

Making IT Work

22-23 May 2017

Proceedings

British Computer Society
Southampton Street, London, UK

The National Museum of Computing
Bletchley Park, Milton Keynes, UK

Contents

Martin Campbell-Kelly

Introduction 5

Doron Swade

The Historical Utility of Reconstruction and Restoration 7

Andrew Herbert, Chris Burton and David Hartley

The EDSAC Replica Project 22

Christopher P. Burton

Maintainability and Sustainability Issues in Restored and Replicated Computer Systems 35

John K. Chilvers

Curatorial Lessons from Other Operational Preservationists 44

Robert Garner

Restoring and Demonstrating 1960s Vintage Computers at the Computer History Museum 54

Robert Lusch, Nicholas Hekman, and Don Manning

IBM System/360 Printer Revitalization at TechWorks! 68

Johannes Blobel and Jochen Viehoff

The New Public Engagement in Computer Museums and in the History of Computing Machines 81

Benjamin J. Trethowan

Capturing, Restoring, and Presenting the Independent Radar Investigation System (IRIS) 91

Peter F. Linington

Constructing and Testing a Replica Store for the EDSAC 97

Nicholas Enticknap

Report of the TNMoC Workshop, 23 May 2017 107

Proc. Making IT Work, 22-23 May 2017

3

Introduction

Martin Campbell-Kelly

Editor

The *Making IT Work* conference, 22-23 May 2017, was jointly organized by the Computer Conservation Society of the British Computer Society and The National Museum of Computing, the UK's two leading institutions for computer conservation. The conference brought together the leaders of national and international computer conservation and restoration projects, including several recipients of the Google-sponsored Tony Sale Award for Computer Conservation. The conference enabled practitioners of computer conservation to share knowledge, and newcomers to learn about their motives and achievements.

The conference took place over two days. The first day consisted of a formal lecture programme held at the headquarters of the British Computer Society, Southampton Street, London, while the second day consisted of a workshop showcasing projects at The National Museum of Computing at Bletchley Park.

Following a welcome from the chair of the Computer Conservation Society, David Morriss, the first day opened with three of the UK's pre-eminent computer conservationists addressing the issues of utility, management, and sustainability. Doron Swade, well known for the construction of Babbage's Difference Engine Number 2, talked about the purpose and value of these activities — whether



Conference speakers and attendees at The National Museum of Computing workshop

Introduction

they be conservation, restoration, or replication. Andrew Herbert, formerly managing director of Microsoft Research in Europe, described issues of project management in building a replica of the EDSAC — easily the most ambitious recent project of its kind. Chris Burton, the doyen of British computer conservators, spoke from long experience about the chequered history of sustaining projects once they have been completed. The afternoon session began with a lecture from John Chilvers, a distinguished conservation architect and heritage consultant, who set computer conservation in the wider context of the preservation of cherished objects, from historic buildings to railway locomotives. The final three lectures of the day came from overseas speakers. Robert Garner described the restoration of a pair of IBM 1401 computers at the Computer History Museum, Mountain View, California — this project won the 2014 Sale Award. Next, Nicholas Hekman of the Vintage IBM Computing Center at TechWorks! in Binghamton, New York, described the restoration of an IBM 1403-N1 line printer originally supplied with IBM System/360 computers — the project was runner-up for the 2016 Sale Award. Lastly, Johannes Blobel and Jochen Viehoff described the building of an ENIAC accumulator for visitor interaction at the Heinz Nixdorf MuseumsForum — the winner of the 2016 Sale Award. For those able to attend, the day closed with a convivial dinner at Bella Pasta on the Strand.

The morning of the workshop at The National Museum of Computing was taken up with a series of presentations from project leaders of five of the museum's major restoration and replication projects. Following a welcome from Andrew Herbert, chair of the museum trustees, Delwyn Holroyd described the restoration of a 1980s-era ICL 2966, the museum's largest exhibit and one of the most powerful British mainframes. Next, Ben Trethowan talked about the IRIS, until recently a key component of the UK's air traffic control infrastructure. He was followed by Phil Hayes who described the rebuilding of the Colossus, the World War II electronic code-breaking computer. Next Peter Linington described the fabrication of a nickel delay-line memory for the EDSAC replica. Lastly, Kevin Murrell described the restoration of the Harwell Dekatron Computer, the world's oldest functioning computer. (Trethowan and Linington contributed papers describing their work to this volume.) During the afternoon attendees were given demonstrations of these five projects. So that they could see systems close-up and ask questions, workshop participants were divided into three circulating groups of comfortable size. A detailed account of the workshop by the computer journalist Nicholas Enticknap appears as the last paper in this volume.

The Historical Utility of Reconstruction and Restoration

Doron Swade

Royal Holloway, University of London

Abstract

During the last several decades there has been an increasing number of initiatives to reconstruct historic computing devices and machines, and to restore original specimens to working order. These projects pose many questions. To what extent is it allowable to compromise the physical integrity of the original to achieve a working system? What status does a “restored original” have as an evidentiary source? In the case of reconstructions, to what category of object do these new-old machines belong — fictitious antiques, new primary sources, public monuments? How important is physical fidelity to historical authenticity? What is the social and historical value of such projects to their makers, to the scholarly community, and to visitors who view them? This paper uses as case studies major reconstruction and restoration projects from the modern era to address a raft of questions agitated by such projects — museological, historiographical and social.

Introduction

During the last several decades there has been an increasing number of initiatives to reconstruct historic computing devices and machines, and to restore original specimens to working order. The starting point here is to take such activity as a given and, without making assumptions about the original motivations or how they might have been justified at the time, seek to reverse engineer the implications and issues that relate to their historical, museological and social meanings. The process is based on a suite of projects drawn from the major flagship projects of the modern era.

A word about “utility” in the title of this presentation. “utility” is not here a reference exclusively to “practical usefulness”. The concept is borrowed from economic theory where Utility is thought of as that property of a thing that makes it desirable to possess (Robinson, 1964: p. 48). So we might say that the utility of a knife is its ability to cut. One might say that the utility of art is its aesthetic appeal or its ability to affect one. Here utility is unrelated to practical usefulness. Entirely untethering utility from practical usefulness, we have the case of the architectural folly. We might say that a folly has cultural utility or socio-cultural utility as a demonstration of wealth or status, but it is not that a folly is practically useful that makes it desirable to possess. It is its very uselessness that confers meaning on the notion of its utility and that allows us to talk meaningfully of the utility of a folly, however paradoxical this might sound.

So what is it about reconstructions or restorations that makes them desirable? What is their value to history, to their makers, to their host organisations where these exist, and to their public and private audiences? None of the individual projects considered here are treated in any detail. Rather each is used as a placeholder for a particular aspect of utility. Many of the individual features are not unique to specific projects but are shared by many.

Case Studies

The pool of machines and projects to be drawn from includes:

Thomas Fowler’s Ternary Calculating Machine (c. 1840)

Charles Babbage’s Calculating Engines: Difference Engine No. 2 (1847-9); Analytical Engine (1836-71)

Konrad Zuse’s Z3 (1941)

Manchester “Baby” (1948)

Colossus (1943)

ENIAC (1946)

Ferranti Pegasus (1959)

Each is used here as a place-holder for particular features it exemplifies.

Thomas Fowler’s Ternary Calculator

Thomas Fowler was an impoverished self-tutored fell-maker who designed and built out of wood a digital three-state mechanical calculator demonstrated in the 1840s and witnessed by no less than Mr. Babbage and Augustus de Morgan amongst others. The poverty of sources here is extreme: not only did the physical original not survive but there are no plans or drawings. Fowler refused to release details of the invention, even for evaluation and endorsement by George Biddell Airy, the Royal Astronomer. Fowler was an outsider to the London scientific establishment and paranoid about his invention being stolen, and with good reason. He had earlier invented a thermosiphon for circulating hot water in heating systems and his unprotected invention had been ripped off by unscrupulous exploiters with no financial advantage to Fowler — something that aggrieved him till his death.

The sources for the reconstruction by Mark Glusker (Figure 1) are partial contemporary (19th century) textual accounts, one by de Morgan, and a memorial stained glass window installed circa 1870 in a church in Devon. (Glusker *et al.*, 2005).

Reconstructing context, purpose and technical detail from such shards is as much archaeology as it is technical history. This speculative reconstruction is a placeholder for the role of sources in reconstruction, and issues of historical authenticity.



Figure 1. Mark Glusker with the Fowler reconstruction (2005)

Babbage's Calculating Engines

Babbage's Difference Engine No. 2 (Figure 2), designed between 1847 and 1849, remained unbuilt until the modern era. No part of it was built before 1986 and the machine in its entirety (that is including the output apparatus) was completed in 2002 (Swade, 2005 and 2000). The sources are a set of 20 main drawings that are detailed enough to describe parts and their relationship to each other, but insufficiently detailed for manufacture. There is also almost no textual description of the design rationale or the logical function of the mechanisms. An understanding of the machine was essentially reverse engineered from the drawings, a process first embarked on by the late Allan Bromley.

The built machine is routinely referred to as a replica or reconstruction. It is neither. There is no original machine of which this is a repeat. It is strictly a "construction" and, being the first physical realisation of the design; it is original in one of the definitional meanings of that word — that of being the first.

Although 19th-century manufacturing techniques were not used, significant care was taken to ensure that the machine looked as though it had been made in the 19th century. Machining, rather than casting techniques, were used to produce the bright finish that was the hallmark of machining in Babbage's time. Contemporary practice was used as a guideline for choice of materials, and composition analysis was made of bronze parts used by Babbage in the 1820s to establish appropriate the grades of bronze available at the time. We welded where Babbage would have forged, but the welds were dressed to appear as those produced by 19th-century processes.

Tim Robinson in the U.S. has constructed in Meccano a difference engine based on Babbage's original designs for Difference Engine No. 2 (Figure 3). He has also made some assemblies from the Analytical Engine designs (Robinson, 2008).

It is clear that physical fidelity is not a *sine qua non* of meaningful realisation. What might we call such derivative constructions?




Figure 2: Charles Babbage's Difference Engine No. 2 (2002)




Figure 3. Tim Robinson's Meccano Babbage Engines (2008). Top: Difference Engine No. 2. Insets: Analytical Engine. Tim Robinson

Piers Plummer (2017) has constructed a difference engine based on Babbage's Difference Engine No. 2. The calculating wheels are made by 3D printing using sintered nylon, false coloured to look like bronze and the construction includes laser-cut stainless steel, unavailable in the 19th century (Figure 4). The machine is driven by a small steam engine. It has fewer columns of differences and lower digit precision than shown in Babbage's design and the machine is stretched vertically and flattened out front-to-back so that internal workings, obscured in the original, can be seen for pedagogic purposes. Again, what might we call such a construction?

The Babbage Engine and its derivatives here serve first as placeholders for terminology — what might we properly call different categories of physical realisation; and second, as placeholders for issues of fidelity — visual fidelity, material fidelity, and fidelity of process.




Figure 4. Piers Plummer Difference Engine (2017)

Konrad Zuse's Z3 (1941)

The Z3 is a general-purpose programmable computing machine developed in Germany during War World II using electromechanical relay components. The original was destroyed in the wartime bombing of Berlin and a reconstruction was completed 1967 by Zuse himself. The reason for inclusion here is not to do with the 1960s rebuild, but with the reconstruction of the machine led by Raul Rojas (2005). The reconstruction looks nothing like the original (Figure 5). Components are laid out on large PCBs to convey the logical architecture of the design and it is operated interactively by the user. Relay activity is indicated by LEDs and data paths are illuminated to show signal routing.

Rojas defines the overriding purpose of the physical reconstruction as pedagogic. In seeking to make the workings, operation and design of the machine intelligible to modern generations, he deliberately untethers the new machine from the traditional constraints of physical fidelity to the original. Rojas's reconstruction explicitly challenges the assumption of visual and material fidelity being an indispensable feature of reconstruction and, in our suite of projects, is a placeholder for the place of fidelity in its various forms in the *desiderata* of reconstruction. Documentary records for the original Z3, as well as the 1960s rebuild, were incomplete and essential design and operational features not well understood. A remarkable feature of this project was the excavation, information recovery, capture and dissemination of technical and design information assisted by both Zuse himself and his son. The research, findings, history of the machine, Zuse's papers, as well as simulations of its operations are captured and are publically accessible through the Konrad Zuse Internet Archive.




Figure 5. Reconstruction of Konrad Zuse's Z3 (part). Front view of processor.

Here physical fidelity is not regarded as essential to the purpose of its reconstruction. Also "reconstruction" here extends beyond the physical to include the reconstruction of lost data and knowledge of technical design and function.

Manchester "Baby" (1948)

The Small Scale-Electronic Machine (SSEM or the Baby) was the first computer to run an electronic internal stored program (June 1948). The Baby did not survive as a relic. It was a developmental machine that morphed through expansion into the Manchester University Mk. I, and was later cannibalised for spares. The machine was reconstructed by a team led by Chris Burton and the project offers us a master-class in systematic reconstruction with forensic attention to sources, and issues of authenticity (Burton, 2005).

Researching the technicalities of the machine and its design was an example of unrepeatable timeliness in that surviving original engineers and practitioners played an invaluable role in authenticating technical and operational features. Recovered memory was a feature of the process. Engagement with the machine demonstrated the power of physical realisation to unlock forgotten and undocumented information in the memories of the original engineers and practitioners — technical, operational, developmental and narrative. The project serves here as a placeholder for uncompromising attention to fidelity in all its aspects.



Figure 6. Manchester "Baby" (SSEM). Left: Chris Burton with reconstruction (1998). Right: original (c. 1948)




Figure 7. Colossus, Bletchley Park. Left: Tony Sale with reconstruction (2008). Right: Original (1943)

Colossus (1943)

Colossus was code-breaking computer built at Bletchley Park during World War 2 using electronic vacuum-tube logic and punched paper tape. The reconstruction was led by the late Tony Sale (Sale, 2005) a now legendary driving force behind the saving of Bletchley Park and the founding of The National Museum of Computing. All but two of the several Colossus machines were deliberately destroyed after the war. None are thought to survive. Documentary sources were somewhere between incomplete and non-existent. A few original components survived in private hands. The operation of certain parts of the machine were known, or could be inferred, but with less than detailed technical documentation reconstruction posed new challenges to historical accuracy. A part, for example, was known to rotate. I remember Tony saying that without a specification for the original motor, a motor from the windscreen wiper of a car would do (this was possibly in the context of the Heath Robinson). In the case of the motor we have functional replication without attempt at physical fidelity. The thing does the same as the original but not by the same means. The project is here used as a place-holder for the notion of functional replication.

ENIAC (1946)

ENIAC, operational in November 1945, was a vast primarily vacuum tube computing machine built at the University of Pennsylvania to calculate artillery firing tables. If one was to select a single machine symbolising the start of the age of electronic computing it is surely this. It had some 18,000 vacuum tubes and its success established to the wider community the viability of the large-scale use of vacuum tubes in electronic computing systems. Various cabinets and panels from the original machine survive but none of the original hardware is operational. Documentary sources are fairly comprehensive. Various reconstructions have been attempted. One of these is ENIAC-on-a-Chip produced to commemorate, in 1996, ENIAC's 50th anniversary. The inclusion of ENIAC here is as a placeholder for the reconstruction at Heinz-Nixdorf MuseumsForum (HNF) of an ENIAC accumulator by Jochen Viehof and Johannes Blobel. The HNF project won the 2016 Tony Sale award jointly with Robert Garner's restoration of the IBM 1401 printer at the Computer History Museum in California.

The HNF reconstruction allows members of the public to program an ENIAC accumulator in the same way as the original — i.e. using plug-in cables and switches. HNF has some original ENIAC hardware and this is displayed static and non-working alongside the reconstruction. The reconstruction itself is built from scratch and is driven by customised solid-state electronics and a microcontroller.



Figure 8. ENIAC. Left: Replica of ENIAC Accumulator, Heinz Nixdorf MuseumsForum (2016). Inset: ENIAC (1946)

What is distinctive here is the explicit aim to create a visual facsimile of the physical hardware as part of the look and feel of operating the original. Extraordinary measures were taken to replicate the colour and texture of the panels and fittings. It is a reconstruction designed as a usable interactive visitor attraction. In this it shares features with the Zuse Z3 project. The HNF ENIAC project seeks to reconstruct the experience of operational practice through operational and visual fidelity of a very high order.

Ferranti Pegasus (1959)

The Ferranti Pegasus is a vacuum tube drum-storage machine that first came into service in 1956. Some forty were built between 1956 and 1962. This particular Pegasus dates from 1959 and is in the Science Museum's collections. It was the flagship restoration project of the Computer Conservation Society when it was founded in 1989. It was successfully restored to working order and the restored machine ran successfully for over 25 years. It was retired a few years ago.



Figure 10. Ferranti Pegasus. Left: Engineer's Console. Right: Bit-level simulation by Chris Burton

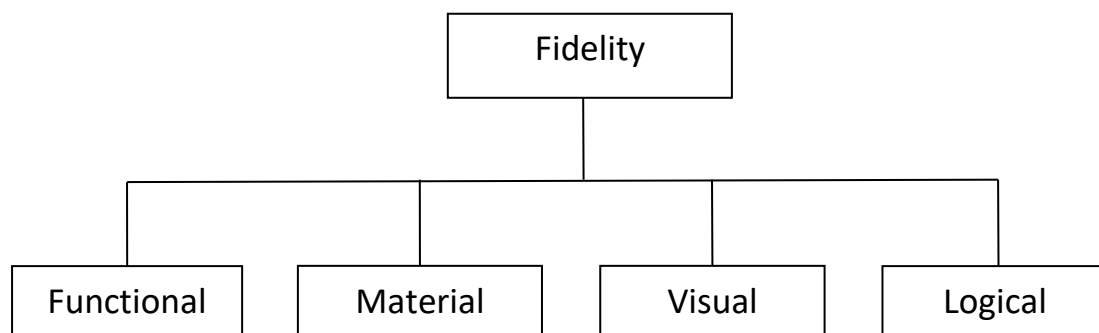
Its inclusion here is as a placeholder for computer simulations of historic machines. Computer simulations can be seen as a form of “logical replication”. In the case of Pegasus a bit-level simulation of the machine was created by Chris Burton. We can see from the screen images that little attempt is made to create a visual facsimile of the console controls beyond recognisable representations of console components and their layout. The functions that correspond to each console-control are scrupulously replicated.

The issue here is not one of visual fidelity but of operational and behavioural exactness of the simulation. The virtual machine can be operated (programs loaded from paper tape, executed, and results punched or printed) by clicking on the controls depicted on screen and the responses are exact, even to the extent of the signal waveforms on the oscilloscope screens. The concern here is functional rather than visual realism. In simulations of this kind we find another example of the abandonment of physical realism in replication and reconstruction (Swade, 2003).

Fidelity

None of the original machines, apart from the Pegasus, was available to act as a physical datum on which to base a reconstruction. Some were deliberately destroyed, others accidentally so. One was simply lost, and another did not exist in the first place.

In the pool of examples there are reconstructions that look nothing like their originals (Z3, Meccano Babbage). There is at least one example of a machine that looks as exactly as possible like the original but works by other means (HNF ENIAC). There are restorations of original machines that are functionally correct but that do not use original or even contemporary hardware. Clearly, fidelity to the original has a wide interpretation in this spread of projects. If we take fidelity to mean “faithfulness to the original” we can identify four forms:



In the case of functional identity the reconstruction does the same thing as the original but possibly by other means (Tony's windscreen motor, modern electronics in the HNF ENIAC and Zuse's Z3). Material fidelity applies where the reconstructed systems or subsystems are made of the same stuff as the original — original or contemporary components and materials (Manchester Baby). Visual fidelity entails facsimile appearance, as in the HNF ENIAC, which in this case also embraces tactile fidelity. Finally, a reconstruction is logically faithful if it implements the same rules as the original as in the case of the Ferranti Pegasus simulation (Swade, 2011). Some of the forms entail or presuppose others and projects are often faithful to the original in more than one respect.

Terminology

We use a variety of words to describe our agency in relation to artefacts. Some of these words are unproblematic. A fake is an object purporting to be an original but is not. We think we probably know what we mean by a simulation. What of "restoration", "replication", "reconstruction", "reproduction", "copy", "model", "rebuild" and so on? May we ask whether this vocabulary contains meaningful distinctions or whether there is a set of precise definitions that are internally consistent when describing these new-old artefacts?

By restoration we convey the sense of returning an object to its original state. In the case of devices or machines, restoration may or may not include returning the object to working order. The livery, upholstery, and paintwork of a vintage car, for example, may be restored, but with the engine left non-operational. The essential point here is that restoration invariably presupposes the existence of the artefact or an identifiable part of it to begin with.



Replicating a machine usually means creating a copy using an existing contemporary object as a reference datum. The etymology of the word (*replicare* — to fold back) implies a sense of "taking an impression". When a replica-maker minimises the physical differences between the original and the copy, the replica is exact. But replicas can be deliberately inexact. Replicas can be scaled copies either larger or smaller than the original, and there are non-working replicas made without any intention of functional fidelity. The case of Guatelli's half-scale non-working replicas of Babbage's original surviving Difference Engine No. 1 is an example. Here again the process of replication usually implies the existence of an original as a source.

So the use of "replica" as a descriptive can usefully be reserved for a copy for which there is an existing machine as a reference datum. We can then reserve "reconstruction" for the creation of a machine for which there is no surviving example. So the project described in the next chapter to realise the Cambridge EDSAC, of which there is no surviving original, would, by this suggested

convention, be called a reconstruction rather than a replica. In cases where part of the original survives and is used, the recreated artefact is a hybrid of original and reconstruction (Nielsen, 2011).

“Reproduction” usually refers to the act of making a copy, sometimes exact, but usually with the implication of imitation, especially in the context of furniture making. There is a pejorative implication that a reproduction is inferior in not being original. In the way we commonly use the word, the original of a reproduction may or may not exist.

Simulation can be seen as “logical replication” i.e. a non-physical form of duplication in which all essential logical predicates are preserved. Where an original exists, simulation can be seen as logical replication; where it does not, as logical reconstruction.

The use of the word “model” is an intriguing one. Of the many definitions and usages the one that is most relevant here is that of “a simplified representation for purposes of demonstration”. The essential feature here is that of a reduction of some kind — of scale (Guatelli’s half-scale replicas for example), capacity (Plummer’s Difference Engine), or reconfiguration (Plummer’s DE; Robinson’s Meccano realisations). The original Difference Engine No. 1, dating from 1832, was assembled for purposes of demonstration and, while full-scale, it is only one-seventh of the capacity of the whole machine as designed. So it could arguably be described as a model, as it possess at least two of the defining features if this attribution — demonstration and reduced capacity.

There is a limit to how far one can pursue these distinctions and retain sanity. It is also the case that reality has an irksome tendency to defy exhaustive taxonomies. But thoughtful use of how we use these descriptors will help avoid historical misconception and also help us to be clear in our own minds about what aspect of originality we are being faithful to.

Utility

It is clear from our suite of projects that their utility is not coextensive with their value as historical evidence in traditional curatorial or conservationist terms. So how might we articulate other forms of utility exemplified by these projects? What makes them desirable to pursue?

I would suggest that two categories are useful here: *Historical utility* and *Social utility*

Historical Utility

The following are aspects of the value to history of ventures exemplified by the suite of projects used as our case references.

Draw one into a level of intimate detail with the machine that rarely occurs by other means

Anyone who has been involved in reconstruction and restoration projects will immediately recognise what is meant. The attention to detail goes beyond general understanding of principles of operation. It is not even that the TTY is cabled to the main cabinet or even that it is a proprietary connector not a DB25 that terminates the cables. It goes down to the level of whether pin 5 at one end is connected to pin 7 at the other and what is the colour of the conductor. Involving oneself in the finest level of granular detail is unavoidable from the top all the way down — architecture, logical organisation, circuit and system functions, specific circuits, components, pin-outs and wiring.

Contingent or unexpected findings not foreseen or foreseeable by analysis or theory

There are countless instances of contingent findings in many of the projects. In the case of the Babbage Engine it was not clear ahead of the build how fast the Engine could be driven and what the limiting factors might be, nor was it practical to try to predict this by analysis. The limiting mechanism turned out to be an intermittent circular motion mechanism — an unexpected

finding. Theoretical prediction, even if attempted, would necessarily have been speculative and would not carry the conviction of an empirical finding.

In another instance the function of two horns on the carry levers was not completely understood before building it and they remained a partial mystery. Adhering faithfully to the original 19th drawings the mechanism was built as drawn and we were dumbfounded at the subtlety of its operation and the ingenuity of its function in ensuring the integrity of calculation (Swade, 2005).

Insights into contemporary ideas and practice

The Difference Engine No. 2 construction completed in 2002 was the first realisation of a complete Babbage calculating engine. No one in the 19th century had the practical benefit of a complete Babbage engine, so there is no operational data on how practical it would have been or whether it was a going proposition at all. With a finished working difference engine to hand we are better placed to explore operational issues: how many people would it take to run; how long can a single operator drive the machine by hand before tiring; how many operators need to be on hand to act as relays for continuous operation; how many attendants are required to mix plaster to replace the stereotype trays in the output apparatus; whether it is possible to phase the mixing sequence so that plaster drying times ensure an uninterrupted supply of trays with exactly the right consistency of plaster; how reliable is the machine; is there any economic case for the machine in the event that it requires more than n attendants to run it, and so on. We can interrogate these objects about key features of contemporary practice in ways that would not otherwise extend beyond speculation, and we can measure contemporary expectations against actual performance. This is experimental history in a very pure form.

Tacit Knowledge — we can learn what it might have been like

To witness Pegasus clacking out square roots on a teletype conveys what it might have been like in ways not communicable by other means. If one cranks the handle of Babbage's Difference Engine there are shock loads at several points in the cycle as locks go in and out. The feel of this is not communicable by anything other than the experience of doing it. Nor is the sound the machine makes fully describable, particularly the rhythmic slapping sound of the locks which is the aural signature of the machine working. The carry mechanism was well-understood but what was not foreseen was the dramatic spectacle of the wave-like rotation of the double helices at the rear of the engine. Sight, sound, and feel. A cocktail of experience that text cannot convey.

Documentary Completeness — new material, recovered memory, e-archaeology

Researching a machine in detail almost invariably involves gathering all known information in one place. Co-locating and collating known material has substantial archival value.

There are instances of new material emerging attracted by the public visibility of the project. The current EDSAC reconstruction is a case in point. One of the EDSAC 2 engineering team had a cache of drawings he removed from a cupboard when the original EDSAC was scrapped. Hearing about the EDSAC project reminded him of their existence and he retrieved them. There are other examples of this — the Babbage family sent priceless material from the family archive when they read about the Babbage Engine construction in the international press.

There are instances of recovering lost knowledge from pioneers or practitioners for whom re-engaging with the machines unlocked memories, information and operational practice they had forgotten they knew. The driver circuits for the Williams Tubes used in the Manchester Baby is one of many in which contemporary practitioners supplied essential information that had not been documented.

There are instances of recreating lost knowledge. The hard disk on the Elliott 401 is an instance. There was no data or metadata about the disk's contents. Tony Sale and Chris Burton designed and made the instrumentation to read the bit-stream from the disk, capture this on a

PC, analyse and recover the content and structure. The data was intact after 50 years. 35 mm magnetic tapes for the Elliott 803, stored in cans in a garage whence they were recovered, were found to be readable 40 years later. Both these findings have implications for digital curation, especially at a time when the impermanence of magnetic media and its longevity remain pressing issues. These two examples are as much part of electronic archaeology as they are computer history (Link, 2016, pp. 79-112).

In the case of restorations there is nothing like working from a circuit diagram to expose discrepancies between the object and its technical description. It is not uncommon for the documentation of the last hardware version of a machine or system not to have been fully updated. There are instances of modifications made that are not recorded, especially in the final developmental stages, as well as local variants of machines that do not conform to the generic descriptions. Restoration to working order and subsequent maintenance will more than often involve a forensic examination of technical description and an assessment of their correspondence with the physical machine.

Physical Completeness

When large and medium-sized computers are acquired they are almost invariably uncabled and dismantled into manageable units for transport. The process of acquisition — the taking of a machine into protective custody is, paradoxically, sometimes the most traumatically destructive process in the life of this supposedly prized machine. Physical reassembly is the first stage of restoration, which in itself creates a physical datum of greater historical probity than leaving the machine in store in pieces. Getting the machine operational tests and verifies electrical completeness. Ultimately, when switched off, one has an exact knowledge of the machine's electrical and physical state of completeness as a direct outcome of restoration activity.

Benchmark for Simulation

It is likely, and probably inevitable, that ultimately, on archaeological time scales, none of the restored or reconstructed machines will be working at some distant or not so distant future time. Bit-level simulations of these historic systems and migrating these from platform to platform hold the prospect of a kind of logical and operational immortality for computing machines. Restoring a machine to working order allows the simulation to be bench-tested against the live original, especially in relation to execution speeds. So “flaring” a machine into working order for even a relatively short time offers the benefit of verifying the conformity of a simulation to the behaviour of a physical original or a restored or reconstructed physical machine.

Social Utility

There is another dimension to utility: the value to those who make up the project teams, to visitors who view and experience the outcomes, and to museums that host the machines and sometimes the activities.

Meaningful Work

Reconstruction and restoration projects provide social and organisational context for veteran experts to share their expertise, exercise their craft, and extend their professional activities in historically and educationally meaningful ways.

Pedagogic Value

For participants in such reconstructions and restorations, as well as for the users and visitors, such projects are the source of both explicit and tacit knowledge as exemplified by the Z3 and HNF ENIAC reconstructions.

Nucleus — focus of visitor attention

Reconstructions and restored machines on display act as a focus for visitor attention and a nucleus for story-telling. Babbage, Pegasus, Colossus, Bombe, EDSAC are large machines, sculptural spectacles in their own right that monumentalise and memorialise episodes in history of which they are an expression. Babbage's Difference Engine No. 2 was built in public view. A commonly asked question was "how do you make this or that?" This was unexpected. The public build prompted conversations about manufacturing, for example, as well as about what the machine was and how it worked.

Cultural objects

Large systems, publicly displayed, memorialise important and often extraordinary episodes and practices for which they act as both placeholders and monuments. They serve as generational bridges to an otherwise lost past.

Conclusion

In articulating the historical and social utility of restoration and reconstruction projects, it is clear that they are desirable to different constituencies for a variety of different reasons. They are of value to museums, curators, archivists and historians concerned to preserve physical artefacts, documentation, narrative histories and operational knowhow. To visitors and users who experience them, they are formidable public attractions that memorialise major events and developments and serve as nuclei for conversation, interaction and speculation. To the practitioners, veterans, experts, and the project teams who are mostly, though not exclusively, volunteers, they provide an arena of meaningful engagement and their activities carry the satisfactions of exercising expertise, knowledge and skill to constructive and valuable ends. The growth of the number and variety of projects concerned with historic computing machines appears to signal that the drive to undertake these projects remains irrepressible, and the satisfactions to those who undertaken them are mirrored in their value to others.

References

- Burton, C. P. (2005). Replicating the Manchester Baby: Motives, Methods, and Messages from the Past. *IEEE Annals of the History of Computing*, 27, pp. 44-60.
- Glusker, M., Hogan, D. M. and Vass, P. (2005). The Ternary Calculating Machine of Thomas Fowler. *IEEE Annals of the History of Computing*, 27, pp. 4-22.
- Glusker, M. (2005). *The ternary calculator of Thomas Fowler* [Online]. Available: <http://mortati.com/glusker/fowler/> [Accessed].
- Link, D. (2016). Enigma Rebus: Prolegomena to an Archaeology of Algorithmic Artefacts. In: Link, D. 2016. *Archaeology of Algorithmic Artefacts*. Minneapolis, MN: Univocal.
- Nielsen, S. (2011). The Sea Stallion of Gendalough: Reconstructing a Viking-age Longship. In: Stabermann, K. (ed.) *Reconstructions: Recreating Science and Technology of the Past*. Edinburgh: National Museums of Scotland.
- Plummer, P. (2017). *Computing by Steam V2* [Online]. Royal Holloway, University of London. Available: <http://babbage.csle.cs.rhul.ac.uk/> [Accessed 20 June 2017 2017].
- Robinson, J. (1964). *Economic Philosophy*. London: Penguin.
- Robinson, T. (2003-2008). *Tim Robinson's Meccano Website* [Online]. [Accessed 20 June 2017 2017].

Rojas, R., Frank, D., Goktekin, C. and Heyne, G. (2005). The Reconstruction of Konrad Zuse's Z3. *IEEE Annals of the History of Computing*, 27, pp. 23-32. *Konrad Zuse Internet Archive* [Online]. Available: <http://zuse.zib.de/>

Sale, A. E. (2005). The Rebuilding of Colossus at Bletchley Park. *IEEE Annals of the History of Computing*, 27, pp. 61-69.

Swade, D. (2000). *The Cogwheel Brain: Charles Babbage and the quest to build the first computer*. London: Little, Brown. US edition: 2001. *The Difference Engine: Charles Babbage and the Quest to Build the First Computer*. New York: Viking.

Swade, D. (2003). Virtual Objects: the end of the real? *Interdisciplinary Science Reviews*, 28, pp. 273-9.

Swade, D. (2005). The Construction of Charles Babbage's Difference Engine No. 2. *IEEE Annals of the History of Computing*, 27, pp. 70-88.

Swade, D. (2011). Reconstructions as experimental history: historic computing machines. In: Staubermann, K. (ed.) *Reconstructions: Recreating Science and Technology of the Past*. Edinburgh: National Museums of Scotland.

Picture Credits

Figure 1: Mark Glusker; Figures 2, 3, 4, 9, 10: Doron Swade; Figure 5: Raul Rojas; Figure 6: School of Computer Science, University of Manchester; Figure 7: National Archives; Figure 8: Heinz Nixdorf MuseumsForum.

The EDSAC Replica Project

Andrew Herbert, Chris Burton and David Hartley

The Computer Conservation Society

Abstract

A description of the EDSAC Replica Project is given, describing the creation and organization of the project and the research undertaken to inform the reconstruction of a functional replica of an important early British computer. A detailed account is given of the design choices made during the reconstruction, the techniques used, and the challenges encountered in balancing the desire to remain authentic to the original whilst producing a demonstrable museum exhibit.

EDSAC

EDSAC — the Electronic Delay Storage Automatic Calculator — was an early electronic digital computer built by the University of Cambridge between 1947 and 1949. EDSAC was conceived by Maurice Wilkes, the director of the Mathematical Laboratory, and built by a team led by Bill Renwick, recruited by Wilkes to be the chief engineer.

EDSAC (Figure 1) ran its first programme on 6 May 1949. By 1950 EDSAC was providing a computing service to the university. Three Cambridge scientists of the period — Kendrew (1962), Ryle (1974) and Huxley (1963) — acknowledged the contribution EDSAC made to research that won them Nobel Prizes.

EDSAC's principal claim to fame is as the first practical machine to provide a regular, open to all, day-to-day computing service to a community of users. Important to the success of EDSAC was the programming system devised by David Wheeler, offering a mnemonic programming code combined with a substantial library of useful subroutines and debugging aids.




Figure 1. EDSAC in 1949

We had two important sources of information about how EDSAC was used in those early days. A short film was made in 1951 showing how programs were written and run. This was re-edited in the early 1970s and provided with an informative commentary by Wilkes. Contemporaneously with the film, Wilkes, Wheeler and Gill (1951) wrote *The Preparation of Programs for a Digital Computer*, the world's first computer programming textbook which describes how to program EDSAC in detail with information about the operating procedures, order code and subroutine libraries.

EDSAC was scrapped in 1958. A few chassis and memory tanks were kept and distributed to museums. An archive of photographs and documents was placed in the University Library, which, in later years, Wilkes was to say would provide sufficient information to reconstruct EDSAC, not that he could see any point in doing so.

The Project Starts

The EDSAC Replica Project originated in an unplanned meeting at a Cambridge University event for benefactors in early 2010 between David Hartley, a former director of the University Computing Service (which grew out of the Mathematical Laboratory) and Hermann Hauser, a Cambridge-based technology entrepreneur. Hauser asked Hartley what it would take to reconstruct EDSAC and Hartley agreed to find out. Being an active member of the Computer Conservation Society (a specialist branch of the British Computer Society), Hartley was aware of the earlier project to reconstruct the Manchester Small-Scale Experimental Machine (the "Baby") and asked the leader of that project, Chris Burton, if he would undertake a feasibility study.

Burton reviewed the surviving documents and photographs and prepared a report. He noted that EDSAC was nearly three times the size of Baby and that a great deal of research and experimentation would be needed to reconstruct the circuit designs. Moreover, unlike Baby where survivors of the original team were able to assist in the reconstruction, sadly none of the EDSAC pioneers was still alive. The good news was that the photographs revealed that the circuits used thermionic valves that were still obtainable from specialist dealers as "new old stock" and there was evidence that Wilkes and his team had used well-known circuits of the period with a great deal of replication of common elements across the machine. A budget of around £250,000 was suggested as necessary, assuming the reconstruction would be built using volunteer labour.

Hartley and Hauser decided Burton's report gave sufficient grounds for believing the project was viable. The first step was to create a charity that could raise funds for the project and own the finished machine. The Computer Conservation Society, The National Museum of Computing (TNMoC), the University of Cambridge and the Hauser-Raspe Foundation (which provided initial funding) were invited to nominate Trustees. Burton was asked to become Chief Designer and the recently retired head of Microsoft's Cambridge Research Laboratory, Andrew Herbert, was appointed as Project Manager.

The National Museum of Computing was chosen as the location where the replica would be displayed. It was decided that from the outset the reconstruction would be a museum exhibit so that visitors could see the project progress and talk to project volunteers to learn what was going on — and indeed by so doing we have recruited several more volunteers!

Fairly quickly, over £150,000 was raised from corporate donors and Cambridge University alumni with links to the Computer Laboratory. A second round of fund raising in 2015 spread the net further to include alumni in Silicon Valley, and succeeded in reaching the desired total of £250,000.

Burton and Herbert found volunteers through the Computer Conservation Society and personal networking. A project website was set up by TNMoC to advertise the project, and a first meeting of interested parties called for 24 March 2012. About 20 people attended, including Trustees, to hear talks about the museum, the feasibility study and the early forensic work done by Burton. At the

end attendees were asked to volunteer to undertake research into different aspects of the machine. The EDSAC Replica Project was underway.

Over time, the volunteer team has grown from the initial core to a membership of about 20 people. and at the time of writing most of the construction of the replica is complete and the team are deeply invested in commissioning major sub-systems with the hope of being able to run simple programs by the end of the year (2017).

We are privileged to have a number of older engineers who started their careers in the thermionic valve era, some working for early British computer companies and others who learned valve electronics from servicing military radar and radio systems during National Service. Alongside them work younger engineers for whom valve circuits have been a new technology to learn. One of the team, Peter Lawrence, knew the original EDSAC — he used it while a Cambridge PhD student in the 1950s.

As the project progressed the volunteers naturally divided into separate teams looking after distinct areas of the machine, such as main control, arithmetic unit, store access and so on. Each team had a lead engineer — who was responsible for the design of the circuits in that area — and others who helped with the construction and testing of chassis, and then with the commissioning of the machine as each area was completed and connected up. During the commissioning phase, team members were encouraged to develop an understanding of further areas of the machine to help spread expertise and avoid delays in fault finding and diagnosis.

To help these informal groups coordinate we set up a simple, formal project structure. Herbert, as Project Manager, took overall responsibility for management and coordination, fund raising, procurement and external communications. In this he was aided by Alan Clarke who took on the role of Storekeeper. Clarke was responsible for finding sources of valves and other components, managing the relationship with the engineering firm that undertook fabrication of racks and chassis and issuing kits of parts to “wireman” volunteers as construction proceeded. Burton, as Chief Designer, was as the “signoff” authority, checking each chassis design was a plausible circuit, as authentic to the original as could be achieved, and consistent with project design and construction principles.

We documented our research and subsequent designs as a series of “Hardware Notes” held in a Dropbox shared file system. Circuit schematics and mechanical drawings were kept as CAD files. The intention is that at the end of construction these documents will be consolidated into a database along with guides for maintenance and fault diagnosis for the finished machine.

As our volunteers were dispersed widely around the UK (and at one time included an Australian member) we relied heavily on email and Dropbox for coordination, enriched by regular “all-hands” meetings held approximately every three months. These gave an opportunity to report progress, demonstrate work in hand, discuss issues and plan next steps. In the early days the meetings were often robust, wide-ranging debates as we tried to pull together the research on individual sub-systems into a coherent understanding of the EDSAC and establish project standards for common circuit elements and construction techniques.

When the project entered the commissioning phase the volunteers’ meetings were augmented by more frequent short “integration meetings” at which issues arising from interconnecting sub-systems for the first time would be resolved — e.g., signal levels, timing, test plans and so forth.

To further provide a record of the project we were fortunate to be approached by David Allen, a retired TV producer who has worked for the BBC, and indeed produced many of the BBC’s early programmes on computers and programming in the 1980s. Allen has videoed many of our volunteer meetings and working parties, and recorded interviews with key members of the project team. His

videos have been put on the project website to keep the wider community informed of progress and to explain some of the engineering challenges we have faced throughout the project.

Technical Approach

Given the incomplete documentation available, it is difficult to determine if our reconstructed EDSAC is indeed a replica of the original — and as EDSAC was subject to continual improvement and modification, it is a moving target. Our aim was to build a *functional replica*, i.e. a machine that would run original EDSAC programs, and would look identical in terms of physical layout of racks, chassis and valves etc., to the photographs of the original. Where we had to fill in gaps, we would use circuit design principles from contemporary practice informed by textbooks of the time and experience of volunteers who had started their careers in the valve era.

The photographic record presented several challenges to interpretation. The photographs were undated and clearly showed the machine at different stages of its life — some were taken during construction, others showed the machine in use.

We initially set our target date as 6 May 1949 when the machine ran its first program¹. We had a few images we knew to have been taken at around this time to provide guidance, and while we knew the early machine could be unreliable we would be replicating a baseline EDSAC against which we could evaluate images and documents relating to later versions. Subsequently we decided to advance the target date to 1951 when EDSAC was providing a day-to-day computing service. This was principally to allow us use of the initial orders written by David Wheeler, to allow use of the library subroutines described in Wilkes, Wheeler and Gill (1951) and the operator's desk shown in the film. By moving to this date we would be able to demonstrate the machine showing how it was programmed and used by the university and better tell the story of EDSAC's contribution to computing history.



Figure 2. Wilkes with EDSAC mercury delay lines

¹ In the early days the Cambridge team used the term 'programme,' but then fairly quickly adopted the American spelling 'program'.

Alongside recreating the electronics of EDSAC, we had the further challenge of replicating the store that was constructed using mercury delay lines. The main store of EDSAC consisted of (eventually) two sets of 16 steel tubes each just over five feet long, each set mounted in a rigid frame (Figure 2). At each end of a tube was a piezo-electric crystal transducer, and the tube was filled with mercury. To reconstruct this would have been difficult for the project to achieve: the tubes and end plates needed to be very precisely machined, mercury is prohibitively expensive and filling the tubes presented health and safety challenges. Moreover we knew from contemporary reports that in the EDSAC the mercury was periodically redistilled to remove impurities, a task that would be impractical in a museum environment.

Fortunately the EDSAC logical design came to our rescue. Wilkes was aware that delay lines were an imperfect solution and lived in hope of something better coming along to replace them. The design of the interface between the computer and the store is very simple: the computer drives a simple recirculation system of RF modulated pulses via an input and output coaxial cable for each tube (or tank in EDSAC parlance). In principle, any sort of delay line can be interposed and so we decided to use nickel delay lines in place of mercury. This is anachronistic in that nickel delay lines are a feature of machines from the mid-1950s onwards, but they work on an analogous principle to mercury delay lines — and importantly we could design a nickel delay line store that connected to the EDSAC store interface without modification. This aspect of the project is described in a companion paper by Peter Linington (2017).

Mechanical Design and Construction

The mechanical design of the replica was established by scaling measurements taken from photographs that reveal the details of the rack construction. We were further aided by having access to a surviving chassis from the Computer Laboratory collection.

The physical layout of EDSAC consists of three rows of racks, each rack containing approximately 12 chassis, comprising a total of 142 chassis. In total there are 71 different types of chassis layout.

The racks consist of a substantial metal frame built from steel angle with thin cladding to the top and sides. The feet of each rack are made of more substantial angle and plate sections to give the structure stability when filled with (typically) 12 chassis weight of electronics.

Each chassis consists of a folded shelf with bolted on end plates to make it rigid, and in some cases additional thin plates for shielding purposes (Figure 3). Each shelf is punched with holes as required to take holders for valves, tuning potentiometers and capacitors, grommets for feeding through leads from valve caps, and so forth.

Along the front of each shelf there are typically two rows of test points, the upper row for measuring voltages at key points in the circuit and the lower for measuring current flow. Underneath each shelf are found two Paxolin tag strips and, following construction principles of the time, passive components are wired between tags on valve bases and tag strips.

On the rear of each shelf is a custom interface tag strip with tags for input and output signals for connecting to other parts of the machine. This wiring is point-to-point, and where wires have to cross from one row of racks to another they pass up to a gantry at the top of the rack where they are looped across to the gantry of the destination rack. The gantries consist of upright studding supporting notched Paxolin beams.

The pair of circuit tag strips required for each chassis was made commercially. The interface tag strips were custom made for each chassis from blank strip and loose tags. The gantries were likewise custom made.




Figure 3. Original and replica EDSAC chassis

A Cambridge-based engineering company manufactured the racks and chassis plates to our order. We provided them with CAD drawings for the racks and a generic chassis. As the project proceeded, we further supplied a CAD template for the hole-punching required on each specific chassis. The racks and chassis were supplied as cut, drilled, punched and painted kits ready for assembly.

On the original machine, power (250 V DC circuit HT and 250 V AC mains) was distributed to further tag strips on the chassis end plates. This is not acceptable in terms of modern standards for electrical safety and so we modified the rack design to include a concealed power wiring with modern fused plugs and sockets.

The layout of the EDSAC gallery at TNMoC (Figure 4) was based on diagrams of the layout of the EDSAC room in a report by Dodd and Glennie (1951) and photographs of the original machine. Being allocated half of the gallery, we were able to have the front and one side of EDSAC exposed to the public and we cut windows into the rear wall to further increase visibility.

A raised wooden floor was constructed to allow cables to mains sockets to be run underneath the machine and to carry the approximately two-ton weight of the completed machine. An area at the back of the machine was set aside as a work area for the volunteer team and a steel and glass barrier erected around the machine to keep the public at a safe distance from the exposed electronics. Information boards about EDSAC were attached to the barrier to explain the exhibit to visitors. We took advantage of the space available to increase the rack spacing compared to the original machine, both in the interests of the safety of our engineers and to give the public a better view into EDSAC.

A lighting frame was hung from the ceiling above the machine to illuminate it. Ducting was also suspended from the ceiling to carry power from the power supplies to each rack. Later further ducting was added to carry cables from the storage regeneration chassis to the boxes holding the delay line stores. The ceiling, lighting frame and ducting were all painted matt black to conceal them as much as possible.

Electronics Design and Construction

The baseline for the electronics design of the reconstruction was a document in the University library entitled *The EDSAC Report* (Anon., n.d.) which gave a detailed description of the principles of

operation for EDSAC along with a large number of logic and timing diagrams and a circuit diagram for a “flip-flop”. This report had been written by a visitor to the Mathematical Laboratory some time during the construction of EDSAC, so we were uncertain how closely the description corresponded to the original. We also found handwritten circuit diagrams for the clock and digit pulse generator circuits. The surviving storage regeneration unit (chassis type 01) showed us the typical assembly of a chassis and its operation was described in an article in the journal *Electronic Engineering* (Wilkes and Renwick, 1948). There was further information in surviving engineer’s notebooks and papers. Finally the preserved EDSAC logbook and “Operating Memoranda” gave us some hints of when various new features and updates to the machine were first introduced.

The first step towards reconstructing EDSAC’s circuits was the development of a logical model and simulator by Bill Purvis, working from *The EDSAC Report*, a task that was carried out in parallel with Burton’s feasibility study. (Purvis was introduced to the project by Burton and they had worked together on the Baby reconstruction.)

Purvis’s simulator, called ELSIE, inputs a gate-level description of the machine and runs a simulation with a clock step of an assumed per stage delay of around 80 ns. The simulated store can be loaded with a suitable test program and the output consists of a graphical trace of selected signals (Figure 5). Using this system Purvis checked the consistency and correctness of the logic elements described in *The EDSAC Report* and demonstrated the execution of simple programs consisting of two or three instructions. In several areas it was necessary to add “black box” elements to simulate aspects of the machine not covered in the report. It was also necessary to add additional delay elements to many circuits to get the overall system timing right.



Figure 4: EDSAC gallery and racks under construction




Figure 5. Typical ELSIE output

Helpfully, the logic diagrams in *The EDSAC Report* (Figure 6) were labelled in a way that suggested the assignment of logical functions to chassis. So once the simulation was working, the corrected logic diagrams could be divided up into logic circuits to be given to other team members as the basis for designing a suitable circuit.

Initially, we had hoped to be able to read chassis labels from the photographs and identify these with the functions described in the report, but sadly the resolution was rarely good enough to make this possible (Figure 7). We had instead to work backwards from circuit designs to match the valves required to those visible on chassis, and where necessary iterate the designs to exactly match plausible chassis locations. Over time the team built up a map of functions to chassis and from chassis to positions in the racks that closely correspond to the photographic evidence. In a few cases we found chassis had more valves than our circuits described — some cases could be explained away as additional amplifiers to drive outputs or clean up inputs, others as elements that perhaps turned out to be redundant once the original machine was commissioned. There was also some evidence of busy chassis borrowing spare space on less crowded ones to further complicate matters.

Working from the few circuit diagrams and by careful forensic analysis of the photographs and bench experiments, Burton produced designs for the common circuit elements used in EDSAC — including gates, flip-flops and amplifiers along with recommendations for signal levels, noise reduction techniques, use of test points and inter-stage coupling.

We took two important design decisions at this stage.

The first was to use modern resistors and capacitors rather than attempt to re-use period examples. Whilst we received many donations of 1940s and 1950s components, their quality was poor and would have made reliable working of the replica difficult to achieve. The practical impact of this choice was a lack of visual authenticity — modern passive components are much smaller than their historical counter parts and coloured differently (Figure 8). But, as in EDSAC they are only found on the underside and rear of each chassis and not as prominent as the valves on the topside, we felt this was acceptable. An alternative approach followed by those who restore vintage wireless equipment would have been to encase modern parts in moulded shells to resemble original parts: the expense and time required to do this was felt to be prohibitive in our case, for little benefit.

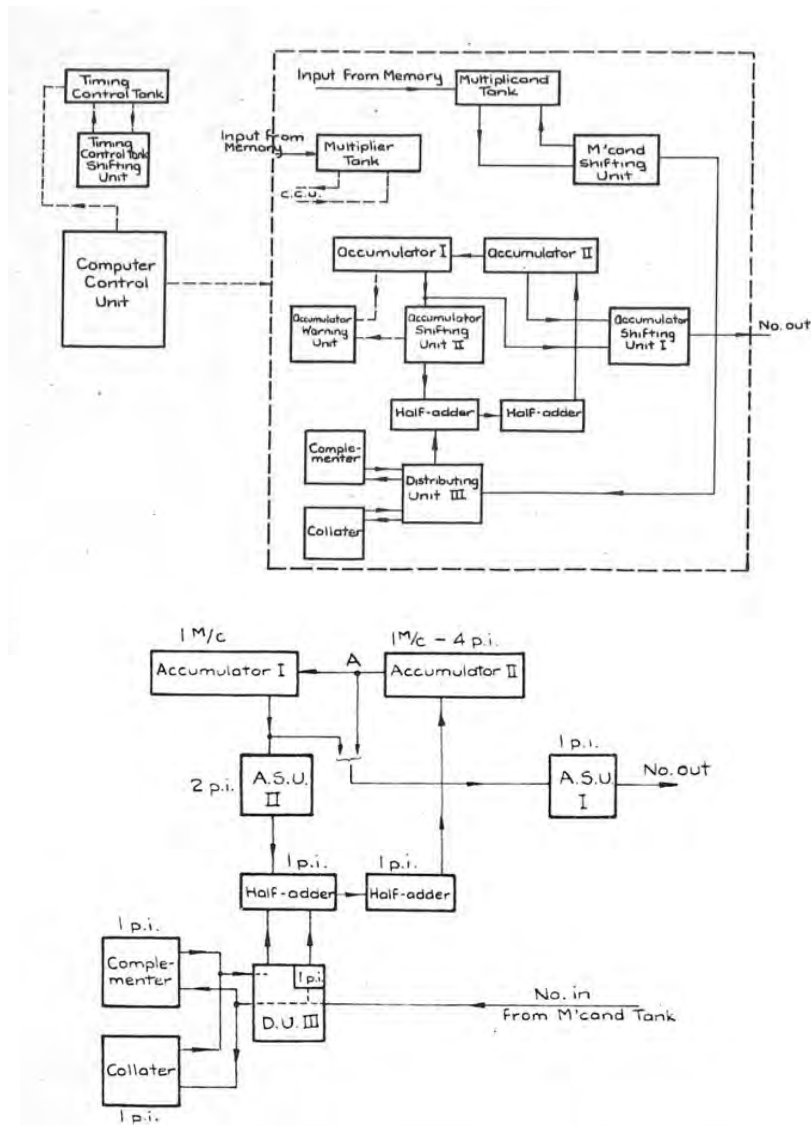


Fig.13. DELAYS IN COMPUTER

Figure 6. Logic and timing diagrams from EDSAC Report



Figure 7. Typical chassis label image





Figure 8. Original versus replica chassis

Our second major design decision was to add a 0 V “signal ground” rail to each chassis rather than grounding through the metal of the chassis as had been done in the original. While this compromised authenticity we felt it would make our circuits safer, particularly for those volunteers who had grown up in the era of low voltage transistor circuits and not experienced in the hazards of working with 250 V DC HT. In retrospect, this has created as many problems as it solved. Running across the length of each chassis, the rails act as an antenna and pick up ambient noise. This was a particular problem for the storage regeneration units — our test equipment, including oscilloscopes, has to be powered from mains earth isolated transformers to avoid tripping the computer power supply.

From testing we found few problems with using vintage thermionic valves: the EF54 and EF55 pentodes, and the EB34 and EA50 diodes used by EDSAC were readily obtainable from specialist suppliers and at reasonable prices when ordered in bulk. Many of these on delivery were found to be in wartime packaging with the oldest date-stamped 1943. Less than 10 per cent were found to be faulty. We were also able to obtain a good number of VCR97 cathode ray tubes for use in the display units on the operator’s desk.

The most problematic component to find were valve holders, in particular the B9G type required for the EF54 and EF55 valves. These have been out of production for many years and while we were able to obtain a small number second hand, it was clear we would have to manufacture our own. We found a Chinese supplier of “Loctal” sockets which are very close dimensionally to the B9G and so sufficient to populate the machine were ordered. It was found on delivery that the quality was poor and the holes for the valve spigot and pins were too tight. This was cured by laboriously dressing the holes in each socket with a dental burr. We had our local engineering company stamp out new rigid flanges for the holders to bring them to B9G dimensions and spot-welded these to the flimsy Chinese article. In use these holders have not been ideal. If the holes are overly opened out valve pins fail to make contact — or worse, if the chassis is inverted during fault finding, the valve

can fall out! If the holes are tight, valves can break on insertion, with a risk of injury to the engineer, and further dressing them in situ is awkward. It is a tedious task to replace a base on a busy chassis. With hindsight we should have gone for custom manufacture to the B9G specification.

Generally construction proceeded chassis-by-chassis. Testing beyond a basic static check of connectivity was often problematic as, being a serial machine, EDSAC circuits are rarely at rest — they are designed to respond to a repetitive stream of data passing through. It was often necessary to build test pattern generators so that the dynamic behaviour could be monitored on an oscilloscope. In some cases, where the length of a test cycle was measured in milliseconds or longer, more elaborate monitoring and recording devices were built. We of course had the advantage of being able to do so using modern microprocessors. How the EDSAC pioneers debugged the original machine without such aids is a marvel.

Chassis were then connected up into the major sub-systems of the machine and the test equipment evolved accordingly to provide system as well unit-test capabilities.

Once chassis had been made according to Burton's initial circuit rules, and some commissioning had started, it was possible to make a judgement about the correctness of those rules. One unnerving revelation was the breakthrough of substantial unwanted signals through the standard gates. Apparently the pioneers had accepted this but to modern eyes it was surprising. Associated with this characteristic it appeared that their standard flip-flop was temperamental when used in anger in a real logic environment. While we were struggling with this, a retired Mathematical Laboratory engineer donated a batch of original circuit diagrams he had retained when the original machine was scrapped. This treasure trove covered only a part of the machine, and furthermore described a machine significantly later and different from the design we were aiming for as a result of improvements made in the early 1950s. Nevertheless two examples of the monostable flip-flop circuit were found in the new diagrams and these enabled us to bring our component values into line with the original, with an improvement in the flip-flop circuit behaviour.

But, much more revealing, it was clear that the flip-flop circuit had been troublesome originally because the new diagrams were full of symmetrical bistables that were more in line with contemporary practice and had a much better resistance to interference. We were reluctant to make wholesale changes to the machine on grounds of authenticity, so have used the monostable flip-flop design where it has been found to work satisfactorily and only adopted the later bistable design in places where the monostable design has been unsatisfactory and there is photographic evidence to support doing so — in particular in the order decoding circuits of main control for which we have photographic evidence of adoption of bistables as early as 1950.

Fairly early on it was clear that we would need some sort of simulated delay line store to allow testing of the circuits for store access and internal registers such as the accumulator and sequence control tank (i.e., program counter). As the nickel delay line store was being researched and developed in parallel with construction of the electronic part of EDSAC, we decided to implement a digital delay line that could be connected in to any storage regeneration chassis to provide the required circulation for either a long (main store) or short (register) delay line. These were designed by Nigel Bennée, the volunteer responsible for the design and construction of the arithmetic unit, using a simple microcontroller to sample the output pulses, store them in a shift register and re-inject, correctly timed to the input side of the regeneration circuit.

Power Supplies

We knew from historical accounts that the original EDSAC was powered from a motor generator set. No pictures or description of this survive. We considered it would be better to build our own power supply from commercially available units and to add a high level of protection against short circuits or other sources of overload.

The power supply system delivers 250 V DC HT and 250 V AC to each rack. The DC is for the machine circuits, the AC for the valve heaters and for local power supplies such as the -6 V required by the digit pulse generators. The power is distributed via enclosed ducts and conduits with master fuses and power indicator lamps on each rack — this enables us to isolate individual racks and gives engineers a visual indication of whether the rack is powered or not. Two wall-mounted red lamps give an indication that the system as a whole is being supplied with power and a klaxon sounds when the power is switched on to the machine to warn those that might be working inside it. The power supply system shows the current and voltage for each row and allows each row to be switched in or out. Each chassis in a rack has local fuses. The power supply continuously monitors current consumption and if a significant variation is detected it trips out the entire machine (because of the design of EDSAC circuits there is significant earth leakage ruling out the use of RCCB devices). The final attribute of our power supply is the provision of a power-on sequence — the supply brings up the valve heater circuits in stages before switching in the HT, also in stages. Our hope is that this will minimize thermal shock to the machine's valves and increase their working lifetime — a similar approach has worked well for the Colossus reconstruction at the museum, whereas in wartime valves were never turned off.

Operator's Desk

From the photographs and EDSAC film it is apparent that the operator's desk was arranged differently at different times in the machine's life. After much discussion trying to date the different configurations, we settled on a layout that is probably later in time than our target period but which would make operating and demonstrating the machine easier in the museum environment.

The plan is for a single large, sturdy, laboratory-style table with a photo-electric paper tape reader, Creed teleprinter, operating panel and two display units — one for key EDSAC registers, the other for inspecting the store — one stacked upon the other. As an interim measure, a slow Creed mechanical paper tape reader will be used for input.

The teleprinter is a Creed device mechanically modified to print the EDSAC 5-hole code and to support the additional circuits needed by the EDSAC F-order, which reads back the status of the teleprinter solenoids to ensure the correct character has been set up.

Summary

Reconstructing EDSAC has been a major enterprise and at the time of writing a working machine is still some months away, although our confidence that we will achieve our goal is growing with every forward step we make. The task was not as easy as Wilkes professed — the documentary record did not provide a blueprint, and a great deal of forensic research and experimentation has gone into the reconstruction. The project has only been possible because of the generosity of the donors who provided the funding and of the hardworking team of volunteers who have given their time and skill to recreating a working replica of an iconic machine in the history of modern computing.

References

Anon. (May 1948). *The EDSAC*. University of Cambridge Archives, COMP A1.

Dodd, K. N. and Glennie, A. E. (July 1951). *An Introduction to the Use of High-Speed Automatic Digital Computing Machines*. Sevenoaks, Kent, England: ARE Fort Halstead, ARE Memo No. 7/51.

Kendrew, J. (1962). Myoglobin and the Structure of Proteins. In: *Nobel Lectures, Chemistry 1942-62*, Amsterdam: Elsevier, 1964, pp. 676-698.

Huxley, A. (1963). The Quantitative Analysis of Excitation and Conduction in Nerve. In: *Nobel Lectures, Physiology or Medicine 1963-1970* Amsterdam: Elsevier, 1972, pp. 52-69.

Linington, P. F. (2017). Constructing and Testing a Replica Store for the EDSAC. In: *Proc. Making IT Work*.

Ryle, M. (1974). Radio Telescopes of Large Resolving Power. In *Nobel Lectures, Physics 1971-80*, Singapore: World Scientific Publishing, 1992, pp. 187-203.

Wilkes, M. V. (1985). *Memoirs of a Computer Pioneer*. Cambridge, MA: MIT Press.

Wilkes, M. V., Wheeler, D. and Gill, S. (1951). *The Preparation of Programs for a Digital Computer*. Cambridge, MA: Addison-Wesley.

Wilkes, M. V. and Renwick, W. (July 1948). An Ultrasonic Memory Unit for the EDSAC. *Electronic Engineering* 20, pp. 208-213.

Maintainability and Sustainability Issues in Restored and Replicated Computer Systems

Christopher P. Burton

The Computer Conservation Society

Abstract

The author's experience with the restoration of the Pegasus and the Elliott 401 computers, together with the project to replicate the Manchester Baby computer, is drawn upon to review various issues associated with maintaining such systems in working order. Some consequential thoughts on sustainability, i.e. for how long will the objects be seen to be useful, will follow. The indispensable value of appropriate volunteers will be highlighted. The discussions are based on the premise that an operational artefact is vastly more useful than a static one.

Introduction

Within the wide field of computer conservation the author has some experience in the segment related to restoring early computer systems to working order and to making re- constructions of defunct early computers. In these types of activity the context typically is that of presenting the resulting working artefacts to the public in a museum, primarily for instruction and education, but also for entertainment and scholarly record. This context is totally different from that of the original machine. Then it was probably in an organisation where it had to earn its keep by offering a timely and cost-effective computing service. Consequently the topics of maintainability and of sustainability may be approached in rather different ways in the two contexts.

Throughout, the vital role of volunteers in the museums' context has to be emphasised. There seems to be two phases in the lifetime of the object. The first is the technical work to bring the object to an operational state. The second is also technical — to keep it in that state, but with the emphasis now on demonstration and explanation to viewers. Typically, retired engineers who have detailed knowledge of the machine, often first hand, will lead the technical work. They will have time to devote to the work and are fired with enthusiasm for the object and for the project. Once the machine is working, volunteering can be widened to include enthusiasts with less technical knowledge, but with skills in interfacing with the public and enthusiastic about the history and “meaning” of the object. Some of these people may be able to fit their volunteering with some other job. As time elapses, it is probably very difficult to replace first-phase volunteers because the original expertise and experience will have died out. But it is possible that new volunteers can be trained to continue the second-phase task, especially if seeing the working object at some earlier time has inspired them.

To set the scene, three projects with which the author has been associated will be described briefly. They are the restoration, demonstration and abandonment of the Ferranti Pegasus system; the partial restoration and abandonment of the Elliott/NRDC 401 computer; and the reconstruction and demonstration of the Manchester Small-Scale Experimental Machine, the “Baby”.

Pegasus Restoration

The Pegasus held by the Science Museum in London is serial number 25 of about 40 examples manufactured by the Ferranti Computer Department in the 1950s (Lavington, 2000).

Pegasus computers were designed to be easily maintainable in the era of potentially unreliable valve (vacuum tube) technology. Strategies such as plug-in circuit package technique, good environmental

control, conservative circuit design, and interference-resistant power supply systems all helped support the overall approach of preventative maintenance. This involved the normally-resident maintenance engineers taking possession of the system for an hour every day to carry out tests to identify individual circuit packages which were performing marginally, replace them with good spares and then fettle the guilty packages off-line after handing the system back to the customer. This took care of the commonest types of failure, namely the gradual drifting out of specification of common passive components and valves. Any sudden failures occurring in customer time had to be diagnosed and could then be repaired quickly by plugging-in replacement packages. The objective was to minimise customer down-time, and the fleet of Pegasus systems in the field easily achieved more than 98 per cent up-time. The consequence of these methods was that a Pegasus never wore out — it was always kept in good order by replacing low-level components or making small adjustments off-line.

The subject Pegasus had had a very chequered history, moving to many different organisations in diverse sites. Its final gainful use was for crystallographic work at University College, London (UCL). In 1984 UCL donated the machine to the Science Museum, which at that time was unable to house it. ICL stepped in and had it installed in a spare room in the computer laboratory at the West Gorton works. The author and a colleague maintained it there in their spare time and occasionally demonstrated it to visiting VIPs. After four years the Science Museum was able to accept it into temporary storage and in 1990 bring it to an area on the South Kensington site to allow the Computer Conservation Society (CCS) to revive it for potential display in a gallery. Each of these moves involved dismantling down to major units and subsequent re-assembly and re-commissioning into working order. Further moves to the Blythe Road overflow store in 1997 and then to a location in the Computing Gallery in 2001 took place, with CCS volunteers continuing to keep the machine operational and demonstrable. So 17 years after it left UCL it found its final home, having been maintained by volunteers throughout (Figure 1).



Figure 1. Pegasus in the Computing Gallery in 2010

In 2009, during a public demonstration, smoke emerged from the power supply unit and the machine immediately shut itself down (Burton, 2010). It is thought that an arc had occurred in the wiring of the direct current supplies, probably as a result of a conductive particle falling on to a whisker of wire in a terminal block. This ostensibly minimal incident caused consternation within the museum's management for a number of good reasons:

- Were volunteers appropriate people to provide live demonstrations of museum artefacts?
- Did those volunteers have appropriate training?
- In this incident was asbestos present in the object and was it released into the area?
- Had the object been correctly conserved and subject to periodic cleaning and conservation?
- Was the object electrically safe to operate.

The investigations and discussions rumbled on for five years. The tiny amount of asbestos present was not involved with the fire and was removed or sealed. The machine was given a thorough inspection and clean by conservators, the power system was inspected and tested by a consulting engineer and declared to be safe to operate, and volunteers were given training in the hazards of asbestos. It seemed to be accepted that no one was more appropriate to operate the machine than the volunteers who had experience of it when it was in customer service. But the damage had been done — the policy decision was that there should be no live operation of these early computers in the Science Museum. By this time the computer gallery itself was planned to be re-vamped as part of the Mathematics Gallery, and Pegasus would be dismantled and moved to the Wroughton long-term storage facility. A final switch-on and photo opportunity was arranged for 5 June 2014, which was the last we saw of this long-lived machine. What a sad end! In a hundred years' time, or even fifty, will there be anyone with the knowledge and competence to re-assemble the machine, including installing the refrigeration system? Will they be able to enjoy the smell of warm transformers, hear the authentic sound of the test program suite running correctly and witness the sight and sound of program tapes running in and punching results? I am not holding my breath!

Pegasus Maintenance

It was explained above that Pegasus was designed to make maintenance easy in a customer's installation, and we would expect those factors to carry forward into a museum context. Any differences, however, might include the relative availability of spare parts, the relative frequency of operation of the machine, and the relative experience of the people tending the machine. In the thirty years that Pegasus has been owned by the Science Museum it has been maintained by a small working party of two to six CCS volunteers, who typically would be retired engineers and former users with experience of the machine. Attempts to recruit and retain younger people in order to effect knowledge transfer were mostly unsuccessful because "the day job" had to take precedence. In the museum environment there might be one or two sessions per month, rarely four, often less. So the concept of daily preventative maintenance is out of the question. The purpose of the sessions is to keep the machine in working order so that it can be demonstrated and for the working party members to develop the demonstrations. There are some types of hardware faults that are affected by being in use, many others are just due to ageing. Consequently, the best maintenance strategy is to keep the machine clean, use it until there is a fault, and then fix that fault. That does mean that sometimes the machine would be working marginally because there would be insufficient time to do marginal testing, and in any case there is a lack of continuity from one session to the next. For many operating sessions there would be at least two engineers with good knowledge of the hardware design, so they could diagnose failures, replace the relevant plug-in package, or just make a repair. This sort of activity is of interest to the public anyway, so is not a detriment to

demonstrating. The Pegasus had accreted a very large stock of spare packages, valves and other parts in its earlier life, so this was helpful. For example a spare could replace a faulty package but there was no urgency to repair the fault which could be carried out at leisure in a later session. One exception to this scenario was a type of package that contains nickel delay lines. The tiny coils on these were susceptible to corrosion and failure due to age, and they were extremely difficult to repair. Informal techniques were developed to mitigate the problem but were never formally instituted.

The input-output peripherals required as much maintenance attention as did the central processor, and could be more difficult because they are subject to wear, albeit slight, and can get out of adjustment.

Elliott 401 Restoration

The Elliott 401 was built by Elliott Brothers with the sponsorship of the National Research Development Corporation. This was a 'one-off' project, the purpose being to demonstrate engineering features such as plug-in standardised packaged circuits and the use of acoustic delay line storage based on nickel wire rather than mercury tubes. The machine was shown working at an exhibition in 1953 (Lavington, 2011, pp. 157ff).

The machine was then installed at Cambridge University where it underwent modifications and enhancements, and then was transferred to the Rothamsted Agricultural Research Establishment for statistical work concerned with horticultural and agricultural plant trials. It was dismantled into separate cabinets in 1965 and donated to the Science Museum, London, where it was put into storage.

Enlightened curatorial management in 1992 permitted the machine to be the subject of a restoration project of the Computer Conservation Society. The units were brought out of store to the CCS area in the museum at South Kensington where the museum's conservators took over and treated the parts where needed. As soon as the magnetic storage drum had been conserved, the project team carefully activated it, by running the drum at half speed, and carefully setting the heads to a very wide spacing, where signals were just about resolvable. Over a period of a few months the team was able to capture the data off all the tracks into a personal computer. Then, using tailor-made programs to analyse the analogue waveforms, a preliminary decode of the Initial Orders was achieved.

Pressure to vacate that assembly area in 1994 resulted in the machine units being moved to the museum's Blythe Road store (Figure 2). It took nearly three years before a dedicated computer room could be prepared and assembly of the machine could now be completed and the team could make a start on re-commissioning. At this point the team comprised four engineers, one of whom had used the machine when it was at Cambridge University. Very complete design and production drawings were available, though the critical logic diagrams were found to be incorrect or incomplete and a considerable time was taken tracing the wiring to produce an authoritative set. But by mid-1999 progress stalled because some team members, including the team leader, found themselves diverted to other projects and activities.

The working party re-formed in 2004 under new leadership and weekly sessions. It is fair to say that this point was the start of good progress with commissioning — the previous periods being concerned with establishing the infrastructure, such as the power control system, inter-cabinet wiring and studying and familiarising with the cabling conventions.



Figure 2. Volunteers working on the Elliott 401 at the Blythe Road store

Clock timing signals for the machine are derived from a special track on the magnetic drum, so effort was concentrated on this area first. This took many months to make progress, given that the team met once a week at most, and often less frequently. In parallel, effort had to be deployed to get the operator's control panels and display units to work, in the face of individual more or less arcane component failures which had to be found and repaired. The magnetic drum system was particularly hard to make work, partly because none of the volunteers in the team had the necessary tacit understanding of the delicate procedure for setting the heads and the amplifiers, despite brief explanations in the documentation. An overriding consideration was the vital need to not damage the drum or the recording heads by a hasty or ill-considered action. By the end of the project we had not got the drum system to work properly, though we were tantalisingly close to getting the various track signals to remain in phase in a coherent way.

The project was abruptly ended after the 2009 incident on the Pegasus on display in the museum. The museum authorities decided that live demonstration of these early computers posed a possible risk for the public and was to be terminated. Consequently there was no point in continuing to restore the Elliott 401 as there was no potential for it to be demonstrated working. It could be displayed as a static exhibit, but restoration was irrelevant in such a case. Thus ended a 17-year project involving 289 sessions by the working party.

Elliott 401 Maintenance

Very little can be said about maintenance of the Elliott 401 while it was in the hands of the working party volunteers, because it never achieved full restoration to working order. While it was being commissioned, failures which held up progress were fixed as required, a process which can be regarded just as a sub-commissioning activity. For example, as part of commissioning the magnetic drum system, it was necessary to select which drum tracks to read. This required control by the hand-switches on the Engineer's Control Console. The relevant circuit was not working, and it was discovered that it relied on a specialised nickel delay line with spaced read coils along the line. A diode associated with one of these was faulty. This is a very unusual technology and many weeks were spent in working out how it worked and how to fix it.

A general problem throughout all the commissioning activity was unreliability of connections via Painton multiway connectors. The log book regularly records hours spent tracing a missing signal from source through cables and connectors to a destination, only to find that the signal had re-appeared. The reason for this behaviour was not always traceable, and it had to be assumed that dirty connector contacts or aged soldering was the culprit. A vivid impression emerges from the log book record of activities, of the extraordinary patience and determination of the volunteer team. Session after session is devoted to curing some problem only to find it repeats again and a different line of attack has to be followed.

Manchester “Baby” Reconstruction

A project to reconstruct the Manchester University Small-Scale Experimental Machine (SSEM or “The Baby”) of 1948 was approved by the CCS in 1994, and work started in 1995 (Burton, 2005). The “Baby” was the first computer in the world to successfully run a stored program. The goal of the project was to construct a replica of the original and run that first program on 21 June 1998, exactly 50 years after the original event. The University of Manchester offered facilities to build the system there. The Museum of Science and Industry in Manchester agreed to accept the completed machine and display it to the public. The computer systems company ICL offered use of workshop services and purchase of materials as exclusive sponsors.

A team of six volunteer engineers familiar with valve technology was recruited who started work in late 1995 to study surviving documents and photographs. Although there were no engineering drawings — there never had been any — it was possible to deduce with a good degree of confidence what the machine had been like in 1948. The original was very much experimental — a sort of laboratory “lash-up”, whose purpose was to exhaustively test the cathode ray tube-based storage system developed by Williams, Kilburn and Tootill. By studying the photographs dated as late as November 1948, the volunteer team was able to replicate the authentic appearance of the original. We understood the type of electronic components used, and the valves, and astonishingly the latter, of original contemporary manufacture, were still available in dealers’ warehouses. During 1996 and 1997 detailed drawings were created, parts were made in workshops and volunteers’ homes, and



Figure 3. “Baby” on display in Museum of Science and Industry

everything brought to the university to be assembled into the complete machine. Some of the original pioneers assisted with advice and recollections. By the end of 1997, programs could run on the “Baby”, and in 1998 it was moved from the university to a bespoke gallery at the museum. Here it was given its final touches ready for Tom Kilburn and Geoff Tootill to re-run the first program on 21 June 1998, exactly as they had done 50 years earlier. After that celebration the machine was handed over into the care of the museum.

Since 1998, the machine has been moved five times, either into temporary storage or to increasingly more prominent areas in the museum. Currently it is located in the entrance area (Figure 3) as an “Icon of Manchester” (Burton, 2013). There are now about 20 volunteers available who demonstrate the machine on four days per week, typically in teams of three. Some of the volunteers are qualified to maintain the machine in working order. Visitor numbers have increased over the years and typically the volunteers interact with 30 to 80 people on a demonstration day.

“Baby” Maintenance

The computer is small — there are only 563 valves of which more than half are simple miniature diodes. On an annual basis a very small number, say 3 to 6, have to be replaced. As many failures appear to be due to poor or intermittent connections or to failing small passive components such as capacitors and resistors.

However, there is a class of faults which are the most troublesome and time-consuming. These are where diagnosis is difficult, or where the fault “goes away” while investigating, only to re-appear next session or a few sessions later. Many of these faults must be due to some signal which is marginal with respect to amplitude or timing and which expresses as a fault depending on the sensitivity of some other circuit. Others are probably due to intermittent connections in the wiring (bad soldered joints) or in connectors, particularly in valve holders. Indeed the B9G holders which are the most frequently-used for pentode valves were always notorious in the 1940s for susceptibility to oxidation and need for re-seating. Evidence for this source of problems comes, for example, from operators who comment that “the fault went away when I stamped on the floor”. Obviously, if the signal can be identified then the fault can be fixed, but identification is the problem. Of our strong team of volunteers, only a few are qualified or experienced enough to pursue this sort of problem. No doubt this is all authentic behaviour of the computer as it was in 1948. After all, it was a lab experiment, built in a hurry. Agreed, we did have better tools and materials in the 1990s, but even so expertise of the construction team must have been variable, and we also had time pressures. Ideally, a structured examination of all connections, connectors, terminals and so on could be carried out, but that would be very time-consuming. The time that volunteers can make available to the museum is obviously limited, and their priority is to demonstrate the machine. However, such an enterprise should not be ruled out in the future.

One particularly embarrassing issue is the chronic state of the main storage cathode ray tube subsystem. For many years this has been more often unreliable than good. This is particularly annoying because the purpose of the original SSEM was to demonstrate how good the CRT store was in a computer! In the construction years of the replica up to the time it was handed over to the museum, the store had worked quite well giving long program run-times of an hour or more. This has to be compared with the behaviour of the original machine in 1948 and later, where it is known that corruption of stored data could occur several times per day. The problem was worked-on vigorously by the pioneers who introduced specialised clean-build CRTs, improved techniques like defocus/focus instead of dash/dot and general good engineering so that the eventual Ferranti Mark I was adequately reliable. But our replica is of the SSEM without those improvements. In our situation, the problem seems to be that the CRTs available to us are extremely hard to set up with the correct signal parameters, and having done that, to make the system retain that set-up. The CRTs were all manufactured in the 1940s and early 1950s, so we suspect that they have deteriorated

through age or misuse. Likely issues might be loss of emission from the cathode, loss of vacuum, loose internal particles or electron gun misalignment. We have tried a slightly different CRT type with much improved focus compared with the standard CV1097, and the few which are available to us do seem to behave better. The store CRT has a 32 x 32 pixel raster which covers much of the face of the tube, so its duty is much harder than the Control and Accumulator CRTs, which only have to store one or two words, and which are consistently reliable. At the moment there is only one volunteer with the know-how to set up the store CRT, so that is an unsatisfactory omen for the future.

At the planning stage for the replica we purchased sufficient spare components assuming we needed to cater for a 25-year life for the machine. The components of concern were the valves and the cathode ray tubes. After a few years it was apparent that 25 years was too modest, and that dealers were winding-up their businesses. A decision was made to increase our spares holding, which seems adequate for another 25 years, though CRTs are scarce. Now in 2017 the machine has already achieved a 19-year life.

Sustainability

The implied objective is to keep these old machines operational “forever” so they can be demonstrated to the curious public and to scholars, in order to provide an experience of what it was like when they were first created.

“Forever” is pretty strong. The reality is that the practitioners are volunteers, doing the work because they want to, not because it earns them a living. If the plan is to restore the object to working order but then to scrap it, there will be very little motivation. There has to be hope that the job is worthwhile so they will be recognised for having done it. The vision might be that the object survives a generation after the volunteer’s lifetime, and 25 to 50 years is acceptable. But it is surprising how 25 years seems to pass so quickly!

Presumably there will be a willing audience who want to see the object working like it did originally. Perhaps for nostalgia, because they used to work with that object, or for historical interest to understand technology development, or perhaps for research for more formal historical purposes. In the extreme case, the willing audience might just be the volunteers themselves.

There are several prerequisites to achieving sustainability:

1. A team of enthusiastic volunteers.
2. Adequate expertise of at least some of the volunteers. Currently many volunteers already have the background technical expertise because they are now at retirement age and this is what they were trained for. It is not clear that new younger volunteers will have the appropriate mind-set to be able to diagnose and fix problems.
3. Adequate documentation for the object. This may well go beyond the original manufacturers’ documents because of the need for explanatory information about the technology and technical background in place originally, and which is not familiar to new generations of volunteers.
4. An appropriate supply of spare parts. Thermionic valves immediately come to mind, and they are critical because they are no longer manufactured nor likely to be except for certain very special types. Cathode ray tubes are especially vulnerable and irreplaceable. For ordinary valves we have the prospect of putting a semiconductor circuit inside an appropriate envelope but that invites questions about authenticity.
5. Commitment by the host institution or owner to continue to display the object.

We can now examine whether these criteria apply to the three cases described earlier. For Pegasus it is fair to say that 1 to 4 had been present up to abandonment. It is likely that 2 would only have lasted for another five years as there was no expertise replacement policy in place. In the event, the host institution stopped the project first.

For the Elliott 401, the team of volunteers was very small, putting a heavy burden on the few. Prerequisite 2 above was just about met but with no backup capability. Item 3 was adequate after some years improving it, and gaining familiarity. Item 4 was also satisfactory, though we had no real need for spares at the commissioning stage. Again the host institution stopped the project, but it is not clear that it could have got much further without an influx of more volunteers and expertise. This had just about started in the few months before abandonment.

The “Baby” replica seems to be in a better situation. The host institution has warmly adopted the project and actively recruits volunteers. So currently there are 20 volunteers operating in teams on at least four days per week. There are a few with good expertise, and interestingly some of the younger recruits have entered that category. Without in any way denigrating their ability, it is fair to say that the “Baby” is an exceptionally simple machine compared with the others, so they may be equipping themselves for key roles in other restoration projects. The spares situation is good on this project, with the exception of cathode ray tubes. There is fall-back to semiconductor-based storage but that is unpalatable bearing in mind that this was the machine that demonstrated the original CRT store. There has been no hint from the host institution other than that it regards the replica as a key item in its collection and an important attractor for visitors.

Conclusions

The expertise to maintain these old computers seems to be available up to around half a century after the objects first existed. After that time it becomes harder to retain enthusiasm because original associations are lost. As a result, keeping the machine working takes longer, particularly in a part-time volunteer regime. The “Baby” is exceptional because of its simplicity, which means the relevant expertise can be acquired more easily.

The three cases discussed in this paper have been hosted by museums. It is well-known that they have to make difficult choices as to which objects they wish to display from the very many which lie in store. Priorities evolve according to changing objectives and resources. Long-term sustainability of operational old computer systems can therefore never be guaranteed.

References

- Lavington, S. (2000). *The Pegasus Story — A history of a Vintage British Computer*. London: Science Museum.
- Burton, C. (2010). “The Pegasus Incident and its Aftermath.” *Computer Conservation Society Resurrection*, No. 53, pp. 15-22.
- Lavington, S. (2011). *Moving Targets — Elliott-Automation and the Dawn of the Computer Age in Britain, 1947-67*. London: Springer.
- Burton, C.P. (2005). “Replicating the Manchester Baby: Motives, Methods and Messages from the Past.” *IEEE Annals of the History of Computing*, 27(3), pp. 23-33.
- Burton, C.P. (2013). “The Teenage “Baby” on Show.” In: *IFIP WG 9.7 International Conference HC 2013 - Making the History of Computing Relevant*. London: Springer, pp. 274-284.

Curatorial Lessons from Other Operational Preservationists: Towards a methodology for computer conservation

John K. Chilvers

Conservation Architect and Heritage Consultant

Abstract

There is a great body of good practice for operational preservation in the fields of transport and industry, and the principles of curatorial ethics are well established in historic building preservation and elsewhere.

If we regard the experience of an environment or process in action as the “artefact” being conserved, and the skills required as another, we can propose a typology of both material and non-material artefacts, with parallel codes of ethics and curatorship.

By reviewing the traditional museum and its practices, alongside the proposed typology, we learn lessons of good practice for operational conservation of computing.

Introduction

Operational preservation and demonstration of industrial and technical equipment is now a well-established pursuit in many fields — from steam railways, via diverse manufacturing processes, to the newer fields of electrotechnics and computing. There have been extraordinary achievements, typically enthusiast-led, and substantial bodies of good working practice have emerged.

However, all this is relatively new. Some continue to regard much of this with suspicion, critical of the compromises involved against the traditional doctrines of conservation for physical artefacts. Nevertheless, it is important that practitioners of operational preservation also recognise that there are conflicts inherent within any such project.

The conflict, it seems, is based on fundamental confusions over what it is that is being “conserved”. It is therefore necessary that a new clarity of thought be found, which can lead to a confidence in curatorial ethics and practice that measures up to established standards and furthers mutual understanding.

The motivation for many operational preservationists is the belief that there is more to an artefact whose nature is a working object than merely its material form. To understand and appreciate it fully, it needs to be brought back to life. An exhibit of non-working computers feels like little more than an exhibit of box designs — clearly important aspects are missing! So the operational preservationist wishes to enable an “audience” to *experience* the life of the working artefact — a non-material aspect of it.

The author therefore proposes a typology of artefact that encompasses both *non*-material as well as material artefacts. This will enable a fuller understanding of the relationships between the preservation of physical objects and the operation of dynamic systems, and enable the non-material to be clearly treated with curatorial ethics of the same standing as those for material objects.

This typology will also help curatorial compromises to be made across classes of object so that a change to the physical material can be made in order that a non-material aspect is better preserved. The typology will enable these compromises to be regarded as ethical to the highest standards, and equivalent to well-established practices in approaching repair, restoration and protection that are solely within the material class.

Table 1. Typology of material and non-material artefacts

Class I:	Traditional physical museum object
Class II:	Dynamic experience of environment or process in action in space and time
Class III:	The skills and practice necessary to demonstrate and maintain the above
Class IV:	Meaning, values and motivation — creativity

The typology therefore allows established principles of curatorial ethics to be applied to the non-material classes of artefact. This can encourage reticent curators to embrace the importance of operational preservation, and give confidence to the operational conservationist that valuable and ethical work is being done. In particular, this paper is intended to show some foundations upon which the community of practitioners of computing conservation can now build with confidence their own detailed codes of ethics and practice and meet the proper scrutiny of colleagues both present and future. The proposed typology of material and non-material artefacts for conservation is shown in Table 1.

Class I is the ordinary class of artefact with which museums have been familiar — the item of material culture that may be displayed. In the context of the conservation of computing, this relates not only to the computers themselves, but also peripherals, ephemera, documentation, physical material on which data and programs reside, etc. There are well-established principles of curatorial ethics and conservation practice for these. The intention of the typology is that these principles should also apply to the other non-material classes in order to produce high quality ethics for operational conservation.

The evolution of the traditional museum and its practices

The evolution of the traditional museum and its practices over the years has led to the principles listed in Table 2. To some, this might seem like stating the obvious. However, it is evident from the sort of experiences that have been discussed at the conference and elsewhere, that it will be a help to state these principles explicitly. Before considering the principles in more detail we take a quick survey of this evolution.

Any respectable Gentleman of the Renaissance would embark on the Grand Tour to the major classical sites and sights of antiquity and would acquire actual ancient artefacts for their private collections. Some individuals granted some public access. Thus the concept of the museum was born — the British Museum dates from the 18th century. Later in the 18th century came a taste for the odd classical temple to enhance a view, or for a nice mediaeval ruin. If there was no ruin, then one would be faked. Significantly, this led to a new interest in Gothic architecture.

So, the 19th century gave rise to a moralistic antiquarianism. Ancient cathedrals were clearly worthy of restoration. But since they were so sure of their understanding of what the ideal mediaeval cathedral *should* have been, they often went on to remodel whole sections of the building, sometimes so thoroughly that hardly anything of the original was actually left to be seen. A halt had to be called. And thus was formed the Society for the Protection of Ancient Buildings (SPAB). The Society continues as guardian and propagator of systematic curatorial methodology as well as of the

Table 2. Conservation principles

-
- 1: Conserve as found
 - 2: Minimum intervention
 - 3: Like for like repair (in materials and manner)
 - 4: Reversibility
 - 5: Clarity and honesty
 - 6: Be sympathetic
 - 7: Documentation of decisions and actions
-

actual techniques of building conservation. Its work has highlighted many of the curatorial dilemmas, and its ways of thinking will help us in computer conservation.

The year 1851 saw the Great Exhibition in the Crystal Palace, showing the latest technological achievements as well as exotica of empire. It made such a financial profit that it enabled the foundation of the great institutions of South Kensington, including the Science Museum. The rest of the century saw the rise of great public collections, mass education and democracy. Academically it saw the rise of many new disciplines. Legally there were the first measures for the Scheduling of Ancient Monuments and the protection of cultural material. The rise of popular tourism, led by the railways, initiated the heritage industry of places to visit.

Coming into the 20th century, professional archivists, the museum disciplines (eg. art restoration) and codes are all well established and represented by the Museums Association, amongst others. We shall return to that, and the enthusiast-led rebellion against it, shortly. But now we can find much to help from the world of historic buildings conservation.

Some principles of curatorial ethics from historic buildings conservation

So, established systems of curatorial thought already exist and are given in Table 2. In the historic buildings world, they are expressed through agencies such as SPAB, English Heritage, and in a history of manifestos, declarations and a complete British Standard — B.S. 7913 (2013) — devoted to their understanding and best practice. And there are the heritage designations made by UNESCO.

Clearly, the first two principles, “conserve as found’ and “minimum intervention” are a variant of “if it ain't broke, don't fix it”. “Fixing” can be a severe temptation out of enthusiasm or misguided perfectionism. There is a strong reminder that the actual intention is conservation as completely as possible of what we have got, and to do anything else is to lose material evidence of importance. The context and arrangement in which a thing is found, as with buried archaeology, also embody information that can be of great value, and so must itself be recorded. When preventive maintenance routines in original service would have called for the destructive removal and replacement with new parts, the alternatives of inspection and non-destructive testing should now be pursued. But occasionally, improvement on an original design fault may have to be made (eg. an inevitably leaky roof detail) to stop the thing from self-destructing repeatedly.

Documentation is the key to many of the points, and is noted as the seventh principle. It links especially to the fourth principle “reversibility” and to the fifth “honesty” — modestly accepting that

one's successors may revisit one's work in the future. It also should clarify what has and has not been done and where the parts left over went to. It is also essential in providing for display interpretation and reliable academic understanding in subsequent research. All the principles work together in an integrated way.

However there can be conflict between some of these strictures, and compromises or creative alternatives have to be found and, of course, recorded. For example, the third principle a “like for like repair”, if too subtle and perfect, may leave no archaeological trace and thereby be misleading. But to be too glaringly honest in showing off the changes made could be seen as *un-sympathetic* and detracting from other qualities. Complete reversibility is, in the ultimate, an ideal and may, in a particular situation be genuinely impractical; but if the “intervention is minimised”, it may be judged acceptable, again if properly argued and recorded. This is why we have the principle six, “be sympathetic”.

If one is concerned with a common run of machines and has more than one example of a particular model, then it would make sense to treat different individuals in different ways. One example may be treated according to the strict traditional museum practices (caricatured as the “glass case” or “white gloves” article). One could present other(s) in different conditions of history, and hold on to yet others to cannibalise for spare parts. This matches the typical treatment of a row of terraced cottages in an open-air museum.

If, however, one's purpose is to conserve, say, Durham Cathedral, then there is only one of it and for the building to continue to function *operationally*, the central task is to address both sides of the conflict simultaneously. There is the archaeological substance of the original to be conserved — both physically and in the continuity of the historical stories and technical achievement. At the same time it has to function, without falling down on the occupants or leaking through the roof. The services and facilities have to operate competently, without giving electric shocks or spreading disease. All this whilst giving aesthetic delight for its artistic and cultural qualities and the sheer *power* of the experience of being there.

All these same considerations are fundamental to computer conservation, including the equivalent *power* of witnessing such extraordinary equipment working. In other words, we are curators of both non-material and material heritage. The chief conflict arises where significant parts of the original physical fabric are reworked or renewed in order to enable the artefact to become operational once again. This is where the traditional “glass case” museum people become very uncomfortable — and rightly so. The role of guardian-in-perpetuity of the material artefact is being compromised. However, it is the proposition of this paper that, by thinking systematically about the distinct requirements of the *non-material* as well as the material artefacts, measured judgements can be made that can stand scrutiny.

The emergence of the dissenting world of operational preservation

By the middle of the 20th century, there was a rapid change in industrialised society. Much of the older plant and infrastructure that had characterised the achievements of the preceding hundred years was starting to be superseded and brought out of use. For some pioneering minds of the time, to present a steam locomotive, for example, as if in a “glass case” and to accompany it with mere photographs of the landscape through which it passed, was to miss the point. Something fundamental was being disregarded by the mainstream museum, architectural and engineering professions; but people were already preserving, restoring and running veteran and vintage cars.

When the romance of steam combined with these instincts and the first operational railway preservation projects were initiated, first in Wales and then elsewhere, something new had stirred. It had to be initiated entirely outside the existing professions, almost entirely by volunteers and enthusiasts — though a few of them may have come from such professions as dissenters outside of

their official positions. There was as yet no corpus of expertise, no co-ordinated representation or methodology, and many influential detractors. Others began to address all manner of unlikely projects, not only in transport infrastructures, but whole industrial installations *in situ* as well as mobile and portable equipment. The whole scene was beginning to emerge with its own experimental approach and a new dimension of excitement of experience had been born. Perhaps this is the stage at which computer conservation is finding itself at the present time.

In a mischievous challenge to the establishment, by coining a term from the *avant garde* art scene of the time, those who had successfully brought back into full operation a steam beam pumping engine for water supply at Ryhope in Sunderland, pronounced the experience as one of “kinetic art”.¹ It was not just the engineering achievement or the historical significance — it was the life force of this immense thing whirling and heaving about you that was the central reality of the case. It was instinctively recognised that this was an artefact of a different philosophical kind that we were setting about conserving — the actuality of the dynamic process in action and in its own true context, at full speed. We need a new level of methodology to be able to handle such exhibits.

The proposed typology of material and non-material artefacts conserved

So we have arrived at a need to recognise the *non*-material aspects of preservation with equal status to that of the material artefact, and the proposal to treat these with the same curatorial approach. Referring back to Table 1, it is proposed to complement the material artefacts with three broad classes of non-material artefact, though there might be subdivisions discernible within each.

Each class is now named and outlined in turn, though time and space constraints mean detailed discussion will need to be laid aside to future opportunities for companion papers, by the current author or others.

Class I Artefact: Traditional physical museum object or archaeological remains

This is the ordinary class of artefact for which there are well-established principles of curatorial ethics and conservation practice, as outlined above.

Class II Artefact: Dynamic experience of environment or process in action in space and time

This can now be distinguished as something different and of its own right. It is embodied in the world of material objects, in that it cannot exist without them, but it inhabits a different sphere, so to speak. It is the “kinetic art”, the actual movement and life of the thing in action. The reality of the *personal experience* is the actual artefact to be conserved — the defining matter. Typically, it involves the passing through the space of a building or environment or exposure to a functioning process in action. It has to be identified for what it is, and thought about curatorially, so that the same rigour is applied.

We can consider allowing appropriate loss or displacement of original physical material in order to enable the conservation of the experience in reality, or as near as may be. That experience as presented must however observe the seven curatorial principles, for example:

Is it being demonstrated in the original way and not reworked as a caricature or travesty of itself?

Are the amendments properly evident and free of concealment, whilst remaining sympathetic to the original?

¹ The author was an industrial archaeologist in the Joint Conservation Team of the Tyne and Wear Metropolitan County Council during the period 1974-1984. During this time he benefitted from collaboration with Dr Stafford M. Linsley, who was Staff Tutor in Industrial Archaeology at the University of Newcastle upon Tyne, and especially attributes the source of this discussion of “Kinetic Art” to him.

Are they executed, so far as is possible, according to contemporaneous and plausible techniques?

Have measures been taken to ensure that there has been no loss of knowledge or other evidence in the recommissioning?

Is there any actual value in showing the process in action, or would some other approach, like simulation or a replica, yield more effective results?

All these are pertinent to operational conservation of computing no less than any steam railway or coal mine.

Class III Artefact: The skills and practice necessary to demonstrate and maintain the above

This section warrants a companion paper on its own as it needs to address the potential conflict of trying to preserve practices that are historically accurate in competition with modern methods that would speedily get the physical artefact working again.

A distinctive lesson from the 1970s and 1980s was the recognition that to operate a unique piece of apparatus would require a clear proactive approach to the passing on of the very particular skills and traditions that had come down over the centuries along with it. To rely on a few retired workers coming back voluntarily would not suffice. Some had already learned the hard way — when their last man died without succession — that a film or recording would not do the job. In some areas, such as steam locos and vintage cars, a large enough corpus of people has become established on a national scale. But unique equipment with non-transferable skills requires different measures. These skills can now be identified as a further distinct class of artefact.

It should be noted that live skills can only be maintained with continued personal use and in the teaching and apprenticeship of one person to another. But they can still function significantly at various levels, ranging from the traditional lifelong Master, through to the occasional hobbyist who can apply more generalised skills from another domain. Occasional practitioners can perform certain more demanding and safety-critical tasks satisfactorily to a limited extent, with special preparation. However, the frequent or continuous performance of such functions must be devolved to those with fully practised instincts that can handle any eventuality and the distractions of multi-tasking.

These skills must be conserved with the same degree of curatorial rigour as the other classes of artefact, using the same ethical principles. And again, the occasional conflict of interest between the classes has to be systematically analysed and judged. For example, some tasks in their unaltered form might be ruled to be just too dangerous — for example, the frequent casual use of unguarded high-voltage switchgear. Extra guards may be added, compromising the physical items. But a little lateral thinking may yield “soft” solutions of another kind — like instituting an interlock key procedure and training of certificated personnel with authority. However, loss or corruption of skills may leave the operational system disrupted and the experience stripped of much of its meaning and value, just as much as does the inappropriate alteration of the Class I artefact.

Equally importantly, we must ask of the conservation of skills — though they be non-material — the carefully judged balancing of: conserve as found, minimum intervention, like for like repairs, reversibility, clarity and honesty, and being sympathetic. Propagation of skills in either direction along this scale of professionalism is now required:

The historical skill precisely preserved;

The historical skill amended to comply with modern safety practice standards and civil expectations;

Modified forms of the original skills, developed progressively to sustain the competent operation of changed, updated or new hardware or environments;

The intellectual re-invention of lost skills, given new life in practice by gifted artisans;

The new skills of the conservers that are required to treat the new materials and classes of artefact, as well as the various existing established museum conservators' skills.

Other aspects of Class III are just summarised here:

Lost skills can be recovered to some degree, but these are conjectural reconstructions.

The onward communication of skills also requires realistic training facilities.

A systematic scheme of learning on the job may help the transfer of skills, and the status, given through formal testing and qualification awards, can be both a good motivator and form a part of the conserved experience.

A support and maintenance infrastructure will be required in sustaining an archaic operational process.

A further sub-set of Class III is systems knowledge. This is that whole body of managerial and operational knowledge, procedural rules and the like, that are both explicit and implicit in the way the whole system runs.

In operational computing conservation, it is also pertinent here to raise the special subject of software. It is suggested that the software itself, as a non-material artefact, is probably best thought of as a dynamic communication problem, rather than as a static knowledge problem.

Design data as manipulated by computer systems are also non-material artefacts — these and all digitised archives are entirely dependent upon a functioning system and the operator's skills for all future access.

Class IV Artefact: Meaning, values and motivation - creativity

These non-material aspects influence the curation of the other classes. The importance of the relative judgement involved is often accepted more readily for Class I than those in Classes II and III.

Historical association is the most obvious example of Class IV. Whether this railway locomotive pulled the Queen Mother's train or that machine was the one that calculated the moon-landing, usually does not really make one molecule of difference in the material object. Nevertheless, such associations bring out strong motivations to honour them, influencing decisions regarding location, conservation regimes and access. Sometimes arguments can become politicised — as when issues of pride of identity are involved. And when objects with religious status are concerned, other issues arise, ranging from the role of religious practices in the handling, to occasional deeper matters of morality. Sometimes people just speak more simply of conserving the “memory”.

As a motivation for museums, the educational ideal is usually invoked. However, motivations can embrace meanings and values from the more philosophical to old-fashioned nostalgia or simple good fun. But there are also questions of partisan interests and propaganda. Also one must be clear about what is the difference between a true museum and other forms of interpretation centre or fun park. For example, where is TNMoC on a scale that goes from Alton Towers to Calder Hall via Westminster Abbey, when, for many, they are all just part of that heritage industry of places to visit?

Table 3. Selected examples of non-material conservation already extant in museums traditionally associated with material objects

Art Gallery	Kinetic Art Installations (see main text); Performance art
Natural History	Living exhibits; Simulation video shows
Archives Office	Audio recordings of aural history, language and folklore
Archaeology	“Experimental Archaeology” i.e. reconstruction of lost skills (e.g. flint knapping, mediaeval construction technology); Lifestyle experiences
Science and Industry	Steam locomotives on the public rail network; Restoration “loans” to enthusiast-led societies; Lighting effects and pyrotechnic simulations; Film, photography and television in action
Textile and Crafts	Working of hand machines; Craft skills activities and training schemes
Applied Art (ie. antiques etc.)	Working clocks, watches and automata; Baroque musical instruments and performance; Educational handling collections that are not part of the “glass case” permanent collection
Open Air*	Working windmills, watermills and various machinery; Training in traditional building skills and architectural restoration; “Realistic” experiences artificially created (cf. Disneyland, amusement parks, etc.)
Conservation/Restoration Workshops	Classical artist and craftsman skills; Development of new techniques; Scientific research and analysis

** Note how these are actually founded in part on the paradox of destructive dismantling and reconstruction involving replacement materials and new workmanship.*

More widely, creativity means the function of being the guardians of the feedstock of ideas and opportunities in science and technology. By conserving not only the material remains, but also the dynamic experiences and the range of skills, it is sought to further:

- Inspiration for the inventions of the future;
- Imaginative reuse and stimulating reinterpretation of what remains;
- Social continuity and understanding;
- Awareness of appropriate and alternative technologies.

Conclusion

Think of the proverbial “grandfather's hammer”, which has had three new heads and five new shafts, but is still my grandfather's hammer. By now, it contains little of the original substance, but when I use it, is it because it does the old job nicely, because it keeps me practised in the skills I learnt with it, or in celebration of the attitudes he inculcated in me? Or, should I hang it on a bracket till the day when the forensics need to prove the family line or speculate on the traces of toxins in his favourite engine oil?

Although some museums have withdrawn from some operational preservation (eg. the demise of Pegasus at the Science Museum), there are many museum departments that have broken these boundaries into the non-material, establishing new forms of practice that are accepted as reputable. A selected list of examples is shown in Table 3. Nor is it being suggested that their existing corpus of conservation skills can be disregarded. To the contrary, not only are they themselves precious artefacts to be treasured in Class III, but they must be extended and added to with newer technologies and new subject matter. Perhaps most notable in this are the counter-intuitive areas of “conservation by use” (e.g. pre-plastic era electric traction motors, electrolytic capacitors, and software). And those already curating retrospective exhibitions of kinetic art or technology-dependent installations may actually find themselves in need of the new expertise of our own speciality.

For computer conservationists, it might seem to be a lonely road conceptually, but it is a road that has been travelled by others. Church and cinema organ builders and railway-signalling engineers have been wrestling with logic machines for a very long time, and mechanical music devices that have handled digital data streams since before the gramophone are playing new arrangements of Beatles songs. They each have their own circles of enthusiasts working avidly on restoration and continued use. There are many related conservationists who could assist with everything from the restoration of textiles, ceramics, wood, slate and early plastics, through to those in the tramway and trolleybus fraternities who have experience of putting all these same materials back under high voltage in use, and doing so in a stringent public safety regime.

Operational preservation is like a sort of “rare breeds gene bank” of technical skills and ideas, urgent to cross-fertilise and seed new inventions into the future, whilst combined with all the expressive dimensions of a performance art. Whether one is *operating* Westminster Abbey, the Flying Scotsman, a complete but defunct UNESCO-designated coal and steel works, or an Olivetti P101 programmable calculating machine, it is suggested that the curatorial issues are in principle similar. So there are all sorts of relevant precedents in related disciplines that can help illuminate and inspire those who are in the vanguard of operational computing conservation.

This paper has proposed a typology of four classes of artefact, both non-material and material, and given a set of principles with which to examine the approach to each of these classes. They come from the experiences, observations and reflections of a longstanding practitioner in various fields of operational preservation and their associated activities. It is hoped they bring some greater clarity of thought and argument that both furthers mutual understanding and suggests a more holistic way of

considering the judgements that are inevitably being made in whatever is done. They are also offered as a contribution to the groundwork for establishing the formal curatorial codes and practices of the Computer Conservation Society, and to an informed confidence in its own intentions and achievements.

Acknowledgement

The author is most grateful to Peter K. Chilvers for editing the post-conference version of this paper.

Reference

B. S. 7913. (1998, rev. 2013). *Guide to the Conservation of Historic Buildings*. London: British Standards Institution.

Restoring and Demonstrating 1960s Vintage Computers at the Computer History Museum

Robert Garner

IBM 1401 Restoration Project Lead, Computer History Museum

Abstract:

Beginning in 2003, the Computer History Museum in Mountain View, California, sponsored long-term projects to fully restore and demonstrate three large 1960s vintage computers: a DEC PDP-1 and two magnetic-tape IBM 1401 systems. By 2005, a small volunteer team of engineers had restored the PDP-1. By 2010, another volunteer team of over a dozen retired IBM customer engineers restored two 1401 systems plus various unit-record equipment. In 2013, the museum opened two public demonstration labs to showcase the restored computers: in the DEC PDP-1 Demo Lab, visitors can pit dueling Spacewar! spaceships against each other on its cathode ray tube display. In the IBM 1401 Demo Lab, docents invite visitors to keypunch their names into punched cards that are then printed on its lively line printer. In this paper, we cover the restoration of the 1960s systems, how they are demoed and used, restoration logistics, ingredients for success, and conclude with thoughts on “Why restore and demo vintage computers?”

Introduction

Museums help shape the identity and character of a community by preserving and curating artifacts, chronicling key accomplishments, and presenting the human stories that underlie them. Our belief is that a technology museum can dramatically enrich its visitor experience by exhibiting and demonstrating authentically restored and operational vintage computers. Nevertheless, museum staff may feel unprepared to take on the challenges of a vintage computer restoration and demo project. What are the ingredients for a successful project? In this paper, we present an account of the successful restoration and demonstration of three 1960s vintage computers at the Computer History Museum (CHM).

The CHM’s mission is “to preserve and present for posterity the artifacts and stories of the Information Age” while bringing the history of computing to life via public exhibits, educational programs, volunteer docent-led tours, speaker series, national media and radio programs, and community outreach. In 1979, the museum, then located in Boston, began assembling what is now the largest collection of computing artifacts in the world.

In 2002, the CHM moved into its current facility at 1401 N Shoreline Blvd, Mountain View, in the heart of Silicon Valley. As the staff began strategizing on goals and exhibits, including plans for a flagship exhibition on computing history, a passionate volunteer petitioned the staff to shepherd a volunteer-driven project to restore and demonstrate a 1960s large vintage computer, believing that it could enrich the visitor experience.

In this paper, we summarize the CHM’s two volunteer-based projects that restored three large 1960s vintage computers: a DEC PDP-1 and two magnetic-tape IBM 1401s.¹ We also discuss how they are regularly demonstrated and their on-going usage. We conclude by recommending

¹ Beginning in 1998, CHM volunteers earlier restored an IBM 1620, “a prototype of a successful artifact intervention” (Spicer, 2005). It is no longer demonstrated to visitors, mainly due to the absence of compelling peripherals. Another team has restored a vintage IBM 350’s fifty-platter stack and servo access arm — the world’s first commercial disk drive — demoed every Wednesday in the Museum’s main Revolution exhibit.

ingredients for a successful restoration and demonstration project and our thoughts on “Why restore and demo vintage computers?”

Restoration of a DEC PDP-1 from Massachusetts

A restored vintage computer should offer a compelling visitor experience. In 2003, volunteers recognized that the DEC PDP-1, an early display-based interactive computer, would fulfill that objective. In particular, visitors could play the legendary Spacewar! video game on its large cathode ray tube (CRT). Spacewar!, an expression of the 1960s MIT hacker culture that shaped the early days of computing, was arguably the inspiration for the video-gaming industry. While only 55 PDP-1s were produced, they had a prominent presence at leading universities, including MIT, Harvard, Yale, Columbia, and Stanford.

Selecting between the three in its collection, the museum elected to restore a PDP-1 it had acquired from a small Boston firm in 1978. Starting in 1962, the firm had rented it out for human-machine research, such as investigating interface designs in early electronic word processing. Beginning in 2003, meeting once a week in a CHM conference room, a small volunteer team of engineers restored, over two years, its transistor-based CPU mainframe, point-plot Model 30 CRT display, paper tape reader, and Soroban console typewriter (Bickley, 2007).

Applying standard practice to its old power supplies, the restoration began by reforming their electrolytic capacitors.² The main challenges were the finicky CRT deflection amplifiers and the console typewriter. The team elected to substitute contemporary power supplies for the CRT’s 10 kV anode and the light pen’s 1 kV photomultiplier tube. A local typewriter repair shop was happy to take on the console typewriter. Only four printed circuit cards — DEC System Building Blocks — were found to be faulty. To enable visitors to play Spacewar!, the team built two button-actuated control boxes mounted on the exhibit railing (Figure 1).

The PDP-1 was fully operational and ready to demonstrate by 2005, when the team began giving weekly demonstrations. The following year, it was moved to a room just around the corner from the main entrance. The back wall has a large vintage photograph of engineers attending to a PDP-1 in DEC’s Maynard, MA facility. A looping video inside conveys the historical significance of the PDP-1 and Spacewar! (Plutte, 2014).



Figure 1. Volunteer and Spacewar! author Steve Russell gives a live demo in the DEC PDP-1 Demo Lab

² After removal, a programmable power supply slowly ramps up the applied voltage on a single capacitor (while safely limiting current), redepositing a natural oxide layer on its positive plate.

The PDP-1 is demoed on the first and third Saturdays of every month and Spacewar! tournaments are held quarterly (Russell, 2011; Verdiell and Bickley, 2017). Volunteers speak to DEC's lab instrument and computer business and demo the mesmerizing Snowflake and Minskytron CRT drawing programs. They next invite visitors to pit dueling Spacewar! spaceships against each other on the ghostly long-persistence phosphor CRT using the push-button control boxes. As an invaluable historic tie-in, two team members are Spacewar! authors. Another of the original MIT hackers presents the PDP-1 playing transcoded Bach scores in four-part harmony. Every December, continuing a vintage MIT tradition, a sing-along features the PDP-1 playing synthesized holiday music.

Restoration of an IBM 1401 from Germany

In 2003, while participating in the PDP-1 restoration, we became aware of an IBM 1401 listing on the German eBay. As many people had strong memories of 1401s — the world's most popular computer in the mid-1960s — the museum staff and volunteers felt that it would offer a compelling vintage computer demonstration. Visitors could keypunch cards, experience its lively chain printer and see spinning magnetic tape drives — the iconic face of computers in the 1960s. With over 15,000 installed worldwide, the business-oriented 1401 Data Processing Systems were an inflection point in computing history, leading the IT industry's transition from decades of unit-record equipment and accounting machines to an era of ubiquitous and more flexible stored-program computing. For many, the 1401 was their first experience with hands-on programming (Garner and Dill, 2010; Garner, 2013a).

When I volunteered to lead the project, my first question was "What is an IBM 1401?" When IBM announced it in 1959, I would not even have been tall enough to reach its control panel. Intrigued with the challenge of bringing a large 1960s vintage computer back to life and, having joined IBM Almaden Research in San Jose, I was curious about IBM's computing heritage. (My nascent computing experiences were on DEC, SDS, Univac, and GE systems.) But first I had to find volunteers with the hands-on instinct and experience to tackle the mechanically rich equipment. Following a tip, I placed an "An IBM 1401 Needs Help" ad in the region's IBM San Jose Retirement Club newsletter. Within a month, over a dozen retired IBM customer and manufacturing engineers stepped forward to help with its restoration (Garner, 2013b; Ross, 2009) — see Figure 2.

The German 1401, which we designated as the "DE 1401", was built in 1964 and operated by an insurance company 24x7 for eight years. In 1972, an entrepreneur procured it for his family service business in Hamm, Germany. In 1977, he mothballed it in a detached automobile garage for 27 years. After arriving at the museum in 2004, the new team of about a dozen volunteers rallied to meet the challenge with a spirit of commitment and determination. Many had been IBM Customer



Figure 2. IBM 1401 volunteer restoration team, vintage attire

Engineers responsible for keeping businesses up and running 24x7 from the 1950s onward. Some regarded the CHM as yet another mission-critical client, but fortunately without the stress and pressure of needing to get a customer's system back up and running ASAP.

The DE 1401 came with over a dozen units: 1401 CPU mainframe, 1406 extended core memory, and its electro-mechanical peripherals — 1402 card reader/punch, 1403 line printer, six 729 tape drives — along with various 1940s unit record equipment — several 026 key punches, 083 sorter, 077 collator, and a 513 reproducing punch. Luckily, the museum's facility had a 1,400 square-foot raised-floor server room that could accommodate the 12-kilowatt load, workbenches, test equipment, and volunteers. From photos taken in Germany, the team felt that the restoration was going to be challenging due to its prolonged storage in a humid climate.

We began the restoration by reforming the electrolytic capacitors in its over two dozen power supplies, replacing several due to deteriorated seals. Given its German pedigree, we searched for a source of 50 Hz, 3-phase power to convert from Europe's 50 Hz, 380 VAC to America's 60 Hz, 208 VAC power levels. We eventually acquired a robust 18-KVA Pacific Power 390-G analog converter. We deemed the alternative, modifying the ferroresonant supplies to run at 60 Hz and altering pulley ratios, as impracticable.

We began bug shooting the 1401 circuits and equipment using only oscilloscopes, schematics, and vintage manuals and tools. Given the project's size, a website scribe was enlisted to steadfastly photograph and record work sessions, biographies, manuals, schematics, specifications, part catalogs, vintage photos and memos, documents, anecdotes, and whatever else was fitting (Thelen, 2004).

While addressing the rusted metallic surfaces, we surprisingly found that corrosion had also impacted the discrete electronic devices. Diode leads had rusted through and even occasionally broken their glass encapsulants. Corrosion had also caused transistor leads to break open or compromise hermetically sealed packages. We learned that iron was used as an ingredient in device leads and packages as it, ironically, made for a better seal. We also found many faulty inductors, capacitors, and resistors, cracked board traces and corroded connector pins, and faulty fuses and circuit breakers. Luckily, there were no defects in the five miles of internal cable bundles or in the Gardner-Denver wire-wrapped backplane.

The 1950s germanium alloy-junction transistor used in the 1401 — the second type of transistor after the world's first point-contact device — had a reputation for being finicky and, since it does not possess a natural oxide layer like silicon, susceptible to surface contaminants. Over three decades of storage, air infiltration and other failings resulted in transistors with excessive leakage currents, low betas, opens, shorts, and unexplained "loopy" I-V curves. These faults led to difficult-to-find inter-signal circuit shorts and perplexing intermittent self-oscillating flip-flops. In the 1950s, solid-state transistors were heralded as a leap in reliability compared to vacuum tubes; now, half a century later, they no longer seem invincible.

One controversy was whether it was safe to treat the Standard Modular System (SMS) printed circuit card edge connector traces with commercial lubricant and corrosion inhibitors. I studied the corrosion literature and enlisted IBM Poughkeepsie's Materials and Process Engineering department to analyze the metallurgy of SMS edge traces. Their X-ray fluorescence analysis revealed a robust Au-Ni-Cu metal stack-up with 100 micro-inches of gold showing no porosity and little surface wear. Given the thickness of the gold layer, unheard of in contemporary connectors, they recommended at most cleaning the edge traces with isopropyl alcohol and cautioned against using commercial contact lubricants.

Restoration of an IBM 1401 from Connecticut

For three years we painstakingly tracked down and repaired over a hundred faulty SMS cards in the DE 1401. With no end in sight, I began to wonder whether it had been afflicted by too many corroded transistors. Luckily, in 2007, I received a call offering another 1401 system, a duplicate twin of the DE 1401. Built in 1961, it too had been operated by an insurance company 24x7, for eleven years. Also in 1972, an entrepreneur bought this 1401 for his family IT services business in Darien, Connecticut, and operated it in their home's humidity controlled basement until 1995, where it remained for another 13 years (without humidity control).

I petitioned the museum staff to procure what we designated as the "CT 1401" as a replacement for the evidently over corroded DE 1401. We raised funds from a dozen generous donors and on a rainy May day in 2008, a rigger extracted the heavy units from the home's basement. Perhaps sensing a twin sibling rivalry, or that its time was up, the week the CT 1401 arrived, the DE 1401 began to work (Figure 3)! Within eight months the CT 1401 was up and running.

One of the key restoration challenges was how to bring up the model 729 magnetic-tape drives and the CPU's Tape Adapter Unit (TAU). We wanted to debug the TAU first, but how without a working tape drive? The tape team decided to embark on what became a six-year project to design a sophisticated 729 analyzer and emulator unit that emulates up to six drives at various tape densities and transfer rates (Figure 4).

Users can remotely load and write virtual tape images using the GUI or web interface, patterned after the switch and lights panel at the top of a 729. The backend unit monitors and controls all 50 signals on the TAU-to-729 interface bus, using tunable op-amp circuits to mimic signal waveforms and cycle-accurate firmware running in an 8-bit, 8 MHz PIC-based DLP-245 microcontroller. The tape team coded thousands of lines of assembly code, C, Java, and HTML.

We first brought up the TAUs and then four tape drives on each system (of various models and ages, out of an original six units per system). Using a volunteer's machine shop, the tape team replaced the bearings and refurbished the magnetic powder clutches in the DE 729s. That resourceful and successful accomplishment notwithstanding, the rest of the team elected not to refurbish the CT 729s (nor the 1402 punched card units or 1403 line printers), electing instead to service individual components as they fail — rather than disassemble, refurbish, and reassemble entire units, each comprising thousands of components.

By 2010, the team had volunteered over 20,000 work hours coaxing both 1401 systems into an operational, maintainable and demonstrable condition. With both systems working, I petitioned the museum staff to retain them both so that live demonstrations could go on when one was down.



Figure 3. CHM's IBM 1401 restoration lab, before remodel; CT 1401 on left and DE 1401 on right



Figure 4. 729 analyzer and emulator

While we had proven that IBM's engineers had designed equipment that is still working after half a century, we also realized that live demonstrations could be easily foiled when any of a system's half-a-million aged components failed.

Luckily, it is has not been difficult to acquire authentic replacement parts. "New old stock" alloy-junction germanium transistors are available on the web and hundreds of spare SMS cards came either with the systems or were donated. While we have an inventory of 260,000 IBM-spec-compliant blank punched cards, we are on the lookout for a new source of cards for every day demo use.

An initial long-term viability concern was the 1403's print-chain assembly; in particular, its plastic band precisely encapsulating many loops of 3-mil wire to which type slugs are attached. Given that print chains had a reputation for breaking, we had elected not to run the infamous 1403 music programs that concurrently fire numerous print hammers. Thankfully, a retired IBMer at the TechWorks! in Binghamton, NY, who had worked on the 1403-N1 printer in the early 1960s, built a custom wire-winding jig and mandrel and crafted a new band and print chain using slugs from an existing broken chain — flawlessly operational for over a year.

The IBM 1401 Demonstration Experience

By 2010, both 1401s were up and running well, but we could not open the lab to the general public because it did not conform to building code's Class A public space. One requirement was for a second door for emergency egress. So we raised funds to remodel the lab and develop a public exhibit with interpretive displays.

In the 1960s, IBM architected theatrical and formal looking computer showrooms viewable through large plate-glass windows. Reflecting that history, the museum designed a showroom for the dual 1401 systems (Figure 5). Visitors enter the IBM 1401 Demo Lab from one side, proceed up onto a raised floor viewing area with equipment just behind an open railing, and depart the opposite side. The walls are colored vintage red-orange and sky-blue, matching a vintage showroom photo. The entry wall has a floor-to-ceiling professional photograph of a street-side showroom with bystanders peering in at a 1401 and its larger cousin. The back wall features interpretive graphics about the 1401 and its era. Looping videos on the wall behind the systems chronicle the 1401's heritage and show scenes from the restoration project (Plutte and Lonquist, 2013). On the way out, visitors can peruse a gallery of vintage marketing and customer 1401 photographs.



Figure 5. The Computer History Museum's IBM 1401 Demo Lab

The IBM 1401 Demo Lab opened to the public with a gala event in June 2013. Live 1401 demonstrations occur every Wednesday afternoon and Saturday morning and also for special events, high school and college class excursions, and 200-plus student Field Trip Days. Due to our redundant systems, docents have presented 1401 demos without interruption since the lab opened.

To prepare for public demonstrations, the museum augmented its volunteer interpreter training program to assist docents in presenting compelling demo experiences to audiences of all ages from around the world. With guidance from the museum's education team, they learn interpretive principles, techniques and universal themes to help deliver cohesive and relevant demonstrations. A pair of docents carry out the demos: the lead does most of the talking while an assistant handles logistics, such as loading punched cards, pushing buttons, and diagnosing machine problems (Madsen, 2017; Ross and Laughton, 2016). Naturally, there are awkward glitches, particularly with the 1402 card readers and 729 tape drives, but they soldier through them, manifesting the challenge of running a half-century-old computer.

The visiting public is generally enthralled and captivated by the live demos, particularly the younger visitors. There are smiles, questions and quizzical looks of astonishment: "This is what a computer was like back then?" Some experience punched cards for the first time and find it hard to believe they were the prevailing way to get information into a computer. Younger engineers come away with a sense of how much more convenient programming is today and better appreciate how applications had to be squeezed into just 16K characters or less of memory. Visitors respond with amazement when shown the two-inch-diameter 1950s equivalent of today's USB cable (Figure 6). They are surprised to hear that, at the 1401's memory's inflation-adjusted cost of \$20 per byte, a mobile phone would set them back over a trillion dollars. They are also surprised that one of our 1401 systems would go for \$2.5 million in today's currency.

Demo sessions are generally half an hour, beginning with the question: "What do your senses say about this room?" The lead docent describes pre-1960s business data processing and demos the 083 sorter as an impressive example of the era's electromechanical unit-record equipment — cards flying through at over 16 per second while sorting on one column. Moving on to the 1401, they discuss how it revolutionized the business IT market by supplanting decades of unit-record accounting machines.



Figure 6. CHM's IBM 1401 demo lab demo session, showing inter-unit cabling

Next the docents invite a visitor to keypunch their name. After adding their punched card to the 'BigPrint' demo deck, the 1403 chain printer noisily hammers it out in big letters: "ALAN TURING VISITED THE COMPUTER HISTORY MUSEUM ON JULY 22, 2015" (Figure 7).

The demo continues with the 'Powers-of-Two' program, rapidly printing a new line every 10th of a second for each number in the series 2, 4, 8, and so on, until all 132 columns are filled. The cadence of print hammers pounding out increasingly wider lines builds to a cacophonous crescendo, particularly with its cover open.

The docents then demo the tape drives, with their reels rapidly starting and stopping and loops of tape jiggling behind glass-paned vacuum columns. The docents demo the tape drives either via the TAU control panel or a program that mimics a tape sort. A high-speed rewind of four tape drives dramatically ends the demo. Visitors are invited to linger, ask additional questions, and keypunch a card. Younger visitors are particularly drawn to keypunches and enjoy composing "secret" messages as patterns of holes in punched cards.

With large school groups, the docent may suggest: "If you ask a bona fide question, you can punch a card and get a printout." Hands shoot up (Figure 8a). For special events, BigPrint name cards can be punched in advance via the lab's PC-controlled 029 keypunch. An 026 keypunch can also be positioned outside the lab to reduce queuing.




Figure 7. Visitor's name on 1403 line printer.



Figure 8.

(a) Left: Asking questions for a chance to keypunch

(b) Right: Trying out an IBM 001 manual keypunch

The lab features a diminutive 1920s IBM 001 Manual Keypunch where motivated visitors doggedly press buttons to punch holes, one column at a time — an engaging way to touch and engage with vintage punched card hardware (Figure 8b).

Ongoing Exploration and Programming of the IBM 1401

The fully operational 1401s have drawn the attention of nearby Silicon Valley engineers, intrigued by how they differ from contemporary computers. Quoting from one: “Studying old computers such as the IBM 1401 is interesting because they use unusual, forgotten techniques. Old computers are also worth studying because their circuitry can be thoroughly understood. After careful examination, you can see how arithmetic, for instance, works, down to the function of individual transistors” (Shirriff, 2015d). Steve Wozniak enthusiastically felt that the 1401 is “cool” because people can explore and learn how computers function from discrete transistors up to basic programming. A precocious ten-year-old visitor, quickly mastering the 1401’s easy-to-learn instruction set, coded and ran a simple “Hello World” program.

Several volunteers and visitors have coded extraordinary non-business programs for the 1401, including an implementation of the SHA-256 cryptographic hash function used by Bitcoin (Shirriff, 2015b), a fractal Mandelbrot graphical printout (Figure 9) (Shirriff, 2015a), and an elementary 3D ray-tracing program (Kesteloot, 2014). Our contemporary studies of the 1401’s architecture include exploration of its qui-binary decimal encoding for arithmetic error detection (Shirriff, 2015d) and an elucidating look at its magnetic core memory circuits (Shirriff, 2015c). Another volunteer authored a 1401 theory of operations manual using contemporary electrical engineering terminology (Fedorkow, 2017).

Another extraordinary use of the 1401 came from a San Jose State University compiler class student who authored a challenging Small-C cross compiler for it. He also wrote an iOS app that re-renders a photograph for the 1403 printer using a vintage grayscale-to-overprinted-characters algorithm. And a media-savvy volunteer has produced several widely viewed videos featuring the restoration team debugging and maintaining the systems (Verdiell, 2015 and 2017).

For creating new 1401 software, an integrated development and simulation environment called ROPE is available for Windows, Linux, and Mac OS (Mak, 2012). ROPE integrates a GUI, a 1401 Autocoder assembler and the SimH 1401 simulator. Users can edit, assemble, and debug a program using breakpoints, while separate windows display the contents of memory, CPU indicators, and line-printer output. ROPE produces an object deck image that can be punched using our PC-controlled 029 keypunch.



Figure 9. Mandelbrot Set fractal on IBM 1403 line printer

We have two PC-based methods to get information in and out of the 1401s. Using the 729 tape analyzer and emulator's web interface, a tape image can be virtually read or written from any IP-connected browser. We regularly run the vintage 1401 diagnostics this way. We also interfaced a PC to the 1401's serial port. (The CT 1401's owner had hooked up an Altair 8800 for this same purpose).

Exploring the 1401's history has connected the team with the IBMers who designed and brought it to market in the late 1950s. For the 50th anniversary of its 1959 announcement, I tracked down its design architect, development program manager, market planner, hardware and software developers, and the co-designer of its Parisian predecessor prototype — the World Wide Accounting Machine. In 2009, I organized a "50th Anniversary of the Legendary 1401" celebratory event with presentations by the original leadership team, recordings of their oral histories, and a commemorative pamphlet (Wichary et al., 2009).

Restoration Logistics

In 2007, the museum formed a committee to formulate a Restorations Policy. This policy outlines restoration project goals, intervention types (conservation, restoration, reconstruction, replication, and model), criteria for selecting candidates, and monitoring of projects (including a restoration team staff liaison). The museum has classified the restored PDP-1 and 1401s as part of its study collection, separate from artifacts accessioned into its permanent collection. In addition to the machine being restored, another instance should be preserved unmodified in the permanent collection.

The museum also issued companion Restoration Guidelines that emphasize safety first and stipulate that two restorers must be present when equipment is open and energized. Its covenants include: maintaining historical integrity and authenticity, reversible alterations, and recording and logging of restoration activities.

To house test equipment and maintenance tools, the museum generously provided a 480 square foot workshop near the demo labs. The shop's equipment, tools, electronic test equipment and spare parts are donated items or personal property. Classic-computing aficionados worldwide also donate vintage parts, tools, and equipment after finding our 1401 restoration website.

Safety of the volunteers and visitors is paramount. In addition to the two-person work rule, a light fixture atop the mainframes indicates when high voltages are present inside the equipment. The museum is in an earthquake-prone region, so sudden shaking of the heavy 1-to-2 ton units is a concern. While some are resistant to toppling because of their low center of gravity, to prevent units from scooting across the floor during an earthquake we placed C-shaped stabilization clips around

their casters. In the 026 keypunches, we bypassed timeworn selenium rectifiers with silicon diodes, which could otherwise emit an odiferous toxic gas when they fail.

Ingredients for Successful Restoration and Demonstration Projects

Based on the CHM's restoration and demonstration project experience, here are my recommended ingredients for a successful project:

Compelling demo: From a project's inception, it needs to offer a compelling demo experience to visitors of all ages. Large and noisy peripherals hold attention and interest. Blinking lights are not so meaningful. Live demonstrations not only explain how a restored computer works, but should also cover why it was developed and how it was used. As a gateway to the past, broader social narratives can also be conveyed.

Restoration team: The size of the team needs to be commensurate with the size of the restoration project. Restorers do not need to have hands-on experience with a particular machine, but must be willing to roll up their sleeves and attack difficult challenges. For the long haul, younger folk need to be enlisted so that elders can pass down skills, empirical knowledge, tricks-of-the-trade, stories, anecdotes and enthusiasm.

Restorer temperament: A key motivator is the satisfaction that comes from interacting with wide-eyed visitors experiencing live "compusaur" and hearing first-person stories and anecdotes. Nostalgically fixing old hardware isn't enough. A playful sense of humor, camaraderie, and sense of adventure helps to keep spirits high. Restorers may enjoy interacting with visitors ambling through the lab during restoration sessions, enlivening tedious work. Social get-togethers with spouses and families are valuable.

Restoration pace: Unlike the frenzied pace of new product development, where value declines the longer a project takes, a restored vintage computer becomes more valuable the longer it takes. Thus, restorers are rarely subject to schedule pressure and can take the time to study all the angles or unintended side effects of a contemplated machine alteration.

Museum liaison: The museum's restoration-staff liaison needs to keep a pulse on the volunteers' concerns, advocate for their needs, and keep them abreast of the museum's activities.

Dual systems: Ideally, if two systems are available, demos can still go on when one is down. Also, the restoration team can compare the behavior of one against the other to resolve issues.

Workshop: A workshop is needed to separately house the potpourri of tools, equipment, spare parts and workbenches for troubleshooting, inspecting, and repairing circuit cards and mechanical items.

Web scribe and documentarian: The project should have a website master and scribe responsible for recording restoration activities, archiving manuals and documents, maintaining team member bios, photos, email exchanges, and so on. Relying on team members to contribute to a generic blog doesn't work.

Why Restore and Demo Vintage Computers?

Visitors appreciate a museum that supports restorers putting their hearts and minds into keeping vintage computers alive and sharing their personal stories. A restoration project can help fulfill a museum's mission of preserving and interpreting computing history by promoting and chronicling the maintenance, operation and demoing of a vintage computer — not possible with conventional static artifact preservation.



Figure 10. A live 1401 demo and time machine experience

When demonstrating the authentic sights, sounds and smells of a living vintage computer, we connect with visitors' minds and imaginations, taking them on a time-machine journey to an era of chattering punched card equipment, cacophonous printers, whirling tape drives, synthesized music, and ephemeral CRT images (Figure 10). As they become immersed in the experience and hear the history and stories, we also help fulfill the museum's education mission. As they see firsthand how computers have dramatically changed over the past half century, they may even ponder: "What will the world of computing be like in the next half century?"

Acknowledgments

For detailed feedback, I especially thank Ron Mak, Chris Beall, Dag Spicer, Katharina McAllister, and my wife Robin Beresford. Also for reviews and suggestions: Michael Albaugh, Lyle Bickley, Carl Claunch, Michael Cheponis, Jennifer De La Cruz, Nora Eibisch, Guy Fedorkow, Iggy Menendez, Stan Paddock, Allen Palmer, Peter Samson, Grant Saviers, Ken Shirriff, John Shoch, Luca Severini, Ken Sumrall, Ed Thelen, and Marc Verdiell.

DEC PDP-1 restoration team: Michael Cheponis (lead), Lyle Bickley, Tim Coslet, Joe Fredrick, Bob Lash, Steve Russell, Peter Samson, Rafael Skodlar, Eric Smith, and Ken Sumrall.

IBM 1401 restoration team: Robert Garner (lead), Stan Paddock (assistant lead); subgroup leads: Ron Williams (1401), Frank King (1403, Bill Flora (1402), Allen Palmer (729), Ron Crane (power/analog), Bob Erickson (unit record), Ed Thelen (web master/scribe); George Ahearn (original 1401 designer), Carl Claunch, Don Cull, Guy Fedorkow, Bob Feretich, Matthias Goerner, Judith Haemmerle, Jim Hunt, Chuck Kantmann, Glenn Lea, David Lion, Don Luke, Doug Martin, Iggy Menendez, Bill Newman, Joe Preston, Grant Saviers, Ken Shirriff, Jeff Stutzman, Milt Thomas, and Marc Verdiell.

Demo Lab design and interpretive media: Kirsten Tashev (CHM VP of Exhibits), Jon Plutte (CHM media director), Jas Lonnquist (screenwriter), and Van Sickle & Rolleri Ltd (exhibit design).

References

- Bickley, L. (2007). *PDP-1 Restoration*. [Online] Available at: <https://youtu.be/O3-DdJPoB18>
- Fedorkow, G. (2017). *IBM 1401 Theory of Operation*. [Online] Available at: <http://ibm-1401.info/IBM-1401-Theory-of-Operation-GF.pdf>
- Garner, R. (2013a). STARS: Early Popular Computers, 1950-1970. *Proceedings of the IEEE*, 101(9), pp. 2134 - 2140.
- Garner, R. (2013b). *Restoring the IBM 1401*. [Online] Available at: <http://www.computerhistory.org/atcm/restoring-the-ibm-1401/>
- Garner, R. and Dill, R. (2010). The Legendary IBM 1401 Data Processing System. *IEEE Solid-State Circuits Magazine*, Winter, 2(1), pp. 28-39.
- Kesteloot, L. (2014). *IBM 1401 Ray Tracer*. [Online] Available at: <https://www.teamten.com/lawrence/projects/ibm-1401-ray-tracer/>
- Madsen, A. R. and S. (2017). *IBM 1401 live demo at the Computer History Museum*. [Online] Available at: https://www.youtube.com/watch?v=_QDCaUh7QYo
- Mak, R. (2012). *1401 Software Development and ROPE*. [Online] Available at: <http://ibm-1401.info/index.html#1401-Software>
- Plutte, J. (2014). *Spacewar!* [Online] Available at: https://www.youtube.com/watch?v=_L1HeZ2kPck
- Plutte, J. and Lonquist, J. (2013). *1401: Dawn of a New Era*. [Online] Available at: <http://www.youtube.com/watch?v=ZPpV8X91neQ>
- Ross, K. and Laughton, P. (2016). *Ken Ross and Paul Laughton demo the IBM 1401*. [Online] Available at: <https://www.youtube.com/watch?v=VQueCt114Gk>
- Ross, P. (2009). Rebuilding the IBM 1401. *IEEE Spectrum*, 30 October, p. 5.
- Russell, S. (2011). *Steve Russell talks about his early video game Spacewar!* [Online] Available at: <https://www.youtube.com/watch?v=PnJvZHegg8I>
- Shirriff, K. (2015a). *12-minute Mandelbrot: Fractals on a 50-year-old IBM 1401 mainframe*. [Online] Available at: <http://www.righto.com/2015/03/12-minute-mandelbrot-fractals-on-50.html>
- Shirriff, K. (2015b). *Bitcoin mining on a 55-year-old IBM 1401*. [Online] Available at: <http://www.righto.com/2015/05/bitcoin-mining-on-55-year-old-ibm-1401.html>
- Shirriff, K. (2015c). *Examining the core memory module inside a vintage IBM 1401 mainframe*. [Online] Available at: <http://www.righto.com/2015/08/examining-core-memory-module-inside.html>
- Shirriff, K. (2015d). *Qui-binary arithmetic: how a 1960s IBM mainframe does math*. [Online] Available at: <http://www.righto.com/2015/10/qui-binary-arithmetic-how-1960s-ibm.html>
- Spicer, D. (2005). IBM 1620 Restoration Project. *IEEE Annals of the History of Computing*, 27(3), pp. 33-43.
- Thelen, E. (2004). The IBM 1401 Demo Lab and Restoration Project. [Online] Available at: <http://ibm-1401.info>
- Verdiell, M. (2015). *Debugging the 1959 IBM 1401 Computer at the Computer History Museum*. [Online] Available at: <https://www.youtube.com/watch?v=PwftXqJu8hs>

Restoring and Demonstrating 1960s Vintage Computers at the Computer History Museum

Verdiell, M. (2017). *IBM 1401 Restoration Team: A Day at the Office*. [Online] Available at: <https://www.youtube.com/watch?v=buN6AzdngZQ>

Verdiell, M. and Bickley, L. (2017). *Lyle Bickley demos the DEC PDP-1 at the CHM*. [Online] Available at: <https://youtu.be/1EWQYAfMYw>

Wichary, M., Spicer, D. and Garner, R. (2009). *IBM 1401 Legendary Data Processing System*. [Online] Available at: <http://www.magcloud.com/browse/Issue/37765>

Picture Credits

Figures 1 - 8, 10: Robert Garner; Figure 9: Ken Shirriff.

IBM System/360 Printer Revitalization at Techworks!

Robert Lusch, Nicholas Hekman, and Don Manning

Center for Technology & Innovation, Inc.

Abstract

A three-generation volunteer team at the Vintage IBM Computing Center at TechWorks! in Binghamton, NY, has revitalized a 1960s high-speed line printer, designed and built for IBM's System/360. Using a three-pronged approach, this 2-plus year effort combined conservation of the printer's mechanical components, reverse engineering of controller hardware in today's technology, and implementation of real-time software to perform complex functions originally controlled by hardware. This System/360 printer is regularly demonstrated for TechWorks! visitors, who can take home personalized souvenir banners printed on classic 20th century fan-fold paper.

Introduction

The mission of the Center for Technology & Innovation, Inc. ("the Center") in Binghamton, NY, is to "document and present in context the inventions and industrial innovations of New York's Southern Tier" — a region in upstate New York that extends from Cooperstown in the northeast to the New York State borders with Pennsylvania and Ohio to the south and west. This area has a two-century legacy of pioneering companies with ingenious workforces, including IBM Corporation, General Electric Aerospace Group, Ansco/GAF, Link Aviation, Corning Inc., Raymond Corporation, Lockheed Martin Mission Systems and Training, BAE Systems, Universal Instruments, and dozens of others. Located within a half-mile radius of the Center are the birthplace of flight simulation; the railroad station that sent the world's first land-based mobile text messages; and the first factories for Nelson Stow flexible shafts, Bundy Manufacturing time clocks, and McIntosh Laboratory audio equipment (see www.ctandi.org/ideashed).

The Center is focused on developing a destination venue called TechWorks!, where visitors can "Experience Innovation - Past, Present & Future." One goal of TechWorks!, as the name implies, is to showcase vintage *technology that works*, because engaging multiple senses helps stimulate visitor interest and learning. For example, hearing the "chunk, chunk" of a keypunch, experiencing the sensations of "flying" a Link General Aviation Trainer, trying to read personalized embossed Braille messages by touch, singing along to music from a Link nickelodeon, seeing images from inside a walk-in pinhole camera, and watching a 50-year-old printer create souvenir banners as fast as (albeit more loudly than) today's desktop printers all contribute to a memorable experience.

The Center is developing TechWorks! from the inside out in a 30,000 square foot former ice cream factory built in 1912. A major exhibit area is the Vintage IBM Computing Center (VICC), where a variety of mid-20th century IBM products are in operating condition or are restoration works-in-progress. Among these is the operational IBM System/360 printer, whose story of revitalization is the subject of this paper.

Background

The early 1960s was an inflection point in the data processing industry as products transitioned from electro-mechanical accounting machines and vacuum-tube-based systems to fully transistorized general-purpose computers. An important example of the latter is IBM's 1401 Data Processing System announced in October 1959 (Garner and Dill, 2010) — IBM's first production use of printed circuitry. That system also introduced a new type of high-speed line printer designated the 1403, which could print 132-character lines at speeds up to 600 lines per minute (Figure 1, top). According to Campbell-Kelly and Aspray (2004, p. 119), the significance of this printer is not to be underestimated:

The 1401 was certainly an excellent computer, but the reasons for its success had very little to do with the fact that it was a computer. Instead, the decisive factor was the new type 1403 “chain” printer that IBM supplied with the system.

Bashe et al. (1986, p. 480) agree: “The timely introduction of an entirely new class of printing devices [chain printers] was an important factor in the success of the 1401 computer system.”

The innovative “print chain” used in the first models of the 1403 printer family was made up of two-character engraved type slugs mounted on a flexible ribbon of wound wire. The chain rotated in a horizontal direction at a constant velocity such that each character passed by each print position at some point in time. When a desired character aligned with its proper print position, a hammer was activated from behind the paper to push it forward, into the inked ribbon and onto the passing type slug. Unlike typewriters and other printers that drove the type into the paper from the front, this was known as “back-printing” because the impetus from the hammer came from the rear of the paper. It was also called “on-the-fly” printing, since the chain was moving continuously throughout the process (Nickel and Kania, 1981).

While the 1403 chain printer was a radical improvement over other printing techniques of the mid-20th century, further increases in performance were constrained by the maximum speed at which the chain could be reliably driven (90 inches per second). In addition, maintenance was a concern, as chain breakage occurred more often than desired. (See Manning (2017) for a procedure to reconstruct broken print chains.) To meet the demand for increased throughput and reliability, IBM modified later models of the 1403 family by replacing the print chain with another innovation called a “print train”. The print train consisted of three-character type slugs riding on a monorail, driven at more than twice the previous speed (206 inches per second) by a gear-toothed rack on the back of each slug (Figure 2). To reduce slurring of the printed characters, the faster type speed required a decrease in hammer impact time. Hence, the hammer unit also underwent a major redesign to reduce hammer mass and increase hammer velocity, while maintaining the same impact energy. These changes nearly doubled throughput to 1100 lines per minute for the 1403 Model 3 and Model N1, which were offered with the IBM System/360. The 1403 Model N1 also introduced acoustically insulated covers to reduce noise levels (Figure 1, bottom) and an adjustable power stacker to improve paper handling (Nickel and Kania, 1981). It is the revitalization of an example of this later, train-driven, higher-performance version of the 1403 family, the 1403-N1, which is described herein.



IBM 1403 Printer

The IBM 1403 printers (top and bottom photos) are permanent-record output devices available to a variety of data processing systems, including the IBM System/360. The 1403 comes in seven models, representing an assortment of speed ranges and printing capacities as follows:

IBM 1403

Model	Print Positions	Maximum Speed	With Systems
1	100	600 lpm	1401 (Models A, B, E, F), 1410, 1420, 7010, 7040, and 7044.
2	132	600 lpm	1401 (Models A, B, C, D, E, F), 1410, 1420, 1440, 1460, 7010, 7040, 7044, and System/360 Models 30, 40, and 50).
3	132	1,100 lpm	1410, 1440, 1460, 7010, 7040, 7044, and System/360 (Models 30, 40, and 50).
4	100	465 lpm	1401 (Model G only).
5	132	465 lpm	1401 (Model G) and 1440.
6	120	340 lpm	1401 (Model G) and 1440.
N-1	132	1,100 lpm	System/360 only.



IBM Printer, Model N-1

Figure 1. IBM 1403 Printers (from (International Business Machines (IBM), 1964, pp. 4-5))

Note in the table in Figure 1 that the 1403-N1 was offered only on the System/360 and not on the 1401 or other prior systems. Therefore, to eliminate confusion with the earlier chain-driven models of the 1403 printer family, the subject 1403-N1 is designated as an IBM System/360 printer throughout the text that follows.

Like many computer peripherals of the era, the 1403 family of printers did not contain the electronic intelligence needed to effect printing; rather, their controllers were external to the printers. Initially,

the controls were provided by circuitry within the 1401 data processing system, including sharing of the system's core memory (Bashe et al., 1986, pp. 492-493). With the advent of the System/360 and its standardized I/O channel interface, however, the control functions moved to a separate box called the IBM 2821 Integrated Control Unit (ICU), to which the printers were attached. Like the 1401, the 2821 ICU was built with IBM's first-generation printed circuits — Standard Modular System (SMS) cards — to interface to the System/360 host and to operate up to three printers and one card reader/punch. In addition, it contained its own magnetic-core buffer memory as well as a "power tower" to supply various voltages to its own electronics as well as to drive the hammers and other circuitry in the attached printer(s). (IBM, 1965c)

Revitalization project description

Revitalization of the IBM System/360 printer began in the spring of 2013 and, slightly more than two years later, made its TechWorks! VICC Action Debut in May 2015. To accomplish this feat, the Center recruited a team from the surrounding community consisting of IBM retirees, current employees or veterans from other companies, and engineering students from Binghamton University. These three generations of computer professionals were essential to the success of the project. The IBM engineers whose careers spanned the 1950s to the 1980s (and who were part of the printer's




Figure 2. Print train schematic with (inset) type slug close-up (from IBM 1964, pp. 69, 79)

original design team) were well versed in electro-mechanical systems and discrete circuits, and provided in-depth understanding of the original design as well as rare documentation from personal archives. Team members with work-experience from the 1970s to the 2010s contributed their expertise in integrated circuits, packaging technology, and microprocessor control of electro-mechanical equipment to meet microsecond-timing requirements. They also helped bridge the terminology gap between the previous generation and student engineers of the millennial generation. The students, in their final year at Binghamton University's Watson School of Engineering, delivered software and hardware to control the printer and built the user interface on a personal computer to manage printer operations. By tapping multi-generational talent to take advantage of their areas of expertise, the synergy of idea exchange between the generations was palpable, expanding insights of all team members and transferring accumulated knowledge across the decades.

The printer was received in reasonably good condition, but with very little documentation and without its required 2821 printer controller. After searching in vain for an extant 2821 ICU and finding only sparse documentation of the device, the revitalization team settled on a three-pronged approach to their work:

1. Conserve the printer's vintage mechanical, electrical, and hydraulic components;
2. Reverse-engineer the hammer driver circuitry in today's technology and create a new card and chassis to interface to the original cables; and,
3. Implement real-time software to perform the complex functions needed to accomplish printing, originally done by hardware in the 2821 printer controller.

Mechanical conservation

For the first step, a mechanical engineer from the original printer design team went over the printer with a fine-toothed comb. The mechanical components required only cleaning and new lubricants, with one critical exception: the idler gear between the train drive motor and the train drive gear was inexplicably missing. This idler gear, in addition to being needed to spin the train, is also used to generate timing pulses to indicate the "home" position on the train. The team contacted the Computer History Museum (CHM) in California looking for a spare gear or an engineering drawing of the part. They did not have either, but they did have a complete duplicate printer in their collection from which careful measurements of the gear were made. Working with this information, plus his own memory of the idler gear design, the team's mechanical engineer recreated detailed specifications for the component. A local machine shop replicated the gear in heat-treated steel.

The printer's hammer unit was opened and inspected, and each hammer and actuator was tested for freedom of movement. One of four tractor belts for moving the paper was missing pins and was replaced with a spare obtained from the CHM. The AC motors driving the print train, the hydraulic unit, and cover raise/lower mechanism were verified to function smoothly. The hydraulic unit was checked for proper oil volume, free-moving valves, and no leaks. Finally, a new ribbon was purchased and installed.

In order to simplify the task at hand, the team made a conscious decision to set aside some printer functions not required for basic printing operations. For example, the carriage-control tape unit and

its associated front-panel control buttons — used for positioning the paper and performing high-speed skips — were disabled. Instead, the operator manually aligns the paper as needed. In addition, various status and error indicator lights, as well as most interlocks, were bypassed; so too was the circuitry used to stop ribbon movement when not printing. Lastly, the SMS amplifier circuitry for the train timing pulses was replaced by LM741 op-amps powered from small ± 5 V regulators. Making these trade-offs eliminated the need to supply multiple voltages to the printer and allowed many interface signals to be ignored, thus reducing the amount of software to be written.

Controller hardware re-creation

The second prong in the revitalization approach involved re-implementing the hammer driver circuitry in today's technology. For early models in the 1403 printer family, the hammer drivers were implemented with SMS single-layer printed circuit cards populated with discrete components to control two hammers (Figure 3, left). Each driver circuit produced a 60 V, 5 A electrical pulse to energize the hammer magnet coil and "fire" the hammer. Other SMS cards containing single-shot circuits controlled the timing of the pulses to about 1.2 ms in duration, a critical parameter for high-quality printing. By the time this hardware had migrated to the 2821 ICU for controlling System/360 printers, it had evolved to denser packaging: four drivers and four timing circuits, all contained on a double-width SMS card.

Rather than attempting to recreate the 2821 design in scarce 1960s technology, the VICC team chose to re-implement the driver functions using today's electronics and packaging. Each new two-sided, 100 mm by 80 mm printed circuit card has drivers for eight hammers and a microcontroller replaces the single-shot timing circuits (Figure 3, right). A simple command bus directs the microcontroller to fire a desired hammer: the embedded software activates the driver circuit, counts




Figure 3. Hammer Drivers: original SMS card (left) and modern replacement card (right)

the required pulse width, then deactivates the circuit. Having software control of this process allowed rapid debug and tuning of the operation. Furthermore, by defining additional commands, this same card is used to control the hydraulic valves in the printer that move the paper-carriage mechanism.

Another significant change was made to the printer to simplify the project. Generally, SMS circuitry used germanium transistors, multiple bias voltages, and favored the P-N-P type of transistor. In this environment, it made sense for the original hammer driver pulses to switch between -60 volts and ground. Modern technology is based primarily on silicon and N-P-N transistors, which work well with positive voltage levels. Hence, it was much easier for the new hammer driver card to switch between ground and $+60$ V instead of -60 V. The hammer magnet coils operate the same way in either case, with current flowing in the opposite direction. However, this necessitated reversing several polarized capacitors within the printer. And, yes, a couple were missed due to scant documentation and ruined capacitors had to be replaced.

The team used a standard 19 inch equipment rack to contain the modern printer controller. At the bottom, an adjustable lab power supply provides $+60$ V for the hammers, and a contactor to control the 3-phase, 208 V input for the printer's AC motors as well as the single-phase 120 V feed to other system components. A tray above the power supply holds the personal computer that acts as the user interface. At eye-level, for visitor viewing, the team created a chassis to hold the new driver cards. One of the Center's partners, Triple Cities MakerSpace, created custom card holders in ABS plastic using a 3D printer. A block of vintage SMS connectors at the rear of the chassis is the system's "blood-brain barrier": one side accepts the SMS paddle cards from the printer's original interface cables and the opposite side connects via cabling to the 21st century hardware (Figure 4). This delineation presents visitors with a striking side-by-side visible comparison of the advances in electronic packaging in the last half-century.




Figure 4. Printer Controller Chassis: original 1960s paddle cards visible at top right

Software implementation

The third portion of the revitalization project replaced printer controller functions previously done in hardware with software running on a 200 MHz chipKIT™ Wi-Fire microprocessor, connected via a USB port to the PC running the user interface. The microprocessor also receives printer-train timing pulses from the amplifier circuitry in the printer. From the PC, the user initializes the USB connection, then selects a .txt file to be printed and starts the operation. The chipKIT™ microprocessor receives the print data from the PC and analyzes the text to be printed. Using the timing pulses, its real-time software determines when each needed character is aligned to its proper print position. With sub-microsecond accuracy, it sends a command on the command bus to a hammer driver card to fire the corresponding hammer at the precise moment of alignment. Note that while it appears to an observer that each line of text is printed all at once, in fact the characters are printed one at a time, with approximately 5 μsec between hammer strikes. Additionally, text generally is not printed in left-to-right order; rather, whenever any needed character aligns with the proper column, its hammer is activated. The software continues working until the entire print line is completed, then issues the command to move the paper to the next line and begins again.

Lines are printed at a fixed pitch of ten characters per inch, which implies that the spacing of the print hammers is one-tenth of an inch. Standard trains use a 48-character array of capital letters only (A-Z, 0-9, plus twelve special characters) and contain five copies of the array (240 total characters on 80 type slugs). The algorithm for determining the printing sequence is based on the hammer spacing (0.100 inches) and the character spacing (0.1505 inches) on the train — every third hammer aligns with every second character in sequence as the train moves 0.001 inches (see Figure 5). Thus, one third (44) of the 132 hammers will have lined up with some character as the train moves 0.043 inches (assuming that the first character is already aligned to the first hammer at time zero). This “ripple” sequence is known as a print sub-scan. After this, there is a short interval to synchronize the second character to the second hammer, then the second sub-scan starts. Three print sub-scans make up a full print scan, during which all 132 hammers will have had the option to print one character (but not the same character). Since there are 48 characters in the array, 48 full print scans are required to expose each character to every hammer, thus completing the print line. Note that each hammer is fired at most once per print line; multiple strikes cannot occur. However, by suppressing spacing between lines, overprints can be obtained if desired.

Perhaps this printing algorithm can be better understood by viewing an animation of it, found in Shirriff (n.d.). This animation was done for the 1403 models using print chains, but it applies to the System/360 model with its print train as well. The only difference is that the chain moves 0.001 inches in 11.1 μs while the faster-moving train covers that distance in only 4.85 μs. Chain printers could print a full line in 80 ms and single-space the paper in 20 ms, resulting in throughput of 600 lines per minute. The System/360 printer accomplishes these two tasks in 55 ms (1100 lines per minute).

In order to implement this algorithm in software and still maintain throughput with the new printer controller embodiment, the chipKIT real-time software written by the VICC team had to be highly efficient, even given the microprocessor cycle time of just 5 ns. Taking too long in code to evaluate whether to exercise an available option to print when presented can result in late hammer strikes and partial or incorrect character imprints. The code was written, simulated, and verified in stages —

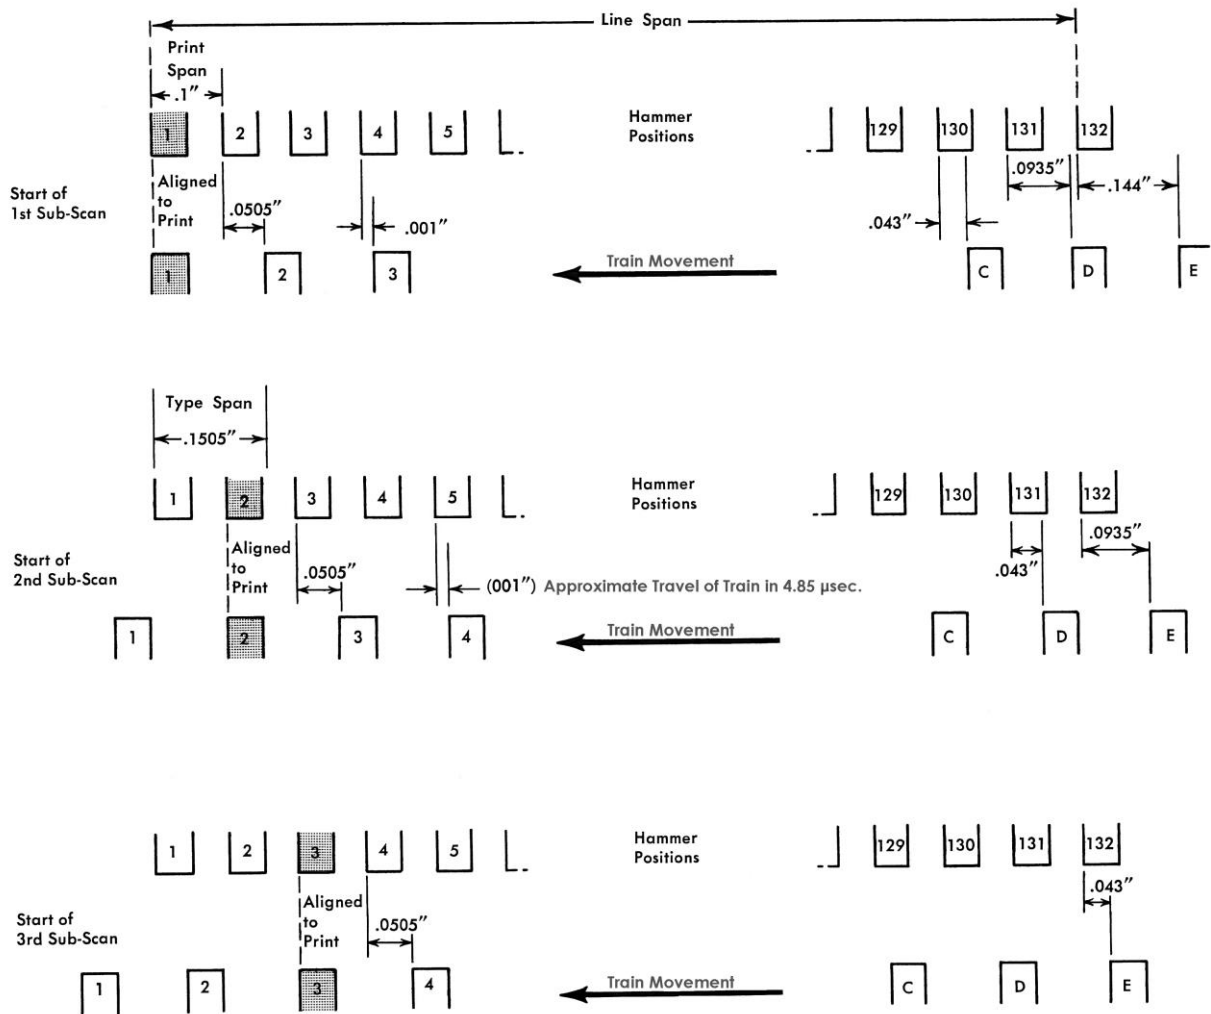


Figure 5. Schematic of type movement in relation to print hammers (adapted from IBM (1964, p. 39))

first to get the basics of the printing algorithm implemented with later revisions adding configurability, control of the paper carriage, and improved performance. Debug, tuning, and maintenance for high-quality printing was relatively rapid using this software approach.

Integration and on-going efforts

The three elements of the project — vintage mechanics, new driver and amplifier hardware, and controller software — were integrated in April/May 2015. The amplifier timing pulses were inspected with an oscilloscope and it was verified that the software could receive and interpret them properly. Then, the hammer driver outputs were scoped to assure the correct pulse width and magnitude when hammers were fired. A spare hammer actuator was then attached and tactilely checked for movement when activated. Finally, one hammer driver card was attached to the printer via its vintage interface cables. Power was applied to the system and nothing untoward or unexpected happened. The operator initialized the controller and commanded it to print some lines, and, with some amazement even among the VICC team, the first characters printed in decades appeared on paper (Figure 6, left). (Note: Each hammer card handles either eight even-numbered or



Figure 6. First printing, 9 May 2015 (left); visitor with personalized souvenir banner (right)

eight odd-numbered hammers, which accounts for the spaces between the characters in the printout.)

After just over two years of effort, the team’s successful revitalization of the System/360 printer was celebrated publicly on 16 May 2015, as part of the Binghamton University graduation activities for the second group of Watson School seniors (Figure 7). The action debut event — *Before Silicon Valley, there was IBM Endicott & Owego* — featured printing of 16 characters per line and remarks by Tommy Hing-K Lam (Lockheed Martin Fellow), Emerson Pugh (IEEE Past President), and Donald Seraphim (IBM Fellow, retired). Since then, work has continued to further tune the software, install more hammer driver cards to add more print columns, and improve the chassis for better visibility. The system now allows the operator to print custom banners in real time from data entered by visitors, who can take away the souvenir banners that they created and watched being printed (Figure 6, right).

Future endeavors under consideration include:

1. Developing an application to print low-resolution images or visitor “selfies”;
2. Re-enabling printer functions previously disabled in trade-offs;
3. When the Center’s 1401 system is operational, demonstrating its 1403 chain printer side-by-side with the System/360 train printer for visitors to experience the difference in speed between the two machines.

As is always the case with mechanical devices, there will be maintenance items to attend to, such as adjusting sticky hammers, tuning of any print quality issues, and repair or replacement of bad hammer coils or other parts that may fail. Critical to the long-term viability of making mid-20th century IT work is nurturing a new generation with the skills of IBM’s Field Engineers and Customer



Mohammed Imran Tommy Lam John Wiseman
Nick Hekman Art Law, Proj. Mgr. Ryan Kulesza
Alena Yampolskaya Emerson Pugh Don Seraphim

Figure 7. The team celebrates success on May 16, 2015

Engineers. Active recruiting and training of mechanically inclined individuals in the “app generation” is a challenge that must be addressed, not just in IBM’s hometown, but in every place, in every culture. This *Making IT Work* conference is an ideal forum for developing and sharing strategies for finding, engaging, and transferring knowledge to the next cadre of computer restorers.

Summary

The printer revitalization project is one of the first revitalization efforts at the TechWorks! Vintage IBM Computing Center to come to fruition. In 2013, having obtained a nearly intact IBM System/360 printer but without its requisite printer controller, the VICC team embarked on a journey to bring the machine back into operation. Crucial to the effort was the teamwork among three generations of volunteers, with experience ranging from the printer’s original designers to student engineers in their last year of undergraduate studies at Binghamton University. The team employed a hybrid approach of mechanical conservation, re-implementation of vintage logic in current technology, and replacing some controller hardware with real-time software. Trade-offs were made to simplify the process, while maintaining basic printing specifications. This solution resulted in a compelling experience that delights TechWorks! visitors. Future work will expand the capabilities of the System/360 printer for live demonstrations, along with activities and displays to place the System/360 printer in the context of printer technology history, the ever-changing relationship between I/O devices and processors, and the human-machine interfaces of yesterday, today, and tomorrow.



Figure 8. Meet the TechWorks! VICC team, from left: Nick Hekman, Fred Petras, Jack Westermann, Bob Lusch, Don Manning, Bill Green, Eric Adler, Jim Ulrich. A full list of VICC team members is found below

Acknowledgments

The authors gratefully thank Susan Sherwood (Executive Director of the Center for Technology & Innovation), Art Law, Tom Schappe, Eric Adler, Bob Arnold, Robert Garner, and Paul Ceruzzi for their assistance in preparing this document and the associated conference presentation.

We also wish to recognize the entire System/360 Printer Revitalization Team (Figure 8) whose talents brought the project to fruition:

Eric Adler, Bob Arnold, Don Blake, Wenna Duan, Peter Haviland, Nick Hekman, Mohammed Imran, Ryan Kulesza, Art Law, Bob Lusch, Don Manning, Tom Schappe, Jim Ulrich, Yuchao Wang, John Wiseman, Alena Yampolskaya and Stanley Yan.

We also acknowledge the technical support to the project from other VICC volunteers and advisors:

The Computer History Museum IBM 1401 Team, Charlie Davis, Gary deBlicke, Will Donzelli, Bill Green, Jack Maynard, Jud McCarthy, Don McCarty, Fred Petras, Donald Seraphim, and Jack Westermann.

Finally, the project could not have succeeded without numerous TechWorks! donors who offered generous financial support, vintage equipment, parts and documentation, and services-in-kind. The Center and VICC team are very thankful for their contributions:

David Adour, Louis Alexander, Rodney Adkins, Buzz Bellefleur, Don Blake, Erich Bloch, Chuck Branscomb, Lewis Branscomb, Robert Bussom, Lawrence Brearley, David Clapp, William Clark, Carl Claunch, Will Donzelli, Endicott Precision, Robert Garner, Geodis, Chuck

Goodwin, Bill Green, Huron Real Estate, i3 Electronics, IBM Almaden Research Lab, IBM Archives, IBM Community Grants Program, IBM Matching Grant Program, IEEE Binghamton Section, Jon Iwata, Scott Jaffe, Kleesh Joshi, Dan Koolish, Tommy H. Lam, Arthur Law, Hochong Lee, Bob Lusch, Don Manning, Jud McCarthy, McCollister's, George Melnyk, Stanley Nemier, Clare Owens Family, David Patch, Family of Frank and Lillian Paul, Fred Petras, Paul H. Pierce, Progressive Tooling, Family of C. T. Reeser, Donald K. Rex, Robert Rosenbloom, Donald Seraphim, Ellen D. and W. Ben Sherwood, James Shurtleff, Sohy Family of Belgium, Eleasa Strauss, Mary Thomas, James Tregaskis, Triple Cities MakerSpace, Jack Westermann, Joseph Wilson III, and Richard Wolff.

References

- Bashe, C., Johnson, L., Palmer, J. and Pugh, E. (1986). *IBM's Early Computers*. Cambridge, MA: MIT Press.
- Campbell-Kelly, M. and Aspray, W. (2004). *Computer: A History of the Information Machine*. 2nd ed. Boulder, CO: Westview Press.
- Garner, R. and Dill, R. (2010). The Legendary IBM 1401 Data Processing System. *IEEE Solid-State Circuits Magazine*, Winter, 2(1), pp. 28-39.
- IBM (1962). *IBM Customer Engineering Reference Manual: 1403 Printer*. Form 225-6493-2 (Major Revision, May 1962). Endicott, NY: International Business Machines Corporation.
- IBM (1964). *IBM Field Engineering Manual of Instruction: 1403 Printers*. Form 225-6492-3 (Major Revision, October 1964). Endicott, NY: International Business Machines Corporation.
- IBM (1965a). *IBM Field Engineering Instruction -- Maintenance: 1403 Printer Model N1*. Form 225-3395-0. Endicott, NY: International Business Machines Corporation.
- IBM (1965b). *IBM Systems Reference Library: IBM 1403 Printer*. Form A24-3073-2 (Major Revision, July 1965). Endicott, NY: International Business Machines Corporation.
- IBM (1965c). *IBM Field Engineering Manual of Instruction: 2821 Integrated Control Unit*. Form 225-3359-1 (Major Revision, December 1965). Endicott, NY: International Business Machines Corporation.
- Manning, D. (2017). *Rebuilding IBM 1403 Print Chains*. [Online PDF] Binghamton, NY: Center for Technology & Innovation. Available at: <http://www.ctandi.org/vicc1403-chain-rebuild> [Accessed 20 July 2017].
- Nickel, T. and Kania, F. (1981). Printer Technology in IBM. *IBM Journal of Research and Development*, 25(5), pp. 755-765.
- Pugh, E. (1995). *Building IBM - Shaping an Industry and Its Technology*. Cambridge, MA: MIT Press.
- Shirriff, K. (n.d.). *Animation of the print chain on the IBM 1403 printer*. [Online] Available at: <http://files.righto.com/ibm1401/printchain.html?m=1> [Accessed 26 March 2017].

The New Public Engagement in Computer Museums and in the History of Computing Machines

Johannes Blobel and Jochen Viehoff

Paderborn University and Heinz Nixdorf MuseumsForum

Abstract

The objectives and complexity of computer conservation projects can vary within a broad range. The Heinz Nixdorf MuseumsForum in Paderborn, Germany, follows different approaches to address the tradeoff between partially opposing restoration goals. On the one hand, a museum's task is to encourage visitors from all ages and backgrounds to understand and reflect on the evolution of information technology from the very beginning to the digital internet era of today; this requires simplification and an entertaining presentation. On the other hand, this simplification must not lead to historical inaccuracy, and a precise reconstruction of historic facts and machines is always an important objective. Also, available manpower, knowledge and budget determine the scope of a restoration project. We present three different projects that were realized by the museum to show the possible tradeoffs between these restoration goals. These include a restoration of a relay-based Turing machine, a rebuild of the first German automatic game machine "Wolf and Sheep", and the award winning build of an interactive ENIAC accumulator.

The Reciprocity between Visitor and Museum

Museums have existed for a long time. Their purpose is to preserve, categorize and exhibit artifacts and knowledge that have a relevance to the culture of a society and are considered to be important enough to be preserved. Since the audience for a museum is typically the general public, museums face the problem of having a very heterogeneous audience with a very wide range of knowledge, interest, time and aspiration. In this section we discuss the challenges and opportunities that are created by this heterogeneity for a computer museum. The expectations and preexisting knowledge that a visitor might have can vary very widely and do not always match the goals of a curator designing an exhibition.

The visitors to the Heinz Nixdorf MuseumsForum range from all ages and all social backgrounds. School classes, tourists, pensioners, families and computer experts visit the museum and expect to learn about the history of computers in an entertaining and accessible way. We therefore cannot make any assumptions about preexisting knowledge and have to ensure that the entry level is very low, so that everyone can benefit from a visit to the museum. But since we don't see ourselves as a purely entertainment leisure facility, we always have to make sure that ease of understanding and an entertaining presentation do not lead to triviality and a distorted view of history. We also do not want to cover the presented topics superficially, but allow educated visitors, who are willing to spend more time on certain artifacts, to gain a deeper understanding of the history and functionality of early computers.

The curators of the museum have their own views of the exhibit and its goals when planning and creating the exhibition. It is important to preserve artifacts and knowledge from the early days of computing up to the present, because computing is one of the most influential cultural domains developed by humans. Understanding the history of this culture and technology enables us to locate ourselves in time and society and to evaluate current discussions and developments in this area.

Since it is a complex task to display the many factors that have influenced history, museums use a wide variety of presentation techniques to show this history from many different angles and to unveil the connections between persons, machines and historic events. The core of each exhibit is the *artifacts* that are carefully chosen from a typically huge set of possible objects. These artifacts can serve as illustrative material and as the entry point to the history that this artifact is a representative of. The actual information can be presented using a wide variety of media — text, images, videos and sounds that all surround and enrich the displayed artifact. A very direct and intuitive presentation technique is an interactive installation, which allows the visitor to touch and use an object. This technique has multiple advantages, because the learning-by-doing approach allows an intuitive understanding. It is much more entertaining than just reading a text and encourages people to learn about the presented topic. Also for families such interactive installations are a good point where different generations to have mutual experiences.

While the presentation techniques for other museum types have evolved over many decades or even centuries, the exhibition of computers and their surrounding techniques is rather new and discussions are ongoing on how to do it best. Software, especially, is a very tricky artifact to preserve and present due to its liminal nature. On the one hand, it is a purely abstract construction like numbers and algorithms; on the other hand, it only becomes a significant thing if it is being executed by real hardware (Dasgupta, 2014). It is also not clear how to preserve and present computers themselves. Some argue that a computer that is not running is nothing but a design object and is therefore not beneficial for understanding the history of computers. There is certainly some truth to it and it is unarguable that the view of a running Colossus, Witch¹ or IBM 1401² is a remarkable experience. A small computer museum in Oldenburg, Germany, takes this principle a step further and allows its visitors to actually use the computers that are displayed (mainly home computers and arcade games from the 70s and 80s), basically using real artifacts as interactive installations (Eddiks, 2017). But of course this has also severe limits and drawbacks. The operation of a computer creates a steady risk of damage due to the limited lifetime of the electrical components. In the worst case, a device can even catch fire due to a short circuit or heating problem which could have disastrous consequences. In 2009 for example the Ferranti Pegasus, which was on display in the Science Museum in London, sustained an electrical fault that “required it to be shut down and Health and Safety considerations subsequently stopped further operation” (Science Museum Blog, 2015). This contradicts a museum’s goal to preserve cultural objects for later generations. Also the additional manpower required for maintenance should not be underestimated.

We can conclude that the requirements for an exhibition come from a wide span of goals and expectations and can even be contradictory. An entertaining presentation and a low entry threshold for visitors, as well as an in-depth presentation of the historic and technical information has to be considered when planning a new restoration project. Also the type of presentation has to be carefully chosen according to different requirements and restrictions. In the following sections we want to outline how the tradeoff between all of these different objectives is handled in the Heinz Nixdorf MuseumsForum.

Different Reconstruction and Restoration Types

When deciding on how history of computers should be presented, different presentation types are possible, as explained in the previous section. In this section we focus on the presentation of computers and hardware, discussing different strategies for reconstruction or restoration.

¹ A rebuild of the 1945 Colossus and the original Harwell Dekatron (WITCH) from 1951 can be seen in operation in the National Museum of Computing in Bletchley Park

² A restored IBM 1401 from the 1960s is in operation at the Computer History Museum in Mountain View

The first question that arises when planning a new project is “What exactly should be reconstructed or restored?” Possible answers are:

- The casing of a machine (non-working original)
- The behavior of a machine (simulation)
- An existing machine (making IT work again)
- A non- or partially existing machine

The first option means that the original hardware is displayed in the museum but is not being operated. This requires little effort if the machine is in an acceptable state. Although if it is already highly corroded, the restoration can also be quite laborious.

Another option is to reproduce the behavior of a computer by means of simulation. For many computers of the first generation simulations exist. Examples are the Zuse Z1 (Mischek, 2013) and the ENIAC (Zoppke, 2013) simulator from the Free University of Berlin, or the EDSAC simulator by Martin Campbell-Kelly (Campbell-Kelly, 2016). These simulations have the advantage that they require little financial expense, as no hardware is required. Additionally they can serve as an interactive presentation, allowing visitors to try out the simulations on their own, thus gaining a deeper understanding. However if a simulation is very close to the original hardware, this is only possible if the user already has a deep understanding of the machine that is being simulated.

An interesting mix of the first two options is to use the casing of an original computer as a pure I/O-device and let the actual functionality be executed by simulation. The BlinkenBone Project, for example, provides a SimH-Simulation of a PDP-machine with either a purely software-based GUI or attachment to a real PDP panel with custom hardware to read in the switches and control the lamps of the panel (Hoppe, 2013). This gives the impression of working with an original historic computer, since the original front-panel is used to control the computer, but it does not have the aforementioned drawbacks of having to restore and maintain old hardware.

The next step in restoration is to take an existing original computer and restore the functionality of the hardware. This can require a lot of effort, depending on the age and condition of the computer and on the available information about it. If successful, however, such a restoration is of course a great achievement and offers a unique insight into the history of early computers. Well known projects include the Harwell Dekatron (WITCH) in the National Museum of Computing, Bletchley Park, and the IBM 1401 in the Computer History Museum in California. In the Heinz Nixdorf MuseumsForum we don't have such large computers in operation, but we do have a fully functioning electromagnetic telephone exchange with a number of telephones connected to it. Watching, and especially hearing, the exchange in operation is fascinating for visitors and encourages them to learn more about it.

The most ambitious project is the reconstruction of a non- or partially existing computer. This requires a huge amount of work, knowledge and resources and is only feasible if all of these resources can be obtained. An artifact that cannot be preserved anymore, can at least be reconstructed in order to keep the knowledge and experience connected with it extant. Of course the Colossus project from Tony Sale and John Harper's Turing-Bombe reconstruction are outstanding projects that used this form of computer reconstruction. The effort for these projects was exceptionally high, but so was the outcome.

The decision as to what kind of reconstruction or restoration should be done is not only determined by the goals and objectives of a project, but to a great extent also by the restrictions imposed on a museum. So, before answering the question of what should be reconstructed and how, the question

of what is available — and therefore possible — has to be answered. As already mentioned before, large reconstruction or restoration of original hardware requires a lot of resources. That is:

- Knowledge
- Components
- Budget and manpower
- Maintenance

The knowledge required to successfully reconstruct or restore a computer includes technical documentation and the knowledge of experts and contemporary witnesses. Even if extensive documentation is available, it has to be read and understood by the people conducting the restoration. It is also beneficial if people who had maintained or built the computer can help with understanding the inner functioning of a machine.

Components for old hardware can also be hard to obtain. These can include electronic valves, relays or transistors. If an original component or a valid replacement cannot be found, one has to make compromises on the historical accuracy and find a suitable replacement. The EDSAC reconstruction project, for example, decided not to use mercury filled delay lines like in the original machine because of expense and operating concerns (The National Museum of Computing, 2010 - present).

Also the required expenses and available manpower limit the extent of a restoration project. The technical staff of a museum is limited and even if external help is available, the budget for paying external experts has to be considered. These financial restrictions determine the scope of a restoration project and can require compromises.

Finally, it is important to consider the maintenance that is required after the restoration is finished. This is even more important than the available resources for building a computer, since running costs can exceed the costs for the reconstruction over time. Robustness is therefore a critical requirement to keep maintenance costs low. This can be achieved by a careful electronic design using electrical surge protection and protective diodes and by modularizing the hardware. Of course this is not always feasible for restoration projects where the hardware design is already predetermined.

Managing the Tradeoff between Opposing Goals

As discussed in the previous sections, different — partly opposing — goals and objectives for computer conservation projects have to be traded. Also the constraints given by available knowledge, components, budget and manpower determine the scope of such a project. We now explain how all these considerations can be weighed and lead to a successful completion of reconstruction projects, as it is done in the Heinz Nixdorf MuseumsForum. Following, we present three examples in order to illustrate our approaches.

It should be clear by now that there is no such thing as *the* solution to the question of how to preserve and exhibit historic computers. In each case several decisions have to be made that determine what should be restored and how it should be done. For our museum there are basically four main approaches, which have proven to be successful in the last decades. These are:

- Exhibit non-functioning artifacts
- Small-scale reconstruction and restoration
- Use multiple media types
- Focus on certain parts and functions of a computer

The first approach should be clear, as it is used by basically every museum and is for many objects a valid solution. In the permanent exhibition area over 2000 objects are shown. Large computers on display include a Zuse Z11 Relay computer and a complete ESER 1055, the Soviet version of the IBM 360 series. The ESER computer, especially, gives a good impression of the dimensions and architecture of such mainframe computers, because the visitors can enter the exhibit containing all the different parts of the computer.

In the following sections we will explain the three remaining approaches on preserving and exhibiting historic computers based on exemplary projects undertaken recently.

Example 1: Hasenjaeger's Register Machine

In 1936 Alan Turing wrote his famous paper "On computable numbers with an application to the Entscheidungsproblem". In this paper he used a novel approach for a proof using a theoretical machine, which would later become known as Universal Turing Machine (UTM). In Germany Gisbert Hasenjaeger was conducting research on mathematical logic at the University of Münster and Bonn, and knew about this work. While Turing's machine was only used as a *gedankenexperiment*, Hasenjaeger thought about how such a machine could actually be implemented. The result was a small Universal Turing Machine with only four states (implemented by 16 telephone relays) and three tapes, which he built in the 1963. After his death in 2006 the family donated his machine to our museum.

The machine itself (Figure 1) was in a rather good state and only needed little work to be made functional again. The tapes, however, were not functional and could only be partially repaired. Also, the exact functioning and the question of what exactly this machine could do was not clear when we got it, since Hasenjaeger never published information about it. Therefore Rainer Glaschick reverse engineered the machine and rebuilt the non-functioning parts of the tapes and finally made the complete machine operational again. The details of this machine can be found in Glaschick's paper (2012).

The small scale of the machine, its uniqueness, and the fascinating underlying questions regarding the universality of the machine encouraged us to get the original machine back into operation. It is an enthralling demonstration of Turing's theoretical concept and we are glad that we had Glaschick's technical expertise to finish this project successfully.




Figure 1: The controlling unit of Hasenjaeger's Turing Machine with 16 Relays

Example 2: The Game of “Wolf and Sheep”

In 1950 a German engineer, Roderich Gräff, built an electronic game using 64 relays to control the machine’s moves. The game implemented is “Wolf and Sheep” (also known as “Fox and Geese”), which is a game for two players and which is played on either the white or the black squares of a chessboard. The first player has four sheep, which are placed on one side of the chessboard. The second player has one wolf, which is placed on a square of his choice on the other edge of the board. In each move, the players can move to a diagonally adjacent square. The sheep can only move forward, while the wolf can move forward and backward. The aim of the wolf is to break through the chain of sheep and reach the other side of the chess-board, while the sheep try to prevent this. If the wolf is in a position where he cannot move anymore, the sheep have won.

The implementation of Mr. Gräff is a wooden box (37 x 47 x 16.5 cm), with a chessboard painted on the upper side (Figure 2, *left*). White and red lamps display the position of the sheep and of the wolf respectively. The positions of the sheep are controlled by the internal logic, implemented by the relays, and the user sets the position of the wolf using cables that are plugged into sockets in the chess-board squares. When taking a closer look, one notices that there are eight columns, but only seven rows. Mr. Gräff stated, that he did not have enough relays back then to control an eighth row. In post-war Germany his only option to obtain relays was to get them at night before closing hour from old telephone relays. Even though this device is clearly not a universal computer, it shows in an interesting way how relays could be used to implement a game’s logic.

After the first inspection of the device and the documentation, it quickly became clear that a restoration to make the device functional again was not practical. The electronics was in a poor shape and the documentation was incomplete. Even though Mr. Gräff is still alive and we were able to ask him about details, it was not possible to reconstruct all details of the implementation. Also the risk of damaging the device when applying power was too high, so we decided on a reconstruction.




Figure 2 Gräffs original version of “Wolf and Sheep” and our working rebuild

The reconstruction was conducted in two stages. First, we build a fully functional replica, using present-day technology inside, but giving it the look-and-feel of the original version (Figure 2, right). This rebuild is used to demonstrate how the game was played. The body is constructed out of plywood that has been treated with a mixture of stain and oil to give it an old look. The chessboard was painted on top with a paintbrush, to get the right texture, and light bulbs were mounted in the holes. The original game has two switches at the top, an ammeter— that does not serve a purpose for the game and was only used for debugging purposes — and a rotary knob to switch to the next move. Similar switches for the reconstruction could be found, but the search for the ammeter and the knob were more tricky. We finally found components that look quite similar, but still did not re-assemble the original in detail. Internally we used an Atmel 8-Bit microcontroller and standard shift-registers to implement the logic and control the light bulbs.

For an interactive installation in a museum, robustness is an absolutely necessary requirement; it should be useable without supervision and is therefore exposed to high stress every day. The reconstructed game, with its light bulbs that can be removed and cables that could be pulled off, is not suitable for a permanent installation in the museum. To still allow the visitors to try out the game, we additionally created a simulation of the game (Figure 3). To create a realistic look for this simulation, we applied a current to each light bulb of the original machine square-by-square and photographed it. From these photographs we then created pixel-precise images for the simulation, for each individual square. Therefore, all the small cracks in the paint and the way the underlying chessboard is illuminated can be seen in the simulation, which gives it a realistic look. While using the original game, there are two sockets for each square (one for the logic, one for illuminating the red lightbulb) and three cables; we reduced the complexity of the simulation to one cable, because it is not clear to the visitor, why multiple cables are necessary. In the exhibit the original game resides in a display cabinet and next to it is a touchscreen with the simulation running. With this setup, we can show the original machine while at the same time allowing the visitor to play the game, using the simulation. Additional information about the inventor and the construction of the machine, as well as an interview with Roderich Gräff can also be obtained via the touchscreen.




Figure 3 Screenshot of the "Wolf and Sheep" simulation

Our approach for this project tries to combine the benefits of different reconstruction techniques. The physical rebuild allows a real hands-on experience during special events, while the simulation can serve as an interactive installation in the permanent exhibit and is also used to display additional information explaining the displayed artifact.

Example 3: ENIAC Accumulator

On the occasion of Ada Lovelace's 200th birthday in 2015 the HNF opened a special exhibit covering the role of women in the history of computing. For this exhibit we wanted to rebuild a part of the ENIAC, one of the world's first computers, as it was mainly operated and programmed by women. The goal for this project was to create an interactive installation that captured the unique look and way of being programmed to let visitors explore how the women back in 1946 worked with that machine (Figure 4).

A requirement for our rebuild was to be simple enough so that all visitors can try it out and configure small programs without supervision, and only with written instructions. But we also wanted to evoke the look-and-feel of the original ENIAC and the difficulties the women had when programming the machine. These requirements lead us to the following design decisions.



Figure 4. ENIAC Rebuild

The complete ENIAC consisted of multiple units, taking up a complete room. For the exhibit we only wanted to rebuild one main unit of the computer: an accumulator. The ENIAC had 20 accumulators, each of which could store a 10-digit decimal number. The accumulator was able to send a number to other accumulators or receive a number from another accumulator. As one accumulator alone is rather useless, we decided to put two of them into one housing and reduce the number of digits to 5 per accumulator. Also a constant transmitter was integrated to get initial numbers into the system. Another simplification was made by omitting the program trays, over which electrical pulses were sent to start the sending and receiving processes, as this is not necessary for the small programs that we had in mind. This design was simplified to only contain the bare essentials of the machine but still be powerful enough to run meaningful programs.

Another requirement was to create a robust installation that could be in operation during the complete opening times of the museum without supervision and with little or no maintenance. A historically correct implementation using electronic valves was therefore not feasible. Instead, we implemented the complete electronics using modern ICs and custom made PCBs to read in the current wiring and display the content of the accumulators. All logic was implemented in software, using an Arduino board with an Atmel microcontroller. To simplify the maintenance in case of hardware defects, all components were strictly modularized.

Beside all these simplifications and robustness requirements, we still wanted to stay as close to the original computer as possible. The displays showing the contents of the accumulators were realized with real neon lamps instead of LEDs, and the control knobs were chosen to match the look of the original ones. Surprisingly, the painting of the housing with black wrinkle paint turned out to be rather difficult, since most paint shops we inquired of were not able to handle this kind of paint anymore. After some failed tests, we finally found a paint shop that could paint the housing, which gave it the finishing touch.

Conclusion

Preserving artifacts and the knowledge from the early history of computers is an important topic for a society where computers have become ubiquitous and that relies heavily on them. However, the question how this preservation should be done is not clear and different approaches exist. They range from displaying non-functioning computers over interactive simulations and multimedia approaches, to large and complex restorations or reconstructions of early computers. In this paper we have argued that all these approaches have benefits and drawbacks and that many objectives and requirements have to be taken into account for planning the scope of a project.

In the Heinz Nixdorf MuseumsForum we endeavour to tradeoff these aspects in order to provide the visitors with an entertaining and informative visit. We presented three projects with different scopes and objectives and explained the decisions that we made during the planning of these projects considering the different objectives and constraints.

References

Campbell-Kelly, M. (2016). *Edsac Simulator*. [online] Available at: <http://www.dcs.warwick.ac.uk/~edsac/>

Dasgupta, S. (2014). *It began with Babbage*. Oxford: Oxford Univ. Press.

Eddiks, T. (2017). *Oldenburger Computer-Museum e.V.* [online] Available at: <http://www.computer-museum-oldenburg.de>

Glaschick, R. (2012). *Alan Turings Wirkung in Münster*. Mitteilungen der Deutschen Mathematiker-Vereinigung, 20(1).

Hoppe, J. (2013). *BlinkenBone*. [online] Available at: <http://www.retrocmp.com/projects/blinkenbone>

Mischek, J. (2013). *Zuse-Z1 | Adder*. [online] Available at: <http://zuse-z1.zib.de/simulations/z1/adders/wgl/>

Science Museum, (2015). *The Pegasus Computer - Science Museum Blog*. [online] Available at: <https://blog.sciencemuseum.org.uk/the-pegasus-computer/>

The National Museum of Computing. *Recreating EDSAC*. [online] Available at: <http://www.tnmoc.org/special-projects/edsac/recreating-edsac>

Zoppke, T. (2013). *ENIAC simulation*. [online] Available at: <http://zuse-z1.zib.de/simulations/eniac/simulation.html>

Capturing, Restoring, and Presenting the Independent Radar Investigation System (IRIS)

Benjamin J. Trethowan

The National Museum of Computing

Abstract

An ongoing project at The National Museum of Computing is to present an early air traffic control system to the public. This paper discusses the importance of capturing an extensive range of information relating to the system at the point of donation. I describe the value of this information within the restoration process, the techniques used in the restoration itself, and the value of expressing the social impact of the system in order to convey its relevance to the public.

Introduction

The Independent Radar Investigation System (IRIS) is an isolated Radar Data Processing (RDP) system, originally used for the investigation of air traffic incidents. It allowed recorded radar and voice (i.e. radiotelephony) data to be replayed together, in synchronisation, for investigative purposes. It was designed and built in the early 1970s at the London Area and Terminal Control Centre (LATCC) in West Drayton, Middlesex, on behalf of National Air Traffic Services (NATS) — the UK's Air Navigation Service Provider for all en-route, and a large proportion of terminal, Air Traffic Services.

The IRIS was in use until the LATCC facility was finally decommissioned in January 2008, when air traffic operations transferred to the New En-Route Centre facility in Swanwick, Hampshire. During its operational lifetime, the LATCC facility handled over 48 million aircraft movements.

Prior to the decommissioning of the LATCC facility, NATS had approached The National Museum of Computing (TNMoC) in order to donate the IRIS. It was felt that the IRIS represented a key element of the history of the UK's transport infrastructure, as well as an excellent example of early British computer engineering. TNMoC accepted the IRIS into its collection in April 2008, following its disassembly, packaging, and transport from the LATCC site in March 2008. Both organisations were keen that TNMoC should restore the IRIS to full working order, such that the system could be demonstrated for the enjoyment of the visiting public.

Between May and September 2008, a small team of volunteers at TNMoC carefully reassembled and restored the IRIS. Restoration activity was finished in October 2008, when the IRIS completed its first successful playback after departing the LATCC facility. The IRIS is shown in Figure 1 during this first successful playback.

Background

The IRIS system is a small, standalone installation of a much larger RDP system, the Processed Radar Display Subsystem (PRDS). The operational requirement of PRDS was to “accept radar plot and label data, flight plan data, etc. and to process and display this data” (Plessey Radar, 1975). In short, the PRDS provided air traffic controllers with their operational situation display, indicating the current position of all aircraft within the selected radar range and system coverage.

The PRDS was not the primary means of radar data processing, but provided a bypass function “intended to provide an alternative radar display service when the Central Computer Complex is unavailable” (Ibid). In brief, the PRDS accepted processed radar data from a primary source



Figure 1. The IRIS

and displayed this to an air traffic controller during normal operations. However, it also provided an alternative means of processing radar data, known as "bypass", should the primary source fail for whatever reason.

The IRIS is composed of a small subset of the PRDS subsystems, representing the smallest and simplest collection of subsystems necessary to generate an operational situation display from recorded radar data. The IRIS is capable of processing a maximum of three radar data channels simultaneously, and displaying a maximum of two of these channels on an operational situation display at any one time.

The PRDS system architecture, and hence the IRIS system architecture, is based upon the use of the Digital Equipment Corporation (DEC) Programmed Data Processor (PDP) 11 range. The IRIS uses four PDP-11 processors, as follows: two PDP-11/34 processors for the Display Equipment Groups (DEGs), used to render vector graphics onto two operational situation displays; one PDP-11/84 processor for the Bypass Equipment Group (BEG), used to process two radar data channels; and one PDP-11/84 processor for the Redundant Bypass Equipment Group (RBEG), used to process one radar data channel and supervise all other systems. These "Equipment Groups" are identified in a detailed view of the IRIS shown in Figure 2.

To achieve playback within the IRIS, recorded radar data is fed into the RBEG and BEG systems for processing. The processed radar data is then passed to the DEG systems for rendering with the vector graphics commands, and then passed to the operational situation display (also known as the Sector Equipment Group, SEG) for drawing onto the screen.

Both the RBEG and BEG systems continually process incoming radar data, regardless of the channel, range, or other selection made at the operational situation display. The DEG systems react to input from the operational situation displays, and are responsible for rendering the display in accordance with the selected channel, dual-channel, range, map overlay, and other user-selectable settings.

Capture

At the time of the original donation, both NATS and TNMoC were keen to ensure that as much supporting material and additional artefacts relating to the IRIS were captured as soon as possible. The LATCC site was due to be demolished shortly after the facility was vacated, so the only opportunity to accept and preserve this supporting material would have to be taken at the earliest

possible stage. The exploitation of this opportunity, and the breadth of the material captured, later proved to be a vital element in achieving the successful restoration of the IRIS.

One of the first issues experienced during the restoration, which was only resolved as a result of the original broad capture, was that of bespoke system design. Despite the fact that the IRIS subsystems utilised a popular computing platform of the period (DEC PDP-11), it transpired that a large amount of the hardware within these platforms was either of a bespoke design, or had been modified. One such example concerned the Radar Data Interface (RDI) hardware contained within the RBEG and BEG systems. Whilst the RDI hardware was modeled on a familiar, full-size UNIBUS card, for obvious reasons the microelectronics on the card were developed specifically for the PRDS by the system's original supplier, Plessey. Initial problems with the RDI hardware, relating to component failure, were only resolved during the restoration process due to the availability of the original Plessey schematics, which had been captured alongside the IRIS at the point of donation. Had this information not been captured when it was, it would have long since been lost, given that the demolition of the LATCC site would have progressed several months by this stage. Considering also that the PRDS implementation was unique, there would have been no alternative source for this information.

A similar issue experienced during the restoration, resolved again by the availability of supporting material, was that of legacy proprietary data communication protocols. The UK Civil Aviation Authority (CAA) had developed a proprietary protocol for the exchange of radar data, based upon a synchronous serial method of communication. The IRIS subsystems were designed to accept radar data only in this format, and were not able to process more recent radar data formats such as the All Purpose Structured Eurocontrol Surveillance Information Exchange (ASTERIX) format. The need to monitor, test, and generate radar data for the IRIS RDP subsystems spawned a further need to first understand the CAA radar data format itself. Fortunately, the original specification for this radar data format, issued by the CAA, had also been captured alongside IRIS at the point of donation.

In addition to the engineering issues mentioned above, the successful restoration was also dependent upon gaining an understanding of the complex operational context within which the IRIS system originally operated. The PRDS subsystems, from which IRIS is built, originally interfaced with a multitude of additional operational systems, performing functions such as aeronautical messaging, code callsign conversion, meteorological data communication, flight plan processing, etc. It was therefore necessary to understand how the IRIS subsystems should operate without the support of these additional systems, and how that might affect the behavior the IRIS would exhibit in its initial



Figure 2. Detailed view of the IRIS

stages of operation following restoration. Again, the initial capture had managed to preserve complete PRDS schematics, including training course material on how the PRDS operated with the additional LATCC systems. Without the availability of this sort of documentation, the successful restoration of the IRIS almost certainly would not have been possible.

Restoration

As well as reaping the benefits of a broad capture at the point of donation, the restoration team also utilised a number of other techniques to combat additional challenges and ensure success.

Unfortunately, it is almost never possible to capture absolutely everything that might be required to facilitate a successful restoration. Setting aside the common issues of storage capacity, transport costs, and underfunding, it is often the case that some supporting artefacts have already been lost prior to the initial offer of donation being made.

For example, the technical documentation associated with a radar data recording and replay system, supplied alongside the IRIS, had already been lost prior to the removal of the IRIS and its associated artefacts in April 2008. This recording and replay system, known as RARE, would later prove to be an essential element in achieving the successful restoration of the IRIS by acting as a source of the CAA format radar data necessary to achieve playback. A number of issues were initially encountered during the restoration of the RARE system, many associated with the total lack of documentation and understanding of how the system should respond in certain situations, or when given certain stimuli. In an attempt to combat these issues, the restoration team contacted the original donor, NATS. The organization made a number of external enquiries and conducted a small amount of research in order to trace a technical contact for the RARE system who might be able to provide the necessary information. After a number of weeks, an individual who had not only worked for the original manufacturer of the RARE system, but who had also been personally involved in its design, contacted TNMoC to offer his assistance. The information that this individual was able to provide, either in hard copy form or through technical interchange meetings, led to the successful restoration of the RARE system, and subsequently the IRIS.

A similar situation was encountered in relation to the IRIS subsystems themselves. In this case, the issue was not that the supporting material had not been captured or lost, but that it simply did not exist. A recurring issue, observed within the power supply units of some of the IRIS subsystems, did not appear to have been documented in any of the material available to the restoration team. After extensive investigations, the restoration team contacted the original donor, NATS, in order to ascertain whether any similar issues had been observed whilst the system was in service at the LATCC facility. Following discussions with NATS, it transpired that the issue we were observing was, in fact, a known one — the result of incorrect technical documentation. Whilst many of the existing NATS engineers were aware of this issue and its cause, this information was not passed on to TNMoC at the point of donation as it had, quite simply, been forgotten about. The support of the original donor and existing third-party suppliers proved to be vital in these circumstances.

Challenges

Whilst many of the challenges encountered during the restoration process were resolved through the use of supporting material, or through contact with the original donor and its third-party suppliers, some could only be resolved through the successful application of engineering best-practice.

One such challenge was that of degraded magnetic tape media, a common feature of many computer restoration projects. An additional radar data recording and replay system, known as the SE7000, had been supplied alongside the IRIS, accompanied by a number of radar data recordings on multi-track magnetic tape. The SE7000 is essentially a multi-track magnetic tape deck, with multiple

read-write heads that are used to record or replay multiple radar channels simultaneously. When inspecting the magnetic media that we had been given, it became apparent that the media was suffering from the common problem of oxide separation. The restoration team was fortunate enough to have been given some additional media, however, which did not exhibit the separation issue. Thankfully, this has meant that we have been able to stop short of attempting to restore the degraded media, using a potentially hazardous process, for the time being.

Similar degradation issues were found to have affected the drive mechanism itself, with numerous rubber capstans, drive belts, etc. having perished. Contemporary replacement parts were found to resolve these issues however, allowing us to operate the SE7000 unit temporarily, in order to capture the radar data from any operable media onto a more suitable modern format.

Interestingly, the more recent RARE system suffered from similar issues, despite its use of closed tape cartridges (Exabyte) instead of open tape reels. A series of recurring issues were encountered with the Exabyte drive mechanisms, which had not been maintained since their withdrawal from service many years before the point of donation. A lack of suitable spares further exacerbated this problem, and led to the decision to immediately retire the RARE unit following the successful capture of the radar data from any operable media onto a more suitable modern format.

Social Impact

After overcoming the many challenges of restoration, and successfully demonstrating the IRIS in full working order since October 2008, it is perhaps worth considering the original motivations behind this project.

The IRIS offers us a unique opportunity to convey the impact that computing technology has had upon society, demonstrating why the history and development of computing is relevant to our visiting public. Our visitors have the opportunity to both study and interact with something that was a vital component of the UK's transport infrastructure for over 25 years — something that they can immediately identify with if they have ever flown away on holiday, or to a business meeting, etc.

It also delivers a significant visual impact as an operational exhibit: its size and "retro" appearance provide a visually captivating display, far removed from the so-called "beige boxes" that many of our visitors expect to see. This visual appeal can be seen in the publicity photograph of IRIS shown in Figure 3.



Figure 3. Publicity photograph of IRIS

Benjamin J. Trethowan

And finally, it presents us with an additional educational opportunity to captivate the many young people that visit our museum, and explain just how the development of computing technology has kept, and continues to keep, us safe in the sky.

Reference

Plessey Radar. (1975). *Processed Radar Display Sub-system for the Civil Aviation Authority, Vol. 1*. Weybridge, Surrey: Plessey Radar, RSL 1631 (CRSS/A1/01).

Constructing and Testing a Replica Store for the EDSAC

Peter F. Linington

University of Kent School of Computing Canterbury, Kent, UK

Abstract

The original mercury delay line store of the EDSAC is described, and the reasons why it would be difficult to make an exact working replica of it are set out. Instead, a technology dating from three or four years later, the wire delay line, is used and this design will be explained, as will the determination of the parameters needed for its construction.

As with any non-trivial technology, a full range of tests is needed and the key tests used for setup and acceptance testing of the individual delay units will be described. Finally, the tests to be used in sub-system integration and for monitoring correct operation in the running machine will be outlined.

Introduction

The developers of the EDSAC needed to minimize its complexity to keep both cost and development time within reasonable limits. Both of these depend critically on the number of subsystems involved, and nowhere is this more evident than in the machine's store. With the technology available at the time, the non-volatile storage of a single bit involved about five valves (vacuum tubes). To apply this to even the 2.2 kilobytes projected for the machine would require more than 90,000 valves, compared with only about 3,000 in the machine as actually produced. It is clear that a store technology was needed in which hundreds of bits would be stored in a single element.

Even at a rate of one separate unit per machine word, the machine would be unwieldy; a storage component was needed that could cope with many words in one unit. This led inexorably to the adoption of a serial architecture, in this case one centred on a delay line store in which hundreds of bits were transiting each delay at any one time; bits were detected as they emerged from the delay line, reclocked and restored to a clean standard pulse shape that could then be reinserted to form a continuously recirculating loop of pulses. Arithmetic was also performed a bit at a time on recirculating bit streams. An overview of the machine can be found in Wilkes and Renwick (1949).

The use of long recirculating loops for storage introduced some extra complexity because the control logic had to wait until the instructions and data needed travelled round the loop to the point where the pulses were being regenerated. However, the logic needed to do this was more than offset by the lower cost per bit incurred in a longer loop. Even so, about a third of the machine's bulk was taken up in addressing and interfacing to the 32 delay lines and nine shorter register stores.

Mercury Delay Lines

Wilkes knew that acoustic delay lines using mercury had been used to correlate successive radar echoes to overcome jamming during the Second World War. He was given the drawings of such a delay by Tommy Gold from the Cavendish Laboratory (Gold's later design for a folded version can be found in Gold (1951). From these Wilkes developed the EDSAC design, consisting of two batteries, each consisting of sixteen 1.67 metre long tanks (see Wilkes and Renwick (1948) and Figure 1, which is taken from it).





Figure 1. Photograph and drawings of the long mercury tanks

Pulses travelling through the liquid mercury were modulated on a 13.5 MHz radio frequency carrier and were introduced and subsequently detected using quartz crystal transducers matched to the carrier frequency. The crystals were about 15 mm in diameter and therefore formed a tight acoustic beam within the mercury so that there were not significant reflections from the walls of the containing tubes, and, as a consequence, no significant echoes or multi-path effects. Renwick's notes (Renwick, n.d.) estimated that the total power loss would be about 68 dB, of which 15 dB would result from attenuation in the mercury, the bulk being loss in the transducers. One of the problems with these transducers involves acoustic impedance matching. To give effective energy transfer into the mercury column, the mercury must wet the quartz surface without any intervening air film. On the other hand, there must be an air gap behind the crystal or about half of the energy will be lost into its mounting.

As a result, the turning of the mounting components had to be done in an oil-free environment, to prevent traces of oil bridging the gap behind the crystal. The crystals were then fixed in a massive steel backing plate and the whole assembly ground flat to a ten thousandth of an inch so that the sending and receiving crystals were parallel to within one acoustic wavelength. The delay tubes were first filled with alcohol to wet the crystals and then the mercury added, displacing the alcohol. Later, glycol was used when filling, in place of alcohol. If an air film subsequently developed, the delay unit failed and had to be flushed and refilled. Filling was a delicate operation because the liquid mercury could slosh in the tubes if it was added too rapidly, shattering the fragile crystals at each end.

Mercury dissolves the more reactive metals, such as zinc, so all the components in contact with it must be made of steel. Over time, it leaches sulphur from the rubber sealing washers and has to be

purified. A high standard of purity is needed if all the elements of a battery are to exhibit the same delay and Wilkes had to have his stock of mercury double-distilled to achieve this (Wilkes, 1985).

Perhaps the main operational problem with mercury delays is their sensitivity to changes in temperature. Mercury has traditionally been used in thermometers just because of the rapid change in its density with temperature, and there is a corresponding change in the velocity of sound, and so in the delay. If the change of delay is large enough that the recovered pulse no longer overlaps with the clock, the regeneration process will fail. In the original design, reclocking was done with a simple *AND* gate and temperature variation had to be less than $\pm 0.5^{\circ}\text{C}$. (In 1951, a more tolerant reclocking design that extended early pulses increased this to $\pm 1^{\circ}\text{C}$.)

If the temperature changed by more than this, the machine clock frequency had to be adjusted to match the actual delay. Each store battery was placed in a light wooden enclosure, called a coffin, that restricted air movements and so limited sudden temperature changes. In 1951, the store was moved into a thermostatically controlled oven that maintained it at a steadier temperature (see Figure 2). However, this was after the target date for the replica.

The combination of likely operational problems, together with issues of cost, and health and safety considerations led to a decision not to use mercury as a storage medium in the museum replica.

Wire Delay Lines

In place of mercury, it was decided to use metal wire delay lines. These began to replace mercury delays within the lifetime of the EDSAC, although they were never used in Cambridge. Their use in Elliott Brothers Research Laboratories was announced in Millership et al. (1951). A complete unit produced by Ferranti is described in Fairclough (1956); Scarrott and Naylor (1956) give an extensive presentation of the design of a larger store, together with the theory underpinning it.

The transducers used in these delay lines exploit the magnetostrictive effect shown by some metals, such as nickel. If a magnetic field is applied to them, their shape changes. In the inverse effect, straining the material induces a magnetic field change. Thus if a current pulse is applied to a coil wound round a nickel wire, the momentary stress produced results in a pair of acoustic pulses travelling in opposite directions along the wire away from the coil. If a pulse then passes through another coil in the presence of a background field, the associated strain modifies the field and induces a current in that coil.



Figure 2. The store oven, introduced in 1951

A short delay line can be constructed using these longitudinal pulses alone, but better performance can be achieved using torsional, or twisting, pulses. These travel more slowly and spread out less as they travel along the wire, allowing longer delays in a more compact format. Longitudinal pulses can be converted to torsional pulses by a mechanical mode transformer (see Figure 3). The input to this is from two metal tapes, each driven by its own coil; the two are connected in anti-phase so that one tape is compressed while the other is stretched. In the mode transformer, these tapes are welded to the opposite sides of a round wire so as to impart a twist to it, launching a torsional pulse. At the far end of the wire another mode transformer turns the torsional pulse into opposed movement of two more tapes, which is then detected by a differential pair of receiving coils.

Figure 4 shows the store from a Marshall HS-100 multi-channel analyser (Marshall, 1958), which had a similar specification to the EDSAC store units. Two delays are interleaved to save space. The sending coils and tapes are at the left and right sides and the sending mode transformers are in the top left and bottom right corners of the unit. From them, the main wires run in a spiral to the inner mode transformers and their tapes lead to receiving coils nearer the centre of the mounting plate. In each case the free ends of the tapes are terminated in absorbent buffers (the silver rectangular element in the figure) which dissipate the pulses passing from the coils away from the mode transformers.




Figure 3. Longitudinal pulses generated by two coils are transformed into torsional pulses. The bias magnet is also shown




Figure 4. Long wire delay line (mid 1950s)

Preserving the Interface

One of the principles agreed on in approaching the replica project was that the main interfaces should be preserved exactly as they were. In this way we have the freedom to revisit any pragmatic variations as and when resources allow, perhaps experimenting later with mercury delays.

In this case, the key interface is the radio frequency feed from the chassis controlling regeneration to the delay unit itself. Care has been taken to ensure that the physical connectors, electrical impedances and signal levels and formats all correspond to our knowledge of the original. No extra signals are provided to the storage units (except for power to the circuits inside the store needed to support the interface). In particular, the store is completely asynchronous without any explicit access to the main machine clock.

Recreating the Store

Although the principles underlying wire delay lines are well documented and understood, many of the fabrication details have been lost and a series of experimental investigations was needed to find the best solutions. First, it was necessary to determine the optimum method of preparation for the nickel tapes to be used in the transducers. Scarrott and Naylor (1956) indicate tapes 0.5 mm wide and 0.125 mm thick would be suitable, and these were prepared by rolling-out commercially available nickel wire. However, the resulting tapes were significantly work hardened and showed weak magnetostriction. Experiments showed that annealing for ten minutes at a bright red heat raised the sensitivity to a plateau. The tapes were heated simply by passing a current through them.

The next parameter was the optimum level of bias field to apply. Experiments showed that the sensitivity increase as the bias was decreased, but that reproducibility required sufficient field to overcome the hysteresis in the metal. In the units produced, there is an offset of 1 cm from a cylindrical AlNiCo 5 magnet with diameter 3 mm and length 12 mm, having a remanence of 1.24 T.

Many experiments were needed to find a satisfactory coil configuration. The various historical descriptions suggest a range of sizes but tests showed that, if the coils were too large, there was a problem with inductive ringing that distorted the pulses in transit. This is probably a result of the use of lower impedance amplifiers than those available in 1950. In the replica, the coils have been made as small as possible consistent with recovering an adequate signal. The coils are wound in 44 SWG enamelled copper wire on a 0.7 mm steel wire former. The opposed pair is made in a single operation, supporting the coils on 0.25 mm thick acetate sheet tab. They are then potted with epoxy and the wires withdrawn. The tab is attached to a printed plastic carriage that also carries the bias magnet and a small connecting socket. This whole assembly is then mounted to slide on brass rails so that the delay length can be adjusted.

It is important in building something as complex as these transducers that the design should be checked by direct observation at each stage wherever possible. A travelling probe was constructed that could position a ten turn, 0.5 mm diameter test coil inside the bore of the transducers to monitor the field actually produced. This confirmed that the field distribution matched the design calculations and that there were opposed induction peaks at the start and end of each pulse. It should be noted that the measures normally used to suppress back-EMF when driving coils, such as using catcher diodes, are counter-productive here because they encourage persistence of a local current loop that extends the coil field and distorts the pulses. Another probe, but with a nickel wire core, allowed monitoring of the field in the metal tapes, as a check of correct assembly.

The main delay wire was made of the nickel steel alloy NiSpan C902, which is commercially available and similar to the formulation given in Scarrott and Naylor (1956). Experiments with heat treatment after drawing showed that heat treatment at 450°C for five hours gave a low

temperature coefficient of delay, being about just 2 per cent that for mercury. The torsion wire used to store 576 bits is 2963 mm long. One practical problem is that the wire needs to be straight before installation if it is to form a regular spiral and show minimum pulse dispersion. This can be achieved by heating it briefly by passing a current while under tension so that it just begins to soften, and then cutting to length, grinding the ends flat and installing the wires immediately.

The next step is to spot weld the longitudinal tapes to the torsional wire. This was carried out using a simple capacitor discharge welder that delivered about 4 Joules when charged to 30 V. The optimum working point needs to be determined empirically so as to provide a firm weld without damaging the tapes.

It is very important that the weld be made right at the end of the tapes and torsion wire. If there is any overhang, the extra material acts as local load, leading to extended ringing at its resonant frequency, often lasting several microseconds. Even as little as a 0.25 mm positioning error can cause problems. However, excess material can be removed by careful grinding.

In addition to launching pulses into the mode transformer, the coils also produce pulses that move away in the opposite direction. To absorb these, the free ends of the nickel tapes are clamped in a buffer assembly, between two neoprene sheets. These are angled so that they apply a gradually increasing pressure to the tapes, absorbing energy without creating sharp echoes. At the back of this unit the tapes are further damped with a short piece of PVC adhesive tape and then grounded to reduce electrical noise.

The main challenge in setting up a delay line is to ensure that there are no sources of echoes in the system. Echo pulses generate errors, and the echoes on paths with different lengths can combine to produce strong pattern sensitivity. The main sources of echoes are poorly adjusted buffers at the tape ends and irregular welds in the mode transformers. The cause of an echo can be determined by moving the transducers and observing the relative change in timing of the echoes with respect to the main signal. This indicates which segment of the line contains the echo source, which can then be localized by using the time offset, together with the known speed of propagation.

Once the store unit has been attached to the regeneration chassis the delay length is adjusted by observing the result of ANDing the delayed pulses with the machine clock (see Figure 5), sliding the transducers as necessary to trim the length and so give a clean output.




Figure 5. Adjusting the delay length for correct regeneration

Matching the Interface

The replica store uses technology that is quite different from the original mercury tanks. It preserves the basic concept of using an acoustic delay based on the signal path length, but the signals are base-band, as opposed to the RF modulated signals in the original. The transducer coils introduce two stages of transformer coupling, so that an incoming pulse emerges in a differentiated form as a central pulse flanked by smaller, negative going pulses. The shape of the original pulse needs to be restored, requiring detection and reshaping. Typical pulses are shown in Figure 6.

Constructing and Testing a Replica Store for the EDSAC

In consequence each delay line needs a significant amount of signal processing. There needs to be a detector for the RF pulses, a pulse shaper and coil driver, a preamplifier for the signal from the receiving coils, a detector, a resaper and finally an RF modulator to restore the pulse to the form expected at the reference interface. This conversion circuitry is comparable in terms of complexity to the regeneration logic in the main machine, which makes up about 30 per cent of its bulk. It was therefore unrealistic to provide these functions in a 1950s style. They are provided using modern technology, which is housed within the store unit enclosures (Figure 7).




Figure 6. Interface and transducer signals. The red trace in (a) and (b) shows the original test pulse at the store input bus. The time scales are 500 ns/div, 500 ns/div, 1 μ s/div and 200 ns/div respectively




Figure 7. The long delay line unit used in the replica

Testing the Store

A subsystem like the EDSAC store needs a number of test tools with different scopes and objectives, ranging from component testing to in-service verification. These need to be applied on timescales from immediate functional testing to extended soak testing with a wide range of data patterns and timing relationships.

It is difficult to determine in advance exactly what patterns of behaviour need to be tested, so it is an advantage for the test environment to be as soft as possible, allowing flexibility for dynamic reconfiguration. Ideally a test environment should handle a large number of signals with complete flexibility with regards to timing and amplitude of each. However, it would be expensive to approach this ideal, so the environment used to test the store makes some compromises. It uses a digital test box based on a Field Programmable Gate Array (FPGA) in combination with a mass-market digital oscilloscope that can be triggered by the test box to gather detail of analogue levels in error situations.

A valve-based system which is under active development must be regarded as a hostile environment; accidents involving a 250V HT rail can have an adverse effect on many test instruments. The test system described here is therefore fitted with optical isolation on all its incoming and outgoing signal lines, protecting both the tester and the associated control and data-logging computer. This adds about a hundred nanoseconds to the various test loops, and this must be allowed for in test designs.

The advantage of using an FPGA is that it gives good time flexibility in signal generation and data capture, without the synchronization and scanning loop delays associated with microprocessor-based tools. This tester handles eight inputs and eight outputs with a 15 nS time granularity and full parallelism. It is based on a Xilinx Spartan 6 development card.

The most basic form of testing is that needed for initial commissioning of the store units and consists of a range of features for pattern generation. Starting with a series of widely spaced single pulses, the path through the unit is traced with the oscilloscope, setting levels and pulse widths at each stage. More complex patterns are then introduced to check for coupling problems that might manifest themselves in the inability to sustain longer sequences of pulses.

The next stage is to check the correct operation of the delay line, ensuring that it preserves the pattern of pulses injected into it. To do this, the tester includes three banks of store of size equal to the target delay time. New test patterns are assembled in the first of these, and, on a command from the user, the pattern is fed into the delay and simultaneously copied bit-by-bit into the second bank. Bits emerging from the delay are compared with the content of the second bank to check that the pattern has been preserved. If any difference is seen, the corresponding bit in the third bank is set and an interrupt raised to cause error reports to be sent to the user interface (Figure 8).

Once the basic store path has been demonstrated, testing is normally applied to the store unit and regeneration chassis in combination, controlling the store via the command signals usually sent from the computer's central control. Different test operations then allow either open-loop testing, in which fresh data continues to enter the store and be checked on exit, or closed-loop testing in which the regeneration and storage of an injected pattern is checked.

At this stage, the data pattern may be static, or it may change systematically so that, for example, different length pulse chains are tried in sequence. Finally, random data patterns may be generated, each being exercised for a short period before moving on to the next.

One of the main problems observed in practice is electrical breakthrough between the RF transmitter and receiver circuits. During regeneration, the bit entering the delay line is generally the same as that which has just left it. They differ only if new data is being written. At other times interference between them has little effect. A more demanding test is therefore one in which a continuously changing bit stream is provided, with each pattern transiting the delay only once. A free-running random generator is provided for this test and provides the most stringent of the family of tests for acceptance purposes.

So far, we have dealt with the testing of single store units, managing the process via the control tags on the regeneration chassis. However, the same testing approach can be applied to the immediately enclosing subsystem, and work is currently in progress to do so. This involves accessing a complete memory rack, consisting of eight store units and the associated addressing and data multiplexing elements. The core of the tester is again a block of RAM, but this time the expected data is inferred by observing the read and write operations from the control logic and checking that data read is consistent with the store's read and write history. Intercepting this interface allows two modes of operation. First, known patterns can be injected and retrieved as in the current tests, but working across more of the address space. Alternatively, however, the tester can be run in an entirely passive mode, tracking the operation of running programs in the machine to see that the store is fulfilling its contract.

Once this is achieved we will have a tool that signals operational errors, indicating whether transient errors originate in the store or in the remainder of the machine. From there, it is but a small step to use the same data streams to check the consistency of operations in the main control and the arithmetic functions.

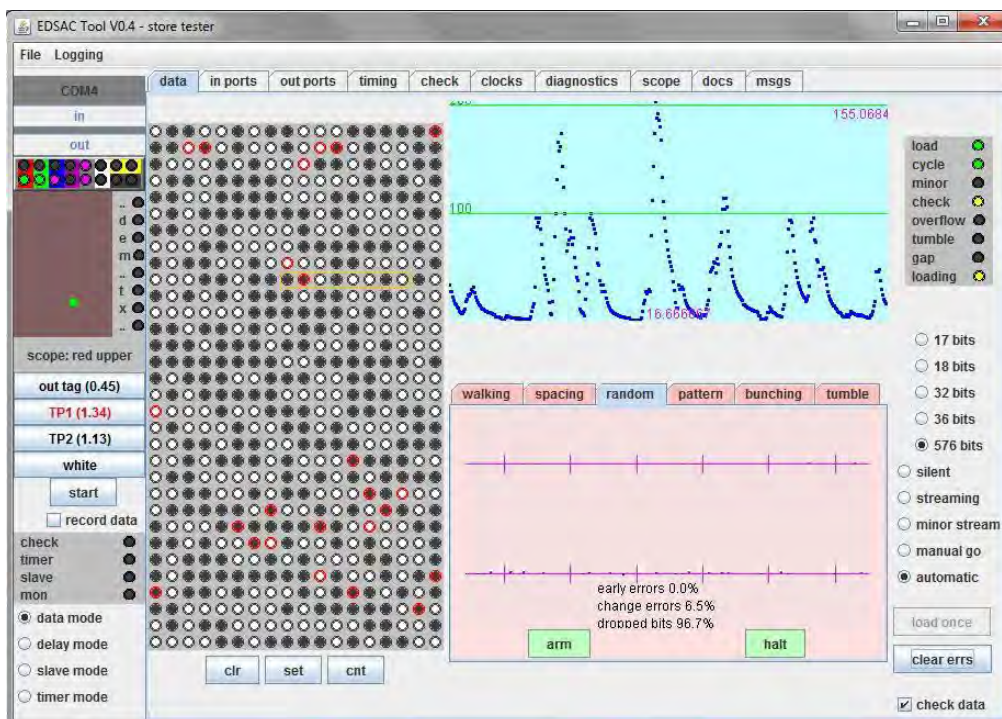


Figure 8. User interface to store tester

Summary

This paper has outlined the technological choices to be made in creating an operational replica of the EDSAC that is capable of being run regularly as a museum exhibit. It has outlined the importance of supporting experiment in determining many details of the technologies not recorded in contemporary literature.

It has also indicated the importance of supporting the effort with a range of testing tools, with scopes varying from the individual component elements to complete subsystems and how these can be configured to support different kinds of problem solving.

References

Fairclough, J. (1956). A sonic delay-line storage unit for a digital computer. *Journal IEE Part B: Radio and Electronic Engineering* 103(3S), pp. 491–496.

Gold, T. (1951). A design of an ultrasonic delay-line. *Philosophical Magazine* 42(330), pp. 787–791.

Marshall (1958). The Marshall Pulse Height Analyser HS 100. *Nature* 182, pp. 989–990.

Millership, R., Robbins, R. C. and Debarr, A. E. (1951). Magnetostriction storage systems for a high speed digital computer. *British Journal of Applied Physics* 2, pp. 304.

Renwick, W. (n.d.), Supersonics. William Renwick Papers, Cambridge University Archives, CSAC 84/2/82.

Scarrott, G. G. and Naylor, R. (1956). Wire-type acoustic delay lines for digital storage. *Proc. IEE*, pp. 497–508.

Wilkes, M. V. (1985). *Memoirs of a Computer Pioneer*. Cambridge MA: MIT Press.

Wilkes, M. V. and Renwick, W. (July 1948). An Ultrasonic Memory Unit for the EDSAC. *Electronic Engineering* 20, pp. 208-213.

Wilkes, M. V. and Renwick, W. (1949). The EDSAC — an electronic calculating machine. *J. Sci. Instr.* 26, pp. 385-391.

Report of the TNMoC Workshop, 23 May 2017

Nicholas Enticknap

Information Technology Journalist

Overview

Day One of the “Making IT work” conference was fundamentally academic. The various speakers analysed why we should conserve, restore or replicate ancient computers, and discussed how to deal with the practical problems that arise when we attempt to do so.

Day Two put some flesh on these bones. The programme consisted of a series of presentations in the morning describing the progress of five specific conservation projects (Figure 1), followed by live demonstrations of the artefacts themselves in the afternoon.

The five projects chosen were an eclectic bunch; the one connecting theme is that all are UK conservation projects whose hardware today can be seen at the National Museum of Computing in Bletchley Park.

Proceeding chronologically, the first was Colossus, whose place in history as the first computer to do useful work anywhere in the world is now assured. Then there were two other first generation valve-based machines — Cambridge University’s justly celebrated EDSAC and the Atomic Energy Research Establishment’s less widely known and more primitive Dekatron. Then we jumped a couple of decades. The fourth working exhibit was a special-purpose configuration, the IRIS air traffic control system. Finally, there was a machine which represents the acme of classic British mainframe design, the ICL 2966.



Figure 1. Phil Hayes explains the Colossus rebuild project to the workshop

Colossus

Colossus (Figure 2) exemplifies a major problem with many ancient computers; at the time the rebuild project started, none of the original machines still existed and minimal documentation had survived. To compound the difficulty, the whole subject of Colossus and its use was shrouded in official secrecy. Still today, over 70 years after the end of World War 2, the people who were involved with the machine at Bletchley Park are reluctant to reveal anything of what they knew.

There were in fact 10 different Colossi of two types. The Mark I was the prototype and the Mark II which followed was five times faster. No two of the 10 were the same. When the war ended eight were dismantled but the two most advanced, numbers 9 and 10, were moved elsewhere in GCHQ and were modified before use on other tasks up until 1955, eventually being destroyed along with all their documentation a few years later.

At this stage Colossus was still covered by the Official Secrets Act and the public knew nothing about it; all histories of computing started with ENIAC over in the US. Then in 1970 the first public information about the machine surfaced in an article by Jack Good, a wartime codebreaker. Brian Randell was one of those to take fire and he started pestering the Government for information, but enjoyed little success; he eventually secured six official photographs but no description of how the machine actually worked. Undeterred, Randell put together what he had learnt from those involved with Colossus, such as “Doc” Coombs and Donald Michie, and published his findings. These matters rested till the 1990s.

Enter Tony Sale, then employed by London’s Science Museum but with a CV that included a spell working at GCHQ from 1957 to 1963 with spycatcher Peter Wright. Tony had been instrumental with Doron Swade in setting up the CCS; he then turned his attention to Bletchley Park and played an energetic part in the formation of the Bletchley Park Trust in 1992. The following year he was ready to make a proposal to GCHQ for the construction of a replica Colossus.

GCHQ was still very reluctant to reveal any information, but allowed the project to go ahead on various stringent conditions, such as the replica could only be a Mark I (the most primitive version), that public access to it could only be via a viewing gallery, and that Tony could not talk publicly about how the machine worked.

Nonetheless, Tony did extract a technical report describing the machine from Don Horwood, again protected by official secrecy — it remained classified till 2003. He also contacted everyone surviving from the period of live operation, including major figures such as Tommy Flowers (the principal designer of Mark I), and Harry Fensom (a key member of Flowers’ team). This allowed him to proceed with reasonable confidence, and by 1996 all was ready for the official inauguration of his replica Colossus by HRH Duke of Kent.

Tony Sale died in 2011 and the project is now run by Phil Hayes, who provided the conference attendees with the information above and demonstrated the machine.

The demonstration involved running a paper tape through a suitably colossal paper tape reader — Hayes says it is the fastest ever produced, operating at 5000 characters per second. The tape contains intercepted cipher text and is formed into a loop which circulates continuously. Colossus performs calculations on the data (indicated by flashing lights) with a view to establishing the start positions of the rotors used by the Germans on the Lorenz machine which created the enciphered text.

There were 12 of these rotor wheels on a Lorenz machine. Initially Colossus was used to determine the starting positions of just five of them, which was enough to make progress, but by the end of the



Figure 2. The Colossus rebuild, with Margaret Sale and Brian Randell in the foreground

war the analysis technique had been refined to allow estimation of the positions of all 12. A single Colossus run at Bletchley Park today takes typically three to four hours.

The result of a run shows what Colossus has calculated to be the most likely starting positions of the rotors. Wheels are then set to those positions on a replica Tunny machine — a logical equivalent of the Lorenz — which, assuming the Colossus calculation has produced the right answer, performs the actual decoding of the message and produces German text output.

EDSAC

EDSAC scarcely needs an introduction (Figure 3). It was one of three influential pioneering computers developed at British academic and research institutions in the 1940s along with the Mark I at Manchester University and the Pilot Ace at the National Physical Laboratory. EDSAC was designed to provide a computational service to the various departments of Cambridge University.

It was also chosen by J Lyons as the basis of the LEO, the world's first computer designed for business data processing, though substantial modification was needed to the I/O arrangements before LEO could perform useful work. Scientific computation of the type EDSAC was designed for typically involves a small amount of data being used in lengthy computations, whereas business data processing involves large amounts of data with minimal computation. So whereas EDSAC started work in 1949 and provided a regular computing service till 1958, LEO did not process its first program till 1951 and did not provide a regular service till 1954.

As with Colossus, the original EDSAC no longer exists, but unlike Colossus there had never been any secrecy about EDSAC and there was a huge amount of published material about it. Building a replica was therefore in theory at least a relatively straightforward exercise. The focus of attention in the presentations at the conference was on the issues presented by the storage system.

EDSAC was built using 3000 thermionic valves. When used to store data, each bit required five valves, so it was clearly necessary to have an external storage system if any useful work was to be done. This challenge confronted all the first generation computer designers. Different technologies were tried in different machines (the ultimate solution to this problem proved to be the ferrite core

store, which became standard in computer systems from the fifties to the seventies, but this was still an experimental technology in the forties and its eventual triumph was not then foreseen).

The EDSAC design team, led by the late Sir Maurice Wilkes (1913-2010), chose mercury delay lines, a technology that had become well understood during World War 2 in radar applications. They built a memory comprising 32 delay lines each capable of storing 576 bits, yielding a total capacity of roughly 18 Kbits. They added nine shorter delay lines, of 18 and 36 bit capacities, to hold working data.

When it came to building a replica, however, snags quickly became apparent. Use of replica mercury delay lines would clearly be more authentic than choosing any alternative, and interfacing them to the rest of the system would be easier. Those considerations however, said presenter Peter Linington, were far outweighed by negative factors, and it quickly became clear to the project team that another solution would have to be adopted.

Negative factors included the extreme temperature sensitivity of mercury (it needs to be kept within a range of $\pm 0.5^{\circ}\text{C}$); the consequent need to keep an oven running constantly; the destructiveness of mercury (which dissolves most metals, forming an amalgam which inhibits proper operation); the unreliability of the surface contacts; the gradual contamination of the mercury in use; its cost; and its weight. These factors together make attaining the necessary degree of reliability well nigh impossible today (Linington told us the store has to preserve circulating patterns of data with an error rate of less than 1 in 10^{10} if any program beyond the non-trivial is to run). In addition, there is modern health and safety legislation which has seen mercury virtually disappear from industrial operations.

The EDSAC replica design team accordingly decided to use wire delay lines instead. This technology emerged in 1951 and was subsequently widely used, but had not been available when the original EDSAC was being designed.



Figure 3. On the dais, Andrew Herbert, project manager of the EDSAC replica, describes the system. Peter Linington (rightmost), designed the nickel-wire-based delay line memory.

It was necessary to keep the original RF interface between the store and the rest of the replica EDSAC, said Linington, both for authenticity and for ease of construction. This was more difficult with wire delay lines than with the mercury delay lines originally used. Implementing the logic using 1940s technology would have added 30 per cent to the overall size of the memory, which would have violated visual authenticity. So modern technology has been used to re-create this logic instead.

The resulting nickel delay lines with their RF components have been housed in two boxes known from their shape as coffins. Each contains 16 individual memories. No photos survive of the coffins used with the original EDSAC, but floor plans have survived and the new coffins occupy exactly the same dimensions as those on the plans. So the wire delay store designed by the EDSAC replica team looks like the store used on the original EDSAC and behaves in the same way in terms of the signals it sends and receives, even if under the covers the technology is not the same. Linington describes this whole exercise as one of “the need to balance authenticity against reliable operation in a museum environment”. If you take too purist a view you will never have anything you can demonstrate.

It is nonetheless worth noting that the mercury delay lines used with the original EDSAC continued to operate satisfactorily — though with a heavy maintenance burden — throughout the operational lifetime of the machine. Despite their disadvantages there were always other projects at Cambridge that were more demanding of attention.

Dekatron

The Harwell Dekatron differs from the two computers discussed above in that it did not need to be re-created; it has been in existence since it was first built and is indeed described by its current curator, Kevin Murrell, as “the oldest still working operational computer” (Figure 4).

Murrell started his story of the machine back in World War II. At that time there was a free interchange of atomic energy research information between the UK and the US, but when the war



Figure 4. Visitors from Germany (Johannes Blobel and Jochen Viehoff) and Italy (Elisabetta Mori) in front of the Harwell Dekatron Computer.

ended that exchange stopped too. This drove the UK Government to set up a centre of expertise of its own. This was the Atomic Energy Research Establishment (AERE), established on the site of a former RAF air base at Harwell.

There a team led by Ted Cooke-Yarbrough and Dick Barnes conceived the need for a computer, went to Cambridge to inspect the nascent EDSAC, put up a proposal to their bosses (who included John Cockcroft and Klaus Fuchs — but what the Kremlin made of it we shall never know) and set to work to build the Dekatron. Construction started early in 1950, and the machine was complete by April 1951.

Physically, the computing part of the machine used electromechanical relays. For storage, the Dekatron was equipped with eight groups of memories (plus one spare) built from dekatron valves, plus a two-part accumulator, also built from dekatrons, and these gave the machine its original name. Dekatrons were 10-state valves as opposed to the more usual diode, triode and pentode valves (with two, three and five electrodes respectively). Program input to the Harwell machine was via paper tape — it was equipped with six readers, five of which are still attached to the machine today. Output was also to paper tape, which could then be used to drive a teleprinter. Logically, the Dekatron was a conventional machine, inasmuch as any computer of this period could be described as conventional.

The Dekatron was in use at AERE for 80 hours a week (55 per cent of the total time available) up till February 1953. Its reliability was perhaps partly due to the exceptionally high build standard, which you can (indeed are firmly instructed!) to admire today.

What Murrell calls the machine's "first retirement" took place in 1957, when AERE had no further use for it. The machine was then acquired by the Wolverhampton and Staffordshire College of Technology, an enterprising establishment which wanted to develop computer technology courses long before the great majority of higher-education institutions. The machine was re-christened WITCH and entered a second phase of existence assisting some of the UK's first computer science students, starting in 1965. The nine students in the first class — five men and four women — were eagerly offered jobs by IBM and have gone on to distinguished careers in industry.

Then came the machine's second retirement, in 1973, as computer technology moved on and the college acquired more modern computing equipment. Dekatron/WITCH migrated to the Birmingham Museum of Science and Industry where it was on display for 24 years up to 1997. After some time in cold storage, it then moved to the National Museum of Computing in September 2009 where it is now enjoying its fourth phase of existence after a period of routine restoration (insofar as restoration of any machine of this age is ever routine).

The Dekatron is the most visually and aurally satisfying of all the computers covered in the day's presentations. It processes data with the relays audibly clicking in synchronisation with the lights flashing on the storage racks. Murrell demonstrated how the machine was driven by programs on paper tape (and so was not a stored program machine in this mode of operation). However, he also showed how the programs could be manually interrupted and additional instructions entered to the memory by hand (so it did have the capability to run as a stored program machine, making it a true computer). This feature must have been invaluable to the staff teaching at the Wolverhampton College. Moreover, the machine is still as reliable today as when it was designed.

IRIS

The fourth exhibit to be discussed and demonstrated during the day was IRIS (Independent Radar Investigation System). This was a special purpose computer system built for National Air Traffic Services (NATS) using DEC processors (two PDP-11/34s and two PDP-11/84s) for investigating air traffic control incidents. Presenter Ben Trethowan did not know exactly when it started operation, but it was in the early seventies. The system was built at the air traffic control centre at West Drayton, and it continued in use there till 2008.

IRIS was part of a larger system called PRDS (Processed Radar and Display System). In its lifetime this system handled over 48 million aircraft movements. Virtually every reader of this report will have flown in an aeroplane whose movement was monitored by PRDS! IRIS was essentially a subset of PRDS, the smallest collection of system elements needed to generate a display of aircraft from radar data.

To achieve playback of radar data on screen, the data was captured on magnetic tape and then fed into one of the PDP-11/84s for processing; the results were passed to one of the 11/34s for rendering into vector graphics commands, and these were then transmitted to the display units themselves. These (there were originally two, although only one is used for demonstrations currently) were giant round green screens flanked by Bakelite dials and switches used to adjust the screen settings, vivid reminders of a bygone era.

In April 2008 IRIS was decommissioned at West Drayton, and the system was transported to Bletchley Park to become an exhibit in the National Museum of Computing. Restoration took a mere seven months (April to October 2008). But although the time between the last operational use of the system and its reconstitution at Bletchley Park was very short, a number of major restoration issues arose.

Although built from off-the-shelf components, IRIS has many unique features, such as use of proprietary protocols and recording formats for radar data which had involved the building of bespoke hardware elements. The technical documentation associated with these has been lost.



Figure 5. A view of the IRIS showing the multiple PDP-11 processors

Getting IRIS back into working order at the museum was thus quite a complex task, involving contact with the original suppliers (which included Plessey) for technical and operational data, and also contact with personnel who had worked with the system at West Drayton, every time an unforeseen problem arose.

Problems of a different nature also arose. One concerned faults with power supplies, which were not documented in any of the paperwork that accompanied the system. Enquiry among NATS personnel revealed that the problem was in fact well known, a result of incorrect technical documentation.

There were also problems with degradation of magnetic tape media, which had suffered from oxide separation, and also the tape drives themselves, with components such as drive belts having perished.

In addition to the issues encountered in the restoration itself, Trethowan stressed the need for emulation of elements of the system in demonstrations, to prevent irreplaceable parts wearing out.

ICL 2966

The last of the five machines demonstrated during the day, the ICL 2966, is a typical representative of the mainframes that dominated computing from the sixties to the turn of the millennium, and dates from roughly the middle of this period. It was a mid-range machine in the 2900 series introduced in 1974, which was described by Restoration Project Leader Delwyn Holroyd as “the last great British computer project”.

The 2966 had been owned by Tarmac Quarry Products of Wolverhampton and had started life as a (less powerful) 2950 when acquired in the seventies. Tarmac upgraded it to 2966 specification in 1988 and then continued to run it till 1999, when it had to be replaced as Tarmac’s software could not be made Y2K-compliant.

As delivered to the National Museum of Computing it was configured with 8 MB RAM and 7 GB total online disc storage, numbers which seem small today but were then more than adequate for running the whole of a corporate business. The disc storage featured drives that used removable packs of 80 MB and 200 MB, and there were also some fixed-disc drives for faster access. The 2966 was also equipped with magtape drives, a card reader and a line printer (which is not original to the system, but is identical to the one that was).

All the hardware was proprietary to ICL, using mostly TTL and ECL circuitry in the processors. The 2966 was the last model to be built before ICL adopted Fujitsu LSI logic circuitry so, says Holroyd, “this is one of the last machines it is possible to consider restoring”. But although technology moved on, the virtual machine architecture of the 2900 series survived through the later ICL Series 39 and Trimetra ranges, and is still in use today in some government applications.

The 2900 series had its own proprietary operating system, VME (Virtual Machine Environment). Its virtual machine architecture however allowed running of other operating systems, and ICL introduced emulator software known as DME (Direct Machine Environment) to permit the use of predecessor ICL and ICT operating systems. Tarmac (along with many other 2900 series users) ran its 2966 using this software to host the George operating system used on the predecessor 1900 series (Tarmac actually ran George 2 while most other users stayed with the later George 3).

The system arrived at Bletchley Park in 2009 (at almost the same time as the IRIS system described above). Restoration to working order has been a slow process. The first EDS 80 disc drive was



Figure 5. Delwyn Holroyd (far right), project leader of the ICL 2966 demonstrates the mainframe system — the museum’s largest single exhibit

working by May 2010, but head crashes occurred frequently and it was soon realised that running these electromechanical products regularly would quickly wear them out, so a disc emulation system was devised using modern hardware (a 64GB Micro SD card), and this was running by September 2014. Two months later George 3 running under DME was achieved; by November 2015 the tape drives were operational; and last year the second system control processor became operational — this is necessary to permit running the native VME operating system.

Visually the Tarmac 2966 is a mishmash of colours ranging from the standard 2900 “hot tango” orange of the 1970s to the standard beige of the 1980s with some blue panelling to complete a rainbow effect. This is less disturbing than it might be for, it seems strange to report, a mainframe of a type totally familiar to the majority of members of the CCS is as unfamiliar to the typical Bletchley Park Museum visitor as the dinosaurs at the National History Museum. “Most visitors have never seen anything like this machine”, says Holroyd. There is no similar machine from ICL or any other contemporary computer manufacturer on public display in the UK today.