Centrifugal Pump Efficiency – What, How, Why & When ?

Joe Evans, Ph.D   http://www.pumped101.com

Introduction

In this tutorial, we will investigate several aspects of centrifugal pump efficiency. First I will define efficiency and give some examples. Next we will examine some of the design criteria that ultimately dictate the efficiency exhibited by a particular pump. We will also try to make that somewhat nebulas quantity, known as specific speed, more meaningful. I will also show its effect on the shape of a pump's performance and power curves. Finally, we will discuss the importance of (or, sometimes, unimportance) of efficiency as it relates to a particular application or process. We will also illustrate the relationship of efficiency, head, and flow as they apply to both steep and flat performance curves and their roles in constant and variable speed applications. We will end with a brief look at the combined efficiency of a pump and its driver.

What is Pump Efficiency ?

When we speak of the efficiency of a any machine we are simply referring to how well it can convert one form of energy into another. If one unit of energy is supplied to a machine and its output, in the same units, is one-half unit its efficiency is 50%. As simple as this may seem, it can still get a bit complex because the units used by our English system of measurement can be quite different for each form of energy. Fortunately, the use of constants will bring equivalency to these, otherwise, diverse quantities.

A common example of such a machine is the "heat engine" which uses energy in the form of heat to produce mechanical energy. This family includes many members but, the internal combustion engine is one with which we are all familiar. Although this machine is an integral part of our every day lives, its effectiveness in converting energy is far less than we might expect. The efficiency of the typical automobile engine is around 20%. To put it another way, 80% of the heat energy in a gallon of gasoline does no useful work. Although gas mileage has increased, somewhat, over the years that increase has as much to do with increased mechanical efficiency as increases in engine efficiency itself. Diesel engines do better job, but still max out around 40%. This increase is due, primarily, to its higher compression ratio and the fact that the fuel,

under high pressure, is injected directly into the cylinder at the top of the compression stroke. Gasoline engines, on the other hand, are limited to lower compression ratios because fuel enters the cylinder prior to the compression stroke.

In the pump industry, much our work involves two extremely simple, yet efficient machines - - the centrifugal pump and the AC induction motor. The centrifugal pump converts mechanical energy into hydraulic (flow, velocity, and pressure) energy and the AC motor converts electrical energy into mechanical energy. Many medium and larger centrifugals offer efficiencies of 75 – 90% and even the smaller ones usually fall into the 50 – 70% range. Large AC motors, on the other hand, can approach an efficiency of 97% and any motor, five hp and above, can be designed to break the 90% barrier.

The overall efficiency of a centrifugal pump is simply the ratio of the water (output) power to the shaft (input) power and is illustrated by the equation below.

$\eta = P_W / P_S$ where η is efficiency, P_W is the water power, and P_S is the shaft power.

In the US, P_S is the power provided to the pump shaft in brake horsepower and P_W is

$P_W = (Q \times H) / 3960$ where Q is flow in GPM and H is head in feet.

The constant, 3960, converts the product of flow and head (foot-pounds or the more politically correct term pound-feet) into BHP. These equations predict that a pump, that produces 100 GPM at 30' of head and is powered by a motor that produces 1 BHP will have an overall efficiency is 75.7% at that point. The second equation will also allow us to compute the BHP required at any point on a pump's performance curve if we know the efficiency. We will see some examples of this in the last section of this tutorial.
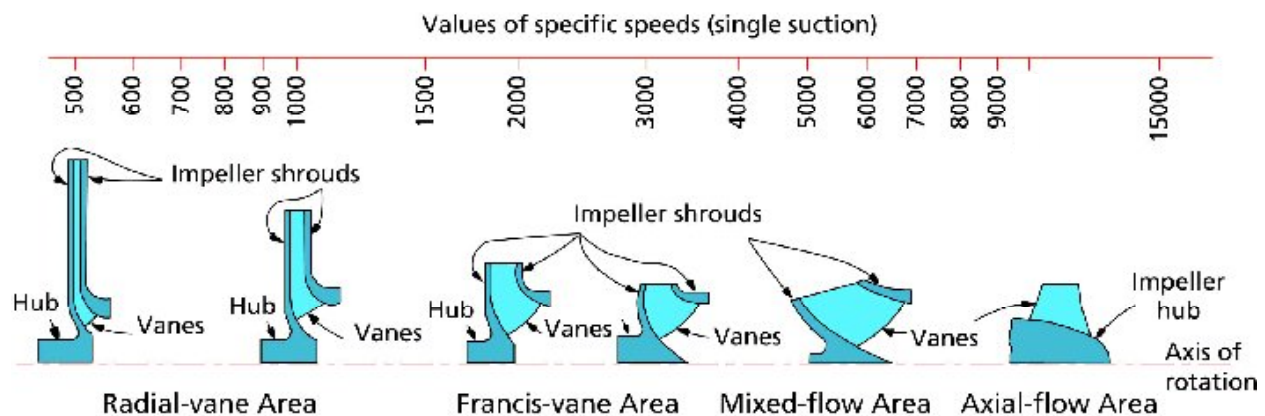
How is Pump Efficiency Attained ?

If you think about it, the centrifugal pump has a lot in common the induction motor when it comes to the design phase. That commonality is that both have only two major components that can be modified by the designer. In the case of the motor it is the rotor and the stator and for the pump it is the impeller and the volute (or diffuser). Of course the friction produced by bearings and other mechanical components (seals, stuffing box, etc) also affect pump efficiency, but the impeller and volute have the greatest influence. Lets start our investigation of centrifugal pump efficiency with the impeller.

The laws of affinity tell us quite a bit about the inner workings of an impeller. We know that, for any given impeller, the head it produces varies as the square of a change in speed. Double the speed and the head increases by a factor of four. If you keep speed constant, the same rule holds true for a change in its diameter. The flow through an impeller follows a similar rule but, in this case, its change is directly proportional to

the speed or diameter change - - double the speed or diameter and flow is doubled. Actually, when we talk about a change in rotational speed or impeller diameter, we are really referring to its peripheral speed or the speed, in feet per second, of a point on its outer most circumference. It is this speed that determines the absolute maximum head and flow attainable by any impeller (see the "UP and Down" Puzzler for an explanation of the falling body equation and how it relates to centrifugal pump head).

The head produced by an impeller is almost entirely dependent upon its peripheral velocity but, flow is influenced by several other factors. Obviously, the width and depth (cross sectional area) of the flow passages (vanes) and the diameter of the impeller eye are important considerations as they determine the ease with which some volume of water can pass through the impeller. Other factors such as vane shape also influence an impeller's performance. But, if you wanted to design an impeller from scratch where the heck do you start? Do you just take a wild guess about dimensions and shapes, make some samples, and then test them? Well, in the early days that is exactly what we did. Today, however, we can draw on years of experience and, at least, find a suitable starting point for our design. And, that starting point is something called Specific Speed.

Specific Speed is often confusing to many of us because when we see the word speed, we immediately think "impeller speed". Actually, it is just a number (often dimensionless like the Reynolds number which is used to predict turbulent flow) that refers to a particular impeller design or geometry without respect to its size (capacity). It uses the knowledge we have gained over the years to categorize the performance of various impeller designs based upon our application requirements. The chart below shows the relationship of the numerical value of specific speed to a particular impeller design.

Values of specific speeds (single suction)

Radial-vane Area   Francis-vane Area   Mixed-flow Area   Axial-flow Area

The lower values (500 to 1000) on the left describe the changing geometry of the radial vane impeller while the higher values (10000 – 15000) on the right equate to true axial

flow impellers. Those in the middle (1500 – 7000) are typical of the Francis vane and mixed flow (which show both radial and axial characteristics) impellers. The cross sectional pictures on the chart show that, as specific speed increases, the impeller inlet or eye diameter increases and eventually approaches or equals that of the vane outlet. The flow passages also increase in size at a corresponding rate.

I think you will agree that while this is a nice comparison, what use is it to the pump designer? Well, there happens to be an equation that relates specific speed and its corresponding geometry to those real application values of head, flow, and rotational speed. That equation is

$$N_s = n \times \sqrt{Q} / H^{0.75}$$

where N_s is the specific speed, n is the pump rotational speed in RPM, Q is flow in GPM, and H is head in feet. We can use this equation to determine which impeller design can best match the requirements of a particular application.

Suppose, for example, we need an impeller that will produce 1000 GPM at 200 feet of head. If we enter these values in Q and H and also enter a motor speed of 3600 rpm we obtain a specific speed of 2140. The impeller would have a geometry similar to the Francis vane impeller seen on the chart at the 2000 point. An 1800 rpm motor would lower the specific speed to 1070 and would have a geometry similar to the radial vane impeller shown beneath the 1000 point. At 1200 rpm specific speed is 714 and the impeller would look like a hybrid of the two impellers seen to the left of the chart.

The chart below shows illustrates how specific speed can provide us with several predictions as to the performance of a particular impeller design. As a rule of thumb,

Values of specific speeds (single suction)
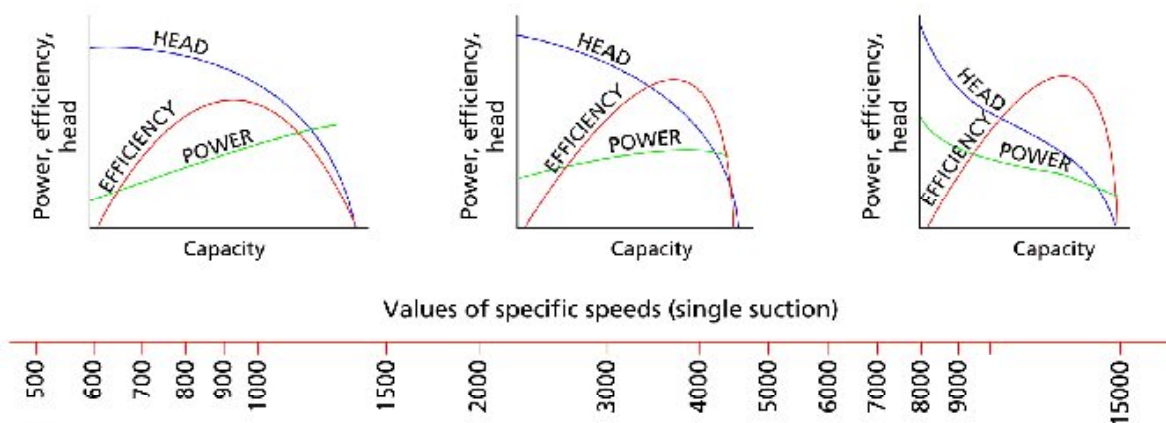
impeller efficiency reaches its maximum at a specific speed between 2000 and 3000 although favorable efficiency can occur at almost any speed. Also the area around the Best Efficiency Point (BEP), or design point, tends to be flatter and broader as specific

speed decreases. (Impeller efficiency also increases with pump rotational speed, especially high speeds, but that increase is not as pronounced at speeds of 3600 rpm and below.) Specific speed also effects the shape of the head-capacity curve. Low specific speeds (500 – 2000) produce relatively flat curves while high speeds (5000 +) produce extremely steep curves. Intermediate speeds produce curves that fall in between these extremes. These results are due to the vane shape (flat versus backwards curved) at various specific speeds. (We will discuss curve shape in more detail in the next section.) Finally, specific speed provides us with one more prediction - - the characteristics of the power curve. At specific speeds below 3500, power drops as flow is reduced and is at its minimum at shut off head. The power curve remains relatively flat, across the entire head-capacity curve, between 4000 and 4500 and rises towards shut off at specific speeds above 5000. At speeds above 9000 the power and head-capacity curves almost parallel one another. Stated differently, power is greatest at shut off and is at its minimum at full flow.

Once a particular impeller geometry is chosen, the pump designer can go through a comprehensive mathematical analysis that will allow him to derive all of the impeller dimensions and angles necessary to meet the design point. To say the least, this is an arduous task. If you would like to review a comprehensive example of how this is done, see pages 2.23 - 2.31 of the second edition of Pump Handbook (McGraw-Hill).

The shape and spacing of the impeller vanes obviously have a large effect upon efficiency. Although the ideal pump would have an infinite number of vanes, the real world limits us to 5 – 7 for typical pumps and even fewer for pumps that handle larger solids. Also, flow would be exactly parallel to the vane surfaces but that doesn't happen either. But oddly enough, if the designer follows some well documented rules, impeller vane efficiency losses remain relatively flat (about 2.5%) across a specific speed range of 500 to 7000. Disk friction, which is caused by contact between the pumpage and the impeller shrouds and hub surfaces, can reduce impeller efficiency another 4 to 15% at specific speeds below 2000 but decreases to 2% or less at 3000 and above. So, depending upon its design, the impeller can reduce overall pump efficiency by as little as 4.5% or as much as 17.5%.

The volute also plays a role in pump efficiency. At specific speeds below 2000, its losses range from 1 to 2.5% but losses can approach 10% at speeds over 5000. Typically, volute design begins with the throat, as its cross sectional area will determine the flow velocity out of the volute. Flow through the throat and other portions of the casing follows the law of constant angular momentum so the designer will try to avoid abrupt changes its nearly circular geometry while gradually increasing its volume.
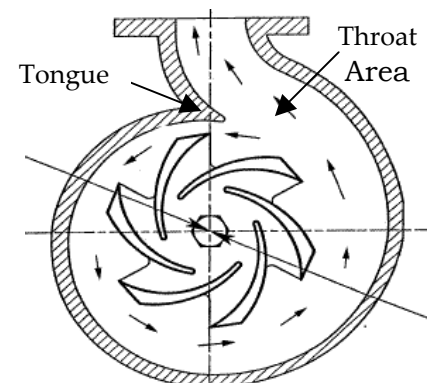
Another critical area of the volute is the clearance between the outer circumference of the impeller and that of the volute tongue or cutwater. As this distance becomes larger, an increasing volume of pumpage escapes entry into the volute throat and is recirculated into the volute case. The smallest distance possible, that does not give rise to pressure pulsations, will produce the best efficiency. As a rule of thumb, 5 to 10% of the impeller radius tends to be a safe value. In the next section we will discuss this in more detail when we compare the efficiencies that result from trimming an impeller versus changing its rotational speed.

It is debatable as to whether the volumetric efficiency of a centrifugal pump is a function of the volute or the impeller (it is probably both) but I will include its effect here. Volumetric efficiency represents the power lost due leakage flow through the wear rings, vane front clearances (semi open impeller), and balancing holes in the rear shroud of an impeller. As a rule of thumb, leakage increases with a decrease in specific speed, flow, or a combination of the two. For example at a specific speed of 500 and a flow of 100 GPM, leakage can account for as much as 7% of the total power consumed. At 2000 GPM it is reduced to about 2%. At higher specific speeds and flows it can be as low as 1%.

The final piece of the pump efficiency puzzle is that of mechanical losses, although some of these losses are not always included in published efficiency curves. In the case of a frame mounted pump, these losses are caused by the shaft bearings and the mechanical seal or packing. For close coupled pumps, bearing losses are figured into the motor efficiency. Again the rule of thumb follows that of volumetric efficiency, and losses increase as flow and / or specific speed decrease. If we use the same values of specific speed and flow, as in the volumetric example above, we could expect losses of 5% and 1% for a frame mounted pump. At higher specific speeds and flows, mechanical losses drop well below 1%.

Why and When is Efficiency Important?

OK, now we are all on the same page as to the definition of pump efficiency and we have some idea of the pump designer's ability to control efficiency during the design phase. But, is it the most important component in pump design? Should we always shoot for the best possible efficiency when we design a pump?

The importance of pump efficiency is entirely related to the use of energy. As the cost of electricity and other energy sources continue to rise, it just makes good sense that we use it as efficiently as possible. Whenever possible, we should select the most efficient pump available as it will usually justify its, potentially higher, first cost during its useful life. Notice that I did say "whenever possible".

Several factors can influence our decision about the importance of pump efficiency. Sometimes it is purely economic - - the market may not be willing to pay the price for higher efficiency. There are also times when a higher efficiency pump may not perform as well as one of lower efficiency. And, there are instances where we just cannot attain a reasonable efficiency based on the head and flow required. Lets take a look at several examples.

A good example of the role of economics is the residential pump. The efficiency of most fractional HP domestic booster and circulation pumps falls into the 50% range (Their motors are not much better, but we will address that a little later). These pumps can be designed to operate at higher efficiencies but, the cost would scare most homeowners away. And, if they are used only occasionally, the energy savings may not justify the additional cost.

Unfortunately this, low first cost, mindset often spills over into sectors that could easily justify better efficiency (certain sectors of the HVAC market come to mind). For example, a typical, "low cost" pump will not incorporate suction wear rings and, as wear progresses, more and more suction recirculation occurs and efficiency decreases. The only way to fix this is to replace the impeller and possibly the volute. Incorporation of simple, flat wear rings into the volute and the impeller adds only a few percentage points to the pump cost but allows lower cost "efficiency" maintenance as it is needed.

Small, "two port" impeller sewage pumps must sacrifice a certain amount of efficiency in order to pass solids without clogging. A three or four vane impeller would provide far better efficiency but the size of the solid it passed would be greatly reduced compared to a two vane impeller. In this case, efficiency becomes secondary to the requirements of the application. The recessed impeller sewage pump also offers some real advantages in certain installations, but is regarded by many as a poor choice because of its very low efficiency (35 – 50%). It uses a two step process (the impeller creates a vortex and the vortex creates flow) that wastes quite a bit of energy but, it will pass stringy material and larger diameter solids than two vane pumps of equivalent size.   In a small, commercial or municipal application (say 5 – 10 HP) which is more costly - - wasting 15 - 20% more energy in the case of the vortex pump or the weekly expense of pulling a standard pump for cleaning?

Finally, there are application design points where reasonable efficiency cannot be attained but a pump is still required. Suppose some million dollar process line cannot use a positive displacement pump but, instead, requires a centrifugal pump that can

deliver 20 GPM at 3000' of head. Would we really care if a single stage pump had to be driven at 23,000 rpm and that its efficiency was less than 25%? Probably not, and there are far more of these types of applications than you might suspect.

When is it important ?

The majority of municipal and commercial, clear water, applications are not restricted by the limitations outlined above and these pumps should be selected with efficiency in mind. If a pump is going to operate at a constant flow and head, it should be selected to operate as close to its BEP as possible. An example of this type of application would be a pump that feeds a municipal water tank. Since elevation and pipeline friction remain constant, the pump can be sized to meet its requirement at the best possible efficiency.

But, not all pumping applications have a constant flow and the shape of the performance curve can often be as important as BEP itself. An example of this type of application is the constant speed booster pump. In these applications, a pressure reducing valve (PRV) throttles the pump when demand decreases in order to maintain some constant pressure. Pressure on the pump side of the valve follows the performance curve pressure while a preset, constant pressure is maintained on the discharge side of the valve. What pump design best fits this application?

The chart below shows the performance curve for a pump that could be used in the booster application described above. The application calls for a pump that can provide a constant pressure boost of 130' over a flow range of 100 to 300 GPM. The number above each of the flow points on the performance curve is the BHP required at that

Constant Speed Pump 1
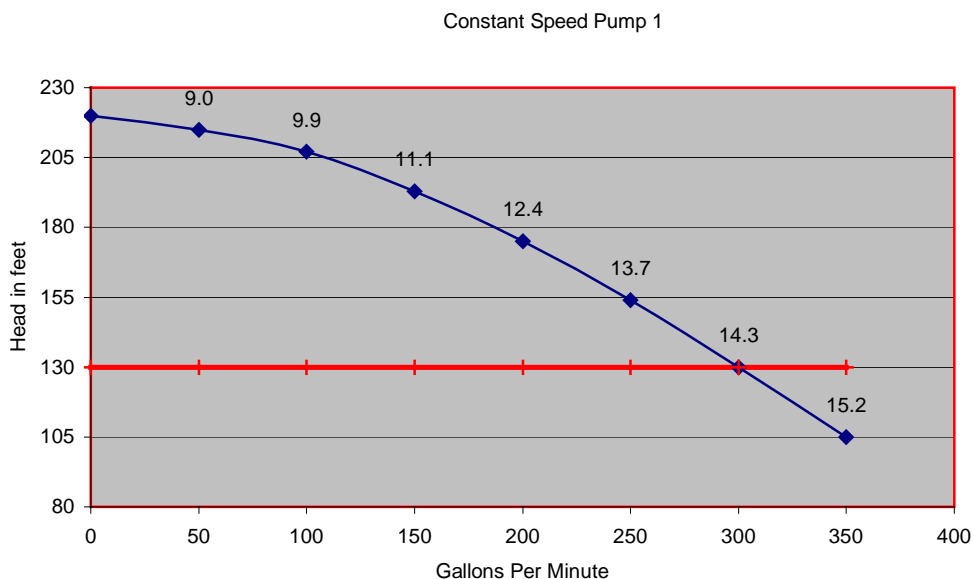
point and the red line is the desired system pressure. The BHP required to operate this

pump ranges from 14.3 hp at full flow to just under 10 hp at 100 GPM and follows the power curve predictions of specific speed we saw earlier.

Below is a chart that shows the performance curve for an alternate choice (Pump 2). The efficiencies for both Pumps 1 & 2 differ by no more than one or two percentage points over their usable range of flow. In the case of Pump 2, the BHP required at full flow is 14.5 hp - - a little higher than that of Pump 1. But, take a look at the power required at lower flows. The "flatter" curve produced by Pump 2 requires less power as flow decreases and can save a significant amount of power at the intermediate flows. Even though it is slightly less efficient at BEP, it is a far better choice for this particular application.

Constant Speed Pump 2
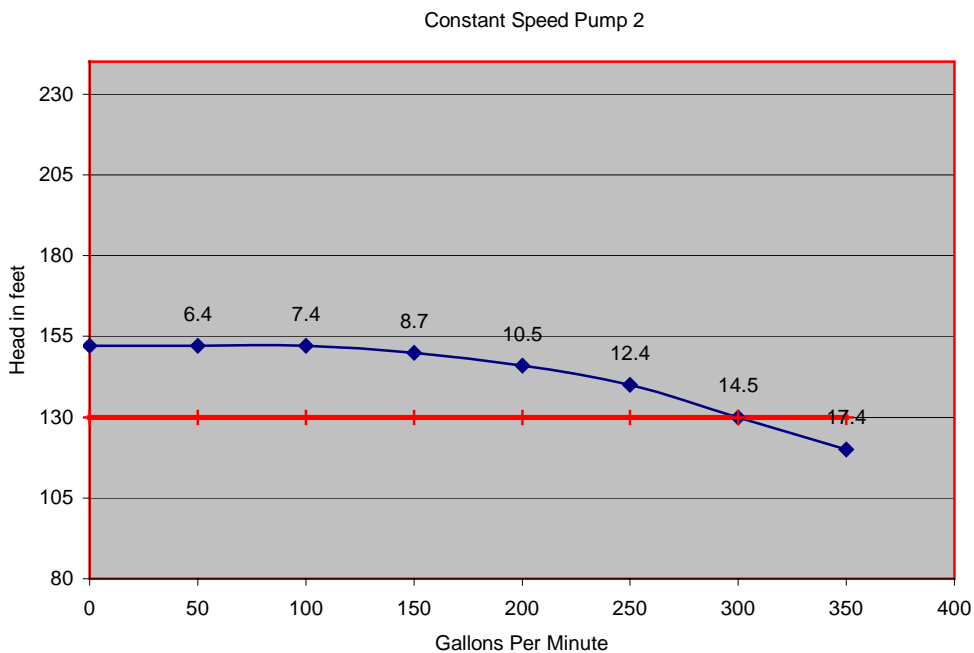
The reason the flatter curve consumes less power than the steeper curve can be illustrated by the following equation.

BHP = (Q x H) / (3960 x efficiency)   where Q is flow and H is head

BHP is directly proportional to both flow and head and inversely proportional to efficiency. Although the flows and efficiencies are the same or similar for both pumps in our example, the higher heads (at flows under 300 GPM) produced by Pump 1 require more power than the heads produced by Pump 2. You can use this equation to compute BHP at any point on a performance curve. See, I told you that there is more to pump selection than just efficiency!

As you review various pump curves for their fit in a constant speed booster application, you will notice another efficiency trait. The efficiencies on either side of BEP are more stable for some pumps than others. For example, a particular pump with a BEP of 77% at 400 GPM is able to maintain 70% efficiency over a range of 250 to 550 GPM. Another pump, with a similar BEP, may drop below 70% much more quickly. The change in efficiency around BEP has a lot to do with the impeller's vane angle at its entrance. At BEP, flow is nearly parallel with the vane but, as flow increases or decreases its entrance angle into the vane also changes. Small changes in the entrance angle design can affect both BEP and the efficiency values on either side. The figure on the right shows how the entry angle changes under various flow conditions.
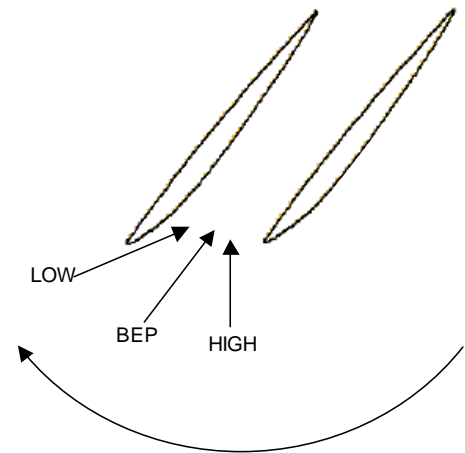
Now, suppose that these same two pumps were candidates for installation in a variable speed booster system (See Variable Frequency 101 if you are not familiar with VFD operation). Would our final selection be different? The chart below shows the results when Pump 2 is operated under Variable Frequency control. The head rise from full flow to shut off is about 15% and would allow just a 4 Hz reduction in speed if we are to maintain a constant pressure of 130'. Unfortunately, this reduction is not nearly enough to achieve a reasonable power savings over that of the constant speed booster. At


VFD Control Pump 2
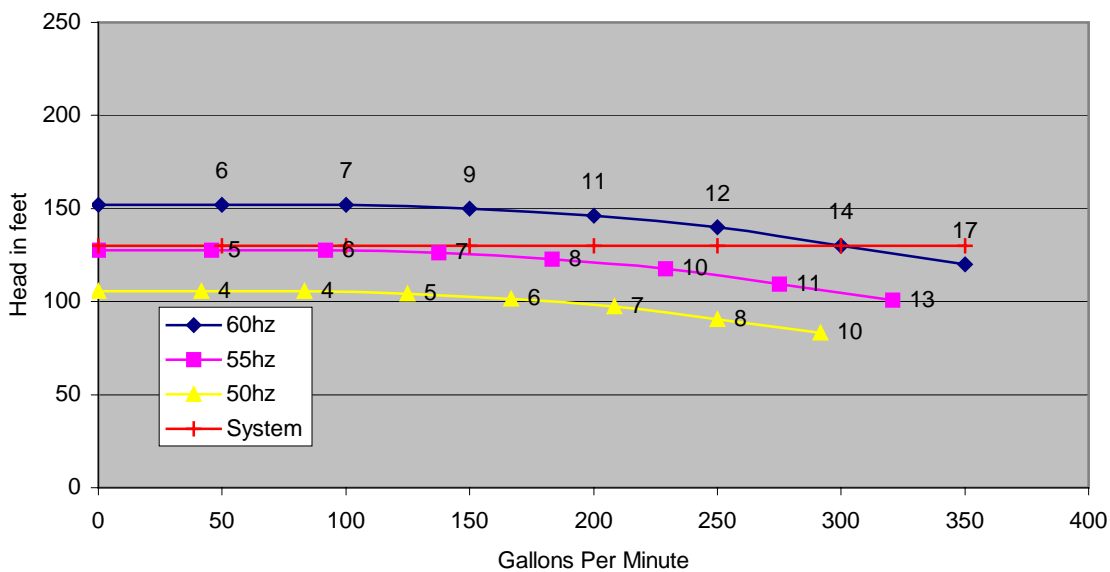
flows of 200 and 100 GPM power savings would be just 2 and 1 HP respectively. Also,

this limited frequency range would not allow precise control of the pump over its rather broad flow range and a significant amount of frequency "hunting" could occur.

But, Pump 1 has a head rise to shut off that is much greater and when operated under VFD control (shown below) it can perform quite well. From 100 GPM to full flow it will operate over a range of 47 to 60 hz and the power savings at each reduced flow point is significantly greater than that of Pump 2 running at constant or even variable speed. Variable speed operation offers an additional benefit. The unbalanced hydraulic forces that exist at lower flows in the constant speed booster are greatly reduced. So the steep curve looses to the flat curve when installed in constant speed boosters but wins in a variable speed application. Keep this in mind when evaluating pumps for booster applications.

VFD Control Pump 1
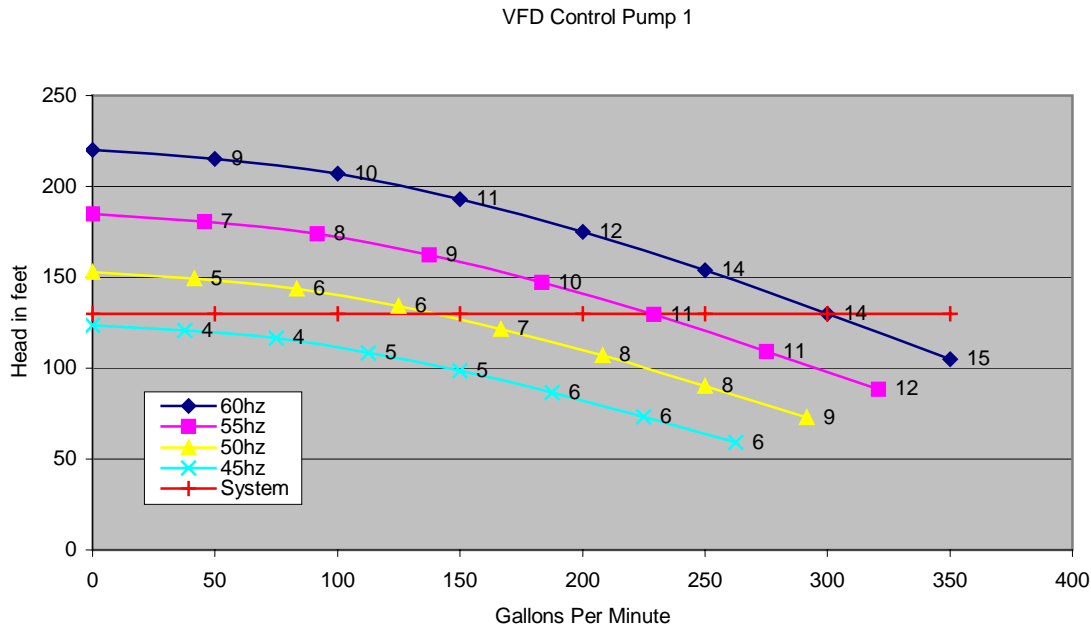
There is, however, a variable speed application where flat curves excel. Closed loop circulation is a very common application in the HVAC market. In these applications the pump sees no head due to elevation and all it has to overcome is the friction in the loop. As flow is reduced, head due to friction tends to fall quickly and pump speed can be greatly reduced.

The chart on the following page shows the energy savings that can be attained by utilizing flat curves in closed loop applications. The system curve (in red) shows that the friction in the loop ranges from about 55' at 1400 GPM to about 15' at 700 GPM. This equates to a speed range of 60 to 31 hertz and results in a power savings of

Closed Loop VFD Controlled
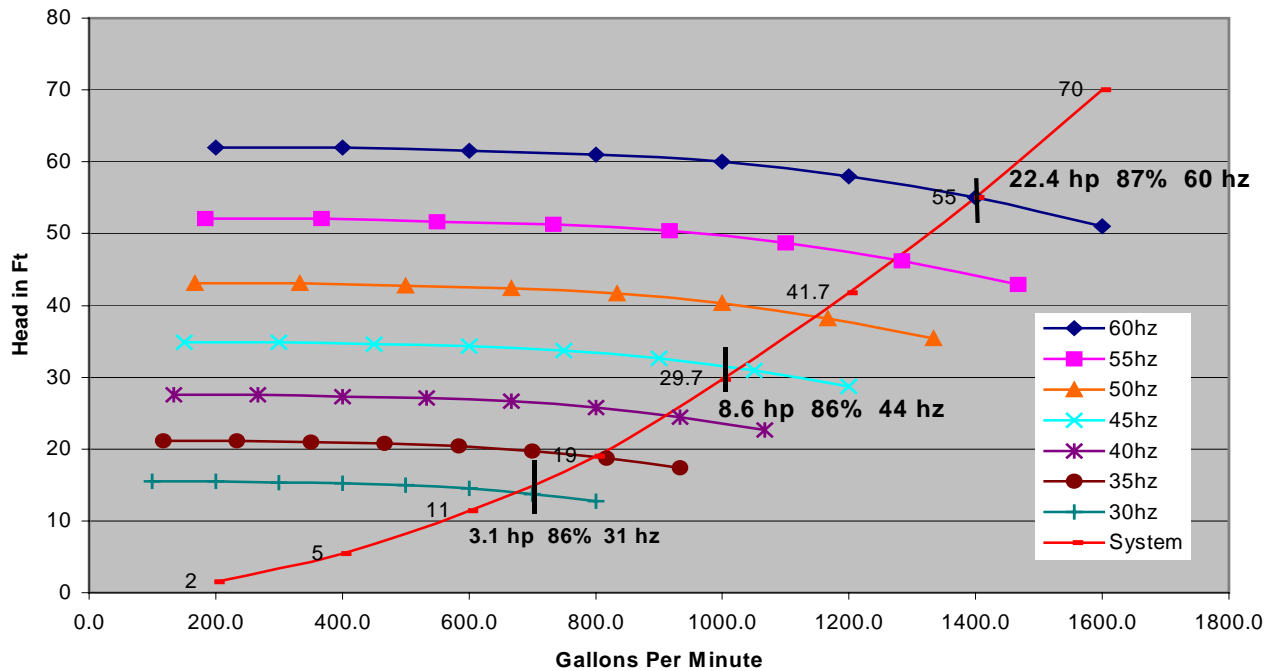
approximately 86% at the minimum flow of 700 GPM. All points in between show a similar savings.

Another centrifugal pump efficiency trait is illustrated by this closed loop example. Notice that the efficiency at 700 GPM is only 1% less than that of the true BEP at 1400 GPM. When the speed of a centrifugal pump is reduced its efficiency, at any capacity point on the 60 hertz performance curve, follows that capacity at the lower speed. In other words, efficiency moves to the left with capacity as speed is reduced. This also occurs in constant pressure – variable flow applications but it is more apparent in applications where both head and flow are variable.

To a certain extent, we will see the same result when an impeller is trimmed - - efficiency will move to the left with capacity. But, at some point, the distance between the impeller periphery and the cutwater causes unacceptable recirculation and efficiency begins to drop. Although small impeller trims can be effective, a change in rotational speed is the most efficient means of changing a pump's capacity and head. And, it for this reason that variable speed pumping systems will continue to evolve.

Combined Efficiency

Finally, lets take a look at something I call combined efficiency as it is probably more important than pump efficiency alone. I define this efficiency as the combination of the

hydraulic efficiency of the pump and the mechanical, heat, or electrical efficiency of the device that is driving it. In the case of an electric motor driven pump, it is called the "wire to water" efficiency and refers to how well the two machines work together to produce hydraulic energy from electrical energy. The reason combined efficiency is so important has to do with the mathematical relationship between the two individual efficiencies.

Suppose we have a pump with a BEP of 80% that is driven by a motor with an efficiency of 90%. If you were to ask the average person to calculate the combined efficiency of the two machines they would typically add the two efficiencies together, divide by two, and give you an answer of 85%. If this were true, combined efficiency would be a non issue but unfortunately, it is not the average but the product of the two efficiencies. Individually, 90% and 80% look pretty darn good but when you multiply one by the other, the resulting efficiency drops to 72%! Still, over the life of an installation, a relatively small increase in the combined efficiency can make a big difference in energy costs.

The Comprehensive Energy Policy Act (EPACT), that became law in 1997, set some minimum efficiencies for general purpose motors from 1 – 200 HP and speeds of 1200 – 3600 RPM. Higher horsepower, lower speed, and definite purpose motors were not required to meet these minimum efficiencies. Interestingly enough, one of these definite purpose motors is the close coupled pump (CCP) motor. These motors are some of the most common pump drivers in use and why they were excluded, I do not know. Their unregulated, efficiencies range from 80% at 3 HP to about 89% at 30 HP. Most manufacturers do, however, offer higher efficiency models that range from 86 to 94% over the same HP range and their use can make a big difference in pump energy consumption over time.

You could probably make a pretty good case that pump efficiency is not too important if the pump is to be driven by a gasoline engine. Although an 80% efficient pump should save quite a bit of energy over one that is 65% efficient, the gas engine (approximately 20%) brings their totals down to 16% and 13% respectively. It may be hard to justify a higher initial pump cost for such a small energy savings, unless the pump is used frequently and for long periods of time.

Joe Evans   November 2005