PAPER FOR CIBSE CONFERENCE 1999

Title : Computer aided design for variable speed pumping circuits

Author: Christopher John Parsloe BSc CEng MCIBSE

Company: BSRIA

Computer aided design for variable speed pumping circuits

1. The Problem

Constant flow heating and chilled water systems are generally well understood by engineers. Full load design conditions can be predicted, and components such as pumps, valves and terminal units can be selected accordingly. However, in systems where there is the facility to vary pump speed, many of the certainties are removed and the engineer is forced to take into consideration part load operating conditions.

By installing variable speed pumps, three port valves which divert flow away from terminal units are replaced by two port valves which modulate terminal branch flows down to zero. For these systems the main problems faced by the engineer are:

Where best to site the pressure sensor to control pump speed?
How to size the two port valves?
Whether or not there is a need for regulating valves and proportional balancing?

This paper addresses these issues based on solutions identified by BSRIA's Pipe Network Analyser software which was published alongside BSRIA Application Guide AG14/99 "Variable Speed Pumping in Heating and Cooling Circuits" in July 1999. The software and publication were part funded by the Department of the Environment, Transport and Regions under the "Partners in Innovation" programme.

This paper summarises the main conclusions. A more complete understanding, and proof for the conclusions can be gained by referring to the Guide and program.

Advantages and disadvantages of variable speed pumping

There are significant benefits to variable speed pumping in terms of both energy cost and capital cost. The main energy savings derive from the pump law relationship that (for a system with fixed resistance) halving the flow rate, would reduce the pump power to roughly one eighth its previous value. Significant capital cost saving opportunities are likely to arise from the potential to take into account diversity of load during pipe and pump sizing, and the possibility that flow regulating devices can be omitted from parts of the system.

However, there has always been uncertainty regarding the best approach to the design of systems with variable speed pumps, with no clear methodology laid down by any industry authority. The consequences of poor design such as noise at control valves, control valves that won't shut off, erosion of control valve seats, poor response of the pump speed controller etc have contributed to making designers hesitant to attempt a variable speed solution.

2. The Approach

Computer simulation

It was decided early on in the research that the best way to gain a clear understanding of the behaviour of pipework systems under varying load conditions would be to model flow distributions using a computer simulation.

The fluid dynamics theory for pipe network analysis is well documented and is available in a number of text books. In simple terms the technique involves making small adjustments to the assumed flow rates in each sub-circuit until the pressure losses in the sub-circuits are in balance, i.e. the pressure losses between any two points are equal for all routes through the pipe network.

A simple computer program was developed which enabled the input of pipework design information, and then sized the pipework and ran the network analysis calculation. Since the calculation was theoretical and depended on pipework resistance data taken from the CIBSE Guide C4, the next step was to test the predicted flow data against real flow information taken from the commissioning results for a real system. The results of this comparison are shown in Figure 1.

Figure 1: Comparison of measured flow rates with computer predicted flow rates

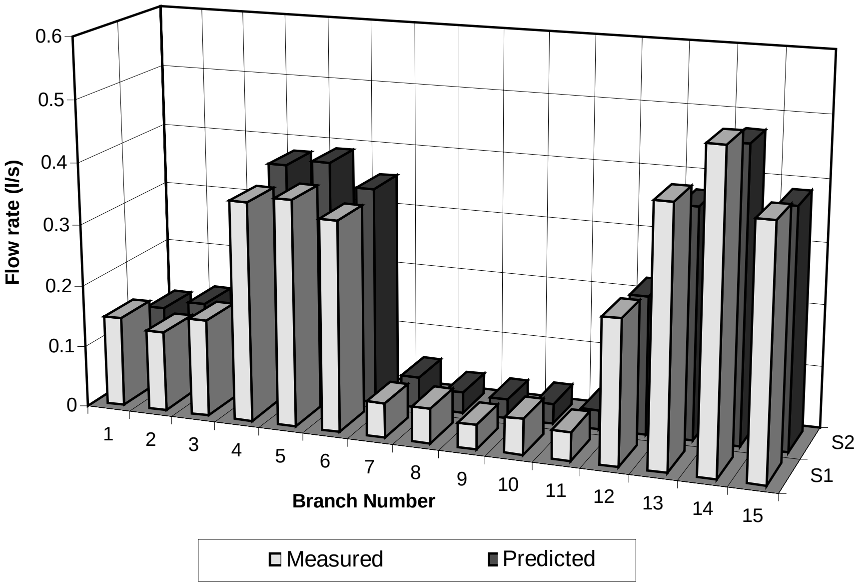

It can be seen from the chart that the flow rates predicted by the computer program were within ±10% of the measured values, apart from in those circuits which had very low flows, i.e. less than 0.03l/s in 15mm diameter pipes. For these the predicted flows could be out by as much as 60%. It was concluded that for these situations, the flow is approaching a transitional phase between turbulent and laminar conditions and the square law relationship between flow rate and pressure loss is breaking down. However, for all turbulent flow conditions the correlation was good.

The computer program facilitated the analysis of a number of situations for which designers had previously relied on rules of thumb. The correct approaches to some of these issues were often hotly disputed between different design engineers.

3. Results

Pump speed control using differential pressure: siting the sensor

In systems where pump speed is to be controlled to maintain a fixed differential pressure between two points in the index circuit, the question arose as to where best to site the differential pressure sensor. American published articles on variable speed pumping recommended that it be located across the furthest branch in the circuit. Common UK approaches included the rule of thumb value of two thirds of the way down the index circuit, whilst others preferred to locate the sensor across the pump.

The computer program quantified the flow distributions suggested by consideration of the pressure gradients for a two pipe direct return circuit. As shown in Figure 2, under part load conditions, flows to sub-branches connected upstream of the sensor would tend to fall, whilst flows to sub-branches connected downstream of the sensor would tend to rise. This is true of all systems for the reasons indicated in Figure 2. Since, to maximise the pump energy saving, over-flow conditions need to be avoided or minimised, the sensor should be located as close to the end of the index circuit as possible. However, placing it too far down the circuit may risk unacceptable under-flows in some upstream circuits.

Figure 2: Typical pressure gradient for a flow/return system under full and part load

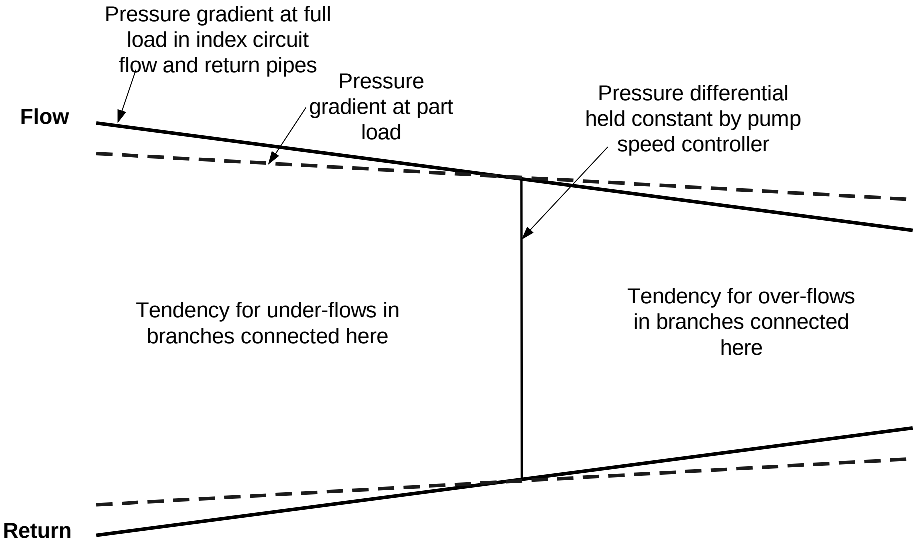

To calculate potential under-flows under part load conditions is normally very difficult and at best involves a laborious manual calculation. The computer program facilitates this exercise by enabling part load conditions to be modelled and the flow consequences of different sensor locations to be quickly determined. The conclusion indicated by the program after testing a number of system layouts, was that none of the aforementioned rules of thumb were universally applicable. However, it is true that to minimise the extent of potential under-flows and over-flows, some mid-point position was usually the best option, with the best energy savings occurring as the sensor moved further towards the index branch. It could therefore be concluded that the two thirds rule of thumb appeared to be a good compromise although, in practice, it might be prudent to install a number of sensors between half way and two thirds down the index so that the alternatives could be tested during commissioning.

Sizing control valves

The sizing of modulating two port control valves was another issue for which design engineers had adopted differing approaches. The terminology used to define valve authority in CIBSE Guides and British Standards was found to have led to different interpretations. The conventional definition was as follows:

$$N = P_1 / (P_1 + P_2)$$

where N = valve authority
P_1 = pressure drop across fully open valve
P_2 = pressure drop across the remainder of the circuit

For constant flow circuits with three port valves, it was straight forward to interpret "the remainder of the circuit" as being the pipework and components between the diverting by-pass connections, since this was the circuit in which flow varied. However, for systems with two port valves, it could be argued that when a two port valve closed down, the flow could vary all the way back to the pump and that therefore, each valve needed to be sized to achieve an appropriate authority (typically 0.5) relative to the pump pressure. Since this approach usually resulted in control valves with unmanageably large resistances, (beyond what was commercially available from control valve manufacturers), the requirement was usually scaled down by accepting authorities as low as 0.25 calculated against pump pressure. Other engineers based their calculation of valve authority on the pressure differential held constant at the sensor whilst others based it on the design pressure across the sub-branch from the index circuit. Some based their calculation of authority on the pressure differential across the local branch containing the control valve itself.

The computer program indicated that this last approach was generally acceptable, but that due to the likelihood of pressure variations across branches, some increase in authority above the normally acceptable value of 0.5 could improve the response of the control valve under some operating conditions. The reason for this was that over-flow conditions in circuits with increased pressure differentials, had to be returned to their design values by the control valves before they were able to begin exercising control over the thermal output of the terminal device. This situation is illustrated in Figure 3. By increasing the control valve's authority, the potential over-flow could be reduced. In effect, by increasing terminal branch resistances the flow increases due to pressure variations are smaller and therefore of less consequence.

Figure 3: The change in control valve characteristic due to increased branch pressures

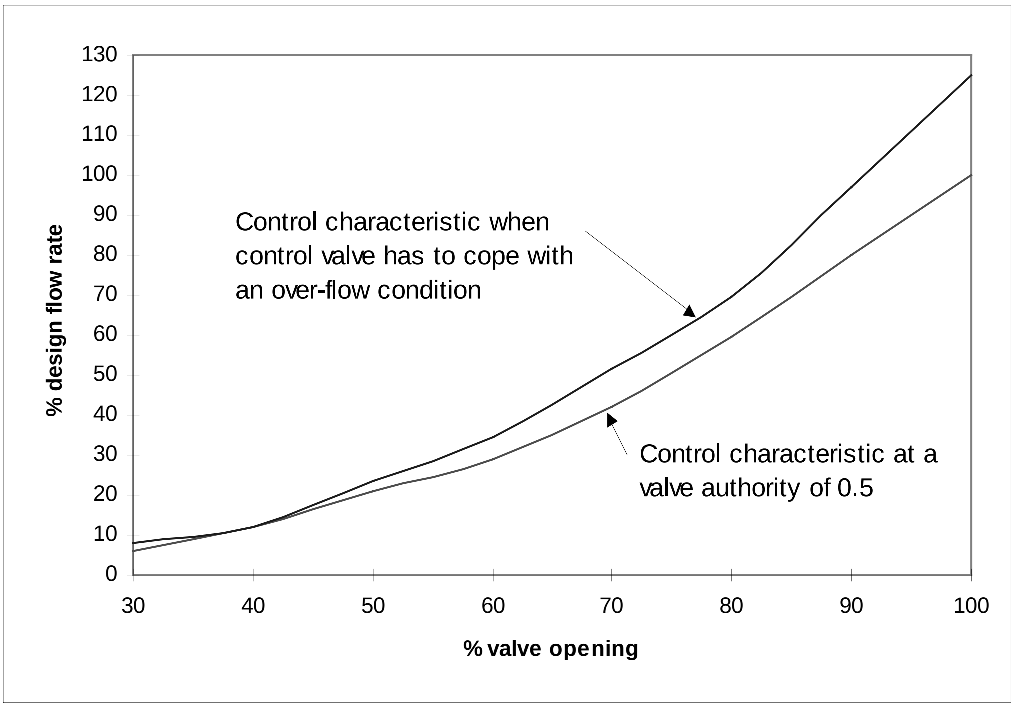

It was concluded that the various rules of thumb used to improve control performance in this way may have been based (knowingly or unknowingly) on a well founded principle, but that they were imprecise in achieving the desired intent and could lead to unnecessarily high resistances across control valves and consequently high system pressure losses. A single clear definition for valve authority was needed, based on branch pressure loss, together with advice explaining the advantages of increasing the authority under certain conditions. The best way to predict these conditions was to create a computer model of the flow distribution between circuits.

Regulating valves and proportional balancing

Another hotly disputed issue relating to the design of systems with variable speed pumping was in relation to the requirement for regulating valves. Articles in American Journals had argued equally strongly both for and against their inclusion in systems.

It was argued that their main function was to ensure that under full load conditions, all terminal units would receive their full design flow rates. However, our computer program proved that as soon as two port control valves begin to throttle, any initial flow balance is destroyed. The consequence is likely to be a larger percentage flow increase to terminals downstream of the throttled one, than to those which are upstream. Since unbalanced systems tend to have a flow pattern which is the reverse of this, valve closures can actually help to bring initially unbalanced flows closer into balance. This is particularly true for the index circuit where a central pressure differential is held constant, as shown in Figure 4. Therefore, unbalanced systems should perform well during operation under part load conditions.

Figure 4: Pressure variations in an unbalanced system

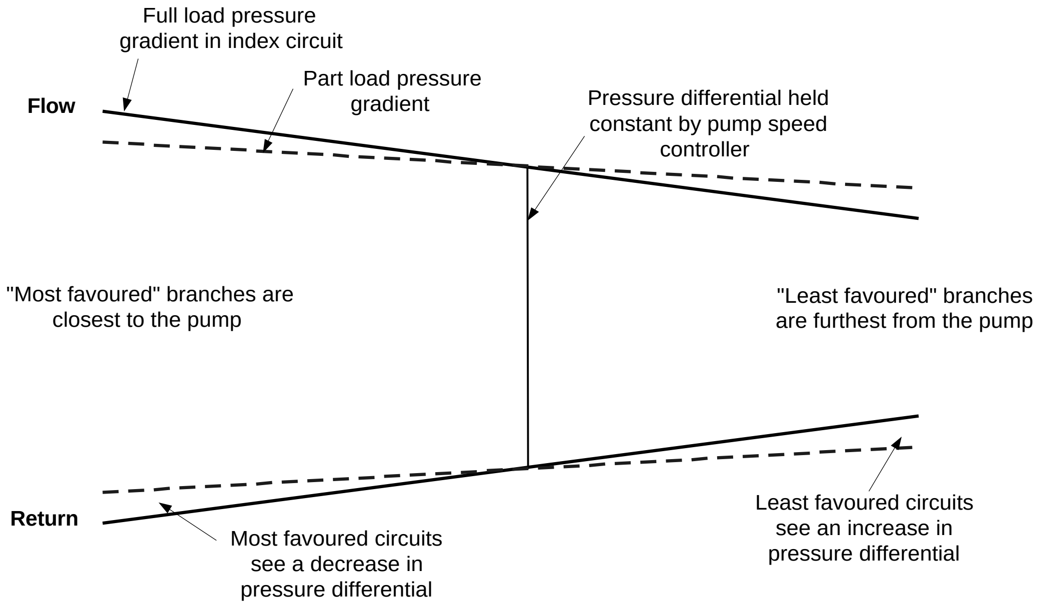

In most buildings, the only time the entire system is likely to call for full load is during pre-occupancy periods when internal temperatures need to be returned to their design values. For some remote parts of the building this could take a long time if flows only reach those areas once other areas are satisfied. The consequence is likely to be a system which runs for longer periods than one where the flows are initially balanced.

However, in modern, well insulated buildings with low air infiltration losses, it can be shown that overall temperature variations are quite small, and for these situations the additional running time does not equate to a significant energy penalty. The concern some building services engineers express is their lack of faith that the builder will achieve low enough air infiltration rates to make the solution viable.

Even when the omission of regulating valves is rejected for this reason, there is still scope to consider this option if it can be shown that potential under-flows are too small to cause a significant reduction in heat transfer. The issue of heat transfer and its sensitivity to flow rate was discussed in BSRIA Application Guide AG 20/95 Commissioning of Pipework Systems - design considerations. Based on the approach explained in this guide, a set of tolerances can be calculated which reflects what is normally permissible in terms of flow rate variations.

Table 1 illustrates a set of suggested acceptable tolerance limits for under-flow conditions. These value are based on the premise that the heating or cooling coils in question have a capacity which is not more than 10% greater than the actual heating or cooling load. Coils which have been over-sized by more than this amount can be treated with greater flexibility.

Table 1: Acceptable under-flow limits

APPLICATION		Acceptable under-flow limit
Heating:	$\Delta T = 11^{o}C$ $\Delta T = 20^{o}C$ $\Delta T = 30^{o}C$	-30% -20% -10%
Cooling:	Sensible only Dehumidification	-10% -0%

By predicting flow distributions between branches using the computer software, it can be determined whether potential under-flows in sub-branches will be within these limits. This will enable an informed decision to be made regarding whether or not to leave out regulating valves.

Conclusion

It can be seen that the use of computerised network analysis techniques are extremely helpful in resolving the various design issues relating to systems incorporating variable speed pumps. With the help of the program, a step-by-step design procedure has been developed which is given in the aforementioned BSRIA Application Guide. If the benefits of variable speed pumping are to be fully realised, there needs to be a greater take-up of the technology which will only come about when the uncertainties of design are removed. The type of approach outlined in this paper provides a solution which will at least improve existing design methods. Whilst the software developed for this project is available from BSRIA, it is hoped that in future, existing providers of pipe sizing software will be able to provide the same analysis options in their programs.