

THE ENGINEERING OF PEGASUS

A talk given to the French Centre of the Institution of Electrical Engineers
on October 5th, 2006 by Ian Merry C.Eng. MIET

(Revised text of a paper originally delivered to a seminar of the Computer Conservation Society of Great Britain, at the Science Museum London, 21 May 1992 on the early computer developmental work of Elliott Brothers London, and the Ferranti development of Pegasus.)

I am uncertain whether to welcome this new opportunity to celebrate the engineering genius of Charles Owen and the conceptual brilliance of Christopher Strachey, or whether like pious Aeneas before Queen Dido of Carthage, I have been asked to awaken ancient and unutterable feelings of regret - regret that, following the outstandingly successful development of Pegasus, the design team was disbanded, and at least to my way of thinking, no worthy successor has ever been developed in Britain.

With this in mind, I will try to place the development of Pegasus in a historical context, and then identify a few of the individual characteristics of the members of the design team which had a significant influence on the design. [Slide 2](#)

First, this is what the original Pegasus looked like. A fully engineered but hand built machine, not a prototype, luxuriating under a restored Adam ceiling at 21 Portland Place. [Slide 3](#) At the desk mid left are the two photo-electric 5-hole fast photo-electric paper tape readers for input, on the desk nearer to the viewer are the output units: a fast 5-hole paper tape punch feeding a standard teleprinter as output media. The three tall cabinets house the logic plug-in packages, the drum-storage packages, the magnetic drum-store itself (partly visible at the base of the nearest cabinet), holding just over five thousand 42 digit words, the various AC and DC power supplies, the fans providing cooling air at the base of the cabinets and the engineers oscilloscopes in the pull-out drawer. It first ran programs from the drum store, with the doors etc closed, of course, on 23 March 1956, and by mid April was providing error free service to programmers and sales staff of up to 39 continuous running hours at a stretch. This was an unheard-of achievement at the time and I will try to explain how it was achieved.

I learnt from my experience with Pegasus that good design requires that:

First.

- + the convictions, even the prejudices of a single design authority require to be both articulated and respected
- + that accurate and extensive problem definition is a vital precursor to problemsolution - I would say a complete definition, but Dr Goedel and others have taught us to be wary of looking for completeness.

Secondly,

- + successful design depends upon solutions with 'designability', amenable to design analysis and calculation. Since design is always constrained by the limitations of materials, successful design goes, so to speak, with the grain.

These were, in effect, the precepts on which both Owen and Strachey based their work. They were not, however, at all typical of the world of electronics of the 1950s.

What was this world like? The enormous impact of radar development during World War II was still much in evidence. A major advance in glass technology dating from 1938 with the appearance of the Philips EF50 valve had fostered a succession of many moderately high -

gain miniature vacuum tubes only about two centimetres in diameter and no more than about a couple of inches in height.

Point contact germanium diodes had been developed during the war as radar demodulators, based on little more than the kitchen science of the 'crystal' wireless-sets of the 1920s, and with the invention in 1947 of the point contact transistor, a hesitant semiconductor industry had arisen, which concentrated on germanium semi-conductor technology, was hampered by the variability of point contact devices, and was only marginally familiar with the technology of silicon and junction devices on which we all now depend. Electronics was still almost entirely the concern of telecommunications and broadcasting. The former with its 18 or 22 inch wide racks of equipment, six or eight feet high and the latter with racks as wide as 24 inches - at the BBC at least - bearing monolithic cadmium-plated steel chassis each weighing tens of pounds with a dozen or more electronic valves or vacuum tubes and associated circuits.

Significantly, most of those involved in the wartime radar development had been graduates in physics without academic engineering backgrounds since, in Britain at least, electronic engineering was widely regarded as a dilettante not to say insecure profession until well into the 1950s. In consequence there were not too many about. Blumlein, the genius who had devised the BBC television video waveform in the mid 1930s was a very rare bird indeed.

I believe it is true to say that Elliott Brothers was, from the late 1940s for about a decade, the only British Company concerned wholly with Electronics.

This circumstance together with wartime pressures had confused the concept of design with the narrower field of circuit design. and established a widespread design tradition of 'suck it and see' whenever a problem arose outside the immediate experience of the designer. Again, while the logical power of electronic digital computation had been clear since the early 1940s to the indoctrinates at Bletchley Park, it was the domain of a smallish band of mainly academic successors, among whom few had studied design as an engineering discipline. We were all, and I believe we still are all extremely fortunate that although a German engineer, Konrad Zuse had foreseen most of what could be done with valve driven electronic computers in 1938/39 and had applied widely for patents in Germany and at least the UK, as I learnt to my astonishment when the Custodian of Enemy Property released the UK patents in 1957 or 58. being someone of moral stature he had made no attempt to bring these to the notice of the Nazi authorities and had, so to speak, kept his head down throughout the war.

Even where, as in the BBC Engineering Departments, numerically-based design was given its full due. on the electronic front this was in the context of high quality analogue audio frequency circuits, made linear with a quasi-statutory 40 decibels of negative feedback, or in the context of radio frequencies, with tuned circuits with a pass band, even for television, not greater than say a seventh of the mid frequency, adjusted by hand and RF instrumentation. Consequently, there was no pressure on the makers of thermionic valves or semiconductor-devices to publish the variances of their device-parameter data. To the contrary, it was rarely if ever clear whether published data represented design targets or achieved median values.

Lastly, despite the marketing 'hype' of the Festival of Britain in 1951, and the governmental initiative of NRDC, the National Development Corporation, there was no longer the recognition of engineering as an important aspect of the British Raj such as existed in the 19th and early 20th centuries, when it was natural for consulting engineers to establish themselves in Victoria Street in order to exert both proper and improper influence on parliament and the Ministries.

Turning now to the major players in the engineering design of Pegasus, these were four in number:

[perhaps here I should say that when this paper was first delivered, I inadvertently upset one or two of the older listeners by not including the chief logical designer C H Devonald whose work was of the highest quality, and who made a major personal contribution, for instance in the design of the multiplier logic. But this was in view of my remit to discuss the engineering design, and only these four individuals had a major degree of autonomy in their own field].

They were:

- 1 Charles Owen
- 2 Christopher Strachey
- 3 Brian Maudsley
- 4 and myself, the only survivor. [Slide 4](#)

Charles Owen had already had a major impact on the restructuring of complex electronic devices through the development of the circuit design and layout techniques required by the then fairly recent concept of small packages. Here I wish to pay tribute to his engineering sagacity, as he was the design architect in chief of Pegasus.

He had the essential vision of a successful engineer, which is to have formed an architectural concept of the finished work from the earliest possible moment; changes to that concept could be made as the design progressed, but any change had to be demonstrably beneficial, and to meet Charles' exacting standards of acceptability.

Given conformity with these standards there was, I found, no need for further persuasion. It was stimulating to work with someone who did not reinforce his prejudices with anything but logic and good sense. In addition, Charles was very ready to pass on his own knowledge, which endeared him to me as I had had no previous experience of digital techniques. With all of that, and though often doggishly witty, Charles was essentially a plain man to whom facts were facts and fancies were fancies.

[Slide 5](#)

Christopher Strachey, on the other hand was a modern 'renaissance man'. Beside his achievements as a mathematical logician - he founded in a later sphere as a prof at Oxford the study of denotational semantics, as he said to counteract what he stigmatised as the "intellectual disreputability of the profession of computing" he was a talented musician - he and his sister and I would sing through Mozart operas to his piano accompaniment - and his skill in technical discussion or general conversation was such as to make everyone else present perform beyond their usual level, as a consequence of his own rather competitive verbal brilliance. The Strachey family had already made its mark on British society, one of his uncles was Lytton Strachey the essayist and the first published biographer of famous individuals. He did not mince his words about his subjects defects. Christopher Strachey's father had been a first World War naval code-breaker, continuing at Bletchley Park throughout WW2. and later finally achieving recognition in a Times obituary when secrecy about Bletchley was just beginning to diminish. Another uncle had been a rather less successful Labour Cabinet Minister.

The third principal member of the design team, Brian Maudsley, of whom sadly I have no photograph was unusual in both background and in personality. A mechanical engineer from Ferranti Edinburgh, fallen, so to speak among intellectuals, he held his own in consequence of

a unique capacity for mechanical innovation, and for mild-mannered unarguable assertions - all supported by a physical stature of six feet eight inches.

As for myself. I joined Ferranti from four and a half years in the BBC Engineering Research Department, where I had worked on a number of electronic and electro-mechanical projects connected with both disk and magnetic recording, and where I had encountered many of the problems that were then still at issue in the Pegasus project. My only previous experience with digital circuits had involved telephone relay switches - of a type that in those days we called Post Office relays - a brief encounter leaving me with no yearnings for further involvement with relay switching.

We come now to the design principles on which the architecture and design of Pegasus were based.

Strachey's major objectives were to reduce the labour of the programmer, and to achieve computation speed through the avoidance of bottle-necks in translating the orders into logical operations.especially:

- * by providing an efficient and consistent order code
- * by freeing the programmer from undue concern with the machine architecture
- * by encouraging a balanced design among the various machine functions
- * by maintaining an autonomous invigilation of all machine functions using odd-parity checking throughout

Owen, for his part, had an intense pre-occupation with machine reliability and availability. His experience led him to believe that while this required conservative design, with care this did not get in the way of achieving elegance and economy, both he and Strachey aiming to build all of the complex control functions without recourse to special purpose circuits. With both prudence and modesty, he took the view that the basic circuit elements used in the earlier Elliott Brothers 401, which had been developed under an NRDC contract, represented the soundest basis for progress. Packages containing several configurations of point contact diode AND gates logically OR'd followed by a cathode-follower direct output, an inverter or a simple pulse amplifier-, retiming- and delay- circuit furnished the logical armoury for the bit level logic; while single word packages using a nickel wire acoustic delay line [Slide 6](#), as the serial storage medium provided immediate access memory for accumulators and registers. In passing, I will recall that the junior engineer responsible for the design of the delay line package John Fairclough crowned his career as Margaret Thatcher's principal scientific adviser, the late Sir John Fairclough.

It is fairly obvious that the more complex a logical function, the more numerous are likely to be the various inputs. In consequence one requires that the number of inputs to individual AND gates should be as little restricted by circuit component deficiencies as is prudent. In upgrading the Elliott Brothers logic circuits, Owen instituted a statistical analysis of the problem, having ascertained the variance of the germanium diode back-leakage resistance, and safely maximised his gate widths, avoiding the Scylla of AND gates with too many leaky diodes exhibiting pattern-sensitive failings and the Charybdis of overly restrictive gate design. At this point, with more time at our disposal than I had in 1992, [Slide 7](#) I have taken the liberty of copying some of Charles Owen's paper of 1956, to show the basic waveforms of the logic circuits, and the structure of the AND gates; and I quote:

>quote

"The timing of digit pulses throughout the machine is controlled by a common 'clock' waveform - a 3 microsecond square-wave in which the positive-going portions define digit positions. [Figure a on slide 7](#).

The digit pulses are generally of the form in [Figure b on slide 7](#); as generated they have their leading edges well in advance of the clock pulse and are of greater amplitude. Considerable distortion of the digit pulse is tolerable since only the portion which coincides with the positive clock pulse is of consequence. Digital pulse trains are 'clocked' (that is to say 'anded' with clock the waveform) only at their entry into a storage system or into a digit-delay circuit.

Inverted pulses are also employed.

Consider the operation A AND-NOT B In the next lower two figures, Pulses A and B are on are of the same nominal timing. Pulse B is inverted forming NOT B and is used to gate pulse A and to prevent its passage. The inverted pulse B will be a little late on B, which may also have been a later than A . Thus when A and NOTB are anded together, a spike may be produced as shown on the lowest line. This spike lies between clock pulses and so will be rejected on clocking.

The maximum number of operations performed in cascade without any loss in nominal timing in Pegasus is five.

Each of the logical packages has more than one circuit unit; A circuit unit is defined as that part of a package which has input and output pins, and no connections to other parts of the package other than supplies.

We may make the following generalisations for the logic units:

Each unit has an AND gate at its input

Each unit has a cathode-follower output

Each unit has an additional output via a germanium diode for making OR gate connections."

>end quote.

and here I should say for those unfamiliar with valve circuitry that a cathode follower is a triode based circuit providing nearly unity voltage gain from a source of high impedance at the grid input to a source of low impedance, at the cathode output. - invented, I believe by the aforementioned Blumlein.

I quote again:

"There are three possibilities for the part of the circuit unit between the input AND gate and the output cathode follower, namely a digit delay, an inverter and a direct connection. [Slide 9](#) The circuit shown is a digit delay with its typical waveforms" End quote.

This is a digit delay with either a three input AND gate or with two such input gates ORed together. Two outputs are provided and the intervening circuit provides an capacitor/inductance delay resulting in a fully clocked output delayed by one pulse time.

Jumping ahead a little, I would add parenthetically that in my part-of-ship, and before my arrival on the scene, Charles Owen had envisaged a nickel plated drum similar to those favoured by the Ferranti Computer Department in Manchester, with a minimal complement of read and write amplifiers using Post Office relays to access individually selected read/write heads. The binary sequences were to be recorded using the Manchester phase-encoding technique, whereby a series of 'ones' was a more or less square wave of magnetisation in

opposite phase to that for a series of 'noughts'; and alternate 'ones' and Slide 10 'noughts' were a squarish wave of twice the period - as shown in the diagram.

In considering the remarkable success in achieving all the design objectives, I would ask you to remember that Pegasus is a serial machine in which the 39 working bits of each word arrive sequentially at any point, or as we now say, at every interface in the machine, and to maintain the economic advantages of this serial approach, interface width has to be kept to a minimum, nearly always only one bit wide. The parts of the machine where static registers hold a number of bits of a word concurrently are thus few in number.

The thinking required in the logical design particularly of the control functions therefore requires the logical designer to envision successive machine states represented by circuit states changing autonomously and much faster, under the inexorable flow of serial data. A duality which is difficult to represent perspicuously in any diagrammatic form, and utterly beyond the descriptive mathematical techniques of the time.

In this regard, I well remember the seminar when the logical designers first gave an explanation of the control architecture. By then, Pegasus was in active use, and the logical design seemed to be consistent, but for my part at least, the description of the various control cycles remained and remains rather baffling.

Coming now to the engineering problems which had to be solved, we can consider them to arise in regard to three design aims:

- * reliability
- * economy
- * performance

And here I should say that among these, from the time of the launch of Pegasus, reliability was what made Pegasus unique. According to the historian of computing, Professor Simon Lavington, for the 8 weeks to 23 March 1957, the fourth production Pegasus returned a figure of just over 99% reliability, and he goes on to say "such evidence astounded users of other contemporary UK computers".

It is true that old programmers still remember the subtleties of the Pegasus order code and its advances on the address modification arrangements devised for the Manchester University machines and used in the Ferranti Mark 1 and Mark 1 Star machines, but I think you would agree that in general buying decisions are made by senior individuals and not programmers. In the mid to late 1950s no previous computers had distinguished themselves by their reliability and as I recall it, it was not until the launch of the IBM 360 series in the early 1960s that computers achieved full and general acceptance among the buyers of industry and commerce.

So reliability...

Past experience had shown that the major areas of transient unreliability were:

- * "pattern sensitivity" of individual logic or storage units, where correct operation fails with the input of some particular sequence of binary digits. This is of course a fault particularly of serial computers
- * drum systems generally
- * package plug and socket electrical contact instability

I have already explained how pattern sensitivity was excluded by Owen's study of gate architecture. Happily, by this time, other electrical component deficiencies did not present as causes of transient problems, and were adequate in terms of operational life given that they were not sourced from suspect quarters such as Government war surplus - a false economy which had bedevilled some other early projects, particularly in Academia.

Except to a mechanical engineer of particular discernment, there is little intellectual stimulus in addressing the problem of erratic plug and socket behaviour - which is why the problem hung around for so long -. In many ways this was the most dangerous of the packaged-circuit problems, since so many package interfaces were at risk. Maudsley tackled and solved it with determination. To keep the contact-resistance of each contact adequately low and stable he insisted that it was insufficient to rely on the comparative incorruptability of noble metal surfaces, and that on each insertion of a package every female contact should scrape its corresponding male and ensure a new metal to metal interface. Given an adequate thickness of noble metal plating of the male contacts, a quite adequate if limited number of insertions could be made. The in-line multi-contact socket used had the necessary female contact geometry with the extra advantage of permitting the easy replacement of a damaged individual socket contact and was engineered to his specification by one of the leading plug and socket suppliers as a modification to one of their standard units.

A concomitant to this solution was the provision of a robust, and stable mounting for package board and socket; this was achieved most economically by the use of aluminium alloy die-castings for the package shelf mountings.

Lastly, the package board itself had to be of an adequate rigidity and stability. Attention to such details had not distinguished previous computer projects.

My own part-of-ship, the drum system, appeared to present problems in every possible area. The geometry of magnetic fields causes the read signal to be rapidly attenuated with an increase in the read head/recording medium gap. The functional relationship had been established by a research worker in Holland while I was at the BBC, and is shown in the diagram. The read signal of a regular waveform is proportional to $\frac{e^{-2\pi d/\lambda}}{d^2}$ to the minus two pi d over lambda, where d is the separation of the head from the magnetic medium and lambda is the wavelength of the recorded signal . This results in 56 decibels loss per wavelength separation or about half the signal lost in geometrical progression for each extra tenth of a bit-length gap.

With a drum diameter of ten inches, and 128 forty two bit words per track, it turns out that an over-all variation in signal of 2 : 1 ensues if the radial run-out of the drum surface due to machining tolerances, variations in coating thickness and bearing shake is no more than about half a thousandth of an inch, and exponentially pro rata. Owen had already decided this was too demanding by a factor of two, and had laid down that alternate bits would be recorded in pairs of tracks, with the write and read diplexing handled by machine logic, a stratagem previously adopted at Elliott's.

However even a total run-out tolerance of half a thou, appeared to have caused difficulties with previous bearing design and bearing life, and there was some disquiet in connexion with the drum under development at Ferranti in Manchester for the forthcoming Mercury computer.

At the BBC, while studying the problems of television magnetic recording I had tackled the problem in the lab by using a horizontally mounted narrow drum or disk, with a little known 360 degree bearing design. This had to be hand lapped to a radial consistency of about one tenth of a thou and the drum surface and its magnetic coating was then sapphire turned. A 360

degree bearing is nominally permanently lubricated and works best with a lubricant like sperm oil, which I introduced with an eye-dropper.

Maudsley very properly would have none of this quasi-magical 19th century engineering, and quickly came up with the solution. Precision boring spindles had to perform with a radial run-out better than a tenth of a thou without adjustment, over a period of many months. Why not approach a manufacturer of such machine tools and give him the problem; and that is what we did. Bill Boneham of Boneham and Turner in Mansfield cheerfully undertook to make a suitable drum mechanism if we provided him with an outline design and details of the motor to be incorporated. The final outcome [Slide 11](#) is shown in the diagram. Boneham entertained no taboos about putting three bearings on one shaft, and he assured constancy in bearing behaviour by using twin sets of angular contact ball bearings under considerable axial pressure.

As I recall it, the original prototype which we used in the first Pegasus cost only £300, say what? £7000 in today's money excluding the electric motor. Please note however that the drum fitted to the machine in the Science Museum is a later and larger version, built on precisely the same principles but with a larger drum.

Other Ferranti drums were run as slaves to the rest of machine, synchronised in a phase-lock with the crystal controlled logic circuits, requiring an elegant servo-system, culminating in a hefty valve power amplifier. This was, in my view, to turn good design on its head, making a cumbersome object slave to a much more pliant system. Our drum needed over a quarter horsepower, making this scheme doubly unattractive, before the availability of power driver transistors. I therefore got Owen to agree that the logic should be driven from a clock track recorded on the drum. The rotational speed of the drum could then be kept within the limits required by the delay lines, using a simple servo loop and a crystal reference. As the drum drive had to be at 150 hertz for the 4000 rpm rotational speed, this loop included the motor excitation of a 150 herz alternator set.

Recording of the clock track was accomplished by first recording an approximately correct but incomplete clock track, which was then very temporarily phase-locked by hand adjustment of the servo, at the correct rotational speed, to a crystal controlled reference using an expanded trace oscilloscope. A burst of clock frequency of the correct length to close on itself was then recorded.

Ferranti drums had previously been nickel-coated, the low coercivity requiring little power output from the write amplifiers, however these coatings had occasional magnetic weak or dead spots, owing, I believe, to stresses in the plating. IBM had similar problems at about the same time with their 650 drums, and eventually overcame them with weird chemicals in the plating baths - Mississippi mud was how one IBM development engineer described it to me in 1957. At the BBC however, I had found that a red iron oxide magnetic oxide dispersion for coating magnetic tape was readily spray painted on to a drum, so we abandoned nickel for red iron oxide. Split rings of low loss ferrite had also worked well as read and write heads up to well over 500 kilohertz, needing only pairs of miniature power pentodes to drive the magnetic oxide to saturation. Better still, correctly formed ferrite recording head blanks had recently come on the scene, and I was fortunate in finding a sub-contractor, Epsilon Ltd, making multi-stacked heads for tape recorders, willing to package banks of ferrite based low impedance heads to my requirements.

Drum system performance and usability was greatly increased by abandoning relay head switching and developing electronic cross-bar switches for writing and reading. These are shown in the two next diagrams but I must refer you to the IEE 1956 Proceedings Volume 103 Part B supplement 2 for a detailed explanation. The write switch was easy, the idea

Slide 13. originated in a technical journal that I cannot now identify, but required tidying with two extra diodes constraining the voltage swing, to limit cross-talk. The read switch using point-contact germanium diodes came before any amplification, Slide 14 was my own and was entirely novel.

Charles Owen had to be convinced by a test lasting several months with a random batch of diodes, that diode noise would not cause errors. However diode noise remained below 250 microvolt, harmless to the unamplified phase modulated signal from the read head of some 3 or 4 millivolt. These switches meant that relay switch settling times of 20 milliseconds were avoided; in fact the read amplifier recovery-after-writing time of about 500 microseconds became the limiting factor while track switching took only about half that time. This both improved performance where more than a track's worth of data was in play, and enabled single word as well as eight word block transfers to be efficient.

Lastly, block addresses in the drum address track were permuted to leave two blocks between blocks with successive block addresses. Within any given block word addresses remain in natural sequence. This gave time for some computation between successive blocks without involving the programmer in fancy addressing of his data, a fearful technique then called "optimum programming", in general use for other drum systems.

In consequence of these innovations, the drum system became uniquely in harmony with the general approach that distinguished the Strachey/Owen design. That of design balance.

There is, I am afraid, only time to mention two other features of the logical design that contributed to overall performance. These were:

- * the provision of multiple accumulators and many more immediate access registers
- * the incorporation in the order code of a comprehensive and logically regular method of handling address modification, using these multiple accumulators and registers

How much more effective would the Pegasus' contemporary the IBM 650 have been with these features, designed as it was with a single accumulator for similar user areas. How sad that Pegasus could not have been equally widely exploited.

That concludes my survey of Pegasus development, except to say as I wrote in the Computer Journal of June 1991 : "It is a matter of record that all of these features were working in the Pegasus pilot by April 1956." Developments begun in 1955 when I purchased two magnetic tape units in Los Angeles added magnetic tape peripherals and punched cards were also added, leading to machine sales of about 40 machines overall, compared to the IBM 650's fifteen hundred or so.

Regardless of these good beginnings, an evident loss of focus on the part of the Ferranti senior management coupled with NRDC's short term financial pre-occupations fostered an atmosphere in which by 1956, the burden of continuing investment was only acceptable at a level requiring a fundamental choice between the Mercury team in Manchester, and the Pegasus team in London.

Few remember 'Mercury' now, but except for the peripheral developments just mentioned, the Pegasus team was largely disbanded, and staff were redirected to work on Ferranti contract and Defence work, or in the case of some of the leading team members, including Charles Owen and myself regrouped under American auspices from September 1956.

This was a blow to the infant British computer industry at a most crucial time from which subsequent events have shown it never wholly recovered, exemplifying how inadequate investment ensures a nett and permanently enduring loss....

The history of Pegasus has always seemed to me to be a paradigm of the British industrial malady, not just the shibboleth of the past 50 years, that Britain is not good at marketing its wares, but more fundamentally, that we no longer recognise or foster and therefore cannot exploit our very real strengths.

In conclusion I thanked the Science Museum and particularly Tony Sale - who went on to ensure the continuing existence of Bletchley Park as a museum for both the art of cryptanalysis and for the history of electronic computing - for the opportunity of making a contribution in celebration of this particular technological Pyrrhic victory. My thanks were also due to the late Professor Elliott for his visionary leadership in assembling and managing the team in which it was a joy to work.