Adiabatic Leaky Integrate and Fire Neurons with Refractory Period for Ultra Low Energy Neuromorphic Computing

Supplementary Material

Marco Massarotto[1*], Stefano Saggini[1], Mirko Loghi[1], David Esseni[1]

[1]DPIA - Polytechnic Department of Engineering and Architecture, University of Udine, Via delle Scienze 206, Udine, 33100, Italy.

*Corresponding author(s). E-mail(s): massarotto.marco001@spes.uniud.it;
Contributing authors: stefano.saggini@uniud.it; mirko.loghi@uniud.it; david.esseni@uniud.it;

Supplementary Note 1: Variability of Synaptic Capacitor Bank

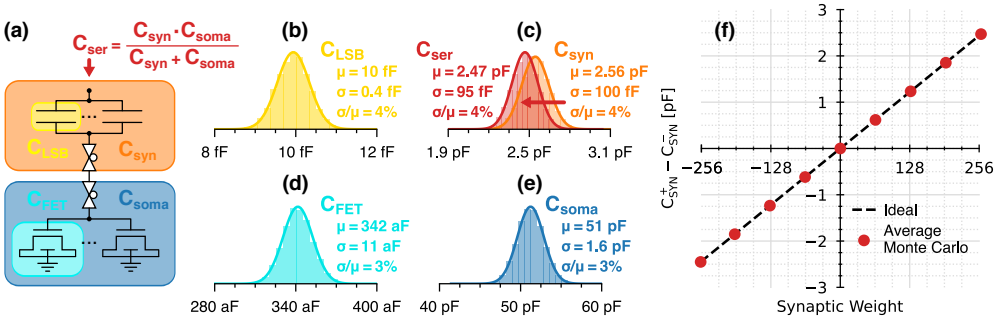

Supplementary Figure 1: Monte Carlo Analysis of the Synaptic Capacitor Bank. a Sketch of the series connection between the synapse and soma capacitors. b to e Gaussian distribution of the capacitors defined in (a). In (c), the series between the total synaptic capacitance (C_{syn} in orange) and the soma (e, C_{soma} in blue) results in a slightly lower effective synaptic capacitance (C_{ser} in red). f Transfer function translating the decimal-encoded synaptic weight into a capacitance difference ($C^{+}_{syn} - C^{-}_{syn}$). Thanks to the limited intrinsic variability of the capacitors, the average transfer function of many Monte Carlo simulations (red dots) is in very good agreement with the ideal case (black-dashed), resulting in an INL=0.01 and DNL=0.001.

The adiabatic implementation of all the neuromorphic functionalities devised in this work rely on the weighting of the charge transferred to the integrating soma, thus making a dependable and accurate programming of the capacitor bank implementing the synaptic weights a most critical requirement. For this reason, the variability of the synaptic weight induced by mismatch and process variabilities has been evaluated by means of Monte Carlo simulations.

First, in fig. 1a it is sketched the series connection between the synapse and the soma capacitors. In particular, we recall that the capacitor bank of the synapse consists of 2^{Nbit}=256 BEOL LSB capacitors (C_{LSB}, yellow), that are arranged by the synaptic weight into either a positive or a negative overall synaptic capacitance ($C^{\pm}_{syn}$). The total synaptic capacitance (C_{syn}=$C^{+}_{syn} + C^{-}_{syn}$, orange) is then connected to the soma capacitor (C_{soma}, blue), which in turn consists of several grounded n-FETs on the active silicon area, each contributing with a capacitance C_{FET} (cyan).

In figs. 1b to 1e are then reported the Gaussian distributions of all the aforementioned capacitors extracted from Monte Carlo simulations, alongside with the average value (μ) and the absolute (σ) and relative standard deviation (σ/μ).

As it can be seen, thanks to the large area of the capacitors, the variability is limited to just a few percents of the average value. On the other hand, the series between the synapse and soma inevitably results in an effective synaptic capacitance (C_{ser}) slightly lower than the nominal value. This is shown in red in fig. 1c, which

reports the maximum deviation occurring for a maximum positive or negative synaptic weight, namely when the entire orange C_{syn} is connected in series with the blue C_{soma} in fig. 1e. In principle, the deviation of C_{ser} from C_{syn} can be arbitrarily reduced by enlarging the C_{soma}.

Next, we have simulated the variability of the capacitor bank by varying the synaptic weight accounting for the distributions in figs. 1b to 1e. The resulting transfer function is reported in fig. 1f. We here recall that the transfer function of the capacitor bank converts the synaptic weight into a difference $(C^{+}_{SYN} - C^{-}_{SYN})$ between the positive and negative synaptic capacitors (see the "Capacitive Synapses and Soma" paragraph of the "Discussion" section in the main paper). Thanks to the limited variability of our capacitors, the average transfer function calculated by means of many mismatch and process-aware Monte Carlo simulations (red dots) is in a very good agreement with the ideal case (black-dashed line).

The linearity of the capacitor bank can be quantitatively expressed in terms of an *Integral Non-Linearity* (INL) and a *Differential Non-Linearity* (DNL), which are defined and evaluated respectively as [1]:

$$\begin{aligned} INL &= \max_{i \in SW} \left| \frac{ideal(i) - real(i)}{C_{LSB}} \right| = 0.01 \\ DNL &= \max_{i \in SW} \left[\frac{real(i+1) - real(i)}{ideal(i+1) - ideal(i)} - 1 \right] = 0.001 \end{aligned} \tag{1}$$

where *ideal* and *real* refer respectively to the ideal (black-dashed line) and non-ideal (red dots) transfer functions in fig. 1f, while the index i spans over all the synaptic weight symbols SW. Indeed, the limited intrinsic variability of all capacitors translates into low INL and DNL values, which ensure a monotonic transfer function without missing codes [1].

Supplementary Note 2: Neuron Comparator

The comparator of each neuron is tasked to monitor the differential membrane potential ΔV_m between the two soma capacitors and then fire an output spike when it crosses a certain threshold. Depending on the outcome of the evaluation, internally to the comparator are also generated the signals that enable the leakage of the neuron and the refractory period.

Single-Threshold Comparator

Our comparator is based on the single threshold comparator sketched in fig. 2a, which resembles a dynamic sense amplifier. The two single-ended membrane potentials $V_m^{\pm}$ are provided as input to the gate terminals of the two branches of the sense amplifier, which are identified respectively in red and blue in fig. 2a. The two branches can be made asymmetric by changing their equivalent conductance. In particular, we express such asymmetry in terms of an *Asymmetry Factor* (AF), which quantifies the imbalance between the W_+/L_+ ratio of the red positive branch with respect to the W_-/L_- ratio of the blue negative side. According to our definition in fig. 2a, the total equivalent area of both branches is independent of the AF.

The comparator is controlled by the EVAL signal (fig. 2b). When EVAL is low the comparator precharges both the Q± nodes of the cross-coupled latch at V_{DD} (fig. 2c). A single one-shot evaluation is then performed when EVAL is raised, which starts the simultaneous discharging of the Q± nodes. The outcome of the evaluation is given by the racing condition between $Q+$ and $Q-$, which depends on the differential membrane potential ΔV_m and on the conductance of each branch. As reported in fig. 2b, the faster-discharging node sets the state of the latch, which will reset only during the following precharge phase (EVAL=0).

Inverting *dynamic level restorers* are added to accelerate the relatively slow V_{DD}→GND transitions of the dynamic nodes (Q±), so as to suppress the otherwise relevant short-circuit currents in the following CMOS logic. As shown in fig. 2d, the OUT± of the level restorers are first reset to GND (EVAL=0) and then, during the evaluation (EVAL=1), one of them raises and marks which side won the racing condition.

All the waveforms shown in figs. 2b to 2d are simulated at the Worst Speed (WS) corner at a temperature of 100 °C. As it can be seen, the chosen evaluation time guarantees a reliable comparison even in the worst case scenario with low-conductive input transistors.

When the two branches are symmetric, which corresponds to AF=1, the racing condition is won by the higher single-ended $V_m^{\pm}$ voltage, and so the comparator is actually monitoring the sign of the differential membrane potential, i.e. by raising OUT+ if $\Delta V_m{>}0$, or OUT− if $\Delta V_m{<}0$.

On the other hand, when the asymmetry factor is greater than one, the node V_m^+ needs to be sufficiently higher than V_m^- in order to win the racing condition. This requirement can be expressed as $V_m^+ > V_m^- + V_{TH} \Rightarrow \Delta V_m > V_{TH}$, which gives rise to a positive threshold V_{TH} on the differential membrane potential. In fig. 2e, the comparator threshold has been evaluated for different AF configurations by means of Monte

Carlo simulations accounting for both mismatch and process variabilities. Indeed, a more pronounced asymmetry results in a higher threshold following an almost linear trend with respect to AF. Moreover, thanks to our definition of asymmetry factor, all the AF configurations occupy the same total area, thus ensuring a fair comparison between the uncertainties on the threshold voltage, which remain limited to just a few mV across all the explored asymmetry configurations.

Hysteretic Comparator

The complete neuron comparator consists of the cascade of two single-threshold comparators, as shown in fig. 2f.

First, the blue "REST" comparator is *symmetric* (AF=1) with a null threshold that monitors the sign of differential membrane potential (fig. 2g) to check if the neuron is in its resting state ($\Delta V_m \leq 0$). As an output, it generates the REST signal (fig. 2i), which disables the LEAK flag (fig. 2j) and in turn the leakage of the neuron. In addition, the "REST" comparator also provides the evaluation signal (EVAL_FIRE, fig. 2k) to the second comparator, which enters the evaluation phase only when ΔV_m is positive, thus saving up dynamic energy.

On the other hand, the second, red "FIRE" comparator is *asymmetric* (AF>1) and determines the firing thresholds of the neuron following fig. 2e. In particular, whenever the membrane potential overcomes the threshold, the comparator raises the FIRE signal (fig. 2l) together with the REFR flag (fig. 2m). The REFR flag triggers the refractory period of the neuron, which is terminated by the REST signal when the neuron returns to its resting state.

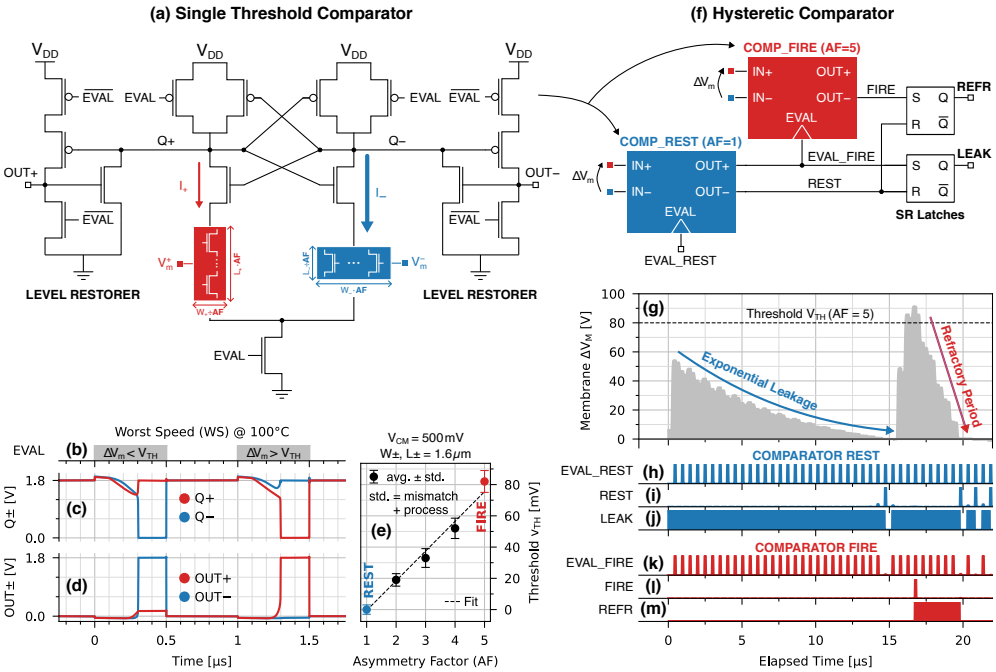

Supplementary Figure 2: Neuron Comparator. a Dynamic single-threshold comparator, whose outcome depends on a racing condition between the positive (red) and negative (blue) branches. The differential membrane potential (ΔV_m, purple) is provided as a differential input between the two branches. b EVAL signal alternates the comparator between precharge (EVAL=0) and evaluation phases (EVAL=1). c,d Voltage waveforms at the Worst Speed (WS) corner at 100 °C of the Q± and OUT± nodes when $\Delta V_m < V_{TH}$ (left) and $\Delta V_m > V_{TH}$ (right). e Simulated threshold V_{TH} as a function of the asymmetry factor (AF) when the common mode of the membrane potential is 500 mV. Error bars indicate the standard deviations extracted from Monte Carlo simulations by accounting for the mismatch and process variability. f Hysteretic neuron comparator consisting in the cascade of a symmetric "REST" comparator, which monitors the resting state of the neuron (blue, AF=1), and an asymmetric "FIRE" comparator, which fires an output spike when the membrane potential overcomes its firing threshold. g to m Example of a differential membrane potential (g) and of the resulting signals generated by the "REST" (h to j) and "FIRE" comparators (i to m).

Supplementary Note 3: Modelling of Exponential Leakage

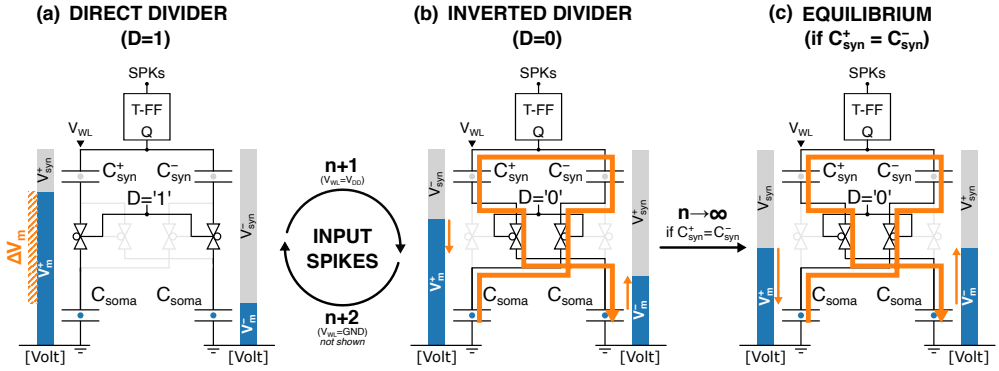

Supplementary Figure 3: **Charge redistribution in the Synapse Capacitive Network.** a Direct divider, set by D=1, connecting C^+_{syn} to V^+_m and C^-_{syn} to V^-_m. The soma (blue) and synapse voltages (gray) add up to the V_{WL}=V_{DD} enforced by the T-Flipflop. b Inverted divider, set by D=0, connecting C^+_{syn} to V^-_m and C^-_{syn} to V^+_m. Upon the transition from the direct divider in (a) to the inverted divider in (b), the charge redistributes to maintain V_{WL}=V_{DD} (orange). During the opposite transition from (b) to (a), V_{WL} is enforced to GND and a similar charge redistribution occurs (not shown). c Asymptotic equilibrium corresponding to V^+_m=V^-_m reached after many commutations between (a) and (b) and for a synaptic weight equal to zero, namely for C^+_{syn}=C^-_{syn}.

As described in the main paper, our implementation of the exponential leakage of the differential membrane potential (ΔV_m) exploits the charge redistribution among the synaptic ($C^{\pm}_{syn}$) and soma capacitors (C_{soma}) that occurs when their dividers are inverted upon each input spike.

We recall that the configuration of the capacitive dividers is controlled by the signal D, which sets the direct divider in fig. 3a for D=1, or the inverted divider in fig. 3b for D=0. If we let n be a time instant corresponding to the direct divider configuration in fig. 3a and, moreover, use the subscripts DIR and INV to denote the voltages $V^{\pm}_m$ respectively in the direct divider (fig. 3a) and inverted divider configuration (fig. 3b), then the charge redistribution can be analytically modelled by a system of finite-difference equations. In fact, upon the arrival of a spike at the instant (n+1) that reconfigures the capacitive network into the inverted divider of fig. 3b, the voltages

$V^+_{m,INV}(n+1)$, $V^-_{m,INV}(n+1)$ in the new divider configuration can be written as:

$$\begin{aligned} V^+_{m,INV}(n+1) &= \frac{C_{soma}V^+_{m,DIR}(n) + C^-_{syn}V^-_{m,DIR}(n)}{C_{soma} + C^-_{syn}} \\ V^-_{m,INV}(n+1) &= \frac{C_{soma}V^-_{m,DIR}(n) + C^+_{syn}V^+_{m,DIR}(n)}{C_{soma} + C^+_{syn}} \end{aligned} \tag{2}$$

namely in terms of the voltages $V^+_{m,DIR}(n)$, $V^-_{m,DIR}(n)$ before the spike. Upon the arrival of a second spike at the instant $(n+2)$ which brings the capacitive network back to the direct divider configuration in fig. 3a, the new voltages $V^+_{m,DIR}(n+2)$, $V^-_{m,DIR}(n+2)$ can be similarly expressed as

$$\begin{aligned} V^+_{m,DIR}(n+2) &= \frac{C_{soma}V^+_{m,INV}(n+1) + C^+_{syn}V^-_{m,INV}(n+1)}{C_{soma} + C^+_{syn}} \\ V^-_{m,DIR}(n+2) &= \frac{C_{soma}V^-_{m,INV}(n+1) + C^-_{syn}V^+_{m,INV}(n+1)}{C_{soma} + C^-_{syn}} \end{aligned} \tag{3}$$

By substituting $V^+_{m,INV}(n+1)$ and $V^-_{m,INV}(n+1)$ from eq. (2) into eq. (3) and then recalling the definition of the membrane potential $\Delta V_m = (V^+_m - V^-_m)$, we readily obtain:

$$\Delta V_{m,DIR}(n+2) = \Delta V_{m,DIR}(n) \cdot \frac{({C_{soma}}^2 - C^+_{syn}C^-_{syn})^2}{(C_{soma} + C^+_{syn})^2(C_{soma} + C^-_{syn})^2} \tag{4}$$

The $\Delta V_{m,DIR}(n+2)$ in eq. (4) depends on the synaptic weight, which is encoded in the difference between C^+_{syn} and C^-_{syn}.

For $C_{soma} \gg C^+_{syn}, C^-_{syn}$, eq. (4) can be simplified at the first order with respect to $(1/C_{soma})$, which leads to:

$$\Delta V_{m,DIR}(n+2) - \Delta V_{m,DIR}(n) \approx -2 \cdot \Delta V_{m,DIR}(n) \cdot \frac{C_{syn}}{C_{soma}} \tag{5}$$

The $\Delta V_{m,DIR}(n+2)$ expression in eq. (5) is now independent of the synaptic weight, in fact we recall that $C_{syn} = (C^+_{syn} + C^-_{syn})$ is the total capacitance of the synapse and it is independent of the synaptic weight.

The proportionality between the $[\Delta V_{m,DIR}(n+2) - \Delta V_{m,DIR}(n)]$ and $\Delta V_{m,DIR}(n)$ expressed by eq. (5) results in an exponential decay of ΔV_m to zero over the arrival of many spikes reconfiguring the capacitive dividers, provided that the SW is zero (i.e. $C^+_{syn} = C^-_{syn}$) so as to avoid the direct influence of the synaptic weight on ΔV_m.

In our design, these conditions are fulfilled by the clock spikes governing the neuron leakage. In fact, the charge transferred between the two C_{soma} every clock period T_{CLK} can be interpreted as an average current equal to $\frac{1}{2}Q/T_{CLK}$. In turn, such a current can be equivalently described in terms of a resistance R_{eq} equal to

$$R_{eq} = 2 \cdot \frac{T_{CLK}}{C_{syn}} \tag{6}$$

between the two C_{soma} capacitors, finally creating an RC circuit with a time constant given by:

$$\tau_{eq} = R_{eq} \cdot \frac{C_{soma}}{2} = T_{CLK} \cdot \frac{C_{soma}}{C_{syn}} \tag{7}$$

This is the behaviour summarized in Fig.6(e),(h) of the main paper.

Supplementary Note 4: Optimization of the Transmission Gates

In our architecture, the resonance between the inductive driver and the capacitive word-line generates a sinusoidal current. Given that each input spike triggers a half-period of the resonant oscillation, the conduction losses per synaptic operation for a minimum length MOSFET can be written as:

$$ESOP_{cond} = \frac{R_{DS}}{W}\frac{I_{pk}^2}{4}\frac{1}{f_{LC}} \tag{8}$$

where R_{DS} [Ω m] is the triode resistance per unit width of the minimum length MOSFET, W is the channel width and I_{pk} is the peak amplitude of the current.

In principle, the conduction losses can be arbitrarily reduced by enlarging W, however, this would also increase the energy required to drive the gate capacitance. In fact, during a half-period of the resonant oscillation, the TGs undergo both a charging and a discharging of their gate capacitance, thus their gate-driving energy is readily given by $ESOP_{driv}{=}W \cdot C_G{V_{DD}}^2$, where C_G [F/m] is the effective gate capacitance per unit width of a minimum length MOSFET.

Because the energy dissipated in the switches is inversely proportional to W, while their conduction losses are inversely proportional to W, there exists an optimum channel width that, for a given frequency f_{LC}, can minimize the overall dissipated energy. This is readily obtained by setting to zero the derivative with respect to W of ($ESOP_{cond}{+}ESOP_{driv}$), which leads to:

$$W_{opt} = \frac{1}{2}\frac{I_{pk}}{V_{DD}}\sqrt{\frac{R_{DS}}{C_G}\frac{1}{f_{LC}}} \tag{9}$$

and corresponds to $ESOP_{cond}^{opt}{=}ESOP_{driv}^{opt}$.
By substituting W_{opt} in the expression for $ESOP_{cond}^{opt}$ and $ESOP_{driv}^{opt}$, the optimum energy per spike event can be written as:

$$ESOP_{tot}^{opt} = ESOP_{cond}^{opt} + ESOP_{driv}^{opt} = V_{DD}I_{pk}\sqrt{\frac{R_{DS}C_G}{f_{LC}}} \tag{10}$$

As it can be seen in eq. (9), for a given f_{LC} and given values of the technological parameters R_{SD} and C_G, the optimum sizing of each transistor is proportional to the peak I_{pk} of the corresponding sinusoidal current. Hence, the size of each transmission gate (TG) was optimized by first following the TG tree sketched in fig. 4a to estimate the corresponding I_{pk}, and then by substituting I_{pk} in eq. (9).

As for the calculation of the I_{pk} of each transmission gate, we recall that the amplitude of the inductor current is proportional to the resonance frequency and can be expressed as $I_{L,pk} = \pi C_{WL}V_{DD}f_{LC}$. Such a current flows through the word-line selector TG-WL and is then distributed among all the synapses served by the WL.

Internally to each synapse, the current gets further divided among the branches of the capacitor bank. In particular, the current through the TG-SYN of the i-th capacitor ($C_{syn,i}$=$2^i C_{LSB}$) is $I_{syn,i}$=$2^i I_{LSB}$, with i ranging from zero (LSB) to $N_{bit}-1$ (MSB). In our design, the synaptic weight has a resolution of 8 bits, hence the LSB current I_{LSB} is 1/128-th of the current flowing through the MSB capacitor. The TG-SYNs currents on the two output branches of the capacitor bank depend on the synaptic weight, and they are finally forwarded to the soma capacitors by a TG-SOMA.

Hence, in order to simplify the optimization of the TG-SYN sizing, we assumed that all synapses have a null synaptic weight. This results in an equal distribution of the current first among all synapses, and then among the TG-SOMAs. Moreover, we also assume that the current gets split evenly also among the n-MOS and p-MOS of each TG.

Figure 4b shows the optimization of the n-MOS of the TG-WL at f_{LC}=100 kHz. Such optimization has been carried out also for all the other TGs and resonance frequencies ranging from 10 kHz to 10 MHz, and the resulting optimum widths of their n-MOS are summarized in fig. 4c as a function of f_{LC}. As can be seen, the TG-WL (red) drives the highest current and so requires the largest optimal width, followed by the TG-SYN of the MSB and by the TG-SOMA (orange). On the other hand, the TG-SYNs of the lesser significant bits (shades of green) rapidly saturate to the minimum channel width allowed by the technology, which in our case is 220 nm.

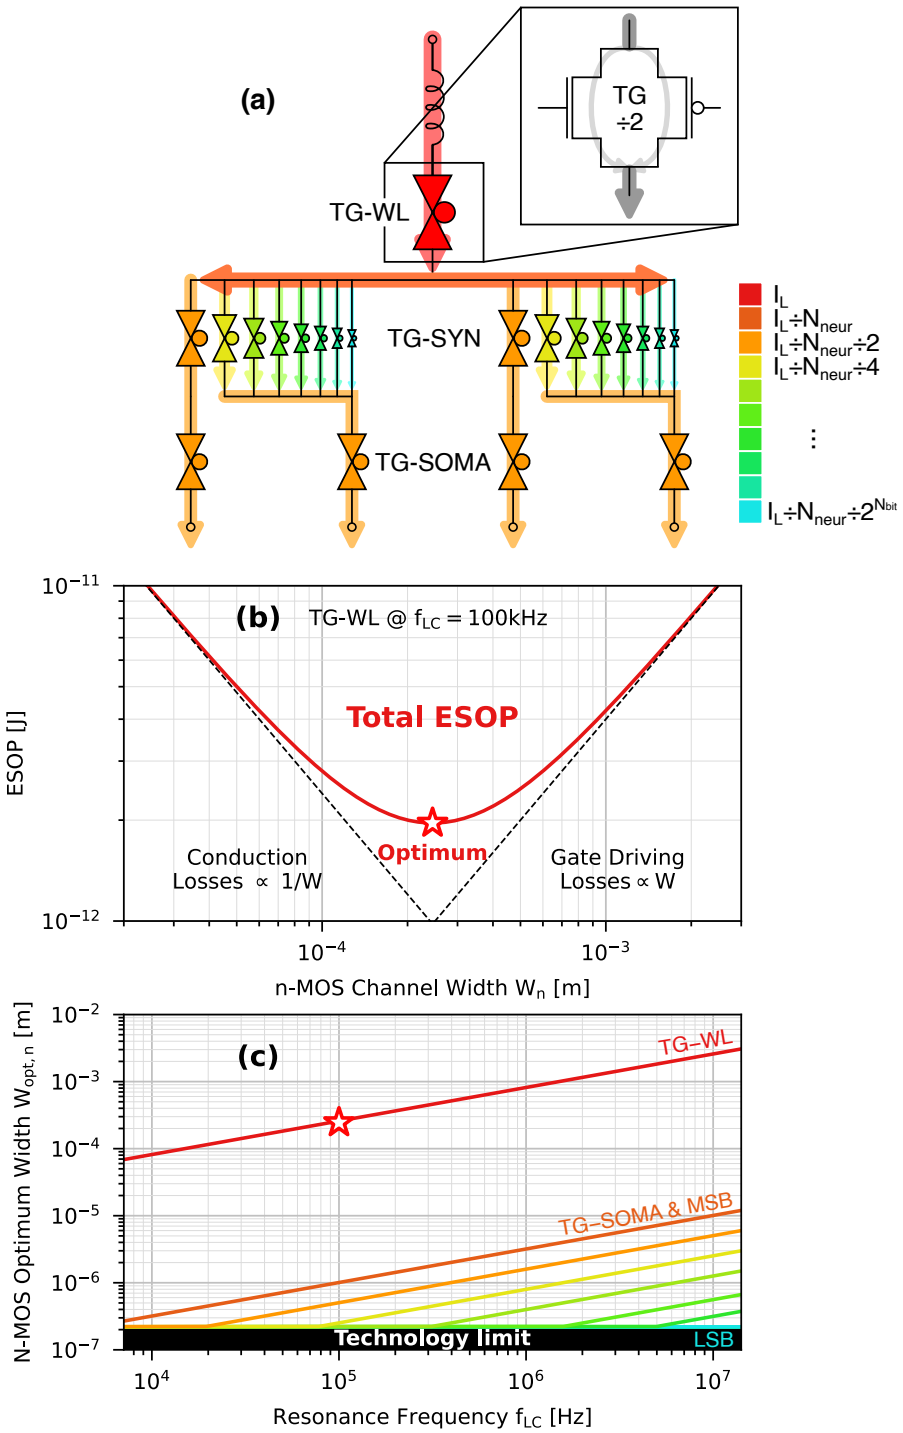

Supplementary Figure 4: Optimization of the Transmission Gates. a Sketch showing how our optimization assumes the inductor current splits among the transmission gates involved in the adiabatic energy transfer. b Overall energy per synaptic operation dissipated by the n-MOS of the TG-WL either due to the conduction losses ($\propto 1/W_n$) or due to the gate-driving losses ($\propto W_n$) plotted versus the corresponding device width W_n, and for a resonance frequency f_{LC}=100 kHz. An optimum device width is clearly observed, which is identified by the W_{opt} in eq. (9). c Optimum channel width for the different classes of TGs as a function of the resonance frequency.

Supplementary Note 5: Energy Breakdown

As already mentioned in the "Methods" section of the main paper, the energy performance of our architecture has been thoroughly evaluated by separating the energy consumption owing to each of its sub-circuits. So, in this section, we report and discuss the breakdown of such energy contributions.

First, we recall that all the energy figures are here expressed in terms of *Energy per Synaptic Operation* (ESOP), namely by normalizing the total dissipated energy to the number of input spikes (both neurons and clock spikes), and then to the number of synapses served by each sub-circuit.

The breakdown of the dynamic ESOP is reported in fig. 5a as a function of the resonance frequency f_{LC}. As it can be seen, the dynamic consumption of the logic in the word-line controller (yellow) and in the crossbar (orange, including the logic controlling the synapses - SYN - and the clock forwarders - CLK-FWD), increases with f_{LC} because the transmission gates (TG) become larger due to the frequency-dependent optimization of the TG sizing described in section Supplementary Note 4:.

The pie charts in figs. 5b to 5e report the breakdown of the contributions to the ESOP at the Minimum Energy Point (MEP), which falls at a resonance frequency 500 kHz as it is shown in fig. 3j of the main paper. In particular, fig. 5b illustrates the proportion between the dynamic consumption due to the conduction losses (c), to the logic (d), and to the static energy (e).

In fig. 5c, the losses in the tristate (pink) account for both the dissipation in the output stage of the Toggle Flip-Flop (T-FF) and those in the TG-TRS (see fig.2c in the main paper). Such losses are non-adiabatic and so depend only on the amplitude of the voltage error caused by the incomplete transition of the WL. Actually, such correction also causes non-adiabatic dissipations in both the TG-SYNs and TG-SOMAs that add up to their adiabatic losses occurring during the integration phase.

On the contrary, the TG-WL (green) does not take part in the correction of the WL, thus the corresponding conduction losses are purely adiabatic. Despite this, the TG-WL always dominates over the other TGs because it carries the entire current that is supplied by the inductor to the WL, as opposed to the TG-SYNs (blue) and TG-SOMAs (purple) which carry just a small fraction of said current, as shown in fig. 4a of this document.

In fig. 5d, the logic consumption of the crossbar (orange) and of the neuron comparator (brown) inevitably depend on the actual spiking activity of the neurons. In fact, first the profile of the membrane potential directly affects the ESOP of the "FIRE" comparator (see section Supplementary Note 2: and fig. 2f in this document), because the comparator is active only when the neuron is in an excited state, namely when we have $\Delta V_m > 0$. Second, both the refractory period (which is related to the number of fired output spikes) and the neuron leakage govern how many times the SR latches switch to toggle the REFR and LEAK signals (see fig. 2f). These, in turn, force the synapse controllers to change their current synaptic weights by re-programming the TG-SYNs of the capacitor banks, thus inevitably dissipating more energy.

Finally, in fig. 5f it is shown the breakdown of the dynamic ESOPs when the system is operated in non-adiabatic mode (refer to the "Adiabatic and Non-Adiabatic

Operating Modes" paragraph of the "Methods" section in the main paper). The overall energy consumption is 4.2 pJ and, as expected, does not depend on the resonance frequency. The energy is clearly dominated by the conduction losses in the transmission gates (marked in pink, purple and blue in the figure), which sum up to 4 pJ. It is worth noting that the TG-WL (green) does not contribute to the conduction losses in the non-adiabatic mode, because the energy is not supplied by the inductive driver but instead by the T-Flipflop, which bypasses the WL selector as shown in fig. 2c of the main paper.

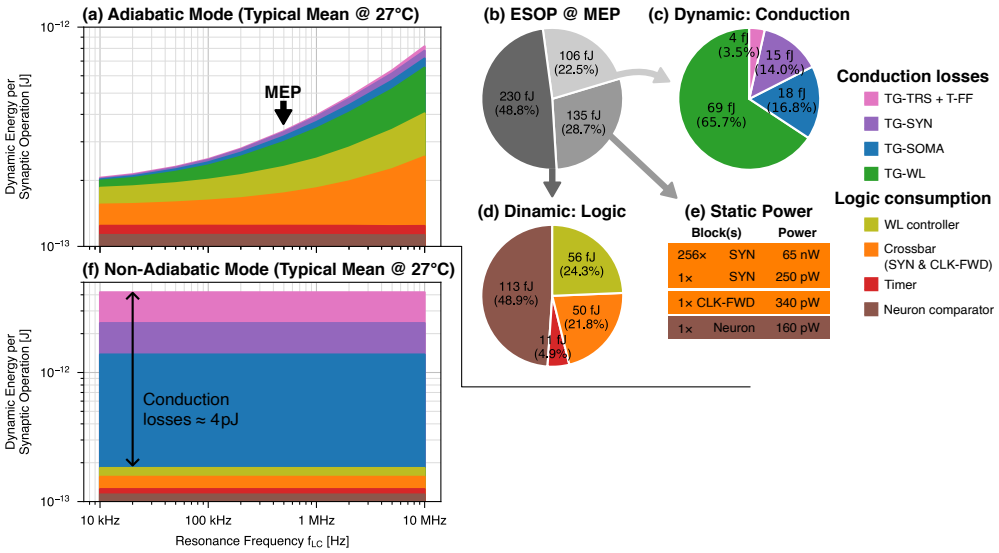

Supplementary Figure 5: **Breakdown of the Energy per Synaptic Operation.** The figure reports all the individual ESOPs owing to the conduction losses in the transmission gates (pink to green), and in the logic serving different sub-circuits (yellow to brown). The ESOP of the crossbar (orange) combines the consumption of both the synapses (SYN) and clock forwarders (CLK-FWD). a Frequency scaling of the dynamic ESOPs in adiabatic mode. The Minimum Energy Point (MEP) falls at f_{LC}=500 kHz. b Breakdown of the ESOP at the MEP into its constituents due to the dynamic conduction losses (c), dynamic consumption of the logic (d) and the static dissipation (e). f Dynamic ESOPs in non-adiabatic mode.

Supplementary Reference

[1] Analog Devices, Inc., 2001, *INL/DNL Measurements for High-Speed Analog-to-Digital Converters (ADCs)*, https://www.analog.com/en/resources/technical-articles/inldnl-measurements-for-types-of-highspeed-adcs.html