Regeneration in GHEtool (part 1)

Author: Wouter Peere – Date: 22/09/2026

Imbalance is one of the most common challenges when it comes to designing geothermal borefields. Although different solutions exist (such as changing your design or opting for a hybrid geothermal system), one common solution is to rely on regeneration to actively counteract the imbalance. In this article, the newly released regeneration module in GHEtool Cloud is discussed.

What is regeneration?

When designing shallow geothermal borefields, you often have to deal with a long-term temperature drift due to the imbalance of your system. Either your borefield will cool down year after year due to having a higher extraction load than injection load, or vice versa. In the temperature profile below, a borefield with a significant extraction imbalance is shown.

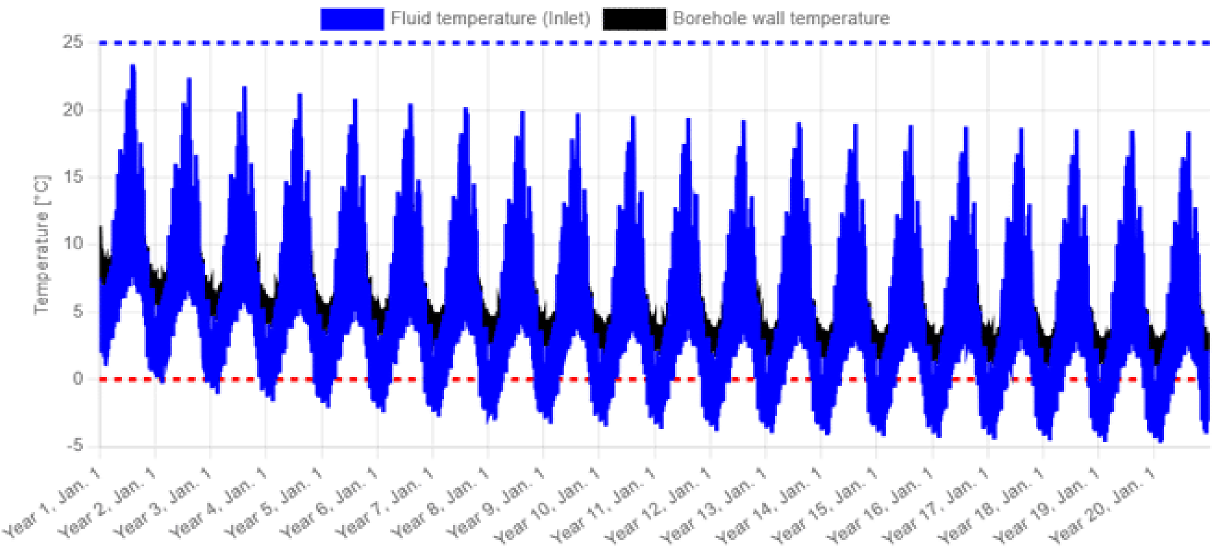

Temperature profile of an extraction-dominated borefield.

Such an imbalance is unwanted since it can lead to a larger required borefield size to ensure that the fluid temperatures stay within the set limits (as discussed in our article on imbalance here).

Conceptually, there are two different ways (besides the architectural or occupational changes to the building) to avoid this imbalance: hybridisation and regeneration. To clearly understand the difference between these two options, let us revisit the Sankey diagram (more information here) of the temperature profile above.

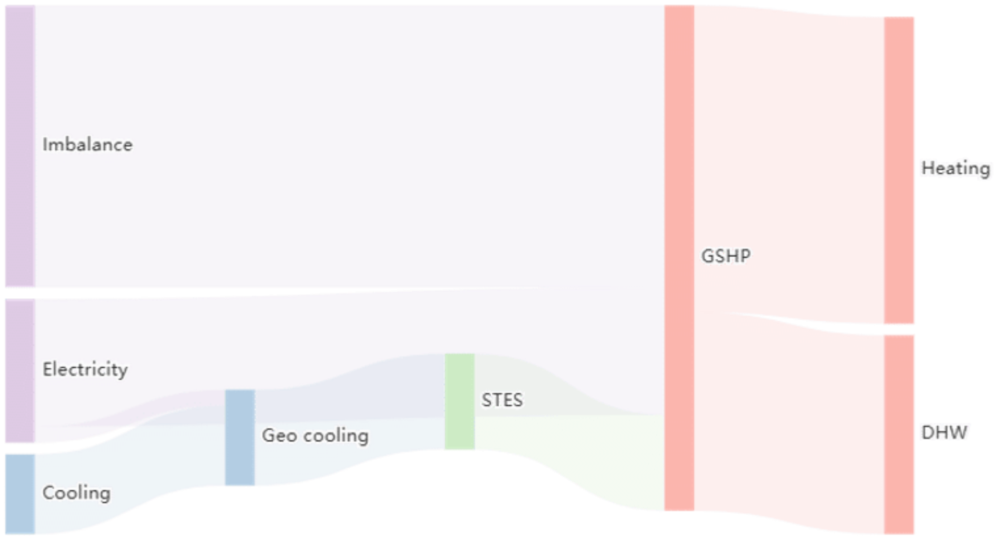

Sankey diagram of an extraction-dominated borefield.

A Sankey diagram is a graphical way to illustrate the energy flow of our geothermal system. In the example above, the heating and domestic hot water demand are on the right, both provided by the ground-source heat pump (GSHP). This GSHP, in turn, also has some different inputs, namely the electricity demand (for the compressor in the heat pump), the imbalance and the seasonal thermal energy storage (STES).

Both the imbalance and the STES account for energy extracted from the ground, but they are conceptually different. Imbalance is using the ground as a source of heat, whereas the STES is using the stored heat from the cooling during the last summer. Since shallow geothermal borefields are essentially a technology to store heat rather than deplete the ground, this green part ideally is as big as possible. In the Sankey diagram above, it is clear that it is mostly the imbalance that feeds the GSHP, explaining the significant downward temperature trend of the borefield.

The first option to overcome this problem of having a high imbalance is by going for a hybrid technology, as shown in the Sankey diagram below.

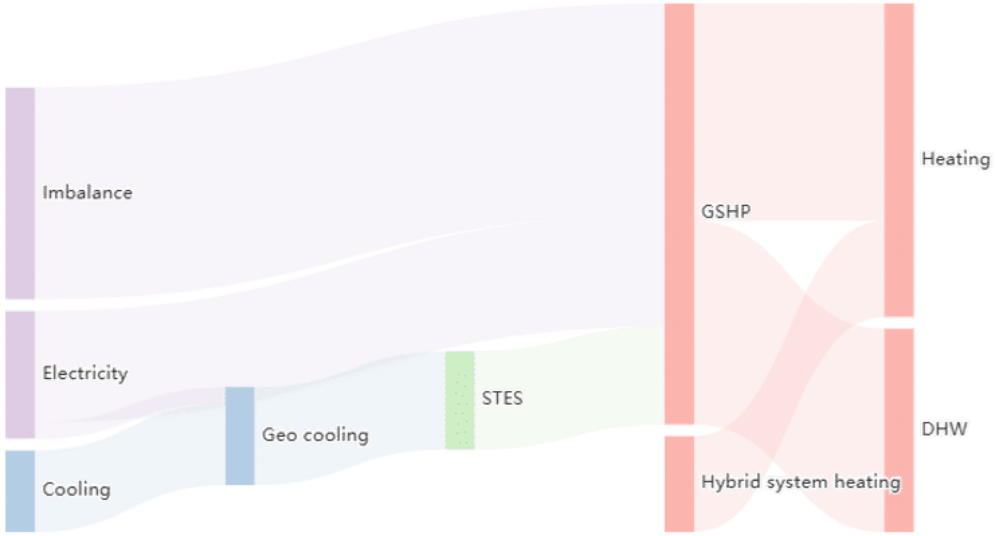

Sankey diagram of the borefield with a hybrid heating system.

Here, part of the heating demand is delivered by a hybrid system (more information here), for example an air-source heat pump. By doing so, the share of the GSHP reduces, and since the cooling and geo cooling demand stays the same, the imbalance reduces significantly. This results in the temperature profile below where now the fluid stays within the predefined limits, but the imbalance is still clearly visible. This also implies that, if the simulation period is extended to, for example, 40 years, the share of the hybrid system should be increased to keep staying within the limits.

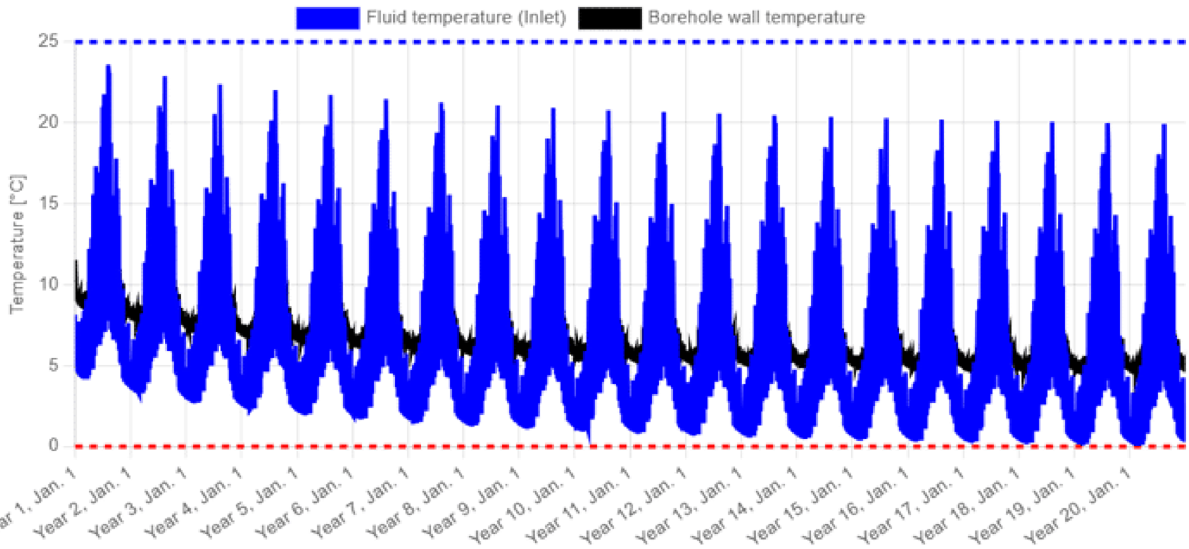

Temperature profile for an extraction-dominated borefield with a hybrid system.

Another option is to go for regeneration of the borefield, which is shown in the Sankey diagram below. Here, the heating (and DHW) load is still 100% provided by the GSHP, but the imbalance is now completely gone. With the help of regeneration, extra heat is injected into the ground, storing it, and using it later by the GSHP.

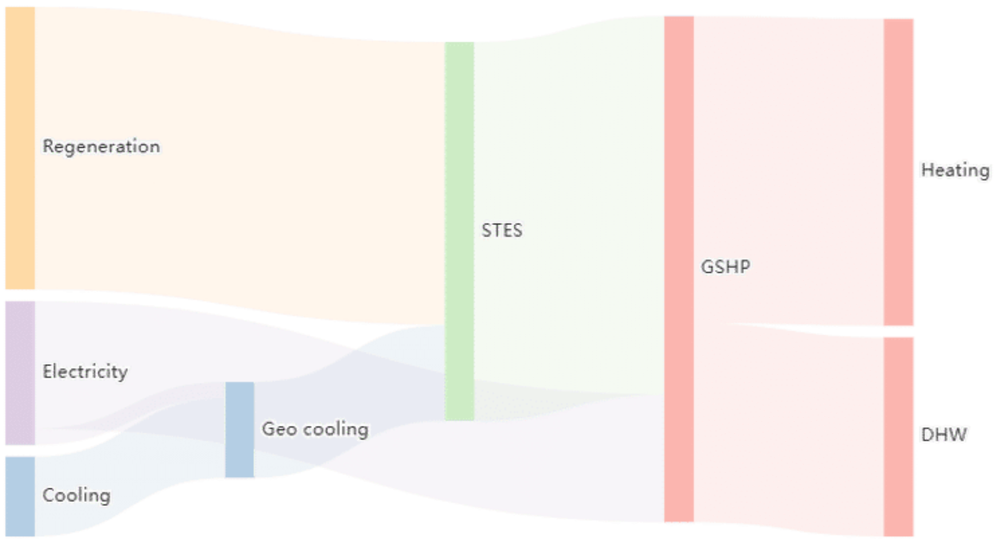

Sankey diagram with regeneration.

The situation above results in the temperature profile below. Here, it is clear that there is no longer any long-term temperature drift, implying that, if the simulation period were extended to 40 years, this would not influence the feasibility of the system.

It is also clear that the profile now looks completely different, since regeneration changes the load profile of the ground, by injecting (or extracting) extra heat into (or from) the ground.

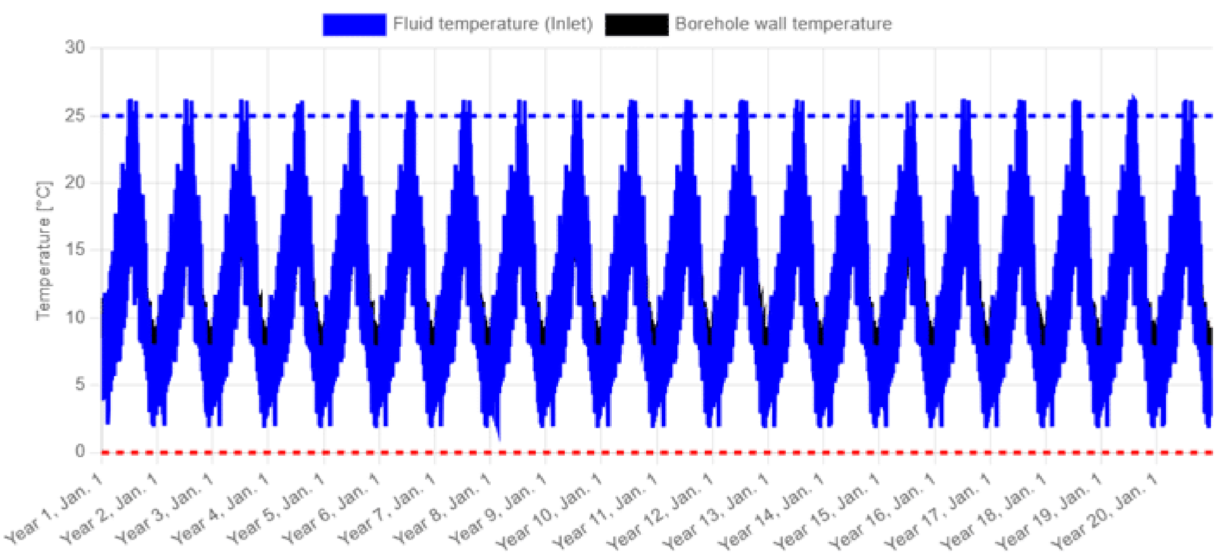

Temperature profile after regeneration.

Below, the different categories of regeneration technologies will be discussed, after which the details of the regeneration algorithm (the rules for regeneration, where to place it and which algorithm to pick) will be explained.

Regeneration technologies

Regeneration can be done with almost any technology as long as it is able to deliver either heating, cooling or both. That, in a sense, is the beauty of regeneration. If you have an extraction-dominated borefield, anything that can be used as a heat source is suitable for regeneration, as well as any cold source in the case of an injection-dominated borefield.

In GHEtool Cloud, four categories of regeneration technologies are implemented in which, to some degree of abstraction, all regeneration technologies can fit. These are discussed below.

Constant power

The first and most straightforward technology category for regeneration is that of a constant power. This implies that we always have a certain amount of power available for either injection or extraction. Think, for example, of an electric heating element that can be turned on at any time to deliver a certain amount of power, or even an air-source heat pump that has a certain installed power that it can deliver in heating and/or cooling.

The idea with constant power technology is that its power is constant and that it has no time-dependency.

Constant temperature

The counterpart of a regeneration technology with a constant power is one that has a constant reference (or regeneration) temperature which we can exchange energy with via a heat exchanger. This is for example a district heating or cooling network, sewage water or exhaust air ventilation.

All these sources have (to a high degree) a constant and time-independent temperature that can be used for regeneration. For example, the exhaust air ventilation temperature is typically around ±22°C year-round. When a heat exchanger is installed here and the fluid temperature is colder than this temperature, heat can be extracted from the air and injected into the ground. When the fluid temperature is warmer, the heat can be blown away, cooling down the ground.

The power $\dot{Q}$ in [W] that such a heat exchanger can deliver is proportional to the temperature difference ΔT between the temperature of the regeneration source and the average fluid temperature in the heat exchanger. When the temperature difference increases, the power the heat exchanger delivers increases as well. In GHEtool, this is assumed to be a linear relationship with a scaling factor k in [W/K]:

$$\dot{Q} = k \cdot \Delta T$$

Solar thermal

One very common regeneration technology is that of solar thermal systems, like solar thermal collectors, solar absorbers (also called swimming pool heaters, like the one shown below) or PVT panels. These are all technologies that convert solar radiance to heat and are typically used in situations where the borefield is extraction-dominated.

Example of solar absorbers (unglased collectors). (Source: AQSOL)

Besides converting (part of) the solar energy to heat, these technologies also exchange energy with the surrounding air. Typically, when solar thermal collectors are used for heat production, this is a heat loss to the environment, but in the case of regeneration, where the fluid temperatures are rather low, it can also be seen as a gain.

The power that solar thermal regeneration technologies can deliver can be written as:

$$\dot{Q} = a_0 \cdot \dot{Q}_{sol} \cdot A + a_1 \cdot A \cdot \Delta T + a_2 \cdot A \cdot (\Delta T)^2$$

Where $\dot{Q}_{sol}$ is the solar power in [W/m^2], A the surface area in [m^2], ΔT the temperature difference between the surrounding air and the average fluid temperature in the panel and a_0, a_1 and a_2 the zeroth, first and second order heat loss coefficients.

The formula above also shows that solar thermal panels can be used as well for the extraction of heat, especially when the fluid is warmer than the surrounding air. This is, however, typically not as efficient as a dry cooler since the least amount of sunlight heats the fluid back up.

> ***!Note***
> *When considering solar thermal panels, their orientation is also important, not only for the absolute value of the heat they can produce but also for when it does so. In GHEtool, currently, for simplification reasons, this is abstracted away by considering 80% of the global horizontal irradiance on the panel. This should be more or less equivalent to a panel in an average orientation.*

Dry coolers

Dry coolers (or in fact any sort of heat exchangers with the outside air) are the last category of regeneration technology. They also depend on the outside environment, just like solar thermal technologies, but they do not harvest any of the solar irradiance. Just like the regeneration with a constant temperature, they act like a heat exchanger with the (now changing) surrounding air and a possible regeneration power that is proportional to the temperature difference.

Dry coolers are very suited for both regeneration of an extraction-dominated borefield as well as an injection-dominated one.

Three rules for regeneration

Although both hybridisation and regeneration solve the same problem of imbalance, they do so in very different ways. Hybridisation takes away part of the load from the borefield, making the temperature threshold less critical. Regeneration does the opposite: it increases the geothermal load such that the imbalance is compensated for, but in doing so, it changes the temperature profile significantly. This can have a significant impact in the future.

Imagine the case above where there was a significant extraction imbalance and we want to regenerate this by injecting extra heat into the ground. Ideally, this can be done with some sort of solar thermal regeneration in summer (when it is warm and there is quite a lot of sunlight), but this would increase the borehole wall temperature as well as the fluid temperatures, reducing the potential for (passive) cooling.

In order to overcome this issue, three rules for regeneration are used in the back-end.

1. **Do not over-regenerate**

   Regeneration is meant to counteract the imbalance, but when you're not careful, it can create another imbalance as well. Imagine the case where regeneration is used to compensate for an extraction-dominated load by injecting extra heat into the ground. If you keep doing that, eventually, the borefield will become injection-dominated. Therefore, regeneration should be limited to, at maximum, the current imbalance.

2. **Do not cross the temperature threshold at the current hour**

   Regeneration changes the temperatures at the current hour and you should ensure that they stay within the set temperature limits. This automatically limits the power you can inject or extract with regeneration.

3. **Do not cross the temperature threshold at a future hour**

   Besides changing the temperature at the hour in which the regeneration takes place, the temperatures do also change in the hours following the regeneration. Imagine for example that there has been some regeneration by heat injection an hour before there is a very high cooling load onto the borefield. It can be that, by rule 2, the fluid temperature was not crossed during the previous hour, but due to the regeneration, the borefield is slightly warmer when handling the cooling peak, crossing its limits.

To overcome this, every time we regenerate, we need to check in theory all future hours to make sure that what we do now will not negatively impact the future. However, having a simulation horizon that equals the simulation period makes the algorithm quite slow without any reason. In practise, having a regeneration period of at maximum 4380 hours (i.e. half a year) is more than enough, and 100 hours will typically also be enough.

!Note
In the following part, some simulations will be done to illustrate the importance and the effect of the simulation horizon on both the speed and the accuracy of the method.

By following the three rules set above, the regeneration can happen without the temperature thresholds being crossed within the simulation horizon.

!Note
It can happen that even with these rules, there are some temperature violations. Imagine, for example, the case where the borefield is extraction-dominated and the maximum peak power is already above the temperature limit. In that case, regeneration can only make this worse, since the temperature limit was not even met before the regeneration started.

Where to place the regeneration?

The next important question is: where do I position the regeneration technology? There are two options implemented in GHEtool Cloud: right before the borefield inlet or right after the borefield outlet. Both of these will be discussed briefly below.

Right before the borefield inlet

When the regeneration technology is placed right before the borefield inlet, it receives the coldest temperatures during heating and the warmest ones during cooling, since it is positioned right after the GSHP. (See also the article on inlet and outlet temperatures.)

This is, for example, interesting in the case when there is an extraction-dominated borefield with still some high cooling peaks. Therefore, regeneration should happen in winter or in the mid-seasons, to stay as far away as possible from the moment when the high cooling peak happens. In that case, we want a fluid temperature that is as cold as possible so it is easier to inject heat into the ground, which is why we place the regeneration right before the borefield inlet.

Right after the borefield outlet

When the regeneration technology is placed right after the borefield outlet, it receives the warmest temperatures during heating and the coldest ones during cooling.

This could be interesting, for example, when we again have an extraction-dominated borefield, but now, there are very few cooling peaks in summer (or they are active, so the maximum fluid temperature does not matter that much). In that case, ideally, the regeneration happens in summer, when the surrounding air is the warmest and the sun shines the brightest. Since in summer, when there is cooling, the outlet position gives back the lowest temperatures, this would be the ideal solution.

This shows that, depending on your regeneration technology as well as your building load, the ideal position can be different from case to case.

Regeneration algorithms

Regeneration has the purpose of counteracting the imbalance, but the question now becomes: what imbalance exactly? Depending on how you look at the problem, four different algorithms can be defined, each of which corresponds with a control strategy in practise. These will be discussed below.

Yearly

The first, yearly, regeneration algorithm has the purpose of compensating the imbalance of every year individually. Imagine the case where there is a 20 MWh/year imbalance. In that case, this method will regenerate at max 20 MWh/year (or lower, if there is not enough power available to regenerate everything) but not more.

When there is enough regeneration capacity available to compensate the entire imbalance, this method will yield a temperature profile that is identical every year.

Simulation period

Another way of looking at this problem is by considering the total imbalance over the entire simulation period. Imagine a simulation period of 20 years is considered, with 20 MWh/year of imbalance. This results in 400 MWh that could be regenerated over the entire simulation period. This method will just start regenerating as much as possible as soon as possible until the entire imbalance is compensated.

This is generally the fastest regeneration method, since it can happen that the borefield is completely balanced before the end of the simulation period. Imagine investing extra in solar thermal panels to compensate for the imbalance. It can be that these panels inject more heat into the ground than the imbalance in a single year. This implies that, for example, after 15 years the entire imbalance of 20 years is already compensated and no regeneration is needed in these last 5 years. This is also shown in the figure below.

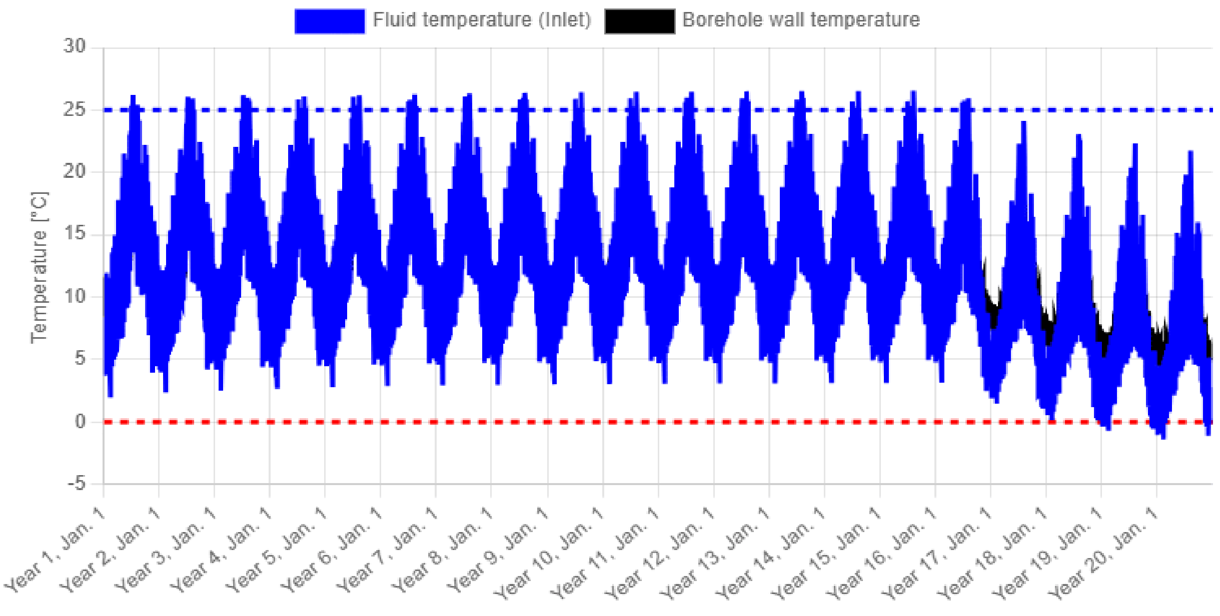

Temperature profile with regeneration and the simulation period algorithm.

> ***!Note***
> *In this method it can happen that (on a short time scale) the imbalance switches from extraction to injection dominated in certain years, but since the imbalance is now considered over the entire simulation period, this is still in line with the first regeneration rule above.*

Proceeding horizon

The two methods above assume some pre-existing knowledge about the imbalance, whereas this is often not known in practise. In the 'proceeding horizon' method, the algorithm can only regenerate the imbalance that has already happened. This means that, for example, after winter, the sum of all the energy exchange with the ground is typically dominated by extraction, so the algorithm can compensate for that. However, if somewhere during summer, the imbalance is 0, it will stop with regenerating, even if there is more imbalance coming in the months to come.

> ***!Note***
> *In this method, it can be that it is impossible to remove all the imbalance from a certain borefield, since the algorithm cannot anticipate future imbalance like the other two methods can. For example, when part of the extraction imbalance happens in November and December, the algorithm will not have the next summer to compensate for this, leaving behind a small imbalance.*

Proceeding horizon +

This last method basically extends the idea behind the previous algorithm, but now, regeneration can happen in both directions. Just as before, the imbalance is calculated at every hour based on the sum of all the previous heat injections and extractions. It can happen that, for some profiles, this net resulting imbalance before summer is extraction-dominated and right after summer injection-dominated. In this case, this method will first regenerate with injection and afterwards with extraction, whereas the proceeding horizon method before will only regenerate in the direction of the overall imbalance (either extraction or injection).

Conclusion

In this article, the new regeneration method in GHEtool is introduced. First, the difference between hybridisation and regeneration for the avoidance of imbalance was discussed, after which the four different categories of regeneration technologies were introduced. Afterwards, the three different rules for regeneration were explained in the context that regenerating now will also change the temperatures in the future. Here, the concept of the simulation horizon was also introduced as a compromise between simulation time and accuracy.

Next, the difference between positioning regeneration at the borefield inlet or after the borefield outlet was discussed, illustrating that depending on your particular project, one or the other could be beneficial. Lastly, four different regeneration algorithms implemented in GHEtool were introduced, each corresponding to a specific possible control strategy in practise.

In the next part, which will be released on the 17th of November, the theory of this chapter will be used on a couple of examples to illustrate the flexibility of this new method in GHEtool. Stay tuned!

References

- Watch our video explanation over on our YouTube page by clicking here.