Preparing for Diagnostics from DDC Data - PACRAT

J. Jay Santos, P.E.
Facility Dynamics Engineering

Jeffrey Rutt, P.E.
National Security Agency

Synopsis

Trended Direct Digital Control (DDC) Data is a valuable resource. This data can be utilized to diagnosis common faults and problems with HVAC systems. This paper will discuss practical considerations owners need to consider in preparing for and implementing a diagnostic program that utilizes DDC Data.

Results will also be presented on two case studies where PACRAT (an Auto-Diagnostic Software Tool) was implemented. The first is a large pharmaceutical company in the Midwest and the other is the National Security Agency (NSA). NSA has been using the software for diagnostics as well as system performance characterization.

PACRAT is a database software tool that allows users to effectively utilize trended data to diagnose system problems, characterize and provide ongoing documentation of system performance, and conduct measurement and verification (M&V) analysis to protect investments in energy projects, commissioning, and contracted maintenance services. In short, PACRAT serves as a continuous commissioning tool.

About the Authors

J. Jay Santos, P.E., is a principal and co-founder of Facility Dynamics Engineering. Facility Dynamics focuses in HVAC controls design, HVAC controls master planning, HVAC Diagnostics and Building commissioning. He teaches HVAC continuing education controls courses (Direct Digital Controls and Building Commissioning) for the University of Wisconsin, North Carolina State, and the Energy Resource Station in Iowa. He developed a DDC User Guide. The Guide covers interoperability, architecture, hardware, and software issues in a standard format and generally assists owners/engineers design and select systems that meet their needs. He is also on the Editorial Board of HPAC Engineering and contributes to a column titled “Control Freaks”.

Jeff Rutt, P.E., is Chief Mechanical Engineer for the Facilities Department at The National Security Agency. He is responsible for technical oversight for all HVAC planning, design, commissioning and operations. He is currently serving as Vice-chairperson of ASHRAE Technical Committee 1.7, Operations and Maintenance.

Using Data from DDC Systems

DDC system capabilities and their associated data collection potential in particular have dramatically increased over the past 10 years. For example:

- DDC systems are much more powerful and facility management personnel are looking to take advantage of this potential.
- Data storage costs have decreased dramatically, removing a barrier for NOT collecting and storing trend data.
- New protocols and standardization practices help make data easier to share.
- Sensors commonly used in HVAC-based control systems are lower cost, more robust, and higher accuracy than ever. Cost for accurate controllers (those with A/D conversions of 10 bit and higher) is dropping and their use is becoming more prevalent.
- New metering technologies – including power, air and liquid flow, and thermal energy - are available, more affordable, easier to implement, and increasingly reliable and accurate.
- Communications capabilities – such as the Internet, Intranets and wireless technologies - allow for real-time transmittal of facility data to those who need it.

Providing Data *Interpretation* – Moving Beyond Data Visualization

There are a large number of services and products available to provide various levels of data visualization. The challenge then becomes achieving timely, accurate, and cost-effective analysis and corresponding interpretation of building data and performance metrics. Ideally, one would like to match the capacity for data analysis and interpretation to the speed and ability to gather the data. A common complaint expressed by many facility managers is that they have plenty of resources that can deliver more data to their desks in a variety of striking formats and graphics, but they lack the resources to determine what the data actually says, and – most importantly – what specifically to do about it. This is not surprising given the relative shortage of qualified people with the required skill set to provide these analyses. Finding resources that have the necessary backgrounds to do HVAC diagnostics is difficult. A good HVAC diagnostician must possess a working knowledge of HVAC systems and applications, facility operations, control systems design and implementation. They must also have the ability to rapidly look through data from multiple systems and find problems at both the whole-building performance level and also diagnose down to the component/device/control loop level.

Preparing for Diagnostics

There are two major steps in preparing for continuous diagnostics when using PACRAT or any other software tool that uses trended DDC data. The first involves setting up the DDC system to trend and save the required data to be used for the analysis. The second involves gathering information about the system's configuration, design parameters, and energy attributes in order to make a meaningful interpretation of the data. As one goes through the process of setting up trends on equipment controlled by the DDC system, there may be practical limitations on how much data can be collected. This may or may not be an issue at any given site. The trending capacity of a given system will depend on the specific manufacturer and version of system as well as the network architecture of the system. Unfortunately, the system architecture of most systems is dictated by the installing contractor and not designed by the engineer.

DDC System Architecture Issues

Many DDC manufacturers make controllers of two distinct levels. Some make only one. These "levels" refer to where the controller resides within the system architecture on the control network. Knowing the difference between these controllers is important because the appropriate controller is application-dependent. Many specifications do not distinguish between these various types of controllers. This becomes an important issue relative to trending capability as well as fundamental control of the various systems.

Typically, higher end controllers live on a higher-level network and communicate in a peer-to-peer fashion. I will call these controllers "primary controllers". Peer-to-peer means that the controllers can share information to other peer-to-peer devices without going through an intermediary device (which I will call a supervisory interface). These primary controllers can range in cost from $1,500 to $4,000 and more. These controllers have more memory, more sophisticated CPU's, higher resolution A/D converters, more accurate clocks, and can store more complex control strategies as well as trends, schedules and alarms. The critical issue for the use of diagnostics is their trending and storage capability.

Manufacturers also make less sophisticated controllers that typically reside on a lower-level polling network. These controllers have more limited memory and processing capabilities, and must utilize a Supervisory Interface device to communicate with all other devices. They come in a number of flavors and can cost from $100 to $1,000. Some are designed for typical terminal applications like VAV boxes or Fan Coil Units. Others may be used for air handling systems with simple to moderately complex sequences of operation. These terminal controllers are typically configured for the number of points required for that application. Some of these controllers utilize a free form of programming (they need a complete set of custom programming) while others have application-specific programs for typical applications. These programs have selectable parameters that can be set-up for each individual application. Since these controllers have more limited memories, they typically do not store historical information (like trends) and rely on the supervisory interface for this function. The secondary polling networks are configured such that one supervisory interface can monitor a limited number of controllers (See Figure below). This limitation varies by manufacturer. A large number of controllers on a secondary controller network can negatively impact the number of trends that can be practically utilized and the amount of data that can be processed and its speed of transmission over the network. How many is too many on this secondary network? Well, that's the classical answer – "It depends". It depends on the manufacturer, the speed of their network and the application in question.

Since many specifications are performance based, either level of controller hardware can meet this type of specification. The Engineer should indicate which controller is best suited for the application. A Pentium III-type processor is preferred for a computer running a CAD program, while an older 486 may still be adequate for running a basic word processor. In HVAC controls, these lower level controllers may be adequate for many simple applications, but a primary controller is appropriate for more critical applications. How does one specify these distinctly different controllers? First, the Engineer must define requirements of various types of controllers and their corresponding interfaces (both network and operator). Once defined, the Engineer can dictate which controller should be used on various applications.

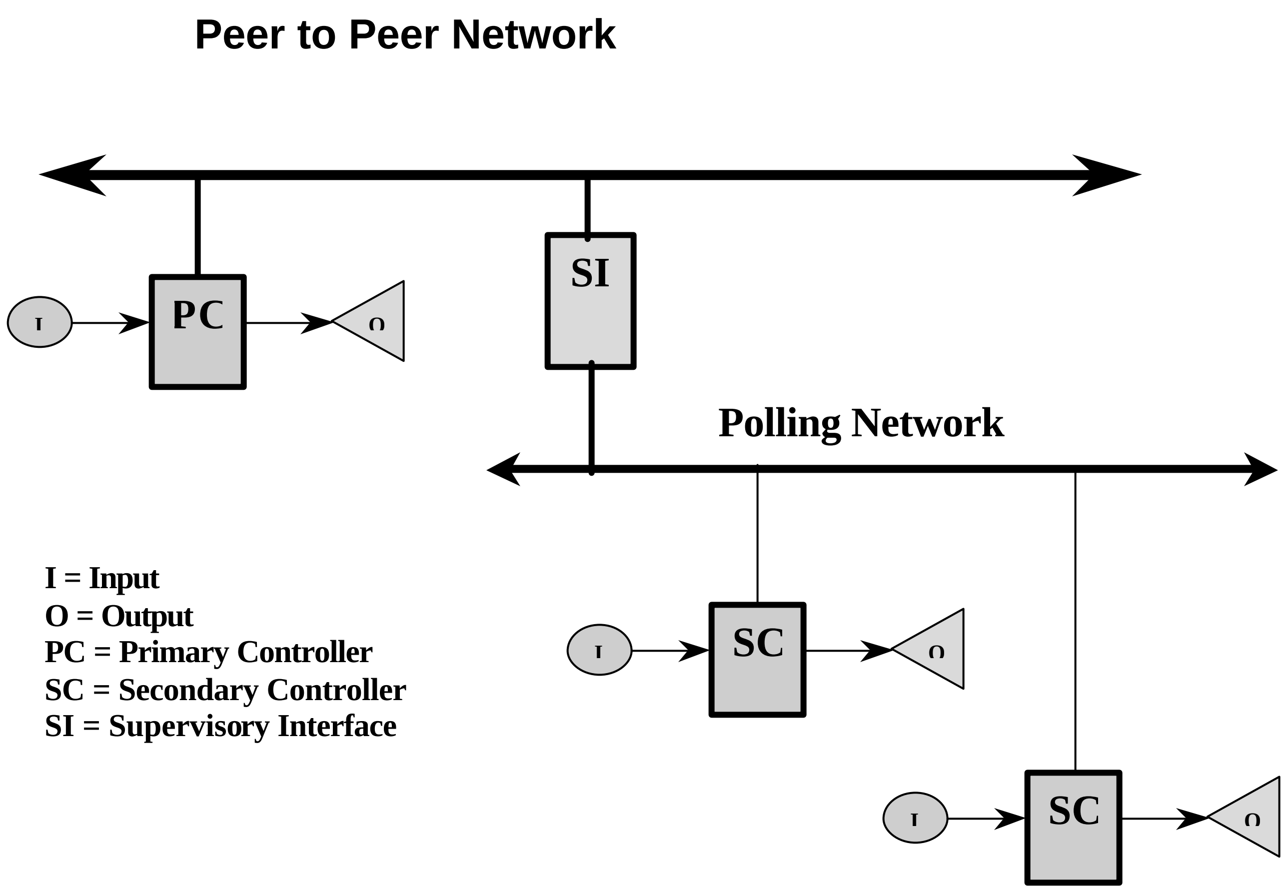

What if the vendor does not make "primary controllers"? There are two options. If in the Engineers' opinion, a primary controller is warranted due to the sophistication of the system, the specification could disallow all but these peer-to-peer devices. Alternatively, the Engineer could allow a dedicated Supervisory Interface for the secondary controllers required for the control sequence. Sometimes, depending on a manufacturer's product line, it is necessary to use more than one controller to accomplish a given complex sequence of operation. This forces the application Engineer programming the sequence to split up the various control loops on a given system. This could result in the control loops for the starting and stopping of the fans and the VSD to be on one controller and the control loops for the mixed air section, cooling and heating coils to be on the other. The potential problem comes from the polling nature of the secondary network. If one controller needs to send information to the other controller, it must first send it to the SI, which then forwards it to the other controller. Any loss of communication between these devices would interrupt the proper operation of the controls for this system. The Engineer must determine whether this is acceptable for the application. If the vendor is allowed to install a large quantity of units on this secondary network, or implement a lot of network heavy functions (like trending) this could slow down critical control communication. For example, if the output to a variable frequency drive is connected to one controller and the controlled variable input is connected to another control, the time delay in transmission may render the loop difficult or impossible to control.

Bottom Line: The design Engineer needs to be more specific in dictating controller type for each application. To accomplish this, the Engineer must possess detailed knowledge of the various available control products. If the Engineer has not kept up with the DDC industry, this is no small undertaking.

In an effort to make this information more readily available, the web-site DDC-online.org was created. This website presents the 20 different product lines in a generic "ladder" by the classification of the device. This ladder lists generic layers, which are labeled on the right hand side of the typical system architecture. The proprietary products are then placed on the appropriate layer. From this system architecture drawing, the user can double-click on a device and it links to a generic cut sheet appropriate for that product. This allows the user to more readily do a comparison of similar products.

Storing Control Information ('Trending') on the DDC System

The capabilities of the different vendors' DDC system vary considerably with regards to trending capabilities. Unfortunately for those attempting to interpret this data, there are currently over 20 DDC systems currently on the market, and almost as many schools of thought by the developers regarding trend programming and its attendant data archiving. Trending features and archiving capabilities have historically been a lower overall priority in the controls industry. It is not uncommon today to find controls technicians that are unfamiliar with their own product's trending capabilities and associated data archiving routines.

Common issues to be considered when designing trends for diagnostic purposes include:

- Selection of specific points to be saved/stored by the system;
- Impact of trend activity on control and communication functions;
- Data time interval required; and
- Robustness of the analysis method being used.

And naturally, each of these issues in fact depends on the others.

In our continuous commissioning and automated diagnostics work with PACRAT, Facility Dynamics has found that "if in doubt, trend it". This approach works given the relatively low cost of data storage and the capability of PACRAT to provide data visualization of all parameters (even those NOT used for diagnostics). Virtually all sensor inputs and control outputs (valves, dampers, on/off commands) are trended and most are used in the analysis. Statuses and setpoints are also useful information and should be trended if they are dynamic in nature (variable and/or calculated). Points of little/no value for analysis include alarm points and maintenance points (such as filter differential pressure and run time).

Regarding the selection of a trend time interval, one needs to consider the characteristic speed of the processes or the rate of change of the variables being analyzed. PACRAT analyses have found that for most common HVAC processes, a 15-minute time interval was sufficient to identify anomalies when analyzing results across a 30-day period. Shorter intervals may result in slightly improved anomaly identification when analyzing on a more frequent basis (say daily or weekly), but would likely do so at the expense of control system functionality and increased analysis time. One-hour data intervals are used by some PACRAT clients with good success, particularly when analyses are being carried out over a period of several months worth of data.

We have found that the implementation of a diagnostic continuous commissioning tool can easily justify the additional cost of sensors not normally associated with the necessary control of a given system. For example, this can be utilized to justify an additional sensor to read "between coil" temperature that is a monitor only point. On an air handler of relatively small size, one can justify the cost of this point the first time the software diagnoses a leaky valve.

PACRAT Case Study at Pharmaceutical Company

PACRAT was tested for a 6-month period at a pharmaceutical campus in the Midwest on 8 air handling units of varying types and applications, totaling about 230,000 cfm. This particular client sought to closely document the anomalies identified by PACRAT versus a field investigation to determine the number of successful reports of problems and to determine if campus-wide use of PACRAT would be cost effective. This case study specifically involved only the automated diagnostics portion of PACRAT, and did not include performance verification, data visualization, or M&V functionality.

Results of Anomaly Tests

The overall accuracy of reported anomalies was found to be 91.6%, and totaled 93 individual anomalies found over the 6-month period. This accuracy rate was 86.4% if one included the anomalies found by the evaluator that PACRAT does not presently diagnose. In addition, the evaluator found the additional information provided – consequence of the anomaly, causes, and potential resolution – to be “accurate in nearly all cases”. The client did not analyze the findings for the accuracy of the energy waste calculations.

Accuracy of reported anomalies was categorized as follows:

- Correct Analysis - reported anomaly was correct
- Leading Analysis - reported anomaly finding was partially correct but easily led to the real problem
- Incorrect Analysis - reported anomaly was not correct or could not be verified by evaluator
- Incorrect Setup - reported anomaly was due to incorrect setup of PACRAT parameters during system configuration
- Not Found by PACRAT - problems that PACRAT is not presently designed to analyze that were found by evaluator

The following table lists the number of anomalies found for each category, and the accuracy percentage of each category:

PACRAT Anomaly Category	Correct Analysis Count	Leading Analysis Count	Incorrect Analysis Count	Incorrect Setup Count (Note 1)	Not Found Count	Percent Accurate (Note 2)	Percent Accuracy Excluding Setup Errors (Note 2)
Simultaneous Heat and Cooling	3	3	1	0	0	85.7	85.7
Sensors	25	7	1	3	0	88.9	97.0
Valves	14	4	0	0	0	100.0	100.0
Stability	8	0	1	0	0	88.9	88.9
Occupancy	2	0	0	2	0	50.0	100.0
Economizer	1	0	0	0	0	100.0	100.0
Space (Return)	4	4	4	0	0	66.7	66.7
Uncategorized	1	0	0	0	5	16.7	16.7
Total (93)	58	18	7	5	5		

Note 1: Setup errors were due to: Inaccurate occupancy hours configured (2), incorrect sensor characterization during configuration.

Note 2: “Accurate” defined as either “Correct Analysis” or “Leading Analysis”

A detailed cost/benefit analysis was performed by the client to determine whether to proceed with further PACRAT deployment. These results are not available, but the client has elected to proceed, extending the types of equipment to be analyzed to include chillers, chilled water pumping stations, and primary/secondary hydronic systems. They will also be including the performance characterization tools and M&V functionality in their deployment.

PACRAT Case Study at National Security Agency

As Government downsizing continues to be the norm for the future, the need to optimize and maximize performance from automation systems as an EMCS will combat declining maintenance staffs and resources. The Facilities Department at The National Security Agency and Facility Dynamics developed this diagnostic methodology and applied it, with great success, to an EMCS controlling 34 house air handlers in a 1,200,000 square foot, multi-use Government building at Ft. Meade MD. This program has identified over $240,000 of energy waste during a 12-month period, by identifying faulty control components and poor control sequencing. Approximately 40% of the energy savings were captured with just a few days of technician time. It also documented significant over-sizing due to "liberal" design assumptions. These design criteria were subsequently modified resulting in $800,000 in cost avoidance on a project under design. In addition, the program uses carbon dioxide and temperature measurements to calculate fresh air quantities to ensure ASHRAE 62-99 ventilation requirements are met.

This methodology allows building operators and engineers the opportunity to perform an examination on the health of their system; eliminate waste due to faulty components and/or programming; and a reality check or optimization of mechanical systems. Typically, design calculations represent a conservative approximation of equipment size and performance. Based on extreme weather data and building usage, the heating ventilating and air conditioning systems operate at part load most of the time. Using EMCS data, building operators are able to predict part load performance and load trends, therefore optimizing energy efficiency and system operation.

Specifics Results for NSA

- Performance characterization module saved $800,000 in construction cost due to oversizing the cooling system. Original system was sized for 10 watts/ft^2; PACRAT determined actual space loads of 2 to 4 watts/ft^2.
- NSA considered installing economizers as an energy saving project. PACRAT identified actual cost savings of $30,000 far below any savings identified by engineering models. Avoided a $2,100,000 construction project. Actual payback would have been 70 years.
- Chiller performance characterization module determined actual chiller efficiencies, which resulted in an energy saving project that interconnected five chiller plants. The resulting installation reduced the number of chillers that operated at part load from four to one. This resulted in an annual energy savings of $400,000.
- Used this chiller data to determine actual operating cost for the chilled water system.
- Used the anomaly reports to prioritize maintenance so that limited resources were used most effectively to repair/replace or calibrate the most critical defective sensors and actuators.
- Used the chiller data to determine actual chilled water loads and compare these to outside air temperatures. This information is used to assess maximum system demands, compare it to system capacity and projected future demands.
- Currently using outside air calculations to validate air handler operation so that ASHRAE 62-99 standards are met.
- Used chiller data to determine that the chillers operate most efficiently at full load. Energy saving strategies such as chilled water reset and condenser water reset had very little impact on chiller efficiency.