THE LANGUAGES OF BUILDING CONTROLS AND AUTOMATION, UNDERSTANDING THE ROOT CAUSE OF CURRENT DESIGN SYSTEM LIMITATIONS

Alan F. Stewart, P.E.
Stewart Engineering Services Corp.

Abstract
The paper traces the early conversion from pneumatic controls to direct digital controls. The first facility automation programs for facility control were written as direct analogues of the pneumatic control systems design. The paper demonstrates how the language then became the limitation on the operation and design of facility control systems. Working examples demonstrate the existing shortfall in facility automation strategies and propose open-ended solutions. The analysis is typical for all automation systems.

Overview
The purpose for this paper is to discuss the current state of control and automation systems. The goal is to have you think differently about the expectations you have for facility automation systems. To accomplish this task we must elevate system controls and automation to a higher level of performance by raising our expectation of these systems. I will discuss the control operational languages, current design practice for control systems and then discuss alternative strategies that more fully utilize current computer capability. Before we address control system strategies, we should discuss the language of control.

The Languages of Control and Automation
What are control system languages? In the most simple terms control system languages are very similar to our own languages of the earth. In early times we had multiple spoken, graphical, pictorial, and written languages. The languages of control systems struggle with simplicity for fast configuration or programmability for the richness required for complex systems. These languages come in graphical, pictorial, and text. We should review how we evolved to the systems of today.

Before the computer systems of today all complex controls were accomplished with pneumatic control systems. I define complex, as the level above simple start/stop or off/on. The pneumatic systems offered a method to convey analog data, such as space temperature or humidity as a pressure typically between 3-15 PSI. Further, pneumatic controls gave the only practical way to develop power to operate large valves and dampers. Complex air handler controls needed to ensure proper operation for multiple rooms served by a single air handler. For all these reasons and more, we developed pneumatic devices to work the logic of control. There are switching relays to choose a night temperature sensor, or load analyzers that have several room sensor inputs. And choose the high or low value to output. The point here is not to offer training for pneumatic controls but to allow the reader to realize that the pneumatic devices of control constitute a language.

Why would I refer to pneumatic devices as a language? The pneumatic devices offered the set of all possible choices for the controls designer to select a system of components, for control of facility air conditioning in this case, to operate in a predictable way. Our language both inspires and inhibits our ability to do work or comprehend ideas. Case in point, as we develop new ideas we develop new words to describe the idea, computers, cars, telephone, and others. There are two key points; our language limits our ability to communicate (and possibly think), until both the sender and receiver have common language. And secondly, the language is in continual flux, adapting to the changes in our learning. Both of these points have impacted control and automation evolution. As the control knowledge resided in the pneumatics specialist, and the electronics knowledge resided in the computer specialist.

As direct digital control (DDC) systems became available, the electronics offered complete freedom to operate facility systems in any matter. Any output device, say an air handler could be programmed to operate under several strategies. However, the people with programming expertise knew nothing about facility systems, or how they should work. The people who understood facility systems knew little about computers. To bridge this impasse the computer programmers built an intermediate language for control system designers. This language allowed the control system designer to develop the control algorithms as with pneumatic design. This approach allowed the control system designer to develop the required control strategies, as before, and implement the strategies on the new platforms with little training. While this method allowed the control system designer to complete his work at acceptable labor cost, the true value of the computer platform was lost. The language of control prevented the quantum leap in new capability.

New capabilities implemented, included the ability to learn from improper control response, although most systems struggled with this concept. Improper control response is defined as not meeting the required conditions for a control system. Some simple routines that were implemented included the start of building equipment at variable times before occupancy to allow proper temperature conditions when the personnel arrived for work. Not addressed was the ability to conduct facility operations in economic terms. There were no energy consumption or economic dispatch, the ability to insure that multiple machines are operated in optimum efficiency, terms in the new language. And today, some twenty years after digital controls the new standards for interoperability still focus on the hardware points, sensors, operators and such. Shipping very low-level data about hardware points. Imagine you are working in a campus environment with 30,000 hardware points, the automation system should be able to respond to problems, start a second chiller if the primary fails. Then report to the system operator that the system handled the problem but the technician is needed to look at the failed chiller. The point I'm making is the system must act and then report. Many automation and control systems in place telemeter the data to the operator, observe and report, alarming an out of tolerance condition, and requiring the operator to take action to solve the problem. The extra workload on the operator prevents doing the higher-level management review functions of how the campus is operating.

An additional example, the facility has 40 rooftop packaged air-conditioning units all are 20 tons each. The loads in each area are comparable. By measuring the parameters from the units, supply air temperature, return air temperature, and kW. We can establish norms for the entire population of 40 rooftop units that begin to predict such things as compressor failure by observing real time behavior out side the reasonable values. The direction here is to have the machine (automation system) do all the routine analysis and leave the human operator to martial resources for those items the machine cannot fix. What about the graphical language solutions?

The graphical language solutions suffer the same fate as above. The graphical language blocks allowed only the same types of physical devices, as did the written language. Additionally, both languages suffer from single loop analysis. The single loop concept for control has one or more inputs and outputs. Time flow is generally left to right on the system schematics. For simple systems, a chilled water valve operation in response to the leaving air temperature of the air handler, this linear time flow, left to right on the drawing, is fine. Complex algorithms controlling multiple chillers in a large central plant need to have both predictive qualities, when to start another chiller, and statistical data to analyze the operation of chillers on line. These types of recursive decision making strain the capabilities of graphical languages by violating the left to right time flow.

Progress is occurring in the control world. Most vendors have embraced a multi-level product approach. From single loop controllers, sometimes refer to as application specific controllers, to progressively more capable controllers programmed by a wide array of languages. The control languages of today, are similar to general purpose programming languages with the added facility for HVAC specific devices. All of these evolutions have helped simplify the platform, however, as we'll discuss later the full opportunities are untapped.

One last comment on the language barrier. A number of groups are hard at work developing standards for communication among the many vendors of equipment. From the topics we have discussed there are multiple problems to be solved. Passing numerical data and physical point names are fairly achievable projects. As we discussed graphical languages, pneumatic root languages, and newer multilevel decision making all add to the challenge of true inter-operability. The control systems are thought of as real property with permanence and longevity. Perhaps we may have to recognize the age of disposable controls is here, in order to embrace the openness we profess to need.

System Wide Control Solutions

The following is a example of control and automation opportunities existing today. In a typical office-cooling environment we have the following components within the HVAC system, beginning in the office and rejecting the heat ultimately to a cooling tower.

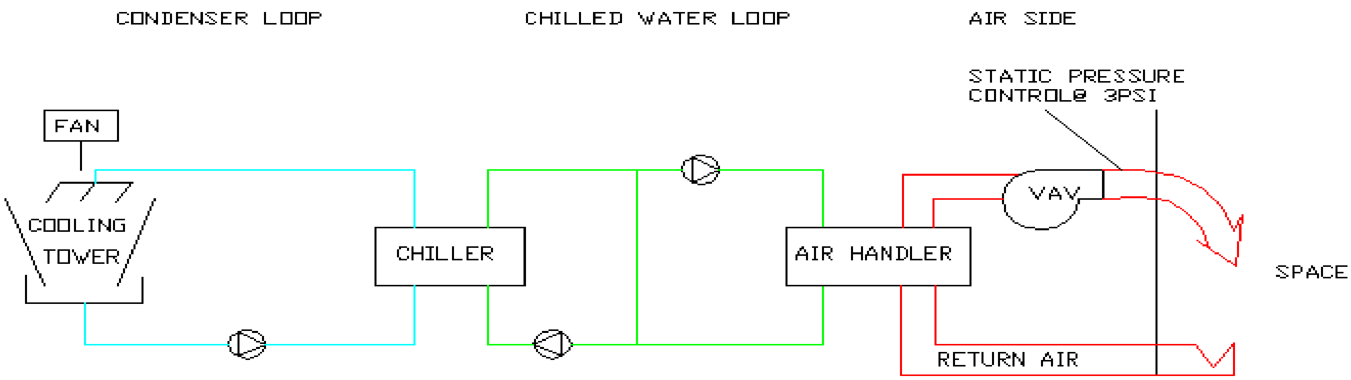

Cooling Tower through Cooling the Space
Working backward in the heat rejection process, from the tower to the space we have the following:

The cooling tower functions by passing water over the fill material, designed to create the greatest water and air contact, and into the tower basin. The tower fan will be operated to control water temperature in the basin to some desired level by forcing air to flow up through the tower. Even without the fan the hot water from the chiller will cause some level of convection airflow, due to the temperature differences. Not shown in the example drawing is the condenser water bypass valve that can divert water into the tower basin in cold weather and low load. Freeze protection is provided by tower basin heaters.

The condenser water pump operates the condenser water loop. The pump is selected to move the design quantity of water, gallons per minute, (GPM) through the condenser water barrel in the chiller. With safety factors the pump can be over sized.

The chiller compressor is sized to transfer the heat from the evaporator to the condenser. This compressor operational working range is usually defined over a narrow set of operating curves, specifying entering water and leaving water temperatures. Very cold condenser water generally limited to 10 degrees F above chilled water (verify requirements from the chiller manufacturer) can cause internal refrigerant flow problems for the compressor.

The evaporator water pump or pumps are designed to get the chilled water to the air handler. There are many possible configurations for the chilled water design.

The air handler is designed to convey the airside heat to the chilled water, and to supply the cooled air to the space.

There are many design variations here also. In this example we consider variable air volume (VAV) as the airside control. The air handler may use a variable frequency drive, inlet air dampers or outlet air dampers to vary the airflow. Typically the air handler controlled to provide a fixed pressure output to the ductwork. The pressure having been established in the design analysis.

The airflow is conveyed to the VAV box, which based on the local sensor will cool as necessary to meet the temperature set point.

Each of the above system components function as an autonomous elements, controlled by individual set points. Let’s go back through my simple example and determine how many control loops exist:

Cooling tower-
- Divert water through the tower or to the basin (1)
- Operate the fan speed(s) (1)
- Operate the basin heaters, if freezing weather (1)

Condenser water pump(s)
- Pump with pressure verification, on/off (1)

Chiller
- Condenser water flow, for chiller shut down
- Compressor inlet vane control to maintain chilled water temperature (1)
- Chilled water flow, for chiller shut down
- Current limit for motor protection

Chilled water pump(s)
- Pump with pressure verification, on/off (1)

Secondary Chilled water Pump(s)
- Pump with pressure verification, on/off (1)
- Variable Flow (1)

Air handler

- Chilled water flow through the cooling coil (1)
- Fan start/stop (1)
- Fan pressure control set point (1)
- Minimum outside air for ventilation (1)

VAV box

- Flow control based on space set point (1)
- May have minimum flow set point with reheat (1)
- Space humidity (1)

Recognizing this to be a simple system we have about 18 control loops total. For those of pure heart we will only consider the loops marked with (1) as controlling energy. Even in this simple example we have 15 loops that all have an impact on the energy consumption for the entire system. Further I submit that if the information available to the individual control loops were available to all 18-control loops the decision-making would further reduce facility energy consumption.

Examples of System Opportunities

1. The air handler fan, for a typical VAV system, is controlled at a preset discharge air pressure (design pressure). The design pressure was derived during the cooling/heating load calculations. And the peak flow may only occur a few hundred hours per year. The balance of the annual hours, the fan wastes energy. In fact, pressure is an alias variable, airflow is the requirement. The airflow (pressure) can be controlled to meet the cooling/heating needs of the VAV box with the greatest load. I have used this technique effectively in many facilities (subject to minimum throw for the terminal diffusers). To avoid conflicting control loop activity between the fan speed and VAV box, I assign the worst-case space temperature plus one degree as the set point for the fan speed control. That way the worst-case VAV box is fully open and I drive the system air pressure from that space. The other spaces operate the VAV normally. A second benefit of this approach is the reduced effect of noise created by throttling the central fan airflow through the VAV box. Savings from this technique can be over 50% of the total fan energy.
2. The chiller supply water temperature can be reset based on knowing the actual sensible and latent load requirements of the space. If only sensible cooling is required then the chilled water temperature can be balanced against the airflow required to meet zone-cooling requirements.
3. Condenser and chilled water loops can be flow/delta temperature controlled, within limits, to reduce operating cost. Typical water flow limits are approximately 80% of rated flow. In practice, the water side flows are 10-50% more than rated, consuming 1-3% more of the total cooling system energy.
4. Actual performance can be measured and compared to manufacturer specification for chillers, pumps, air handlers, cooling towers and other components. For example, if some of the cooling tower nozzles are fouled the performance degradation would show up over time, as the tower would not continue to meet peak load performance. Using the methods of today the tower will degrade until even running the fans on high will not meet condenser water set point.
5. Even overloaded air handlers can be aided by developing lower chilled water supply temperatures. The reason for the overload should be determined. This approach can solve seasonal air handler overload due to ventilation cooling loads, at least until permanent solutions are developed.

Two inter-related principles have been discussed above. The existing approach of multiple single loop control where the loops are not working together has tremendous energy savings opportunity. Secondly, shortcomings in the existing subsystems may be overcome by operating the system as a single loop reset with data from each subsystem. In practice other sub-systems can cover some or all, of the shortfall from parts of the larger system. Some examples should help.

Software Solutions to Hardware Problems

1. Office Renovation- Many office environments have been subdivided to smaller offices. This is done without regard for HVAC zoning leaving some offices hot and others cold, the only satisfied person has the thermostat. It is possible to add dampers and thermostats to some of the small-overcooled areas (less than 25 percent total air flow for DX cooling to avoid coil freeze up) and pinch down on flow to minimize overcooling. All of the minimum airflow criteria should be evaluated for heat and outside air. My experience is that just having the lights on will provide enough heat-load to accommodate minimum airflows. If the air handler is on chilled water the automation can balance leaving air temperature, total airflow, and input from the new zone controller to provide fairness to all occupants.

2. Process load- There is a new process being installed that has 100 tons of cooling load from the 100 percent outside air handler. The piping for the chilled water is short as the chiller is located near the air handler. You calculate the chilled water capacity at 200 gallons. Checking with the chiller vendor they recommend the minimum chilled water quantity at 1000 gallons. The chiller vendor says the chiller takes three minutes to unload from full load conditions. Based on the information you now have if the chiller is well loaded and the process is shut down, the cooling load could swing to almost nothing. And the chiller would cool the mass of the water you have to the

freeze limit shutdown before the compressor can unload. The options are:

- Install a tank for 900 gallons of water to assure time for the chiller to unload.
- Go to a Glycol solution to avoid freezing the tubes in the chiller
- Monitor the process loads so the chiller can unload safely when the process is unloaded.

3. Demand Load Shed- The opportunity to shed loads for dollar savings are few indeed. Early in my military career I had identified some facility categories considered, by me, to be of lesser importance in the grand scheme. Such as air-conditioning for facilities not occupied, at least not to 78 degrees F. When I shed those air-conditioning loads during a couple of hot August days, I soon found one of life's little secrets, "There are no insignificant facilities." Seems even though the base theater was not in operation, nor any personnel inside the candy bars were at "parade rest" waiting to be called to service. By shedding the cooling to the theater the candy bar "parade rest" became more of an "at ease." That was a valuable lesson in Load Shedding 101.

4. Demand Load Shed II- Next time I looked at load shed as an opportunity I had more respect for the tool. Turns out I had a process load of 200kW that ran randomly to pump resin into feeder hoppers from the main storage. Typical times of operation ran from 10 to 20 minutes two or three times per day. This industrial plant had a large chilled water system. I reasoned that I could monitor peak load with the energy management system and when the resin transport began I could unload one of the 450-ton chillers. After a test, the program worked like a charm allowing the chiller to recover the chilled water system temperature. As memory serves me the excursion for the chilled water was between 2-3 degrees F. Oh, I never lost another candy bar to heat stroke.

5. I was working the hydronics solution to a large chilled water plant. The facility was originally planned to be the fist phase of a facility with the second phase doubling the total cooling load square footage. As luck would have it the second phase never came into existence. Without going into all the analysis it came to pass that I identified some automatic butterfly valves that would be critical to have operating a semi-synchronous fashion. Trouble was the valves were 20" and 26" in diameter and the flows were around 4000 GPM. At first I considered paralleling the valves with appropriate sized valves for this flow, but there was no space available. It occurred to me that I could software range the valve so that they only opened enough to appear like the small valves. This solution was part of a larger software scope of work, but I had struggled as to how to solve the valve flow problem. It now seemed quite simple to me, cycle the large valve from closed to partially open (but with the slightest pressure drop across the partially open valve). I reasoned this to be the range of operation for my new algorithm. Interestingly, when I asked the programmer of the facility site automation to program the valve for this operation, he informed me "It won't work." I asked where he had tried this solution, to which he replied he'd never heard of this approach but was sure it would not work. Actually, I had not used it on water valves although I had employed this procedure on steam coil valves to stop the over shoot in a critical laboratory environment for the preheat coil in the outside air. Wind up, as it is called in control parlance, is the travel of a controlled device not producing desired effect, and when the direction of travel is reversed the PID algorithms drive too far before the system can recover. After we determined the new full open "Soft-Stop" (software limit) for the valves and recognized that flow to percent open would have a new curve the solution worked quite well. The huge valves now performed as if 12 inch diameter valves and at a large cost savings.

Summary

The intent of this paper is to encourage automation platform use in ways not possible a few years ago. In some respects we have moved from pneumatic controls to DDC making the controls better but not exploiting the new potential. Some of the problem may rest in that the software developed for the DDC emulated the capabilities of the pneumatic component, restricting growth of the computational value of the facility information. Lastly we must require more supervisory functions be imbedded in the automation platforms. We can make good management decisions when we have data in a useful format.

About the Author

Owner, Stewart Engineering Services Corp, develops Solutions Engineering projects for a broad range of customers desiring to reduce facility ownership costs through technology and technical services. A ten-year veteran of Honeywell Inc, he developed a wide variety of facility and process solutions for fortune 100 companies. He is a Honeywell President's Club member and recipient of the Johnson and Johnson Quality award. He served twenty-one years in the U.S. Air Force in a variety of engineering and technical roles both in flight simulation systems and facility operation and design. Alan has published several technical articles from Energy Auditing, Facility Simulation to Fiber Optic local area network design. He is active in AEE, ASHRAE and NSPE. Alan is a licensed engineer in four states. His bachelor of electrical engineering degree is from Rutgers University and his graduate studies in mechanical engineering were through the Air Force Institute of Technology.