11 Appendices

11.1 Optical waveguides

For ease of installation and laying into ducts, with low-enough optical attenuation, the only known satisfactory answer will be to use a glass fibre as the transmission medium. At one limit of its constitution a 'glass' becomes pure silica; and one form of such a silica, called 'InfrasiT, has been measured to have a bulk-property attenuation of not more than 10 dB/km at the normal gallium arsenide-laser wavelength of 0.93 microns. But silica is more difficult to draw into thin fibres than the usual glasses, on account of direct and side effects of its higher melting point. But on basic grounds it is believed that a fairly normal glass, with a suitable lower melting point, should be producible in as transparent a form as the Infrasil. Research is now in progress on this matter, in several countries. For frequency bands of up to a few gigaherz the first answer is likely to be a fairly accurate circular-section glass fibre of not more than several microns in diameter, operating in the HE_U mode. To minimise other-moding problems it is planned to fuse round this inner core another glass of a slightly lower refractive index, to a diameter of about 100 microns, this outer sheath having an outer layer that is optically lossy.[18] A suitable bundle of such composite fibres will be drawn down, from a much-larger diameter sample, to the required thickness. It will not be easy to learn how to keep the needed tolerances in mass production, at economic prices: but this will be done. The outer sheath is necessary to make the tolerances practical. But any energy flowing into it is (or should be) completely lost, adding to the total attenuation. When in such fibre form, therefore, a total attenuation of about 20 dB/km must be expected.

11.2 Wideband optical fibres

In multi-channel cables having large total bandwidths, an important economy can be obtained by using wideband repeaters fed by wideband fibres, thus minimising, in particular, the number of repeaters. The most promising practical wideband optical waveguide that can be foreseen now is the 'quadratically-graded refractive index' type. If the fibre is made such that its refractive index varies radially as $n=n_0\,(1— kr^2)$, where r is the distance from the axis of the fibres, n_0 is the refractive index at the axis, and k is a positive constant, then a focussing action occurs by which most of the optical energy is prevented from impinging on the walls of the fibre.[19]

If a beam of light could be launched into the fibre in an ideal manner, then the beam would continue to propagate uniformly, the effects of diffraction being compensated by the non-uniform refractive index of the fibre. If the launching is imperfect, as is inevitable in practice, the beam will follow an undulating sinusoidal path along the fibre, and, moreover, the diameter of the beam also will vary sinusoidally. The longitudinal periods of these sinusoidal oscillations are, respectively, 2pi/sr rt(2k) and pi/sq rt(2k) This is illustrated in Fig 5.

Although this system is greatly over-moded, the modes have equal propagation velocities, and this type of fibre therefore offers a much wider bandwidth than in the narrower-type, single-mode fibres. It should be capable of carrying at least the 10 gigabits/sec rate that we have already foreseen. Any bends in the fibre will, statistically, cause an increase in the amplitude of the undulation of the beam, eventually resulting in energy being lost from the surface of the fibre. An automatic compensator of a simple, cheap and reliable form which would be installed at intervals along the fibre to prevent this, is now being studied.[20] There is good hope that such fibres will operate successfully down to about 50 microns in diameter, enabling both them and the light they convey to bend round the needed radii.

Quadratic-law glass fibres have already been made in Japan,[21] and in the USA; but these early types are unlikely to meet the tolerances needed here. A fairly promising new way of producing the fibres, based on successive deposition of thin glass layers from liquid glass having a time-varying refractive index, has been proposed.[22] Successful development of any manufacturing method will again be a stimulating challenge. A main problem is the tight tolerance on the quadratic law governing the refractive index. An easing of this tolerance on new lines is now being studied. It is based on a balancing method, together with the introduction into the Second Order differential equation, representing the fibre characteristics, of a spatial damping term. Because of the smaller optical power striking the walls it is hoped that such quadratic-law wideband guides will have less attenuation than the single-mode types. The only safe guess now, though, is that it will be about the same, i.e. 20 dB per km.

In the general type of glass that now seems likely to meet best all the requirements, one effect of the positive dispersion will be to extend the effective widths of 10 picoseconds-duration signal pulses by about 40 picoseconds per km length of the fibres. At rates of about 1 gigabit/sec and above the corresponding extensions cannot usually be tolerated without compensation. Suitable dispersion compensators on conventional lines seem quite unpractically large and expensive. A small, robust design, however, using different principles but still equivalent to a passive linear network, is now being studied.[23]

11.3 Repeater powers and spacings in optical cables

The early versions of room-temperature lasers will probably produce a pulse peak power of about 300 milliwatts when at a rate of 300 megabits/sec, when the duty cycle is 30 per cent. At the normal gallium arsenide wavelength of 0.93 microns (a frequency of 3.2×10^{14} Hz), the energy per photon is 2.14xlO~[19] joules. At a PRF (pulse repetition frequency) of 300 MHz, suitable for future single-channel TV working when using PCM, we thus get 1.4×10^9 photons in each pulse. The basic limit to the analogue signal-to-noise ratio (S/N) in each pulse of a digital optical system is set by the square root of the number of photons within it, as their times of arrival are random.

If we take 20dB, a normal figure, for this S/N needed, and allow for predictable further and fairly basic deterioration, we find we need about 400 photons per pulse. But at small cable attenuations this is such a small fraction of the photons per pulse available, 9.4 $\times 10^8$, that a second source of noise will entirely predominate — a noise which though basically determinate must in practice be treated as gaussian, as it contains many independently variable components. This is the noise due to the sum of many small, stray reflections from and between a large number of slight impedance irregularities in the fibre waveguide. It would probably not be realistic to foresee now that this total reflection noise will be more than 30dB down on the signal peak power, though this figure may prove pessimistic. So whatever the laser power or the attenuation in the fibre, we cannot get a S/N ratio of more than about 30dB. If we have 60dB cable attenuation the ratio of signal to photon noise is also about 30dB, which will give a total S/N ratio of about 27dB, providing a good safety factor against further losses. With a cable attenuation of 20dB/km, we thus get a repeater spacing of 3 km, for the 'narrow-band' fibres.

If we now use a wideband fibre suitable for an information rate of 10 gigabits/sec, with a pulse duty cycle reduced to 10 per cent, each pulse peak now being at 900 milliwatts, we get 4.2 $\times 10^7$ photons per signal pulse. In this case the waveguide attenuation must therefore be reduced to about 45dB for the equal-noise condition that gives a total S/N ratio of 27dB. The repeater spacing is then 2.25 km. Some further considerations in the optical repeater designs have been given elsewhere.[24]

11.4 Design factors in sub-ocean optical cables

Several severe constraints will have to be met that are additional to those in the long overland optical highways. Perhaps the chief is the component reliability, mainly in the lasers at the repeaters. There are likely to be about 20 fibres in each cable, each fibre having its own repeaters. Assume that in each we need a half-life of 20 years, i.e. a probability of not less than 0.5 that in this length of time all repeaters will continue in correct operation. Let us first calculate the value of p_r, the maximum allowable failure probability of each repeater in that time span, when a new cable is switched into service. For a cable 4 500 km long, likely to be about the longest route, there will be about 2 000 repeaters. The probability of each repeater *not* failing is $(1 - p_r)$. So the probability of all the 2 000 repeaters not failing is $(1 - Pr)^{2000}$, which we have set at 0.5. Therefore

$$1-(1-p_r)^{2000} = 0.5; \text{ from which}$$

$$p_r = 1-(0.5)^{\frac{1}{2000}} \approx 0.0003.$$

We need a standy-by type of redundancy in all the components at all likely to fail, for any practical solution. Suppose this redundancy is 6-fold, 5 new groups of components being available to be switched in successively, reliably and automatically, by a monitor at the repeater on a fault condition within it — a design feature that is certainly feasible. All the components in the operating group are ageing together. Let the probability of a group's failing within 20 years be p_u. Then the probability of failure of the whole repeater, $p_r = (p_u)^6$, from which

$$p_u = (p_r)^{\frac{1}{6}} = (0.0003)^{\frac{1}{6}} \simeq 0.25.$$

Considering now the structure of each repeater unit containing the necessary group of components, each unit will probably contain 3 devices of type a, fully driven gallium arsenide laser oscillators, with 3 devices of type b, which will be either under-driven gallium arsenide lasers or of a passive 'Q-switch' design. Failure in any one will cause failure in the repeater. From the meagre evidence so far it is likely that each device will be subject to two kinds of failure, the first, type-1, giving a failure probability in time *dt* given by $p_{ia}t\text{-}dt$ or *p^t-dt,* and the second failure type by $p_{2a}\text{-}dt$ or *p^b'dt.* Type-1 failure is thus a true deterioration due to use, while type-2 is equally likely to occur at any time provided the device is in operation. As all these probabilities are small, even when integrated over 20 years, the following formula applies:

$$p_u \simeq 3\int_0^{20\ years} (p_{1a}+p_{1b})t + (p_{2a}+p_{2b})\cdot dt$$

$$= \frac{3T^2}{2}(p_{1a}+p_{1b}) + 3T(p_{2a}+p_{2b}), \text{ where } T=20.$$

p_2b are likely to be about equal; and there is no reason so far why their effects, considered alone, should not give a 1 000-year statistical half-life, as in a good silicon transistor. We shall assume therefore that ^3$_{2a}$=j)$_2${,=0.0005, the corresponding value. As a further best guess, we shall assume that $p_1h = l/3p_{ia}$. We can now solve for p_{la} ; we get a value of 0.00024, from which we find that the type a devices, the fully driven gallium arsenide lasers, at their sub-ocean temperatures, must have a half-probability life of about 65 years, a figure that it will be a real challenge to attain, though quite achievable within the 30-year period before these still better lasers will be needed. As an alternative, a better and cheaper way of switching in stand-by lasers is now just beginning to look quite feasible, even up to a 100-fold redundancy figure, thus making usable gallium arsenide lasers of much shorter average lives than 65 years. It is based on the optically pumped variety that is referred to in Appendix 11.12.

The next design point is the power supplies to these repeaters. As we have seen, each repeater point will probably contain 20 sections, one for each fibre, with still better gallium arsenide lasers consuming about 2.4 watts per section and giving a total of 48 watts at each repeater point. Individual isotope heater units at each, probably using strontium- 90 and energising thermo-electric elements, are already being considered for lower-power, lower frequency submarine cable repeaters. But for this optical version, unless the use of large atomic reactors, from which the fuel is a byproduct, spreads much more rapidly in the next thirty years than is now expected, they would probably add at least R60 000 to the price at each repeater point, as the main factor is the cost of the nuclear fuel (for a 20-year half-life). The world's supply of strontium-90, too, will probably not be adequate for this method.

Conventional repeater powering, by + dc volts at one end of the cable with an equal — dc at the other, could be practical, however. At each repeater point the total dc supply could be at 9.6 V with 5 amps, a perfectly feasible cable current (through a metallic conductor additional to the glass fibres). If the dc resistance loss equals the repeater power consumed, the supply at each end would he +19.2 kV, a quite feasible value; and the total dc power needed would be 192 kW, larger than now used, but adding only a negligible amount to the annual costs. The usual high-voltage, expensive blocking condensers at the repeaters could be dispensed with, giving an appreciable reduction in total cost. Looking ahead beyond about 15 years, it has become clear very recently that we may reasonably expect room-temperature super-regenerative gallium arsenide lasers having drive powers of only about 100 milliwatts, with efficiencies when used on continuous wave of up to 50 per cent —• and of either the junction or the optically pumped varieties. If the latter were used, the over-all efficiencies would be 25 per cent, and the optical outputs would be not more than lOdB down on the figures given above. The repeater spacings would then be reduced by 7/9 to give the same signal-to-noise ratio as before. The cable dc drive current could then be reduced to only about 0.7 amps, at a total cable voltage 9/7 times as large. 9/7 times the number of repeaters would be needed; but the increased, cheap laser redundancy available by optical pumping would more than compensate for this slight increase in the total number of lasers from the standpoint of reliability.[29]

11.5 Cost factors in optical networks

Just as in most present networks, it is nearly certain that at least 80 per cent of the total expense will be in the subscriber's local and extended-local areas. The main factors here are: (a) the interest and depreciation on the capital invested; (b) the labour charges for repairs and maintenance; and (c) the labour charges for changed and new subscriber connections. At present factors (b) and (c) together are often more than half the total; and other things being equal they will of course increase as the result of higher pay packets. But there will be counteracting trends. The expected demands will be so large and reasonably certain that it will become economic to build-in subscriber terminals in almost all new homes and offices while new groups of houses are being erected, the cheapest time to do it, to be used and paid for when required. No architect even now would dream of planning a normal new house or block and leave out the electricity supply. It will be the same for its communication facilities. The capital cost, moreover, in installing at the start a subscribers' network having much larger reserve capacities than at present will be cheaper by optical methods, reducing the cost of extensions later. A further factor will arise from the much faster procedure at cable joints and drop-off points; instead of soldering techniques as at present, a plug will simply have to be put into a socket —• a method, too, that will be proof against human errors that now occur. Almost all fault diagnosis will be carried out, largely automatically, in the exchange buildings, not in the field; all the field man will have to do will be to change a plug-in assembly at the place directed.

Now a more basic point: nearly all experience so far has shown that when the bandwidth of an electronic component rises, so does its cost. Will this still be true when we change to guided optical communication? — The answer is not necessarily. The two basic components are (a) the future-type gallium arsenide lasers (or some new kind equally good and still cheaper); (b) the optical fibre waveguides. As to (a), so little raw material is needed for each laser, about 5 milligrammes, that the cost will be negligible compared with other factors, when we know not only how to make satisfactory types but also how to get a device yield rate of, say, at least 50 per cent in mass production. The over-riding factor will be in getting back the quite large research and development expenses. Similar arguments apply to the future costs of the encased optical fibres; the completely over-riding factor will again be the research and development expenses before we know how to tool-up for reliable mass production.

As to installation first costs, when suitable methods have been fully developed it should be no more than that of present-day multi-pair voice-frequency types, as the cost-reducing factors already mentioned should roughly balance the price difference between the installed, cheap, mass produced extra repeaters and the man-hole loading coils now used. The local-area maintenance costs, though, will rise somewhat, as there are more active components in the network, but by not more than about 50 per cent, eventually, due to the labour cost-reducing factors already mentioned, together with the use of a stand-by repeater unit at each repeater man-hole, automatically switched in from the exchange when a fault condition occurs, thus enabling most of the maintenance work to be done on a cheaper, pre-planned schedule. The total operating annual cost per subscriber terminal, in a given local area, will probably not be more than three times what it is at present.

11.6 Optical switching methods

It is most likely that a combination of, (a) time switches and, (b) space switches will continue to be used, as is the present plan for base-band PCM tandem exchanges. Time switches operate by moving a channel from one time slot to another within a given TDM channel group, by adding a suitable delay. The space switches are much faster, non-mechanical versions of the older rotary or cross-bar types.

To obtain the adjustable delays for the time switches, an optical circulating digital delay-line technique will often be the best answer. The delay medium would be, for example, a wideband glass fibre — or a multiple reflection version without a wave-guide, between parallel plates, at a slight angle to the normal. The active element to keep the optical pulses circulating would probably be a gallium arsenide laser preferably optically pumped, in integrated circuit form — a uniform film of gallium arsenide, between suitably reflecting layers, giving separate laser actions wherever a real image from a larger, pumping laser impinges on it. It is probable that at least 3 000 separately adjustable delays, each with 1 000 equal steps, will eventually be obtainable in this way in a volume of not more than 1 cubic metre.[12] These devices would be suitable for a rate of 10 to 20 gigabits/sec. An alternative is to write onto the gallium arsenide film a '!' or a '0' digit by the pumping pulse only of a similar optically pumped laser, at a level suitably below the starting threshold for lasing. The information would then be read out at the later time desired by a suitably timed optical signal pulse, of an energy and wave shape adapted to overcome the above starting threshold. This version would be smaller; but the storage time would be limited by the maximum spontaneous decay time of the gallium arsenide as a laser, at present about 10 nanoseconds. A third method would be a crystal store but the write-in time would then be slower, at the practical write-in powers, probably needing at least 10 pulse periods at a rate of 10 gigabits/ sec. Still further methods are being examined.

For the space switches the use of optical carrier frequencies gives us a new dimension in facilities. A focussed laser beam, rapidly tilted in two planes by electro-optical means, takes the place of the mechanically moving arm of the old-type rotary switches. Several ways of achieving this have been studied experimentally.[13] Most such existing methods are capable also of handling at least a small group of TDM channels by each laser beam — but only at bit rates much lower than the total that will be used in our future optical channel groups and master-groups. A breakthrough in technique is needed. The key will be the use of purely optical gates of the AND and other varieties, much faster analogues of the well-known devices in the electronics art, in many examples producing focussed beams tilted at the correct angles. The most promising line so far lies in employing the nonlinear amplitude characteristics of gallium arsenide laser devices. One version has already been described.[14] Improved varieties are being studied, including also all-optical analogues of waveshape regenerators, pulse time-noise removers, etc.[15] This type of research is the beginning of a new technology; it is truly fascinating in its possibilities, for future still higher-speed computers as well.

An altogether different approach, now being studied, embodies a miniaturised optical version of SYNSOL described in the text. It will resemble cross-bar in its mode of operation, but will be more flexible, and within wide limits self-adaptive to changing traffic demands.

11.7 Holographic crystal stores

There are a number of lucid articles on the elements of holography. One of them has recently been published.[25] The basic principle is quite simple; it goes back to the Fresnel zone plate, often explained in elementary physics courses. If consists of an opaque plate, A, (Fig 6) having concentric annular transparent regions, with their radii and widths quadratically graded. If a spot of monochromatic light is placed at 0, with suitable distances for OD and DI, the path lengths OBI, OCI, ODI, etc., differ consecutively by the wavelength of the optical source. The zone plate therefore acts like a convex lens, producing at I a real image of the point source at 0. By a suitable coaxial combination of a parallel beam with one that is diverging, both floodlighting a sensitised surface, such a zone plate can be made photographically by the interference fringe produced — if the two beams are sufficiently space-coherent, e.g. by producing them from a single laser by conventional beamsplitting arrangements. The resulting graded opacity gives in fact a better result than when using the completely opaque-or-transparent method of Fig 6.

If we now use *n* diverging beams, from bright spots, again in combination with a single parallel beam, the result is *n* real images, in the configuration of the originating *n* sources. It is in this way that a hologram stores a picture — or a series of binary-digital words, in terms of bright or dark spots.

If the parallel writing-in beam is now tilted, a different picture can be stored. In this way, within limits, a fairly large number of pictures can be stored, each capable of being read-out separately, if for each picture the parallel readout beam is at the correct angle in relation to the hologram plane. If a parallel stack of *m* photographic plates is used for the write-in, *mn* separate pictures can be stored, provided that the resulting average total opacity of each plate is small enough to leave enough light for the plate that is furthest away in depth from the floodlighting source, and that the total scattered, defocussed light from the other images is not too large.

Although such photographic holography is reasonably efficient for a fairly permanent store, it cannot be used in the many telecommunication applications where the stored information has to be quickly erasable, to be replaced by an equally fast writing-in of new information. Neither, by present techniques, does it lend itself to efficient storing in three dimensions rather than two. Owing to an accidental discovery, however, a different method is now being explored by Bell Telephone Laboratories.[26] The authors of this article are duly cautious, suggesting only tentatively a theory to explain their highly interesting and important experimental results. But from the evidence they quote it seems to me that their explanation is at least highly plausible, even if not yet completely proved. The predictions that follow in this lecture concerning many useful applications of their method assume that at least their basic theory *is* correct. If this assumption proves to be wrong, we may have to think again — at least for the quickly changeable types of store. It is nearly certain that if necessary a good alternative technique *could* be found; so far we are only on the fringe of the many possibilities.

The method consists in replacing a stack of photographic plates or films by a crystal of an electro-optic material such as lithium niobate or lithium titinate. The split-beam writing-in wavelength used so far was 4 880 angstroms, which they also used for the single-beam read-out. The resulting solid hologram that was observed is believed to be produced as follows: (a) At the antinodes of the write-in 3-dimensional interference fringes, electrons captured by 'traps' in the crystal lattice receive enough extra energy to free them from these traps, (b) These freed electrons are then pulled through the crystal by an electric field, believed to be due to surface effects at the crystal edges, (c) The electron drift continues until each is captured by an adjacent trap (which the experiments have shown to have an average linear spacing of a small fraction of one micron), (d) The result is a non-uniform electric field within the crystal, on the microscopic scale corresponding to the wavelengths used — while before the electrons were selectively moved from one trap to another, this field by the same criterion was substantially uniform, (e) The non-uniform field gives rise to corresponding changes in refractive index of the crystal, by the well-known electro-optic effect in the crystals used. Each refractive-index change is small; but as the effects of many of them add up together at (say) each bright spot on the picture read-out, the over-all picture that results is fairly good.

There is, of course, a basic difference between this crystal hologram and its photographic counterpart: whereas in the latter true opacity is produced at the interference-fringe antinodes, by absorption, in the former there is merely scattering. The result, though, is the same — except that the extra scattered light from the crystal store will give rise to a general illumination background on the picture read out, which is absent in the photographic version. There is clear evidence from the article[26] that the information storage on at least a few superimposed stored pictures will eventually be good enough, particularly by the digital methods that will guffice.

The crystal-store techniques are now only in their infancy. For example, though the shortest write-in time (for one hologram) so far reported, 35 sees at 1 watt/cm^2 optical input, would be short enough for certain permanent-store applications, in the majority of telecommunications cases it would be far too long. The authors' erasing method too, by heating the crystal to 170°C, would not be nearly fast or selective enough for most practical uses. But thermal erasing immediately suggests a much better alternative: to re-excite optically only those trapped electrons concerned with a single picture to be erased, by displaying it again via the write-in device, and then to re-establish statistical uniformity in spacings among these electrons by a suitable externally-applied electric field — probably in the time waveform of a few periods of a damped sine wave, as conventionally used for demagnetising a piece of iron, but at about 20 GHz.

In the experiments reported,[26] the authors used a cubic crystal of 1 cm side. For really good-qualit}' picture storage not more than from about 7 to about 50 pictures will usually be superimposed — though a much greater number is now suggested. When the write-in can be relatively slow, as in a permanent store, the over-riding limit to the number is set only by mutual picture deterioration due to scattering. In this case a single crystal cube of about 0.25 mm side could certainly store 7 superimposed frames of a 700-line TV picture in PCM form (7 X10^7 bits) when we can achieve an average trap density of about 10^{17} per cm^3, about half of them being filled by electrons (1 per trap). This is quite probably feasible, as the basic limit is of course as high as about 10^{23} per cm^3, the number of atoms in 1 cm^3 of the crystal. A word of caution, though: we *may* have to think again, both on this point and as to means for achieving uniform depths. Exact understanding of the theory of trapping, particularly in crystals of the lithium niobate type, is at present just beyond the borderline of scientific knowledge. This is true even as to the different kinds of trapping mechanisms involved, though many are known (e.g. the 'F-centres').

Apart from mutual interference between stored pictures, the above trap and filled-trap densities would enable 10^5 such frame-pictures to be super-imposed. But if we use only 7, we then store 7 X10^7 bits per cube, (which is 4.3 X10^{12} per cm^3). For access, if we have a slice (or slices) of the crystal 0.25 mm thick, and use this slice as if it were 14 000 separate cubes (each of 0.25 mm side), we can store a total of 10^{12} bits (i.e. one high-quality film in PCM form, running for one hour). The storage is in a volume of 0.9 cm^3.

For writing it in, we could use a single pair of laser beams. The first of each pair would be unmodulated, collimated to a circular cross-section of -J mm, projected so as to scan successively all the J-mm cubes, and be tilted at 7 or 8 different angles (in two planes) while stopping at each. The second would follow the first in position and moments of tilting, but be modulated (in suitable form) by serial PCM signals from the film being stored. To write-in each frame, at the usual (interlaced) speed of 25 frames per second, we can use a pulsed input optical peak power of 2.5 W from each laser (at a wavelength of 1 micron) — if we can achieve an average capture efficiency of as high as 10 per cent, of the writing-in photons by the trapped electrons at the density stated of 0.5 X10^{17} per cm^3. This is probably quite feasible; but it may take quite a bit of research to obtain it, as some types of trap can be said to be so electrostatically screened by outer electron shells that their capture efficiency would be much less than 10 per cent.

If however we need to write-in only a single 1 000-bit word, with this same capture efficiency of 10 per cent laser pulses of only 1 nanosecond duration are required, obtainable from an incident optical output (for 1 micron wavelength) of about 1 watt. There is a square law connecting the bits per frame and the needed energy to write it in. In this case, in the limit the collimated laser-beam cross-section would have to be about 35 microns — certainly obtainable eventually from gallium arsenide lasers, though not at present. To write-in a 1 000-bit word, the beam scanning could conveniently be done by using 10 stages of a binary-tilting unit, each comprising an electro-optic crystal in combination with a naturally birefringent material such as calcite.[13] Alternatively, the 10-stage tilting could be achieved in practice by a suitably placed group of 10 lasers, switched on one at a time, to replace the single pair of lasers referred to so far.

be a gallium arsenide laser preferably optically pumped, in integrated circuit form — a uniform film of gallium arsenide, between suitably reflecting layers, giving separate laser actions wherever a real image from a larger, pumping laser impinges on it. It is probable that at least 3 000 separately adjustable delays, each with 1 000 equal steps, will eventually be obtainable in this way in a volume of not more than 1 cubic metre.[12] These devices would be suitable for a rate of 10 to 20 gigabits/sec. An alternative is to write onto the gallium arsenide film a '!' or a '0' digit by the pumping pulse only of a similar optically pumped laser, at a level suitably below the starting threshold for lasing. The information would then be read out at the later time desired by a suitably timed optical signal pulse, of an energy and wave shape adapted to overcome the above starting threshold. This version would be smaller; but the storage time would be limited by the maximum spontaneous decay time of the gallium arsenide as a laser, at present about 10 nanoseconds. A third method would be a crystal store but the write-in time would then be slower, at the practical write-in powers, probably needing at least 10 pulse periods at a rate of 10 gigabits/sec. Still further methods are being examined.

For the space switches the use of optical carrier frequencies gives us a new dimension in facilities. A focussed laser beam, rapidly tilted in two planes by electro-optical means, takes the place of the mechanically moving arm of the old-type rotary switches. Several ways of achieving this have been studied experimentally.[13] Most such existing methods are capable also of handling at least a small group of TDM channels by each laser beam — but only at bit rates much lower than the total that will be used in our future optical channel groups and master-groups. A breakthrough in technique is needed. The key will be the use of purely optical gates of the AND and other varieties, much faster analogues of the well-known devices in the electronics art, in many examples producing focussed beams tilted at the correct angles. The most promising line so far lies in employing the nonlinear amplitude characteristics of gallium arsenide laser devices. One version has already been described.[14] Improved varieties are being studied, including also all-optical analogues of waveshape regenerators, pulse time-noise removers, etc.[15] This type of research is the beginning of a new technology; it is truly fascinating in its possibilities, for future still higher-speed computers as well.

An altogether different approach, now being studied, embodies a miniaturised optical version of SYNSOL described in the text. It will resemble cross-bar in its mode of operation, but will be more flexible, and within wide limits self-adaptive to changing traffic demands.

11.7 Holographic crystal stores

There are a number of lucid articles on the elements of holography. One of them has recently been published.[25] The basic principle is quite simple; it goes back to the Fresnel zone plate, often explained in elementary physics courses. If consists of an opaque plate, A, (Fig 6) having concentric annular transparent regions, with their radii and widths quadratically graded. If a spot of monochromatic light is placed at 0, with suitable distances for OD and DI, the path lengths OBI, OCI, ODI, etc., differ consecutively by the wavelength of the optical source. The zone plate therefore acts like a convex lens, producing at I a real image of the point source at 0. By a suitable coaxial combination of a parallel beam with one that is diverging, both floodlighting a sensitised surface, such a zone plate can be made photographically by the interference fringe produced — if the two beams are sufficiently space-coherent, e.g. by producing them from a single laser by conventional beamsplitting arrangements. The resulting graded opacity gives in fact a better result than when using the completely opaque-or-transparent method of Fig 6.

If we now use *n* diverging beams, from bright spots, again in combination with a single parallel beam, the result is *n* real images, in the configuration of the originating *n* sources. It is in this way that a hologram stores a picture — or a series of binary-digital words, in terms of bright or dark spots.

If the parallel writing-in beam is now tilted, a different picture can be stored. In this way, within limits, a fairly large number of pictures can be stored, each capable of being read-out separately, if for each picture the parallel readout beam is at the correct angle in relation to the hologram plane. If a parallel stack of *m* photographic plates is used for the write-in, *mn* separate pictures can be stored, provided that the resulting average total opacity of each plate is small enough to leave enough light for the plate that is furthest away in depth from the floodlighting source, and that the total scattered, defocussed light from the other images is not too large.

Although such photographic holography is reasonably efficient for a fairly permanent store, it cannot be used in the many telecommunication applications where the stored information has to be quickly erasable, to be replaced by an equally fast writing-in of new information. Neither, by present techniques, does it lend itself to efficient storing in three dimensions rather than two. Owing to an accidental discovery, however, a different method is now being explored by Bell Telephone Laboratories.[26] The authors of this article are duly cautious, suggesting only tentatively a theory to explain their highly interesting and important experimental results. But from the evidence they quote it seems to me that their explanation is at least highly plausible, even if not yet completely proved. The predictions that follow in this lecture concerning many useful applications of their method assume that at least their basic theory *is* correct. If this assumption proves to be wrong, we may have to think again — at least for the quickly changeable types of store. It is nearly certain that if necessary a good alternative technique *could* be found; so far we are only on the fringe of the many possibilities.

The method consists in replacing a stack of photographic plates or films by a crystal of an electro-optic material such as lithium niobate or lithium titinate. The split-beam writing-in wavelength used so far was 4 880 angstroms, which they also used for the single-beam read-out. The resulting solid hologram that was observed is believed to be produced as follows: (a) At the antinodes of the write-in 3-dimensional interference fringes, electrons captured by 'traps' in the crystal lattice receive enough extra energy to free them from these traps, (b) These freed electrons are then pulled through the crystal by an electric field, believed to be due to surface effects at the crystal edges, (c) The electron drift continues until each is captured by an adjacent trap (which the experiments have shown to have an average linear spacing of a small fraction of one micron), (d) The result is a non-uniform electric field within the crystal, on the microscopic scale corresponding to the wavelengths used — while before the electrons were selectively moved from one trap to another, this field by the same criterion was substantially uniform, (e) The non-uniform field gives rise to corresponding changes in refractive index of the crystal, by the well-known electro-optic effect in the crystals used. Each refractive-index change is small; but as the effects of many of them add up together at (say) each bright spot on the picture read-out, the over-all picture that results is fairly good.

There is, of course, a basic difference between this crystal hologram and its photographic counterpart: whereas in the latter true opacity is produced at the interference-fringe antinodes, by absorption, in the former there is merely scattering. The result, though, is the same — except that the extra scattered light from the crystal store will give rise to a general illumination background on the picture read out, which is absent in the photographic version. There is clear evidence from the article[28] that the information storage on at least a few superimposed stored pictures will eventually be good enough, particularly by the digital methods that will gufifice.

The crystal-store techniques are now only in their infancy. For example, though the shortest write-in time (for one hologram) so far reported, 35 sees at 1 watt/cm^2 optical input, would be short enough for certain permanent-store applications, in the majority of telecommunications cases it would be far too long. The authors' erasing method too, by heating the crystal to 170°C, would not be nearly fast or selective enough for most practical uses. But thermal erasing immediately suggests a much better alternative: to re-excite optically only those trapped electrons concerned with a single picture to be erased, by displaying it again via the write-in device, and then to re-establish statistical uniformity in spacings among these electrons by a suitable externally-applied electric field — probably in the time waveform of a few periods of a damped sine wave, as conventionally used for demagnetising a piece of iron, but at about 20 GHz.

In the experiments reported,[26] the authors used a cubic crystal of 1 cm side. For really good-quality picture storage not more than from about 7 to about 50 pictures will usually be superimposed — though a much greater number is now suggested. When the write-in can be relatively slow, as in a permanent store, the over-riding limit to the number is set only by mutual picture deterioration due to scattering. In this case a single crystal cube of about 0.25 mm side could certainly store 7 superimposed frames of a 700-line TV picture in PCM form (7 x 10^7 bits) when we can achieve an average trap density of about 10^{17} per cm^3, about half of them being filled by electrons (1 per trap). This is quite probably feasible, as the basic limit is of course as high as about 10^{23} per cm^3, the number of atoms in 1 cm^3 of the crystal. A word of caution, though: we *may* have to think again, both on this point and as to means for achieving uniform depths. Exact understanding of the theory of trapping, particularly in crystals of the lithium niobate type, is at present just beyond the borderline of scientific knowledge. This is true even as to the different kinds of trapping mechanisms involved, though many are known (e.g. the 'F-centres').

Apart from mutual interference between stored pictures, the above trap and filled-trap densities would enable 10^5 such frame-pictures to be super-imposed. But if we use only 7, we then store 7 x 10^7 bits per cube, (which is 4.3xl0^{12} per cm^3). For access, if we have a slice (or slices) of the crystal 0.25 mm thick, and use this slice as if it were 14 000 separate cubes (each of 0.25 mm side), we can store a total of 10^{12} bits (i.e. one high-quality film in PCM form, running for one hour). The storage is in a volume of 0.9 cm^3.

For writing it in, we could use a single pair of laser beams. The first of each pair would be unmodulated, collimated to a circular cross-section of ^ mm, projected so as to scan successively all the J-mm cubes, and be tilted at 7 or 8 different angles (in two planes) while stopping at each. The second would follow the first in position and moments of tilting, but be modulated (in suitable form) by serial PCM signals from the film being stored. To write-in each frame, at the usual (interlaced) speed of 25 frames per second, we can use a pulsed input optical peak power of 2.5 W from each laser (at a wavelength of 1 micron) — if we can achieve an average capture efficiency of as high as 10 per cent, of the writing-in photons by the trapped electrons at the density stated of 0.5 X10^{17} per cm^3. This is probably quite feasible; but it may take quite a bit of research to obtain it, as some types of trap can be said to be so electrostatically screened by outer electron shells that their capture efficiency would be much less than 10 per cent.

If however we need to write-in only a single 1 000-bit word, with this same capture efficiency of 10 per cent laser pulses of only 1 nanosecond duration are required, obtainable from an incident optical output (for 1 micron wavelength) of about 1 watt. There is a square law connecting the bits per frame and the needed energy to write it in. In this case, in the limit the collimated laser-beam cross-section would have to be about 35 microns — certainly obtainable eventually from gallium arsenide lasers, though not at present. To write-in a 1 000-bit word, the beam scanning could conveniently be done by using 10 stages of a binary-tilting unit, each comprising an electro-optic crystal in combination with a naturally birefringent material such as calcite.[13] Alternatively, the 10-stage tilting could be achieved in practice by a suitably placed group of 10 lasers, switched on one at a time, to replace the single pair of lasers referred to so far.

To maintain the read-out resolution when a hologram ' is used appreciably in depth, the read-out and write-in optical wavelengths must be substantially the same. This constraint blocks the way, unfortunately, from the almost complete immunity of the stored information from read-out beam deterioration that the use of two different wavelengths would achieve. Where this stored information has to be read out a large number of times before changing it, therefore, it will probably be necessary to use a second storage crystal. At suitable intervals each crystal in turn would be erased, and re-written from the other via a noise-removing PCM-type regenerative digital repeater. In this way the predicted storage time of at least 100 years should still be achievable however many read-outs are required. Some probable features of the crystal-store read-out devices are given in Appendix 11.12.

11.8 Calling system within buildings

We could use, for example, a TDM group of 1 000 channels, one for each occupant of the building (allowing for visitors). We could use, too, audio calling via each man's pocket-phone, the calling time being (say) 250 milliseconds, and give each occupant while he is in the building a temporary personal code of 10 binary digits, enough to identify each. If for each call these 10 digits were sent serially by 10 pulses having duty cycles of 30 per cent, each pulse would be 7.5 microseconds wide, needing an RF bandwidth of about 26 per cent if they are used to modulate a 1 MHz calling carrier by double sideband.

Each occupant would have a very small, integrated-circuit radio transmitter-receiver plugged into his pocket-phone, visitors being lent one on arrival. The receiver side would respond in the band from 875 kHz to 1 135 kHz. The transmitter side could use the same band, by interlaced TDM, enabling a single 1 MHz oscillator to be used in both directions, at the pocket-phones and the calling stations in the building. The channel synchronising at each end could use now-conventional methods. In practice, the simplest and smallest device to achieve the 2-way operation would be a self-quenching super-regenerative circuit. Contrary to popular belief, such circuits can easily be made quite reliable, by adding two negative-feedback loops. A higher-than-usual super-regenerative percentage bandwidth would be needed; but this too is easily feasible. The pocket-phone circuit can be designed to draw negligible power from its small storage cells except when actually transmitting or receiving an RF pulse, the mean drive power being thus reduced to about 0.2 per cent of its peak value. The pocket-phone transmission, giving the person's temporary building call-code, serves to up-date the reporting of his position.

There will be a fixed transmit-and-receive station in every office and at other suitable points, to enable every occupant to be always within a few metres of one of them. This provides the first of three means to prevent mutual interference between adjacent such buildings — for the total radio power emerging from each can and will be correspondingly small, from both the fixed and the pocket-phone transmitters. The receivers will have stabilised fixed gains. This first means also enables an additional method to be used: a maximum path length for possible interference that does not exceed about one radio-carrier wavelength, thus giving nearly an inverse-cube law of field strength with distance (as for the field of a permanent magnet), instead of the usual more penetrating inverse-square radius of action.

The third means for minimising the calling interference distance is, of course, the use of a digital pulse system. The low-level clippers on the receiving sides give a capture effect, as with FM methods. Attached to each fixed station in (and just outside) each building, there will be one or more plugs or sockets through which, when rung, the calls can be completed via the pocket-phones and the land networks. When any two buildings are very near together, the pair will have to be treated as a single unit.

11.9 Optical stores for global personal telephone numbers

To see the order of magnitude of the needed access to the store, assume that in any 160 000 km^2 area handled by such a storage centre there are a maximum of 50 million subscribers, and that at times half of them are away from their nominal, fixed subsets. Assume further that at times 10 per cent of the mobiles are on the motorways at a speed of 360 km/hr, by automated driving, needing a location-reporting rate of 0.1/sec — that 10 per cent are on other roads at 180 km/hr with a reporting rate of 0.05/sec, and the remainder needing an average reporting rate of only 0.01/sec because they are not in vehicles. The total average reporting rate in the area is thus about 6 x 10^s per second. Each incoming (optical) pulse group at the store need last only about 0.1 microsecond, so when combined in a single wideband fibre at the store-access input only 6 per cent of the time-space available need be used. If one special crystal is used for the mobile-handling hologram store, a single group of laser beams, movable by suitable techniques, can be used for its write-in and read-out — by the general method explained in Appendix 11.12, in this case only a small fraction of one hologram frame being used for each write-in or read-out.

11.10 The SYNSOL method for congested mobile networks

For cars, the (roof-mounted) aerials will usually have horizontal polar diagrams with equal major lobes 180° apart, to enable chain links to be formed via them in two opposite directions. Their normal orientations will be forward and backward. On motorways, and other roads in at least the second-grade category, the multichannel ground stations transducing the mobile channels to the fixed optical networks will be spaced not more than about 1 km apart on straight, line-of-sight sections — and suitably nearer together when bends or convex slopes make this necessary to preserve line-of-sight paths. They will normally be mounted on lamp-posts, or the equivalent, about 5 metres high. In more open areas that are sufficiently used to justify it, the ground transducers will be in the form of a 1 km spaced square matrix, where necessary adding further ground stations away from the roads in the region. Where such more open areas have a smaller normal-peak usage, the matrix spacings will be increased — in this case the 1 GHz band being often employed even without heavy-rain conditions. When off the communication-installed roads, the car aerials will either be rotated when necessary away from their normal fore-and-aft polar diagrams, or a second, all-round-looking aerial will be employed, as in the pocket-phones when used alone. Commercial security prevents the disclosure of further details of the Synsol scheme, particularly concerning technical strategy for the fastest way of forming new chain links that minimise the times needed for readjusting the chains that are already established.

11.11 Automatic access to a lower frequency band in heavy rain

The relevant circuit plan is explained in outline in general terms, not limited to its best form for SYNSOL. Consider any 2-way radio link between two stations A and B, the signals being, as an example, in the form of a carrier of high frequency modulated by pulses, as shown in Fig 7 (a). The carrier frequency could be about 60 GHz, the pulse durations T being 1 nanosecond. The carrier frequencies could be either the same, the channels being separated by a TDM method, or slightly different in the two directions. After rectification at either receiver, the base-band waveshape shown in Fig 7 (b) would be obtained.

In Fig 8, (1) is a receiving antenna giving input to the 60 GHz radio receiver (2), connected to the PCM decoder (3) and headphones (4). (5) is a 60 GHz oscillator, digitally modulated in suitable form by (6), and feeding the transmitting antenna (8) via the gate (7). (6) also digitally modulates a 1 GHz-band pulse generator (9) giving the waveform of 1 (b) which also feeds the antenna (8) via the gate (7) — but in this case via the time-differentiator (10) in addition. The gate (7) is biased so that in the presence of input from the wire (11). also of waveform 6 (b), it passes only the digitally modulated output from (5), cutting off the baseband output from (9). This is the normal situation, where there is no heavy rain, or the need for diffraction round an obstacle that the 1 GHz band can deal with but which the 60 GHz band cannot. When however there is heavy rain, or the need for such extra diffraction, the gate (7) will fail to open in the direction (5) to (8), but is arranged to open instead in the direction (9) to (8), thus causing the 1 GHz band itself to be transmitted (in the waveshape of 6 (c)), without any 60 GHz carrier. It is arranged that the bistability of gate (7) is assisted by the then-present baseband output of a second, 1 GHz-band radio receiver (13) and time-integrator (14), giving input into the line (11), (as well as into the PCM decoder (3)). The aerials (1) and (8) are both designed to pass the bands at 1 GHz, as well as at 60 GHz — and without inadmissable mutual deterioration of their characteristics. In another form, by 2-channel TDM the same aerial may be used for both transmitting and receiving. The bandwidth needed by the 1 GHz aerial is reduced by the time-differentiator (10).

Let the wave form Fig 7(b) be time-differentiated, becoming then the waveshape shown in Fig 7 (c), each element being a dipulse. The 7 (c) wave contains no do component, the amplitude of each harmonic in the Fourier Series representing it being reduced, in relation to its value in the waveshape of 7 (b), linearly and inversely in proportion to its harmonic frequency. The resulting spectrum will depend, of course, on the exact wave-shape of 7 (b), but in the general form shown the spectral energy per dipulse falls off fairly rapidly below frequencies equal to $^1/_a$!T.

11.12 Access to permanent crystal stores

A future solution will be to use a 1 200 X1 200 matrix of integrated-circuit gallium arsenide lasers for the hologram reference beams, only one laser being in operation at a time, and stepping from one to another to change the frames at a rate of 6 000 per second — 240 times the usual analogue TV frame speed. This would would cater for an 8-digit PCM code when allowing for 30 simultaneous groups of users, seeing the same film at 30 optional starting times, 2 minutes apart. The optical receiver is again a similar matrix of gallium arsenide integrated-circuit lasers, but containing about 700 x 700. Each reference-beam laser produces internally 10 picoseconds pulses at a repetition rate of about $6x10^9$ per second. They are all optically pumped (from a higher-power laser, doped with gallium phaphide to give the required optical frequency rise). There are no junctions, electrical drives, or physical connections, a laser point being formed wherever the focussed pumping impinges on a homogeneous thin film of suitable gallium arsenide, having Fabry-Perot sandwich reflectors on each side to increase the internal optical field strength by resonance.

By two binary chains of electro-optic crystals, one for each dimension across the matrix and a conventional logic unit, the pumping beam can be made to move in jerks, by any desired route until the whole matrix has been covered, each lasing point being scanned *n* times if the logic programme fits *n* simultaneous groups of users. The pumping beam, jumping at only 600 steps per second, allows time for about 10^6 ten-picosecond pulses at each stop. The operation is described diagrammatically in Fig 9. At the hologram pick-up end, each of the 700 x 700 matrix lasers receive in cyclic order up to 30 groups of pulse trains, each group representing a spot on a different frame, in PCM form. Each train is divided into 8 sub-groups, one for each digit of the PCM code. As weU as being scanned by a second pumping laser, each spot on this second matrix as sampled by a further single gallium arsenide laser giving $^x/_6$ nanosecond pulses, impinging as a focussed beam on to each spot in turn at a frame speed of 6 000 per second. In the pick-up matrix, each laser changes digitally according to a received '0' code or *'I'* code, (either in PRF phase or in optical phase), imparting a similar digital phase change to the single sampling laser. The free PRF's of the input pumping laser, the reference-beam matrix lasers, pick-up matrix lasers, and pick-up sampling laser are consecutively slightly lower, in that order. This causes the average PRF of the final sampling laser to lock into step with the reference-beam PRF, by a one-way action, while stiD responding to the digital signal modulations stored in the hologram. The total pulse rate and information rate of the sampling laser is thus 6 gigabits/sec suitable for transmission through a future single wide-band fibre, before being demultiplexed to the various groups of simultaneous film users. The diverging, write-in, pulsed signal beams are as in conventional holography. Their optical phases will be obtained, before signal modulation, from those of the corresponding reference beams.

VOTE OF THANKS

Professor W. Cormack (Past President): Mr President, Mr Vice Chancellor, Ladies and Gentlemen,

At the outset of his lecture tonight, Mr Reeves stated very clearly that he was looking only into the future. In thanking him on behalf of us all, I hope I shall be pardoned for starting in the opposite direction and having a quick look into the past.

We are told that one of the characteristics which distinguished early man from the other animals was his capacity for making tools. In the first instance these were probably weapons to supplement his own strength which was puny compared with that of some of the other animals. Soon the tools he devised needed more than his own strength — the early plough is probably a good example — and he turned to other sources of energy. The domestication of animals and probably the enslavement of his own kind provided the motive power he needed.

With the exception of developments in the harnessing of wind and water, in windmills and watermills, there was little change in the situation, until the invention of the steam engine in the eighteenth century. This opened the way to the exploitation of the vast store of energy contained in fossil fuels and made possible the construction and operation of machines quite unthought of in previous times.

The application in the nineteenth century of the principles of electricity to the distribution and control of this energy led to the development of the modern power system, which makes the energy available to all in both large or small quantities.

In case any are wondering whether I am giving the wrong vote of thanks to the right lecture, may I hasten to explain that I mention the power development because Dr Bernard Price was closely associated with the electricity supply industry. He was also closely associated with both the Institute on the one hand and the University on the other, and, as has been pointed out by the Vice Chancellor, it is most fitting that the joint meeting between the two bodies should be devoted to a lecture which commemorates his name.

Returning to the distant past, another characteristic which separated man from the other animals was his ability to communicate. A coded system of grunts and squeaks and possibly whistles became a language which when developed, permitted the most abstract thoughts to be expressed in words. We are told that systems of communication do exist between other living creatures — bees discovering honey can pass on the information to the hive by means of 'long and weary dances' — but these systems, although probably not clearly understood, appear to be relatively simple compared to that developed by mankind.

Not content with communication on a person to person basis, communication over distances became an aim. The development of writing not only allowed a permanent record to be kept but messages could be sent over distances greater than that reached by shouting. However, the time of transmission was slow and other methods, with more limited information capacity perhaps, were used as well, i.e. signal fires, smoke signals and in this country the mysterious drums of Africa.

The eighteenth century saw the establishment of a semaphore telegraph system in France, by which short messages could be transmitted a distance of some 150 miles in 13 minutes, weather permitting!

As with energy transmission, the principles of electricity were applied to the transmission of information. The invention of the electric telegraph, the telephone and the radio led to the development of communication networks such as we know them today. In this country direct dialling between the Witwatersrand and other centres has recently been introduced and, as a user of the system, my amazement at the rapidity with which I can be connected to my friends in Cape Town, is only equalled Mr President, by the frustration *I* sometimes experience trying to talk to those in Johannesburg.

Mr Reeves has now shown us the possible future development. He draws an amazing and possibly even frightening picture of what can be expected. Instant communication between any two persons anywhere in the world is the aim, and we have been told that this is not only quite feasible from a technical point of view, but can possibly be done at a reasonable cost. In this connection, however, I think we in this country should note one point, that:

> "the chief single expense item is a home subscribers terminal; the future good but cheap picture-type solid-state TV screen and associated information processor will be there in any case in nearly all developed areas, for entertainment"!
>
> Communication between persons, either from fixed stations or from mobile ones is not the only thing which will be accomplished but the system will serve to give access to information centres, thus doing away with libraries, and there is also the possibility of the complete remote control of machines.

The social implications of this technological development will be considerable. A former philosopher at this University once pointed out to me that in multi-storey buildings which had been constructed in France at the time of the Revolution, the best accommodation was always on the ground floor, and the worst on the top floor. In modern buildings, he went on, the best is very often at the top and this social inversion came about because of the development of lifts. This is a simple example of the impact of technological development on social life and I leave it to you to imagine the impact the new communication system is likely to have. The changes are unlikely to occur without the creation of problems and we must be grateful to Mr Reeves for introducing us to them as well as for outlining the technology which brings them into being.

I can only conclude then by thanking Mr Reeves on your behalf, not only for the lecture itself, but for the ideas he has loft with us and for the attitude which he urges us to take up with regard to the problems of the future, which is best summed up in his own words 'life would not be worth living if every rock face climbed, failed to show a still more difficult pitch ahead'.