Heating system behaviour

The system

The heating system works by injecting a large volume of air into the upper part of the nave, which slightly pressurises the building. The hot air is spread to all part of the church by two effects: The speed at which it is injected causes it to swirl around, causing a large pool of hot air. The slight pressurisation of the building causes air to leak through joints all round the outside (mainly through the stained glass, but also round doors) which draws the hot air out towards the edges of the building. The swirling pool of hot upper air mixes with the air below, where people are sitting and standing.

The building structure acts as a heat reservoir. When it is cold it cools the air and when it is warm it heats it. This thermal mass makes any temperature change in the building much slower than it would otherwise be.

Measurements

We can gain some insights into the behaviour of the heating system (and the building of which it is part) from the temperature records gathered so far[1]. The graphs below show the temperature at three places:

- By the pulpit (orange) – This is the closest to the pew level temperature, but because it is next to the stonework it will read lower and/or respond more slowly than a free-air reading while the heating is running in cold weather. Conversely, when the heating is not running, it may read higher than a free-air reading – for example overnight in Figure 2 it is warmer than the top of the screen.
- At the top of the screen (red) – This is within the pool of hot air distributed around the church, and thus gives an indication of whether the heating is running. A sharp rise shows the hearing coming on, and a sharp fall shows it going off. We don't yet know enough about the behaviour to interpret other features of the trace. For example the extent to which it varies depending on how hard the heating is running (ie fraction of the time during which hot air is being pumped in).
- On the tower roof (blue) – This represents the external air temperature, although being in a sheltered position near the tower structure it will not show the same extremes as a free air measurement would.

Figure 1 shows Sunday 6 February 2011, which was mild, with overnight external temperature of 11C and residual internal temperature of 13°C. The heating came on some time about 6 am (shown by the rise in high-level temperature) and the low-level temperature rose steadily until the thermostat cut in between 8 and 8.30am. The reduced high level temperature reflects the lower heat input needed to maintain the more or less steady pew level temperature for the rest of the morning.

The heating will have switched off completely at some point, and the upper and lower temperatures became very similar, but the exact point is hard to determine because the upper level temperature was already relatively low, reflecting the low level of input needed to maintain the temperature. The building was still quite warm when the heating came on again around 5pm, the low-level temperature was back to its peak in about an hour, and the thermostat again reduced the heating input. When the heating went off, the cooling was again very gradual, and by midnight the church was still warmer than at the start of the day. The temperature only reached its starting point shortly before the heating came on some time after 8am the following day (off the graph) .

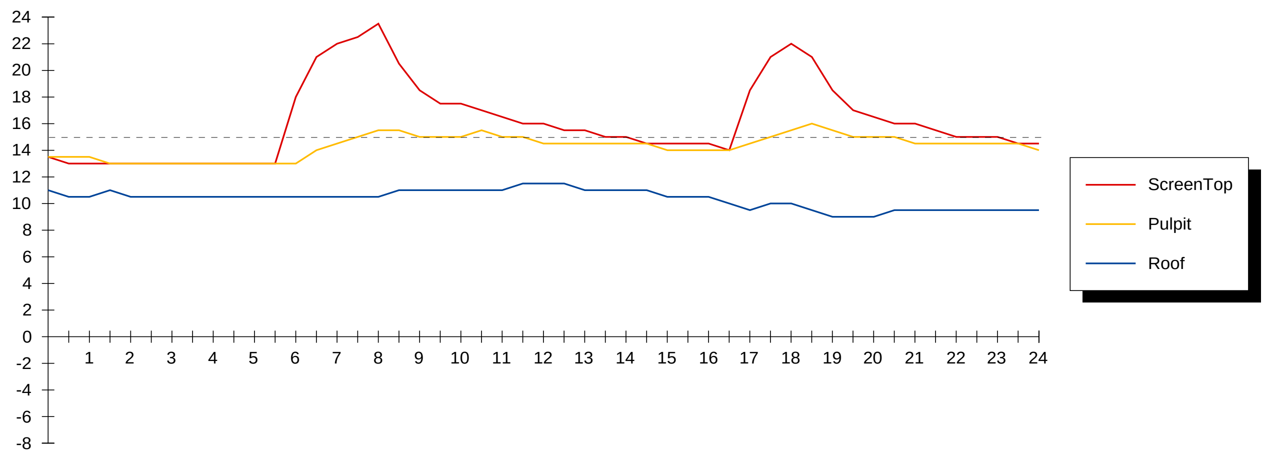

Figure 1: Sunday 6 February 2011

1 Measurements were recorded every half hour. The graphs show all points on the hour and half hour, but the actual time varies slightly between them, since it took several minutes to set up all the sensors. The offsets are: pulpit +1m, screen top +9m, roof +11m).

Figure 2 shows Sunday 19th December, which was a much colder day, with the outside temperature well below freezing all day. Overnight the low-level internal temperature fell to 9°C. (Note that the high level temperature was even cooler, which suggests that the air temperature in the building was lower than the residual temperature of the stonework by the pulpit sensor.) When the heating came on, the low level temperature began to rise at about the same rate as in Figure 1, but starting from the lower base it didn't reached the set point of the thermostat before the heating turned off around 11.30am The high level temperature confirms this.

The building cooled down during the afternoon, but not much more than half way to its starting temperature. When the heating came on again, the low-level temperature rose as before, but did not have time even to reach the morning's high before the heating went off. By midnight the temperature was nearly down to its starting level.

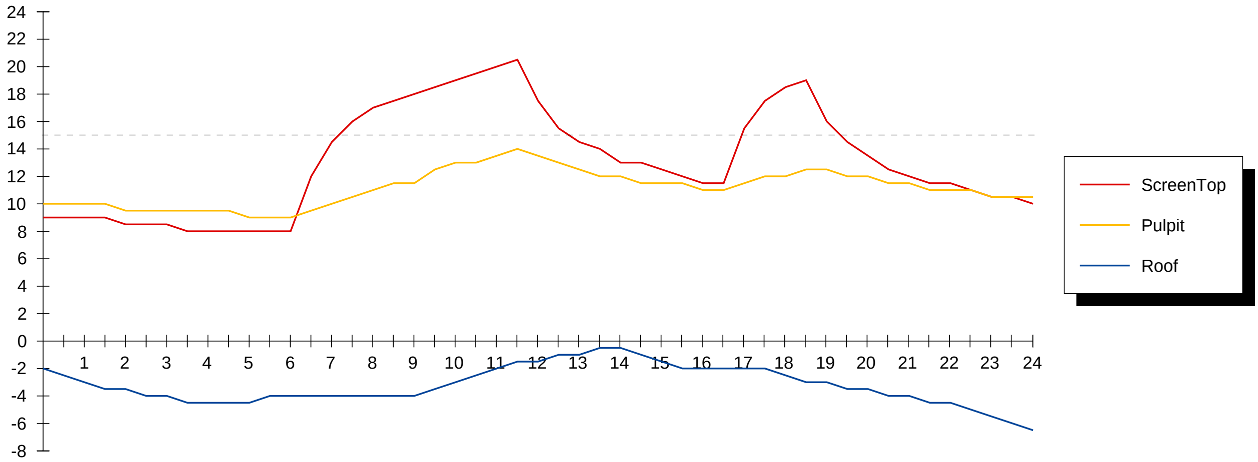

Figure 2: Sunday 19 December 2010

Figure 3 shows the response to a single, fairly short, spell of heating. Wednesday 8 December 2010 was also cold. The outside temperature wasn't quite as cold as in Figure 2, but the building was cooler. The heating seems to have run continuously to cover both Morning Prayer and Wednesday Mums, from shortly after 7.30am until just after 10.30am. From a low of 7°C, the low-level temperature rose steadily to 10°C when the heating went off, and it then cooled steadily for the rest of the day, reaching its starting temperature shortly before midnight.

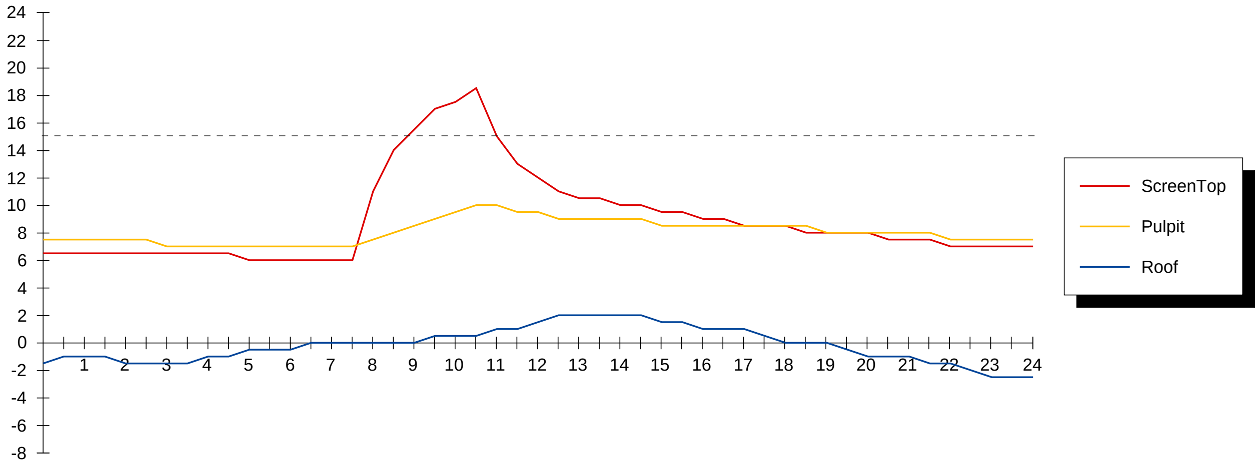

Figure 3: Wednesday 8 December 2010

Figure 4 shows the effect of a sustained heating over Christmas Eve and Christmas Day 2010. During a 26 hour period, the heating was on for 60% of the time. It came on around 7.30 on Christmas Eve and ran for about 11 of the next 17 hours before going off some time after midnight. It came on again about six hours later and ran until around 11am. The low level temperature curve shows two effects, with small undulations of a few degrees that correspond to the heating coming on and off superimposed on a steady overall rise from the building's overnight low of 8°C to the thermostat set point of 15°C, first reached in the afternoon, and then re-gained during each subsequent heating period. The initial peak of the high-level temperature curve is similar to those shown in Figures 1-3, but the peaks later in the day have distinct flat tops. These occur when the heating is still running, but when the low level temperature has reached it working level, so heat input is regulated by the thermostat. The apparent shift between the two curves is discussed below.

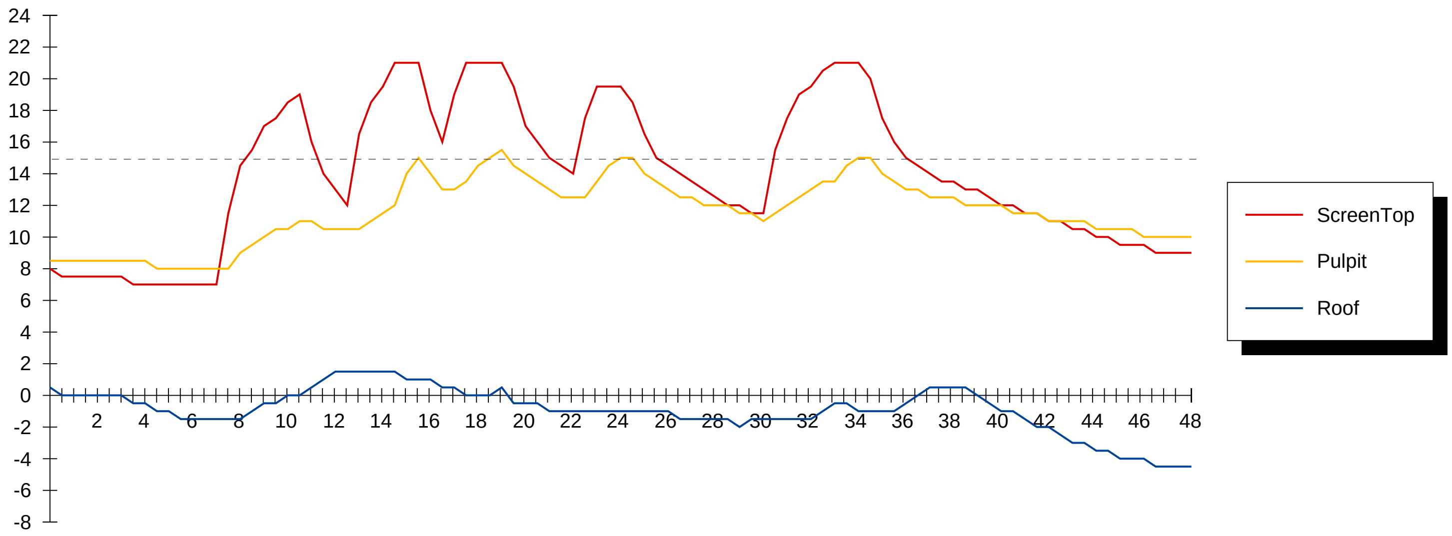

Figure 4: 48 hours over Christmas 2010

Figures 5 and 6 show the same information as Figure 4, but with each day presented at the same scale as all the other figures. Compare Figure 6 with Figure 1. Even after cooling all night at Christmas, and even with the outside temperature below zero, the internal temperature was higher than it was on the much warmer February day shown in Figure 1.

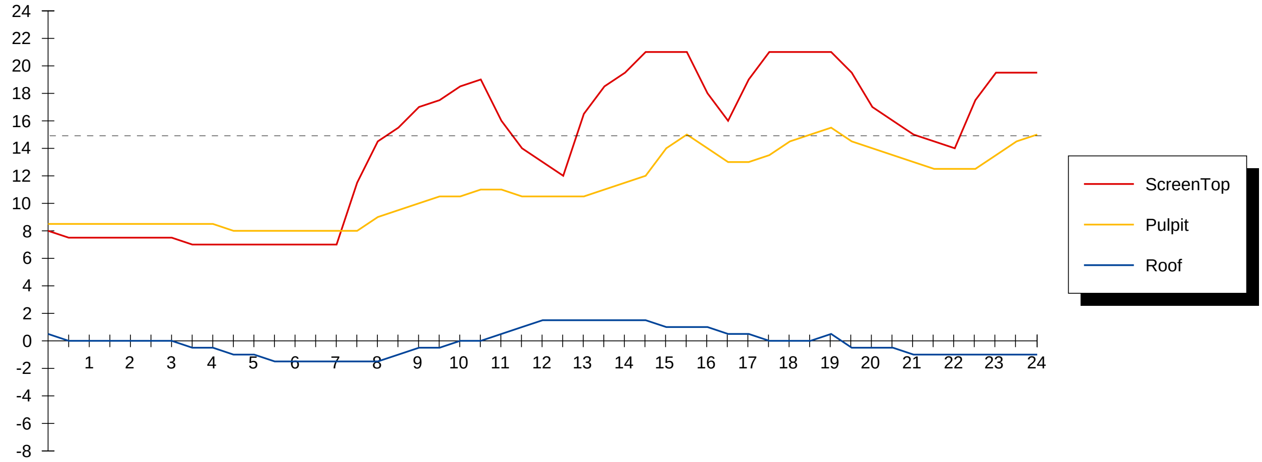

Figure 5: Christmas Eve – Friday 24 December 2010

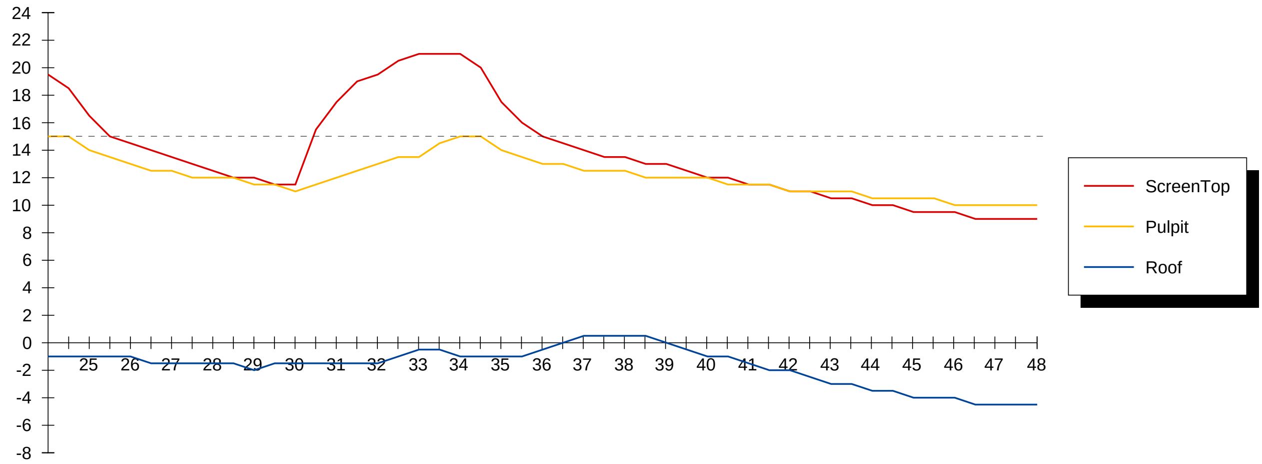

Figure 6: Christmas Day – Saturday 25 December 2010

Interpreting the detail

The graphs show the dynamic response of the building to changes of heat input. When the heating comes on, the low-level temperature rises at around 1°C per hour until it reaches the set point of the thermostat. The rate of cooling when the heating switches off seems to be more variable (between 0.2°C and 0.6°C per hour), which could be due to several factors. The time from switch-on to reach working temperature is proportional to the difference between the desired and starting temperatures. To reach 15°C starting at 12°C takes about 3 hours, but to do it from 8°C takes around 7 hours. This may seem slow compared with a domestic house, but the building is very different. Traditional church heating systems took much longer to heat up. In cold weather, the boiler wasn't fired up a couple of hours before the building was needed, but the day before .

Looking at the high-level temperature, three features are apparent. When the heating comes on, the temperature rise quickly (4°C or more in an hour). This period my be dominated by getting the pool of hot air in the church up to 'working temperature'.

After an hour or so, the temperature continues to rise, but much more slowly. This is very noticeable in the longer morning heating periods, especially in Figure 2, where the high level temperature rises at almost the same rate as the low-level temperature, maintaining a difference of 7-8°C between them. This period may represent a balance between heat being fed in and heat being transferred to the solid structure (walls, pillars and probably furniture) which only warms very slowly because of its greater thermal capacity.

The flat tops that appear on the high-level temperature curve during the extended heating over Christmas come later in the day. They almost certainly represents the system being regulated by the thermostat rather than running at maximum power. There is a slight anomaly in that the flat tops slightly precede the low level temperature reaching its set point. This is probably because the sensor near the pulpit, which is leaning on the stonework, reaches the set point after the rest of the air.

Some calculations

The descriptions above are qualitative, but to get a feel for the numbers involved, here are some calculations based on what we know. The volume of the church[2] is ~5500m^3. The heater has a rated output[3] of 226kW. The blower is rated at 15,000m^3/h (~4.2m^3/s)[4]. Most of the other figures used are standard physical constants.

Air supply temperature – The specific heat of air is ~1kJ/kg/°C = 1.3kJ/m^3/°C[5]. So running at full power and full flow, the air will be heated ~42°C above the inlet temperature. An internal control will limit the temperature to a safe working level, but I assume[6] this will not cut in during normal operation. On 19 December 2010, the external air was at -4°C, so the air in the duct will have been around 38°C and on 6 February 2011, the external air was 11°C, so the air in the duct would have been around 53°C. In both cases, this assumes that the heater and blower are running at maximum, which it might not have been if there is a limiting mechanism.

These temperatures may sound high, but the temperature away from the grills into the nave will rapidly reduce because the injected air mixes with the ambient air. The combined area of the two grills is ~3m^2 so the average exit velocity will be around 1.5m/s (~5fps) which will entrain a lot of the surrounding air.

High-level air temperature – The hot pool of upper air is formed by mixing hot air from the heater with cooler air that is already in the church. The inlet temperature is 42°C above the outside temperature (see above), and the air in the church outside the hot air pool will be at the temperature of the structure, assumed to be similar to the low level temperature. This can be checked by reading sample points from the curve, calculating the temperature for a given air mix, and comparing the result with the measured high level temperatures. Figure 7 is based on 9 sample points at the end of heating periods in Figures 1-6. The actual readings (yellow squares) are plotted against outside temperature. The black crosses represent a mix of ~20% of air at 42°C above outside temperature and ~80% air at low-level temperature. Two points don't match well, and the small errors on the others are in opposite senses for high and low external temperature, so there must be other effects as well as mixing.

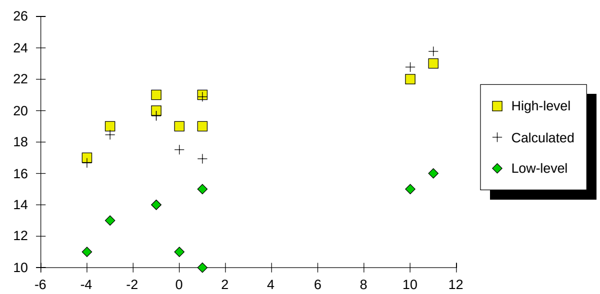

Figure 7: High-level temperature calculated with 15% hot air mix v actual

2 Previous calculation, not recently repeated.
3 The rated input (ie the gas used) is 293kW, so the burner + heat exchanger has an efficiency of 77%.
4 The blower may deliver slightly less than its rated when working into the back pressure of the church with doors closed, but I suspect the effect is small, since the the back pressure is quite low. I once tried to measure it using an altimeter with 10m resolution (equivalent to ~1mbar). Sometimes did, and sometimes I didn't,get a different reading between the pressure inside and outside.
5 The density of air is ~1.3kg/m^3
6 The manufacturer's current products quote 95°C.

Heat input to the building – The usable heat injected into the building depends on the temperature difference between the injected air and the building (structure and initial air temperature). Thus when the outside air is cooler than the building, the usable heat input is less than the total heat supplied by the heater. For example, on 6 February 2011, with the outside air ~2°C below building temperature, the useful input power would be ~215kW, but on 19 December 2010, with the outside air ~13°C below building temperature, the useful input power would be ~156kW, about 30% less.

Heating the air in the church – When the heating first starts, the church is full of cool air, and the system starts to pump in hot air. An air flow of 15,000m^3/h could fill the ~5500m^3 volume of the church with hot air in about 20 minutes if all the cool air obligingly left before any of the heated air. In practice the air will not form an orderly queue, the hot and cool air will mix, and it will take longer than that before most of the air has come from the heater, rather than being there in the first place. Of course the air never reaches the exit temperature of the heater, because there is continual loss of heat to the stonework. So an hour does not seem an unreasonable period for most of the air to have been heated, and for the system to enter a phase where the behaviour depends on a balance of heat between the source and various losses. This would be consistent with the change in the rate of rise of the high level temperature that is observed.

Heat loss to the walls – A very crude estimate[7] of the area of stonework is around 1500m^2. The transfer of heat from air to a surface by free convection is 5-25W/m^2°C. For want of anything better, assume a mid value of 12. The overnight air temperature should be in equilibrium with the walls, so the walls start at that temperature. The high-level air typically rises quickly to between 6 & 8°C above this. Assume the lower value, ie a temperature difference between air and wall of 6°C. On that basis, heat transfer to the walls would be between a little over 100kW, which is between half and two thirds (depending on outside temperature) of the heat being injected into the church. Given the uncertainties in the calculation, this is close enough to be in the same ball park.

Heat lost in air leakage – Air leaks out of the church at the same rate that it is pumped in (~4.2m^3/s). If the air leaving is at an average temperature of say 6°C above the temperature of the building, then the rate of effective heat power leaving is ~25kW, ie 10-15% of the heat power entering the building.

Warming the building – To find out whether the relatively linear rise of both high and low level air temperature after the initial hour corresponds to the building itself warming up, we need to solve the heat equation:

$$\frac{\partial u}{\partial t} = \frac{k}{c\rho} \frac{\partial^2 u}{\partial u^2}$$

The solution would need to make various assumptions and simplifications. Finding an analytic solutions is non-trivial, and initial attempts at a numerical solution haven't yet yielded anything useful. If any results emerge, I will add them later.

Some provocative thoughts and tentative conclusions

The above interpretations appear to account for the measured behaviour. So assuming that they are correct (enough) here are some speculative thoughts about possible changes to the system and/or the way it is operated.

- **A 'comfortable' temperature?** – Is the desired temperature (at pew level) fixed, or does it vary with external temperature? In cold weather, many people (notably those who have walked) arrive at church wearing outdoor clothes. When they come into church they expect the temperature to be warmer than outside, but if they sit in their outdoor clothes, then the temperature that they find comfortable will vary depending on what they are wearing, which in turn will vary with the outside temperature. This is a crude generalisation, with several confounding factors. In wet weather people may want to remove their outer layer. People who drive to church may not dress for the outside conditions.
- **More takes longer** – It is unreasonable to expect a building with a long thermal time constant to heat up quickly if it is allowed to cool down between use to well below the desired working temperature. With the heating running, pew level temperature rises at around a degree per hour. If a given temperature is required at a given time, then the heating should come on at an earlier time, corresponding to how much cooler the initial temperature is. In extreme weather, this will mean long lead times.
- **De-stratification?** – Is it possible to increase the mixing between the hot pool of upper air and the lower air, without unacceptable hotspots. Last time I looked it did not seem promising, but might be worth revisiting. The air inlet grills should be optimally set, but there might be scope for fine-tuning them. Separate de-stratification units are available commercially, but they might prove unacceptable on the grounds of visual aesthetics or noise.

[7] Very crude, with many approximations. It may be out by a factor of 2, but it should be adequate for order of magnitude calculations.