

Seamlessly Fused Digital-Analogue Reconfigurable Computing using Memristors

Alexantrou Serb, Ali Khiat, Themistoklis Prodromakis

Corresponding author address: A.Serb@soton.ac.uk

University of Southampton, Highfield campus, Southampton, UK, SO17 1BJ

As the world enters the age of ubiquitous computing, the need for reconfigurable hardware operating close to the fundamental limits of energy consumption becomes increasingly pressing. Simultaneously, scaling-driven performance improvements within the framework of traditional analogue and digital design become progressively more restricted by fundamental physical constraints. Emerging nanoelectronics technologies bring forth new prospects yet a significant rethink of electronics design is required for realising their full potential. Here, we lay the foundations of a design approach that fuses analogue and digital thinking by combining digital electronics with analogue memristive devices for achieving charge-based computation; information processing where every dissipated charge counts. This is realised by introducing memristive devices into standard logic gates, thus rendering them reconfigurable and capable of performing analogue computation at a power cost close to digital. The versatility and benefits of our approach are experimentally showcased through a hardware data clusterer and an analogue NAND gate.

Realising the rapid expansion of the Internet of Things (IoT) relies on the availability of energy-, area- and computationally-efficient yet affordable and often-reconfigurable hardware platforms that could allow for bespoke customisation¹. Within the fully digital design paradigm that still dominates modern electronics, downscaling of integrated circuits² has been the main driver for lowering power dissipation; a process now reaching its physical limits. Simultaneously, reconfigurability has continued to rely for its physical implementation on dedicated memory blocks and progressively bottlenecked by power-hungry data transfers between physically separate memory and processing elements³. This has occurred despite work on Field Programmable Analogue Arrays (FPAAs) based on floating gate MOSFETs⁴, which still require high programming voltages (typ. 10V+)⁵ and their own, dedicated area on the chip, thus forcing the memory element and the circuit using said memory to remain far apart (with all parasitic capacitance and energy dissipation consequences this entails).

The rich landscape of modern electronics design became even more diverse with the steady introduction of memristive devices⁶ into the family of standard electronic components^{7,8}. The ability of memristors to act as thresholded electrically tuneable⁹, multi-level¹⁰, non-volatile resistive loads¹¹, combined with their inherently scaling-friendly^{12,13}, low power¹⁴ and back-end integrable¹⁵ fabrication processes has rendered them a highly promising candidate for use in future electronics applications^{9,16–18}. These properties promote memristors as ideal candidates for achieving in-silico reconfigurability in a post-Moore context, i.e. without relying on front-end integration density for performance and operating on the principle of separate, dedicated memory and processing elements.

In this work, we lay the foundations of a design approach that amalgamates the analogue non-volatile memory capacity of metal-oxide memristors with the fundamental building blocks of digital design: logic gates. This fusion occurs at a fundamental component level, enmeshing memristors and transistors in order to achieve collocation of memory and computation. Our approach is thus distinct from conventional mixed-signal design, whereby the analogue/digital parts remain separate entities, interacting purely at the signal level. This true coalescence of paradigms engenders a distinct set of fundamental building blocks: analogue reconfigurable gates featuring embedded memory. We first demonstrate the reconfigurability modes of a memristor-enhanced inverter by delineating how tuning the memory states of individual devices enables the control of the transfer characteristics of the gate. The proposed design paradigm is completed by introducing appropriate read-out circuits that make our modified gates interoperable with standard digital gates. We envision this emerging concept becoming a staple in numerous emerging applications and showcase the versatility of the proposed paradigm by experimentally demonstrating two applications: a hardware analogue gate and an analogue domain template matcher.

RESULTS:

Reconfigurable analogue gate concept, operation and performance

Much like multi-valued logic is a generalisation of standard Boolean logic, the proposed analogue gates are inspired as generalisations of standard logic gates. This is realised through the topology shown in Figure 1a for the analogue inverter (comparison with a standard Boolean inverter is provided in Supplementary Figure 1). Every current path to/from the output node of the gate is regulated by the presence of a tuneable resistive element, in our case a metal-oxide memristor (see Supplementary note 1 for the fully general approach). This architecture has the same inputs and outputs as a standard gate but receives analogue inputs and generates an analogue output. Overall, the proposed analogue inverter serves as a potential divider comprising two transistors and two memristors (2T2R). Depending on the precise levels of the input voltages, the output voltage behaviour may be dominated either by the states of the transistors (standard logic gate operation) or by the memristive components and their interrelations (divider operation). The former is obtained when input voltages are clear binary values, whilst the latter at intermediate levels. This occurs because at each edge of the input voltage range one of the transistors always exhibits a source-drain impedance that is sufficiently high to dominate the entire divider and lead to standard Boolean inverter operation. Simultaneously, at intermediate values of input voltage both transistors are open and the memristive potential divider becomes dominant. This introduces a plateau in the transfer characteristic of the inverter, visible in Figure 1d, which controls the shape of the mapping between input and output voltages whilst maintaining the fundamental inverter nature of the circuit (0 maps to 1 and vice versa). Controlling the resistive states of R_{UP} and R_{DN} allows this soft mapping to be reconfigured. Whilst our design approach is technology agnostic, the reconfiguration quality is defined by the characteristics of the employed memristive technology. Most importantly, the programming of R_{UP} and R_{DN} can be achieved by a time-shared scheme via periphery circuitry that minimises the pixel's power and area constraints to a minimum; details provided in Supplementary note 2.

A key feature of the proposed design is its power efficiency that can be illustrated by comparing an analogue vs a digital inverter. For any given input voltage both inverters can be described to some

approximation as two component potential dividers, illustrated in Supplementary Figure 2b. Whenever the input voltage changes from some value $V_{IN,1}$ to $V_{IN,2}$ the corresponding outputs must change from some $V_{OUT,1}$ to $V_{OUT,2}$; an operation that requires changing the amount of charge stored on capacitor C_{out} through a capacitor current i_{cap} while keeping leakage current i_{leak} low (illustrated in Supplementary Figure 2 with full derivation in Supplementary note 3). In the case of a standard inverter the only possible input (output) voltages are GND and VDD for logic 0 and 1 respectively, which guarantees that one of the transistors M1 and M2 will always be OFF, i.e. at very high source-drain impedance. This, in turn, minimises i_{leak} allowing the inverter to operate closer to the theoretical lower bound energy $C_{out} \frac{VDD^2}{2}$ for each state change¹⁹ (see caveats in Supplementary note 3). The analogue inverter is governed by the same fundamental dynamics with the exception that M1 and M2 may both be partially ON at the same time. As a result, the inverter becomes capable of performing analogue-in/analogue-out computation at a fraction of the energy consumption of its digital counterpart (e.g. factor of ~4-5 see Supplementary note 4 for details).

Given any fixed set of resistive states for R_{UP} and R_{DN} , an analogue gate will implement a specific soft input/output mapping, with each memristor constituting a design degree of freedom (dof). Since every memristor augments the impedance seen from the output node to either supply or ground along a unique current path, these design degrees of freedom are linearly independent. Thus, the proposed analogue inverter features two degrees of reconfigurability freedom, whose span is bounded by the range of resistive state values that the corresponding memristor can attain. The design of the CMOS counterpart of the system, and in particular the settings of the key design parameters of transistor aspect ratio $\frac{W}{L}$ and power supply voltage VDD, may then be tailored so as to optimise gate functionality given the chosen memristor technology's inherent resistive state ranges (See Supplementary Figure 3 and Supplementary Table 1). The 2-dof reconfigurability space of the inverter is illustrated in Figure 1b, showing two useful, orthogonal analogue inverter mapping control modalities. In both cases, the values of the two memristors R_{UP}, R_{DN} are altered simultaneously first under the constraints $\frac{R_{UP}}{R_{DN}} = c$ (ratio-fixed modality) and then under the constraints $R_{UP} + R_{DN} = c$ (sum-fixed modality), where c is a suitable constant in each case. The two control modalities exert orthogonal effects on the plateau in the inverter's transfer characteristic (see Figure 1c). The ratio-fixed modality controls the breadth of the plateau by altering the balance between the total memristor-transistor (source-drain) impedance whilst the sum-fixed modality controls the altitude of the plateau by altering the voltage distribution within the memristive potential divider. Provided that the distribution of voltage between memristors is not affected by the overall voltage drop across them, these modalities do not interact with each other (see Supplementary Figure 4). Figure 1c presents experimental evidence of an analogue inverter whose memristors have been successively programmed into four configurations: $\{R_{UP}, R_{DN}\} \in \{HH, HL, LH, LL\}$ where H, L stand for high and low resistive state respectively. Results demonstrate operation under an ON/OFF resistive state ratio of approximately 10, leading to significant changes in the input/output transfer characteristic of the inverter. The HH - LL pair illustrates modulation of plateau width independent of altitude whilst the HL - LH pair illustrates altitude modulation sans width modulation.

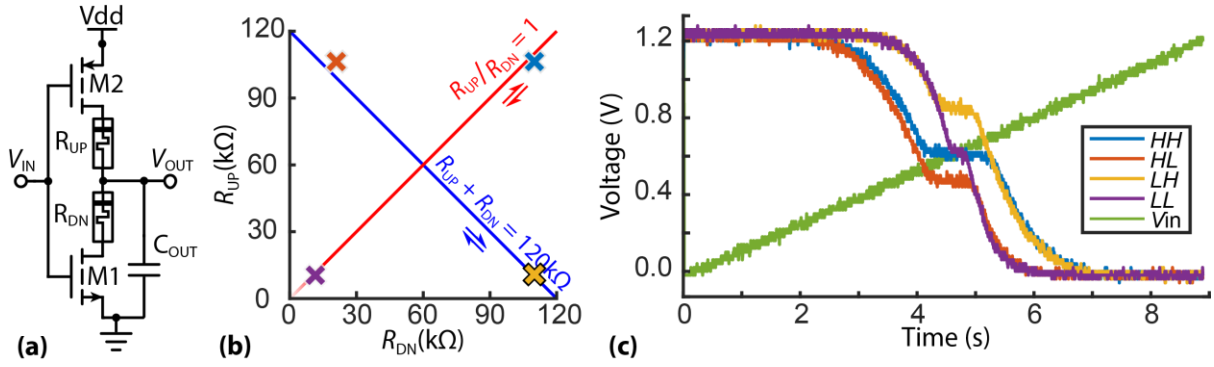


Figure 1: Reconfigurability modalities in an analogue inverter gate. (a) Memristor-enhanced analogue inverter topology. (b) Changing the resistive states of the memristors R_{UP} , R_{DN} in the inverter so as to keep their sum (along blue line) or ratio (along red line) constant offers flexibility in controlling the inverter's transfer characteristics. The constant sum modality allows independent control of transfer characteristic's plateau height whilst the constant ratio modality allows for independent control of the plateau's width (see Supplementary Figure 4). Colour-coded crosses correspond to the R_{UP} , R_{DN} configurations used in the results of panel (c) (see Supplementary Table 2 for details). (c) Four measured examples of analogue inverter transfer characteristics corresponding to the cases where R_{UP} and R_{DN} are both high (HH), high and low (HL), low and high (LH) and both low (LL) respectively. The measured input voltage during the HH trial is shown in green as V_{in} (similar for all trials). Note independent modulation of plateau width and altitude by the sum and ratio between R_{UP} , R_{DN} .

Interoperability with standard digital electronics

No electronic system can become commercially competitive vis-a-vis standard CMOS technology if it cannot be both read in a simple and efficient manner and modularly chained, which in our case means that the output of an analogue gate has to be a suitable input for the next one. Since the proposed gates employ analogue voltages as both inputs and outputs this compatibility is ensured. We note however, that with each memristor-based gate acting as a non-linear amplifier, each device maps a relatively restricted range of input voltages to the full power supply range, thus the chaining cannot be continued ad infinitum. Rather it is envisaged that short chains (perhaps 2-5 links) of these analogue gates ending in a digitisation stage, where a 1 or 0 answer will be committed to, will be used. In this manner the ability of digital to perform implicit error-correction by collapsing all answers to a binary space is balanced with analogue's ability to discern between those two levels. The chaining ability is exemplified in Supplementary Figure 5, where a NAND gate operates on the basis of input received from an analogue inverter.

Transferring from analogue logic to Boolean can be easily achieved using a read-out circuit consisting of a simple inverter fed through a mirror supply as shown in Figure 2. The analogue-to-Boolean link rests on the fact that the read-out inverter will be characterised by a switch-point voltage, i.e. an input voltage level at which both transistors are simultaneously ON and the inverter output voltage is close to the middle of the supply. Any input voltages above switch-point will be digitised to 0 whilst values below switch-point will digitise to 1 (Figure 2, V_{OUT1}). A small range of input values very close to switch-point, however, will lead to an unclear digitisation that may stochastically result in digital 1 or 0 values. Notably, changing the resistive states of the memristors in any analogue gate can alter the input variable space regions which lie above/below the read-out inverter switch-point and thus indirectly tune the overall mapping from analogue gate inputs to digital output. The switch-point of the read-out inverter is determined by the aspect ratio (W/L) of its constituent transistors and is a design parameter; though there is no reason memristors cannot be used to render the read-

out inverter reconfigurable, too. On the other hand, converting Boolean-to-analogue requires no conversion as any Boolean input level is automatically a valid analogue input.

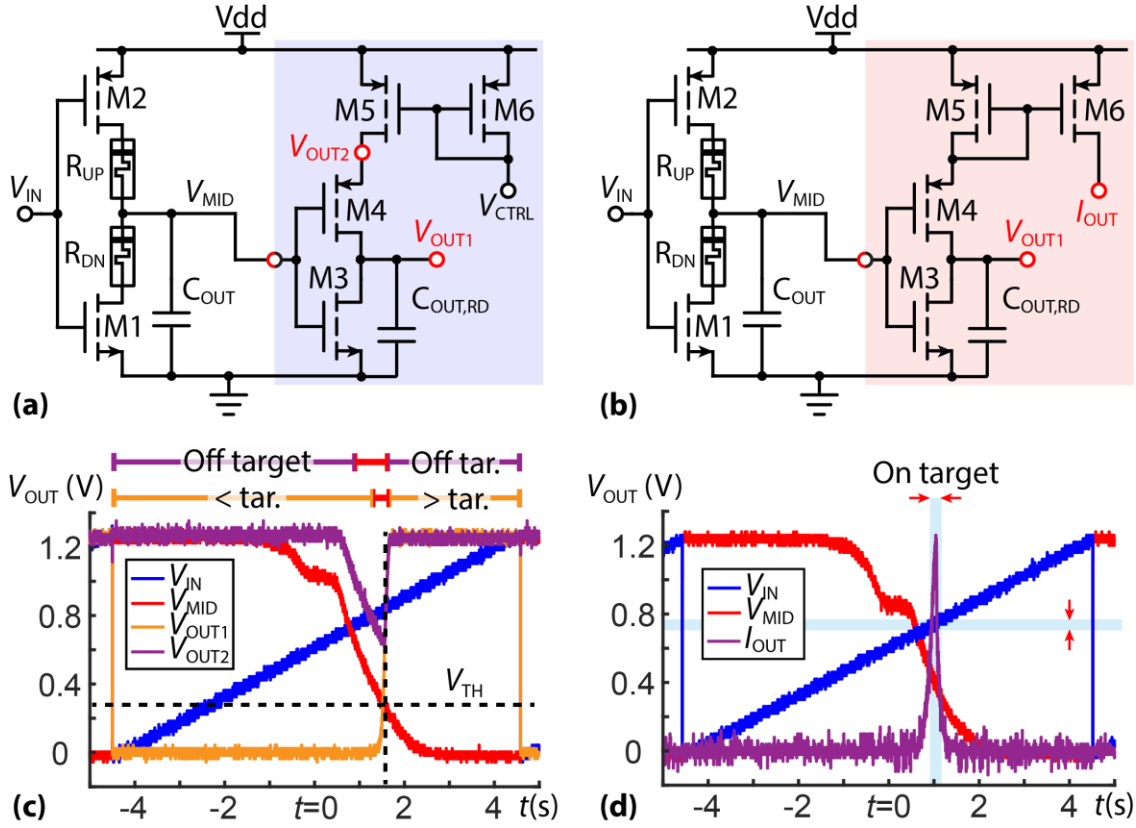


Figure 2: Reading analogue gates. (a) Read-out circuit (shaded blue) designed to receive an analogue gate output (V_{MID}) and then digitise to 1 or 0 based on whether V_{IN} is above or below some threshold V_{TH} (at V_{OUT1}) and also indicate whether V_{IN} is close or far from V_{TH} (at V_{OUT2}). (b) Alternative read-out circuit (shaded red) where V_{OUT2} has been replaced by I_{OUT} , which now indicates proximity of V_{IN} with V_{TH} by sourcing a large current only if $V_{MID} \approx V_{TH}$. (c) Measured results from the read-out circuit in (a). V_{OUT1} successfully digitises the analogue inverter output V_{MID} through most of its input voltage range (V_{IN}) as shown by the orange bar above the plot (red segment indicates potentially ambiguous digitisation). V_{OUT2} , on the other hand exhibits a dip only when V_{MID} is sufficiently close to V_{TH} , directly indicating whether V_{IN} is on-target (i.e. close to V_{TH}) or off-target. This is illustrated by the purple bar above the plot. (d) Measured results from the read-out circuit in (b). Only I_{OUT} shown for clarity. I_{OUT} peaks within a narrow range of input voltages satisfying $V_{MID} \approx V_{TH}$ (on target).

The mirror supply shown in Figure 2a is not strictly necessary for the digitisation strategy described previously, but offers an interesting alternative approach. The principle of operation relies on the observation that the read-out inverter conducts most current at the switch-point voltage, when both its transistors are maximally ON simultaneously and digitisation at node V_{OUT1} is unclear. The mirror supply exploits this by attempting to force a reference current into the inverter. If the reference current is chosen appropriately, the voltage on node V_{OUT2} from Figure 2a is driven towards a digital 0 only when the inverter is sufficiently close to its maximally conducting state (both transistors simultaneously open). This circuit therefore offers a way of performing digitisation by mapping a very specific analogue gate output level to a digital 1, i.e. allowing that particular value to act as a target and the analogue gate to output an off-target/on-target response (see 2c, purple bar). Modifying the magnitude of the reference current will tune the analogue gate output level range for which digitisation using this strategy will return a digital 1. A slight modification of the mirror supply leads to the circuit depicted in Figure 2b, which mirrors the read-out inverter current directly to a circuit terminal. Providing the digitisation output in the form of a current allows for easy summation

of digitisation results from many analogue gates. The resulting, summed current can then easily be provided as input to an integrate and fire neuron as used in neuromorphic engineering²⁰. Notably, both variants in Figure 2 indicate that the on-target range of the input voltage is of the order of 100mV as evidenced by V_{OUT1} voltages that are no longer a clear digital 1 or 0 (Figure a,c) or I_{OUT} currents significantly above baseline (Figure 2b,d). The on-target range will depend on the resistive state values of the memristors as is immediately evident by observing that the traces in Figure 1c will intercept the threshold shown in Figure 2c at different points along the x-axis (and consequently different input voltage levels).

Case studies: an analogue gate and a template-matching pixel

The introduced design paradigm forms a generalised framework that can be used to develop a broad range of applications. Here, we showcase two cases, namely: analogue gates and template matching. Soft logic relies on analogue-in/analogue-out gates in order to perform computation on a continuum between the extreme values of 0 and 1. This implements a soft input/output mapping that can be used for function approximation. The example of an analogue NAND gate is shown in Figure 3. An analogue NAND can be implemented via a three-way divider consisting of four transistors and three memristors (4T3R). Figure 3b shows the transfer characteristics from the two inputs to the output, which now define a surface. Notably, when input A is fully ON (digital 1), then the analogue NAND reduces to an analogue inverter from B to the output with a mapping determined solely by memristors M_B and M_C . We shall term this the analogue inverter reduction of B. The same holds when input B is a digital 1. In the case where either of the inputs is fully OFF (digital 0), the output of the analogue NAND will always be a digital 1. Changing the resistive states of the memristors controls the shape of the analogue function surface whilst retaining its inherent NAND nature. The measured results shown on Figure 3b denote an essentially multiplicative interaction between the analogue inverter reductions of A and B, suggesting that the analogue function surface may be controlled in a reasonably orthogonal way by varying M_A and M_B at the cost of restricting the possible shapes it may assume. Similar conclusions may be drawn for other types of gates (see Supplementary Figure 6), although analogue NANDs are already functionally complete (in the sense that using multiple analogue NANDs, mappings corresponding to any other analogue gate can be constructed).

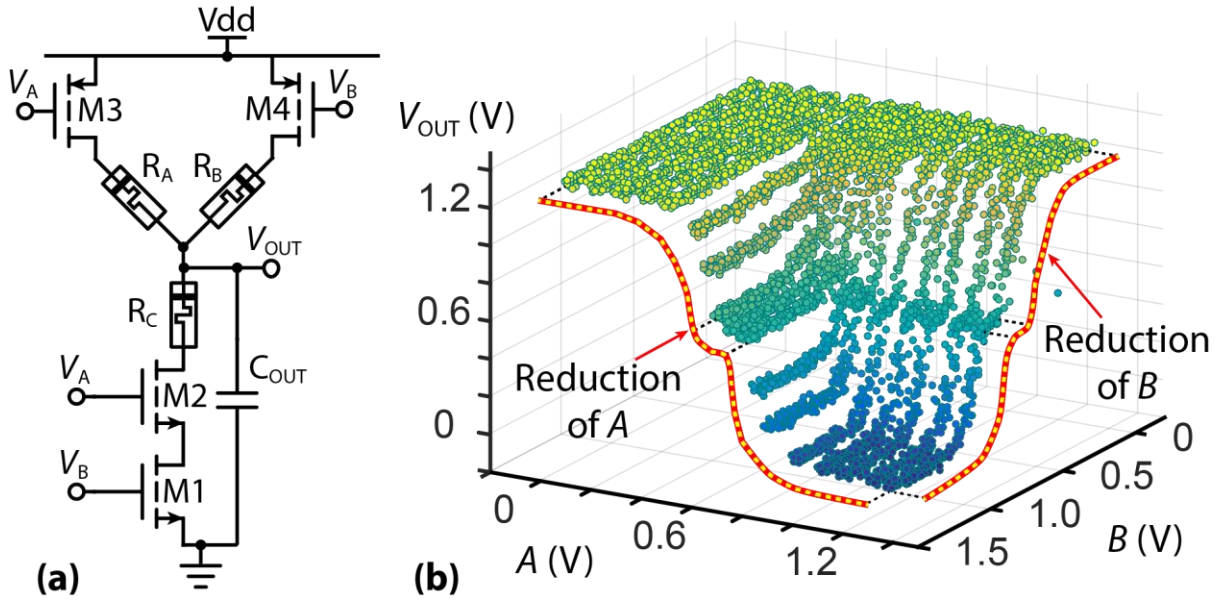


Figure 3: Analogue NAND architecture and basic behaviour. (a) Analogue NAND topology. (b) Measured analogue NAND transfer characteristics. The inverter reductions of A and B (see text) are shown as red/yellow lines. The overall transfer characteristic arises from an essentially multiplicative relation between the reductions of A and B.

Template matching is a technique whereby a small part of a signal (audio/electrical waveform snippet or even image segment) is compared against a stored template. Specifically in the emerging field of bioelectronics, template matching facilitates neural spike sorting for electrophysiological studies²¹, whereby electrical waveforms recorded from neural cell assemblies are template-matched in short snippets of typically 10-20 samples²². Its strength stems from the fact that whenever a match is found the system registers the occurrence of a spike and the matching template ID, thus simultaneously providing spike timing and identification information. A simplified version of our proposed analogue inverter with the read-out circuit from Figure 2b is used herein, shown in Figure 4a, for demonstrating spike detection/sorting. We refer to this circuit as a texel (template-matching pixel) and its operation can be understood as searching for an input voltage value V_{IN} that matches stored value determined by the resistive state of memristor R1. This process is shown in Figure 4b-d by illustrating measured transfer characteristics of a discrete texel. This implementation, being modular, power efficient and truly scalable, allows for aggregation into arrays; a proof of concept 4-texel array is further shown in Figure 4e. The array is fed with appropriately selected samples (see methods) from nine neural spike waveforms from the same database²³ (Supplementary Figure 7) and summing the current outputs of each texel down a common load resistor, as shown in Figure 4e,f. Three spike instances were chosen from each available class of spikes: a low (L), a medium (M) and a high (H) instance corresponding to spikes exhibiting lower than, similar to or higher than class-average voltage levels (see methods and Supplementary Figure 7). The voltage level at the system's V_{OUT} terminal is linked to the degree of matching between the input vector \mathbf{k} and the stored template and was directly used as a matching degree metric. Due to the similarity between the H instance of class 1 and the L instance of class 2 and the limited resolution of our instrumentation, the experiment for these two instances was ran only once with a common input vector \mathbf{k} (see Supplementary Table 3). Results in Figure 4g show a texel array set up to discriminate for class 2 spikes. Even using only four samples from each waveform (marked in Figure 4e) strong discrimination between templates is clearly achieved. The memristor resistive states were confirmed to remain stable before and after the experiment (Supplementary Table 4).

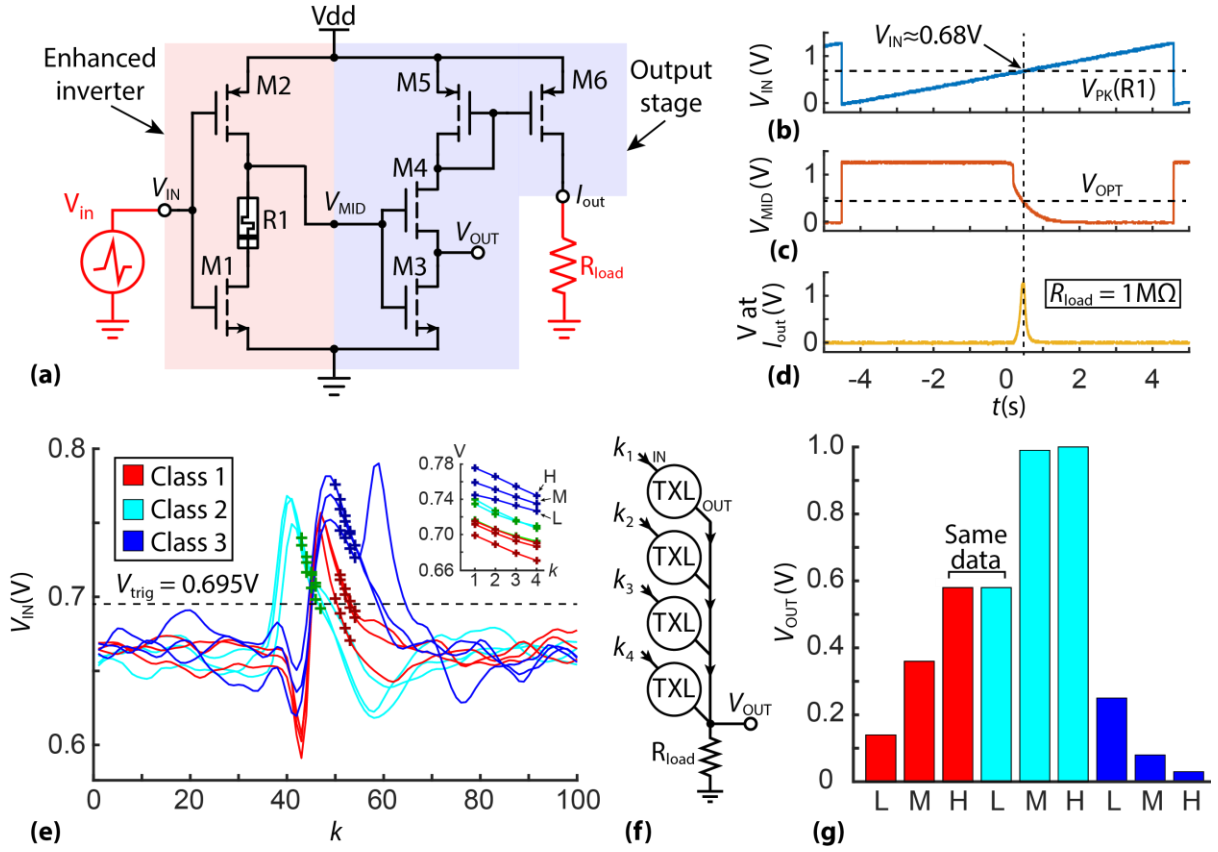


Figure 4: Analogue domain template matching enabled by memristive technologies. (a) Schematic of texel circuit illustrating consisting of an analogue inverter and the read-out stage (red and blue shading correspondingly). (b-c) Texel transfer characteristics from input voltage (b), through mid-point voltage V_{MID} (c) and voltage at the output node (c) marking V_{IN} input voltage level (V_{PK}) and V_{MID} voltage level (V_{OPT}) at which the output stage sources its maximum current. (e) Selected spike waveforms used as input to the test texel array. Crosses indicate the sample points used to feed the array. k : sample number. V_{trig} : texel array sampling trigger level (see methods section). Inset: close-up of the chosen sample points. L, M, H and arrows: Low, medium and high voltage instances of spikes in class 3. (f) Schematic of 4-texel array used to carry out experiments. (g) Measured output voltage when spike samples from (e) are applied to the texel array in (f). L, M, H versions of spikes in each class shown. Higher V_{OUT} voltage means greater degree of matching between input data and stored template. The texel array was programmed to respond best to class 2 spikes. Colours as in (e). Class 1-H and class 2-L results refer to the same experiment (see methods section).

DISCUSSION:

Our design approach provides a simple, powerful and generic tool for truly fusing the analogue and digital design worlds. We note that this concept is optimised first for low energy dissipation and then for speed, reflecting the pertinent needs of modern ubiquitous computing hardware. We thus envision the proposed technology to complement and coexist with standard CMOS implementations. The transistors and memristors in Figure 1a can in fact be any elements that modulate a resistance (e.g. floating gate MOS), however, our proposal of embedding memristors as analogue resistance tuning elements yields a power efficient implementation. In every case, the basic concept remains the same: signal-controlled tuneable resistance elements intended to operate as ON/OFF switches are combined with initialise & operate continuously tuneable resistance elements. This combination results in systems where digital and analogue behaviours coexist, with the ON/OFF and continuously tuneable components tasked to sustain each type of operation respectively. Importantly, we consider the reconfigurability aspect of our approach to be of fundamental significance as it can endow embedded hardware with the capabilities of (1) trimming to compensate for technology

(both CMOS and memristors) imperfections (variation and mismatch), (2) tuning their functionality for efficiently addressing applications involving drifting specifications (e.g. texel - searching for a variety of spike templates) and (3) repurposing the core operation of a system.

At a technical level there are some crucial points to be noted: First, though it may seem counter-intuitive to use canonical circuits such as the one in Figure 1 that include a current leakage term it turns out that because of the sheer simplicity of the circuit some tasks may be carried out more efficiently than even using traditional digital approaches, as shown in Supplementary note 4. This holds so long as each of these circuits is operated only so long as it is computing something useful and coming to a stable answer (i.e. for the correct number of time constants as remarked in Supplementary note 3). For the rest of the time the circuit may be either power gated off, or simply parked in a digital state where its inputs are forced to a binary value and the system reverts to standard Boolean operation, relying on a CMOS comparable leakage.

Second, the issue of variance often plagues analogue electronics, leading to the use of transistor sizes much larger than technically necessary given the rules of how mismatch scales²⁴. Introducing trimming elements in our circuits we have a degree of freedom previously not obtainable in pure CMOS alone. Especially if we consider the case of program rarely and read often, using a trial and error-based programming circuit, as described in Supplementary note 2, the possibility of automatically compensating for mismatch and variation through tuning of resistive states becomes viable. This is extremely important as it removes one of the biggest barriers analogue electronics face when downscaling. In this work we have kept transistor sizings conservatively large in order to ensure that even without resistive state tweaking (all devices assumed programmed at nominal) the variance shown by Monte Carlo simulations (shown in Supplementary figure 8) remain under check. The overarching aim of this work is to push the limits of scaling by introducing performance-check-based programming. The exceptionally high resolution of memristor resistive states¹⁰ works in our favour by supporting the idea of very fine resolution, ultra-compact, few-component memristive trimming (no need for digital registers and switching banks). Finally we note that proving dominance over purely digital approaches is by no means trivial. The fact however that our proposed circuit architectures consist of few relatively large transistors, as opposed to many fully downsized ones, implies that there may still be room at the bottom for additional improvements.

At the application level we note that our template matching example might appear niche to begin with, however pattern matching in general is one of the most fundamental operations of signal processing. In the case of spike sorting the raw signal itself is such that a template matcher can identify a spike class directly from a simple set of texel inputs. Nevertheless, the input signal need not be a waveform and need not be raw (unprocessed by some other system before entering the template matcher). A simple example could be an output vector from a neural network layer (e.g. a few levels deeper than the input layer).

In conclusion, this work lays the conceptual foundations and provides experimental proof of a design approach for electronics formed by marrying standard, digital circuitry with analogue electronics and rapidly emerging memristive devices to offer analogue computation at close to the power/area price of digital with the added benefit of reconfigurability. This is achieved by collocating and enmeshing memory and computation whilst maintaining full interfacing compatibility both with standard digital logic and internally (ability to chain analogue gates as is done with digital gates). Finally, the

versatility of this paradigm is illustrated by two independent applications in analogue logic and template matching for bio-signals that given the power constraints of alternative technologies would have been impossible to realise.

Methods

Memristive device fabrication and specification

Memristive device fabrication and specification: All the memristors used in the experimental setups are in $3 \times 3 \text{ mm}^2$ chips that are wire-bonded to PLCC68 packages. Each memristor is a $20 \times 20 \text{ }\mu\text{m}^2$ cross-point of top and bottom electrodes (TE & BE). The BEs were first fabricated on 6 inch Silicon wafer that was thermally oxidised to grow 200 nm thick SiO_2 , which serves as insulating base layer. Using conventional optical lithography the BEs were patterned with negative-tone resist AZ-2070 followed by a low power and short reactive ion etching (descum) to remove any residual resist on the patterned areas. Then, 5 nm Titanium (Ti) and 10 nm Platinum (Pt) film were deposited with electron beam evaporation at low rate ($0.5 \text{ }\text{\AA}\text{s}^{-1}$), with Ti serving for adhesion purposes. Leybold_Lab700eb tool that has high crucible-wafer distance (greater than 1m) was used for evaporation to guarantee parallel deposition. The combination of parallel evaporation and negative-tone resist which has undercuts after development ensures a good liftoff process, resulting in well-defined electrodes without wings (fences) that would affect the subsequent thin layer, thus becomes harmful for the final device. Next, similar photolithography was carried out using the active-layer mask that allows depositing the metal-oxide active bi-layer everywhere except on the BE pads, followed by 1 min descum. Lambda controlled plasma assisted reactive magnetron sputtering (Leybold Helios Pro XL) was used to deposit the active layers; 25 nm TiO_2 followed by 4 nm Al_2O_3 , at room temperature. TiO_2 was sputtered from Ti metal target with 8 sccm O_2 , 35 sccm Ar flows and 2 kW at the cathode, and 15 sccm O_2 flow and 2 kW at an additional plasma source. Al_2O_3 was sputtered from Al metal target with 15 sccm Ar flow and 100 W at the cathode, and 25 sccm O_2 flow and 1.5 kW at the additional plasma source. Before any sputtering, the additional plasma source was used for an extra clean of the substrate with 8 sccm O_2 , 10 sccm Ar and 2 kW. Optiwet-ST30 tool was then used to perform lift-off process with the following parameters; 3 mbar pressure and 60°C hot NMP for 1 hour. This tool ensures 100% lift-off yield even after sputtering which deposits material everywhere making the lift-off difficult, and even at large features (pad areas) that tends to stick to the surface after dissolving the resist. 1 min descum is needed to clean the surface before TE lithography. The 10 nm Pt TEs were patterned and defined in similar fashion to the BEs. No sonication was used in this process. Finally, $3 \times 3 \text{ mm}^2$ chips were diced for wire-bonding.

Experimental set-ups and procedures

All experiments illustrated in Figures 1-4 were carried out on circuits prototyped on breadboard or strip-board. External power supplies and signal generators were used to supply both signal inputs and power whilst results were gathered exclusively by oscilloscope. For these experiments packaged devices were used, connected to the set-ups via breakout boards. We used custom made 64-pin breakout boards connecting the devices to the rest of the system using jumper wires. This setup is significant because it demonstrates functionality after wafer dicing and wire-bonding. All experiments were ran under a power supply between 1.2-1.3V (single decimal figure precision power supply). In each case, the memristive devices used were placed in the required resistive states using an ArC ONE instrumentation board (ArC instruments, UK). All devices used for all experiments

were located on the same die, i.e. only one memristive device package containing a total of 32 memristors was sufficient to carry out all the work presented here.

Additional notes on analogue NAND experiment: Results for the analogue NAND experiment were taken strip-by-strip by setting input *A* to a succession of fixed values and sweeping input *B* for each of those values by use of ramp (saw-tooth) signals. Input *A* was stepped with 100mV resolution throughout the entire power supply range except between 0.5-0.8V where it was stepped with 50mV resolution for enhanced visibility.

Additional notes on spike sorting application experiment: A 4-point texel array with common load resistance of 300k Ω was implemented. The signals fed into this system arrived from two, dual-channel benchtop power supplies with two significant decimal digits resolution. The benefit of using synthetic neural recording input data is that it contains ground truth information on spike identification and timing. On that basis an automatic sample selection script was ran on each spike instance available in the dataset in order to choose which data-samples from each instance are to be fed into the texel array for matching against a stored template. The script operated as follows: the data-points in each spike instance were read sequentially and once a trigger threshold V_{trig} was exceeded for the first time the script skipped six samples and then choose the subsequent four as candidate inputs for the texel array set-up. This methodology was chosen because it rendered the three classes of spikes visibly distinguishable despite the use of only four template points. The overall trigger and sample approach is similar to the work by Restituto-Delgado et al²⁵. In a more mature system implementation a larger texel array containing more than four samples would be used. Next, the extracted candidate four-texel sample sets were separated by single unit-template class. From each class, three texel sets were chosen for further processing: one featuring typical (M), one featuring lower than usual (L) and one featuring higher than usual (H) voltage values (selection shown in Supplementary Figure 7). Waveforms where the presence of more than one spike within each instance had corrupted the output of the sample selection script were automatically excluded from the selection. The voltage range of all nine selected sample sets (L, M, H instances for each of the three classes) was then adjusted by application of a common pair of gain and offset settings (Gain: 0.1; Offset: 0.66V). The adjusted texel data-point voltages were then suitable for working with the input voltage values the texel circuits were built to discriminate between. These adjusted values are shown in the inset of Figure 4e and were used as the input to the texel array after being rounded to 10mV precision (two significant decimal digits). This procedure caused the rounded texel voltages of the H instance of class 1 and the L instance of class 2 to completely overlap, hence that experiment was conducted only once for both cases.

Data availability: All relevant data is available from the authors. Furthermore this is deposited at TBD with DOI TBD.

References:

1. Biason, A. *et al.* EC-CENTRIC: An Energy- and Context-Centric Perspective on IoT Systems and Protocol Design. *IEEE Access* **5**, 6894–6908 (2017).
2. Moore, G. E. Cramming more components onto integrated circuits, Reprinted from *Electronics*, volume 38, number 8, April 19, 1965, pp.114 ff. *IEEE Solid-State Circuits Newsl.* **20**, 33–35 (2006).

3. Wulf, W. A. & McKee, S. A. Hitting the memory wall. *ACM SIGARCH Comput. Archit. News* **23**, 20–24 (1995).
4. Hall, T. S., Twigg, C. M., Gray, J. D., Hasler, P. & Anderson, D. V. Large-scale field-programmable analog arrays for analog signal processing. *IEEE Trans. Circuits Syst. I Regul. Pap.* **52**, 2298–2307 (2005).
5. Kang-Deog Suh *et al.* A 3.3 V 32 Mb NAND flash memory with incremental step pulse programming scheme. *IEEE J. Solid-State Circuits* **30**, 1149–1156 (1995).
6. Yang, J. J., Strukov, D. B. & Stewart, D. R. Memristive devices for computing. *Nat. Nanotechnol.* **8**, 13–24 (2013).
7. Chua, L. Memristor-The missing circuit element. *IEEE Trans. Circuit Theory* **18**, 507–519 (1971).
8. Strukov, D. B., Snider, G. S., Stewart, D. R. & Williams, R. S. The missing memristor found. *Nature* **453**, 80–83 (2008).
9. Gupta, I. *et al.* Real-time encoding and compression of neuronal spikes by metal-oxide memristors. *Nat. Commun.* **7**, 1–16 (2016).
10. Stathopoulos, S. *et al.* Multibit memory operation of metal-oxide bi-layer memristors. *Sci. Rep.* **7**, 17532 (2017).
11. Waser, R. & Aono, M. Nanoionics-based resistive switching memories. *Nat. Mater.* **6**, 833–40 (2007).
12. Khiat, A. *et al.* High Density Crossbar Arrays with Sub- 15 nm Single Cells via Liftoff Process Only. *Sci. Rep.* **6**, 32614 (2016).
13. Xia, Q., Yang, J. J., Wu, W., Li, X. & Williams, R. S. Self-Aligned Memristor Cross-Point Arrays Fabricated with One Nanoimprint Lithography Step. *Nano Lett.* **10**, 2909–2914 (2010).
14. Torrezan, A. C., Strachan, J. P., Medeiros-Ribeiro, G. & Williams, R. S. Sub-nanosecond switching of a tantalum oxide memristor. *Nanotechnology* **22**, 485203 (2011).
15. Shulaker, M. M. *et al.* Three-dimensional integration of nanotechnologies for Computing and Data Storage on a Single Chip. *Nat. Publ. Gr.* **547**, 74–78 (2017).
16. Prezioso, M. *et al.* Training and operation of an integrated neuromorphic network based on metal-oxide memristors. *Nature* **521**, 61–64 (2015).
17. Sheridan, P. M. *et al.* Sparse coding with memristor networks. *Nat. Nanotechnol.* (2017). doi:10.1038/nnano.2017.83
18. Serb, A. *et al.* Unsupervised learning in probabilistic neural networks with multi-state metal-oxide memristive synapses. *Nat. Commun.* **7**, 12611 (2016).
19. Baker, R. J. *CMOS : circuit design, layout, and simulation.* (Wiley, 2010). doi:10.1002/9780470891179
20. Indiveri, G. *et al.* Neuromorphic silicon neuron circuits. *Front. Neurosci.* **5**, 73 (2011).
21. Rey, H. G., Pedreira, C. & Quiñ Quiroga, R. Past, present and future of spike sorting techniques. *Brain Res. Bull.* (2015). doi:10.1016/j.brainresbull.2015.04.007

22. Navajas, J. *et al.* Minimum requirements for accurate and efficient real-time on-chip spike sorting. *J. Neurosci. Methods* **230**, 51–64 (2014).
23. Quiroga, R. Q., Nadasdy, Z. & Ben-Shaul, Y. Unsupervised Spike Detection and Sorting with Wavelets and Superparamagnetic Clustering. *Neural Comput.* **16**, 1661–1687 (2004).
24. Pelgrom, M. J. M., Duinmaijer, A. C. J. & Welbers, A. P. G. Matching properties of MOS transistors. *IEEE J. Solid-State Circuits* **24**, 1433–1439 (1989).
25. Rodriguez-Perez, A., Ruiz-Amaya, J., Delgado-Restituto, M. & Rodríguez-Vázquez, Á. A low-power programmable neural spike detection channel with embedded calibration and data compression. in *IEEE Transactions on Biomedical Circuits and Systems* **6**, 87–100 (2012).

Acknowledgements: This work has been supported by the Engineering and Physical Sciences Research Council (EPSRC) grants EP/K017829/1 and the EU FP7 RAMP project.

Author contributions: AS and TP conceived the study, AS ran the experiments, processed the resulting data, AK developed the process flow and fabricated the memristive devices. All authors contributed to writing and refining the manuscript.

Competing financial interests: The authors declare no conflict of interest of any kind.

Figure legends:

Figure 1: Reconfigurability modalities in an analogue inverter gate. (a) Memristor-enhanced analogue inverter topology. (b) Changing the resistive states of the memristors R_{UP} , R_{DN} in the inverter so as to keep their sum (along blue line) or ratio (along red line) constant offers flexibility in controlling the inverter's transfer characteristics. The constant sum modality allows independent control of transfer characteristic's plateau height whilst the constant ratio modality allows for independent control of the plateau's width (see Supplementary Figure 4). Colour-coded crosses correspond to the R_{UP} , R_{DN} configurations used in the results of panel (c) (see Supplementary Table 2 for details). (c) Four measured examples of analogue inverter transfer characteristics corresponding to the cases where R_{UP} and R_{DN} are both high (HH), high and low (HL), low and high (LH) and both low (LL) respectively. The measured input voltage during the HH trial is shown in green as V_{in} (similar for all trials). Note independent modulation of plateau width and altitude by the sum and ratio between R_{UP} , R_{DN} .

Figure 2: Reading analogue gates. (a) Read-out circuit (shaded blue) designed to receive an analogue gate output (V_{MID}) and then digitise to 1 or 0 based on whether V_{IN} is above or below some threshold V_{TH} (at V_{OUT1}) and also indicate whether V_{IN} is close or far from V_{TH} (at V_{OUT2}). (b) Alternative read-out circuit (shaded red) where V_{OUT2} has been replaced by I_{OUT} , which now indicates proximity of V_{IN} with V_{TH} by sourcing a large current only if $V_{MID} \cong V_{TH}$. (c) Measured results from the read-out circuit in (a). V_{OUT1} successfully digitises the analogue inverter output V_{MID} through most of its input voltage range (V_{IN}) as shown by the orange bar above the plot (red segment indicates potentially ambiguous digitisation). V_{OUT2} , on the other hand exhibits a dip only when V_{MID} is sufficiently close to V_{TH} , directly indicating whether V_{IN} is on-target (i.e. close to V_{TH}) or off-target. This is illustrated by the purple bar above the plot. (d) Measured results from the read-out circuit in (b). Only I_{OUT} shown for clarity. I_{OUT} peaks within a narrow range of input voltages satisfying $V_{MID} \cong V_{TH}$ (on target).

Figure 3: Analogue NAND architecture and basic behaviour. (a) Analogue NAND topology. (b) Measured analogue NAND transfer characteristics. The inverter reductions of A and B (see text) are shown as red/yellow lines. The overall transfer characteristic arises from an essentially multiplicative relation between the reductions of A and B.

Figure 4: Analogue domain template matching enabled by memristive technologies. (a) Schematic of texel circuit illustrating consisting of an analogue inverter and the read-out stage (red and blue shading correspondingly). (b-c) Texel transfer characteristics from input voltage (b), through mid-point voltage V_{MID} (c) and voltage at the output node (c) marking V_{IN} input voltage level (V_{PK}) and V_{MID} voltage level (V_{OPT}) at which the output stage sources its maximum current. (e) Selected spike waveforms used as input to the test texel array. Crosses indicate the sample points used to feed the array. k : sample number. V_{trig} : texel array sampling trigger level (see methods section). Inset: close-up of the chosen sample points. L, M, H and arrows: Low, medium and high voltage instances of spikes in class 3. (f) Schematic of 4-texel

array used to carry out experiments. (g) Measured output voltage when spike samples from (e) are applied to the texel array in (f). L, M, H versions of spikes in each class shown. Higher V_{OUT} voltage means greater degree of matching between input data and stored template. The texel array was programmed to respond best to class 2 spikes. Colours as in (e). Class 1-H and class 2-L results refer to the same experiment (see methods section).