The ability to achieve automated stepless control of pressure and flow rate in fluid power systems has undergone major development in the past twenty-five years. Electrohydraulic servo valves were invented in the late 1930s as a high tech, though high cost, solution to motion control needs. The mid-1980s saw the practical introduction of proportional valves as a viable and reasonably priced alternative to servo valves. This article will explore the technology used in these proportional and servo valves as well as attempting to shed some light as to what type of valve may be most appropriate for a given application.

If true, stepless control of pressure or flow wasn't that critical to the operation of a machine, preset pressure or flow control valves could be achieved by having a bank of preset valves. The appropriate valve would be connected into the circuit via the actuation of a solenoid valve. For example, three discrete pressures could be achieved by having two pilot relief valves connected in parallel to the vent port of a ventable pilot operated relief valve. These two valves would be isolated from the pilot operated relief valve by two-way normally-closed solenoid valves. By individually actuating the 2-way valves, three different pressures could be achieved. But what if stepless pressure control was required? What if the pressure increase or decrease needed to follow a specific rate? What if that rate change wasn't constant? Then what could the machine designer do?

Until the invention of servo valves, if the need existed to achieve varying pressure to an actuator in an effort to control force or torque, one needed to either have the machine operator turn an adjustment knob, stroke a lever, or a mechanical means needed to be designed to have a mechanical input device or linkage vary the setting of the valve. The same need held true if flow needed to be varied. Human control of a valve could be fairly inconsistent. Mechanical control of the valve, while possibly being more consistent and repeatable, might not offer much flexibility for different adjustment rates.

Most electrical machine control systems were not all that well developed until the introduction of the microprocessor in the early 1980s. As most machines that had electrical automation used relay logic, the sequence of operations of the machine was not easily revised. Relays are digital, or on-off, devices. The introduction of the microprocessor, and hence, the PLC (programmable logic controller) opened the door to a great amount of control options to machine designers. The operation sequence of a machine was no longer hardwired in relay logic. While Boolean operations were possible in relay logic, it was inconvenient to do so, as well as difficult, expensive, and time consuming to accomplish. The common introduction of PLCs and proportional valves greatly expanded the control options available to machine designers.

The initial proportional valves that appeared on the market are what are now commonly termed "open-loop" valves. In contrast to mechanical feedback (MFB) servo valves, a feedback link does not exist between the coil assembly and the valve spool. Since a feedback loop between the input command and the valve output does not exist, the feedback loop is "open" rather than "closed." In an effort to improve the performance of proportional valves, relative to the performance of servo valves, manufacturers added linear variable displacement transducers (LVDTs) to proportional valves in order to sense the spool (or poppet) position. The output signal from the LVDT was fed back to the amplifier card. A summing amplifier on the amplifier card calculated the difference between where the spool was supposed to be and where it actually was, and the output to the coil was changed in an effort to position the spool to achieve the desired output based on the input. These enhanced proportional valves are termed "closed loop" proportional valves. Since the means of feedback is electrical, not mechanical, this gave rise to the term "electrical feedback" (EFB).

How Proportional Valves Operate:

An electrical input from some source is wired to an amplifier card, which in turn, controls a coil on the proportional valve. Since the electrical input from most sources is generally much lower in power than the amount of current required to operate the coil, the input current must be amplified. This function is fulfilled by an amplifier card. The amplifier may be mounted on the valve, sometimes termed OBE (on board electronics) or be remote from the valve. The source of the input may come from

Proportional and Servo Valve Technology

several devices, including a potentiometer controlled by the machine operator, from preset potentiometers, a joystick, or from a PLC.

The amplifier card drives the valve coil with a current signal. As the current flows through the coil, electromotive force is developed, causing the armature of the solenoid assembly to move. The armature, in turn, inputs force to the valve spool, in a flow control, pressure reducing, or directional control valve, or the poppet in a pressure relief valve. The spool or poppet is offset by a spring. Therefore, the force input by the solenoid assembly is opposed by the force of the spring.

Many proportional directional control valves, such as the valve illustrated in Figure 1, have two solenoid assemblies, one solenoid being located at each end of the valve. Proportional directional control valves provide control of direction as well as of flow. This particular valve includes an LVDT. For the most part, dual coil proportional valves are based on standard on-off directional valves. The major differences between direct acting on-off (sometimes termed bang-bang) and direct acting proportional directional valves are:

1) The centering springs in proportional directional valves are stronger than the centering springs used in on-off directional valves.

2) Proportional solenoids are engineered to produce more force than do on-off solenoids.

3) Proportional valves always use DC solenoids.

4) While the bodies of on-off valves and proportional valves are almost always identical (the same body part is used for both versions by most manufacturers), the spools in proportional valves:

a) are tailored to the flow rate the valve is designed to control;

b) are available in a range of flow rates within a valve size;

c) and have metering notches that provide for a variety of flow rate profiles vs. the electrical command input to the valve.

5) Spools for directional controls are available 1:1 and 2:1 flow ratios between the two work ports to allow for control of hydraulic motors and double rod cylinders and 2:1 area ratio cylinders.


*Figure 1 - Courtesy of Denison Hydraulics*

Figure 1a - Courtesy of Denison Hydraulics

Some proportional directional valves have only one coil. These valves typically have four, rather than three, position envelopes. Figure 2 shows a schematic for a 4/4 (4-way, 4-position) single coil proportional valve. Note that in the de-energized condition (extreme left flow envelope), all ports are closed. In order to shift the spool to the "center" position, the spool must travel through a condition in which flow from the pressure port will connect to one of the working ports, as the opposite working port is connected to tank. Though the spool passes through this active flow condition very rapidly, its affect on the system still must be considered. As the solenoid assembly is completely de-energized, either as part of the machine control sequence or in the case of an electrical failure, the possibility of unwanted actuator movement must be considered. Solenoid operated blocking valves located in the working lines between the proportional valve and the actuator may be used to prevent unwanted actuator movement. These single coil 4-way directional valves tend to be high performance valves, relative to dual coil proportional valves. These valves are termed, by some manufacturers, as "servo-proportional" valves, indicating their higher dynamic performance. This higher performance capability stems from the fact that the displacement of the spool from the "center position" of the valve is not influenced by the hysteresis typical of the centering springs found in a dual coil proportional valve.


Figure 2 - Courtesy of Denison Hydraulics

So, how does a proportional coil actually work? All coils used on proportional valves are direct current (DC) coils. AC coils have an inrush current approximately five times greater than their holding current. If the armature in an AC solenoid assembly is not allowed to completely shift into position, its current draw will remain high. The coil will overheat and burn out; an AC is not designed to handle a sustained amp draw five times its holding amp draw. DC solenoids don't exhibit inrush current, so the

armature can remain partially shifted indefinitely without an increase in amp draw. By virtue of being able to partially shift the armature, the spool or poppet may be partially shifted as well, resulting in a partial output from the valve.

The simplest means of varying the current used to drive a proportional coil would be to locate a rheostat (adjustable resistor) between a DC power supply and the coil. The problem with this solution is that any current that isn't directed to the coil will be changed into heat. This is analogous to using a fixed displacement pump at less than its rated flow to supply an actuator. Just as the excess pump flow is directed to tank over a relief valve at full pressure drop, generating heat, excess amperage will dump over the third leg of the rheostat, generating heat as the power supply delivers full amperage at full current to the rheostat. A more efficient means of controlling the proportional coil is needed.

A much more effective method of partially shifting an armature is to send a pulse width modulated (PWM) current signal to the coil. PWM is a technique in which an on/off transistor located on the amplifier card turns current to the coil on and off very rapidly. Since the switching transistor turns off the unneeded current, unneeded current does not need to be dissipated, thereby reducing heat generation. Low frequency PWM is in the range of 100 to 400 hertz (Hz, cycles per second) while high frequency PWM is in the range of 4000 to 5000 Hz. As the pulse rate remains constant, the duration of the pulse is varied. For example, if the width of the pulse is 30% of its maximum duration, theoretically the valve should shift enough to deliver a 30% output of flow or pressure, whichever is being controlled. By the same token, if the width of the pulse is 80% of its maximum duration, the valve output should be at 80%. By varying the "on time" of the coil, the displacement of the spool or poppet is controlled.

Several of the factors influence the difference between the input current to the coil and the output from the valve. Spring hysteresis and spool or poppet frictional losses, generally termed "sticktion," as well as losses in the coil itself, are among these factors. In an effort to overcome sticktion and inertia, a low-amplitude high-frequency sine wave is oft times superimposed on the PWM signal. This extra signal is called "dither." The effect of the dither is to keep the spool or poppet in constant motion in an effort to overcome the response losses caused by inertia and the

sticktion. Ideally, oscillation caused by dither will not alter the output of the valve.

Solenoid controlled, pilot operated proportional directional valves, sometime called two-stage proportional directional valves, use either of two schemes to control the position of the main spool. If the valve is an open loop valve, the pilot valve is actually a dual pressure reducing valve; two proportional pressure reducing valves are contained in one body. The main spool is positioned as a function of reduced hydraulic pressure acting on one end of the spool as the spool is balanced against the opposing centering spring. Closed loop valves, in contrast, use an open loop proportional directional control valve as a pilot valve. The pilot valve is used to position the main spool. The LVDT connected to the main spool sends a feedback signal back to the amplifier card. The feedback signal is then analyzed by the summing amp, which processes any position error of the main spool; the pilot valve will be commanded to a null (centered) position after the main spool has been displaced the correct amount so as to produce the desired output flow. Figure 3 shows a two-stage open loop valve while Figure 4 shows a two-stage closed loop valve.


Figure 3 - Courtesy of Denison Hydraulics

Figure 4 - Courtesy of Denison Hydraulics

The name Moog is virtually synonymous with the wire feedback MFB servo valve design. Still, Moog has kept pace with the market and manufactures a wide range of EFB design valves. Moog's EFB valves may be broken down into valves based upon Atchley's jet pipe design and Moog's linear force motor design.

A linear force motor (LFM) contains two coils, centering springs (located to the right side of the illustration) and a means to adjust the position of the armature relative to the springs. An LFM is capable of developing approximately 45 pounds of force.

Figure 6 illustrates Moog's DDV (direct drive valve). The LFM is located on the right end of the valve and the LVDT and onboard electronics package is located at the left end of the valve. As with the dual coil proportional valves that were previously discussed, proportional closed loop control of direction and flow is achieved.

In contrast to most servo valves which conform to the ISO 10372 port circle mounting interface, Moog's DDV valves


Figure 6 - Courtesy of Moog

conform to the ISO 4401 mounting interface (D03 and D05). Unlike the proportional valves discussed earlier, the spool of a DDV valve does not directly contact the body of the valve. Instead, in keeping with traditional MFB servo valve design, the spool is contained in a precision machined sleeve. The inclusion of the sleeve allows for simpler machining of the sleeve lands in comparison to machining the internal body lands of a typical proportional valve. It is easier to optimize the spool and sleeve profiles in relationship to each other with this construction method. This allows the spool and sleeve to be cut to provide underlap, zero lap, or overlap (the lap of spool and sleeve relate to where the lands of the spool line up in relation to the lands of the sleeve; this affects internal leakage, and the control of the actuator).

The null position of the spool relative to the sleeve is factory adjusted by way of positioning the armature relative to the centering springs.

Moog's other EFB valve is based upon the jet pipe design, illustrated in Figure 7. In a jet pipe servo valve, system flow is directed through a jet pipe, which is basically a tube with a nozzle on the end. The flow exiting the jet pipe's nozzle is directed toward a receiver. The receiver has a hole into which fluid from the nozzle is directed. Inside the receiver, the hole branches into two passages. Each passage is connected to an end of the main spool. A force motor, which is sometimes referred to as the electrical bridge, responds to the electrical input from the control system. Force motors require a much lower amount of current than do proportional coils, and therefore, may be driven directly by a PLC. Once the force motor is actuated in one direction or the other, the angle of the jet pipe is changed, thus directing the flow toward one edge of the receiver. The flow from the nozzle is thus directed more toward one receiver passage than the other passage, creating a higher pressure in one of the passages. This higher pressure then acts upon that passage's spool end, shifting the main spool. The displaced spool then connects system pressure to one of the working ports while, at the same time, the opposite working port is connected to tank. The angle of the jet pipe is proportional to the input current applied to the force motor. The pressure in one receiver passage rises proportionally to the angle of the jet pipe, and the resulting spool displacement is proportional to the rise in pressure in the adjoining passage. An LVDT is used to close the loop, making a jet pipe an EFB servo valve.

Moog's highest flow valves use the two-stage jet pipe design. Figure 8 illustrates the cross-section of a two-stage jet pipe servo valve.


Figure 7 - Courtesy of Moog

Figure 9 shows a cross-sectional view of a two-stage MFB servo valve. As with the jet pipe design, a torque motor receives an electrical input. The torque motor armature moves in response to the electrical flux created by the current flowing through the

Figure 8 - Courtesy of Moog

coils of the force motor. A flapper is connected to the armature by a thin walled flexure sleeve. The flapper is positioned between two opposing nozzles. Fluid at system pressure flows through these nozzles. Branching off from the inlet passage of each of the nozzles is a connection to each end of the main spool. When no input current is applied to the torque motor, the flapper remains centered between the two nozzles, and pressure between each of the nozzles and the flapper remains balanced, and the pressures at the ends of the spool remain balanced. Also connected to the flapper is a thin length of stainless steel wire, termed a feedback spring. The free end of the feedback spring rests in a groove in the main spool. In this way, the spool position is mechanically fed back into the flapper nozzle assembly.

In figure 10, a current has been applied to the force motor, rotating the armature counterclockwise. This moves the right side of the flapper closer to the right nozzle, creating higher pressure in the right passage and lower pressure in the left passage. The


Figure 9 - Courtesy of Moog

higher pressure in the right passage acts on the right end of the spool applying force to displace the spool to the left. At the same time, lower pressure is acting on the left side of the spool, creating a force imbalance, facilitating the spool's movement to the left. As the spool moves to the left, the small jewel

on the end of the feedback spring pulls the feedback spring to the left, thereby inputting a feedback force on the flapper, counteracting the force generated by the torque motor, re-centering the flapper between the nozzles. Once the flapper is centered between the nozzles, pressure between each of the nozzles becomes equalized, and so does the pressure acting on the ends of the spool. Once the pressure at the ends of the spool equalizes, the spool stops moving, yet the spool remains displaced, controlling flow to and from the working ports. Figure 11 shows this actuated condition.

This is all simple enough so far. So, how does one decide which valve to use in a given application? The decision is ultimately based upon the performance and flow rate required from the valve. Performance is generally measured in frequency response. In short, frequency response is a measure of


Figure 10 - Courtesy of Moog

Figure 11 - Courtesy of Moog

how quickly a valve can control flow, for example, how quickly a valve can open to some given value, such as 80% or 100% and then close back a lower value, such as 20% or 0%. Even though all proportional and servo valve manufacturers publish Bode plots documenting the performance of their valves, each manufacturer seems to test to a different set of parameters making direct comparisons difficult, if not impossible. Additional measures of performance include step response, hysteresis, leakage, deadband, and linearity. A discussion of these topics would require an article in itself.

One question that often arises is whether or not a closed loop proportional valve will

perform "better" in a system than will an open loop valve. While it is true that some closed loop valves exhibit greater frequency response, examination of data sheets indicates that some of the closed loop valves aren't much faster than the open loop valves. This, of course, compares valves with equivalent flow ratings, for example, 10 gpm D03 valve to another 10 gpm D03 valve. One thing a closed loop valve will pretty much guarantee is that if the spool is commanded to 35% displacement, that's where the spool will be positioned. That doesn't necessarily mean that the valve will be delivering 35% of its flow rate. In the end, what matters is what the actuator is doing. If the need is for a given force, velocity, or position, the machine designer should design in an external feedback loop. Closing an external loop on a closed loop valve may create a "cross-talk" condition in which the closed loop of the valve fights the closed loop on the valve/actuator combination. The required performance may often be achieved with an open loop proportional valve and closed loop control of the actuator without incurring the expense of both a closed loop valve and a closed loop control system.

One might wonder, "Why, for a given flow rate category, do servo valves have higher frequency response?" Note that in a dual coil type of proportional valve, whether open or closed loop, the coils work against centering springs as they position the spool. Even if the coil is optimized to act very quickly, if the current to the coil is lowered, in order to lower the flow rate of the valve, the component that moves the spool to a position of less displacement is the spring. In the opposite condition, if the coil acts very rapidly to move the spool into a position of greater displacement, and the spring cannot provide enough opposing force, the spool will overshoot. The maximum performance envelope of the valve is somewhat limited by the dynamics of the centering springs.

In contrast, MFB and jet pipe servo valves are controlling the spool displacement with high pressure fluid. In a servo valve system, one third of the system pressure is used for control of the spool. In a 3000 psi system, 1000 psi is used to control the spool and 2000 psi is left over for doing work. This represents a fairly significant loss of pressure with which to do work, but what is gained is system response. The high pressure fluid at the ends of the spool functions as a set of high force springs.

As mentioned earlier, the linear force

motor used by Moog DDV valve in their DDV valve will develop up to 45 pounds of force. This allows for the use of relatively high force centering springs, in comparison to the spring rates used in dual coil proportional valves, which in turn, allows for crisper control of the spool displacement (position) than if springs of lower force were used.

Smaller valves tend to have a higher frequency response because their spools have less mass. A small mass is easier to control, accelerate and decelerate, than a large mass.

Still, in the end, the decision comes down to choosing a valve that will satisfy the motion profile requirements of the machine.

Up to this point we've examined directional valves, but what about pressure and flow controls? Proportional technology has been applied to these valves as well. In the case of pressure controls, most valves are based on open and closed loop proportional technology.

For low or pilot flow, direct acting pressure controls are commonly available. For higher flows, as with conventional pressure control valves, a two-stage design is used. Figure 12 illustrates a pilot operated relief valve with closed loop proportional control and a maximum pressure override.

Figure 13 illustrates a closed loop flow

Figure 12 - Courtesy of Denison Hydraulics

control valve.

Several sources for further study of the technology used in proportional and servo


Figure 13 - Courtesy of Denison Hydraulics

valves are:

"Electrohydraulic Valves… A Technical Look," Moog Industrial Controls Division, CDL6566 Rev D 500-170 402

"Electrohydraulic Proportional Valves and Closed Loop Control Valves Theory and Application," Bosch, 1987764002, 10.92

"Electronics in the Mobile Industry," HydraForce, Inc., www.hydraforce.com/Electro/Elec-pdf/3-561-1.pdf

"Principles of Proportional Valves," Eaton Corporation, GB-9042A-4.0-11.94 SCS-DGV. ●

Typical Proportional and Servo Valve Control Systems

how the components are arranged

Proportional and servo valves are used to control the position, velocity, or force of an actuator. In some cases, two or more of those parameters are controlled. For example, as a cylinder is extended, its velocity might be controlled. At the end of the cylinder's stroke, the servo system may be used to either position the cylinder or to control the force it exerts on the work piece. The layout of the control systems is the same for both pneumatic and hydraulic circuits.

"Electrohydraulic Valves… A Technical Look," published by the Industrial Controls Division of Moog, does an excellent job of detailing these three types of control systems. These circuits illustrate closed loop systems. Open loop systems would, of course, not include the feedback components; since a feedback loop would not exist, an open loop circuit would not really be considered to be a "servo" system.

A Typical Position Servo System:

As with any servo control system, the servo amplifier receives a command input. The amplifier sends an output to the valve, energizing the appropriate valve coil, thereby actuating the valve which provides fluid flow to the actuator. A position transducer attached to a cylinder, or a rotary encoder if a fluid motor is being controlled, sends a signal back to the servo amp. The servo amp compares the position of the actuator to the value of the command input. Any difference between the two values produces an error signal which is used to change the input signal to the valve until the actuator is positioned per the amplifier's input signal.
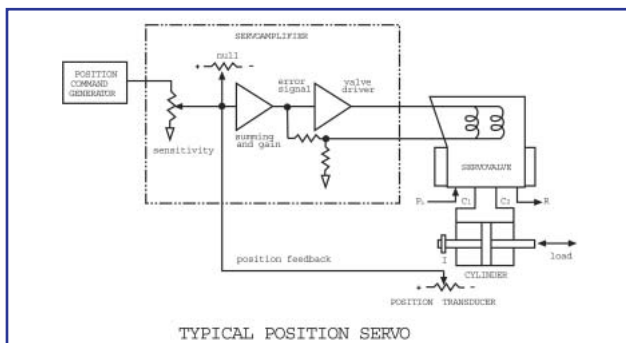

TYPICAL POSITION SERVO

Illustration Courtesy of Moog

A Typical Velocity Servo System:

Velocity control systems are used with both fluid motors as well as with cylinders. The servo amplifier must be capable of calculating velocity (distance/time) in order to provide an error signal for further processing. As with a position control system, the servo amplifier controls the position of the valve to provide the desired output. One difference in the spool output of the valve in a velocity circuit, relative to the position of the spool in a position circuit, is that in a velocity circuit, when the error is zero, the spool is shifted, metering flow to and from the actuator. In a position servo system, when the actuator is in the desired position, the valve spool is centered, thereby blocking flow to the actuator.


TYPICAL VELOCITY SERVO

Illustration Courtesy of Moog

A Typical Force Servo System:

Force or torque is a function of the load on the actuator. A servo pressure control valve is commonly used to control pressure in a circuit, though a directional valve may be used in certain instances. The force or torque may be sensed by a load cell or by a pressure transducer. The feedback signal from the sensor is analyzed by the servo amp which controls the valve. As with a velocity circuit, the valve is actuated (shifted) in order to achieve the desired pressure. ●
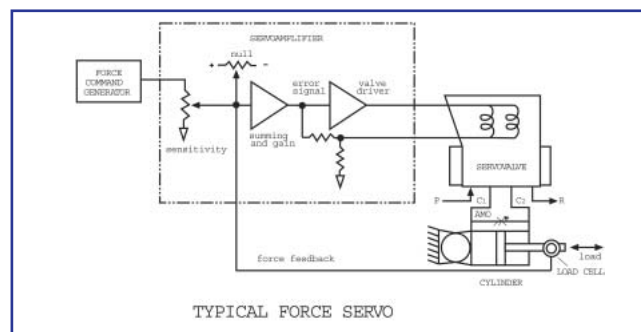

TYPICAL FORCE SERVO

Illustration Courtesy of Moog