PRECISION COFFEE

There is a clear trend in consumer appliances, such as coffee machines, towards improved quality and increased personalisation. Consequently, these appliances require higher control accuracy, while costs need to remain low. This article focuses on water temperature control in coffee machines. Varying types of heaters each have their own physical limitations that influence both start-up effects and steady-state behaviour. Accuracy is therefore largely bound by heater choice, but control can still be used to get the most out of the machine. Guidelines are presented for selecting optimal product architectures and control strategies.

NEAL MEIJERS, MIEKE VAN DEN BELT AND HENK VAN DER WULP

Introduction

A key trend in our daily food and beverage consumption is an increasing demand for convenience and personalisation. This, in combination with consumers' growing awareness and knowledge of quality food, creates a market for affordable home appliances that allow for easy personalisation. A good example of this trend is coffee preparation, where appliances nowadays offer a large variety of consistent- and high-quality coffee recipes in convenient ways.

Achieving this level of quality and personalisation requires accurate control of many parameters, especially of temperature and water flow. Optimal coffee extraction requires water between 91 and 96 °C. Lower temperatures will lead to underextraction, causing weak and tasteless coffee, whereas higher temperatures close to the boiling point will cause bitterness [1]. Furthermore, consumers may want to adjust their extraction profile to be slightly sweeter or bitterer by changing the temperature within the 91-96 °C range. Making different coffee flavour profiles that can be easily discerned and reproducibly produced within this range will require a very high degree of accuracy, in the order of +/–1 °C.

As well as ensuring that coffee is extracted at exactly the right temperature, hot-water control also plays a large role in creating steam, which is used to make the milk foam for cappuccinos, etc. However, steaming brings its own challenges and therefore requires other design strategies. This article only considers hot-water control.

Controlling water temperatures and flows might not seem complicated at first glance; however, given strict boundaries of cost price, new legislation and sustainability (material use and energy consumption), it becomes a significant challenge.

PCV Group has extensive experience in the development and realisation of home and professional appliances. From this experience arose the need to develop guidelines to support the selection of the architecture of high-precision hot-water control systems depending on product requirements and the aforementioned boundaries. This article presents and explains the fundamentals behind these guidelines.

Boilers and flow-through heaters

When considering hot-water control in consumer appliances, many types of heaters can be distinguished. Thermal mass, flow velocity and power are only a few of the many properties that influence the required control strategy. The optimal control strategy is not only affected by architectural variations, but is also largely influenced by a wide variety of specifications, requirements and boundaries. Controlling heaters is therefore a trade-off between parameters, e.g. control rate versus reaction rate, or start-up effects versus stability.

The two major sources of temperature error are start-up effects and steady-state errors. Start-up effects refers to the first phase of a dispense, not to preheating the device. For example, a standard espresso is only 30 ml, and thus has a short dispense time. Start-up effects in this case are more dominant than for a typical lungo, which has a volume of

Two types of heaters. (Source: www.espresso-expert.com/boiler-or-thermoblock-which-heating-system-to-choose)
(a) Flow-through heater.
(b) Boiler.

AUTHORS' NOTE

Neal Meijers (senior engineer), Mieke van den Belt (engineer) and Henk van der Wulp (senior consultant) all work at PCV Group, a product design and development company located in Enschede (NL). PCV stands for People Creating Value.
This article is based on a presentation that was accepted for the DSPE Conference on Precision Mechatronics 2020, which has been postponed till 2021.

mieke.vandenbelt@pcvgroup.com
www.pcvgroup.com

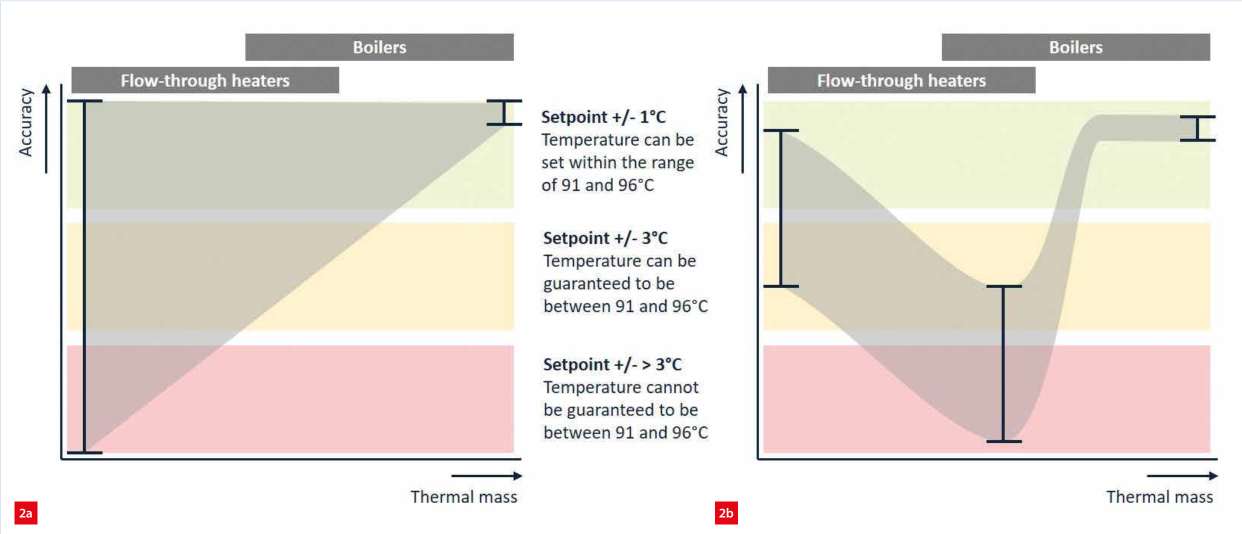

Temperature accuracy that can be achieved. (Source: PCV research)
(a) During steady-state behaviour.
(b) Start-up behaviour.

50-150 ml. For the latter, steady-state effects are much more significant. However, both start-up effects and steady-state behaviour play a large role in every coffee machine and therefore both need to be taken into account.

Heaters can be roughly divided into two groups: boilers, in which water is heated and stored until used; and flow-through heaters, in which water is heated while it flows through the heater for direct use. Examples are shown in Figure 1.

Figure 2 shows the achievable accuracies for boilers and flow-through heaters when used to dispense a typical lungo at typical flow velocities. The grey areas show the bandwidth that can be achieved with different control techniques. These are physical limits, dependent on system architecture and design choices, meaning that even perfect control cannot improve accuracy further.

As can be seen, for obtaining high temperature accuracy in both start-up and steady state behaviour, a large boiler is the optimal choice, as it can heat up water in advance and keep it at a very constant temperature until needed. Very basic control is therefore sufficient, which keeps costs low. However, boilers take up a lot of space and the preheat time is long. The high thermal mass also leads to a high energy usage, especially if only a single cup of coffee is made before cooling down. This makes the larger boilers mostly suitable for professional or catering industry use, rather than for household consumer appliances. The relation between preheat time and thermal mass is valid for all heaters, as thermal mass always leads to a longer preheat time. This is most apparent in the energy consumption of large boilers.

In the lower thermal mass range, i.e. for flow-through heaters and for boilers that are too small to dispense one entire serving, a contradiction occurs. Overall, higher thermal mass typically shows better steady-state results, in so far that high accuracies are easier and cheaper to achieve. However, lower thermal masses, which are a lot harder to accurately control, show much better start-up behaviour.

Limiting physical principles

When establishing guidelines, the first step is thus to select the type of heater that is the best fit for the application. This decision is a trade-off between accuracy, available space, material and development costs. Here, it should be understood that physical principles limit the extent to which accuracy can be improved by control.

The influence of thermal mass on steady-state accuracy is relatively straightforward: the larger the thermal mass, the less effort it takes in control to obtain high accuracy. However, regardless of the quality of the controls, errors in the output temperature of water will occur in the transition from a static (pre-heated) state to an operating state during dispense. Such start-up effects, common in many controlled systems, quickly become highly relevant in a coffee machine, where a single dispense is short. These start-up effects are heater-architecture-dependent and will be discussed separately for boilers and flow-through heaters.

Flow-through heaters

In a flow-through heater, water passes through a long slim heater tube, heating up in the process, much like in a heat exchanger. The driving force of heat transfer in such a system is convection between tube wall and water. This parameter is highly dependent on flow velocity: the convection is at least one order of magnitude higher when water flows compared to when it stands still. This means that there are two equilibria possible: a static one, where the system is pre-heated, but before dispense occurs; and a dynamic one, where water flows through the heater during dispense. In a static situation, the convection to the water is relatively low and consequently the temperature of the heater is homogeneous. Depending on how long the heater has been turned on without flow, the water in the heater will either be the same temperature as the wall of the heater or colder.

When flow starts, convection to the water becomes a more dominant effect, compared to internal conduction of the heater block. This causes a temperature gradient to occur in the heater between heating element and tube surface. Additionally, most heaters are designed such that in order for enough heat to be transferred to the water in a short time, the surface temperature of the heater tube has to be higher than the desired water temperature. This allows for smaller heaters, up to a point.

What this means in practice is that as soon as flow starts, the heater will start to cool down because of convection to the water, and will therefore have to heat up again. At a minimum, the core of the heater will have to heat up to compensate for the heat distribution inside the heater. For smaller heaters, this means having to heat the tube surface to higher levels. This cannot be prepared for in the static situation, as overheating the heater without flow will cause unwanted boiling behaviour, generating steam and thus pressure.

This is where thermal mass, contact area, power and flow rate come into play. A higher thermal mass decreases the speed with which the heater can change its temperature, thereby significantly increasing the time in which the start-up effect takes place. Having surplus power to heat up more quickly, or more contact area will help minimise this effect. This is what causes the limitations as shown in Figure 2. The exception to this effect is having a heater with a high thermal mass and a large surface area, but with a relatively low water flow. In such cases, the start-up effect can be minimised, as the heater barely has to heat up further when flow starts.

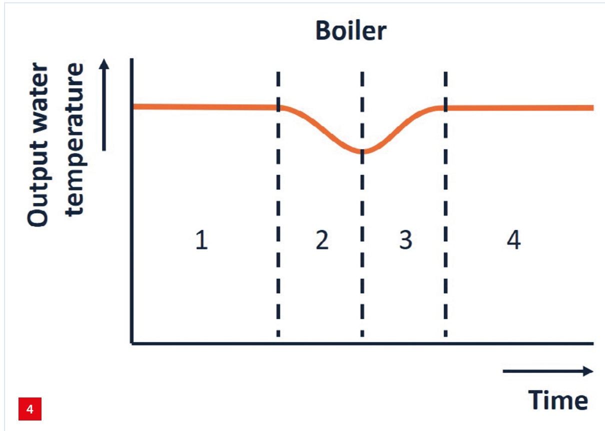

Boiler behaviour over time.
1 Only hot water reaches the outlet, while cold water enters at the bottom.
2 First mixed water is reaching the outlet. The heater is on, but the dynamic steady state has not yet been reached.
3 The boiler is reaching the dynamic steady state.
4 The boiler now functions as a flow-through heater with relatively large thermal mass.

Boilers

Boilers function slightly differently than flow-through heaters. The thermal mass of a flow-through heater is concentrated in the metal mass of such a heater, whereas the thermal mass of a boiler is concentrated in the water itself. The specific heat capacity of water is about five times higher than that of aluminium, and water has a thermal conductivity (when still) that is orders of magnitude lower than that of aluminium. When a boiler is turned on, the water inside starts to move due to the mixing of hotter and colder water. This way, heat spreads through the water. This is a relatively slow effect, however, only relevant when there is no flow. If flow is produced with a pump, and kept laminar, cold water can be pumped into the heater from the bottom, forcing hot water up and thus out of the boiler (Figure 3). This way, the temperature of the water that exits the boiler is not directly affected by the cold water entering the system, because of its low thermal conductivity.

Figure 4 shows how output water temperature in a boiler changes over time. As long as the volume of the boiler is sufficiently large, one dispense can take place fully in phase 1, thereby making the output temperature independent of the input temperature. For small boilers, this is no longer

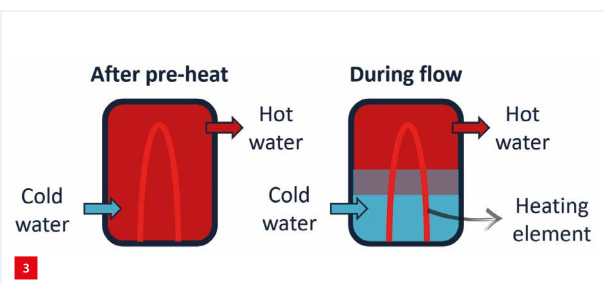

Schematic of boiler behaviour.

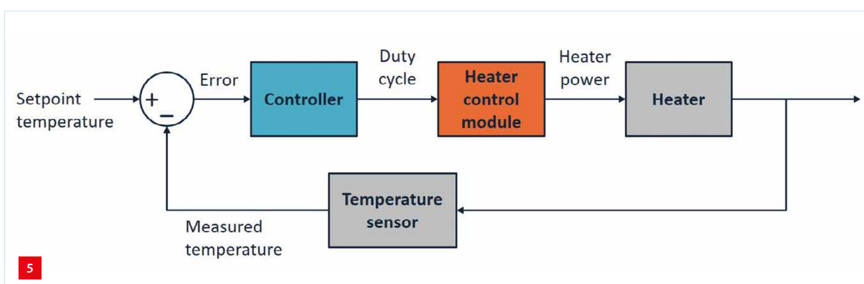

Control diagram for a heater.

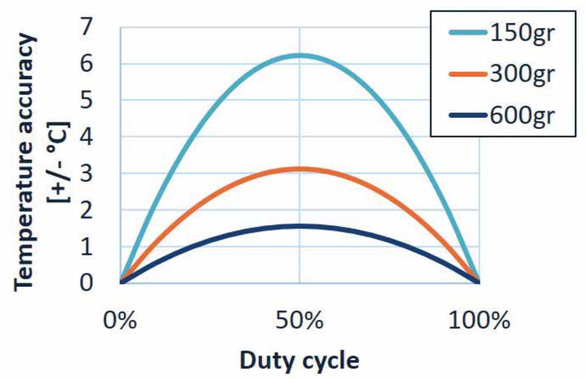

6

Temperature accuracy as a function of duty cycle and thermal mass, as determined by the flicker regulation.

the case, and cooler water will reach the output before it has fully heated up. At this point, the boiler behaves like a flow-through heater with a relatively large thermal mass, causing a significant dip in temperature while transitioning from static to dynamic behaviour.

Control strategies

There are two main contributors to optimising temperature accuracy with controls: the controller and the heater control module. This is shown in Figure 5.

The controller itself, combined with the location of the sensor(s), determines the final accuracy of the output. However, the heater control module can function as a limitation on this controller. Costs for the controller are mainly determined by the development costs: higher accuracy requires a more advanced controller and therefore more development work. The heater control module, on the other hand, mainly affects hardware costs, impacting the appliances' bill of materials. Note that this only considers control based on temperature. Flow-based control is also possible, but variations in flow have an impact on extraction similar to that of changes in temperature.

Under normal circumstances, a heater does not require its maximum power, meaning a control module is needed to control the power output of the heater module.
This can be done in several ways, such as:
- Zero-cross on/off control.
- Phase-cutting control.
- A hybrid combination of on/off and phase-cutting control.
- PFC (power factor correction) control.

These methods are all limited by EMC (electromagnetic compatibility) compliancy codes, namely the flicker and higher harmonics codes [2, 3].

By far the cheapest solution is zero-cross on/off control, which switches the heater on and off at a set interval. Because of the flicker code, the control interval is limited to an average of 12.5 switches per minute for a 1,400-W heater, which is a very standard power level for this application. This causes a temperature fluctuation depending on thermal mass, as can be seen in Figure 6. In the worst case, which is at a 50% duty cycle, a 600-g heater is required to obtain a temperature accuracy of +/– 1.5 °C. In practice, this limit can be achieved with a lower mass for two reasons, namely that heaters generally run far above the 50% power level and that due to less switching actions during start-up more than 12.5 switches per minute are allowed during the steady state of a dispense. The resulting behaviour is shown in Figure 7.

More advanced heater control modules like PFC and phase-cutting solutions are capable of accurately controlling the duty cycle of the heater without breaking flicker regulations. When using one of these solutions, the heater control module no longer acts as a bottleneck in the water temperature performance of the system. These solutions do cause higher harmonics, however, which need to be filtered in order to pass regulation. Such filters are quite expensive, which is a severe limitation for their implementation in many (consumer) devices.

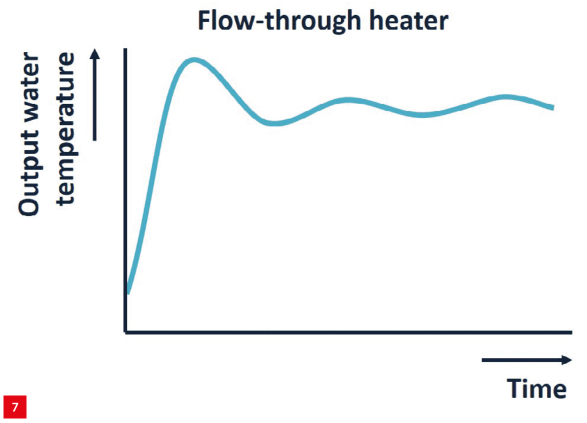

7

Oscillation in output temperature.

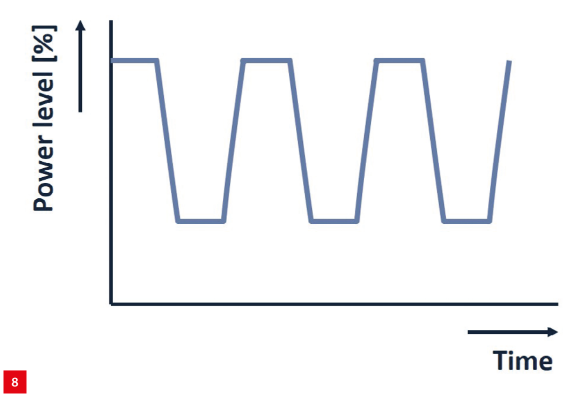

8

Power level ramp up/down using phase-cutting.

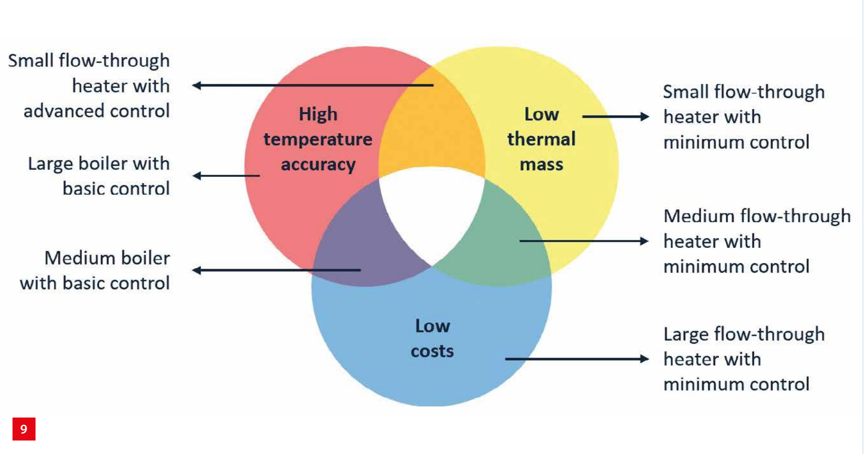

Trade-off between high temperature accuracy, low thermal mass and low costs.

The use of a hybrid combination of on/off control with phase-cutting is also a possibility. This method uses an on/off method but ramps up and down at every on/off cycle with phase cutting, as shown in Figure 8. This allows for a higher control interval while staying within flicker regulation, leading to improved accuracy. The designer can tune this method to have control intervals that match performance requirements, where higher control intervals lead to higher performance, but also to more harmonics, a larger filter and therefore higher costs.

The chosen design of the heater control module has a significant impact on control costs. The chosen water system and the level of effort put into the controller need to be matched with the chosen water control module. Note that a lower thermal mass design requires the most advanced heater control module to achieve high levels of steady-state accuracy.

Causes of variation

Optimising the controller itself is the last step in achieving high accuracy. There are many sources of variation that impact the heating behaviour of the water system. The most important of these are the variation in mains voltage and the input temperature of the water. Use of a closed-loop controller that takes the physics of the system into account is therefore highly recommended. This step is not detailed further in this article.

Conclusion

As presented in this article, selecting a heater and a control strategy is not an easy and straightforward job, because of the many contradictions. A decision will always be a trade-off between three main aspects, namely temperature accuracy, costs and thermal mass-related properties, such as preheat time and size. This trade-off is shown in Figure 9 and summarised below.

Temperature accuracy

High temperature accuracy is achieved by going for either a very high thermal mass in the form of a large boiler, or a low-thermal-mass flow-through heater. The latter requires advanced control to obtain high accuracy, therefore making it an expensive option. Large boilers, on the other hand, have disadvantages in their preheat time, energy usage and build-in volume. However, they are a very efficient solution for systems that need to dispense a lot of coffee, for example in professional and catering industries.

Costs

Costs are mainly determined by two aspects: hardware and development. Hardware costs are hard to generalise, as both small flow-through heaters and very large boilers are usually expensive. Large boilers, however, can do with much simpler control, making development costs significantly lower than for small flow-through heaters. The cheapest option, being a medium-to-large-thermal-mass flow-through heater with minimum control, makes high accuracy close to impossible.

Preheat time, build-in volume, energy usage

These effects are directly correlated with thermal mass. A higher thermal mass leads to higher preheat times, more build-in volume and a higher energy usage. Regarding these aspects, a lower thermal mass will always be better, however the obvious disadvantages need to be kept in mind: a very low thermal mass requires a lot of control to gain sufficient accuracy, and is therefore expensive, whereas medium thermal masses do not allow for high accuracy. Overall, boilers are slightly more space-efficient than flow-through heaters.

REFERENCES

[1] Kingston, L., *How to make coffee*, Ivy Press, 2017.

[2] Koninklijk Nederlands Normalisatie-Instituut, "Electromagnetic compatibility (EMC) - Part 3-2: Limits – Limits for harmonic current emissions (equipment input current =16 A per phase)", 2019.

[3] Koninklijk Nederlands Normalisatie-Instituut, Electromagnetic compatibility (EMC) - Part 3-3: Limits - Limitation of voltage changes, voltage fluctuations and flicker in public low-voltage supply systems, for equipment with rated current <= 16 A per phase and not subject to conditional connection", 2019.