

A Short History of the Data Network in the Department of Computer Science and the School of Informatics at the University of Edinburgh

George D M Ross
gdmr@inf.ed.ac.uk, gdmr@dcs.ed.ac.uk

November 30, 2021

This document relates the evolution of the data networks in the Department of Computer Science and the School of Informatics at the University of Edinburgh. It attempts to say why things ended up the way they did, as well as looking at what emerged. Note that diagrams may be simplified slightly for clarity, though significant details are included. Related documents can be found on the Edinburgh Computer History site: <http://history.dcs.ed.ac.uk/>.

The Author was a member of the Department/School from 1974–1976 and from 1979–2022.

At the time of writing, the “new EdLAN” project (see page 18) deployment was only part-way through, and its full potential effects on our network had not really been felt, having been delayed and had its schedule altered for a number of reasons, not least Covid-19.

1 Departmental Link Boards

A significant step towards the development of a high-speed data network in the Department of Computer Science was the design in the early 1970s of the “Link” board. This provided independent RX and TX serial connections, often but not always used as a pair. The Links ran at up to 1Mbps over standard 75Ω coax cable. Each connection was 8-bit clean and flow-controlled, with a positive ACK signalled to the sender for each octet transmitted. Implementations existed for the Interdata backplane, the DEC unibus (PDP-11 and VAX-11) and the DEC PDP-9 and -15 machines.

By 1975 the Department had two Interdata clusters interconnected by Links, as well as a PDP-8, PDP-9 and a PDP-15/20. Machines were gradually added, so that by about 1980 the following were interconnected by Links (see figure 1):

- the “Filestore,” an Interdata processor with two 70MB disc drives and a number of Link boards, the purpose of which was to act as a central file storage repository for all the Department’s other small machines
- the Interdata clusters
- the PDP-9 and PDP-15/20, now using the Filestore as virtual DECtape drives, as this turned out to be faster than using real physical hardware
- a PDP-15/40
- an ICL 7502 discless editing station, which ran VECCE and very little more
- the “LEGOS” systems: standalone Interdata machines with rather more memory than the small cluster machines




Figure 1: Network diagram from around 1980

- the Edinburgh Remote Terminal Emulator (“ERTE”), a research project into performance testing
- a DEC VAX-11/780 “ECSVAX” running VMS, which was the Department’s main multi-user machine at the time
- the “Data Curator”: a Perkin-Elmer 3220 used as part of another research project
- a lineprinter.

The Filestore was the hub of the network, and almost all traffic over the Links was filestore protocol (described elsewhere¹), either directly manipulating the files in the case of discless machines, or to transfer them back and forth in the case of ECSVAX and the Data Curator. There were only two real exceptions to this.

One of these was the lineprinter. This was provided with a Link receiver (no need for a transmitter) and connected to its host by this means. Initially this was to one of the Interdata clusters, and latterly to the Filestore. The Link’s inherent flow control meant that the despooler could simply throw characters down the line as space in the printer’s buffers permitted.

The other exception was that there was a direct connection between the Data Curator machine and ECSVAX. This was originally intended to allow direct data transfer between the two machines. In fact, it was not used for this purpose, with file transfer being staged via the Filestore instead, but it did run a home-brew bidirectional remote login protocol for a while.

2 2Mbps Departmental “Ethernet”

There were a couple of drivers for the development around 1980 of the Computer Science Department’s 2Mbps ethernet-like network. The Metcalfe & Boggs paper² had come out, showing what

¹For example, <http://history.dcs.ed.ac.uk/archive/os/APM-filestore/FS.1976/FSDOC.LAY.html>

²Linked from <https://en.wikipedia.org/wiki/Ethernet>

could be done; and the sheer number of connections back to the Filestore were becoming unmanageable — indeed, latterly there was a multiplexor machine attached to the Filestore, mainly to provide additional connectivity.

The Departmental ether was different in a few ways from standard ethernet: it ran at 2Mbps rather than 10Mbps; it had a slightly different addressing structure; and by default it presented as a *reliable* datagram service between configurable end-station ports, in contrast to standard ethernet’s unreliable datagrams between stations, which allowed filestore protocol to be layered directly on top. Ether stations were connected to their hosts (Filestore, LEGOS, ECSVAX) over serial Departmental Links, now 9-bit-clean and running at 2Mbps, and were designed to offload the ether protocol processing from the hosts. Like standard ethernet, the Departmental ether ran over 50Ω yellow “thick ether” cable, and used the same “vampire” clamps, though interfaced through purpose-built transceivers.

The arrival of 32-bit microprocessors, in particular the Motorola 68000 series, prompted several members of the Department to construct an experimental testbed, known as the “Fred Machine.” This was a modular system, with processor cards, memory cards and graphics cards all plugging into a common backplane. In particular, there was an ether card, with a revised version of the Departmental ether station. These machines proved so successful that they were relabelled the “Advanced Personal Machine” (“APM”) and gradually rolled out as the Department’s main workhorse for staff and students.³

Unfortunately, the original version of the ether station as used by the Filestore turned out to be under-buffered, and unable to cope with the increasing load. Packets would be dropped quite frequently, resulting in broken connections, rebooted machines and lost work. It was decided to reimplement the Filestore using a 160MB Winchester disc attached to a Fred machine. This took place from 1983 onwards, in two phases:

1. as the Interdata Filestore’s ether connection was too unreliable for the intensive traffic involved alongside its normal load, the filestore protocol interpreter was initially implemented on ECSVAX, calling out to VMS’s native filesystem; then
2. having uploaded an image of the necessary file structure to the ECSVAX “filestore,” the APM-based Filestore was developed using the relatively more reliable facilities it provided.

Once the APM-based Filestore was brought into service, the original Interdata Filestore was retired, and run as a printer-despooler driving the original lineprinter over its Link interface. In addition to the 160MB Filestore, two 300MB Filestores were deployed, one mainly for staff use and the other mainly for student use, along with four based on 70MB drives for development and small projects.

3 10Mbps Ethernet

Economies of scale kicked in and researchers started to buy machines, mostly from Sun Microsystems, rather than designing and building them in-house. These machines came with standard ethernet, as had lfc’s VAX-11/750 and Pyramid machines (“cstvax” and “cheops”). A LANCE-based ethernet card was built for the APMs and installed in a couple of the Filestores. Initially 10base5 (“thick ether”) was used, but as clusters of Suns appeared in an *ad hoc* manner they were linked with segments of 10base2 (“thinnet”) and routed by adding second interfaces to one or more machines in each cluster.

The Department’s machine halls at the time also housed a Gould PowerNode (“itspna”), shared with the Artificial Intelligence (“AI”) and Electrical Engineering (“EE”) Departments, and a couple of GEC machines for ERCC’s commercial arm.⁴ AI accessed the Gould by serial lines connected to their South Bridge site using a statmux pair and a BT kilostream link, while EE had their own ethernet bridged to ours over a fibre pair run through one of the KB site’s

³<http://history.dcs.ed.ac.uk/archive/scans/gordon/apm/>

⁴One of these eventually became the original “tardis” machine. See <http://www.tardis.ed.ac.uk>




Figure 2: Network diagram from around 1986

heating ducts (figure 2). This link in fact remained in service until building work in 2000 meant the removal of the fibre.

Although there was this common infrastructure, there were actually at first a few separate islands of connectivity: all the BSD-based machines spoke TCP/IP; the GECs spoke Coloured Book protocols, for political reasons; and the Filestores and ECSVAX initially spoke neither. A TCP/IP stack was therefore written for these latter machines, including RIP-based routing, rlogin and rcp clients and servers, and bidirectional lpr-based printing.

Though perhaps not strictly part of the data network, it's worth mentioning that at this time the Department also ran serial lines to offices and a terminal room in 2501 from ECSVAX, `cstvax`, `cheops` and `itspna` using a Gandalf multiplexor.

Gradually as the APMs reached the end of their service life they too were replaced by Sun workstations, and the thinnet structure was necessarily rationalised. Discless Suns were heavy network loads, and so the network was explicitly subnetted into several staff segments and several student segments, with the Sun file servers doubling up as routers. The "CS1 lab" in Appleton Tower was initially separate, before it too was linked to the rest of the Department at JCMB, first using UUCP and then SLIP over a 4800bps EPS8 BT line (fortunately both sites were within the same 667 exchange area). All subnets were /24, and routing information was distributed using Routing Information Protocol ("RIP").

4 Wider: RConet and EdNet

The wider University network is described elsewhere⁵, so only those aspects which were CS-related are mentioned here.

For the first period of time covered by this note (mid-1970s to roughly 1980) the Department used serial terminals attached to the University's multi-access EMAS systems for its teaching lab in Appleton Tower, until that was replaced by Sun workstations. At JCMB there were also EMAS terminals in staff offices and Block O/P, but the Departmental and ERCC networks were initially almost completely separate. The only connection, such as it was, was the ERTE system, which could be used for small-scale file transfer between EMAS and the Filestore when it wasn't required for its research project work.

When the ERTE project finished the hardware was mostly disposed of. Its HDLC interface (a DQS11) was, however, redeployed in ECSVAX, with a locally-written driver and protocol modules (written by Gordon Brebner and Paul McLellan⁶). This provided a full set of remote login, file transfer, printing and email applications.

⁵<https://www.wiki.ed.ac.uk/display/iticis/University+of+Edinburgh+Networking+Overview>

⁶Initially the driver kept crashing the machine for no obvious reason. Eventually after rummaging through the fiche Paul realised that the IPL being used by the DQS11 didn't actually have any handler assigned to it by VMS.

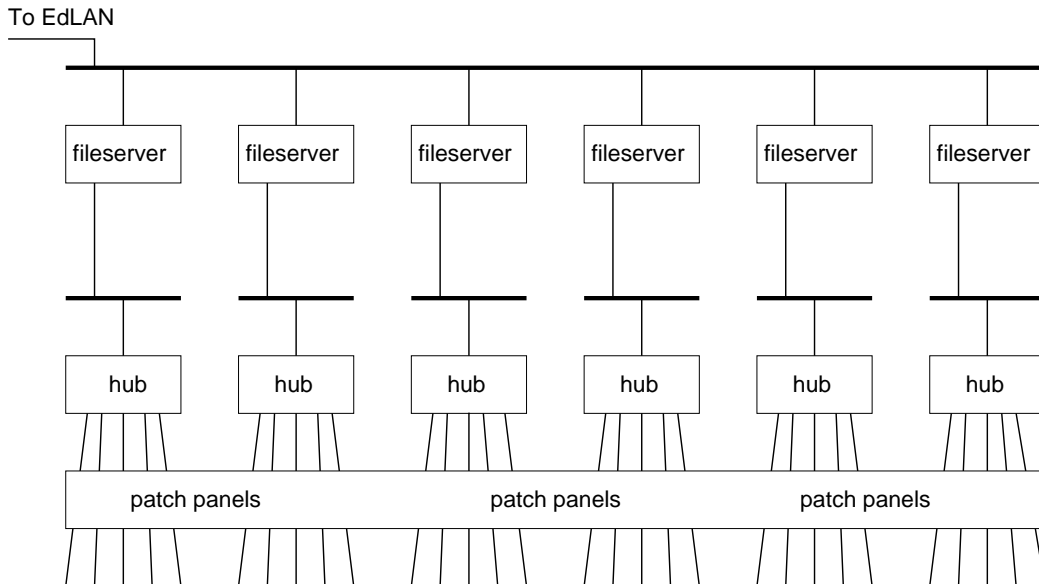


Figure 3: Outline diagram after the First Great Rewiring

`cstvax` also had a connection to EdNet, using an LSI-11 as a front-end (a “York Box”).

5 TP and the First Great Rewiring

As the number of Suns in the Department increased, it was necessary to control the allocation of machines to subnets ever more carefully. Discless Suns paged heavily, and the applications used by some groups of users were quite network-intensive. As a result, we had numerous 10base2 segments strung around the various parts of the Department’s accommodation to spread the load.

Physically, 10base2 isn’t very robust. It’s disruptive to add new connections, as the cable has to be cut and connectors carefully fitted. Connections are then subject to strains as machines are moved and feet tangle in the cables. Corrosion sets in (not helped by drips from wet bicycles parked in offices!). Hum-loops can be created as connectors come in contact with building metalwork. Faults might affect everyone on the wire, or just the odd machine. Debugging is tedious and time-consuming, and can often “fix” problems for no obvious reason. Finally, we had too many subnets in use to be able to run them physically through all the offices, simply because of the size of the cable bundle, which put a severe constraint on the allocation of people to offices!

By 1991 we had concluded that we couldn’t carry on with 10base2. We decided to have the First Great Rewiring, introducing 10baseT on cat-3 cables. Physically this is much more robust than 10base2, and gives better, but not perfect, fault isolation. Crucially it made it possible to patch any office to any subnet (figure 3), considerably easing the room-allocation problems of the admin staff. It was also now even easier to split subnets to add bandwidth, and eventually we had around 25 subnets in use for various purposes. We were back to point-to-point connections, which had proved problematic with the Links and the Filestore, but we weren’t now limited by backplane space as it was easy to add new hubs as required. We had active electronics instead of a passive piece of wire, but the overall robustness more than made up for the added complexity. Wiring was installed by the Department’s technicians, with one hub site in an alcove in room JCMB-2501, covering the offices, and one in the South Machine Hall covering the servers. A few hubs were also spread around under the Machine Halls’ floor for the students’ client machines. At its peak this network had over 30 hubs in 2501 and another 15 or so in the Machine Halls.

This was a considerable improvement over what we had had before, but there were still issues. Although we had individual connections from machines and offices to the hubs, those hubs were

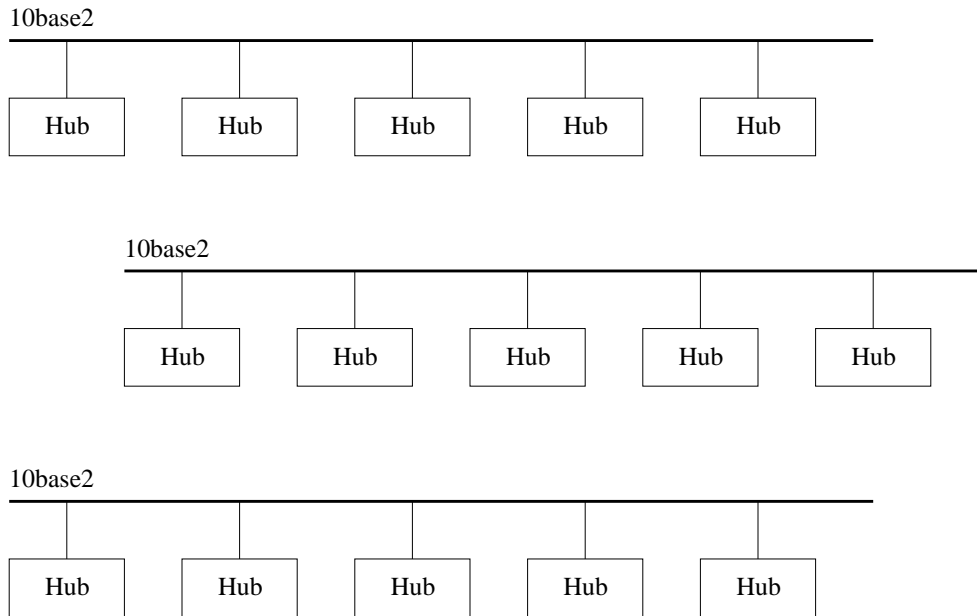


Figure 4: Hubs connected by 10base2

interlinked using 10base2 (figure 4). We still had a shared medium — hubs are simply repeaters, and copy everything they receive out onto all their ports — and bandwidth was still under pressure. The hubs were connected together using the 10base2 cabling, which was rather old and fragile. And although faults were generally much more contained, hub faults could affect everyone on an entire subnet, often in strange ways, and were hard to diagnose.

And we now had the issue of how to manage the large number of connections at the 2501 hub site, particularly as people moved offices or had new machines to connect. Churn resulted in the 10baseT patch areas getting completely out of hand over time, so that on a couple of occasions we just ripped the whole lot out and repatched again from scratch.

6 EdLAN

When EdNet was replaced by EdLAN the Department connected both its JCMB and Appleton Tower sites, and retired the SLIP line as soon as we could. Up until then we had been using IP addresses chosen somewhat arbitrarily, we but renumbered both sites onto the officially allocated addresses as they were connected to EdLAN.

7 Switches and the Second Great Rewiring

The arrival of switches (marketing-speak for multi-port bridges) opened up a couple of avenues for improving the network. These were intelligent devices (well, slightly...) with much better fault diagnosis and reporting than was available from the hubs. They also implemented VLANs, allowing for several subnets to be carried over the same physical wiring, and could group several links together into one “fatter pipe” to increase aggregate bandwidth.

We introduced a stacked pair of 3Com switches in the 2501 hub area and another pair in the machine halls server area. Hubs were now connected as individual leaf nodes off the switches, immediately reducing contention for bandwidth, and inter-switch traffic was carried over dual 100baseTX links (figure 5). Faults now had a much more localised effect, and the switches’ diagnostic tools allowed for faster resolution and better forward planning. Once again we had




Figure 5: Hubs interconnected using switches

complex electronics replacing simple wires, but the overall result was generally much more reliable. However, the patch areas were still completely out of hand!

Although the 3Com switches improved things, there were still a few rough edges. The firmware didn't initially support VLANs; and the switches would sometimes crash and lose their configurations, though unstacking them did help with this. The biggest issue, though, was that we still had servers connected to shared segments at 10Mbps, albeit now at least single-hub segments.

The introduction by HP of their Procurve 4000M modular switches, at an affordable price and with a lifetime warranty, allowed us to rethink our network connectivity. Each machine could be given its own private 100Mbps link to one of the switches, and the switches themselves could be connected together using 1Gbps fibre and copper links. Ports could be assigned to VLANs based on configuration files, so there would be no more need ever again for physical repatching and the mess that resulted. Redundant links could be added, under Spanning Tree Protocol control.

Unfortunately, in order to make use of these opportunities we had to undertake a Second Great Rewiring (in 1998). Our initial TP cable installation had too many over-length lines, which passed the 10baseT electrical requirements but not the 100baseTX timing requirements. We had also doubled up many connections using the "spare" cable pairs, which 10baseT could tolerate but which caused unacceptable crosstalk for 100baseTX. Finally, most of our initial cable installation was cat-3. Once again the Department's technicians installed the new cat-5 cable, now terminated in two sites (rooms 1415 and 1619), and a pair of 4000M switches was installed at each. Further switches covered the teaching machines in the Machine Halls, and three more switches provided "core" connectivity as well as links to the servers. Most users had two outlets installed for them, with mini-hubs being used where more machines were needed, the realisation that the panels above the JCMB office doors were asbestos-based having meant that only existing holes could be reused for wiring! Though expanded later by the addition of the IPAB robot labs and some student workrooms, this design remained essentially unchanged until we moved from JCMB to the Informatics Forum in 2008 (figure 6).

Initially we didn't entirely trust this new hardware, so we ran one of the server-area switches as a hot-spare, with a shadow configuration. Later we concluded that this wasn't really necessary, and some of its ports were put to "real" use. We did, however, still try to make sure that we had enough spare ports overall to carry all the most important servers if necessary.⁷

Routing was still initially performed by the file servers, and then by dedicated Suns with multi-port 100baseTX cards. For some reason these believed by default that all interfaces on a machine should use a common MAC address, which did not interact well with switches which used a common forwarding table for all their VLANs and believed that a MAC address should only be in use on one port. This was solved by running a script early on in the boot sequence to reconfigure

⁷We use bonding to achieve resilience these days.




Figure 6: JCMB network using switches throughout

the Suns to use separate MAC addresses for their interfaces.⁸ (The Linux routers we run now do not have this quirk. Instead they believe by default that any IPv4 address configured on a machine can be used on any of its interfaces, which strangeness we also have to reconfigure away early in the boot sequence.)

Eventually we added a ProCurve 5308xl at the core, able to route traffic at wirespeed, in order to handle increasing traffic levels. As a beneficial side effect of this, the fact that we didn't have to bother about which "side" of a fileserver we were speaking to from any given machine meant we could considerably simplify our automount maps. This marks the point at which we could really start to decouple host configuration from the ever increasing complexity in our network topology, and is probably one of the more significant moments in the development of the Department's network.

8 Configuration and Monitoring

While we still had only a small handful of switches connected as a backbone for the hubs, it was reasonable to configure them using their CLI or web interfaces. Once we had decided to roll out "switch ports for all," however, it was immediately apparent that some kind of configuration tool would be needed to allow us to devolve sufficient network management features in a simple way. The requirements were that it should support both our 3Com and Procurve switches with a common user interface; and that interface should expose only those items which were necessary, and in such a way that a simple subset could be made available for the technicians and support staff to allow them to make routine changes. The tool should run on a commodity machine, with the configuration stored in its usual filesystem and backed up in the usual way.

At the time there wasn't a suitable tool commercially available, so we developed our own,⁹ based around Tcl, `expect`, and the Tnm SNMP package. A Procurve or 3Com switch would be

⁸<https://history.dcs.ed.ac.uk/archive/os/lcfg/ether-object.txt>

⁹<https://svn.lcfg.org/viewvc/source/trunk/ProCurve/> is the HP/ProCurve/Aruba version. This was originally checked into CVS in September 2002. The 3Com version is no longer maintained, and not on-line.

queried using SNMP for its current state, which would then be compared against the switch's configuration file and any updates applied as required. For the Procurve switches those updates could also be done using SNMP. The 3Com switches, however, had such a baroque VLAN implementation that it was simply not safe to try to use SNMP there too, and so updates for those were applied by sending commands to the CLI using an `expect` script. The tool was deployed using the “`rfe`” remote file editing tool and `make`. CogSci also used the tool, with a front-end to generate the configuration files. (A 2009 development project took another look and still couldn't find anything more suitable for our use than this in-house tool. It is, however, expected that things will change for the edge at least once the new EdLAN is finally rolled out.)

We also needed better routine monitoring, for error and basic statistical trend reporting. Again, this was written in house, using Tcl and Tnm. All the switches are polled every 5 minutes, and their traffic counters extracted for every port, whether configured or not. These are stored using `rrdtool` for future use. Index files are generated for configured ports several times per day, as required, and the graphs are available as necessary using cgi-bin script wrappers around `rrdtool`'s facilities. `rrdtool` is also used to produce temperature and power-utilisation graphs, as well as composite network graphs based on several ports' data.

9 Traffic filtering

While we had an isolated data network, the issue of allowing or blocking traffic from outside didn't, of course, arise, though not all our internal users were as benign as might have been desired. When JANET's “shoestring” project came along and we connected to the Internet, we initially relied on the underlying strength of the systems we were running to keep outsiders out. As the Internet grew and connectivity improved, this position became clearly unsustainable.

We were using Suns as routers at the time, so the natural tool to use was Darren Reed's `ipfilter`.¹⁰ We reorganised our subnets slightly so as to have one “outside” subnet, linked onwards through EdLAN, with all the rest being “inside.” We did consider setting up a DMZ in the usual way, but concluded that it didn't really fit into the way our network was used. We configured `ipfilter` initially by hand, and later under `lcfg` control, so as to block inbound traffic by default but permit outbound traffic by default. The more restrictive position of blocking everything by default would have resulted in an excessive management load then, as it would now. Later as we moved to using Linux perimeter routers we ported our rules to `ipchains` and then `iptables`, effectively unchanged in the first instance, as this was the native filtering system available.¹¹ We did consider using the ProCurves' ACLs, but these were, and remain, less flexible.¹²

Blocking inbound and permitting outbound by default has remained the configuration to date. It's not ideal, but it is still the best fit for us from the well-known paradigm of “choose any two from ‘secure,’ ‘cheap’ and ‘easy to use.’”

10 Formation of Informatics

With the formation of the Division (later School) of Informatics in 1998 came the requirement to integrate the constituent Departments' networks.

- The situation in CS was as described above: we had our main site at JCMB with the “CS1 lab” in Appleton Tower as a satellite operation. All machines were connected to a 100Mbps switch port, either directly or through a mini-hub. Hard patching had been replaced by soft configuration of the switch ports under control of the configuration tools.
- Artificial Intelligence (“AI”) was split across two sites: Forrest Hill and South Bridge. The Forrest Hill site had been rewired with cat-5. South Bridge was still mostly 10base2.

¹⁰<https://en.wikipedia.org/wiki/IPFilter>

¹¹This was very much a case of “build something three times” to get it right.

¹²We're taking a closer look at them at the time of writing (late 2021), with a view to making some use of them to segment our network a little more, but they aren't an ideal fit for the problem.

- Cognitive Science (“CogSci”) was sited at the East end of Buccleuch Place. The buildings had been rewired with cat-5 and a switch-based installation was in use, very similar to CS’s.

The requirements were as follows: Informatics machines should be treated as part of one unified network which happened to be spread across the five (initially) sites; and the network should be robust, there should be transparent fast paths between sites, and all external traffic should be filtered. We decided that it would be best to keep subnets and their associated VLANs within sites as much as possible, to minimise switch table use, and inter-site and EdLAN dependencies, a position we have maintained to date; we introduced a “transit” subnet for inter-site traffic, with only routers and a small number of network infrastructure machines on it; and we gradually introduced switch/routers throughout, to improve intra- and inter-site routing throughput. Inter-site traffic was carried by tagged VLANs across EdLAN, for the most part. We did have some dark-fibre links between AT and FH and between AT and BP for a while, but instabilities introduced by being part of the EdLAN spanning tree meant that we had to keep them mostly disabled.

Routing was still controlled using RIP, initially using `gated` and latterly `quagga`. Traffic was directed through preferred primary and secondary routers by adjusting the metrics advertised. We realised quite soon that the use of RIP at our privately-interconnected sites did not interact well with the OSPF that EdLAN was running — in particular, metrics were lost in the conversion, and so taking routes from EdLAN at one site, propagating them internally to another site, and then forwarding them from that site to a different EdLAN router was not safe — and so we had to introduce heavy filtering of the RIP routes we sent and received. The realisation that a large amount of Physics’s traffic was transiting across our networks as a result of RIP/OSPF interaction helped to increase the priority of this work.

11 South Bridge fire and relocation to Appleton Tower

The South Bridge fire in 2002 did not actually have much of an immediate impact on the network, other than getting rid of all remaining 10base2. Our Appleton Tower core was upgraded by the addition of a switch/router, and temporary edge switches in level-3 and level-5 were connected to it. The subsequent relocation of the core from level-5 to the basement server room was done essentially transparently to the users by relying on the redundant inter-switch links to keep the traffic flowing while relocating the main links to new fibre. As floors were refurbished, a pair of edge switches was installed in each and linked back to the core. (When floors were later upgraded we moved to the “Forum” model of a hierarchy of switches, as a more cost-effective solution.)

The Appleton Tower network was designed in the light of experience at our other sites (figure 7 on page 11). Core and edge switches were clearly distinguished by rôle, and models were chosen for their particular rôle, not as a compromise for mixed use. The initial design had one switch/router at the core, alongside one non-routing switch, with Linux routers to provide redundancy. The edge switches, two per floor, had a “main” 1000baseSX link to `atba` and a backup 100baseFX link to `atbb`. A PoE switch was added, with one link to each of the floor’s edge switches. Those original core switches have since been replaced by a pair of more capable switch/routers, with one 20Gbps trunk between them and another to the Forum core, and with one edge switch from each floor connected to each AT core switch.

12 Informatics Forum

The design of the network in the Informatics Forum once again drew on experience at our other sites. As the building was being constructed from scratch we were able to get a network design which was reasonably close to the way we wanted it, rather than having to be retro-fitted to an existing building. The funding included provision of new switches, as the ones we had installed in our previous sites would have been due for replacement anyway.

At the core of the building network was a cluster of three ProCurve 5412zl switch/routers (since upgraded), each connected to the other two by 20Gbps trunked links. Each of these switches also

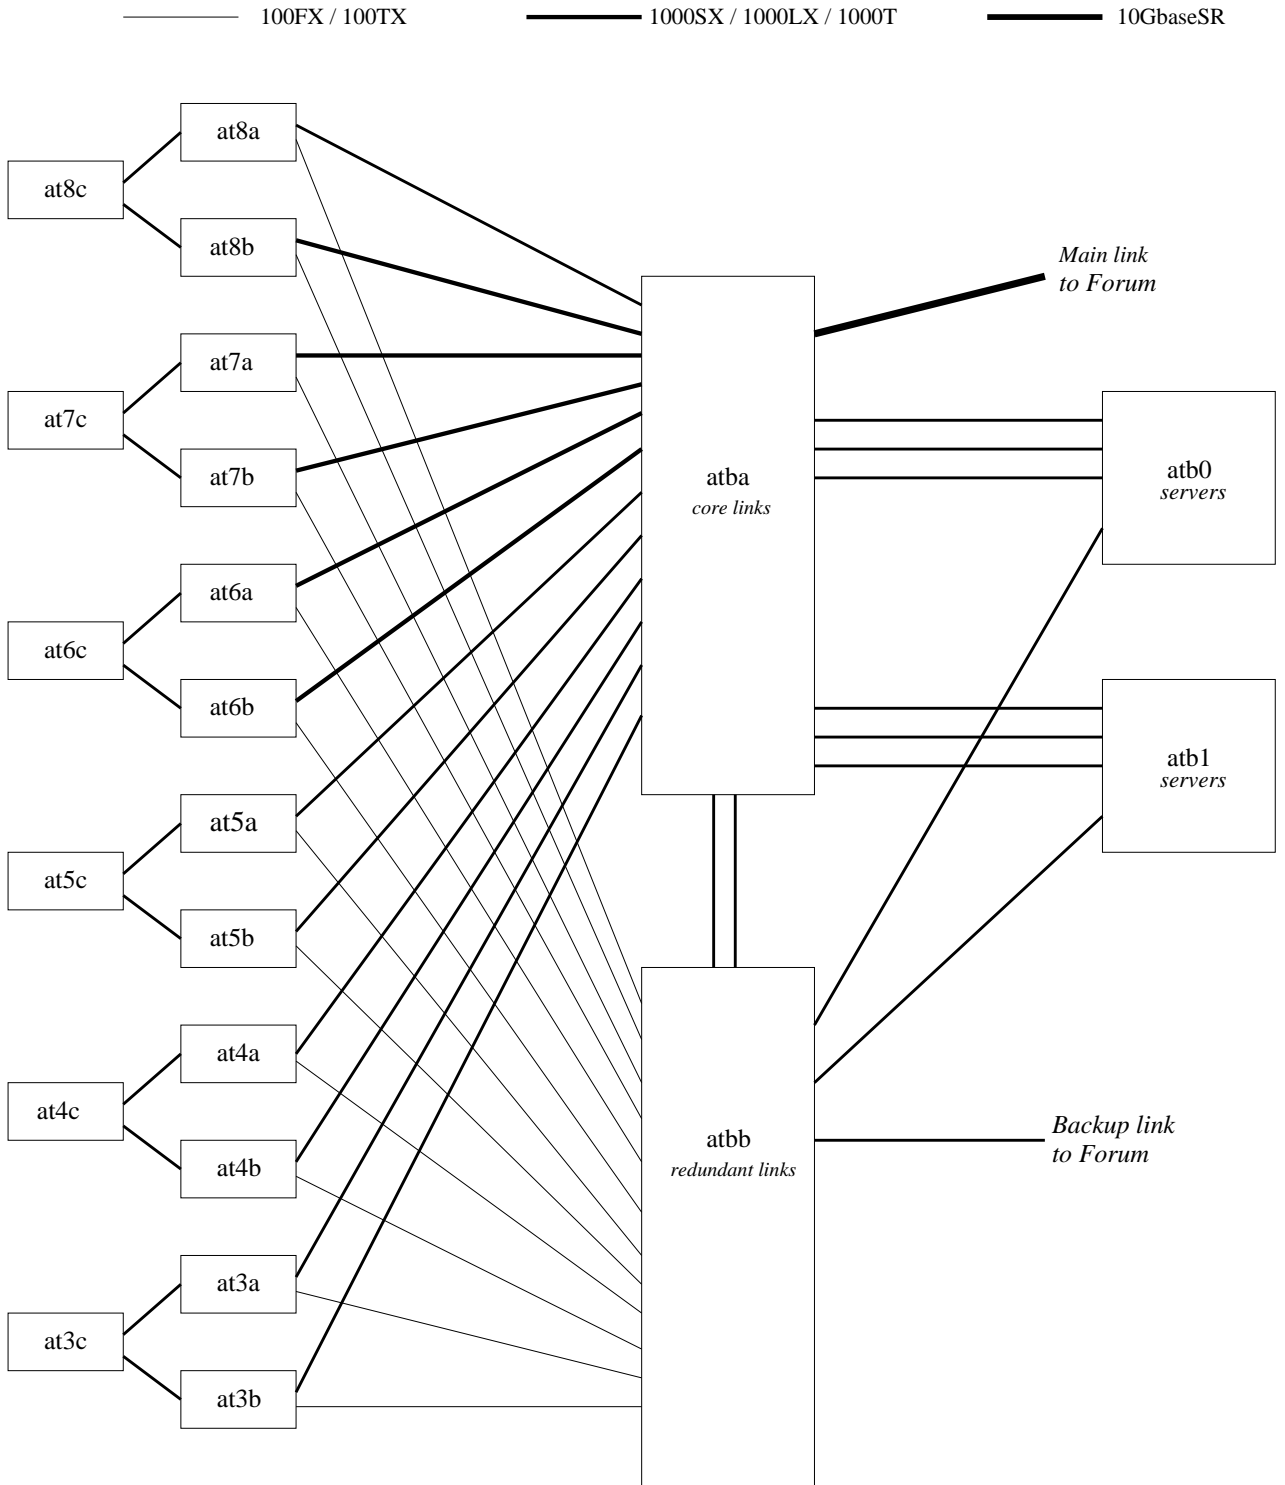


Figure 7: Appleton Tower network

has a link to one of the Appleton Tower core switches, the whole forming one tree under Rapid Spanning Tree Protocol control¹³ with priorities and link costs chosen so that the tree is stable and traffic is contained within the two sites and passed preferentially over the desired links. These switches have four power supplies, connected to three different power circuits and two bottom-of-rack UPSes, for robustness, but the intention of the design is such that one or even two of the core switches can fail without the service for the rest of the building being much affected. (This also proved very useful during the upgrade process, as we could take a switch down for replacement without the users noticing much effect, other than briefly while the spanning-tree rebalanced.)

The physical cable installation was designed in detail by the building's engineers, more or less to our specifications. Each floor has from one to three "IT closets," located to fit in with the building's other services. In the upper floors the closets are sited in three columns, vertically one above the other, with a direct cable path between. Unfortunately the ground floor and basement closets are offset slightly from this arrangement. Fixed copper wiring radiates out from each closet to a number of distribution boxes under the raised floors. Floorboxes, with six data and four power sockets, are then connected to these under-floor distribution points. The intention was to allow flexibility for floorboxes to be added or moved later without the disruption of having to wire them all the way back to the closets. This has proved to have been very useful over the years, as more and more occupants have been squeezed into the building and rooms reconfigured to suit. There are also a number of fixed data outlets spread around the walls of the building, for printers, wireless access points and the like.

Of the six data points in each floorbox, four were initially patched in. Each user was provided with one 1000baseT link, two 100baseTX links and one VoIP phone link. As with our other sites, these could be assigned to VLANs under the control of the configuration system. This arrangement was chosen as being a reasonable compromise between cost and speed. Each closet then has fixed-format ProCurve switches installed, sufficient to service the required number of each type of port. At the time of writing there were more than 150 switches in the Forum, providing for almost 4000 data connections and over 500 phone connections, patched in 17 closets and three server rooms. The 100Mbps switches were later upgraded, so that all data outlets are now patched to 1Gbps ports.

Unlike Appleton Tower, where each main edge switch had its own direct link back to the core, not all the Forum switches are directly connected, due to the cost of doing so being prohibitive. Instead (see figure 8) all the Gbps switches in each column of closets are interconnected vertically using 10GbaseT links (initially 10GbaseCX4). Roughly every third switch in each column then has its own direct 10GbaseSR link back to the core, with the remainder being indirectly connected. 100Mbps switches were then connected indirectly via one or other of the Gbps switches. In all cases, connections are arranged so that each switch has at least two independent paths back to the core, all still under RSTP control, so that the failure of one switch does not result in a loss of service for any other. There are also some cross-links between the columns, partly for robustness and partly because it turned out in some cases that this was the best way to provide dual paths back to the core. Each closet also has one PoE switch, to power wireless access points. The phone network effectively parallels the data network, to minimise the chances of either impacting on the service of the other. In all cases "close" switches are linked back to different core switches for added robustness. This arrangement has shown itself to be easy to extend: additional Gbps switches can be added into the existing columns, with additional uplinks to the core as necessary, with new 100Mbps switches connected to the existing Gbps and 100Mbps switches. The general arrangement has been retained, though all of the original Forum switches have since been replaced with newer models.

We have not stacked switches since our 3Com days, for a couple of reasons: firstly, it means that we do not have to have one single model in a closet, which makes upgrades and end-of-life replacements easier and with the cost spread to suit the available budget; and secondly, our configuration management tools mean that there is no particular advantage in stacking. And in

¹³We filter BPDUs where we connect to EdLAN, so that neither side is affected by the stability or otherwise of the other. This does of course limit bridged traffic to use only one of those links, though routed traffic can use whatever is available to it. This is an area we expect to review following the new-EdLAN roll-out.

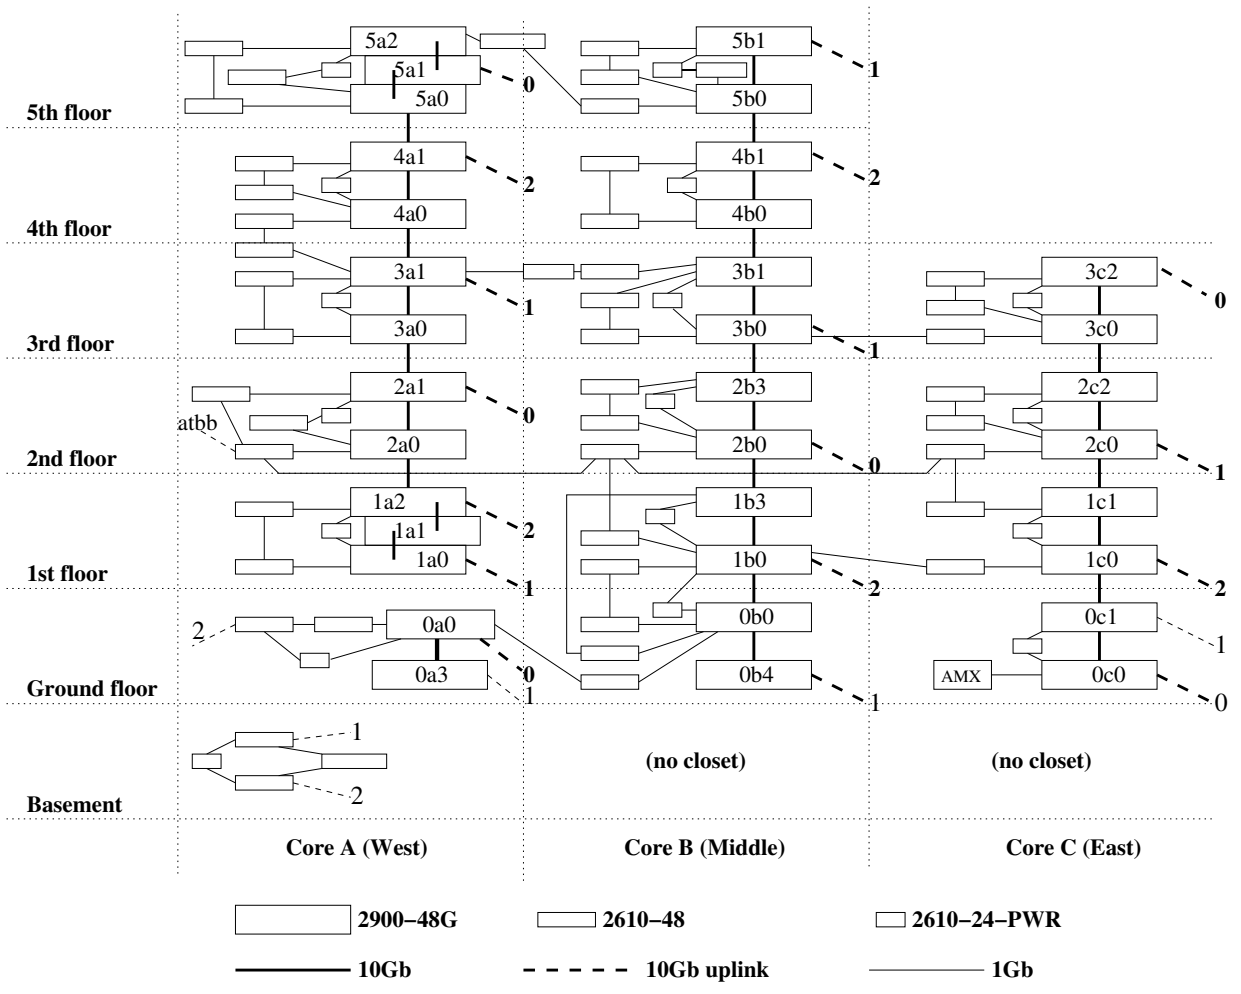


Figure 8: Forum edge data network

a few closets we have too many switches to be able to combine them all in one stack anyway, so the simplicity case is lost there. However, the new-EdLAN edge is based around switch stacks, so there is expected to be some change as a result of its roll-out.

The main server room in the Forum is fairly straightforward. Each group of three racks has a pair of switches installed at the top along with a redundant power supply. The pair is connected together using 10GbaseT, and each switch is connected back to a different core switch, again using 10GbaseT. Finally, each switch has a 1000baseT link to a switch somewhere else in the room, chosen so that it links back to the third core switch. The self-managed server room are similarly arranged.

13 Routing using OSPF

We had been using RIP as our routing protocol ever since we had started running TCP/IP, in spite of its awkward interface with EdLAN. However the move to the Forum prompted us to re-evaluate. A larger building suggested that sticking with /24-sized subnets throughout would make client management rather more complicated than it really needed to be, as well as preventing fully efficient use of the (scarce) address space. We did consider RIPv2, but as EdLAN was using OSPF as its native routing protocol at the time it made much more sense for us to adopt that too.

In conjunction with Information Services (“IS,” formerly UCS, formerly ERCC), it was decided that all of Informatics should be in one OSPF area (area 64). The Area Boundary Routers would be the EdLAN routers at JCMB (`kb6` at the time) and Appleton Tower (`at5` at the time; both have since been upgraded). The two “external” subnets (42 and 160) would each have one EdLAN router and all our Linux perimeter routers on them, to provide for better control of external traffic as well as robustness against router failures. The transit subnet remains, but now has only switch/routers (and a handful of infrastructure machines) but no Linux routers. Each site has one subnet where Linux routers and switch/routers speak OSPF, and routes for almost all other subnets are injected by the switch/routers acting as ASBRs. Traffic is directed preferentially by setting the routers’ path costs — this is necessary to ensure that external traffic is routed symmetrically, so that `iptables`’s connection-tracking works properly. Figure 9 shows this in outline form: indicative switch/routers are labelled “Sxx” and Linux routers “Lxx.” As well as allowing for non-/24 subnets (we now have a variety of /22–/27-sized subnets in use) this arrangement provides improved robustness, by allowing any site to act as a standby perimeter router for any other site. New-EdLAN has prompted a revision of this arrangement; see section 20 below.

The transition from RIP to OSPF was handled in several stages. The first step was to enable OSPF on all the core switch/routers, but with the administrative distance set so that RIP continued to be preferred. Once that was in place, OSPF was introduced on all the Linux routers. Finally, RIP was disabled everywhere, and we had a pure OSPF site.

Moving to OSPF also meant that we had to change the way routes are distributed to non-routers. While we had RIP this was done by those non-routers snooping on the RIP traffic. We didn’t want non-routers participating in the OSPF exchanges, as this would have loaded the DRs and BDRs unnecessarily as well as weakening security. We considered RIPv2, but concluded that as we only really needed to distribute default routes it would have no advantage over router-discovery. We have, however, tuned the router-discovery parameters to provide a faster failover time. Those machines which are not able to listen to router-discovery are expected to honour the values sent to them by DHCP. (Similar arrangements are now in place for IPv6 using RA.)

14 AT Basement Refurbishment

The Appleton Tower basement server rooms were refurbished in 2009, and we had to decant all our kit out, mostly to Forrest Hill, for the duration. The comms rack was able to stay *in situ*, but behind a temporary wall and with no power or ventilation. Although spread across three locations




Figure 9: Outline OSPF routing arrangements

(Appleton Tower upper floors, Forrest Hill server room and the Forum server room), we decided to maintain the “AT network” identity throughout, to make management of clients, servers and self-managed systems easier.

We installed a pair of temporary switch/routers in the Forum server room, initially running in parallel with the AT core. They were connected to the Forum core switches, which had had the AT VLANs added. Once those switches were up and running we were able to re-patch the upstairs switches to them using the fibre we had had installed between the buildings when the Forum was built. The decanted servers in Forrest Hill were catered for by installing a pair of Gbps switches, connected back to the Forum over single-mode dark-fibre. Figure 10 on page 16 shows the arrangement, with the four switches above the dotted line being in the Forum and the two below in Forrest Hill.

Once the refurbishment was complete, the AT core switches were reinstalled and the upstairs floors re-patched. A pair of new switches was installed in the server racks, connected back to the core using three 1Gbps lines trunked together as the distance was too far for CX4 and the expense of fibre wasn’t really justified for the switches involved.

When these server room switches were later upgraded we did also upgrade their uplinks to use 10GbaseSR. We decided on fibre rather than using 10GbaseT for a couple of reasons:

1. the links were run on cable baskets shared with power and other unidentified data lines, and we weren’t confident about 10GbaseT’s noise immunity under this arrangement; and
2. we needed SFP+ cards in the switches anyway, for links to the Forum and to EdLAN, and this allowed us to share resources.

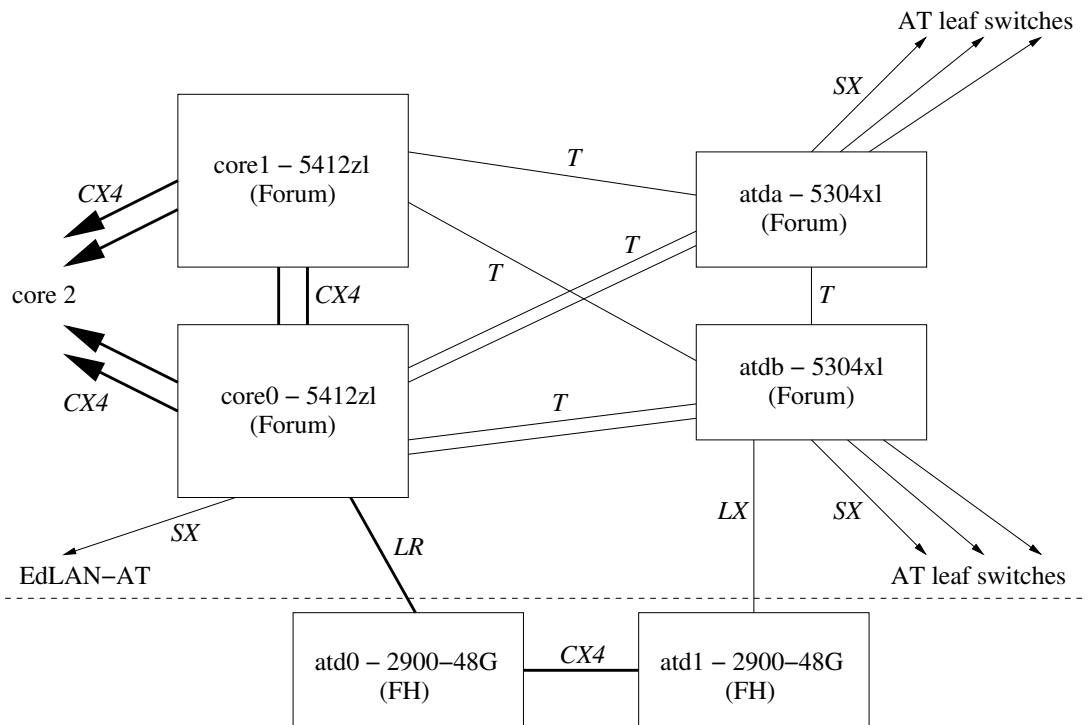


Figure 10: Appleton Tower decant arrangements

15 Moving out of Buccleuch Place and Forrest Hill

After we had moved out of the Buccleuch Place and Forrest Hill offices into the Forum we still had machines in the server rooms there for quite a while afterwards. In conjunction with IS, therefore, we decided to continue to manage the edge networks for the new occupants of the offices in the buildings using the existing switch infrastructure. This was reasonably straightforward on the whole, as there was not nearly as much VLAN churn as would have been produced by our own users. Once all our machines were out of Buccleuch Place we reset all the switch passwords and SNMP communities to temporary transitional values and handed the entire network installation over to IS to manage from then on. The same procedure was followed when we moved out of Forrest Hill, though this was delayed by the Appleton Tower basement server room decant.

16 Decanting out of Appleton Tower again!

In 2015 the Appleton Tower recladding project¹⁴ required us to decant out of the upper floors of the building while the work was in progress. Teaching went back to Forrest Hill, which had been extensively refurbished to suit our needs as much as possible, while non-teaching moved to the Wilkie Building. We moved out in Summer 2015, with the intention of moving back in Summer 2016. Additional building work, however, resulted in this being delayed by a year.

For simplicity it was decided to run both Forrest Hill and Wilkie as “virtual floors” of Appleton Tower, connecting back to the core using 10GbaseLR links over dark fibre. There were a couple of advantages to this:

1. machines could be taken off their AT desks, installed and connected in at their new locations, and would then continue to work with no further reconfiguration required; and

¹⁴<http://www.ed.ac.uk/projects-estates/current-projects/appleton-tower-recladding>

2. there was a cost saving, as core infrastructure was not required at either site.

The disadvantage, of course, is that it introduced outside dependencies, but given the expected duration of the decant this was felt to be an acceptable compromise. We did finally return to Appleton Tower in 2017.

17 IPv6

Recognising that IPv6 support was becoming more and more necessary, a development project¹⁵ was started in 2015 to investigate how this might be provided on the Informatics network, and to set up an initial service for managed machines. We joined the IPv6 Internet in February 2016, and by September 2016 all of our managed Linux subnets were IPv6-enabled.

As of August 2019 there were over 3500 AAAA entries in `inf.ed.ac.uk`, with another couple of hundred expected to be created soon after.

We decided to use BIRD¹⁶ as our IPv6 routing mechanism, and having found it easy to use and well supported we moved all of our IPv4 routing to use it too, ahead of the expected new-EdLAN changes.

18 Bayes Centre

The Bayes Centre occupies the remaining North-East corner of what was the “Potterrow car park” site.¹⁷ The building was completed in September 2018, with Informatics occupying three of the floors. For the most part, we manage the network on “our” floors, and IS manage the network on the rest.

The initial question was to what extent we should manage the building as a separate entity, or integrated with the Informatics Forum. The decision was the latter, for a couple of reasons:

1. Cost: we wouldn’t need core switches and the three Linux servers that a full “site” of this size would require; and
2. Shared VLANs: there would be many VLANs which would naturally be common across the two buildings, and quite a number which strictly needn’t be but would be simpler to implement if they were regarded as Forum offshoots.

The original intention had been to configure the network hierarchically, as we do in the Forum and Appleton Tower, but unfortunately some of the inter-closet copper wiring was not installed. Instead we operate each floor independently, with the switches connected in a 10Gbps ring, with connections back to the (Forum) core every second or third switch. The switches are not stacked, but instead connections are under spanning-tree control. Again, this arrangement may change with the new-EdLAN edge roll-out.

We’re not quite completely independent of IS, however. Because high-level and soffit-mounted data outlets are connected to the switches on the floor above, we do share one closet with IS. Fortunately the way the wiring is installed means we don’t really intermingle with their patch leads and switches.

There are also services in the building (such as BEMS, locks, café tills, CCTV and a bus-timetable display) which are attached to our part of the network but naturally connect back to IS’s part. As a result there are several VLANs which span across between the two. It’s not ideal, and does take a bit of coordination, but the alternatives (all managed by us or all managed by IS) were less appealing.

¹⁵<https://history.dcs.ed.ac.uk/archive/docs/Informatics-IPv6.pdf>

¹⁶<https://bird.network.cz/>

¹⁷<http://dtisite.inf.ed.ac.uk/>

19 New EdLAN

After what many would see as years of chronic under-investment by the University in EdLAN, the decision was eventually made that the active kit should all be replaced. By August 2019 the procurement phase had been completed and the project had moved on to the detailed design phase. Although input from Informatics and other Schools was provided into the process, it remains to be seen to what extent the features (mainly automation, devolved management and network login) will meet our needs.

Unlike the “old EdLAN,” which was mostly a bridged network, the new EdLAN is designed to be a fully-routed network, divided up into separate “VRFs”¹⁸ for overall robustness. VLANs can still be stretched across the network where necessary, but are carried in IP tunnels rather than as native bridged traffic. The core comprises two routers in AT and two at KB, cross-connected in a mesh, with “distribution routers” attached to the core, and edge switch stacks then connected to the distribution routers. There are separate appliances for network management, wireless control and firewalling. BGP¹⁹ is used as the underlying routing protocol. Note that even for those sites where there is a core-router presence there are still separate distribution routers, so as to keep the functionality clearly distinct as well as to allow for different connectivity options.

20 BGP for routing

Because EdLAN is transitioning away from OSPF to BGP, it was necessary for us to do likewise, at least so far as our interaction with it goes. Unlike OSPF, where we can basically set a few parameters and leave the routers to get on with it themselves, BGP requires a bit more in the way of explicit configuration. On the plus side, though, it does give us rather more control over which routes we send to and receive from EdLAN.

At the time of writing (November 2021), we were still using EdLAN essentially as an ISP.²⁰ We have direct 20Gbps connections to the Informatics Forum and Appleton Tower distribution routers, and an indirect connection to the Kings Buildings distribution router. We take only the IPv4 and IPv6 default routes, and we announce only our Forum routes to the Forum distribution router, our AT routes to the AT distribution router and our KB routes to the KB distribution router. As a transition arrangement, we do still (November 2021) speak OSPF to the old-EdLAN routers. We anticipate that this will no longer be necessary once the migration to the new distribution is complete.

Having found BIRD easy to set up and configure, we use that for both IPv4 and IPv6 routing, speaking OSPF internally as well as BGP to EdLAN.

People

The following were involved in the events described above:

- Peter Lindsay designed the Link boards
- Rainer Thönnies and Paul McLellan wrote the original Interdata Filestore system, with input from Hamish Dewar
- Rainer Thönnies designed the 2Mbps ether stations and wrote the Fred machine system
- Fred King designed the Fred machines (APMs)
- Gordon Brebner wrote the WAN software for ECSVAX; Paul McLellan wrote the DQS11 driver

¹⁸https://en.wikipedia.org/wiki/Virtual_routing_and_forwarding

¹⁹<https://en.wikipedia.org/wiki/Bgp>

²⁰It's still unclear how this will develop in future.

- the Author wrote the APM-Filestore system, based heavily on the “1976” Interdata version and with contributions by Rainer; the TCP/IP stack for APM and VMS; and the switch configuration and monitoring tools; and designed and implemented the Department’s and School’s data networks from the First Great Rewiring to date.

The following made some very useful comments on an early draft: John Davis, Hamish Dewar, Gordon Hughes, Mike Hurst, Nick Rothwell, Frank Stacey, Rainer Thönnies, Graham Toal. Thanks to Alan Boyd for the EdLAN documents.