Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Projected Disc-type Phase-change Memory for Analog In-Memory Computing

Ghazi Sarwat Syed

ghs@zurich.ibm.com

IBM Research - Zurich https://orcid.org/0000-0002-2054-4887

Loris Coccia

IBM Research https://orcid.org/0009-0005-2515-1504

Vara Jonnalagadda

IBM Research - Zurich

Antonio Mio

CNR-IIM

Asit Ray

IBM Research

Matthew Breitwisch

IBM

Abu Sebastian

IBM Research Europe https://orcid.org/0000-0001-5603-5243

Article

Keywords: Phase Change Materials, Memory Devices, In-Memory Computing

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11453571/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Projected Disc-type Phase-change Memory for Analog In-Memory Computing

Ghazi Sarwat Syed,[1] Loris Coccia,[1] Vara Prasad Jonnalagadda,[1] Antonio Massimiliano Mio,[2] Asit Ray,[3] Matthew BrightSky,[3] and Abu Sebastian[1]

[1)] *IBM Research – Europe, Säumerstrasse 4, 8803 Rüschlikon, Switzerland*

[2)] *Istituto di Microelettronica e Microsistemi, Consiglio Nazionale delle Ricerche, strada VIII 5, I-95121, Catania, Italy*

[3)] *IBM Research– Yorktown Heights, 10598 NY, USA*

Resistive memory technologies offer a compelling advantage for in-memory computing. However, realizing a device architecture that simultaneously achieves high computational precision, efficiency, and density has remained elusive due to inherent trade-offs among these performance metrics. Here, we introduce a compact phase-change memory device architecture that combines an ultra-confined active switching volume for enhanced electro-thermal efficiency with a non-insulating thin film that suppresses temporal conductance fluctuations. We analytically model and back-end integrate these devices into crossbar arrays. Even with conventional, undoped phase-change materials, the device architecture enables a computational precision approaching 6 bits, low conductance values below $50\ \mu S$ with an adequate conductance window, and a viable pathway toward sub-$100\ \mu A$ programming currents under nominal operating voltages.

Keywords: Phase Change Materials, Memory Devices, In-Memory Computing

INTRODUCTION

Phase-change memory (PCM) has traditionally been developed for storage class memory and embedded memory applications[1,2]. Recent research efforts, however, have shown significant potential of this technology for analogue in-memory computing (IMC) particularly for performing deep neural network inference[3–5]. In deep neural networks, matrix-vector multiplication (MVM) operations dominate computation and can be implemented using crossbar arrays of synaptic unit cells, each containing one or more PCM devices alongside selector devices such as field-effect transistors[6,7]. Analogue IMC with PCM offers the promise of unprecedented compute density and weight capacity compared to approaches based on static random access memory. Realizing this potential, however, requires improvements in the integration density of PCM devices, which hinges on reducing the write or programming current that determines the size of selector devices and, consequently, unit cells. Another essential device-level feature is maintaining low conductance across all PCM phase configurations to reduce energy consumption during MVM operations and to mitigate non-idealities, such as parasitic voltage drops due to crossbar wire resistance. Additionally, intrinsic temporal non-idealities, specifically conductance drift and read noise, limit the achievable numerical precision during programming and computation. Suppressing these non-idealities is therefore critical for high-precision IMC.

In efforts to enhance PCM devices for analogue IMC, much research has concentrated on optimizing the composition of phase-change materials[8–10]. More recently, significant progress has been made in exploring phase-change superlattices[11–13]. Although such approaches have demonstrated promise in micron-scale devices, the most effective and industrially proven strategy—one that accounts for process integration compatibility—remains the reduction of the material volume that undergoes the phase transition, thereby lowering the required write current[14]. However, analogue IMC requires a continuum of conductance states at low absolute conductance values. Simply reducing the volume of phase-change material by reducing the span of phase change channel to achieve can degrade analogue behavior, as well as increase fabrication complexity, and—at extreme scales—promote elemental segregation under strong geometrical confinement. On the other hand, to suppress the non-idealities, a promising alternative is to project the amorphous volume of the phase-change material onto a non-insulating projection layer (liner), which effectively dictates the read-out conductance[15,16]. Projection efficiency is to be highest when the liner is fully and uniformly aligned in parallel with the active phase-change material. This configuration is difficult to achieve in conventional device architectures, such as mushroom-type vertical devices[17], which are primarily optimized for high integration density rather than optimal projection alignment.

In this article, we introduce and demonstrate a phase-change computational memory concept that elegantly fulfills both requirements. While retaining a compact vertical geometry, our design minimizes the active phase-change volume by reducing both the phase-change material thickness and its contact area with the bottom electrode[18], while

simultaneously enhancing projection through a continuous physical connection to a non-insulating projection liner. We characterize these devices on a back-end-integrated crossbar array with MOSFET selectors. The devices enable low programming currents and conductance values, suppresses temporal conductance fluctuations, and remains compatible with conventional fabrication processes, making it ideally suited for analogue IMC applications.

DEVICE CONCEPT

The proposed device architecture, illustrated in Figure 1A, consists of an ultra-thin phase-change material film sandwiched between a confined bottom electrode and a dielectric spacer. Between the bottom electrode and the phase-change layer lies a ultra-thin projection liner, which spans the full lateral extent of the device, corresponding to a radius of r_{pcm}. Electrical contact is made laterally via a top electrode that wraps around the edge of the phase-change film, enabling current to flow outward from the central bottom electrode. A short, high-current RESET pulse amorphizes a region of radius u_{a} of the crystalline phase-change material adjacent to the bottom electrode. The overall device conductance is determined by the geometry and the extent of the amorphous region given by its radius u_{a} (colored red), which is electrically connected in parallel with the projection liner. Geometrically, both the amorphous region and the liner take the form of cylindrical discs, benefiting from the physics of ultra-confinement. Furthermore, this configuration yields an inverse logarithmic dependence of conductance on $G_{\text{device}} \propto 1/\log(u_{\text{a}})$. Consequently, in this device concept the conductance is highly sensitive to small variations in the amorphous radius when u_{a} is close to r_{BE}, whereas this sensitivity decreases for larger amorphous radii. Moreover, the lowest achievable conductance state scales with the ratio $r_{\text{pcm}}/r_{\text{be}}$. Consequently, reducing the bottom-electrode radius permits a proportional reduction in device foot-print, which aligns naturally with advances in the lithographic scaling of minimum feature sizes.

The device in a RESET state can be modeled (see inset of Figure 1A) with the following conductive components: G_{amor} (conductance arising from the amorphous volume), G_{crys} (conductance from the crystalline volume), $G_{\text{liner,amor}}$ (conductance from the liner underneath the amorphous volume) and $G_{\text{liner,crys}}$ (conductance from the liner underneath the crystalline volume). Two additional conductance components determine the device functionality: $G_{\text{liner-pcm}}$ (contact resistance between the liner and the phase-change material) and $G_{\text{be-liner}}$ (contact resistance between the liner and the bottom electrode). The electrical conductivities of the phase-change and liner materials satisfy $\sigma_{\text{amor}} < \sigma_{\text{liner}} < \sigma_{\text{crys}}$, where the conductivity contrast is defined via $\sigma_{\text{liner}} = \alpha\,\sigma_{\text{crys}}$ for a tunable scaling parameter $\alpha \in (0,1)$. Together with the geometrical parameters (film thickness), this conductivity contrast enables optimization of the memory window and projection efficacy. In the limit of a small $G_{\text{be-liner}}$, and for arbitrary values of $G_{\text{liner-pcm}}$, the minimal achievable conductance approaches $G_{\text{be-liner}}$. We are able analytically describe these device characteristics by incorporating the relevant bulk and interfacial material properties, as well as geometric parameters. The model accurately captures both experimental observations and finite-element method (FEM) simulations (Extended Data Fig. 1). We find that the contact resistance between the liner and the phase-change material plays a dominant role in determining projection efficacy[19]. An increase in this resistance degrades analogue programmability and limits the correction of the amorphous-phase non-idealities, as will be discussed in a following section.

EXPERIMENTAL VALIDATION

Prototype devices were fabricated using the conventional, undoped phase-change material $Ge_2Sb_2Te_5$ (GST) with a thickness approaching 5 nm and an amorphous-carbon (aC) based liner of 3 nm. The devices were integrated with n-MOS transistors serving as selectors, with bottom electrodes formed from chemical mechanical polished metal nitride (Ti_xN_y) of r_{be} 20 nm. Additionally, a 10×8 array of these devices was designed and fabricated, incorporating designated word lines for transistor control, along with separate drain and source lines for programming and read-out operations (see Figure 1B-D). To further minimize the bottom-electrode to PCM contact area, devices with unplanarized bottom electrodes penetrating directly into the device stack were also fabricated.

We programmed the devices using successive RESET pulses with increasing amplitude, which melt-quench and amorphize the phase-change material, followed by SET pulses that crystallize the amorphous regions created by any of the RESET pulses. The device conductance was measured at a 0.1 V bias after each SET or RESET pulse, generating a programming curve based on the amplitude of the current programming pulse needed. Figure 2A shows four representative programming curves from the device, together with reference lines indicating the SET conductance and onset current of a conventional mushroom device with 80 nm GST. The device achieves a continuum of analogue states, with a reduction in the onset current, from approximately 500 μA (see Extended Data Fig. 2) in the mushroom-type device to 90 μA. Note that we define the onset current as the point at which the device conductance

begins to decrease. This point corresponds to the onset of melting and helps avoid ambiguity in resistance scaling behavior, which can depend on the device geometry. In addition to these improvements in programming, the devices also achieve over 10x lower SET conductance , while preserving a large programming window of more than 10x. These results validate that physical scaling of the device dimensions and contacts leads to strong electro-thermal confinement, resulting in more efficient Joule heating within the phase-change film and a reduction in conductance values.

We are able to accurately and reliably program the devices to target conductance states using a standard iterative programming scheme. Figure 2B illustrates the programming error associated with encoding synaptic weights. In this experiment, eighteen target conductance states (G_{Target}) are programmed ten times within a $\pm 2\%$ margin. The weight error is defined as $|G_{\text{Target}} - G|/G_{\text{max}}$, where G is the device conductance measured after convergence, and G_{max} is the maximum reliably programmable unit-cell conductance. Our ability to program within a narrow margin is enabled by the suppressed conductance fluctuations (read noise), which we will discuss shortly. Furthermore, we find that the programmed states exhibit excellent retention characteristics. The programmed conductance traces exhibit stable retention behavior across the measured conductance range, even at elevated temperatures (e.g., 85 °C, as shown in Figure 2C). Furthermore, we find that even with single shot (open-loop) programming, the RESET and SET conductance states are well separated across the crossbar array, as shown by the complementary cumulative RESET distribution and the cumulative SET distribution (see Figure 2D). Additionally, the RESET states show a characteristic crystallization behavior that depends on the amplitude and duration of the SET pulse. To demonstrate this, we applied box-type SET pulses with fixed falling and trailing edges of 8 ns, but varying amplitude and width (see Figure 2E). The color map shows the device conductance as a function of the pulse parameters, revealing a SET speed of approximately 900 ns. Figure 2F demonstrates reliable programming of the device under a burst of RESET pulses over 1 million cycles. This endurance exceeds the requirements for IMC applications, where frequent reprogramming is unnecessary. Notably, it represents a $100\times$ improvement over undoped-GST compositions, which we attribute to the physical confinement of the phase-change material. This confinement likely suppresses elemental segregation mechanisms.

Figure 2G–H presents the characterization measurements across multiple devices in the crossbar array. The left panel of Figure 2G shows a heat map of the dynamic on-resistance (R_{ON}) measured across 64 devices. A higher R_{ON} corresponds to high Joule heating efficiency. Notably, we find that compared to the mushroom-type device, R_{ON} is improved by $\sim$10x. The right panel shows the corresponding statistical distribution in R_{ON}, highlighting the device-to-device variability. Similarly, the left panel of Figure 2H illustrates the heat map of the programming onset current (I_{Prog}). Notably, the higher R_{ON} values leads to lower I_{Prog} values in the devices, in part explaining the improved programming efficiencies. The right panel presents I_{Prog} distribution across the same 64 devices. Note that both parameters exhibit a tight distribution. State-dependent threshold voltages for the SET transition were similarly characterized across the array (see Extended Data Fig. 3), featuring forming-free operation at sub-2V voltages.

We evaluated additional critical parameters for IMC, namely the temporal conductance variations caused by conductance drift and read noise. Drift refers to the progressive decrease in conductance within the RESET state, while read noise constitutes the random, typically $1/f$ fluctuations superimposed on the resting conductance value. Together, drift and read noise determine the weight error and, consequently, the accuracy of MVM operations. Figure 3A shows the conductance versus time for multiple non-volatile states in a device. By fitting these data with the standard drift model, we extract the effective state-dependent drift coefficients (ν_{eff}). Evidently, conductance drift is significantly suppressed in these devices. For RESET states, the drift coefficients are reduced by up to a factor of 20 relative to mushroom-type GST devices, highlighting the effectiveness of the projected-cell design. Similarly, for intermediate states, the devices exhibit drift coefficients lower than those of doped GST. The latter is a current state-of-the-art material engineered to suppress conductance values to the scale we otherwise achieve through our device geometry. Figure 3B shows a comparison of the state-dependent drift coefficients across different conductance states. For this analysis, each target conductance state (G_{T}) was programmed using the iterative programming scheme, and its conductance versus time behavior was fitted over ten separate trials. We can accurately model the state-dependent conductance of our devices using the analytical description described earlier. This state dependency originates from the underlying phase configurations. For larger values of u_{a}, the liner begins to dominate the resistor network, leading to decreasing ν_{eff}. As with improvements in the drift behavior, we find that our good projection efficacy significantly suppresses temporal fluctuations, leading to a reduction in read noise (σ_{G}), which was extracted by subtracting the model-fitted drifting conductance from each individual measurement and computing the standard deviation of the residual noise. Figure 3C plots the state-dependent signal-to-noise ratio (SNR), defined as $\sigma_{\text{G}}/G_{\text{T}}$, for various target conductance states that we have iteratively programmed. Notably, the SNR improves by a factor of $4\times$ compared to to mushroom-type GST devices.

In Figure 3D, we plot the heat map of the drift coefficients in the full RESET states of 64 devices in the crossbar array. Notably, the drift coefficient are suppressed in all devices. Conductance change due to ambient temperature fluctuations is another important device state-variable that impacts temporal stability of the states. This effect is more prevalent in the RESET states, due to the non-trivial activation energies (E_a) for charge transport in the amorphous volume. For $\sim$5 nm GST we measure $E_{a,amor} = 0.29\,eV$. Projection is expected to suppress the temperature dependence of the device, with the effective activation energy of the device, $E_{a,eff}$, approaching that of the liner (i.e., $E_{a,liner}, 0.10\,eV$). To verify this, we measured $E_{a,eff}$ across 27 randomly selected devices in the crossbar array. For each device, we created a RESET state and measured its conductance at different ambient temperatures. The resulting data were fitted to the Arrhenius equation, that scales conductance exponential to E_a. Figure 3E illustrates the distribution of $E_{a,eff}$, evidently showing a mean value close to $E_{a,liner}$.

Collectively, these results indicate that these devices can improve computational precision in analogue arithmetic operations. To validate this, we performed fully hardware-implemented MVM experiments. We first investigated the state-dependent contributions to computational accuracy. To this end, multiply–accumulate (MAC) operations were performed by programming all devices connected to a common bit-line to conductance states sampled from a uniformly distributed weight matrix W. A close loop iterative programming procedure with 1% error tolerance was used to map the weights in conductance states. A total of 250 independent measurements, each consisting of a constant voltage amplitude of 0.1 V vector, were applied to the word lines. The resulting device currents, representing the scalar multiplication operations, were accumulated along the shared bit-line. Note that this measurement approximates the pulse-width modulation scheme commonly used in AIMC, where input values are encoded as constant-amplitude voltage pulses. The MAC error was defined as $\frac{y_{ideal} - y_{experimental}}{y_{ideal}}$, which quantifies the linear error arising from the hardware implementation, including both the PCM devices and the associated interconnect network. The state-dependent MAC errors obtained from 250 independent measurements are shown in Figure 3F. Notably, conductance states exhibiting strong projection achieve computational precision exceeding 6 bits (also see Extended Fig. 4). In contrast, intermediate conductance states closer to the SET state exhibit larger errors. This behavior is consistent with the analytical model and is expected to be further mitigated through electrode engineering. Next, we evaluated complete MVM operations by programming the crossbar array with a uniformly sampled weight matrix W. The resulting computational precision is shown in Figure 3G. Notably, the hardware achieves more than 5 bits of effective computational precision. These results are further compared with an emulator incorporating the experimentally observed device non-idealities. Importantly, this level of precision represents a non-trivial improvement over state-of-the-art computational memory systems, which typically report 3–4 bits of effective computational accuracy[16].

DISCUSSION AND OUTLOOK

The geometry of our proposed device introduces an additional design parameter for the placement of the bottom electrode to minimize contact area. When the bottom electrode fully penetrates the phase-change material–liner stack, the contact geometry transitions from a planar circular interface to a cylindrical sidewall interface. Based on our devices parameters, this configuration increases the local power density at the contact, enhancing Joule heating efficiency, and raises the effective thermal resistance for heat dissipation, improving thermal confinement; thereby enhancing the programming efficiency. For the geometric dimensions of our devices, where the heater radius is comparable to the phase-change material thickness, this reduction in contact area approaches $1/2$. To validate this design strategy, with the same material stack, we fabricated and integrated devices in which the bottom electrode penetrates completely through the stack, establishing an edge-contact configuration similar to the contact made by the side electrode (see Figure 4A (i)). Figures 4B-C show HAADF-STEM micrographs of a cycled device in a RESET state, and in SET state, respectively. Crucially, we observe clear evidence of amorphized and recrystallized volumes in the devices, which is consistent with the melt-quench and crystallization behavior of GST, as also confirmed by the corresponding Fast-Fourier Transform (FFT) patterns. Also note that the structural integrity of the confined bottom electrode and the stoichiometry of the phase-change material remain intact. Atomic-resolution HAADF-STEM imaging of the phase change material channel in both the RESET and SET states (see Extended Data Fig. 4), reveals a highly ordered layered structure. Distinct pseudo van der Waals (vdW) gaps, appearing as dark horizontal lines parallel to the phase change channel, indicate a pronounced texture with preferential alignment of the crystalline grains. The measured distances between neighboring vdW gaps are approximately 1.05 nm and 1.35 nm, corresponding to the characteristic thicknesses of Sb_2Te_3 and $Ge_1Sb_2Te_4$ structural blocks, respectively. The observation of the same layered motifs in both conductance states suggests that the overall crystalline framework remains preserved during electrical switching. The structural models superimposed on the HAADF-STEM images corroborate the identification of these layered building blocks and their ordered stacking within the phase-change material channel (see Extended Fig. 5).

Moreover, as shown in Figure 4D, the programming efficiency itself is significantly improved. The programming current drops by nearly half to 50 μA, while low conductance states are reliably maintained across all programmable levels. For the 64 devices in the crossbar array, the inset presents a heat map highlighting the spatial distribution of the programming onset current. Figure 4E–F show the corresponding distributions of the drift coefficients and RESET conductance values. The extracted drift coefficients exhibit a median value of $\nu_{\text{eff}} = 0.006$, indicating limited temporal drift across devices. These results demonstrate that the characteristic signatures of good projection efficacy are consistently preserved in these device structures. Similarly, further decreasing the thickness of the GST film can also provide a means to more significantly decrease the device conductance[18]. Together, these avenues validate improvements that can be enabled to our proposed device concept.

We would also like to point out a specific feature of the projection is that it reduces the apparent memory window, typically defined as the resistance ratio $R_{\text{RESET}}/R_{\text{SET}}$. In IMC, however, the quantities of primary interest are conductances rather than resistances. The resistance-to-conductance mapping acts like a squashing function, already compressing large resistance values. Moreover, the number of reliably usable discrete levels scales with the conductance separation (not their ratio) and inversely with the degree of temporal fluctuations, i.e., it is proportional to $\frac{G_{\text{max}}-G_{\text{min}}}{\sigma_{\text{total}}}$, where σ_{total} represents the temporal conductance variability arising from both read-noise and conductance drift[20]. Consequently, variations in G_{min} alone have a less significant impact on the effective operational window for mapping synaptic weights. As an example, conventional mushroom-type PCM devices may exhibit memory windows exceeding 10^3, yet still accommodate fewer stable discrete levels than our proposed devices. This improvement arises primarily from the suppressed temporal variations achieved in our architecture[21]. As next steps, exploring alternative material compositions presents promising opportunities for performance enhancement, such as ultra-thin superlattice and heterostructure films that offer natural confinement properties, and GeSb, nanocomposite films for lower threshold voltages and improved retention[22,23]. For commercial integration, achieving wafer-scale film uniformity is essential, especially because the material properties are highly sensitive to thickness variations in sub-10 nm films. In this proof-of-concept study, we already demonstrate the fabrication feasibility using the industry-preferred sputter deposition method. In cases where tighter control over film thickness is required, especially for thicknesses approaching sub-3 nm, more conformal deposition techniques, such as atomic and chemical layer deposition may become required. These deposition schemes are already widely optimized in the fabrication of confined-type devices. We would also like to highlight that while some advanced materials and device structures show indications of similar improvements, such structures largely limit analogue capability due to the small material volume required for melting and/or pose challenges for large-scale integration with back-end-of-line processes.

CONCLUSIONS

In summary, we have introduced a PCM device concept for analogue in-memory computing that achieves low programming current, low conductance values, and high compute precision, while being well-suited for large-scale back-end-of-line integration. We have experimentally demonstrated this concept by fabricating prototype CMOS integrated crossbar array with the standard GST material, which can reversibly switch at programming currents around <100 μA while maintaining a conductance values below 50 μS, and suppressed drift and read noise, allowing more than 5-bit compute precision. Further enhancements may be possible by reducing the phase-change material thickness, minimizing the electrode contact area, and using optimized phase-change materials. These results mark a significant advancement toward enabling low-power, high-density resistive memory based in-memory computing.

MATERIALS AND METHODS

Device fabrication

Deposition of the PCM stack on 180 nm node n-MOS front-end-of-line wafers was initiated by in situ etching (ISE) of the heater electrode, drain-contact region, source, gate, and ground pads using Ar plasma for 1 min at 100 W. A 4 nm doped amorphous carbon (a-C) projection layer containing <20 sccm Ar dopant with H_2 was subsequently deposited, followed by the *in situ* deposition of a 5 nm $Ge_2Sb_2Te_5$ (GST225) layer by 50 W DC magnetron sputtering. Finally, a 6 nm SiO_2 capping layer was deposited by 600 W RF sputtering. The cell structures were defined in a e-beam lithography step. A 70 nm-thick negative-tone hydrogen silsesquioxane (HSQ) resist layer was spin-coated and exposed, and the patterns were transferred into the underlying layers by ion milling. Immediately after etching, and to minimize oxidation of the PCM sidewalls, the wafer was transferred to a sputtering system,

where the exposed sidewalls were cleaned using Ar plasma and subsequently encapsulated with a 50 nm W layer. For the bottom-heater devices, edge/top electrode structures were defined on the device in a fourth e-beam lithography step using negative-tone ARN 7520 resist. Following development, the patterns were transferred by reactive ion etching (RIE). For the top-heater devices, the devices were patterned adjacent to the transistor drain rather than directly on the drain. Following sidewall encapsulation with W, edge/top electrodes were defined on the PCM structures, and the devices were encapsulated with a 50 nm SiO_2 layer deposited by atomic layer deposition (ALD). Subsequently, vias to define the heater region on the PCM stack and to the transistor drain were opened using an additional e-beam lithography step with PMMA resist. After pattern transfer by RIE and resist removal, the wafer was transferred to a physical vapor deposition (PVD) chamber, where the exposed PCM sidewalls and drain-contact surfaces were cleaned using Ar plasma. The wafer was then transferred to the ALD system, and the vias were conformally filled with 60 nm TiN to form the top-heater structures. In a subsequent e-beam lithography step employing ARN 7520 resist, contact openings between the TiN heater and the drain electrode were defined and transferred by RIE, thereby electrically connecting the heater to the drain contact. Following the formation of both bottom-heater and top-heater structures, all devices were encapsulated with a 50 nm SiO_2 layer deposited by ALD. Contact vias to the edge/top electrodes, source, gate, and ground pads were opened using e-beam lithography with PMMA resist. The developed patterns were transferred by RIE. Prior to metallization, the exposed electrodes and contact pads were cleaned *in situ* using Ar plasma in the sputtering chamber. Subsequently, a 150 nm W layer was deposited by PVD to form the contact metallization. To define the 10×10 crossbar arrays and connect the corresponding gate, source, drain, and ground terminals to dedicated routing pads, an additional e-beam lithography step was performed using ARN 7250 resist. After development, the patterns were transferred by RIE.The completed crossbar arrays were encapsulated with a 70 nm SiO_2 layer deposited by RF sputtering at 600 W. Finally, Au contact pads on W routing pads were fabricated using an optical bilayer resist process consisting of LOR 5B and AZ 1512. Following resist development, contact openings to the underlying W routing pads were created by RIE. The wafer was then transferred to the PVD chamber, where the exposed electrode surfaces were cleaned using Ar plasma. A 5 nm Ti adhesion layer (300 W, 20 sccm Ar), followed by a 200 nm Au layer (200 W, 20 sccm Ar), was deposited *in situ* and subsequently patterned by lift-off to form the final contact pads. The devices were fabricated with two different r_{PCM} spans, 200 nmand 700 nm.

TEM studies

Transmission electron microscopy (TEM) lamellae were prepared by focused ion beam (FIB) - scanning electron microscope (SEM), milling in a Thermo Scientific™ Helios™ 5 UC DualBeam system using 30 keV Ga+ ions, followed by low-energy (2 keV Ga+) polishing to minimize FIB-induced amorphization. A JEOL ARM200F Cs-corrected TEM, equipped with a cold-field-emission gun and operating at 200 keV, was used to analyze the TEM lamellae. Micrographs were acquired in Z-contrast mode by High-Angle Annular Dark Field scanning TEM (HAADF–STEM). A GIF Quantum ER system was used for Electron Energy Loss Spectroscopy (EELS) measurements in Spectrum Imaging (SI) mode. Low-magnification Bright Field (BF) STEM were acquired by the STEM module of the FIB-SEM at 30 keV.

Electrical characterization

The electrical measurements were conducted using a custom-built multi-probe probe station. DC measurements of the device state were performed using a Keithley 2606B System SourceMeter. AC signals were applied to the device with a Tektronix AFG 31102 arbitrary function generator. A Tektronix oscilloscope (DPO 5104B), recorded the voltage pulses applied to and transmitted by the device. Switching between the DC and AC measurement circuits was achieved using mechanical relays in a Keithley 707A switching matrix mainframe equipped with two Keithley 7173-50 matrix cards. A custom-made relay board allowed a 50 Ω resistor to be connected from each of the 24 outputs to ground for impedance matching during programming. The chip stage was heated using a light bulb as the heating element, driven by a DC power supply (EA-PS 3150-04B) and controlled by a temperature controller (Eurotherm 2416). The programming curves were obtained by applying voltage pulses in the sequence RESET-READ-SET-READ. The RESET pulse used leading and trailing edges of 8 ns, and a width of 50 ns, while the amplitude was varied to create amorphous discs of increasing size. The SET pulse had fixed amplitude with leading and trailing edges of 500 ns and 5000 ns, and a width of 10 ns. For MVM operations, the weights are programmed using an iterative write-and-verify scheme, targeting a conductance accuracy within a convergence tolerance of 1 %. All READ operations were performed at 0.1 V. The activation energies of the amorphous, crystalline, device, and liner contributions, $E_{\mathrm{a,amor}}$, $E_{\mathrm{a,cry}}$, $E_{\mathrm{a,device}}$, and $E_{\mathrm{a,liner}}$, were extracted from linear Arrhenius fits of the resistance vs. temperature data according to $\log R(T) = \log R_\infty + \frac{E_\mathrm{a}}{k_\mathrm{B}} \frac{1}{T}$, where $R(T)$ is the resistance at temperature T, R_∞ is the pre-exponential resistance factor, and k_B is the Boltzmann constant. The activation energy was obtained from the slope of the fit. The crystallization speed measurements shown in Figure 2 were performed by sweeping the pulse width and amplitude of box-shaped pulses. The trailing and falling edges for the pulse was 8 ns. The BOX pulse width varied from 10 to 1200 ns. To extract the drift coefficient, the resistance vs. time traces were fitted using the resistance drift relation $R(t) = R_0 \left(\frac{t}{t_0}\right)^\nu$. Equivalently, the drift coefficient is given by $\nu = \frac{\log(R(t)/R_0)}{\log(t/t_0)}$. Here, R_0 is the device resistance at the reference time t_0, and $R(t)$ is the resistance at time t.

REFERENCES

[1] T. Kim and S. Lee, "Evolution of phase-change memory for the storage-class memory and beyond," *IEEE Transactions on Electron Devices*, vol. 67, no. 4, pp. 1394–1406, 2020.

[2] P. Cappelletti, R. Annunziata, F. Arnaud, F. Disegni, A. Maurelli, and P. Zuliani, "Phase change memory for automotive grade embedded NVM applications," *Journal of Physics D: Applied Physics*, vol. 53, no. 19, p. 193002, 2020.

[3] W.-S. Khwa, Y.-C. Chiu, C.-J. Jhang, S.-P. Huang, C.-Y. Lee, T.-H. Wen, F.-C. Chang, S.-M. Yu, T.-Y. Lee, and M.-F. Chang, "A 40-nm, 2M-cell, 8b-precision, hybrid SLC-MLC PCM computing-in-memory macro with 20.5-65.0 TOPS/W for tiny-Al edge devices," in *International Solid-State Circuits Conference (ISSCC)*, vol. 65, pp. 1–3, IEEE, 2022.

[4] M. Le Gallo, R. Khaddam-Aljameh, M. Stanisavljevic, A. Vasilopoulos, B. Kersting, M. Dazzi, G. Karunaratne, M. Brändli, A. Singh, S. M. Mueller, *et al.*, "A 64-core mixed-signal in-memory compute chip based on phase-change memory for deep neural network inference," *Nature Electronics*, pp. 1–14, 2023.

[5] M. Lanza, A. Sebastian, W. D. Lu, M. Le Gallo, M.-F. Chang, D. Akinwande, F. M. Puglisi, H. N. Alshareef, M. Liu, and J. B. Roldan, "Memristive technologies for data storage, computation, encryption, and radio-frequency communication," *Science*, vol. 376, no. 6597, p. eabj9979, 2022.

[6] G. S. Syed, M. Le Gallo, and A. Sebastian, "Phase-change memory for in-memory computing," *Chemical reviews*, vol. 125, no. 11, pp. 5163–5194, 2025.

[7] A. Sebastian, M. Le Gallo, R. Khaddam-Aljameh, and E. Eleftheriou, "Memory devices and applications for in-memory computing," *Nature nanotechnology*, vol. 15, no. 7, pp. 529–544, 2020.

[8] W. Zhang, R. Mazzarello, M. Wuttig, and E. Ma, "Designing crystallization in phase-change materials for universal memory and neuro-inspired computing," *Nature Reviews Materials*, vol. 4, no. 3, pp. 150–168, 2019.

[9] S. G. Sarwat, "Materials science and engineering of phase change random access memory," *Materials science and technology*, vol. 33, no. 16, pp. 1890–1906, 2017.

[10] D. Dragoni, J. Behler, and M. Bernasconi, "Mechanism of amorphous phase stabilization in ultrathin films of monoatomic phase change material," *Nanoscale*, vol. 13, no. 38, pp. 16146–16155, 2021.

[11] S. Prili, V. Bragaglia, V. P. Jonnalagadda, J. Luchtenveld, B. J. Kooi, F. Arciprete, A. Sebastian, and G. S. Syed, "Understanding the growth and properties of sputter-deposited phase-change superlattice films," *Advanced Materials Interfaces*, vol. 12, no. 12, p. 2500058, 2025.

[12] K. Ding, J. Wang, Y. Zhou, H. Tian, L. Lu, R. Mazzarello, C. Jia, W. Zhang, F. Rao, and E. Ma, "Phase-change heterostructure enables ultralow noise and drift for memory operation," *Science*, vol. 366, no. 6462, pp. 210–215, 2019.

[13] A. I. Khan, A. Daus, R. Islam, K. M. Neilson, H. R. Lee, H.-S. P. Wong, and E. Pop, "Ultralow–switching current density multilevel phase-change memory on a flexible substrate," *Science*, vol. 373, no. 6560, pp. 1243–1247, 2021.

[14] M. Boniardi, A. Redaelli, C. Cupeta, F. Pellizzer, L. Crespi, G. D'Arrigo, A. L. Lacaita, and G. Servalli, "Optimization metrics for Phase Change Memory (PCM) cell architectures," in *2014 IEEE International Electron Devices Meeting*, pp. 29.1.1–29.1.4, 2014.

[15] W. W. Koelmans, A. Sebastian, V. P. Jonnalagadda, D. Krebs, L. Dellmann, and E. Eleftheriou, "Projected phase-change memory devices," *Nature communications*, vol. 6, no. 1, p. 8181, 2015.

[16] G. S. Syed, K. Brew, A. Vasilopoulos, V. P. Jonnalagadda, B. Kersting, T. Philip, V. Bragaglia, S. Ambrogio, J. Büchel, J. Giannopoulos, *et al.*, "In-memory compute chips with carbon-based projected phase-change memory devices," in *2023 International Electron Devices Meeting (IEDM)*, pp. 1–4, IEEE, 2023.

[17] S. Ghazi Sarwat, T. M. Philip, C.-T. Chen, B. Kersting, R. L. Bruce, C.-W. Cheng, N. Li, N. Saulnier, M. BrightSky, and A. Sebastian, "Projected mushroom type phase-change memory," *Advanced Functional Materials*, vol. 31, no. 49, p. 2106547, 2021.

[18] G. S. Syed, T. Philicelli, J. Luchtenveld, V. P. Jonnalagadda, S. Gautam, K. Brew, B. J. Kooi, V. Narayanan, N. Saulnier, and A. Sebastian, "Disc-type phase-change memory devices for analog in-memory computing," in *2025 IEEE International Electron Devices Meeting (IEDM)*, pp. 1–4, 2025.

[19] B. Kersting, V. Ovuka, V. P. Jonnalagadda, M. Sousa, V. Bragaglia, S. G. Sarwat, M. Le Gallo, M. Salinga, and A. Sebastian, "State dependence and temporal evolution of resistance in projected phase change memory," *Scientific reports*, vol. 10, no. 1, p. 8248, 2020.

[20] S. G. Sarwat, M. Le Gallo, R. L. Bruce, K. Brew, B. Kersting, V. P. Jonnalagadda, I. Ok, N. Saulnier, M. BrightSky, and A. Sebastian, "Mechanism and Impact of Bipolar Current Voltage Asymmetry in Computational Phase-Change Memory," *Advanced Materials*, vol. 35, no. 37, p. 2201238, 2023.

[21] N. Li, C. Mackin, A. Chen, K. Brew, T. Philip, A. Simon, I. Saraf, J.-P. Han, S. G. Sarwat, G. W. Burr, *et al.*, "Optimization of projected phase change memory for analog in-memory computing inference," *Advanced Electronic Materials*, vol. 9, no. 6, p. 2201190, 2023.

[22] H. Kwon, A. I. Khan, C. Perez, M. Asheghi, E. Pop, and K. E. Goodson, "Uncovering thermal and electrical properties of Sb_2Te_3/GeTe superlattice films," *Nano Letters*, vol. 21, no. 14, pp. 5984–5990, 2021.

[23] Y. Chen, C. Rettner, S. Raoux, G. Burr, S. Chen, R. Shelby, M. Salinga, W. Risk, T. Happ, G. McClelland, *et al.*, "Ultra-thin phase-change bridge memory device using GeSb," in *International Electron Devices Meeting*, pp. 1–4, IEEE, 2006.

COMPETING INTERESTS

The authors declare no competing interests.

ACKNOWLEDGMENTS

This work was supported by the European Research Council Grant INFUSED (Grant No. 101222715) and by the IBM Research AI Hardware Center. A.M. Mio acknowledges financial support from the project BEYOND NANO Upgrade (CUP G66J17000350007) and from NextGenerationEU, M4C2, within the PNRR project NFFA-DI, CUP B53C22004310006, IR0000015, having benefited from the access provided by CNR-IMM@CT in Catania. We thank Timothy Philicelli, Stephan Menzel, Siddharth Gautam, Andrea Cassini, Jesse Luchtenveld for valuable technical discussions and contributions, and Vijay Narayanan for his management support. We also acknowledge support from Urs Egger toward electrical characterization setup. We also acknowledge the cleanroom operations team at the Binnig and Rohrer.

FIGURES

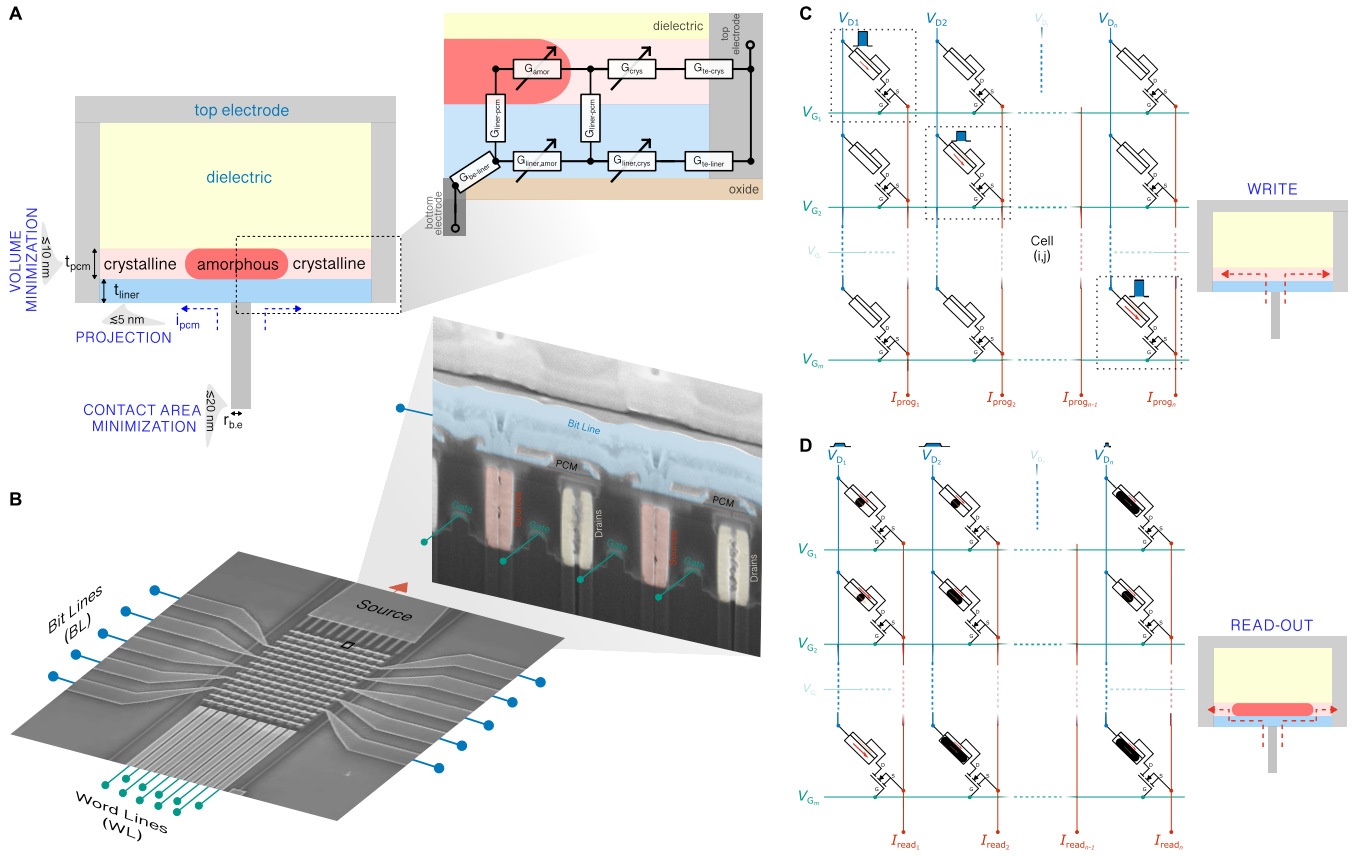


FIG. 1. **Device concept**. (A) Cross-sectional schematic of the proposed computational phase-change memory device. The device uses an ultra-thin ($\lesssim 10\,$nm) phase-change layer and a projection liner ($\lesssim 5\,$nm) positioned between the dielectric and the bottom electrode. During programming or reading, current flows radially from the bottom electrode across either the phase-change layer or the liner to the edge-contacted top electrode. The device achieves high compute -efficiency (via reduced volume and contact area), -precision (through liner suppressing temporal fluctuations), and -density (from vertical geometry and efficient programming). The inset shows the equivalent electrical circuit model, represented by lumped elements, including the contact interfaces. Here, G_{ele} represents the conductance contributions from the various interfacial and bulk elements which are critical for understanding and designing the device behavior. (B) SEM micrograph of an crossbar array; each cross-point contains an nMOS selector in series with a phase-change device. Inset: TEM micrograph of a single device. (C) Illustration of programming (write) operation: devices are programmed to analogue states (phase configurations), with current flowing through the higher-conductance phase-change segment, generating Joule heating and triggering the amorphization or crystallization. (D) Illustration of readout operation: devices of varying conductance (phase configurations) can be read individually or used collectively for MAC operations. In this mode, current bypasses the amorphous region and flows through the non-drifting projection liner segment parallel to the liner.

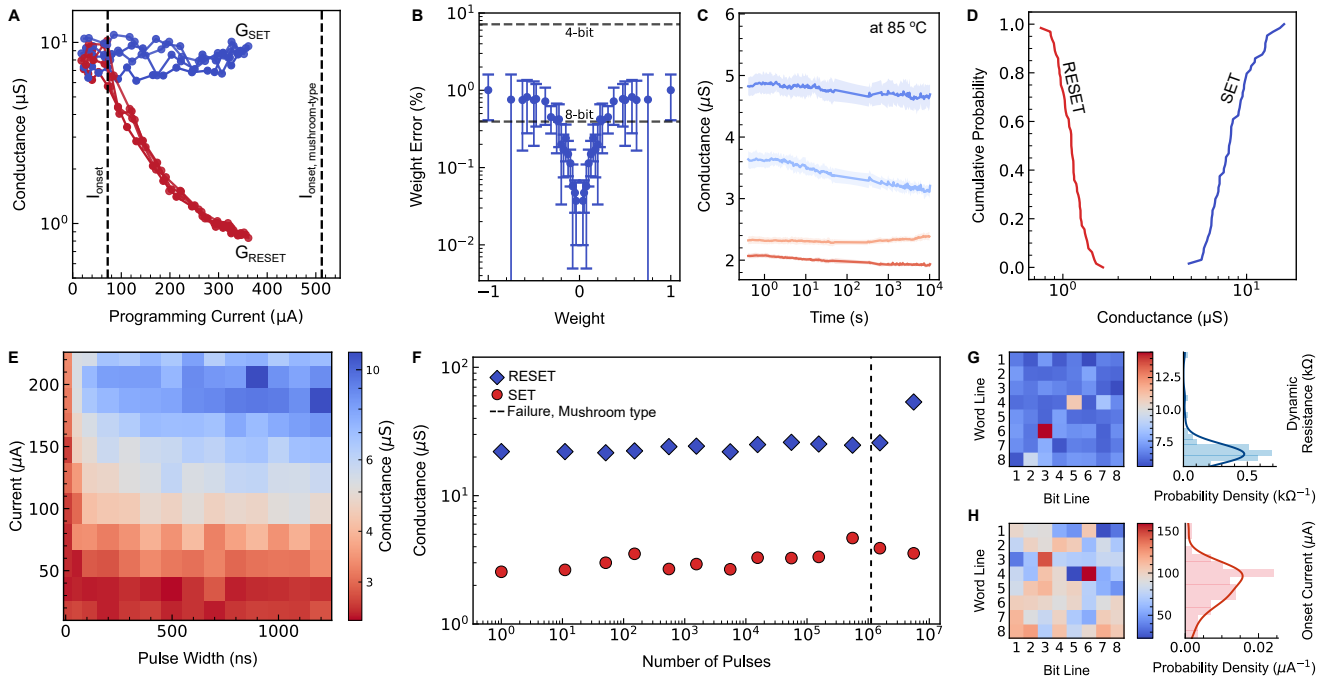


FIG. 2. **Programming characteristics**. (A) A plot illustrating the programming curves of a device showing a notable reduction in programming currents. (B) A plot showing the weight error obtained from mapping synaptic weight values onto the analogue conductance states of the device. The error bars in represent one standard deviation. Also indicated are the 4-bit and 8-bit effective quantization noise. (C) The temporal evolution of conductance for a two representative states of the device at elevated temperature (85 °C), showing stable retention behavior. (D) Reliability plot showing the complementary cumulative distribution of RESET conductance values and the cumulative distribution of SET conductance values measured across 64 devices in the crossbar array, obtained from single-pulse programming. (E) Measurement of the crystallization speed in the device using box-type SET pulses of varying pulse amplitude and duration. The color map shows the device conductance as a function of the pulse parameters. (F) A typical endurance plot, where a device is repeatedly stressed by bursts of RESET pulses. It can be seen that it is possible to achieve reliable reversible switching between SET and RESET states in excess of 100 million cycles. (G) Heat map and corresponding distribution illustrating the dynamic ON-resistance values of 64 devices. Higher R_{ON} enhances Joule heating efficiency, resulting in a reduced onset current. (H) Heat map and corresponding distribution of the onset current across the 64 devices.

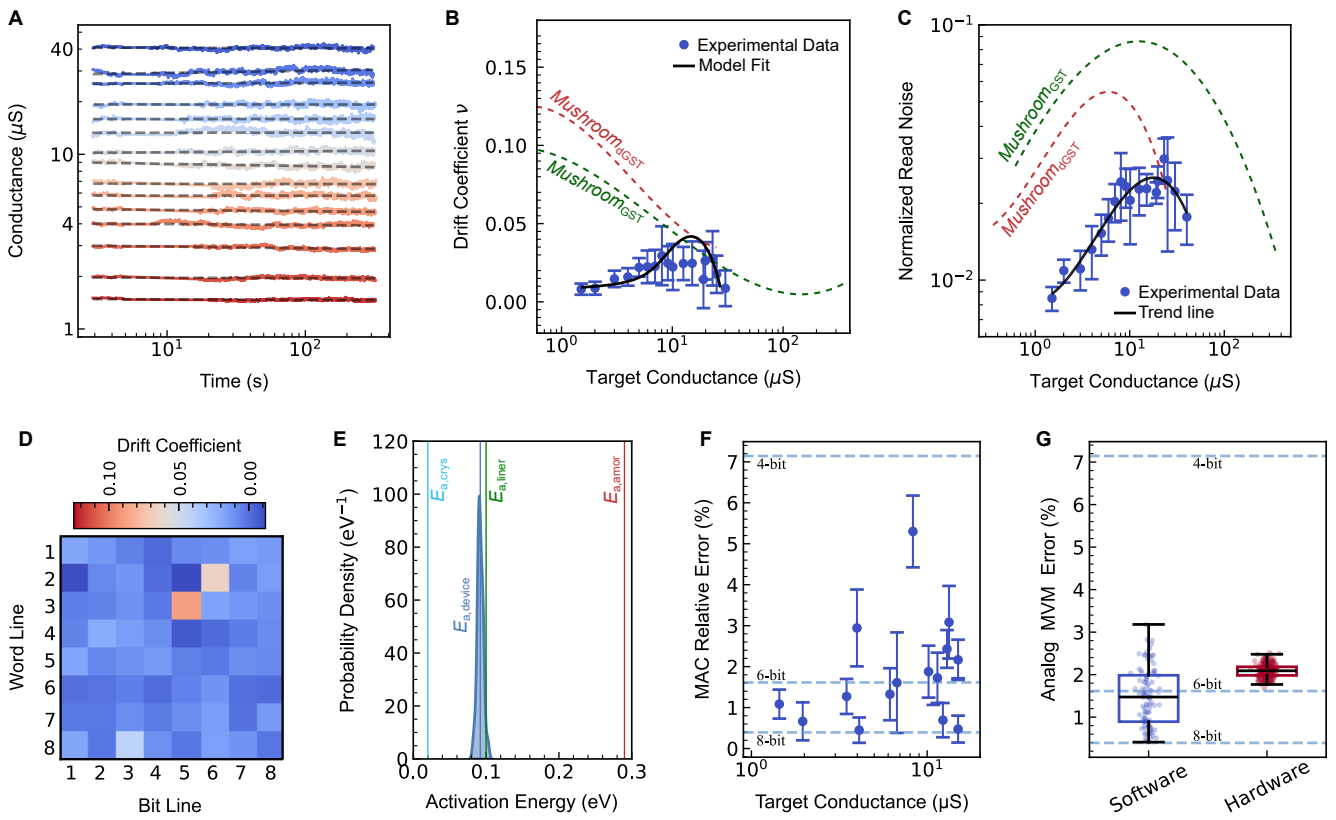


FIG. 3. **Read-out characteristics**. (A) Conductance versus time traces of many programmable states in a device. The data is fitted with the conductance drift equation, $G(t) = G(t_0)(t/t_0)^{-\nu_{\text{eff}}}$ (black dotted traces). (B) A plot comparing the proposed computational memory device with undoped and doped mushroom-type devices on state-dependent conductance drift coefficients. The black trace is the fit to the data using the analytical model. (C) Comparison of the state-dependent normalized read noise of the proposed computational memory device with that of undoped and doped mushroom-type devices. The normalized read noise is calculated as the standard deviation of the residual raw conductance traces after subtraction of the drift component, divided by the target conductance. Data in (B) and (C) are averaged over 10 repetitions per iteratively programmed state; error bars indicate one standard deviation. (D) Heat map illustrating the spatial distribution of drift coefficients for full RESET states across 64 devices in the crossbar array. (E) Distribution of activation energies extracted from measurements of 27 devices in the RESET state. (F) State-dependent relative error of individual multiply–accumulate (MAC) outputs, where each MAC corresponds to the scalar dot product between one crossbar row and the applied input vector, with all devices participating in the MAC programmed to the same target conductance state. Measurements are shown for different target conductance states, together with the error limits associated with different digitally quantized precisions. (G) Analogue matrix–vector multiplication errors, quantified as the normalized (l_2) difference between the complete measured and ideal output vectors, comparing accumulation performed in software with full hardware accumulation on the crossbar array.

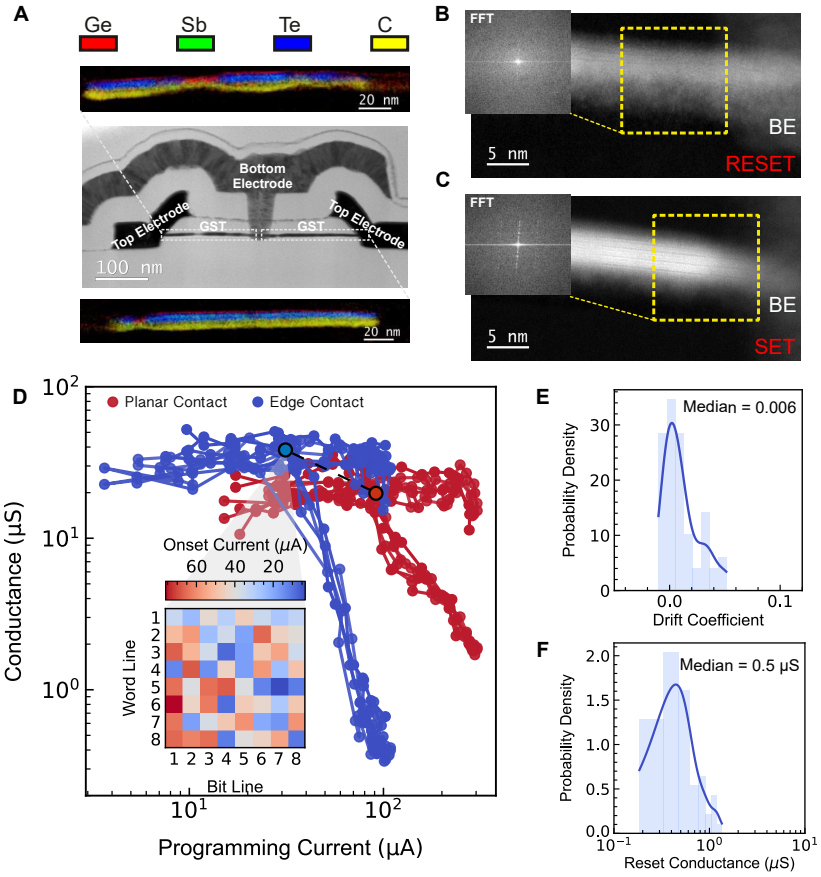
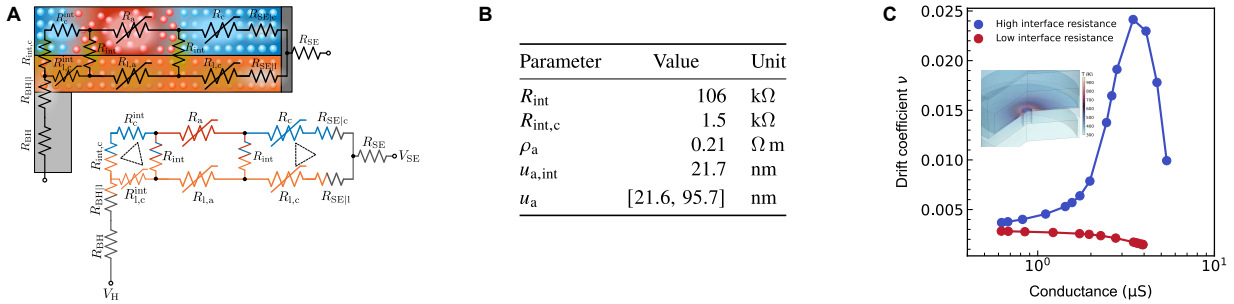


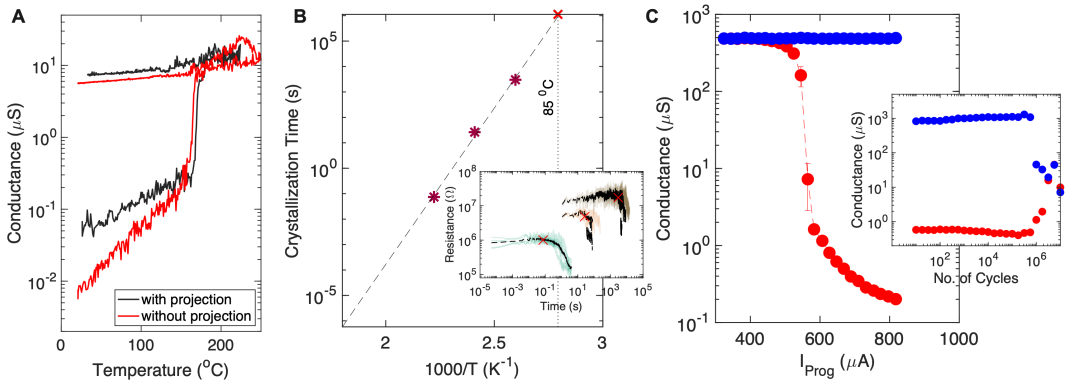
FIG. 4. **Prospects for further contact area minimization**. (A) Bright-Field scanning transmission electron microscopy (BF-STEM) micrograph of a device with a penetrating bottom electrode, together with the corresponding elemental composition maps. (B) HAADF-STEM image of a device in the RESET state, showing an amorphous region adjacent to the bottom electrode (BE). The corresponding FFT pattern confirms the amorphous character of the region. (C) HAADF-STEM image of a device in the SET state, showing a crystalline region adjacent to the BE. The corresponding FFT pattern confirms the crystallinity in the region. (D) Comparison of programming curves between planar- and edge-contacted devices, showing improved performance when the bottom electrode fully penetrates the device stack. Inset: heat map of the onset programming currents measured across 64 devices in the array. (E) Distribution of drift coefficients measured across 64 devices, highlighting significant drift reduction. (F) Corresponding RESET conductance values across the measured devices.

EXTENDED FIGURES

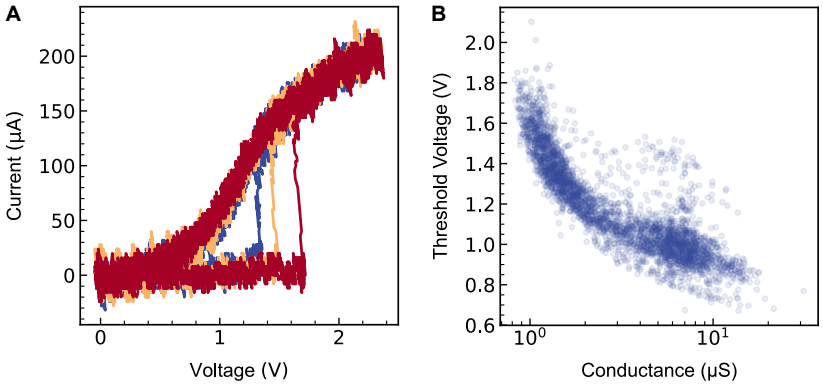


Parameter	Value	Unit
R_{int}	106	kΩ
$R_{\text{int,c}}$	1.5	kΩ
ρ_{a}	0.21	Ω m
$u_{\text{a,int}}$	21.7	nm
u_{a}	[21.6, 95.7]	nm

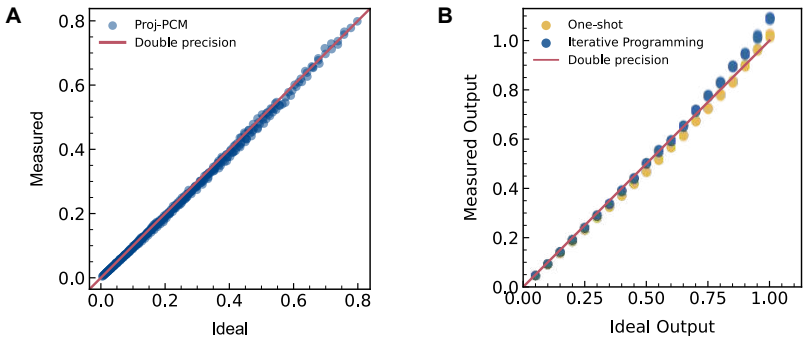
Extended Data Fig. 1. **Analytical Description of State-Dependent Device Operation.** (A) Radial cross-section of the device indicating the dominant resistance components, together with the equivalent circuit model used to describe the readout characteristics. We distinguish between resistors representing bulk materials (fully colored) and resistors representing interfaces between materials (half-colored). Two delta configurations can be identified in the leftmost and rightmost parts of the circuit, as highlighted by the dashed triangles. (B) Parameters extracted from a global fit of the analytical model to the experimentally measured traces. For each initial conductance state, measurements from repeated programming iterations are combined to obtain a median resistance trace, and the traces corresponding to all conductance states are fitted simultaneously. The outer amorphous radius, u_{a}, is fitted independently for each initial conductance state, whereas the liner–PCM interface resistance (R_{int}), the liner–PCM interface resistance associated with the crystalline region ($R_{\text{int,crys}}$), the amorphous-phase resistivity (ρ_{a}), and the radius of the inner crystalline disk ($u_{\text{a,int}}$), are constrained to be shared across all states. The remaining geometrical, material, and electrical parameters are fixed to experimentally measured values or values taken from the literature. (C) FEM simulations comparing bottom-heater devices with low and high liner–PCM interface resistance, highlighting the impact of the interface resistance on projection efficacy. The inset shows the three-dimensional device model used for the FEM simulations. The drift coefficients are extracted by linearly fitting the simulated resistance over time in log–log space, assuming an experimentally observed power-law dependence and an amorphous drift coefficient of $\nu_{\text{amor}} = 0.10$. All material parameters are either measured experimentally or taken from the literature.



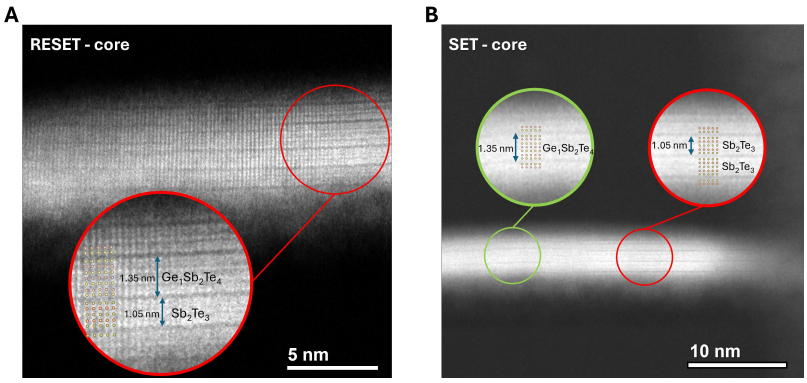
Extended Data Fig. 2. **Crystallization behavior and programming characteristics of mushroom-type devices.** (A) Resistance–temperature (R–T) characteristics of the edge-contacted devices with and without the liner. The crystallization temperature is approximately 165 ^{0}C, consistent with the characteristic behavior of GST.(B) Crystallization activation energy extracted for the fully RESET state of the device. The inset shows the raw traces. The extracted activation energy is 2.43 eV, consistent with reported values for GST. This corresponds to an estimated data retention time of approximately 11 days at 85 ^{0}C, which is sufficient for inference applications. (C) Programming curves of a mushroom-type device. The device comprises an 80 nm GST phase-change material layer with a bottom electrode radius of $r_{\text{be}} = 20$ nm. The inset shows the endurance characteristics, indicating device failure after approximately 1 M SET–RESET cycles. The incorporation of dopants, such as oxide additives, is well known to improve both the endurance and retention of GST-based materials.



Extended Data Fig. 3. **State-Dependent Threshold Voltage.** (A) Experimental current–voltage characteristics demonstrating threshold switching for three RESET states corresponding to different amorphous volumes. (B) Extracted threshold voltages as a function of device state, highlighting the state-dependent scaling of the threshold-switching behavior.



Extended Data Fig. 4. **Scalar and MAC Operations.** (A) Scalar multiplication performed using a single projected disc-type device programmed to different conductance states. For each programmed state, positive read voltages were applied to the device and the measured output current is plotted against the ideal current computed in software using double-precision arithmetic from the target conductance value. The agreement with the ideal line corresponds to an equivalent scalar-multiplication precision greater than 6 bits. (B) Experimentally measured MAC operations plotted against the exact result for one-shot and iterative programming schemes. In the iterative case, one row of six devices was programmed to the same target conductance, with the procedure repeated for different target values. A fixed input voltage of 0.1 V was applied to the row, and the measured accumulated current was compared with the ideal MAC result computed in software. In the one-shot case, the devices were programmed with a single pulse and measured individually; their measured conductance values were used as reference values for the software calculation. This represents a best-case comparison, as the accuracy is limited only by the hardware accumulation step, yielding an equivalent precision greater than 6 bits.



Extended Data Fig. 5. **Details of the Atomic Structure. (A)** HAADF-STEM image of a device in the RESET state. The inset shows a magnified view of the phase change material channel. Van der Waals (vdW) gaps, visible as dark horizontal lines, are oriented parallel to the phase change material channel, evidencing the preferential alignment of the crystalline grains. The regions enclosed by adjacent vdW gaps exhibit thicknesses of approximately 1.05 nm and 1.35 nm, corresponding to the lengths of Sb_2Te_3 and $Ge_1Sb_2Te_4$ structural blocks, respectively. A schematic representation of the corresponding atomic arrangement is overlaid for clarity. **(B)** HAADF-STEM image of a device in the SET state. The inset shows a magnified view of the phase change material channel. vdW gaps as before are visible as dark horizontal lines and oriented parallel to the phase change material channel, corresponding to the lengths of Sb_2Te_3 and $Ge_1Sb_2Te_4$ structural blocks.