Preserving Computer-Aided Design

Digital Preservation Coalition Report

Alex Ball

21 October 2013

Abstract

The Digital Preservation Coalition Technology Watch report 'Preserving Computer-Aided Design' explains the particular challenges of working with CAD models, and sets out some of the standards and techniques that are used to help preserve them. In this talk, Alex Ball expands on the themes of the report with examples from his previous project work in this area.

I work for the Digital Curation Centre (DCC). We are a collaboration between the Universities of Edinburgh, Bath and Glasgow . . . funded as a centre of expertise helping the UK Higher Education sector. While these days we concentrate mainly on research data management practice, we remain interested in the full range of digital curation activities including digital preservation.

I was first introduced to the challenges of preserving CAD models in the KIM Project . . . EPSRC Grand Challenge . . . ESRC funding . . . 11 different universities . . . aerospace, defence and construction. Looked at challenges of firms providing products as services . . . information organisation, selection policies for data, and finding ways of learning design lessons from in-service data and knowledge.

Subsequently I co-wrote a report on data curation in engineering disciplines as part of the DCC SCARP (Sharing, Curation, Reuse and Preservation) project . . . involved in two Jisc-funded projects looking at research data management in the University of Bath's Mechanical Engineering Department.

¶ Most recently the DPC asked me to write a Technology Watch report on Preserving CAD, and that was published earlier this year. It provides an introduction to CAD, explains some challenges of preserving CAD and suggests possible approaches. In this presentation I'd like to draw out some of the themes from the report and relate them back to some of my earlier work. . . .

Contents

house.jpg

View image here: http://blogs.umass.edu/bct420-schreyer/files/2009/01/2_854_main3_b.jpg

minster.png

View image here: http://guides.archaeologydataservice.ac.uk/attach/Cad_1-3/minster_3_large.png

deansway.png

View image here: http://guides.archaeologydataservice.ac.uk/attach/Cad_Cs1/cad_fig04_thumb.gif

Figure 1: CAD is important in many different fields including transport engineering, architecture and archaeology.

1 Curation challenges for CAD

Digital curation is all about making sure digital objects are as useful as they can possibly be regardless of where they are in their lifecycle, and preserving the objects of lasting value long into the future.

¶ And there is a lot of valuable information tied up in CAD models (Figure 1). Planes, trains and ships may be in service for anything from 30 years to 100, while with care buildings and bridges may last centuries. The plans need to be available throughout that time in case these products and constructions need to be modified as needs change, or to help accident investigators should anything go wrong.

Furthermore, CAD has become very popular as a way of recording archaeological dig sites. And because archaeology is a destructive science, the information going into those CAD models is irreplaceable.

¶ The problem is that the CAD models don't last anywhere near that long. New versions of CAD systems typically come out every six months, and it's by no means certain that the new version will read old files correctly, because how the model renders is highly sensitive to the modelling kernel in use. And after about ten years, it's likely they'll stop making new versions and the CAD system will become obsolete.

This constant turnover of CAD systems is a symptom of intense competition between the CAD vendors for customers. In order to keep revenue up and costs down, vendors are striving to produce innovative new features and retire their older systems, while at the same time trying to stop customers being lured away to alternative systems.

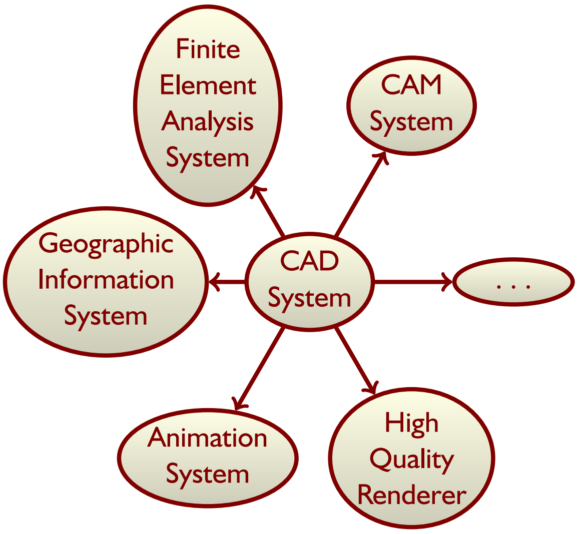

Figure 2: CAD models are used by more than just CAD systems

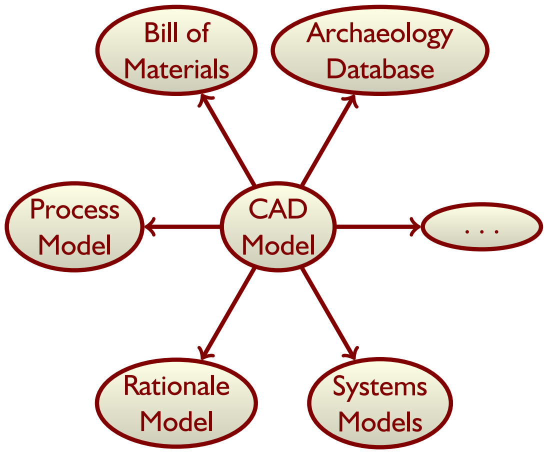

Figure 3: CAD models are just part of the story

Anecdotal evidence suggests CAD vendors pay much more attention to getting their import filters right than their export filters.

¶ It's not just the CAD systems themselves that are the worry. The trend is for ever tighter integration between CAD and other systems (Figure 2), so you'll find CAD models being used in manufacturing systems, geographic information systems, ray tracers, and many more besides.

¶ Just moving the data around is one thing. Sometimes it needs to be supplemented with additional data from, say, a database of building material properties or wiring diagrams (Figure 3). There's also the matter of understanding it: why things were modelled at all, why they were modelled in a particular way, and how one might go about modifying them. For that, you need to access to client specifications, rationale models, reports and meeting minutes, and other forms of contextual metadata. How do you link it all together?

Over the next twenty minutes I'll show you some attempts to solve these problems.

2 The CAD format problem

I've already mentioned that CAD models don't stay readable for as long as we need them to. And that's a shame, because some of the advanced features of CAD models, like feature-based modelling, procedural modelling, and built-in part libraries don't translate well to other formats. So what to do?

¶ You could try software emulation. Pros

- CAD model unchanged *so that keeps the provenance trail simple and makes long-term preservation largely a bitstream problem*

Cons

- Licence may *be time-limited, and therefore* not allow it
- Need to preserve expertise in the system
- Hard to maintain integration with newer systems

¶ You could try continuously migrating the model to match the latest versions of the CAD software. Pros

- Models stay usable by current designers and systems

Cons

- Incremental data loss/corruption
- Each migration needs to be validated, problems resolved*, and that can be costly in time and money; one of the reasons the Airbus A380 took so long to finish was because they had a CATIA upgrade to cope with in the middle of it.*
- Migration path may run out . . .

¶ So that's not terribly satisfactory. A more sustainable approach is normalisation to a standard format or two. What makes it distinct from the previous approach is that when you load up an old model, you know it has only been migrated twice: once when it went into the archive and once when it came out. Pros

- Limited data loss
- Support for archival format likely to improve rather than degrade over time

Cons

- Some data loss/corruption
- Migration needs to be validated, problems resolved

¶ What formats could you choose for normalisation?

DXF has a published specification and is supposed to be able to express all the information encoded by AutoCAD's native DWG format. By all accounts, though, recent specifications have not been detailed enough for full support to be implemented in third party tools.

STEP is the Standard for the Exchange of Product Model Data, or ISO 10303. It has the opposite problem: it has been in continuous development for almost 30 years, and at 600 parts is the largest of all the ISO standards. This makes it very hard for non-experts to get their heads around.

There are various ways of expressing CAD data in STEP. Most CAD systems support AP 203 Edition 1, which is quite basic. But Edition 2 is more expressive, and AP 214 extends support in a different direction. And soon there will be AP 242 that supersedes both of them. Some vendors are already working on supporting it through the CAx Implementer Forum, which bodes well for the future.

STEP is also the basis of other international standards such as the Industry Foundation Classes for Building Information Models, though I think it's fair to say that it is still early days for IFC. LOTAR is a work in progress but one of the really useful things it does is help you navigate the STEP standard, so you know which of the many ways of doing things is right for a particular type of information.

These formats are our best shot for preserving the most information, but because they are highly complex there's a lot to go wrong, and the resulting files can be unwieldy to store or move around.

¶ That's why normalisation formats are useful.

Standards

U3D Universal 3D – ECMA-363 *is fairly basic and is used in 3D PDFs*

X3D ISO/IEC 19775, 19776, 19777 *is aimed at virtual reality applications, so it's good for animation and interactivity*

IGES Initial Graphics Exchange Specification – ANSI Y14.26M-1981 . . . ANS US/PRO/IPO-100-1996 *was the first exchange standard, now withdrawn as it has its problems, but remains widely supported*

PRC Product Representation Compact – ISO/PRF 14739 *more advanced format used in 3D PDFs; can manufacture from it*

JT ISO 14306 *widely used in aerospace and automotive industry, can manufacture from it, and is linking up with STEP*

Non-standards, *generally tied to single CAD system, and unclear how long or widely they will be supported*

DWF/DWFx AutoCAD DWG Web Format

3D XML CATIA visualisation format

¶ So those are the options. Which do you choose? Well, there was a project at MIT called FACADE that developed and methodology and set of tools for archiving architectural CAD models. They recommended that CAD models should be stored in *four different formats:*

1. the original format *for the reasons already described*
2. a full exchange format, e.g. STEP *in order to preserve the maximum amount of information for the long term*

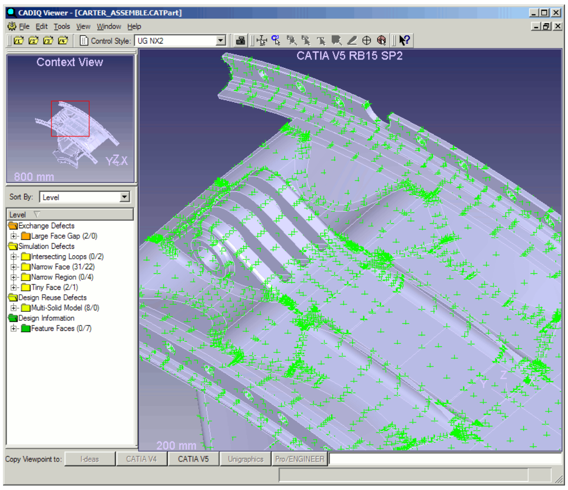

© ITI/SNECMA

Figure 4: Validating the survival of shape data using a point cloud: the point cloud is generated using the original software, and loaded into any subsequent software rendering the model, where it should line up with the displayed geometry. (This technique was developed by LOTAR International.)

3. a desiccated format, e.g. IGES *so that the basic shape data isn't lost if there should be a problem with the full exchange file*
4. an access format, e.g. 3D PDF *for in-browser previews of the data*

But if you are viewing the model in something other than the original system, how do you know what you're seeing is right? ¶ A good technique is to use validation tests. You pick an aspect of the design that you're interested in, run the same test on it in the original CAD software and a new system, and if you get the same answer both times you know that information has survived. On the slide (Figure 4) is a test devised by LOTAR International to test shape data using a point cloud.

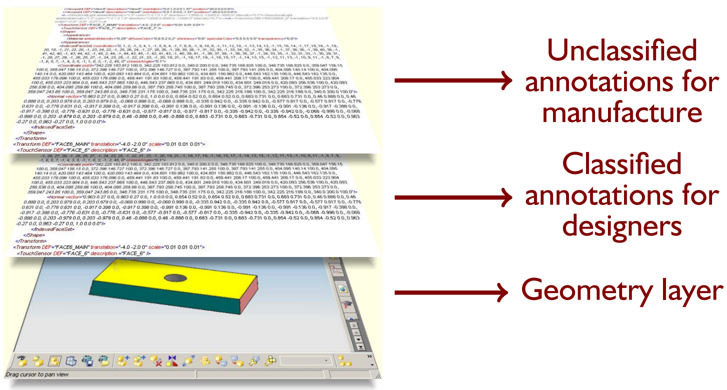

Figure 5: Lightweight models with multilayer annotations

If you know there are things you are going to lose, there might be other ways to save the information. ¶ One idea we experimented with in the KIM Project was a system of lightweight models overlaid with annotations stored in external files (Figure 5). The idea was to store information like feature semantics in an XML annotation file. Using named entities or point clouds, the annotations could be related to parts of the geometry, so you could apply them to original model or any normalised version. We anticipated having different annotation files for different purposes and security levels, and layering them up as necessary.

3 The rationale problem

I'd like to move on now to consider how to make sure the rationale behind a CAD model is recorded.

¶ One way is through feature-based modelling. Features explicitly record the engineering significance of certain shapes in the design. So a curved blend between two surfaces might anticipate a manufacturing process or remove a concentration of stress; the feature semantics would tell you which the designer intended. Incidentally, features tend to be viewpoint-specific. Take this shape (Figure 6).

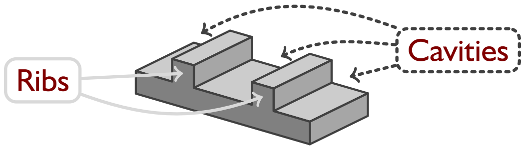

Figure 6: Sample model with an ambiguous feature set

The designer might need this marked up as rib features applied to a base surface in order to calculate the optimum distribution. The manufacturing engineer, though, might need it marked up as cavities milled out of a thick surface. It's not usually possible to have both in the same model.

1. Insert cylinder $l = 20$ $r = 1.0$

   ⟵ Change cylinder $l = 40$ $r = 0.5$

2. Insert sprocket $r = 3.0$
3. Fit sprocket to cylinder
4. Group cylinder and sprocket
5. Scale group by 1.75×

. . .

Figure 7: Construction history modelling

¶ A lot can be gleaned from the construction history of the model (Figure 7). Some CAD systems record how a model was put together so designers can step backwards and forwards through the process and insert or edit some of the steps, but it can also help other designers to understand the model.

¶ In the KIM Project we went a step further (Figure 8). The team at Heriot-Watt University coaxed a CAD system to output a log of user interactions, and ran analysis software on it to infer rationale and detect inefficiencies in the design environment.

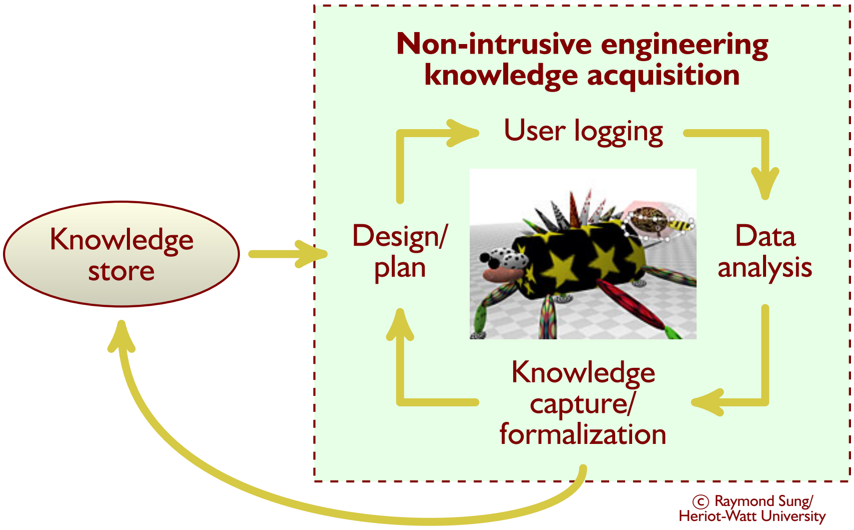

Figure 8: Automated information capture

The research focussed initially on virtual reality design environments, but in order to apply the same principles to a CAD environment, the team used a modified version of the Bamzooki software produced by the BBC.[1]

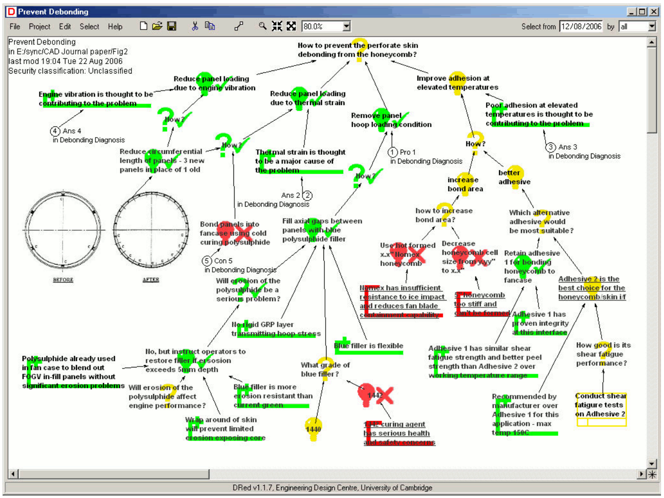

© Rob Bracewell/University of Cambridge

Figure 9: A rationale map produced using Rolls Royce's Design Rationale Editor

Meanwhile the team at Cambridge developed a tool called DRed, which could record the decision-making process more explicitly.¶ Here (Figure 9) is a DRed map showing how a particular design issue was solved. The map is hyperlinked so you can jump back and forth to other maps, reports, spreadsheets, and designs. This tool is now in common use in Rolls Royce in place of traditional reports.

[1] `http://www.bbc.co.uk/cbbc/bamzooki/`

4 The contextual problem

So far I've been looking at rationale in fine-grained terms, but just as important is the rationale behind entire ventures and approaches, and other contextual matters such as how a set of records all fit together. This is an area where Building Information Models and their near equivalent in industrial engineering, Product Lifecycle Management systems, can help. But in the projects I was involved in we were also interested in how SMEs and researchers could get by without them.

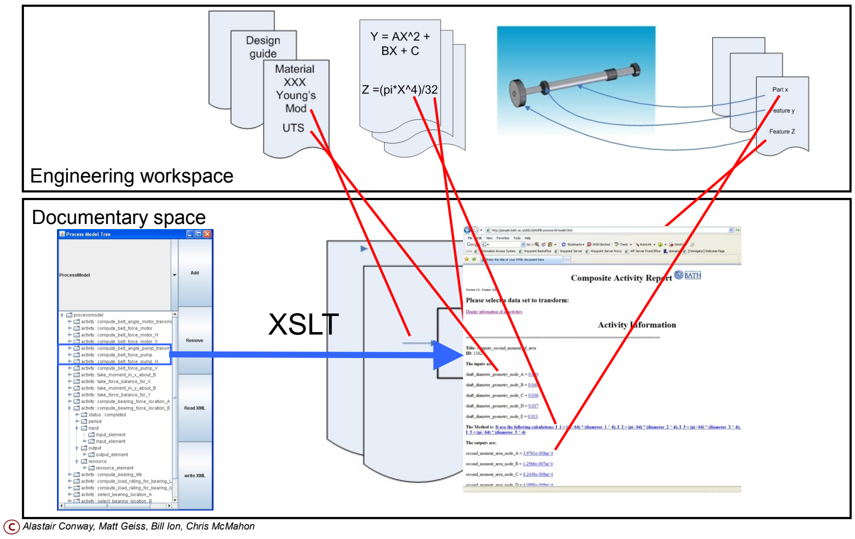

Figure 10: Recording the associations between transactional records

¶ The University of Bath KIM team developed desktop tools to help designers record how they are using information – for example, which figures they use in their calculations, or what product catalogues they consult (Figure 10). These relationships are stored in XML files, which can be transformed into IDEF$_0$ diagrams or textual reports through the magic of XSLT. They also experimented with linking concepts in different documents using topic maps to reveal information dependencies.

¶ When it came to the ERIM and REDm-MED projects, we saw in microcosm among the researchers at Bath many of the problems we'd seen affecting industry as a whole.

The university and department hadn't set out best practice for how researchers should be managing their data, so everyone was doing their own thing. The closest thing we had to a storage policy was a naming convention for project areas on the university's shared research drive, and some plain text indices saying who was responsible for each folder and when the retention period was due to expire. Metadata that might be needed for preserving the data just wasn't being collected.

We found that researchers were generating all sorts of data. CAD was only a very small part: there were photos, videos, process flow diagrams, materials data sheets, activity models, product models, topic maps, bills of materials, questionnaire responses, interview transcripts, 3D laser scans, flight path data, thermal data profiles, surface roughness data . . . There tended to be many, many small files instead of a few large ones, and trying to understand how they all related to one another was nightmarish. We interviewed some researchers about their data and they were sometimes hazy about it themselves, and it didn't help that some of it was missing, as it was only on the hard drives of researchers who had left.

So we felt that in addition to setting up some data management planning guidance, we needed to make sure that enough contextual metadata was being collected to explain the data.

¶So we defined a Minimum Mandatory Metadata Set; REDm-MED version at: `http://opus.bath.ac.uk/30372/`

- Drew on PREMIS for preservation metadata
- Drew on DataCite for descriptive metadata
- *We wanted it keep it as simple as possible so we* Excluded metadata that could easily be generated later
- *and* We looked for ways to collect it automatically *if at all possible*

(NB: an extensive specification of what metadata to collect for archaeological CAD can be found at: `http://guides.archaeologydataservice.ac.uk/g2gp/Cad_Toc`)

We also felt it was important to create a manifest of records produced by the research and note the associations between them, as an aid to understanding the collection and working out which were the most important records.

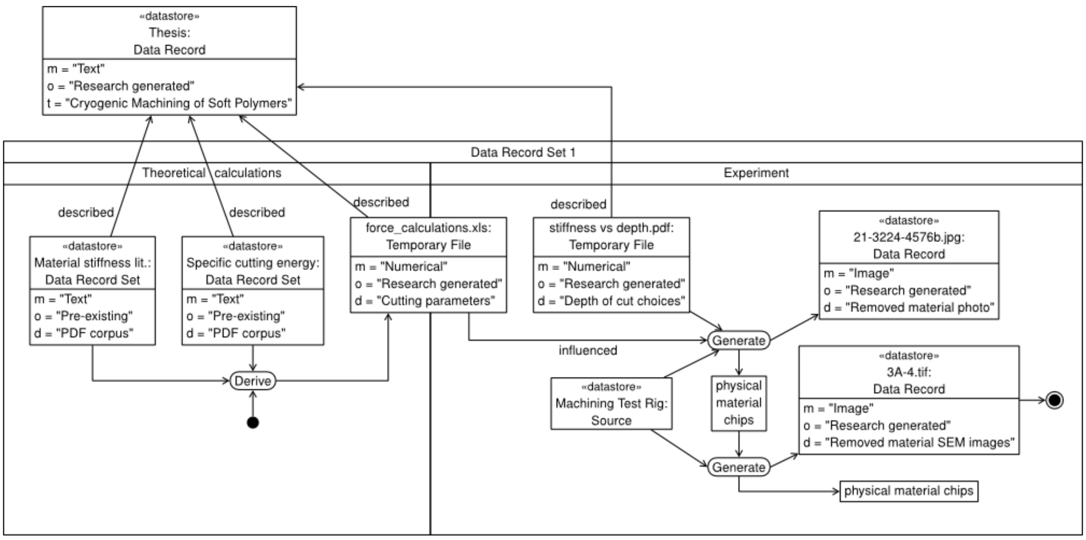

`http://sourceforge.net/p/raidmap/`

Figure 11: A Research Activity Information Development (RAID) map

¶ So in ERIM we developed a systematic way of recording this, based on UML so it was intuitive for humans to look at and could be made machine readable with a little effort (Figure 11). We called it RAID, or Research Activity Information Development, mapping. And in REDm-MED we developed a RAIDmap tool which would partially automate the process of creating these maps and collecting the metadata for each record.

Even though this is tuned slightly differently, you can see it is answering the same basic question as the transactional record tracking from KIM. If a mistake was introduced you could use this to track down where it came from. If you were selecting records to add to an archive you could use this to assess which records held unique or vital information, and which could safely be discarded. And if you wanted to repurpose the data you could use this to understand how to start adapting and integrating it with your own work. None of this is peculiar to CAD, but it becomes very important where CAD models are concerned because they are intrinsically fragile, and hard to fathom if you are unfamiliar with them.

5 Final thoughts

So, in conclusion then, CAD files can be used for many different purposes, so it is important for an archive to establish why models are being kept and use that as the basis for deciding a preservation strategy.

If particular properties are important, create validation tests that can prove whether they have survived.

There is information that cannot be exported from native CAD formats, so they may as well be preserved for as long as they are readable, but since they will be unreadable one day, they should also normalized to a full exchange standard like STEP, and to a visualisation format for backup.

It is important not to neglect supporting documentation, metadata and contextual records, which can be vital for future users to be able to understand the models.

And if we are serious about preserving CAD models we need to impress on regulatory bodies and CAD vendors how important support for standards is. The information locked up in CAD models is too valuable for its survival to be left to the whims of third party corporations whose interests may lie elsewhere. It needs to be put into open, standard formats where it can be exchanged and preserved for the benefit of society as a whole.

Alex Ball. DCC/UKOLN, University of Bath. http://alexball.me.uk/

The DCC is supported by Jisc.

For more information, please visit http://www.dcc.ac.uk/

Preserving CAD DPC Technology Watch Report: http://dx.doi.org/10.7207/twr13-02