USING REAL-TIME AND HISTORICAL DATA THROUGH THE INTERNET AND ON THE DESKTOP AT MIT

CHRISTOPHER J. RUSSO
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
CAMBRIDGE, MA 02139
CRUSSO@ALUM.MIT.EDU

Keywords: process information, historical and real-time data, internet applications

ABSTRACT

The Department of Facilities at the Massachusetts Institute of Technology owns and operates a cogeneration and district energy facility used to supply energy to its campus in Cambridge, Massachusetts. Several years ago, MIT, using a combination of commercially available software and custom programming, integrated its control system with the Internet and made live and historical information available online to management, academic and research personnel. Additionally, MIT is in the process of connecting energy metering systems in all its buildings through the Internet, allowing closer monitoring of campus energy usage over a wide geographical area.

Some of the benefits realized have been new collaborative research arrangements with overseas and domestic colleagues on plant performance improvements, improved insight into plant operation and performance, new outreach programs with the MIT academic community and improved compliance with governmental metering regulations. The concept of a "virtual engineer" is becoming a reality with our use of distributed information.

While integration of plant data with the Internet and the computer desktop has posed new challenges, particularly in the area of network security, we have seen significant improvements in our operation and academic relations through new methods of distributing our data.

DESCRIPTION OF THE CONTROL INFRASTRUCTURE AND PI SYSTEM

The MIT Cogeneration Project was conceived in the early 1980s and brought online in September of 1995. There were two main reasons for the decision to construct the plant: first, MIT's energy needs were (and still are) growing rapidly, and there was a need for an economical, scalable and environmentally sustainable energy supply. Secondly, a stable and disruption-free electrical and thermal power supply was essential to the research conducted at MIT; previous years had revealed inherent inadequacies with the electrical and thermal distribution systems on campus and a near-total redesign of these systems was a part of the plant construction. With contributions from throughout MIT, the existing boiler and chiller plants were repowered with a gas turbine from ABB Stal of Finspång, Sweden (with a combustor designed in part by the MIT Combustion Research Facility) and an advanced heat recovery boiler, leading to overall thermal efficiencies in excess of 80%. In addition to the main cogeneration plant, operators also monitor (via remote computer control) a satellite chilled water plant on the opposite end of campus as well as the high-voltage transmission system for the campus.

With the installation of new plant equipment came a corresponding increase in the complexity of our control system; the amount of instruments that we now had available was an order of magnitude larger than what we had been able to monitor before. The control system upgrade presented several new challenges:

- Because many of our personnel are not physically located at the facility, there was a need for them to be able to remotely monitor and check on the performance and operations of the facility. Furthermore, this connection needed to be implemented securely to mitigate any possible risk of interference with plant operations.
- Increased process complexity demanded greater amounts of data to properly tune and design control processes and calculate plant performance. In particular, combustion tuning for emissions required that a large amount of data be stored for model analysis and simulation. Additionally, because MIT lacked in some specialized areas of expertise, efficient outside collaboration was a necessity.
- Regulatory requirements mandated that MIT keep close track of consumption and production totals for billing.
- The MIT academic community had great interest in the plant, and we wanted to make the facility a resource for them as well.

MIT considered several options for providing these services, including developing our own software, extending the existing control network across campus and several commercially available packages. We selected a commercial package from OSI Software of San Leandro, California, and began a collaboration to extend its capabilities and implement it at our plant.

OSI's PI[1] system was selected to fulfill our needs in June of 1997 and installed shortly thereafter. The PI system is in truth a series of applications that allow MIT-wide access to our data through the Internet. The first layer is comprised of the PI interfaces to our control systems. These computer programs communicate with our process control systems and transmit the data, via internal TCP/IP networks, to our PI server. The PI server then collects and stores the data[2] and makes it available to computers connected to MITnet and the Internet over the TCP/IP protocol. The third layer is the end user layer, where clients use various display, analysis and web tools to access real-time and historical data from the server.

By consolidating information from several control systems into one server accessible over the Internet, we have created a central "data warehouse" for historical process information with equal access to all users on campus. As far as end users are concerned, there is only one interface to all campus process data; through this type of integration, we are able to present a unified picture of the campus energy flow to our end users.

---

[1] PI stands for "Plant Information." More information on OSI may be found at www.osisoft.com

[2] Data storage in PI is not, as is often the case, stored in an relational database. Instead, the PI system uses a published algorithm to perform time series compression. This algorithm, which is essentially a variable sampling frequency, uses adaptive compression to achieve high amounts of data storage and high data resolution in relatively little storage space.

Figure 1 – Central Utility Plant Live Overview As Available Through Internet

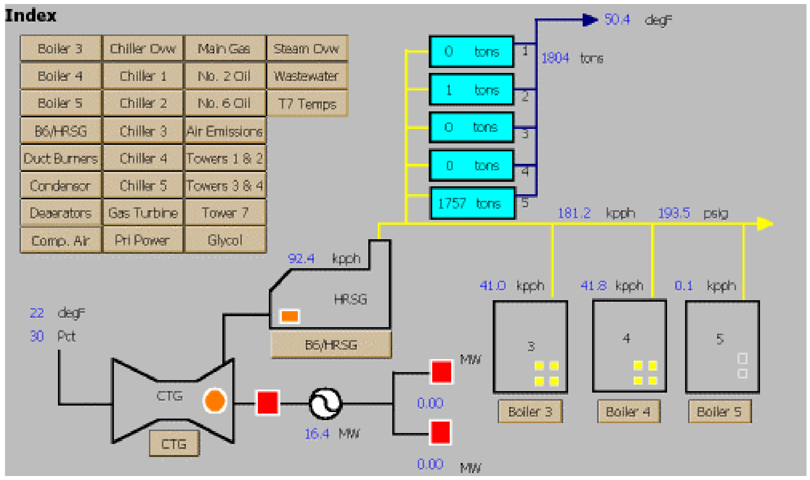

<u>USING THE INTERNET TO COLLECT CAMPUS ENERGY DATA</u>

In addition to our efforts at the Cogeneration Plant, the Department of Facilities has initiated a separate, parallel program to collect and analyze information from energy metering equipment located in different buildings around campus. The Department of Facilities manages over 90 buildings on campus, each with its own set of energy control hardware. MIT has utilized the Internet to connect these pieces of equipment to each other as well as to centralize data storage and distribution through our PI server.

The primary impediment to development of a system to collect and analyze data was not necessarily technical in the sense of combining systems, but rather logistical; multiple installations over a wide geographical area with high data requirements presented challenges in terms of physical connections. While methods such as radio transmission were considered for energy monitoring systems, MIT decided to utilize MITnet[3] to connect these systems.

Internet connectivity was implemented on two levels. Firstly, MIT leveraged the capabilities of newer PLC-type controllers to communicate via the ethernet protocol to literally make the Internet a part of our control network. In addition to this type of integration, we also utilized MITnet to connect each building controller to a computer that transmitted data back to the central PI server through the Internet. The primary advantages of this integration were not only performance or flexibility, bur rather the low cost and ease of implementation. At a location such as MIT, and elsewhere in the corporate world, the cost of connecting another ethernet cable to the network is miniscule next to the cost of a dedicated control network.

---

[3] MITnet is the portion of the Internet that is located at MIT. In this sense, MITnet *is* the Internet, and the terms are used interchangeably throughout this paper. MITnet has the unique advantage of being one of the original Internet networks, and thus network performance and reliability are exceptionally good.

Figure 2 – Central Utility Plant Network Topology

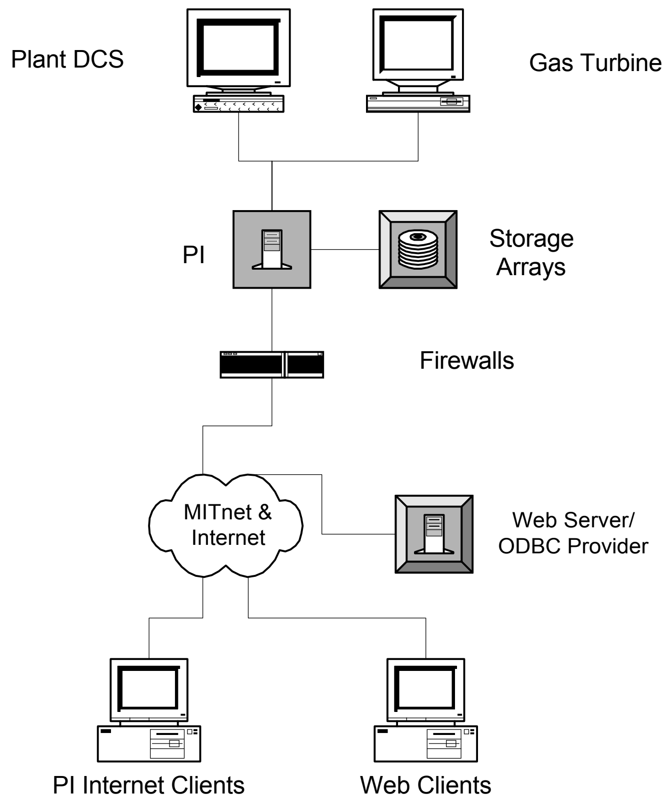

Development of the campus energy monitoring system network is still in the early stages. In the future, we plan to extend the capabilities of the system to use the more detailed data provided in order to better comply with regulations requiring accurate tracking of departmental energy usage. This will, eventually, lead to much more accurate energy usage billing and possibly time-of-use and variable energy pricing.

ACADEMIC OUTREACH EFFORTS

While the Cogeneration Plant is primarily tasked with delivering energy services to the campus, it has, over time, grown to become a resource for the larger MIT community. The academic community has been a part of the project from the start, from developing the advanced combustion techniques used in the turbine to representing the plant in regulatory hearings, and there has been strong mutual interest in working together to make the plant an educational resource. We have succeeded, we believe, in satisfying both the research and pedagogical interests of the Institute by making much of our plant data and information available through MITnet.

From the plant's inception, we have made an effort at outreach efforts to the community, including frequent plant tours and lectures to research and teaching staff members. Technical challenges and security concerns represented the largest challenges to our goals; MIT is a highly heterogeneous computing environment, and there is no one computer standard that permeates the campus, and so making online data available was a considerable challenge. In 1998, we utilized the ODBC[4] capabilities of our web server with new software from OSI and brought our online data to the web. With the use of some CGI[5] scripts and standard desktop tools, we have also allowed trends of data to be generated through web browsers.[6] While the available web data is textual-only at the moment, we are working on technology that will allow us to display live process graphics through any web browser through the Internet.

Our cooperation with the Department of Aeronautics and Astronautics has led to a productive arrangement where information from the plant is used to teach undergraduate engineers about general engineering concepts. The Unified Curriculum class is an class for undergraduates that covers several fields of engineering, from principles of combustion, to control system design, to materials engineering.[7] Because of the multi-faceted engineering challenges of a cogeneration and district heating plant, it is a perfect complement to their classroom studies. In 1998, the author was invited to work with the department to develop a series of lectures about the plant as well as assignments that involved real process data from the plant. Using the previously described mechanisms, we tied our web server into the PI database and assigned the students several problems which required that they look up and use data from the web.

While the biggest problem was confusion over units that are standard in the American power industry (students typically use a fully-SI curriculum), students generally thought the experience was a positive one. There were no glitches on the technical side, and the instructors were enthused with the results. This type of collaborative education continued again last year and is expanding to other departments in the Institute. Through web-based education, we can provide information to a large number of users regardless of their computing platform.

EMISSIONS MODELING, COLLABORATIVE RESEARCH EFFORTS AND "VIRTUAL ENGINEERING"

In response to stringent environmental requirements, we selected for the plant an ABB gas turbine with an advanced combustor, water injection system and catalytic oxidation systems. This turbine utilizes a method called premixed combustion that delivers very low emissions through sophisticated aerodynamic mechanisms and fluid injections. Advanced combustion, however, comes at a price, and the cost is a combustion reaction that is highly sensitive to control inputs and that must be tightly optimized to yield low emissions with an aerodynamically stable flame. The gaseous emissions are affected by up to as many as twelve different parameters, most with non-linear reaction kinetics. After we combined process data with standard desktop tools through the use of OSI software, we were able to use PC based modeling tools to perform advanced multi-variable data analysis on the combustion properties and reduce our emissions and simultaneously increase the aerodynamic stability of the flame, leading to greater reliability.

Gas turbines have sometimes been referred to as "R&D in the field;" while this may be an exaggeration, the rapid development cycle for gas turbines compels manufacturers to closely

---

[4] Open Database Connectivity Standard – a protocol developed to provide uniform programmatic access to databases. More information is available at www.microsoft.com

[5] Common Gateway Interface – a standard for programs to interact with a user through a web page

[6] This information is available through the web at http://cogen.mit.edu/unified/

[7] Commonly known as "unifried" because of its high demands, the class is nevertheless highly rated by students.

monitor and analyze field units, not only to ensure that they meet guarantees, but also to collect information for future development. ABB engineers would come to our site to collect data, communicate with their colleagues in Sweden running models, make changes to the control logic and return several months later to check on the effect of their modifications.

Figure 3 - Web View of Gas Turbine Data

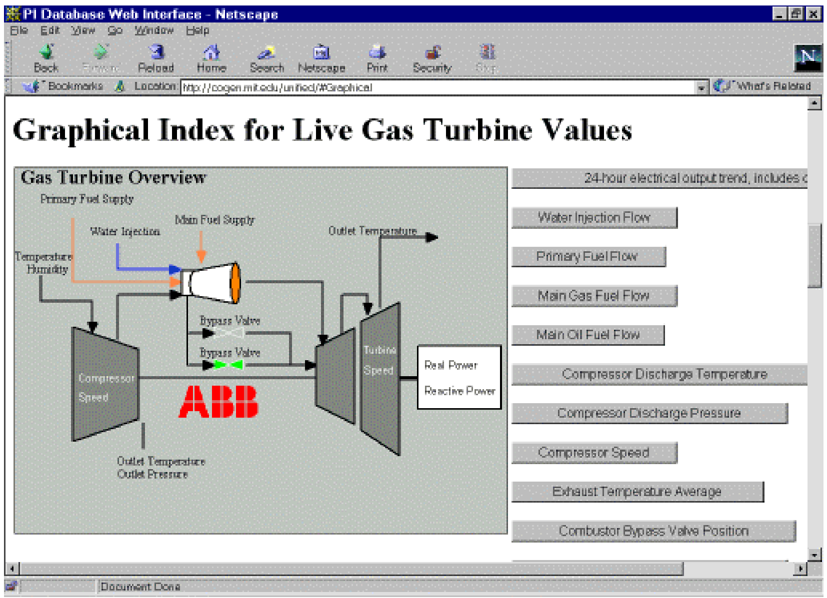

While the tuning methods were effective, they involved an enormous cost in terms of travel and time. Furthermore, there was little historical data available for the engineers when they arrived on site and the effectiveness of their analyses was limited by a relatively small sample size of data. Since the installation of the PI archive, we have distributed client software to our Swedish colleagues to allow them to see exactly the same data that we see in the plant. Using these client tools, they have been able to download the large amounts of data necessary to simulate the combustion process to their systems and execute these models at a considerable cost savings over the old method. Recently, our construction of a web interface to our data archive, , shown in Figure 3, has allowed even wider access to the data. Through this type of long-distance monitoring, we have been able to reduce the amount of time that ABB engineers have had to spend on site, while simultaneously increasing the performance of our plant.

In the last several months, we have also begun a project with a California-based consulting firm[8] to analyze the campus chilled water system in order to improve cooling performance. The hydraulic analysis necessary required a substantial amount of process data and the geographic distance made site visits difficult. After sending the firm a copy of the client software on a compact disc, they were able to log into our server remotely over the internet and have access to the same data they would were they to be sitting in the control room.

While the two examples above are notable for their technical details, the principal benefit of these trials may not be necessarily realized in their contribution to technical operation at MIT, but rather for their introduction of a new way of engineering. While the term “Virtual

---

[8] Rogers & Associates, www.rogersengr.com

Engineering" smacks of current buzzwords, its promise is that disparate resources can be brought to bear on a problem regardless of location. Especially for facilities such as MIT, where specialized expertise may be less widespread than for a energy-focused company, the idea of being able to collaborate more easily with these external resources promises great gains.

PERFORMANCE EVALUATION AND ANALYSIS

Through the use of distributed information, we have also been able to gain greater insight into our process through integration with various other tools. Traditionally, we had conducted plant performance analysis through manually taking measurements and entering them in prebuilt spreadsheets. With the installation of the PI system over MITnet, we have now been able to work with live data in order to more accurately determine performance and take corrective measures.

Of particular interest in our case was our ability to use a plug-in program to Excel to load live data into spreadsheet cells. By doing this, we could then integrate with our own calculations as well as make use of existing libraries for such things as enthalpy and psychrometric calculations. One of the primary performance indicators of gas turbines is a loss of compressor efficiency due to blade fouling; atmospheric particulates are ingested into the turbine and coat the compressor blades, disrupting the airflow and decreasing performance. Performance degradation is a subtle, but measurable process, and plants must determine when the performance increase to be gained by washing outweighs the maintenance cost of a unit outage. With increased integration capabilities, we are now able to open up prebuilt spreadsheets containing performance corrections and combine them with live data in situ to yield real-time performance and efficiency calculations. Because of the networked infrastructure, this information is available not only to engineers at the plant, but to all authorized users throughout the Institute and via the Internet.

In order to meet regulatory requirements for measurement of fuel and energy flows, MIT developed applications using spreadsheets and desktop client software to provide simple integrated totals of flows to different buildings and equipment. Through incorporation of these formulas into a spreadsheet template provided to users and through the use of a few macro scripts, users are able to calculate daily, monthly or yearly totals for energy usage, fuel flows and electrical usage in several seconds for any measurement on campus.

The advantage of using a network to link this data together is that there now exists one canonical source for process data; all data resides on our server and is dynamically generated for the user upon request. This significantly reduces older problems of corrupted and modified data that would often manifest themselves when plant data was passed from person to person.

SECURITY CONSIDERATIONS

While a highly accessible and networked data archive confers many benefits, the issue of information security is always present, especially at an institution such as MIT. Our plant computer security logs show frequent cracking attempts, from both within and outside of MIT. In order to successfully deploy plant data through PI over the Internet, we had to seriously consider the assumed risk and our responses to it. In truth, concerns about network security posed the greatest impediment to realization of our plan for universal data access. Through consultation with OSI programmers and MIT's Information Security office, we were able to develop and implement a security strategy. PI had never been deployed on an "open" TCP/IP network, and this posed a new challenge for us.

Our control system network is configured, at the hardware level, to allow read-only access to the PI system. While PI is capable of and is often used to return sophisticated calculations to the host control system, we allow our system to only read data from the control network. Furthermore, we employ several security measures within PI to allow most users limited access to different data values.[9] In addition, both hardware and software firewalls exist between our PI server and the Internet. Through packet-level and port-level filtering, we have developed a highly secure system. While one can never unequivocally declare that a system is absolutely secure, we believe that the enormous advantages of data deployment outweigh the much smaller potential risks.

FUTURE PLANS

Future plans for network integration center around providing more services to campus consumers of energy. While we have already realized our initial goal of providing "self-serve" information to MIT users, there is clearly much more that can be done. Web integration, in particular, can create multiple avenues for development in the area of customer services. Using web-based real-time data, we may be able to create an electronic mechanism to give campus personnel information about current and pending services and disruptions.

A particularly appealing possibility is the idea of using our process data integration results to help develop an elastic mechanism for energy pricing for our consumers. If sufficiently accurate methodologies can be developed to make the leap from process data to financial costs and this data made available through the Internet, one can quite easily envision, for example, a web page which shows users the approximate cost for energy at any point in time. Historically, energy suppliers have assumed an inelastic demand for energy consumers, but this assumption has been made in the absence of a mechanism for communication with end users. If time-of-use rates become electronically accessible to us, it is a relatively simple process to combine them with our unified data infrastructure and provide them to the campus.

Overall, our net experience of integrating our IT and control networks has been very positive. At our plant in particular and in the industry in general, the concept of an "industrial desktop" is increasingly blurring the lines between conventional desktop applications and process data. Providing information to our employees and customers is becoming as important as providing them with energy services; it may well be that the most critical wires in the energy industry may carry not the power, but rather the data.

[9] One of our most effective strategies to data has been to confine access to certain sensitive points to only certain network "subnets." A subnet represents a physical or topological Internet subdivision, and so can effectively limit access to certain points to, for instance, only users on the MIT campus or even users within a single building or department.