. The silicon rubber coating can be handled normally, and also serves to locate the solder slug inside the feedthrough tube.

The assembly procedure is therefore as follows:

- The brass ferrule is bonded to the end of the fibre, and the fibre end is cleaved and polished (Figure 7.7.3);
- The silicon rubber coating is moulded onto the fibre (Figure 7.7.4);
- The silicon slice is located in the package;
- The solder slug is slid onto the free end of the fibre;
- With the brass ferrule pointing upwards, the fibre is slid into the feedthrough tube and optical connector (Figure 7.7.5);
- Epoxy is applied to the end of the fibre protruding through the optical connector;
- The fibre is rotated, and the brass ferrule is located and epoxied to the photosite;
- Both sets of epoxy are cured;
- The cathode pole-piece is located and welded on;
- A furnace cycle solders the fibre to the feedthrough tube;
- The outer end of the optical fibre is cleaved and polished.

The assembly procedure is critical, since if the fibre is misaligned or broken after being bonded to the slice, the whole thyristor could be wasted. For this reason, the use of an epoxy which can be dissolved is very desirable, but even so, the cost of reworking could be considerable if there was a high reject rate. It may be necessary to revise the above procedure in the light of experience.

7.8 DESIGN OF THE OPTICAL MIXER.

The function of the optical mixer is to take light from several sources, mix it together, and feed it to a number of output fibres. This must be achieved with the minimum optical loss. Commercial optical mixers are available, but they are generally expensive, limited in the number of inputs and outputs, and have poor uniformity across outputs. It was therefore decided to develop a mixer specially for this application.

There are two main sources of optical loss in a mixer: interface loss and geometry loss. Interface loss occurs at all interfaces in the optical circuit, and arises from reflection, dirt and misalignment of optical axes. It can be minimised by good mechanical design and handling practices. Geometry loss occurs when, for example, a rectangular mixing block is coupled to circular fibres, so that some light is lost from the system; see Figure 7.8.1. There can also be a loss in radiance arising from mismatches in area and NA, but this was less critical than achieving the necessary power density, since lasers with a very high radiance were being used.

The prototype optical mixer was designed to have 6 inputs and 24 outputs. The laser array sources needed a fibre with a core diameter of at least 400um diameter for good coupling efficiency, and the output fibres could not have a core diameter greater than 300um, to be compatible with the thyristor. Computer analysis of the mixer geometry showed that minimum loss was achieved with a simple rectangular array of output fibres, as shown in Figure 7.8.1, rather than a close-packed hexagonal array and/or a circular geometry.

The prototype design was originally based around the use of plastic fibres. Although known to be lossy at 905nm wavelength, they offered several advantages:

- They have a thin cladding layer (eg 5um rather than 20-50um for glass fibres), which has a considerable beneficial impact on the geometry loss, as can be seen from Figure 7.8.1.
- They are available with outside diameters of 250 and 500um, which meets the system constraints and offers a simple 4:1 area ratio for the mixer design.
- They are much easier to handle and terminate than glass fibres.

Unfortunately, when the plastic fibres were procured they were found to have a loss several times their datasheet value at the required wavelength; this made them unusable for the mixer. Baking in an oven to dry them out gave no improvement, and so glass fibres had to be used.

The mixing rod was rectangular, and to enable the mixer to be reworked in case of fibre damage, the arrays of fibres were butt-coupled to the ends of the rod without epoxy. In order to achieve accurate alignment, a precision dowel system was used, as shown in Figure 7.8.2.

A butt-coupling system was also used to connect the laser arrays to the input fibres. Suitable spring-loaded connectors were not available at the time, and so an accurate mechanical assembly was needed to bring the end of an optical connector close to the output of the array without actually touching it. Active alignment of the array to the fibre was used.

In a commercial system, monitoring of the laser powers would probably be triplicated. Since the output of the mixer is orders of magnitude

more powerful than a photo-diode can measure, considerable attenuation is required in the monitoring circuit. This means that the three most lossy outputs of the mixer could be used for monitoring purposes.

An important operational consideration is that it should be possible to change a laser input array whilst the system is operating. This affects the mechanical and electrical design of the front end of the mixer system.

The prototype mixer was successfully built and tested. The main difficulties experienced were in the handling of the fibres without breakage, and in constructing the rectangular arrays of fibres. The latter was because the fibres naturally fall into hexagonal arrays, and constructing rectangular arrays was troublesome. This could be overcome with experience and some simple jigs. Plastic fibres were originally selected to avoid breakage during the construction of the mixer, and so the difficulties with glass fibres were not unexpected. This again would improve with experience. Plastic fibres with lower loss at 905nm are being developed by several manufacturers, and when these are available they will considerably facilitate construction of the mixer.

7.9 DESIGN OF THE LASER DRIVE CIRCUIT.

The pulsed laser arrays used in the prototype system had a maximum pulse length of 200ns and a peak drive current of 40A. This implies a di/dt in the drive circuit of the order of 1000A/us, which places considerable constraints on the electrical and mechanical design of the drive circuit.

Many applications of these lasers use a simple sinusoidal current pulse, generated by a thyristor triggering and discharging an LC circuit into the laser. However, in order to achieve the rapidity of pulsing for the dc gating mode whilst remaining within the duty cycle limits of the laser (maximum 0.1%), a variable pulse length was needed: 200ns pulses would be used for triggering, and 125ns pulses for dc gating. This meant that an LC discharge was not suitable for driving the laser, and instead a switched "constant-current" source was needed. This was achieved by using a high source voltage (210V nominal) and resistors to limit the current. An FET was used as the switching element, as shown in Figure 7.9.1.

In order to achieve a low inductance in the drive circuit, several measures were employed. The FET, capacitors and resistors were all located as close as possible to the laser array. Parallel arrangements of four capacitors and eight resistors were then used to reduce self-inductance effects, and wide tracks and ground planes were employed on the PCB to minimise circuit loop areas. Finally, several CMOS output stages were parallelled to drive the FET. By using these techniques, current rise-times of 25ns or less were achieved, corresponding to a di/dt of 1700A/us.

The high reliability target for the system had an impact on the way the drive system was designed. A fibre-optic link was used to convey firing signals from the main control circuit to each laser driver, so that failure of a driver could not affect the main control circuit. In addition, the timing of the length of the laser pulses was generated separately for each driver, so that that this did not represent a common-mode failure point.

Each driver had its own switched-mode power supply to generate the high and low voltage rails for the drive circuit. The use of a common power supply for all the drivers was considered, with resistive current limiting so that supply to the others would not be interrupted if one driver failed. However, the recharge time required for the driver capacitors during dc gating meant that high pulse currents were needed; a resistor which would permit this would dissipate excessive amounts of heat if the driver went short-circuit, and there might also be a fire hazard in the driver itself. The use of a fuse in the supply to each driver was rejected on the grounds that the fuse would probably be the most unreliable component in the driver. Individual power supplies were therefore used.

The prototype drivers were built, operated and life-tested successfully. The only practical problem encountered was that the 20kHz switching action of the power supply generated a significant amount of interference; this did not affect the driver itself, but could be picked up by oscilloscopes electrically or physically close to the supply.

In addition to the drivers for the prototype system, several individual drivers were produced with one laser array coupled to a single fibre.

The laser was controlled by a 5V CMOS input. These individual light firing systems were used extensively for life-testing of the lasers and laboratory triggering of thyristors during evaluation work.

7.10 DESIGN OF THE CONTROL SYSTEM.

The control system has three main functions:

- to generate normal trigger pulses;
- to generate repetitive pulses for dc gating operation;
- to monitor the state of the laser diodes and take appropriate action.

Each of these functions is discussed below.

a) Normal Triggering.

For the majority of its working life, the control system will simply be required to receive Fire commands from the valve control system, and transmit suitable trigger signals to the laser drivers. When all the lasers in the system are new, there will be a considerable excess of optical power available for normal triggering. It would therefore be possible to maximise laser life in this operating mode by using short trigger pulses at the start of life, and then increasing the optical pulse length as the lasers deteriorate. (The pulse amplitude is fixed by the driver high-voltage supply). However, this approach was rejected on the grounds of reliability for two reasons; firstly, in order to achieve maximum reliability the circuit involved in generating the trigger commands should be as simple as possible, and secondly, apart from any effects on the reliability of the system itself, failure could result in a short gate pulse being sent when a long one was required. If the thyristors were vulnerable to damage from being turned on by a weak gate pulse (as was thought at the time), then all the thyristors in the valve could be destroyed. This was clearly an unacceptable risk, and so a fixed length trigger pulse was used. As stated above, the

length of the gate pulse was generated by a separate timer circuit for each laser driver, so that no one fault could result in all the lasers sending a short pulse.

b) DC Gating.

The dc gating mode requires that a minimum of 5mW average optical power be delivered to the thyristor gate, and that the pulses involved should not be separated by more than 20us. The simplest way to implement this would be to fire a laser every 20us, and if the pulse energy was below a certain threshold (necessary to give 5mW average power), the next laser in sequence would be fired immediately. However, this approach could lead to a "self-destruct" mode, as follows. The laser outputs fall as their temperature increases, and so if several lasers were near the minimum threshold, then when one fell below it, the resulting increase in firing rate for the others could make them warm up and fall below threshold. This would further increase the rate of firing of the remainder, and so on.

The only way to avoid this scenario is to make use of the optical energy output of each laser, however low it may be. This approach avoids any step changes in the duty cycle on the lasers, thereby reducing the possibility of the "thermal runaway" described above. There are two methods of implementation: the first is to use a fixed pulse spacing, and vary the length of the optical pulse delivered by the laser. The pulse length would therefore increase progressively as the laser aged. The second method is to use a fixed pulse length and vary the pulse spacing; as the laser ages, the delay in firing the next one in sequence reduces, so that the same average power is maintained. Both these approaches minimise the duty on the lasers during the dc

gating mode, but from the point of view of the electronics design, the 20us gap between pulses can be controlled much more accurately than the pulse length of 200ns or less, and so the second approach was chosen.

Even with this technique for dc gating, a positive feedback mechanism for laser degradation is present, since the pulsing frequency increases as the laser powers fall. However, the degradation should happen more slowly than with the simple technique, and so it should be possible to detect and alarm the condition before the system performance is seriously affected.

c) Laser Monitoring.

There are two types of laser failure mode; progressive and catastrophic. Progressive failure normally occurs as a result of lattice defects propagating from the electrical contacts into the lasing region of the crystal, and causing dark spots. Deterioration of the facets and the epoxy contact to the front facet can also cause a gradual loss of laser power. Catastrophic failure usually occurs due to facet burn-off, where the optical energy in a single pulse is so high that the facet coating is vaporised. This could result from a control error, or the failure to short-circuit of the FET in the laser driver, so that the optical pulse length is much longer than intended.

Sudden loss of laser output could also be caused by a failure in the laser drive electronics, even though the laser was still intact.

The lasers represent a major component of system cost, and therefore the maximum use must be made of their working life. A number of simple monitoring and replacement strategies were considered, but they resulted in lasers being replaced more often than was actually necessary. A

fairly complex system was therefore developed to reduce laser replacement costs to the theoretical minimum, and this is described below.

The monitoring philosophy was that no single event should result in a weak gate pulse being sent. If the system state had deteriorated to the point where a single event could cause this, then a "Trip" signal would be generated. This would shut down the converter, so that one or more lasers could be replaced. In order to minimise the number of forced shutdowns of this nature, it was further required that no single event should result in a Trip signal being generated without prior warning. An "Alarm" signal is therefore generated if a single event could result in a Trip.

Minimising the frequency of laser replacement requires that the power of each individual laser be measured. The laser outputs are combined in the optical mixer, and so each laser must be fired separately for its power to be measured at the output of the mixer. Sequential firing of the lasers is already used in the dc gating mode of operation, and so it is convenient to measure the laser powers during dc gating operation.

In case dc gating is only used occasionally, after every 1024 Fire pulses a single cycle of dc gating is initiated by the control system during the valve conduction interval. This enables the laser powers to be measured whilst having a negligible effect on laser life and no effect on valve operation.

Having measured the laser powers, the system then calculates Alarm and Trip levels to be applied to the output of all the lasers when fired simultaneously. The minimum permitted gate pulse power is programmed

into the logic. The worst event that could occur is that the most powerful laser in the system fails catastropically, and so to ensure that this does not result in a weak gate pulse being sent, the system adds the output of the most powerful laser to the minimum permitted pulse power to obtain the Trip level. If the output of all the lasers falls below this level, then failure of the most powerful laser would result in a weak gate pulse, and so the system is shut down; see Figure 7.10.1.

The Alarm level is calculated in a similar way, except that the outputs of the two most powerful lasers in the system are added to the minimum permitted power, as shown in Figure 7.10.2.

When the Alarm and Trip levels have been calculated, every subsequent Fire pulse is compared against them, so that the appropriate action can be taken. The Alarm and Trip levels are revised every 1024 Fire pulses, or every time dc gating is used, whichever comes first.

7.11 CONCLUSIONS.

A design for the valve light triggering system was reached after consideration of several types of light source and several system configurations. This design offers the lowest cost, best utilization of the light sources and greatest flexibility to make use of improvements in semiconductor laser technology. In particular, it incorporated the following improvements over systems previously developed by others:

- The system uses a light source technology which is being continuously developed world-wide for several major markets;
- The use of lasers permits an efficient design of thyristor optical gate structure to be achieved;
- The light sources are commercially available from multiple sources;
- The high laser radiance means that standard diameter multimode fibres can be used to transmit optical pulses to the thyristors. Fibres and connectors of this type are cheap and readily available with low optical attenuation;
- Complex tapered hockey-sticks are not required in the thyristor package;
- Monitoring of the laser powers is achieved with minimum penalty to system performance;
- Utilization of laser lifetime is maximised through the use of an "adaptive" monitoring system;
- The level of laser redundancy can be tailored to the system requirements, thereby minimising system cost.

A prototype system using 6 fibre-coupled laser arrays to drive 24 thyristors was built and tested. The normal triggering and dc gating modes of operation were successfully demonstrated, and valuable exper-

ience was gained during the construction of the prototype.

The light triggering system was developed in parallel with the light triggered thyristor, and this required that some assumptions be made about the performance of the thyristor. The prototype system that was built met the technical specification, but the design did not quite meet the economic targets set for it. However, the light-triggered thyristor turned out to have better sensitivity than anticipated and also to be immune to damage from weak gate pulses, and making use of these improvements would result in a light triggering system that was economically as well as technically acceptable.

CHAPTER EIGHT: SUMMARY, CONCLUSIONS AND RECOMMENDATIONS.

8.1 INTRODUCTION.

This project was concerned with the development of a new generation of HVDC valve technology, and as a consequence has covered a wide range of technical fields. The initial proposition was that mature HVDC technology would embody outdoor thyristor valves, which would only require remedial maintenance rather than the regular preventative maintenance needed for the present generation of equipment. The objective of the project was to investigate and develop the technologies necessary to bring about this level of maturity.

For outdoor valves to be a realistic proposition, a substantial improvement is required in the reliability of thyristor valves, which can only be achieved through the use of light-triggered, self-protecting thyristors. Since thyristors of this type had never been developed before, the major part of the project concentrated on progressing thyristor technology to the point where the principles of a light-triggered, self-protecting thyristor had been demonstrated.

Although the thyristor development is the main technological breakthrough needed for an outdoor valve to become practical, an optical system capable of triggering such thyristors is also beyond the scope of commercially available technology. The development of a valve light triggering system is therefore an integral part of the overall project objectives, and a prototype optical system was produced which incorporated several improvements over the systems previously developed by other workers.

Finally, the mechanical, insulation and cooling system designs that might be used in an outdoor valve were considered.

This chapter summarises the conclusions that have been reached in the various chapters and appendices of the thesis, and makes recommendations for further work. Some conclusions are also drawn concerning the management of the project, since effective management is clearly vital to the success of a development project of this nature.

8.2 THE THYRISTOR.

The development work carried out towards the goal of building a light-triggered, self-protecting thyristor was described in three chapters, as follows: Chapter Four reviewed the work carried out at MEDL before the author became involved, Chapter Five covered the investigation of the original prototype design carried out by the author, and in Chapter Six the design and evaluation of the test structures was described.

Many detailed technical conclusions were reached during the course of the work, but only the main ones are highlighted below.

a) Light Triggering.

The critical trade-off between optical sensitivity and dv/dt capability was examined, and it was shown that as the diameter of the secondary optical well is reduced (ie the area actually illuminated by the optical fibre), a better trade-off is achieved. This therefore led to the use of a laser to trigger the device, which permits a very small secondary well to be used. This preference for a laser source is contrary to most previous practice, where large-area sources such as LED's and flash-lamps have been used for the light triggering system.

It was also shown that for the type of device investigated, the small secondary well made the gate structure immune to di/dt failure caused by weak triggering. If a weak gate pulse (either electrical or optical) is sent to a device with a large gate structure, only part of the gate turns on, and so in order to use the full capability of the gate structure, "overdrive" is therefore needed to ensure that all of the gate is activated. However, for a small gate structure, all of the gate is activated by the time the latching current is reached, regardless of

the strength of the gate signal. There is therefore no need for gate overdrive, which for light-triggered thyristors is expensive and technically difficult to provide. This important advantage of a small secondary well in an optical gate structure has not previously been recognised in the literature.

The use of control resistors to limit current in the optical gate during turn-on was investigated in some detail. Anti-modulation rings were required, as shown by other workers, and these were also used to distribute current around an interdigitated amplifying gate. In fact, the resistors were so effective in supporting voltage during turn-on that passivation was needed to prevent surface breakdown.

It is economic to etch the control resistors at the same time as the optical well, and it was shown that there is an optimum depth for this etch. Up to a certain point, deepening the etch makes the optical well more efficient, so that sensitivity improves. Beyond this optimum depth, however, the resulting high values of control resistance start to affect the early stages of the turn-on process, so that optical sensitivity deteriorates. Once this optimum depth has been established, it can be used to maximise the performance of the thyristors in production.

Single and double amplifying gate structures were investigated, and their limits established for different values of control resistors. Computer analysis had shown that a third amplifying gate would give little benefit. Measurements of current and voltage in each part of the gate structure were made, giving valuable insights into the turn-on process. Significantly, the technique of measuring currents in the gate structure during turn-on has not previously been reported.

A 56mm thyristor with the preferred values of control resistors was shown to have impressive repetitive turn-on capability with only threshold optical gating. A series of tests were also carried out to determine the optimum design of optical well, given the limitations of processing and assembly technology, and this yielded a useful "library" of information. All of these results taken together mean that, for any future light-triggered thyristors, the optical gate structure can be designed with confidence in achieving both the required performance targets and the efficient utilization of optical energy from the light triggering system.

b) Over-voltage Protection.

The "etched well" technique for producing built-in over-voltage (VBO) protection was investigated, using several procedures for making the well and controlling its surface finish. Some interesting phenomena were discovered during this investigation, from which several novel approaches to VBO protection could have been explored. However, it was concluded that the etched well approach to VBO protection is too demanding (in terms of process technology) to be economic. The use of a simple external BOD operating through an electrical gate contact to the thyristor offers excellent performance at low cost. It was therefore decided to suspend work on VBO protection, and instead concentrate resources on those areas where technical feasibility had yet to be proven.

c) Dv/dt Protection.

Built-in dv/dt protection was successfully demonstrated on a 56mm thyristor. Indeed, it was concluded that if the thyristor can be made

immune to di/dt failure from weak optical triggering, then dv/dt protection follows automatically. This is because excessive dv/dt creates conditions in the gate structure very similar to those produced by weak optical gating.

The original prototype thyristor design used dv/dt compensation to improve its performance. However, it was shown that the topology of compensation is incompatible with the use of control resistors to prevent di/dt failure. This is because dv/dt compensation requires the presence of surface diode junctions in the gate structure, and these diodes will have breakdown voltages of a few tens or at most hundreds of volts. If the control resistors in the gate are of high enough value to prevent di/dt failure under weak triggering, then thousands of volts will appear transiently in the gate structure, destroying the surface diodes. In any case, adequate device performance was achieved without the need for dv/dt compensation.

d) Forward Recovery Protection.

The original thyristor design which aimed to provide forward recovery protection was generally unsuccessful, although a proper assessment of the design was difficult due to the number of ambiguous test results obtained. Work was therefore undertaken to design, manufacture and test a new set of development thyristors. This set incorporated an improved version of the original design, as well as several other alternative approaches.

The alternative approaches were based on the principle of allowing forward recovery failure to occur at a random location in the main cathode (as for normal thyristors), but using a special cathode design

to improve the likelihood of surviving forward recovery failure without damage. Several different versions of the special cathode designs were implemented using existing processing steps, so that no extra cost was incurred in manufacture; however, none of these proved to be successful in practice. It was therefore concluded that, although the principles involved might still be valid, extra processing steps would be needed to implement them, which would prejudice the commercial viability of the thyristor.

In contrast, the improved version of the original (Selective Failure Zone) design showed a very promising performance. The main change, compared to the original, was to approximately halve all the horizontal dimensions of the structure, giving higher levels of charge diffusion from the main cathode. This design of Selective Failure Zone functioned correctly even with very low differential radiation doses. The gate voltage was monitored to demonstrate that the Zone was turning on the main cathode under forward recovery failure conditions. Unfortunately, the length of the gated edge was not adequate to cope with the inrush conditions, and device failures occurred either due to breakdown of the small control resistor used, or melting of the main cathode turn-on edge. However, there are established techniques for overcoming this type of problem, and recommendations for improvements are given in section 8.6.

8.3 THE VALVE LIGHT TRIGGERING SYSTEM.

Previous implementations of valve light triggering systems had varied widely, with different manufacturers using flash-lamps, large LED's, multiple small LED's and semiconductor lasers as the light source. The optical transmission systems also varied from single fibre bundles to multiple interwoven fibre harnesses and quartz slabs. In every case, there was at least one part of the system that was costly and difficult to manufacture, and not available commercially. It was therefore necessary to develop a new system that would overcome the deficiencies of previous designs, whilst not leaving GEC dependent on its competitors for the supply of critical components.

An analysis of thyristor behaviour showed that laser-based triggering systems were to be preferred because of their high optical power densities (section 8.2). This feature, common to all types of laser, gives the following advantages:

- The thyristor sensitivity is improved for the same dv/dt rating;
- Conventional single-core fibres can be used, which are cheaper than fibre bundles;
- Losses at connectors can be 1dB or less, compared to at least 3dB for a fibre bundle connector;
- There is no need for a tapered fibre "hockeystick" inside the thyristor package;
- The margin between radiance available and the minimum radiance required is such that redundant sources and monitoring points can be included. Large-area sources (such as LED's and flash-lamps) are operating on the limits of the radiance required, so that both redundancy and monitoring are difficult to provide.

All the various types of laser available were reviewed. Semiconductor lasers were selected in preference to gas and solid-state lasers for reasons of cost, replacement interval and wavelength compatability with silicon thyristors.

Several configurations of optical system were considered, covering both "one system per thyristor" and "one system per valve" approaches. It was decided to adopt the "one system per valve" technique, where the outputs of several lasers were combined in an integrator, and then distributed to all the thyristors in a valve. This approach gives the following advantages:

- As higher-power lasers are developed, this can be used to reduce the number of lasers in the system, and thus improve cost-effectiveness.
- The degree of redundancy in the system can be tailored to the system requirements.
- One monitoring point can be used to measure the outputs of all the lasers in the system.
- Pulsed lasers can be fired cyclically to achieve a "pseudo-dc" output.
- As high-power CW lasers become available, these can be added on to the system to provide a dc output, thus reducing the duty on the pulsed devices.

A prototype light triggering system was built to trigger 24 thyristors, and this was successfully demonstrated. The only proprietary part of the system was the optical integrator, and even in this case, conventional components were used. The novelty lay in the detailed mechanical design of the integrator.

As well as the 24-thyristor system, several scaled-down versions of the system were built for triggering single thyristors in the laboratory. These were extensively used in the thyristor evaluation program.

When the specification for the light triggering system was being drawn up, it was believed that a weak optical gate pulse might cause the thyristor (or in this case, all the thyristors in a valve) to suffer di/dt failure. A fairly sophisticated monitoring system was therefore used to ensure that no single-contingency fault could result in a weak gate pulse being sent. The subsequent achievement of "controlled turn-on", such that the thyristors are immune to di/dt failure caused by weak gating, means that the monitoring system could be simplified. The prospects for further development of the system are described in section 8.6.

8.4 THE VALVE INSULATION, COOLING AND MECHANICAL DESIGN.

In the design of an outdoor valve, careful attention must be paid to achieving high reliability and relatively easy access for maintenance. This therefore has an important bearing on the insulation and cooling systems used for the valve, as well as the mechanical design of the installation. These factors are considered in Appendices One ("Valve Insulation and Cooling") and Two ("Valve Mechanical Design"). Since the main thrust of the project was concerned with the thyristor development, these areas were not investigated in great depth. However, the following preliminary conclusions were reached:

- The valves should be insulated by SF_6, in accordance with current trends in the industry. The only other serious possibility for the insulant is mineral oil, which suffers important disadvantages related to cost, flammability, reprocessing and inconvenience for maintenance.
- The SF_6 should be used at a fairly low pressure, to avoid pressure-vessel regulations and the potential problems of high dc field stresses on solid insulation systems.
- The valves should be cooled by a separate liquid coolant, preferably water/glycol but otherwise a Freon or other SF_6-compatible fluid.
- The valves should be enclosed in earthed steel tanks. As many valves as realistically possible should be included in one tank, up to a limit of 12 (or 24 for a rectifier/inverter group).

- Maintenance should be based on the philosophy of replacing components *in situ*. Each tank should have access for an operator to gain entry, carry out diagnostic tests, and replace any failed components. The drawbacks of this approach are small compared to those of replacing complete tanks to carry out maintenance.
- GEC's existing indoor valve mechanical design is suitable for an outdoor valve without significant modification.

The preferred approach for the insulation, cooling and mechanical design of the valve will depend fairly significantly on the actual installation being considered. Factors such as the voltage rating of the valves, climate, availability of land area, local legislation and customer preference will be important considerations in determining the optimum configuration.

8.5. THE MANAGEMENT OF LARGE DEVELOPMENT PROJECTS.

a) Introduction.

Whilst the managerial aspects of the project have not been discussed in the main body of the thesis, they are clearly an area of interest and importance. Several conclusions have been reached during the course of the work regarding the management of large development projects, and it is likely that these conclusions will have a more general application.

At the outset, there were a number of features of the project which would place greater than average demands on the project management, and these are discussed in the following sections. They are all related to the interactions between the technology areas in the project, which were described in Chapter Three.

b) Technological Complexity and Diversity.

There were three main areas of technology involved in the project:

- Valve electrical design;
- Semiconductor thyristor design;
- Lasers and optical systems design.

Each of these areas is a highly complex subject in its own right, and at the start of the project, there was no one person involved who was familiar with more than one of these technology areas. This would not necessarily have been a problem, except for the high level of interdependence between the different areas of the system design. For example, the nature of the self-protection required in the thyristor is critically dependent on the rest of the valve electrical design (section 3.3). This leads almost inevitably to the situation where an

expert in one area has to take a decision regarding technology strategy, compromise etc, with only a limited appreciation of the impact of that decision on the other parts of the system. Conversely, he might decide not to pursue a particular line of investigation, without realising the benefits it could have offered to other parts of the system.

c) Geographical Separation.

The problems of technological complexity and diversity were compounded by the fact that the experts in the different areas were geographically separated. The centre of valve design expertise was at Stafford, the centre of thyristor design expertise was at Lincoln, and the centre of laser and optical systems expertise was at Cambridge. This prevented the normal cross-fertilisation of ideas and information that would occur if the people involved had been working in close proximity. In the early stages of the project, the three groups only met at formal project progress meetings, which were monthly or sometimes even less frequently. Due to time restrictions at meetings like this, each group "filtered" the information it presented, on the basis of its understanding of what was and was not important. The other participants in the meeting were not always experienced enough in the area to appreciate the implications of what they were being told, or to suggest alternative strategies which might have benefits to their area of the system.

d) "Past history"

As will often be the case, the participants in the project were not "starting from scratch" at the beginning of the work, but had previously been carrying out development in the same or related

areas. This previous work had not necessarily been undertaken with the same objectives and constraints as the present project, but nonetheless it could influence the priorities set and the approach adopted. Even the technology consultants responsible for the development of the valve light triggering system, who had had no previous experience of HVDC or power thyristors, exhibited a predisposition to certain types of solution. This arose from their experience in other areas, but it was not always the optimum for this project.

Since it was recognised that the technology consultants had no relevant prior experience, regular meetings and contacts were maintained by TDPL to ensure that the work being carried out was consistent with the overall project objectives. This meant that any misunderstandings that arose, or approaches that were adopted which were sub-optimal in the context of the whole project, were picked up and corrected fairly quickly.

The situation was different with MEDL for several reasons. Firstly, MEDL had collaborated with TDPL for many years in the development of thyristors for HVDC applications. Secondly, MEDL had already carried out substantial work in the fields of light triggering and, to a lesser extent, built-in over-voltage protection. Finally, the work at MEDL was being funded internally rather than through TDPL, so that there was a level of implied autonomy which was not present with the consultancy.

In the course of the project, it became apparent that although MEDL had collaborated with TDPL over thyristor design in the past, the degree of knowledge of valve design required for this thyristor project was far greater than had ever been the case previously. This was because much of the expertise in valve design, ie triggering and protection, was

being moved from the valve electronics (TDPL's domain) into the thyristor (MEDL's domain). In effect, the amount of valve technology being distilled into the thyristor was being increased several times over, so that a much higher level of interaction between thyristor design and valve design was necessary in order to achieve a viable and optimised result.

It was to meet this need for greater interaction that the author spent four months working full-time at MEDL, followed by weekly visits for the rest of the project duration. This had a valuable effect on the "efficiency" of the development work, by ensuring that all the parties concerned were more aware of the global technology context in which they were working. Any decisions on strategy or technological compromise could therefore be optimised more effectively against the overall project goals. By the end of the project, the author had become conversant with all three technology areas, so that much better coordination and decision-making was achieved.

e) Conclusions.

With the benefit of hindsight, it seems that the extent of the need for "up-grading" of the knowledge about valve design (at MEDL) and thyristor design (at TDPL) was not fully appreciated at the start of the project. This might not have been too important, except that the geographical separation of the parties involved prevented a normal exchange of ideas and experience. The secondment of the author to MEDL, followed by very close technical supervision, led to a substantial improvement in the situation. The following general conclusions can be drawn from this experience:

- If a project involves experts from different technology areas working together to achieve an optimum solution, the degree of mutual education required should be considered at the start of the project (and reviewed as it progresses). Ideally, this will be achieved by both formal exchanges and informal contact in the working environment. If informal contact would not normally take place (for example, due to geographical separation), then temporary or permanent secondment should be seriously considered. There is no substitute to working alongside someone for gaining an understanding of their field of expertise. Because of restrictions of time and the inevitable "filtering" of information that will take place, formal project review meetings are very limited in the extent to which they can permit efficient interaction.

- If the parties involved in the project have existing development programs in related areas, then it is important that the objectives and priorities of the present project are clearly stated. If this is not done, then work may tend to continue along existing lines, even though this might not be appropriate for the project in hand.

- It is widely recognised that for a project to be successful, there must be a project "champion" who is committed to driving the project towards its objectives. If different parties to the project have independent resources, this gives an implicit level of autonomy in decision-making which may undermine the effectiveness of the project champion. It may therefore be preferable to have all the resources controlled by the party who is the "main contractor" for the project.

In spite of the initial shortcomings described above, it is clear that in the later stages of the project these deficiencies were dealt with, and a very high level of coordination and effective interaction was achieved.

8.6 RECOMMENDATIONS FOR FURTHER WORK.

At the start of the project, various technology areas were identified where progress was necessary to make an outdoor HVDC valve a realistic proposition. Several important advances have been made during the project towards this goal, the most significant of which is that the principle of built-in protection against forward recovery failure in a thyristor has been demonstrated for the first time. However, it is clear that there is still significant work to be undertaken before an outdoor valve becomes a commercial reality, and this is discussed in the following sections.

a) The Thyristor.

The light triggering and dv/dt protective functions were so successful in the 56mm development samples that these aspects of the thyristor can be considered as fully developed. One more design/build/test cycle could be carried out, but this would simply be fine-tuning of the design. An area where "productionising" is still needed is in the choice and method of application of the passivation on the control resistors, but this should not be a significant problem.

The main area of thyristor technology where further development is required is in the forward recovery protection mechanism. The principle was demonstrated in the development samples, but the turn-on capability of the selective failure zones must be improved to achieve protection under all conditions. This should be possible using existing gate design techniques such as amplifying gates and extended turn-on lines (possibly combined with control resistors), although the need to minimise the distance between the zone and the main cathode may require

some novel topologies. If all goes well, a commercial design could be arrived at after two more design/build/test cycles.

For reasons of cost and performance, it is recommended at present that over-voltage protection be provided by an external BOD. However, it is recognised that this is an inconvenient solution to the problem, and that built-in protection would be preferable. Unfortunately, none of the approaches normally considered for this offer the prospect of an attractive cost/performance trade-off, even if they are technically successful. Novel approaches to the problem must therefore be developed if built-in over-voltage protection is to be viable, and some suggested avenues for investigation were put forward in Chapter Five.

Although the packaging design for the thyristor was considered in the project, this may need further work. The package is inevitably rather complex and delicate, and taking this from development into production could present difficulties. The proposed design does offer a fairly efficient solution to the requirements of the device, but some modifications may be needed in the light of experience. The technique for making an electrical gate connection to the device (for an external BOD) could involve ultrasonic bonding, which is being used increasingly in the manufacture of power devices

b) The Valve Light Triggering System.

Semiconductor laser technology is continually developing, with "demand pull" coming from the market for compact disk players and other optical read/write systems. It is therefore sensible to wait until commercial implementation is being approached before carrying out significant further development of the system, since better devices are continually becoming available.

Once development has restarted, the balance between pulsed and CW lasers will need to be evaluated again. The most fundamental choice to be made is whether the existing format of the light triggering system should be retained, ie using several devices feeding a complete valve via an optical mixer, or whether this should be dropped in favour of one or possibly two lasers driving each thyristor. This choice is necessary because the requirements of the system (regarding optical overdrive, the avoidance of weak gate pulses etc) and the technology available have both changed significantly since the prototype system was developed. Medium-power CW lasers (40-50mW) with built-in monitoring are now relatively cheap, and so a very simple system using one laser feeding each thyristor via a single fibre could be the optimum solution. Alternatively, a mixer-based system could still be attractive if there are "economies of scale" in laser manufacture, so that it is cheaper to use a small number of high-power devices driving a whole valve.

If a one-laser-per-thyristor system was adopted, then development of the driver and monitoring circuits would be fairly straightforward. The choice between series and parallel connection of lasers would be made on the grounds of reliability, cost and laser life. Many commercial systems increase the drive current to a CW laser as it ages, to maintain constant optical output; this maximises laser life, but increases the complexity of the drive circuit, and would require separate control for each laser. Connecting all the lasers in one series string driven at constant current may be the optimum solution for the valve light triggering system.

If a mixer-based system is retained, the comments in the following paragraphs are relevant.

It may be cheaper to provide the "dc gating" function of the system with a small number of high-power CW devices feeding into the optical mixer, rather than multiplexing the pulsed lasers as in the prototype. As the cost of CW lasers continues to fall, they may be preferred for both triggering and dc gating functions, since they are considerably easier to drive than pulsed devices.

There is substantial scope for simplifying the control system, since the thyristors are no longer vulnerable to damage by weak gating, and therefore no protective action is required to prevent a weak gate pulse being sent; the worst that could happen is that the valve would simply fail to trigger. Therefore, it would no longer be necessary to measure the laser outputs and predict the effect of the most powerful device failing catastrophically. Faults can be alarmed after they have occurred, and there would be no call for a compulsory "Trip" function, since no equipment is in danger. Further simplification of the control system could be achieved if CW devices were used for the "dc gating" mode, since the fairly complex variable pulse spacing control for the pulsed lasers would no longer be required.

The elimination of the need for optical overdrive of the thyristors means that fewer lasers are required at the input of the optical mixer. This will allow smaller output fibres to be used, so that standard 200/280um multimode fibres could be used to transmit gate signals to the thyristors. This would be cheaper than the specialised fibre used in the prototype system. Manufacturing of the mixer could be made easier through the use of lower-loss plastic fibres currently under

development, and the performance of the system might be improved by the use of square-section fibres in the mixer assembly, to improve packing efficiency.

c) The Valve Insulation, Cooling and Mechanical Design.

Many of the technologies and components necessary for an SF_6-insulated, liquid-cooled outdoor valve have already been researched and demonstrated in prototype form by others. However, since none of these exercises have gone beyond the "demonstrator" stage, the technology cannot be regarded as being fully commercialised, and therefore further development is needed. Whilst the work required on the mechanical design of the valve will be fairly conventional in nature, development of the insulation and cooling systems could be both costly and time-consuming, since it will require the manufacture of full-scale equipment for life-testing at realistic voltage levels. Although the general form of the outdoor valve has been described in the thesis, its detailed implementation will require a significant amount of design and testing. If fluids other than the relatively familiar SF_6 and water/glycol or Freon are used, the amount of evaluation work required will increase substantially. In either case, it is likely that the justification for development will be conditional on achieving further progress with the light-triggered self-protecting thyristor, such that a high probability of a successful outcome is achieved.

8.7 CONCLUSIONS.

GEC TDPL's goal of developing an outdoor thyristor valve for HVDC converters is an ambitious one, encompassing a wide range of technologies. To achieve this goal, several technological objectives have to be reached, the most challenging of which was recognised to be the development of a light-triggered self-protecting thyristor. The work described in this thesis has made a major contribution to achieving these objectives, including the first demonstration of the principle of built-in forward recovery protection in a thyristor, and an increased understanding of the processes involved in optical thyristor turn-on. Other technological developments have been in the areas of the light triggering system for the thyristors, and the conceptual design of the insulation, cooling and mechanical arrangements of the valve. In addition to these technical advances, the need for a "multi-disciplinary" project member to coordinate research in different areas has also been highlighted, as this has made a valuable contribution to the overall efficiency of project execution.

Further work is still required towards the overall project objective, the most important aspect of which is to take the concept of built-in forward recovery protection from its demonstration in principle to a commercial reality. Success in this area will provide the spur to full-scale development of the other technologies required for the outdoor valve.

REFERENCES.

CHAPTER ONE.

1.1 A Kimbark, "High Voltage Direct Current Transmission, Volume 1". Published by Wiley Interscience, John Wiley and Sons, 1971, Chapter One.

CHAPTER TWO.

2.1 G Moraw (Ed), "Voltage and Current Stresses on HVDC Valves." A report prepared by Task Force 14.01.01 of the CIGRE Working Group 14.01, 1987.

CHAPTER THREE.

3.1 A A Jaecklin and I Bajan, "Novel Gate Concept Improves Performance of Light Fired Thyristor." IEEE Publication CH1504-0/79, 1979, Paper 10.5, pp 254-257.

3.2 V A K Temple, "Advanced Light-Triggered Thyristors for Electric Power Systems", International Conference on Thyristors and Variable Static Equipment for AC and DC Transmission, London, Nov/Dec 1981, pp 86-91.

CHAPTER FOUR.

4.1 F E Gentry et al, "Semiconductor Controlled Rectifiers: Principles and Applications of p-n-p-n Devices". Published by Prentice Hall, New York, 1964, Chapter Two.

4.2 J L Moll et al, "p-n-p-n Transistor Switches", Proceedings of the Institute of Radio Engineers, Vol 44, Sept 1956, pp 1174-1182.

4.3 V M Volle et al, "High-current thyristor switched by microsecond light pulses", Sov. Phys. -Tech. Phys. 30(8), August 1985. Published in the United States by the American Institute of Physics, 1986, pp909-912.

4.4 A A Jaecklin, "Light-Activated, High-Voltage Thyristors", "From Electronics to Microelectronics", W A Kaiser and W E Proebster (Eds), North Holland Publishing Co, 1980.

4.5 "Electric Forum", Volume 9, No 3, 1984. Published by the General Electric Corporation, USA.

4.6 V Temple and H Mehta, "Advanced Light-Triggered Thyristor". Fourth International Conference on AC and DC Power Transmission, London, September 1985. IEE Conference Publication No 206, pp310-314.

4.7 D J Page, "Light-Triggered Thyristors for VAR Generator Applications", Proceedings of the 7th IEEE/PES Transmission and Distribution Conference and Exposition, April 1979, pp222-226.

4.8 N Konishi et al, "High-Power Light-Activated Thyristors". Proceedings of the International Power Electronics Conference, Tokyo, 1983, pp559-570.

4.9 H Ohashi, T Ogura and Y Yamaguchi, "Directly Light-Triggered 8kV 1.2kA Thyristor". International Electron Devices Meeting, 1983, IEEE Publication CH1973-7/83/0000-0210, pp210-213.

4.10 A Tada et al, "4kV, 1500A Light Triggered Thyristor". Proceedings of the 12th Conference on Solid State Devices, Tokyo 1980. Published in the Japanese Journal of Applied Physics, Volume 20 (1981), Supplement 20-1.

4.11 A Ekstrom and E Albertsson, "Electrically-Triggered Thyristors Still Have the Edge". "Modern Power Systems", October 1986, pp65-69.

4.12 "A Great Impact on HVDC Transmission - The Light-Triggered Valve". Toshiba Publication No KSP-E1490, 1984-5(1).

4.13 T Kano et al, "Development of Light-Triggered Thyristor Valve for HVDC Transmission". Proceedings of the International Power Electronics Conference, Tokyo, 1983.

4.14 G Addis et al, "Advanced HVDC Valve". Fourth International Conference on AC and DC Power Transmission, London, September 1985. IEE Conference Publication No 206, pp340-345.

4.15 V A K Temple, "Controlled Turn-on Thyristors". IEEE Transactions on Electron Devices ED-30, No 7, July 1983. IEEE Publication No 0018/9383/83/0700-0816.

4.16 H Ohashi et al, "Design Consideration for High Power, Overvoltage Self-Protected Thyristor". Proceedings of the International Power Electronics Conference, Tokyo, 1983.

4.17 J Przybysz, "Laser Trimming of Thyristors to Add an Overvoltage Self-Protected Turn-on Feature". IEEE Publication No 0275-9306/85/0000-0463, 1985, pp463-468.

4.18 O Harding, P Taylor and P Frith, "Recent Advances in High Voltage Thyristor Design". Fourth International Conference on AC and DC Power Transmission, London, September 1985. IEE Conference Publication No 206, pp315-319.

4.19 Y Yamaguchi et al, "Basic Characteristics of High Voltage VBO Free Photothyristors". Japanese Institution of Electrical Engineers, National Conference, Showa 59. In Japanese.

4.20 Reference 4.1, p198.

4.21 Reference 4.1, p92.

CHAPTER FIVE.

5.1 N Konishi et al, "High-Power Light-Activated Thyristors". International Power Electronics Conference, Tokyo, 1983.

5.2 N Konishi et al, "Light Triggering Phenomena on Thyristors with an n-Emitter Layer Over the Light-Sensitive Gate Area". IEEE Transactions on Electron Devices ED-33, No 1, January 1986, pp34-40.

5.3 M Woodhouse et al, "The Control and Protection of Thyristors in the English Terminal Cross-Channel Valves, Particularly During Forward Recovery." International Conference on Thyristors and Variable Static Equipment for AC and DC Transmission, London 1981. IEE Conference Publication No 205, pp158-163

5.4 Reference 4.1, p113.

5.5 V A K Temple and H Mehta, "Advanced Light-Triggered Thyristor". Fourth International Conference on AC and DC Power Transmission, London, September 1985. IEE Conference Publication No 206, pp310-314.

5.6 V A K Temple, "Turn-on Reliability Break-Through in Light-Triggered Thyristors." International Electron Devices Meeting, 1984, IEEE Publication No CH2099-0/84, pp428-431.

CHAPTER SIX.

6.1 J P Ballad, "Further Deliberations on the Proposal for a Thyristor Structure Resistant to Forward Recovery Damage." GEC Engineering Research Centre Report NGu 975, Nov 1985.

6.2 K Glockner and M Fullmann "On the Turn-on Behaviour of Thyristors with Field-Initiated Gate." Solid-State Electronics, 1977, Vol 20, pp476-477. Published by the Pergamon Press, Great Britain.

6.3 N Diseko et al, "Application of a Digital Computer Program to Transient Analysis and Design of HVDC and AC Thyristor Valves." International Conference on Thyristors and Variable Static Equipment for AC and DC Transmission, London 1981. IEE Conference Publication No 205, pp158-163.

6.4 Y Kao et al, "A New Triggering Mode for Light-Triggered Thyristors." International Electron Devices Meeting, 1984, Paper 16.2. IEEE Publication No CH2099-0/84, pp 432-434.

6.5 V A K Temple, "Comparison of Light-Triggered and Electrically-Triggered Thyristor Turn-on", IEEE Transactions on Electron Devices, Vol ED-28. No 7, July 1981. pp860-865.

CHAPTER SEVEN.

7.1 "Light Triggering for HVDC Transmission." Electronic Engineering, October 1984, p13.

7.2 "Package for light-triggered semiconductor device." UK Patent Number 1 529 145.

7.3 A Yahata et al, "Optimization of Light Triggering System for Directly Light-Triggered High Voltage Thyristors." Proceedings of the International Power Electronics Conference, Tokyo 1983, pp571-577.

7.4 Toshiba Datasheet for Infra-red LED S4672, October 1983.

7.5 T Kano et al, "Development of Light-Triggered Thyristor Valve for HVDC Transmission." Proceedings of the International Power Electronics Conference, Tokyo 1983, pp87-92.

7.6 T Kano et al,"Development of Key Components for the Light-Triggered Thyristor Valve." Hitachi Review, Vol 32 (1983), No 2.

7.7 A Tada et al,"4kV, 1500A Light Triggered Thyristor." Japanese Journal of Applied Physics, Vol 20 (1981), Supp 20-1, pp99-104.

7.8 K Niwayama et al, "A High Voltage High Power Light Triggered Thyristor." Proceedings of the Power Semiconductor Conference, 1984, pp15-18.

7.9 T Tanabe et al, "Development of ±500kV DC SF6 Gas Insulated Switchgear." Proceedings of the International Conference on DC Power Transmission, 4-8 June 1984, Montreal, Canada, pp 141-147.

7.10 A Ekstrom and L Eklund, "Importance of the Gotland Link for HVDC Valve Development." United Nations Economic Commission for Europe, Seminar on High Voltage Direct Current (HVDC) Techniques, Stockholm, 6-9 May, 1985.

7.11 A Jaecklin, "Light-Activated, High-Voltage Thyristors " 'From Electronics to Micro-Electronics', W Kaiser and W Proebster (eds), North-Holland Publishing Company, 1980, pp 360-363.

7.12 A Jaecklin and I Bajan, "Novel Gate Concept Improves Performance of Light Fired Thyristor." IEEE Publication CH1504-0/79, 1979, Paper 10.5, pp 254-257.

7.13 H Mitlehner, "Light-Activated Auxiliary Thyristor for High-Voltage Applications " Siemens Forsch.-u. Entwickl. -Ber Bd 14 (1985), No 2, Springer-Verlag 1985, pp50-55.

7.14 D Page, "Light-Triggered Thyristors for VAR Generator Applications." 7th IEEE/PES Transmission and Distribution Conference and Exposition, April 1979. IEEE Publication No CH1139-5/79, pp 222-226.

7.15 L Lowry et al, "Field Testing of Light-Triggered Thyristors in an Electric Utility Application." International Conference on Thyristors and Variable Static Equipment for AC and DC Transmission, London 1981. IEE Conference Publication No 205, pp 92-95.

7.16 Y Kao et al, "A New Triggering Mode for Light-Triggered Thyristors." International Electron Devices Meeting, 1984, Paper 16.2. IEEE Publication No CH2099-0/84, pp 432-434.

7.17 O Hashimoto and Y Sato, "A High Voltage, High Current Light Activated Thyristor with a New Light Sensitive Structure." Power Electronics Specialists Conference, 1981. IEEE Publication No CH1652-7/81, pp 226-231.

7.18 D Silber, "Light Activated Triacs and High Voltage Thyristors." International Electron Devices Meeting, 1980, Paper 25.1, pp 638-641.

APPENDIX ONE: VALVE INSULATION AND COOLING.

A1.1 INTRODUCTION.

The ultimate goal of GEC's HVDC development program is to develop valve technology to the point where an outdoor valve becomes a realistic proposition. Such a valve would be treated in the same way as a power transformer, ie it would be situated outdoors and only maintained if it failed. This would eliminate much of the costly civil works associated with present indoor valve designs.

The light-triggered self-protecting thyristor is necessary for this goal to be achieved, since the reliability of conventional valves is not sufficiently high to be compatible with an outdoor valve design. Assuming the thyristor becomes available, decisions will be required as to the insulation and cooling technologies to be used for the outdoor valve.

Several HVDC converter stations have been built using outdoor valves, but these cannot be regarded as wholly successful since no manufacturer has ever repeated the exercise. Whilst the unreliability of conventional valves must be largely responsible for this (since costly extra redundancy must be built into the valves), it emphasises the attention that must be paid to reliability, cost and maintainability in the event of failure. This chapter reviews the work that has been carried out in the field of outdoor converter stations and similar equipment, and then discusses the various options for the insulation and cooling systems for an outdoor valve. Some conclusions are drawn, but further work would be required to enable definite recommendations to be made.

A1.2 LITERATURE REVIEW.

a) HVDC Encapsulated Equipment.

The earliest example of an outdoor valve was the thyristor unit installed by GEC (then English Electric) at the Lydd converter station of the first 160MW Cross-Channel link [A1.1]. Little has been published about this valve, but it is known that it was insulated by oil and cooled by bulk circulation of the oil over finned heatsinks. The valve was in three identical parts, each enclosed in a very large hollow ceramic insulator. The valve was installed in 1971, and the scheme was de-commissioned in 1984.

The first two schemes to be fully equipped with outdoor valves were Cabora-Bassa (Africa) and Shin-Shinano (Japan), both commissioned in 1977. The Cabora-Bassa valves, built by the HVDC Working Group, were insulated with transformer oil and cooled by pumping the oil through heatsinks which vented into the tank [A1.2, A1.3, A1.4]. A number of problems arose during commissioning of this scheme, but these were associated with conventional equipment rather than the valves themselves [A1.5]. In particular, no problems were encountered with the valve insulation or cooling systems, although due to its being the first of its kind, the scheme was undoubtedly cautiously designed.

The oil-insulated valves for the Shin-Shinano link were constructed using technology licensed from the HVDC Working Group, although in practice the implementation differed in some areas [A1.6, A1.7, A1.8, A1.9]. In particular, the thyristors were cooled by bulk circulation of the oil rather than by pumping of oil through enclosed heatsinks. As with Cabora-Bassa, no problems have been reported with the insulation or cooling systems of these valves.

After these early outdoor valve installations, all manufacturers adopted indoor air-insulated technology using either air or water cooling. This was largely due to the convenience of indoor air insulation. However, in the mid 1970's, the Electric Power Research Institute (EPRI) in conjunction with General Electric, both of the USA, undertook development of an outdoor valve insulated by SF_6 and cooled with Freon [A1.10, A1.11]. Although the design of many components for this had to be undertaken virtually from scratch [A1.12, A1.13, A1.14], a compact HVDC converter was in fact built using this technology, which was completed in 1978 [A1.15]. However, no commissioning details have been published since then, and the scheme was still not in service in 1983 [A1.16].

Although there have been no further publications concerning work on SF_6-insulated valves, GE have continued to develop Freon cooling systems because of their inherent insulating properties [A1.17]. Leaks from a water-cooled valve could cause flash-overs in a high-voltage environment, whereas Freon does not suffer from this disadvantage [A1.18]. As part of this development, a prototype Freon-cooled valve was installed in the Sylmar terminal of the Pacific Intertie HVDC link, and the experience gained has been reported [A1.19]. Several significant problems were encountered, which cannot bode well for the adoption of this technology.

Mitsubishi Electric developed an oil-insulated, oil-cooled thyristor valve in the early 1970's [A1.20]. However, this did not result in a commercial application, and the development appears to have been abandoned. More recently, they have entered into partnership with the Kansai Electric Power Co for the joint development of an SF_6-insulated, SF_6-cooled thyristor valve [A1.21, A1.22]. Two test installations have been built, comprising a prototype system and an insulation test system. The SF_6 is used at a pressure of 0.5MPa (five times atmospheric pressure), and at this pressure it has sufficient thermal capacity to enable the thyristors to be cooled with open finned heatsinks. There is therefore no separate cooling circuit, and problems of coolant leakage are avoided. It is envisaged that the valves would be constructed and sealed at the factory, and shipped to site complete.

b) Power Transformers.

In terms of physical size, voltage rating and the need for cooling, power transformers bear a number of similarities to HVDC valves when considering insulation and cooling systems. Development work in the field of power transformers is therefore included in this literature survey.

Conventionally, power transformers are insulated and cooled with mineral oil. However, some transformers are situated in environments where the flammability of oil makes it unacceptable as an insulant; for example, in the basements of buildings. PCB's (polychlorinated biphenyls) were once widely used as a coolant in these applications, but they are now recognised to be environmentally damaging, and therefore alternatives have had to be found [A1.23]. A number of synthetic replacement fluids are available [A1.24] which have good fire resistance, long-term stability and dielectric strength, but these are all substantially more expensive than mineral oil. The large volume of insulant required by power transformers means that the conventional, less costly, mineral oil is still preferred for general applications.

For a given volume of equipment, liquid insulation is generally significantly more expensive than gas insulation. Some work has therefore been done on developing transformers insulated by gas, with the more expensive liquid only being used for direct cooling of the windings [A1.25, A1.26]. These designs have employed SF_6 gas, since this has offered the best combination of cost, dielectric strength and other properties. The coolants employed include $C_2Cl_3F_3$ (Freon 113), C_2Cl_4 and $C_8F_{16}O$, with both closed coolant circuits and evaporative "spray" systems being used. In addition, a mixture of C_2Cl_4 and transformer oil was found to yield cost and fire-resistant properties that seem promising for more conventional liquid-immersed designs.

One interesting form of gaseous insulation is the "vapour mist" system being developed by EPRI and Westinghouse [A1.27]. Mists of liquids such as $C_8F_{16}O$ can have a dielectric strength several times that of SF_6, as well as good cooling properties. The principles of vapour insulation and cooling are discussed in [A1.25].

c) Other Gas Insulated Systems.

SF_6 is widely used as the insulant in metal-clad switchgear and bus-work for high-voltage AC installations. Despite being significantly cheaper per unit volume than liquid insulants, the cost of the make-up gas can still be as much as 20% of the total installation cost at high voltages [A1.28]. In addition, at common operating pressures, the SF_6 liquifies at around -30°C, which is unacceptable in some applications. Research has therefore been carried out into mixtures of SF_6 and other gases which potentially offer lower cost, better low-temperature performance and in some cases improved dielectric strength as well [A1.29]. Mixtures of SF_6 and N_2 offer the first two advantages without sacrificing too much in terms of dielectric strength, and this mixture has been applied in commercial circuit breakers [A1.30]. Mixtures with fluorocarbons such as C_4F_8 and other gases can offer higher dielectric strength than pure SF_6, although features such as toxicity and undesirable breakdown products need fuller evaluation [A1.29].

The application of metal-clad technology to DC equipment is more difficult than to AC for two main reasons: firstly, the effects of small conducting particles (eg swarf) which are already serious for AC are even worse under DC stress, and secondly, achieving even voltage grading across insulators is much more difficult. When an alternating voltage is applied across an insulating system, the voltage distribution is determined by the permittivities of the various elements of the system, whereas a direct voltage is distributed according to the resistivities of the elements. The resistivities of common insulating materials are much more difficult to control than their permittivities [A1.11, A1.13, A1.14, A1.31, A1.32]. In addition, with direct voltage, surface charge can build up on an insulator which can only decay with a time constant measured in hours; if the direct voltage is suddenly reversed, this charge may distort the voltage distribution enough to cause flash-over [A1.33, A1.34]. Because of these problems, even more care must be taken in the design, manufacture and installation of metal-clad equipment for DC than is already taken with AC equipment.

d) Conclusion.

It is apparent from the literature that there is a wide variety of insulants and coolants that can be considered for an encapsulated valve. The next section reviews the criteria that must be used in assessing the feasibility of the different liquids and gases.

A1.3 CRITERIA FOR ASSESSMENT OF INSULATING MEDIA.

In addition to dielectric strength, there are many other attributes of insulating fluids which are important for this type of application. These are listed and discussed below. In some cases, the requirements for the coolant are different to the insulant, and so they are discussed separately where appropriate.

a) Dielectric Strength.

This is important for the insulant since it affects both the physical size of the valve and the mechanical design of the conductors and support insulators. For most insulators there is no single figure for dielectric strength [A1.35, A1.36], but it varies between AC, DC and impulses of different shapes, as well as being dependent on temperature, pressure and electrode design and surface finish. For a given electrode configuration, the DC and impulse withstands are normally polarity-sensitive. The effects of chemical and particle contamination can be very important for all types of applied voltage, and in any real engineering environment it is impossible to avoid these. Stress conditioning can also occur, where the withstand level varies with the duration of a pre-applied voltage below the breakdown level, and breakdown conditioning can take place, where the breakdown level can rise or fall after a breakdown depending on the prevailing mechanisms.

The valve coolant, if separate to the insulant, needs to have reasonable dielectric strength since it will be in contact with points of differing potential. However, peak fields will generally be lower than those in the insulant, and for liquid coolants, the need to keep leakage currents to acceptable levels may be a more important constraint on cooling circuit design.

b) Corona Inception Stress.

Depending on the dielectric and the prevailing test conditions, corona inception can occur almost coincidentally with breakdown or at a significantly lower stress. Even if breakdown is not imminent, operating with corona discharge present is normally avoided since in the long term it may degrade the dielectric and any surrounding materials.

c) Leakage Resistance.

The effects of leakage currents are important in the following areas:

- Electrolytic corrosion of metal components;
- Dielectric heating;
- Dissipating electrostatic charge built up by circulating fluids;
- Voltage distribution in mixed dielectric systems.

Leakage currents are determined by the electric field and the bulk resistivity of the material, where the latter can be very sensitive to factors such as temperature, purity and moisture content. If leakage

currents are too high, then electrolytic corrosion and dielectric heating can become problematic; if they are too low, then build-up of electrostatic charge can cause difficulties. If necessary, the bulk resistance can sometimes be modified by additives.

The last factor listed above relates to mixed dielectric systems (such as oil and paper), where dc voltages are distributed according to the resistivities of the various components. Care must be taken to ensure that the system performs as required, given the wide fluctuations that can occur in resistivity. Where an ac or impulse voltage (distributed according to permittivities) is super-imposed on an existing dc stress, highly non-uniform voltage profiles can result.

d) Permittivity.

The permittivity of the insulant determines the distributed stray capacitance in the valve, which will tend to be high anyway if a metal-clad design is used. High levels of stray capacitance incur costs in the valve design, and so a low value of dielectric constant is preferable. For mixed dielectric systems, the permittivities are also important for the reasons described in the previous section. For a separate coolant, however, it is unlikely to be a critical parameter.

e) Cost.

Cost is clearly a critical factor, since it has a direct bearing on the commercial viability of the outdoor valve. It is the total system cost that matters, rather than a simple cost per unit volume, and so attention must be paid to:

- the initial cost of the fluid;
- the cost of replacement or reprocessing;
- the cost of any special facilities required as a consequence of using that fluid.

The initial cost considerations will generally favour gas/vapour insulants rather than liquids, which tend to be more expensive per unit volume. Some fluids will be more prone to leakage than others, and so the cost of replenishment will be higher. Fluids such as transformer oil degrade with time, and so need potentially costly reprocessing. Finally, any special precautions required, for example due to toxicity or flammability, will add to the cost of using a fluid.

f) Toxicity.

Adverse toxicity could be a barrier to customer acceptance, as well as introducing extra costs into the design. The toxicity of some of the more exotic chlorinated and fluorinated compounds has not been fully evaluated, and so in some cases data may not be available.

g) Stability.

It is desirable that the insulant and coolant are as stable as possible over time and temperature. However, such stability can make them environmentally unacceptable (see below), and therefore some degree of compromise has to be accepted. For fluids that do degrade, there are three main causes of degradation:

- Time;
- Temperature;
- Electric field.

Degradation with time (such as occurs in mineral oil) can sometimes be slowed down by the use of additives. Normally, however, it cannot be eliminated altogether, and so reprocessing or replacement will have to be allowed for. The other two degradation mechanisms can be avoided by good design practice, ie avoiding temperature hotspots and local areas of high electric stress. The by-products of degradation can be conducting, toxic or chemically aggressive, and in these cases, regular monitoring of by-product levels will be required as a minimum precaution.

A flashover inside the tank would cause "instantaneous" degradation of the insulant. Any by-products arising from such an event also need careful consideration, since although the immediate damage from an arc would be localised, the by-products could contaminate the whole tank and possibly the cooling system as well.

h) Environmental Acceptability.

Polychlorinated biphenyls (PCB's) were once widely used as insulants because of their dielectric strength, non-flammability and excellent stability. However, their resistance to degradation meant that when released into the environment (by accident or design), they accumulated in the food chain. This process is known as bio-accumulation, and has led to them being banned in many countries. In many cases, equipment insulated with PCB's is having to be drained, purged and re-filled with a more acceptable insulant. This illustrates the importance of not only meeting present environmental legislation, but also anticipating future legislation. In this context, the present debate over Freon suggests that it too may be banned in the foreseeable future due to its effects on the ozone layer.

i) Non-flammability.

Several different standards exist for defining non-flammability. Whilst most gases would pass as non-flammable, there are very few liquids that are truly non-flammable under all conditions [A1.37]. Even excluding mineral oil, most liquids will burn or release toxic or explosive substances under intense fire conditions. It is therefore difficult to apply a rigid "non-flammability" criterion to the insulant and coolant. Indeed, since the majority of converter stations will use transformers cooled by mineral oil, it would be anomalous to apply a more stringent requirement to the valves than is used for the transformers. Only in cases such as compact urban HVDC converters would a fire resistance

better than that of mineral oil be required both for the transformers and the valves.

j) Operating Temperature Range.

The questions of stability and flammability, relevant at high temperatures, have already been discussed. In addition, the dielectric strength of many liquids falls with temperature. Boiling must normally be avoided unless it is part of the cooling technique, in which case the dielectric strength of the two-phase system is important.

At low temperatures, freezing (of liquids) and liquefaction (of gases) are important considerations. Freezing of a liquid coolant would almost certainly be catastrophic, although freezing of an insulant would not automatically be dangerous unless the dielectric strength fell or mechanical damage resulted. When insulating gases liquefy, the vapour pressure is usually low enough for the dielectric strength to be significantly impaired. Operating pressures must therefore be chosen accordingly.

k) Heat Transfer Properties.

These considerations are mainly relevant to the coolant, where features such as specific heat capacity, film heat transfer coefficient and possibly latent heat of vaporization are important. If the insulant had good properties in this area, this would be a bonus.

l) Density and Viscosity.

As above, these are of secondary interest for the insulant, but are important properties of the coolant, since they directly affect the pumping power and operating pressure of the cooling circuit. Low levels of viscosity are therefore desirable up to a point, but very low values may make it difficult to construct a leak-free system.

m) Chemical Compatability.

The fluids used, and any breakdown or degradation products they produce, should not react with materials commonly used in electrical equipment. If special construction materials are needed, this will be reflected in the cost of using that fluid.

n) Refractive Index.

If the refractive index of the insulant is not close to that of the optical fibres used in the valve, then special care must be taken to ensure that it does not enter the fibre-optic connectors and introduce an index matching loss.

o) Development Timescale and Cost.

Apart from the final cost of the system, the cost and time to develop it are important commercial considerations. This naturally favours the use of established technologies, where cost and uncertainty can be minimised.

p) Customer Acceptability.

Customer acceptability is the ultimate criterion for the development, since the combination of cost, reliability, maintenance requirements, risk etc must be saleable. Electric utilities tend to be risk-averse due to the high consequential costs of equipment failure, and so even though the manufacturer may be convinced that the technology is viable, he must be able to persuade the customer of this. Once again, this favours the use of established technology, although this does not necessarily require previous experience in the particular field of HVDC. Reference could be made to good experience in other fields.

A1.4 DISCUSSION OF INSULATING AND COOLING MEDIA.

The insulating and cooling media to be considered can be grouped into the following classes:

Liquid Dielectrics:

Natural liquids:
- Deionised water.
- Napthenic oil
- Paraffinic oil.
- Highly purified oil.

Synthetic liquids:
- Organic compounds:
  - Esters.
  - Hydrocarbons.
- Organic compounds including a foreign atom:
  - Silicones.
  - Chlorinated compounds.
  - Fluorinated compounds.
  - Refrigerants.

Mixtures of liquids: - Various.

Gaseous Dielectrics:

Gases:
- Simple gases.
- Oxide gases
- Electro-negative gases
- Hydrocarbon gases.

Vapours of dielectric liquids: - Various.

Mixtures of gases: - Various.

Vacuum.

This section discusses the properties of these individual media.

a) Liquid Dielectrics.

Natural Liquids.

Deionised Water.

Deionised water is an excellent insulant provided that its purity is maintained. This normally requires continuous processing, which is routinely provided in HVDC converters using water-cooled valves. It has many obvious advantages, being low cost, readily available, environmentally neutral and having excellent cooling properties. Its major disadvantage is in the area of materials compatability; many common metals such as copper, mild steel, tin and aluminium will dissolve in water in the presence of quite low levels of leakage currents. At higher currents, even stainless steel will dissolve. Furthermore, mild steel will rust if there is oxygen present in the water.

Since freezing of the water would have to be avoided, any practical installation would require the use of an additive such as monoethylene glycol. This would have the added advantage on the one hand of preventing the growth of colonies of bacteria which might otherwise occur, but on the other it would lead to a deterioration in the cooling properties of the water due to its effect on the specific heat and film heat transfer coefficient.

All of the difficulties of using water as a coolant in a high voltage environment have been overcome in practical engineering systems. However, using water as the insulant for the whole valve would encounter the same problems but in much greater measure. Although these problems are well understood, the cost of overcoming them might outweigh the benefits of water as an insulant.

Napthenic Oil.

Napthenic mineral oil to BS 148 has for many years been the standard insulant for high-voltage transformers, circuit-breakers and cables. There is therefore a substantial amount of long-term experience in many applications, including those involving dc stress. Thus whilst it has drawbacks in terms of flammability and degradation, these (and the techniques for overcoming them) are well understood. In addition to its excellent dielectric strength (2-3 times better than SF_6 at atmospheric pressure), mineral oil also has reasonable cooling properties, and there is wide-spread experience in its use as a coolant.

One question that is likely to become increasingly important is the long-term availability of this oil. Napthenic oils, so-called because of their relatively high proportion of napthenic carbons compared to paraffinic and aromatic carbons, are only available from a limited number of sources [A1.38], some of which are in politically sensitive areas; the supply of these oils is therefore not guaranteed in the long term. As a precaution against this, research has been carried out into the use of paraffinic oils (see below).

Paraffinic Oil.

The majority of the world's oil reserves are paraffinic, so that there is little concern about the long-term availability of this type of oil. The problem with their use for insulation is that they have a cloud-point, and that it is even more difficult than with napthenic oils to obtain a satisfactory compromise between oxidation stability and stability under electrical stress and ionization [A1.39].

"Clouding" happens when the paraffinic carbons crystallize, and this occurs naturally at not much below 0°C. The effect of clouding is to considerably increase the viscosity of the oil, which makes it unacceptable as a coolant. By refining, the cloud-point can be reduced fairly easily to around -20°C; however, to achieve values lower than this (as would be necessary in many parts of the world) an additive is required to lower the cloud point [A1.38]. It is only recently that the long-term performance of oils with such additives has been investigated, and consequently there is only limited field experience.

The compromise on stability mentioned above arises from the fact that it is desirable that the oil absorbs any gas that may be evolved during ionization, so that complete breakdown can be prevented. However, if the oil absorbs gas well, it will also absorb any oxygen it comes in contact with, which will degrade the oil due to the ensuing oxidation of its components.

For the reasons discussed above, it is likely that in the long term paraffinic oils (suitably modified with additives) will replace napthenic oils as the standard electrical insulating liquid. It is therefore probably safe to assume that if transformer oil is chosen for use in the outdoor valve, a suitable type of oil will always be available.

Highly Purified Oil.

By refining and purifying mineral oil, it is possible to considerably improve its fire-resistant properties (eg improving the fire-point from 145°C for BS 148 oil to 300°C). The resulting liquid is known as High Molecular Weight Hydrocarbon (HMWH) oil, and it can be used as a direct replacement for transformer oil. The draw-backs of HMWH oil are its cost, due to the extra processing required, and its increased viscosity due to the molecular structure of its constituents [A1.40]. It has a pour-point of around -24°C compared to -40°C for BS 148 oil.

Synthetic Liquids.

All of the liquids in this category are significantly more expensive than mineral oil due to the manufacturing processes required, since many of them are energy-intensive. Much of the information in this section is derived from reference [A1.36].

Organic Compounds.

Esters.

Esters are oily liquids which can be made with exact proportions of paraffinic or aromatic characteristics; this enables them to be "tailor-made" for a particular insulation system while avoiding the flammability problems of mineral oil. They can also have very good stability and electrical characteristics, but are typically 3 to 4 times the cost of mineral oil.

Hydrocarbons.

Hydrocarbon liquids can be manufactured by the polymerization of suitable gases under very high temperature and/or pressure. Some types can have extremely good stability and heat-resistant characteristics, but costs are similar to those of esters.

Organic Compounds Incorporating a Foreign Atom.

Silicones.

The inclusion of a silicon atom in a suitable organic compound can yield a dielectric that has good insulating and fire-resistant properties. Thus silicone liquid is commonly used as a fire-resistant insulant for small and medium-sized power transformers. Its main disadvantage is poor heat transfer properties arising from a relatively high viscosity and, in particular, a slow fall in viscosity with temperature. This property can lead to thermal runaway in some designs of equipment.

Chlorinated Compounds.

The most well-known members of this group are the polychlorinated biphenyls (PCB's) which are relatively cheap (twice the cost of mineral oil), very fire-resistant and stable. However, their toxicity and tendency for bio-accumulation has led to them being banned in many countries. Other chlorinated compounds such as C_2HCl_3 and C_2Cl_4 have similar attractions to pcb's, but problems of stability, materials compatability and environmental effects would need thorough investigation.

Fluorinated Compounds.

Some fluorocarbons have excellent insulation and fire-resistant properties, but are very expensive at present (40 to 50 times the cost of transformer oil). There are also questions about the nature and toxicity of the gases produced on discharge.

Refrigerants.

The main attraction of refrigerants (such as Freon 113) is in their use as evaporative coolants; the refrigerant can be selected to boil at a temperature consistent with achieving acceptable life for solid insulating components. The presence of the electro-negative chlorine or fluorine gives them excellent discharge suppression properties, by the same mechanism as does the fluorine in SF_6. However, the presence of these constituents also gives rise to the same problems with discharge products, ie the toxicity and chemical aggressiveness of the hydrochloric/hydrofluoric acid vapours produced in the presence of moisture. They are also fairly expensive compared to mineral oil, and have recently been the subject of debate over their effects on the ozone layer. For this latter reason, it is possible that they may eventually be banned in some countries.

Mixtures of Liquids.

Given the number of liquid insulants presently in use, the number of mixtures that could be evaluated is enormous. However, by understanding the mechanisms that give rise to desirable properties, the majority of mixtures can be discarded as not offering any worthwhile

advantages. Some of the relatively few mixtures that have been evaluated are discussed below.

Strictly speaking, mineral oil should be considered under this heading, since it is a complex mixture of hydrocarbon liquids. However, in evaluation and use it is generally treated as a single entity rather than a mixture. The same also applies to some commercial synthetic fluids, such as the ester-based MIDEL's.

The general philosophy of mixing is to improve the properties of a low-cost fluid by means of additives. A mixture of two expensive liquids would have to have exceptional properties to be worth investigating. Of the liquids discussed above, only three types are sufficiently inexpensive to form the basis of a viable mixture: water, mineral oil, and some of the chlorinated compounds, in particular C_2Cl_4. These are considered below.

As discussed earlier, the major problem with deionised water is that of freezing. This can be overcome by adding monoethylene glycol in relatively high proportions (up to 60% for some applications), at the expense of reduced cooling efficiency. The problem of dissolving of metals is more intrinsic to the liquid, and cannot be overcome by the conventional use of additives.

The use of additives to inhibit oxidation in mineral oil is long-established. Additives are also used to remove acids formed by decomposition of the oil, and to depress the cloud-point when necessary. Unless the more volatile components of the oil are removed by processing, however, it is not possible to solve the problem of flammability with only low levels of additives.

The chlorinated compound C_2Cl_4 is less than twice the cost of mineral oil, with comparable dielectric strength. Its main use is as an industrial cleaning agent, giving rise to its relatively low cost. The disadvantages of its use as an insulant are that it freezes at -23°C. and it has a strong solvent action on materials. However, it happens that a 50/50 mixture of this liquid with transformer oil overcomes not only the freezing and solvent problems of C_2Cl_4, but also the flammability of the oil. This mixture is therefore quite attractive as an oil substitute, and it is being commercially exploited by Westinghouse under the trade name "Wecosol".

The behaviour of mixtures is not always easy to predict. For example, experiments have been carried out on transformer oil containing a high level of dissolved SF_6 [A1.41]. It was thought that the electronegative properties of the SF_6 would inhibit breakdown in the oil; however, in practice the field intensification caused by the SF_6 ions at the anode resulted in a reduction in electric strength. This illustrates the complexities that can be encountered when studying mixtures.

The only viable mixtures of practical interest are water/glycol as a coolant and oil/C_2Cl_4 as an insulant or coolant. Indeed, it is unlikely that other mixtures will be discovered that are both better and cheaper than these two; however, it is quite possible that mixtures with different combinations of properties may come to light which are more desirable for particular applications.

b) Gaseous Dielectrics.

Gases.

Simple Gases.

Due to their natural occurrence, gases in this category are all relatively cheap. However, none of them offer better dielectric strength than air, and so they are not really contenders for this application. A brief discussion of these gases is given below for completeness.

As one would expect, nitrogen has a dielectric strength virtually equal to that of air. It is highly stable, inert and abundant, and is of most interest as an additive to SF_6 to reduce the minimum operating temperature.

Hydrogen is used as a coolant in turbo-generators due to its low viscosity and good heat capacity, but its dielectric strength is only about half that of air. Safety is a major consideration in its use, since concentrations between 6% and 67% in air are highly explosive.

Helium is a very poor insulant (less than 10% of the dielectric strength of air) and is of interest only in cryogenic applications, where its dielectric strength in the liquid phase is substantially better than air.

Oxygen has a dielectric strength approaching that of air. However, it is not used as an insulant, presumably because of the consequential fire hazard and the problems of oxidation.

Oxide Gases.

Carbon dioxide is an inert gas that has been used (pressurised) as the insulant in high-voltage, low-loss capacitors. Its dielectric strength is somewhat less than that of air.

Sulphur dioxide falls in this category, but is not used as an insulant.

Electro-Negative Gases.

Electro-negative gases generally have good dielectric strength because they allow the formation of a stable corona layer around areas of high electric stress. This "stress-relieving" mechanism makes them tolerant of features such as poor surface finish on electrodes.

SF_6 is by far the most significant gas in this category. It was originally applied as the insulating medium in switchgear, since its electro-negative properties and good thermal conductivity make it an excellent arc-quenching medium. However, its good dielectric strength (2 to 3 times that of air) meant that its use then spread to metal-clad buswork, where its arc-quenching properties were no longer utilised. Today it is universally used for substations of 132kV and

above, and the break-even voltage above which the technology becomes desirable is gradually falling. Its properties under all conditions are thoroughly documented, and there is world-wide experience in its use.

The Freons have excellent dielectric strength, especially in the liquid phase, but they are not generally used as bulk insulants because all of them undergo the liquid-gas phase change within the range of normal operating temperatures. They also have a strong solvent action, as evidenced by their use as industrial cleaning fluids. Their boiling points make them suitable for two-phase cooling systems, most notably refrigeration equipment; they are also being used increasingly for semiconductor cooling applications. However, recent concern over their effect on the ozone layer places a question-mark over their long-term acceptability for industrial applications.

The gases C_4F_8 and CH_2Cl_2 also come into this category. They both have good dielectric strength, but are expensive due to their being relatively exotic. Their use as insulants is limited.

Hydrocarbon Gases.

Hydrocarbon gases such as methane, ethane, propane, butane and hexane have been investigated as dielectrics, but their high flammability makes them unsuitable for most applications.

Vapours of Dielectric Liquids.

Some types of transformers have been designed to be insulated by a Freon in the gas phase. Under cold start conditions, the Freon is predominantly liquid, but as the equipment warms up, the vapour pressure of the Freon increases until the full dielectric strength is achieved. SF_6 is usually added to provide some degree of insulation under cold start conditions.

The potentially most interesting item in this category is the vapour mist system being developed by EPRI and Westinghouse [A1.27]. This was first developed for transformer insulation, but could be applied to many other types of equipment including HVDC valves. "Climatic control" will obviously be required to ensure that the mist is maintained.

Mixtures of Gases.

As with liquids, the usual purpose of mixing gases is to take a low-cost "base" gas and improve its properties by means of additives. The most well-known example of this is the addition of a small amount of SF_6 to air; the strong electro-negative properties of SF_6 mean that a low concentration of SF_6 (eg 5%) produces a large increase in dielectric strength. Further additions of SF_6 produce progressively smaller improvements in dielectric strength, until eventually the conventional breakdown strength of SF_6 is achieved at 100% concentration.

The use of a mixture of SF_6 and N_2 represents a slightly different approach. Nitrogen is a poor insulant compared to SF_6; however,

pressurised SF_6 liquefies at inconveniently high temperatures, and the addition of some (low-cost) N_2 gives a useful overall improvement in low-temperature performance. The dielectric strength of SF_6 can also be improved by the addition of some fluorocarbons and other gases, but only limited evaluation of these mixtures has been carried out.

c) Vacuum.

Vacuum insulation is widely used in circuit-breakers for low to medium voltages. However, providing such insulation for a complete valve would be a significantly more complex exercise, due to the need to avoid leaks in a much larger structure. Whereas leaks in pressurised equipment result in some loss of operating pressure, there is no automatic contamination of the equipment. However, leaks in a vacuum system result in the ingress of uncontrolled foreign materials, which is intrinsically more serious.

Most high-power transformer tanks built today are capable of withstanding some degree of vacuum, so that the unit can be evacuated in the later stages of manufacture to speed up the removal of moisture absorbed in the paper insulation. However, it is likely that substantial development would be required before such tanks (similar to those which would be needed for an outdoor valve) could maintain vacuum conditions over life. All of the valve components would also need to be evaluated for their suitability for use in vacuum.

A1.5 SELECTION OF THE INSULATING MEDIA.

A rigorous appraisal of the costs and benefits of the different insulating and cooling media has not been carried out. However, in compiling the information presented so far, the following general impressions have been gained:

- Transformer oil is a very good insulator. However, its disadvantages (mainly flammability and the need for reprocessing) mean that considerable effort has been expended over many years to find alternatives to it. To develop a new generation of equipment (such as the outdoor valve) insulated by transformer oil would be going against this trend in the industry. The main problem with alternatives to oil has been to achieve adequate cooling efficiency at the same time as good insulation. However, an HVDC valve would require a separate cooling system anyway, so that some of these problems would be avoided.

- Gas insulants are significantly cheaper than liquids. During the development phase they are also generally more convenient to work with.

- SF_6 is by far the most widely used gas insulant, and new applications are still being developed. These factors (along with their attendant advantages such as cost, familiarity, availability, service experience etc) make SF_6 a very strong contender as a gas insulant.

- Water is a very good coolant, with its nearest rivals being several times less efficient. This holds true even when glycol is added to prevent freezing. Only a phase-change Freon cooling system can achieve similar levels of heat transfer, with considerably greater complexity. There are therefore strong arguments for using water as a coolant.

In practice, it is not possible to choose the insulant and coolant independently of each other. The possibility of leakage from one into the other (and its consequences) means that insulant/coolant pairs must be evaluated. On the basis of the above impressions, the ideal combination would be to use SF_6 as the insulant and water as the coolant. This combines the proven advantages of SF_6 as the industry-standard gas insulant with the excellent cooling performance of water. However, this system does suffer from the hazardous by-products that are produced if an arc occurs in SF_6 in the presence of moisture. In particular, problems would occur if a water leak caused a flash-over. The number of flash-overs that have occurred by this mechanism in commercial air-insulated, water-cooled installations has been low; however, this may be due to the relatively low field stresses that are used with air insulation. Making use of the properties of the SF_6 by going to higher field stresses may make small leaks more serious. It may be possible to circumvent these disadvantages by careful design; for example, one could ensure that leaking water was guided along regions of low field stress. Components vulnerable to hydro-fluoric acid could be positioned away from regions where water might leak, or shielded or coated to protect them from moisture. An extreme measure would be to ensure that the gas was always at a higher pressure than

the coolant, so that leaks would result in gas dissolving into the water system rather than vice-versa. In addition, it would be necessary to persuade the customer that satisfactory measures had been taken. In practice, it could be argued that a flash-over would cause extensive damage inside the valve tank, and therefore dealing with any breakdown products might be regarded as fairly straightforward compared to other aspects of the repair process. Whatever the difficulties might be, the advantages of using water are substantial.

If the problems associated with the SF_6/water combination proved insurmountable, then the next best approach would be to use SF_6 with a compatible coolant. Freon is the most obvious choice, having already been investigated in some depth. There are other possibilities; the main requirement is that the coolant does not contain hydrogen which could react with fluorine in the event of an arc. A detailed engineering evaluation of the various possibilities would be needed to choose a candidate for development.

A1.6 CONCLUSIONS.

The criteria for choosing an insulant and coolant for the outdoor valve have been discussed, and the wide range of possible fluids has been considered. It is not possible to choose an insulant or coolant on the basis of one or two simple parameters such as dielectric strength or cost; instead, one must take into account all the engineering implications of using a particular combination of insulant and coolant in an outdoor valve, such as life-time cost, materials compatability, environmental constraints, customer acceptability and so on. From the limited survey carried out, the conclusion is that SF_6 in combination with either water or a more compatible coolant is to be preferred. A more detailed evaluation would be needed to justify these conclusions.

A1.7 REFERENCES.

A1.1 J Weaver et al, " Development of a thyristor valve for HVDC transmission." IEE Conference on Power Thyristors and Their Application, IEE Conference Publication No 53, Part 1, 1969, pp 339-346.

A1.2 M Danders et al, "Results of the type tests and production test procedure for HVDC thyristor valves." Direct Current, Vol 1, No 4, 1970, p 133.

A1.3 J Hengsberger et al, "An oil-cooled HVDC thyristor valve for outdoor installation." CIGRE conference, Paris, 1970, paper 14-04.

A1.4 V Beriger et al, "Cabora-Bassa HVDC transmission: oil-cooled outdoor thyristor valves." IEEE Transactions on Power Apparatus and Systems, Vol PAS-94, No 6, Nov-Dec 1974, pp 1061-1072.

A1.5 E Raynham et al, "Apollo inverter station HVDC operating experience." CIGRE conference, Paris, 1982, paper 14-05.

A1.6 K Yoshida et al, "Development and field tests of 125kV oil-immersed thyristor valves for HVDC." CIGRE conference, Paris, 1974, paper 14-01.

A1.7 M Yasuda et al, "Shin-Shinano frequency converting station". CIGRE conference, Paris, 1978, paper 14-02.

A1.8 M Yasuda et al, "Outline of the Shin-Shinano frequency converting station of the Tokyo Electric Power Co, Inc." Hitachi Review, Vol 28 (1979), No 3.

A1.9 T Horiuchi, "300MW Frequency converter." Toshiba Review, No 119, Jan-Feb 1979.

A1.10 N Hingorani et al, "Compacting DC terminals." EPRI Journal, Vol 1, No 10, 1976, pp 18-22.

A1.11 N Hingorani et al, "Compact gas insulated HVDC terminals - EPRI prototype DC link." CIGRE conference, Paris, 1978, paper 14-06.

A1.12 H Jaster et al, "Dielectric behaviour of the coolant for a gas-insulated HVDC thyristor valve." IEEE Transactions on Power Apparatus and Systems, Vol PAS-97, No 3, May-Jun 1978, pp 828-837.

A1.13 R Nakata, "Controlled particle scavenging technique for use of HVDC SF_6 gas bus." IEEE Conference Paper A79 425-0, IEEE Publication No 76 CH1135-3-PWR.

A1.14 R Nakata et al, "Practical gas bus design based on particle motion studies." IEEE Conference Paper A77 198-5, IEEE Publication No 77 CH1190-8-PWR.

A1.15 F Fischer et al, "HVDC link to test compact terminals." Electrical World, 1 Feb 1977, pp40-42.

A1.16 J Douglass, "New solid-state valves for HVDC." IEEE Power Engineering Review, Sept 1983, pp 4-7.

A1.17 R Scaringe et al, "Forced-vaporization cooling of HVDC thyristor valves." EPRI Research Project 1207, Final Report EL-2710, October 1982.

A1.18 J Fink, "Solid-state valve cooling techniques." Electric Forum (General Electric), Vol 9, No 3, 1984, pp 43-44.

A1.19 B Damsky et al, "Design, Installation and Initial Operating Experience of an Advanced Thyristor Valve", IEEE Transactions on Power Delivery, Vol PWRD-2, No 4, Oct 1987, pp1103-1109.

A1.20 E Ohno et al, "Development of oil-immersed HVDC thyristor valves." Mitsubishi Denki Giho, Vol 13, No 1-4, 1972, pp 77-96.

A1.21 S Matsumara et al, "New approach to SF_6 gas insulated HVDC converter station." IEEE Transactions on Power Apparatus and Systems, Vol PAS-102, No 9, Sept 1983, pp 2871-2880.

A1.22 S Matsumara et al, "Development of an HVDC thyristor valve insulated and cooled by compressed SF_6 gas." IEEE Transactions on Power Apparatus and Systems, Vol PAS-102, No 9, Sept 1983, pp 3243-3253.

A1.23 D Duckett, "Environmentally acceptable insulating fluids may replace Askarel." Paper presented at the General Meeting of the Edison Electric Institute, Transmission and Distribution Committee, Minneapolis, Minnesota, 8 May 1975.

A1.24 R Popple, "'MIDEL' insulating fluids." GEC Report S/NKu 735, 10 Aug 1979.

A1.25 C Moore et al, "Recent developments in gas/vapor and fire-resistant transformers." Proceedings of the IEEE Power Engineering Society Transmission and Distribution Conference and Exposition, 1-6 Apr 1979, pp 59-67.

A1.26 J Morris et al, "A new concept for a compressed gas-insulated power transformer."Proceedings of the IEEE Power Engineering Society Transmission and Distribution Conference and Exposition, 1-6 Apr 1979, pp 178-183.

A1.27 "EPRI, Westinghouse receive I-R 100 award for developing new gas insulation process." IEEE Power Engineering Review, Feb 1985.

A1.28 A Cookson, "Compressed gas insulated transmission systems: the present and the future." Proceedings of the International Symposium on Gaseous Dielectrics, Knoxville, Tennessee, USA, 1978, pp 286-308.

A1.29 N Malik et al, "A review of electrical breakdown in mixtures of SF_6 and other gases." IEEE Transactions on Electrical Insulation, Vol EI-14, No 1, Feb 1979, pp 1-13.

A1.30 A Buttkau et al, "Siemens SF_6/N_2 circuit-breaker for service at low temperatures." Siemens Power Engineering VI (1984), No 1, pp 30-35.

A1.31 R James et al, "Insulation problems in transformers and reactors for DC transmission." IEE Conference Publication No 22, pp 339-346.

A1.32 R Alvinsson et al, "The world's first GIS for HVDC applications in commercial operation." Paper presented at the United Nations Economic Commission for Europe, Seminar on High Voltage Direct Current (HVDC) Techniques, Stockholm, May 1985.

A1.33 K Nakanishi et al, "Surface charging on epoxy spacer at DC stress in compressed SF_6 gas." IEEE Transactions on Power Apparatus and Systems, Vol PAS-102, No 12, Dec 1983.

A1.34 C Cooke et al, "Compressed gas insulation for advanced HVDC transmission equipment." CIGRE conference, Paris, 1982, paper 15-14.

A1.35 I Adamezewski, "Ionization, conductivity and breakdown in dielectric liquids." Published by Taylor and Francis, London, 1969.

A1.36 A Bradwell (ed), "Electrical insulation." Published by Peter Peregrinus, London, 1983. Chapters 1, 2 and 7.

A1.37 J Bastian, "Evaluation criteria for distribution transformers at fire endangered locations." Elektrizitaetswirtschaft (Germany), Vol 83, No 20, Sept 1984, pp 878-880. (In German).

A1.38 B Fallou et al, "Assessment of mineral insulating oils from paraffinic crudes." CIGRE conference, Paris, 1982, paper 15-11.

A1.39 A Mills et al, "Mineral insulating oil: its present and future role." IEE Colloquim, 6 Jan 1976.

A1.40 S Northrup et al, "Cold start performance of transformers filled with High Molecular Weight Hydrocarbon liquid." IEEE Transactions on Power Apparatus and Systems, Vol PAS-103, No 11, Nov 1984, pp 3373-3378.

A1.41 A Nosseir et al, "Effect of dissolved sulphur hexafluoride on electrical breakdown in mineral oil." IEEE Proceedings, Vol 113, No 2, Feb 1966, pp 359-363.

APPENDIX TWO: VALVE MECHANICAL DESIGN.

A2.1 INTRODUCTION.

Appendix One considered the insulating and cooling media that would be used for an outdoor valve, and this Appendix goes on to look at the way in which an outdoor valve would be implemented in practice. The features to be considered here include the valve structure, the number of valves per tank, and the provisions for maintenance and replacement of valves. These are generally treated as the mechanical aspects of the valve design.

The mechanical design of the valve is important in two areas. Firstly, it has substantial influence on the cost of the converter station, both in terms of the direct capital cost and the cost of time lost during maintenance. Secondly, it determines the seismic capability of the valves, which is usually specified by the customer, and can be an important design constraint in some parts of the world.

The mechanical structure of the valve itself generally comprises all parts of the valve except the electrical components and conductors, the cooling system and the fibre-optic system. However, it is possible for some components to perform dual functions, such as the central air-duct of an air-cooled valve also supporting the valve modules. The particular mechanical design of the valve will depend on a number of factors, including the following:

- The number of thyristor levels in the valve;
- Whether or not the valves are to be stacked;
- The seismic requirements for the installation;
- The cooling system chosen (in particular, whether liquid or gas);
- The insulation system chosen;
- The maintenance philosophy.

With these general considerations in mind, the following section reviews the mechanical designs that have been employed in existing outdoor valve installations.

A2.2 TECHNOLOGY REVIEW.

Existing valve mechanical designs fall naturally into two classes: those that are gas-cooled (normally with air) and those that are liquid-cooled. In a gas-cooled valve, the volume of gas that must be circulated to achieve the required thermal mass flow rate means that a large duct is required somewhere in the system. It is often convenient to put this duct in the centre of the valve and arrange the heat-dissipating components round it; gas can then either be sucked or blown across them [A2.1]. Apart from the HVDC Working Group, all HVDC manufacturers have built valves to this design. However, liquid cooling offers a number of advantages over gas cooling, an important one being that the same thermal mass flow rate can be achieved in a much more compact system due to the greater heat capacity per unit volume of liquids. Eliminating large ducts and transporting the coolant in flexible pipes allows much more freedom in the mechanical design of the valve. This freedom can be used to optimise the valve against other criteria such as compactness, ease of maintenance and seismic withstand capability. All manufacturers therefore have moved or are moving to liquid-cooled systems; the only exception to this is Mitsubishi, whose high-pressure tanked valve design uses circulating SF_6 to cool the valve [A2.2].

In order to minimise the area of the valve hall, indoor valves are usually stacked four high. This arrangement, known as a "quadrivalve", is convenient for electrical connections and clearances, but it often results in a structure that is tall and thin, and therefore vulnerable to seismic damage. To avoid this problem, ASEA have designed a flexible quadrivalve which hangs from the ceiling of the valve hall [A2.3]. The valves can move in a controlled manner in response to seismic loading, and are therefore intrinsically more tolerant to earthquakes than rigid structures.

The maintenance philosophy chosen for an installation has a substantial impact on the mechanical design of the valve. Three different approaches have been adopted to valve maintenance:

- Valve replacement;
- Module replacement;
- Component replacement.

Valve replacement was universally employed on mercury arc schemes, since the converter could not be out of service for the length of time it took to carry out routine maintenance on a valve. The valves were mounted on wheeled platforms to facilitate their movement between the valve hall and the maintenance area. The same philosophy was also employed in the tanked-valve schemes at Cabora-Bassa and Shin-Shinano, due to the length of time needed to drain, repair and refill a valve if maintenance was required. Elaborate valve transporters were needed to move the tanked valves from their outdoor location to the maintenance building.

One of the major attractions of indoor air-insulated valves is that there is easy access for maintenance, so that it is feasible to maintain the valve *in situ*. This avoids the substantial cost of a complete spare valve. The type of *in situ* repair that can be carried

out depends on the design of the valve and the ease with which faults can be located; a widely-used approach is modular replacement, whereby a module contains one or more thyristor levels, and is the smallest replaceable item in the valve. The faulty module is located while the valve is operating, and at the next shutdown it is replaced and taken away for diagnosis and repair. The advantage of this approach is that only a small number of spare modules are required, rather than a complete valve. If a quadrivalve arrangement is used, this effectively dictates *in situ* maintenance, since individual valves are no longer removable.

Further advantages are gained by using *in situ* replacement of components. If modules are used, the valve must have two independent structures: a module structure to support the components, and a valve structure to support the modules. If there are no removable modules, the whole valve structure can be integrated and simplified. However, this does require that, during a maintenance period, it is possible to locate individual failed components rather than just a failed module. This can be made easier by the use of thyristor-level status monitors which indicate if the level is faulty whilst the valve is operating, and also by careful valve layout and the use of dedicated test equipment. Component replacement gives the minimum possible cost of spares holding.

For an outdoor valve, the number of valves per tank is an important consideration. Cabora-Bassa has two valves per tank, whereas Shin-Shinano and the new Mitsubishi design have only one. Increasing the number of valves per tank reduces the number of bushings required (ie interfaces between the tank internal insulation system and the outside world), which are an important source of unreliability in outdoor installations. Figure A2.1 shows the number of bushings required for a group of 12 valves versus the number of valves per tank. However, increasing the number of valves per tank also incurs the following disadvantages:

- The tank becomes larger and heavier, making handling and shipping more difficult;
- Achieving a leak-tight tank structure becomes more difficult as the tank becomes bigger;
- If maintenance is based on tank replacement, then the cost of the spare tank increases.

A final consideration in the design of an outdoor valve is whether the tank, if made of a conducting material, should be at earth potential ("dead") or at some other potential in the circuit. For Cabora-Bassa, the tank formed one of the connections to the valves inside, thus saving one bushing per tank. However, this means that the tank has to be stood on post insulators and fenced off, resulting in poor seismic capability and increased land area requirements. The use of the more compact "dead-tank" approach has become standard in the industry for conventional metal-clad equipment.

A2.3 DESIGN CONSIDERATIONS.

The technology review in the previous section points to a number of factors which must be taken into account in the mechanical design of an outdoor valve, as follows:

- Tank material;
- Connections between tanks;
- Maintenance philosophy;
- Number of valves per tank;
- External environment;
- Fault monitoring;
- Valve mechanical structure.

These factors will now be considered in more detail below.

a) Tank Material.

Steel is almost universally used for the construction of tanks for "metal-clad" electrical equipment (hence the name). However, the use of an engineering plastic of some kind does have some attractions, including:

- Lower weight;
- Reduced stray capacitances between the valve and the tank;
- Easier design of bushings;
- General insulating properties.

A large steel tank could weigh several tens of tons, and therefore the total weight of the tanked unit could be considerably reduced by the use of a plastic tank. However, unless there are specific transport limitations, the actual cost saving (in site foundations) from this would be small. The use of a steel tank would result in high stray capacitances between the valve components and earth (particularly if a liquid insulant was used), which is undesirable for the valve electrical design, whereas this would not be the case for a non-conducting tank material.

Bushing design would also be simpler with a plastic tank. This is because for a conventional structure, the high-voltage connections have to pass through a solidly-earthed plate, resulting in high field stresses and therefore a complicated (and potentially unreliable) insulation system. With an insulating tank, these problems would be considerably alleviated.

Finally, if the tank is insulating, it may be possible to reduce internal electrical clearances compared to a metal design, since the tank surface could float at some intermediate potential rather than being solidly earthed.

To have the necessary strength, any non-metallic material would have to be a fibre-reinforced composite; however, a steel base-frame would still be necessary to enable the tank to be lifted. If a cylindrical tank was used, then it could be possible to make it as one part; cylinders up to 6m in diameter can be produced in one piece on a

mandrel. A non-cylindrical tank would probably have to be fabricated from sections.

The last three advantages discussed above (lower stray capacitance, easier bushing design and reduced internal clearances) all arise from the fact that the tank walls are not earthed conductors, and therefore electric and magnetic fields can penetrate the walls. However, this then encounters the problems faced by a live tank, namely that the tank must be supported on insulators and fenced off. Furthermore, surface charge would build up on the tank, which would increase field stresses in the event of a polarity reversal. The surface of the tank would have to be shedded or protected from the environment to prevent surface tracking in regions of high field stress. Finally, the tank would not screen radio interference from the valves, so this would have to be provided separately. It is believed that these disadvantages would outweigh the benefits above, and therefore a conductive tank is to be preferred. Steel is the economic choice of material for such a tank.

b) Connections Between Tanks.

The method by which connections are made to the tank also includes the question of whether the tank should be "live" or "dead".

High-voltage bushings to the open air are bulky (up to 8m long for a 500kV DC bushing), expensive and failure-prone. In some environments, regular cleaning is required to remove pollution. Given that the converter transformers and the valves are in metal tanks, it is clearly undesirable to bring a connection out from the transformer tank through one bushing, and then back into the valve tank through another. Using metal-clad buswork would avoid the substantial drawbacks of these bushings. The use of SF_6-insulated buswork is now standard practice for high-voltage AC installations, and the technology has already been applied to HVDC converters (See Appendix 1, section A1.2).

If metal-clad interconnections were used, this would effectively preclude the use of a live-tank design. There are two main advantages to a live-tank design:

- the number of bushings is reduced if the tank is used as a connection;
- the insulation requirements inside the valve tanks are minimised. Allowing the tank to be live means that for some valves in a converter, part of their insulation can be provided by low-cost post insulators outside the tank. When a dead-tank is used, all the insulation has to be provided inside the tank.

If metal-clad buswork was used, then the first advantage largely disappears, since the open-air bushings are not needed. The second advantage is also lost, but it could be argued that an insulation system in the protected environment inside the tank will be more reliable than one exposed to the atmosphere. Given the strong industry preference for dead-tank equipment, it is recommended that this approach should be adopted. This gives the advantages of compactness, safety, protection from lightning strikes and shielding of radio interference.

The insulant used in the buswork could be chosen independently of that in the valves and transformers, since interfaces between common insulants (eg oil and SF_6) are standard practice. The choice between single-phase and three-phase conductor systems would be determined by detailed mechanical and economic analysis of a particular design.

c) Maintenance Philosophy.

This aspect will be a key feature of the design and marketing of an outdoor valve. The maintenance strategy must have the following features:

- A low requirement for maintenance (demonstrated by calculation and preferably by service experience as well);
- Relatively easy maintenance when it is required;
- A fault monitoring system to indicate the status of the valve.

If sufficient good service experience is gained, this last requirement could be dropped, but it will almost certainly be needed at first.

The use of light-triggered, self-protecting thyristors is an essential requirement for achieving the levels of reliability that will be needed. It has also been industry practice that, for radical innovations, a prototype valve has been installed in an existing HVDC converter in order to gain service experience, and this would almost certainly be required for an outdoor valve.

It is inevitable that over the 25-35 year life now being specified for HVDC converters, some valve maintenance will be required. This must therefore be made as easy as possible, subject to the economic trade-offs involved. The maintenance philosophies discussed in section A2.2 (valve, module or component replacement) are considered below.

A valve replacement strategy (or tank replacement, if there is more than one valve per tank), is attractive in that converter operation can continue while a faulty valve is being repaired. Since the tank will normally be a sealed insulation system, gaining access for repair and then resealing the tank afterwards will take longer than the corresponding operations on an indoor valve. Valve replacement removes the pressure for rapid access, diagnosis, repair and recommissioning. However, this must be set against the substantial disadvantages of this approach, namely:

- The high cost of the spare valve;
- The difficulty in transporting complete valve tanks around the converter site;
- The need for mechanical and electrical interchangeability of valve tanks.

To minimise the cost of the spare tank, it will be necessary to have as few valves per tank as possible, but this will increase the number of connections between tanks, with the resulting effects on cost and reliability.

In practice, with good diagnostic systems, it may be quicker to enter the valve tank and repair it *in situ* than to disconnect the tank and

replace it with the spare. Informal discussions suggest that this is now the practice adopted at Cabora-Bassa, even though the scheme was originally designed with valve replacement in mind.

If the valves are repaired *in situ*, then this allows maintenance to be carried out on more than one tank at a time. This may be necessary if a "wear-out" mechanism is encountered, whereby significant numbers of a component could fail in a short space of time. Because of this and the other factors listed above, a valve replacement strategy does not seem attractive.

If repairs are to be carried out inside the tank, this has implications for the mechanical layout of equipment. For example, it may be necessary to leave larger clearances for access than are necessary for electrical reasons. However, the cost implications of this would not be substantial unless a very expensive insulant was being used. There will also be the need for strict cleanliness during maintenance, and this will be more difficult to achieve on site than in a purpose-built repair facility. A temporary tent could be erected over the tank being repaired, or alternatively the valves could be housed in a permanent (but low-cost) "barn" similar to those often used for metal-clad switch-gear. The degree of cleanliness needed will depend on the insulation system being used; a high-pressure SF_6 system using forced circulation of the gas (as in the Mitsubishi design [A2.2]) will be much more sensitive to contamination than a low-pressure system using a separate coolant.

Another implication of *in situ* maintenance is that replaceable items should be kept to a size and weight that can be conveniently handled by one person. If lifting gear is required for maintenance (as is the case with conventional modular replacement), this will considerably complicate the overall mechanical design of the system. Where a separate liquid coolant is used, it is a definite advantage to be able to replace a thyristor without disturbing any cooling circuit connections, and therefore component-level replacement (at least for actively cooled components) is preferable in this case.

The maintenance philosophy recommended for an outdoor valve is therefore *in situ* repair based on component replacement.

d) Number of Valves Per Tank.

The considerations relating to the number of valves per tank are as follows:

- The number of connections between tanks;
- Electrical operation;
- Cross-contamination of equipment;
- Whether the tanks are to be shipped complete or in pieces;
- The mechanical problems of constructing a large, leak-free tank;
- The thermal expansion and distortion of a large tank;
- The scheme compactness required.

These points are discussed below.

Firstly, it should be noted that the number of interconnections between

tanks falls as the number of valves in a tank goes up, giving the advantages already discussed in section A2.2.

Regarding electrical operation, all schemes built today are designed only for 12-pulse operation, rather than the 6-pulse operation used in early schemes. Since failure of any one valve will therefore require that a whole 12-pulse group be taken out of service, there is no operational disadvantage in having all twelve valves in one tank. For a back-to-back link, it would be acceptable to have a rectifier and inverter group (24 valves in total) in one tank, thus eliminating all DC bushings.

If three-phase transformers are used (rather than separate single-phase units), then having six valves per tank is a convenient arrangement, since one valve tank will be associated with one transformer tank.

The potential problems of cross-contamination will depend on the insulant and coolant used. If a flashover leads to undesirable by-products, then the more equipment that is located inside one tank, the more work will be required in decontaminating the tank.

There are several mechanical considerations which will affect the largest size of tank that can be used in practice. There are important advantages in shipping tanks as complete sealed units (for reasons of cleanliness in assembly, ease of commissioning etc), but there are clear limits in the size of unit that can be shipped. Achieving freedom from leaks will also be inherently more difficult as the tank size increases. The effects of thermal expansion become more pronounced as size increases, although the use of a "barn" to shield from direct sunlight would prevent uneven heating from solar gain. Finally, if space is at a premium, then increasing the number of valves per tank will reduce the amount of interconnecting buswork and therefore the site area required.

Generally speaking, therefore, it is desirable to have as many valves in one tank as possible (up to a maximum of 12 or 24 for a rectifier/inverter unit) so long as the tank size remains manageable. The actual number of valves per tank for a scheme will depend mainly on the number of levels per valve, as well as the other factors discussed above.

e) External Environment.

If valves are to be maintained *in situ*, then some means will be necessary to stop contamination of the valve tank when it is opened. If the valves are located outdoors, a temporary or permanent "air-lock" will be required to create a clean environment around the tank opening. Alternatively, a permanent barn structure over the whole converter would be convenient, although possibly more costly. The approach used will depend on the climate and environment at the converter site.

f) Fault Monitoring.

The increased difficulty of inspection for a tanked valve means that, initially at least, some form of automatic fault monitoring will be required. However, there is an intrinsic problem with this: conventio-

nal monitoring systems (indicating the status of each thyristor level) require a fairly significant amount of active electronics, which would be substantially less reliable than the light-triggered thyristor being monitored. This would defeat the object of using LTT's in the first place. To be realistic, therefore, a monitoring system would have to utilise one of the advanced passive sensing technologies now being developed.

Such a sensing system would be fibre-optic based, with either one fibre per thyristor or a distributed system with one fibre covering several devices. The voltage on each level could be sensed with a simple piezo-electric transducer causing micro-bending loss in a fibre. Alternatively, the damping resistor temperature could be used as a surrogate measurement; "too hot" implies repetitive BOD firing, "too cold" implies a short-circuit thyristor. Current in the damping network could be measured, and passive current sensing is already at the laboratory stage. The development of this sensing system would be an important exercise in its own right, with potential commercial application in other areas.

g) Valve Mechanical Structure.

The mechanical structure of the valve is strongly affected by the type of coolant used (gas or liquid), and the maintenance philosophy adopted. Since the recommended design for an outdoor valve does not differ from GEC's approach to indoor valves in these areas, ie liquid cooled with *in situ* component replacement, then the outdoor valve can be very similar to the present indoor design. This has obvious attractions for standardization and minimizing development costs.

The existing layout of a valve tier is determined primarily by mechanical rather than electrical clearances. The spacing between tiers is also dictated by the need for access to change a thyristor It is therefore unlikely that any significant advantage could be taken of a better insulating gas in making the valve more compact. The main gain would be in reducing the clearances between valve stacks (if there is more than one in a tank) and between the valves and the walls of the tank. Even in these areas, however, the need for human access for maintenance may restrict the extent to which clearances can be compressed.

Where advantage is taken of the better insulating properties of the gas in the valve tank, more efficient corona shields may be required to control maximum field stresses. If the valve has to withstand seismic disturbances, the deflection of components inside the tank will also need to be studied to ensure that adequate clearances are maintained.

A2.4 CONCLUSIONS.

Previous practice in the mechanical design of outdoor valves has been reviewed, and the considerations involved discussed. Given that Appendix One recommends a design based on a gas insulant and a separate liquid coolant, the following conclusions have been reached:

- The tank should be made of steel;
- Tanks should be earthed ("dead-tank");
- Metal-clad connections should be used between tanks;
- Maintenance should be carried out by entering the tank and replacing components *in situ;*
- As many valves as possible should be located in one tank, as long as the tank size remains practical, and independent 12-pulse groups are not located in the same tank.
- A "barn" structure over the whole converter offers advantages in maintaining cleanliness and allowing easy access in all weather conditions.
- Fault monitoring using passive optical-fibre sensing should be developed for early installations.
- The existing indoor valve mechanical design could be used without major modification.

These recommendations represent a departure from previous practice in this area, but should offer the best combination of cost, convenience and customer acceptability.

A2.5 REFERENCES.

A2.1. S Goddard et al, "The new 2000MW interconnection between France and the United Kingdom." CIGRE conference, Paris, 1982, paper 14-09.

A2.2. K Morii et al, "Development of totally metal enclosed SF_6 gas insulated HVDC converter station." Electrical Engineering in Japan, Vol 104B, No 1, 1984, pp 40-49.

A2.3. S Berggren et al, "A new concept for mechanical design of thyristor valves for HVDC." CIGRE conference, Paris, 1982, paper 14-07.

Scheme	Transmission Distance (km)	Voltage x No of circuits (kV)	Capacity (MW)	Commissioning Date
Gotland-Swedish Mainland	96	± 150	30	1954/1970
Cross-Channel 1 (GB-F)	65	± 100	160	1961[1]
Volgograd-Donbass (USSR)	470	± 400	720	1962-1965
Konti-Skan (DK-S)	180	250	250	1965
Sakuma (J)	50/60Hz tie	125 x 2	300	1965
New Zealand (NZ)	609	± 250	600	1965
Sardinia-Italian Mainland	385	200	200	1967
Vancouver Pole 1 (CDN)	74	+ 260	312	1968/69
Pacific Intertie (USA)	1362	± 400	1440	1970
Nelson River Bipole 1 (CDN)	890	± 450	1620	1973-77
Kingsnorth (GB)	82	± 266	640	1974

[1] Decommissioned 1984.

Table 1.1: HVDC Schemes in service (Mercury-arc valves).

Scheme	Transmission Distance (km)	Voltage x No of circuits (kV)	Capacity (MW)	Commissioning Date
Eel River (CDN)	As	80 x 2	320	1972
Skagerrak (DK-N)	240	± 250	500	1976/77
David A Hamil (USA)	As	50	100	1977
Cabora Bassa-Apollo (MOC-ZA)	1414	± 533	1920	1977-79
Vancouver Pole 2 (CDN)	74	- 280	370	1977/79
Square Butte (USA)	749	± 250	500	1977
Shin-Shinano (J)	50/60Hz tie	125 x 2	300	1977
Nelson River Bipole 2 (CDN)	930	± 250	900	1978
CU Project (USA)	710	± 400	1000	1979
Hokkaido-Honshu (J)	168	250	300	1979/80
Acaray (PY-BR)	50/60Hz tie	26	50	1981
EPRI Compact (USA)	0.6	100/400	1000	1981[1]
USSR-Finland	As	± 85 x 3	1070	1982
Inga Shaba (Zaire)	1780	± 500	560	1982
Durnrohr (A)	As	± 145	550	1983
Gotland 2-Swedish Mainland	98	150	130	1983
Eddy County (USA)	As	82	200	1985
Itaipu (BR) Bipole 1	783	± 600	3150	1985
Chateaugay (CDN)	As	140	1000	1984
Pacific Intertie Upgrade (USA)	1362	± 500	400	1985
Highgate (USA)	As	56	200	1985
Oklaunion (USA)	As	82	200	1985
Blackwater (USA)	As	56	200	1985
Miles City (USA)	As	82	200	1985
Madawaska (CDN)	As	144	350	1985
Sidney (USA)	As		200	1986
Cross-Channel 2 (GB-F)	72	± 270 x 2	2000	1985/86
Intermountain (USA)	794	± 500	1600	1987
Quebec-New England (CDN-USA)	1000	± 450	2000/2090	1986/92

[1]Completed but not in service.

As = Asynchronous tie

Table 1.2: HVDC Schemes in service (Thyristor valves).

Device	Radiation Dose	Mask
EL300 C22-B1-29	0	None
EL300 C22-B1-33	40 kRad	None
EL300 C22-B1-12	40 kRad	#1
EL300 C22-B1-48	80 kRad	#1

Table 5.3.1: Radiation doses and masks used for the first set of 30mm devices.

Device	Radiation Dose	Mask
EL300 C22-B2-16	23 kRad	None
EL300 C22-B2-29	23 kRad	None
EL300 C22-B2-51	11.5 kRad	#1 on main cathode
EL300 C22-B2-33	11.5 kRad	#1 on main cathode
EL300 C22-B2-62	23 kRad	#1 on main cathode
EL300 C22-B2-04	23 kRad	#2 on satellite well
EL300 C22-B2-25	23 kRad	#2 on satellite well
EL300 C22-B2-05	23 kRad	#3 on satellite well
EL300 C22-B2-13	23 kRad	#3 on satellite well
EL300 C22-B2-02	75 kRad	#2 on satellite well
EL300 C22-B2-03	75 kRad	#2 on satellite well
EL300 C22-B2-08	75 kRad	#2 on satellite well
EL300 C22-B2-14	75 kRad	#3 on satellite well
EL300 C22-B2-07	75 kRad	#3 on satellite well
EL300 C22-B2-10	75 kRad	#3 on satellite well

Table 5.3.2: Radiation doses and masks used for the second set of 30mm devices.

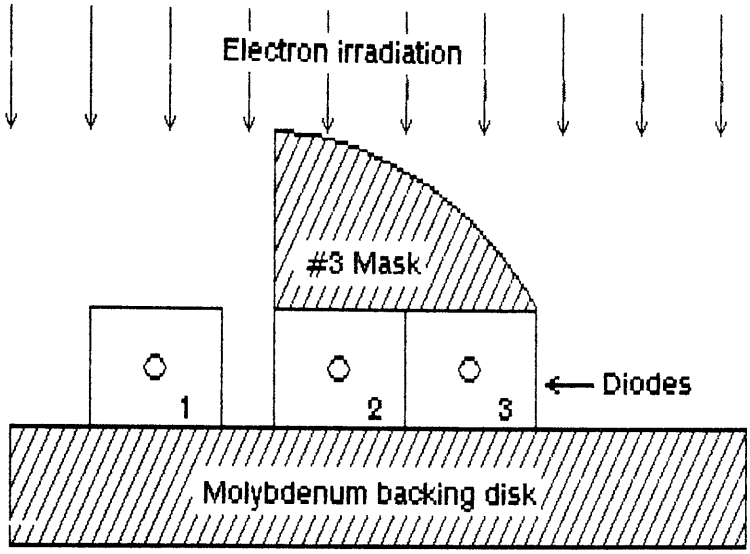

Mask Position	Dose Received
1	75 kRad
2	7 kRad
3	20 kRad

Table 5.3.3: Penetration of radiation through the #3 mask.

Temperature	dv/dt	Peak ramp voltage
110'C	4 kV/us	3 kV
125'C	3 kV/us	3 kV
140'C	1 kV/us	3 kV

Table 5.5.1: Dv/dt test results for the optical well.

Number of amplifying stages:	One	Two	Three	
Control resistor values:				
Pilot	105	100	80	R
Aux 1			20	R
Aux 2		5	5	R
Total control resistance	105	105	105	R
Peak currents:				
Pilot	44.7	44.7	44.7	A
Aux 1			32.3	A
Aux 2 (first peak)		99.0	90.0	A
Main	1960	1843	1744	A
Thyristor energies:				
Pilot	18.3	17.9	17.6	mJ
Aux 1			10.9	mJ
Aux 2		80.2	74.3	mJ
Main	2.82	2.38	2.20	J
Control resistor energies:				
Pilot	148	110	80	mJ
Aux 1			33	mJ
Aux 2		254	310	mJ
Total gate energy dissipated:	166	462	525	mJ

Table 6.3.1: Results of computer studies for thyristors with one, two and three amplifying stages, fixed total control resistance.

Control resistor values:			
Pilot	65	100	R
Aux	40	5	R
Total control resistance	105	105	R
Peak currents:			
Pilot	44.7	44.7	A
Aux	35.9	99.0	A
Main	1944	1843	A
Thyristor energies:			
Pilot	17.7	17.9	mJ
Aux	11.9	80.2	mJ
Main	2.77	2.38	J
Control resistor energies:			
Pilot	67	110	mJ
Aux	113	254	mJ
Total gate energy dissipated:	210	462	mJ

Table 6.3.2: Results of computer studies for double amplifying gate thyristors with different splits in control resistance.

Control resistor values:			
Pilot	2	100	R
Aux	1	5	R
Total control resistance	3	105	R
Peak currents:			
Pilot	260	44.7	A
Aux	149	99.0	A
Main	1590	1843	A
Thyristor energies:			
Pilot	221	17.9	mJ
Aux	199	80.2	mJ
Main	1.61	2.38	J
Control resistor energies:			
Pilot	123	110	mJ
Aux	508	254	mJ
Total gate energy dissipated:	1051	462	mJ

Table 6.3.3: Results of computer studies for double amplifying gate thyristors with different total control resistances.

Device type	Turn-on voltage causing damage			
	Sample 1	Sample 2	Sample 3	
P+ zone, 1mm	1850			V
P+ zone, 1.5mm	2300 *	2150 *	1650	V
P+ zone, 2mm	1800	1700	2000	V
P+ grid	1600	1600	<1200	V
Gated turn-on	1700	1650		V
Lateral field	1500	<1250		V
Control	1900	1450	1600	V

(* Time of ramp application initially set long and then reduced, rather than set short and then increased. Damage might therefore have occurred for turn-on from voltages lower than this.)

Table 6.4.1: Ungated turn-on voltages causing damage in forward recovery test structures (excluding Selective Failure Zone devices).

Radiation dose	Control resistor shorted?	Peak gate voltage
80 kRad	No	530 V
80 kRad	Yes	>120 V
40 kRad	Yes	48 V
20 kRad	Yes	100 V

Table 6.4.2: Peak gate-cathode voltages generated for different configurations of Selective Failure Zone device.

Device	Control Resistors			Maximum test conditions	Test Results
	Rc1	Rc2	Total		
#0	76	18	94 R	1.5kV 20'C	Accidentally damaged
#43	47	10	57 R	3.0kV 30'C	Resistor surface tracking
#11	485	40	525 R	3.0kV 30'C	di/dt failure on main cathode
#46	110	25	135 R	5.0kV 30'C	Survived test
				4.7kV 105'C	Survived test
#45	58	16	74 R	5.0kV 30'C	Voltage failure during test
#16	162	38	200 R	5.0kV 30'C	Survived test
				4.5kV 90'C	di/dt failure, main cathode

Table 6.5.1: Summary of results of repetitive turn-on tests on 56mm devices.

Device	Total control resistance	Threshold sensitivity (60V,22'C)
#0	79 R	28 nJ
#46	135 R	18 nJ
#17	127 R	26 nJ
#49	541 R	22 nJ
#21	284 R	13 nJ
#13	414 R	16 nJ
#25	80 R	30 nJ
#45	74 R	26 nJ
#36	546 R	18 nJ
#16	197 R	8 nJ

Table 6.5.2: Control resistance and optical sensitivities for the 56mm devices.

Device	Temperature	Dv/dt withstand	Notes
#0	110 'C	> 6.0 kV/us	
	125 'C	3.0 kV/us	
	140 'C	1.0 kV/us	
#46	110 'C	2.5 kV/us	Fully shorted structure
	110 'C	3.0 kV/us	Optical well isolated
	110 'C	3.3 kV/us	Optical and auxiliary isolated
#49	66 'C	> 5.0 kV/us	
	85 'C	4.2 kV/us	
	95 'C	3.1 kV/us	
	110 'C	2.2 kV/us	
	115 'C	1.9 kV/us	
	130 'C	0.5 kV/us	
	150 'C	0.28 kV/us	
	90 'C	> 5.0 kV/us	Optical well isolated
	115 'C	1.9 kV/us	Control resiistor Rc1 shorted
#17	105 'C	> 5.0 kV/us	
	120 'C	3.0 kV/us	
	135 'C	1.2 kV/us	
	150 'C	0.24 kV/us	
#36	62 'C	3.8 kV/us	
	92 'C	2.8 kV/us	
	110 'C	2.0 kV/us	
	125 'C	1.4 kV/us	
	150 'C	0.18 kV/us	
#13	92 'C	> 5.0 kV/us	
	110 'C	3.3 kV/us	
	125 'C	2.1 kV/us	
	150 'C	0.23 kV/us	
#16	90 'C	2.0 kV/us	
	110 'C	1.4 kV/us	

Table 6.5.3: Dv/dt test results for 56mm devices, 3kV ramp voltage.

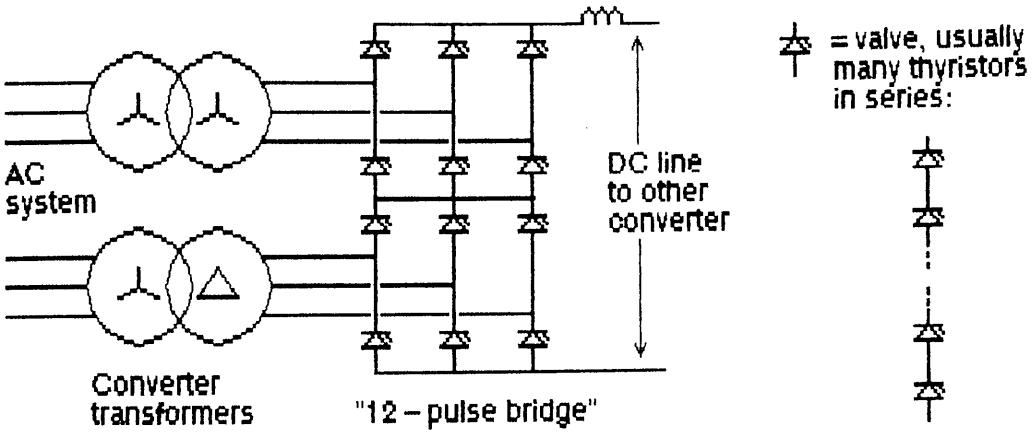

Figure 2.1.1: Normal arrangement of an HVDC converter 12-pulse bridge.

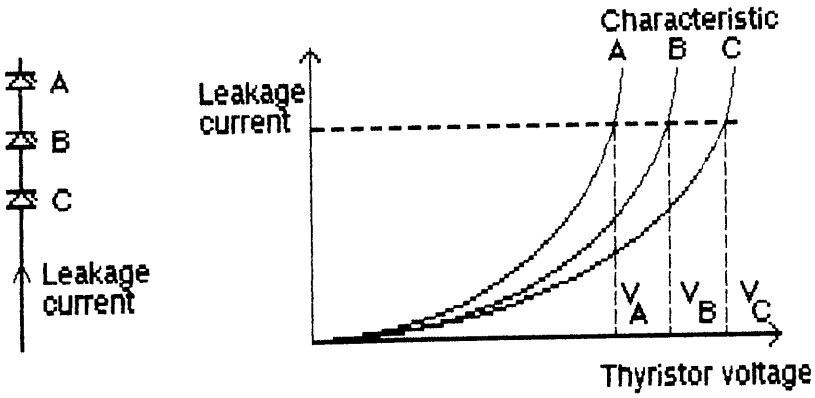

Figure 2.2.1: Voltage sharing unbalance due to unequal leakage current characteristics.

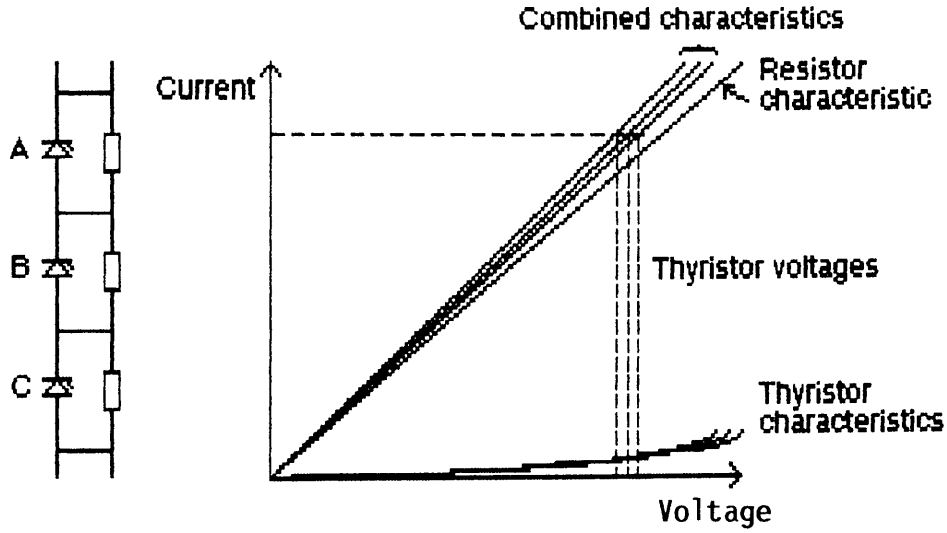

Figure 2.2.2: The use of grading resistors to enforce even voltage sharing.

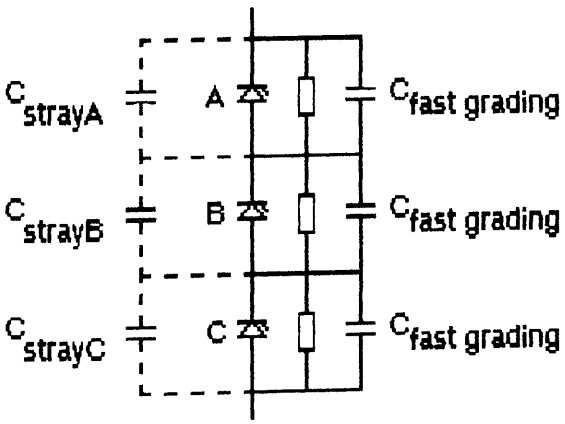

$C_{strayA} \neq C_{strayB} \neq C_{strayC}$ $\qquad$ $C_{fast\ grading} >> C_{stray\,A,B,C}$

Figure 2.2.3: The use of Fast Grading Capacitors to overcome differences in stray capacitance between levels.

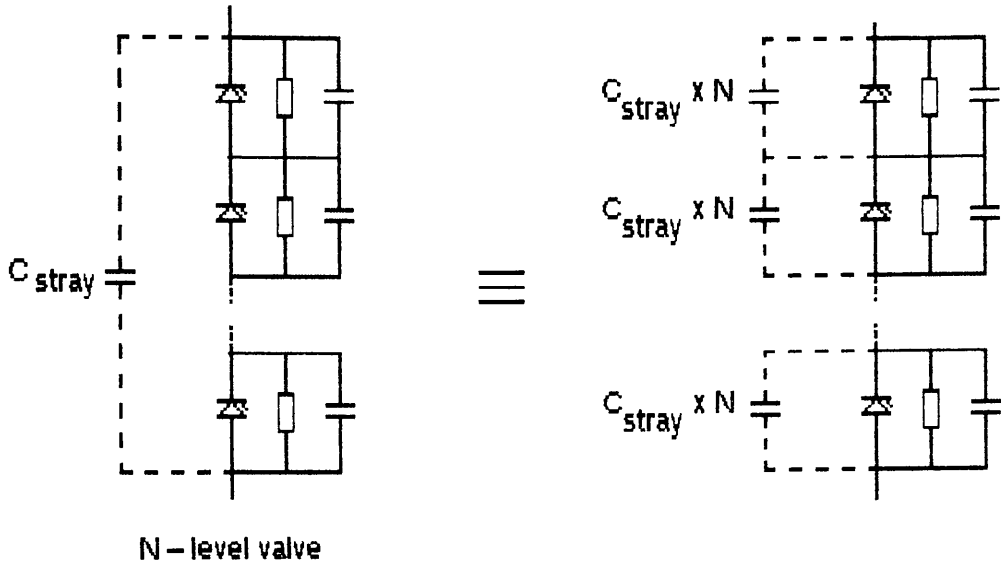

Figure 2.2.4: The effect of stray capacitance across the whole valve when referred to one level.

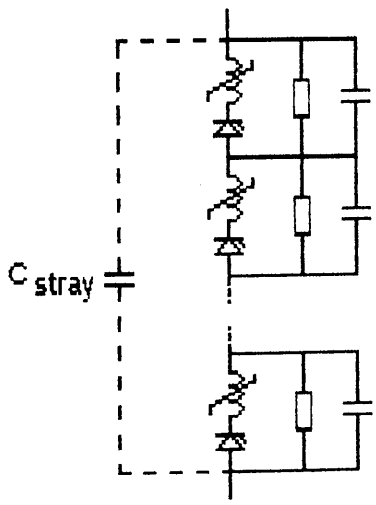

Figure 2.2.5: A typical thyristor-level circuit showing a saturable reactor used to limit inrush current in the early stages of turn-on.

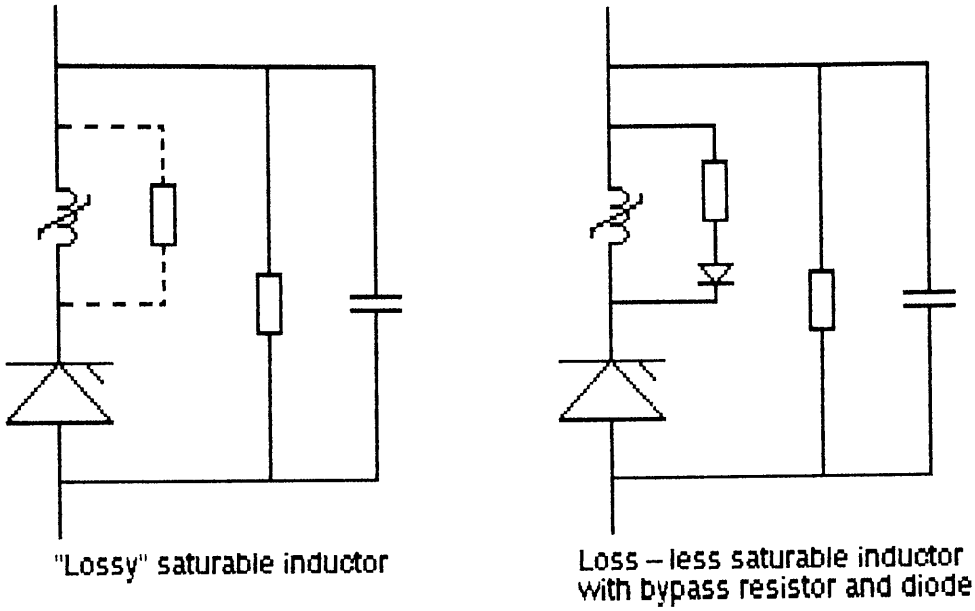

Figure 2.2.6: The use of a "lossy" saturable reactor or a bypass resistor and diode to damp the inrush current oscillation.

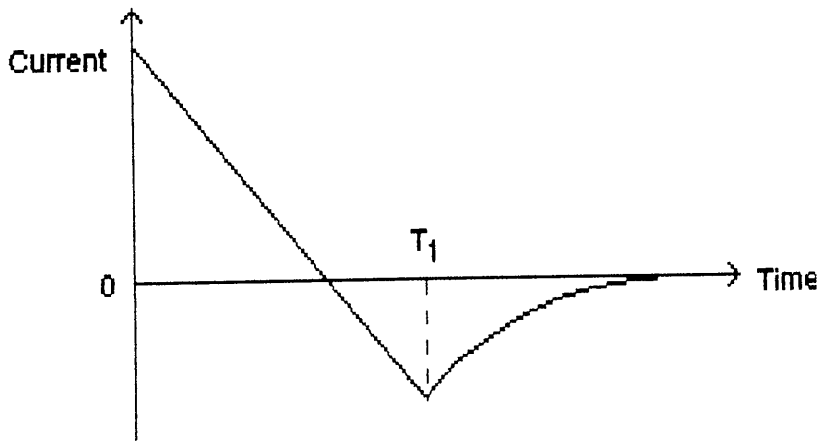

Figure 2.2.7: Thyristor current flow at recovery.

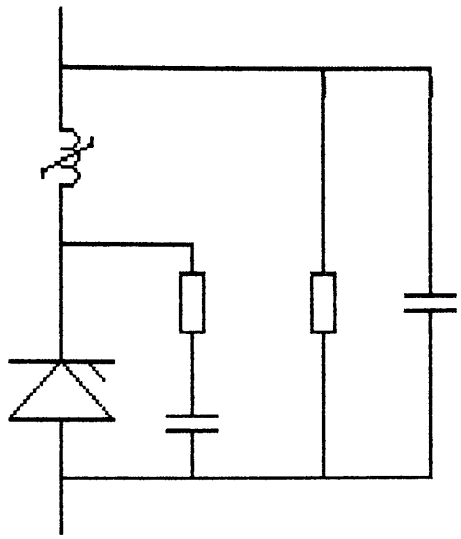

Figure 2.2.8: Thyristor-level circuit showing the RC damping network across the device.

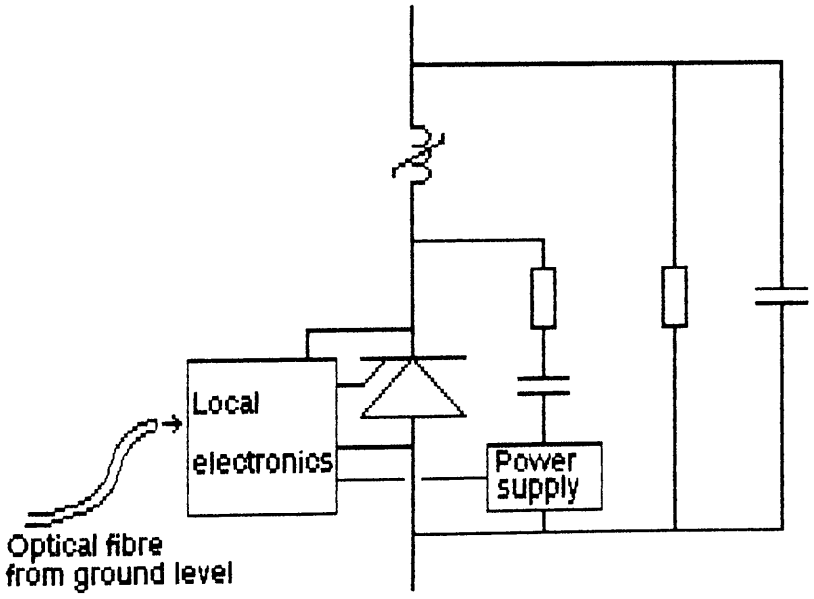

Figure 2.2.9: Arrangement of the local electronics and power supply in the thyristor level.

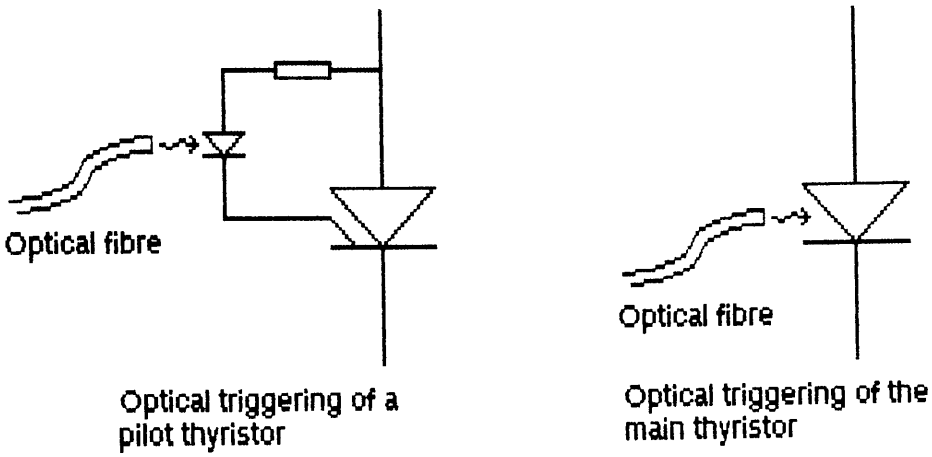

Optical triggering of a pilot thyristor

Optical triggering of the main thyristor

Figure 3.2.1: The two techniques for optical triggering of an HVDC valve.

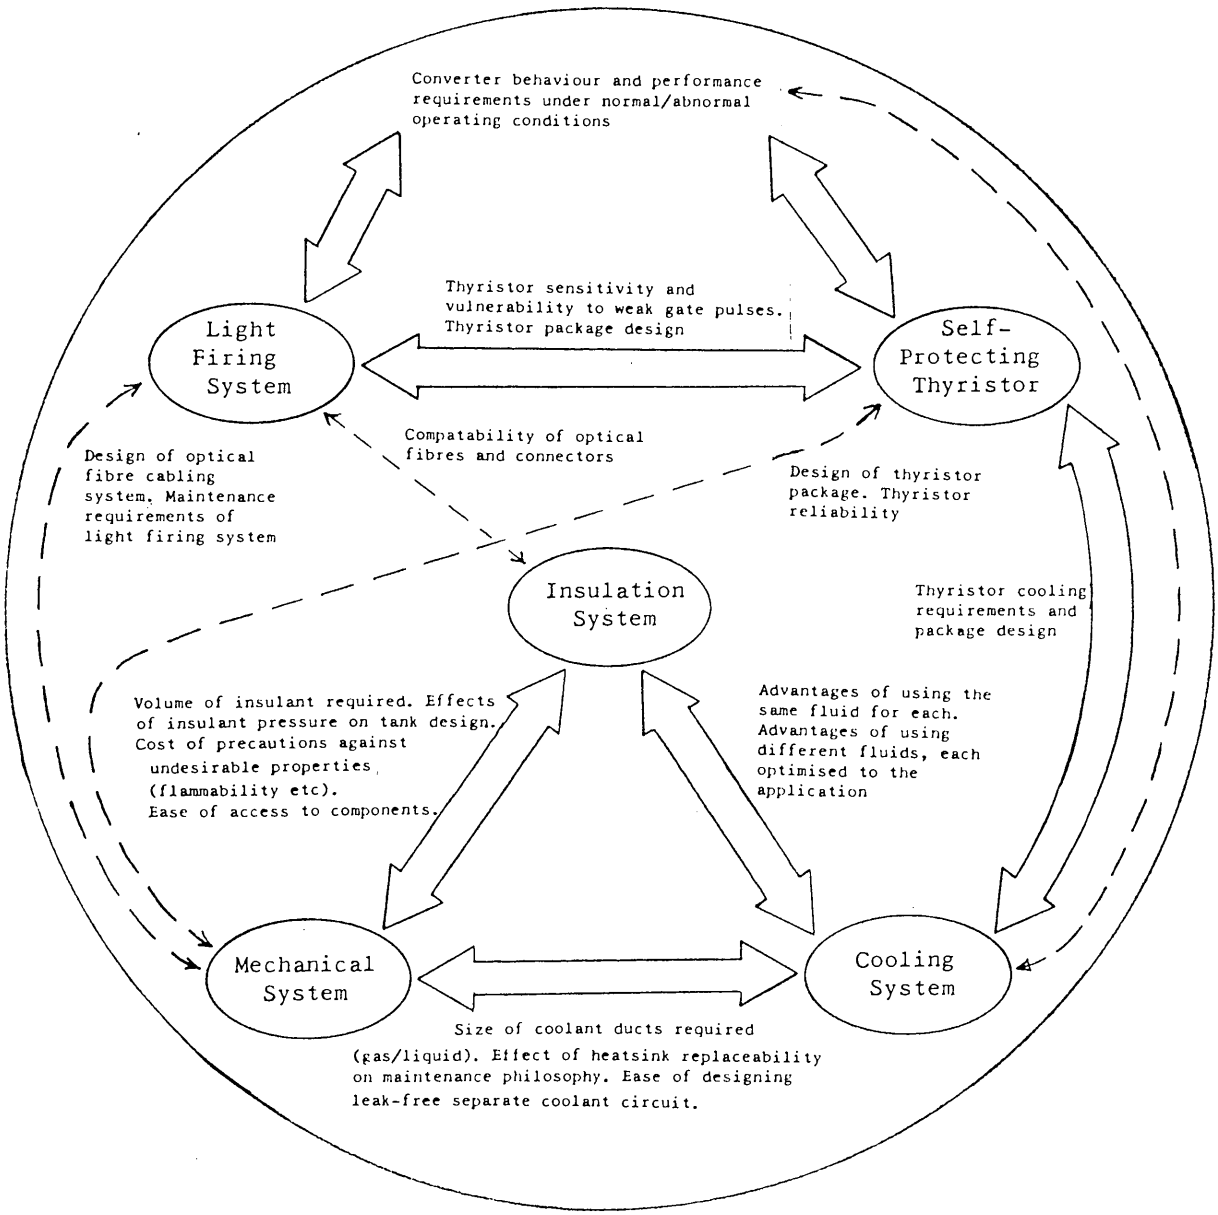

Global Constraint: System Cost.
- Capital cost
- Cost of losses
- Cost of spares
- Cost of lost revenue due to - planned maintenance
  - forced maintenance

Figure 3.3.1: Interactions between the technology areas in the project.

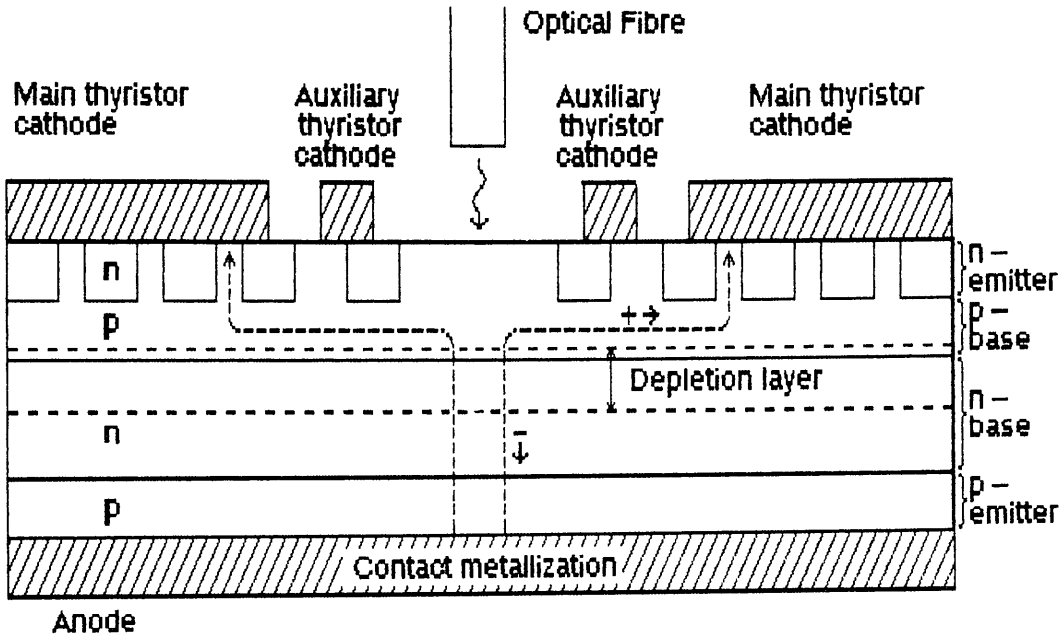

Figure 4.2.1: Optically-induced current flow.

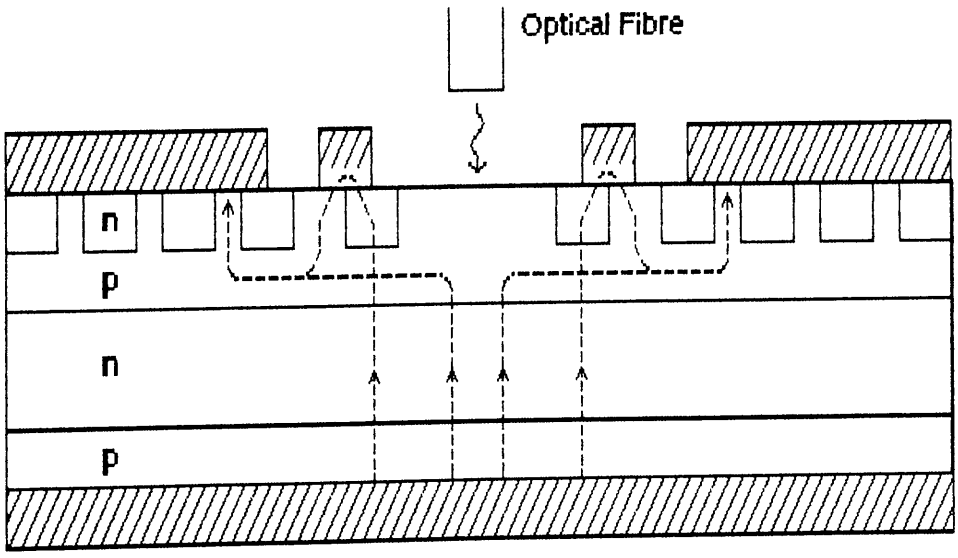

Figure 4.2.2: Optically-induced auxiliary turn-on.

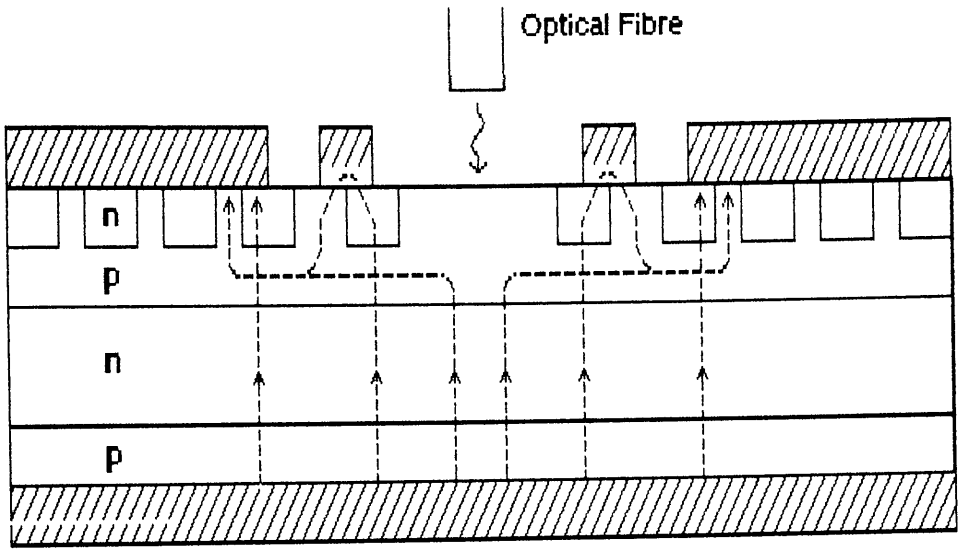

Figure 4.2.3: Optically-induced turn-on of the main cathode.

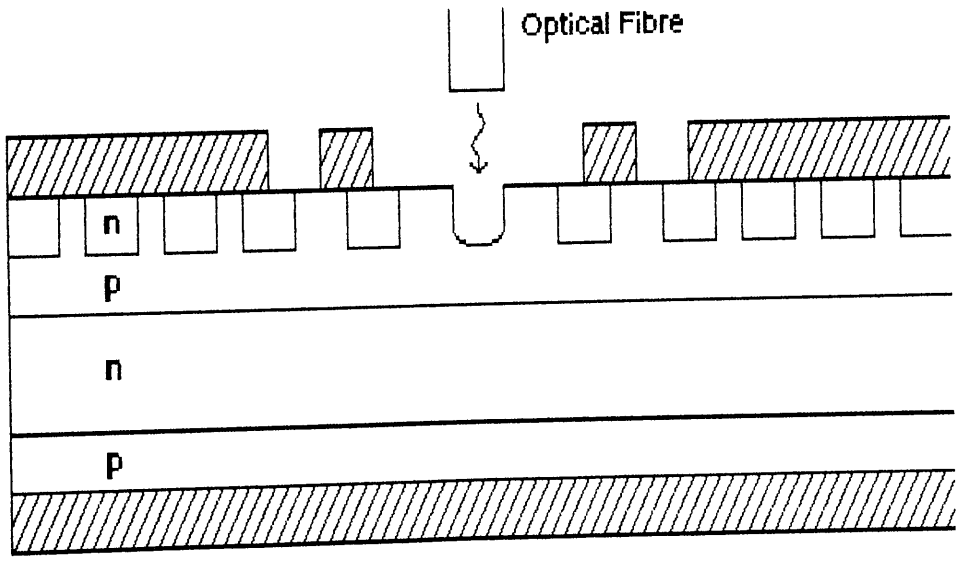

Figure 4.2.4: Etched well for improved optical sensitivity.

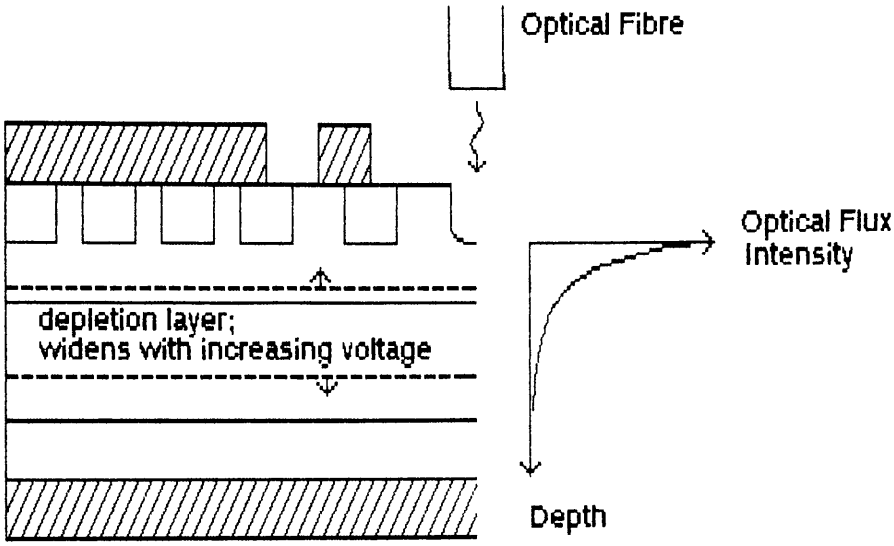

Figure 4.2.5: Optical flux density vs depth.

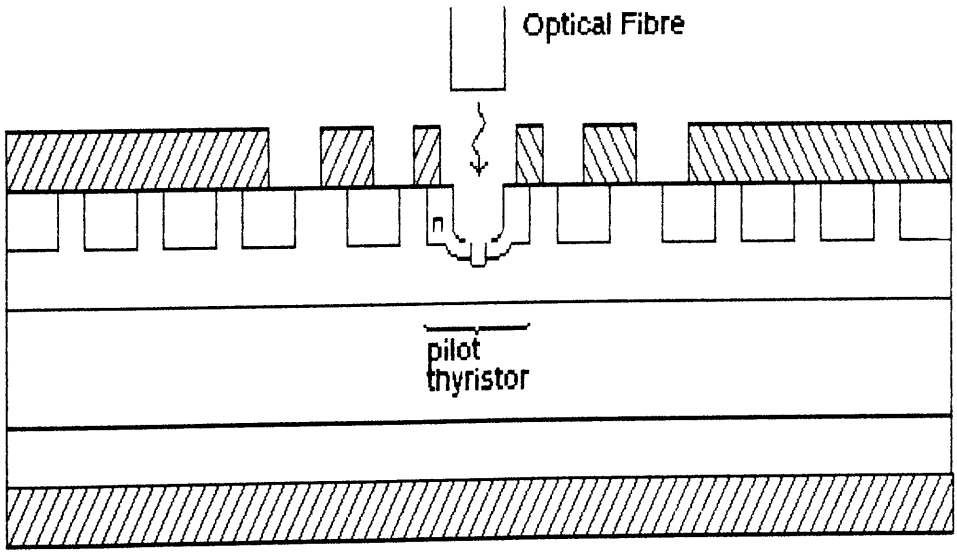

Figure 4.2.6: Structure with extra 'pilot' thyristor.

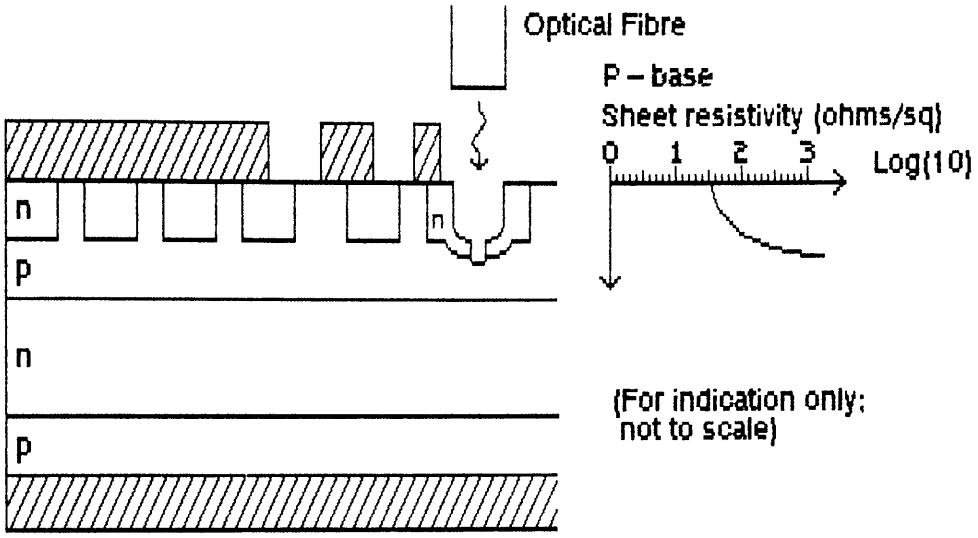

Figure 4.2.7: Variation of p-base sheet resistivity with depth.

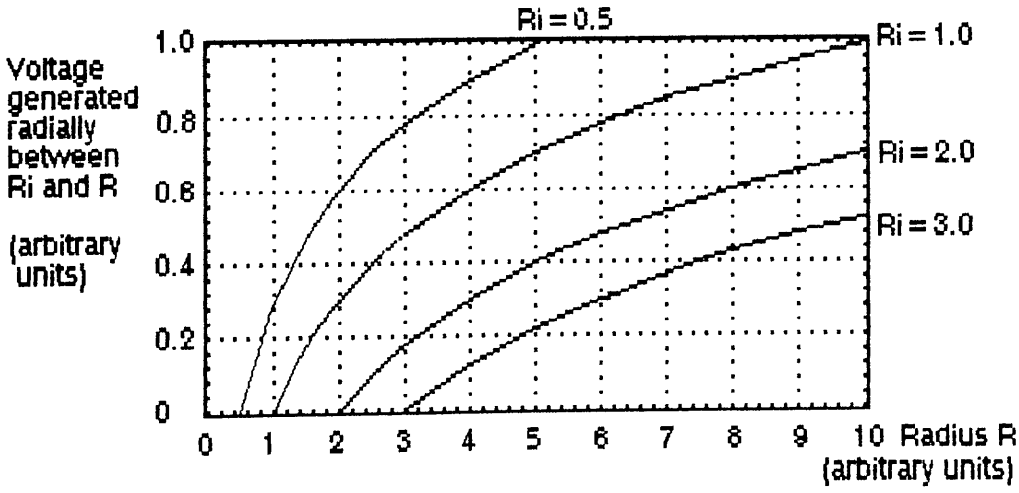

Figure 4.2.8: Voltage generated along the p-base by a fixed optically-induced current flow, for different internal radii (Ri) of the pilot thyristor.

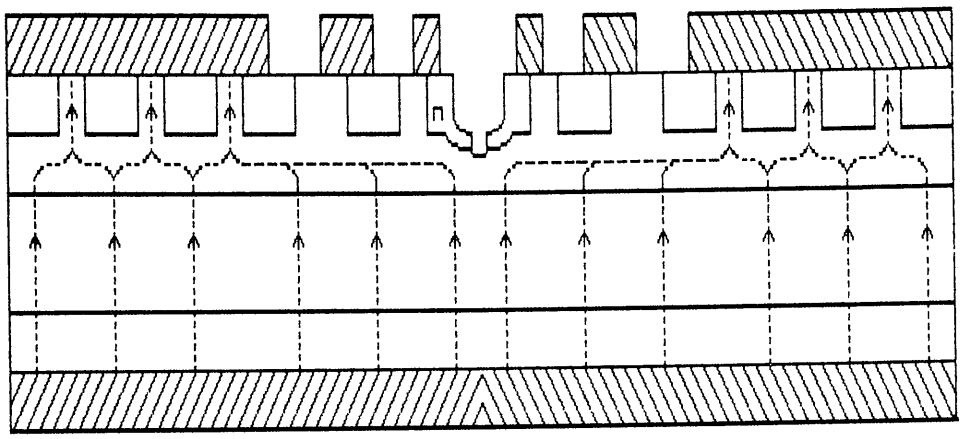

Figure 4.2.9: Dv/dt current flow in the off-state thyristor.

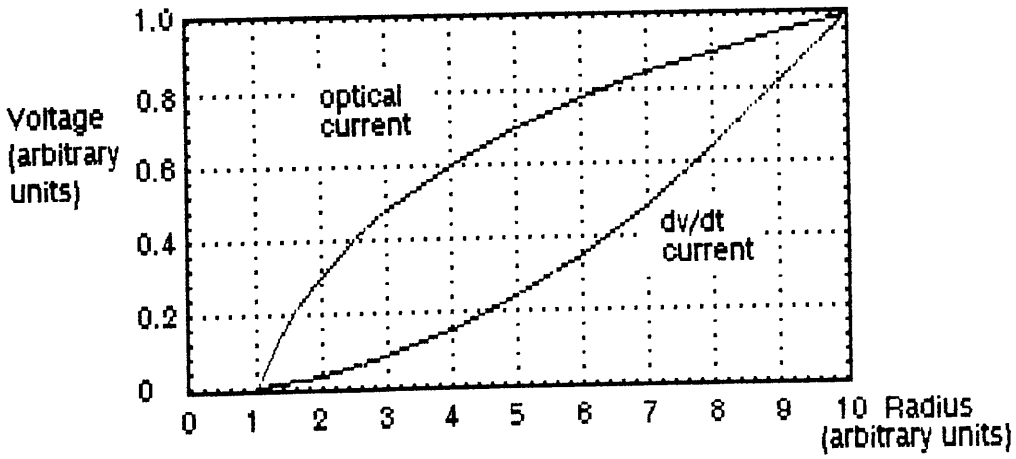

Figure 4.2.10: Voltage profiles generated by optical and dv/dt current flow.

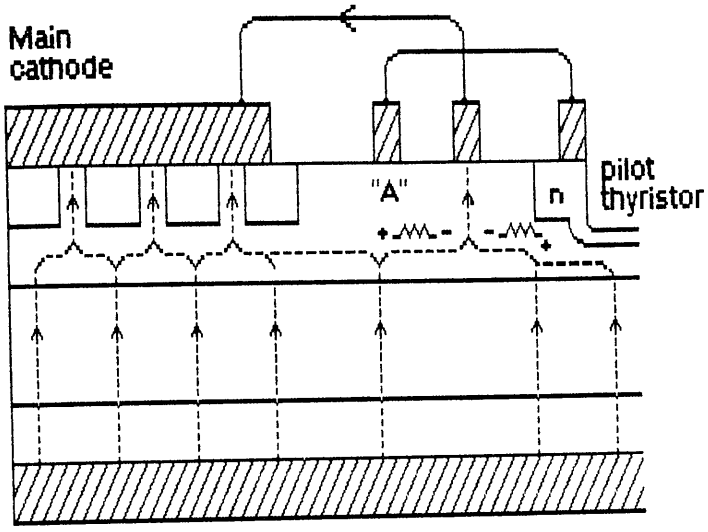

Figure 4.2.11: Connections for dv/dt compensation.

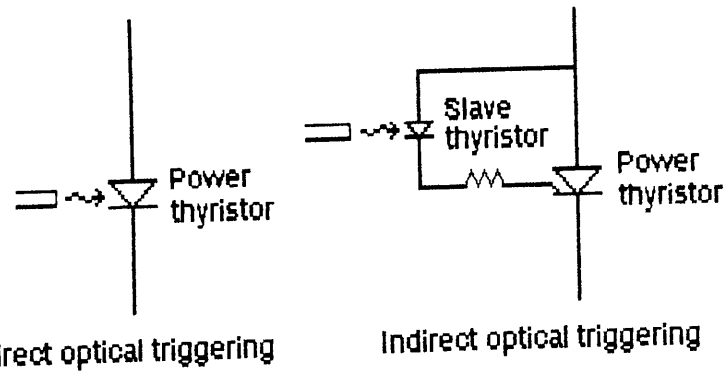

Figure 4.2.12: Comparison of direct and indirect optical triggering.

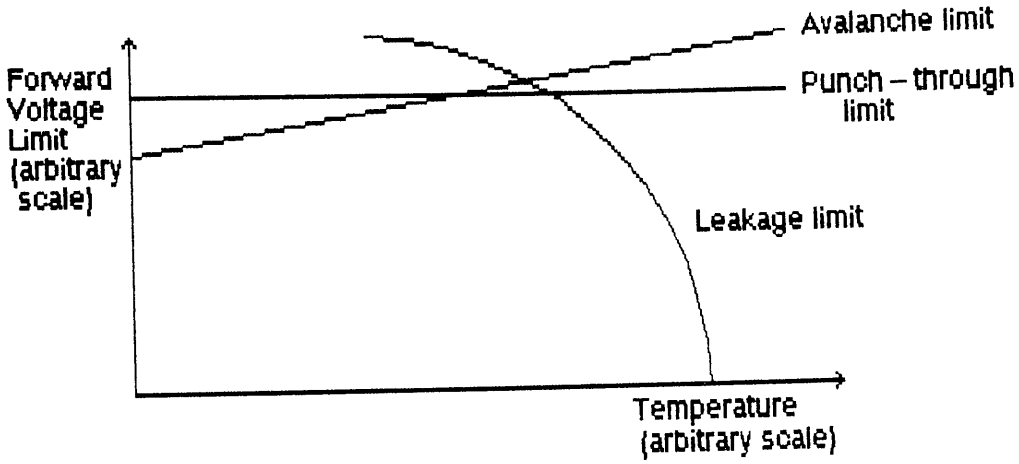

Figure 4.3.1: Limits to thyristor forward voltage capability.

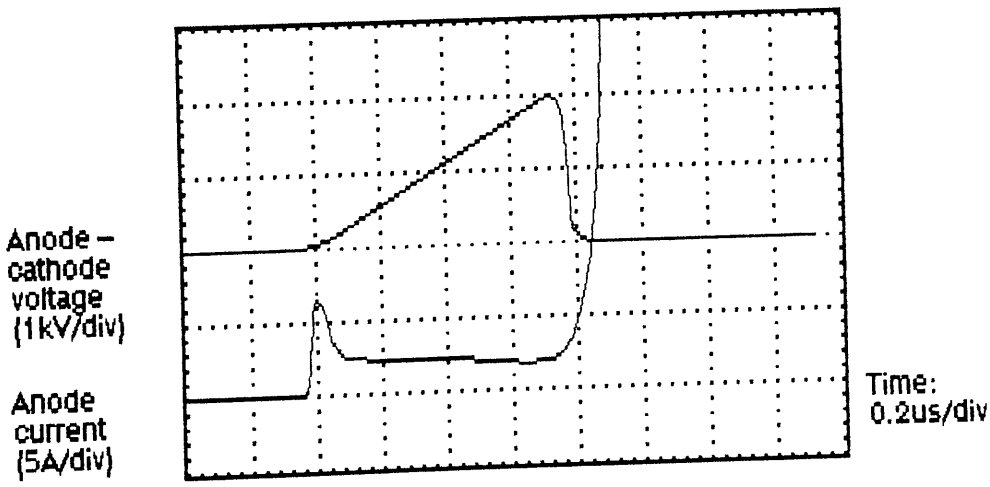

Figure 4.3.2: Typical record of marginal dv/dt failure.

Figure 4.3.3: Typical record of marginal forward recovery failure.

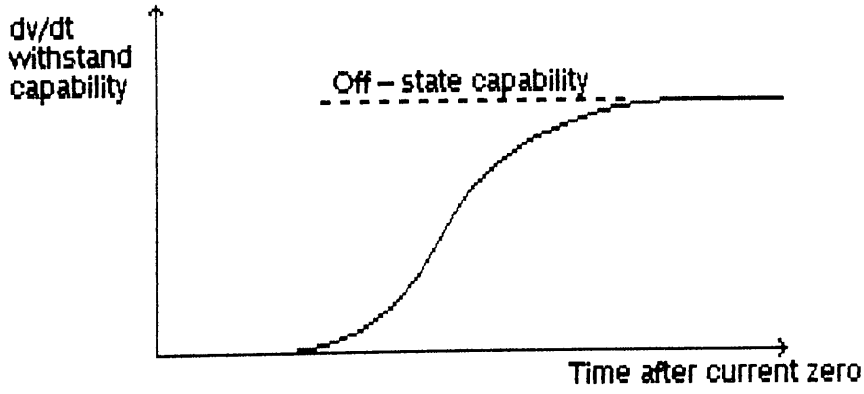

Figure 4.3.4: Thyristor dv/dt withstand capability as a function of time after current zero.

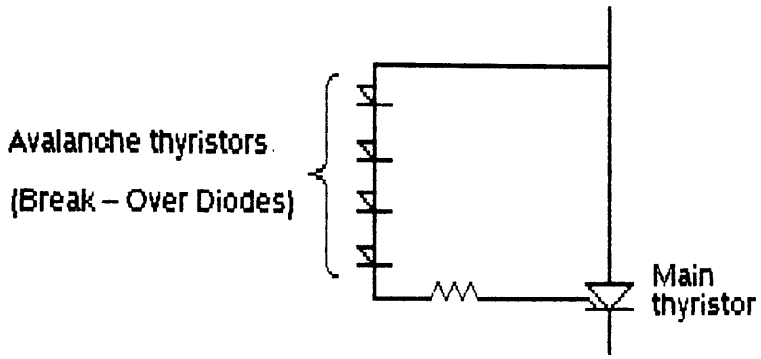

Figure 4.3.5: Break-Over Diode assembly for forward overvoltage protection.

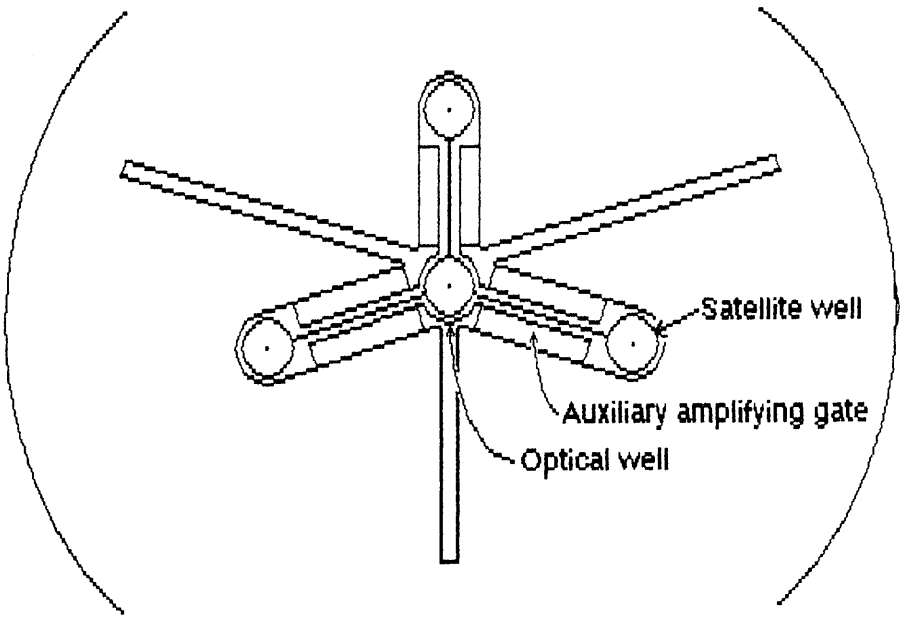

Figure 4.6.1: Outline of prototype thyristor (part of main cathode omitted for clarity).

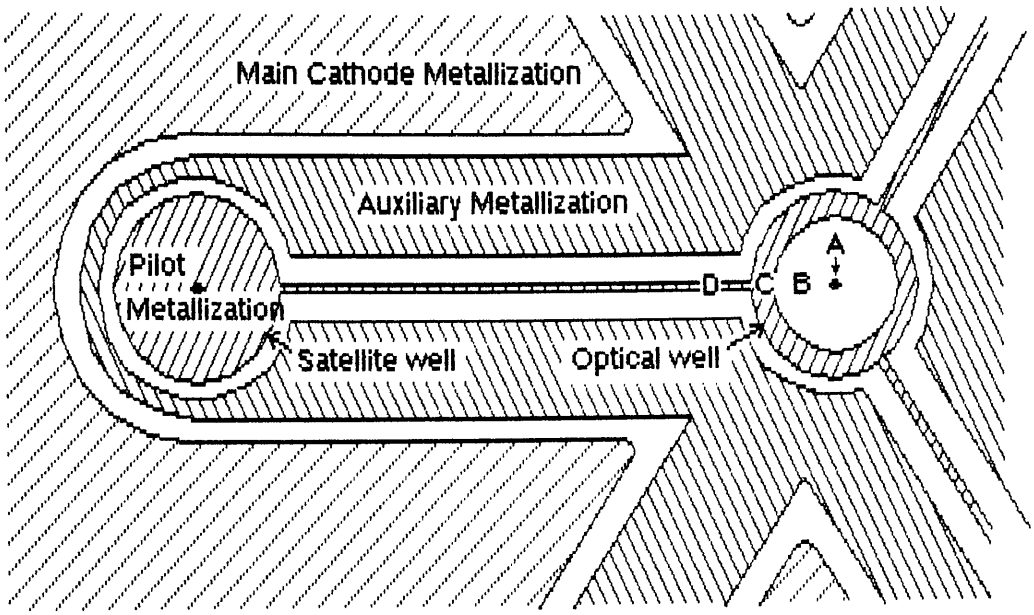

Figure 4.6.2: Detail of part of prototype gate structure.

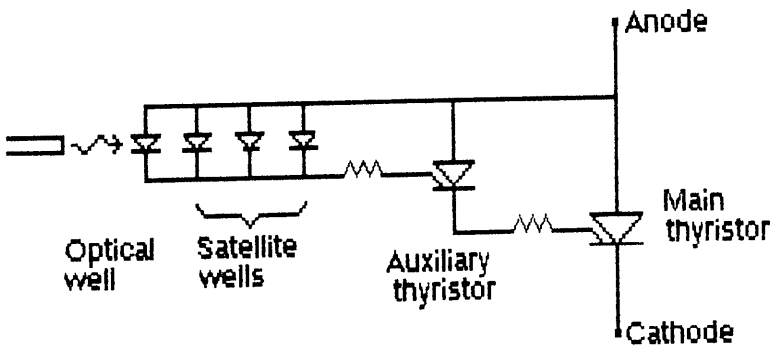

Figure 4.6.3: Equivalent circuit of the prototype thyristor.

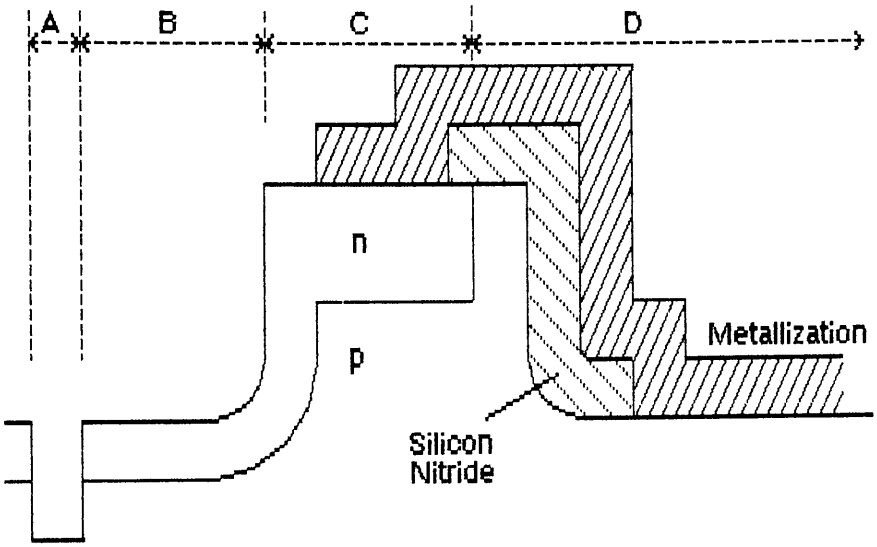

Figure 4.6.4: Cross-section through part of the optical well.

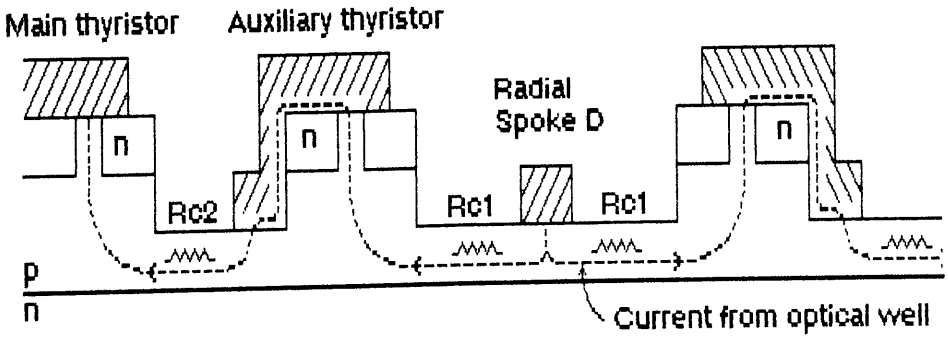

Figure 4.6.5: Control resistors Rc1 and Rc2.

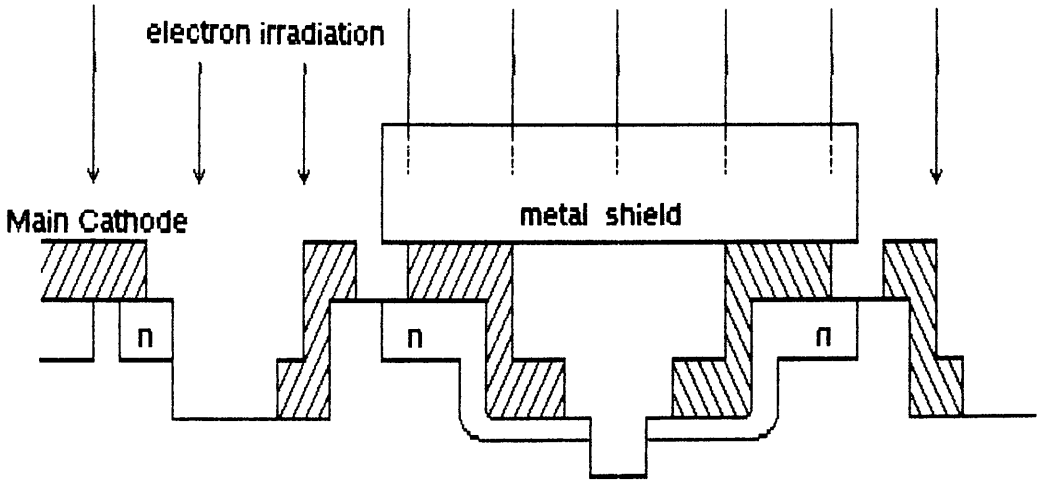

Figure 4.6.6: Shield over satellite well during irradiation.

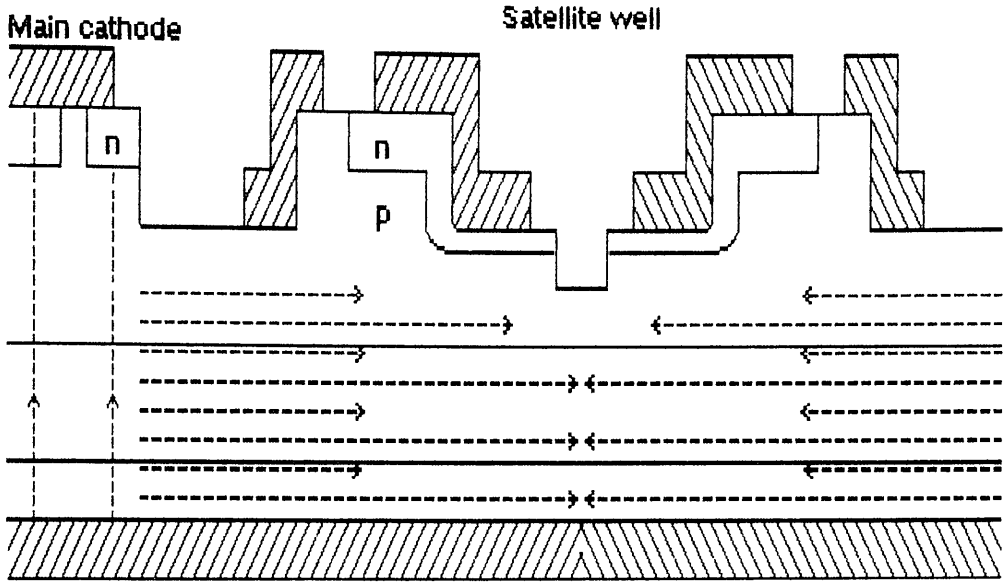

Figure 4.6.7: Charge diffusion under a satellite well from the main cathode.

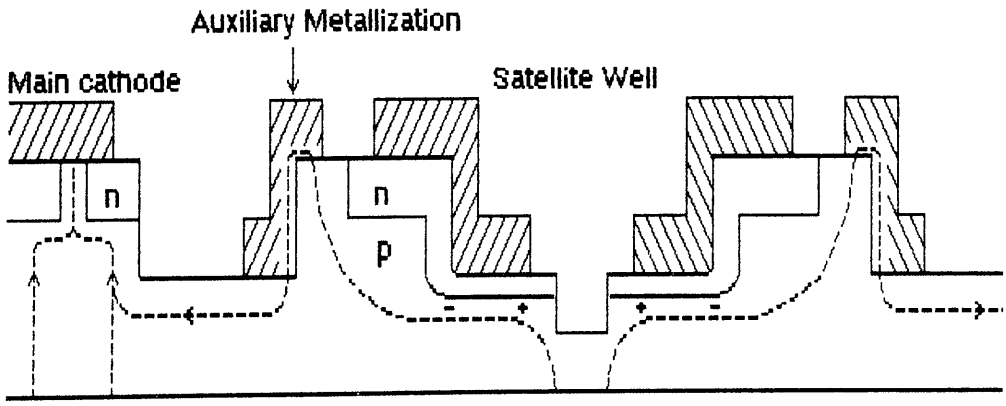

Figure 4.6.8: Dv/dt current flow under a satellite well.

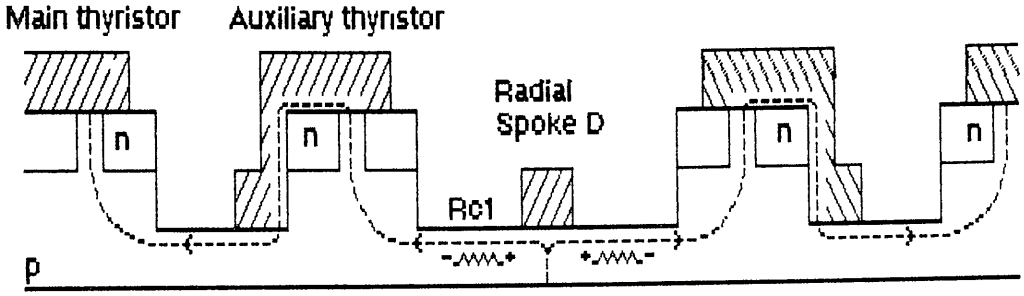

Figure 4.6.9: Compensation voltage generated by the dv/dt current flow under Rc1.

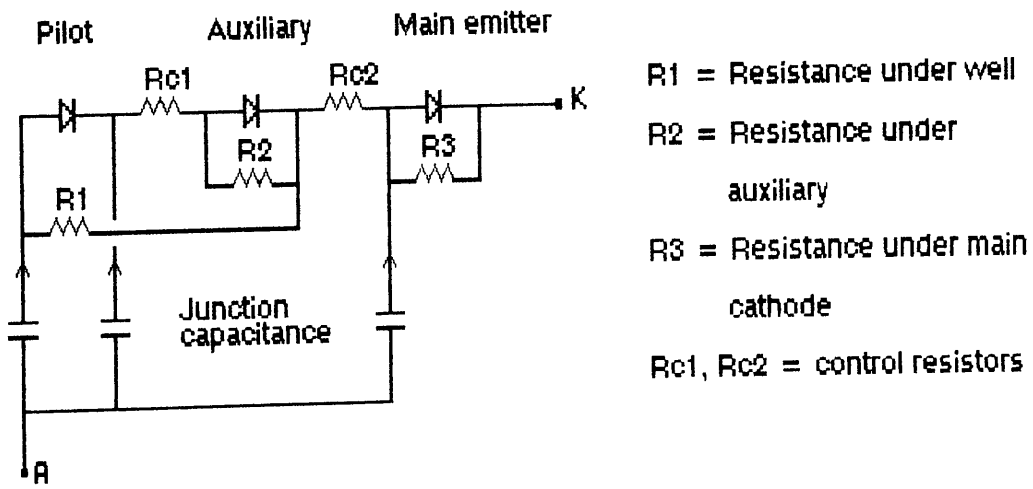

Figure 4.6.10: Equivalent circuit of the prototype thyristor showing dv/dt compensation.

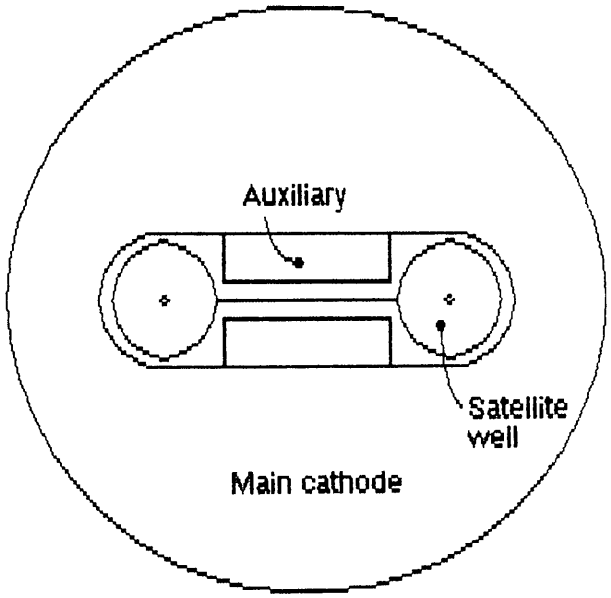

Figure 4.6.11: 30mm version of the prototype thyristor.

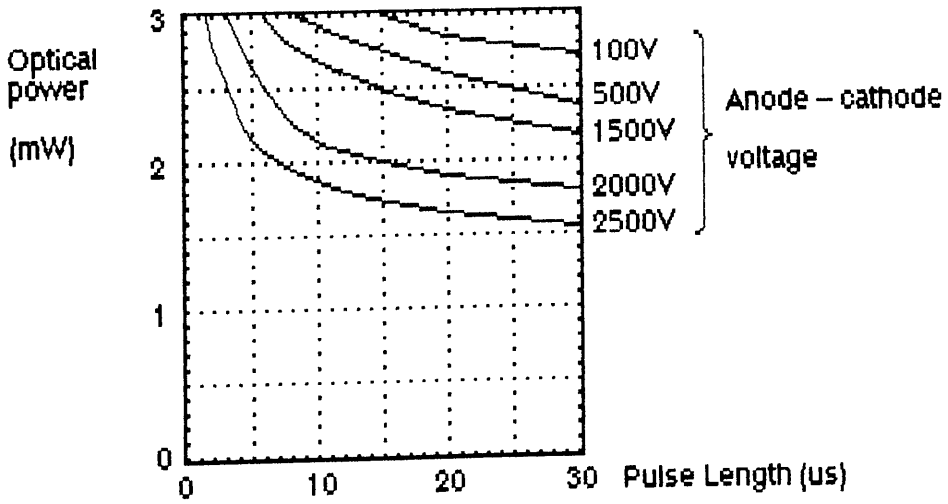

Figure 4.7.1: Optical sensitivity curves, LED triggering.

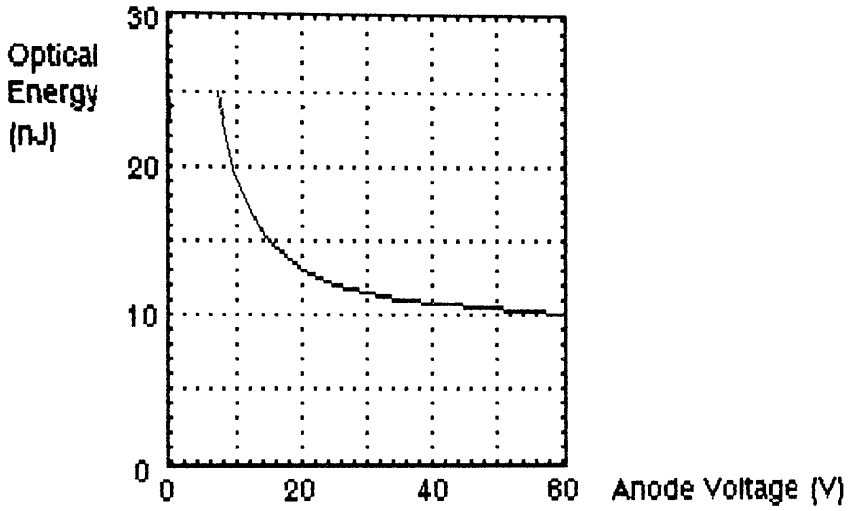

Figure 4.7.2: Optical sensitivity curve, laser triggering.

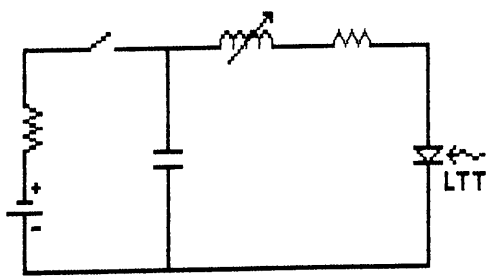

Figure 4.7.3: Di/dt test circuit.

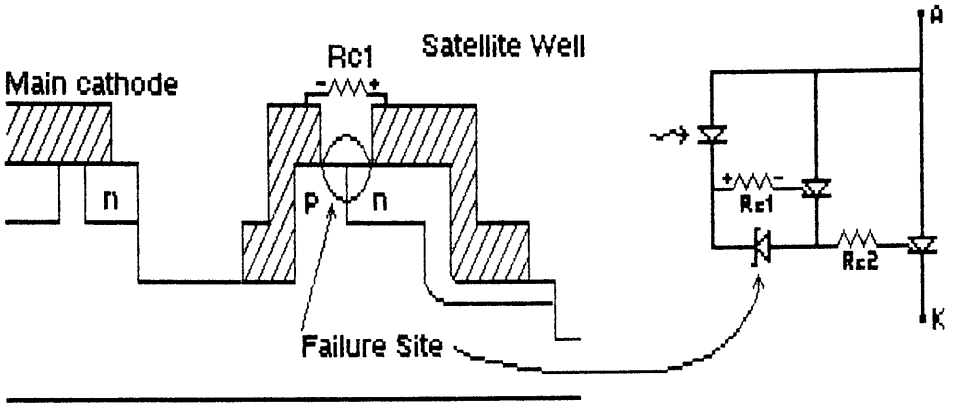

Figure 4.7.4: Turn-on failure site at parasitic diode.

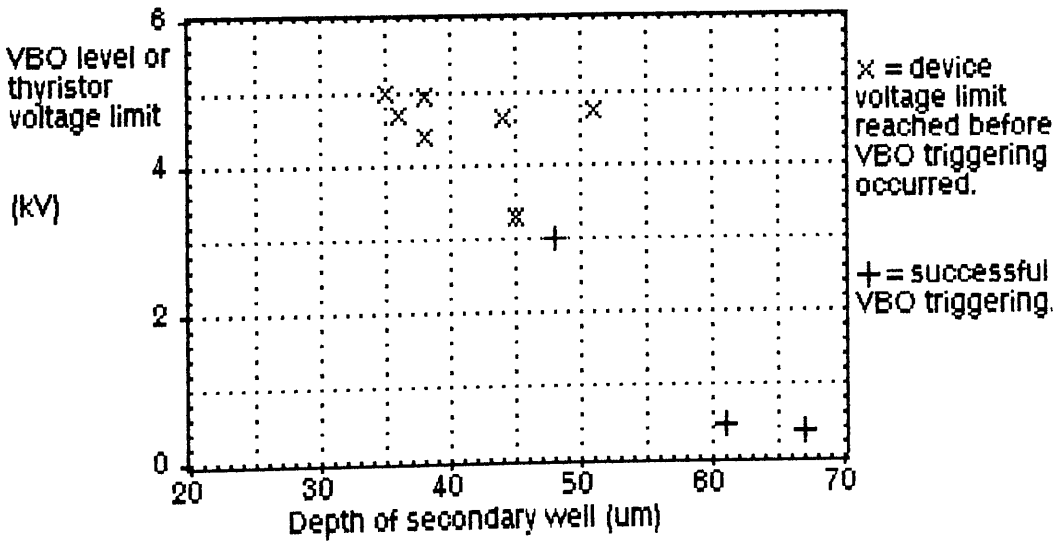

Figure 4.7.5: Variation of VBO level (or device voltage limit) with depth of secondary well.

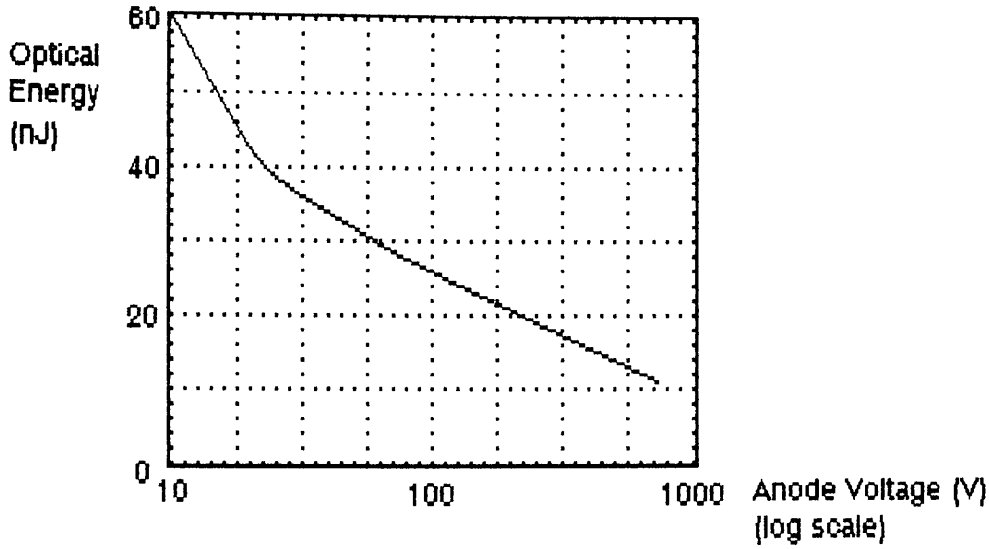

Figure 5.2.1: Optical sensitivity versus anode voltage, laser triggering.

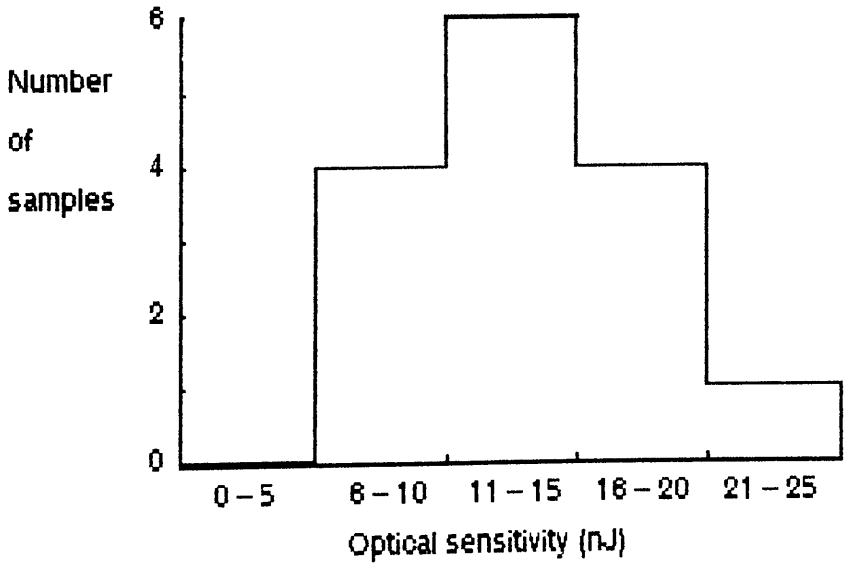

Figure 5.2.2: Distribution of optical well sensitivities for the 100mm devices.

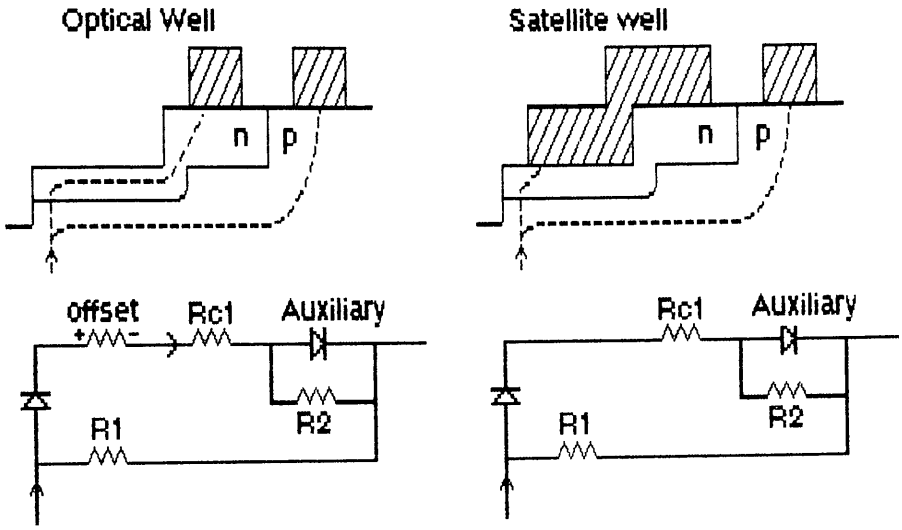

Figure 5.2.3: Comparison of turn-on with and without metal in the well.

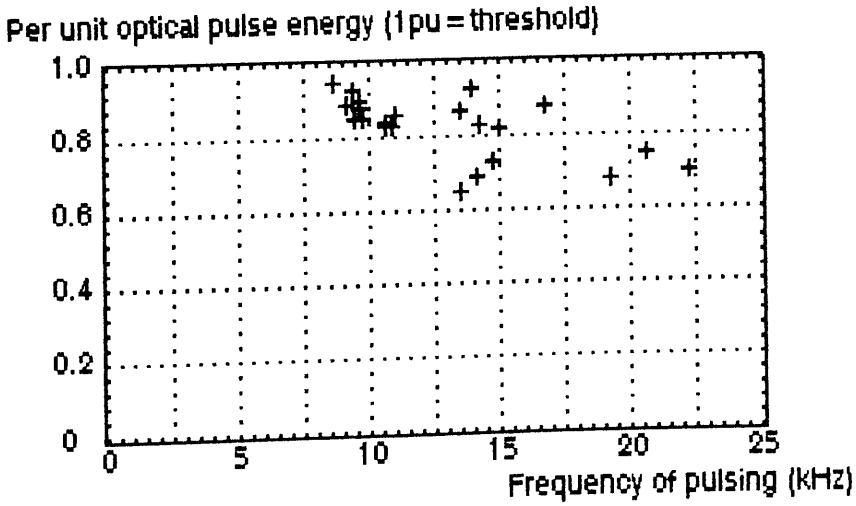

Figure 5.2.4: Per unit optical pulse energy versus frequency required for triggering.

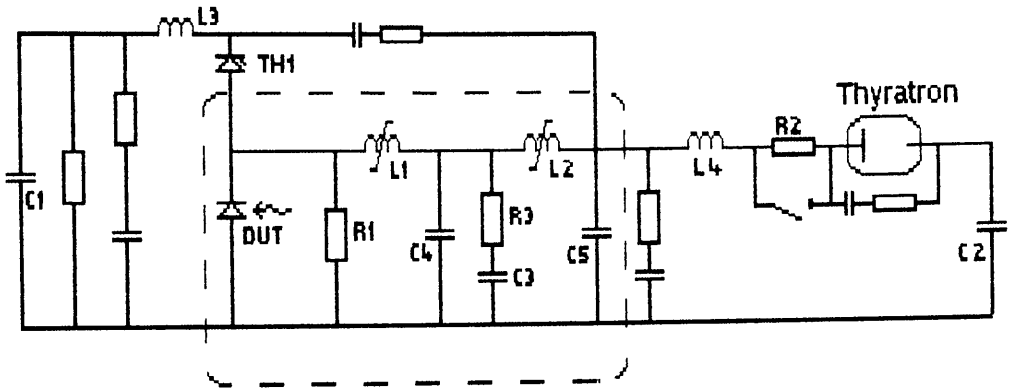

Figure 5.3.1: Forward recovery test circuit.

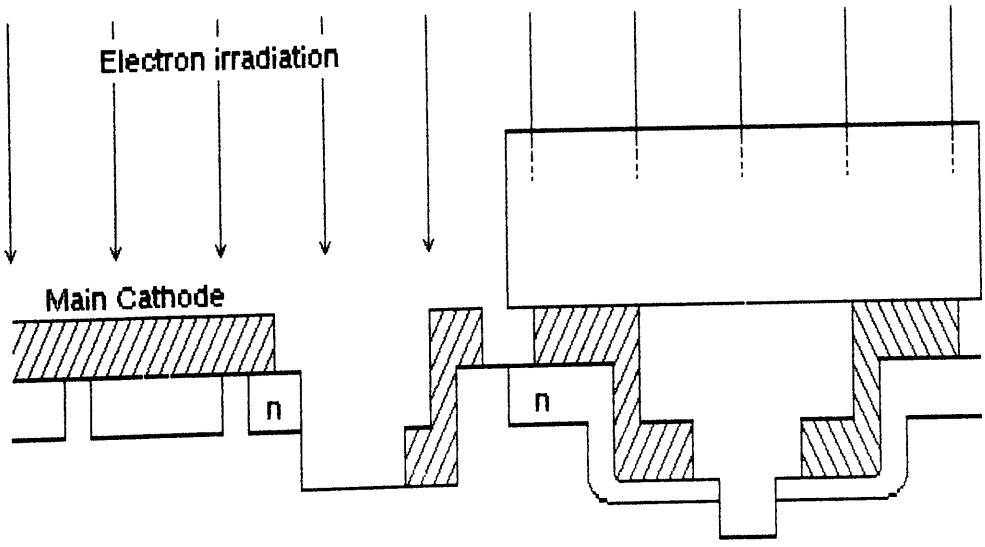

Figure 5.3.2: Design of Mask #1.

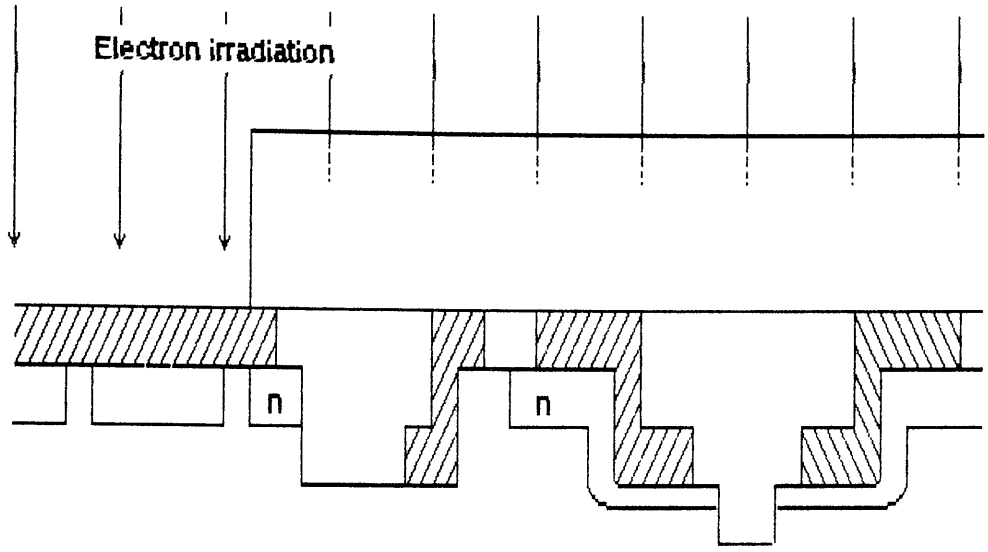

Figure 5.3.3: Design of Mask #2.

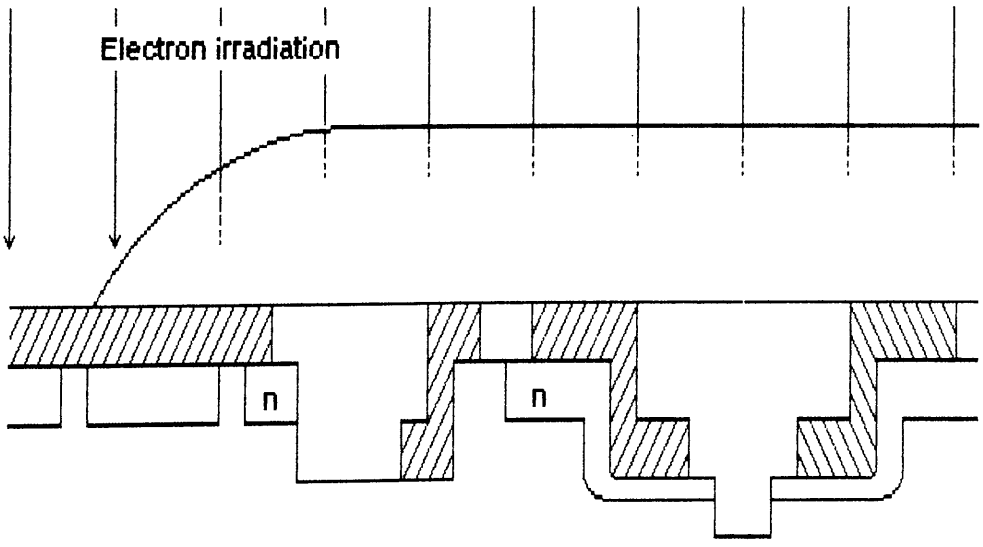

Figure 5.3.4: Design of Mask #3.

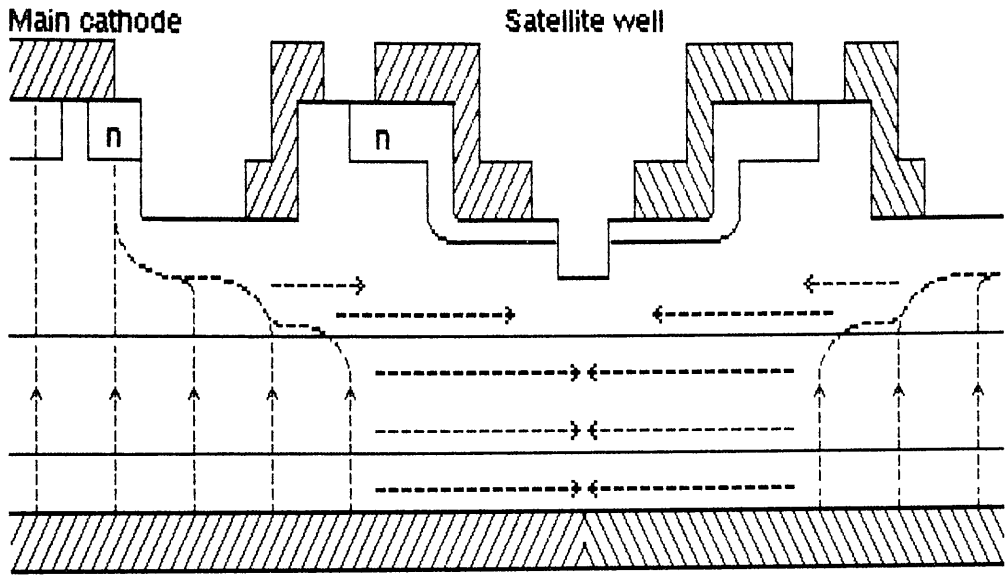

Figure 5.3.5: Lateral current flow under a satellite well caused by the Vf gradient.

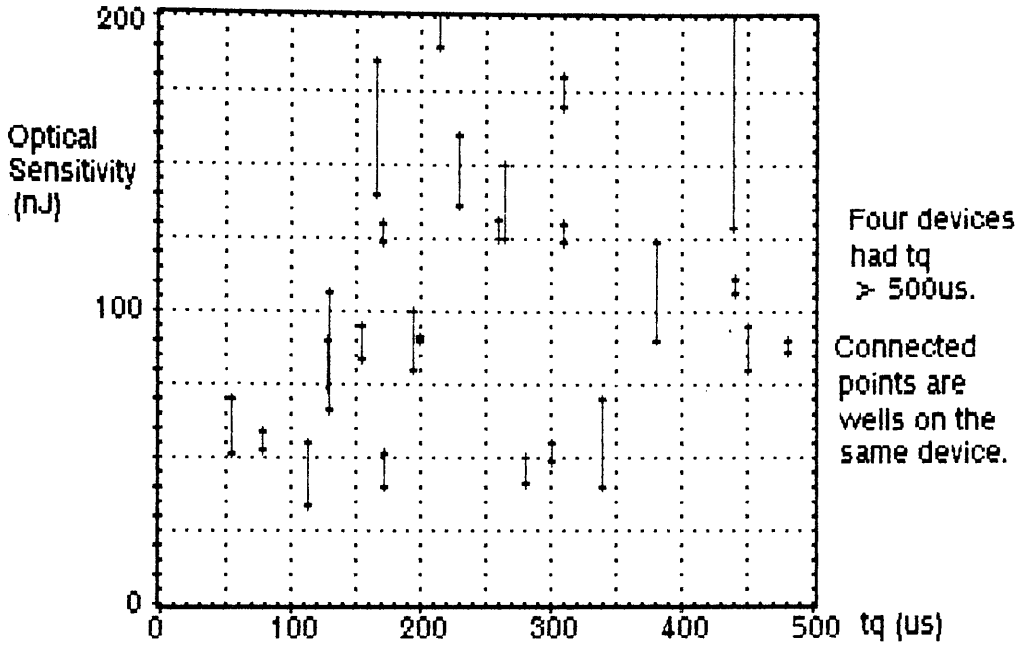

Figure 5.3.6: Optical sensitivity versus turn-off time (tq) for the 30mm samples.

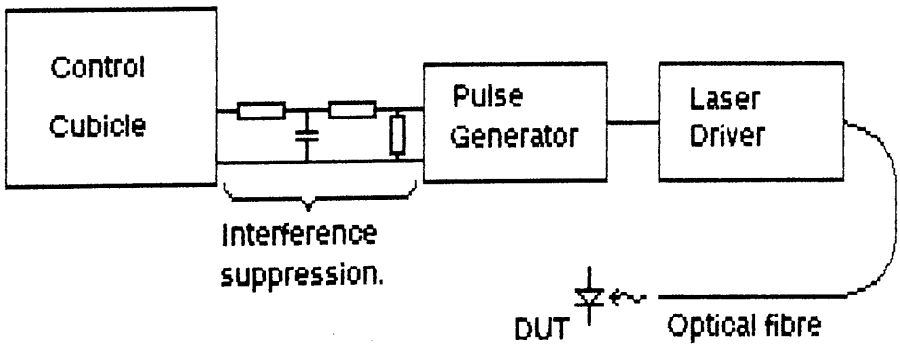

Figure 5.3.7: Circuit arrangement for optical triggering.

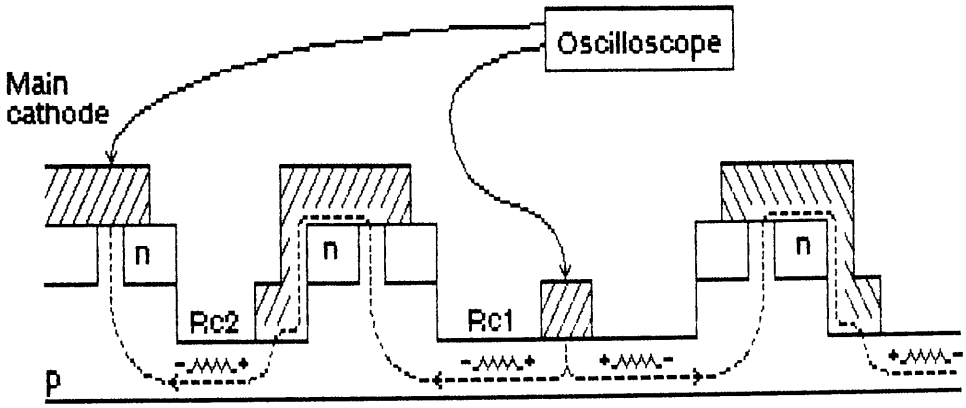

Figure 5.3.8: Arrangement for monitoring the voltage generated across the control resistors during turn-on.

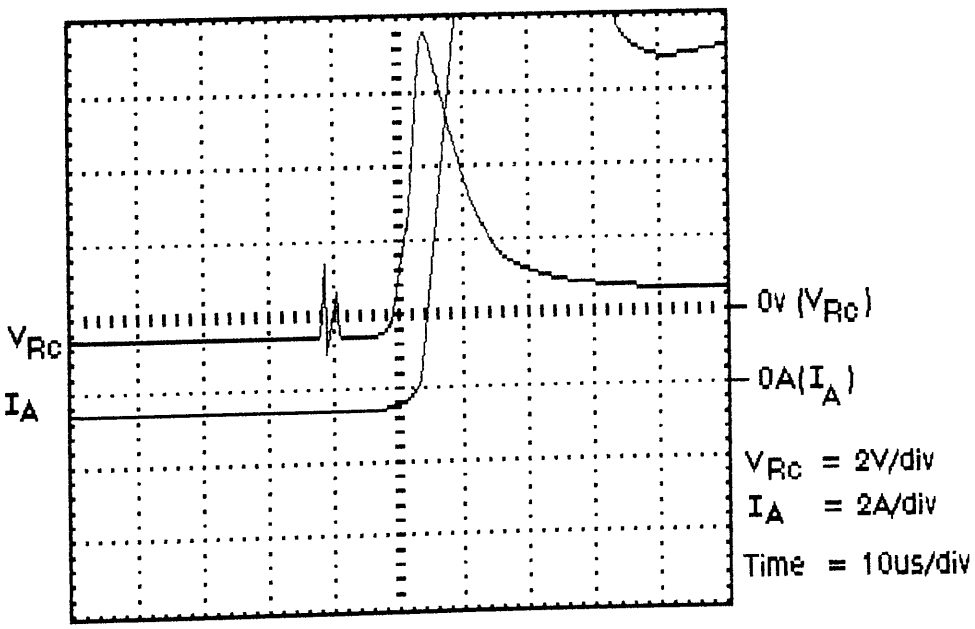

Figure 5.3.9: Voltage waveforms during forward recovery failure.

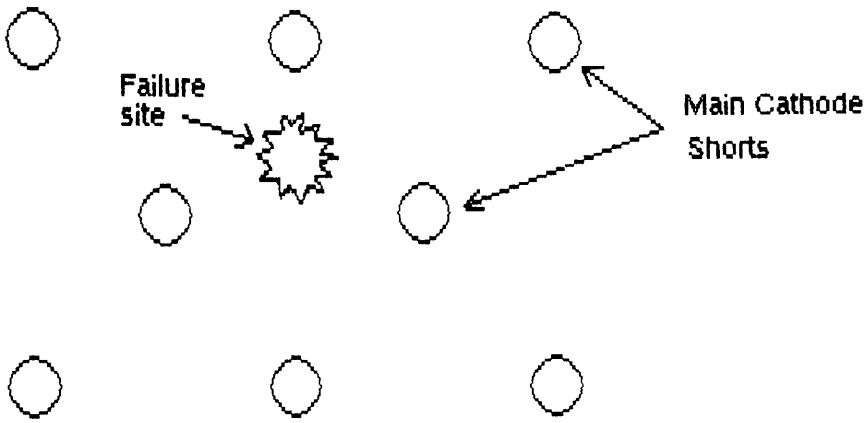

Figure 5.3.10: Location of failure site in relation to main cathode shorting pattern.

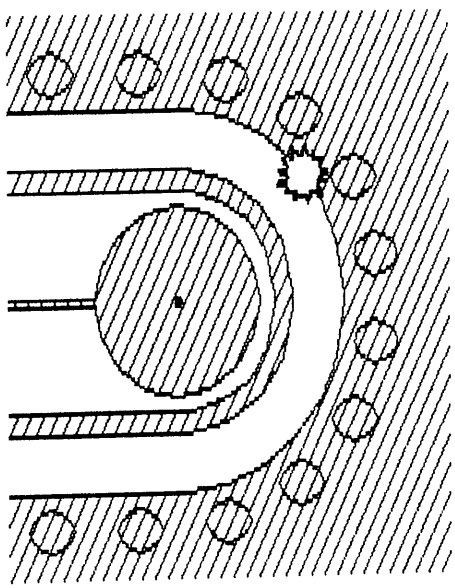

Figure 5.3.11: Typical failure location for devices with #2 and #3 masks.

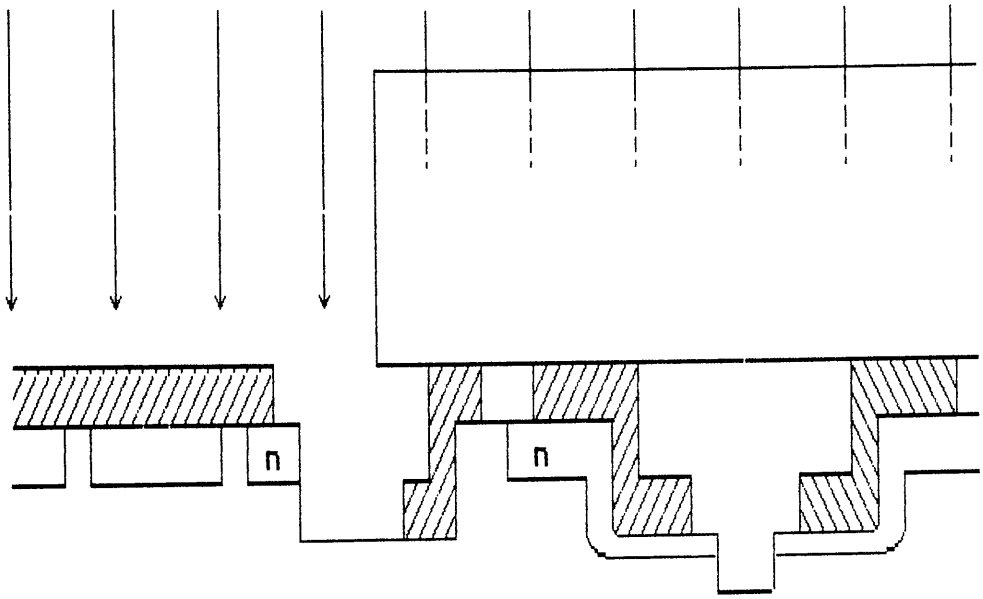

Figure 5.3.12: Design of Mask #4.

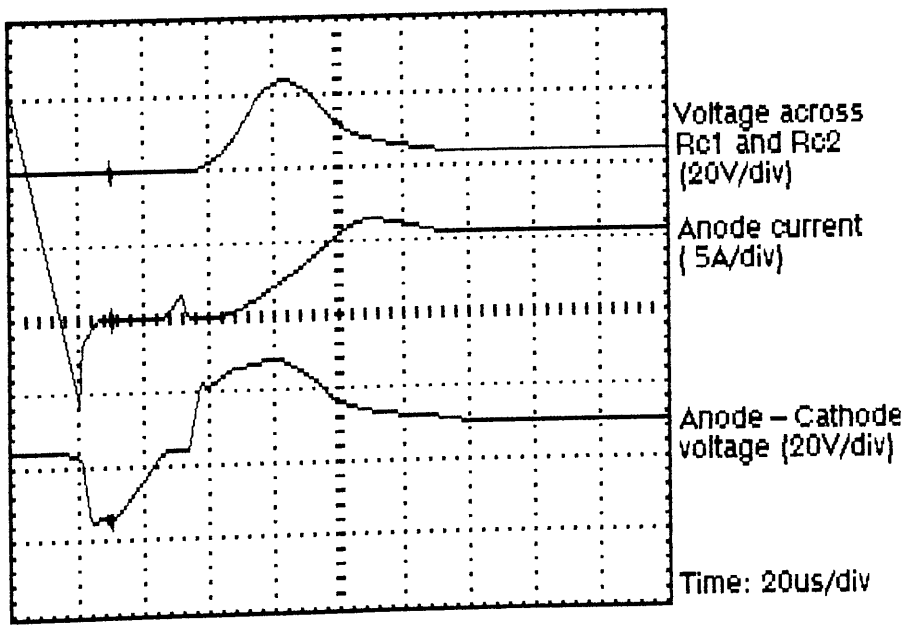

Figure 5.3.13: Recovery failure waveforms for the first 100mm device, low dv/dt.

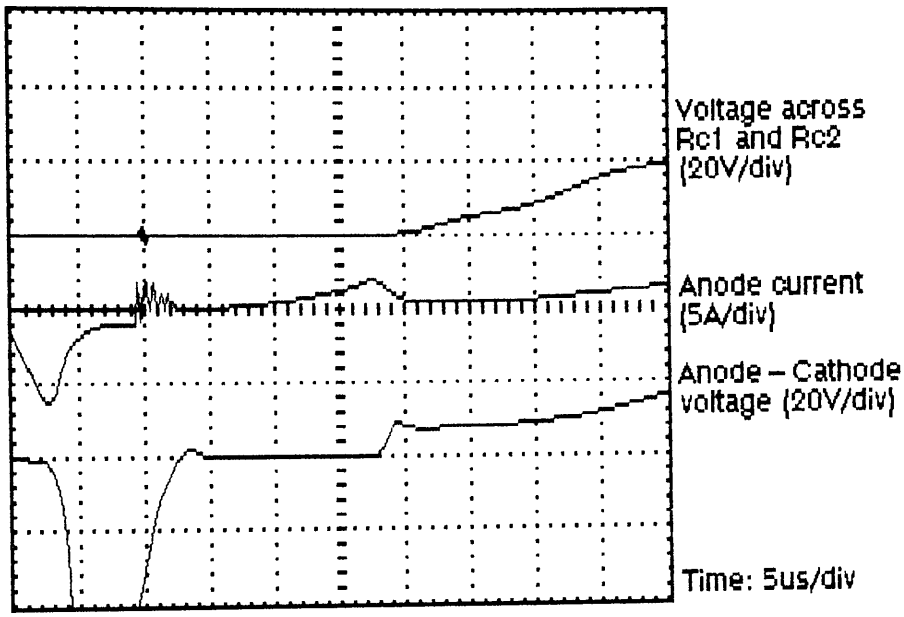

Figure 5.3.14: Recovery failure waveforms for the first 100mm device, low dv/dt, with expanded time-base.

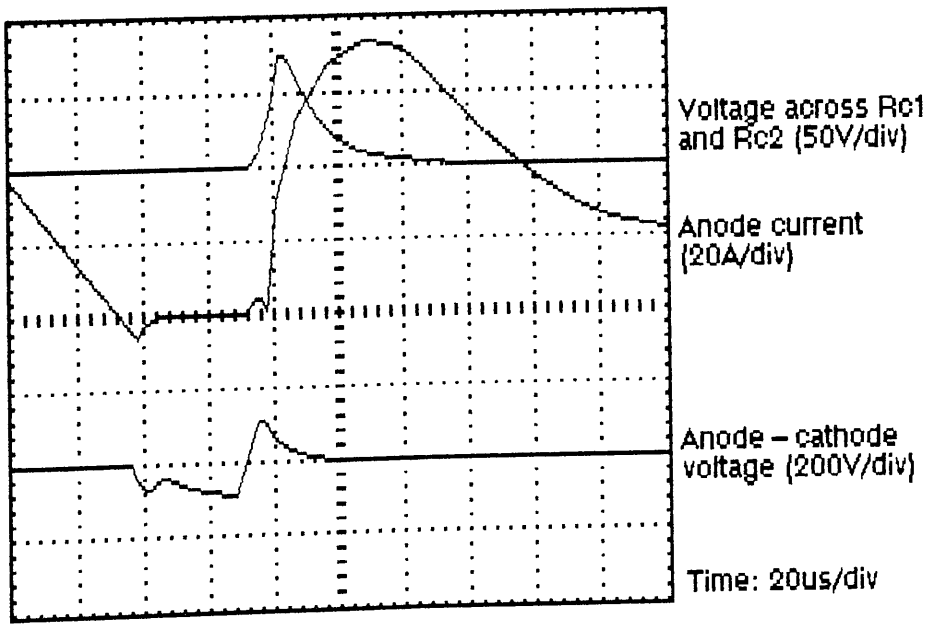

Figure 5.3.15: Recovery failure waveforms for first 100mm device, high dv/dt.

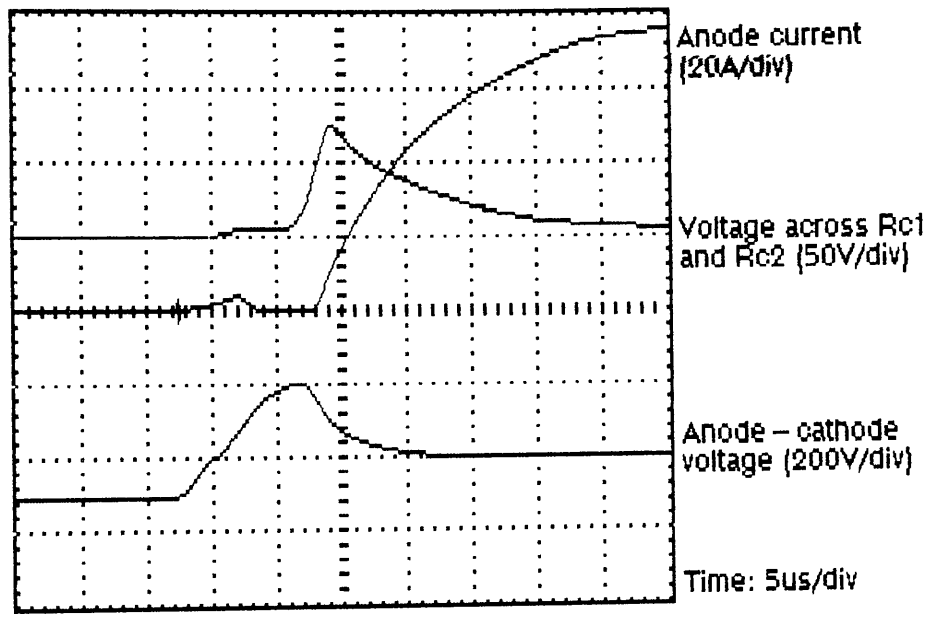

Figure 5.3.16: Recovery failure waveforms for the first 100mm device, high dv/dt, with expanded time-base.

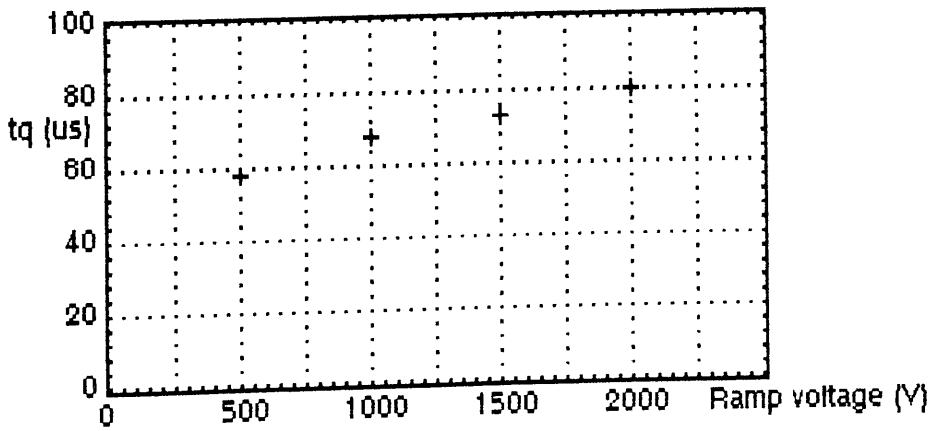

Figure 5.3.17: Recovery time versus ramp voltage.

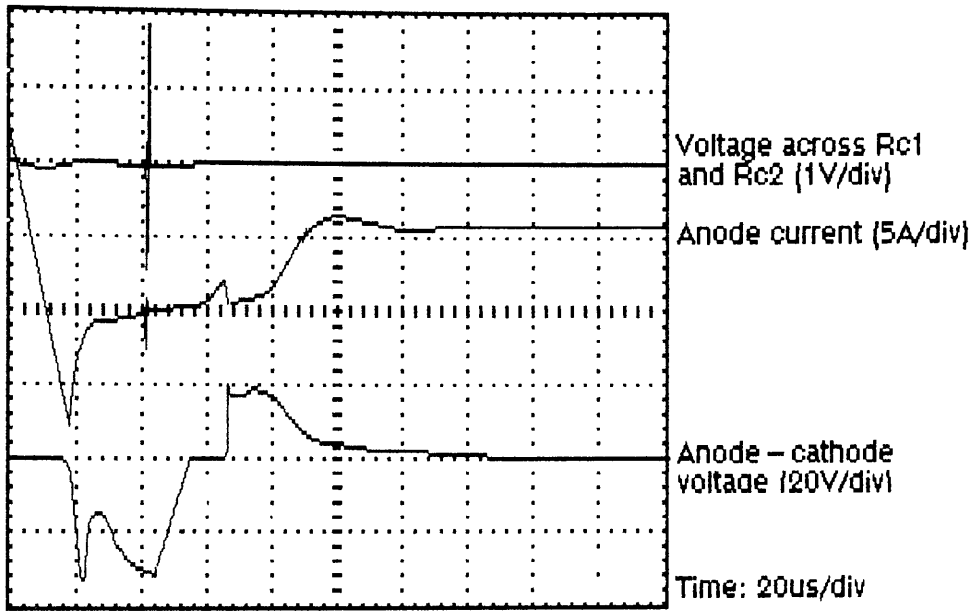

Figure 5.3.18: Recovery failure waveforms for the unirradiated device, low dv/dt.

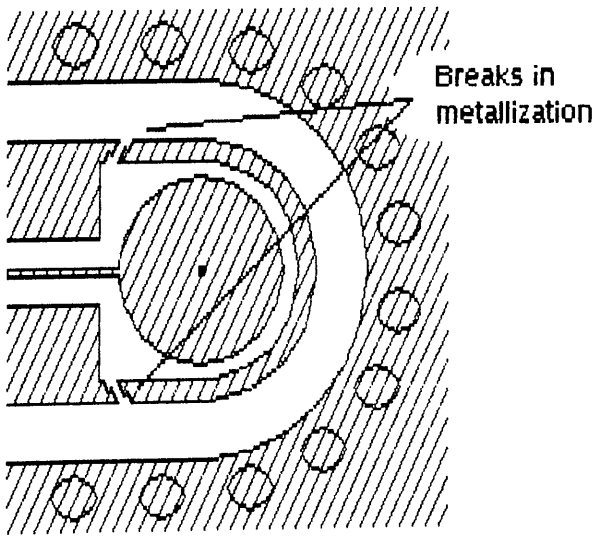

Figure 5.3.19: Breaks in auxiliary metallization to prevent scorching.

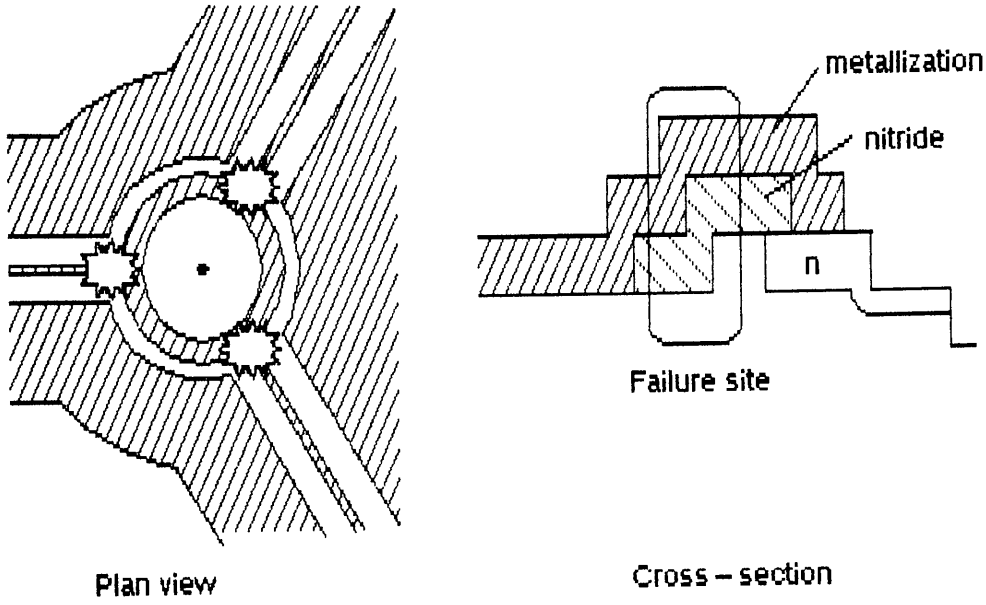

Figure 5.3.20: Failure sites round the optical well.

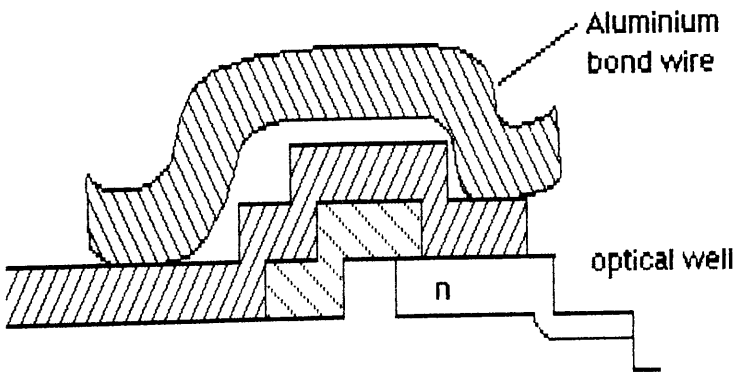

Figure 5.3.21: Reinforcement of steps using aluminium bond wire.

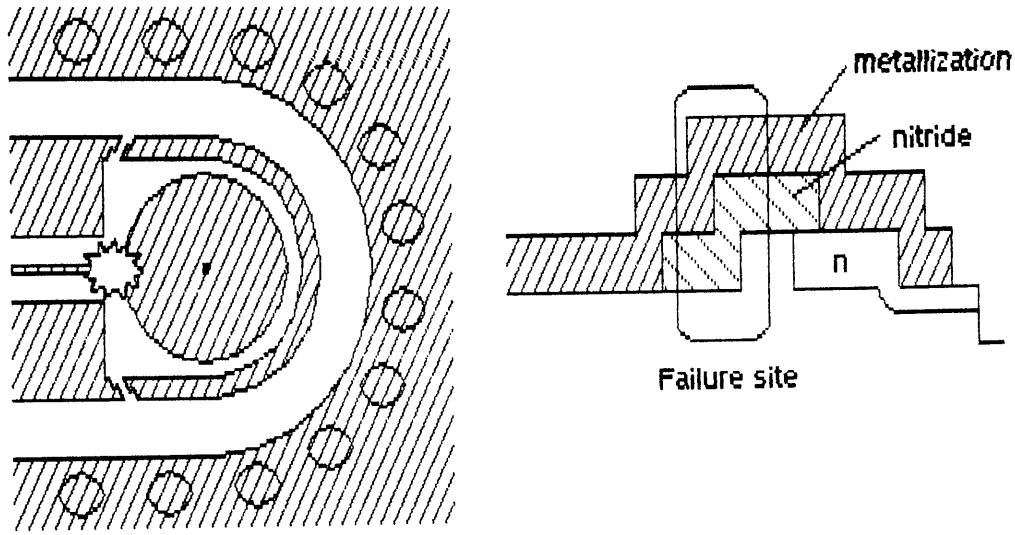

Figure 5.3.22: Failure site where the radial spoke leaves the satellite well.

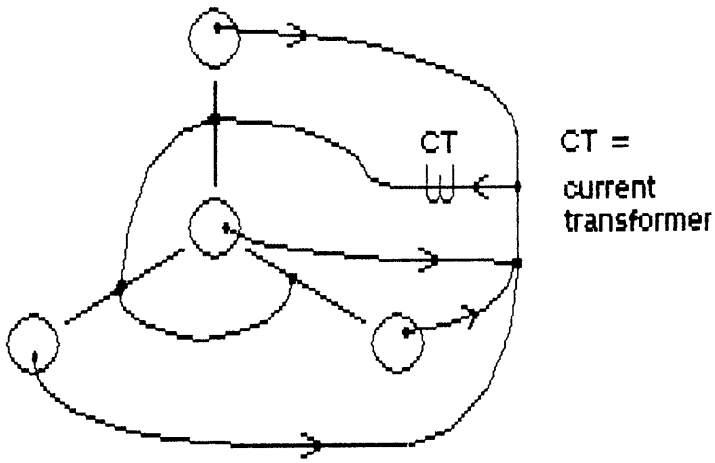

Figure 5.3.23: Arrangement for monitoring current through the wells.

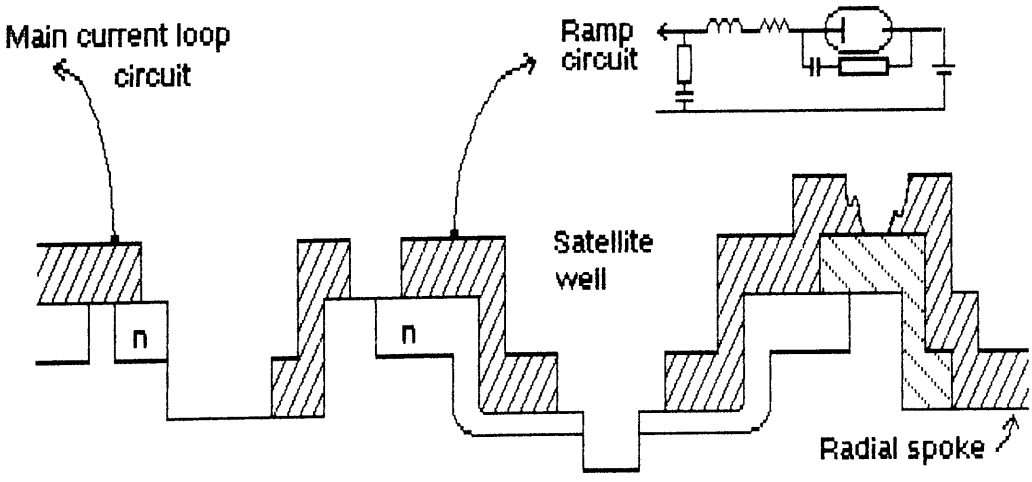

Figure 5.4.1: Ramp circuit applied to isolated satellite well.

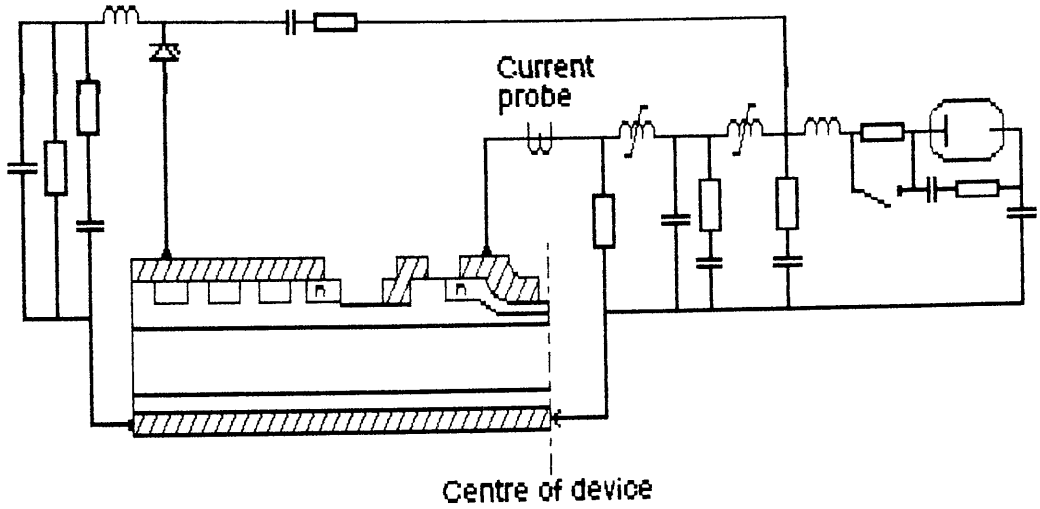

Figure 5.4.2: Test circuit for measuring charge extracted from the well.

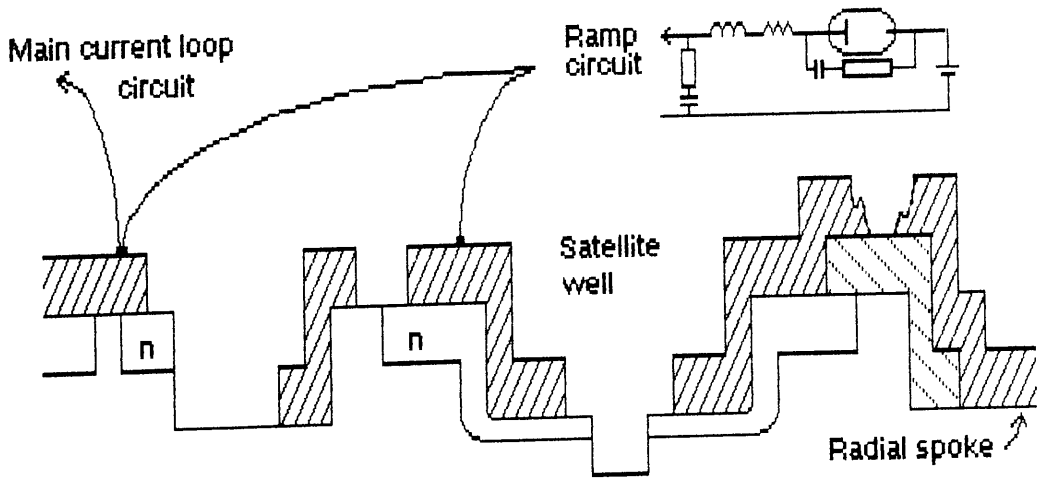

Figure 5.4.3: Ramp circuit applied to satellite well and main cathode.

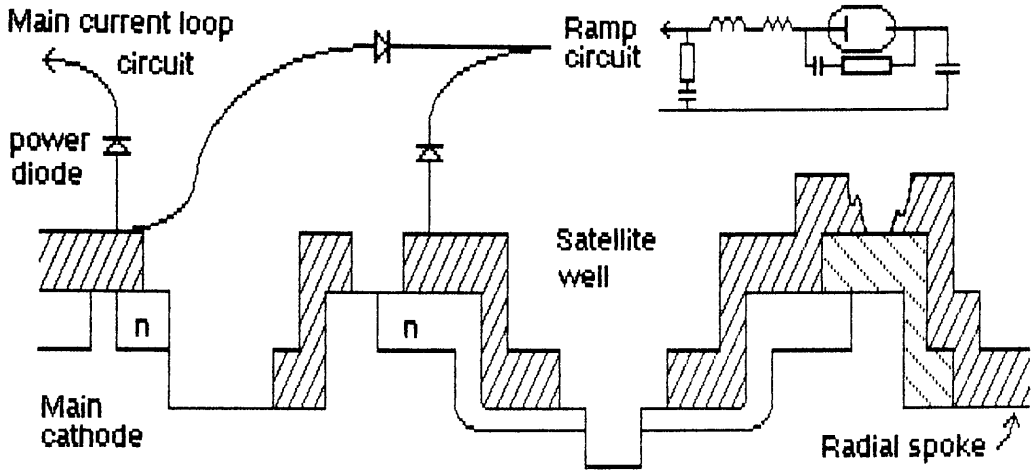

Figure 5.4.4: Arrangement of diodes in the charge extraction circuit.

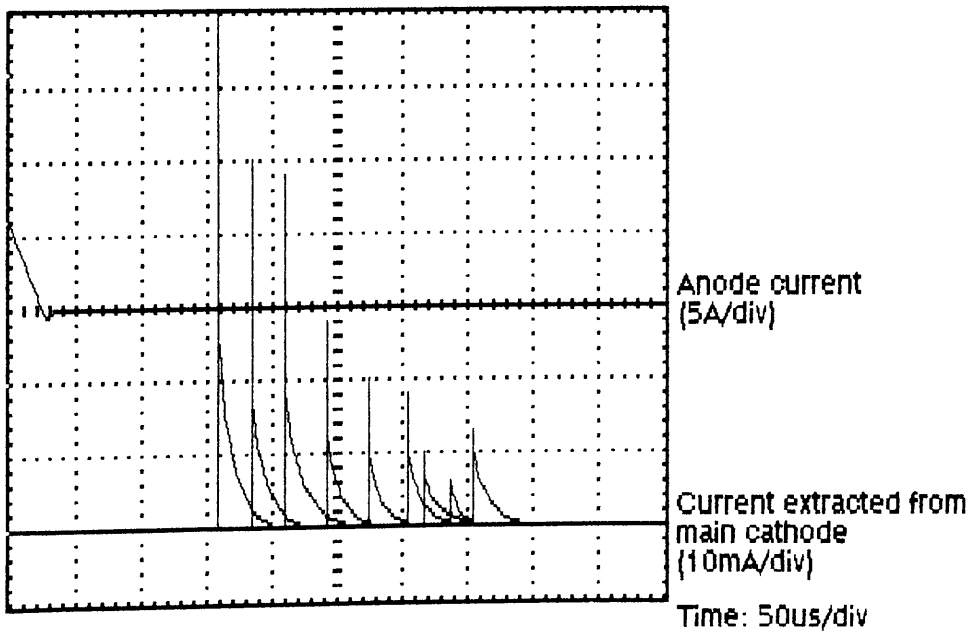

Figure 5.4.5: Charge extracted from main cathode of 30mm device.

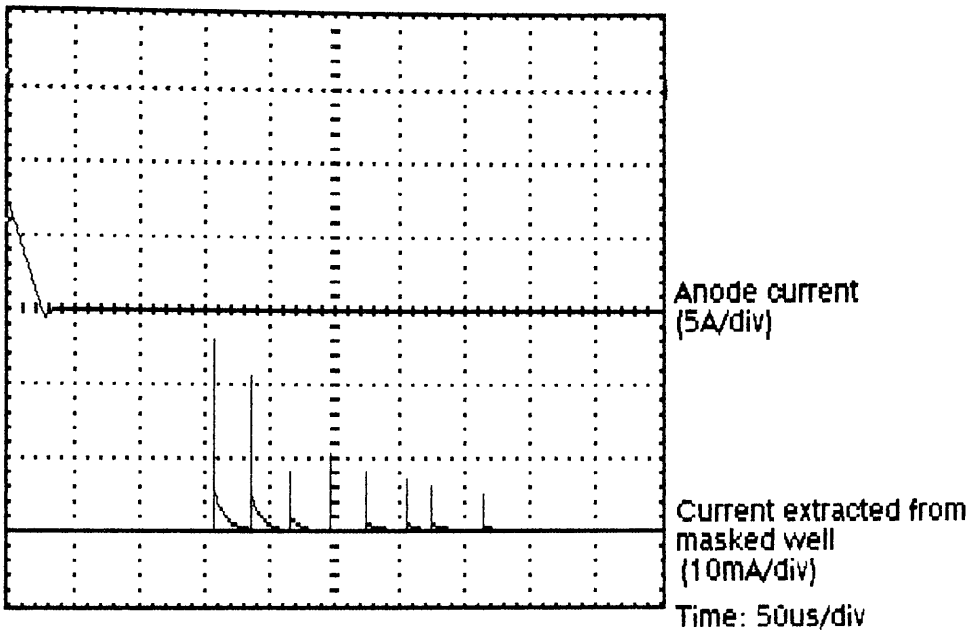

Figure 5.4.6: Charge extracted from the masked well on the 30mm device.

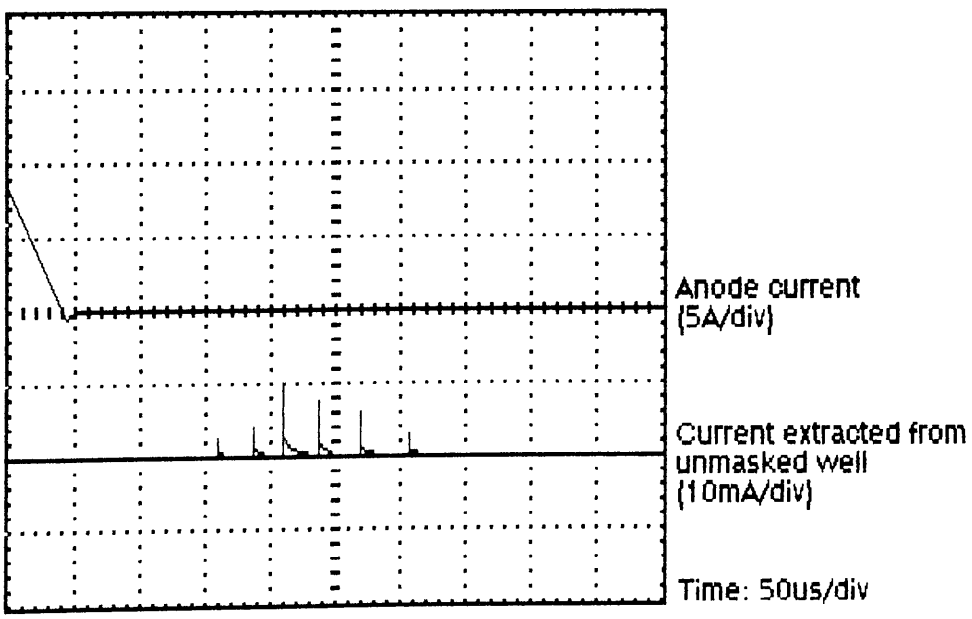

Figure 5.4.7: Charge extracted from the unmasked well of the 30mm device.

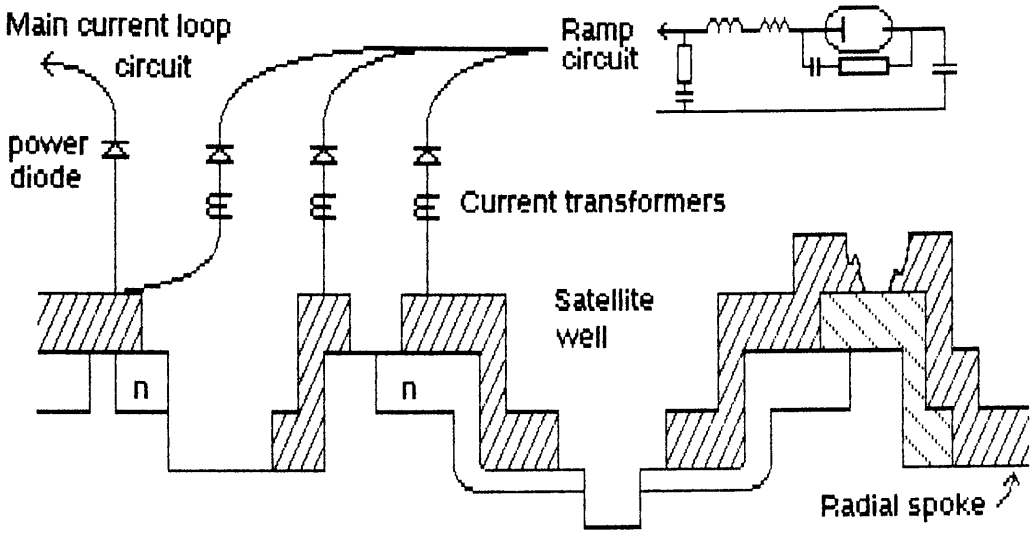

Figure 5.4.8: Arrangement for 100mm device measurements.

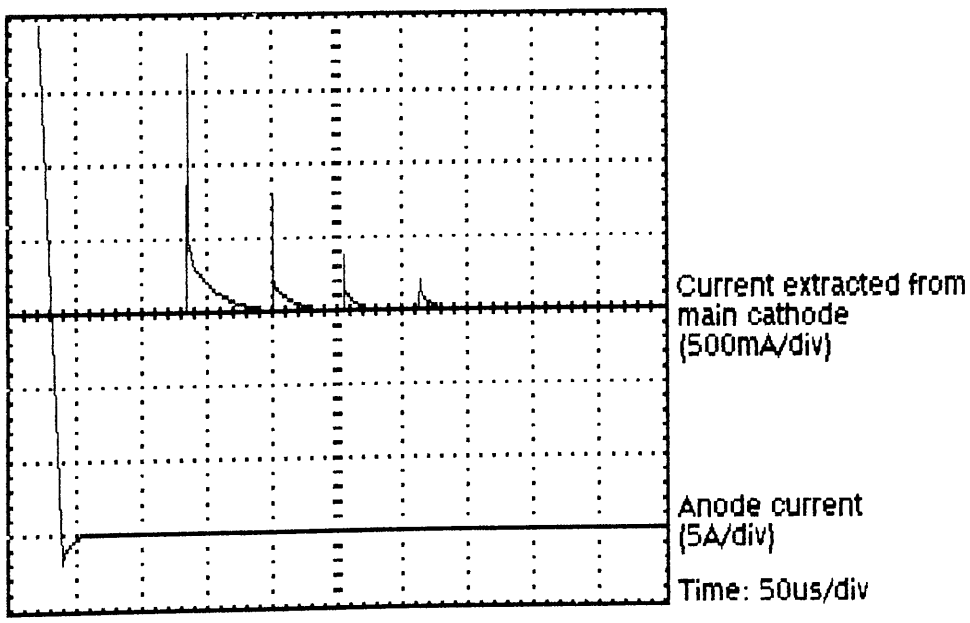

Figure 5.4.9: Charge extracted from the main cathode of the 100mm device.

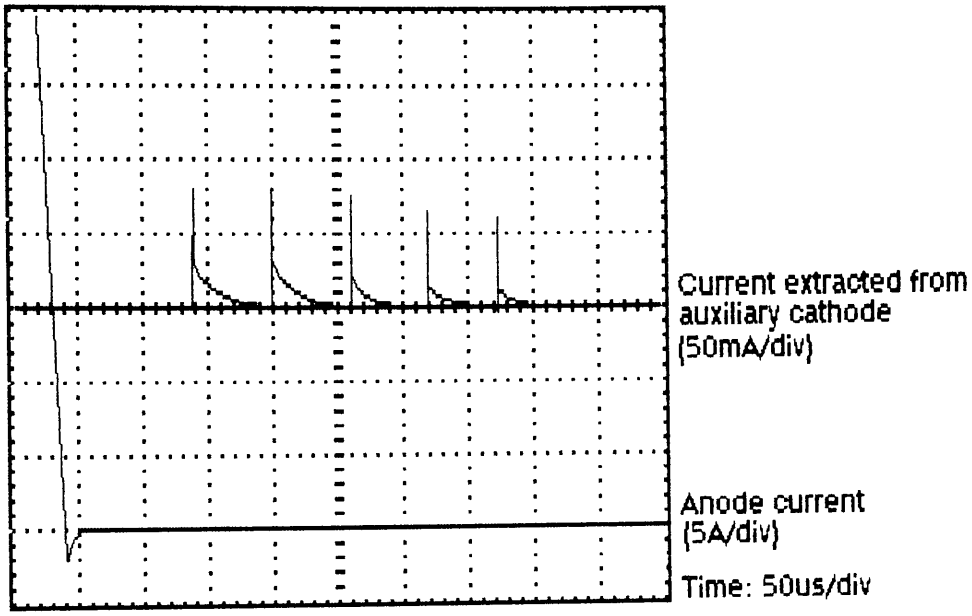

Figure 5.4.10: Charge extracted from the auxiliary region of the 100mm device.

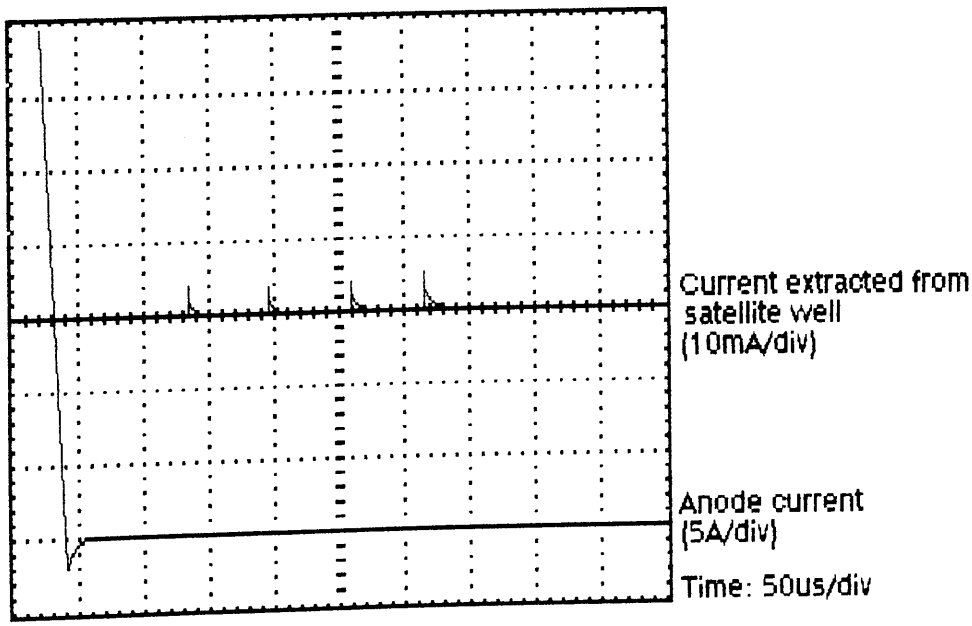

Figure 5.4.11: Charge extracted from the satellite well of the 100mm device.

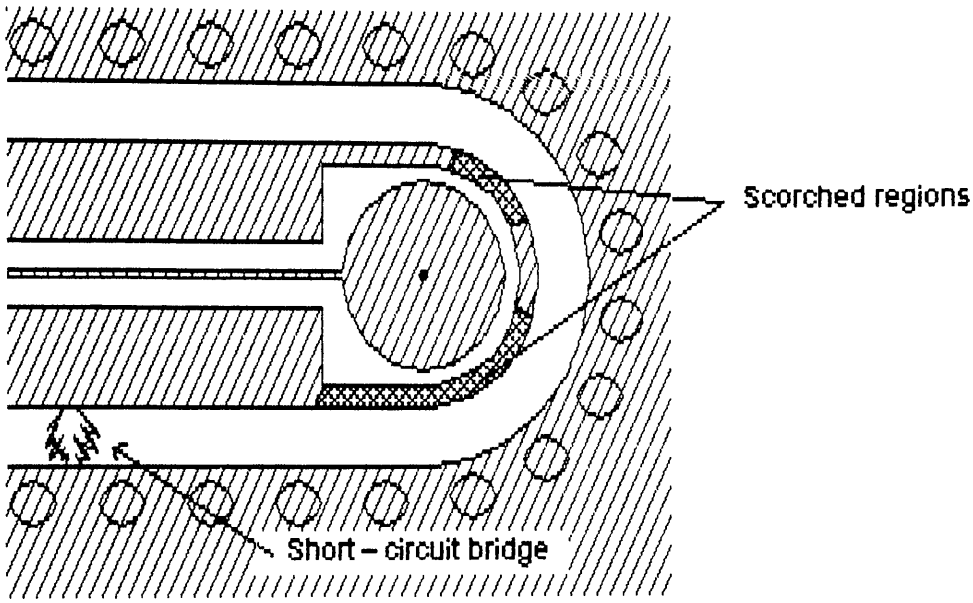

Figure 5.4.12: Damage to the 100mm device.

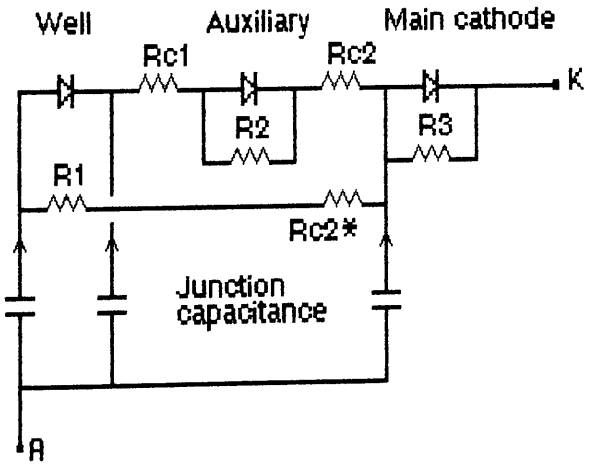

R1 = Resistance under well

R2 = Resistance under auxiliary

R3 = Resistance under main cathode

Rc1, Rc2 = control resistors

Rc2* = region of Rc2 round well

Figure 5.4.13: Dv/dt compensation with the equipotential point at the main cathode.

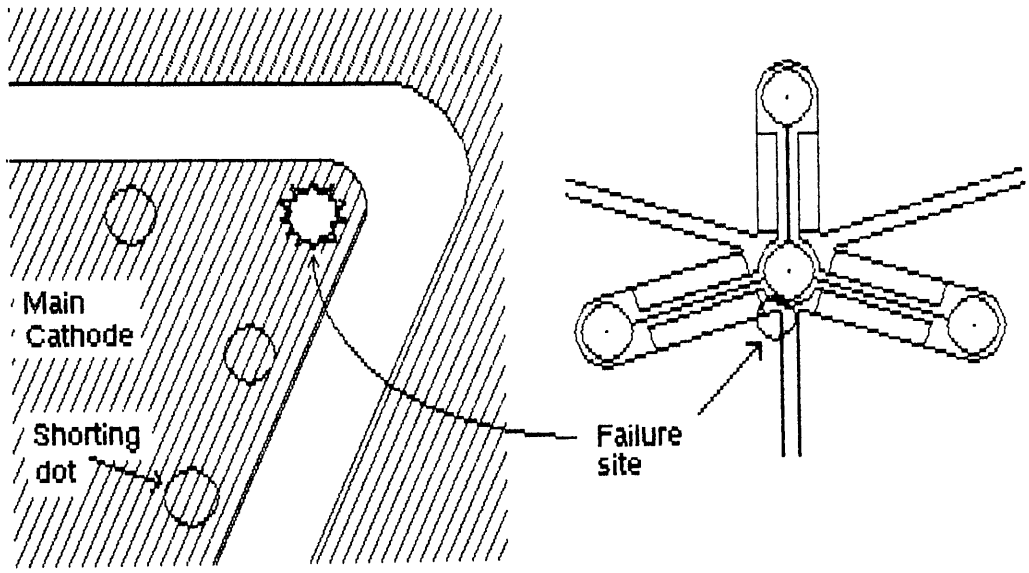

Figure 5.5.1: Location of the failure site for dv/dt testing.

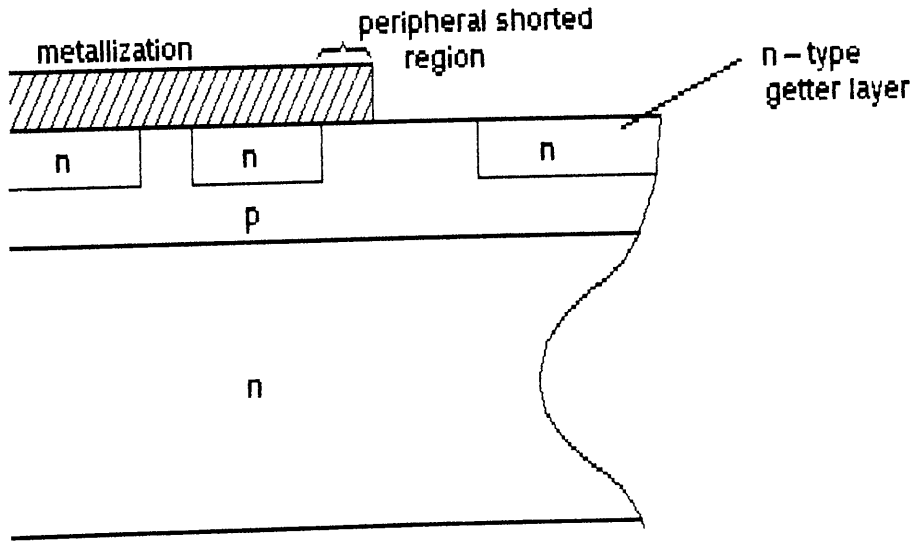

Figure 5.5.2: Cross-section through the edge bevel showing the peripheral getter layer.

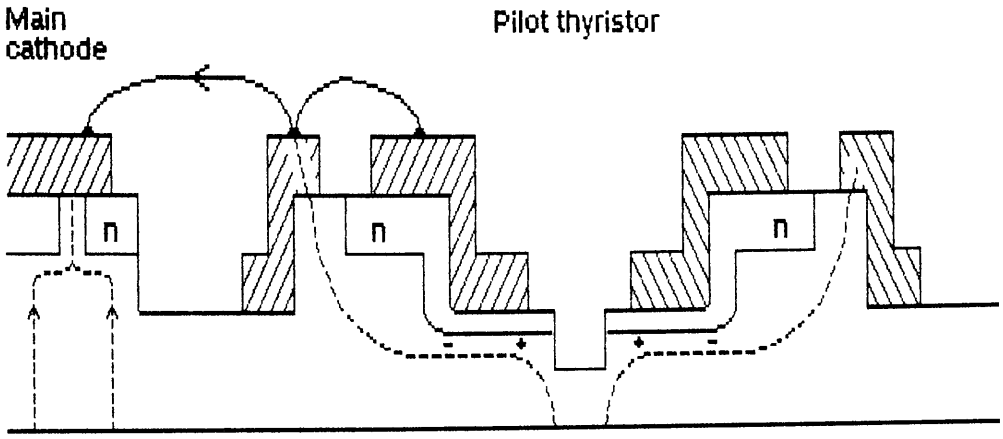

Figure 5.5.3: Shorting of the gate structure for dv/dt tests.

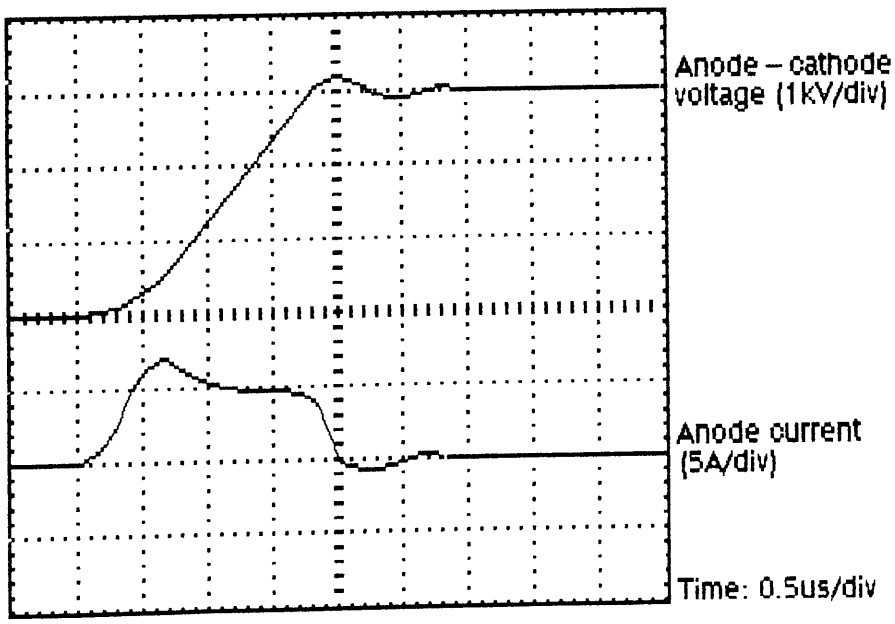

Figure 5.5.4: Typical dv/dt test results.

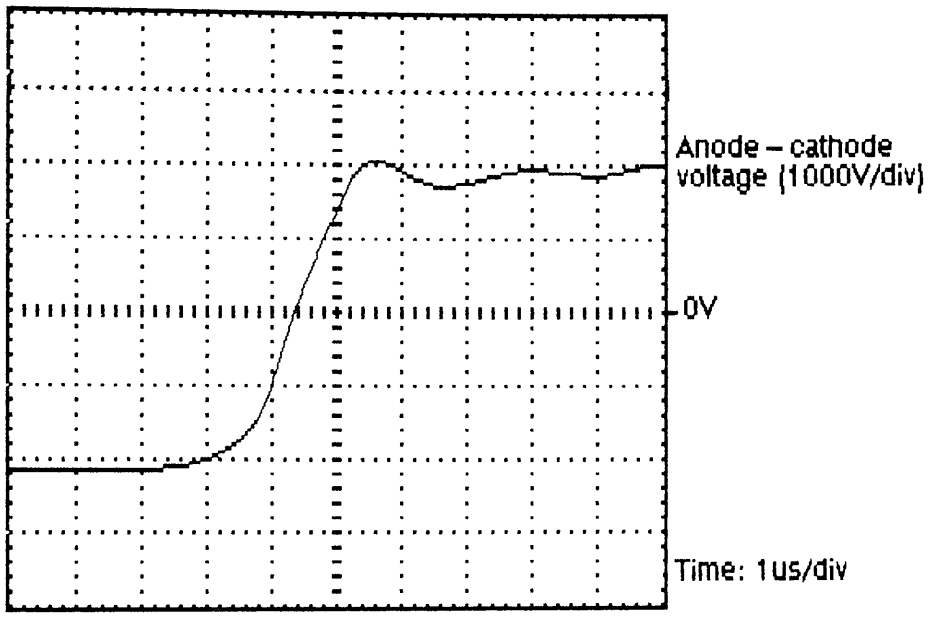

Figure 5.5.5: Typical worst-case transient applied to a thyristor in a valve when another series valve turns on.

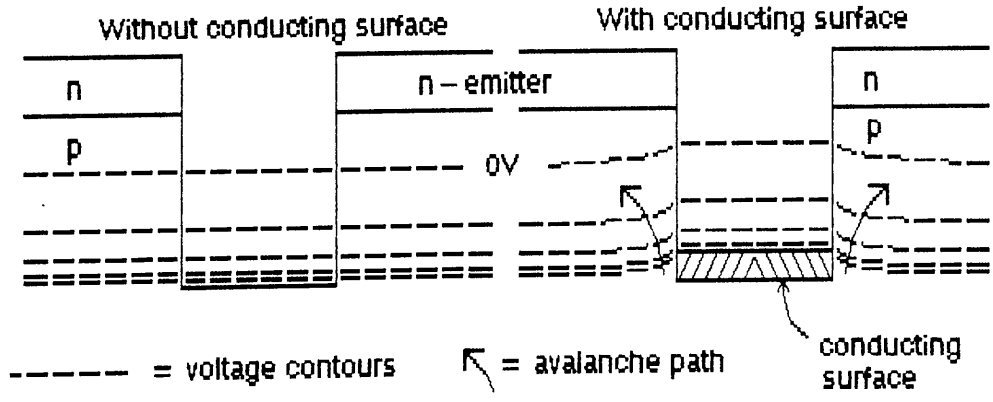

Figure 5.6.1: Estimated electric field patterns with and without a conducting layer in the bottom of the well.

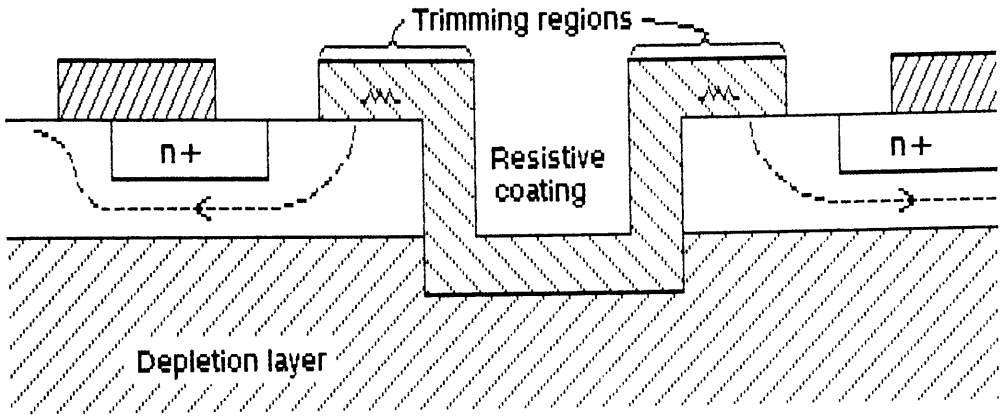

Figure 5.6.2: VBO triggering with an "over-etched" well.

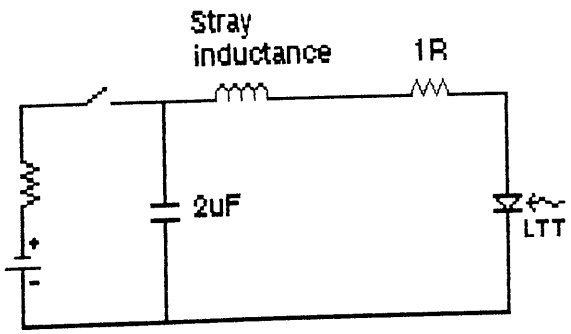

Figure 5.7.1: Preliminary turn-on test circuit.

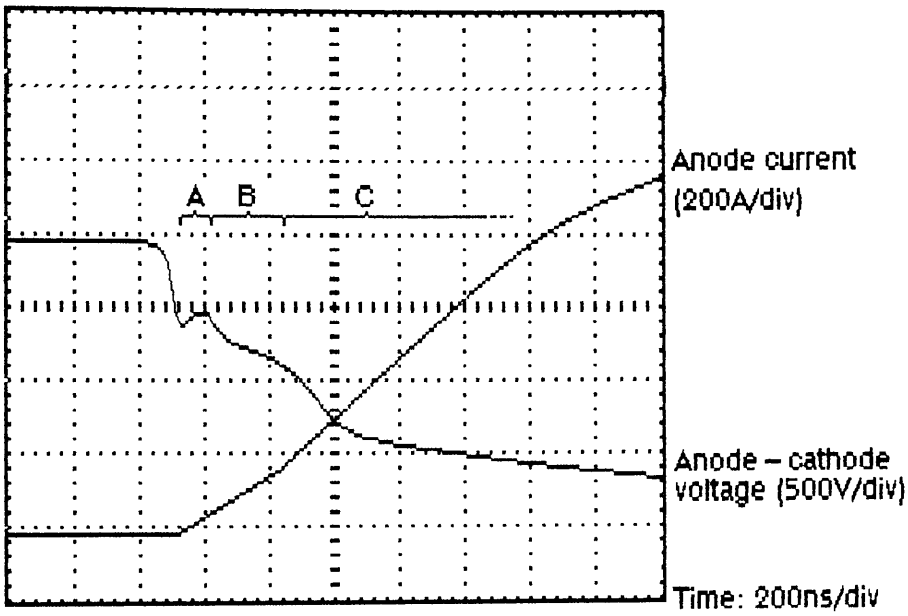

Figure 5.7.2: Turn-on waveforms for the 30mm device.

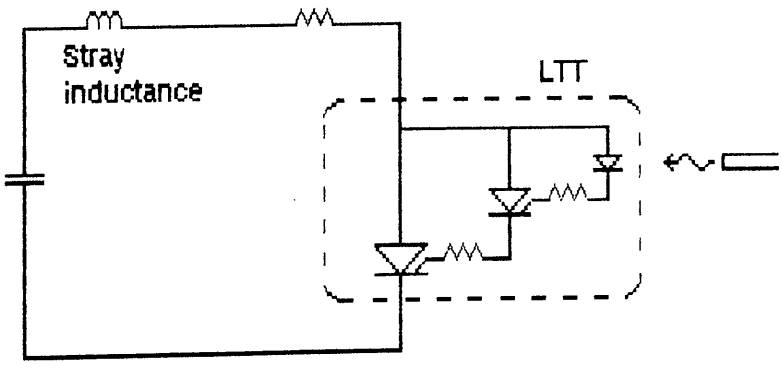

Figure 5.7.3: Equivalent circuit for preliminary turn-on tests.

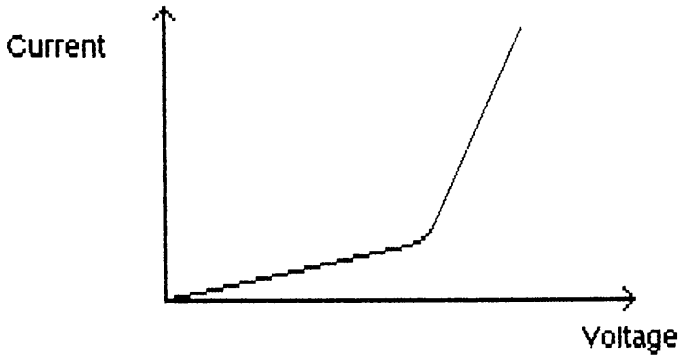

Figure 5.7.4: Shape of V-I characteristic of the control resistors.

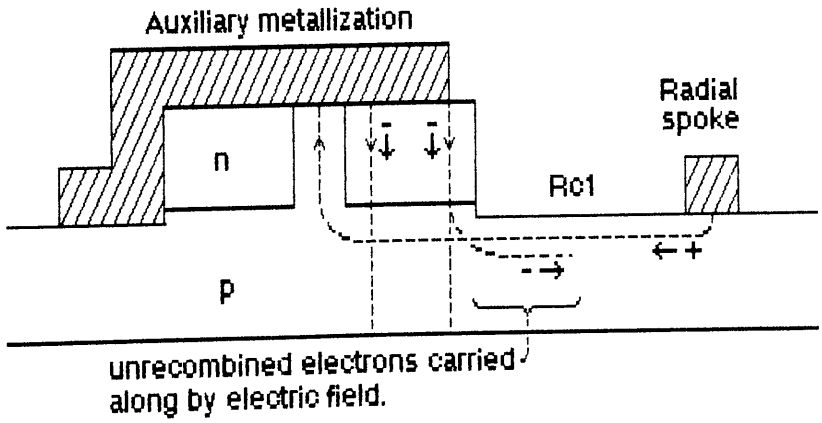

Figure 5.7.5: Injection of minority carriers causing resistor modulation.

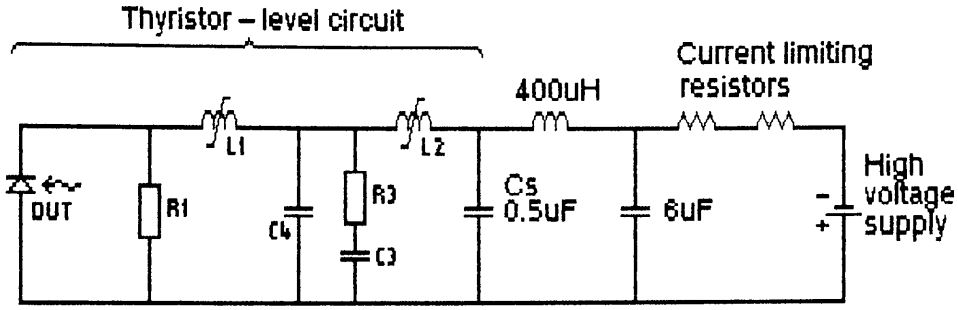

Figure 5.7.6: Turn-on test circuit for 100mm devices.

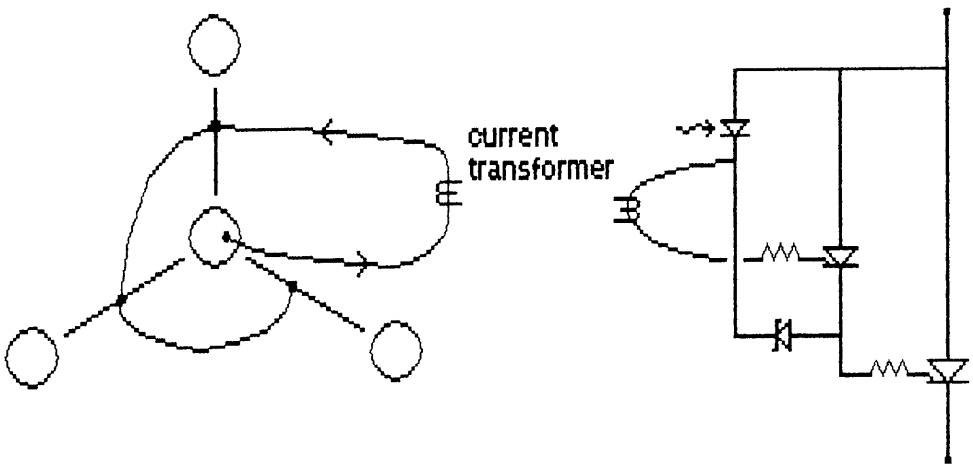

Figure 5.7.7: Arrangement for measuring current in the optical well.

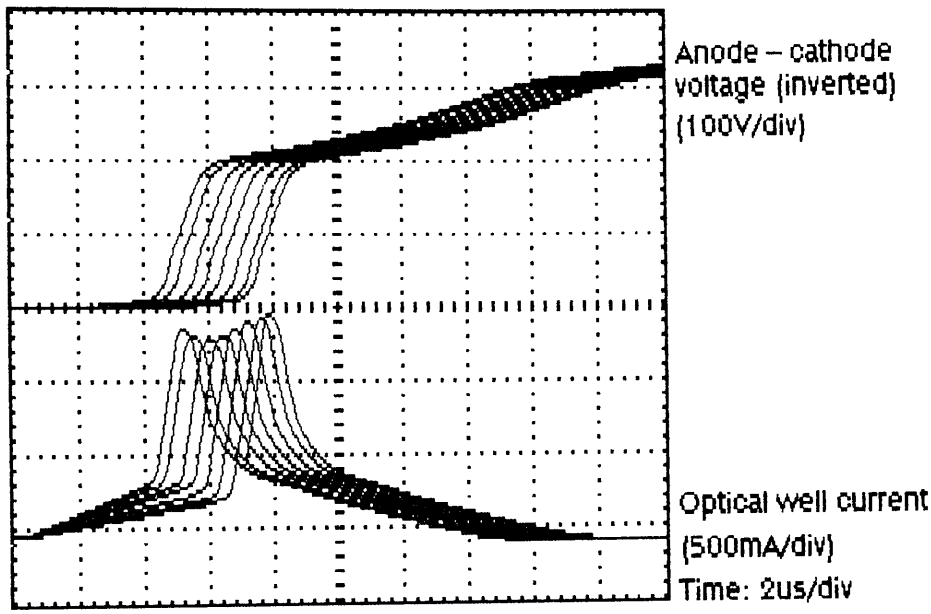

Figure 5.7.8: Turn-on waveforms for varying degrees of optical overdrive (400V turn-on).

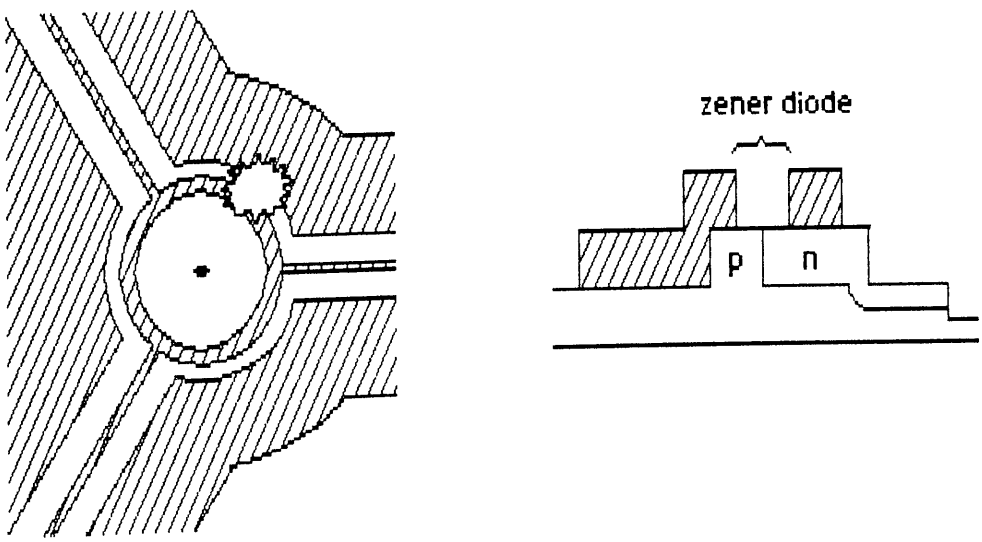

Figure 5.7.9: Failure site during turn-on testing.

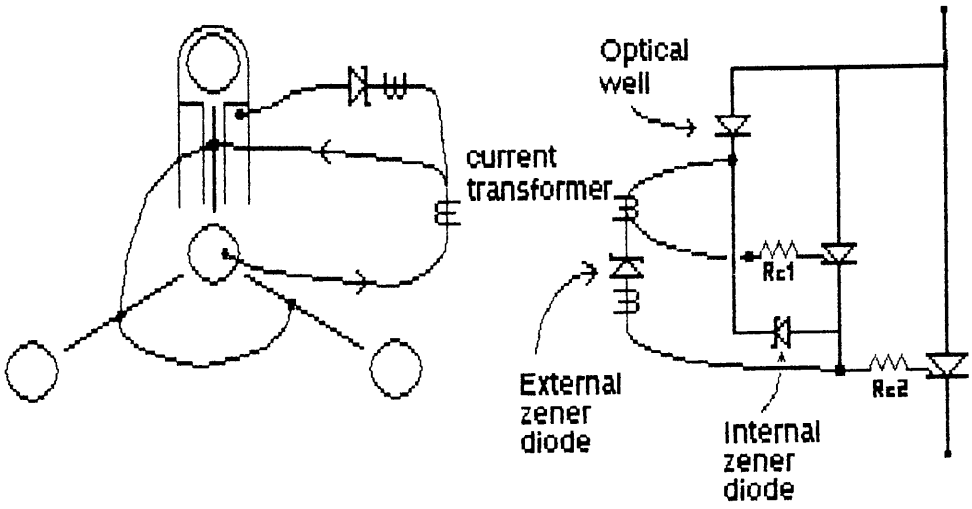

Figure 5.7.10: Connection of the external zener diode.

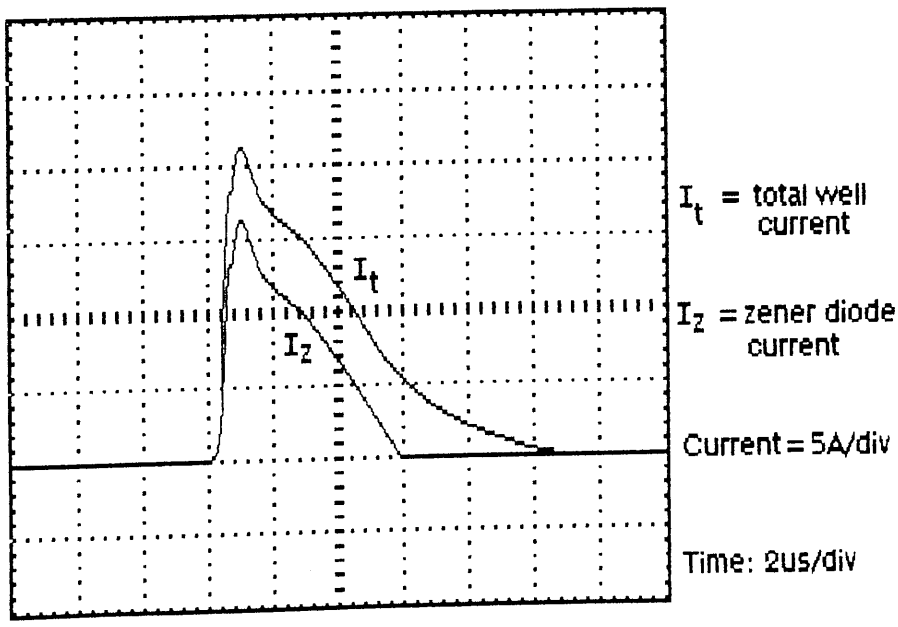

Figure 5.7.11: Turn-on current waveforms with external zener diode connected (1500V turn-on).

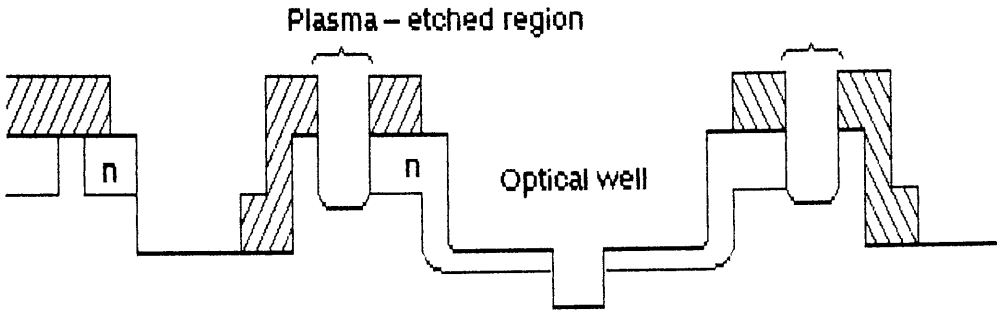

Figure 5.7.12: Plasma-etched region around the optical well.

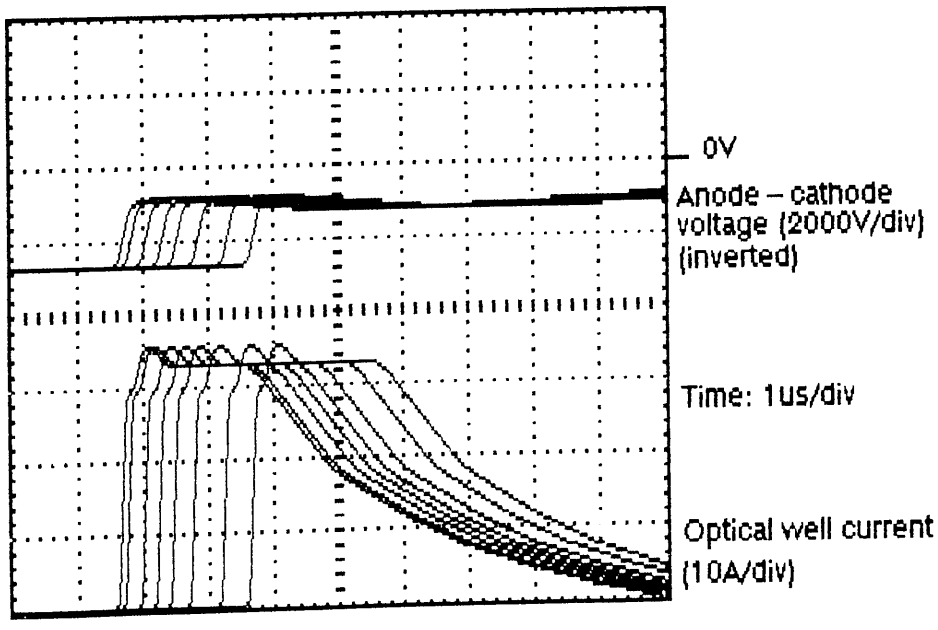

Figure 5.7.13: Turn-on waveforms for the plasma-etched device.

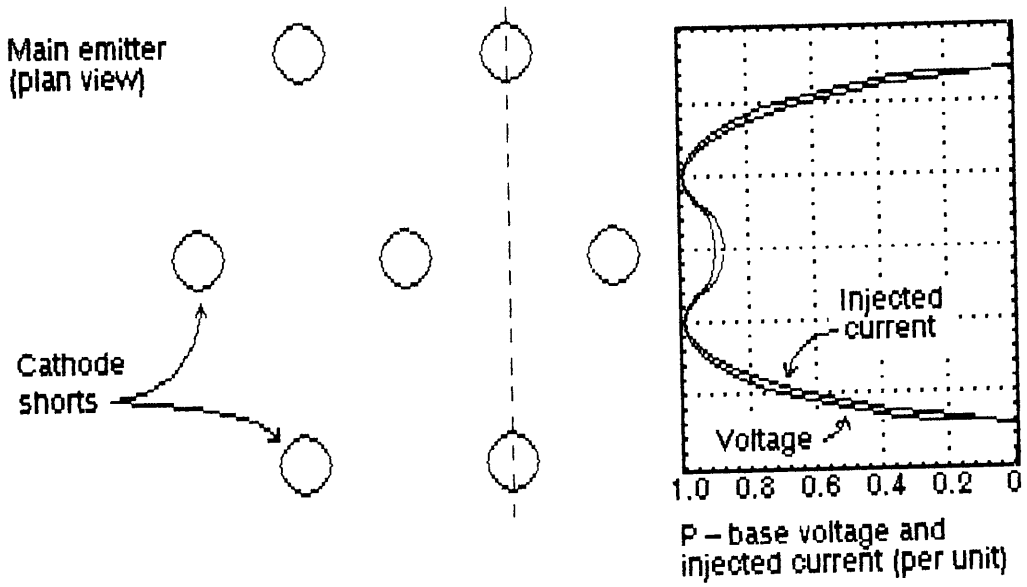

Figure 6.2.1: P-base voltage and injected current during a forward voltage ramp. (Graph shows variation along dashed line).

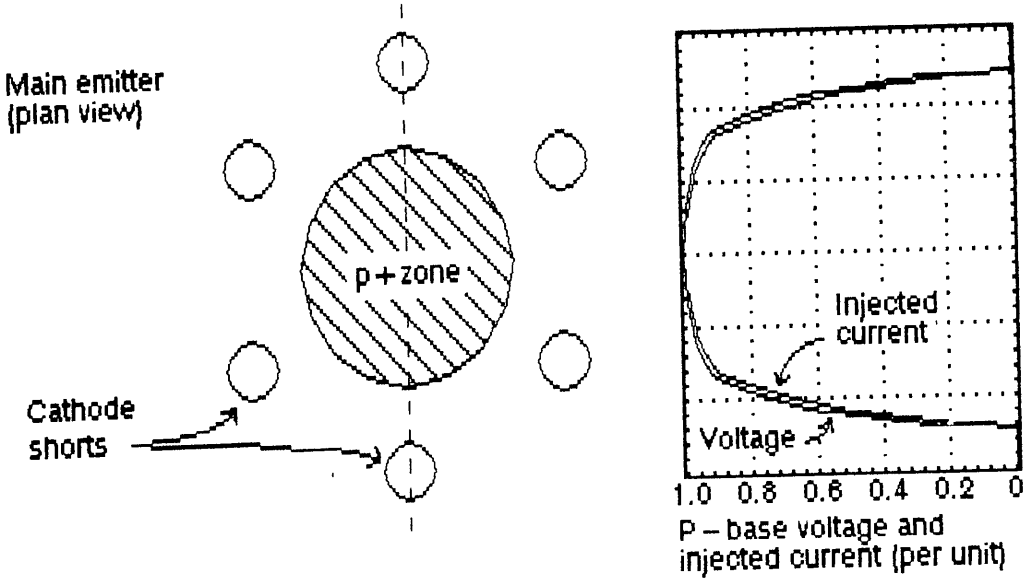

Figure 6.2.2: P-base voltage and injected current during a forward voltage ramp, as modified by the p+ zone. (Graph shows variation along dashed line).

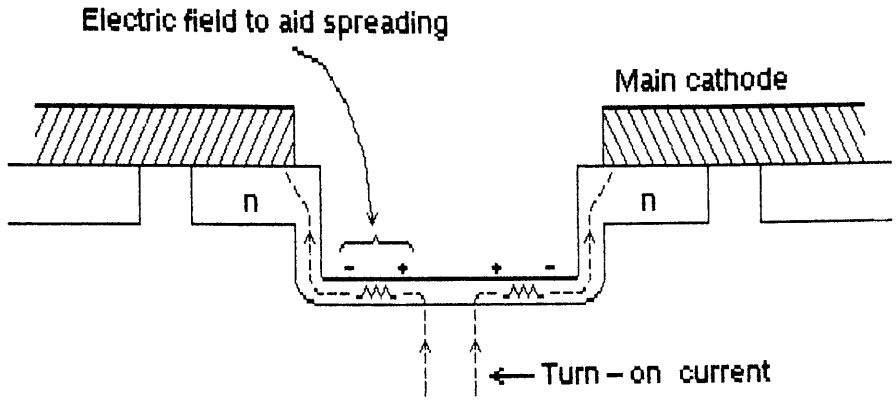

Figure 6.2.3: Lateral field created in a resistive region during turn-on.

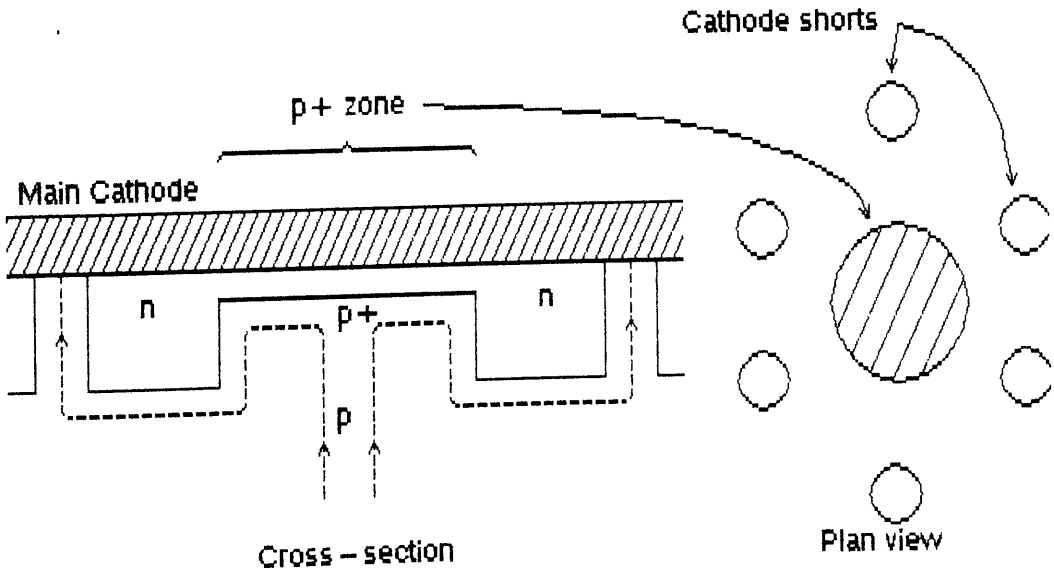

Figure 6.2.4: Design of the "p+ zone" devices.

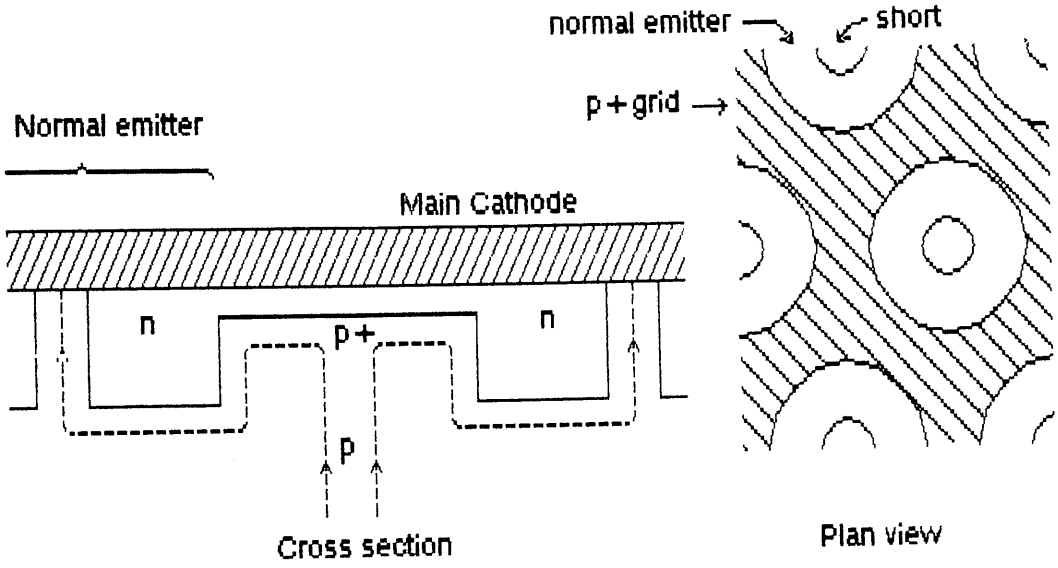

Figure 6.2.5: Design of the "p+ grid" device.

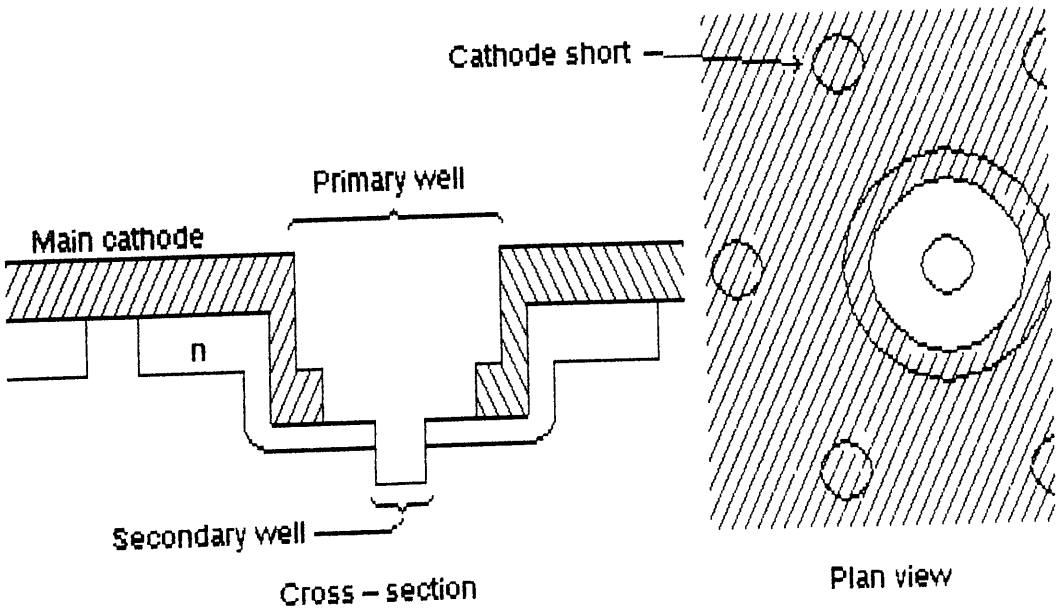

Figure 6.2.6: Design of the "Lateral Field" device.

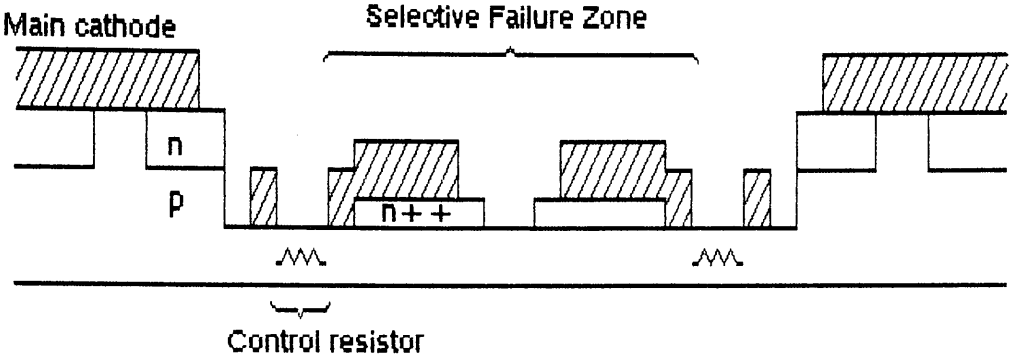

Figure 6.2.7: Design of the "Selective Failure Zone" device.

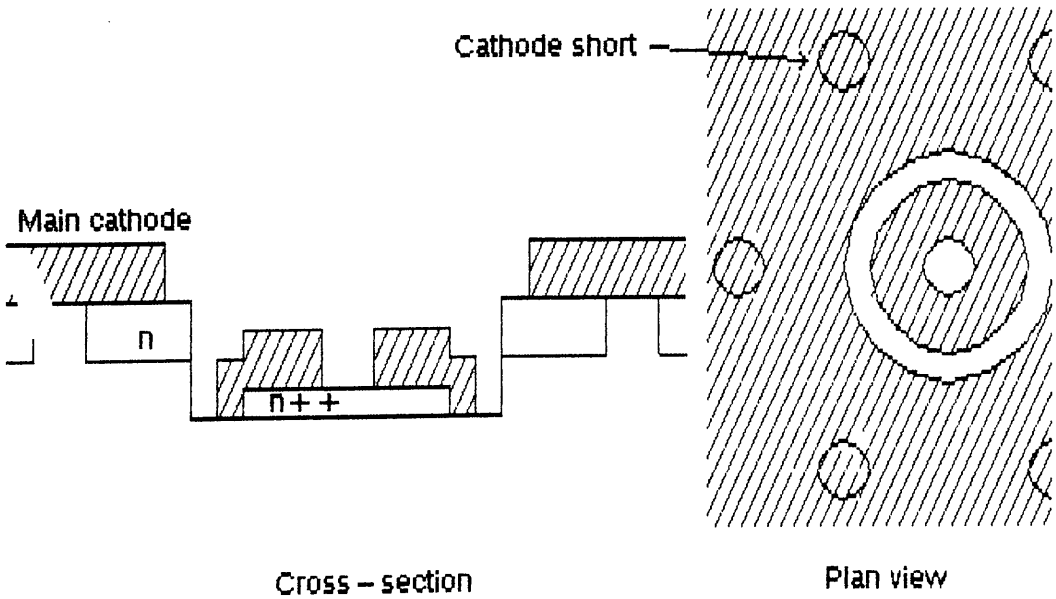

Figure 6.2.8: Design of the "Gated Turn-on" device, with amplifier.

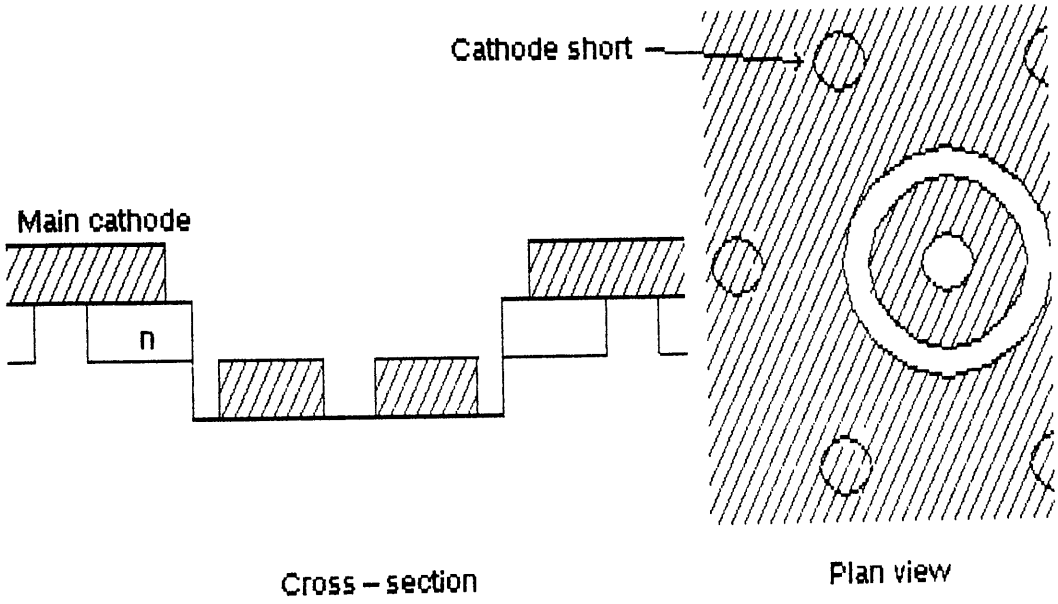

Figure 6.2.9: Design of the "Gated Turn-on" device, without amplifier.

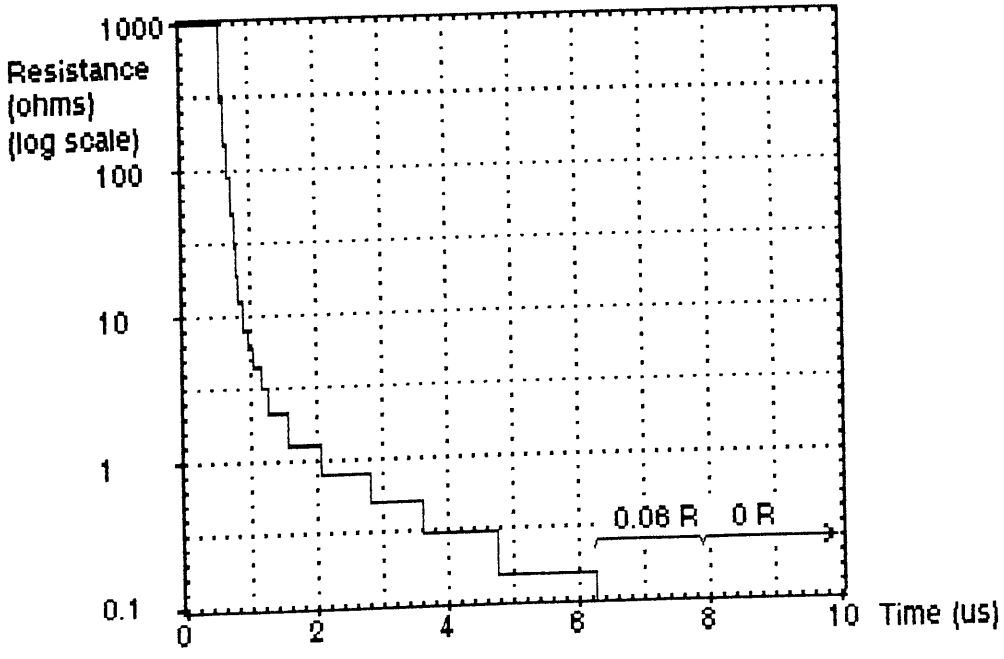

Figure 6.3.1: Resistance versus time characteristic used to model thyristor turn-on.

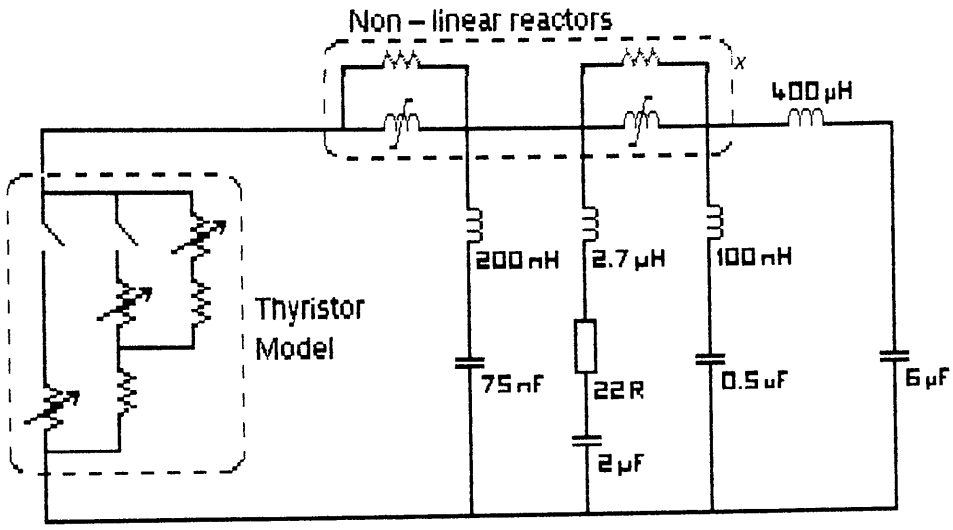

Figure 6.3.2: Basic circuit used for computer modelling.

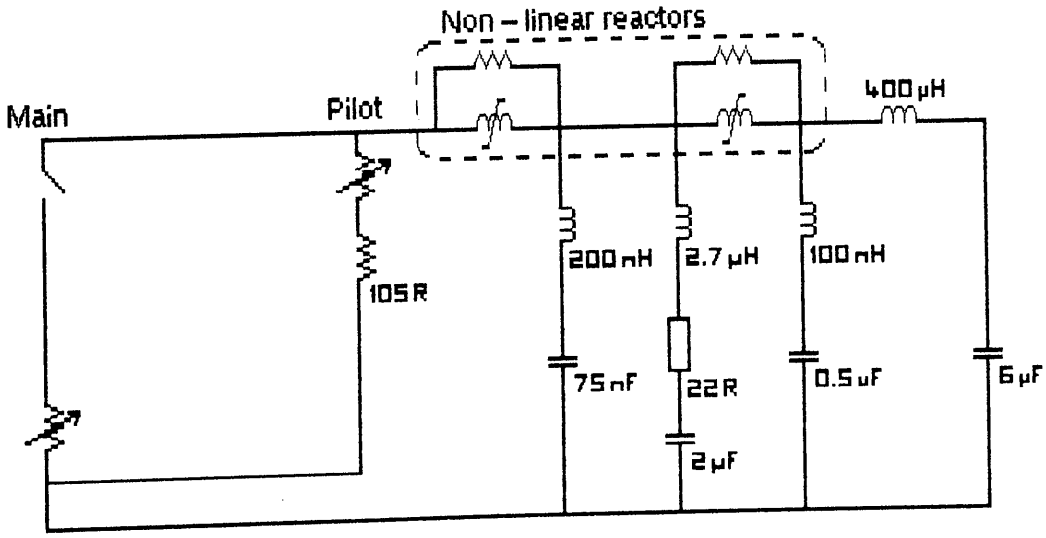

Figure 6.3.3: Single amplifying gate, 105R control resistor.

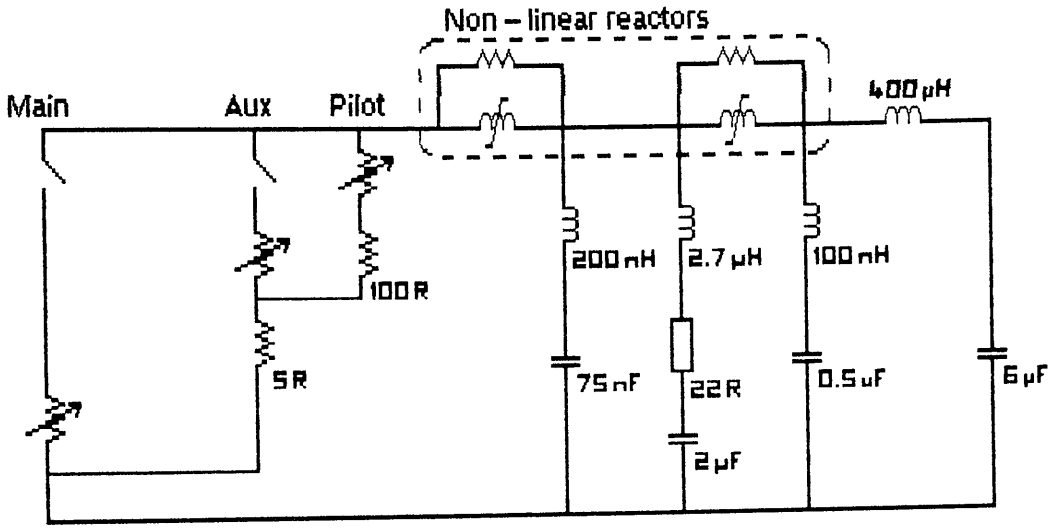

Figure 6.3.4: Double amplifying gate, 100R and 5R control resistors.

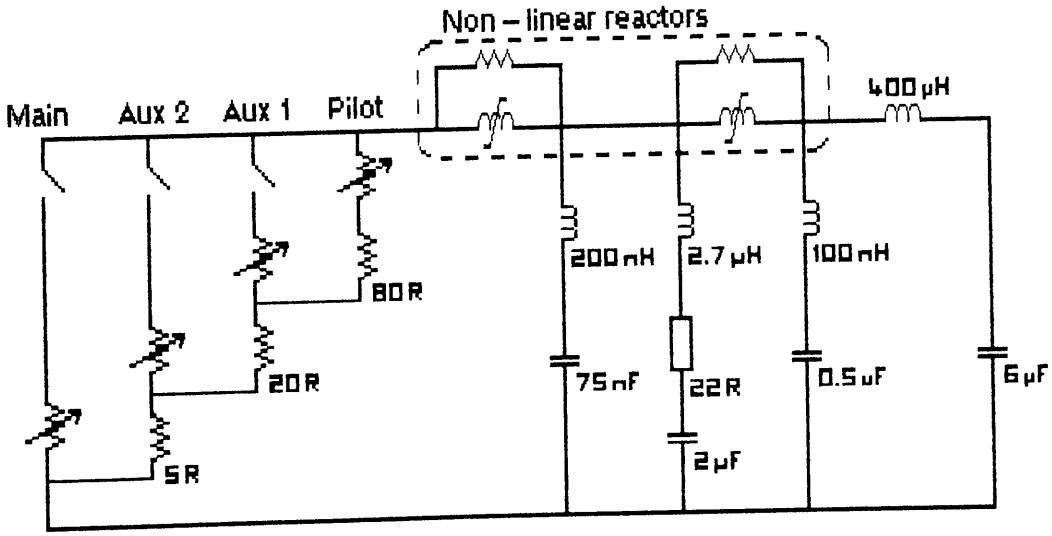

Figure 6.3.5: Treble amplifying gate, 80R, 20R and 5R control resistors.

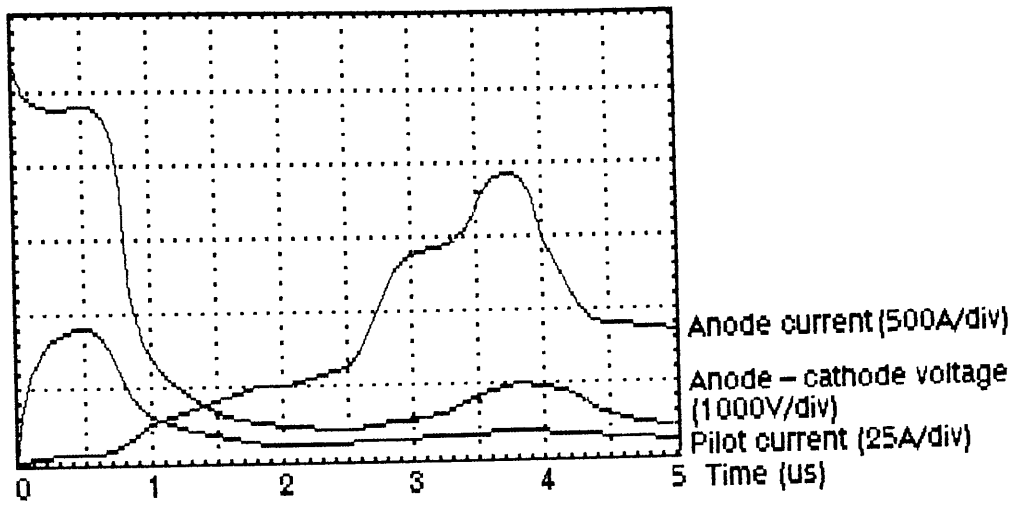

Figure 6.3.6: Turn-on waveforms for single amplifying gate model.

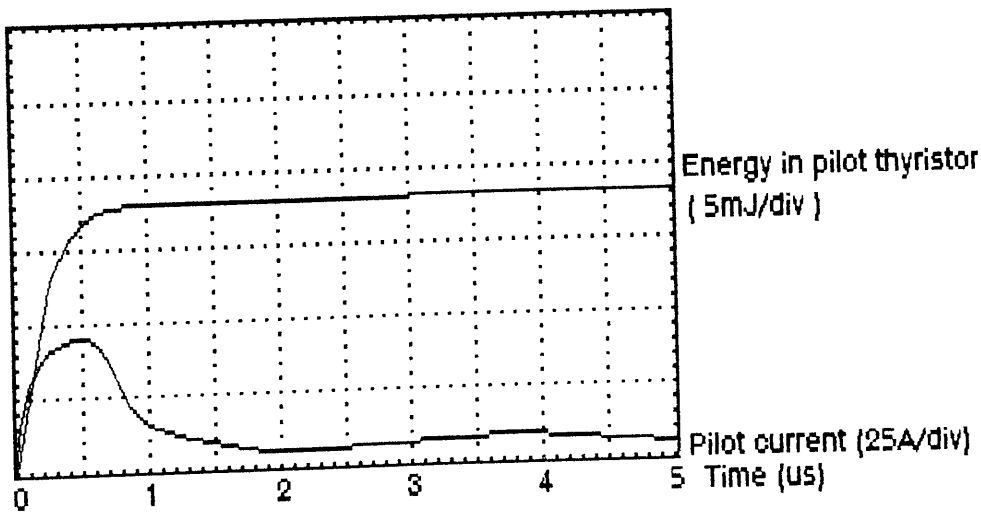

Figure 6.3.7: Current and energy in the pilot thyristor.

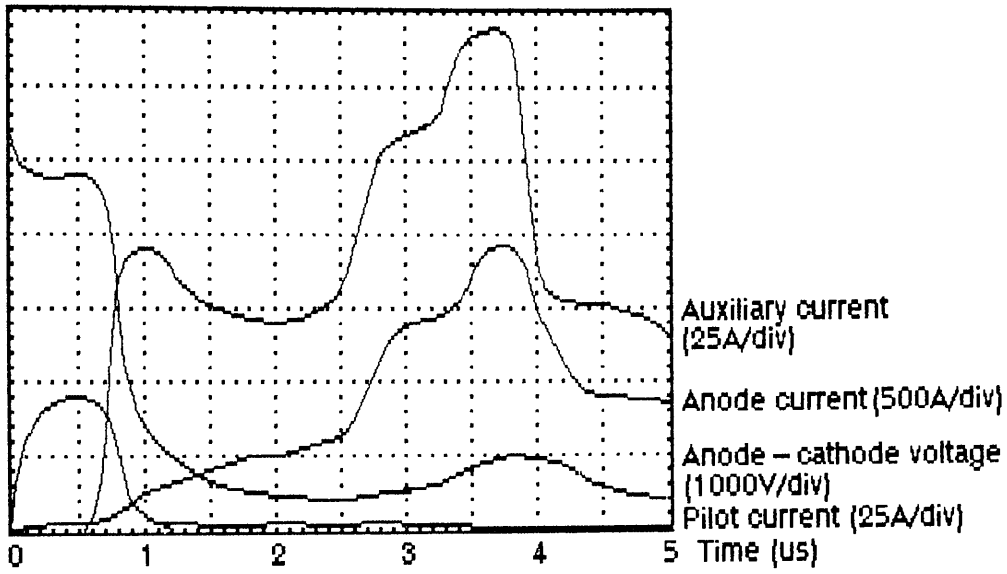

Figure 6.3.8: Turn-on waveforms for double amplifying gate.

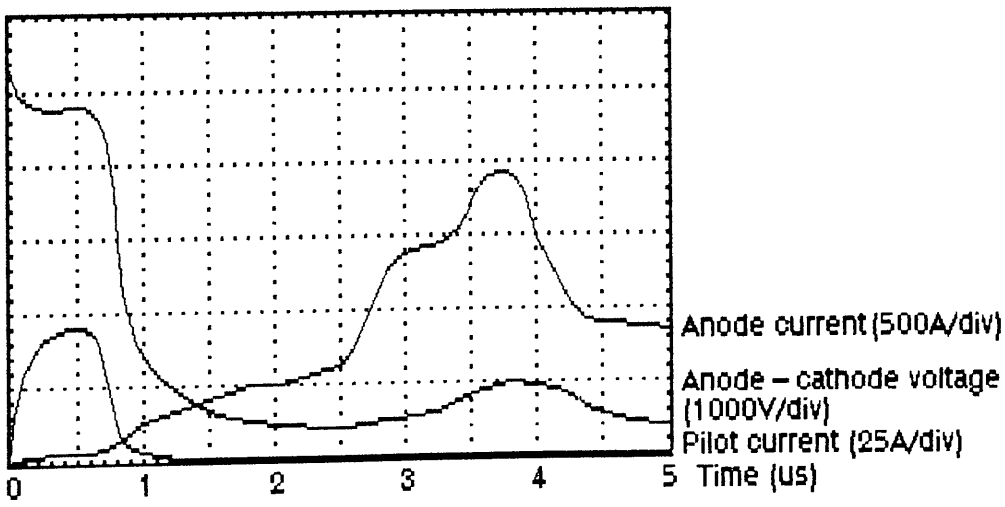

Figure 6.3.9: Turn-on waveforms for treble amplifying gate.

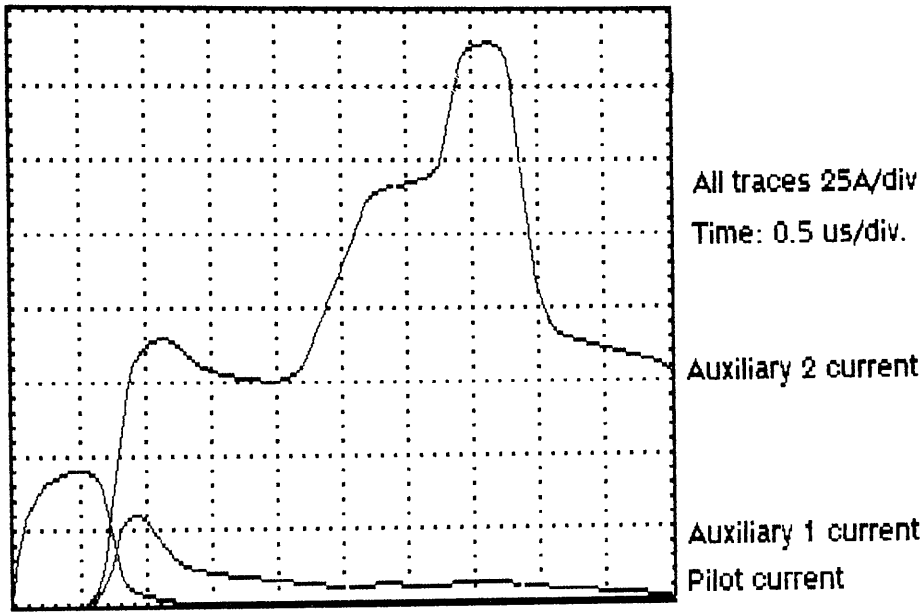

Figure 6.3.10: Gate current traces for the treble amplifying gate.

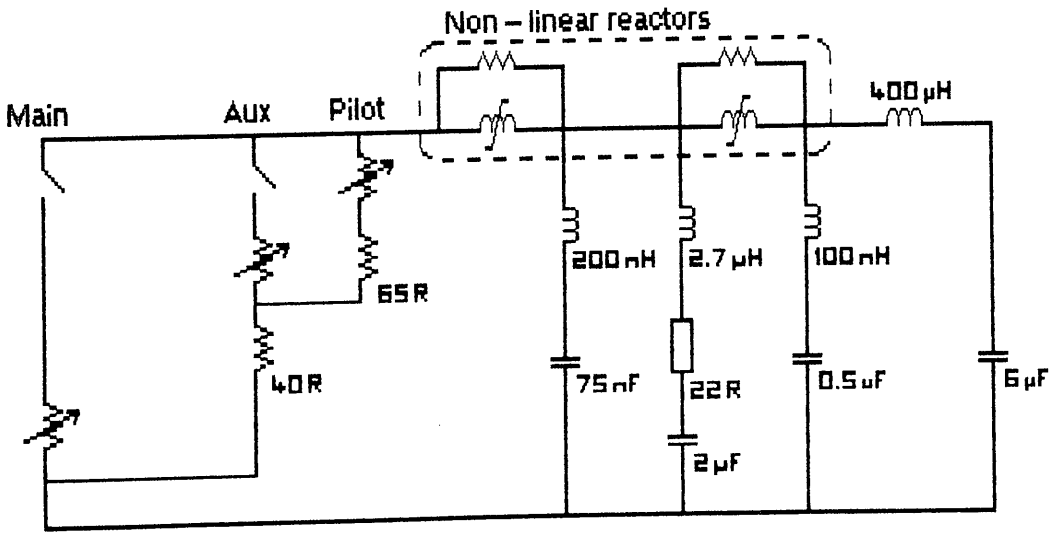

Figure 6.3.11: Double amplifying gate, 65R and 40R control resistors.

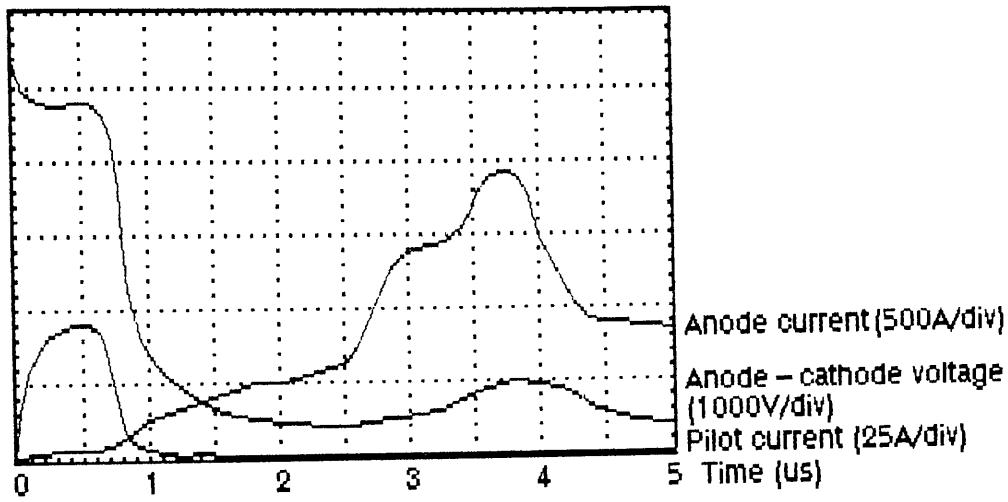

Figure 6.3.12: Turn-on waveforms for double amplifying gate, 65R and 40R.

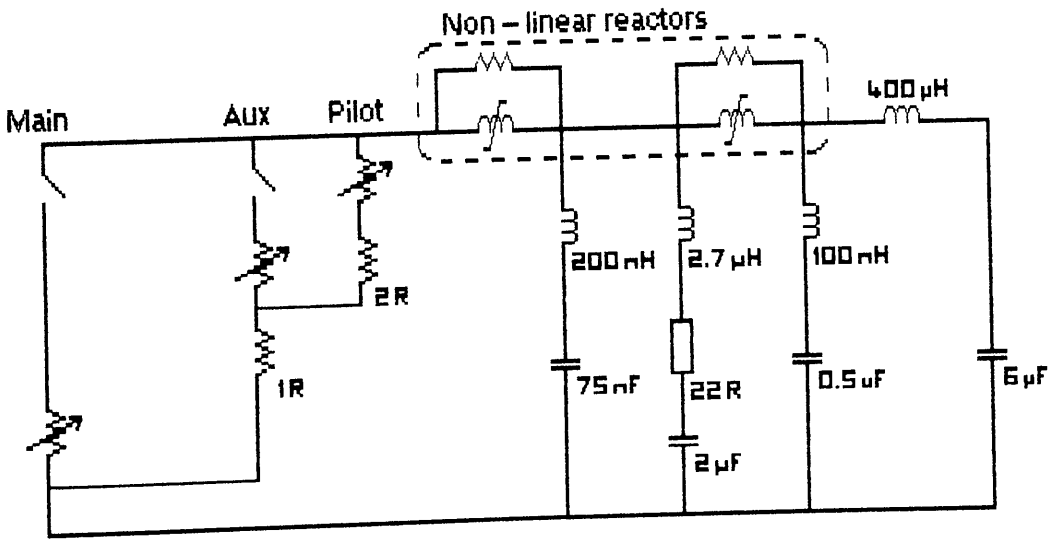

Figure 6.3.13: Double amplifying gate, 2R and 1R control resistors.

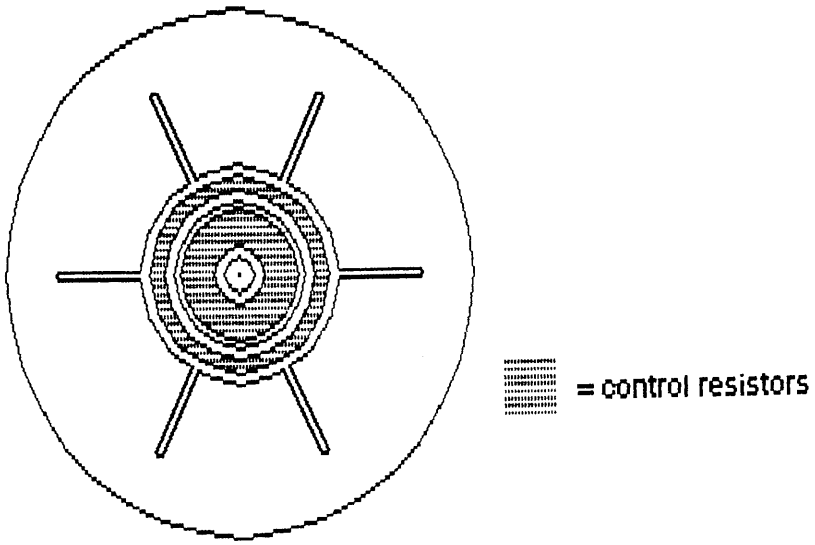

Figure 6.3.14: Overview of the 56mm device.

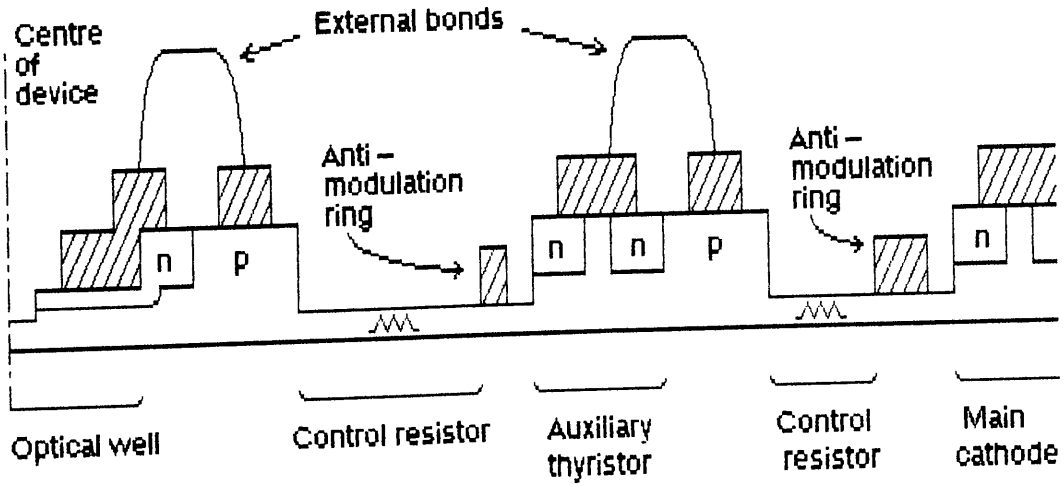

Figure 6.3.15: Cross-section through the 56mm device.

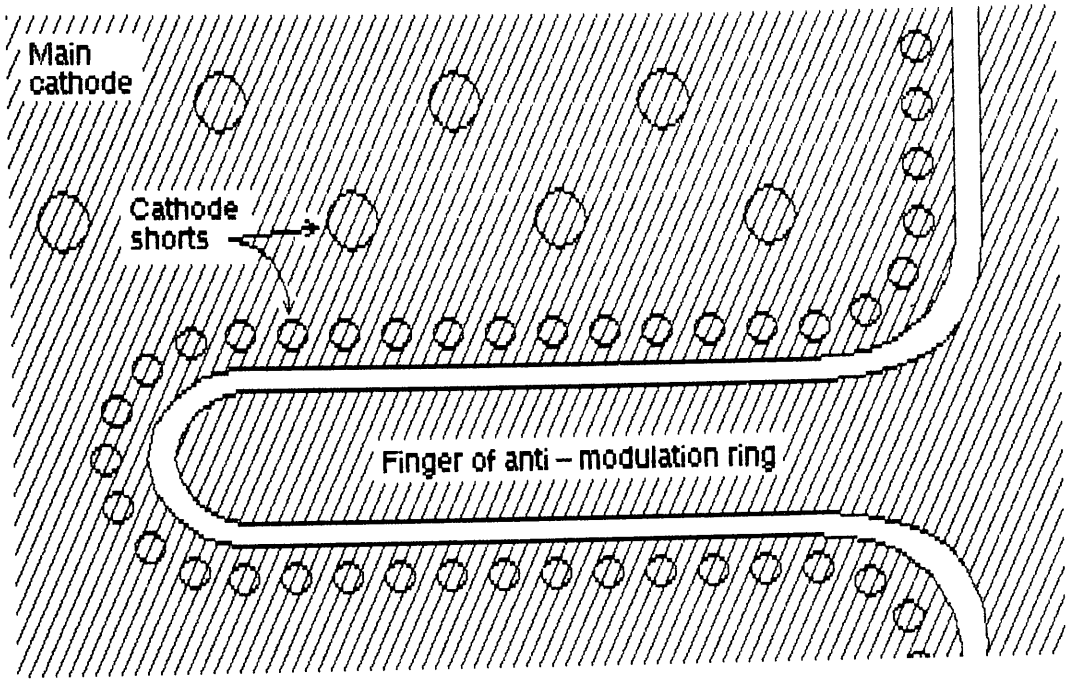

Figure 6.3.16: Design of the shorting pattern at the corners of the inside edge of the main cathode.

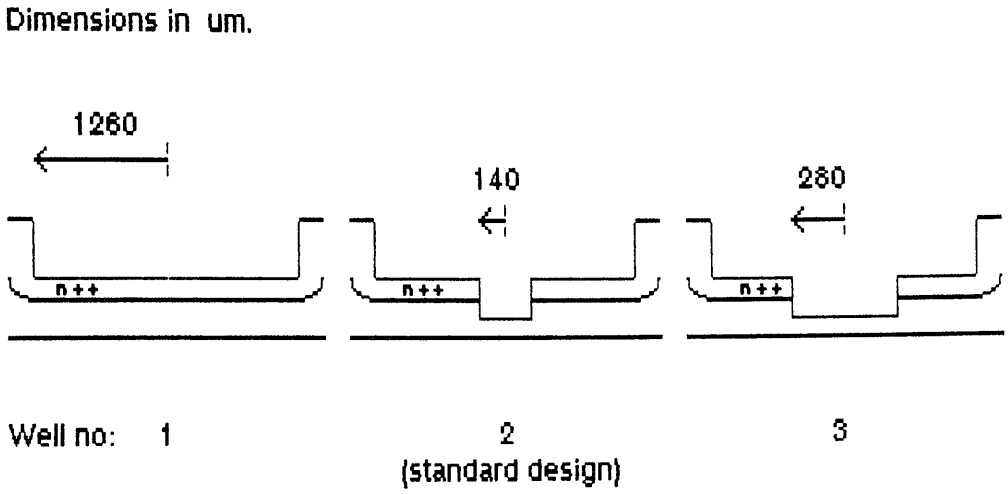

Figure 6.3.17: Design of the #1 set of optical test structures, showing the variation in diameter of the secondary wells. (Not to scale)

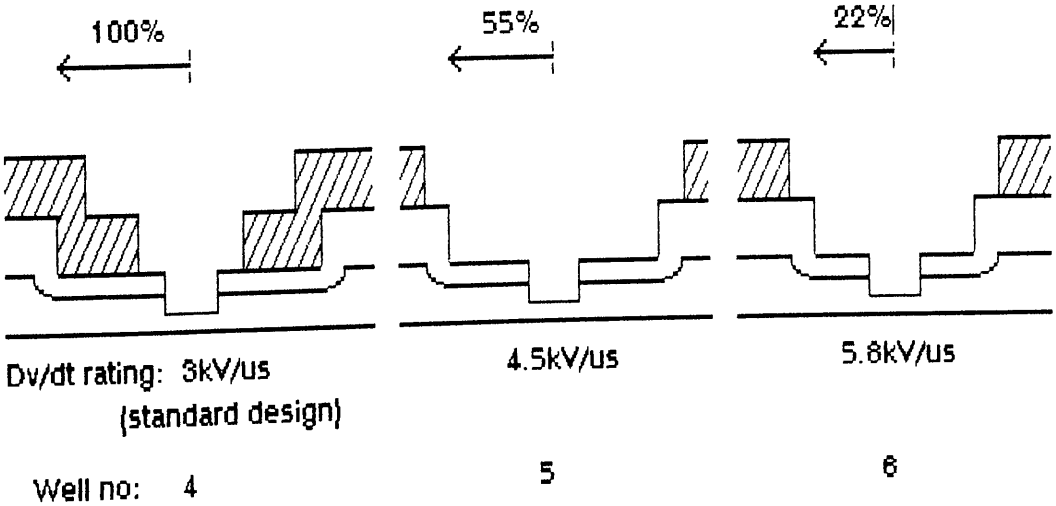

Figure 6.3.18: Design of the #2 set of optical test structures, showing the variation in diameter of the primary wells. (Not to scale)

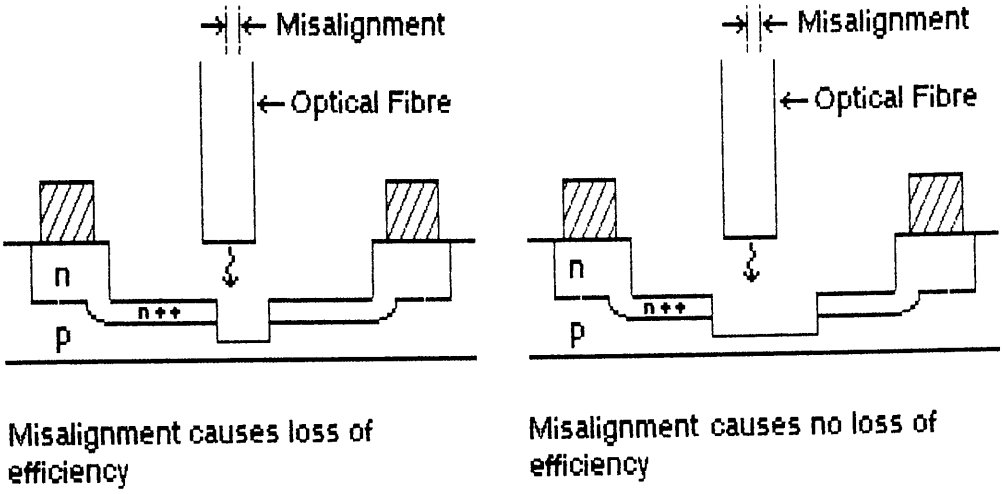

Figure 6.3.19: Effect of fibre misalignment on optical sensitivity achieved in practice.

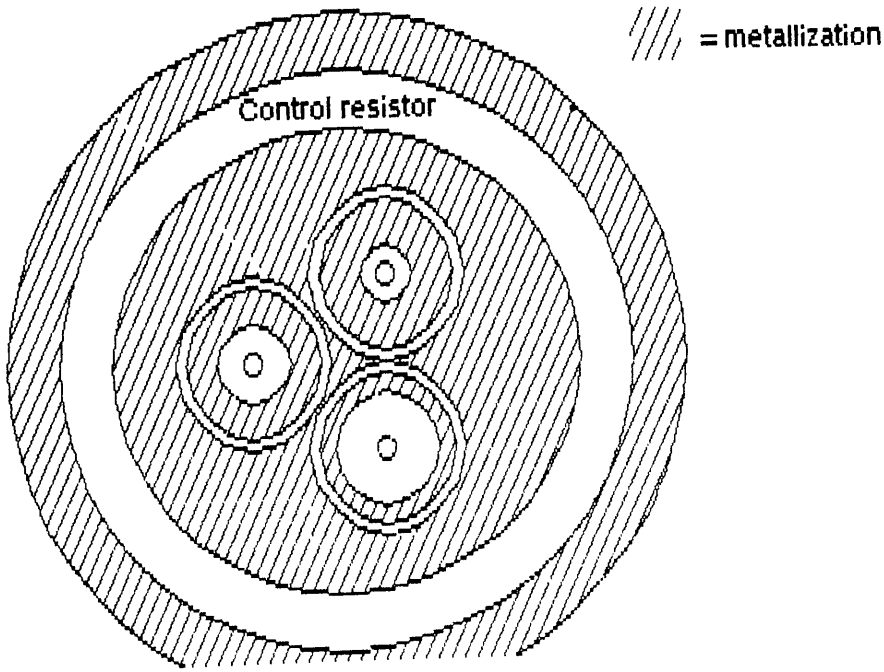

Figure 6.3.20: Outline of the 30mm optical device design.

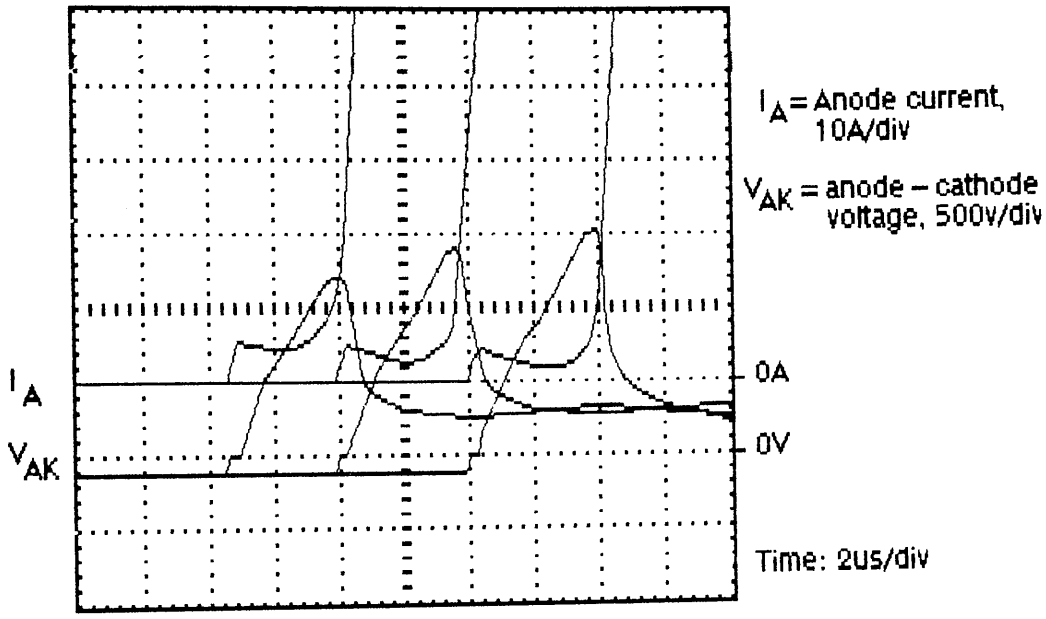

Figure 6.4.1: Typical forward recovery failure waveforms.

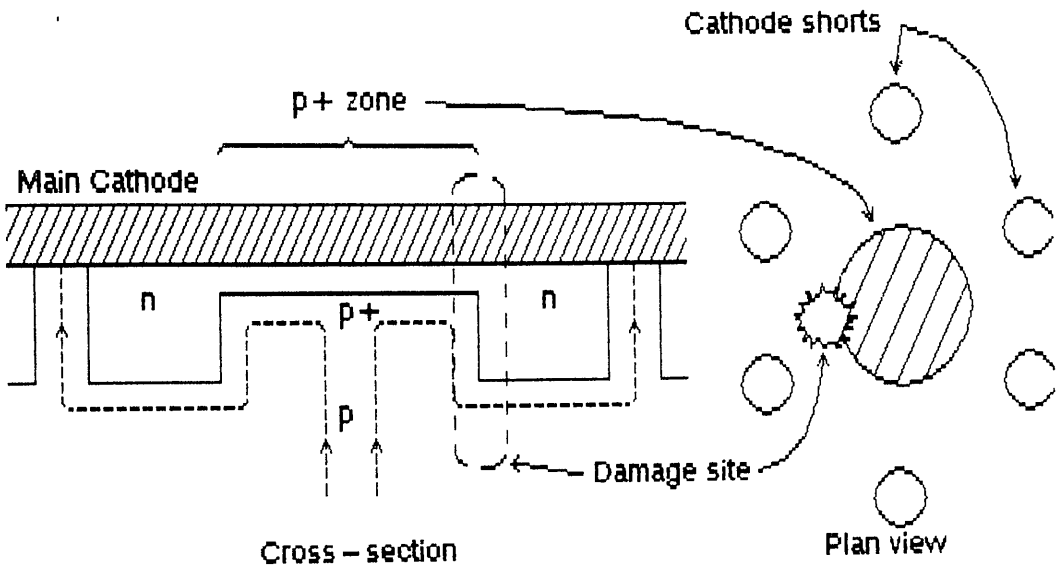

Figure 6.4.2: Failure site for the p+ zone designs.

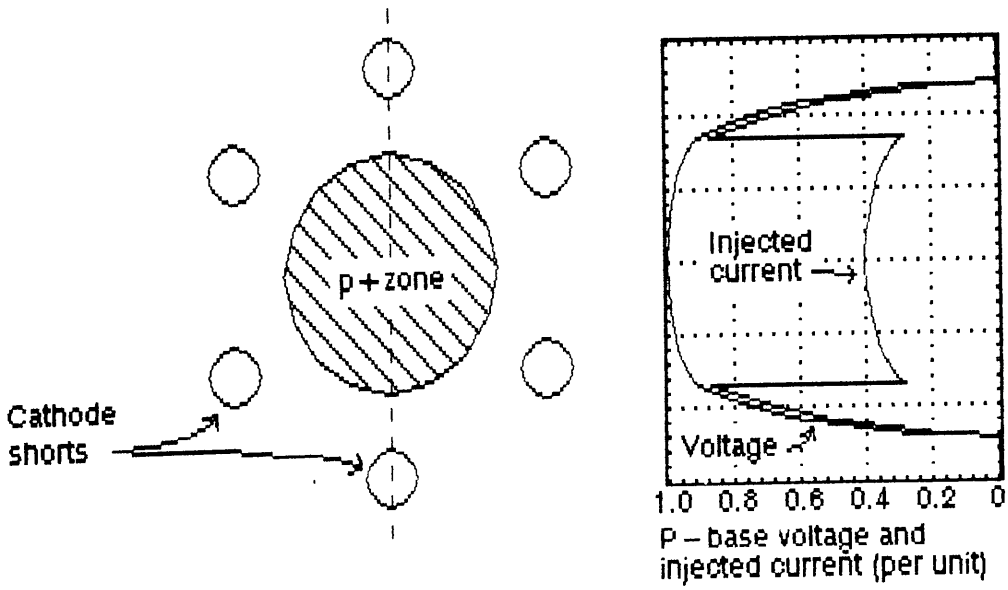

Figure 6.4.3: P-base potential and current injection during a forward recovery ramp, with a higher built-in potential in the p+ zone.

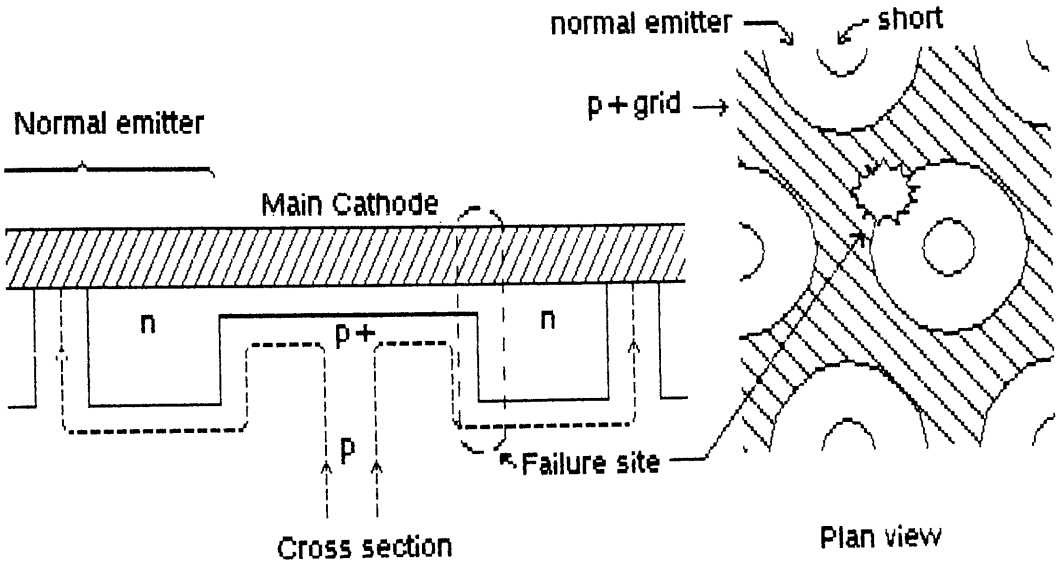

Figure 6.4.4: Failure site for the "p+ grid" devices.

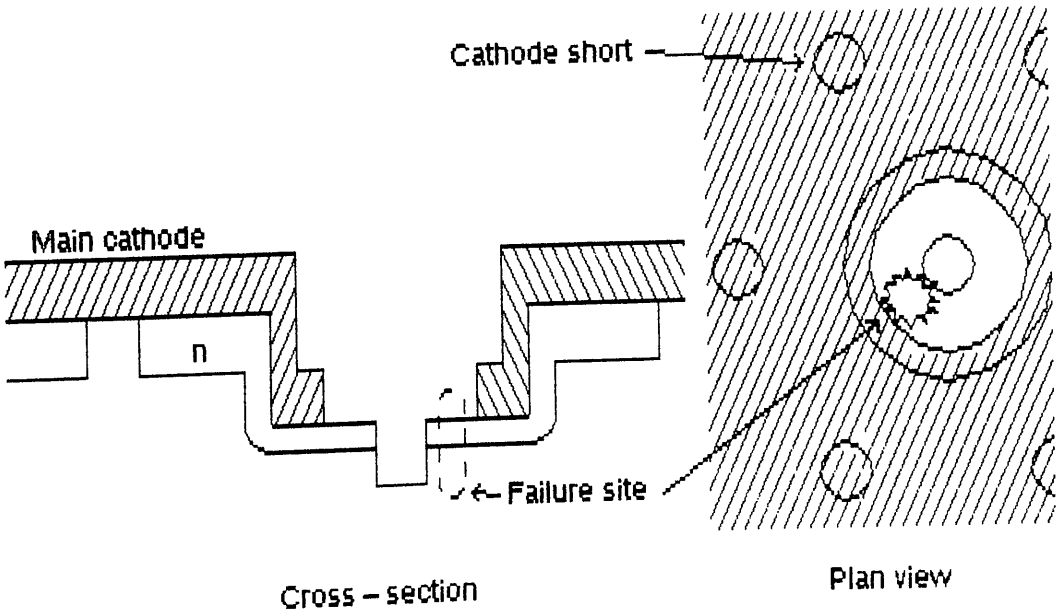

Figure 6.4.5: Failure site for the lateral field devices.

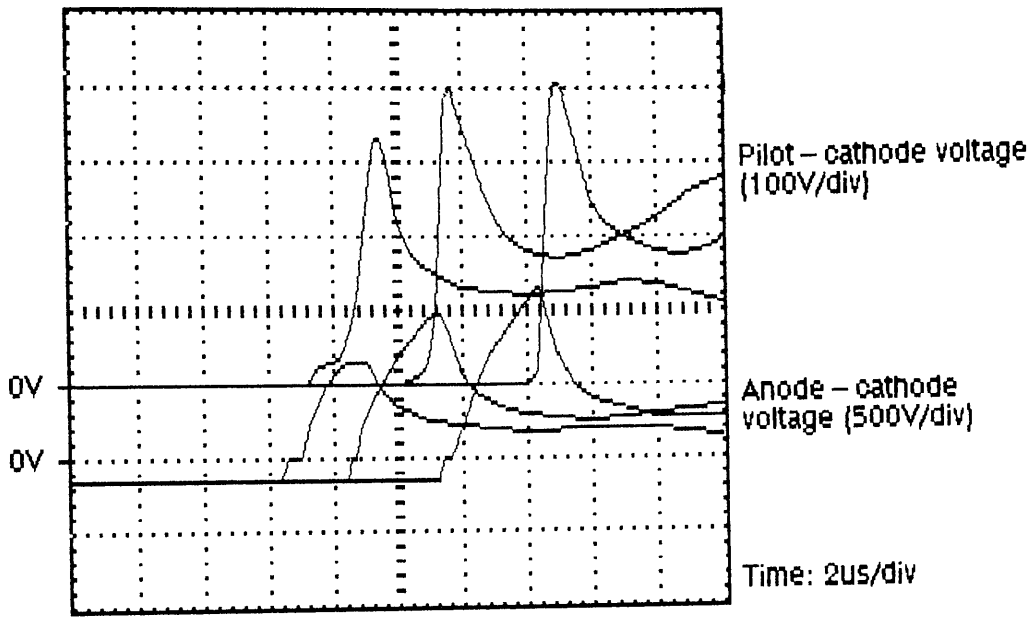

Figure 6.4.6: Forward recovery failure waveforms for the Selective Failure Zone device.

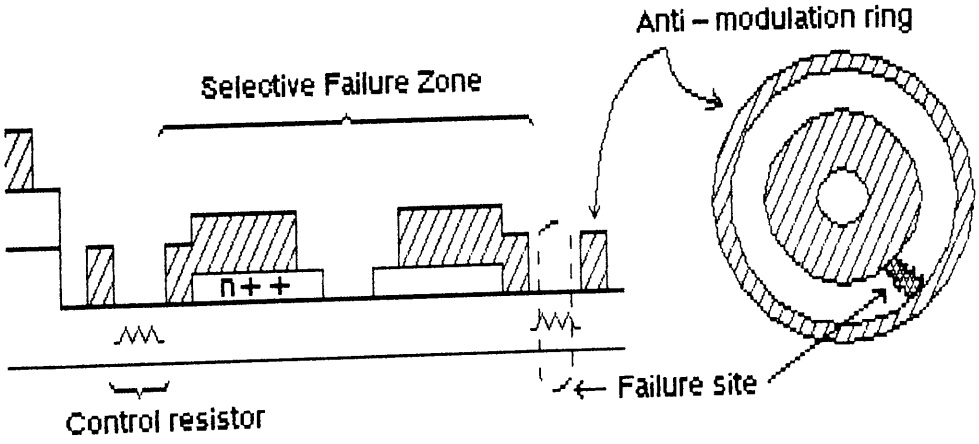

Figure 6.4.7: Failure track across the control resistor on the Selective Failure Zone device.

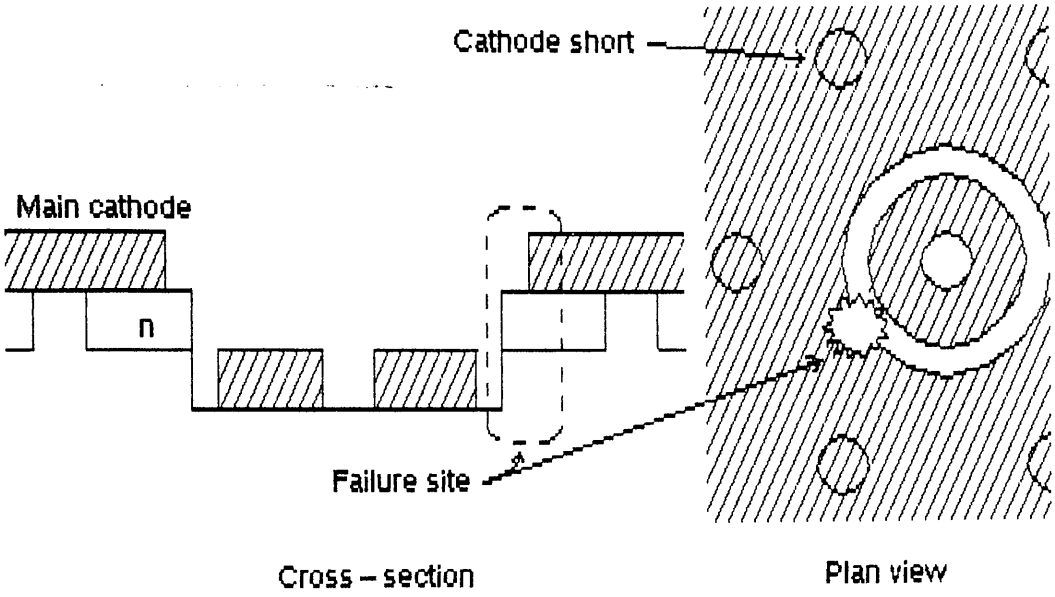

Figure 6.4.8: Failure site on the Gated Turn-on devices.

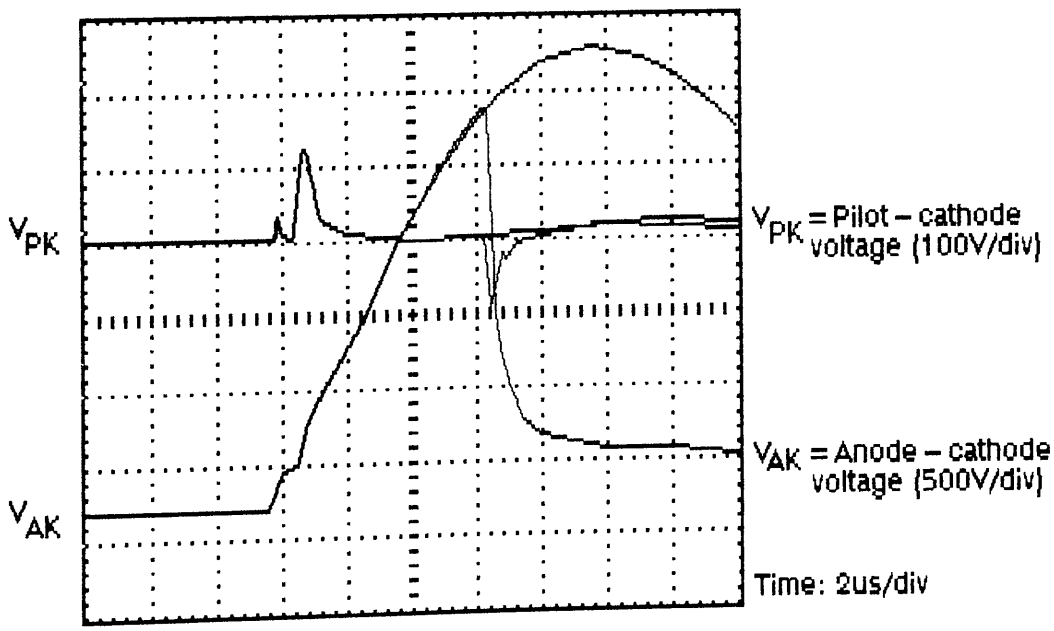

Figure 6.4.9: Waveforms for successful and unsuccessful forward recovery for the 56mm device.

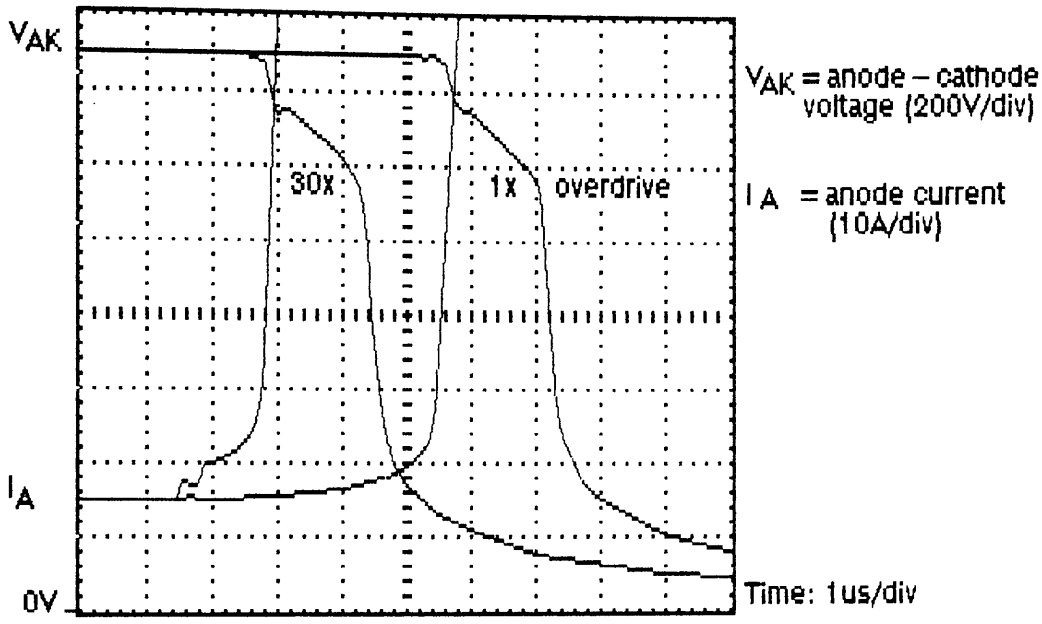

Figure 6.5.1: 1.5kV turn-on waveforms for the #0 56mm device.

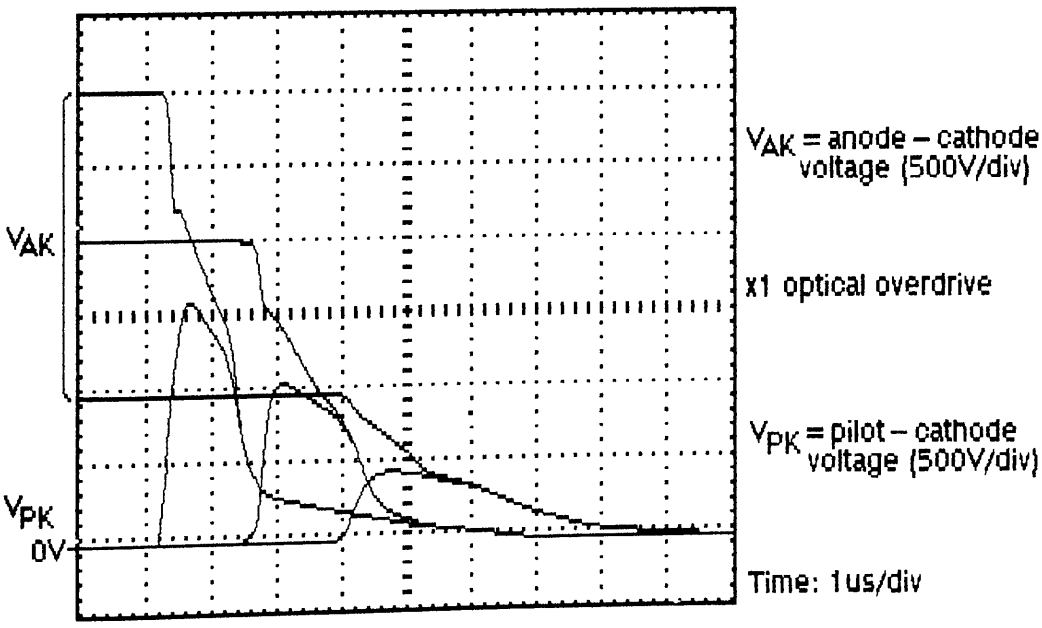

Figure 6.5.2: Turn-on waveforms for device #43 at various voltages.

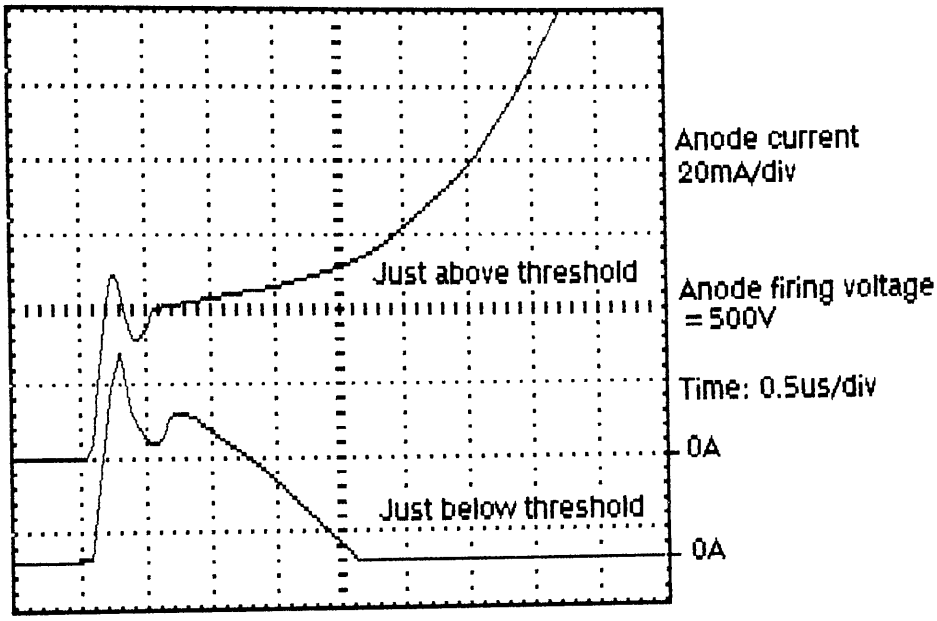

Figure 6.5.3: Anode current for optical triggering just below and just above threshold, device #43.

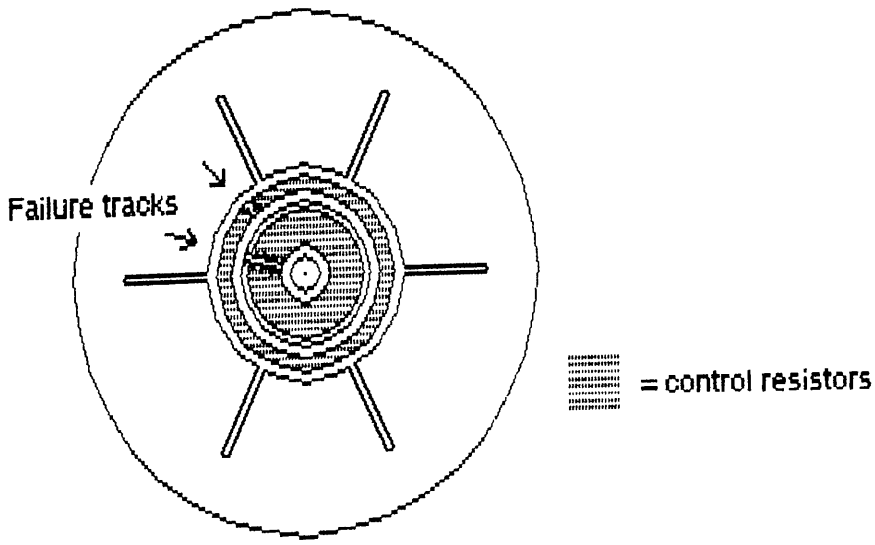

Figure 6.5.4: Failure tracks across the control resistors on device #43.

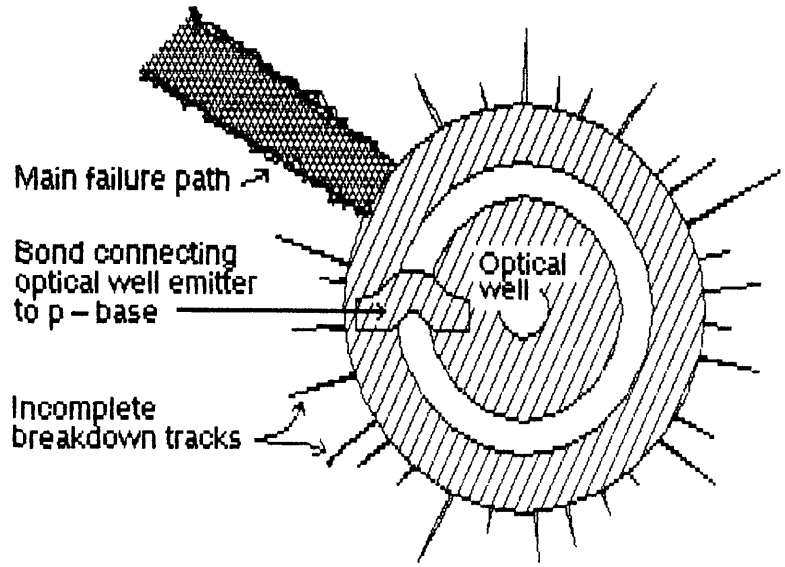

igure 6.5.5: Detail of the inside edge of control resistor Rc1, showing incomplete failure tracks (device #43).

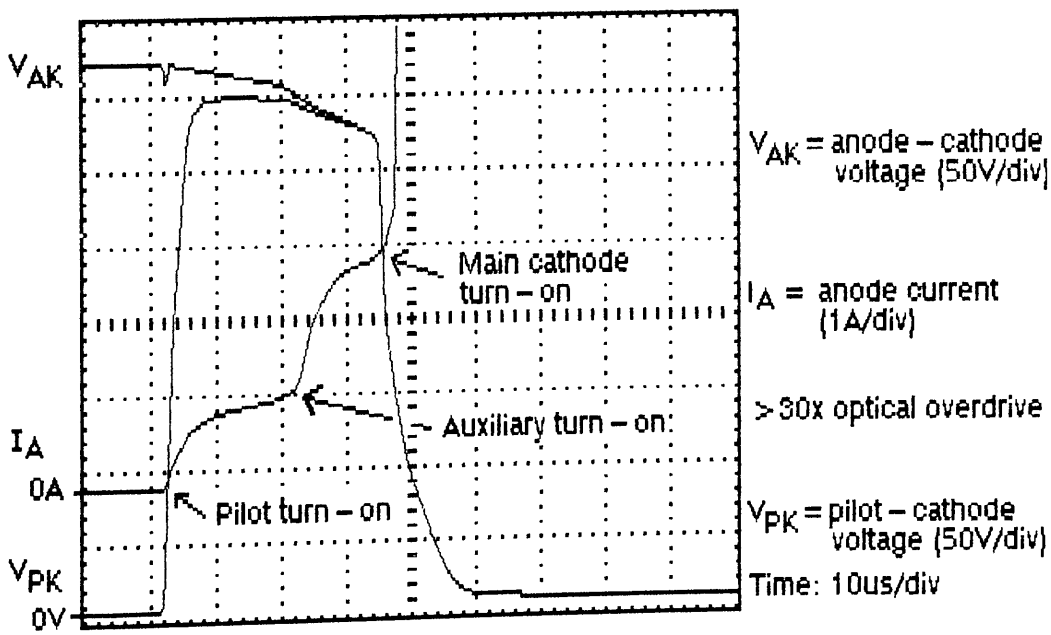

igure 6.5.6: Low-voltage turn-on waveforms for device #11.

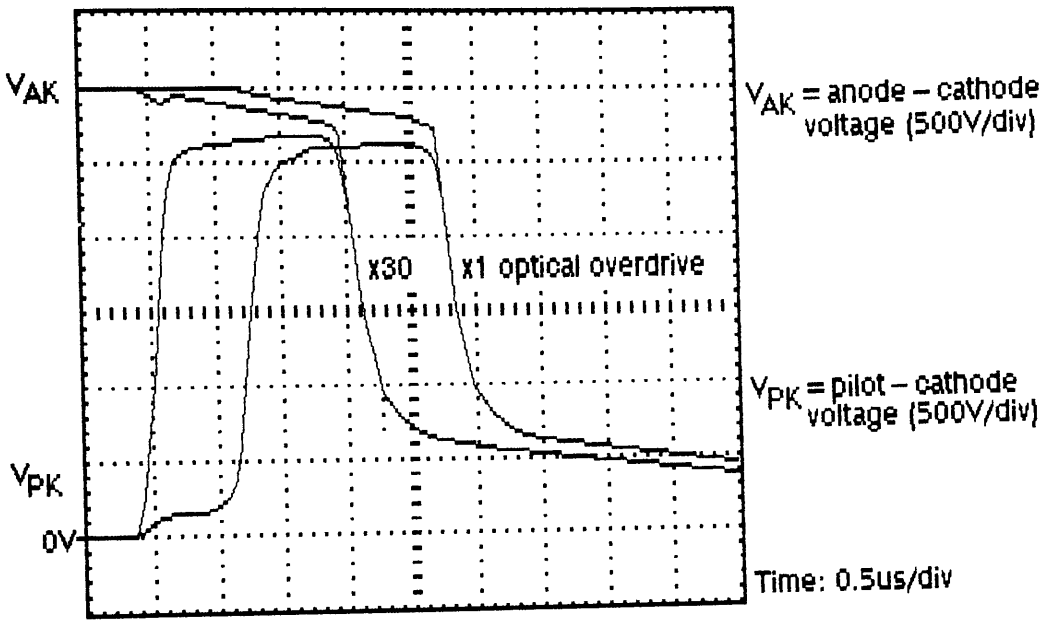

Figure 6.5.7: 3kV turn-on waveforms for device #11.

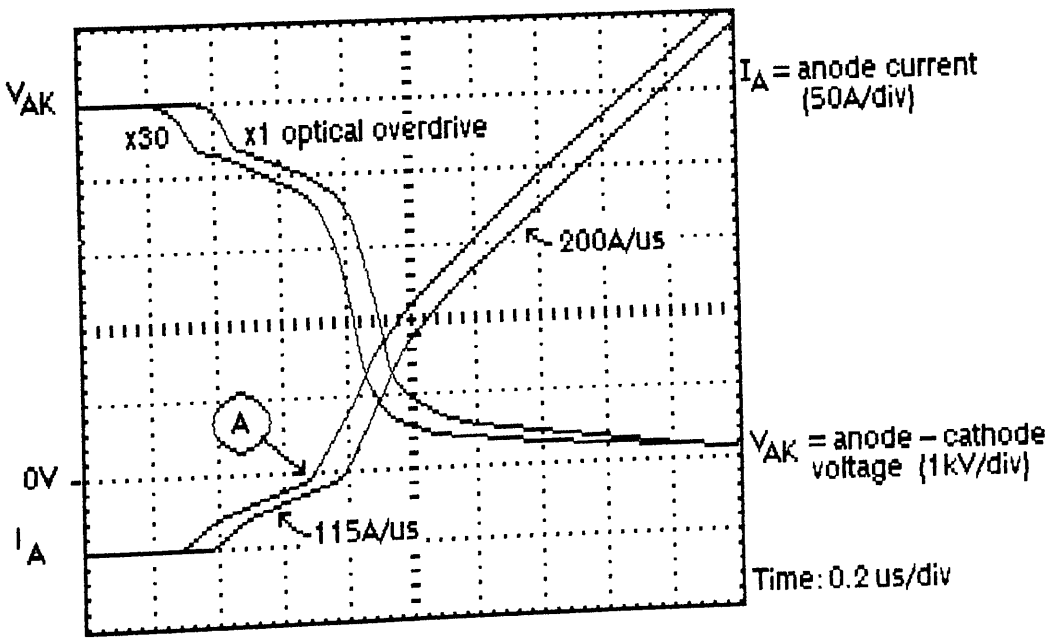

Figure 6.5.8: 5kV turn-on waveforms for device #46, 30°C.

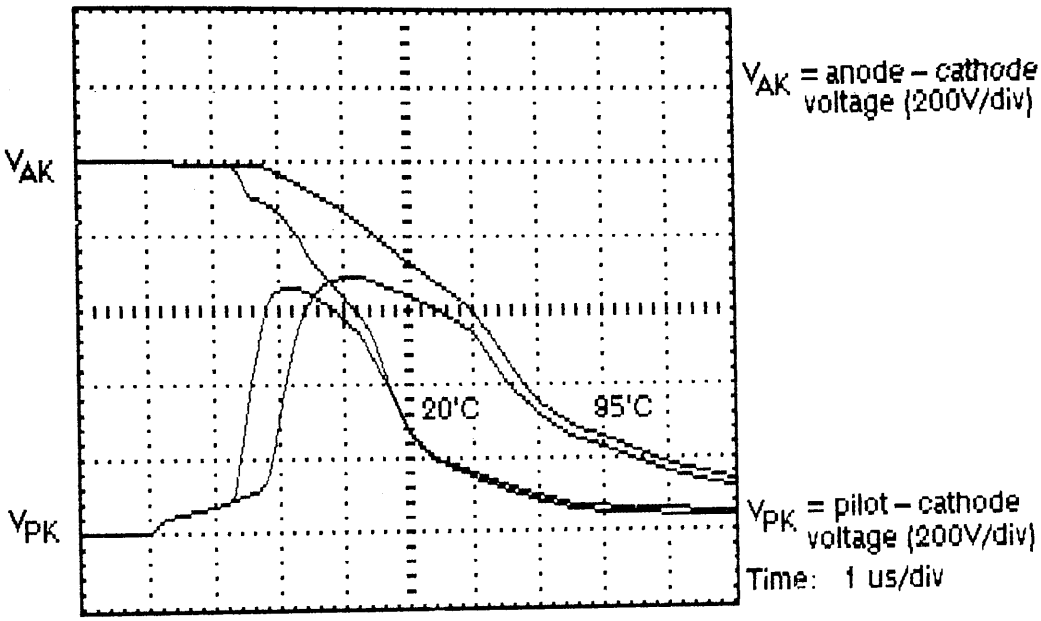

Figure 6.5.9: 1kV turn-on waveforms for device #46 at 20'C and 95'C.

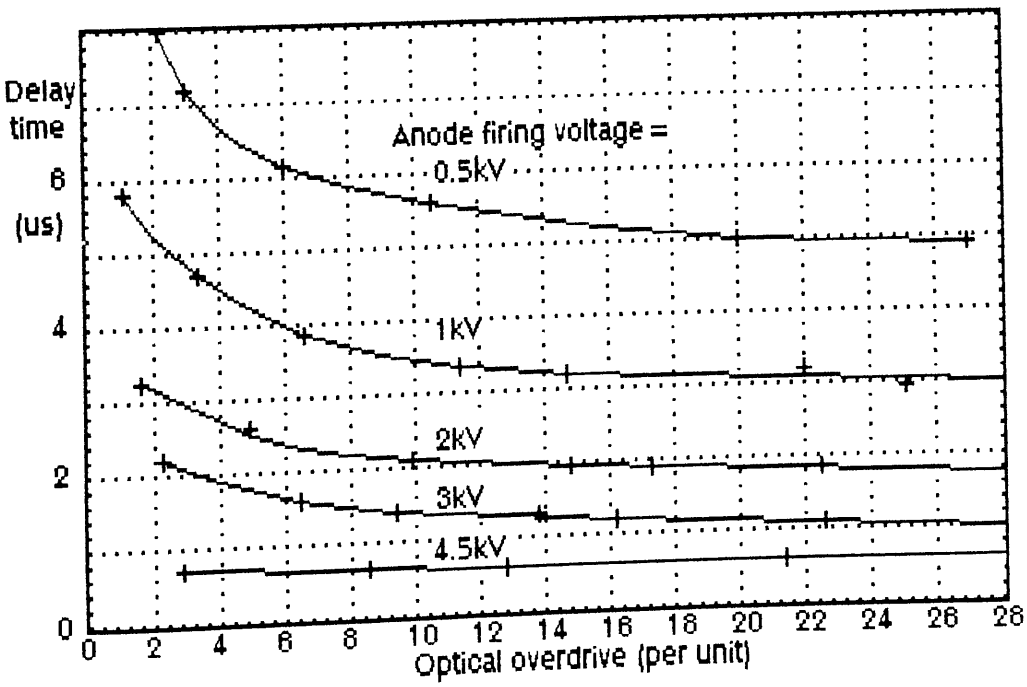

Figure 6.5.10: Turn-on delay time for device #46, 95'C to 105'C.

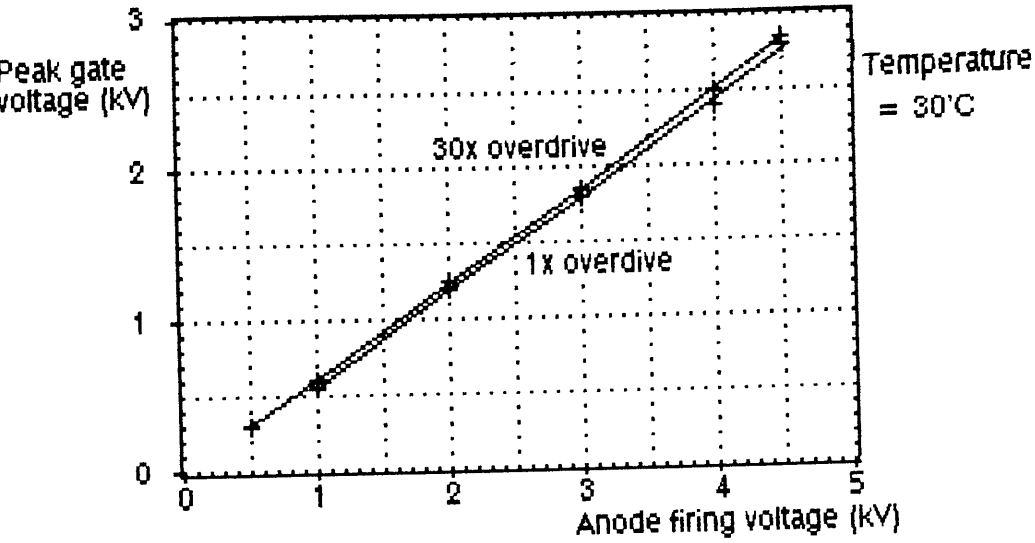

e 6.5.11: Peak gate voltage as a function of anode turn-on voltage and optical overdrive, device #45.

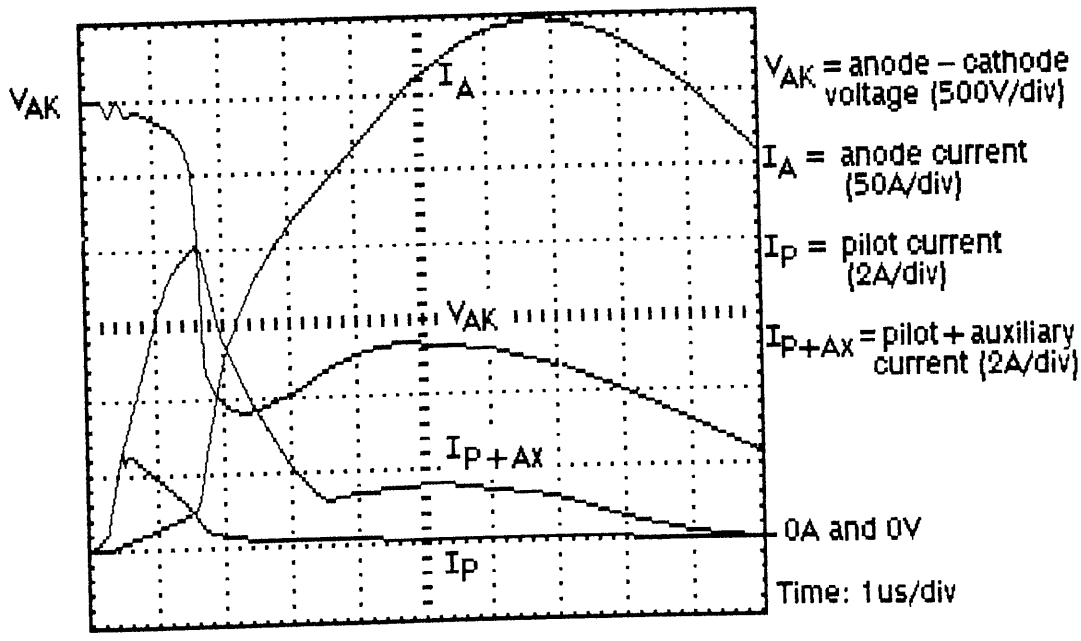

ure 6.5.12: Turn-on waveforms for Device #49, 3kV, room temperature.

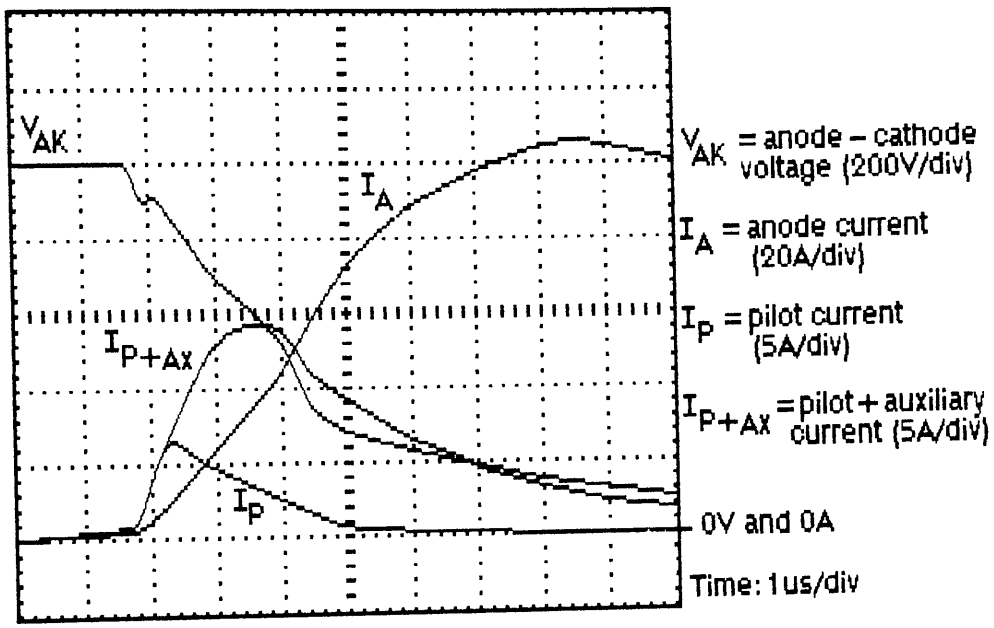

gure 6.5.13: Turn-on waveforms for Device #25, 1kV, room temperature.

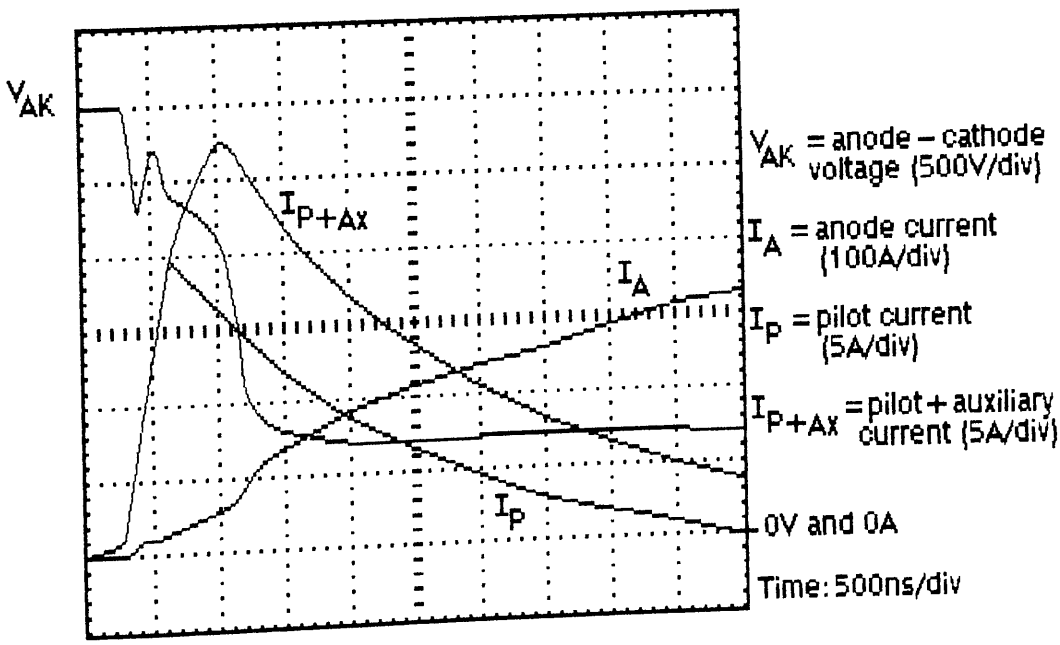

igure 6.5.14: Turn-on waveforms for Device #25, 3kV, room temperature.

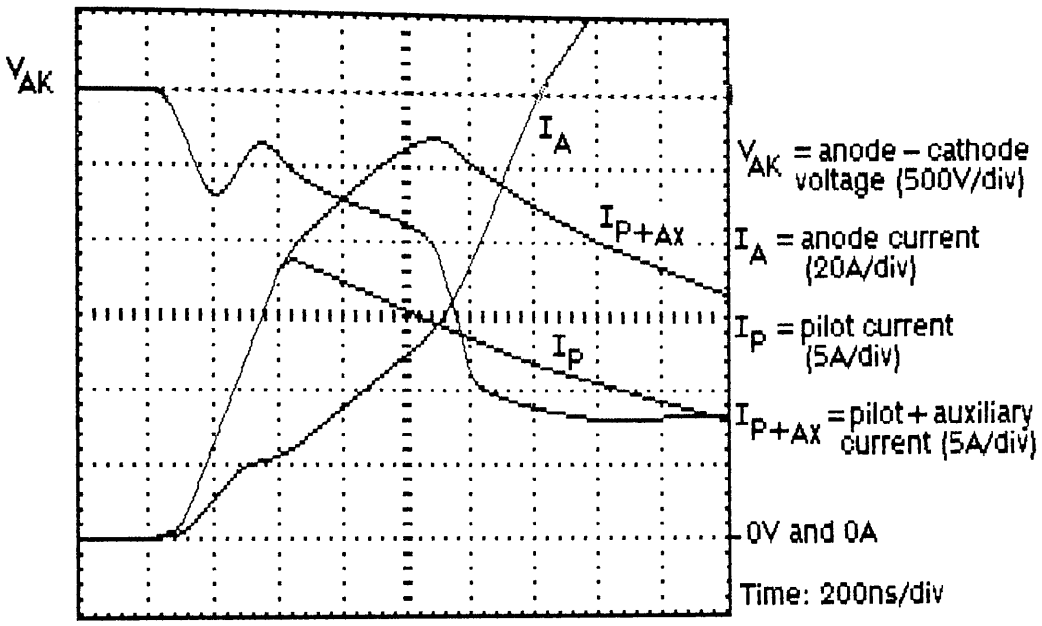

Figure 6.5.15: Turn-on waveforms for Device #25, 3kV, room temperature.

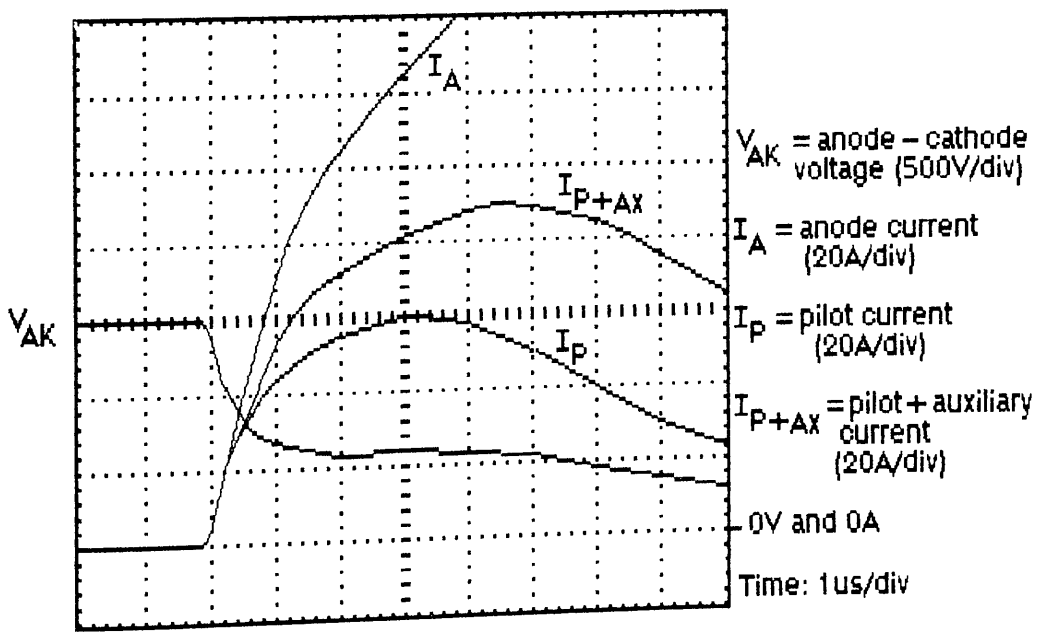

Figure 6.5.16: Turn-on waveforms for Device #4, 1.5kV, room temperature.

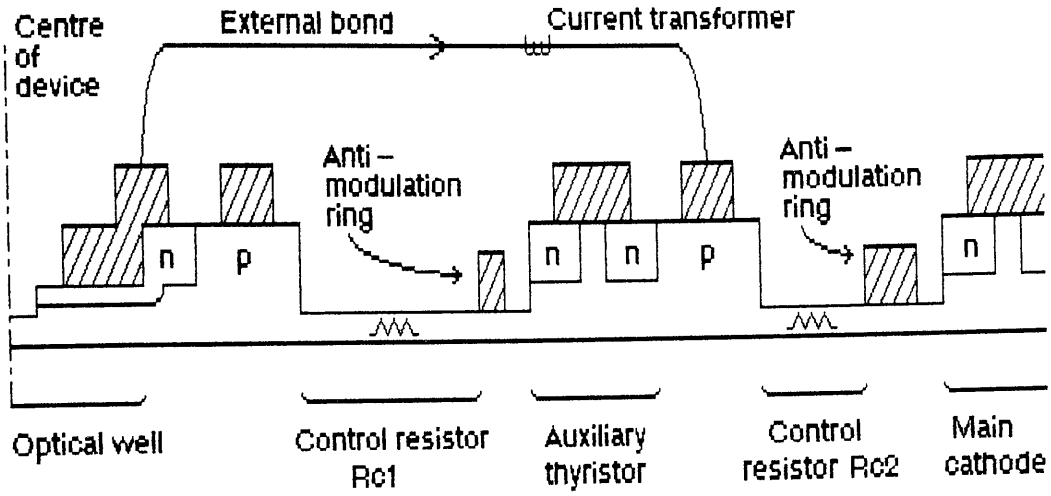

Figure 6.5.17: Arrangement for by-passing the auxiliary amplifying gate.

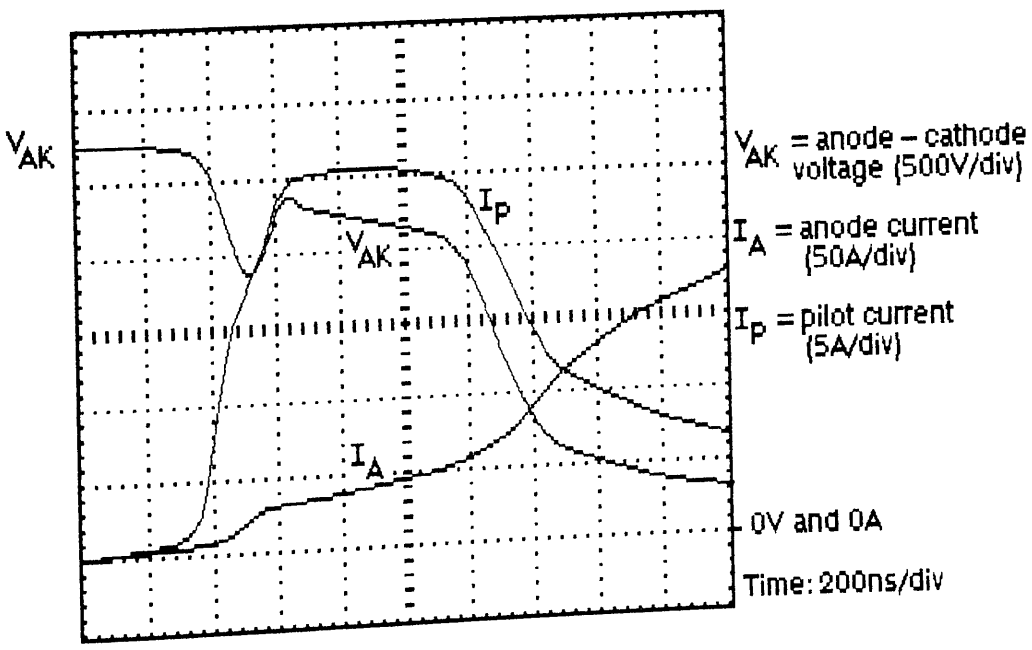

Figure 6.5.18: Turn-on waveforms for Device #35, 3.25kV, 90'C.

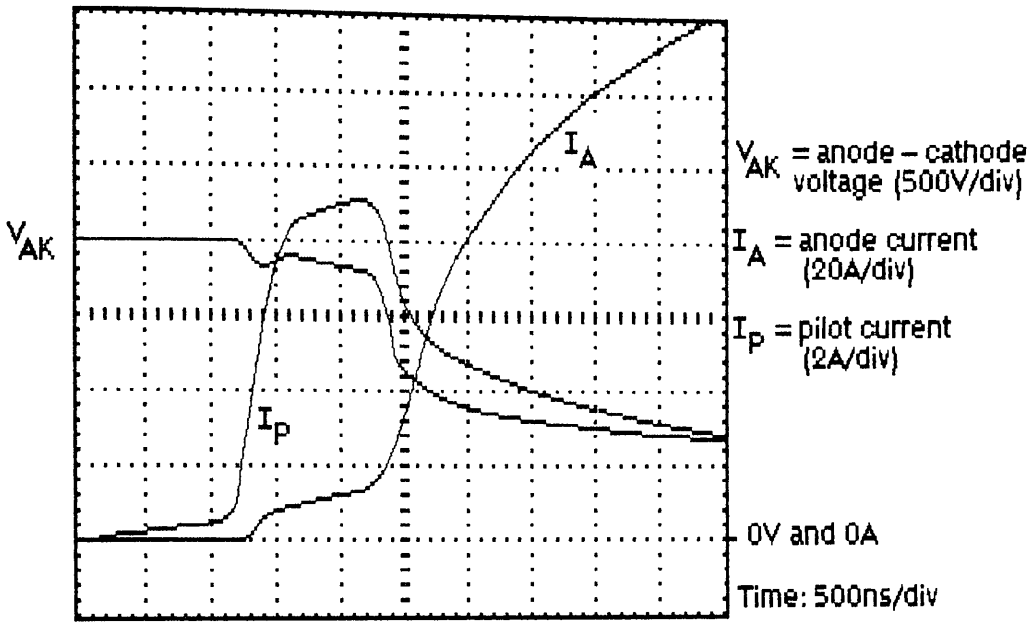

Figure 6.5.19: Turn-on waveforms for Device #15, 2kV, 90'C.

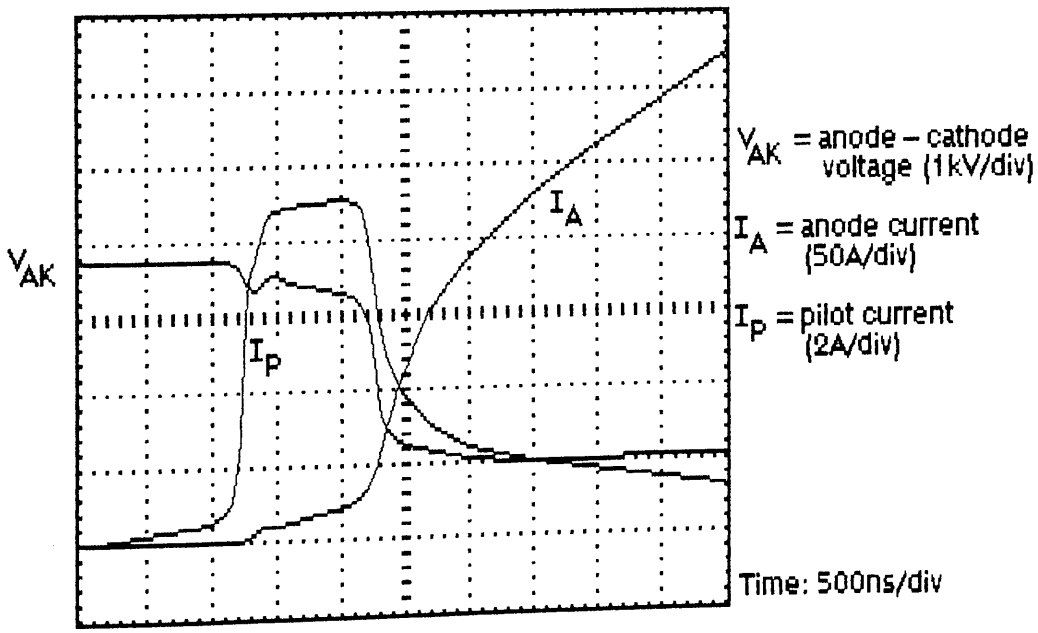

Figure 6.5.20: Turn-on waveforms for Device #49, 3.75kV, 90'C.

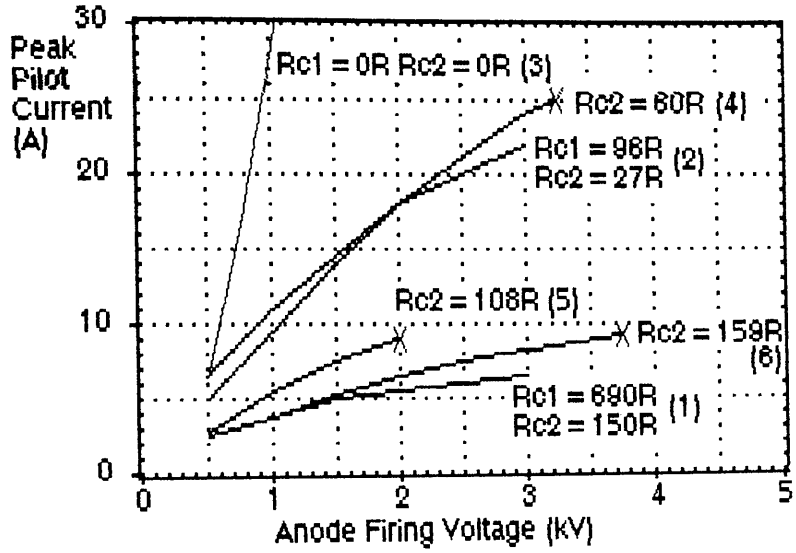

Figure 6.5.21: Peak pilot current as a function of turn-on voltage for the various devices tested. (Bracketed numbers indicate the test referred to).

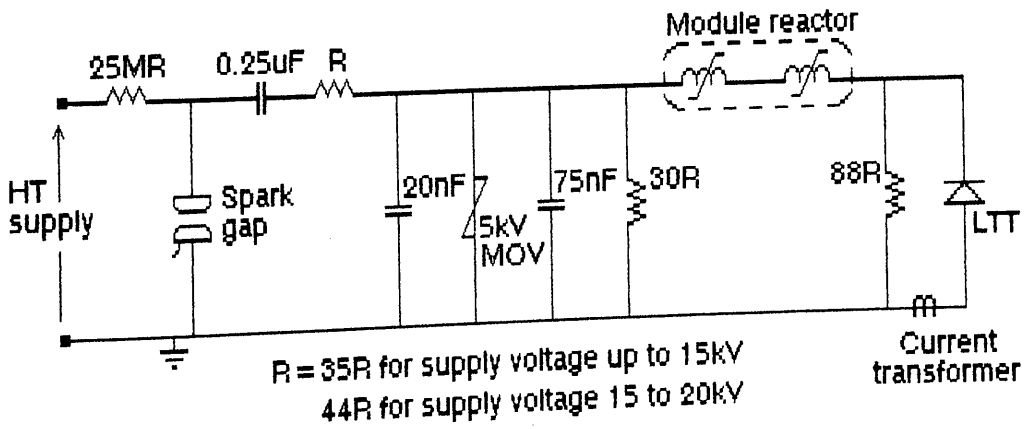

Figure 6.5.22: Dv/dt test circuit.

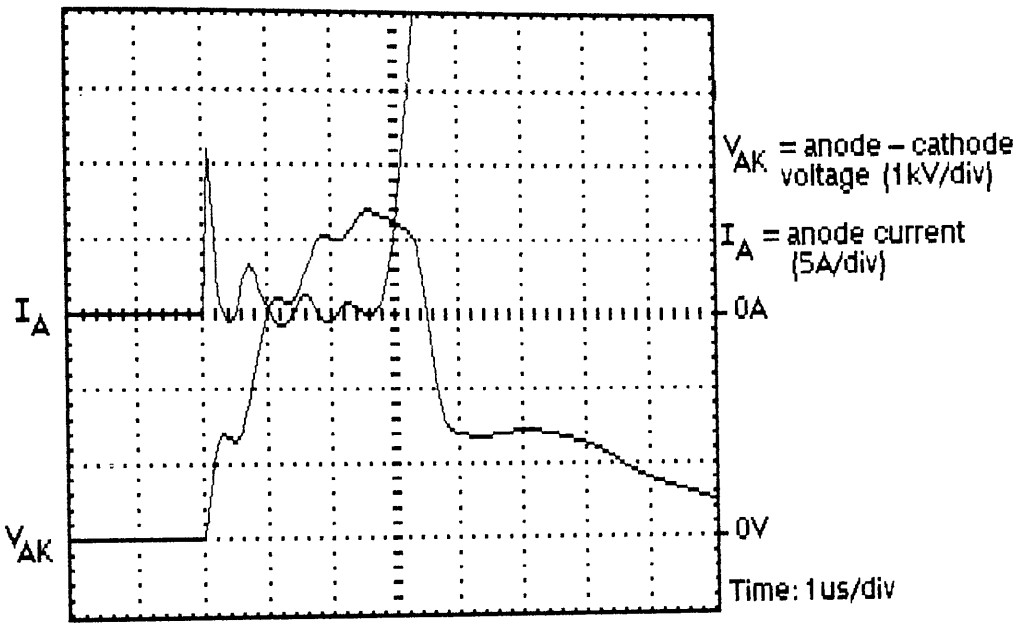

Figure 6.5.23: Dv/dt turn-on waveforms for device #21, 110'C.

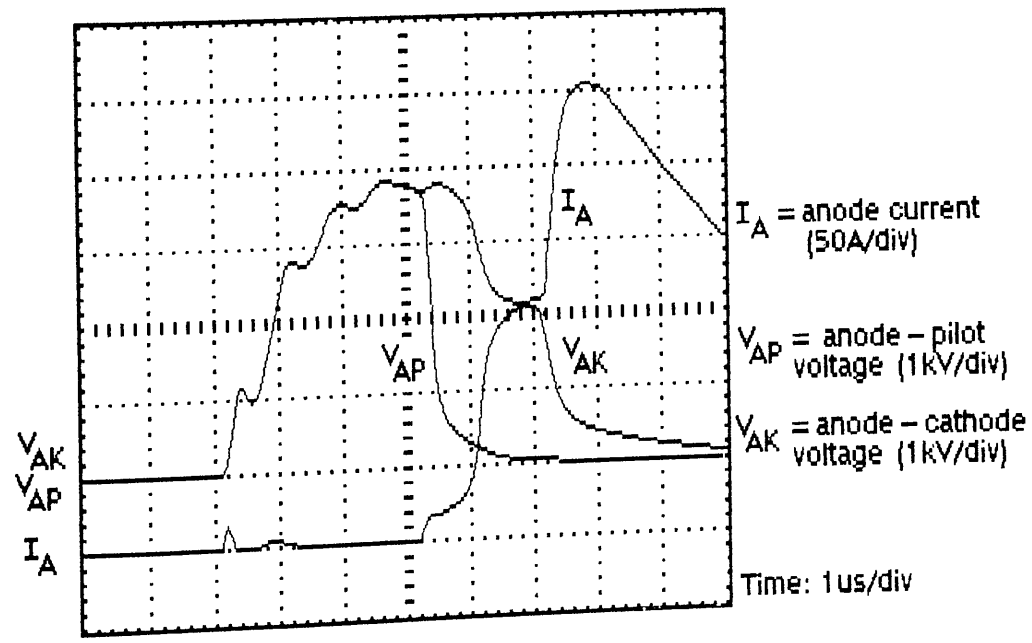

Figure 6.5.24: Marginal dv/dt turn-on waveforms for device #46, 110'C.

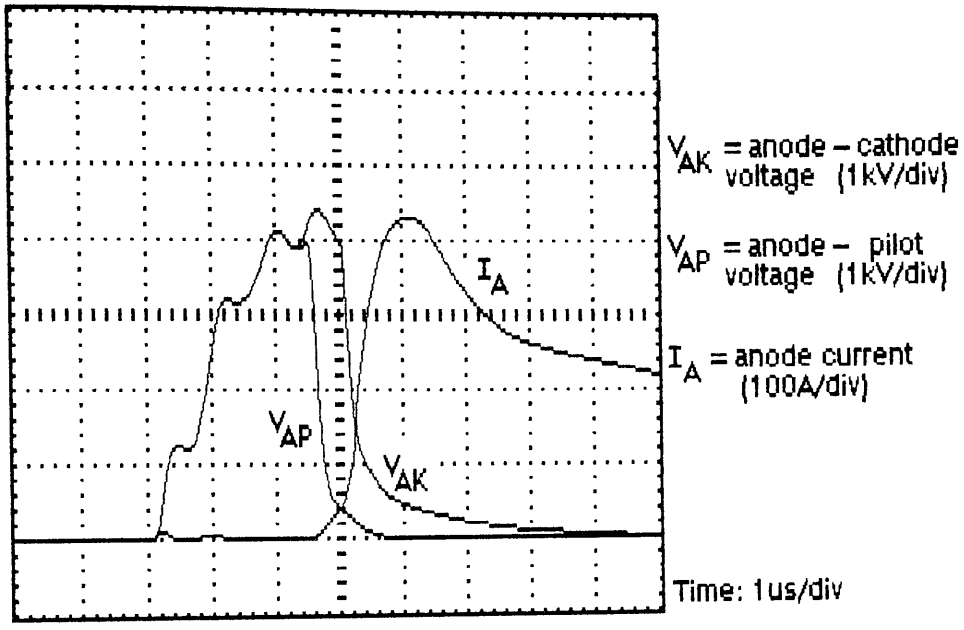

Figure 6.5.25: Strong dv/dt turn-on waveforms for device #46, 110'C.

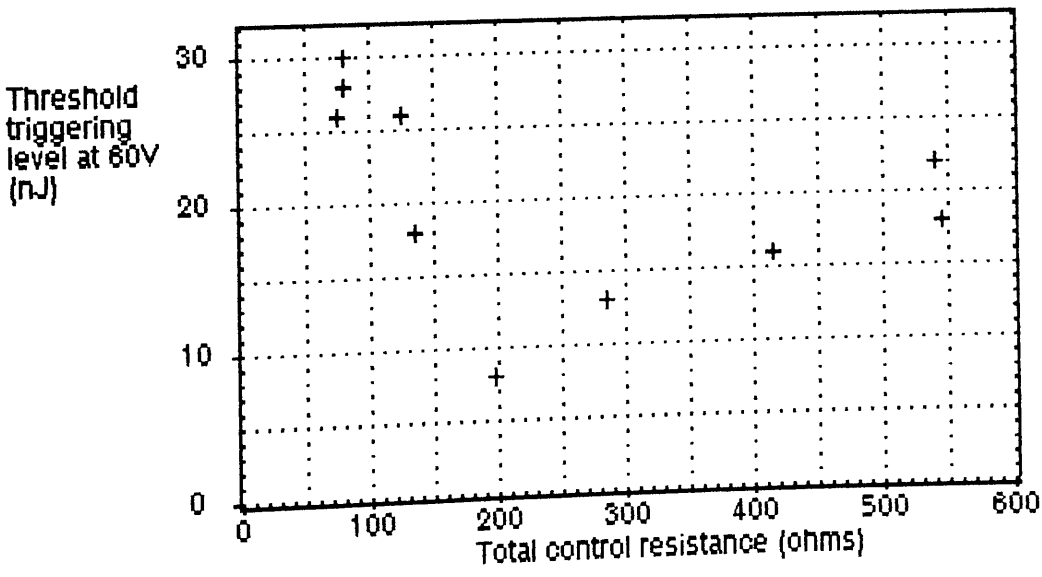

Figure 6.5.26: Optical sensitivity versus total control resistance, for various devices.

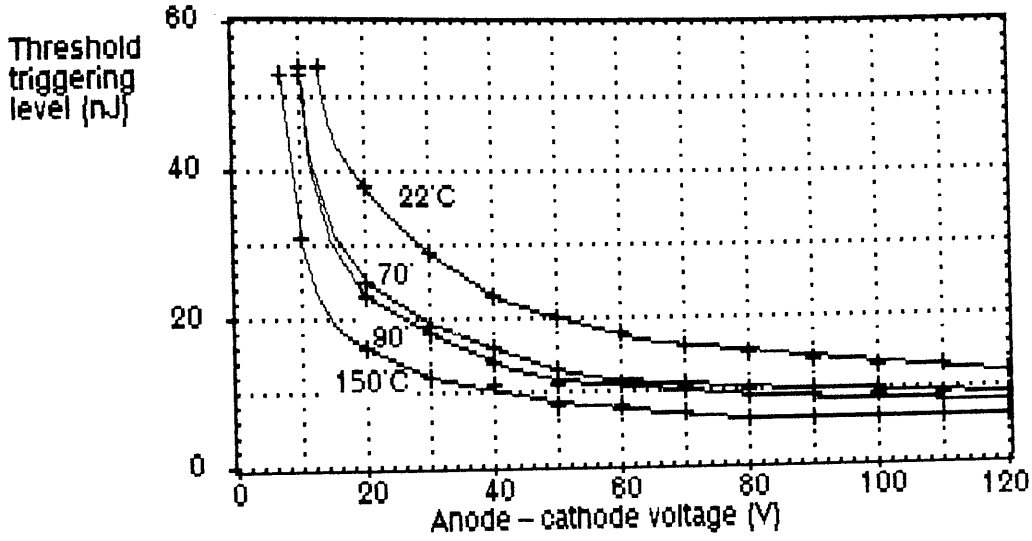

Figure 6.5.27: Optical sensitivity of device #46, as a function of voltage and temperature.

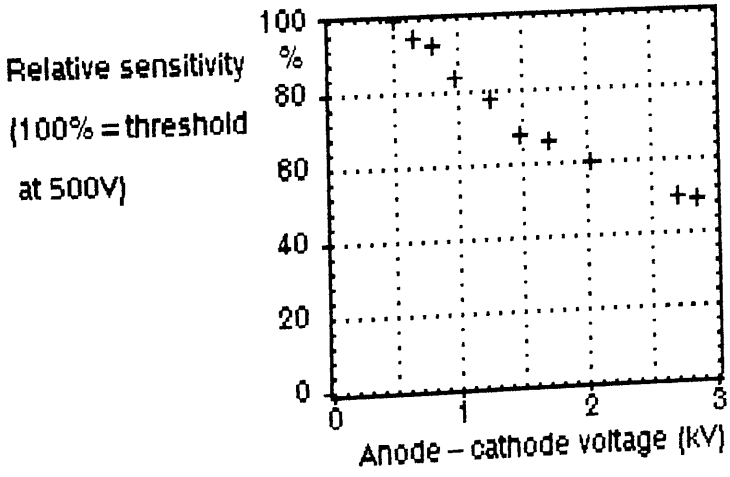

Figure 6.5.28: Relative variation of optical sensitivity with voltage for device #46, 20°C.

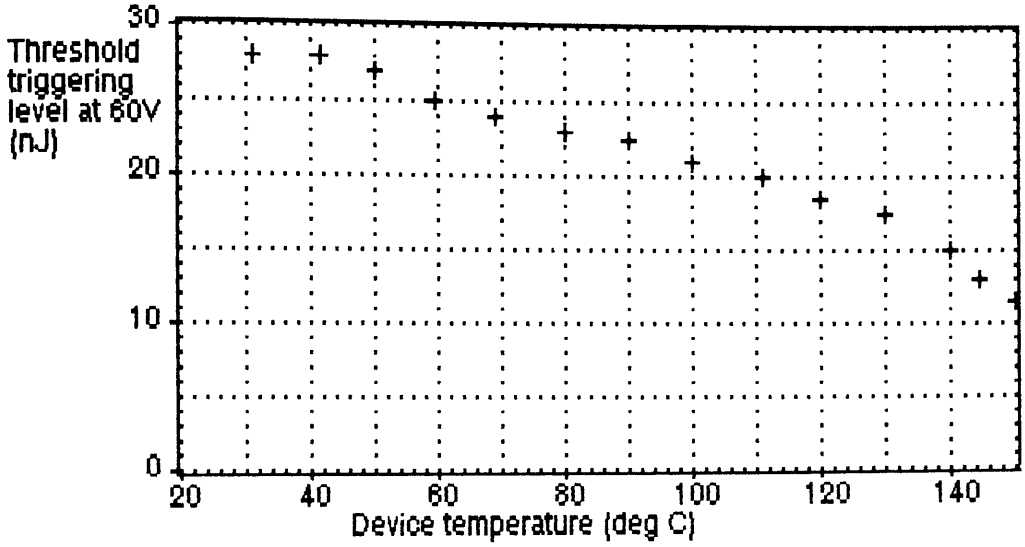

Figure 6.5.29: Optical sensitivity of the #0 56mm device as a function of temperature.

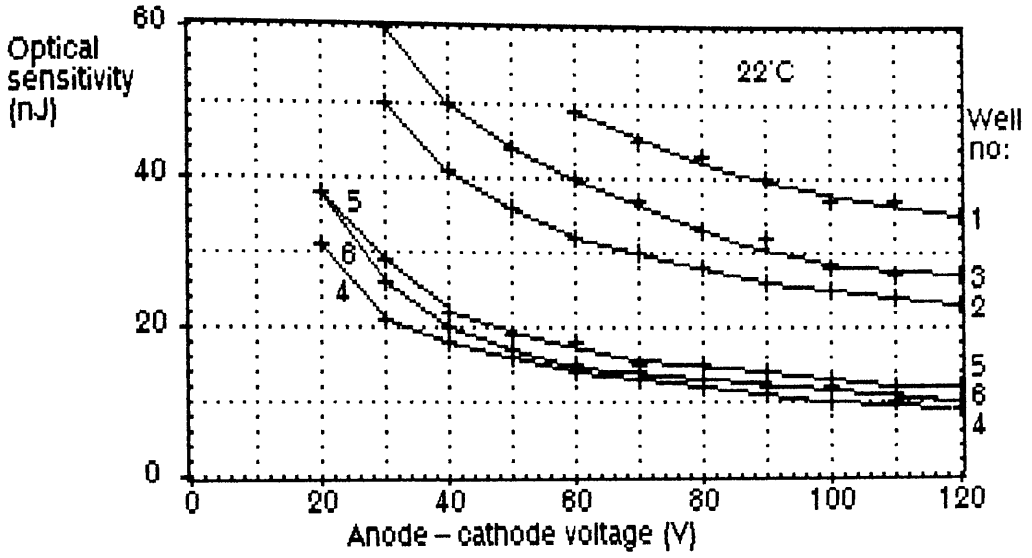

Figure 6.5.30: Optical sensitivity versus voltage for the 30mm devices from slice #43, 22'C.

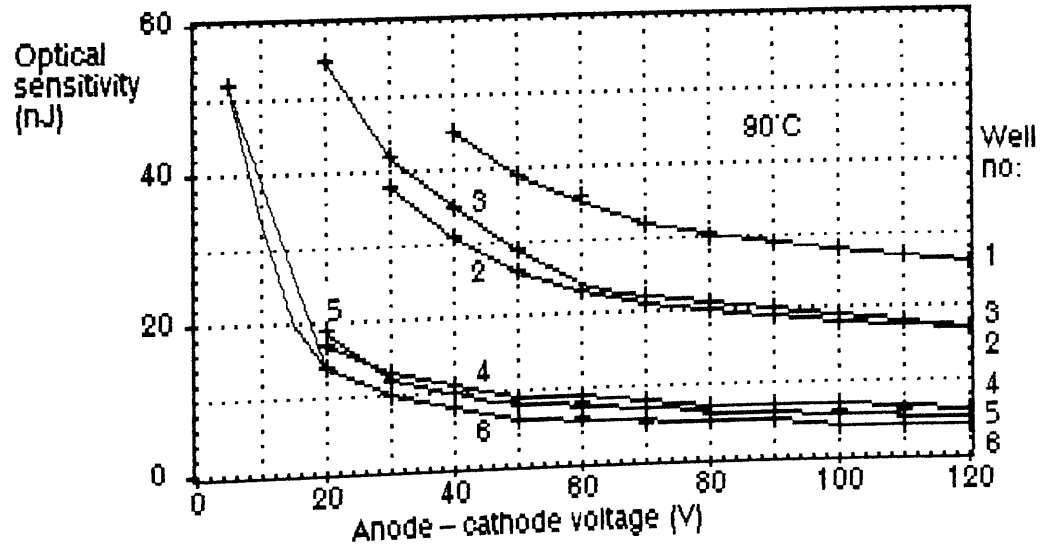

Figure 6.5.31: Optical sensitivity versus voltage for the 30mm devices from slice #43, 90'C.

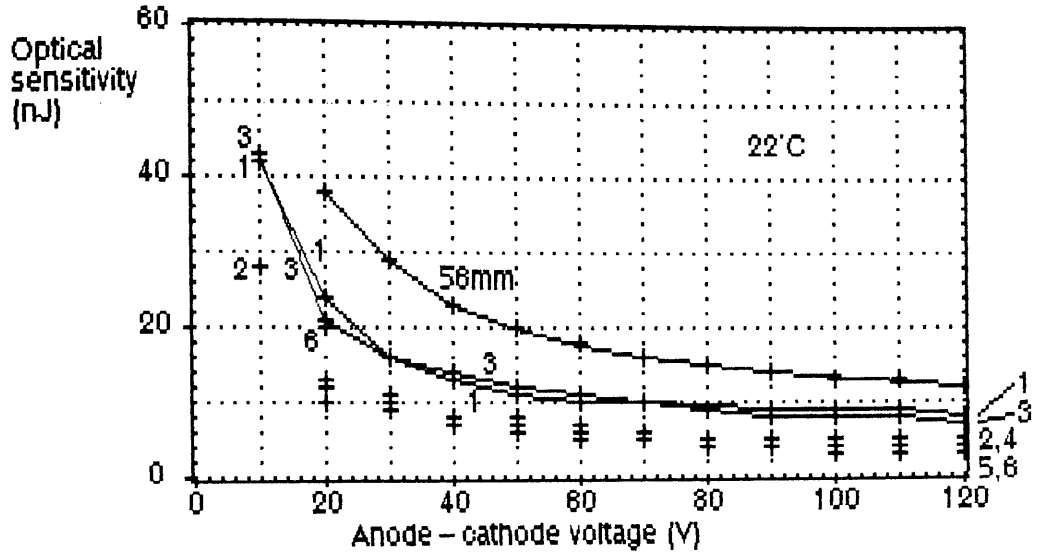

Figure 6.5.32: Optical sensitivity of the various devices from slice #46, 22'C.

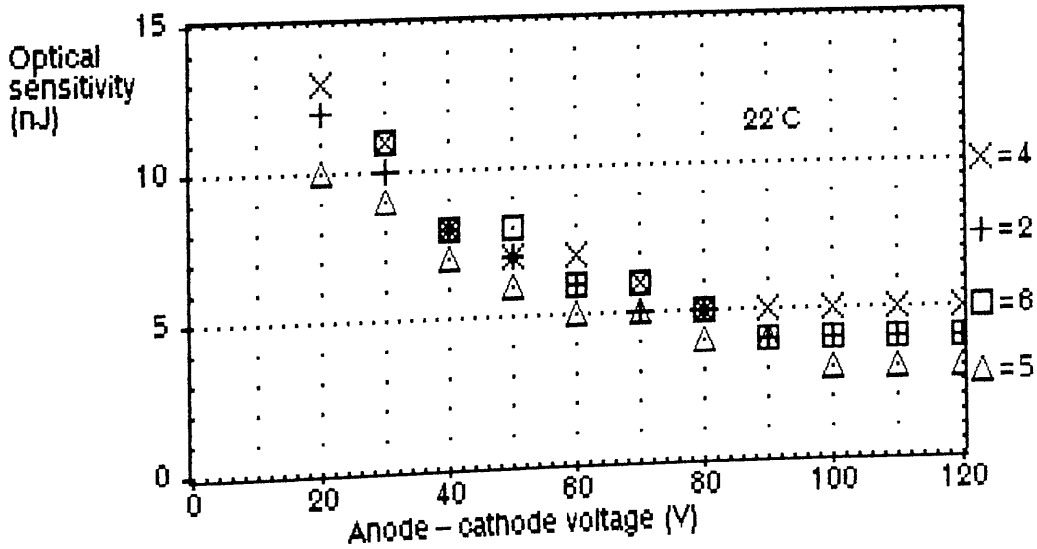

Figure 6.5.33: Optical sensitivities of wells 2,4,5 and 6 from slice #46, 22'C.

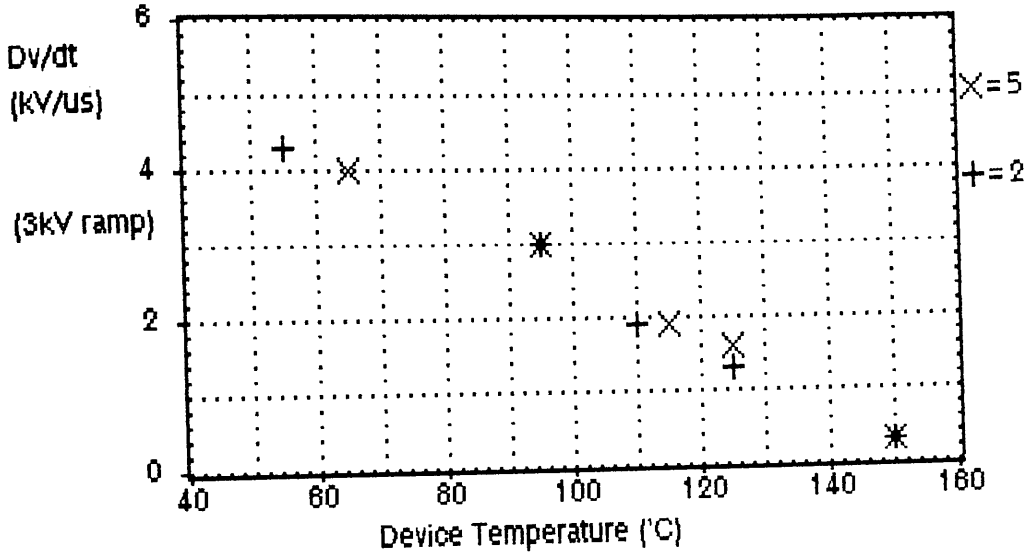

Figure 6.5.34: Dv/dt switching levels for wells 2 and 5 from slice #13, as a function of temperature.

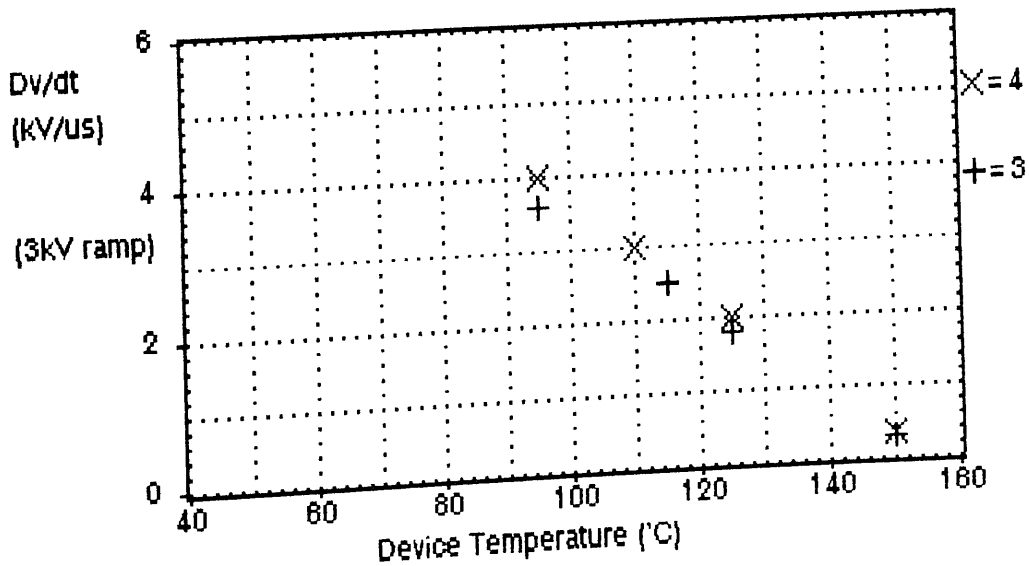

Figure 6.5.35: Dv/dt switching levels for wells 3 and 4 from slice #5, as a function of temperature.

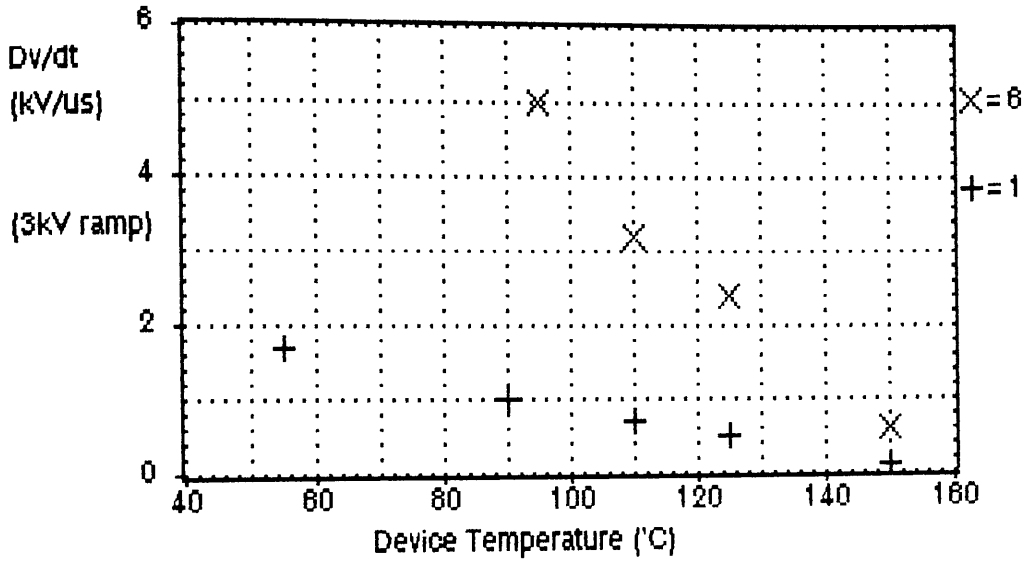

Figure 6.5.36: Dv/dt switching levels for wells 1 and 6 from slice #16, as a function of temperature.

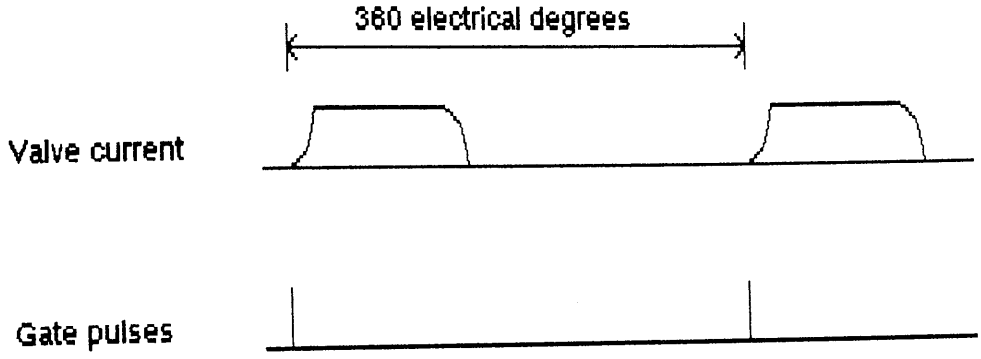

Figure 7.2.1: Gate pulses required for normal operation.

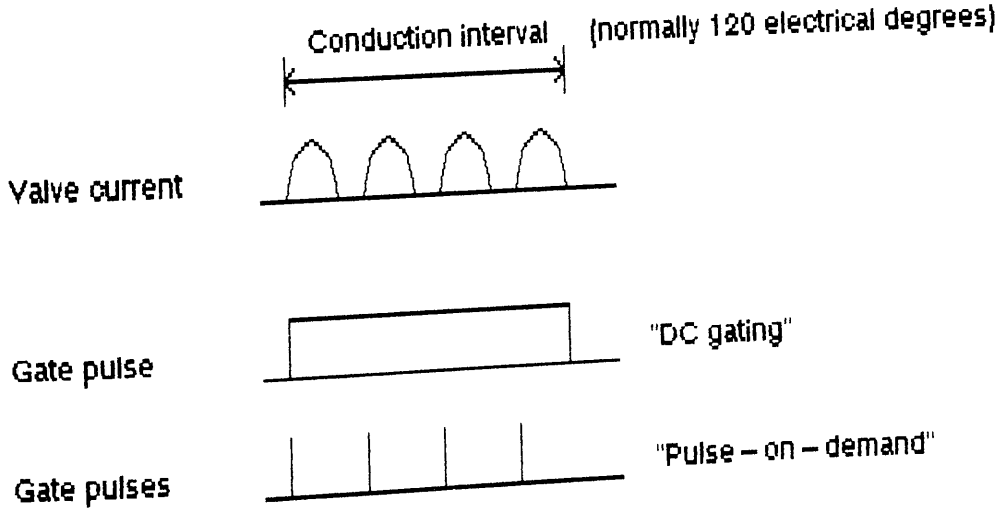

Figure 7.2.2: Gate pulses required for intermittent conduction.

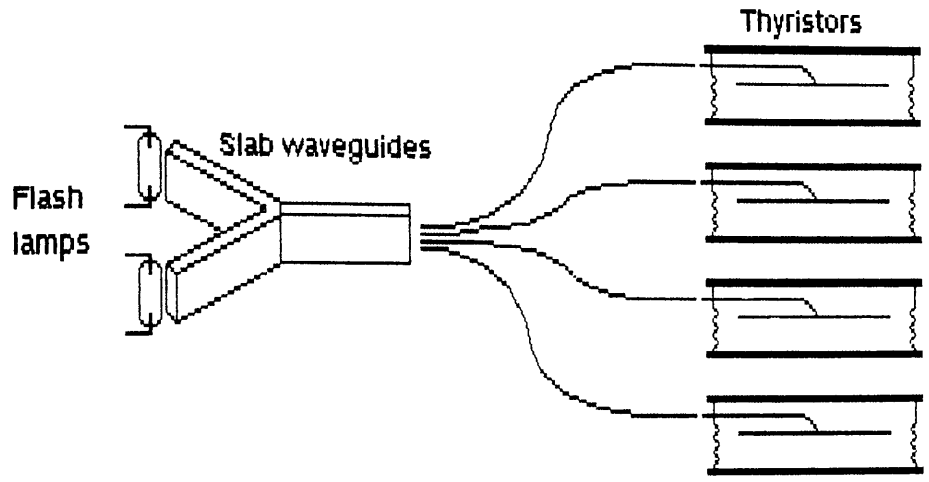

Figure 7.3.1: The GE/EPRI system.

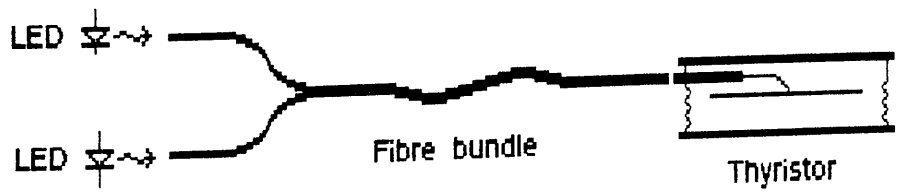

Figure 7.3.2: The Toshiba system.

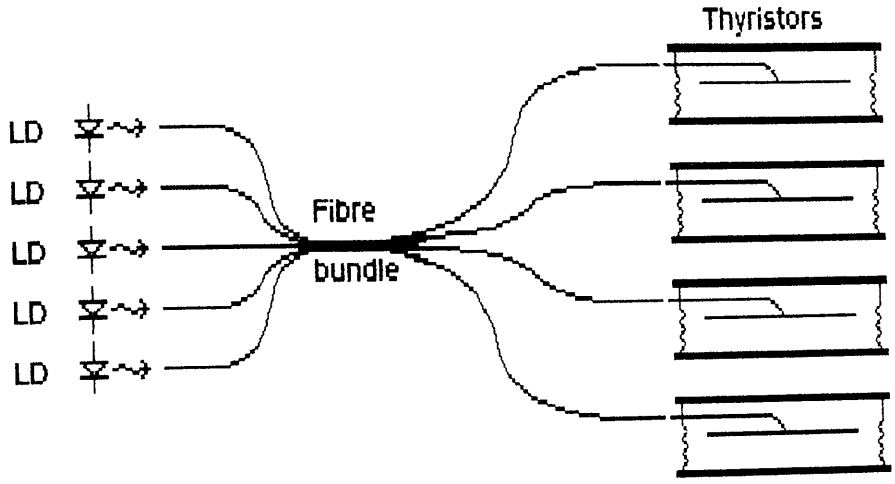

Figure 7.3.3: The Westinghouse system.

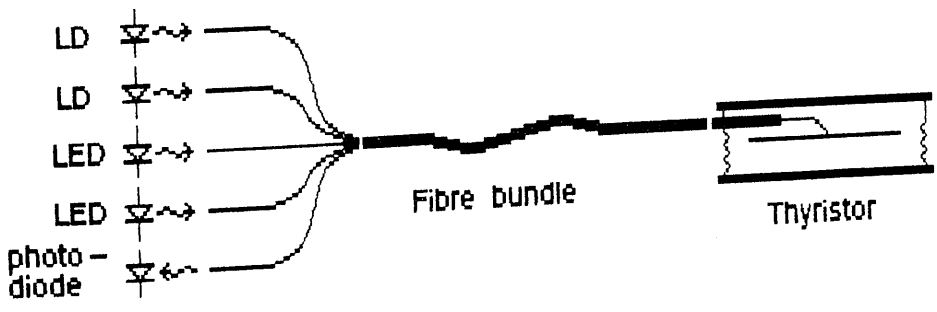

Figure 7.6.1: Design of System 1.

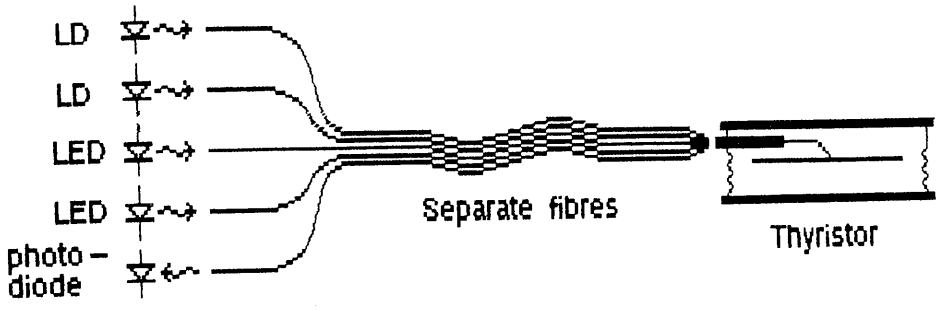

Figure 7.6.2: Design of System 2.

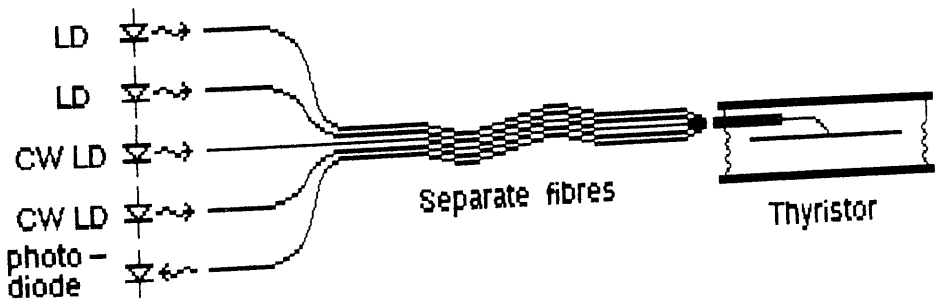

Figure 7.6.3: Design of System 3.

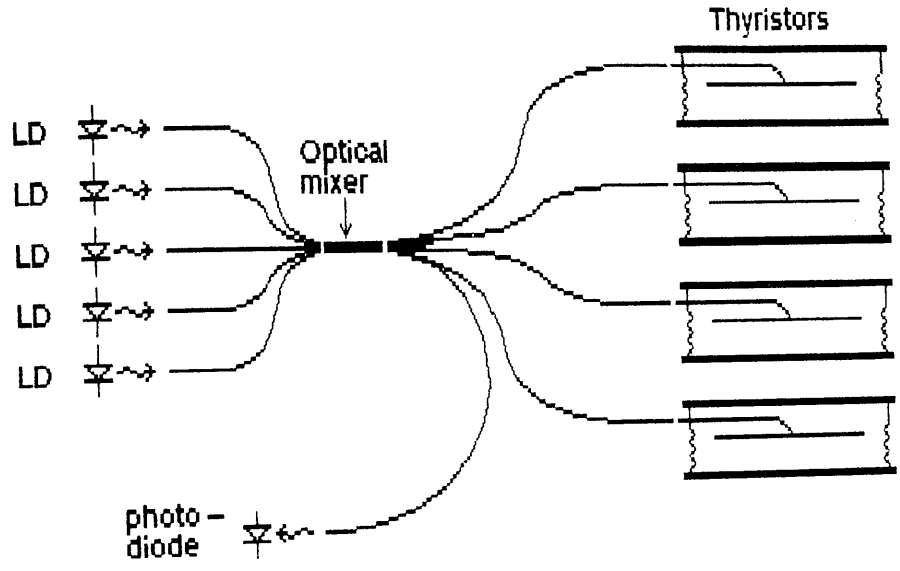

Figure 7.6.4: Design of System 4.

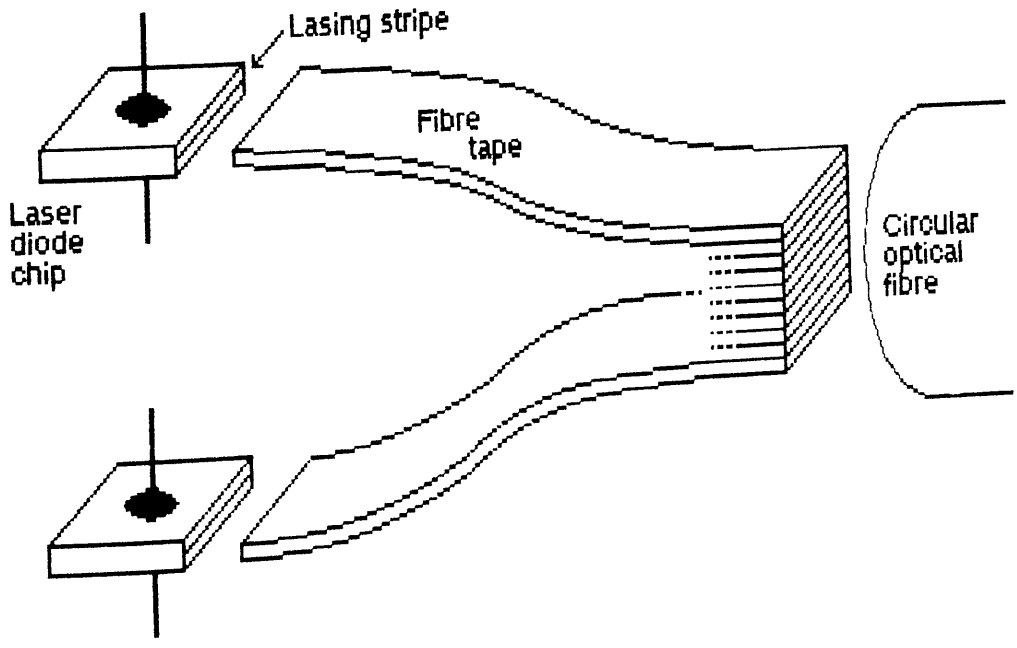

Figure 7.6.5: Arrangement of fibre coupling in a laser array.

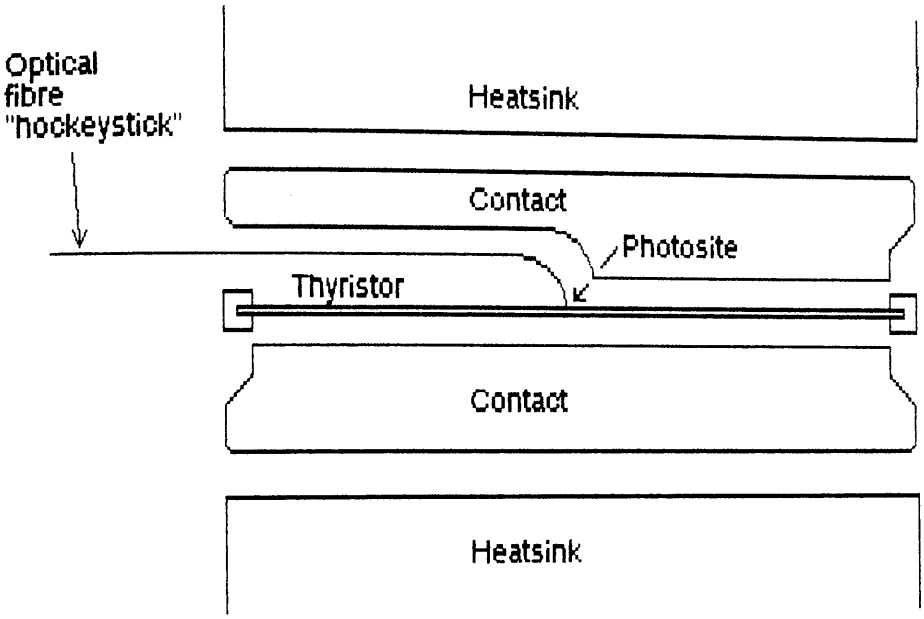

Figure 7.7.1: General arrangement of the thyristor package, with the optical fibre entering the side of the package.

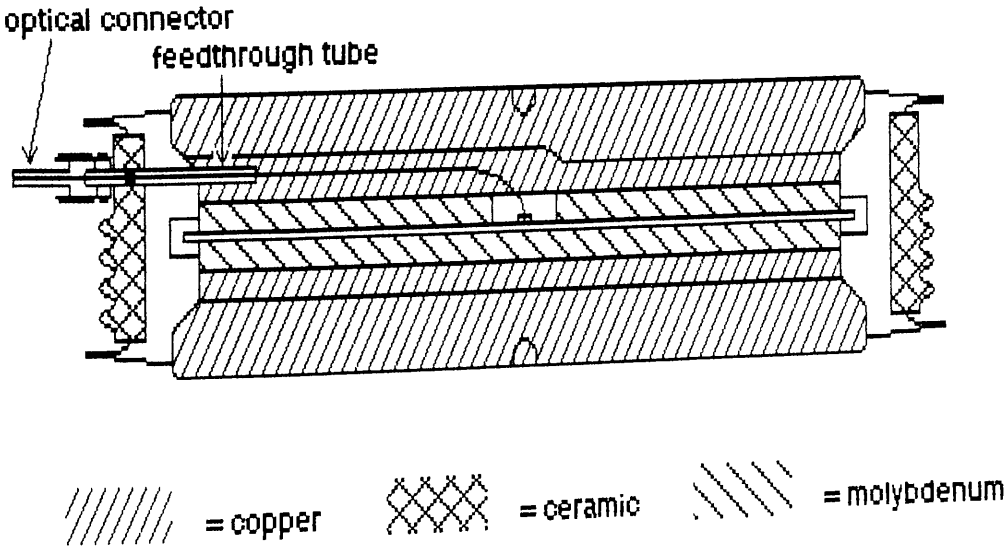

Figure 7.7.2: Proposed package design.

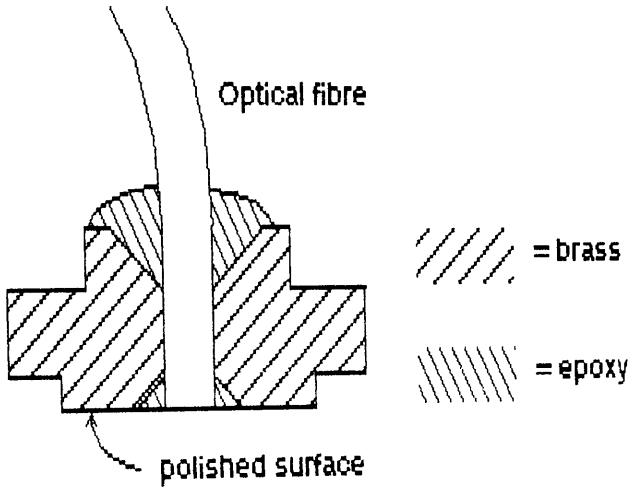

Figure 7.7.3: Brass ferrule attached to the photosite end of the optical fibre.

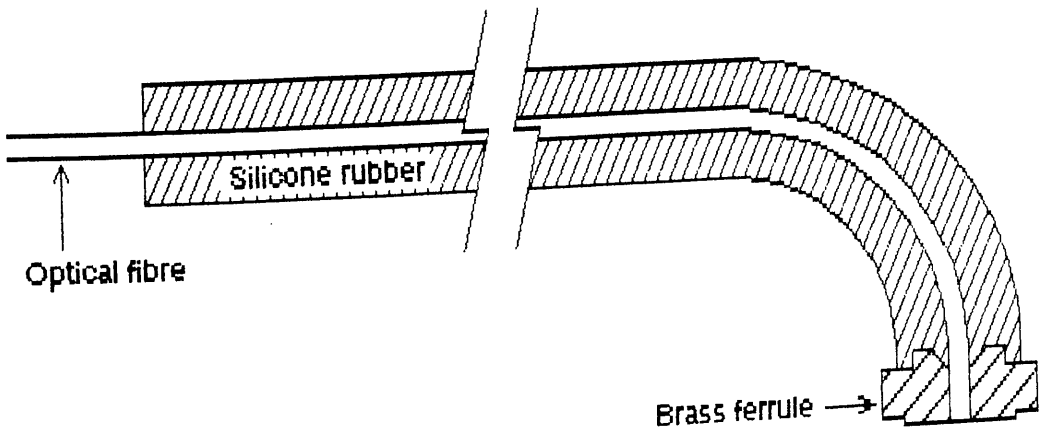

Figure 7.7.4: The optical fibre with the silicone rubber coating in place.

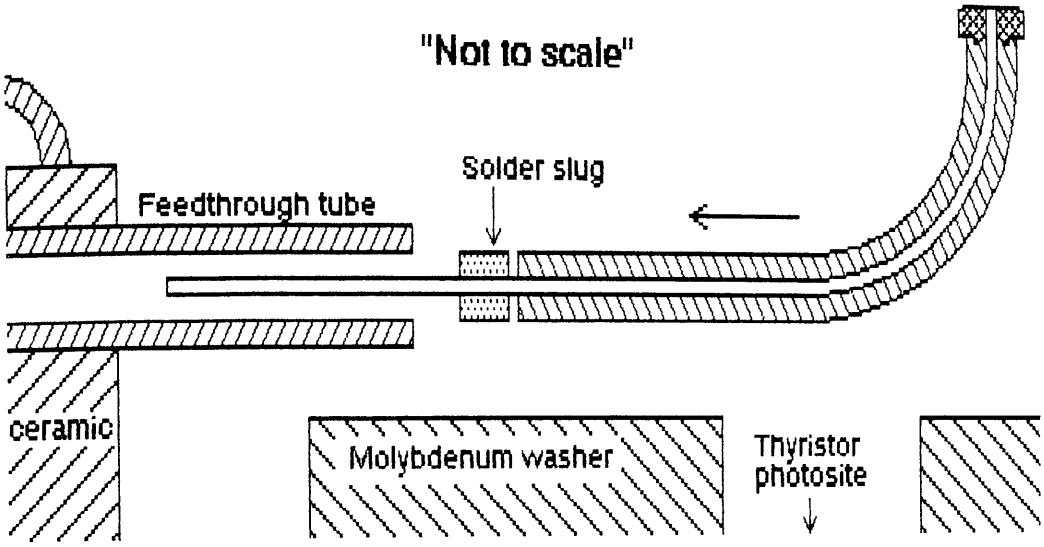

Figure 7.7.5: Sliding the optical fibre into the feedthrough tube.

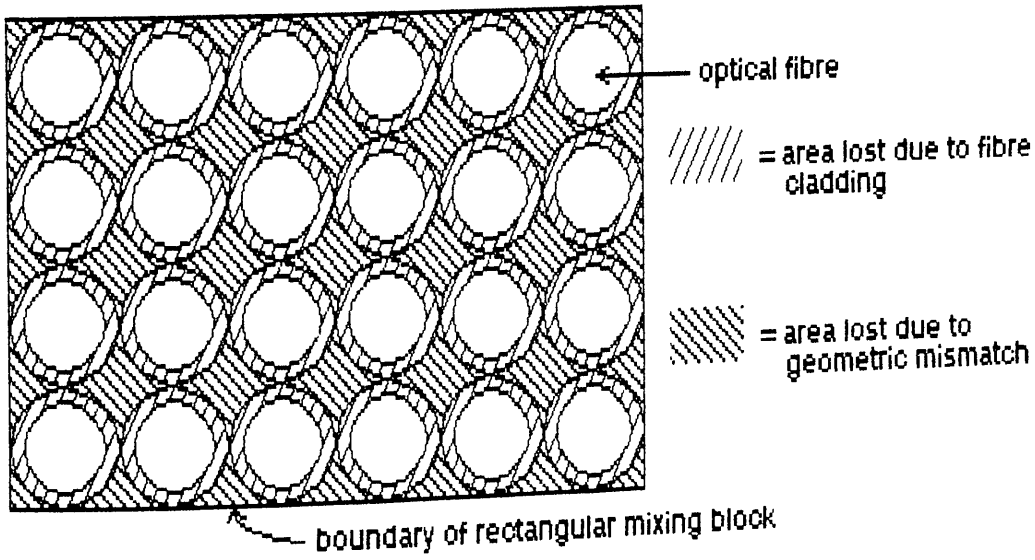

Figure 7.8.1: Rectangular array of output fibres from the optical mixer.

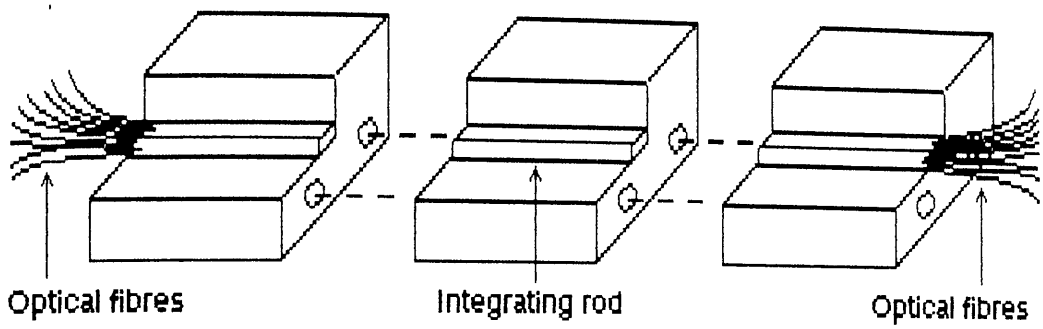

Figure 7.8.2: Arrangement for coupling the fibre arrays to the optical mixer rod.

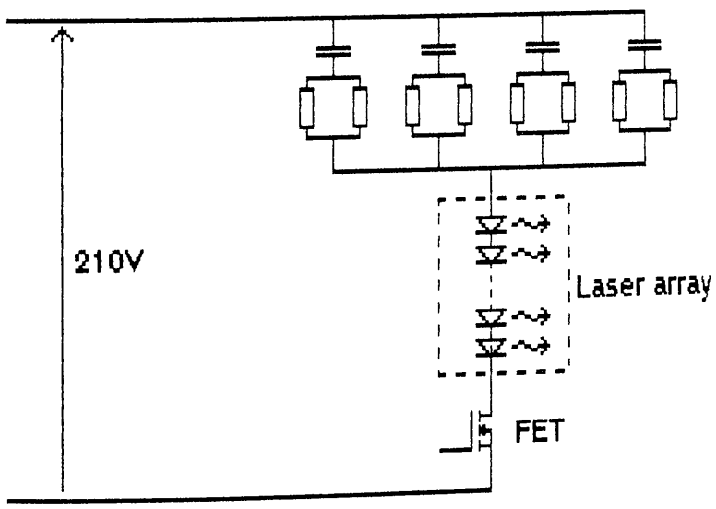

Figure 7.9.1: Basic design of the laser drive circuit.

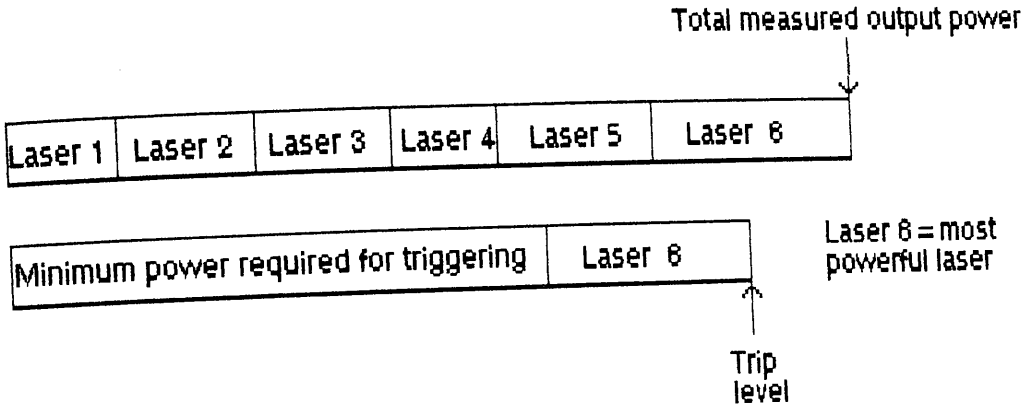

Figure 7.10.1: Derivation of the system Trip level.

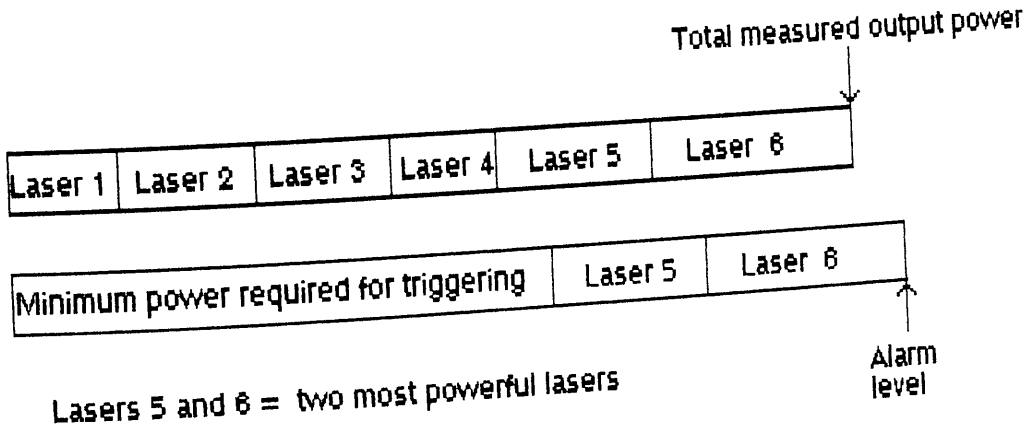

Figure 7.10.2: Derivation of the system Alarm level.

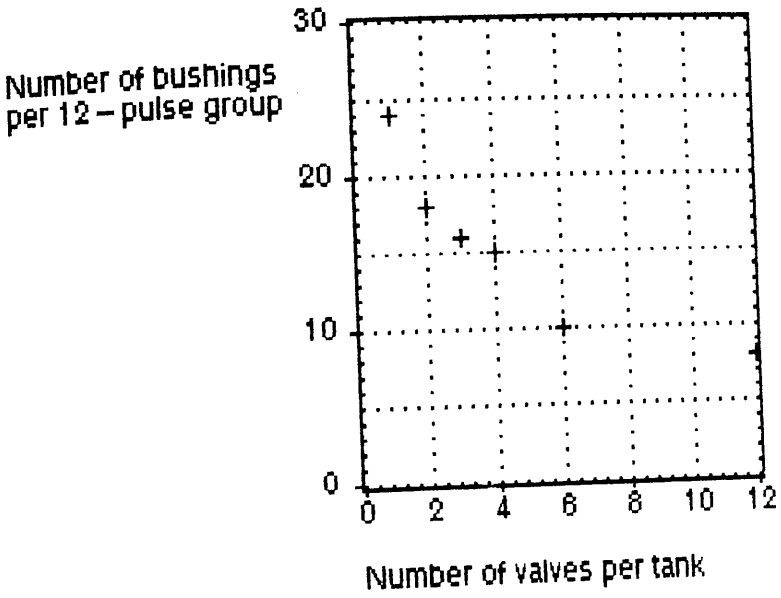

Figure A2.1: Number of bushings per 12-pulse group as a function of the number of valves per tank ("dead-tank" design).