Avoiding overdesign

This month we look at how we can use statistics to determine suitable design parameters for efficient systems

If we want building systems to be efficient in their use of materials and operational energy we must avoid over-designing them. There is constant pressure on the building services engineer to over-design, which must be countered at every turn. This comes in part from a culture of risk avoidance in the industry and in part from a client or developer being unclear on how the building is to be used with the engineer responding by designing for all eventualities. However, operational flexibility beyond a provision that can be reasonably expected always leads to over-design and inefficiency.

As intelligent engineers we are able to exercise judgment based on experience or precedent. However, when presented with a new circumstance, which all projects do to some extent, we need to be extremely careful to avoid trying to eliminate all risk. It is not possible to completely avoid the risk that, under extreme circumstances, the systems will not be able to perform to the desired level. What we need is a way to establish an acceptable level of risk given the consequences of failure of the system to perform as required.

The consequence of occasionally being unable to deliver the full design flow of hot water is very different to the consequence of a major international airport being unable to function because of a predictable weather occurrence, and so we must evaluate the risk and service levels accordingly. In order to solve these questions we need a basic understanding of statistics, but the reward can be considerable savings in terms of initial investment and running costs with a small, acceptable level of risk of non-performance.

These types of statistical analyses are applied widely in the building services field, although we often don't realise it. The 'demand units' method for sizing piped services is based on a statistical analysis of the simultaneous occurrence of a number of independent events. If we understand the fundamental approach we can estimate diversity factors for a wide range of design problems.

If, for example, we were asked to design an office for a firm of estate agents, should we design the air conditioning (assuming that we have not achieved our design by passive means) for 100% occupancy or some lesser number based on the knowledge that estate agents are frequently out of the office.

If we discuss the working patterns with the client we might discover that on average each viewing lasts for half an hour and that visits are entirely randomly distributed between the agents and throughout the day. So we can be fairly certain that the office will not be fully occupied, but the derivation of occupancy is not >

Doug King

Even in an entirely conventional office it is unlikely that there will be continuous 100% occupancy due to absences, breaks and meetings. How, therefore, should we determine the optimum design level for building services without over-designing?

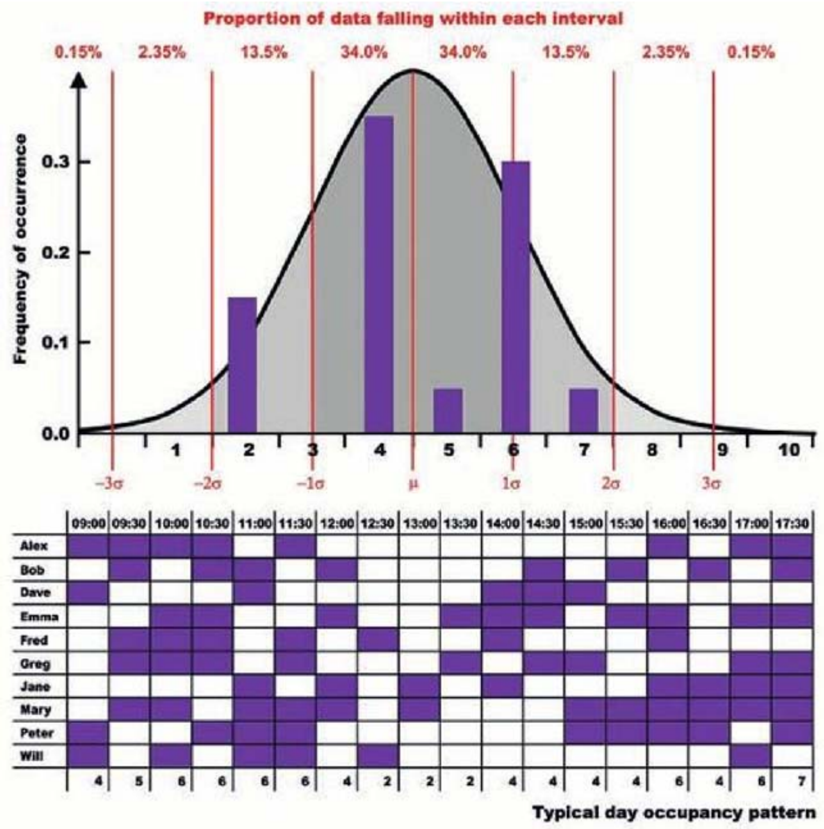

Figure 1: Generating a random pattern of occupancies for a notional office of 10 estate agents reveals that 100% occupancy is an unlikely occurrence. In order to determine suitable design occupancy we need to delve further into the statistics. The Gaussian (or Normal) Distribution curve represents the frequency of occurrence of events randomly distributed about a mean (_). The standard deviation (_) is a measure of the spread of the data. Knowing the rule of thumb, that 95% of the data fall within 2_ of the mean and that 99.7% fall within 3_, allows us to quickly establish suitable design parameters and understand the significance of events occurring outside those parameters.

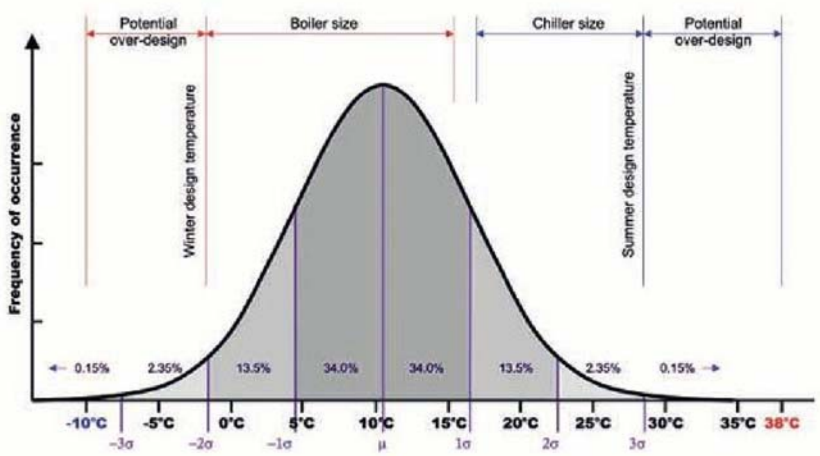

Figure 2: When we apply statistical analysis to weather data, such as the widely available monthly averages and extremes, like this for London, we can begin to appreciate the significance or otherwise of occasional extreme events. To design for these events would result in severe over-sizing of plant and equipment compared to that designed for a more rational balance of cost and risk.

> as simple as designing for either 100% occupancy or 50% occupancy.

Simply generating a random pattern of occupancy (see Figure 1) shows us that it is unlikely that the office will be fully occupied at any time. However, this simple approach does not give us a suitable basis for determining a design level of occupancy, as it is still easy to argue that random chance can give rise to maximum occupancy. So it can, but with what probability?

If we plot the occupancy data as a histogram, we can fit a Gaussian Distribution curve to it by calculating the mean and standard deviation. It is then helpful to know that, to a close approximation, events more than three standard deviations from the mean occur with a probability of less than 0.15%.

So we can now examine the data again and determine that the occupancy is unlikely to exceed eight people for more than about 0.5% of the time, or eight working hours per year. Conversely, this means that designing the air conditioning for the very few occasions on which the occupancy will exceed eight people will increase the plant size and cost by 25% (assuming the plant size is directly proportional to the occupancy).

Understanding these sorts of diversity factors is critical in determining the appropriate conditions to which we should be designing. Consider, for example, that this simple diversity applies not only to the fresh air volume and therefore external air gain, but also to the occupancy gain and the casual gains and electrical consumption from computers and task lights. If we don't get our diversity factors right then the consequences can quickly snowball.

These days, of course, we would probably design the plant for 100% occupancy and then justify the decision by adding variable volume control with some form of occupancy sensor. This will reduce some of the operating inefficiency but we are still stuck with a system which is essentially over-sized for the use to which it is typically put.

Now let's apply a similar statistical method to gain an understanding of external design temperatures: The Met Office weather record for London shows extremes of –10C and 38C with a mean of 10.5C and a standard deviation of about 6C. If we assume a Gaussian distribution for the temperatures (which is a good approximation), and fit these parameters we can immediately identify some of the key points that have historically been used to define London design temperatures (Figure 2).

Temperatures more than three standard deviations from the mean are likely to be exceeded less than 0.15% of the time, or just 13 hours in an average year. In the summer the majority of these extremes will occur during the working day and so 28C became adopted as the temperature for cooling design. In the winter the majority of the cold extremes will occur at night and so a lower deviation is accepted for commercial buildings leading to the –2C rule of thumb for heating design.

If we were to select our cooling plant based on the extremes of temperature that we sometimes experience, we will be significantly over-sizing the plant based on a very small risk of occurrence. Designing for 38C would result in a cooling system nearly twice the capacity of one sized for the design temperature. This situation of over-design is often exacerbated by the value engineering that inevitably results from the selection of over-sized plant in the first place. During value engineering the third chiller that the engineer selected to provide better part-load efficiency will be eliminated in favour of two machines and a cost saving. This now means that each machine will effectively be sized to 100% of the normal design condition and therefore the operation will be at part-load for the vast majority of the time, with the other machine largely redundant.

Although the *CIBSE Guides* now recommend a more sophisticated analysis to determine suitable heating and cooling design conditions, it is still vital that we understand the significance of the statistical probabilities revealed by this first approximation. As climate change is beginning to impact on our weather it becomes increasingly difficult to resist adopting design temperatures closer to the extreme events. Even with these increasing extremes, the probability of such events is still small.

> **"Operational flexibility beyond a provision that can be reasonably expected leads to over-design"**

CIBSE, the Met Office and a consortium of universities are presently working on the UKCIP probabilistic future weather data to develop sample weather years that can be used in analysis software and to determine new design temperatures. However, until then, we can still determine useful design information from the historic weather records, provided they are up to the present date, but we should not allow our judgment to be driven by the occasional bad weather experience.

The purpose of defining diversity factors and design temperatures is to avoid expensive over-design of the systems while balancing this with an acceptable risk of exceedance. It is essential to apply these kinds of analyses at the briefing stage in order to discuss risk and cost with the client. This allows informed decisions about system design criteria, which will result in a more robust and cost effective design.

This approach can also often avoid the round of value engineering that occurs when the cost of over-designed systems exceeds the client's expectations, as you can demonstrate that the design is the most cost effective for the actual anticipated need with an acceptable level of non-conformity. ●

Doug King is principal of King Shaw Associates and visiting professor of Building Physics at Bath University