


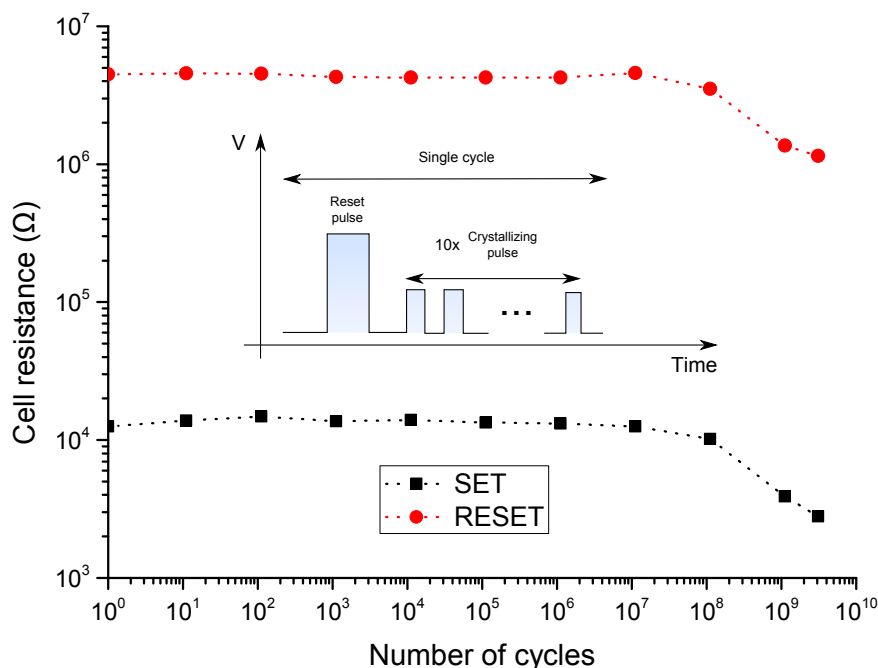
Stochastic phase-change neurons

Tomas Tuma, Angeliki Pantazi, Manuel Le Gallo,
Abu Sebastian and Evangelos Eleftheriou

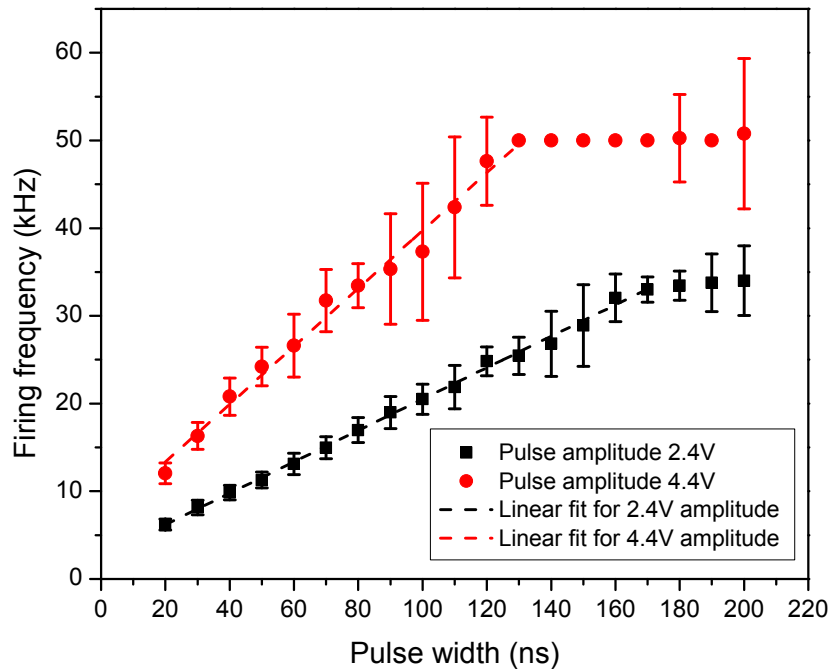


Supplementary Figure 1: Conductance curves of phase change devices for different time scales.

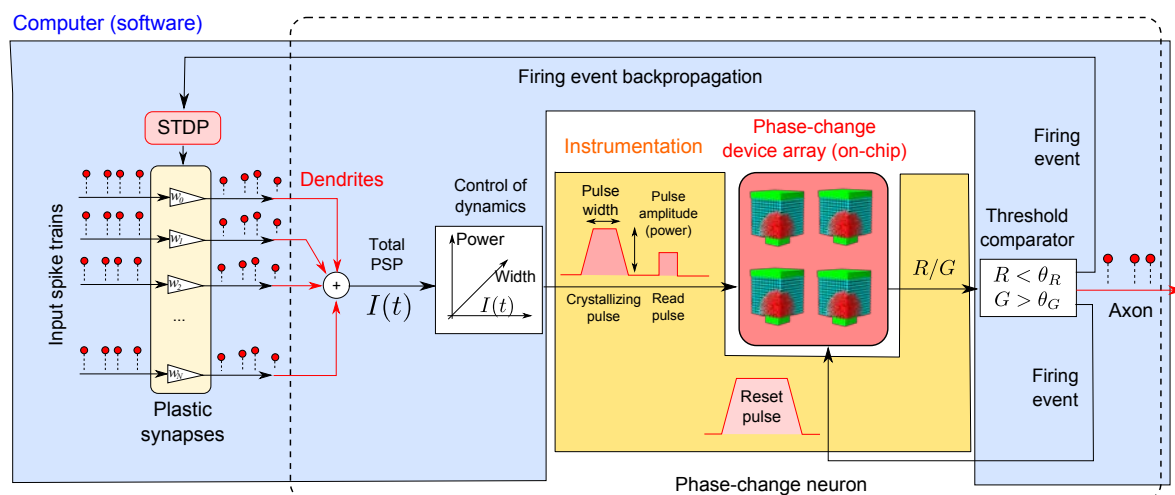
We applied a reset pulse and a subsequent sequence of crystallizing pulses with a period ranging from 100 ns to 10⁸ ns (cf. Figure 2(b) of the main text which shows the pulse sequence for a period of 100 ns). For each pulse period, the figure shows the device conductance as a function of the number of crystallizing pulses applied. We see that the basic characteristics of the device response are consistent across multiple orders of magnitude in time. This enables one to tailor the firing response of the phase change neuron according to the needs of a particular application, ranging from biologically relevant firing rates to MHz-rates. The irregularities in the conductance curves are due to the stochastic nature of the device response (see Supplementary Note 6 and the main text).



Supplementary Figure 2: Endurance of a phase change neuron. The endurance of the phase change neuron was verified when exposed to the full operating cycle, i.e., resetting the neuron to the resting potential and applying a number of crystallizing pulses that will cause the neuron to fire. In each operating cycle, we applied a RESET pulse of 6 V amplitude and 100 ns duration, and subsequently applied 10 crystallizing pulses of 1.65 V amplitude and 32.5 ns duration. The crystallizing pulses were applied every 132.5 ns (i.e., with a frequency of approx. 7.5 MHz). The figure shows the evolution of the RESET resistance (in red) and SET resistance (in black) as a function of the number of operating cycles. More than 10^9 operating cycles (more than 10^{10} pulses) were achieved before breakdown of the normal device characteristics. The slight decrease of both the RESET and SET resistance that occurs after 10^8 operating cycles is characteristic of the doped $\text{Ge}_2\text{Sb}_2\text{Te}_5$ phase change material used.



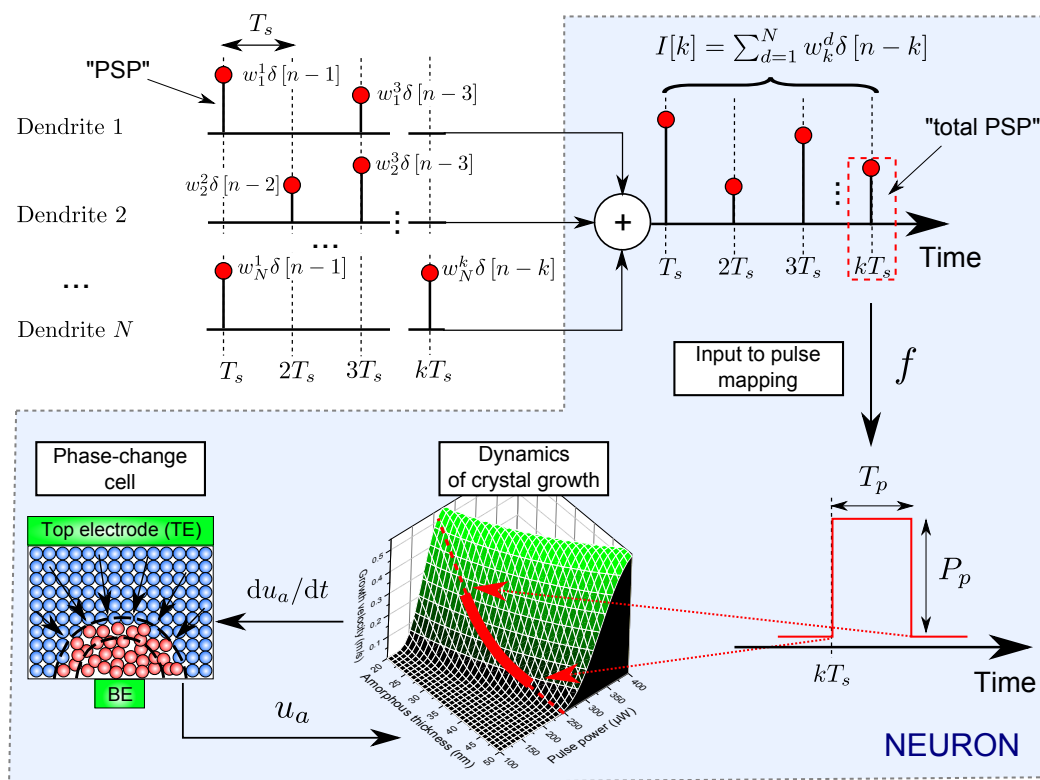
Supplementary Figure 3: Single neuron frequency response to a periodic input. We measured the average firing frequency of the phase change neuron in response to a sequence of pulses of fixed duration, amplitude and period; such sequence of pulses corresponds to a fixed-rate train of equal PSPs arriving at the neuronal input I (see Supplementary Note 1). The average frequency is determined by the inverse of the mean interspike interval (cf. Figure 4(a) of the main text). The rate response is approximately linear for a range of input pulse widths, and saturates at a level determined by the pulse power and the characteristics of the device. Each data point corresponds to the mean of 20 experimental measurements; the error bars represent the standard deviation from the mean firing frequency.



Supplementary Figure 4: Schematic of the experimental setup used for phase change neuron-based algorithms. In the experiments presented in Section 2 and Section 4, the neuronal algorithms were implemented in an experimental platform that combines a software-based environment (shown in blue), high precision instrumentation (signal generators and measurement equipment, shown in yellow) and arrays of phase change devices (shown in red). All these components were enclosed in a feedback loop. The generation of the signals, the synaptic dynamics (if present) and the translation between the neuronal input and the pulse power/width has been done in software (C#-based custom code, DSP code). The generation of the nanosecond-time-scale pulses and the measurements were performed by automated instrumentation. The membrane potential was stored and read out from arrays of phase change devices (implemented in hardware), and the logic related to the firing threshold detection and firing events was implemented in software again. As opposed to an on-chip realization of a tailored neuronal circuitry, by using a computer-driven feedback loop and yet storing information exclusively in the phase configuration of the phase change devices, we could conduct a variety of well-controlled experiments and algorithms, and estimate the fundamental power and energy requirements. Supplementary Note 3 discusses a circuit level realization of the phase change neuron.

1 Supplementary Note: Dynamics of a phase change neuron

In this supplementary note, we develop the link between the neuronal input and the evolution of the membrane potential in a phase change neuron. In doing so, we revisit the dynamics of crystal growth and elaborate on the control of the crystallization process. Subsequently, we present experimental measurements that demonstrate the theoretical considerations in a real phase change device.



Supplementary Figure 5: Link between neuronal input and membrane potential The neuronal input $I[k] = I(kT_s)$ is formed as the instantaneous sum of the postsynaptic potentials (PSP) arriving at the dendrites at time kT_s where T_s is the sampling period. Each PSP is mapped to a pulse with duration T_p and power P_p that is applied to the phase change device. In the device, the membrane potential is represented by the amorphous thickness u_a that dynamically evolves according to the applied pulse power and duration.

In the phase change neuron (Supplementary Figure 5), the input from other neurons in the network arrives at dendrites $1, \dots, N$ in the form of postsynaptic potentials (PSP). For simplicity, we neglect the dynamics related to the generation of the action potential, the process of neurotransmission and the synaptic response, and model each PSP as an all-or-none event arriving at a discrete time instance kT_s , where T_s denotes the sampling period and $k = 1, 2, \dots$. Assuming that $w_k^d \in \langle 0, 1 \rangle$ is the magnitude of the PSP at dendrite d and time kT_s (e.g. w_k^d is the weight applied to the corresponding presynaptic potential at the synapse, and is zero if there is no PSP), the PSPs received at dendrite d can be represented as the weighted spike train

$$D_d[n] = \sum_k w_k^d \delta[n-k] \tag{1}$$

where

$$\delta[n] = \begin{cases} 0 & n \neq 0 \\ 1 & n = 0 \end{cases} \quad (2)$$

is the unit sample sequence, the discrete analog of Dirac's delta function. Subsequently, the total neuronal input $I[k] = I(kT_s)$ is the instantaneous sum of the incoming PSPs (i.e., the total PSP) at all dendrites,

$$I[k] = \sum_{d=1}^N D_d[k]. \quad (3)$$

A key step in the realization of the neuronal dynamics is the link between the neuronal input $I[k]$ and the evolution of the membrane potential u . In the phase change-based neuron, the membrane potential is stored in the phase configuration of the material within the device. The phase configuration is represented for simplicity by the thickness of the amorphous portion of the material (referred to as "amorphous thickness", u_a) that upon application of electrical current evolves according to the differential equation

$$\frac{du_a}{dt} = G(u_a, P_p) \quad (4)$$

with the thickness- and power-dependent function

$$G(u_a, P_p) = -v_g(R_{th}(u_a) \cdot P_p + T_{amb}) \quad (5)$$

where v_g denotes the crystal growth velocity at the crystalline-amorphous interface, R_{th} denotes the effective thermal resistance, T_{amb} is the ambient temperature and P_p is the power applied to the electrodes. In addition to the power, also the duration of the crystallizing pulse affects the effective change in u_a . For a pulse with duration T_p and constant power P_p , the effective change in amorphous thickness is

$$\Delta u_a(T_p, P_p) = \int_{t_0}^{t_0+T_p} G(u_a(t), P_p) dt, \quad u_a(t_0) = u_0 \quad (6)$$

where t_0 denotes the initial time when the power was applied to the device. Notice that $u_a(t)$ on the right hand side of the equation evolves in time according to Equation 4 and hence, solving the integral generally corresponds to the solution of a closed-loop feedback system.

When both the power and the duration of the crystallizing pulses are controlled, the neuronal input $I[k]$ is linked to the membrane potential by the map

$$f : I[k] \rightarrow (P_p, T_p) \quad (7)$$

which effectively maps the input strength to the change in amorphous thickness,

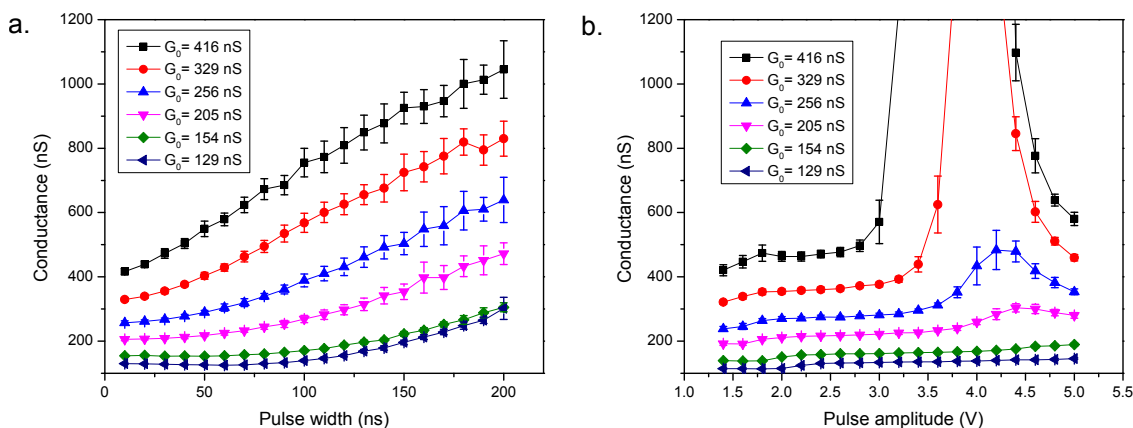
$$f_{\text{eff}} : I[k] \rightarrow \Delta u_a(P_p, T_p). \quad (8)$$

The properties of the f -mapping are key to the temporal integration properties of the phase change device: upon the arrival of input $I[k]$, the membrane potential u_a changes by $f_{\text{eff}}(I[k])$. A simple and particularly useful mapping f is the one in which the pulse power is kept constant for all $I[k]$ and the pulse width is modulated linearly with the input strength,

$$f(I[k]) := (P_{\text{const}}, \alpha I[k] + \beta). \quad (9)$$

By modulating the duration of the pulse at constant power, the effect of the nonlinearity of G on the temporal integration properties can partially be alleviated. In particular, with the linear dependence between $I[k]$ and T_p , f_{eff} is monotonous: at any fixed initial condition $u_a(t_0)$, $f_{\text{eff}}(i_1) < f_{\text{eff}}(i_2)$ for $i_1 < i_2$. This follows from Equation 6 and the fact that always $G < 0$.

We note that the pulse durations T_p must be shorter than the sampling period, $T_p < T_s$. A certain margin is preferable in a physical realization to allow for implementation-specific overhead such as leading and trailing pulse edges, and the read-out of the device conductance for the implementation of feedback functionality. Since T_p is typically on the order of tens to hundreds of nanoseconds, sampling rates up to the order of MHz should be achievable. We have used $T_p < 200$ ns and $T_s = 10$ μ s in the experiments shown in Section 2 and Section 4 of the main text, which gives a margin of two orders of magnitude between T_p and T_s .



Supplementary Figure 6: Evolution of the phase change device conductance for pulse-width and pulse-power updates for varying initial states. For a range of intermediate initial device states, the plot shows the device conductance after application of a pulse of 4 V amplitude and varying width (panel (a)), and 100 ns width and varying amplitude (panel (b)), respectively.

Supplementary Figure 6 shows the effect of choosing different f -mappings between the neuronal input $I[k]$ and the evolution of the membrane potential in a real phase change device. In these experimental measurements, we applied a single reset pulse and a single fixed crystallizing pulse of pre-determined power and width to bring the device into one of its intermediate states. We chose intermediate states in which the amorphous thickness varies from a maximum (corresponding conductance G_0 of approx. 129 nS) to a minimum (corresponding conductance G_0 of approx. 416 nS). Subsequently, for each of the intermediate states, a crystallizing pulse of given power and width was applied and the resulting conductance was measured. In Supplementary Figure 6(a), pulses of fixed power and with widths between 20 ns and 200 ns were applied. In accordance with the crystal-growth dynamics and our discussion above, the change in conductance is monotonous in T_p and for operating regions where G is approximately constant (here, regions with comparably smaller amorphous thickness) also approximately linear in T_p . For purposes of comparison, Supplementary Figure 6(b) shows the change in conductance for pulses of fixed width and with voltage varying between 1.4 V and 5 V. Here, the change in conductance is clearly linked to the shape of the function G (cf. Figure 2(a) in the main text of the paper) and is non-monotonous in the control variable. In particular, for low amorphous thickness (high initial conductance G_0) and a range of voltages between approx. 3 V and 5 V, the crystal growth velocity is high and leads to a significant drop in the amorphous thickness that can be observed as the excessive “peak” in conductance in the figure.

2 Supplementary Note: Depolarization and hyperpolarization in phase change neurons

Most integrate-and-fire neuronal models enable the realization of both excitatory and inhibitory connections between neurons. In a generic nonlinear integrate-and-fire neuron, the gradient of the membrane potential

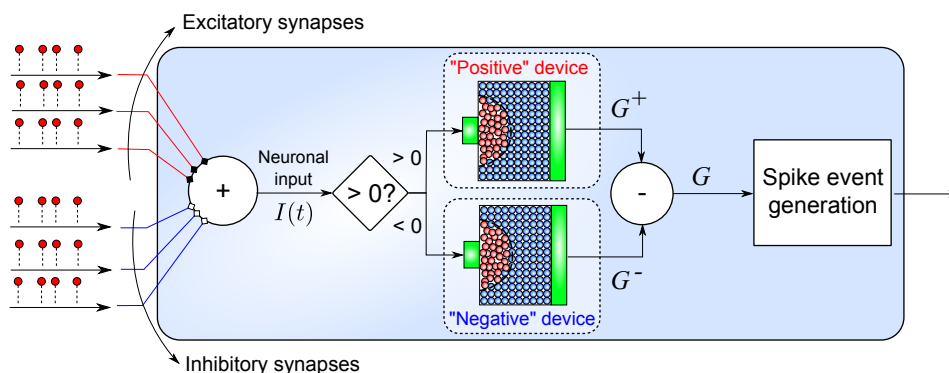
$$du/dt = F(u) + G(u) \cdot I \quad (10)$$

can be driven by either positive input $I > 0$, which corresponds to *depolarization* of the neuronal membrane due to an excitatory postsynaptic potential (PSP), or by negative input $I < 0$, which corresponds to *hyperpolarization* of the neuronal membrane due to an inhibitory PSP. While this is a crude approximation of what happens in real (biological) neurons, the ability to reverse the evolution of the membrane potential is often an important component of neuronal dynamics.

In phase change neurons that use a single phase change device to store the membrane potential, simultaneous realization of both depolarization and hyperpolarization might be costly in terms of power and device degradation. This is because the evolution of the membrane potential is determined by the dynamics of the interfacial crystal growth,

$$du_a/dt = -v_g (R_{th}(u_a)P_p + T_{amb}), u_a(0) = u_0 \quad (11)$$

with the crystal growth velocity $v_g > 0$ and power $P_p > 0$. Hence, there is no straightforward way to map the neuronal input I to the power or duration of the crystallizing pulse to realize $du_a/dt > 0$. A conventional approach that enables both positive and negative updates to u_a is to reset the device before an update to a maximum amorphous thickness and then to reach the desired u_a by a subsequent crystallizing pulse of an appropriate magnitude and duration (see e.g. phase change-based synapses by Kuzum et al. [1]). Although this is a viable solution, reset pulses require relatively high current and thus might lower the endurance of the device compared to crystallization pulses.



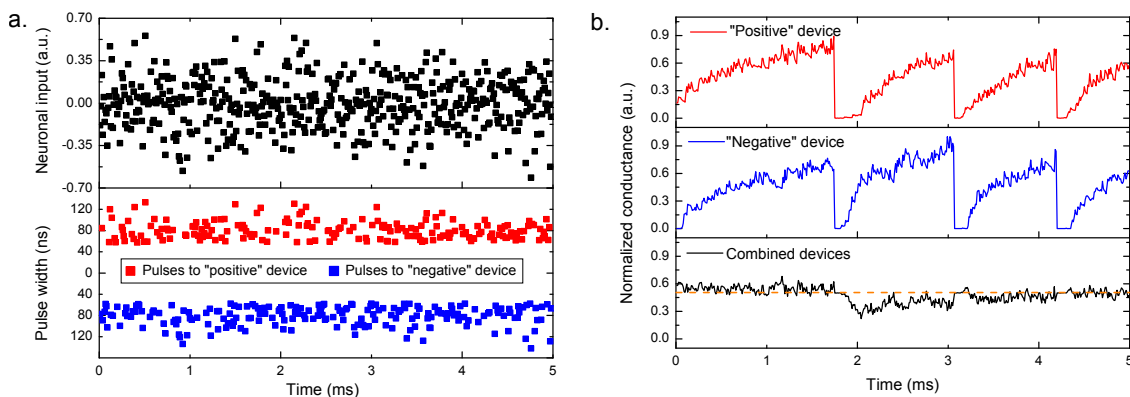
Supplementary Figure 7: phase change neuron with two phase change devices in a differential configuration

Another approach is to use multiple phase change devices to represent the state of the membrane potential in a single neuron. By doing so, the resolution of the representation can be increased (by being able to store more intermediate states), and both positive and negative updates can be realized using the crystal growth dynamics without a significant power or endurance burden. Moreover, the footprint of an individual phase change device is small relative to the rest of the circuitry associated with a phase

change neuron (see also Supplementary Note 3). In what follows, we present a particular instance of this approach in which two phase change devices are connected in a differential configuration to represent the membrane potential (Supplementary Figure 7). Denoting the conductance of the first device G^+ ("positive device") and the conductance of the second device G^- ("negative device"), the membrane potential corresponds to the combined conductance

$$G = \frac{G^+ - G^- + 1}{2}. \quad (12)$$

The normalization constants ensure that for $G^+, G^- \in \langle 0, 1 \rangle$, also $G \in \langle 0, 1 \rangle$. In this configuration, G^+ is increased upon a positive update to the membrane potential ($I > 0$) by applying a crystallizing pulse, and G^- is increased upon a negative update to the membrane potential ($I < 0$) by applying a crystallizing pulse. This ultimately enables positive and negative updates to the combined conductance G . Once either of G^+ or G^- saturates (the device is close to reaching its minimum amorphous thickness), the positive and the negative device must be reset and reprogrammed such that the overall G is preserved. A similar approach has recently been demonstrated in the context of phase change-based synapses [2, 3]; the extension of this idea to neurons is new.



Supplementary Figure 8: Experimental results: Maintaining membrane potential equilibrium in a neuron with two phase change devices in a differential configuration

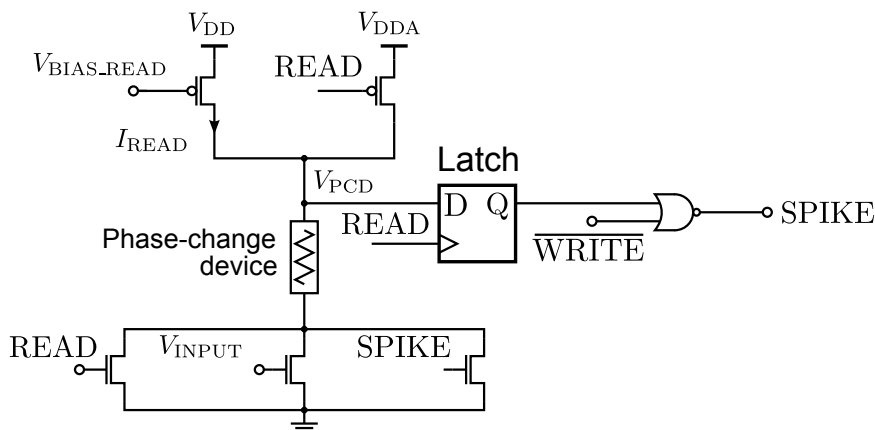
We verified the viability of this architecture in an experimental setting. In the experiment, we first generated bipolar neuronal input I by sampling a white Gaussian process with zero mean and standard deviation of 0.2 (Supplementary Figure 8a, upper display). Such input should cause the membrane potential to evolve around an equilibrium point, with frequent positive and negative updates. Subsequently, at every $t = kT_s$ ($T_s=10 \mu\text{s}$) I was translated into a crystallizing pulse of fixed amplitude and duration linearly proportional to the magnitude of I . The mapping was chosen such that for $|I| > 0.05$, the pulse durations were linearly mapped into the range between 50 ns and 200 ns. For $|I| < 0.05$, no pulse was applied. The pulse of given duration was applied to the positive or negative device, respectively, depending on the sign of I (Supplementary Figure 8a, lower display).

Supplementary Figure 8b shows the evolution of the measured conductances of the positive device, the negative device, and the combined conductance, respectively. In the plots, all conductances were normalized into the interval $\langle 0, 1 \rangle$. Because the input I is balanced around zero, we can see that the rate of increase of the device conductance is approximately the same for both the positive and the negative device, and hence the membrane potential remains at equilibrium as desired. We can also clearly notice the reset events at approx. 1.8 ms, 3 ms and 4.1 ms. These happen whenever either of the conductances

G^+ , G^- reach the threshold of 0.75, and ensure that the system does not freeze out due to saturation of the conductances. While this adds certain small circuit complexity to the implementation of two-device neurons, it enables both depolarization and hyperpolarization of the membrane potential while leveraging the benefits of low-power, high-endurance crystal-growth dynamics.


3 Circuit-level realization of phase change neurons

One of the key advantages of the phase change neurons is their potential for scalability and integration in highly dense neuromorphic systems. In this supplementary note, we discuss the circuit elements necessary for a full-fledged operation of a phase change neuron and estimate the real-estate requirements of such circuitry.



Supplementary Figure 9: Circuit realization of a phase change-based neuron

One potential circuit implementation of the phase change neuron is shown in Supplementary Figure 9. We opted for a hybrid circuit that combines analog elements, namely the phase change device and the access devices, with digital elements, namely the latch and the “NOR” logic gate, to implement the integrate-and-fire functionality of the neuron. A key advantage of the presented circuit is that all building elements can be realized with small requirements on the silicon real estate which renders the circuit particularly suitable for the implementation of highly dense populations of neurons and seamless integration with dense synaptic arrays.



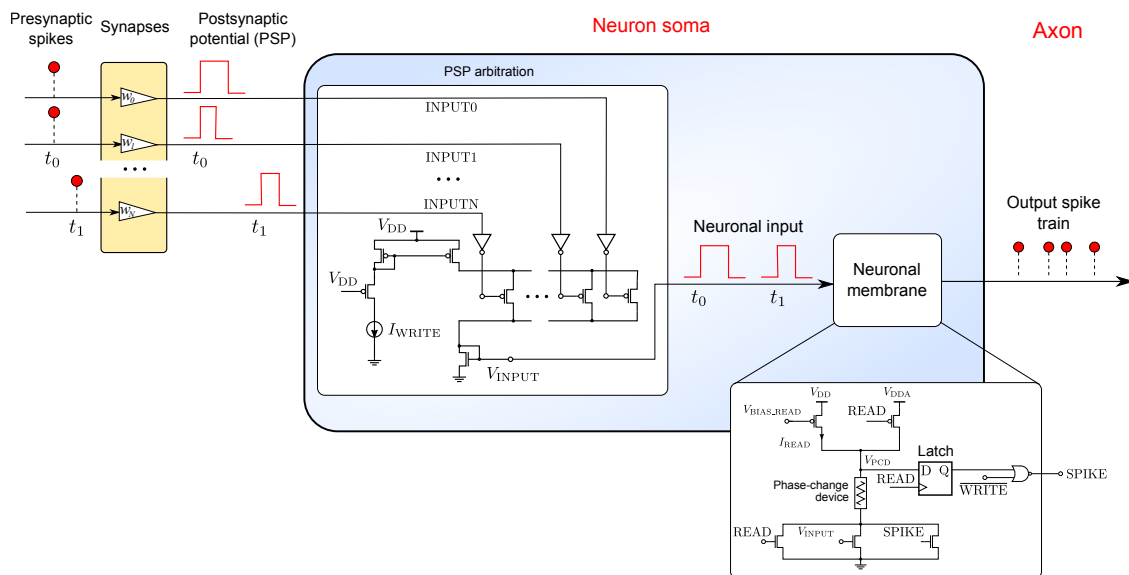
Supplementary Figure 10: Typical signal waveforms in a circuit realization of the phase change neuron

Large on-chip realizations of neuronal circuits typically operate synchronously using a sufficiently fast clock. In our circuit, the neurons are synchronized by means of two global periodic digital signals (READ and WRITE) that change the neuron state between reading the actual membrane potential

(READ) and an update of the membrane potential (WRITE). The READ and WRITE signals are assumed to be provided by a global (shared) signal generator (not shown in the schematic), with a typical period of approx. 200 ns to 500 ns, and we assume that WRITE signal always precedes the READ signal.

The neuronal input I , the equivalent of the total PSP (see Supplementary Note 1) is provided by the analog signal V_{INPUT} . The output (spike) of the neuron is represented with the digital signal SPIKE. Both the input and the output signals are synchronized with the global WRITE signal. The state of the phase change device is changed only when the V_{INPUT} signal is active. In the READ state, when the conductance of the phase change device increases over the firing threshold and the voltage V_{PCD} falls below the threshold θ_V , the firing event occurs and is stored in the latch. In the subsequent period of the global WRITE signal, the actual spiking occurs (SPIKE signal is active). The same SPIKE signal is used for resetting the neuron. The time evolution of typical signal waveforms is shown in Supplementary Figure 10.

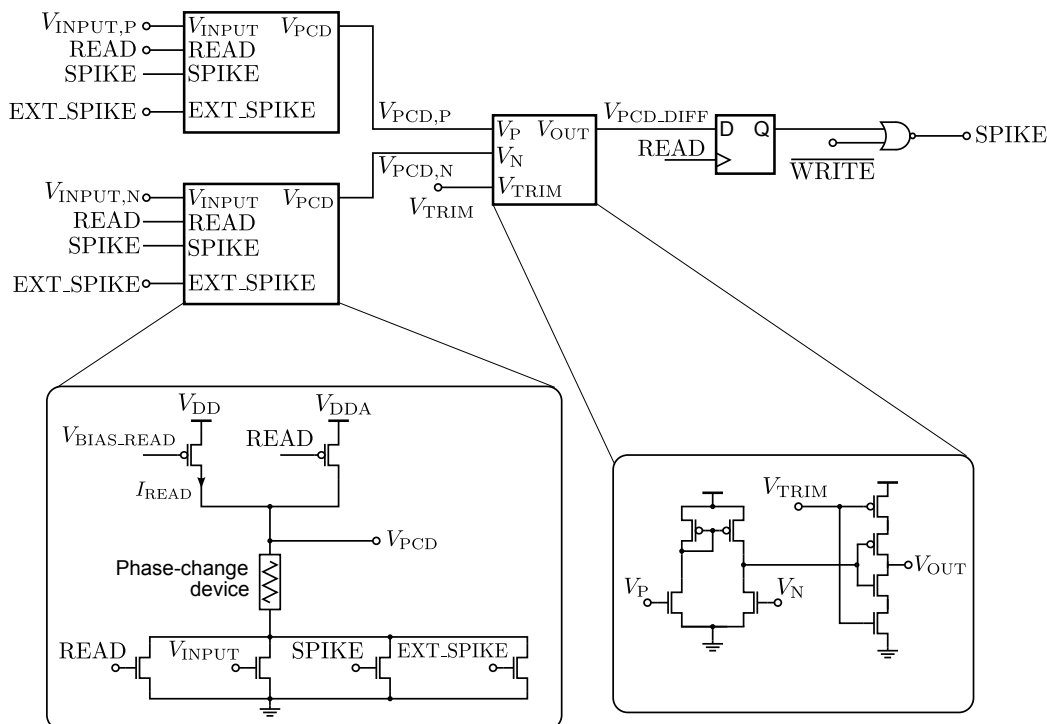
To estimate the areal requirements of the circuitry, we have performed a simplified layout and a simulation (Cadence, California, US) on a 14 nm technology node using a mushroom-type phase change device. With standard layout rules used, the estimated area required for the circuit realization is approximately $0.5\text{-}1\ \mu\text{m}^2$. Such area is orders of magnitude smaller than it would be in a capacitor-based implementation. For example, a 1 pF NFET capacitor would typically consume more than $50\ \mu\text{m}^2$, about $100\times$ more than the circuit presented. The power consumed in the circuit is approximately less than $10\times$ the power needed to update the phase change device. More functionality or functionality that cannot be supported by the phase change device natively can be added at the expense of larger circuit area and increased power/energy consumption.



Supplementary Figure 11: Generation of neuronal input and arbitration of PSPs in an artificial neuron

When integrated in a network of neurons and synapses, the basic neuron circuit depicted in Supplementary Figure 9 can be extended in various ways. One useful extension is the generation of the neuronal input signal V_{INPUT} based on the total instantaneous PSP arriving through a plurality of synapse-driven dendrites. Supplementary Figure 11 shows an example of such extension in which the translation be-


tween a presynaptic spike and a PSP signal is performed locally in each synapse. To enable efficient updates of the membrane potential, the PSP takes form of a pulse-width-modulated voltage signal that is representative of the internal weight of the synapse w_i . Subsequently, the arriving PSPs are arbitrated in a current-mirror-based circuit that applies an OR-like function to simultaneously arriving spikes. The output of the arbitration circuit is used as V_{INPUT} in the rest of the neuronal circuit. In this example, a key advantage is small complexity because only a single transistor per dendrite is required in the neuron. Where needed, a more complex arbitration circuit implementing the “summation” operation may be used instead.



Supplementary Figure 12: Circuit realization of a phase change-based neuron using a couple of paired phase-change devices in a differential configuration

Another important extension is the use of two paired phase change devices in a differential configuration to enable both excitatory and inhibitory neuronal inputs (see also Supplementary Note 2). Supplementary Figure 12 shows a schematic of the corresponding circuit realization. Here, each of the two devices (the “positive” device and the “negative” device) is embodied within a corresponding basic circuit as shown in the bottom left inset of Supplementary Figure 12. The excitatory inputs are applied at $V_{INPUT,P}$ and the inhibitory inputs at $V_{INPUT,N}$. The output voltages, $V_{PCD,P}$ for the positive device and $V_{PCD,N}$ for the negative device, respectively, are input into a current-mirror-based difference circuit (bottom right inset of Supplementary Figure 12). Using the V_{TRIM} voltage, the threshold of the last inverter and hence of the entire difference circuit can be set. Depending on the implementation, the non-ideal characteristics of the current mirror might be exploited to automatically reset both devices when they approach their maximum conductance without using the additional signal EXT_SPIKE (see also Supplementary Figure 8b. in Supplementary Note 2).

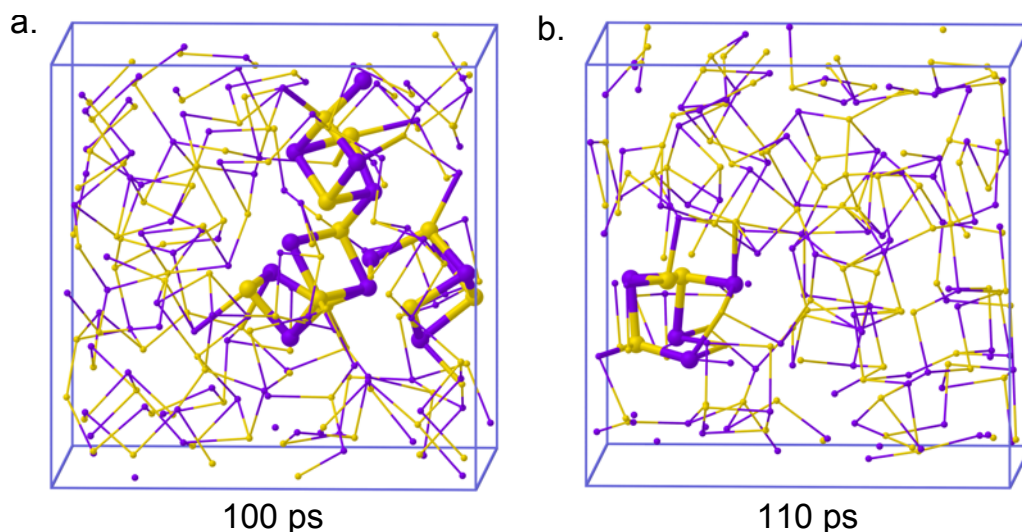
Typical signal waveforms in the two-device circuit realization are shown in Supplementary Figure 13. In the example shown, both excitatory inputs ($V_{INPUT,P}$) and inhibitory inputs ($V_{INPUT,N}$) are



Supplementary Figure 13: Typical signal waveforms in a circuit realization of the phase change neuron using two phase change devices in a differential configuration

applied during the first and the third WRITE cycle. This causes a drop in the corresponding voltages $V_{PCD,P}$ and $V_{PCD,N}$, respectively. Since the excitatory input is stronger than the inhibitory input (the width of the pulses corresponding to the excitatory input is larger), the voltage in the difference circuit $V_{PCD,DIFF}$ drops (the neuron membrane potential increases). However, the increase is not sufficient to trigger the spike event. The neuron fires only after an additional strongly excitatory input in the fourth write cycle when $V_{PCD,DIFF}$ drops below the threshold voltage θ_V . In the subsequent WRITE cycle, the SPIKE event is generated and the neuron is reset back to the resting potential.

4 Supplementary Note: Evidence of randomness in atomic configurations of the reset state from molecular dynamics (MD) simulations



Supplementary Figure 14: Evidence of randomness from MD simulations. Amorphous GeTe structures: Ge and Te atoms are depicted as dark yellow and violet spheres, respectively. Bonds between these atoms are shown as cylinders. Panel (a) shows a configuration taken after 100 ps; panel (b) shows a configuration taken after 110 ps.

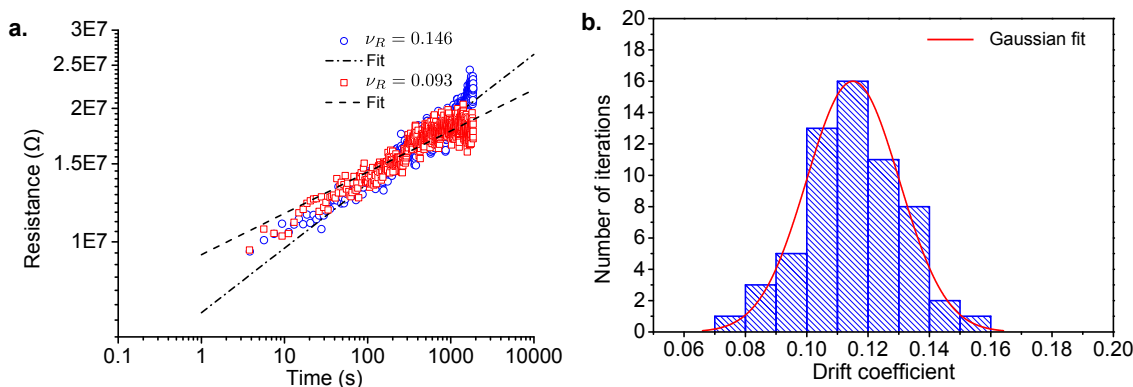
For every reset pulse applied, the amorphous region created via the melt-quench process is likely to have a slightly different atomic configuration. The high mobility of the atoms in the molten state ensures that even the slightest variations in the initial conditions or pulse characteristics will result in the formation of a new glass state during the glass transition. We performed extensive molecular dynamics (MD) simulations to test this hypothesis. Classical MD simulations on amorphous GeTe phase change material were used to generate structures by quenching at short time scales. The interactions between Ge and Te atoms were modeled using the Tersoff-based potential that Zipoli and Curioni recently developed [4] for GeTe and that was successfully validated against experimental data as well as density-functional theory (DFT) calculations. Liquid GeTe was equilibrated at 1000 K for 100 ps via a Langevin thermostat using the classical potential. Thereafter the thermostat was switched off and two configurations were taken after 100 and 110 ps from constant-energy MD (NVE). Subsequently, the two configurations were quenched by ramping down the temperature, controlled via a Berendsen thermostat, linearly from 1000 K to 300 K over 1 ns time duration.

The resulting two amorphous geometries were optimized using DFT at the PBE level [5]. Wannier centers were used to identify bonds between the atoms [4]. From Supplementary Figure 14, it is clear that the atomic configurations of these two amorphous states differ substantially, even though they are created in an identical manner except that in one case the system spent an additional 10 ps in the liquid state. Experimental observation of the variability in the structural relaxation process inherent in these highly stressed amorphous states provides further evidence of this hypothesis (see Supplementary Note 5). These amorphous regions are also likely to have different distributions of crystalline nuclei [6]. Given the small number of atoms used in the MD simulations, it is difficult to visualize the crystal nuclei.

However, the highlighted bond configurations in Supplementary Figure 14 can be assumed to be the seeds for crystal growth, as these cubes are the building blocks of the crystalline phase of GeTe. Here again we can see a clear difference in the number and distribution of these bond configurations. Also note that, unlike in GeTe, in GST the crystallization mechanism is nucleation-dominated and hence a significantly higher number of quenched-in crystalline nuclei will be present in the amorphous matrix of melt-quenched GST.

5 Supplementary Note: Evidence of randomness in atomic configurations of the reset state from structural relaxation

As the molten phase change material in a phase change device is quenched rapidly, the atomic configurations are frozen into the less ordered amorphous state, which is also highly stressed. The subsequent structural relaxation that occurs in such a glass is assumed to be responsible for the temporal evolution of the low-field resistance of the device. At constant ambient temperature, the resistance exhibits a temporal dependence characterized by $R(t) = R(t_0)(t/t_0)^{\nu_R}$, where $R(t_0)$ is the resistance measured at time t_0 [7, 8]. Experimental measurements show that a device exhibits slightly different resistance-drift behavior even for comparable reset states. This variability in the structural relaxation is indicative of the different atomic configurations achieved during the glass transition.

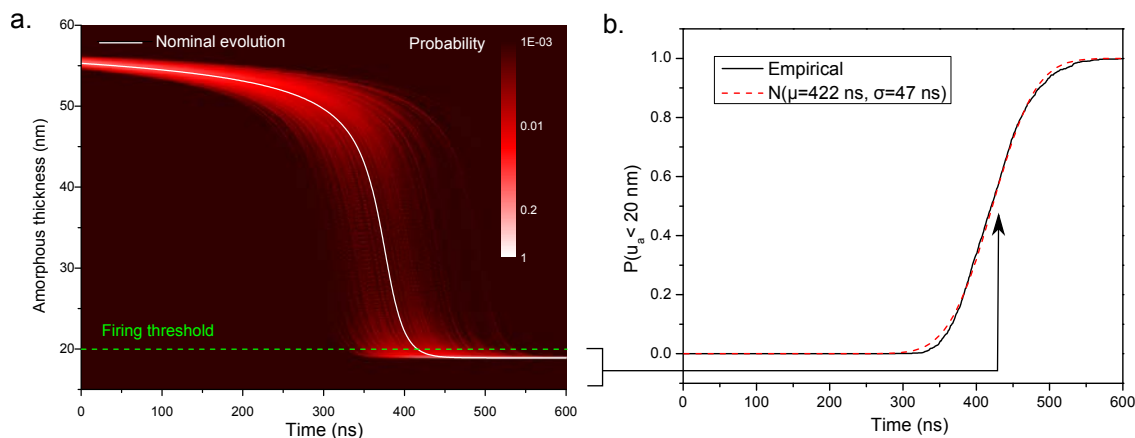


Supplementary Figure 15: Evidence of randomness in the atomic configurations of the reset state.

Supplementary Figure 15a shows the temporal resistance-drift values and the associated ν_R corresponding to two such reset states. After each reset operation, the resistance is monitored for 2000 s. The drift coefficient ν_R is derived from this experimental data. Supplementary Figure 15b shows a distribution of the ν_R values obtained for 60 such iterations. It exhibits a mean of 0.12 and standard deviation of 0.015.

6 Supplementary Note: Stochastic dynamics in phase change neurons

In phase change neurons, evolution of the membrane potential is inherently stochastic, and so are the resulting firing dynamics of the neuron. In Supplementary Notes 4 and 5, we argued that a major contributor to this stochastic behavior is the fact that the structure of the amorphous dome is reconfigured in a different way each time the neuron fires and the phase change device is reset. In this supplementary note, we discuss the implications that this has on phenomenological models of the neuronal dynamics and in particular, on the effective interspike intervals observed experimentally.

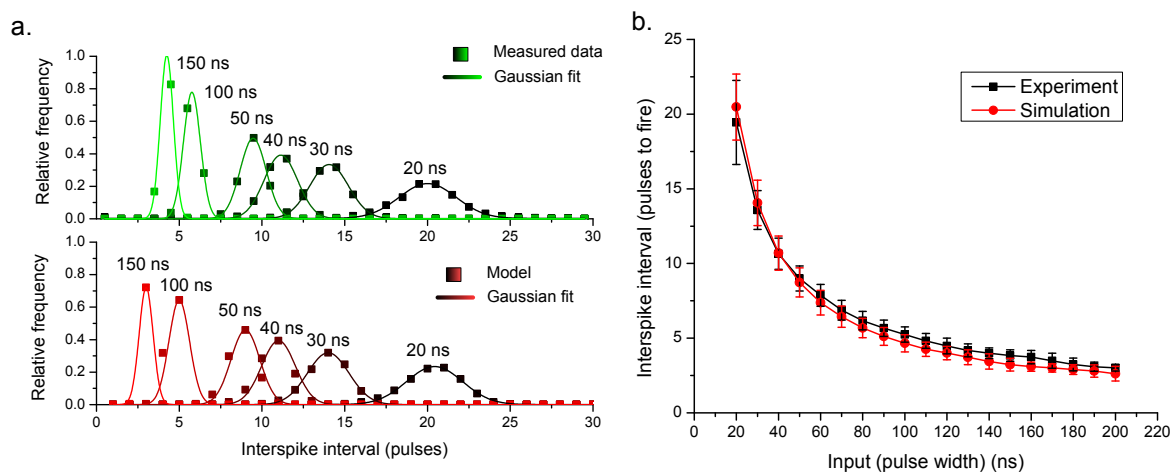


Supplementary Figure 16: Evolution of the membrane potential trajectory. In (a), distribution of $u_a(t)$ is shown assuming a random Gaussian distribution of the initial conditions. In (b), the probability of $u_a(t)$ being below the spiking threshold of $\theta = 20 \text{ nm}$ is shown.

Besides variations in the amorphous thickness $u_a(0)$, also the atomic configuration of the amorphous dome including the distribution of crystalline nuclei differ after each reset event. Hence, the effective growth velocity as well as the effective area of the crystalline/amorphous interface are subject to local variations during the evolution of the crystal front. In a phenomenological model, the effect of these random variations on the evolution rate of the amorphous thickness can be captured by adding a noise term to the crystal growth equation in addition to the randomness associated with $u_a(0)$:

$$du_a/dt = -v_g(R_{\text{th}}(u_a)P_p + T_{\text{amb}}) + \xi(t), u_a(0) = U_0 \quad (13)$$

where U_0 is a random variable whose distribution depends on the reset pulse applied to the phase change device. In such model, the evolution of the amorphous thickness is viewed as a Langevinian noisy relaxation process. The macroscopic (slow) variable is the amorphous thickness u_a and the microscopic (fast) variable is the stochastic noise term $\xi(t)$. This approach is conceptually similar to the conventional models of noisy integrate-and-fire neurons based on the Ornstein-Uhlenbeck process with the difference that Equation 13 is intrinsically nonlinear in u_a which affects the resulting trajectory and the firing characteristics of the neuron. While simulations based on Equation 13 are promising and capture the experimentally observed interspike intervals (data not shown), the physical underpinnings of such a model are yet to be established. In particular, a rigorous experimental study is needed to quantitatively distinguish the stochasticity arising from variations in the initial conditions $u_a(0)$ and from the $\xi(t)$ noise term.



Supplementary Figure 17: Distribution of interspike intervals. In (a), the experimentally obtained distribution of ISI is compared to the distribution obtained using the model based on randomization of the initial conditions. In (b), the mean ISI and the standard deviation for the range of inputs from 20 ns to 200 ns are compared between the experiment and the simulation.

A simplified approach that captures the observed neuronal dynamics, albeit physically less motivated, is to model both the variations due to U_0 and the randomness associated with the atomic configuration collectively by variations in an *effective* initial amorphous thickness, $u_a(0)$. Here, the actual value of $u_a(0)$ is viewed as a realization of a Gaussian random variable, $u_a(0) \sim N(\mu_0, \sigma_0)$ whose mean and variance are determined by the characteristics of the reset pulse. With $u_a(0)$ being stochastic we account collectively for variations in the internal structure across the entire amorphous dome, and the neuronal dynamics evolve deterministically thereafter,

$$du_a/dt = -v_g (R_{th}(u_a)P_p + T_{amb}), u_a(0) \sim N(\mu_0, \sigma_0). \tag{14}$$

Supplementary Figure 16a compares a noise-free trajectory of u_a (shown by a white line) to the trajectory with normally distributed initial conditions with mean of $\mu_0 = 55.3$ nm, standard deviation $\sigma_0 = 0.3$ nm and under fixed-power input of 280 μ W (shown in shades of red). We see that due to the nonlinear character of Equation 14, the spread in the initial amorphous thickness has a significant effect on the trajectory. Importantly, the distribution of u_a “along” the spike threshold ($\theta = 20$ nm) is approximately normal. Supplementary Figure 16b shows that the probability of u_a dropping below the firing threshold can be approximated with the normal cumulative distribution function with $\hat{\mu} = 422$ ns and $\hat{\sigma} = 47$ ns.

An approximate distribution of the interspike intervals (ISI) can be derived based on the knowledge of the dynamical equations, or of the empirical distribution of the (stochastic) trajectory u_a . We denote T_p a fixed pulse width and $N(T_p)$ the random variable that describes the number of pulses of width T_p that cause the neuron to fire, starting from the reset state. To know the ISI distribution, it is sufficient to compute $\Pr(N(T_p) = k)$ where k denotes the number of pulses. We will demonstrate that such a distribution matching our experimental measurements can be obtained analytically.

The probability $\Pr(N(T_p) = k)$ is given by the joint probability distribution

$$\Pr(N(T_p) = k) = \Pr(u_a(0) > \theta, u_a(T_p) > \theta, \dots, u_a((k-1)T_p) > \theta, u_a(kT_p) < \theta) \tag{15}$$

which is by approximation

$$\Pr(N(T_p) = k) \approx \Pr(u_a(kT_p) < \theta) \prod_{j=0}^{k-1} \Pr(u_a(jT_p) > \theta) \quad (16)$$

The individual probabilities can be computed analytically from the dynamical equations or empirically obtained from the $u_a(t)$ trajectory. When using Equation 14, $\Pr(u_a(t) < \theta)$ can be approximated by

$$\Pr(u_a(t) < \theta) \approx \Phi\left(\frac{t - \hat{\mu}}{\hat{\sigma}}\right) \quad (17)$$

Hence, referring to Equations 14 and 17, the ISI distribution can be obtained analytically as

$$\Pr(N(T_p) = k) \approx \Phi\left(\frac{kT_p - \hat{\mu}}{\hat{\sigma}}\right) \prod_{j=0}^{k-1} \left[1 - \Phi\left(\frac{jT_p - \hat{\mu}}{\hat{\sigma}}\right)\right] \quad (18)$$

Supplementary Figure 16a compares the experimentally obtained ISI distribution (cf. Figure 4a of the main text) with that obtained analytically using Equation 18). Given the actual complexity of the crystal growth process in a real phase change device and the simplicity of the model, this match is remarkably accurate. Supplementary Figure 16b shows the ISI distribution obtained experimentally and in simulation, respectively, for the entire range of inputs used in our experiments.

7 Supplementary Note: Modeling the dynamics of populations of stochastic phase change neurons

Populations of stochastic phase change neurons exhibit fairly complex dynamics that are based on the stochastic response of individual neurons as well as on the inter-neuron variability within the population. To predict the aggregate performance of the entire population in representing the input stimulus, we introduce a high-level population model based on noisy integrate-and-fire neurons. The neuron models are tailored to the essential response characteristics observed in phase change neurons.

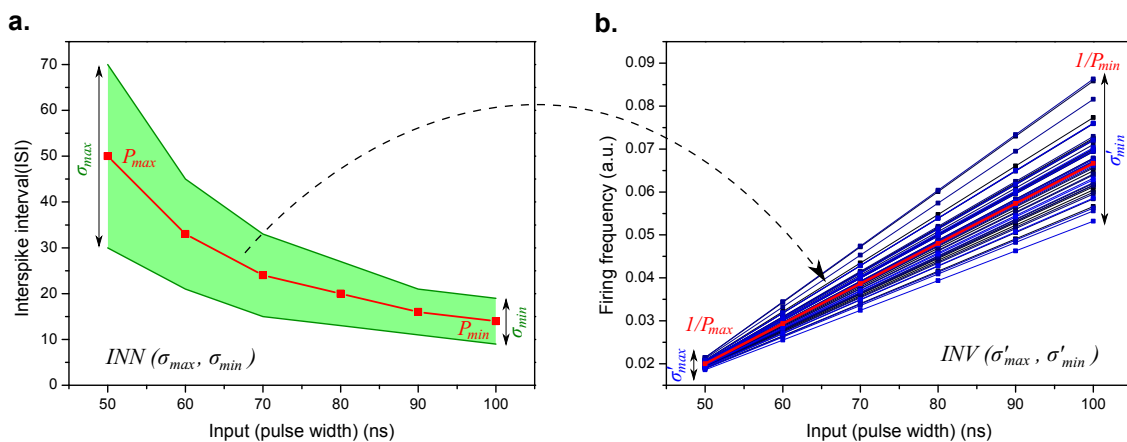
In a single neuron, the interspike interval (ISI) at constant input I can be modeled as a stochastic process with normal distribution $N(P(I), \sigma(I))$, where $P(I)$ is the mean ISI for input I and $\sigma(I)$ is the standard deviation for input I . A Gaussian distribution of ISI can be observed experimentally (cf. Figure 4(a) of the main text), and methods for model-based estimation of the ISI distribution parameters can be found in Supplementary Note 6. Based on the dynamics of the crystal growth in the phase change material, we can let $P(I)$ evolve as $1/f(I)$ (cf. the fit in Figure 4(a)), where $f = \beta + \alpha I$ is a linear function of I . In this way, the product of $P(I)f(I) = \text{const}$. Specifically, we set

$$P(I) = \frac{1}{\beta_P + \alpha_P \times I}, \tag{19}$$

where $\beta_P = 1/P_{\text{max}}$ and $\alpha_P = 1/P_{\text{min}} - 1/P_{\text{max}}$. Similarly, the standard deviation is

$$\sigma(I) = \frac{1}{\beta_\sigma + \alpha_\sigma \times I}, \tag{20}$$

where $\beta_\sigma = 1/\sigma_{\text{max}}$ and $\alpha_\sigma = 1/\sigma_{\text{min}} - 1/\sigma_{\text{max}}$. The membrane potential of the neuron, V , is set to the resting potential of $V_r = 0$ and is incremented by $1/n(I)$, where $n(I)$ is drawn from $N(P(I), \sigma(I))$ whenever input I is passed to the neuron. Upon reaching $V = 1$, the neuron spikes, and the membrane potential is reset to the resting potential.



Supplementary Figure 18: Models of a single neuron and a neuron population. Panel **a** shows the simulated response of a single noisy integrate-and-fire neuron with stochasticity $INN(\sigma_{\text{max}}, \sigma_{\text{min}})$ (cf. the experimental data presented in Figure 4(a) of the main text). Panel **b** shows a population of 50 stochastic neurons with inter-neuron variability $INV(\sigma'_{\text{max}}, \sigma'_{\text{min}})$.

Supplementary Figure 18a illustrates the mean ISI (in red) and the standard deviation (in green) obtained by using the above noisy integrate-and-fire model with $P_{\text{max}} = 50$, $P_{\text{min}} = 15$, $\sigma_{\text{max}} = 20$ and

$\sigma_{\min} = 5$. For convenience, we denote the stochastic distribution parameters for a single neuron with fixed mean ISI as $INN(\sigma_{\max}, \sigma_{\min})$.

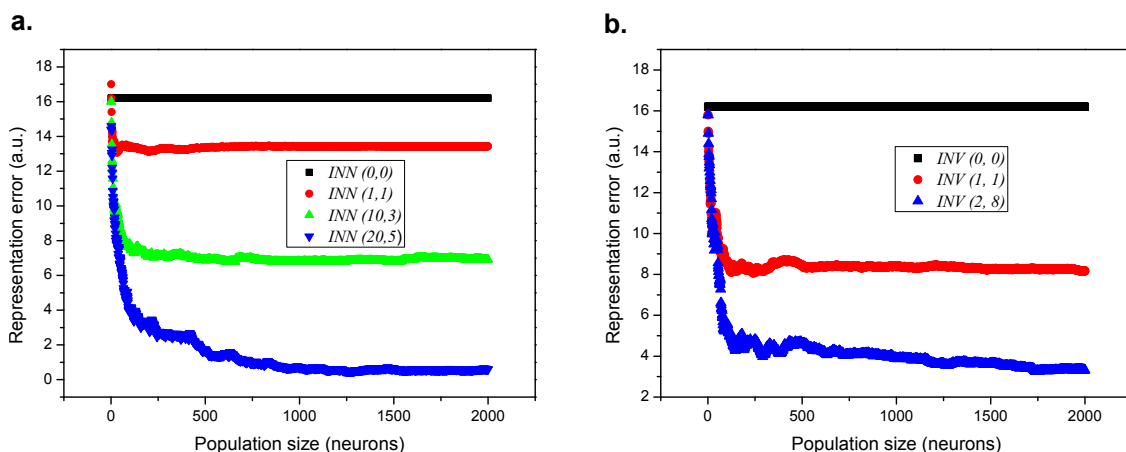
In a population of neurons, the mean ISI of each neuron varies. We model this by drawing P_{\max} and P_{\min} for each neuron from a normal distribution, namely,

$$P_{\max} \sim N(P_{\max}^0, \sigma'_{\max}) \quad (21)$$

$$P_{\min} \sim N(P_{\min}^0, \sigma'_{\min}). \quad (22)$$

For convenience, we denote the stochastic distribution parameters for a population with fixed mean ISI as $INV(\sigma'_{\max}, \sigma'_{\min})$.

Supplementary Figure 18b illustrates the frequency response for a population of 50 stochastic neurons with inter-neuron variability $INV(2, 8)$ in which each individual neuron has the ISI distribution shown in Figure 18a. We see that the average frequency response of each neuron (shown in blue) is the inverse of the mean ISI for the corresponding input; the red line shows the average frequency response of the population. Note that in the frequency domain, the distribution of the spiking frequencies across the population for fixed I is a reciprocal normal distribution and as such does not have its theoretical moments defined. However, for truncated distributions with $P_{\max}, P_{\min} \gg 0$, the mean and standard deviation of the population response at any I can be approximated as $1/P(I)$ and $\sigma'/(P(I)(P(I) + \sigma'))$.



Supplementary Figure 19: Effect of intra-neuron noise (INN) and inter-neuron variability (INV) on the representation error. Panel **a** shows the representation error for four populations of stochastic neurons with no inter-neuron variability and different levels of intra-neuron stochasticity. Panel **b** shows the representation error for three populations of stochastic neurons with no intra-neuron stochasticity and different levels of inter-neuron variability.

The population model allows us to understand the fundamental performance limits in the representation of an input stimulus, in particular the effects of the intra-neuron stochasticity, the inter-neuron variability and the population size on the representation error. In what follows, a triangular waveform input stimulus of fundamental frequency $f_B = 10$ kHz is applied to neurons with $P_{\max} = 50$ and $P_{\min} = 15$. With sampling frequency $f_S = 100$ kHz, the equivalent average firing frequencies range between 2 kHz and 6.6 kHz and are below $2f_B = 20$ kHz.

Supplementary Figure 19a shows the representation error, measured as the 2-norm of the difference between the input stimulus and the population code, as function of the population size and the intra-neuron stochasticity (INN) when all neurons in the population have equal characteristics (INV =

(0,0)). We see that the performance is worst for fully deterministic neurons INN(0,0) (black squares), and does not improve with the number of neurons as there is no stochasticity in or between the neurons. By increasing the intra-neuron stochasticity, the representation error can be decreased dramatically, although there is no variability in the average response of individual neurons. The error improves with larger standard deviation, but to reach the best achievable error floor also a larger population is required. The performance can be improved by randomizing the initial conditions V at the start of the algorithm; however, the effect of this improvement becomes relatively negligible with increasing INN.

Supplementary Figure 19b shows the representation error as function of the population size and inter-neuron variability (INV) when all neurons in the population are deterministic. Similarly to the result for the INN, the performance is worst when there is no variability between the deterministic neurons. With increasing INV, the representation error can be reduced. A similar trend is observed in populations of stochastic neurons, however, the effect is smaller for populations with large INNs.

References

- [1] Kuzum, D., Jeyasingh, R. G. D., Lee, B. & Wong, H. S. P. Nanoelectronic programmable synapses based on phase change materials for brain-inspired computing. *Nano Lett.* **12**, 2179–2186 (2012).
- [2] Suri, M. *et al.* Phase change memory as synapse for ultra-dense neuromorphic systems: Application to complex visual pattern extraction. *2011 Int. Electron Devices Meet.* 4.4.1–4.4.4 (2011).
- [3] Burr, G. W. *et al.* Experimental demonstration and tolerancing of a large-scale neural network (165 000 synapses) using phase-change memory as the synaptic weight element. *IEEE Trans. Electron Devices* **62**, 3498–3507 (2015).
- [4] Zipoli, F. & Curioni, A. Reactive potential for the study of phase-change materials: GeTe. *New J. Phys.* **15** (2013).
- [5] Perdew, J. P. J., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
- [6] Lee, B. S. *et al.* Distribution of nanoscale nuclei in the amorphous dome of a phase change random access memory. *Appl. Phys. Lett.* **104** (2014).
- [7] Ielmini, D., Sharma, D., Lavizzari, S. & Lacaíta, A. L. Reliability Impact of Chalcogenide-Structure Relaxation in Phase-Change Memory (PCM) Cells Part I : Experimental Study. *IEEE Trans. Electron Devices* **56**, 1070–1077 (2009).
- [8] Sebastian, A., Krebs, D., Le Gallo, M., Pozidis, H. & Eleftheriou, E. A collective relaxation model for resistance drift in phase change memory cells. *IEEE Int. Reliab. Phys. Symp. Proc.* MY.5.1–MY.5.6 (2015).