Isolation of a hidden heat load between two identical server racks from a single shared temperature reading with Gaussian and sampled posterior representations

Ricardo Marquez, Ph.D.
rmarquez123@gmail.com

September 2026

Abstract

Diagnosing a thermal fault from building sensors is often ambiguous, because the few sensors a facility carries aggregate heat from many sources, and distinct faults can predict the same readings. This report studies the smallest such case: a server room of two identical racks of 300 W base load, cooled through a building loop by chilled water at 290 K and watched by one room-air temperature sensor with 0.5 K noise, in which a hidden 250 W load sits on one rack. The question is whether a diagnosis can report that the fault is on one of two racks without naming a false compromise. It is an inference study: the rack and room energy balances, the reading and a prior that each rack is either healthy or carries about 250 W are combined into a posterior over the two faults, from the same model and the same reading of 294.641 K, in two representations. The Gaussian representation, a maximum-a-posteriori estimate with the Laplace approximation, places 125.0 W on each rack with a posterior standard deviation of 98.3 W, a state to which the prior gives almost no density. The sampled representation, 1,500 fault pairs drawn from the prior and weighted by the reading, gives probabilities of 0.51 and 0.49 of a fault beyond 100 W on the two racks, two peaks in each rack's belief, 0.967 of the weight on configurations with exactly one faulted rack, and no weight at the midpoint. The reading fixes the sum of the two faults to a standard deviation of 81.4 W and leaves their difference at its prior spread of 179 W. The price is 1,500 physics solves against one; the effective sample size is 798 of 1,500.

1 Introduction

A data center converts nearly all of its electrical power into heat that has to be removed, and cooling is a large and persistent part of the energy it uses; Masanet, Shehabi, Lei, Smith and Koomey (2020) estimate global data-center electricity use and attribute the flat trend of recent years to efficiency gains in servers and in cooling. When a rack draws more power than its meter or its model says — a misconfigured server, a failed fan, equipment installed without being recorded — the extra heat warms the room, the cooling responds, and the cause has to be found. Instrumentation is sparse: a room typically carries a handful of air temperature sensors shared by many racks, so a reading constrains the heat the racks produce together rather than the heat of any one of them. The diagnostic problem is therefore not only to detect that something is wrong but to say honestly which components could be responsible when the data cannot single one out.

Fault detection and diagnosis has a long literature. Katipamula and Brambley (2005) review the methods used in building systems, and Venkatasubramanian, Rengaswamy, Yin and Kavuri (2003) and Isermann (2006) the quantitative model-based methods of process engineering, in which

measurements are compared against a physical model and the residuals are attributed to faults. In model-based diagnosis, de Kleer and Williams (1987) made explicit that the answer to a diagnostic question can be a set of candidate explanations, ranked by probability, rather than one. On the statistical side, the Laplace approximation (Tierney and Kadane, 1986) summarizes a posterior by a Gaussian around its most probable point, which is cheap and accurate when the posterior has one peak; a mixture of Gaussians (McLachlan and Peel, 2000) is the standard way to state a belief with several peaks. Sequential Monte Carlo methods represent a posterior by weighted samples instead: the bootstrap filter of Gordon, Salmond and Smith (1993) and the Monte Carlo filter of Kitagawa (1996) propagate samples through a model and weight them by the readings, the effective sample size of Kong, Liu and Wong (1994) and Liu and Chen (1998) measures how many of the weighted samples carry the information, and Douc and Cappé (2005) compare the resampling schemes used when it falls; Doucet, de Freitas and Gordon (2001) collect the methods. Loeliger (2004) introduces the factor graphs in which such models are written. None of these constructions is new here, and none is claimed as new.

TERRA is a general-purpose engine for probabilistic inference on physical networks. It treats a network as a graph on which the physics is written locally: a conservation law at every node, a constitutive law on every edge, and, attached to the same graph, whatever the sensors report and whatever is known in advance. Each of these is a relation that holds only to within a stated tolerance, and the engine weighs them together to find the most probable state of the system, the uncertainty of every quantity in it, and the relative plausibility of competing explanations of the same readings. Domains enter through plugins that supply their laws, units and file formats; the inference and the bookkeeping are shared. This study uses its data-center plugin to write the energy balances of a room with two racks and a latent fault term on each, and represents the posterior over the two faults given one room-air reading in two ways: as a Gaussian around the most probable state, and as weighted samples, each a full solve of the room's physics. Section 3 states the framework and each of these constructs precisely.

This report performs a check on the smallest system in which the ambiguity is exact. Its question is whether a diagnosis can report that a hidden heat load is on one of two identical racks, when a single shared sensor sees only their sum, without naming a compromise that is false for both. The construction gives the identical model, prior knowledge and reading to a Gaussian and to a sampled representation of the posterior and compares what each reports, from the ranked summaries an operator reads down to the densities themselves. It is an inference study, and it adds a quantified account of where the Gaussian summary fails, why, and at what cost the sampled one succeeds. Section 2 states the room, the data and the model, Section 3 the two representations and the measures, Section 4 the results, Section 5 their limits and Section 6 the answers; the appendices hold the reproduction recipe and the provenance.

2 System and data

The room. Figure 1 draws the system. Two identical racks each draw a base electrical load of 300 W, all of which becomes heat, plus a small temperature-dependent leakage of 0.1 W per kelvin of rack temperature above 300 K. Each rack also carries a hidden fault: an additional heat load that no meter records and that is zero on a healthy rack. The racks reject their heat into the shared room air through the air stream that crosses them; the room air rejects it to a building cooling loop, and the loop to chilled water held at 290 K (about 17 °C). Every link is a thermal conductance, the heat it carries per kelvin of temperature difference. The conductances are sized for the heaviest rack the study considers, the base load plus a 250 W fault, 550 W: the rack's air

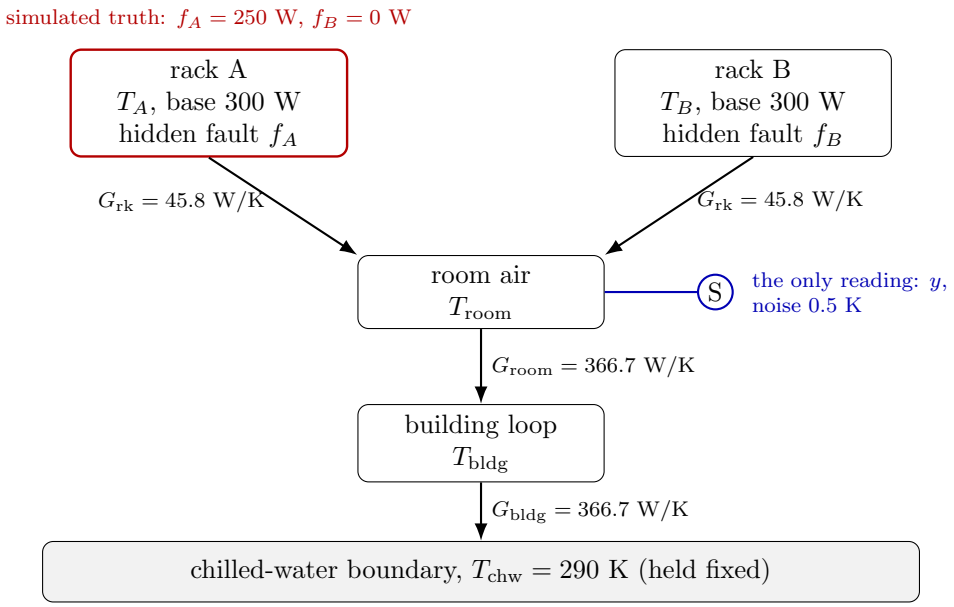


Figure 1: The room of the study. Two identical racks of 300 W base load each carry an unknown additional heat load (the hidden fault) and reject their heat into the shared room air, which is cooled through a building loop by chilled water at a fixed temperature. Arrows follow the heat; each carries its conductance. The single sensor S reads the room-air temperature. The simulation that produces the reading places 250 W on rack A and none on rack B; neither inference method is given this.

stream for a 12 K temperature rise and the room and building links for a 3 K approach at twice that heat. Table 1 lists every input.

The data. Everything is synthetic. A simulation fixes the fault on rack A at 250 W and on rack B at 0 W, solves the room, and reads the room-air temperature, 294.6406 K, to which it adds one draw of Gaussian noise of standard deviation 0.5 K from a generator seeded with 7. The single reading $y = 294.6413$ K is the only datum; neither inference method sees the simulated faults. A healthy room, both faults zero, would read 293.273 K.

The model. Figure 2 shows the model as a factor graph. The six unknowns are the two faults, the two rack temperatures, the room-air temperature and the building-loop temperature. Four energy balances tie them together, the reading attaches to the room-air temperature, and a prior attaches to each fault. The figure makes the difficulty visible: the reading reaches the faults only through the room balance, in which the heat of the two racks is added.

Assumptions. The following are held throughout.

- Steady state; the temperatures are in kelvin and the heat flows in watts.
- Linear heat transfer across each conductance, and rack leakage linear in temperature about 300 K.
- The chilled-water temperature and the base loads are known exactly; the conductances follow the sizing rule.

Table 1: Every input of the study. The conductances follow from the design rack by the sizing rule (5)–(7); everything else in the report is computed.

quantity	value	note
racks	2, identical	each with a hidden fault
base load per rack	300 W	$\alpha = 0.1$ W/K about 300 K
design rack	550 W	base load plus the faulted prior mean
rack conductance G_{rk}	45.83 W/K	550 W at a 12 K rise
room conductance G_{room}	366.67 W/K	1,100 W at a 3 K approach
building conductance G_{bldg}	366.67 W/K	1,100 W at a 3 K approach
chilled-water temperature	290 K	held fixed
sensor	room air	Gaussian noise, $\sigma_y = 0.5$ K
simulated truth	250 W on A, 0 W on B	hidden from both methods
prior per rack	mixture	weight 0.5 on each of $\mathcal{N}(0, 20^2)$ and $\mathcal{N}(250, 20^2)$ W
samples	1,500	one physics solve each, seed 7
fault threshold θ	100 W	a fault beyond it counts as real

- The sensor noise is Gaussian with standard deviation 0.5 K and known.
- Each physics relation is trusted to a residual standard deviation of 10^{-3} in its own unit.
- Each rack is a priori healthy or faulted with equal probability, independently of the other, and a fault is known to be about 250 W give or take 20 W.

3 Methods

3.1 The room as residuals

Each rack r, A or B, balances the heat it produces against the heat its air stream carries to the room; the room balances the heat of both racks against what it passes to the building loop; the loop passes the same heat to the chilled water. With P_b the base load, α the leakage slope, $T_{\mathrm{ref}} = 300$ K and f_r the hidden fault,

$$P_b + \alpha\,(T_A - T_{\mathrm{ref}}) + f_A = G_{\mathrm{rk}}\,(T_A - T_{\mathrm{room}}), \tag{1}$$

$$P_b + \alpha\,(T_B - T_{\mathrm{ref}}) + f_B = G_{\mathrm{rk}}\,(T_B - T_{\mathrm{room}}), \tag{2}$$

$$G_{\mathrm{rk}}\,(T_A - T_{\mathrm{room}}) + G_{\mathrm{rk}}\,(T_B - T_{\mathrm{room}}) = G_{\mathrm{room}}\,(T_{\mathrm{room}} - T_{\mathrm{bldg}}), \tag{3}$$

$$G_{\mathrm{room}}\,(T_{\mathrm{room}} - T_{\mathrm{bldg}}) = G_{\mathrm{bldg}}\,(T_{\mathrm{bldg}} - T_{\mathrm{chw}}), \tag{4}$$

each written as a residual, the left side minus the right, with the weight of a relation trusted to 10^{-3}. The conductances follow the sizing rule for a design rack of power $P_d = 550$ W, a rack temperature rise $\Delta T_r = 12$ K and an approach $\Delta T_a = 3$ K across the room and building links, each of which carries the heat of both racks,

$$G_{\mathrm{rk}} = P_d/\Delta T_r, \tag{5}$$

$$G_{\mathrm{room}} = 2P_d/\Delta T_a, \tag{6}$$

$$G_{\mathrm{bldg}} = 2P_d/\Delta T_a, \tag{7}$$

which give 45.83, 366.67 and 366.67 W/K. The reading adds one residual on the room-air temperature, weighted by σ_y^{-2} with $\sigma_y = 0.5$ K,

$$r_y = y - T_{\mathrm{room}}. \tag{8}$$

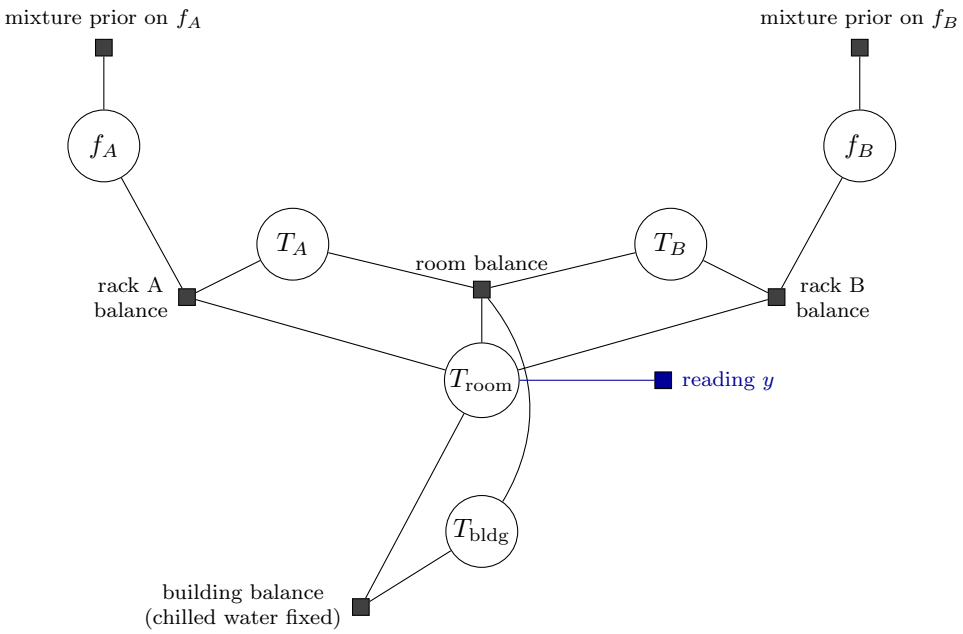


Figure 2: The model as a factor graph. Circles are the unknowns: the two hidden faults and the four temperatures. Black squares are relations, each written as a weighted residual: the energy balance at each rack, at the room air and at the building loop, the last with the chilled-water temperature fixed, and a prior on each fault. The blue square is the reading, which touches only the room-air temperature. Every path from the reading to a fault passes through the room balance, where the two racks' heat enters as a sum.

Adding the two rack balances (1) and (2) and substituting the room balance (3) eliminates the rack temperatures, and the faults enter what remains only as $f_A + f_B$. Because the model is linear, the room temperature is an affine function of that sum,

$$T_{\text{room}} = T^0_{\text{room}} + s\,(f_A + f_B), \tag{9}$$

with $T^0_{\text{room}} = 293.273$ K the healthy room and $s = 5.472$ mK/W, read from the simulated 250 W fault, which warms the sensor by 1.368 K, 2.74 times its noise. One reading therefore fixes the sum to about $\pm\sigma_y/s = \pm 91$ W and says nothing about how it is split between the racks.

The prior. What rules out a split is prior knowledge: a rack is either healthy or faulted, and a fault is about 250 W. Each fault receives, independently, a two-component Gaussian mixture with weight $\pi = 0.5$, means $\mu_0 = 0$ W and $\mu_1 = 250$ W and a common standard deviation $s_0 = 20$ W,

$$p(f_r) = \pi\,\mathcal{N}(f_r;\,\mu_1, s_0^2) + (1-\pi)\,\mathcal{N}(f_r;\,\mu_0, s_0^2). \tag{10}$$

Its density at the midpoint 125 W is 6.57×10^{-11} per watt, against 9.97×10^{-3} at either peak.

3.2 The probabilistic model

The residuals of the previous subsection are all of one kind, and the engine treats them as such; this subsection states the general model they belong to. A physical network is a graph, and the engine writes its physics on that graph. The state is a vector z of unknowns: one potential per node, one

flow per edge, and whatever else is uncertain, such as demands or parameters. Every relation the state must satisfy is a factor, written as a residual $r_k(z)$ that vanishes when the relation holds and carrying a weight w_k, the inverse variance of how far it may be off. A conservation law at a node, a constitutive law on an edge, a sensor reading and a prior on an approximately known quantity are all factors of this one kind; they differ only in their residual and their weight, a physics relation being trusted to a small tolerance, a reading to its sensor's noise, a prior to its stated spread. Treating each factor as an independent Gaussian belief about its residual gives the posterior density of the state given the readings y,

$$p(z \mid y) \propto \prod_k \exp\!\big(-\tfrac{1}{2}\, w_k\, r_k(z)^2\big). \tag{11}$$

The state that maximizes it, the *maximum-a-posteriori* (MAP) state z^*, minimizes the weighted sum of squared residuals $C(z)$ and is found by Gauss–Newton iteration on the sparse Jacobian J of the residuals, $J_{ki} = \partial r_k / \partial z_i$, with W the diagonal matrix of weights, from a flat start until the gradient norm falls below a tolerance ϵ:

$$C(z) = \tfrac{1}{2} \sum_k w_k\, r_k(z)^2, \tag{12}$$

$$z^* = \arg\min_z C(z), \tag{13}$$

$$z \leftarrow z - (J^\top W J)^{-1} J^\top W\, r(z) \quad \text{until } \|J^\top W\, r(z)\| < \epsilon. \tag{14}$$

Near the minimum the cost (12) is approximately quadratic, so the posterior is approximately Gaussian about z^* with information matrix Λ. This is the *Laplace approximation*, and the *posterior standard deviation* σ_i of any unknown is read from the diagonal of Λ^{-1}, from the same sparse factorization the last Gauss–Newton step used:

$$p(z \mid y) \approx \mathcal{N}\big(z^*, \Lambda^{-1}\big), \tag{15}$$

$$\Lambda = J^\top W J, \tag{16}$$

$$\sigma_i = \sqrt{(\Lambda^{-1})_{ii}}. \tag{17}$$

When several factor graphs compete to explain the same readings — a leak here, a leak there, no leak — each is compared by its *evidence* $Z = \int \prod_k \exp(-\frac{1}{2} w_k r_k(z)^2)\, dz$, the probability of the readings under the model once the extra freedom of its unknowns has been paid for. In the Laplace approximation the evidence has a closed form, and normalizing every factor by its own weight lets graphs with different numbers of unknowns d and factors compare as densities,

$$\log Z = -C(z^*) + \tfrac{d}{2} \log 2\pi - \tfrac{1}{2} \log\det \Lambda + \sum_k \tfrac{1}{2} \log \frac{w_k}{2\pi}. \tag{18}$$

Differences of $\log Z$ between hypotheses, from (18), are in *nats*, natural-log units of odds; five nats is odds of about 150 to 1. The posterior probability of a hypothesis is its evidence normalized over all hypotheses under equal prior odds. Everything that follows is a choice of variables and factors on this graph, followed by these three operations.

3.3 The Gaussian representation

The maximum-a-posteriori iteration and the Laplace approximation of Section 3.2 require a Gaussian prior, so the mixture (10) is replaced by the Gaussian with the same mean and variance,

$$\bar{\mu} = \pi \mu_1 + (1 - \pi) \mu_0, \tag{19}$$

$$\bar{s}^2 = s_0^2 + \pi(1-\pi)(\mu_1 - \mu_0)^2, \tag{20}$$

here $\bar{\mu} = 125$ W and $\bar{s} = 126.59$ W. The fit (13) of all residuals gives the most probable state, and (15) its covariance, from which the two faults' standard deviations, their covariance and the standard deviations of their sum and their difference follow. Each fault's 95 % interval is its mean ± 1.96 standard deviations, and the joint 95 % region of the two faults is the ellipse

$$(f - \hat{f})^\top \Sigma_f^{-1} (f - \hat{f}) \le \chi^2_{2,\,0.95}, \tag{21}$$

with $\hat{f}$ the estimate, Σ_f the two faults' covariance block and $\chi^2_{2,\,0.95} = 5.99$.

3.4 The sampled representation

The sampled representation keeps the mixture. It draws $N = 1,500$ fault pairs $f^{(i)}$ from the prior (10), independently for the two racks. For each pair it solves the room: the fit (13) with both faults held at their drawn values by priors of standard deviation 10^{-6} W and no reading, which returns the temperatures the physics implies for those faults. It then weights the pair by how well its room-air temperature matches the reading and normalizes the weights,

$$\log \tilde{w}_i = -\tfrac{1}{2} \left(T_{\text{room}}(f^{(i)}) - y \right)^2 / \sigma_y^2, \tag{22}$$

$$w_i = \tilde{w}_i \Big/ \textstyle\sum_j \tilde{w}_j. \tag{23}$$

This is one step of the bootstrap filter of Gordon, Salmond and Smith (1993): for a static problem it is importance sampling with the prior as the proposal, and the prior needs no approximation. The effective sample size

$$N_{\text{eff}} = 1 \Big/ \textstyle\sum_i w_i^2 \tag{24}$$

counts how many equally weighted samples the weighted set is worth; if it fell below $N/2 = 750$ the samples would be resampled systematically, and every quantity below is computed from the weighted samples before any resampling. Each rack's belief is summarized by its weighted mean, its weighted median and 2.5 and 97.5 percentiles, the mass beyond the fault threshold

$$m_r = \textstyle\sum_i w_i \, \mathbf{1}\big[\, |f_r^{(i)}| > \theta \,\big], \tag{25}$$

with $\theta = 100$ W, and the number of peaks in a weighted histogram of 20 equal bins over the range of the samples, smoothed by a three-bin moving average, in which a peak is a bin higher than both neighbors and at least 0.4 of the highest. The four configurations the prior allows — neither rack, rack A only, rack B only and both racks faulted beyond θ — each receive the weight of the samples that fall in them. Finally, the belief itself is recorded as a weighted density on fixed bins B_k of width Δ over -150 to 400 W, a range set by the prior (7.5 lobe standard deviations beyond each peak) rather than by the samples: 40 bins for each fault and 30×30 for the pair,

$$\hat{p}_k = \Delta^{-1} \textstyle\sum_{i:\, f^{(i)} \in B_k} w_i, \tag{26}$$

stored to 4 significant digits; no sample lies within 1.0×10^{-9} W of a bin edge.

Table 2: The Gaussian representation: the maximum-a-posteriori fault on each rack, its posterior standard deviation, the interval of ±1.96 standard deviations and the ratio z of mean to standard deviation. The prior it starts from is the moment-matched Gaussian of mean 125 W and standard deviation 126.6 W.

rack	mean [W]	std [W]	95 % interval [W]	z
A	125.04	98.3	[−68, 318]	1.27
B	125.04	98.3	[−68, 318]	1.27

3.5 Metrics

The measures are grouped by the question each answers. Whether the Gaussian estimate is plausible is read from its location: the estimated fault on each rack against the two values the prior allows, 0 and 250 W, and the prior density (10) at the estimate against its density at a peak. What the reading has taught is read from the covariance along the two directions that (9) separates: the standard deviation of the sum $f_A + f_B$, against its prior value $\sqrt{2}\,\bar{s} = 179.0$ W, and of the difference $f_A - f_B$, against the same prior value, which a reading that cannot see the difference must leave unchanged. Whether the sampled representation localizes the fault honestly is read from the mass m_r (25), against the 0.5 that the symmetry of the room requires of a correct answer; from the number of peaks, against the two the prior and the physics imply; and from the mass on the one-fault configurations, against the 0.50 the prior gives them. Whether the samples are adequate is read from the effective sample size (24), against the threshold $N/2$. Whether the simulated states are physical is checked against fixed bounds: every rack below 320 K, the room air above 291.15 K, and every energy balance closed to a relative residual below its bound. Cost is counted in physics solves, not timed, so that it does not depend on the machine.

4 Results

The reading. The room-air sensor reads $y = 294.6413$ K, against 294.6406 K without noise and 293.273 K for a healthy room. By (9) the 1.368 K warming from a 250 W fault is 2.74 times the noise, so the reading fixes $f_A + f_B$ to about ±91 W.

The Gaussian representation. Table 2 lists the estimate for each rack: the fault, its posterior standard deviation, its 95 % interval and the ratio z of estimate to standard deviation. Both racks receive 125.04 W, the prior mean, 125 W from the faulted value and 125 W from the healthy one, with a standard deviation of 98.3 W and $z = 1.27$. The two racks are indistinguishable in every column, so the ranking gives no order of inspection, and the prior density (10) at the estimate is 6.57×10^{-11} per watt, against 9.97×10^{-3} at a peak: the estimate names a state the prior regards as essentially impossible, and nothing in the table signals it. The covariance of the two faults is −6,360 W^2 against variances of 9,670 W^2, a correlation of −0.66. Along the sum the standard deviation is 81.4 W, down from the prior 179.0 W; along the difference it is 179 W, the prior value 179.0 W to the printed precision. The reading has narrowed the total and taught nothing about the split, as (9) requires; the Gaussian represents that correctly, but it can only place its single peak at the center of the split.

The sampled representation. Table 3 summarizes each rack's weighted samples. The weighted means, 126 and 122 W, sit near the midpoint like the Gaussian estimate, but the medians, 206 W

Table 3: The sampled representation: the weighted mean and median of each fault over the 1,500 samples, the interval between the weighted 2.5 and 97.5 percentiles, the mass m_r beyond the 100 W threshold (25) and the number of peaks in the weighted histogram.

rack	mean [W]	median [W]	95 % interval [W]	m_r	modes
A	126	206	[−29, 283]	0.51	2
B	122	39	[−33, 281]	0.49	2

Table 4: The four configurations the prior allows, with a rack counted as faulted when its fault exceeds $\theta = 100$ W: the mass each holds under the mixture prior and under the weighted samples.

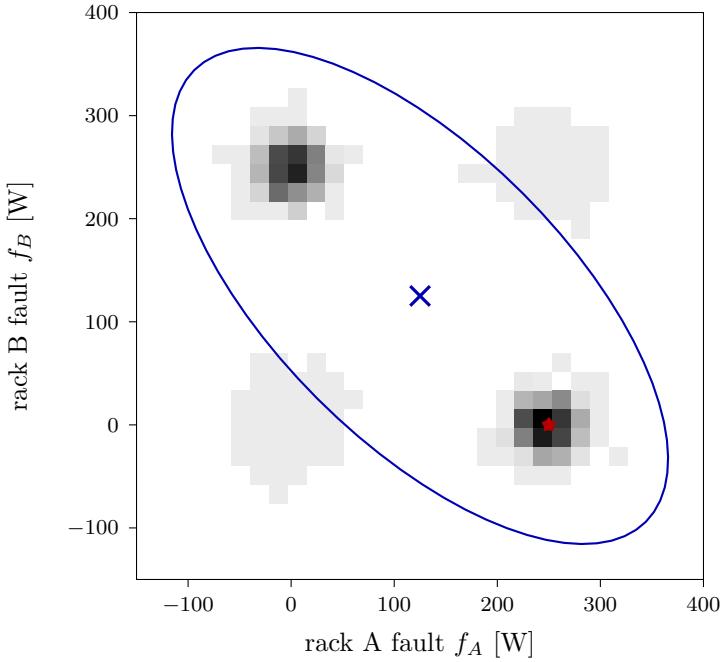
faulted beyond θ	prior mass	weighted sample mass
neither rack	0.250	0.017
rack A only	0.250	0.489
rack B only	0.250	0.478
both racks	0.250	0.016

on rack A and 39 W on rack B, fall on opposite lobes, and the masses beyond 100 W are 0.51 and 0.49, close to the 0.5 the symmetry requires. Both racks show 2 peaks, the number the prior and the physics imply. The mass is slightly higher on rack A, which carries the simulated fault; resting on one noise draw and 1,500 samples, the difference is not treated as a finding, and the ranking says in effect that either rack is equally likely. The effective sample size (24) is 798 of 1,500, above the 750 at which resampling would be applied, so none was applied. Table 4 lists the four configurations. Under the prior each holds a quarter of the mass. After the reading, rack A alone holds 0.489 and rack B alone 0.478, together 0.967; no fault keeps 0.017 and both racks 0.016. These two configurations predict a reading one fault-sized warming away from the observed one, 2.74 noise standard deviations, and are down-weighted by that distance rather than eliminated.

The joint belief. Figure 3 draws the joint density of the weighted samples over the two faults, rack A across and rack B up, with the Gaussian estimate, its 95 % region (21) and the simulated truth overlaid. The samples occupy two islands, one around the truth at (250, 0) W and one at its mirror image, with 138 of the 900 cells occupied. The cell containing the Gaussian estimate has density 0. The Gaussian region is an ellipse elongated along the difference $f_A - f_B$, as the covariance requires, and it contains both islands; it also contains the empty space between them, and its center lies there.

The belief about each rack. Figure 4 splits the same beliefs by rack. The weighted sample density on each rack keeps the two peaks of the prior, in the bins centered at −5.62 and 255.62 W, and every bin from 70.00 to 166.25 W is empty on both racks. The Gaussian density has its single peak inside that empty range, and the simulated truth, 250 W on rack A and 0 W on rack B, lies in its shoulders, while it lies on a peak of the sampled density on both racks.

Physical checks and cost. In the truth and healthy-room solves the racks lie between 299.8 and 306.7 K and the room air at or above 293.27 K; leakage stays below 0.22 % of base load; the heat and flow balances close to relative residuals below 1.0×10^{-12}. The Gaussian estimate's state has both racks at 303.9 K and passes the same checks. Across all 1,500 sample solves the racks lie



gray cells: weighted sample density (darker is denser); blue cross and ellipse: Gaussian estimate and its 95 % region; red star: simulated truth

Figure 3: Joint belief over the two rack faults. The gray cells are the weighted sample density on a fixed 30 × 30 grid of 18.33 W cells over −150 to 400 W, shaded in proportion to the density; empty cells are white. The blue cross is the Gaussian estimate and the blue curve bounds its 95 % region, clipped to the plot. The red star is the simulated truth. The samples form two islands, one around the truth and one at its mirror image; faint cells near (0, 0) and (250, 250) W hold the little weight left on no fault and on both racks faulted. The Gaussian estimate lies in the empty space between the islands.

between 297.9 and 309.5 K, below the 320 K limit. The Gaussian representation costs one fit; the sampled one 1,500 solves, one per sample.

5 Discussion

The Gaussian representation does not fail in its arithmetic. Its estimate and covariance are the correct answer to the problem it was given, and its covariance even encodes the ambiguity: a large, unchanged uncertainty along the difference between the racks. The loss happens before the reading is used, when the two-peaked prior (10) is replaced by its moment-matched Gaussian (19)–(20). A one-peaked family has to place its peak somewhere, and with symmetric evidence it places it between the explanations. A reader who sees only the ranked table has no way to tell that the midpoint is excluded by the prior.

The sampled representation succeeds because it never approximates the prior, and because in this network the samples inside each lobe carry nearly equal weight: the 20 W spread of a lobe moves the sensor by about 0.11 K, well inside the 0.5 K noise, so the effective sample size stays at 798 of 1,500. Its weighted mean is as uninformative as the Gaussian estimate; the median, the threshold mass, the peak count and the densities are what carry the answer.

The result has clear limits. It rests on one noise draw and one symmetric two-rack geometry;

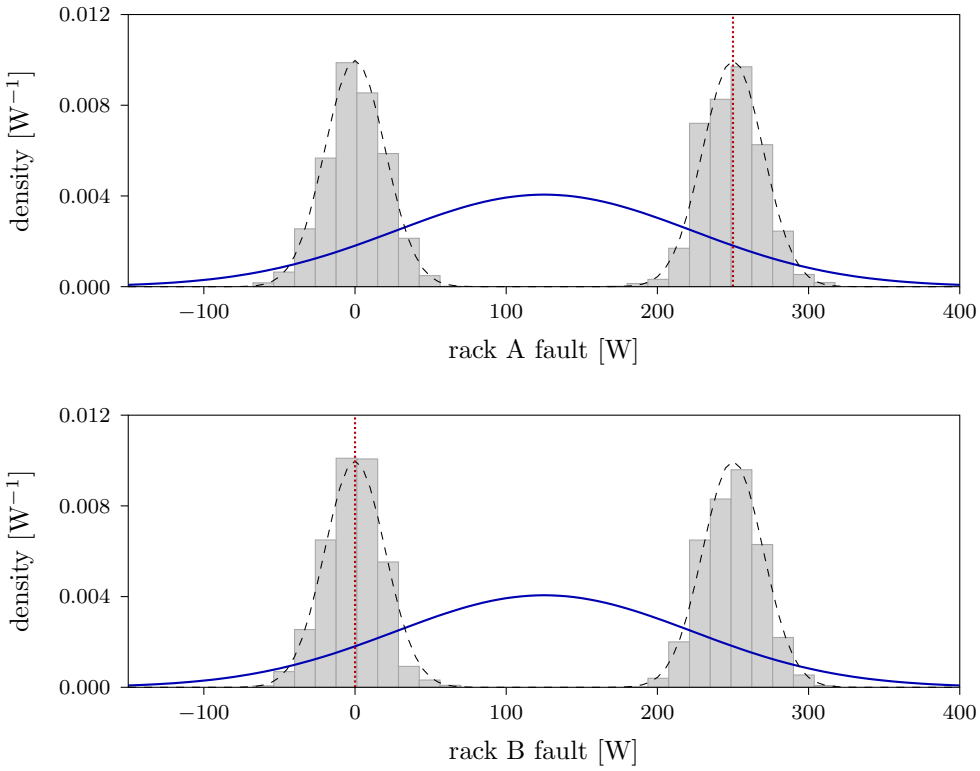


Figure 4: Belief about each rack's fault separately, rack A above and rack B below. The gray bars are the weighted sample density on 40 fixed bins of 13.75 W; the blue curve is the Gaussian estimate's marginal density; the dashed curve is the mixture prior; the dotted red line is the simulated truth. The samples keep the prior's two peaks, at 0 and 250 W, on both racks; the Gaussian has one peak where both the prior and the samples have none.

the small excess of mass on the faulted rack is not a finding. The two peaks come from the prior, not from the data: with a one-peaked prior the samples would show one peak as well, and the prior's assertion that a fault is about 250 W is what excludes the midpoint. Sampling from the prior scales poorly: with many racks, or with readings sharp enough to separate lobes, most samples would carry negligible weight and a better proposal or resampling would be needed. The peak count depends on its histogram settings (20 bins, a 0.4 floor) and is a heuristic, not a test. The model is steady and linear, the chilled-water temperature and base loads are known, and no dynamics or second sensor are considered. Finally, the Gaussian's covariance is printed to three significant digits and the densities to four, because the least significant digits of the covariance depend on the linear-algebra kernel of the machine. A second sensor that sees the racks differently, or a prior that allowed both racks to fail with a different magnitude, would change the conclusions; the first would remove the ambiguity, the second would move the belief toward the midpoint.

6 Conclusion

This report asked whether a diagnosis can report that a hidden heat load is on one of two identical server racks, when the room's single temperature sensor sees only their combined heat, without naming a compromise that is false for both. The reading of 294.641 K fixes the sum of the two

faults to a standard deviation of 81.4 W and leaves their difference at its prior spread of 179 W. Given the same model, prior knowledge and reading, a Gaussian representation of the posterior places 125.0 W on each rack, a state the prior all but excludes, with nothing in its output to mark it. A representation by 1,500 weighted samples reports probabilities of 0.51 and 0.49 of a real fault on the two racks, two peaks in each rack's belief, 0.967 of the weight on one-fault configurations and none at the midpoint, at the cost of 1,500 physics solves against one. The answer is yes, provided the belief is represented by a family that can hold two peaks; the caveat is that the peaks come from the prior. The next step is to add a sensor that separates the racks and to replace the prior as proposal with one that scales to rooms of many racks.

A Reproduction

```
cd docs/terra_graph_engine/case_studies/level2_room_hall/bimodal_fault_isolation
PYTHONPATH=../../../../../src python run.py
PYTHONPATH=../../../../../src python make_report.py
cd ../../../../..
python -m pytest -q tests/plugins/datacenter/test_bimodal_fault_isolation_case_study.py
```

`run.py` simulates the reading, forms both representations, runs the physical checks and writes `results.json`; `make_report.py` reads only `results.json`, writes `report.tex` and builds `report.pdf`. One generator seeded with 7 draws the sensor noise first and then the fault pairs: for rack A, 1,500 lobe choices, 1,500 draws from the faulted lobe and 1,500 from the healthy lobe, and then the same for rack B. The run is deterministic; the test in the last line regenerates the results, compares them with the committed file and rebuilds the tex.

B Data and provenance

No measured data enter this report. The reading is generated by `run.py` from the model of Section 3, and `results.json` is the only input the report reads. Conductances follow the sizing rule of Eq. (5)–(7). The model at commit `5be377ed` used $G_{\mathrm{rk}} = 10$ W/K, $G_{\mathrm{room}} = 20$ W/K and $G_{\mathrm{bldg}} = 100$ W/K; its figures are not comparable. The study leaves two open problems: a proposal for the sampled representation that keeps the effective sample size high when a room has many racks, and a placement of a second sensor that separates the racks at the least cost.

References

[Masanet et al. 2020] Masanet, E., Shehabi, A., Lei, N., Smith, S., Koomey, J. (2020). Recalibrating global data center energy-use estimates. Science 367(6481), 984–986. doi:10.1126/science.aba3758.

[Katipamula and Brambley 2005] Katipamula, S., Brambley, M. R. (2005). Methods for fault detection, diagnostics, and prognostics for building systems — a review, Part I. HVAC&R Research 11(1), 3–25. doi:10.1080/10789669.2005.10391123.

[Venkatasubramanian et al. 2003] Venkatasubramanian, V., Rengaswamy, R., Yin, K., Kavuri, S. N. (2003). A review of process fault detection and diagnosis. Part I: Quantitative model-based methods. Computers & Chemical Engineering 27(3), 293–311. doi:10.1016/S0098-1354(02)00160-6.

[Isermann 2006] Isermann, R. (2006). Fault-Diagnosis Systems: An Introduction from Fault Detection to Fault Tolerance. Springer, Berlin. doi:10.1007/3-540-30368-5.

[de Kleer and Williams 1987] de Kleer, J., Williams, B. C. (1987). Diagnosing multiple faults. Artificial Intelligence 32(1), 97–130. doi:10.1016/0004-3702(87)90063-4.

[Tierney and Kadane 1986] Tierney, L., Kadane, J. B. (1986). Accurate approximations for posterior moments and marginal densities. Journal of the American Statistical Association 81(393), 82–86. doi:10.1080/01621459.1986.10478240.

[McLachlan and Peel 2000] McLachlan, G., Peel, D. (2000). Finite Mixture Models. Wiley Series in Probability and Statistics, Wiley, New York. doi:10.1002/0471721182.

[Gordon et al. 1993] Gordon, N. J., Salmond, D. J., Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings F (Radar and Signal Processing) 140(2), 107–113. doi:10.1049/ip-f-2.1993.0015.

[Kitagawa 1996] Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. Journal of Computational and Graphical Statistics 5(1), 1–25. doi:10.1080/10618600.1996.10474692.

[Kong et al. 1994] Kong, A., Liu, J. S., Wong, W. H. (1994). Sequential imputations and Bayesian missing data problems. Journal of the American Statistical Association 89(425), 278–288. doi:10.1080/01621459.1994.10476469.

[Liu and Chen 1998] Liu, J. S., Chen, R. (1998). Sequential Monte Carlo methods for dynamic systems. Journal of the American Statistical Association 93(443), 1032–1044. doi:10.1080/01621459.1998.10473765.

[Douc and Cappé 2005] Douc, R., Cappé, O. (2005). Comparison of resampling schemes for particle filtering. In: Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis (ISPA 2005), pp. 64–69. doi:10.1109/ISPA.2005.195385.

[Doucet et al. 2001] Doucet, A., de Freitas, N., Gordon, N. (eds.) (2001). Sequential Monte Carlo Methods in Practice. Springer, New York. doi:10.1007/978-1-4757-3437-9.

[Loeliger 2004] Loeliger, H.-A. (2004). An introduction to factor graphs. IEEE Signal Processing Magazine 21(1), 28–41. doi:10.1109/MSP.2004.1267047.