Chapter 11

Controls

Contents of Chapter 11

Instructions
Objectives of Chapter 11
11.1 Introduction
11.2 Controls Basics
11.3 Typical Control Loops
11.4 Introduction to Direct Digital Control (DDC)
11.5 Direct Digital Control of an Air-Handler
11.6 Architecture and Advantages of Direct Digital Controls
The Next Step
Summary
Bibliography

Instructions

Read the material of Chapter 11. Re-read the parts of the chapter that are emphasized in the summary and memorize important definitions.

Objectives of Chapter 11

Chapter 11 starts off by describing the basics of control and introducing you to some of the terminology of HVAC controls. After this introduction, we consider the physical structure and software of Direct Digital Control (DDC) systems. In this section, we demonstrate some of the control possibilities that are available with DDC by revisiting some of the references to controls in earlier chapters. Finally, there is a brief introduction to the architecture of DDC systems and their advantages. After studying the chapter, you should be able to:

Explain the following terms: normally open valve, modulating, proportional control, controlled variable, setpoint, sensor, controller, and controlled device.

Describe an open control loop and a closed control loop and explain the difference between them.

Explain how the DDC system replaces conventional controllers.
List the four main DDC point types and give an example of each one.
Explain how the knowledge in a DDC system can be put to good use.

11.1 Introduction

Every piece of equipment that we have introduced in this course requires controls for operation. Some equipment, such as a rooftop package unit, will likely come with factory-installed controls, except for the thermostat. The thermostat has to be mounted in the space and wired to the packaged unit. In other built-up systems, every control component may be specified by the designer and purchased and installed under a separate contract from the rest of the equipment.

Whether the controls are a factory package or built-up on site, well-designed controls are a critical part of any HVAC system. The controls for a system may differ from project to project for a number of reasons. Design considerations for controls choices include availability of expertise in maintenance and operations of the controls, repair and maintenance expense budgets, and capital costs of control equipment.

To elaborate, one should always choose controls that are suited to the available maintenance and repair expertise and availability. Find out how the client will be arranging maintenance of the system. As an example, it is generally unwise to choose the latest and greatest high-tech controls for a remote school, unless the school has a maintenance system in place to support the controls. It is generally better to aim for simplicity and reliability in this type of situation.

On the other hand, if the client has experienced, well-trained, controls staff available, on site or by contract, there is an opportunity to specify something quite sophisticated. As always, economics plays a controlling role and the challenge is to demonstrate how the sophisticated computerized system will perform better and save energy compared to a simple off-the-shelf option.

There are several types of controls and each has specific features that make it the best choice in particular circumstances. The following is a brief introduction to the main types.

Control Types

Controls fall into broad categories based on a particular feature.

Self-powered Controls require no external power. Various radiator valves and ceiling VAV diffusers have self-powered temperature controls. These units are operated by the expansion and contraction of a bellows that is filled with a wax that has a high coefficient of expansion. As the temperature rises, the wax expands, lengthening the bellows. This closes the radiator valve (cuts back on heating) or opens the VAV diffuser (increases the cooling). The advantage of these units is that they require no wiring or other connection so installation cost is minimal.

Electric Controls are powered by electricity. We will introduce two types of electric controls in this course:

On/off Electric Controls are used in almost every system to turn electrical equipment on and off. The electric thermostat is the most common example.

Modulating Electric Controls are based on small electric motors and resistors that provide variable control.

Pneumatic Controls are controls that use air pressure. The signal transmission is by air pressure variation and control effort is through air pressure on a diaphragm or piston. For example, a temperature sensor may vary the pressure to the controller in the range of 20 kPa to 100 kPa. The controller will compare the thermostat line pressure with the setpoint pressure and, based on the difference, adjust the pressure to the heating valve to open, or to close, the valve. The heating valve will typically have a spring to drive it fully open and the increasing air pressure will close the valve against the spring. The valve is called a "**normally open**" valve, since failure of the air system would have air pressure fall to zero and the spring would open the valve. A "**normally closed**" valve is the opposite, with the spring holding it closed until the air pressure opens the valve.

Pneumatic controls require a continuous source of compressed air at 100 kPa for sensing and controlling. When considering the total cost of the pneumatic system, the provision of the compressor(s), the operation and maintenance cost, and the energy lost with leaks have to be factored into the total cost. However, the pneumatic system does have the advantage of relatively inexpensive and powerful actuators (a device that moves a valve or damper) and it is relatively easy to learn to maintain and service.

Electronic Controls, or more correctly **Analog Electronic Controls**, use varying voltages and currents in semiconductors to provide modulating controls. They have never found great acceptance in the HVAC industry, since Direct Digital Controls offered much more usability at a much lower price.

Direct Digital Controls (DDC) are controls operated by one or more small computer processors (microprocessors). The computer processor uses a software program of instructions to make decisions based on the available input information. The processor operates only with digital signals and has a variety of built-in interface components so that it can receive information and output control signals.

There are many instances where the types of controls are mixed. For example a DDC system could have all electric "**sensors**," the units that measure temperature, humidity, pressure, or other variable properties. This same system may also have pneumatic actuators on all the valves, since pneumatics provide considerable power and control at low cost. A "**transducer**" creates the interface between the electrical output of the DDC system and the valve. The transducer takes in the DDC signal, say a voltage between zero and ten volts, and converts it to an output of 20 kPa to 100 kPa. Thus, at zero volts the output will be 20 kPa, rising to 100 kPa at ten volts.

We will spend considerable time on DDC controls later in the chapter. For now, let us consider the basics of controls—what makes them work.

11.2 Controls Basics

You instinctively know about control. You control all sorts of actions in your daily life. In this section we are going to introduce the basic ideas of controls. Your understanding of the rest of the chapter depends on your being really comfortable with the ideas in this section. Take the time to think about the ideas presented and how controls operate.

We are going to start with the simplest of controls, "on-off." As the name implies, the element being controlled is either "on," or "off."

> Consider a domestic hot water tank with a thermostatically controlled electric heating element near the bottom. Water becomes less dense as it is heated above 4.1 °C, so as the element heats the water, hot water will rise to the top. When the water at the thermostatically controlled element is hot, all the water above it is hot and the thermostat will turn off the heating element.
>
> Now, let us assume someone runs a little hot water. Cold water enters the bottom of the tank and cools the thermostat. The thermostat switches the element "on" and soon the tank is filled with all hot water once more. Suppose that later, one person runs a shower as someone else is running hot water for washing clothes. Although the element will come on, it can't keep up with this large load. Very soon all the hot water is gone and the tank is full of cold water. Both users turn off the taps. The element at the bottom of the tank will slowly heat the whole tank back up to the required temperature.

In order to achieve a quick recovery of hot water, we need a second element near the top of the tank. An element near the top of the tank only has to heat the water above it, so it will get a small amount of water up to temperature much more quickly. Now we have two elements. Do we need them on at the same time? No. If the top element is needed, the bottom element is not needed. So, when the top element turns on, for quick recovery, it also breaks the circuit to the bottom element. Once the water above the top element is hot, that element switches "off" and the bottom element switches "on" to heat the rest of the water in the tank. This gives us an "either/or" control decision—either the top or the bottom element can be "on."

The result is a tank that heats a little water quickly, and, in a much longer time, heats the full tank, with the electrical load of one heating element.

This is a simple example of how "on-off" controls can be cascaded to produce simple, but very effective, control. With some ingenuity, quite complex and extremely reliable electric controls can be developed.

Now let us move from "on-off" to "**modulating**" controls. Modulating means "variable." One type of modulating control is **proportional control**. This is best explained with a demonstration.

Take a jug and fill it with water.

Take a tumbler, and place it in a spot that won't be damaged if water overflows (like in a sink).

Next, see how quickly you can use the jug to fill the glass so that the water level is right up at the rim of the glass – so full you'd need a very steady hand to drink it!

If you didn't actually do the task, take a few moments to relax, and visualize the empty glass with the full jug on the table beside it. The jug is heavy as you pick it up and start to pour. You hear the water flowing in, feel yourself tipping the jug, see the level rising and feel that tension as you slow the flow to drips, to make it just reach the top, and then you stop pouring.

What happens? When you see the glass empty or just starting to fill, the water level is a long way from the rim. Naturally, you start pouring quite quickly. As the glass fills, you slow the flow until you're just *dripping* the water in, to get the glass quite full, to the rim, without going over. The change in rate at which you pour is roughly proportional to the distance of the water from the rim. This change in rate is called the **gain**. You are acting as a "**proportional controller**."

Proportional control is the basis of the majority of control loops—the rate is proportional to the distance from the target—the setpoint.

Now imagine a more complicated scenario of proportional control. In this scenario, we will be demonstrating **offset** and **overshoot**.

> Someone has attached a hose to the bottom of your glass and runs it to a tap downstairs, out of your sight. Your job, now, is to keep the glass full. You fill the glass, and then you notice the level is dropping slowly. In response, you start to pour slowly, just keeping the glass near full. Then, you realize the level is dropping faster, so you tip the jug. Suddenly the glass is full–it is overflowing!
>
> What happened?
>
> Initially, the hose tap was opened just a little, and it was easy for you to pour slowly to keep the level near the rim. When the tap was opened wide, though, the jug had to be tipped a lot to keep up, so when the tap was suddenly closed, it took a moment for you to realize that the water level in the glass was rising rapidly. It took another moment—too long—to straighten the jug and stop the flow, and the water overflowed over the top of the glass.

Just like you, a control system has an easy time with slow steady changes. But you had to notice that the water level had dropped before you started to pour. This created a time delay. Note also that you attempted to keep the glass almost full rather than totally full. This represented an "**offset**" from target, the setpoint. Then, when the glass started to drain rapidly you poured faster, to keep it from being empty, rather than trying to maintain the level just at the rim—even more offset.

Finally, the drain on your glass stopped, and you were too slow to straighten the jug and stop the flow. The water overflowed—serious "**overshoot**"! This overshoot could have been reduced if you had been restricted on how fast you could pour. If you had less **gain** you would have had less ability to keep up with sudden changes and the overshoot would have been much less.

You now have some feel for controls and what they do. There are many added refinements to controller action that are explained in the ASHRAE Learning Institute course, *Fundamentals of HVAC Control Systems*.[1]

Now consider some real HVAC examples.

There are two types of controls: "**closed loop**" and "**open loop**." Let us start by considering the main components of a closed loop control as shown in *Figure 11-1*.

Closed Loop

The top half of *Figure 11-1* illustrates a simple air heating control loop. A temperature sensor measures the temperature of the heated air and sends that information to the controller. The controller is also provided with the

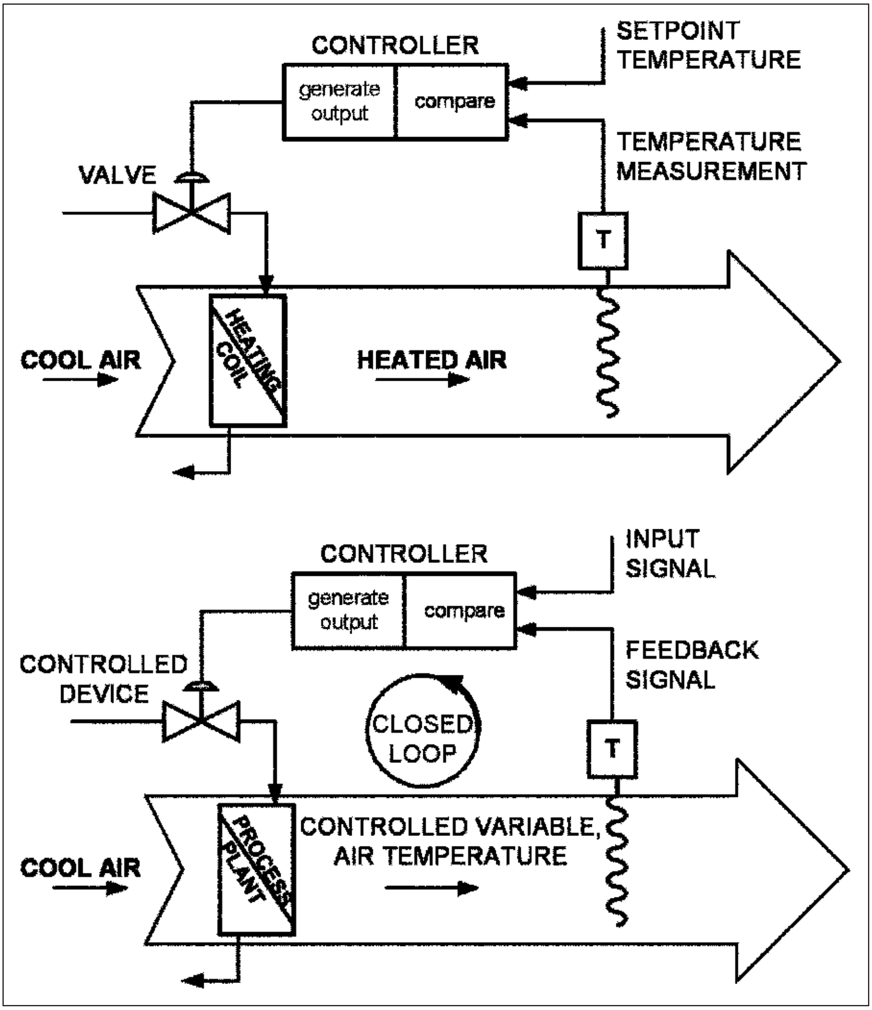

Figure 11-1 Closed Loop Control

required setpoint (similar to the setting on the front of a room thermostat). The controller first compares the measured temperature with the setpoint and, based on the difference (if any), generates an output signal to the valve. If the sensed temperature were a little higher than the setpoint, the controller would generate an output to close the valve a little. The valve would close, reducing the heating coil output. The air would be warmed less and the temperature sensor would register a lower temperature and send that information to the controller—and so on round and round the closed control loop.

The lower part of the figure is the same process with the generic names for the parts of the control loop.

The "**controlled variable**" is the variable, in this case temperature, that is being controlled. Controlled variables are typically temperature, humidity, pressure, and fan or pump speed.

The "**setpoint**" is the desired value of the controlled variable. In this example it is the air temperature that is required.

The "**sensor**" measures the controlled variable and conveys values to the controller. In this case the sensor measures temperature.

The "**controller**" seeks to maintain the setpoint. The controller compares the value from the sensor with the setpoint and, based on the difference, generates a signal to the controlled device for corrective action.

Note that a room thermostat contains the temperature setpoint, which is your adjustment of the setting on the front of a room thermostat. It also contains the room temperature sensor and the controller. A humidistat is the same, except that it is sensing relative humidity.

The "**controlled device**" responds to signals received from the controller to vary the process—heating in this example. It may be a valve, damper, electric relay, or a motor driving a pump or a fan. In the example it is the valve controlling hot water or steam to the coil.

To make sure you understand the above definitions, think about the simple example of pouring water in to fill the glass. What do you think were the following?

Setpoint
Sensor
Controller
Controlled device
Controlled variable

Answers are included at the end of this section.

So far, we have been discussing closed loop control—based on feedback, the controller makes continuous adjustments in order to maintain conditions that are close to the setpoint.

Open Loop

Another type of control is called "**open loop**" control, where there is no feedback. Consider a simple time clock that controls a piece of equipment. The time clock is set to switch "on" at a specific time and switch "off" at a later time. The time clock goes on switching "on" and "off" whether the equipment starts or not. In fact, it will go on switching even if the equipment is disconnected. There is no feedback to the time clock, it just does what it was set to do.

In Chapter 8 we introduced the idea of "outdoor reset." Outdoor reset is a method of adjusting the temperature of a heating source, or cooling source, according to changes in outdoor temperature. This is an example of open loop control.

We are going to add outdoor reset to our air heating system, as illustrated in *Figure 11-2*, below.

Figure 11-2 illustrates the same closed control loop as in *Figure 11-1*, but with outdoor reset added. The ambient (outdoor) temperature sensor provides Controller #1 with a signal, and the setpoint is provided as a variable according to the outdoor temperature. This is illustrated as a little graph, in the top right hand corner, showing a falling supply setpoint temperature (Y axis) as the temperature rises (X axis). The output of Controller #1 is the setpoint for our

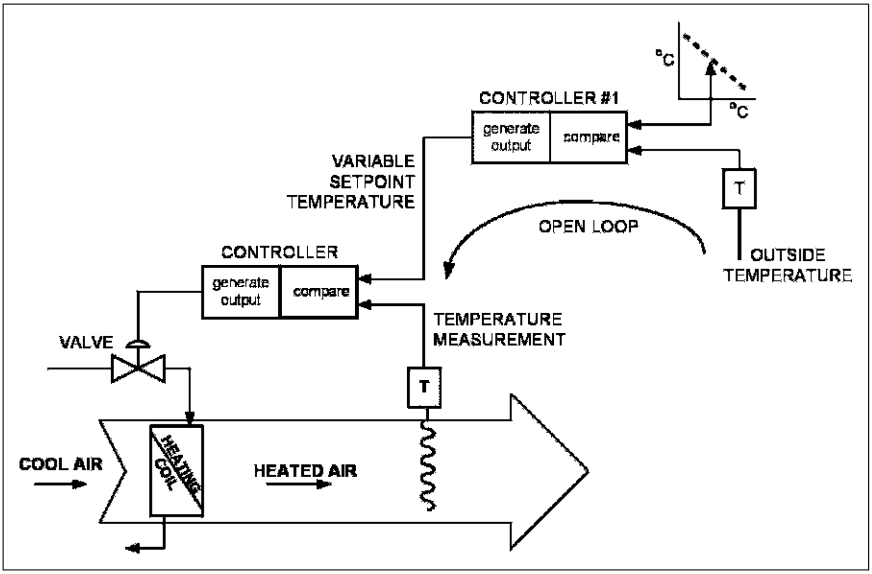

Figure 11-2 Open and Closed Control Loops

closed loop controller. The open loop measures temperature and provides an output. It has no involvement with the result; it just does its routine—open loop control—no feedback.

Alternatively, we could have chosen to use a chilled water system and to use outdoor reset to raise the chilled water temperature as the outside temperature dropped.

Outdoor reset is a common requirement, so manufacturers frequently package the two controllers into one housing and call it a "**reset controller**." This packaging of several components of the control loop is similar to the thermostat package where the setpoint, the temperature sensor, and the controller are packaged in one little box.

11.3 Typical Control Loops

Having considered the basics of control loops in the previous section, now look at some real, more complete, control loops. We will start by adding time control, another open loop, to our previous example, as *Figure 11-3*.

A time clock now provides power to the controllers according to a schedule. Typical commercial thermostats include the **5-1-1** time clock function. 5-1-1 means that they have independent time schedules for the 5 weekdays, 1 for Saturday, and 1 for Sunday.

The system shown also has a manual-override switch that allows the occupant to switch the system 'on' when the time clock has it "off." There is an

---

Answer to previous page's example:

Setpoint – top edge of the glass, **Sensor** – Your eye, **Controller** – your brain, **Controlled device** – the jug of water, **Controlled variable** – depth of water in the glass

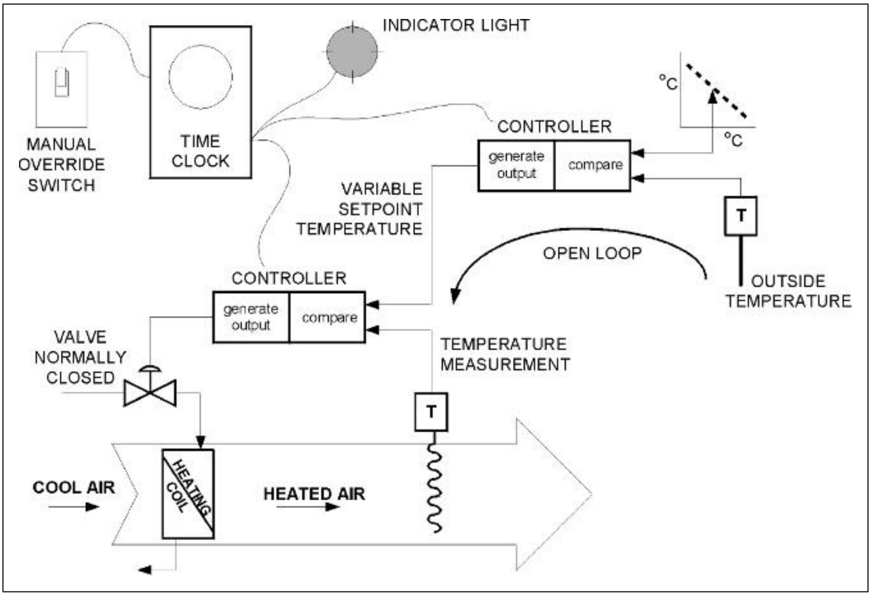

Figure 11-3 Controls with Time Clock Added

obvious energy waste issue here, since the occupant may forget to switch back to the time clock. In most time clocks, the manual switch is part of the unit, rather than a remote switch as shown in the figure.

In addition, there is an indicator light to show that the system is "on." When the time clock switches "on," it provides power to the lamp and power to the controllers. It has no idea whether the controllers are "on" or even whether they are connected! The lamp does not indicate that the system is working. What it indicates is that power from the time clock is available. This type of open-loop indication is very common. If you are involved in trouble shooting equipment, think about the real information provided. Even if the lamp "off," it does not mean there is no power to the controllers—the lamp could have burned out!

In our diagram there is just one heating coil being controlled. In many packaged units, there will be two stages of cooling and two stages of heating. A single 5-1-1 thermostat will provide full control, turning "on" one stage of heating/cooling, and then the second stage. The really good feature of a single, packaged thermostat is that there cannot be any overlap of control. For example, if a separate thermostat were used for heating and another for cooling, they could mistakenly be set so that the first stage of heating was "on" when the first stage of cooling starts—a real waste of energy.

Minimizing Energy Use

The issue of staging controls so that energy use is minimized is important in many areas. An example is the sequencing of control in a VAV box with a reheat coil. A VAV system provides cold air for cooling and ventilation. Should a zone require less cooling than is provided at the minimum airflow

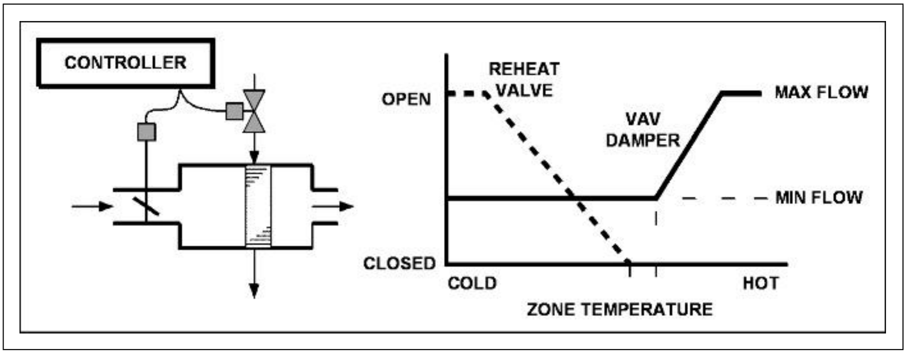

Figure 11-4 VAV Box with Reheat

for ventilation, then the reheat coil is turned on. In the control system for the box, there are two important requirements:

1. The heating coil must only be activated at minimum airflow.
2. There must be minimal cycling between "coil on" and "coil off."

To achieve this, a single controller is used to control both the airflow and the coil, in sequence. The heating valve and volume damper are normally closed. The volume damper has a minimum setting for minimum ventilation.

As an example, the box and controller actions are shown in *Figure 11-4*. Starting on the left, when the space is cold, the controller opens the heating valve fully. As the zone warms up, the controller closes the heating valve. Once the heating valve is closed, there is a dead band of temperature change (no heating, no additional cooling) before the controller starts to open the volume damper to increase the cooling up to maximum.

In addition to the simple control loops we have discussed, there are more complex loops that have many inputs. Staying with VAV for a moment, there are many systems where the fan speed is controlled by the requirements of the VAV boxes. A system, for example, might have 50 or more boxes. We want each to have enough air but we don't want to run the fan any more than needed. To manage this, we need to know when each box has adequate air flowing through it.

A VAV box has enough air if its damper is not fully open. Thus, we would be very confident that if every box has its damper at less that 95% open, there is enough air pressure in the system. However, determining if every box meets this condition is only practical in a DDC system. We will begin to examine these DDC systems in the next section.

11.4 Introduction to Direct Digital Control (DDC)

As briefly mentioned earlier, small computer processors (microprocessors) operate Direct Digital Controls. "Digital" means that they operate on a series of pulses. In the DDC system, all the inputs and outputs remain, however, they are not processed in the controllers, but are carried out in a computer, based on instructions called the "**control logic**."

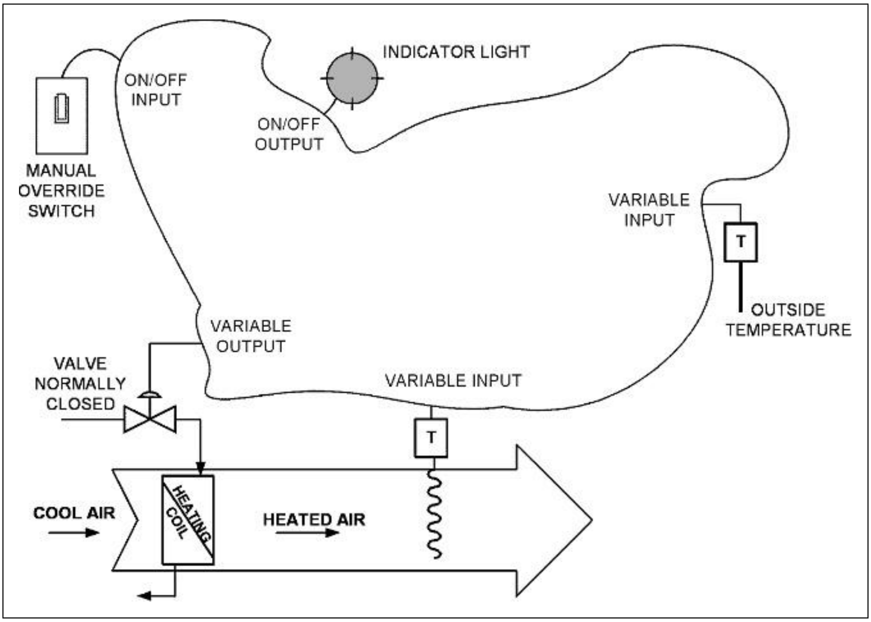

Figure 11-5 Control Scheme (from *Figure 11-3*) without Controlling Components

Figure 11-5 shows the same control diagram that we saw in *Figure 11-3*, but with the controlling components, (the time clock, and the two controllers), blanked out. All the system that has been blanked out is now replaced by software activity in the computer.

In *Figure 11-5*, each input to or output from the DDC computer has been identified as one of the following

On/off input – **manual switch**
On/off output – **power to light**
Variable input – **temperature from sensor**
Variable output – **power to the valve**

These are the four main types of input and output in a control system. Let's consider each one briefly in terms of a DDC system.

On/off input. A switch, a relay, or another device closes, making a circuit complete. This on/off behavior has traditionally been called "digital." Therefore in DDC terms it is generally called a "**Digital Input**," or **DI**.

The term "digital" is not considered technically correct, since there is no series of pulses, just one "on" or "off." Thus, for on/off points the term "binary" is considered more correct, and the term is being encouraged in place of "digital." So, "binary input," BI, is the officially approved designation of an "on/off" input.

On/off output. The on/off output either provides power or it does not. The lamp is either powered, "on," or not powered, "off." In a similar way, this is called a "**Digital Output**," **DO,** or more correctly, binary output, **BO**.

Variable input. A varying signal, such as temperature, humidity, or pressure, is called an "analog" signal. In DDC terms, the input signal from an analog, or varying, signal is called an "**Analog Input**," or **AI**.

Variable output. In the same way, the variable output to open or close a valve, to adjust a damper, or to change fan speeds, is an "**Analog Output**," **AO**.

You might think the next step is to connect these DI, DO, AI, AO points to the computer. Things are not quite that simple. A sensor that measures temperature produces an analog, varying signal and our computer needs a digital signal. So between each AI device and the processor there is an "**A/D**," "**analog to digital**," device. These A/D devices, for AIs, are usually embedded with the computer.

Similarly, for AO points there is a "**D/A**," or "**digital to analog**," device that converts the digital signal to a 0–10 volts or 4–20 milliamp electrical signal. This signal has too little power to operate a valve or damper. If, for example, the controlled device is a valve that is powered by compressed air, the analog electrical signal will go to a "**transducer**" in which the electrical signal will be converted to an air pressure that drives the valve. If the valve is electrically powered, the transducer will convert the low power, analog signal to a powerful electric current.

Only standard telephone cable is required to carry the analog electrical signal, hence the transducer is often separate from the processor and close to the controlled device. This is because it is far less effort, and cost, to run standard telephone cabling to the transducer rather than to run the air line (or electric power cable) to the processor location and back to the valve.

Naming Conventions

In a DDC system every input and every output must have a unique name. There are a variety of naming conventions depending on personal preference and the size and complexity of the system. Many are based on a hierarchy of elements such as

Type – Building – System – Point – Detail

"Detail" allows for a number of identical points, VAV boxes for example. If we assume our build is called "NEW," our points list might be:

AI NEW AH1 OAT	AI, in NEW, on air-handler1, outside air temperature
AI NEW AH1 DT	AI, in NEW, on air-handler 1, duct temperature
AO NEW AH1 DT	AO, in NEW, on air-handler1, duct temperature control
DI NEW AH1 MAN	DI, in NEW, on air-handler1, manual control
DO NEW AH1 IND	DO, in NEW, on air-handler1, indicator light

Sequence of Operations

Now look back at *Figure 11-5*. As we noticed earlier, the controllers and time clock are all blanked out. In a DDC system, all the actions of the controllers and time clock are carried out through software in the small computer processor. The software is a set of ordered operations, which is often called the "**sequence**

of operation." What do we require our software to do? The following is a very simple "English Language" sequence of operations.

```
Do the following things:

    If the time is between xx:xx a.m. and yy:yy p.m. run mode is 'ON', otherwise run mode is 'OFF'
    If the manual switch is closed, DI, run mode is 'ON'
    If run mode is 'OFF' close heating valve
    If the system is in run mode 'ON', do the following commands

        Check ambient temperature, AI, and remember the value as 'ambient'
        Using 'ambient', lookup required setpoint from (graphic) schedule to find required air setpoint temperature. Remember this value as 'setpoint'
        Check air temperature in the duct, AI, and remember it as 'temperature'
        If 'temperature' is less than 'setpoint' increase output to valve, AO
        If 'temperature' is greater than 'setpoint' decrease output to valve, AO

Go back to the beginning
```

These instructions are typically written into the DDC processor using a standard personal computer, PC. The programming may be a more formal version of our little example, or may use graphic symbols instead. The DDC processor can also be programmed to sound alarms, issue warnings, write messages, plot graphs, and draw graphics through the PC.

So now we can redraw *Figure 11-5*, to show the DDC system, *Figure 11-6*.

As it is shown, there is no way of accessing the processor. In a real system, there is a communication connection providing access from a computer, typically a desktop PC or a laptop computer. The PC has many names including "the operator interface," "front end," or "operator machine interface" (OMI). Assuming that this panel has everything it needs to run the system, it is called a "**standalone panel**." Standalone means it has everything to keep running on its own.

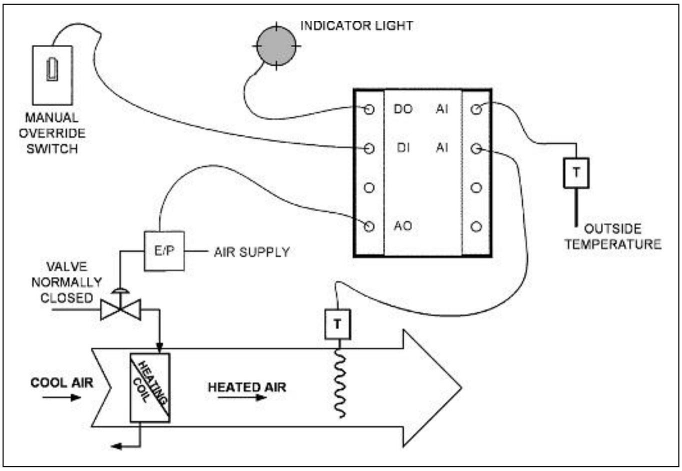

Figure 11-6 DDC Control Schematic

A really important thing to understand is that the DDC controller can record what happens over time and either directly use that information in useful ways or provide it to the operator.

In our simple system, for example, the DDC system could check how many hours the manual switch had been on. If it had been on more than three hours, it could issue an alarm to the PC, asking the operator if it should still be in manual. This alarm could repeat every two hours to remind the operator to change back to the schedule.

In addition, there are some faults it could be programmed to detect. The heating valve is normally closed when the system is "OFF." When the system starts, the duct temperature should be no hotter than the building or the outside ambient temperature. Now, our system does not know the building temperature, but we could assume it would be no higher than 27 °C. Thus, we could have a software routine that checked, on startup, that the duct temperature was both no higher than 2 °C above ambient, and no higher than 29 °C. If the duct temperature were above both these two checks, it could issue a warning that the heating valve may not be shutting off completely.

It is this ability to collect information about every point and to process it, that makes DDC so powerful. Treating it as only a controller replacement is to miss out on the real power of the system.

Let's consider a very simple illustration of this power of knowledge that can be written into a DDC system. We are going to consider two offices served by a single VAV box as illustrated in *Figure 11-7*.

The objective is to provide the occupants with conditions that are as comfortable as possible. If we connect an occupancy sensor and temperature sensor in each office, the DDC system will know if the office is occupied and the current temperature in each office. When both offices are occupied, the system can average the temperatures of the two offices and keep the average as close to setpoint as possible. Now, when the occupancy sensor detects that one of the offices is vacated, the controller can wait a few minutes to avoid annoyance and then slowly change to controlling based on just the occupied office temperature.

In addition to improving the temperature control, the occupancy sensors also allow the system to modify the amount of outside air being brought in. If one office is vacated, the outside-air volume can be reduced by the assigned volume for the empty office.

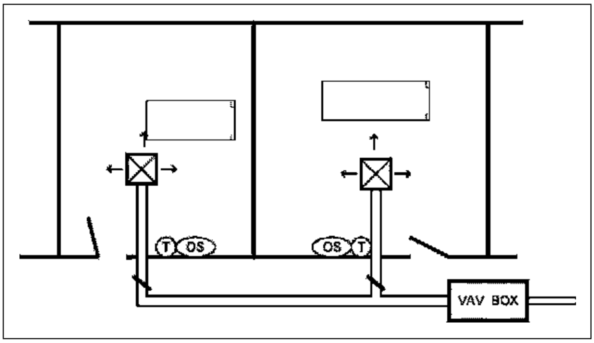

Figure 11-7 Two Offices Served by One VAV Box

Finally, when both offices are empty, the system does not need to maintain the temperature to the same tight limits, and there is no requirement for ventilation air, so, if there is no thermal load, the VAV box can be completely closed. Similar to the example of CO_2 control in Chapter 4, Section 4.5, the system only provides service to occupants who are present.

There is one more advantage. The system can be designed to prevent the lights being left on for long periods when the office is unoccupied. One method is to provide power to the lighting circuits (not switch them on, just provide power) when the room is sensed as being occupied. The occupant can switch the light on and off when they like, but when they leave, it will soon go out. The system delays turning the light off for several minutes, to avoid annoyance when the occupant is only away for a few minutes.

This section has introduced you to basic ideas of DDC.

- The sensors and actuators stay, but all the control logic is in the software.
- There are four types of input and output, DO (BO), DI (BI), AO, and AI.
- The software is a set of instructions that the DDC system can interpret and act upon.

In the next section, we are going to consider the points and sequence of operation of an air-handler. Then in Section 11.6 we will consider how DDC units can be interconnected, and can share information with each other and the operators to make a full-scale control system, rather than a collection of control loops.

11.5 Direct Digital Control of an Air-Handler

In this section we are going to consider a constant-volume air handler serving a single zone, designated "001." The air handler uses space temperature for control, with no mixed air control, unlike air handlers that we have discussed before. This is a design choice, unless there is a local code that requires a specific method. Where *ASHRAE Standard 90.1-2004, Energy Standard for Buildings Except Low-Rise Residential Buildings,*[3] is incorporated into the local building code requirements, the use of mixed air control is not allowed. We will discuss this standard in the next chapter.

To specify a DDC control system, ideally, one produces three things:

1. A schematic of the system with the control points labeled, *Figure 11-8*
2. A list of control points with their characteristics, *Figure 11-9*
3. A schedule of operations

The schematic with the points labeled is not always provided, but it can avoid arguments about the location of points after installation, and it provides the maintenance staff with a map for locating points.

Sequence of Operation

Schedule: Provide calendar/time schedule with minimum of three occupied periods each day.

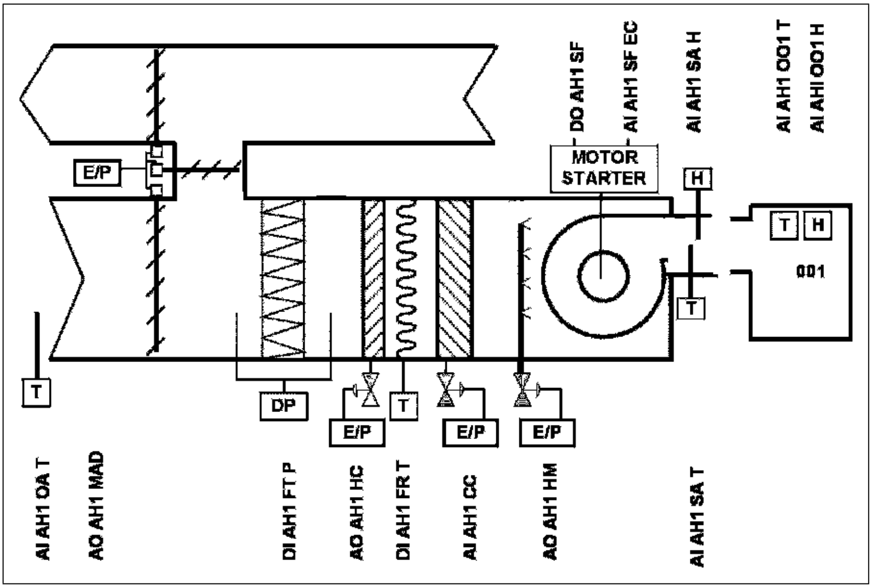

Figure 11-8 System Schematic

Unoccupied: When calendar schedule is in unoccupied mode, and if space temperature is above 16 °C, the fan shall be off, heating valve closed, cooling valve closed. If space temperature falls below 16 °C, then the outside dampers and cooling valve to stay closed, heating valve to 100% open, and start fan. When space temperature reaches 18 °C, turn fan off and heating valve closed.

Occupied: When calendar schedule is in occupied mode, the fan shall be turned on and after 300 seconds, the heating valve, outside air dampers and cooling coil shall be controlled in sequence to maintain space temperature at 22 °C.

The control sequence shall be: heating valve fully open at 0% and going to fully closed at 33%, at 34% the dampers will be at their minimum position of 20% and will move to fully open at controller 66%, the cooling valve will be fully closed until 66% and will be fully open at 100%.

Economizer control: When the outside temperature is above 19 °C, the outside air dampers shall be set back to minimum position of 20%, overriding the room controller requirement.

Fan Control Alarm: If the fan has been commanded on for 30 seconds, and the fan current is below alarm setpoint 85% of commissioned current, the fan shall be instructed to stop, outside air dampers closed, and heating and cooling valves closed. An alarm of "low fan current" shall be issued.

If the fan has been commanded off for 10 seconds, and the fan current is above the low limit, the fan shall be commanded off, and dampers, heating coil and cooling coil shall be controlled as in occupied mode. An alarm of 'fan failing to stop' shall be issued.

Filter alarm: If the filter pressure drop exceeds 75 Pa, the filter alarm shall be issued.

			Inputs							Outputs										
			Analog				Digital			Analog			Digital			Alarms				
System: Air-handler 1	**Point designation**	Device number	Temperature	Humidity	Flow	Electric current	Freeze thermostat	Differential pressure	Flow switch	Transducer	Current 4–20 ma	Voltage 0–10 Vdc	Contact	Solenoid valve	Relay	Contact open	Contact closed	Value greater than	Value less than	Comments
Outside air temperature	AI AH1 OA T	1	**X**																	
Mixed air dampers	AO AH1 MAD	7								**X**										
Filter pressure	DI AH1 FT P							**X**									**X**			Filter change alarm
Heating coil	AO AH1 HC	7								**X**										
Freeze thermostat	DI AH1 FR T						**X**										**X**			Freeze alarm
Cooling coil	AO AH1 CC	7								**X**										
Humidifier	AO AH1 HM	7								**X**										
Supply fan on/off	DO AH1 SF												**X**							
Supply fan electric current	AI AH1 SF EC	6				**X**												105%	80%	Fan current high alarm or low alarm
Supply air temperature	AI AH1 SA T	2	**X**																	
Supply air humidity	AI AH1 SA H	4		**X**														85%		Supply air high humidity alarm
Space 001 temperature	AI AH1 001 T	3	**X**															85	53	Space temp high or space temp low alarm
Space 001 humidity	AI AH1 001 H	5		**X**														60%		Space humidity high alarm

Figure 11-9 Points List for AH1

Freeze Alarm: If the supply air temperature drops below 7 °C, hardware freezestat operates, system changes to unoccupied mode and issues "freeze" alarm.

Manual override: If the manual override is sensed, run in "occupied mode" for 3 hours.

System status: 280 seconds after entering "occupied mode" the room temperature, supply temperature, and ambient temperature shall be recorded along with current date and time.

Note that, in this case, the point names are given in full. It really helps future maintenance if a point-naming convention is established and enforced, including having the contractor label every input and every output device with its point name. It also discourages the contractor from accidentally dropping into the naming convention of the last project!

The convention used here is only an example, chosen to make this text easy to understand. Many naming conventions do not include the spaces and many do not include the AI, AO, DI, and DO, but make the name self-explanatory. For example, instead of AO AH1 CC, they might use AH1 CCV, meaning AH1 Cooling Coil Valve.

The column "Device number" refers the contractor to the specification for the device. In this case, device number 1 is an outdoor air temperature sensor, device number 2 is a duct temperature sensor, and device number 3 is a room temperature sensor.

Most of the sequence of operations shown here can be achieved with any control system. Two DDC-specific routines have been included, to aid maintenance and to help avoid energy waste:

The *first* DDC-specific routine is to start the fan, leaving all controls alone. The fan will circulate air from the space, so after 280 seconds the sensors should have stabilized. The space temperature sensor should record the same temperature as the supply air temperature, except for the small rise in temperature due to fan energy that occurs as the air goes through the fan. Let's suppose this rise is normally 0.5 °C on this example system.

One cool day, when the chilled water system is shut down, the operator checks the startup temperature rise. It is surprising to note that it is −2 °C, so something has gone wrong. It is cool outside, so the outside dampers could be letting in cold air, even when they are controlled to be fully closed. It is also possible that the space temperature sensor or supply air temperature sensor is providing the wrong reading. The operator does not know which is the actual problem but will probably start by checking the dampers.

The erroneous temperature difference will provide different possibilities for what is wrong under different weather conditions, depending on whether the chilled water was available, and whether the temperature difference was positive or negative. The designer can fairly easily work up a written decision tree of possible problems to help the operator. As all the information is available in the DDC system, the designer can also program the system to work through the decision tree and present the operator with the possible problems to check.

This level of sophistication is becoming available on factory produced standard products. On larger systems, and for remotely monitored sites, this type of self-analysis is becoming a valuable feature of high-level DDC systems. However, it is generally not warranted on a small, simple system where the programming is being written for that one project.

The *second* DDC-specific routine is to use a current sensor on one of the cables to the fan to provide a measure of fan current. Our example is a constant-volume system, so the load on the fan will be relatively constant. It will not be completely constant, since the pressure drop across the filters will rise as the filters become loaded with dirt. Based on the actual fan current when the system is commissioned, a high alarm point and a low alarm point can be chosen. Then, if a bearing starts to fail, the load will typically rise, and this can be detected before bearing failure and possible destruction of the fan. Also, if the fan is belt driven and the belts slip or break, the fan current will drop substantially. This will also be detected. Finally, if the fan starter or motor fails, there will be no motor current, again sensed as low current and alarmed.

These are just two examples of how a small change in the arrangement of the DDC controls can provide better control and maintenance.

Now that we have considered the basics of DDC and a sample system, we will move on to how systems are interconnected and built up into networks serving a whole building or many buildings.

11.6 Architecture and Advantages of Direct Digital Controls

So far we have considered the controls of a single, simple system connected to a single DDC panel. In many buildings, there will be several systems, often with many more points controlling air-handlers, VAV boxes, heating valves, pumps, boilers, and chillers. Wiring from a single huge DDC panel is not a practical option for two reasons. First, failure of the unit means failure of the entire system, and second, the wiring becomes very extensive and expensive. Instead, the system is broken down into smaller panels that are linked together on a communications cable, called a "communications network."

All of this is fairly simple if the system uses equipment from only one manufacturer. However, when more than one manufacturer is involved, it is not as simple. There are three communication issues that create problems. Let us identify them in terms of human communication first.

Languages

The problem is very similar to the problems of human language. In order for people from different countries to communicate, interpretation or language translation is required.

Similarly, in the controls world, different companies have worked up different control languages. The languages differ both in terms of the words and in terms of sentence structure. There are two ways of enabling communication so that one manufacturer's equipment can communicate with another manufacturer's equipment that uses a different programming language. The first is to have an interpreter, called a "**gateway**," between the two units. The second way is to program an additional, common language into both manufacturers' units.

Vocabulary and idea complexity

Different people learn different sets of words in the same language. For simple, everyday things, like bread and water, everyone learns the words in each language. In addition, different people are trained in different skills. Consider, for example, when an engineer and an accountant want to discuss the long-term value of a project. They can find themselves having great difficulty communicating, because they have different vocabularies and different thinking skills in the same language.

Transmission method and speed

Finally, people send messages over long distances by a variety of methods at various speeds. For example, consider a letter being faxed to a remote recipient. It first goes through the fax machine (gateway) to be converted into telephone data. The telephone data is routed through various telephone exchanges (routers) till it reaches the receiving fax machine (gateway) that converts the data back into the original text letter.

In addition to the method, there is an issue of speed. Faxing is a quick and easy way of sending a letter, but if a whole book of text is to be sent, the much higher speed available on the Internet is considerably more attractive.

The issues of language, vocabulary and idea complexity, and transmission method and speed are very much the same in DDC systems.

Typically, a DDC panel includes software that provides the sequence of control activities and software for communicating with other panels. The internal software is generally proprietary to each manufacturer, and the communications software can be proprietary or public. There are several good, reliable communication languages, called "**protocols**" for simple information such as "the temperature is 40 °C," "open to 60%." The problems arise as soon as higher level communications, including any form of logic, are required.

In an attempt to eliminate the cost and challenges of no communication or expensive and limited gateways, ASHRAE produced a communications standard called "**BACnet**®." This is a public protocol that is designed to allow communication at all levels in a DDC system. It is documented in *ASHRAE Standard 135-2004, BACnet®, A Data Communication Protocol for Building Automation and Control Networks.*[2]

BACnet is particularly aimed at facilitating communications between different vendors' products at all levels. This allows buyers to have more vendor choice. It is important to note, though, that while the BACnet standard establishes rules, the designer still has to be very careful, since the number of rules used by different manufactures can make "BACnet compatible" systems and components unable to communicate. However, with careful specification, one can obtain units and components from a variety of manufacturers that will communicate with each other.

The ability of different manufacturers' equipment to work together on a network is called "**interoperability**." To assist in ensuring interoperability and the use of BACnet, a BACnet interoperability association has been formed to test and certify products.

System Architecture

Let us now consider a DDC system and how it might be arranged—the system architecture. Consider the system illustrated in *Figure 11-10*.

Across the top of the figure is a high-speed network connecting main standalone panels and the operator terminal. In this example, the standalone panel on the left uses a different communication protocol from the protocols used by the other two panels and the operator workstation. Therefore, a gateway

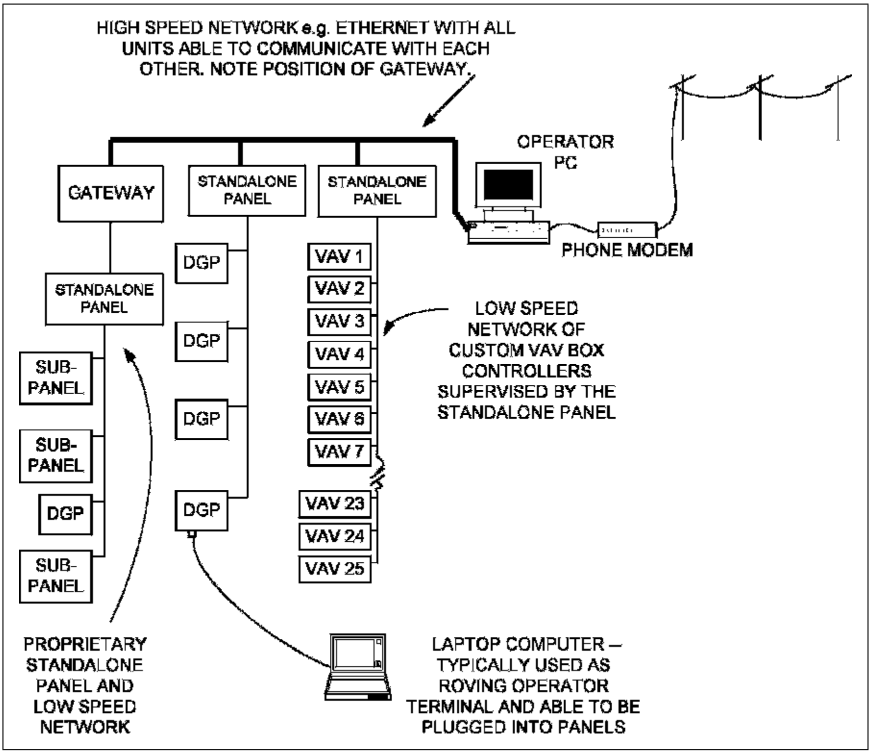

Figure 11-10 DDC System

(translator) connects the standalone panel on the left to the network. A "gateway" is a processor specifically designed to accept specific information in one protocol and send out the same information in another protocol.

Note that gateways are specific in terms of "protocol in" and "protocol out" and are often not comprehensive. By "not comprehensive," we mean that only specifically chosen information, not all information, can be translated (think of it as a translator with a limited vocabulary and limited intelligence).

The standalone panel on the left has a lower speed network of devices connected to it. The sub-panels might be small processors dealing with an air-handling unit, while the "**data gathering panel (DGP)**," may be simply gathering outside temperature and some room temperatures and transmitting them to the other panels.

The central standalone panel does all the processing for its branch of the system, with remote DGPs to collect inputs and drive outputs. A laptop is temporarily connected to one of the DGPs to allow the operator/maintenance staff to interrogate the system. The use of a laptop allows the operator/maintenance staff to have access to every function on that network branch, but it may not allow access through the standalone panel to the rest of the system.

The right-hand standalone panel is shown as having numerous VAV box custom controllers connected to it. These controllers are factory-produced, with fixed software routines built in to them. Programming involves setting setpoints and choosing which functions are to be active. These custom controllers are attractive because they are economical, but they are restrictive, in that only the pre-written instructions can be used.

In *Figure 11-10* there are a variety of devices in various arrangements with an operator PC as the local human interface. In addition, a phone "modem" is shown allowing communications with the system via a telephone from anywhere in the world. The modem is a device that converts the digital signals from the PC to audio signals, to allow them to travel on the telephone lines. There are three strikes against modems: they are slow, telephone charges can be prohibitive, and only one connection can be made to the modem.

These restrictions are now being removed by adding a "**web server**." A web server is another computer! The web server connects between the high-speed network and the Internet. It is programmed to take information from the DDC system and to present it, on demand, as web pages on the Internet. This enables anyone who has the appropriate access password to access the system, via the Internet, from anywhere in the world, at no additional cost.

Within the facility, web access allows any PC with web access to be used as an operator station, instead of only specifically designated operator stations. This is much more flexible than having to go to the operator's terminal to access the system. For example, the energy manager can use an office PC to access energy data on the machine that is used for normal day-to-day office work.

This chapter has done no more than introduce you to some of the basics and general ideas of DDC. The system has *advantages* including:

- Increased accuracy and control performance
- System flexibility and sophistication that is limited only by your ingenuity and the available financial resources.
- The system ability to store knowledge about the internal behavior over time and to present this information in ways that assist in energy saving, monitoring, and improved maintenance.
- Remote access to the entire system to modify software, alter control settings, adjust setpoints and schedules via phone or via the Internet.
- With increased use and the falling price of computer systems in general, DDC is often less expensive than conventional controls.

Then, there are the *disadvantages*:

- DDC systems are not simple. Qualified maintenance and operations people are critical to ongoing success. They must be trained so that they understand how the system is designed to operate.
- Extending an existing system can be a really frustrating challenge due to the frequent lack of interoperability between different manufacturers' products and even between upgrades of the same manufacturer's products.

For fairly detailed information on the specification of DDC systems, *ASHRAE Guideline 13-2000, Specifying Direct Digital Control System,*[4] is available.

The Next Step

In Chapter 12 we move on to consider energy conservation. We will review the subject in general before a brief discussion of the *ASHRAE/IESNA Standard 90.1-2004, Energy Standard for Buildings Except Low-Rise Residential Buildings,* and some heat recovery and evaporative cooling energy saving methods.

Summary

Chapter 11 has been an introduction to the ideas behind controls. This is a vast field and we have only provided a glimpse of the subject. A more technical and detailed introduction to controls is available as a Self-Directed Learning Course in this series, *Fundamentals of HVAC Control Systems.*

The chapter started off with some general discussion on control types: self-powered, electric controls, pneumatic controls, electronic (analog electronic), and direct digital controls. Each of these types has a niche where it is a very good choice but there is a general trend toward DDC controls. We then considered a very simple electric control of a two-element hot water heater to show how controls can be considered in a logical way. Next we introduced the control loop and the difference between open loop control (no feedback) and closed loop control (with feedback). The parts of a control loop that you should be able to identify are: setpoint, sensor, controller, controlled device, and controlled variable.

To illustrate the main issues with modulating controls, we had you imagine pouring water into a glass. The ideas illustrated were

Proportional control is control in which the control action increases in proportion to the error from the setpoint

Offset is the change of apparent setpoint as the control action increases in a proportional controller

Gain is the ability of the controller to make a large change in control signal

Overshoot is the result of applying too large a control signal and being unable to reduce it in time to prevent overshooting the control point

Speed of reaction is the time it takes for the controller to initiate a significant change

Having considered the standard control loops we went on to consider the four main types of DDC points:

Digital/Binary Input: a circuit such as a switch closing.

Digital/Binary Output: providing power to switch a relay, motor starter, or two-position control valve.

Analog Input: typically a signal from a temperature, pressure, or electric current sensor.

Analog Output: providing a variable signal to a valve, damper, or motor speed controller, often via a transducer that changes the low power signal to a pneumatic or electric power source with the necessary power to drive the valve or damper.

Having introduced the four main point types, we introduced the concept of using a point identification scheme. Then we considered a very simple example of a sequence of operations which are the logical instructions for the DDC controller to execute, to provide the required control. The required information to specify a DDC system control was then illustrated with a single air handler. The list of control points and schedule of operations is always required, but the schematic can be omitted, though doing so can lead to misunderstandings at the time of installation.

In addition to accuracy, a major advantage of DDC is the ability to record data and either use it for more intelligent control or as information for the operator.

Finally, we considered DDC architecture, introduced BACnet and interoperability, and listed the pros and cons of DDC.

Bibliography

1. ASHRAE. 1998. *Fundamentals of HVAC Control Systems*. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.
2. ASHRAE. 2004. *ASHRAE Standard 135-2004, A Data Communication Protocol for Building Automation and Control Networks.* Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.
3. ASHRAE. 2004. *ANSI/ASHRAE/IESNA Standard 90.1- 2004, Energy Standard for Buildings Except Low-Rise Residential Buildings.* Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.
4. ASHRAE. 2000. *ASHRAE Guideline 13-2000, Specifying Direct Digital Controls Systems.* Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.