

Ferranti Ltd. and minicomputers.

General background.

The electrical engineering company Ferranti Ltd. was founded in 1882 by Sebastian Ziani de Ferranti, an inventor who was born in Liverpool. Ferranti Ltd. soon achieved substantial success in the fields of electrical generation and supply, transformers and electricity meters. This was followed by work on electrical measuring instruments and, in due course, wireless technology and radar.

In 1948, when the government was looking for a company to produce a re-engineered production version of a pioneering digital computer produced at Manchester University, Ferranti was given the contract. This led to the Ferranti Mark I which, when it was delivered in February 1951, was the world's first production computer to have been installed at a users' site. Throughout the 1950s and 1960s Ferranti Ltd. continued to build computers, initially aimed at the scientific and engineering sectors of the market. These mainframe machines are described in other sections of the *Our Computer Heritage* site.

The mainframe Ferranti Computer Department merged with ICL in 1963 but the process-control activities continued under the Ferranti name at various locations including Bracknell and Wythenshawe. Ferranti continued to manufacture smaller computers for industrial process control and for real-time military applications, particularly the Argus range as described in the new contribution – based on a paper supplied by John Steele, who worked for Ferranti from June 1963 until about 1969 - below. The company began to lose its separate identity in 1987, owing to legal and financial difficulties over the purchase of an American defence company. Ferranti was forced into bankruptcy in 1993 and various computer-related parts of the business were acquired by GEC-Marconi and, in due course, by BAE Systems.

Simon Lavington May 2015, amended 2nd March 2016.

Ferranti Argus Computer – Evolution & Architecture

1 Argus - later renamed Argus 200

This system was originally developed prior to 1963 as the launch control computer for the Bloodhound Mark 2 anti-aircraft missile.

The renamed Argus 200 computer was expensive and had limited program and data memory but its architecture made it faster than Argus 100 (see section 2 below).

Technology	Germanium transistor
Construction	Standardised logic assemblies installed into a card cage. The card cages were rack mounted into standard 19-inch racks, 5ft. high ¹ .

The logic basis for the design was NOR logic gates.

Word length 12 bit data, 24 bit program.

Program store Ferrite pegs plugged into holes in a tray organised as 64 rows of 24 bits each plus 3(?) parity bits. Trays were installed horizontally into a rack mounted arrangement.

The trays contained horizontal loops, one around each row (program instruction), with vertical loops to read the data. Presence of a peg indicated a 1 bit, absence indicated a 0 bit.

Working storage 1024 12 bit words

Internal architecture Serial/parallel 2 bits at a time

1.1 Initial Commercial process control applications

A personal note

Prior to 1963 an Argus computer was installed at ICI Fleetwood and was used to control a Soda Ash plant. It was widely thought (believed?) in Ferranti that this was the World's first true on-line process control application². It was experimental at the time and ran for a few years until the plant was closed down.

The only other commercial Argus installation that I recall (and I did visit this site during my first 2 weeks at Ferranti while acceptance trials were running) was at West Thurrock power station. Installation took place in 1963. In this case the application was to control the start-up and shut-down of one of the five turbines and to monitor boiler temperatures.

These limited industrial applications led directly to the development of the Process Control Division at Ferranti and the development of a range of computers specifically designed for this purpose².

¹ As far as I can recall. ² Perhaps it should be noted that Elliott introduced process control systems in 1959

2. Ferranti Argus 100, 300, 400 and 500 systems

2.1 Evolution

The Argus 100, 300, 400 and 500 computers were a family starting with the Argus 100. This was a machine developed using the experience gained from the commercial industrial applications of the original Argus computer (renamed as Argus 200) as an industrial process control computer.

Process Control v. Data Processing

Process Control applications, when compared with Data Processing, are characterized by a need to address a large number of external devices. In the 1960s, Data Processing required only paper tape and card readers as input devices and teleprinters and line printers as output devices, together with some form of bulk storage devices such as magnetic drum/disk and tapes. Process Control, however, required these devices plus the ability to directly access a range of plant sensors and controls, such as :-

1. Temperature sensors (thermocouples) read in as voltage through an A/D converter. Many plants had hundreds to thousands of these.
2. Contact closure sensors read in as single-bit inputs. These were often provided as a spare set of contacts on an actuator.
3. Single bit outputs to control relays or light indicator lamps.
4. Analogue outputs to act as inputs into servomechanisms controlling the plant.

For economy a single storage technology was used for both program and data stores. Non-volatile memory was desirable so that the computer would hold program and data in the event of a power failure. Core store was therefore the technology chosen.

It was logical to build on the Argus 200 experience and the same germanium transistor logic boards and card cages developed for the earlier Argus 200 were used for the Argus 100 and 300 computers.

2.2 Argus 100

The 12-bit word length of the original Argus 200 was extended to 24 bits and the first computer in the family, the Argus 100, was implemented as a serial computer. The instruction set used was a subset of 30 instructions derived from the Ferranti Pegasus computer designed a few years earlier. The first production machine went to Jodrell Bank to control its Mark 2 radio telescope in 1963. It is interesting to note that an Argus 400 computer was later supplied to Jodrell Bank to retrofit into its Mark 1 telescope. That computer was last seen in the Jodrell Bank Museum and was personally commissioned by myself.

The Argus 100 computer was (I believe) designed by David Senior and Mike Eyres with assistance from Stan Redshaw.

Even as the Argus 100 computer was going into production it was realised that it was perhaps a little slow and that an alternative offering was needed. Development of a parallel version was started in 1963/4. This was the Argus 300.

2.3 Argus 100, 300 construction

The basic unit was a plug-in printed circuit board with a gold-plated edge connector. These were known as a 'package', and each was roughly the same size as a single Eurocard. Each card had a grey plastic handle at the front so the card could easily be removed. The circuit cards were single sided and all produced in-house from layout through bare board manufacture to component assembly.

The logic cards plugged into a card cage and used soldered wire connections at the rear to form the interconnections between logic cards in the same card cage. A Ferranti standard pink wire was used throughout. As the insulation was PVC great care was required when making connections, particularly for modifications, as the heat radiated from a soldering iron would melt nearby insulation if one lingered too long!

To enable card cages to be pre-wired on the bench connections which passed between card cages went to wire wrap posts with one row at the top and one at the bottom. These could be laid flat for bench assembly to make access easier.

Inter cage wiring in the racks was made to wire wrap posts which were a static fixture in the rack.

When a card cage was installed in the rack a "U" link wire wrap completed the connection. This made it possible to remove a complete card cage in a few minutes for modification. Depending on where the wires were and how many changes were required a skilled wireman would often choose to make changes in situ without removing the card cage. Removal and reinsertion, although easy, took between 30 minutes to an hour. Removal was quick – a pair of side cutters would be used to snip all the U links. The time consuming task was individually unwrapping each of the wraps from the posts. This was a chore often left to the designer who had asked for the modification! There was a good relationship between the skilled craftsmen and the designers that made such things possible. There was not much evidence of job demarcation.

The logic cards were based on diode transistor technology. By convention logic one was ground and logic zero was -6 volts. This was very logical in the process control world where a "one" could be used with an open collector power driver to switch on an indicator lamp or close a relay. The logic function we used throughout was NOR. As the outputs were single ended it was possible to connect two outputs together producing a logical OR function. Connecting two standard outputs together however reduced the signal driving capability so we had open collector versions of the logic gates to be used for this purpose.

As logic designers we had a choice of configurations. These varied according to the number of inputs into a single gate. From memory there were four outputs and up to

12 inputs. The design was common and the final assembly selected which combination was supported by the circuit board. It was never done but it would have been possible to reconfigure a standard card to a different combination.

Registers were built up with flip-flop packages. These were latched on the clock and I think both phases of input were needed. We had two forms available. A "Triple" which would latch data from 3 different sources depending on which clock signal was used and a "Single" which provided 4 flip-flops on a single card.

The core store electronics was built using the same technology but required a double card cage. Circuit boards containing the core matrix drive circuits and the read amplifiers were fitted at the front. The actual core store matrix was mounted on the back.

The core store was designed as a 6 microsecond access time module.

Argus 100 and 300

Technology Germanium transistor (Argus 100, 300)

Construction Standardised logic card assemblies installed into a card cage.

Most of these logic cards were identical to those used for the Argus 200 computer.

Card cages were rack-mounted into standard 19-inch racks. The main computer used 3-foot high racks with a Formica topped work surface. The control console was mounted on the work surface.

Logic Design based on NOR logic

Architecture common to models 100, 300:

Word length 24 bit program and data

Storage 6 microsecond core store (Argus 100) organised in modules of 4096 24 bit words. In fact the matrix had 26 bit planes with the intention that the two additional bits would be used for parity. It was however discovered that the reliability of the parity generation/decode was less than the core store and the parity logic, although designed, was never installed. The two additional planes however did provide a useful spare that was used at least once.

The Argus 300 had 2 microsecond core store modules.

Working storage 1 to three modules of 4096 * 24 bit bit words or 12 to 36 kilobytes

Internal Architecture Large I/O address space (0-7777 octal)
Memory space 10000-37777 octal

7 Accumulators held in main storage (octal addresses 10001 to 10007)

The technology and construction of the Argus 400 and 500 is covered in detail below.

2.4 Argus 400 – New integrated circuit technology

In parallel with the development of the germanium transistor based components the whole World was meanwhile investigating the benefits of silicon solid-state technology. Ferranti had developed a rather weird Diode/Transistor logic family called Micronor 1. This suffered from too many deficiencies to be useful and a parallel development of a new integrated circuit family, Micronor 2, and a computer to exploit it, the Argus 400, was started in 1963. This development was just starting as I joined Ferranti in June 1963 and the development of the logic design of Argus 400 was my first ever design.

Development of an integrated circuit based computer in 1963 was an exciting adventure. It was evident that single sided circuit boards would not be adequate and that multilayer boards and some means of joining layers together would be needed. The technology for bonding multilayer boards was bought in and the techniques for plating through were developed initially in-house.

Meanwhile the drawing office was developing techniques for laying out multilayer boards and gaining experience. They used tape on Mylar film at 4 times full size for each of the layers. All camera work involved in producing the masters for etching were naturally (being Ferranti) also done in-house.

Various topologies were tried with differing numbers of ICs on each circuit board. Eventually a size of 39 TO5 cans mounted in three rows of 13 was adopted. The Argus 400 required 20 such boards. The boards were interconnected by a backplane. As one of the initial markets was the RAF a rugged design was required and it was decided to use wire wrap U links to connect the individual cards to the backplane. There were two rows of 35 pins making 70 U links in all. The principle was the tried and tested technique used for connecting the Argus 100 card cages to the frame wiring but the wire wraps were considerably smaller.

Core store technology was used again. The core store cycle time was 2 microseconds. This was built in a similar manner to the processor backplane and provision was made, as in the Argus 100, for up to three 4k word store modules to be fitted.

There was a final backplane that was originally planned to contain the specialised I/O equipment required for aircraft use. When it was realised that the hoped-for airborne market was not ready for such a machine this became the operating console interface plus device drivers for interfacing to Process Control I/O equipment.

The 3 to 5 backplane modules (depending on how many core store modules were fitted) were interconnected via a flexible printed circuit so that the whole assembly could be opened out for maintenance. The whole assembly was designed to fit into a standard aircraft instrumentation case.

2.5 Argus 500

The Argus 500 used a parallel computing architecture but used the same integrated circuits as the Argus 400. It used 2 microsecond core store but the architecture was extended to enable memory bank switching to address four times as much memory.

The architecture also included a similar technique for mapping the accumulator addresses providing eight sets of accumulators. This greatly reduced the time taken to process external interrupts.

The Argus 500 was packaged into a larger vertical plug-in module rather than the aircraft case style used by the original Argus 400. The memory banks were also made into a similar form factor and plugged into processor frame. The Argus 400 was then repackaged in this style making the two machines plug compatible.

The I/O interface was identical between the Argus 400 and 500 so a processor upgrade was a simple swap of the processor module.

3 Integrated Circuit Details

3.1 Micronor 1

This was the first integrated circuit family produced by Ferranti. It predates 1963. The conceptual circuit diagram (reproduced from memory) is illustrated in Figure 1 below.




Figure 1 – Micronor 1

It looks unconventional in that the output stage contains the diodes forming the logic gate. These are connected to the input transistor base. This has a major benefit in that it is impossible to connect the output stage to more places than it is capable of driving. This is however the only benefit and the noise immunity of the circuit was poor due to the transistor base being connected to an external pin and to potentially several tracks on the PCB. It would be therefore unsuitable for building a large design such as a computer.

Encapsulation was in a multi-pin (probably 8 pin) TO5 metal can. It was probably designed to run off a 4.5 volt rail but I cannot confirm that. As far as I am aware it was never produced in commercial quantities.

3.2 Micronor 2

Recognising the limitations of Micronor 1, and looking at alternative experimental devices that were appearing in 1963 such as a range of integrated circuits from Fairchild (probably Resistor Transistor Logic) Ferranti decided to embark on the development of a new logic family. This was to be fast, for its day, and have high noise immunity to ease circuit layout. Micronor 2 was born. Conceptually the circuit configuration is shown in Figure 2 below.




Figure 2 – Micronor 2

The design, in its original form (actually close to that shown above) was designed to drive four loads. An alternative device, a power gate could drive 25 loads. It should be noted that there is some degree of optimisation here. The design was to be used for a 24 bit computer and hence one power gate could drive all the elements of a 24 bit register eliminating any additional problems of skew that could arise from clocks arriving from different sources.

Noise immunity was provided by the two diodes in series feeding the base. The combination provided (from memory) about 1.5 volts of noise immunity on the inputs. Far better than the 74 series TTL from Texas Instruments that came later. The family was designed to run off a nominal voltage of 4.5 volts. The commercial specification devices would run over a range of 4.25 to 4.75 volts. The military versions would run from 4 volts to 5 volts. The temperature range for commercial was 0 to 50 degrees C, for the military versions -55 to 70 degrees C. The typical delay through a gate was better than 7 nanoseconds but one had to be careful about layout to minimise capacitance with the resistive pull-up. The value was 750 ohms from memory.

A JK flip-flop was produced using a number of such circuits cross-connected together internally. This device caused early production problems as it required too high a gain (beta of 25) from the internal transistors. With the silicon purity that was then achievable we were lucky to get one working device from a diffusion of 200 devices on a single wafer. At the time a single JK flip-flop cost 1.5 times that of a graduate engineer's weekly wage.

A eureka moment happened when the designer realised that the diodes closest to the transistor base could be converted into a transistor this forming a Darlington Pair. With a beta of 5 for each transistor, which was readily achievable on an integrated circuit, the required composite beta of 25 was now feasible and yields increased. At the same time the first design of the Argus 400 had been completed and it was realised that a greater fan-out than 4 (3 for JK flip-flops) would significantly reduce the chip count. The increased gain of the Darlington Pair permitted this to be raised to 8 (6 for JK flip-flops) and this is the final specification for the basic IC.

During the Argus 400 design there was also a realisation that one particular grouping of interconnections was occurring repeatedly. See Figure 3 below.




Figure 3 – Small Scale integration

The initial function identified enabled the designer to select one of two sources to feed a JK flip-flop. Both polarities of output were provided to ease this process. A section from the main register of the Argus 400 is illustrated in Figure 4 below. The design opened the register to the core store data when memory was being read. Otherwise the register was connected as a serial register.




Figure 4 – Section of main register

Much later another use for this SSI device was found. By sequencing reset and set signals this useful device could be turned into a simple latch. In many cases this was more convenient than a JK flip-flop and it was certainly cheaper. This is illustrated in Figure 5 below.




Figure 5 – SSI simple latch

As one of the target markets for the Argus 400 and hence the Micronor 2 family was the military and alternative manufacture was required to second source the devices. Marconi was chosen and did manufacture some devices. Texas Instruments, I was told then approached Ferranti for a license to manufacture and market the devices. Ferranti said no. This led to TI developing 74 series TTL and another marketing opportunity lost. 74 series TTL had significantly worse noise immunity than Ferranti DTL and did nasty things to the power supply when it switched state as a result of its totem pole output stage. It was also slower at nominal 10 nanoseconds compared with Micronor's 7 nanoseconds.

Micronor 2 was almost indestructible. It would survive its output being short-circuited to the Vcc rail permanently without apparent damage. I have received a number of round burn marks on fingers particularly from power gates that were

designed to carry 25 loads. In one case it took about 2 weeks of debugging on a new design before an artwork fault was discovered where such an output was directly connected to the 5-volt rail. The aluminium can had by then lost its bright shiny appearance but when the offending track was cut the gate carried on working.

When it became obvious that TI 74 series TTL was taking over the world Ferranti started to manufacture the range under license. It then became more cost effective for us to use the 74 series family but there was a problem – 74 series used a nominal voltage of 5 volts compared with 4.5 volts for Micronor 2. We agonised for a time about this dilemma until the chip designer decided to see just how well Micronor 2 would work at 5 volts. He went to the production line and reset the test criteria upper limit to 5.25 volts from 4.75 volts and ran a large production batch through the test equipment again. There were no failures. The problem was solved by retrospectively changing the voltage specification to be compatible with 74 series. Over a period of time all machines were adjusted to the new range and, as far as I am aware no failures occurred.

The first versions of Micronor 2 were packaged in TO5 8 pin metal cans. Later 14 pin DIL versions were also made.

Development of this IC and pilot production was undertaken at Ferranti in Wythenshawe on a prototype production line. Large scale production eventually moved to Ferranti Gem Mill with Marconi being licensed as a second source. Ferranti suffered from being unable to produce pure enough silicon wafers and were left behind when TI was able to improve the purity and the size of the wafers. Ferranti used the smallest size chips for this reason – it was their only way of obtaining adequate yields.

The chip designer was Peter Bagnall who eventually left Ferranti to work for Motorola.

Document authorship info.:

John Steele, submitted 2000, revised John Steele & D. Bew 4th March 2016