Circuit and Electromagnetic System Design Notes

Note 19

May 1975

'LARK' -
A MODEST REPETITIVE PULSE GENERATOR

C. Edwards

Imperial College of Science and Technology
London, England

M. D. Hutchinson
J. C. Martin
T. H. Storr

Atomic Weapons Research Establishment
Aldermaston, England

"Lo! Hear the Gentle Lark........"

Originally published as
SSWA/JCM/755/99

INDEX **Page**

INDEX

INTRODUCTION

This note summarises the progress made after 5 weeks in building a nominal 600 kV 10 pps repetitive short pulse generator 'LARK' (low amplitude repetitive kit, for those interested in acronyms). The system has not yet been run at its nominal output levels, but the measurements made to date and reported below suggest that its design goals could be met if required. Since there may be interest in the future in rep. rate systems, it seems just worthwhile describing the approach adopted in this case and the areas of difficulty experienced (which were surprisingly few). Ideally, the writing of this note should have been postponed until further work with and on the system had been completed, but owing to the rapidly fading memory of one of the authors (JCM), it was deemed desirable to record the data while it was still fresh (for want of a better word) in the mind. Details of the various calculations (again, for want of a better word) are not given, as nearly all of the numerical methods used are given in earlier notes. As regards accuracy, the uncorrected voltage waveforms are probably good to 2 to 3 per cent, while many of the other values quoted (such as heating rates, etc.) are good to no better than 10 per cent.

GENESIS

The only begetter of the system is Professor Dan Bradley of Imperial College and any errors and mistakes we made are, of course, entirely his responsibility. The AWRE pulse group had felt for some time that it should do some work on repetitive pulses, but we have had no valid internal application for such machines. Owing to innate laziness and pressure of other work, we had procrastinated and this work was done directly as a result of the stimulation and encouragement afforded us by Professor Bradley, who had (or arranged to have) a valid requirement. In addition, he very kindly made it possible for one of the authors (CE) to come down from Imperial College and help in the construction and testing of LARK. Thus any merit in the work reported is also directly his responsibility.

STRATEGY

As one of the authors declines into his dotage (again JCM), he has become obsessed with the reasons (if any) why things are done: and this section and the next are the result of this obsession.

There are a number of earlier nanosecond multi-pulse generators which have been built. Notable among these are those used for streamer spark chambers, which application was pioneered by Dr Al Odian and his colleagues at Stanford Research Institute, some 7 years ago. In addition to these, there is a 250 kV 5 pps pulser built by Physics International, which uses a DC-charged single gap approach. Additionally, there are large output systems such as that made by Maxwell Laboratories, which provide a burst of ten pulses at high rep. rate from a series of separate modulators. Within the time and effort constraints placed on the programme, none of these approaches seemed directly applicable to LARK. In particular, the required output of 600 kV mitigated against a single spark gap DC-charged approach. In addition, it was felt that it was desirable to obtain experience with a pulse-charged output gap. However, with the period of only 5 weeks allowed for construction and testing of the pulser, it was felt that attempting to solve the problems of -repping a low inductance Marx generator at 10 pps might be overly ambitious.

Thus it was decided to employ an air-cored pulse transformer to pulse-charge a home-made high voltage capacitor. A self-closing pulse-charged output gap then connects this to the load. (Figure 1 shows a schematic of the system.) This approach had the virtue of only 2 moving parts (one 'DC'-charged gap; 1 pulse-charged gap) and meant that we could fairly quickly get to the areas of uncertainty. The longer-term disadvantage of the approach adopted is that it is limited in output to a couple of kilowatts or so of delivered power. It is interesting that earlier we had found our pulse transformers were limited to a delivered energy per pulse of a few kilojoules, due to mechanical effects, etc. The limitation in reasonably sized transformers to a few kilowatts is due to thermal effects. Thus, in any more substantial output system it would probably be necessary to go to a Marx generator as the main voltage multiplier, just as it was earlier found necessary to do for single pulse work.

Consideration was given to obtaining some extra voltage gain in the high speed section by using the lumped element version of the Blumlein. However, this would have involved an extra pulse-charged switch and as the pulse-charged energy for each shot was quite low (~ 100 joules), it was feared that the effect of stray capacities, etc., would lead to a disappointing gain. Thus it was decided to obtain the voltage gain needed (700 kV for about 600 kV out) entirely from the transformer.

We have built plus and minus pulse transformers with total outputs of 3 MV and hence a single-sided transformer of the required output was obviously a possibility. However, the higher the single-sided output voltage is pushed, the more prone to flash-over the transformer becomes, if its output is allowed to ring on. Such a fault condition is bound to occur in useage, as the output gap has to operate close to the peak of the charging waveform in order to give a reasonable energy transfer efficiency. In addition, the leakage inductance of a single-sided transformer is 4 times that of an equivalent double-sided one and this leads to reduced gain and hence to reduced energy transfer efficiency. Thus despite the difficulty that the 'DC'-charged condenser bank excurses in potential to half the pulsed output voltage, it was decided to use a double-sided transformer, as shown in Figure 1.

In addition to the major decisions outlined above, a number of other more minor ones were taken. Most important of these was the one to run the system under atmospheric pressure freon. This was decided because of ease of access and maintenance, compatibility with our home-made high voltage capacitor, and the low weight of the system compared with, say, using transformer oil as the general insulant. There was also a minor gain to be had, from the user's point of view, with regard to fire risks, at the expense of having to discharge mildly disintegrated freon from the building. Repeated discharges at the edges of the capacitors did lead to a small amount of noxious ex-freon products, but the major effluent irritant came in practice from the air coming from the spark gaps. A second minor decision was to run the gaps with air in them, since it was felt that a fair rate of flow would be necessary and compressed air has the user advantage of cheapness and availability. It was also decided, in a rare moment of kindness to computers, to take significant pains to make the system reasonably RF tight. It had been found that a small unshielded test rig had radiated badly and because of the rep. feature was easily locatable as the source of interference. (Single pulse systems can get away with murder in this respect.)

In concluding this section it should be emphasised that nearly all the above decisions were dictated by questions of convenience and speed and that the solutions adopted for LARK are almost certainly only applicable to this particular system and should not be taken as general guidelines. If this note has any interest to others, it is probably as a record of the difficulties (or lack of them) that were encountered in the work and the performance of some of the components such as the output gap and the high voltage pulse charged capacitor. Having said this, it does, however, appear that most of the decisions were sound ones, since all the limitations to the approach observed were all expected beforehand and mostly correctly estimated ahead of time.

TACTICS

The main difference between single pulse systems and rep. rate ones would seem to lie in 3 areas. Firstly, will the components last for several million pulses irrespective of the rate at which it is fired? Secondly, will they operate at the required rep. rate for brief bursts? Thirdly, will the system operate for prolonged periods at the required rep. rate?

These questions are not completely decoupleable, but are to a useful degree. With regard to the large required life of the components, regardless of the rep. rate, this mainly concerns the behaviour of the solid dielectric in the DC and pulse-charged capacitors and the transformer, but also involved erosion in the spark gaps. It was felt that the behaviour under a lot of shots of the freon dielectric would be alright, as it would be changed reasonably frequently in use. If the required capacity (~ 400 pf) could have been obtained with an atmospheric pressure gaseous dielectric, this would have been used. However, the use of liquid or solid dielectrics was dictated by stored energy density considerations and it was considered that as there was highly stressed solid dielectric in the transformer, it was better to stick to this in the pulse-charged capacitor, even though a solid dielectric would have a finite pulse life at any given stress. In order to provide data on this point, some small-scale tests were performed at elevated stresses, to give estimates of the allowed stresses in the pulse-charged capacitor and to check that the corona discharges at the interfaces in this did not significantly degrade the life. The results of these tests will be briefly summarised in the sections dealing with the components involved.

The second question relates primarily to spark gap deionisation time. The start (or 'DC'-charged gap) would not be expected to give any trouble at 10 pps when fed by a suitably designed power pack. However, we had to use resistive charging because of the available power packs and it was feared that the arc might not extinguish. In addition, aerodynamic considerations of the movement of the hot gas in the large spacing of the output pulse-charged gap suggest that this too might not recover between shots.

The third question mainly relates to heating problems. Normally, in a single pulse system delivering 100 joules per pulse to a load, overall efficiency is of very little concern. However, in a rep. rate system all the energy not delivered to the load finishes up somewhere else in the pulser and can easily lead to thermal problems. For instance, a loss of 5 joules per shot in a component is insignificant for single shot work, but at 10 pps this leads to the dissipation of 50 watts in that component. This may not sound too bad, but try feeling a 50 watt bulb while it is running. The high voltage requirements lead to nearly all the components being largely made out of good dielectric materials and these are both good thermal insulators and cannot in general be allowed to heat to temperatures much above 80° C. The solutions adopted in LARK are to use forced gaseous cooling of the transformer and gaps, natural convection cooling of the pulse-charged capacitor and monitor chains and the heat sink properties of the DC-charged capacitor.

In order to get a feel for the above problems, and many more, a roughly half-scale system was quickly built by Tommy Storr and tested under various conditions. These tests showed that the half-scale transformer and capacitor could withstand a 500 kV ringing waveform and was pushed to 560 kV with a switched-out pulse for a few shots without any trouble. It was also used to show that both gaps would operate at 10 pps (at reduced voltages) and also gave information about the energy dissipated in the various components. The half-scale model was also run for some 300, 000 shots at stresses significantly above those applying to the full-scale system and thus gave assurance that the life would be at least a couple of million shots. This work took some 5 weeks of Tommy's time before Chris Edwards came down to help build the full system and provided an invaluable guide as to what things were worth bothering about. Again the results obtained in the half-scale tests (HST) will be referred to below, where they provided information of relevance.

If the reader will indulge a positively decrepit author (JCM, yet again), I feel that the tactics employed as described above were an almost essential part of the development and without the results of the quick and dirty experiments we would have under-designed some parts and, almost as bad, over-designed others. For instance, the design of the transformer might have had its high voltage properties compromised because of heating worries, or vice versa. In addition, when we came to build LARK, we knew the approximate nature of all the components and none of these had to be altered in a major way. The man-hours spent on the preliminary experiments could easily have been lost if a major redesign and rebuild of any of the 6 main components had been necessary. Additionally, much valuable experience was gained on more minor details, such as monitoring, layout, operation, fault modes, etc., ahead of time and this greatly shortened the final construction and testing phases of the full system.

PERFORMANCE SPECIFICATION

There were 2 distinctly different approaches to this. On our part we inclined to the view that the required output should be that which the pulser provided. Dr Henry Hutchinson of Imperial College (who has provided much appreciated moral encouragement and a power pack) inclined to the view that the output should be twice any figure that we tentatively mentioned. Eventually a compromise was reached, in that the aim should be a system capable of eventually working at 10 pps and providing a peak output voltage of 600 kV into 200 ohms. The rise time of the pulse should be a few ns and the pulse duration at 400 kV should be around 50 ns. This corresponds to about 70 useful joules per pulse. An exponentially decaying pulse would be satisfactory, but if the fall of the volts below 400 kV could be hastened, this would be desirable. If the system could work with a lower impedance output load (with a corresponding reduction in pulse duration) this, too, would be desirable. No mention was made of prepulse, but we resolved to design the system with a prepulse of less than 20 kV, but also with the option to provide a significant prepulse in case the switch-on time of the cathode proved excessive.

GENERAL DESCRIPTION OF LARK

The pulser is contained in a freon-tight box, 64" long by 42" wide by 28" tall. As the delivery date of the proper metal box was found to be longer than 5 weeks, a temporary box, made up of perspex, mylar, and aluminum baking foil, was knocked up. That this was possible was an incidental advantage of using freon and not transformer oil. Print 1 gives a photograph of the system in the temporary box, with the transparent gauze-covered wall removed. The 'DC'-charged capacitor is on the left and from the top of this a strip line connects it to the primary turns of the transformer. The spark gap is located between the capacitor and the transformer in the underside of the strip line. This somewhat bizarre location was dictated by the desire to use negative charging volts and to avoid having the gap stick up above the line towards the roof. In use, a probably unnecessary perspex cover is placed over the top of the capacitor and the feed to the transformer, but this has been removed for the photograph.

Next to the capacitor bank is the transformer. The positive output of this is taken down the central axis of the transformer and earthed to the bottom of the box. The negative output (the outside turn) is connected to the pulse-charged capacitor which occupies the centre of the print. The bottom of the capacitor stages are raised about 1 inch off the base by an aluminum foil-covered foam base. There are 12 stages, or flat pads, of overall height 8". Above these there is a rather carefully contoured 3-D high tension bung made of foam covered with metal foil. The capacitor stages are assembled partly by DC pressure from 140 lb of lead blocks on top of them and partly by electrostatic forces from charge separation after it has been fired a number of times.

To the right of the pulse-charged transfer capacitor is the pulse-charged output gap. This has a number of fins around it to prevent outside tracking during its pulse charging. On the right-hand side of this gap is an elliptical HT bung which is so 3-D, it is practically 4-D. Opposite to this, on the output face of the pulser, is an additional raised field shaper with a 12" OD hole in it. One function of these two bungs is to grade the diode when it is in place. In the present experiments, as is shown in the print, a copper sulphate load and output voltage monitor replaced the diode.

The design of a 600 kV+ adjustable load, to dissipate 1 kilowatt and be no longer than 22 cm, would seem to be a challenging problem. In practice, these requirements were easily met and the design also allowed for the inclusion of a 2 ns rise time high impedance (16 kΩ) voltage monitor. Also visible in print 1, below the output gap, is the 32 kΩ capacitor voltage monitor resistor chain. Not shown in the photograph is the cooling arrangement for the copper sulphate solution of the load, external to the temporary box.

The final box, which has now been received, will have the 2 long sides transparent. These sides are covered externally with blackened copper gauze, which will be electrically bonded to the rest of the metal box. The gauze is blackened to improve visibility into the box. It was feared that uv from the repping gaps might be an eye hazard. However, while the light output of the gaps is substantial, the radiation has to pass through at least 1/2" of perspex and this effectively cuts off all radiation below 3500 Å. According to the literature, radiation above 3200 Å presents no additional hazard to the eyes compared with any other intense light source and the main hazard of conjunctivitis arises mainly from radiation below 2800 Å. However, there may be some risk to epileptics at 10 pps, as the system is rather luminous and makes the laboratory look rather like a cluttered Disco without music. In the proper box, an area over the output end will be semi-permanently installed and the thermosyphon cooling radiator will be mounted on this, along with a fan. The rest of the lid will be removable for access and will also have a motor on it to drive the transformer cooling fan, rather than the elegant structure shown in print 1.

Before dealing in detail with components within the box, mention should perhaps be made of the components required outside the box. These comprise a 0-30 kV power pack of at least 2.5 kVA capacity, if the desirable constant current charging is used. If resistive charging were used, a pack of about double the wattage would be needed. Also required is an up to 60 psig air supply, capable of delivering about a litre a second. Also there are 2 reducing valves and pressure gauges, to supply the two gaps independently. In addition, freon is, of course, required and an effluent line to discharge outside the laboratory the used air and spill-over freon. The rate of useage of necessary freon has not yet been determined, but is of the order of one box change per couple of hours, or quite probably even less frequently.

The next sections will deal in some detail with the design and performance, so far, of the major components, in the order listed above. However, before this, the approximate values of the electric components will be derived.

APPROXIMATE COMPONENT VALUES REQUIRED

In order to provide the required pulse length with a 200 ohm load, a pulse-charged capacity of about 420 pf is required. Not all of this capacity has to be provided in the pulse-charged capacitor, as a fraction of the capacity of the transformer is available to the load during the useful phase of the pulse (~ 30 pf). In addition there is some 10 pf stored in the strays in the freon. Thus the pulse-charged capacity needs to be around 380 pf. If this is charged to 700 kV, some 70 per cent of the stored energy is available as useful pulse energy and this will provide about 70 joules to the assumed 200 ohm constant impedance load, by the time the volts have fallen to about 60 per cent of the starting volts. This is providing the discharge time of the capacitor is reasonably short compared with its charging time. Such a capacity stores about 103 joules at 700 kV and allowing for the energy stored in the output capacity of the transformer, which is not available during the useful pulse, about 120 joules per pulse has to be provided by the transformer. A reasonably sized transformer can be made up to 60 per cent efficient, hence requiring an energy stored in the 'DC'-charged capacity of about 200 joules, provided the output gap can be closed pretty near to peak volts. This raises an important point: it is very desirable that both gaps have a rather small voltage jitter at closure. For instance, if both gaps have a 5 per cent standard deviation and it is required that the output gap fails to close no more frequently than 2 per cent of the time (this is for heating reasons), the mean firing level of the output gap must be set some 20 per cent below peak volts on the capacitor and an initial stored energy of some 300 joules is required. While much of this extra energy will finish up in an ohmic load, this may not happen to the same extent when the pulser is driving a diode. In either event there will be a considerable extra dissipation of energy in the components, leading to significantly and perhaps considerably increased heating problems. Thus it is much more important to strive for overall electrical efficiency in a rep. rate system than it is in a small single pulse machine.

Having established the approximate stored energy in the bank, the actual value of the capacitors used depends on what is available and what gain can be built into the transformer without losing too much efficiency. Calculations suggested that double-sided transformers with gains of around 20 to 25 could be made with the above quoted efficiency. This implied starting 'DC' voltages of 30 to 37 kV. Fortunately we had some suitable capacitors which could be arranged easily to give 0.5 microfarad and hence a maximum working voltage of 30 kV resulted. The rough circuit values are given in the schematic in Figure 1.

'DC' CHARGED CAPACITORS AND VOLTAGE FEED

To make the store, 2 capacitors were used. These were nominal 0.25 μF, 75 kV capacitors made by CSI and obtained from Hartley Measurements Ltd., in the UK. The capacitors have nominal dimensions of 3.75" x 8.5" x 13" and a life of 10^5 shots. Because of the potentially large number of shots they might have to provide at internal temperatures of up to 60° C, it was felt that operating them at a maximum voltage of 30 kV was a desirable, if not necessary, derating. The particular capacitors obtained have their output electrodes along the 8.5" dimension, a necessary arrangement to provide a low inductance and resistance. When tabbed with a wide strip line output feed, the inductance of the pair was about 15 nH and the resistance 15 milliohms, and the measured total capacity was 467 nF.

During the charging of the capacitor and rapid discharge, about 1 1/2 per cent of the stored energy is dissipated in dielectric losses. In addition, about 4 1/2 per cent of the stored energy is dissipated in the metal conductors during the discharge phase and subsequent low amplitude damped oscillation. Thus about 6 per cent of the stored energy was calculated as being deposited as heat in the capacitor in each shot, where the output gap closes. In those cases where this does not happen and the circuit rings on damping slowly, about 30 per cent of the stored energy is dissipated internally in the bank. Thus it is desirable to restrict the number of shots on which this happens to a per cent or two.

In order to check the above calculations, the system was run for about an hour at 6 pps, with the capacitors as thermally insulated as possible. In fact the ambient freon temperature rose at almost exactly the same rate as the temperature of the case and the cooling correction was negligible. Using a mean figure for the thermal capacity of the capacitor material of 2 joules per °C per cc, the total temperature rise gave the energy deposited in the capacitors, and hence the fraction of the energy stored which was deposited in them per shot. This came out at 6 per cent, in accidentally exact agreement with the calculations. If these capacitors are used at the nominal maximum conditions (200 joules per shot and 10 pps), the rate of rise of temperature at the centre (and essentially at the outside) is 20° C per hour. Hence, a two-hour run can be allowed.

Clearance between the cases and the holder were provided, as well as ventilation holes, but the natural convective cooling in freon would be very poor. In any case, even if forced cooling was arranged, so as to keep the cases at ambient temperature, the centre temperature of the capacitor at equilibrium would be roughly 150° C above ambient, and obviously unacceptable.

In the present application the maximum running time is limited to 2 hours and the capacitors act solely as a thermal sink. After the 2 hour run, the capacitors will have to be left for a day to cool down, according to calculations, but this has yet to be checked experimentally.

The thermal restrictions on the capacitors represents an important limit on any pulsed system, whether it be a fast Marx or a transformer one, as in the case of LARK. The capacitors are already being run at one-seventh of their nominal stored energy and yet they cannot be taken anywhere near their equilibrium central temperature. It is possible that in a new capacitor design the internal resistance can be reduced by an order of magnitude, by using much thicker metal foils and connectors, even though the present condensers have a pretty low internal resistance. A different impregnant and dielectric may reduce the dielectric loss by about a factor of 2. Both of these changes would reduce the stored energy, however; the result would be that with them the present capacitors could be run to thermal equilibrium at the 100 joules stored in each per pulse. This is for a capacitor of nominally 750 joules storage for 10^5 shots.

It seems even improved rep. rate capacitors will have to be severely derated, from a thermal point of view as well as a life one, and the cost per joule will be most of a magnitude up on present costs for fast discharge capacitors, unless capacitor manufacturers can play some quite clever tricks. Obviously the preferred shape of a rep. rate capacitor is as thin as possible and of as large an area as possible. However, apart from the extra cost of making such capacitors, the allowable working stress through the capacitor will limit how far this approach can go. Attempts to cool the capacitors internally would appear to be non-starters, even using heat pipes, since these would have to cover essentially all the area of the pads, to be effective, and hence be very expensive as well as possibly introducing voltage breakdown problems.

The capacitor bank rises to approximately half the output voltage of the transformer, that is to about 350 kV maximum. The connections to the power pack and earth from the bank have to have impedances in them to withstand this voltage, of a suitably high value so that the energy loss down them is small. Two uniformly wound 1.4 mH inductances were carefully located in the fringing field of the capacitor and transformer so that they were also pulse graded uniformly. These inductors are shown most clearly at the rear, right, in print 2. In addition, two resistances are included in the connections (100 ohms in the earth and 50 ohms in the HT lead) in order to increase the long period damping of the ringing waveform when the output gap fails to fire. In addition, a filter capacitor of 10 nF was added at the end of the HT cable, to reduce the voltage transients appearing on this in the non-fire case. The HT feed cable has a fairly onerous duty cycle on it of up to 30 kV at 10 pps and any additional pulsed voltages could lead to a serious reduction of its life.

Figure 2 shows the capacitor isolation and feed circuit.

The present power pack arrangement is a temporary expedient and consists of a transformer feeding a 6 μF capacitor. The smoothed output from this is fed by a high wattage resistor network to charge the capacitor. The resistance is arranged so that the charging RC is approximately one-third of the time between pulses. Thus for 10 pps its value would be 70 kilohms. As soon as time permits, it is intended to change the power pack to a constant current charging arrangement.

START GAP

The first version of this that was used on the half-scale test rig was a gap we had to hand which consisted of a pair of rail electrodes some 8" long, in a 1 1/2" OD perspex tube. The spacing used between the electrodes was about 3 mm and its inductance was slightly over 10 nH. The brass electrodes were flattened over their length, as well as being contoured at the ends, so that the outer regions of the electrodes acted as a pressurised gas feed to the central region where the spark self-closed on the charging waveform. The gap operated well for some 3×10^5 shots and after this an estimate was made of the small amount of material removed by erosion. The main disadvantage of this gap was that its jitter was significant, being of the order of ± 4 per cent, although there were periods when it was running when the jitter was much less than this. There was a considerable deposit of debris on the inner walls of the gap, but at no time did it show any signs of wall tracking. Because of the importance of reducing the jitter, it was decided to build a new gap which was to include two improvements - corona irradiation and a vortex spent gas extraction system.

The corona irradiation was arranged by introducing a 10 thou. platinum wire into a hole drilled in one of the electrodes (see Figure 3). This corona wire was charged negatively via a 2 gigaohm resistor chain from the appropriate feed to the gap. At atmospheric pressure the voltage of the point charges linearly to about 3 kV and then collapses to about 1 1/2 kV as a burst of corona forms at its tip. The stray capacity of the needle and the chain, and the resistance of the latter, determines the repetition rate of the corona, which in this case was about 1 kilocycle. Each corona pulse provides a weak burst of uv and possibly injects some electrons into the gap. In a small test gap working on the bench with a very small energy load, the standard deviation was about ± 0. 03 per cent in the breakdown voltage, when it was running between 1 and 10 pps.

When the new start gap was freshly installed, the repetition rate was very constant. However, as the gap eroded, the performance became less constant, as the arcs moved on occasion away from the irradiated central region. However, the jitter was still better than the unirradiated gap and it is intended to recontour the electrodes and change the central electrode material to a tungsten alloy, to reduce the erosion rate, hopefully. The point will also be advanced, so as to irradiate a larger volume of gas.

One difficulty was experienced after the gap had been left a couple of weeks sitting around, unused. It was noticed that the debris in the gap had changed its nature and had become much lighter in colour. As the air is extracted through the electrode and past the area where the corona point feed was made through the wall of a perspex tube, debris was almost certainly deposited in this region, too. Anyway, after leaving the gap around for a couple of weeks, it was found that a resistance of only a hundred megs or so had developed between the needle feed and the brass electrode. This, of course, prevented the point charging up to the corona inception voltage. Presumably the debris had chemically reacted and become deliquescent, causing a low impedance film to occur, linking the two points. A new feed-through bush was made to support the needle with a much greater insulation length and also with ridges on the surfaces, so that there were regions on which the debris would be unlikely to settle. A very quick way to check whether the corona point is working is to hook an electrostatic voltmeter to it, when this will read the mean voltage of the needle. Table I gives the values of this voltage as a function of the pressure in the gap.

TABLE I
FEED RESISTOR 2×10^9 OHMS

Gap Pressure (psig)	Mean Corona Point Voltage (kV)
0	2.3
10	2.9
20	3.6
30	4.3
40	4.8

It is seen that the mean voltage goes as $p^{1/2}$.

At the same time that corona irradiation was installed, it was decided to introduce a vortex spent air removal system. This was largely done to humour one of the authors (guess who?) and was almost certainly unnecessary, as the original gap had shown no tendency to stay struck when tested at 10 pps. The idea here was to remove the hot plasma column as quickly as possible, by extracting the air through holes in the electrodes. In addition, a vortex was set up so that prior to this extraction by the shortest possible path, the hot column of low density air would be stretched by the velocity shear existing in the vortex. Now, in a gap with axial symmetry it is easy to establish a vortex and, provided the flow is non-turbulent, to calculate its approximate properties. With a rail gap, things are geometrically and aerodynamically more complicated. A simple perspex model was constructed and run with water at an appropriate rate. It was not possible to scale Reynolds number, but as the flow proved highly turbulent in the water model and the gas flow case has a Reynolds number ten times higher than the water model, real life conditions would be even more turbulent than the model showed. The vortex was not a strong one, but it formed at all times, provided the flow down the two feeds was within a factor of 5 of each other. The water circulated about twice before exiting through the electrodes, so there was not as much shear as was hoped. However, tests with potassium permanganate solution showed that the flow was very highly turbulent and a region of hot gas would mix in a cone with a 20° semi-angle, or so. Thus a plasma column of 1 mm diameter would be mixed in with ten times its mass of cold air by the time it had moved about half a centimetre. As linear gas velocities of 2×10^3 cm can easily be achieved, the hot gas should be well mixed by about 300 microseconds. The water model was also used to show that the pressure drop through the vortex was about 0.1 psi at flow rates around a litre a second: in other words, negligible.

Figure 3 shows the mode of injection of the air from two opposing sides of the spark gap body and roughly sketches a mean flow line. The flow of gas also, of course, keeps the spark gap body cool; not that it warmed very greatly at 3 pps in the original gap. Debris still deposits on the wall, but only in the immediate vicinity of the sparks. This is probably because of sub-vortices caused by the rail electrodes, but may be due to debris being flung out directly from the arc roots.

Judging from results to date, the erosion is rather faster than in the original unblown gap, possibly due to the removal of more molten metal per shot. So the erosion calculations on which the original electrode bump shapes were based are incorrect and the gap has widened by about 1 mm from its original spacing of 3 1/2 mm and has had to be closed up. As was mentioned, the electrodes will be recontoured to provide more metal where the sparks occur and this metal will be made of a tungsten alloy. If this fails, the vortex blowing will be dropped and a return made to a rail gap with about four corona irradiation sites along it.

The inductance of the new gap is a bit higher than the old one, because there is a gas feed slot between the gap body and the return metal feed from the transformer. Its calculated inductance is 17 nH. This, plus the inductance of the condensers (15 nH) and about 3 nH in the feed to the transformer, gives an overall inductance of the bank plus gap plus feed into the transformer of about 35 nH.

Concluding this section, while it has not been demonstrated, it is reckoned that the start gap can have a standard deviation of closure voltage of less than ± 1 per cent and a life, between maintenance, of one million shots or more. Print 3 shows a rather poor view of the start gap under the feed to the transformer.

TRANSFORMER

Physical Description

The transformer has a two-turn primary and a total of 53 turns, giving a turns ratio of 26 1/2 to 1. The insulation is 10 thou mylar and the secondary turns are wound in 3 thou copper. The primary turns and the feeds are made out of 5 thou copper. The thickness of the winding is 1.95 cm and the thickness of the inner perspex former is 0.5 cm; that of the outer cylinder is 0.32 cm. The mean radius of the transformer is 8.5 cm. The windings taper from a primary turn width of 16 cm down to an output turn width of 1.5 cm. The width of the 10 thou mylar insulation is 22.5 cm and the overall height of the transformer is 25 cm. The resistivity of the impregnating copper sulphate solution was 20 kilohm cm, but could with advantage probably be twice this.

In the early dielectric tests it was shown that with dilute copper sulphate solution grading, the life of 10 thou mylar was greater than 2 x 10^4 shots with 50 kV pulses on it. As the maximum voltage per turn is 15 kV, the pulse life of the mylar in the transformer is 10^8 shots nominally.

The HST showed that the transformer would take a ringing waveform shot at 500 kV and went to 560 kV for a number of shots with the output gap firing. A bit above 500 kV on a ringing shot the small transformer tracked down the outside, under the transformer casing and up to the inside terminal. This happened a number of times at lower voltages, when the transformer leaked $CuSO_4$ solution slightly from the output terminal, owing to overheating. In no case was any damage done to the transformer. The track was initiated from the bottom of the transformer housing to the earthed cable running up to the inner output knob. In the final transformer, a number of perspex tubes of different sizes were sleeved around a rather well insulated cable core which earthed the transformer. In order to provide a support for the transformer, the inner tube on which it was wound was extended down to a plate. If a track were now to occur from the earthy cable as it entered the transformer, it would be likely to punch through this perspex tube at the bottom of the windings. Consequently, a small slot was cut through the tube under the bottom of the transformer, below the output knobs, to allow the track to complete around the transformer.

In practice, the transformer has only been taken to 700 kV on a ringing waveform and as the expected tracking voltage is a little over 1 MV, it is reassuring that no transformer tracking has occurred at any time and hence the safety slot has not been tested.

As was mentioned above, when the HST transformer got very hot in the early tests there was a slight leak from the outside output knob. This is because expansion of the case opens this joint, but not the inner one. In the proper transformer, a made-in-place silicon rubber seal was added between the brass connector and the perspex case, in addition to the other normal attachment procedures.

Prints 3 and 4 show the transformer, print 3 with the cooling freon ducting in place outside the transformer, and print 4 with this removed. The dark bands at the top and bottom of the transformer are the paper spacers which are wound in to aid with the impregnation. In the prints, the nest of perspex tubes surrounding the thick insulated earth connection can be seen. Also, at the base of these can be seen the freon spreader: this reduces the mixing of the freon and the air as the box is filled. Even with this, the flow of freon in should be quite slow to start with and the rate of filling increased as the freon fills the box. It takes about 15 minutes to fill the box fully, without wasting freon by undue mixing.

Electrical Characteristics

Using the dimensions and the standard approximate relations, the transformer's leakage inductance comes out at 35 nH. To this has to be added the capacitor switch and feed inductance of 35 nh, giving a total series inductance of 70 nH. From the high frequency component of the ringing waveform an experimental value of 70 nH was obtained. The primary inductance was calculated as 470 nH, while the value obtained from the slow period of the ringing waveform was 380 nH. The discrepancy of 20 per cent is partly explained by the shorting effect of the feeds, which had to be bent over rather close to the transformer. However, it does seem from this and other transformers that the approximate relation over-estimates the primary inductance for longish transformers by some 15 per cent. This error typically alters the gain calculations by less than 1 per cent. The value taken for calculation purposes is 400 nH.

The series resistance of the bank plus spark gap, etc., deduced from the damping of the waveform, is 40 milliohms, in agreement with an earlier value determined directly. The stray capacity of the transformer is calculated to be 90 pF, of which some 30 pF is available to feed energy into the useful part of the output waveform. The pulse-charged capacitor with strays has a value of 390 pF, giving a total load reflected into the primary of 340 nF, of which 300 nF is useful. The parallel resistance comprises the effects of the inductors at the midpoint of the transformer, the capacitor monitor (32 k), the feed to the output gap irradiator (26 k) and dc/dt effects in the transformer and the transfer capacitor. The parallel effect of all these is about 5000 ohms. This reflected into the primary is 7 ohms. These circuit values are summarised in Figure 4.

The value of the gain without damping is 27.8. Allowing for the effect of the resistances in the circuit, the gain is 25.4. The observed gain was 24.8, averaged over several series of shots. In some test shots the transformer output was loaded with various extra resistors of values down to 1.3 kΩ and the reduced gains observed were in agreement with the calculated ones.

Thermal Effects

The natural convection cooling properties of freon are a factor of about 3 worse than air when gas is pure. However, there is frequently a gradient of density in the freon gas, due to impurities, and this serves to suppress convection cooling further. Thus it is of some importance to keep the gas reasonably pure and even then resistor wattages, etc., have to be derated. It was no surprise, therefore, that the HST transformer heated rapidly when just sitting in freon. A fan (closely modelled on the Wright Brothers' one) was therefore installed to keep the wall temperature of the transformer under control. The freon is sucked up from the floor of the box and flows in a thin annulus over the inside of the transformer until it enters the fan region. It is then passed over the top of the transformer and discharged in a 1 1/2 cm annulus down the outside of the transformer. It is intended in the final box to duct the freon exiting from the bottom of the transformer and to blow it over one of the metal walls of the box, to try and control the rise in ambient gas temperature in the box. In the temporary box, all the walls were of perspex and this was not worth doing. Of course, because blowing gas about the place is so easy, we have had more trouble with this aspect than any component, other than the power pack. On various occasions a nylon bolt fatigued, broke, and the fan fell off; a flexible drive decoupled, some dirt got into a home-made insulated bearing, the perspex drive rod overheated and developed a 45° bend. In the final set-up we hope to do a bit better, but it is rather important to ensure that the cooling gas flow is being maintained, otherwise the transformer will start to boil after about 30 minutes at maximum rating. On the occasion when the fan fell off, it was estimated the temperature at the centre of the HST transformer went to over 80° C. A few small bubbles appeared at the top of the winding (but these disappeared overnight) and the gain dropped about 5 per cent. Apart from these effects, no damage was done, but we would not advise repeating that particular uncontrolled test.

The watts being dissipated in the transformer under ordinary running conditions and under output gap no-fire conditions were measured by running the system with the fan not operating. Small cooling corrections were made to the data and from the thermal capacity the energy dissipated derived. Table II gives the joules per shot dissipated, expressed as a fraction of the energy stored in the capacitors, for both the half-scale transformer and the LARK one.

TABLE II
JOULES DISSIPATED PER SHOT ÷ ENERGY IN BANK

	HST Transformer	LARK Transformer
Normal conditions	3.5%	6.8%
O/P gap no-fire	30%	24%

The ratio of values are roughly consistent with the calculations. For normal running conditions, the volume of mylar stressed is double, and also dc/dt effects in the resistive film grading nearly doubles for the LARK transformer. With regard to the o/p gap no-fire case, extra damping was introduced by the resistors in the HT and earth lead circuits, which led to more rapid damping of the ringing waveform in the LARK experiments compared with the HST. The crudely calculated losses are also in rough agreement with the observed ones and suggest that under normal running conditions the main heating effect is the dc/dt one (~ 75 per cent) and dielectric dissipation and resistive heating in the primary turns are about equal and account for the rest. By increasing the resistivity of the copper sulphate impregnating solution, the energy lost per pulse can be decreased a little, as the transformer seems over-safe as regards turn to turn tracking.

A second conclusion is that as the transformer heats, the grading length increases and the energy dissipated rises and under some conditions this could lead to thermal run away. Table II also shows that the energy dissipation is considerably worse under the no-fire condition and the number of shots in which the output gap fails to fire should be restricted to a few per cent, for this as well as other reasons.

When cooled by fanning, a temperature gradient develops outwards from the primary turn to the walls of the transformer. There is a temperature drop between these and the cooling freon, controllable by the velocity of the freon. Finally, there is the rise in ambient temperature of the freon in the box, again nominally controllable.

The time the temperature drop in the transformer takes to reach internal equilibrium depends weakly on the watts passing through it (because of conduction along the transformer) but is of the order of 20 minutes. The value of this drop was measured by running the transformer at 6 pps and 125 joules per shot on the capacitors for about an hour with the fan on. The rise in wall temperature and ambient temperature were both monitored and at the end of the run the fan was turned off. The excursion of the wall temperature as the transformer thermally equilibrated (time constant ~ 5 minutes) was measured and corrected for natural convection cooling. From this data, the centre-to-wall temperature could be estimated and agreed well with that calculated from the watts deposited and the thermal conductivity of the transformer materials. This equilibrium temperature difference, for running times greater than a half hour or so, is given by

$$\Delta\theta_{\text{centre}} = 1.8 \times 10^{-3}\ W\ ^{\circ}\mathrm{C}$$

where W is the watts into the capacitor bank. Thus, under the nominal maximum running conditions, the centre temperature will be about 36° C over the wall temperature. The temperature difference between the wall and the flowing gas can hopefully be held to + 5° C with improved fanning and it is hoped that the rise in ambient freon temperature can be held to about the same temperature difference. The centre of the transformer will then reach some 70° C, which is a little close to its useable maximum.

For any future transformers, the heat dissipated per shot can probably be halved and, of course, a larger transformer can be built, but it would appear that it will be difficult to raise the delivered power to the load to much more than a couple of kilowatts with a load of about 400 pF, unless some bright idea is forthcoming. Circulation of the impregnant is out of the question, as the liquid films are only 4 thou thick, of course. Heat conduction up metal sheets is pretty useless in any allowed thickness, which leaves only thin cylindrical multiple heat pipes, which are probably useless, for the reasons given earlier with regard to capacitor cooling plus some more when applied to the transformer. If the impregnant could be an insulant, the heat dissipation could be reduced by like a factor of 4 and, indeed, this possibility was looked at very quickly in the very early tests. However, both castor oil and transformer oil were spectacularly poorer than copper sulphate solution, giving lives of a few thousand shots at 20 kV per turn, due to the onset of arcing at the edges of the windings. However, this could be looked at again.

Summarising, the electrical performance of the transformer has been very good, as was expected from the over-tests done in the HST. Its performance closely agrees with calculation and in view of the relatively small energy per pulse, is quite efficient (~ 60 per cent into the full capacitive load). The thermal loading is nearing the danger point at the maximum operating conditions: however, if the output capacity were raised, it would be alright. With the present load of about 400 pF, any future design would appear to be limited to about 2 kW delivered to a load. However, if the capacitive load were increased considerably, about 5 kW might be deliverable at voltages around 600 kV.

PULSE CHARGED CAPACITOR

Physical Description

The pulse charged capacitor consists of 12 stages of an overall height of 20 cm. The dimensions of the metal foil covered compliant blocks which form the spacers between the mylar stage insulant are approximately 45 cm by 30 cm by 1.6 cm thick. The construction of these mildly bendable, slight compressible, blocks has been improved a bit recently, so it is probably just worth describing their present mode of construction.

The basic strength (such as it is) comes from a central sheet of 0.3 cm thick perspex, to which two sheets of 0.6 cm spongerubber are evostuck. These layers provide the compressibility but have rather poor surface finish in fine detail, so the outside surfaces are covered by 5 thou mylar sheets to which have been stuck 2 thou aluminum foil which is polished flat to quite a good finish before these are stuck to the rubber. The two 5 thou mylar sheets are stuck on one aluminum sheet and the continuous edge is arranged to face the output. The other edges of the aluminum are joined by double folding and taping down. Double folding has been shown to make very low resistance joints and can carry large currents for short times and many pulses without sparking apart, so making the output edge unjointed was probably a needless precaution.

The resulting spacer blocks are reasonably sturdy in this size but have to be handled a bit carefully to prevent local denting due to finger pressure and wrinkling due to stretching the aluminum with large radius bending.

The insulant per stage comprises 2 layers of 5 thou, 1 layer of 10 thou, plus 1 layer of 5 thou and 2 layers of 5 thou. The two 5 thou layers are of different widths and are joined by tape along their edges to their mates in the stage above, enclosing the spacer in a double mylar envelope. This is done so that the high fields in the gas, as the mylar leaves the metal covered spacers, is split into a number of electrically isolated zones to inhibit tracking. In the HST, the volts per stage were taken to over 90 kV under freon with no signs of trouble.

The stack sits on a 4 cm high base and is crowned with a 3D contoured HT bung. Figure 5 shows a sketch of these parts and also the clearance to the wall.

The stack is assembled by about 140 lb of lead blocks and is also assembled by electrostatic forces. Inspection of stacks of mylar sheets compressed in such a capacitor shows that over ~ 70 per cent of the area any given pair of sheets is in optical contact, while there are thin air films in the remaining areas. Determination of the mean air film thickness by means of capacity measurement suggests that the eventual average thickness of air per layer is around 2×10^{-4} cm, which is consistent with the above visual observation (but see the section below on the dielectric constant of mylar). The stack takes about a day to assemble and leak out the air, and measurements of the capacity as a function of time and loading are all consistent with calculations involving air films of this general thickness.

During a voltage pulse, corona occurs around the edges of the stage spacers (and partially in the thicker air films within the mylar stack) leading to some freon decomposition products. After a number of pulses, charge deposition has occurred and the corona-ing is considerably reduced, but not eliminated entirely. Judging by the eyes and nose, the freon decomposition products of even prolonged runs are less annoying than the effluent from the spark gaps. Prints 3 and 4 show some details of the capacitor.

Dielectric Constant of Mylar

All the measurements of capacity were made with rather accurate Wayne Kerr capacity bridges operating at a frequency corresponding to $\omega = 10^4$ radians/sec. One bridge was a new one which had been calibrated by our Instrument Section and this compared to better than .1 per cent with our old bridge. In addition, some 1 per cent capacitors were measured and gave the quoted values to about 1/2 per cent. Thus it is not felt likely that the capacity measurements are in error as much as 1 per cent.

The dielectric constant of the mylar we buy from ICI Films Group is quoted in Melinex Bulletin MX 102 as 2.941 at 20° for the frequency of the bridge. In experiments in developing the new method of making spacers, the air film thickness was calculated by measuring the capacity and comparing this with the value expected from a perfectly assembled stack of mylar films. Using the above value for the dielectric constant, negative air thicknesses were fairly easily obtainable. A sub-series of experiments was then instituted to try and find out what dielectric constant seemed to apply for large areas. Various things were found to be necessary to obtain the very highest values. One was that the aluminum foil had to be carefully cleaned and degreased, as some batches had rolling oil on them. Another smallish effect was that the mylar had a thin film of some sort of oil on it, possibly to prevent it sticking to itself in rolls, and this had to be removed. Especially flat and mirror-finished polished aluminum foils on flexible 1 mm thick perspex backings, compressed by a compliant layer with several lead bricks, then gave values up to 3.40 for the dielectric constant. Some measurements with thin layers of water between the mylar sheets and also between these and the electrodes gave values slightly higher at around 3.43. Care had to be taken to make sure the water squeezed out of the electrode assembly was dried off, so that the area of the capacity measured was the area of the electrodes.

This is not the place to go into a detailed discussion of this point and it is always possible that the author making these measurements (JCM) may have made some silly error, but this does not appear to be too likely, as there are some additional rougher measurements which tend to confirm the higher value of the dielectric constant for large areas suggested by these observations. However, in this note a value of 3.40 is taken in determining the capacity of the capacitor at $\omega = 10^4$ with no air films in it. This value is the one used to deduce the mean thickness of the air films quoted in the preceding section.

In order to obtain the capacity at the charging frequency ($\sim$ 1 megacycle), the published curves show a drop in ϵ of 3 per cent. In addition, (a) the internal air films break down; and (b) corona occurs around the edges of the capacitor. As has been explained, charge separation occurs which tends to suppress the corona-ing. Thus in determining the effect of the air films on the pulsed capacity, it is assumed that they only hold some half of the volts that they do at low fields. At low voltages, the air films inside the LARK capacitor account for a 4 per cent drop in capacity and this is taken to halve at higher voltages. Thus in the absence of corona at the edges of the capacitor, the pulsed value of the capacity would be about 1 per cent below the bridge value for our particular set-up. However, edge corona does occur and a crude estimate of the effect of this suggests that the pulsed voltage high speed value may be a per cent or two above that measured by the bridge. In the calculations, the value for the pulsed capacity given is essentially the bridge value plus an estimate of the capacity of the bung to the rest of the box. Not quite all this capacity is available to deliver energy quickly to the load, but this has been ignored, as it is a per cent or two effect.

Mylar Life

The quick and dirty mylar life tests were performed in order to arrive at a working stress in the pulse charged capacitor and to check that the edge corona was not influencing this unduly. Most of the tests were done with 100 cm^2 area samples and four of these were run in parallel, to speed things up. The test pulse voltage had about the right frequency ($\sim$ 1 Meg) and rang on with about a 15 per cent damping factor. The maximum voltage was around 50 kV peak, but most of the tests were done with 30 to 40 kV. In general, there were three layers of 3 thou or four layers of 2 thou and the stressed volume was around 2 cc.

Summarising the findings, before considering some in more detail, it was determined that a There was no observable difference between air and freon.

b There was no statistical tendency for the breaks to appear near the edge of the test sample.

c Intentionally poor metal electrodes did lower the life a bit (presumably because of big air voids).

d For reasonably well assembled samples, the life law was

$$\text{Life} = \left(V_{\substack{\text{single shot}\\ \text{breakdown}}} \div V_{\text{operating}}\right)^{12} .$$

e A few tests under H_2O showed that the life was considerably greater than for gaseous environment operation.

f Around the edges of the electrode spacers there was significant surface marking but no breakdowns and no carbon under freon.

g There were areas of lighter mottling on the mylar sheets, both under electrodes and between sheets. From the shapes of these it appeared that these were areas where there were air films.

h There was some slender evidence that, in general, the breakdown sites were not particularly correlated with air voids when good electrodes were used.

Figure 6 plots the life data obtained in all but the poor electrode shots against a nominal field in the mylar. In these the average air film thickness was calculated from capacity measurements (based on a value of 3.40 for the dielectric constant) and the field on the mylar calculated on the assumption that the air did not break down. This is almost certainly a pessimistic assumption and therefore the real mylar fields are probably a little higher than the nominal ones plotted; however, the error is not large ($<10\%$). The single shot breakdown value of about 3.8-4.0 MV/cm lies midway between the value for thick mylar sheets (3.4 MV/cm) and that applying to 2-3 thou sheets (4.5 MV/cm) plotted in 'Pulse Breakdown of Large Volumes of Mylar in Thin Sheets' (Dielectric Strength Note 14, March 1967). This observation is of dubious certainty but is consistent with the life law of 1/12, which again is about that expected for thick mylar sheets (1/8th) and 2 to 3 thou sheets ($>1/20$th).

It is tentatively suggested by one of the authors (JCM) that these observations, if real, may be a proof of the conjecture on the last page of the above note that 'If the thin sheets of plastic were wrung together very well, so as to exclude effectively all the liquid or air, the avalanche process might well carry on as if the plastic were a monolithic block and the volume dependency come back with a vengeance.' This contention is weakly supported by the fact that with water films (which were an order of magnitude thicker physically, let alone from a mass point of view), considerably better lives were obtained with 10 thou sheets. Also observation f above, that the breakdowns did not seem to occur preferentially where erosion had been going on in air films, implies that optically wrung together mylar could break at the observed fields and lives.

Not too much reliance should be put on the above observations, which are really included to annoy someone enough (IDS?) to prove them wrong.

However, if well wrung together sheets are behaving partly as if they were monolithic blocks, it does make the deduction of the life of a much larger volume of mylar than the test one rather more speculative. The difficulty is that if one takes the lines drawn in the above note for 5 to 10 thou mylar and uses them to obtain a single pulse breakdown field, one might be being optimistic.

Seizing our courage with both (?all) hands and employing the above arguments, a single pulse breakdown field for ~ 1000 cc of mylar is estimated to be around 2.4 MV/cm. At the maximum field to be used in the LARK condenser (700 kV peak volts) this corresponds to 0.63 MV/cm and a mean life of about 10^7 shots results. For the HST condenser, where the fields averaged around 0.84 MV/cm for 3×10^5 shots without breakdown, the expected life would be $\sim 5 \times 10^5$ shots, which is uncomfortably near the number used without any breaks but not inconsistent with it. However, it seems likely, either from field scaling the HST results or from the above calculations, that the pulse life of the condenser is 10^7, or perhaps greater. Since the transfer capacitor is home-made, it is a relatively cheap and quick job to reclaim it, if the mylar does break. It takes about one day to replace the insulation and requires about 4 lb weight of mylar.

If a load could be placed right up against the face of the spacers (which it can't), the inductance would be 3 nH. However, a suitable load could be placed an average of about 8 cm away and then the total inductance would be around 40 nH. The ringing frequency would then be about 40 megacycles, possibly just adequate for a home-made 700 kV capacitor.

Heating Effects

The energy deposited in the corona-ing at the edges of the spacers warms freon, which can circulate. However, the dielectric losses in the mylar, and any energy deposited by corona in the thin air films remaining within the capacitor, leads to heating of this. This has not been directly measured, as it involves surgery on one of the spacers to insert a thermometer and the calculations strongly suggest that the heating is less serious than that occurring in the 'DC' charged capacitors. At the nominal maximum operating conditions, the rate of rise of temperature is calculated to be about 70° C per hour, one-third of the rate in the bank. The cooling time constant of the capacitor is not too easy to estimate, but is of the order of 6 hours. Thus 24 hours rest should suffice to see the temperature back to ambient. In the HTS experiments, six stages were run at 3 pps for 2 hours at a significantly higher stress than the maximum used in the LARK capacitor (equal to 6 pps at full volts on LARK) and no ill effects accrued to the mylar. Indeed the original mylar used in these tests is used in the new capacitor.

All things considered, the transfer capacitor has performed very well and should have an entirely acceptable life, granted it can be easily rebuilt. It also has the desirable property of having its capacity changeable over a significant range of values.

PULSE CHARGED OUTPUT GAP

One of the main objectives from the AWRE group's point of view was to investigate the behaviour of a 600 kV gap at up to 10 pps. Summarising the results obtained, it proved delightfully easy after uv irradiation had been introduced into the gap. That some sort of irradiation would be likely had been foreseen, although, to be honest, it was felt that the start gap would be the one that would need it most. As it turned out, it was desirable to have irradiation for both gaps, in order to keep the energy transfer efficiency as high as possible. However, in some circumstances the output gap could be driven to nearly three times its irradiated breakdown level at atmospheric pressure for a few shots. The output gap design was to some extent dictated by the decision to use air and at a relatively modest pressure, as the working dielectric. Consequently the design of the LARK gap is not necessarily the one that would be used in other circumstances (such as transformer oil).

Physical Description

The LARK output gap body is made of a perspex cylinder 13.3 cm OD and 70 cm long. The wall thickness is 0.6 cm and one end is removeable. As is shown in prints 5 and 6, the body is covered with anti-tracking fins. These were perpetrated by one of the authors (wrong - it was THS) in a positive frenzy of twisted perspex ingenuity. The LARK electrodes are made from 1 3/8" OD copper tube, which is flattened mainly on the back side in a vice (it is a joy to see that the old crafts aren't dying out). Away from the centre the copper tube is also slightly flattened and the last 2" at either end are made of solid brass which is contoured as usual.

The overall length of the electrodes is 40 cm and the maximum thickness of the electrodes in the central region is about 2.9 cm and here the gap is 6.3 cm, which is over half the ID of the containing tube. This is a rather unusually large gap-to-tube ID ratio for a rail gap and it is not easy to ensure, even with two-way contoured electrodes, that the breakdowns will occur in the central region, so some extra field shaping is done outside the spark gap body. Figure 7 shows a couple of cross sections of the gap, indicating how this was done. The field shapers outside the gap were made of aluminum foil-covered plastic foam. This trick was also played in the HST gap and here the outside shielding electrodes overlapped the inside electrodes by a rather larger factor and were poorly contoured at the ends. As a consequence of this and poor simplexing, punch-through occurred a couple of times at the base of the fin nearest the external shield on the pulsed capacitor side and where the fields were greatest. This didn't happen immediately, but after 3×10^4 shots, or so. In both cases when this happened, a track over the outside surface grew over a fair number of shots and then punched through the wall of the spark gap body. In neither case did the gap body disintegrate and it was repaired.

In the LARK gap the contouring of the outside electrode was more carefully designed, the fin spacing decreased on the transfer capacity side and the simplexing much more carefully done. To date, no trouble has been experienced from this cause, but the HST gap experience is a warning that external field shapers need to be carefully designed and can only be used to assist in the control of the field on the internal electrodes, not substitute for contouring these as well as possible.

While the electrodes are symmetrical, for reasons of speed of manufacture, the field is significantly crunched over towards the transfer capacity before the gap fires. Thus a slightly larger gap could have been obtained by making the inside electrodes of two different squashed thicknesses; however, it was not felt worthwhile to do so in this case. Allowing for the asymmetrical field and the contouring of the electrodes, the field enhancement factor (FEF) of the LARK gap was estimated to be 1.47, giving a uniform field equivalent gap $d_{eff} = 4.3$ cm.

The HST gap used 3" OD perspex tubing as the spark gap envelope and had electrodes made out of mildly squashed 1" OD copper tubing. The physical gap was 2.8 cm and the d_{eff} about 2.25 cm. The HST electrodes were also used initially in the LARK gap body, where they looked absurd. In this case the physical gap was 8.5 cm and gave a d_{eff} of approximately 4.9 cm. It was never seriously intended to run the LARK gap with the HST gap electrodes, but it was thought worthwhile to attempt to do so and, indeed, some mildly interesting results were obtained (which will be described later in this section).

Performance of HST Gap

When the 3" OD gap in the HST was first operated, its performance was very close to that predicted using the curves in 'High Speed Breakdown of Pressurised Sulphur Hexafluoride and Air in Nearly Uniform Gaps' (Switching Note 21, February 1973). Also the standard deviation of the pulsed breakdown voltage was about a per cent or two. However, after the first substantial run ($\sim 10^4$ shots), the gap was left overnight and in the morning was tested single shot, and the breakdown curve was up over the previous curve by a roughly constant voltage equal to about 40 per cent of the breakdown at atmospheric pressure. In addition, the jitter was considerably bigger, especially at atmospheric pressure, where it was like ±10%. At higher pressures of about 4 atmospheres, the jitter was more like ±5%. The gap was run for a few hundred shots and left again. On the day following this the effect was even more pronounced. The mean breakdown strength at one atmosphere had almost doubled and the jitter was ±15%. A few breakdowns were recorded which were at voltages 3 times the subsequently measured irradiated breakdown voltage at one atmosphere. When the gap was repped, the breakdown voltage decreased somewhat and the jitter considerably decreased, but was still about ±5% over the operating range.

While the gap had had some 10^4 shots, it is unlikely that all the whiskers which provide initiating electrons had been knocked off by sparks: however, when the gap was left overnight with sparked air in it, it had rotted the whiskers good and proper. When, after a few hundred further shots during which the gap was not flushed and was left sealed up over night, the process went even further. Whether this is a particular property of copper is not clear, but there was a tendency for the breakdowns to move towards the end brass caps, which might suggest that copper may be particularly prone to whisker rotting.

The solution to the problem was obvious: some sort of uv or electron irradiation was needed to start the streamers out as soon as the Townsend multiplying condition was reached. A few moments' thought provided at least six different solutions. What then ensued was most frustrating. With this number of possible solutions, at least one and usually two or three would normally have proved to be easily realisable in practice. As it turned out, they all suffered from sometimes subtle practical disadvantages and it was only after a series of fruitless attempts to get round these difficulties that one of the authors (wrong again - it was THS) suggested decoupling both the uv irradiator electrodes from the existing main electrodes, and a practical solution resulted.

One of the earlier attempts is worth recording, because of its possible incidental implications. This evolved from an attempt to provide a uv irradiator externally to the gap body. It was realised that the perspex wall would probably prevent sufficiently hard uv penetrating to the electrodes, so it was somewhat surprising when it appeared to work. It was only after some further investigation that it was realised that it was the presence of the body of the irradiator at earth potential away from the local spark gap body potential that was producing the effect. It was then found that a length of coax cable inner attached to the earthy electrode, taken over the adjacent fence, run along the outside of the gap body and returned to the earthy electrode, could reduce the breakdown voltage and the jitter. In addition, the breakdowns occurred between the electrodes, under the region where the insulated wire ran. Since no external corona was evident in the insulating freon, and in any case it was highly unlikely that uv of a suitably short wavelength could penetrate the perspex wall, it was deduced that incipient wall discharges were occurring inside the spark gap wall and uv from these was irradiating the gap. A run of a few thou shots was done with this arrangement and, sure enough, discharge markings occurred on the inside wall of the spark gap body, under where the external earthy insulated cable ran. These markings were across the body and were about 3 cm long. Thus wall discharges were happening during the pulse charging but before they could complete, the main gap closed. Now, admittedly the gap wall field had been significantly perturbed by the added auxiliary electrode, but the fields estimated to be present were scarcely large enough to lead to electron emission from the plastic and avalanching. Consequently, it is possible that the effect was due to the deposition of electrons, or, less plausibly, negative ions, during the preceding spark. These deposited charges can add to the local fields produced by the auxiliary electrode. The charges are driven out because after the spark channel closes, it and both electrodes are at a high potential with regard to the surrounding earth of the system. Indeed, it is a good question why, in a wide electrode spaced gap such as the present one, the arc channel doesn't initiate a spark travelling outwards through the gap body and down to earth. Certainly if the spark channel were replaced by a wire of the same diameter, such an event would be very likely to occur. The explanation is left as an exercise for the student. However, if the above tentative arguments have any validity, it implies that charges can easily be deposited on the inside wall of a spark gap body and if the resistance of this is high, the fields on subsequent pulses can be very far from those calculated using Laplace and only considering the metal bits. It is believed that a rash of drop-outs (very low voltage, long path, breakdowns) in a DC-charged rail gap Marx which has been recently observed at AWRE is connected with wall charging phenomena.

As a uv irradiator, the simple solution described above suffered from the defect that it became less effective as time passed and also did not reduce the jitter of the gap to the levels we required. It only improved the very bad situation after whisker rotting had taken place: it wasn't a really satisfactory solution. This was quite apart from worries about what would happen in the full-scale gap at about twice the voltage, not to mention the thought that on some shot the incipient wall track would complete over an erosion-weakened plastic surface.

UV Irradiator

Once it had been suggested that the uv irradiator be a small entirely separate gap, rapid progress was made. Figure 8 shows a sketch of the arrangement. A length of Tekmatic razor blade was mounted a short distance from the rounded edge of a brass sheet, to make an edge-plane gap. A fast rising pulse was applied to this sub-gap, which had a small capacity built into it. When the razor edge was positive, the breakdown voltage was slightly lower than with a negative edge, but more erratic. Consequently the negative edge polarity was selected and a gap of 0.1 cm was chosen as being suitable, after some tests. This broke at 11 kV in air at atmospheric pressure. Some quick life tests were done with a 1 cm length of razor blade, when 50,000 shots were put on it at 20 pps without changing its breakdown voltage. As the proper system has about 10 cm of blade in it, this corresponds to about half a million shots life test. The irradiator gap is mounted inside the earthy electrode of the gap and is fed by a 26 kilohm chain from the negatively pulse charged electrode. A series of holes is drilled in the earthy electrode (which has a reduced field on it compared with the other one) and the edges of these dented in slightly. The irradiator razor blade is cut away, except opposite these holes.

This arrangement has operated most satisfactorily and after 100,000 shots in the HST gap was removed and retested on the bench, where its breakdown voltage was found to have increased by less than 10 per cent. Unfortunately, it is not easy to monitor its breakdown voltage as a function of pressure in situ in the gap, but the breakdown voltage of the irradiator should go as $p^{2/3}$. The little sparks produced by it are quite bright and obviously are very luminous, since the light output only lasts a fraction of a microsecond. The capacity across the irradiator and the value of the charging resistance are arranged so that the irradiating spark occurs well before the voltage on the main gap reaches half its breakdown value.

It can be checked that the uv irradiator is working by observing the location of the sparks which cluster around but are slightly outside the holes. The sparks are also much straighter than those occurring without the irradiator. The jitter of the irradiated gap is less than 1 per cent of the breakdown voltage and occasionally two-channel closure was observed. The breakdown voltage curve against pressure has remained constant and agrees well with the curves given in the 'High Speed Breakdown of Nearly Uniform Gaps' note, both for air and SF_6 (at low pressures). It may be wondered why this should be so, since the gaps in that work were not irradiated. The reason is that in those experiments the electrodes were of brass, of large area, and quite rough in finish, and each gap probably had no more than 500 shots on it during the period of the measurements. Thus there were always plently of whiskers to provide electron emission before the avalanching field was reached.

Performance of LARK Gap

When the LARK gap was first tested, the HST electrodes were installed in it. The aspect ratio (gap to electrode thickness) was ludicrously big and even with uv irradiation the sparks fairly frequently closed at the ends of the rail electrodes at atmospheric pressure. As the pressure was raised, they completely transferred to the ends of the electrodes. In addition, the breakdown voltage curved over and failed to increase, as calculated, with pressure. The interesting observation was that there was considerable 'fizzle'. 'Fizzle' is the spark gap terminology for passage of significant current before plasma channel closure. (Laser people should read 'high E/p uniform discharge pumping' for fizzle.) Even at 3 atmospheres absolute air pressure, some 1/2 kA would flow for 40 ns or more, and this not in an edge plane gap!

In view of the observations, we were encouraged in our original intent and the new larger electrodes were made. The LARK gap with these electrodes in it performs closely as calculated and has a jitter of less than 1 per cent. There are still slight signs of fizzle, but it amounts to only a very small effect at the operating pressures and is probably only visible because of the rather high value of the output load (200 ohms). Figure 9 shows the breakdown curve for the gap operating at peak volts, where the 89 per cent time width is about 70 ns.

LOAD AND OUTPUT MONITORS

Load

As can be seen in prints 5 and 6, the copper sulphate load comprises two parallel 1" bore 9" long tubes. These are joined at the HT end by a disc of solution, so that flow through them can occur. External to the box, 1" plastic piping runs vertically about 2 feet to a finned radiator. This is forced-air cooled by an office fan very kindly loaned to us by Mrs Vikki Horne. (Vikki became mildly alarmed when typing the earlier portions of this note and the authors would like to assure her that the remarks that 'the fan fell off and the drive shaft bent through 45°' referred to another fan. Her fan behaved in an exemplary manner and has now been returned intact, with many thanks.)

The solution in the load thermosyphons and its temperature rise in equilibrium is given approximately by

$$\Delta\theta_{CuSO_4} = 0.5\ W^{1/2}$$

where W is the watts into the storage capacitors.

Capacity Monitor

Partly to avoid wasting energy and partly because of heat dissipation requirements, the monitors are of rather high impedance. They are made from chains of old-fashioned 1 watt carbon composition resistors and are carefully located in the best possible fringing field distribution, à la 'Nanosecond Pulse Techniques'. The capacitor chain consists of 68 resistors of 470 ohms value, giving a measured overall impedance of 31.3 kΩ. The tap-off pad at the bottom is arranged so that the pulse volts and wattage dissipation in the bottom resistor (which has a value of 725 ohms) are both within 15 per cent of the values applying to the rest in the chain. The response of the monitor to a unit step pulse is estimated to be a fast rise to half amplitude and then a climb up with a time constant of about 40 ns. This leads to an under-estimate of about 2 per cent when monitoring the pulse charged capacitor ringing waveform, where $\sqrt{LC}$ = 115 ns. Including this factor, the overall attenuation ratio with the tap-off pad is 1300 and, with the 'scope sensitivity of 267 volts/cm, gives an overall value of 347 kV/cm. The monitor can be seen best in print 5.

Output Monitor

This consists of a chain of 48 resistors, each of 220 ohms, giving an overall resistance of 10.5 kΩ. This resistor is graded by being located between the two limbs of the copper sulphate load. The attenuation factor of this monitor is 970 and the overall monitor channel sensitivity is 257 kV/cm. The time response of this monitor is estimated to be a little less than 2 ns.

Both of these monitors were compared with each other and with an earlier monitor used in the HST and internal agreement to within a couple of per cent obtained. After the test series on LARK, the monitors and the 'scope were recalibrated and answers within a couple of per cent of the previous ones confirmed.

As regards their high voltage performance using the values obtained with them, the gain of the transformer remained constant at all levels up to

700 kV. The old HST monitor in air displayed some super-linearity, a point referred to again in a later section on system flashover results.

The 'scope response was determined and consists of approximately a 2 ns time constant rise to about 80 per cent, followed by a rise time for the rest of the pulse of about 10 ns, due to the effects of a signal delay cable.

Apart from one unusual test, where the load resistors were emptied of solution, no monitor flashover occurred and the temperature rise in them at about 40 per cent of maximum watts was satisfactorily low.

OUTPUT BUNG AND DIODE GRADING

The output bung design is rather more complicated than is usual, since not only does it have a major effect on the diode grading, but it also has to be designed so that the field on it will not lead to flashover during the 600 kV 80 ns decay pulse. The overall shape can be seen from prints 5 and 6 and really defies drawing in any simple way. Together with the field shaper on the output face of the box, the diode is uniformly graded to within ± 20 per cent for cathode stalk diameters between 1 and 2 inches in diameter. As would be expected with the 2" stalk, the maximum field is at the first (earthy) stage. With the 1" stalk the field is maximum three-quarters of the way up the diode. The diode design consists of eight 1" perspex stages and is 9" long and overall 12" OD. The hole inside the aluminum plates is 6" OD.

At the full output volts (where the output bung rang up briefly to 750 kV), it did not track with a 250 ohm load on the generator.

L/R PREPULSE AND WAVE TAIL CONTROL

The output circuit values are shown in Figure 10. In the absence of any L/R circuit and with a 10 kΩ load, a prepulse of 26 per cent of the voltage on the transfer capacity was measured, where the calculated value was 25 per cent. The addition of a quickly made 30 μH inductor in series with a 44 ohm load reduced this to an observed level of ~ 1 1/2 per cent. The calculated level was 3 per cent. This is not a practically significant difference, but is by far the worst disagreement between calculation and observation in the whole series of tests. The inductor (not the final version) can be seen in print 6.

This series combination also has an effect on the fall of the voltage across the load and the later reverse voltage tail. The section on output waveforms discusses these effects.

In the proper application, the behaviour of the cathode will not be like a constant ohmic impedance and the values of the L/R combination will be adjusted in the light of the cathode impedance history. However, it is quite likely that the impedance of the cathode will rise as the volts fall during the tail of the negative pulse, so significant energy is likely to be dumped into the resistance and this will therefore be one using a thermosyphoned copper sulphate solution. It has been installed in the proper box and consists of a vertical tube running from top to bottom of the box with the hot electrode midway up. It will normally be connected to the forced air cooled radiator, but can be disconnected from this when the copper sulphate load monitor cell is in use instead of the diode.

This completes the description of the major components and their performance so far. The remaining sections deal with the system performance and outline some possible developments.

FLASHOVER RESULTS

There was one obvious change in going from the HST to LARK which gave rise to thought and some worry. This was that while the HST just sat on a bench, either in air or freon, LARK would be enclosed in a metal box. Thus while the HST results were very encouraging, they were not really representative in respect of flashover levels. Flashover during normal operation with the output gap firing was not a worry; indeed it is considered that the flashover level in this case is around the 1 MV level for LARK. However, when the output fails to fire (nearly always because of a slightly low start gap firing), the system rings on and late time flashover becomes much more likely.

One odd thing that had been observed in the HST with the system in air was that the old monitor (which was cast in simplex) started to go superlinear. This effect set in fairly sharply and was tentatively attributed to the top of it getting bathed in fizzle from the transfer capacitor bung. In the HST, the effect went away when the system was run under freon and did not reappear at the highest levels to which the system was tested.

When LARK was assembled up to the transfer capacitor, but without the output gap, the same effect was found with the old resistor at a level corresponding to about 400 kV. In addition, small incipient streamers were seen moving out from the sides of the transfer capacitor. Moreover, the slow component of the ringing waveform was considerably more damped than at low levels. This damping corresponded to an added load of about 800 ohms, obviously due to electrical 'pumping' of the air. Normally, for few μsecond time scale systems, the introduction of freon raises the breakdown strength by a factor of about 2.6. If the same sort of factor applied to LARK, it was obviously very safe against flashover, even in a no-fire case. However, when a roughly 30 per cent freon mixture was used in the box, the system flashed over at 440 kV, late on in the ringing waveform. Some quick tests were done with reduced spacings and 100 per cent freon, as well as roughly 35 per cent freon, and the following explanation emerged.

For long pulses (fast rise 25 μsec tail) in wide gaps, work done a fair time ago had established that air/freon mixtures gave breakdown fields proportional to the mixture ratio (in distinction to SF_6). Figure 11 plots this data, which is reasonably accurate. The results from the quick and rough measurements of the late time breakdown in air and freon mixtures are also plotted. The results are rather rough as they depend in part on the FEF for 2D curved surfaces and also the percentage of freon was not at all well established. However, they tend to support the contention made before the results were obtained as to the shape of the curve for faster pulses. The warning is that for fast systems, the breakdown improvement factor practically applying when the air is replaced by freon is not 2.6, but can be significantly lower, and in this case was more like 1.7. This is not because the freon has got weaker; it is because the air is considerably stronger. It raises the vaguely amusing thought that if the same thing happens with SF_6, there is some fast pulse waveform where the breakdown strength is linearly related to the mixture ratio. It also suggested that under pure freon the late time no-fire flashover voltage would be around 700 kV, for a 6" spacing between the bung and the box. The temporary and final box widths were both increased by 4" overall and some ways of raising the sparking voltage further were investigated. The flashovers occurred along the length of the sides of the bung, but tended to concentrate at the 3D rounded ends. Mylar flash guards raised the breakdown level by about 10 per cent and then the sparks only occurred at the rounded ends of the capacitor.

When the output gap was installed and tested under one-third freon, there were late time tracks, in the no-fire case, around the removeable end, at around 400 kV. Some of these involved tracks running 2 feet up the outside of an air feed pipe. As the removeable end plate of the output gap was held on by brass bolts, these were replaced with nylon bolts and the flanges cleaned up a bit. In addition a pair of elbow pad shaped perspex sheets were installed at either end of the spark gap. These can be seen in print 5. At this point it was necessary to press on with other tests and also our nerve failed and we discontinued the one-third freon tests.

After the other tests were completed and the IC power pack installed, the system was taken up to 29 kV on the capacitors, with the output gap both firing and not firing. At 700 kV under pure freon, in the no-fire case, the system tracked at late times near, but not at, the site of the new capacitor voltage monitor. This had had some extra mylar placed around the capacitor bung, in the region where the resistor chain was attached. It did not track at the elbow pad covered 3D regions, which was satisfying. The late time tracks ran down from the underside of the capacitor bung to the earth plane under the centre of the output gap, very closely following the line of the capacitor monitor, which was also satisfying. With mylar shielding all along the output gap edge of the capacitor bung, it is estimated that the late time no-fire flashover can be raised to about 750 kV, at which voltage tracking should start at the other unshielded corners of the bung. However, with the results of the cathode studies to hand, it might well be decided to raise the value of the pulse charged capacity by about 20 per cent and in this case the present system will operate without any further changes, at the maximum levels that the heating effects allow. Another possibility is to instal an edge plane safety gap in series with a damping resistor, so that in the event of a slightly low firing of the start gap, the system won't ring on. Under these conditions (a sort of ensured fire case) the flashover voltage could well be around 900 kV.

MAXIMUM RUNNING CONDITIONS, SO FAR

LARK's longest run so far is 65 minutes at 6 pps with 125 joules per pulse stored in the 'DC' charged capacitors. This, of course, is significantly below the objectives. However, the run was performed without trouble, within a few days of the system being finally assembled. In the test, the voltage on the capacitors was about 23 kV and the capacity was pulse charged to 560 kV. The load was 250 ohms, dropping to around 180 ohms during the run. The output voltage is complex, as will be explained in the next section, but briefly the output pulse reached a little over 600 kV and by the time the voltage had dropped to 300 kV at 50 ns, some 40 joules of 'useful' energy had been delivered per pulse. The total energy delivered per pulse to the output load was about 80 joules and this was absorbing 1/2 kW.

During this run, the transformer of the power pack ran significantly hot and the pack circuitry will have to be improved before the output watts can be raised, and indeed, to meet the objectives, constant current charging will have to be arranged.

Despite the relatively modest performance so far demonstrated, the results with the HST set-up and the thermal measurements and breakdown data achieved to date lead us to believe that the design objectives can be fairly closely approached. However, until we have some details as to the diode characteristics, the exact component values and operating conditions cannot be specified and it seems more sensible to move onto these investigations as soon as LARK is installed in its proper box, before pushing it up to its final levels.

OUTPUT WAVEFORMS

Figures 12 and 13 show the output waveform for a load of 250 ohms, and Figures 15 and 16 show the same for an output load of 125 ohms. The waveforms are shown with the output volts normalised to the capacitor volts at the time of output gap closure. The nominal gain through the system is experimentally 24.8, but in practice the output gap would close a few per cent before peak capacitor volts; thus a gain of 24 practically applies. The details of the reversed polarity tail depend fairly strongly on the exact phasing of the output gap closure time and that of the peak of the charging waveform. The long time tail curves given in Figures 13 and 15 are for closure at peak volts. In actual operating conditions, the amplitudes of tail would be significantly lower (~ 20 per cent less) and smoother.

Concentrating on the important fast negative going phase of the output waveform (Figures 12 and 14), these show a damped oscillation on the front. This is the transfer capacitor ringing into the capacity of the bung and load (25 pF) via the inductance of the spark gap and the output feed (250 nH). This gives a $\sqrt{LC}$ of 2.5 ns and an impedance of 100 ohms. Both of these agree well with the 'scope and monitor response corrected waveforms. Most of the energy to heat the spark channel comes from the energy stored in the capacity of the output gap (~ 8 pF), but the first half-cycle amplitude as corrected is lower than would be the case with a perfect switch, so it is likely that some energy is still being taken by the forming spark channel. In addition there are some EMP-like velocity of light transit effects which affect the amplitude of the first peak.

The 'scope records shown in Figures 12 and 14 have been corrected for 'scope and monitor response in two stages. The first and more reliable way is to take the approximate smoothed curve, ignoring the output bung oscillations and to correct this. When this is done the theoretical early decay time is obtained within 10 per cent (which gives confidence that the unfolding of the recording response has been correctly done). An attempt has been made to correct the oscillations for recording effects, but this is, of course, significantly more uncertain; however, the corrected waveform is believed to be on the pessimistic side and the amplitude of the curves obtained seem reasonable.

The back extrapolation of the corrected smoothed output is within a few per cent of the level expected after capacity sharing with bung output capacity and allowing for the observed fizzle.

The rise time of the gap is a complicated question, but as applied to the smoothed output pulse the resistive phase time is like 1.7 ns and L/R like 1 ns, for Figure 12, giving a maximum slope rise time of about 2.7 ns for this case.

The observed curves shown are slightly smoothed versions of the original records which contain ~ 5 per cent amplitude additional ripples, due in part, it is believed, to a small mismatch in the 'scope.

Also shown on the figures are the waveforms recorded with the L/R combination in circuit. These curves are essentially identical for the oscillating portion of the waveform and have been omitted in that region, in a desperate attempt to simplify already overcomplicated graphs.

However, the effect of the L/R combination is to reduce the duration of the tail significantly at the expense of a couple of per cent of the energy delivered during the useful phase of the pulse.

From a practical point of view, the amplitude and phase of the bung oscillation can be changed by adding inductance between it and the output gap: somewhat better results can be achieved by adding an inductance in parallel with a resistance (the clipped Marx). However, the effect of these changes can only be worked out when the impedance history of the cathode is known.

In the case of the 125 ohm load, the smoothed curve gives an output of 600 kV for 29 kV on the storage capacitor and the pulse width will be around 30 ns, by which time about 50 joules of useful energy will have been delivered. However, if a cathode of impedance as low as this is used, it will be energetically desirable to increase the value of the transfer capacity as far as possible, in which case 60 joules per pulse of useful energy may be obtainable. However, again the cathode impedance history must be known before reliable estimates can be made.

Summarising, the output pulse seems to be behaving as would be expected with a small, quite usual, shortfall and providing the jitter on the start gap can be reliably held to 1 per cent, the 'useful' energy delivered per pulse can be expected to approach those aimed at in the objectives with a 200 ohm cathode, and be not too much less at 100 ohms.

ENERGY BALANCE

Doing an energy balance on the output is one of the more testing calculations that can be done. Equally the accuracies achievable are normally at best 10 per cent. However, it seemed worthwhile to make the effort in this case. Table III gives the energy determined by integrating the response corrected smoothed waveforms in Figures 12-15 as fractions of the energy stores in the capacitor. The various energies are E_{fast} and E_{slow}, the energies in the negative and positive portions of the waveform (without the L/R on, of course). E_{tot}, the sum of these is the fraction of the stored capacitor energy delivered to the load. E_{useful} is 70 per cent of E_{fast} and represents the energy delivered by the time the voltage has fallen to 60 per cent of the smoothed peak value.

TABLE III

Load Value	E_{fast}	E_{slow}	E_{tot}	E_{useful}
250 ohms	42	24	66	30
125 ohms	47	14	62	33

While some of the difference between the figures for the two load resistances is due to accumulative errors, it is reasonable that E_{fast} should be larger for the 125 ohms load, as the energy is extracted more quickly from the transfer capacitor and hence less returns to the transformer circuit. Also it is very reasonable that there is less energy in the long tail, for the same reason. If the high speed pulse oscillations had been included in the integration, the E_{fast} values would have been raised by about 2 per cent.

Taking the average of the last two columns, the mean E_{tot} is 64 per cent, and the mean E_{useful} is 32 per cent.

Table IV lists the losses measured or estimated in the various components. In a few cases the estimates are rather rough ones, but fortunately most of these items are fairly small and the larger terms are reasonably well determined.

TABLE IV
ENERGIES EXPRESSED AS FRACTIONS OF THE INITIAL STORED ENERGY

	Per Cent
Losses:	
'DC' charged capacitor	6
Start gaps and feeds	~6
Transformer	7
Bank:	
Inductors and resistor chains	4
O/P Gap capacity and Bank capacity	2
Pulse charged capacitor	~3
Fizzle (1% - 3%)	2
Charging bung capacity	6
	36
Delivered to load:	64
	100

So help us, that is the way it came out. We seriously thought of altering the numbers so that it didn't happen, but that seemed to be unscientific, too. Needless to say, the 100 per cent value is entirely fortuitous, agreement to ±10 per cent would have been considered good going. Anyway, providing all the losses have been considered, it would appear that the observations are in good agreement with the energy available. Also it can be seen that to improve the situation significantly, improvements would have to be made in a large number of separate components and therefore a lot of work would be involved.

FUTURE DEVELOPMENTS

As has been mentioned earlier, it may well be worthwhile increasing the capacity of the transfer capacitor, as far as this can be done with due regard to the life of the mylar.

In a new system, with a modestly enlarged box, the transfer capacity could be increased to 2 nF by placing a second one above the present one and increasing the linear dimensions of the spacers by 60 per cent each. The output gap would be made wider and with tailoring of the box near the output end, it is considered a 4-5 ns rise time could still be maintained with a 40 ohm load. Either three feed through diodes would be needed or a single racetrack multi-insulator tube, like WEB, could be used, in which case the pulse could then be fed to the cathode without significant additional degradation of the rise time. With 700 kV on a 2 nF transfer capacitor, 500 joules per pulse would be stored in it and a DC charged store of some 800 joules should suffice, as the energy transfer would be better with the use of a physically larger transformer. This would be about twice the diameter of the present one and use higher resistance impregnating solution. Under these conditions it is thought that the thermal effects in it would be controllable at 10 pps at the above levels. Thus it is possible to consider a 40 ohm load system delivering 5 kW to the load, of which some 3 kW would be useful.

The above speculations should not be taken too seriously, but are only included because one of the authors (yes, JCM) cannot afford to live in the present and has always to be thinking of the next system on, or preferably the one after that.

CONCLUSIONS

The design, construction and preliminary testing of a modest repetitive pulser are described. Some of the limitations to repetitive pulsed systems are outlined on the debit side, while the relative lack of trouble with the major components is described on the credit side.

ACKNOWLEDGMENTS

The authors would again like to express their gratitude to Professor Dan Bradley and Dr Henry Hutchinson for their encouragement and help, and in conclusion the authors would like to thank any readers who have reached this far.

....."Up with the lark and to bed with a wren[‡]"
("Sweetest and Lowest" - a wartime revue)

[‡] US terminology: wren ≡ wave

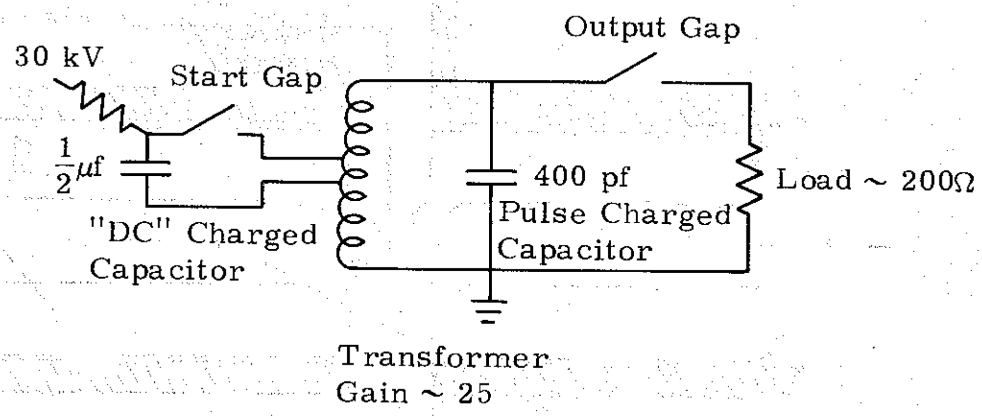

Figure 1. Schematic and Approximate Circuit Values

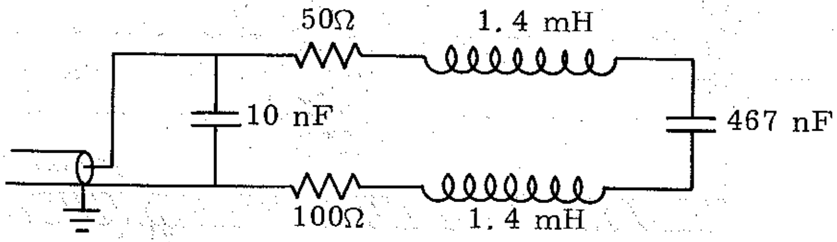

Figure 2. "DC" Charged Capacitor Feed Circuit

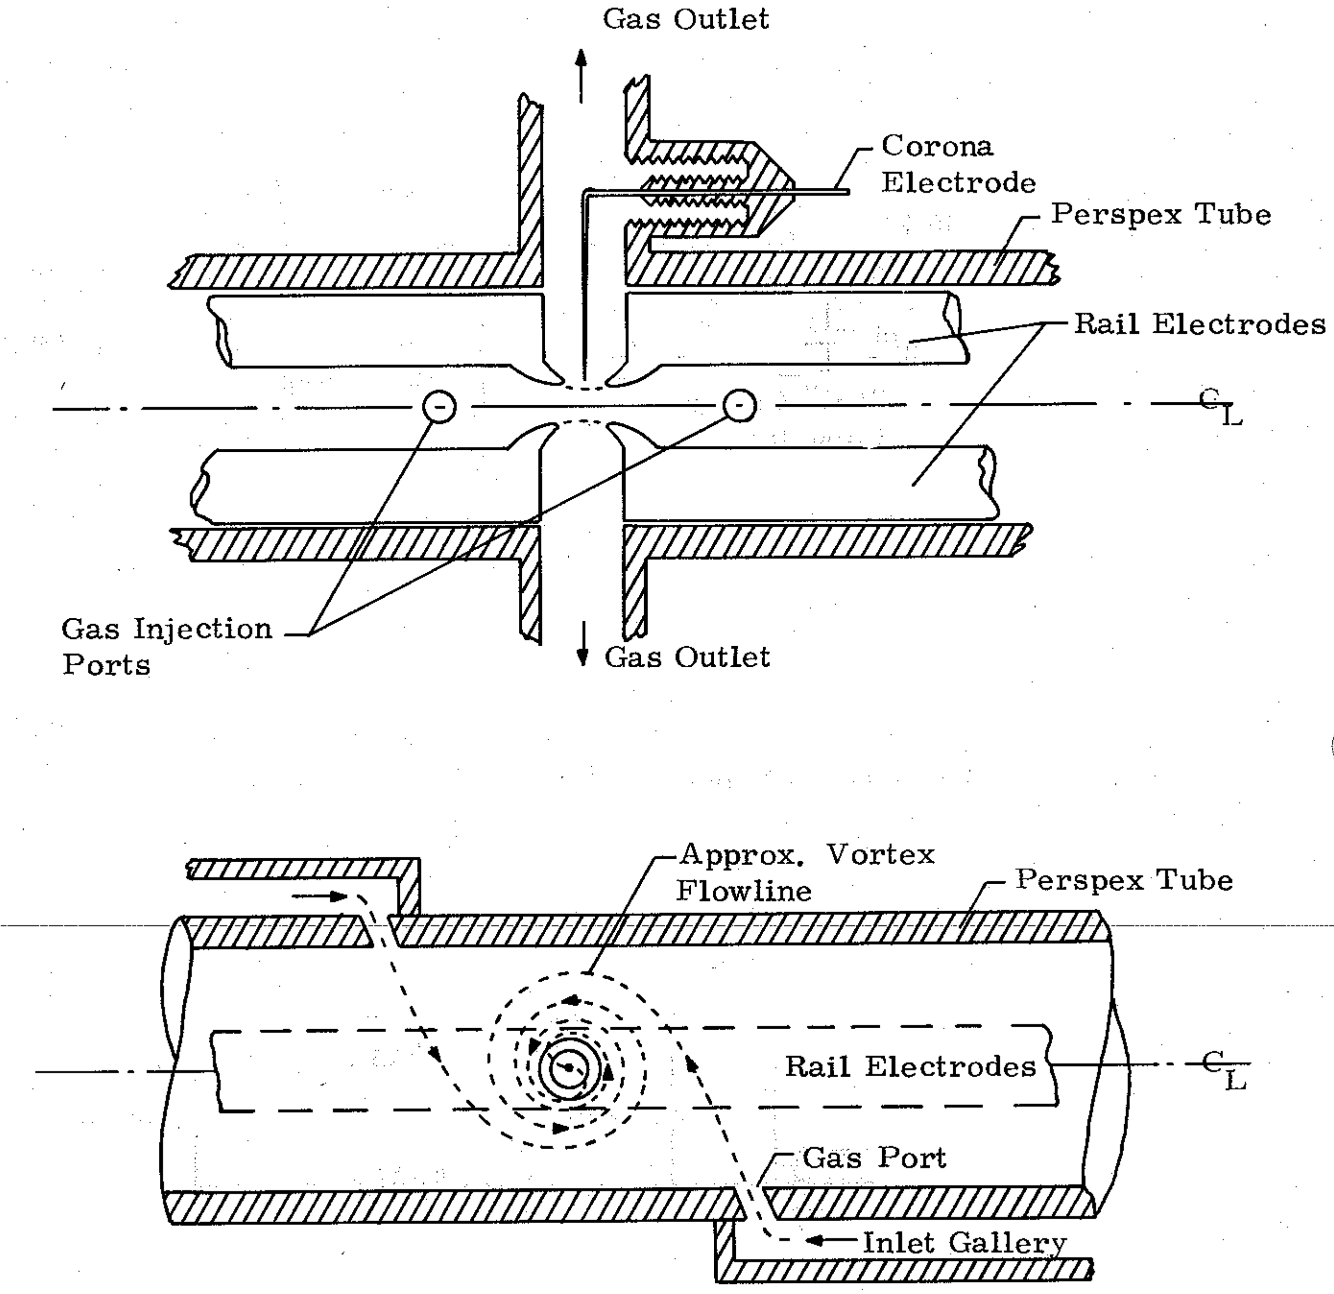

Figure 3. Start Switch

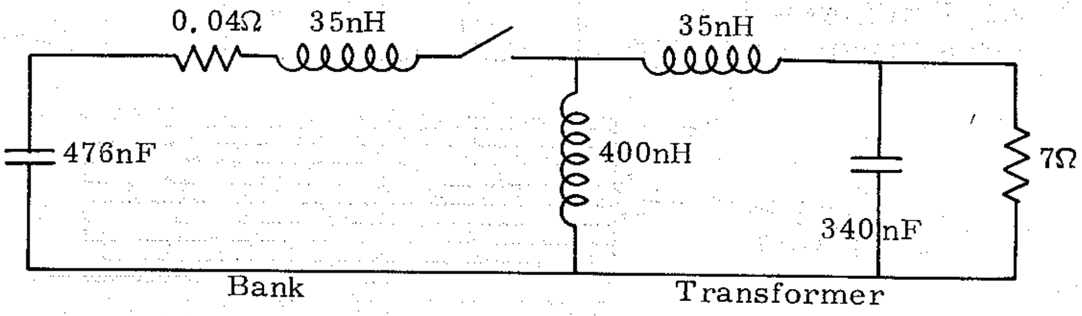

$\sqrt{LC}_{fast}$ = 117ns

$\sqrt{LC}_{slow}$ ~ 580ns

$\sqrt{L/C}_{fast}$ = 0.60Ω

$\sqrt{L/C}_{slow}$ ~ 0.72Ω

Gain = 25.5 Calc

= 24.8 Observed

Figure 4. Transformer Circuit Values Referred to Primary

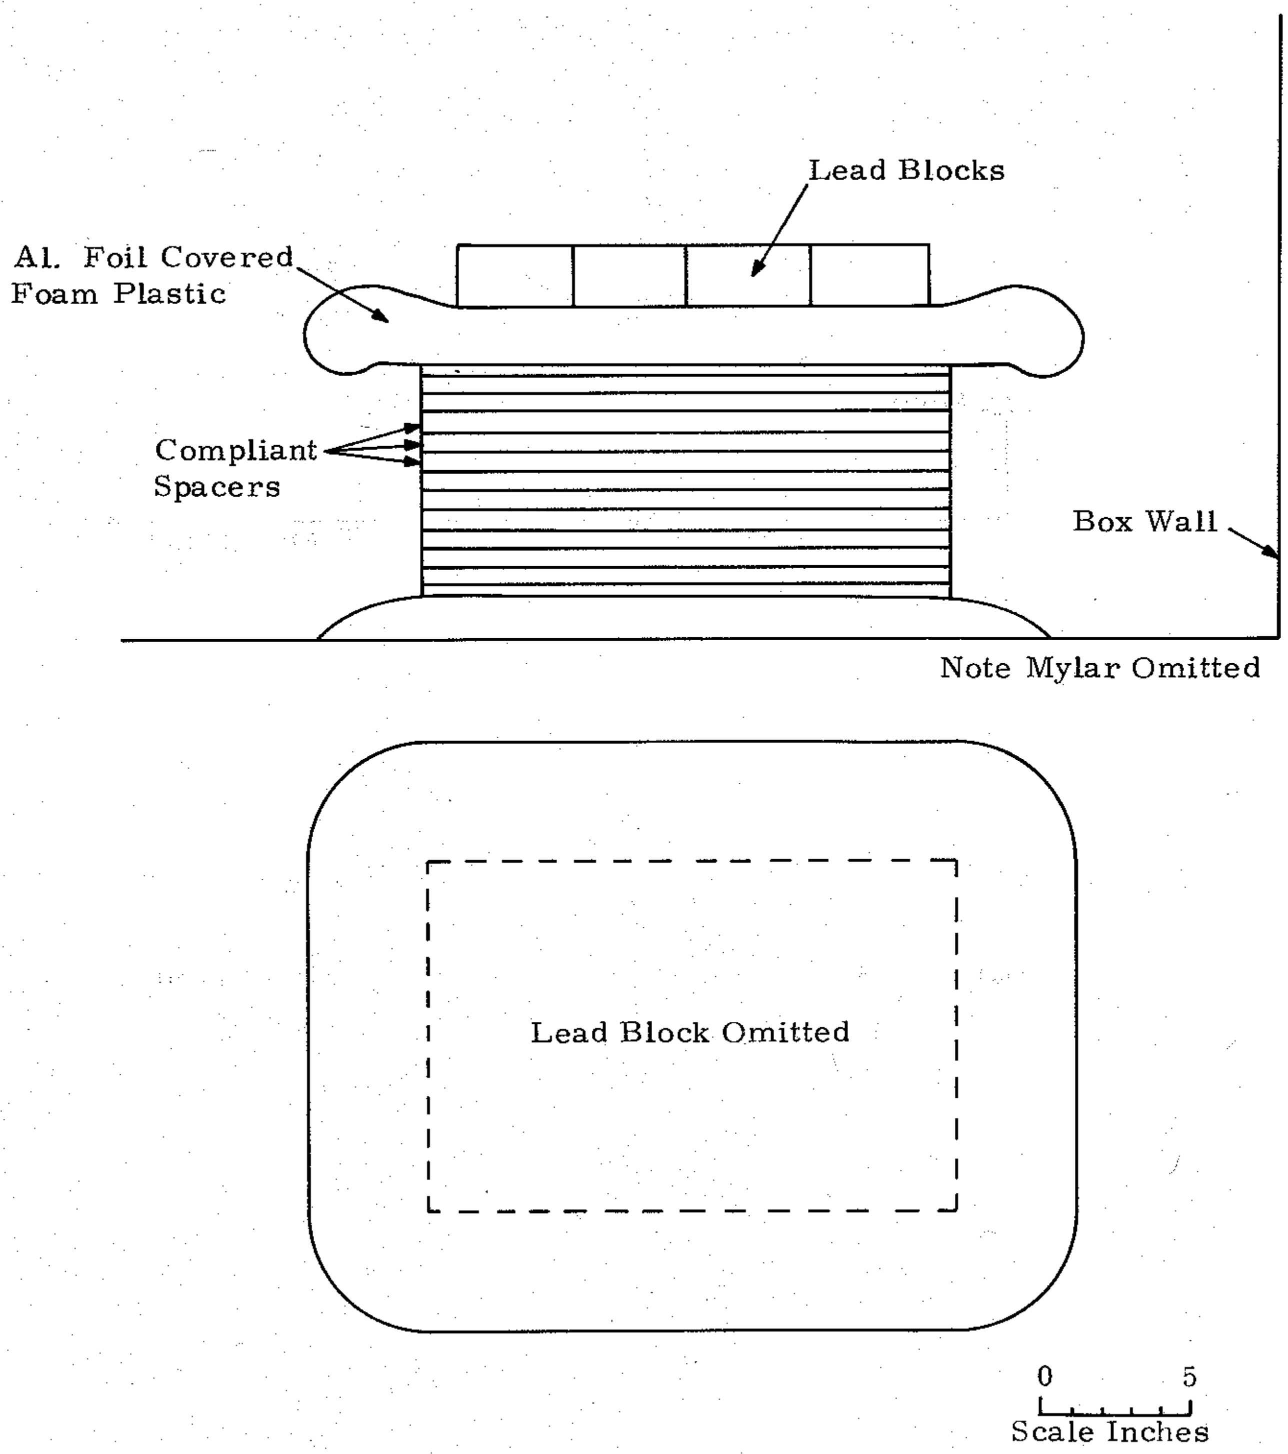

Figure 5. Pulse Charged Transfer Capacitor

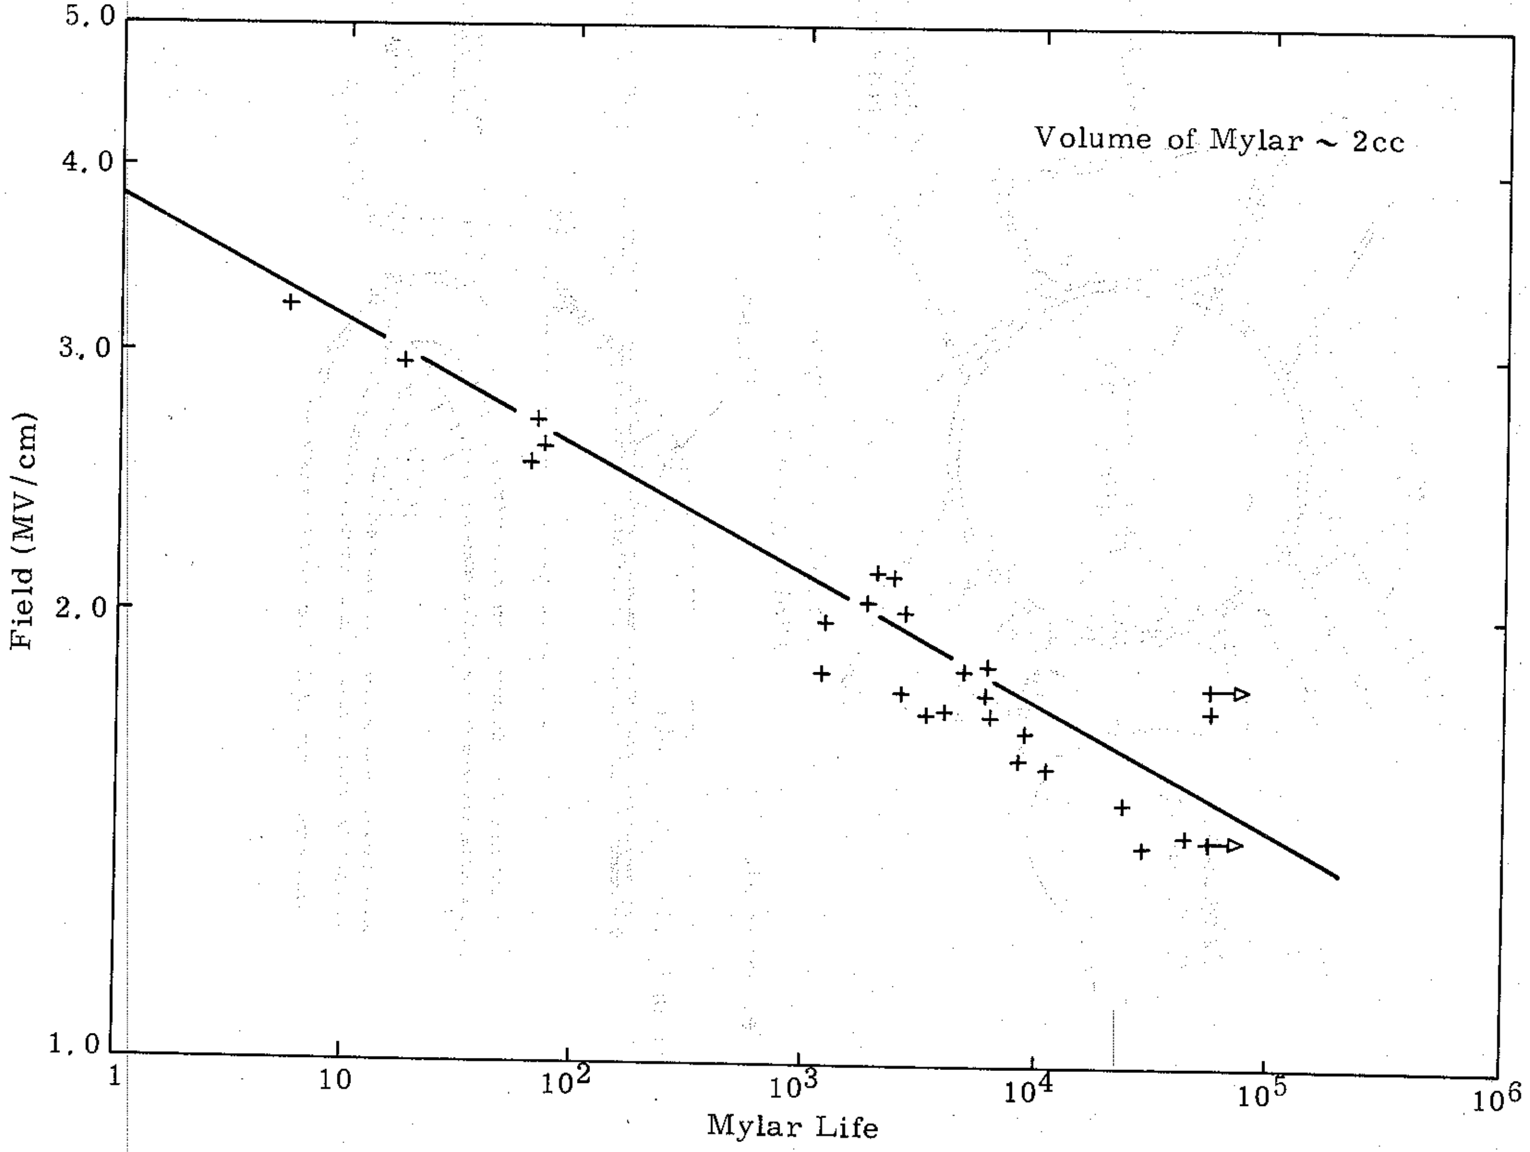
Volume of Mylar ~ 2cc
5.0
4.0
3.0
2.0
1.0
Field (MV/cm)
1
10
10^2
10^3
10^4
10^5
10^6
Mylar Life

Figure 6

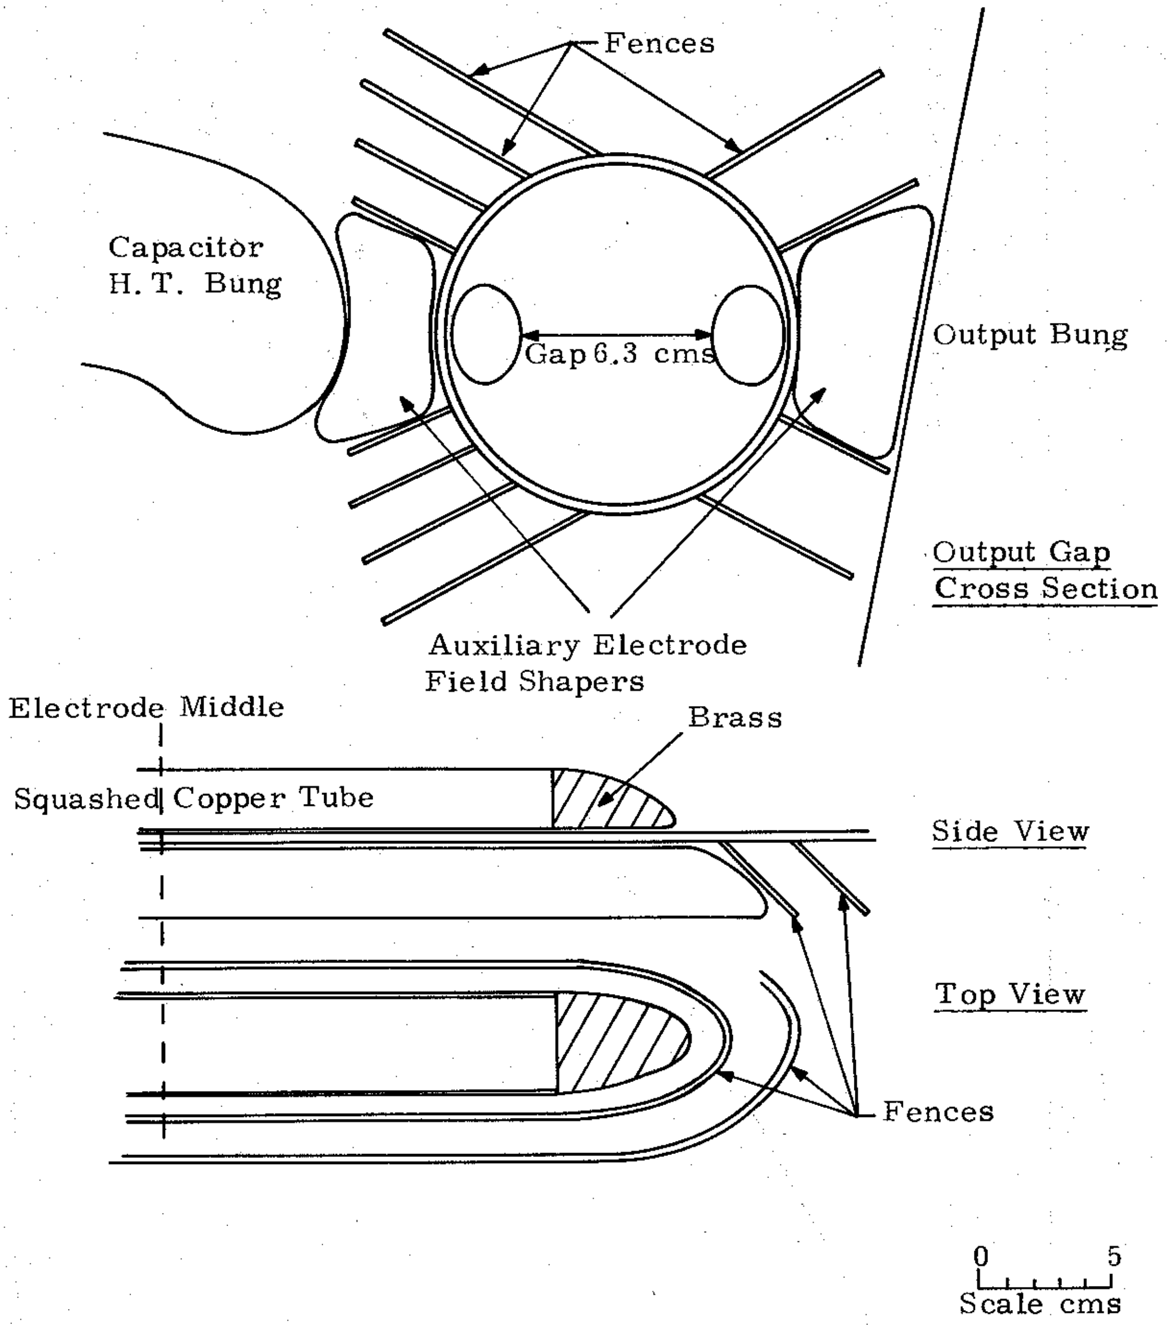

Figure 7. Sketch of Pulse Charged Output Gap Showing Approximate Location of Fences and Auxiliary Electrode Field Shapers

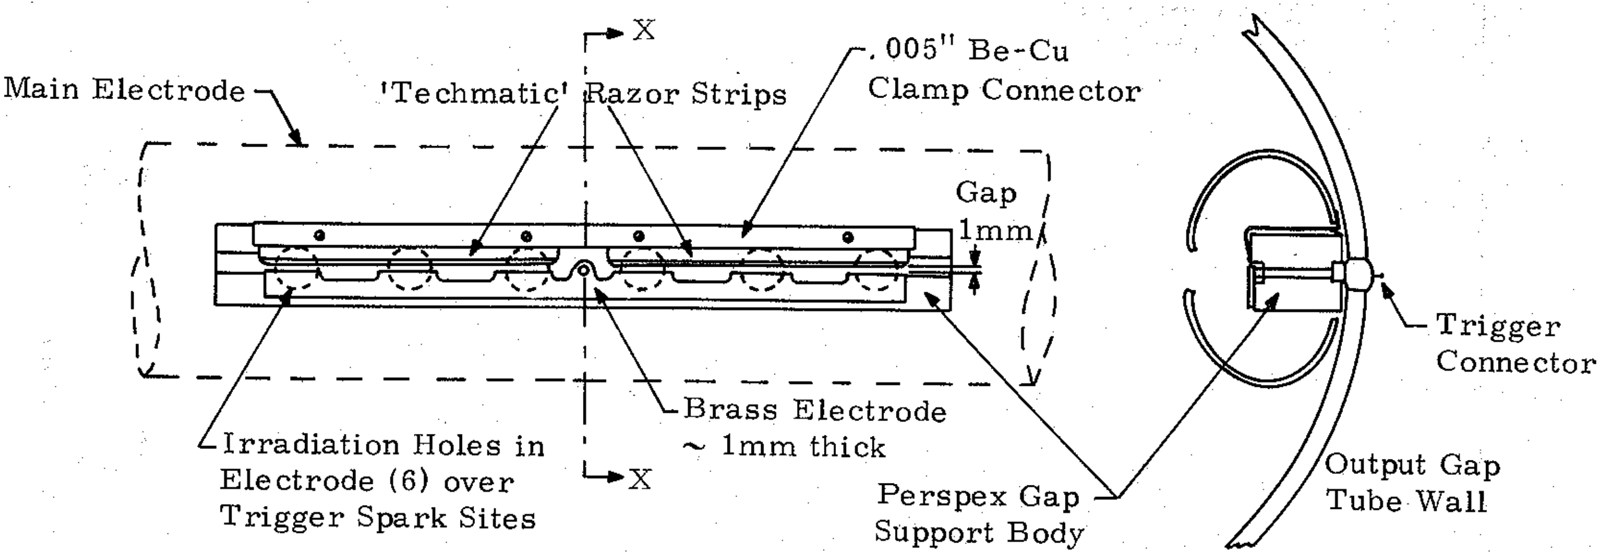

(Full Scale)

Plan of Irradiator Gaps

Cross-Section X-X

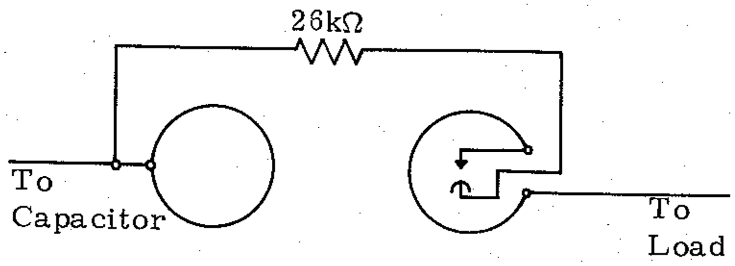

Output Gap Circuit

Figure 8

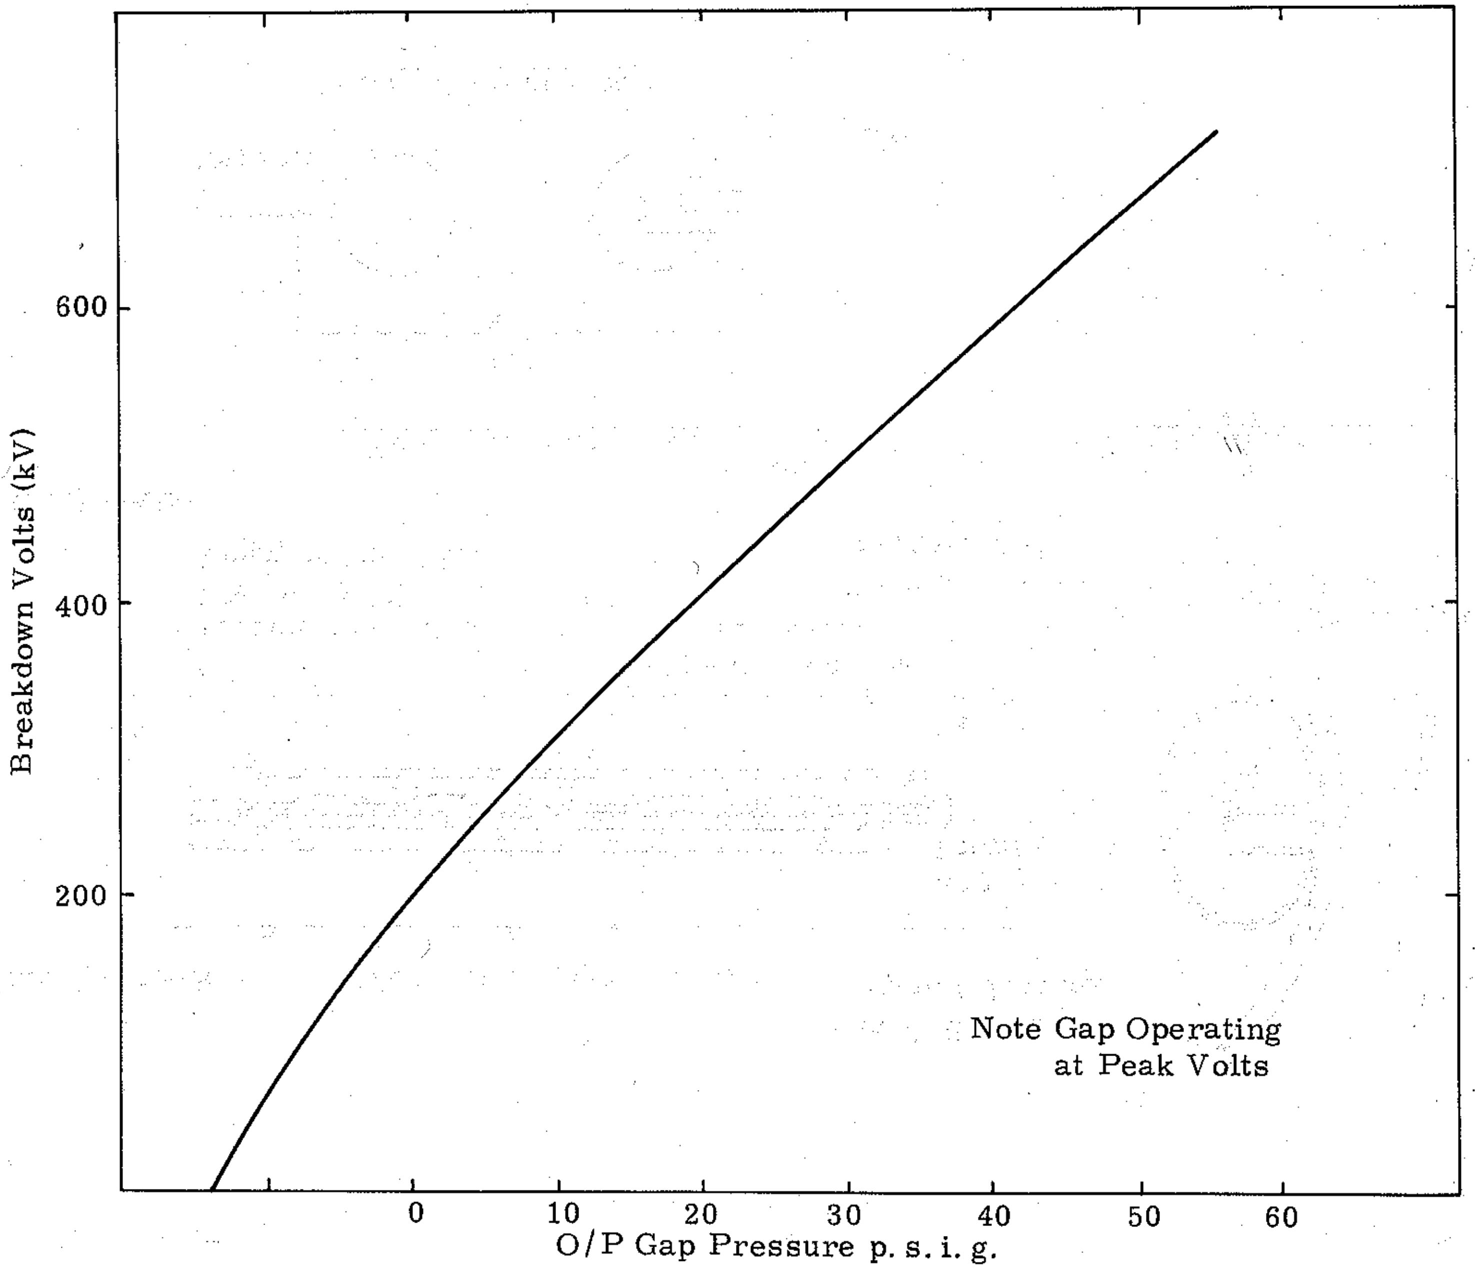
Breakdown Volts (kV)
600
400
200
0
10
20
30
40
50
60
O/P Gap Pressure p. s. i. g.
Note Gap Operating
at Peak Volts

Figure 9

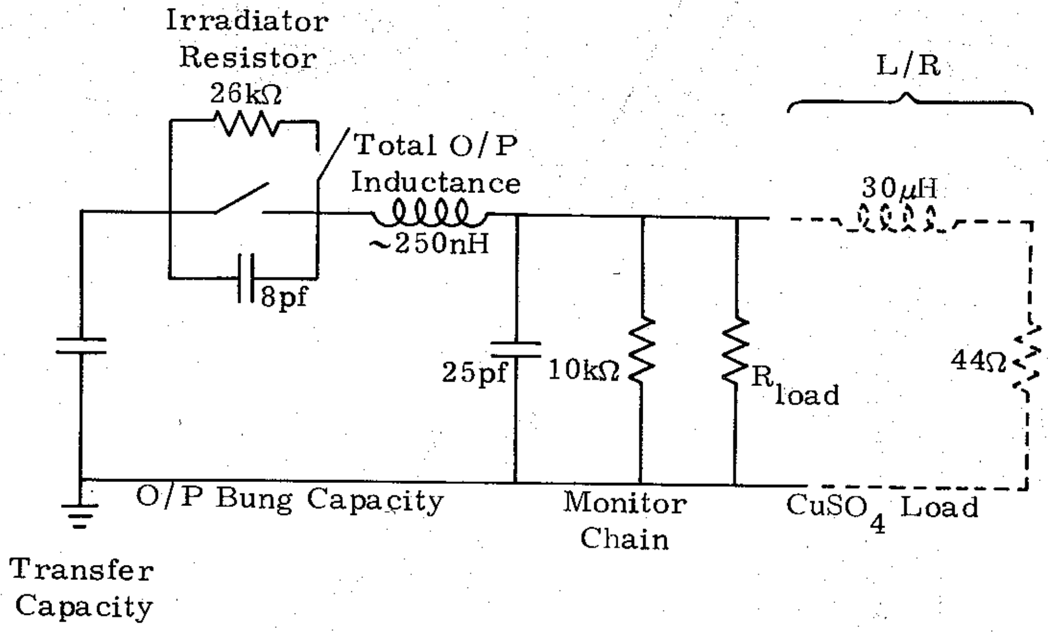

The transfer capacity is the fast available value including 30 pfs of the transformer self capacity

Figure 10. Output Lumped Constant Circuit and L/R

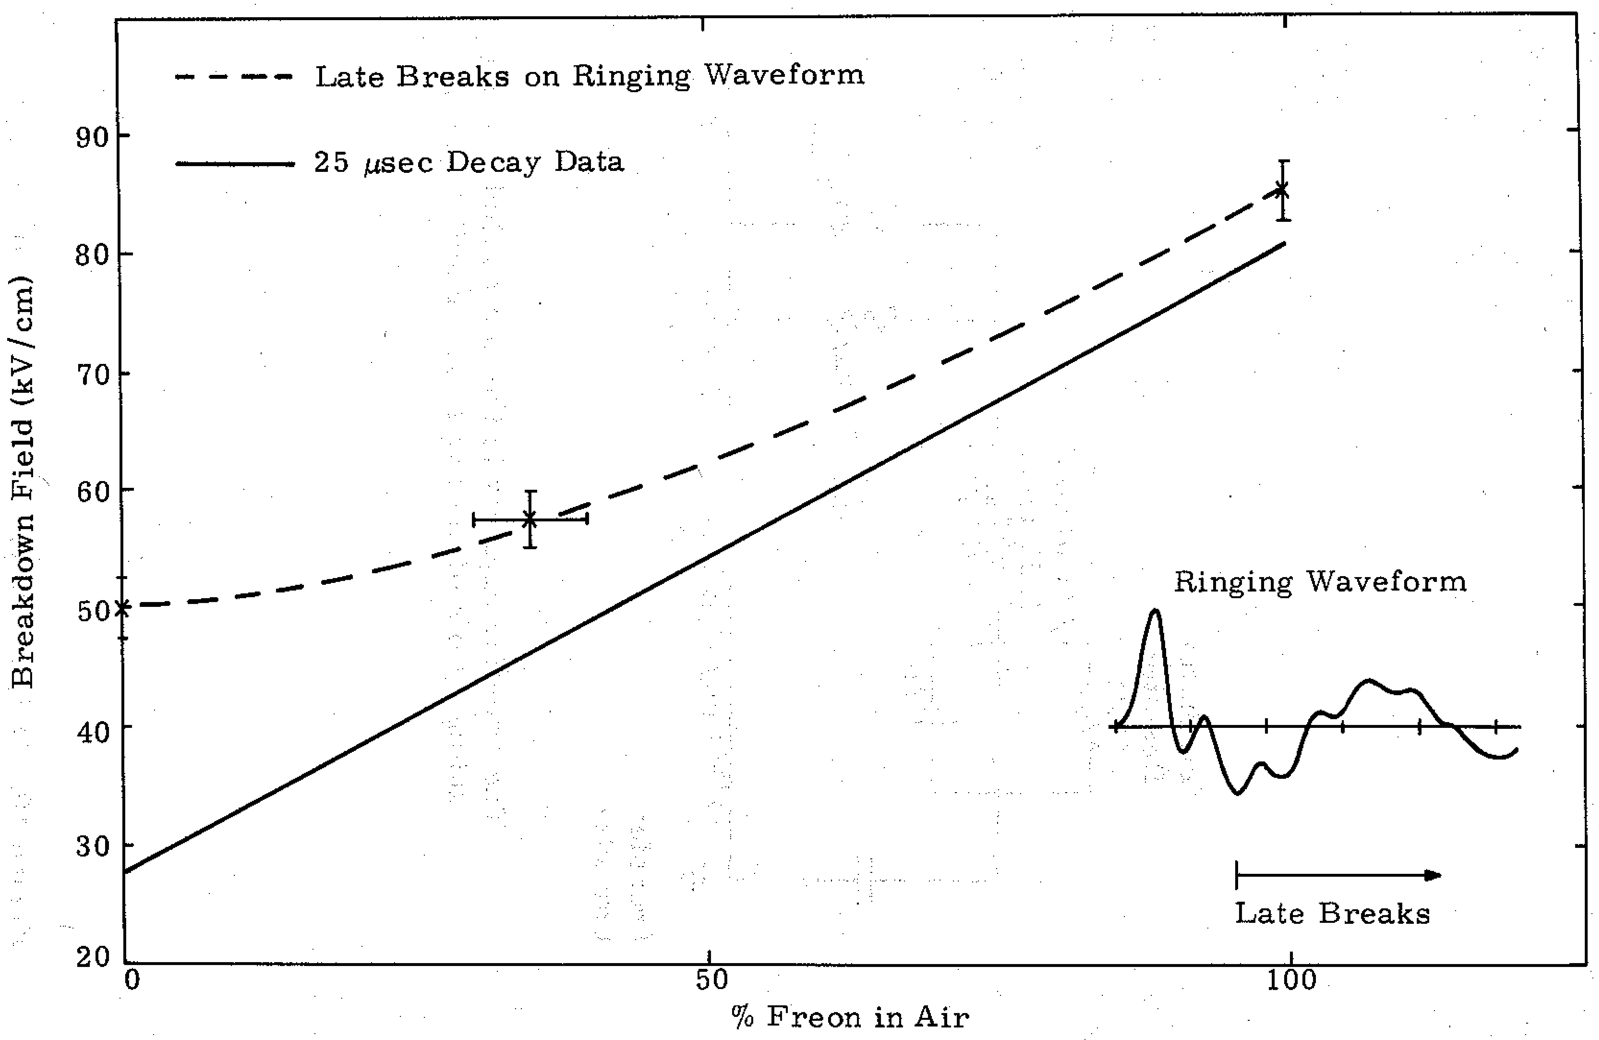
Late Breaks on Ringing Waveform
25 μsec Decay Data
Breakdown Field (kV/cm)
90
80
70
60
50
40
30
20
0
50
100
% Freon in Air
Ringing Waveform
Late Breaks

Figure 11

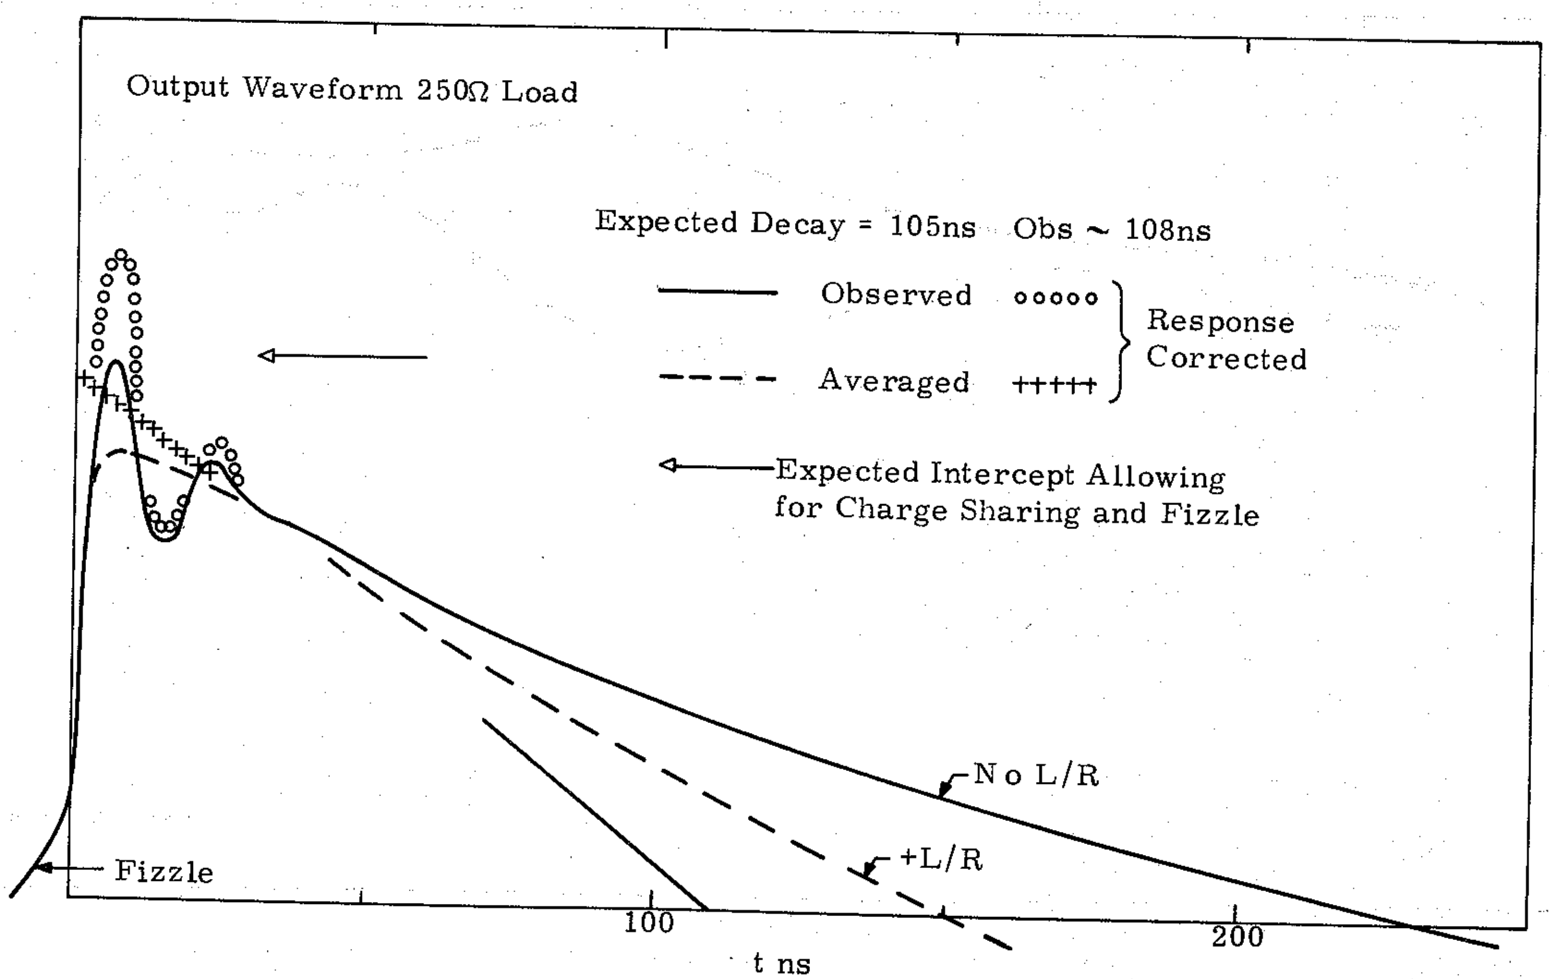
Output Waveform 250Ω Load
Expected Decay = 105ns Obs ~ 108ns
Observed ooooo
Averaged +++++
Response Corrected
Expected Intercept Allowing for Charge Sharing and Fizzle
No L/R
+L/R
Fizzle
100
200
t ns

Figure 12

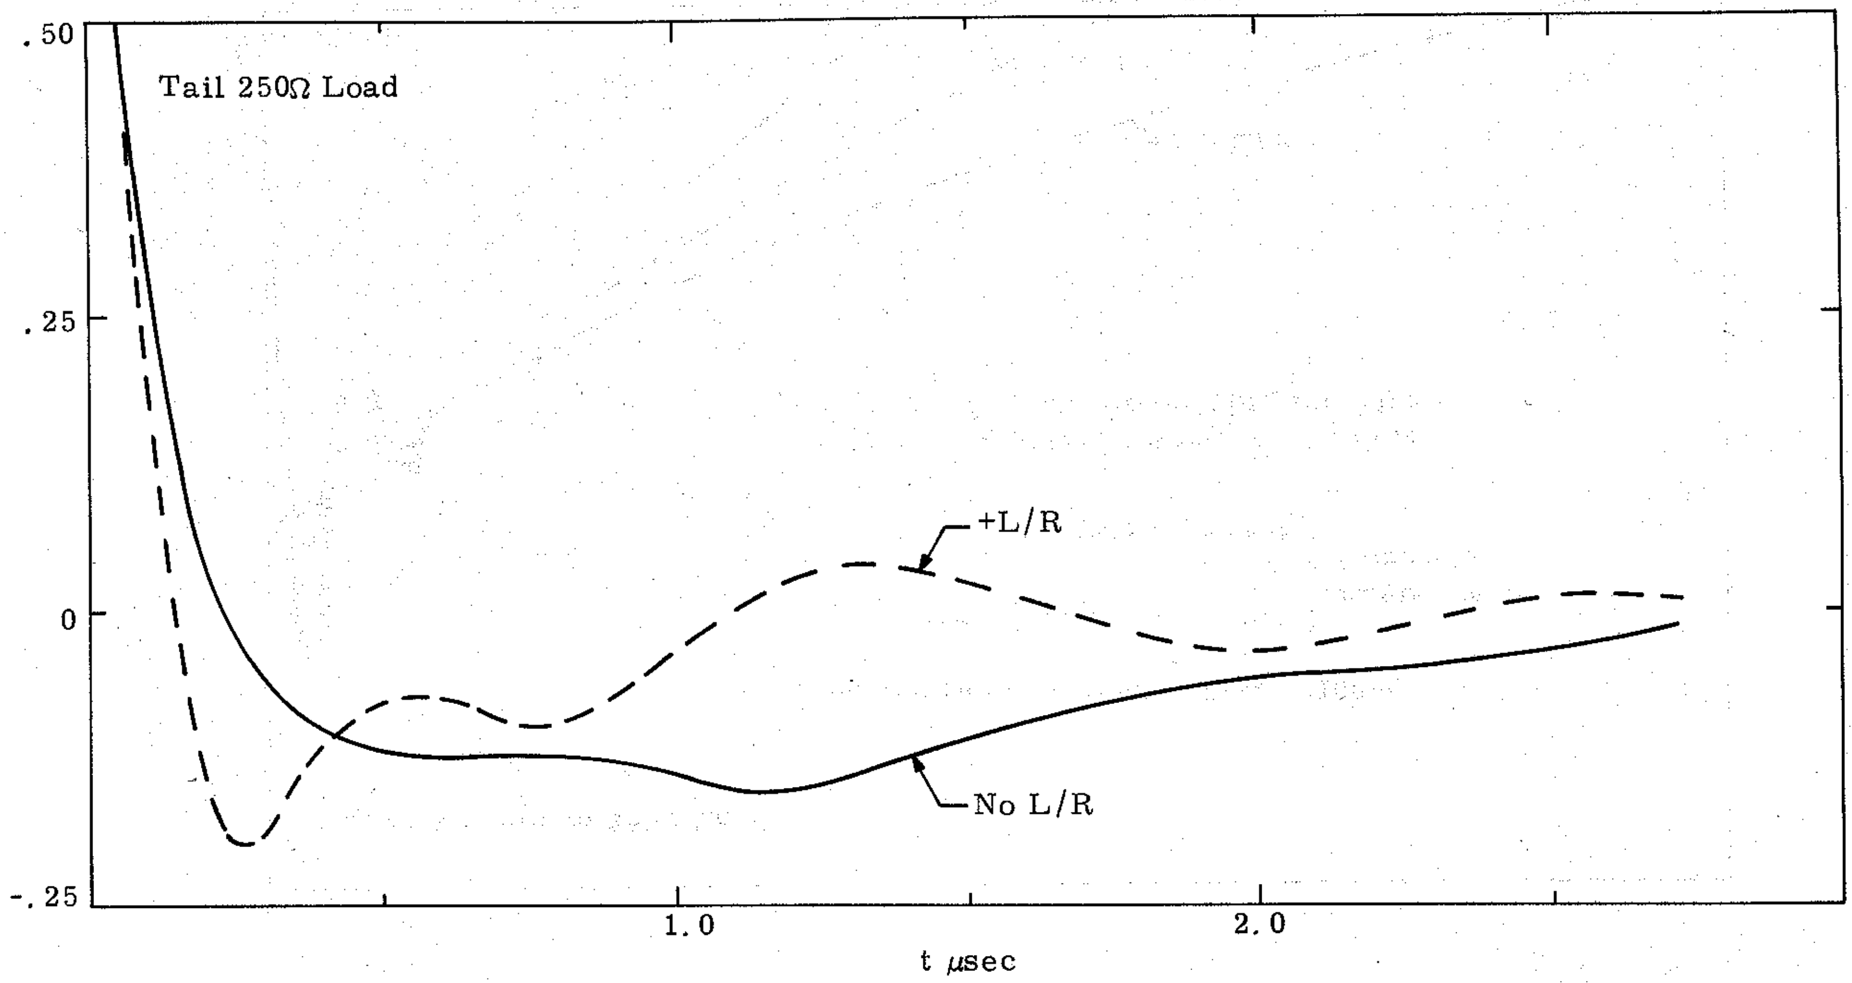

Figure 13

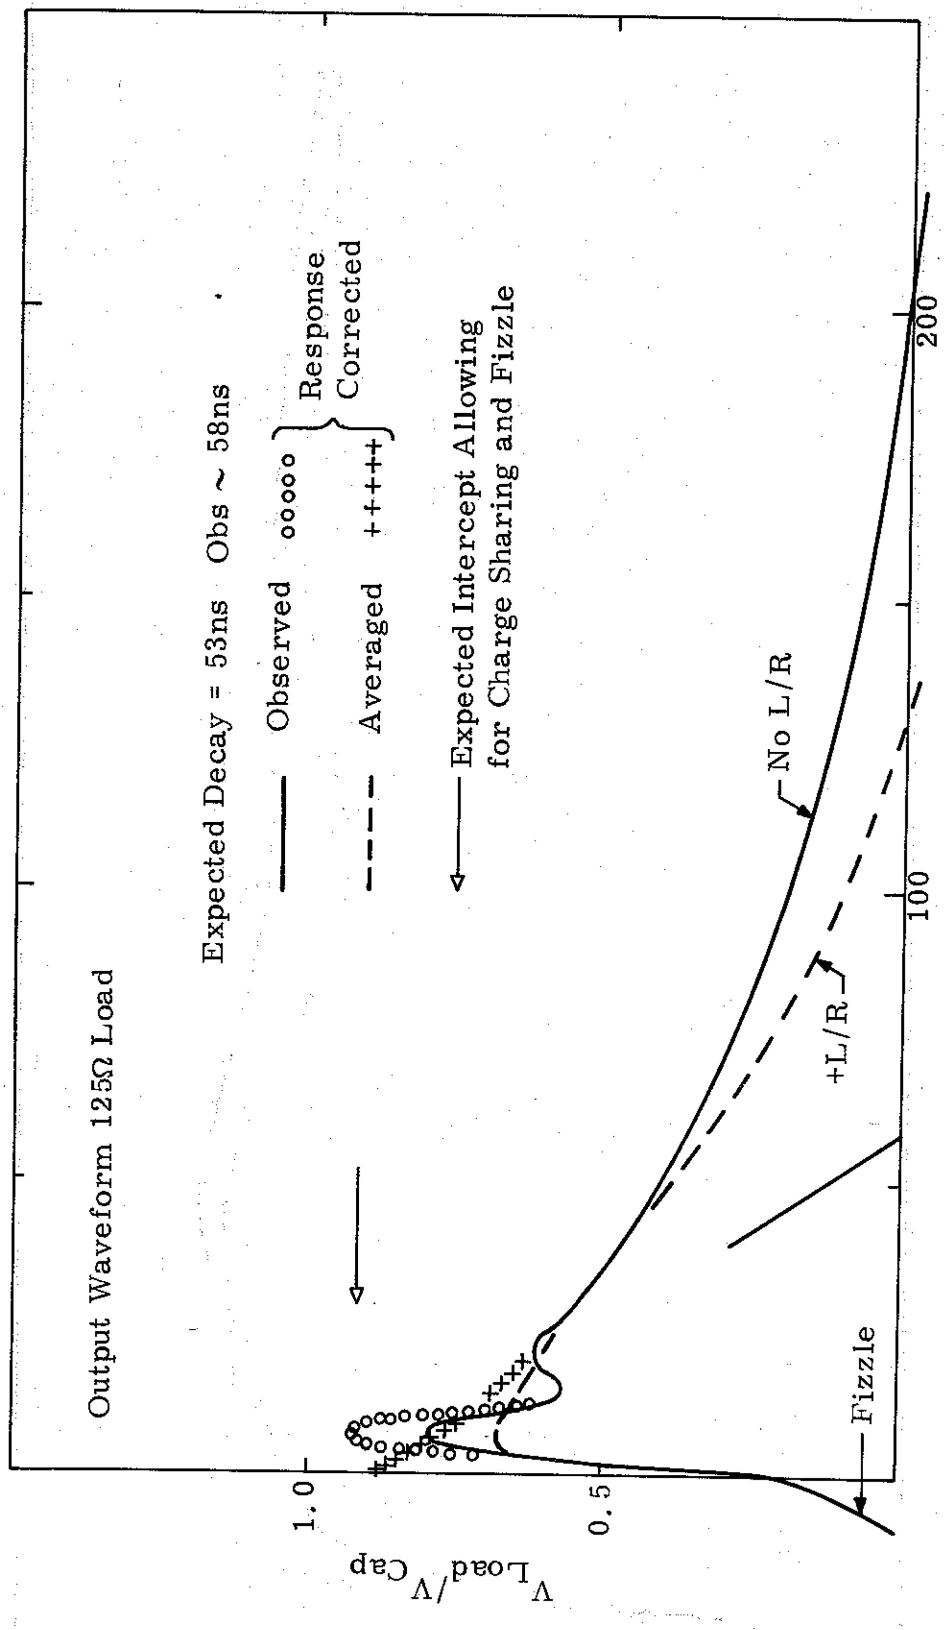
Output Waveform 125Ω Load
Expected Decay = 53ns Obs ~ 58ns
Observed ooooo
Averaged +++++
Response Corrected
Expected Intercept Allowing for Charge Sharing and Fizzle
No L/R
+L/R
Fizzle
1.0
0.5
V_Load/V_Cap
100
200
t ns

Figure 14

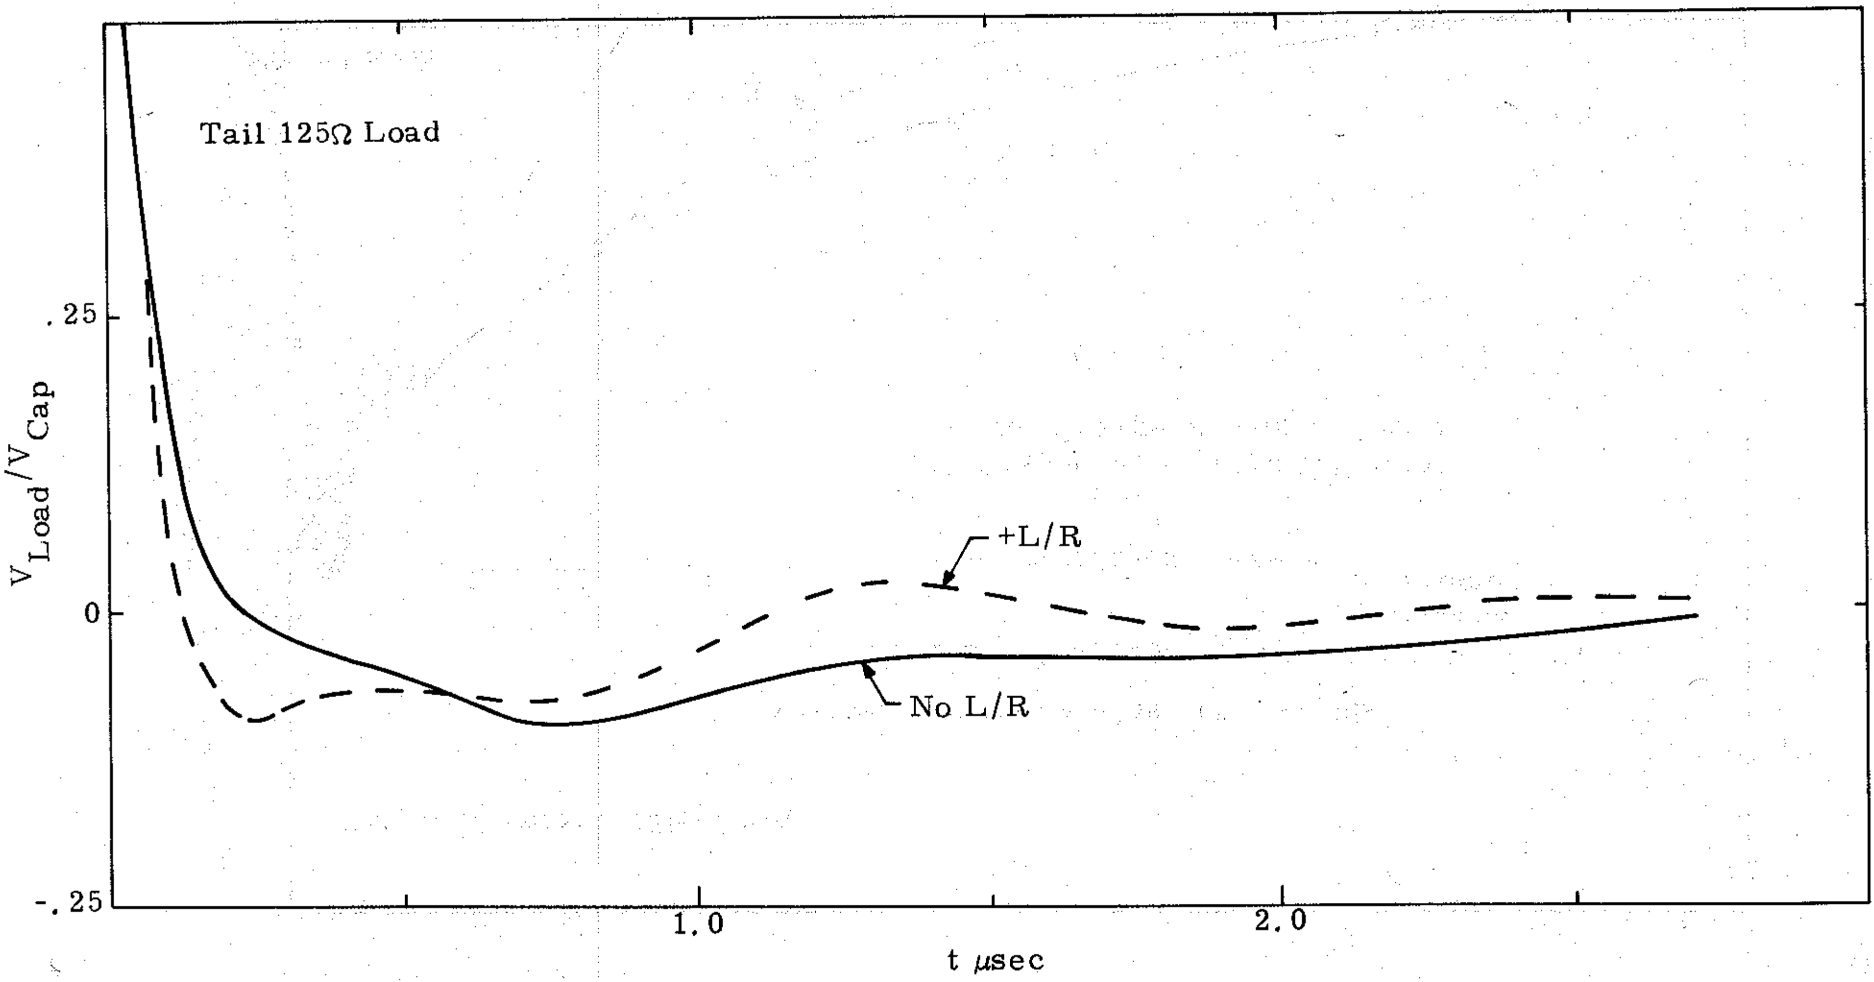

Figure 15

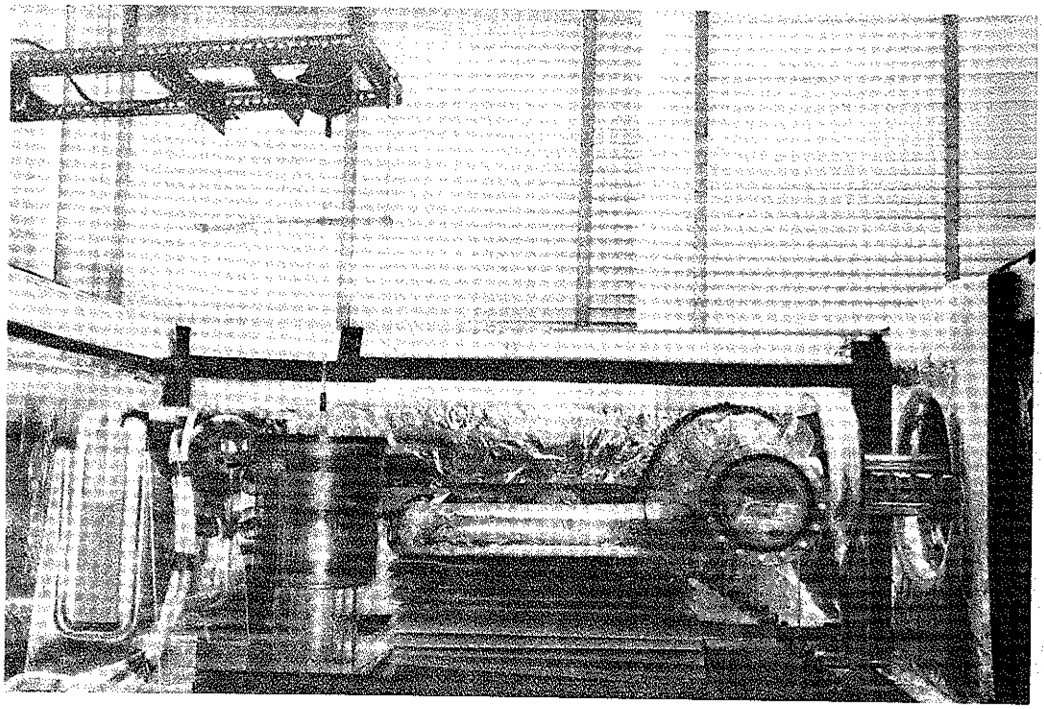

PRINT 1 - GENERAL VIEW OF LARK

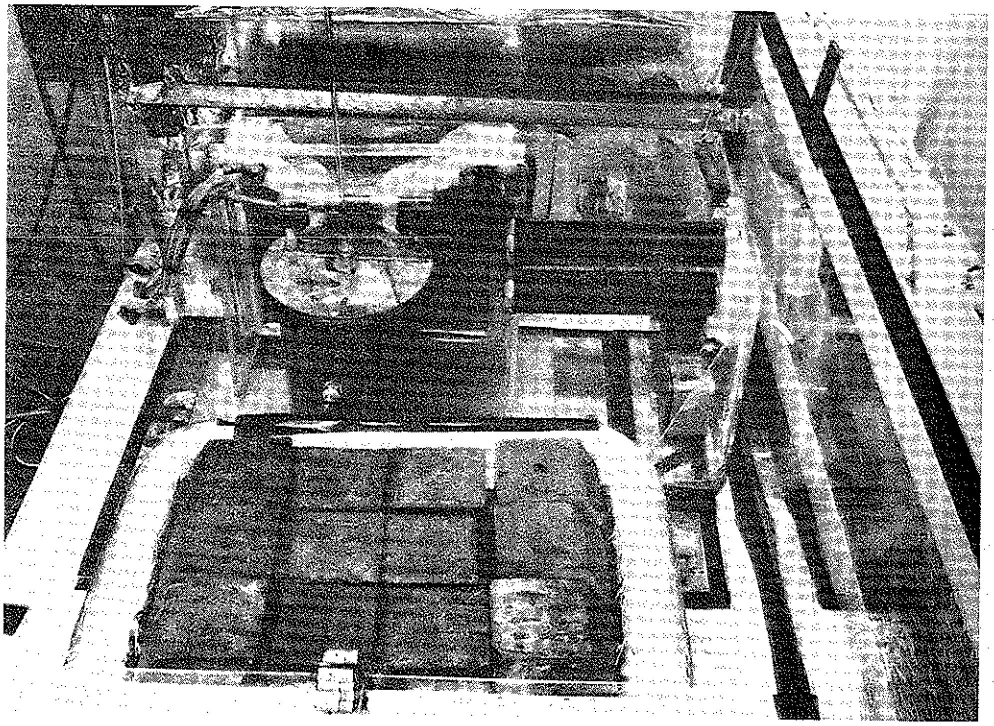

PRINT 2 - VIEW SHOWING ISOLATING INDUCTORS

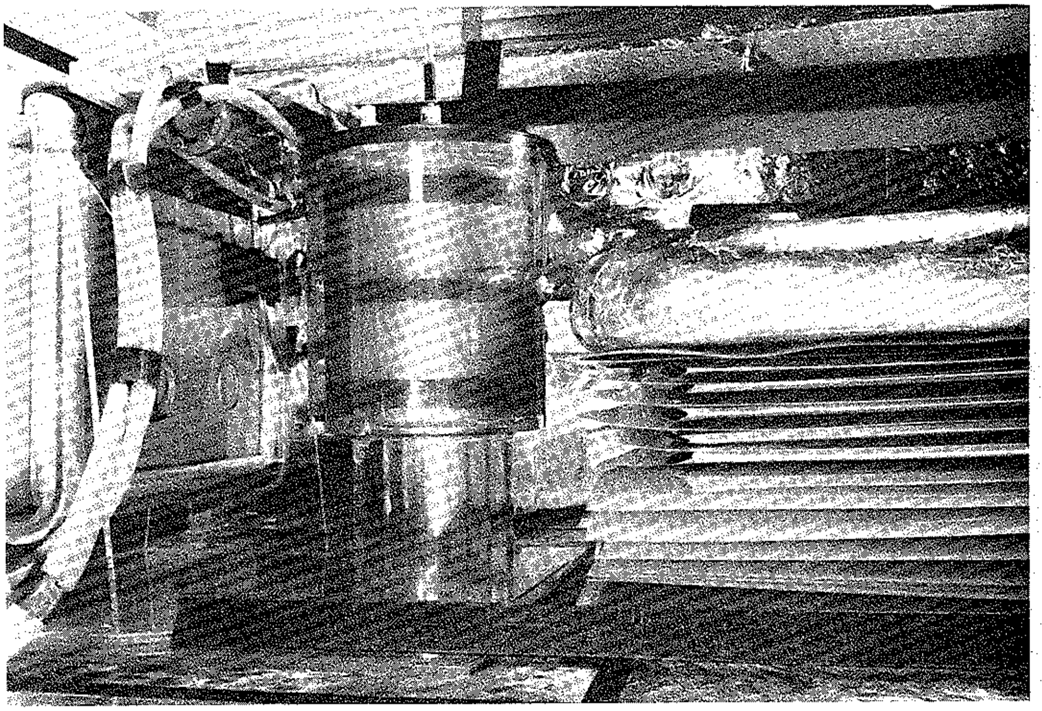

PRINT 3 - VIEW OF START GAP and TRANSFORMER WITH SHROUD

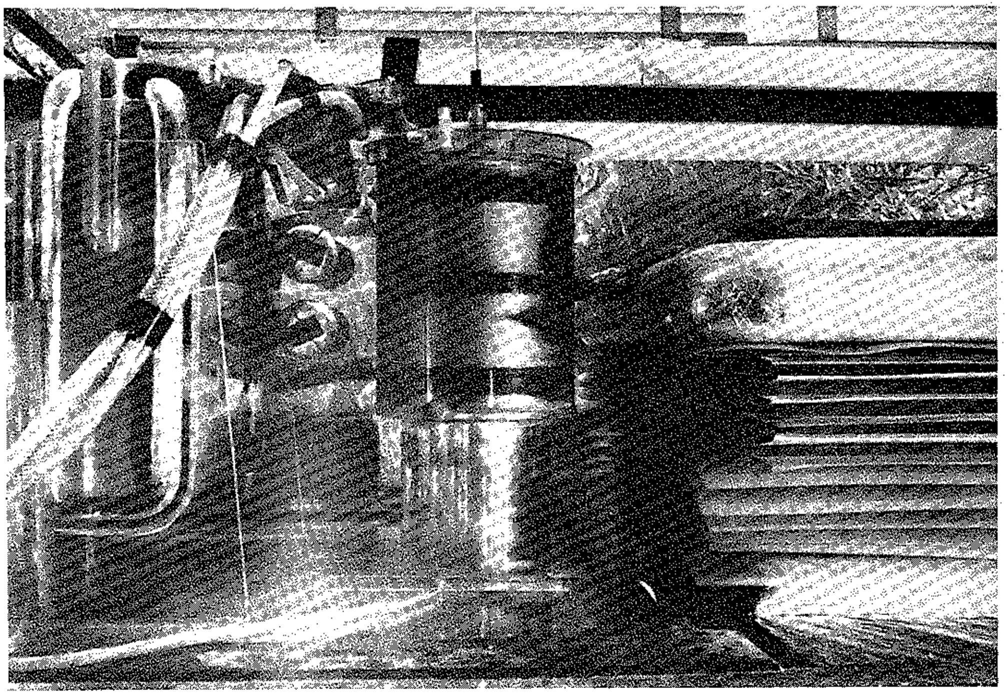

PRINT 4 - VIEW OF SHROUDLESS TRANSFORMER and CAPACITORS

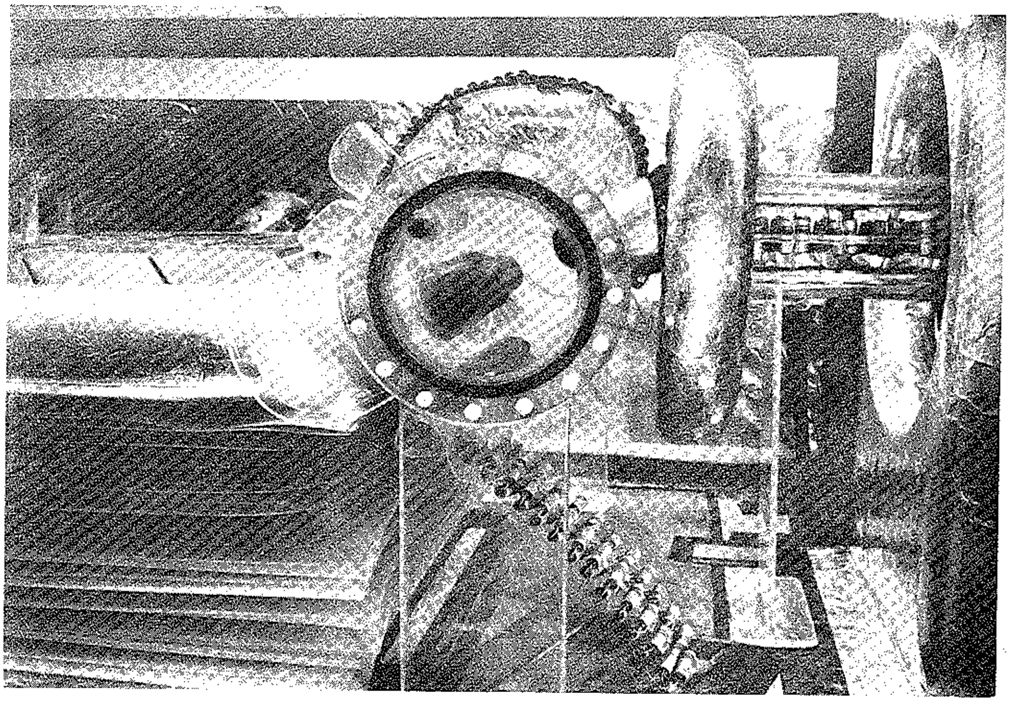

PRINT 5 - OUTPUT GAP, LOAD, and MONITORS

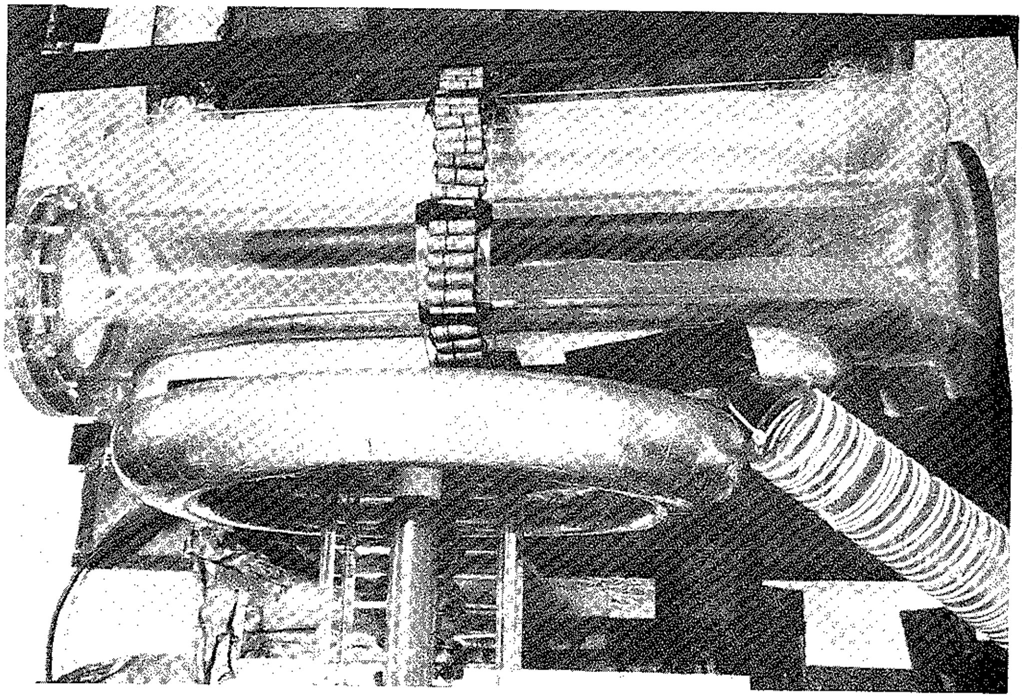

PRINT 6 - TOP VIEW OF OUTPUT BUNG and INDUCTOR