

Taming nonequilibrium thermal fluctuations in subthreshold CMOS circuits

Nahuel Freitas,^{*,†} Geremia Massarelli[Ⓜ], Jeremy Rothschild[Ⓜ], Dylan Keane[Ⓜ], Ethan Dawe,
Sewook Hwang[Ⓜ], Akhil Garlapati^{Ⓜ,†} and Trevor McCourt^{†,‡}
Extropic Corporation, Cambridge, Massachusetts, USA



(Received 4 July 2025; revised 16 January 2026; accepted 12 February 2026; published 18 March 2026; corrected 3 April 2026)

As CMOS technology scales down, thermal fluctuations increasingly impact circuit behavior, posing challenges to conventional circuit design. However, the inherent stochasticity introduced by thermal noise is now being explored as a potential resource in the emerging field of probabilistic computing. This work presents a fully CMOS experimental platform that enables direct control over its intrinsic thermal fluctuations. These devices function as programmable multivariate Gaussian samplers, offering a hardware primitive for energy-efficient stochastic computing and serving as an experimental platform for studies in electronic noise and stochastic thermodynamics.

DOI: [10.1103/5I2z-3f1h](https://doi.org/10.1103/5I2z-3f1h)

The aggressive miniaturization of CMOS technology is a double-edged sword; as a transistor circuit is scaled down in size, it becomes faster and more energy-efficient [1], but it also becomes noisier [2]. As device dimensions and supply voltages continue to be pushed to their thermodynamic limits, the noise that is intrinsic to transistors will begin to upset traditional deterministic computer architectures [3,4].

Recent proposals suggest that this apparent roadblock imposed by noise may be circumvented by building intentionally stochastic *probabilistic circuits* (p circuits) into a computing system [5–11]. P circuits harness the noise present in all devices to generate samples from computationally useful probability distributions.

Compared to purely deterministic approaches, probabilistic computer architectures can be more energy efficient when running specific algorithms that rely heavily on random sampling. Random sampling is expensive on a deterministic computer: one must employ complex circuitry involving thousands of transistors that generate pseudorandom numbers [12]. These large circuits consume similar amounts of energy to those implementing computationally richer operations, such as addition.

For a p circuit to be practically useful in computing, it must be predictable, programmable, and performant. Predictability means that the stochastic dynamics of the p circuit obey some simple physical model that can be used to engineer it. A p circuit is programmable if the distribution it samples from can be tuned at runtime. The performance of a p circuit is measured by its correlation time, energy consumption, and physical size, which jointly quantify the space-time-energy cost of using the circuit for random sampling.

Current implementations of p circuits fall short in at least one of these key areas. P circuits built using magnetic tunnel junctions are a promising long-term direction as they obey well-characterized stochastic dynamics [13] and are straightforwardly programmed [14–16]. While these devices are, in principle, compatible with CMOS processes [17], work on integration is still underway [18,19], which limits near-term performance [20,21]. To avoid these integration issues, practitioners have also considered p -circuit architectures involving only transistors [22,23]. However, current approaches generally rely on extremely sensitive phenomena, such as the randomness of bistable latches at power up [24,25], or are based on crude treatments of noise [22], making programmability and prediction difficult.

Circuits of transistors operating in the deep subthreshold regime are a natural candidate for building practical p circuits. Charge transport in these circuits is predominately thermally activated, leading to shot-noise dynamics [26]. Recently, shot-noise models have been generalized from the single-transistor level to networks of transistors via a technique based on Markov jump processes (MJPs) [27], suggesting that the stochastic dynamics of subthreshold circuits may be rigorously predicted. Subthreshold networks are also straightforwardly programmable via control

^{*}Contact author: nfreitas@extropic.ai

[†]These authors contributed equally to this work.

[‡]Contact author: trevor@extropic.ai

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

voltages. Additionally, small transistors leak more current, increasing the performance of subthreshold circuits [28].

This paper experimentally demonstrates that practical p circuits can be built using subthreshold transistors. Specifically, we introduce and characterize a family of circuits that sample from voltage-programmable multivariate Gaussian distributions. These circuits enable the control of correlations between multiple degrees of freedom and represent a powerful primitive that can be used either as a standalone p circuit or as part of a larger circuit performing a non-Gaussian sampling operation (p bits being the most direct and prominent example [11]). Our work paves the way for near-term, large-scale computers that leverage p circuits; we have already applied the principles outlined here in designing a probabilistic computing system that runs diffusionlike models [29].

In addition to the potential of our work for probabilistic computing, it also substantially advances our physical understanding of shot noise in subthreshold transistors by validating the MJP model presented in Ref. [27]. While a large body of work studies shot noise at the single-transistor level [26,30,31], studies of noise physics in larger circuits are limited and often application-specific [32]. Our work tests a simple and generic noise model for subthreshold transistor networks against experimental data, allowing the model to be confidently used in future engineering pursuits.

The remainder of this paper will detail our Gaussian sampling circuits, which we refer to as nonequilibrium, adjustable temperature resistor networks (NEAT-RNs). We will first introduce NEAT-RNs and explain (as suggested by their name) how voltage controls can be used to program their steady-state voltage distribution. Then, we will demonstrate that the MJP model of Ref. [27] accurately predicts experimental measurements of the behavior of a simple 2D NEAT-RN. Next, we will demonstrate that the programmability of NEAT-RNs can be significantly enhanced by introducing additional degrees of freedom that mediate interaction. Finally, we will discuss the speed and energy efficiency of NEAT-RNs and demonstrate that they are sufficiently performant to be useful as part of a modern probabilistic computing system.

NEAT-RNs are composed of multiple CMOS inverters with independent gate control, where mutual capacitances couple the outputs of different inverters. The simplest non-trivial example in the family is shown in Fig. 1(a). The experiments discussed in the following involve the measurement of the statistical properties of the signals $v = (v_1, v_2)$ for different operating voltages [see Fig. 1(c) for examples of time series]. Most of the variability in those signals is explained by two sources of noise: (i) thermal fluctuations associated to the transport of charge through individual transistors, and (ii) low-frequency fluctuations in the overall conductivity of each transistor. While the first contribution is of fundamental nature, as it has a

purely thermodynamic origin [26,27], the second one is usually understood to arise due to material defects that act as fluctuating charge traps [33–37]. The obtained data is compared to an analytical model that is able to incorporate both sources of noise. However, in what follows we restrict the discussion to the analytical description of the thermal fluctuations only, which are anyway the dominant ones as they contribute around 70% of the total signal power in the frequency range given by the measurement duration and sampling rate (see Appendix D).

After linearization around an operating point, NEAT-RNs become equivalent to a noisy RC network [38–40], where the effective conductance and temperature of each resistor can be controlled at will [see Fig. 1(b)]. When the effective temperatures are equal, the system attains an equilibrium steady-state probability distribution $P(v)$ that is entirely determined by the capacitances, which cannot be controlled in nanoscopic circuits. However, if the effective temperatures differ, the system attains a nonequilibrium steady-state distribution that becomes dependent on the effective temperatures and conductances. In this way, the steady-state distribution can be controlled by manipulating the voltages V_{dd}^i and $V_{n/p}^i$ [see Figs. 1(f) and 1(g)].

Specifically, combining the MJP formalism of Ref. [27] with the Enz-Krummenacher-Vittoz (EKV) model for subthreshold transistors [41], the steady-state fluctuations of a NEAT-RN can be found to be approximately Gaussian,

$$P(v) \propto \exp -\frac{1}{2} (v - v^*)^T \Sigma^{-1} (v - v^*). \quad (1)$$

Here, v is a vector of voltages for each node in the circuit, v^* is the dc level of each voltage, and Σ is the steady-state covariance matrix. Σ is given by the solution to the Lyapunov equation,

$$G \Sigma C + C \Sigma G = 2k_b G^{1/2} \mathcal{T} G^{1/2}, \quad (2)$$

where G is a diagonal matrix of effective linear conductances associated with each inverter, $\mathcal{T}$ is a diagonal matrix of effective temperatures, and C is the circuit's Maxwell capacitance matrix. The matrix elements of $\mathcal{T}$ and G can be written explicitly in terms of the circuit's control voltages; see Appendixes B and C.

In the limit that each inverter is balanced such that $v_i^* = V_{dd}^i/2$ and the subthreshold slope parameter of each transistor approaches 1 (see Appendix C), the expressions for the matrix elements of G and $\mathcal{T}$ take an intuitive form,

$$G_{ii} = 2 \frac{I_0^p}{V_T} \exp \left(\left(V_{dd}^i/2 - V_p^i \right) / V_T \right), \quad (3)$$

$$\mathcal{T}_{ii} = T \frac{1 + \exp (V_{dd}^i/2V_T)}{2}. \quad (4)$$

Here, I_0^p is the EKV model current parameter for each transistor. T is the physical temperature of the system




FIG. 1. Predictable and programmable Gaussian sampling with a NEAT-RN. (a) A simple NEAT-RN: two independently controlled, capacitively coupled CMOS inverters. (b) The fluctuations around the dc operating point of a NEAT-RN are equivalent to those of a linear RC circuit with controllable effective conductances and temperatures. (c) Snapshot of the measured time series of voltages v_1 and v_2 , for the operating point indicated by the arrow. For different points along the voltage sweep mentioned in the text. (d) Aspect ratio of the covariance matrix of signals v_1 and v_2 , (e) angle of the main principal component of the same matrix, and (f) normalized gyration coefficient l (integration time $t_l = 200$ ns). (g) Example of the delayed autocorrelation (left) and cross-correlation (right) functions. (h) Example centered histograms of the signals v_1 and v_2 for two operating points. Dashed lines indicate the principal axes. In all cases we fixed the powering voltages $V_{dd}^1 = 137 \pm 1$ mV, $V_{dd}^2 = 142 \pm 1$ mV and NMOS gate voltages $V_n^1 = V_n^2 = 76 \pm 1$ mV, while PMOS gate voltages satisfy $V_p^1 + V_p^2 = 200 \pm 1$ mV.

and $V_T = k_b T / q_e$ the associated thermal voltage (q_e is the electron charge). We chose to eliminate the control voltage V_n^i from this expression using the balance constraint. This decision was arbitrary, and the expressions could be written just as easily with V_p^i eliminated.

This limiting case makes it clear that NEAT-RNs are highly programmable; the effective conductance of each branch can be set by varying the gate control voltage with respect to V_{dd} , and the effective temperature can be elevated with respect to the actual ambient temperature by increasing V_{dd}^i above ground. The potential values that G_{ii} and T_{ii} can take are limited by the onset of moderate inversion, at which point the shot-noise modeling assumption used to derive Eq. (2) becomes invalid. As long as the devices are in subthreshold, T and G can be independently controlled to program the p circuit with a wide range of covariance matrices.

To empirically establish the predictability and programmability of our p circuit, we fabricated the circuit shown in Fig. 1(a) using an advanced TSMC FinFET [42–44] process and fit Eq. (2) to its steady-state voltage distribution. Specifically, we fixed all voltages except for the PMOS control voltages, which were swept such that

$V_p^1 + V_p^2 = 200 \pm 1$ mV. Voltage time series were measured at each operating point, from which a covariance matrix was estimated (the specifics of the fitting procedure are outlined in Appendix E).

We characterized the measured covariance matrix via its spectrum. Specifically, we computed the eigenvalues λ_k and eigenvectors q_k , satisfying $\Sigma q_k = \lambda_k q_k$ with $\lambda_0 \geq \lambda_1$. From the eigenvalues and eigenvectors of Σ , we define the aspect ratio of the eigenvalues α and the angle of the principal component θ ,

$$\alpha \equiv \sqrt{\lambda_0 / \lambda_1}, \quad \theta \equiv \arctan((q_0)_1 / (q_0)_2). \quad (5)$$

Examples of the observed steady-state distribution are shown in Figs. 1(f) and 1(g).

Figures 1(c) and 1(d) compare the observed dependence of α and θ on V_p^1 to the best fit of Eq. (2). These results indicate that the fluctuations in the signals v_1 and v_2 can be reliably controlled by changing the operating voltages $V_{n/p}^{1/2}$. Additionally, we observe that the simple modeling based on Eq. (2) effectively captures the main features of the data. Thus, the produced signals are both controllable and predictable.

To establish that the MJP model can predict the dynamic properties of a NEAT-RN, not just the steady-state distribution, we measure the circulation of the voltage state and compare it to the prediction of the model. This circulation can be quantified via the coefficient $L(\delta t) \equiv \langle v \times \delta v \rangle$, where $\delta v = v(t + \delta t) - v(t)$ is the displacement in time δt . At steady-state conditions, it is possible to rewrite L as

$$L(\delta t) = \mathbb{X}_{1,2}(\delta t) - \mathbb{X}_{2,1}(\delta t), \quad (6)$$

where $\mathbb{X}_{j,k}(\delta t) \equiv \langle v_j(t) v_k(t + \delta t) \rangle - \langle v_j(t) \rangle \langle v_k(t + \delta t) \rangle$ is a delayed correlation, which can be obtained from time series. The previous expression is known as the cross-correlation asymmetry and is a measure of time-reversal symmetry breaking [45–47]. For example, in Fig. 1(g) we show the normalized delayed correlations $\chi_{j,k} \equiv \mathbb{X}_{j,k} / \sqrt{\text{var}(v_j) \text{var}(v_k)}$ for a particular operating point. We observe that the delayed correlations are indeed asymmetric, indicating a breakdown of time-reversal invariance due to the nonequilibrium conditions.

By analyzing the stochastic differential equation that describes the dynamics of a NEAT-RN, it is possible to derive the following expression for time averages of delayed correlations (see Appendix F):

$$\begin{aligned} \bar{\mathbb{X}}_{j,k} &= \frac{1}{t_I} \int_0^{t_I} d\tau \mathbb{X}_{j,k}(\tau) \\ &= \frac{1}{t_I} \left[\Sigma \left(\mathbb{1} - e^{-t_I G C^{-1}} \right) C G^{-1} \right]_{j,k}. \end{aligned} \quad (7)$$

Figure 1(e) shows the normalized circulation coefficient $l = (\mathbb{X}_{1,2} - \mathbb{X}_{2,1}) / \sqrt{\text{var}(v_1) \text{var}(v_2)}$ for an integration time of $t = 200$ ns and the comparison with the model best fit. As before, the agreement shows that the statistical properties of the signals can be reliably controlled and predicted.

Additionally, these results show that the device in Fig. 1(a) can be considered an electronic implementation of a Brownian gyrator, an elementary heat engine [48], using CMOS circuits. Previous implementations of Brownian gyrators have varied the environmental temperature directly by heating or cooling the circuit [49], which is more cumbersome than our electronic implementation.

NEAT-RNs with additional nodes allow for richer programmability than the minimal example considered thus far. In such circuits, one can split the nodes into visible and latent nodes. Then, the correlations among visible nodes can be shaped by their common interaction with latent ones, which naturally enhances the range of achievable distributions.

To experimentally establish this enhanced programmability, we built a 3-degrees-of-freedom (DOF) NEAT-RN that utilizes a latent inverter to cover a larger range of two-dimensional (2D) distributions, as shown in Fig. 2(a). The




FIG. 2. Increasing the programmability of a NEAT-RN using latent variables. (a) Extending the circuit in Fig. 1(a) by adding an extra degree of freedom. The self-capacitance of each free node is omitted. (b) Observed points in the plane angle-aspect ratio for 8000 iterations of the control space exploration procedure, with $V_{dd}^{1/2} \simeq 5.5V_T$ and $V_{dd}^0 \simeq 7V_T$. (c) Results of an analogous protocol applied to the circuit in Fig. 1(a).

additional DOF mediates the coupling between the two outputs, allowing for richer control of their correlation.

We studied the programmability of our NEAT-RNs by exploring their control spaces for extreme covariance matrices. Specifically, for fixed powering voltages $V_{dd}^{0/1/2}$, we generated gate voltages $V_{n/p}^i$ (in the range $[0, 0.175$ V]) via a sampling procedure designed to find configurations with extreme aspect ratios. The results are shown in Fig. 2.

By comparing Figs. 2(b) and 2(c), we can see that our 3-DOF NEAT-RN is much more programmable than the 2-DOF version. For any given value of θ , the 3-DOF circuit allows α to be controlled within a much larger range than the 2-DOF circuit. In fact, for values of θ near $\pi/4$, α can barely be controlled for the 2-DOF circuit, whereas the 3-DOF circuit still features full programmability.

For a circuit to be an efficient entropy source in a probabilistic computer, it must generate fluctuations large enough to affect downstream devices. For thermally driven electrical systems, this means the noise generator must

produce voltage fluctuations comparable to V_T . This interaction voltage scale applies to a broad class of physical phenomena, including semiconductors and electrochemical reactions [50].

However, realizable passive circuits generate weak fluctuations and are not efficient in practical applications. Specifically, in a one-dimensional (1D) RC circuit,

$$\frac{\text{var}(V)^{RC}}{V_T^2} = \frac{C_T}{C}, \quad (8)$$

where $C_T = q_e^2/k_b T$ is the thermodynamic capacitance scale, which is approximately 6 aF at room temperature. Circuits fabricated using the latest transistor processes like the one used here feature capacitances much larger than this, $C \gtrsim 300$ aF.

As such, the noise amplification provided by NEAT-RNs is crucial to their utility in room-temperature systems. For the 1D case of our NEAT-RN operating at the balance point of Eqs. (3) and (4), the variance of the voltage fluctuations is enlarged compared to a passive RC circuit,

$$\text{var}(V) = \frac{1 + \exp(V_{dd}/2V_T)}{2} \text{var}(V)^{RC} \quad (9)$$

the amplification factor depends exponentially on V_{dd} and is around 10 for $V_{dd} \approx 6V_T$.

Because this amplification is intrinsic to the inverter and no complex external circuitry is used, it is achieved efficiently. Specifically, we define the energy consumed by the circuit per random sample as $E = P\tau$, where τ is the correlation timescale of the output signal, and P is the dc power consumption $P = V_{dd}I^*$ given the dc current I^* . At the balance point, τ is

$$\tau = \frac{C}{G}, \quad (10)$$

where G is given by Eq. (3). The energy per sample follows,

$$E = \frac{V_{dd}C^2}{q_e} \text{var}(V) \tanh\left(\frac{V_{dd}}{4V_T}\right), \quad (11)$$

where $\text{var}(V)$ is as in Eq. (9).

Equations (9) and (10) elucidate a practically useful property of our system: the timescale and variance of the noise produced by our circuit are independently controllable. τ can be made arbitrarily large or small (within practical limits) by appropriately manipulating the gate control voltages with respect to V_{dd} . In contrast, the variance depends only on the powering voltage. From Eq. (11), we can see that the energy consumed per sample does not depend on τ and scales linearly with the desired variance (in the limit $V_{dd} \gg V_T$).

The parameters found from the previous fitting can be used along with Eqs. (10) and (11) to find that NEAT-RNs could be used as a performant entropy source in a probabilistic computing system. Namely, taking $C \approx 1000$ aF, $I_0^p \approx 1$ nA, and $\Delta V = 5.5V_T$, we find that $E \approx 15$ aJ. In the same scenario, the practical minimum value of τ will be achieved when $V_p = 0$, at which point $\tau \approx 1$ ns.

Overall, we have shown that predictable, programmable, and performant p circuits can be built using networks of subthreshold transistors. These subthreshold networks can now be easily integrated with other circuitry to build probabilistic computers using advanced transistor processes. We have already proposed one such architecture in Ref. [29] and hope to report on its implementation in a future contribution.

Moreover, our results motivate a deeper exploration into *probabilistic integrated circuit design*. This burgeoning subfield of analog design is mostly unexplored, offering a rich landscape for new, useful circuit topologies and scientific discoveries. These discoveries could include other theoretical developments, building on the work in Ref. [27], or further experimental work to validate these theories.

ACKNOWLEDGMENTS

T.M. would like to thank Isaac Chuang for his invaluable advice on experimental design and the writing of the manuscript.

DATA AVAILABILITY

The data that support the findings of this article are openly available [51].

APPENDIX A: FABRICATION, CONTROL, AND MEASUREMENT OF THE DEVICES

The experiments conducted in this work utilized a test chip specifically designed to characterize noisy subthreshold circuits in detail. Our chip featured 8-bit DACs with a dynamic range of [0, 0.175 V] that were used for manipulating control voltages. The output signals of each experiment were measured using high-bandwidth and high-input-impedance amplifiers, which were also implemented on the same die. This amplification chain allowed the analog signals to be routed off chip for measurement using a 1-GHz oscilloscope.

APPENDIX B: SHOT-NOISE MODELS AND THEIR DIFFUSIVE LIMIT

We consider the stochastic description of nonlinear electronic circuits developed in Ref. [27], where each conduction device in the circuit exhibits shot noise. The state of the circuit is described by the net number of elementary charges in the free nodes (i.e., the nodes that are not

regulated by voltage sources), given by a vector $n \in \mathbb{N}^d$, where d is the number of free nodes. Pairs of free nodes can be connected via conduction devices (diodes, tunnel junctions, transistors, etc.) that are identified by an integer index $\rho > 0$. If a conduction device is connected between two free nodes, then elementary charges can jump between them in both directions. For each device ρ , we assign a transition rate $\lambda_{+\rho}(n)$ to forward conduction events $n \rightarrow n + \Delta_\rho$, and a transition rate $\lambda_{-\rho}(n)$ to backward conduction events $n \rightarrow n + \Delta_{-\rho}$, with $\Delta_\rho = -\Delta_{-\rho}$. The forward direction is arbitrary, and the vectors Δ_ρ encode the change in the state n corresponding to each jump or transition.

For any state n , the voltages of the free nodes can be computed as $v = q_e C^{-1} n + v_r$, where C is the Maxwell capacitance matrix of the free nodes, q_e is the charge of the elementary charges, and v_r is a constant vector that depends on the regulated voltages. Let $P_t(v)$ be the probability to observe voltages v at time t . Given the previous description, $P_t(v)$ evolves according to the master equation:

$$\begin{aligned} \partial_t P_t(v) = & \sum_{\rho} \lambda_{\rho}(v - q_e C^{-1} \Delta_{\rho}) P_t(v - q_e C^{-1} \Delta_{\rho}) \\ & - \sum_{\rho} \lambda_{\rho}(v) P_t(v), \end{aligned} \quad (\text{B1})$$

where we abuse notation by considering the rates $\lambda_{\rho}(v)$ to be now functions of the voltages v . A diffusive approximation of this Markov jump process can be obtained by a second-order truncation of the Kramers-Moyal expansion of Eq. (B1), which corresponds to the limit of large capacitances $\|C\| \gg C_T$. In that case, Eq. (B1) reduces to a Fokker-Planck equation, which implies that the dynamics of the system can be approximately described by an Itô stochastic differential equation (SDE) of the form:

$$C \cdot dv = \mu(v)dt + \sqrt{2K(v)} \cdot dW, \quad (\text{B2})$$

where W is a vector of independent Wiener processes. This approximation is uncontrolled and fails to capture large fluctuations [52,53], but properly describes the first and second moments of $P_t(v)$ in the limit of large capacitances [54]. The drift vector $\mu(v)$ and diffusion matrix $K(v)$ can be obtained from the transition rates $\lambda_{\rho}(v)$. In turn, the two transition rates $\lambda_{\pm\rho}(v)$ associated to the conduction device ρ can be related to its phenomenological IV curve $I_{\rho}(\Delta v)$ via the thermodynamic consistency relations [27]:

$$\begin{aligned} \lambda_{\rho}(v) - \lambda_{-\rho}(v) &= I_{\rho}(\Delta v_{\rho})/q_e \\ \lambda_{\rho}(v) + \lambda_{-\rho}(v) &= \coth(\Delta v_{\rho}/2V_T) I_{\rho}(\Delta v_{\rho})/q_e. \end{aligned} \quad (\text{B3})$$

Using Eq. (B3), it is possible to obtain the following expressions for the drift vector and diffusion matrix:

$$\begin{aligned} \mu(v) &\equiv \sum_{\rho} I_{\rho}(\Delta v_{\rho}) \Delta_{\rho}, \\ K(v) &\equiv \frac{q_e}{2} \sum_{\rho} I_{\rho}(\Delta v_{\rho}) \coth(\Delta v_{\rho}/2V_T) \Delta_{\rho} \cdot \Delta_{\rho}^T, \end{aligned} \quad (\text{B4})$$

where Δv_{ρ} is the voltage drop across device ρ in state v .

For circuits with deterministic fixed-point attractors v^* , for which $\mu(v^*) = 0$, the stochastic dynamics in Eq. (B2) can be linearized to

$$C \cdot dv = -G \cdot (v - v^*) dt + \sqrt{2k_b G T} \cdot dW, \quad (\text{B5})$$

where we have defined the effective conductance matrix G with elements $G_{jk} \equiv -\partial_{v_k} \mu_j(v^*)$ and the effective temperature matrix $\mathcal{T} \equiv G^{-1} K(v^*)/k_b$. The steady-state distribution of the previous dynamics is given by Eqs. (1) and (2) in the main text.

Finally, we note that for a one-dimensional case the linearized dynamics of Eq. (B5) has the following steady-state autocorrelation function:

$$\mathbb{X}(\delta t) = \frac{k_b \mathcal{T}}{C} e^{-(G/C)|\delta t|}, \quad (\text{B6})$$

which leads to the following Lorentzian one-sided power spectral density (PSD):

$$P(f) = \frac{P_0}{1 + (f/f_0)^2}, \quad (\text{B7})$$

with $P_0 = 4k_b \mathcal{T}/G$ and $f_0 = G/(2\pi C)$.

APPENDIX C: EKV-BASED MODEL OF AN INVERTER

We now consider a single inverter [the left or right pair of transistors in Fig. 1(a)]. According to the extension of the EKV model including DIBL effects as presented in Ref. [55], the current through the NMOS transistor is given by

$$\begin{aligned} I_n &= I_0^n \log^2 \left(1 + e^{(V_{gb} - V_{th}^n)/2n_n} e^{-V_{sb}/2} e^{\gamma_n V_{ds}/2} \right) \\ &\quad - I_0^n \log^2 \left(1 + e^{(V_{gb} - V_{th}^n)/2n_n} e^{-V_{db}/2} e^{-\gamma_n V_{ds}/2} \right), \end{aligned} \quad (\text{C1})$$

where I_0^n , V_{th}^n , n_n , and γ_n are model parameters and the voltages $V_{gb} = V_n$, $V_{sb} = 0$, $V_{db} = V_{ds} = v$ are here expressed in units of the thermal voltage V_T . The same expression gives the current I_p through the PMOS transistor, this time in terms of parameters I_0^p , V_{th}^p , n_p , and γ_p , by just replacing $I_p \rightarrow -I_n$ and using the voltages $V_{gb} = -(V_p - V_{dd})$, $V_{sb} = 0$, $V_{db} = V_{ds} = -(v - V_{dd})$.

It is useful to take a number of approximations in order to make analytical progress. In particular, we consider the subthreshold regime where $e^{(V_{gb}-V_T)/2n_{n/p}} \ll 1$, and neglect DIBL effects (i.e., we set $\gamma_n = \gamma_p = 0$). Then, the currents I_n and I_p reduce to the following functions of the output node voltage v :

$$\begin{aligned} I_p(v) &= \underbrace{I_0^p e^{-(V_p-V_{dd})/n_p}}_{c_p} (1 - e^{v-V_{dd}}), \\ I_n(v) &= \underbrace{I_0^n e^{V_n/n_n}}_{c_n} (1 - e^{-v}), \end{aligned} \quad (C2)$$

where we have now incorporated factors $e^{-V_{th}^{n/p}/n_{n/p}}$ into the respective constants $I_0^{n/p}$. The deterministic output voltage v^* satisfies $I_p(v^*) = I_n(v^*)$, which under the above approximations leads to the expression:

$$e^{v^*} = \frac{e^{V_{dd}}}{2} \left[\left(1 - \frac{c_n}{c_p}\right) + \sqrt{\left(1 - \frac{c_n}{c_p}\right)^2 + 4 \frac{c_n}{c_p} e^{-V_{dd}}} \right]. \quad (C3)$$

Using this, we obtain that the effective conductance is given by

$$G = V_T^{-1} \sqrt{(c_p - c_n)^2 + 4c_p c_n e^{-V_{dd}}}. \quad (C4)$$

Also, the effective temperature is given by

$$k_b T = \frac{q e I^*}{2G} \left[\coth((V_{dd} - v^*)/2) + \coth(v^*/2) \right], \quad (C5)$$

where $I^* \equiv (c_p + c_n - V_T G)/2$ is the stationary current. The balanced conditions considered in the main text are achieved when $c_n = c_p$. The expressions in the main text also assume that the subthreshold slopes are 1 for all transistors (i.e., $n_{n/p} = 1$ for all inverters).

APPENDIX D: SPECTRAL ANALYSIS

Figure 3(a) shows the PSD of the signal v_1 , obtained numerically from the time series, at different operating voltages. The dashed black lines over each PSD show the result of fitting them with the following model:

$$P(f) = \frac{P_0}{1 + (f/f_0)^2} + \frac{L}{f^\alpha} + B. \quad (D1)$$

The first term corresponds to the Lorentzian spectrum, which is expected from the theory in the previous sections. The second term aims to capture the pink noise dominating at low frequencies, while the last term accounts for a flat spectral contribution possibly associated to extrinsic noise sources (for example, the amplification stage).




FIG. 3. (a) Power spectral density of the signal v_1 for different values of the control voltage V_p^1 . (b) Total variance of the signal v_1 , as well as its different spectral contributions, as a function of V_p^1 . Same settings as in Fig. 1.

This decomposition allows to split the total power or variance of the signal into different contributions, as shown in Fig. 3(b). Finally, we note that the effective temperature can be computed from the parameters P_0 and f_0 as $k_b T = C f_0 P_0 \pi / 2$, where C is the total output capacitance of the inverter in question [for example, for the left inverter it is $C = C_1 + 1/(1/C_2 + 1/C_m)$, in terms of the self-capacitances $C_{1/2}$ and the mutual capacitance C_m].

APPENDIX E: DATA ANALYSIS AND MODEL FIT

The comparison of the experimental data with the theoretical model proceeds in three stages: (i) determination of intrinsic parameters for each individual inverter, (ii) effective modeling of $1/f$ noise sources, and (iii) determination of global capacitance matrix based on observed correlations. The raw time series and the code used for the analysis can be found in Ref. [51].

In the first stage, for each inverter we simultaneously fit the average output voltage $\langle v \rangle$, the frequency f_0 , and the effective temperature $k_b T$ based on Eqs. (B4), (B5), and (C1). The results are shown in Fig. 4. Panels (a) and (b) show that the dynamics of the system is well captured by the model based on Eq. (C1), while panel (c)




FIG. 4. For each individual inverter, (a) average voltage output, (b) frequency f_0 , and (c) effective temperature $\mathcal{T}$ as obtained from the spectrum [see Eq. (B7)]. The dashed lines indicate the theoretical results from Eqs. (B4), (B5), and (C1) after fitting the three quantities simultaneously. Same settings as in Fig. 1.

shows that the shot-noise modeling assumption leading to $K(v)$ in Eq. (B4) is able to describe the intrinsic thermal fluctuations for the operating voltages employed.

In the second stage we aim to describe the variability of the low-frequency pink-noise variance [see Fig. 3(b)], which is the second largest contribution to the total variance in the signal. Low-frequency noise in MOS transistors is usually modeled as stochastic fluctuations in the threshold voltage V_{th} , caused by the presence of fluctuating charge traps in the gate oxide, which collectively give rise to a $1/f$ -like spectrum [33–37]. A small-signal analysis of the dc model of each inverter based on Eq. (C1) offers a simple approach to understand how fluctuations in V_{th} are mapped to the measured output. For this, we just need to consider the gains $g_{n/p} = |\partial v^* / \partial V_{th}^{n/p}|$, indicating how the deterministic output voltage v^* changes with small perturbations of each of the two threshold voltages. These gains can be easily computed from the model using the parameters fitted in the previous stage [see Fig. 5(a)]. Then, if $S_{n/p}$ is the total power of the fluctuations of the threshold voltage $V_{th}^{n/p}$ in the frequency range corresponding to the total observation window and sampling rate, the total power of




FIG. 5. (a) Power gains $|g_{n/p}|^2$ for each inverter, computed from the model with the parameters obtained from the fit in Fig. 4. (b) Total variance of the pink noise for each inverter. Dashed lines correspond to the best fits according to the model in Eq. (E1). Same settings as in Fig. 1.

low-frequency fluctuations in the output signal in the same frequency range is given by

$$S_{1/f} = g_n^2 S_n + g_p^2 S_p + S_0, \quad (\text{E1})$$

where S_0 accounts for external sources of low-frequency noise, and it was assumed that the three sources of noise involved are independent. Figure 5 compares the total variance due to low-frequency noise as a function of the operating point, as well as the result from the fit using the model in Eq. (E1). We see that, despite its simplicity, this approach allows most of the variability to be explained.

In the last stage we compare the covariance matrix obtained by solving Eq. (2) in the main text with the experimentally observed covariance matrix for different operating points, which allows estimation of the self and mutual capacitances of the inverter outputs. This is done as follows. We compute the conductance G and the effective temperature $\mathcal{T}$ for each inverter at each operating point, using the model parameters obtained in the first stage. Then, given a capacitance matrix C constructed from proposed values of C_1 , C_2 , and C_m , we solve the Lyapunov equation in Eq. (2) to obtain an expected covariance matrix Σ . This covariance matrix is compared to the one estimated from the experimental data, which has the Lorentzian component of the variances as the diagonal elements and the observed correlation $\mathbb{X}_{1,2}(\delta t = 0)$ as the




FIG. 6. (a) Lorentz variance and (b) Pearson correlation coefficient of the signals $v_{1/2}$ as a function of the operating point. Same settings as in Fig. 1.

nondiagonal elements. The results from this procedure are shown in Fig. 6.

Finally, Figs. 1(d) and 1(e) in the main text show the aspect ratio and angle of the covariance matrix of the raw time series (that is, including all noise contributions). The dashed lines in the same plots were computed from the fitted model by considering a covariance matrix $\Sigma_T = \Sigma_L + \Sigma_{1/f}$, where Σ_L is obtained by solving the Lyapunov equation in Eq. (2) and $\Sigma_{1/f}$ is a diagonal matrix with elements $S_{1/f,1}$ and $S_{1/f,2}$ computed according to Eq. (E1).

APPENDIX F: GYRATION COEFFICIENT

We now sketch the derivation of Eq. (7) for the delayed correlations. We first note that, at steady-state conditions, the delayed correlations $\mathbb{X}_{j,k}(\delta t)$ can be written as

$$\begin{aligned}
 \mathbb{X}_{j,k}(\delta t) &= \int dv \int dv' (v'_k - v_k^*)(v_j - v_j^*) P(v', t + \delta t | v, t) P(v) \\
 &= \int dv (v_j - v_j^*) \left[e^{-\delta t C^{-1} G} \cdot (v - v^*) \right]_k P(v) \\
 &= \left[e^{-\delta t C^{-1} G} \Sigma \right]_{kj} = \left[\Sigma e^{-\delta t G C^{-1}} \right]_{j,k}, \quad (\text{F1})
 \end{aligned}$$

where we have used that, for the linear process in Eq. (B5), the conditional probability $P(v', t + \delta t | v, t)$ is a Gaussian distribution with an exponentially relaxing mean value $\langle v' \rangle = v^* + e^{-\delta t C^{-1} G} \cdot (v - v^*)$, and that the stationary

distribution $P(v)$ is given by Eq. (1). Equation (7) follows by just using that $\int_0^{t_f} d\tau e^{-\tau A} = A^{-1}(\mathbb{1} - e^{-t_f A})$ for any invertible matrix A .

- [1] R. Dennard, F. Gaensslen, H.-N. Yu, V. Rideout, E. Bassous, and A. LeBlanc, Design of ion-implanted MOSFET's with very small physical dimensions, *IEEE J. Solid-State Circuits* **9**, 256 (1974).
- [2] H. Nyquist, Thermal agitation of electric charge in conductors, *Phys. Rev.* **32**, 110 (1928).
- [3] L. B. Kish, End of Moore's law: Thermal (noise) death of integration in micro and nano electronics, *Phys. Lett. A* **305**, 144 (2002).
- [4] E. Rezaei, M. Donato, W. R. Patterson, A. Zaslavsky, and R. I. Bahar, Fundamental thermal limits on data retention in low-voltage CMOS latches and SRAM, *IEEE Trans. Device Mater. Rel.* **20**, 488 (2020).
- [5] B. R. Zink, Y. Lv, and J.-P. Wang, Review of magnetic tunnel junctions for stochastic computing, *IEEE J. Explor. Solid-State Comput. Devices Circuits* **8**, 173 (2022).
- [6] S. Chowdhury, A. Grimaldi, N. A. Aadit, S. Niazi, M. Mohseni, S. Kanai, H. Ohno, S. Fukami, L. Theogarajan, G. Finocchio *et al.*, A full-stack view of probabilistic computing with p -bits: Devices, architectures, and algorithms, *IEEE J. Explor. Solid-State Comput. Devices Circuits* **9**, 1 (2023).
- [7] J. Kim, J.-K. Han, H.-Y. Maeng, J. Han, J. W. Jeon, Y. H. Jang, K. S. Woo, Y.-K. Choi, and C. S. Hwang, Fully CMOS-based p -bits with a bistable resistor for probabilistic computing, *Adv. Funct. Mater.* **34**, 2307935 (2024).
- [8] I. Ahmed, P.-W. Chiu, W. Moy, and C. H. Kim, A probabilistic compute fabric based on coupled ring oscillators for solving combinatorial optimization problems, *IEEE J. Solid-State Circuits* **56**, 2870 (2021).
- [9] W. Lee, H. Kim, H. Jung, Y. Choi, J. Jeon, and C. Kim, Correlation free large-scale probabilistic computing using a true-random chaotic oscillator p -bit, *Sci. Rep.* **15**, 1 (2025).
- [10] J. Kim, S. S. Woo, and R. Sarpeshkar, Fast and precise emulation of stochastic biochemical reaction networks with amplified thermal noise in silicon chips, *IEEE Trans. Biomed. Circuits Syst.* **12**, 379 (2018).
- [11] W. A. Borders, A. Z. Pervaiz, S. Fukami, K. Y. Camsari, H. Ohno, and S. Datta, Integer factorization using stochastic magnetic tunnel junctions, *Nature* **573**, 390 (2019).
- [12] N. S. Singh, K. Kobayashi, Q. Cao, K. Selcuk, T. Hu, S. Niazi, N. A. Aadit, S. Kanai, H. Ohno, S. Fukami *et al.*, CMOS plus stochastic nanomagnets enabling heterogeneous computers for probabilistic inference and learning, *Nat. Commun.* **15**, 2685 (2024).
- [13] R. Faria, K. Y. Camsari, and S. Datta, Low-barrier nanomagnets as p -bits for spin logic, *IEEE Magn. Lett.* **8**, 1 (2017).
- [14] K. Y. Camsari, R. Faria, B. M. Sutton, and S. Datta, Stochastic p -bits for invertible logic, *Phys. Rev. X* **7**, 031014 (2017).
- [15] K. Y. Camsari, S. Salahuddin, and S. Datta, Implementing p -bits with embedded MTJ, *IEEE Electron. Device Lett.* **38**, 1767 (2017).

- [16] Y. Shao, C. Duffee, E. Raimondo, N. Davila, V. Lopez-Dominguez, J. A. Katine, G. Finocchio, and P. K. Amiri, Probabilistic computing with voltage-controlled dynamics in magnetic tunnel junctions, *Nanotechnology* **34**, 495203 (2023).
- [17] D. Vodenicarevic, N. Locatelli, A. Mizrahi, J. S. Friedman, A. F. Vincent, M. Romera, A. Fukushima, K. Yakushiji, H. Kubota, S. Yuasa *et al.*, Low-energy truly random number generation with superparamagnetic tunnel junctions for unconventional computing, *Phys. Rev. Appl.* **8**, 054045 (2017).
- [18] K. Yang, Q. Dong, Z. Wang, Y.-C. Shih, Y.-D. Chih, J. Chang, D. Blaauw, and D. Sylvester, in *2018 IEEE Symposium on VLSI Circuits* (IEEE, Honolulu, HI, USA, 2018), pp. 171–172.
- [19] D. Edelstein, M. Rizzolo, D. Sil, A. Dutta, J. DeBrosse, M. Wordeman, A. Arceo, I. Chu, J. Demarest, E. R. Edwards *et al.*, in *2020 IEEE International Electron Devices Meeting (IEDM)* (IEEE, San Francisco, CA, USA, 2020), pp. 11.5.1–11.5.4.
- [20] J. Daniel, Z. Sun, X. Zhang, Y. Tan, N. Dilley, Z. Chen, and J. Appenzeller, Experimental demonstration of an on-chip p -bit core based on stochastic magnetic tunnel junctions and 2D MoS₂ transistors, *Nat. Commun.* **15**, 4098 (2024).
- [21] J. Kaiser, W. A. Borders, K. Y. Camsari, S. Fukami, H. Ohno, and S. Datta, Hardware-aware in situ learning based on stochastic magnetic tunnel junctions, *Phys. Rev. Appl.* **17**, 014016 (2022).
- [22] P. Korkmaz, B. E. Akgul, K. V. Palem, and L. N. Chakrapani, Advocating noise as an agent for ultra-low energy computing: Probabilistic complementary metal–oxide–semiconductor devices and their characteristics, *Jpn. J. Appl. Phys.* **45**, 3307 (2006).
- [23] K. Palem and A. Lingamneni, Ten years of building broken chips: The physics and engineering of inexact computing, *ACM Trans. Embed. Comput. Syst. (TECS)* **12**, 1 (2013).
- [24] J. Holleman, S. Bridges, B. P. Otis, and C. Diorio, 3 μ W CMOS true random number generator with adaptive floating-gate offset cancellation, *IEEE J. Solid-State Circuits* **43**, 1324 (2008).
- [25] J. Bae, J. Koo, C. Shim, and B. Kim, in *2024 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, San Francisco, CA, USA, 2024), Vol. 67, pp. 284–286.
- [26] R. Sarpeshkar, T. Delbruck, and C. A. Mead, White noise in MOS transistors and resistors, *IEEE Circuits Devices Mag.* **9**, 23 (1993).
- [27] N. Freitas, J.-C. Delvenne, and M. Esposito, Stochastic thermodynamics of nonlinear electronic circuits: A realistic framework for computing around kT , *Phys. Rev. X* **11**, 031064 (2021).
- [28] E. Morifuji, T. Yoshida, M. Kanda, S. Matsuda, S. Yamada, and F. Matsuoka, Supply and threshold-voltage trends for scaled logic and SRAM MOSFETs, *IEEE Trans. Electron. Devices* **53**, 1427 (2006).
- [29] A. Jelinčić, O. Lockwood, A. Garlapati, G. Verdon, and T. McCourt, An efficient probabilistic hardware architecture for diffusion-like models, arXiv [arXiv:2510.23972](https://arxiv.org/abs/2510.23972), Submitted to npj Unconventional Computing.
- [30] G. Reimbold and P. Gentil, White noise of MOS transistors operating in weak inversion, *IEEE Trans. Electron Devices* **29**, 1722 (1982).
- [31] S. Tedja, H. H. Williams, J. Van der Spiegel, F. M. Newcomer, and R. Van Berg, Noise spectral density measurements of a radiation hardened CMOS process in the weak and moderate inversion, *IEEE Trans. Nucl. Sci.* **39**, 804 (1992).
- [32] H. Tian, B. A. Fowler, and A. El Gamal, in *Sensors, Cameras, and Systems for Scientific/Industrial Applications* (SPIE, San Jose, CA, United States, 1999), Vol. 3649, pp. 177–185.
- [33] A. G. Mahmutoglu, A. Demir, and J. Roychowdhury, in *2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)* (IEEE, San Jose, CA, United States, 2013), pp. 500–507.
- [34] B. Stampfer, A. Grill, and M. Waltl, in *Noise in Nanoscale Semiconductor Devices* (Springer, 2020), pp. 229–257.
- [35] H. Tian and A. El Gamal, Analysis of $1/f$ noise in switched MOSFET circuits, *IEEE Trans. Circuits Syst. II: Analog Digit. Signal Process.* **48**, 151 (2001).
- [36] C. Jakobson, I. Bloom, and Y. Nemirovsky, $1/f$ noise in CMOS transistors for analog applications from subthreshold to saturation, *Solid State Electron.* **42**, 1807 (1998).
- [37] M. Kirton and M. Uren, Noise in solid-state microstructures: A new perspective on individual defects, interface states and low-frequency ($1/f$) noise, *Adv. Phys.* **38**, 367 (1989).
- [38] S. Ciliberto, A. Imparato, A. Naert, and M. Tanase, Statistical properties of the energy exchanged between two heat baths coupled by thermal fluctuations, *J. Stat. Mech.* **2013**, P12014 (2013).
- [39] M. Baiesi, S. Ciliberto, G. Falasco, and C. Yolcu, Thermal response of nonequilibrium RC circuits, *Phys. Rev. E* **94**, 022144 (2016).
- [40] N. Freitas, J.-C. Delvenne, and M. Esposito, Stochastic and quantum thermodynamics of driven RLC networks, *Phys. Rev. X* **10**, 031005 (2020).
- [41] C. C. Enz and E. A. Vittoz, *Charge-Based MOS Transistor Modeling* (John Wiley & Sons, Ltd, 2006).
- [42] T. Sekigawa and Y. Hayashi, Calculated threshold-voltage characteristics of an X MOS transistor having an additional bottom gate, *Solid-State Electron.* **27**, 827 (1984).
- [43] S.-Y. Wu, C.-H. Chang, M. Chiang, C. Lin, J. Liaw, J. Cheng, J. Yeh, H. Chen, S. Chang, K. Lai *et al.*, in *2022 International Electron Devices Meeting (IEDM)* (IEEE, San Francisco, CA, USA, 2022) pp. 27.5.1–27.5.4.
- [44] J. Liu, S. Mukhopadhyay, A. Kundu, S. Chen, H. Wang, D. Huang, J. Lee, M. Wang, R. Lu, S. Lin, Y. Chen, H. Shang, P. Wang, H. Lin, G. Yeap, and J. He, in *2020 IEEE International Electron Devices Meeting (IEDM)* (IEEE, San Francisco, CA, USA, 2020), pp. 9.2.1–9.2.4.
- [45] N. Ohga, S. Ito, and A. Kolchinsky, Thermodynamic bound on the asymmetry of cross-correlations, *Phys. Rev. Lett.* **131**, 077101 (2023).
- [46] S. Liang and S. Pigolotti, Thermodynamic bounds on time-reversal asymmetry, *Phys. Rev. E* **108**, L062101 (2023).
- [47] J. Gu, Thermodynamic bounds on the asymmetry of cross-correlations with dynamical activity and entropy production, *Phys. Rev. E* **109**, L042101 (2024).

- [48] R. Filliger and P. Reimann, Brownian gyrator: A minimal heat engine on the nanoscale, *Phys. Rev. Lett.* **99**, 230602 (2007).
- [49] K.-H. Chiang, C.-L. Lee, P.-Y. Lai, and Y.-F. Chen, Electrical autonomous Brownian gyrator, *Phys. Rev. E* **96**, 032123 (2017).
- [50] S. Arrhenius, Über die Dissociationswärme und den Einfluss der Temperatur auf den Dissoziationsgrad der Elektrolyte, *Z. Phys. Chem.* **4**, 96 (1889).
- [51] Extropic Corp., Dataset for “Taming non-equilibrium thermal fluctuations in subthreshold CMOS circuits”, Zenodo (2026), <https://doi.org/10.5281/zenodo.18636981>
- [52] P. Hanggi, H. Grabert, P. Talkner, and H. Thomas, Bistable systems: Master equation versus Fokker-Planck modeling, *Phys. Rev. A* **29**, 371 (1984).
- [53] D. Roberts, T. McCourt, G. Massarelli, J. Rothschild, and N. Freitas, Improved diffusive approximation of Markov jump processes close to equilibrium, *Phys. Rev. Res.* **8**, 013161 (2026).
- [54] A. Gopal, M. Esposito, and N. Freitas, Large deviations theory for noisy nonlinear electronics: CMOS inverter as a case study, *Phys. Rev. B* **106**, 155303 (2022).
- [55] A. Low and P. Hasler, in *42nd Midwest Symposium on Circuits and Systems (Cat. No. 99CH36356)* (IEEE, Las Cruces, NM, USA, 1999), Vol. 1, pp. 141–144.

Correction: A typographical error in the text following Fig. 2 has been fixed.