

EVALUATING THE POTENTIAL OF COMMERCIAL GIS FOR ACCELERATOR CONFIGURATION MANAGEMENT*

T.L. Larrieu, Y.R. Roblin, K. White, R. Slominski
Jefferson Lab, Newport News, VA 23606, USA

ABSTRACT

The Geographic Information System (GIS) is a tool used by industries needing to track information about spatially distributed assets. A water utility, for example, must know not only the precise location of each pipe and pump, but also the respective pressure rating and flow rate of each. In many ways, an accelerator such as CEBAF (Continuous Electron Beam Accelerator Facility) can be viewed as an “electron utility.” Whereas the water utility uses pipes and pumps, the “electron utility” uses magnets and RF cavities.

At Jefferson lab we are exploring the possibility of implementing ESRI’s ArcGIS as the framework for building an all-encompassing accelerator configuration database that integrates location, configuration, maintenance, and connectivity details of all hardware and software. The possibilities of doing so are intriguing. From the GIS, software such as the model server could always extract the most-up-to-date layout information maintained by the Survey & Alignment for lattice modelling. The Mechanical Engineering department could use ArcGIS tools to generate CAD drawings of machine segments from the same database. Ultimately, the greatest benefit of the GIS implementation could be to liberate operators and engineers from the limitations of the current system-by-system view of machine configuration and allow a more integrated regional approach.

The commercial GIS package provides a rich set of tools for database-connectivity, versioning, distributed editing, importing and exporting, and graphical analysis and querying, and therefore obviates the need for much custom development. However, formidable challenges to implementation exist and these challenges are not only technical and manpower issues, but also organizational ones. The GIS approach would crosscut organizational boundaries and require departments, which heretofore have had free reign to manage their own data, to cede some control and agree to a centralized framework.

INTRODUCTION

Jefferson Lab

Jefferson Lab is a U.S. Department of Energy laboratory consisting of the main CEBAF (Continuous Electron Beam Accelerating Facility) accelerator and a Free Electron Laser. The physical plant supporting the two experimental facilities is extensive. For CEBAF alone there are 7 km of beamline, more than 2000 magnets and power supplies, 338 5kW klystrons, 42 cryomodules each containing 8 superconducting niobium cavities, a 400,00L low conductivity water (LCW) system, and a 2K helium refrigeration plant. The control system based on EPICS is equally extensive consisting of approximately 65,000 I/O Control Points, 250,000 EPICS Records, 140 IOCs, and 80 UNIX Machines.

The Utility Analogy

Much as a water utility owns and operates a network of pipes, pumps, and flow meters to deliver water, or an electric utility operates generation plants, transmission lines, and transformers to deliver electricity to its business and residential customers, so the CEBAF infrastructure is operated to deliver a high quality electron beam to the experimenters who are the lab’s “customers”. Is it possible that we can borrow ideas, techniques and tools used by utility companies to manage and run their large and complex systems and apply them to better managing our “electron utility?”

*NOTICE: This work was supported by DOE contract DE-AC05-84ER40150 Modification No. M175, under which the Southeastern Universities Research Association (SURA) operates the Thomas Jefferson National Accelerator Facility.

With tight budgets forecast for the future and the inevitable ensuing demand to *do more with less*, Jefferson Lab is in a situation much akin to where many utility companies found themselves in the late 1990's. Deregulation and the consequent erosion of monopoly protections for their businesses forced these companies to make determined efforts to become more streamlined and efficient in order to fend off potential competitors. Many re-evaluated how their information systems functioned and integrated to find ways to stop duplicating data and functionality.

During the preceding decade as different software technologies evolved, departments had developed computerized solutions for various tasks (Outage Management, Mobile Dispatch, Network Analysis, to name examples), but most often these solutions were *islands of automation* with scant connectivity among them [1]. *Islands of automation* is an apt description for the current state of computer applications at Jlab. Various departments have evolved information resources to perform their own allotted tasks. The Engineering department uses Computer aided design tools such as ME10, AutoCAD, and IDEAS to create as-built drawings and model assemblies. Accelerator Operations uses EPICS to monitor control the accelerator hardware. Accelerator physicists use tools such as OptiM or DIMAD to investigate and model the beamline. Facilities Management, Environmental Health & Safety (EH&S), and other departments use yet other commercial or internally developed software.

The drive to re-engineer their IT infrastructure led many utility companies to take advantage of advances made in GIS (Geographic Information System) software to broaden its scope from a departmental endeavour into an enterprise resource (Figure 1) [2]. At the beginning of the decade, GIS had primarily been a tool for performing spatial analysis: practitioners would use specialized "coverage files" and "shape files", often overlaid with CAD data or satellite imagery or joined with database tables, to make analyses such as environmental or demographic studies. As relational database software matured in capability, new versions of GIS software began to store spatial data as tables in a *geodatabase* rather than in specialized file format. The *geodatabase* could then be connected with other enterprise applications such as Customer Relationship Management (CRM) and Enterprise Resource Planning (ERP) software from companies such as Peoplesoft and SAP that organizations were simultaneously adopting during the 1990s.

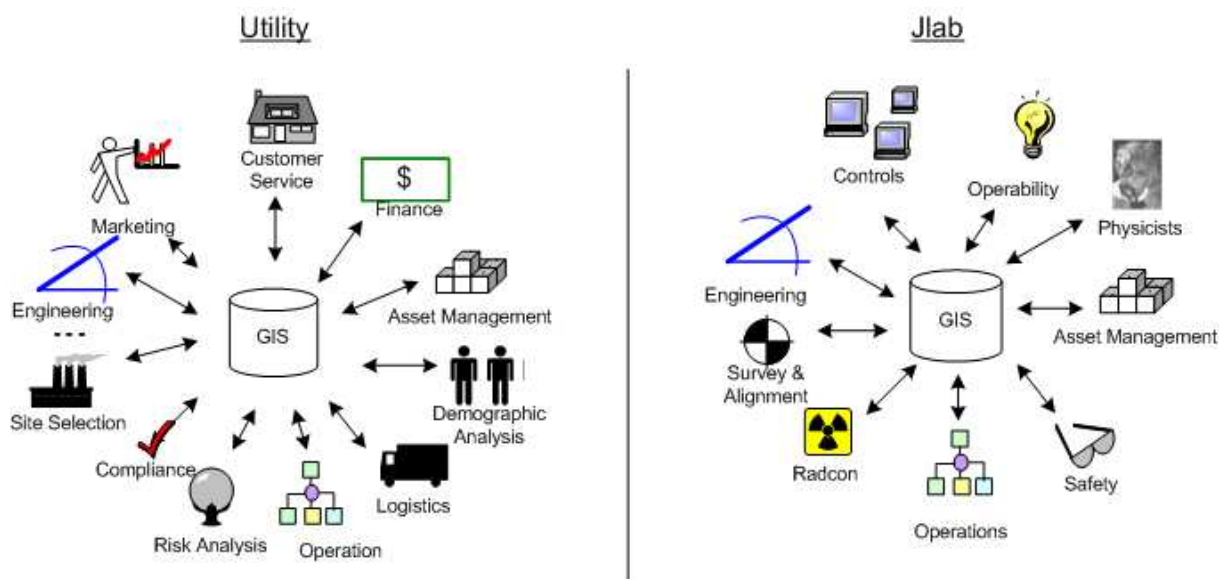


Figure 1. The GIS is a centrepiece of a modern utility company's IT infrastructure that facilitates information sharing throughout the organization[2]. It is possible to conceive that an implementation at Jlab could also have broad reach.

The result of centralizing GIS, for the utility companies that did it, is an enterprise-wide repository of spatial and operational information that changes the way business is done. Engineers can incorporate environmental data early in project planning; Marketing departments can study the

distribution of customer complaints; operations managers can comment on design plans before construction begins [1].

One can readily envision similar scenarios at Jefferson Lab (right hand side of figure 1). Accelerator Physicists could comment on the 12GeV upgrade design as it progresses and use the most recent data for their modelling. The Operability team could analyze the spatial distribution of downtime events. Work planners could easily access site hazard information maintained by EH&S (Industrial hazards, Oxygen deficiency hazard levels) or RadCon (Radiation hazard or contamination areas). Furthermore, a variety of individuals and committees have begun to suggest the need for more regional perspective in a wide-variety of control system tools [3]. The current suite of software and (M)EDM screens, tends to have a system-centric perspective: tools for phasing RF cavities, tools for checking and setting correctors, etc., but operators often need work from the perspective of a particular region such as the Injector or the North Linac when setting up or troubleshooting. A location-aware GIS database could be central to the effort of developing the tools that implement this vision of control system evolution.

GIS SOFTWARE

It is of course possible to build a central accelerator information system without using GIS software, the typical approach being a relational database back-end, with purpose-built interfaces written in one or more high-level programming languages such as C++ or Java. The IRMIS¹ (Integrated Relational Model of Installed Systems) project being pursued by various members of EPICS collaboration is an example of such an approach [4]. Might it be possible to augment if not replace some custom programming for such a project with off-the-shelf tools? For example, the ArcGIS desktop client implements quite closely the desired features of the IRMIS project's yet-to-be-implemented, "Installed Component Search/Viewer/Editor" tool. The next sections of this paper describes some GIS capabilities that might be useful with specific reference to the ArcGIS package from ESRI and its associated toolkits.

CAD Integration

A wealth of information is present in as-built CAD drawings (informally referred to at Jefferson Lab as *songsheets*) of the accelerator. Staff in the Engineering department make great efforts to maintain these drawings as an accurate reference. There are limitations on the usefulness of these drawings however:

- The engineering staff who maintain the drawings must rely on communication, typically via email, from technicians who install/replace equipment or from the Survey & Alignment team who precisely determine component positions in order to incorporate updates.
- Changes to the as-built drawings are effected separately from changes to software configuration and physics models.
- The CAD drawings are tiled. Several must be laid side-by-side to encompass an area.
- The CAD annotation is static. Only element names are shown.

The ESRI ArcGIS software can utilize Computer Aided Design (CAD) drawings produced by commonly-used CAD software such as AutoCAD or ME10 in ways that overcome the aforementioned weaknesses. ArcGIS can use unprocessed CAD drawings as a backdrop for the presentation other GIS data layers (For example, beamline coordinates generated by running the OptiM beamline modelling code can be readily plotted overtop of the songsheets). Going a step further however, one can process the CAD files to create GIS database point, line, and polygon object classes. Once imported into the GIS database, the CAD-derived objects can be extended with additional attributes, can be related to other database tables, and made generally accessible to control system software.

Going in the other direction, ArcGIS can generate CAD drawings from the contents of the geodatabase. This means it is possible to produce drawings that contain elements of interest selected via database query which can be labelled and annotated using database fields.

1. It is in fact our expectation that we could use the IRMIS schema and tools such as the PV crawler to build a *Controls* layer in a geodatabase

Drawings generated from the GIS need not be constrained by the tiling boundaries or annotation conventions of the original CAD drawings.

Geodatabase

The primary repository for GIS information is in a *geodatabase* which utilizes an underlying personal database (Microsoft JET) or enterprise RDBMS (Relational Database Management System such as Oracle or Sybase). In the latter case, the geodatabase inherits the scalability and multi-user concurrency of the mature commercial RDBMS as illustrated in figure 2.

In much the same way that software versioning works in CVS, ArcGIS geodatabase versioning allows users to create named branches of GIS data with which to work. The branched-off version can, if desired, later be reconciled with the main branch. Branches can also be kept as historical records. There are several potential applications of versioning to accelerator configuration management. Versions could be instantiated to support what-if modelling of different configurations.

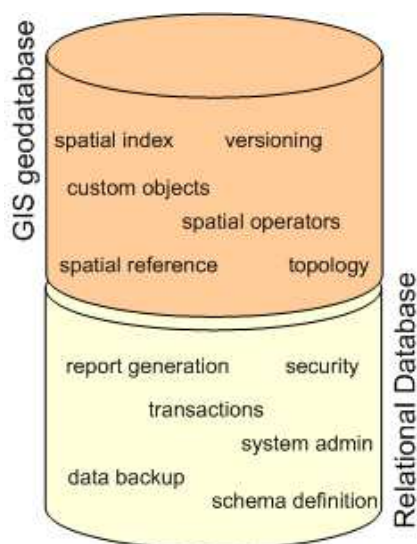


Figure 2. The geodatabase inherits from and adds to the functionality of the underlying RDBMS [5].

Planned updates to software and hardware could be incorporated into a non-production version which then gets integrated into the main production version following a maintenance period. Finally, it would be possible to save “snapshots” of the accelerator at key intervals such as at the completion of an experimental run, thereby making it possible at a future date to refer back to the historical configuration of the machine. One benefit of consolidating information into an enterprise geodatabase could be consistent versioning of drawings, models, and software configuration.

Spatial operators and indexes implemented at the geodatabase level allow efficient queries based on geographic relationships such as proximity, adjacency, and overlay. These could allow users to retrieve accelerator configuration information in ways that would not be easily available with a pure RDBMS approach, such as finding elements within a given distance of each other.

Desktop Client

The ESRI GIS toolkit includes ArcGIS desktop software (ArcView, ArcEditor, ArcInfo, listed in order of increasing feature set) for working with geodatabases. This off-the-shelf software includes functionality for both tabular and graphical editing, visualizing, analyzing, and managing of the contents of the geodatabase. Non-programmers can perform sophisticated query and display operations. New objects can be added to the database by “drawing” them into place using CAD-like tools and then right-clicking to open an attribute editor and specify the non-spatial properties.

Visual Basic and Python can be used to customize the interface and to script complex multi-step operations such as importing and identifying data from CAD drawings.

Development Tools

A Software Development Kit is available for Windows, Linux, and Solaris that supposedly provides a consistent set of data access components and developer controls available as ActiveX, JavaBeans, .Net controls, and Motif Widgets. These tools can be used to create or extend custom applications to use the Geodatabase.

Tracking Server & Tracking Analyst

A pair of intriguing ArcGIS software extensions originally developed by Northrop Grumman, but now available directly from ESRI are *Tracking Server* and *Tracking Analyst*. The former allows processing and display of real-time data on top of other data layers selected from the Geodatabase. The latter can be used to visualize time-series data on a base-map. Usage examples cited by the vendor include real-time tracking of a trucking fleet or mapping the temporal distribution of earthquake activity. Although it is far too early to seriously consider feasibility of implementing these software, we can imagine writing a Channel Access plug-in that could allow intriguing possibilities for displaying PVs either in real-time or from the archiver:

- The PVs to be displayed could be selected spatially. For example select PVs from all BPMs within 10m of a given magnet, or from all dipoles of a particular manufacture in the East Arc.
- The PVs could be displayed on top of the as-built drawings to give the operator or technician better contextual awareness.

AN EXPLORATORY PROJECT

The general discussion of the software capabilities is well and good, but will it live up to its promise in our environment? A Utility company transitioning to Enterprise GIS can find numerous success stories for reassurance. GIS vendors such as ESRI provides industry-specific data models [6] and reference implementations, and numerous third-party vendors offer add-on products (outage tracing, service dispatching, etc.) and consulting services by staff with utility industry GIS experience. In contrast, references to GIS deployment at accelerator facilities are scant (one possible example is at CERN [7]) and reference implementations are non-existent. Without examples of what a final implementation might look like, we have undertaken an exploratory project of limited scope.

Motivation/Scope

The impending rollout of the new Art++ online model server serves as an opportunity to experiment with the capabilities and application of some of the ArcGIS tools. Art++ currently reads its configuration from sets of OptiM files that describe the layout and physical properties of elements in the optics lattice, such as magnets and RF cavities. The intention before Art++ goes into production is to replace the OptiM files with a relational database and then from that database generate AML (Accelerator Mark-up Language [8, 9]) configuration files as needed. We therefore have an opportunity to try to use a GIS database for that backend and to use ArcGIS tools to build and manage it.

Figure 3 presents an overview of how various pieces will fit together for the first phase of the project. At the end of the first phase we desire to have a consistent set of optics configuration stored in a database which we can export via AML to the Art++ model server. If the experience of using the GIS tools and methodologies proves beneficial during the first phase we will consider committing further investment to license the ArcSDE framework and the ArcGIS SDK.

The enumerated goals of the project are as follows:

1. Build a viable database to maintain Optics data previously maintained in OptiM and DIMAD files and interfaces for making the data available to Art++.
2. Forge a collaboration among the Controls Group, Mechanical Engineering, Accelerator Physicists, and the Survey & Alignment team for maintaining the model-server database.
3. Assess how useful GIS tools were to during the completion of steps 1 & 2.

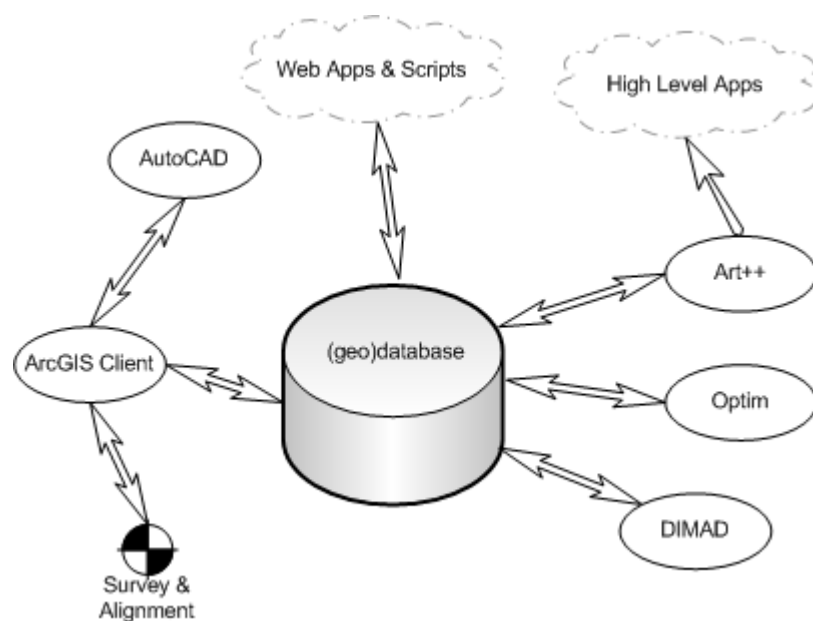


Figure 3. The geodatabase will consolidate data from various sources for the benefit of the Art++ model server.

Work to-date

As our first step in the construction of the geodatabase, we imported data from recent OptiM files into database tables. Next we imported element locations data provided by the Survey & Alignment team (Magnet Centers, RF Cavities) and extracted others (BPMs, Correctors) from DIMAD. When we turned our attention to importing data from the accelerator as-built diagrams (songsheets), we began experiencing difficulty. It turns out that it is difficult to apply automated processing tools to import our CAD file data for several reasons:

- Each Songsheet is produced in arbitrary drawing units. In the end, we must open each of the 37 files, find pair of matchpoints, and provide those matchpoints to ArcGIS so that the CAD file will be displayable using the standard Jlab accelerator coordinate system.
- The Songsheets make no use of layers. Cross-section, plane-view, annotation, and page template are all intermingled in a single drawing layer, thwarting automated extraction of elements.
- Elements are always drawn as multiple lines, never as polygons. Processing is required to merge related lines into a single geodatabase polygon object.
- Annotation is not attached to drawing elements as attributes (AutoCAD xdata), but simply written onto page.

The common thread to all these difficulties is that the *songsheets* date from an era when the only desired use of an engineering drawing was to look at it. Unfortunately, it does mean that considerable effort will be required to integrate the data from these CAD files into a broader information system.

Next Steps

Once the geodatabase begins to hold updated lattice information, we will obviously need to be able to export the contents of the database into the the required input formats of specific modelling codes such as OptiM or DIMAD. With this goal in mind, one of the authors (Y. Roblin) has begun working with and contributing to the Universal Accelerator Parser project, which aims to provide a framework for converting a lattice described in AML into model-specific codes [8,9]. We therefore have concentrated on writing software to import and export AML to and from the geodatabase, expecting the Universal Parser to handle translation between AML and various model formats.

We will need to work with the Mechanical Engineering group to devise updated standards for the as-built drawings that address the shortcomings listed in the previous section if we wish to make use of them. More time-consuming will be the work of retrofitting the standard to the existing drawings. In the meantime, we have begun using ArcGIS “heads-up” digitizing tools to perform operator-guided

electronic tracing of a subset of the as-built data so that we can experiment with the process of exporting from ArcGIS to CAD. Being able to demonstrate the process of creating a couple data-driven song-sheets on-the-fly may help justify the effort of modernizing our current drawings.

So far we have used a Jet database (MS Access format, what ArcGIS refers to as a “personal” geodatabase) to store ArcGIS feature classes, because without the ArcSDE middleware, the ArcInfo permits limited “read-only” OLE connection to Oracle tables. In order to store all data in Oracle and be able to use versioning, enforce access permissions, and enable editing via the desktop client, we must license, install, and configure the ArcSDE layer. The documentation for this product makes it clear that there will be a significant learning curve to become adept at its administration and able to make optimum use of ArcSDE features such as versioning.

We can work on building more tables of non-spatial configuration information. The IRMIS approach of working backwards from the control system software configuration files to try to identify and document our installed base of VME modules, GPIB devices, and the like seems promising. This non-spatial data can later be joined to geodatabase objects that store the spatial representation of equipment racks, connecting cables etc. that we hope to be able to extract from CAD drawings of service buildings, power distribution conduits, and so forth.

CONCLUSIONS

GIS software is used by utility companies as the foundation for managing the configuration of infrastructure every bit as complex (if not more so) to what we must contend with at Jefferson Lab. Our review of the overall capabilities of ArcGIS leads us to believe it is a viable toolset which at very least can complement non-GIS efforts to build an accelerator configuration management database. Our first-hand efforts have not yet matured to the point where we can test the functions which interest us most such as geodatabase versioning, or the possible use of geometric networks to model connectivity of hardware and software components.

REFERENCES

- [1] B. Porter, “Interfacing AM/FM/GIS with enterprise and Operations Systems.” Proceedings of 22nd Annual Geospatial Information Technology Association Conference, Charlotte, NC, April 25-28, 1999
- [2] C. Harder, “Enterprise GIS for Energy Companies.” 107pp., Environmental Systems Research Institute, Redlands, CA, 1999.
- [3] A. Hutton, “High Level Applications: The Way Forward.” TJNAF, Newport News, VA, Presented to Accelerator Controls Group, March 2004.
- [4] <http://www.aps.anl.gov/epics/irmis/Primer.pdf>
- [5] M. Zeiler, “Modeling our World.” 199pp., Environmental Systems Research Institute, Redlands, CA, 1999.
- [6] M. Zeiler, “Patterns for building Utility Data Models.” Proceedings of 22nd Annual Geospatial Information Technology Association Conference, Charlotte, NC, 25-28 April, 1999.
- [7] Carneiro, C. “The GIS of CERN: Specification and Development.” Proceedings of the 4th National Conference in Cartography and Geodesy, Lisbon, Portugal, 10-11 March 2005.
- [8] <http://www.lns.cornell.edu/~dcs/aml/>
- [9] Y. Roblin, “Integrated Accelerator Data Source Management.” Proceedings of ICALEPCS 2005.
- [10] S. Murphy, “Integrating CAD into an Enterprise GIS.” Proceedings of 23rd Annual ESRI International User Conference, San Diego, CA, USA, 7-11 July, 2003.