Placement of Shared Electric Kettles Under Thermal-State-Dependent Service Times

Callum J. Ferrie
Department of Electrical and Electronic Engineering
Northgate Polytechnic Institute
Glasgow
United Kingdom
c.ferrie@eee.northgate-poly.ac.uk

Priya Raghunathan
Department of Electrical and Electronic Engineering
Northgate Polytechnic Institute
Glasgow
United Kingdom

Elspeth M. Dunsmore
Estates and Facilities Division
Northgate Polytechnic Institute
Glasgow
United Kingdom

capstract:2611.00036v1 [eess.SY] 19 Nov 2026

***Abstract*—Shared electric kettles in open-plan offices are placed where the plumbing allows, not where they minimise the time occupants spend getting hot water. We formulate kettle placement as a p-median problem with congestion, in which the service time of each kettle depends on its thermal state. Water left from the previous boil cools by Newton's law, and the next user's boil time depends on how warm it still is. We calibrate the model on a 72-occupant floor using smart-plug metering of 5,241 boils over nine weeks, bench cooling tests and 412 observed trips. We then solve the placement problem with a simulation-based swap search. Moving the floor's single kettle from the kitchenette to a central riser reduces mean trip latency from 171 s to 138 s. A second kettle reduces it by a further 8.7 s, almost all of it during the four daily demand peaks. Between peaks, the walking time saved by the second kettle is offset by longer boils, because each kettle is warm less often, and in the quietest half-hours mean latency rises. A third kettle saves 3.4 s and a fourth saves nothing, while energy per drink rises from 40.9 Wh with one kettle to 54.9 Wh with four. When users fill only the water they need, the thermal penalty of splitting demand largely disappears and a second kettle saves about twice as much.**

***Index Terms*—Facility location, queueing, thermal modelling, discrete-event simulation, building services, energy use**

I. Introduction

Most open-plan offices in the United Kingdom have at least one shared electric kettle, and most occupants use it several times a day. Where it goes is usually decided once, at fit-out, by whoever is routing the water supply. Nobody checks afterwards whether that location makes sense for the people who use it.

The question looks like a textbook facility-location problem. Place k servers to minimise the mean travel time of a population of users [1], [2], with some allowance for congestion when demand is concentrated in time [3], [4]. Kettles differ from the servers in that literature in one way that turns out to matter. Their service time depends on their state. A kettle that boiled two minutes ago still holds water close to 100 °C, and the next user, after topping it up, needs only a short boil. A kettle that has been idle for an hour holds water near room temperature. The service time of each kettle therefore depends on the gaps between its uses, which depend on how demand is split between kettles, which depends on where the kettles are.

This coupling means that adding a kettle does two things. It shortens walks and queues, as any extra server would. It also halves the rate at which each kettle is used, so each one is warm less often and boils take longer. The net effect is not obvious in advance, and it changes with the time of day.

The contributions are:

- A placement model for shared kettles in which service times depend on thermal state, with boil sharing and a time-varying demand profile (Section III).
- A simulation-based swap search for the placement problem, using common random numbers, checked against exhaustive enumeration on a real floor (Section IV).
- A calibration on a 72-occupant floor and an answer, for that floor, to how many kettles it needs and where they should go (Section V, Section VI).
- A sensitivity analysis showing that users' fill habits, more than the cooling rate, decide whether a second kettle is worth installing (Section VI).

The practical answer for our floor is to move the existing kettle before buying another one. Relocation saves about four times as much time as a second kettle.

II. Related Work

The p-median problem was introduced by Hakimi [1], and the vertex-substitution heuristic of Teitz and Bart [5] remains a standard way of solving it approximately. Daskin [2] and Owen and Daskin [6] review the large literature that followed. Congestion enters location models in two main ways: through queueing at each facility, as in the M/G/1 location model of Berman et al. [3], or through dispatch among interacting

servers, as in the hypercube model of Larson [4]. In both, service times are drawn independently of the server's history. To our knowledge no location model has service times that depend on the time since the server was last used.

Batch service, in which customers who arrive during a service can join it, was analysed by Bailey [7]. Boil sharing is a version of this: a user who arrives while a kettle is boiling and finds enough water in it waits for that boil rather than starting a new one.

On the thermal side, the cooling of a partly filled kettle is a standard lumped-capacitance problem [8]. Smart-plug studies of small appliances in commercial buildings have measured kettle energy use [9], and Kastberg and Ferro [10] reported hot-drink demand profiles in open-plan offices with peaks at mid-morning and mid-afternoon similar to ours. We are not aware of work that connects the thermal state of a shared appliance to the time users spend waiting for it.

III. System Model

A. Floor and demand

The floor is a set of workstations $i \in \mathcal{I}$ and candidate kettle sites $c \in \mathcal{C}$, with walking distance d_{ic} along the floor's circulation routes. Occupant i generates drink requests as a non-homogeneous Poisson process with rate $\lambda_i(t)$. We take $\lambda_i(t) = \bar{\lambda}\, g(t)$, with a common profile g estimated from metering (Fig. 1 a) and $\bar{\lambda}$ set so that each occupant makes 2.4 drinks per working day on average. Arrivals are generated by thinning a homogeneous process [11]. Each user goes to the nearest kettle,

$$j(i) = \arg\min_{c \in S} d_{ic}, \tag{1}$$

where $S \subseteq \mathcal{C}$ is the set of sites with a kettle. We tested assignment rules that send some users past their nearest kettle to balance load; for this floor none improved on Eq. (1) by more than the simulation noise, so we keep the simpler rule.

B. Thermal state of a kettle

After a boil ends at time t_0, the residual water at temperature $T_0 \approx 100$ °C cools towards the ambient temperature T_a:

$$T(t) = T_a + (T_0 - T_a)\, e^{-(t - t_0)/\tau_c}, \tag{2}$$

with time constant τ_c depending on the residual volume [8]. A user who arrives at time t finds residual volume V_r at temperature $T_r = T(t)$, tops it up with tap water at T_{tap} to a fill volume $V_f \geq V_r$, and switches the kettle on. The mixed temperature is

$$T_{\mathrm{mix}} = \frac{V_r T_r + (V_f - V_r)\, T_{\mathrm{tap}}}{V_f}, \tag{3}$$

and the boil time and electrical energy are

$$t_b = t_0' + \frac{\rho c_p V_f (T_b - T_{\mathrm{mix}})}{\eta P}, \qquad E_b = \frac{\rho c_p V_f (T_b - T_{\mathrm{mix}})}{\eta} \tag{4}$$

where ρc_p is the volumetric heat capacity of water, $T_b = 100$ °C, P the rated power, η the heating efficiency and t_0' a dead time for the element to reach operating temperature. After pouring a drink of volume V_c the user leaves $V_f - V_c$ behind at T_b, which becomes the next user's residual.

C. Boil sharing and queueing

A user who arrives during a boil joins it if the water being boiled, less what earlier sharers will take, leaves at least V_c for them. Otherwise they queue and start a new boil when the kettle is free. Boil sharing is a batch-service mechanism [7]; on our floor 31% of drinks with one kettle were served from a shared boil.

For a single kettle with Poisson arrivals at rate λ and independent service times S, the Pollaczek–Khinchine formula gives the mean wait

$$W_q = \frac{\lambda\, \mathbb{E}[S^2]}{2(1 - \lambda \mathbb{E}[S])} \tag{5}$$

[12]. It does not apply here, because S depends on the gap since the previous boil and so on the arrival process. We use Eq. (5) only to check the simulator in the limit $\tau_c \to 0$, where every boil starts cold and the service times become independent. In that limit, with boil sharing disabled, simulated and analytical waits agreed to within 2% at all loads we tried. Stationary formulas of this kind are also known to mislead when demand peaks sharply [13], which is a second reason to simulate.

D. Latency and the placement problem

The latency of one trip, from leaving the desk to returning with a drink, is

$$L = \frac{2\, d_{i\,j(i)}}{v} + W + t_b + t_{\mathrm{prep}}, \tag{6}$$

where v is walking speed, W the wait at the kettle (zero if the kettle is free), t_b the boil time (zero for a user who shares a boil that has already started, in which case W is the remaining boil time) and t_{prep} the time to fill, pour and add milk. The placement problem for k kettles is

$$S_k^{\star} = \arg\min_{S \subseteq \mathcal{C},\, |S| = k} \mathbb{E}\big[\bar{L}(S)\big], \tag{7}$$

where $\bar{L}(S)$ is mean latency over all trips in a working day. We estimate the expectation by discrete-event simulation [14].

IV. Placement Algorithm

Eq. (7) has no closed-form objective, so we solve it with a swap search in the style of Teitz and Bart [5], evaluating each candidate set by simulation (Algorithm 1). All candidate sets are simulated with the same random number streams for arrivals, fill volumes and preparation times. This common-random-numbers scheme makes differences between candidates far less noisy than their individual estimates [14]. A swap is accepted only if it improves screening latency by more than δ,

which we set to 2 s, about twice the half-width of a 95% interval on ten screening days.

Algorithm 1 Simulation-based swap search for kettle placement

Require: candidate sites $\mathcal{C}$, number of kettles k, simulator Sim, screening days D_s, evaluation days D_e, tolerance δ
Ensure: site set $S^\star$ with $|S^\star| = k$

$S \leftarrow \emptyset$
for $m = 1$ **to** k **do**
  $c^\star \leftarrow \arg\min_{c \in \mathcal{C} \setminus S} \mathrm{Sim}(S \cup \{c\}, D_s)$ ▷ greedy addition
  $S \leftarrow S \cup \{c^\star\}$
end for
repeat
  $\mathit{improved} \leftarrow$ false
  for each $s \in S$ **and** $c \in \mathcal{C} \setminus S$ **do**
    $S' \leftarrow (S \setminus \{s\}) \cup \{c\}$
    if $\mathrm{Sim}(S', D_s) < \mathrm{Sim}(S, D_s) - \delta$ **then**
      $S \leftarrow S'$, $\mathit{improved} \leftarrow$ true ▷ common random numbers
    end if
  end for
until not $\mathit{improved}$
return $S^\star \leftarrow S$ **and** $\mathrm{Sim}(S, D_e)$ on fresh days

Each call to Sim on ten screening days takes about 0.2 s on a laptop, so the search is cheap. On our floor there are six candidate sites and 63 nonempty subsets, and we enumerated all of them to check the search. Algorithm 1 returned the enumeration optimum for $k = 1$, 3 and 4. For $k = 2$ it returned $\{S5, S6\}$, while enumeration on the screening days ranked $\{S4, S6\}$ first by 0.1 s. On 40 fresh evaluation days the two are indistinguishable (129.7 ± 1.1 s and 129.8 ± 1.1 s). On larger floors, with tens of candidate sites, enumeration is not practical and Algorithm 1 is what we would use.

V. Calibration

A. The floor

The floor is a 64 m by 26 m open-plan office in the institute's administrative building, with 72 workstations and one kettle, at site S1 in a kitchenette at the west end. The building services drawings identify six sites with a cold-water supply and a socket within 2 m: the existing kitchenette (S1), the toilet core on the south side (S2), the print room at the east end (S3), two risers on the north side (S4, S5), and a riser at the centre of the floor (S6). Walking distances were measured along circulation routes on the floor plan and multiplied by a detour factor of 1.15, estimated from 40 observed trips.

B. Metering and observation

The existing 3.0 kW kettle was metered for nine weeks (45 working days) with a smart plug sampling power at 1 Hz. Each boil appears as a block of near-constant power; we extracted 5,241 boils, a mean of 116 per day, with their duration and energy. For two of those weeks, the third author logged 412 trips in person: arrival and departure times, fill level read from the kettle's gauge, and whether the user shared a boil. Fig. 1 (a) shows the demand profile, with peaks at 09:00, 10:30, 13:00 and 15:00 to 16:00.

C. Bench tests

The cooling constant was measured on an identical kettle in the laboratory with a type-K thermocouple, at residual volumes of 0.3, 0.6 and 0.9 L, three runs each (Fig. 1 b). Fitting Eq. (2) gave $\tau_c = 18.7$, 24.1 and 29.0 min respectively, with run-to-run standard deviations below 0.8 min. We used $\tau_c = 24$ min, the value at the median residual volume of 0.6 L. A volume-dependent τ_c changed mean latencies by less than 1 s in every configuration, so we kept the constant. Heating efficiency, from 30 bench boils with measured electrical energy and temperature rise, was 0.85.

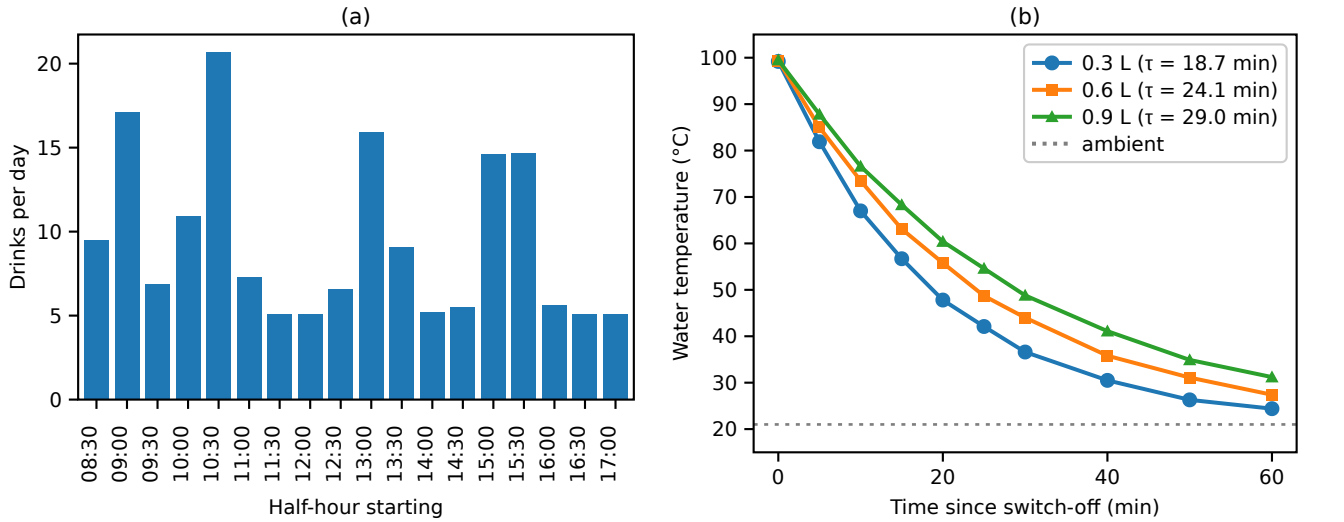


Fig. 1. Calibration data. (a) Mean drinks per day in each half-hour, from 45 days of smart-plug metering and two weeks of observation. (b) Cooling of residual water in the floor's kettle model at three residual volumes, with time constants from fits of the Newton cooling law.

Table I lists the calibrated parameters. Observed walking speed, 1.1 m/s, is below the comfortable gait speed of 1.3 to 1.4 m/s reported for healthy adults by Bohannon [15]. Most of our observed trips back to the desk involved carrying a full mug.

TABLE I
CALIBRATED MODEL PARAMETERS.

Symbol	Quantity	Value	Source
v	Walking speed	1.1 m/s	412 observed trips
τ_c	Cooling constant, 0.6 L	24 min	Bench tests
η	Heating efficiency	0.85	30 bench boils
P	Rated power	3.0 kW	Nameplate
T_{tap}	Tap water temperature	12 °C	Measured
T_a	Ambient temperature	21 °C	Building log
V_f	Fill volume	0.55 to 1.25 L	Observed trips
V_c	Volume per drink	0.30 L	Observed trips
t'_0	Dead time per boil	8 s	Plug traces
t_{prep}	Preparation time	25 s	Observed trips
$\bar{\lambda}$	Drinks per occupant per day	2.4	Metering and observation

A. Relocation

Moving the single kettle from S1 to S6 reduces mean latency by 32.7 s, or 19%. Almost all of the saving is walking time, which falls from 68.8 s to 37.1 s per trip. Waiting and boiling are unchanged, as they should be, since the same demand reaches the same kettle. Over a year of 230 working days this is about 355 occupant-hours.

B. Adding kettles

A second kettle at S5 saves a further 8.7 s per trip. The saving comes from waiting, which falls by 10.1 s, and walking, which falls by 6.1 s. Boiling, however, rises by 7.4 s. With demand split between two kettles, each is used half as often, so a user is less likely to find warm water and less likely to share a boil (20% of drinks against 31%). A third kettle saves 3.4 s more, and a fourth makes mean latency slightly worse, although the difference is within the simulation noise. Energy per drink rises with every kettle, because more boils start from cooler water. Relative to one kettle at S6, the second kettle costs about 215 kWh per year in electricity and saves about 94 occupant-hours.

D. Validation

With the kettle at S1, the simulator reproduces the metered data reasonably well. Simulated boils per day were 116 against 116 metered, which is partly by construction since $\bar{\lambda}$ was calibrated on drinks and boil sharing then determines boils. Simulated mean boil duration was 80 s against 81 s metered. The more independent check is latency: the 412 observed trips had a mean latency of 168 s (95% CI 161 to 175 s), and the simulator gives 171 s. The simulated share of drinks served from a shared boil, 32%, compares with 29% in the observed trips. As a check on the observation log itself, the mean number of people waiting at the kettle during the 10:30 peak was 0.4, consistent by Little's law [16] with the observed arrival rate of 0.69 per minute and mean wait of 35 s.

VI. RESULTS

All results are means over 40 simulated working days with fresh random numbers, with 95% intervals across days. Table II gives the best configuration for each number of kettles.

TABLE II
BEST PLACEMENT FOR EACH NUMBER OF KETTLES, WITH THE CURRENT INSTALLATION FOR COMPARISON. LATENCY IS MEAN TRIP LATENCY WITH 95% INTERVAL OVER 40 SIMULATED DAYS; WALK, WAIT AND BOIL ARE PER-TRIP MEANS; P95 IS THE 95TH PERCENTILE OF TRIP LATENCY. ENERGY IS ELECTRICAL ENERGY PER DRINK. BEST VALUE PER COLUMN IN BOLD.

Kettles	Sites	Latency (s)	p95 (s)	Walk (s)	Wait (s)	Boil (s)	Wh/drink	Boils/day
1	S1 (current)	171.1 ± 1.5	265	68.8	22.9	**54.5**	**40.9**	116
1	S6	138.4 ± 1.4	222	37.1	21.8	54.6	**40.9**	117
2	S5, S6	129.7 ± 1.1	207	31.0	11.7	62.0	46.4	135
3	S4, S5, S6	**126.3 ± 0.7**	199	25.2	7.7	68.3	51.2	146
4	S3, S4, S5, S6	126.7 ± 0.6	**198**	**22.3**	**6.3**	72.9	54.9	150

Fig. 2 shows the three components of latency across configurations. Walking and waiting fall with each kettle; boiling rises. From three kettles onwards, the rise in boiling time cancels what is saved elsewhere.

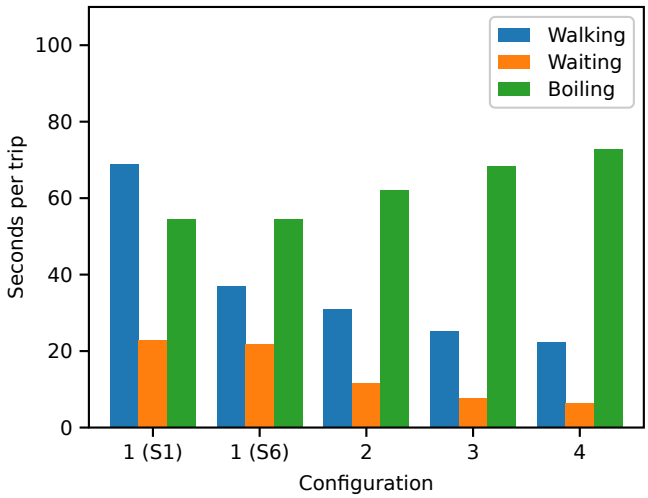


Fig. 2. Components of mean trip latency by configuration: walking, waiting at the kettle and boiling. Preparation time (25 s per trip) is the same in every configuration and is not shown.

C. Time of day

The daily mean hides a split between peak and off-peak hours. Fig. 3 shows mean latency in each half-hour for one kettle at S6 and for two at S5 and S6. In the five half-hours with the highest demand, the second kettle saves 16.6 s per trip, mostly in queueing. Between peaks the picture is different. In the half-hours from 11:30 to 12:59, 14:00 to 14:59 and 16:00 to 17:29, demand is low enough that there is rarely a queue, and the second kettle's shorter walk is more than offset by longer boils. Mean latency in those half-hours is 2.3 s higher with two kettles than with one (95% CI 0.4 to 4.2 s). We did not plan this comparison. We report it because the crossing is clear in Fig. 3 and because it is the mechanism behind the small daily gain.

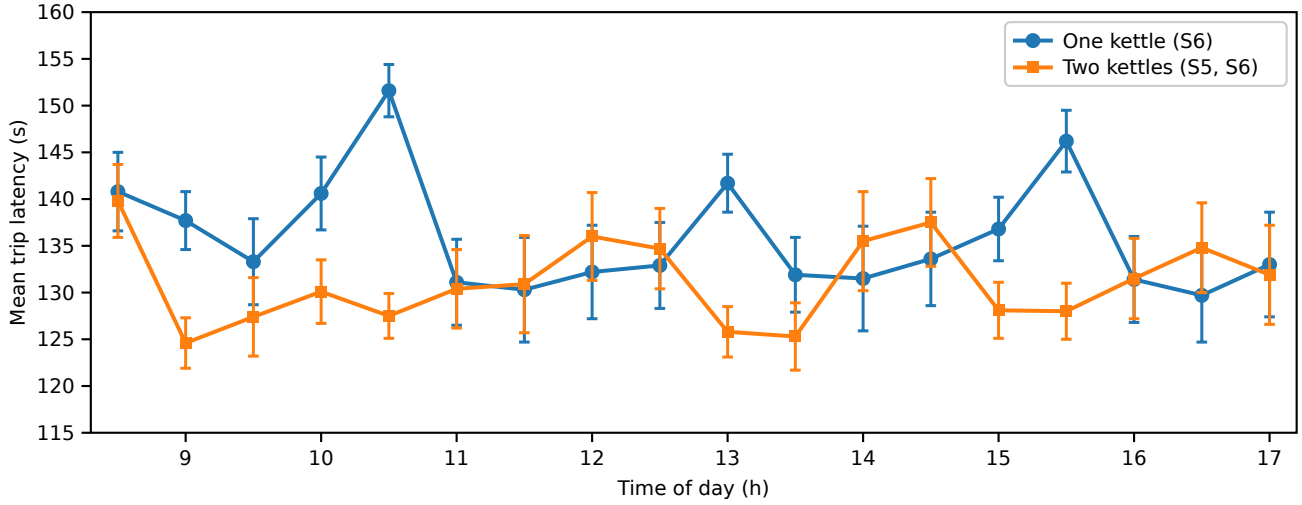


Fig. 3. Mean trip latency by half-hour for one kettle at S6 and two kettles at S5 and S6, over 40 simulated days, with 95% intervals. The second kettle helps at the demand peaks and slightly hurts in the quietest half-hours, when boils from cooler water outweigh the shorter walk.

D. Sensitivity

The thermal penalty of splitting demand depends on two things we can only partly measure: how fast the water cools, and how much water users leave behind. Table III varies the cooling constant by a factor of two in each direction. It also replaces the observed fill volumes with a fill-to-need habit, in which users fill 0.32 to 0.40 L for a single 0.30 L drink.

TABLE III
SENSITIVITY OF MEAN LATENCY (S) TO THE COOLING CONSTANT AND TO FILL HABIT, FOR ONE KETTLE AT S6 AND TWO AT S5 AND S6. BOIL IS THE MEAN BOIL TIME PER DRINK WITH ONE AND TWO KETTLES.

τ_c (min)	Fill habit	One	Two	Δ	Boil, one/two
12	Observed	143.8	137.6	−6.2	58.1 / 68.2
12	To need	133.6	117.4	−16.2	52.5 / 53.4
24	Observed	138.4	129.7	−8.7	54.6 / 62.0
24	To need	132.9	116.4	−16.5	52.0 / 52.7
48	Observed	135.1	124.0	−11.1	52.3 / 57.4
48	To need	132.5	115.8	−16.7	51.7 / 52.2

Slower cooling makes the second kettle more useful, as expected, since a less-used kettle stays warm for longer. Fill habit has a much larger effect. When users fill only what they need, there is little residual water, every boil starts nearly cold whatever the configuration, and the thermal penalty of a second

kettle almost vanishes. The second kettle then saves 16.5 s per trip, about twice the saving under observed habits. Filling to need also shortens boils with one kettle, because there is less water to heat, and lowers energy per drink from 40.9 to 36.7 Wh.

VII. Discussion

The habit of overfilling a kettle is usually discussed as an energy problem. In our model it also does something useful. The water one user leaves behind shortens the next user's boil, which makes the kettle behave as a thermal store shared by everyone nearby. Splitting users between kettles divides that store. This is why a second kettle helps so little off-peak, and why it helps twice as much when there is no store to divide.

The model leaves out several things. Users do not choose kettles strategically. We tried rules in which a user walks past their nearest kettle to one that is warm or free, which would need some way of seeing kettle state from a distance, and none helped on this floor. On a floor with larger distances, or with kettle state visible, informed choice might matter. Demand is also taken as fixed, although a closer kettle may well increase the number of drinks people make. If it does, the latency savings would shrink and the energy costs grow. Finally, the cooling constant was measured on one kettle model, and the institute's other buildings use at least two others.

For the estates team the recommendation is simple. Moving the kettle to S6 costs a water connection and saves 355 occupant-hours a year. A second kettle at S5 saves another 94 hours at a cost of 215 kWh, and the third and fourth are not worth installing.

VIII. Conclusion

Shared kettles are servers whose service times depend on how recently they were used, and this changes the usual reasoning about adding servers. On our floor the dominant decision is where to put one kettle, not how many to buy. The model and Algorithm 1 apply to any shared appliance that keeps some state between uses. Coffee machines that need a warm-up cycle are the next case we intend to measure.

Acknowledgments

We thank the occupants of the third floor for allowing an observer in their kitchenette for two weeks, and the institute's water hygiene contractor for confirming which risers could take a new connection. This work was funded by the Estates and Facilities Division of Northgate Polytechnic Institute.

References

[1] S. L. Hakimi, "Optimum Locations of Switching Centers and the Absolute Centers and Medians of a Graph," *Operations Research*, vol. 12, no. 3, pp. 450–459, 1964.

[2] M. S. Daskin, *Network and Discrete Location: Models, Algorithms, and Applications*. Hoboken, NJ: Wiley, 2013.

[3] O. Berman, R. C. Larson, and S. S. Chiu, "Optimal Server Location on a Network Operating as an M/G/1 Queue," *Operations Research*, vol. 33, no. 4, pp. 746–771, 1985.

[4] R. C. Larson, "A Hypercube Queuing Model for Facility Location and Redistricting in Urban Emergency Services," *Computers & Operations Research*, vol. 1, no. 1, pp. 67–95, 1974.

[5] M. B. Teitz and P. Bart, "Heuristic Methods for Estimating the Generalized Vertex Median of a Weighted Graph," *Operations Research*, vol. 16, no. 5, pp. 955–961, 1968.

[6] S. H. Owen and M. S. Daskin, "Strategic Facility Location: A Review," *European Journal of Operational Research*, vol. 111, no. 3, pp. 423–447, 1998.

[7] N. T. J. Bailey, "On Queueing Processes with Bulk Service," *Journal of the Royal Statistical Society, Series B*, vol. 16, no. 1, pp. 80–87, 1954.

[8] T. L. Bergman, A. S. Lavine, F. P. Incropera, and D. P. DeWitt, *Fundamentals of Heat and Mass Transfer*. Hoboken, NJ: Wiley, 2011.

[9] K. Murdoch-Ansah and D. Pelletier, "Smart-Plug Metering of Small Appliances in Commercial Buildings," Centre for Building Energy Research, Harrowgate Institute, Tech. Rep. CBER-TR-2021-07, 2021.

[10] M. Kastberg and L. Ferro, "Hot-Beverage Demand Profiles in Open-Plan Offices," in *Proc. International Conference on Building Services Engineering*, 2022, pp. 212–219.

[11] P. A. W. Lewis and G. S. Shedler, "Simulation of Nonhomogeneous Poisson Processes by Thinning," *Naval Research Logistics Quarterly*, vol. 26, no. 3, pp. 403–413, 1979.

[12] L. Kleinrock, *Queueing Systems, Volume 1: Theory*. New York: Wiley, 1975.

[13] L. V. Green, P. J. Kolesar, and W. Whitt, "Coping with Time-Varying Demand When Setting Staffing Requirements for a Service System," *Production and Operations Management*, vol. 16, no. 1, pp. 13–39, 2007.

[14] A. M. Law, *Simulation Modeling and Analysis*. New York: McGraw-Hill, 2015.

[15] R. W. Bohannon, "Comfortable and Maximum Walking Speed of Adults Aged 20–79 Years: Reference Values and Determinants," *Age and Ageing*, vol. 26, no. 1, pp. 15–19, 1997.

[16] J. D. C. Little, "A Proof for the Queuing Formula: $L = \lambda W$," *Operations Research*, vol. 9, no. 3, pp. 383–387, 1961.