

Circuit Improvements to Q2LAL, S2LAL, and SCRL

Zettaflops, LLC Technical Report ZF012 v2

Erik P. DeBenedictis
Zettaflops, LLC
Albuquerque, New Mexico, USA
August 2, 2022

Abstract—This document identifies a circuit optimization for Static 2-Level Adiabatic Logic (S2LAL), Quiet 2-Level Adiabatic Logic (Q2LAL), and with some applicability to SCRL. When these logic families are used to create shift registers, about 30% of the transistors are redundant. The basic concept is explained, including how to eliminate the redundant transistors and discussion of the effect on simulation and layout.

Keywords—*adiabatic computing; reversible computing; CMOS; cryo-CMOS; Static 2-Level Adiabatic Logic; S2LAL; Quiet 2-Level Adiabatic Logic; Q2LAL; Split-rail Charge Recovery Logic*

I. BACKGROUND ON ADIABATIC LOGIC

This document is applicable to at least three existing adiabatic logic families, notably SCRL [1], S2LAL [2], and Q2LAL [3]. Interested readers are referred to the references.

II. REDUNDANT GATES

The observation is that when these logic families are used to create shift registers, about 30% of the transistors are redundant. Since shift registers can be used for data storage, which is a fundamental resource, removing the redundant transistors may shrink practical circuits considerably.

The observation and result are illustrated in Fig. 1.

Fig. 1a shows the reversible circuit framework that is the basis of SCRL, S2LAL, and Q2LAL. Structurally, it is like a chain.

Fig. 1b is the result of taking Fig. 1a and horizontally shifting the bottom row to the right until the rectangular transmission gates line up vertically. The triangular adiabatic amplifiers [4] are then rotated so they all point down. Switches are added. If the switches are in the up position, the circuits in Fig. 1a and b are the same.

The reader will see that the adiabatic amplifiers on each branch of the switch have the same data input and clock. The output of an adiabatic amplifier is determined solely by its input—i. e. there is no tri-state output whose value is determined by an external source. This opens the possibility that a switch can move between positions without changing the circuit's functional behavior.

For the possibility to be realized, the adiabatic amplifiers must have the same function and operate on the same data. For example, if F_5 was an NAND function while R_3 was the XOR

function, the output voltages would differ at times, resulting in an electrical short circuit and different functional behavior. If the circuit is a shift register, the F's and R's are either all inverters (SCRL) or all non-inverting buffers (S2LAL and Q2LAL). Hence flipping the switch will not affect a shift register.

In fact, Fig. 1b shows two copies of a hybrid or generalized stage. Adjacent stages have a second adiabatic amplifier when the switch is up but an extra wire when the switch is down.

Fig. 1c and d show two renditions of the alternative circuit. The two circuits are functionally identical, but flipping alternate stages upside down eliminates the crossovers. The circuit in Fig. 1d could be described as alternating inverted T's.

III. RESOURCE CONSUMPTION

The observation in the previous section can be exploited to reduce resource consumption. There will be no change in circuit function since the method described is simply the elimination of redundant transistors. This implies that the original circuit in Fig. 1a could be intermixed with the circuits in Fig. 1b or c on a single chip. For example, a circuit could comprise all members of a universal logic gate set intermixed with shift register memory.

There will be a tradeoff regarding complexity and distance. Two instances of the original circuit connect using two wires or rails, such as S and $-S$, but the alternate circuit would have four wires, S , $-S$, Q , $-Q$. The alternate circuit is favored when cells are tightly packed. In this case, the wires may be short or nonexistent (such as when standard cells abut).

While the observation yields circuits of identical logical behavior, electrical behavior could be different. The two adiabatic amplifiers will need to drive the same total current at the same times, creating an argument that combining two adiabatic amplifiers should yield a single amplifier with double-size transistors.

However, the decision to combine is only reasonable when the interconnecting wires are short. This leads to a counterargument that the combining will usually occur in situations where the transistors are of minimum size anyway—so there will be a tendency to combine circuits with minimum size transistors into a circuit with minimum size transistors. In this situation, the alternate circuit would have fewer transistors,




Fig. 1. Circuit simplification applicable to shift registers. (a) The familiar circuit diagram. (b) Shift the top row relative to the bottom row so the transmission gates line up. This will cause F_{n+2} to have the same inputs as R_n , so these circuits are redundant and merged, resulting in downward pointing triangle with labels $F_{n+2}R_n$. This is only valid if F_{n+2} the same function as R_n and the inputs are the same – which is true for shift registers. (c) If desired, the crossovers can be removed by vertically flipping every other stage.

consume less area, have less capacitance to drive (and hence have lower power), and would have a tendency (which we observe) to have less wiring congestion.

IV. DESIGN CONSIDERATIONS

The first step would be to construct the functional definition of a circuit using the circuits in Fig. 1b instead of Fig. 1a. Connections between circuits would be lines representing four wires $\pm S$ and $\pm Q$, rather than just two $\pm S$ wires.

In the second step, the designer or an automatic design tool would search for redundant adiabatic amplifiers and then

assess whether the savings due to reduced transistor count would be more or less than the cost of the extra wire. The switch would be moved to the down position if the alternate circuit was correct and advantageous.

One option for a third step would be to replace all circuits with one of two standard cells based on the switch setting. This would allow shift register storage arrays to have the highest layout density without constraining arbitrarily laid out circuitry.

A second option for the third step would be to a CAD tool to circuit optimization based on the setting. In other words, the

CAD tool would search for unused adiabatic amplifiers and wires and eliminate them automatically. The CAD tool could then perform automatic layout of the simplified circuit.

V. S2LAL SIMULATION TEST

An S2LAL spice netlist was enhanced to support the alternative circuit. The author had previously constructed an S2LAL [2] simulation to assess dissipation and sought to add a software-type flag that would control the merging of adiabatic amplifiers at netlist read-in time. Comparing the dissipation of the original and alternate circuits is described below.

The enhancement was tedious but understanding may be enhanced by the two-step development in Fig. 2.

Fig. 2a is a dual-rail version of [2, Fig. 5 and Fig. 6], although the reader should understand that the numerical subscripts can be shifted mod 8. The **green** dotted lines are in anticipation of the enhancement and should be taken as wires.

However, it should be noted that there are two basic spice subcircuits. The first contains left-pointing adiabatic amplifiers and clamps to ground. The second contains the right-pointing adiabatic amplifiers and basically has all nFETs replaced with pFETs, power with ground, and vice-versa (Q2LAL and SCRL would only require one basic circuit).

The circuit diagram in Fig. 2b illustrates the enhancements by the replacement the **green** dotted lines with **red** wires and switches.

Basically, the interconnect between stages was originally one S_n signal per rail. The connections were enhanced to two signals $\{S_n, Q_n\}$ per rail. The Q_n 's transport the output of the rightward-pointing (computing) adiabatic amplifier backwards to the previous stage.

The connection between the leftward-pointing (uncomputing) adiabatic amplifier and its transmission gate had to be replaced by a switch of sorts. While spice has a switch component, the need is for a switch that operates during netlist read-in time. The method is to specify both the adiabatic amplifier output and the transmission gate it connects to as subcircuit parameters and then connect them differently depending on which circuit is needed.

The vertical switch position illustrated in Fig. 2b is used when the parent cell needs the alternate circuit. In this case, the parameter connecting to the pass gate in the uncomputing direction (bottom two rows) is connected to the output of the adiabatic amplifier of the next stage in the computing direction (top two rows).

If the parent cell needs the original circuit behavior, it shorts the two parameters together.

The parameters are always shorted together in the top two rows and routed to the previous stage as the Q_n signal.

It should be noted that either the **red** wiring (except the switch) or the ~~adiabatic amplifiers~~ with strikethrough can be deleted.

A 10 μ s simulation with a 1 MHz clock was performed with the transistor models in the Sky130 PDK. The simulation

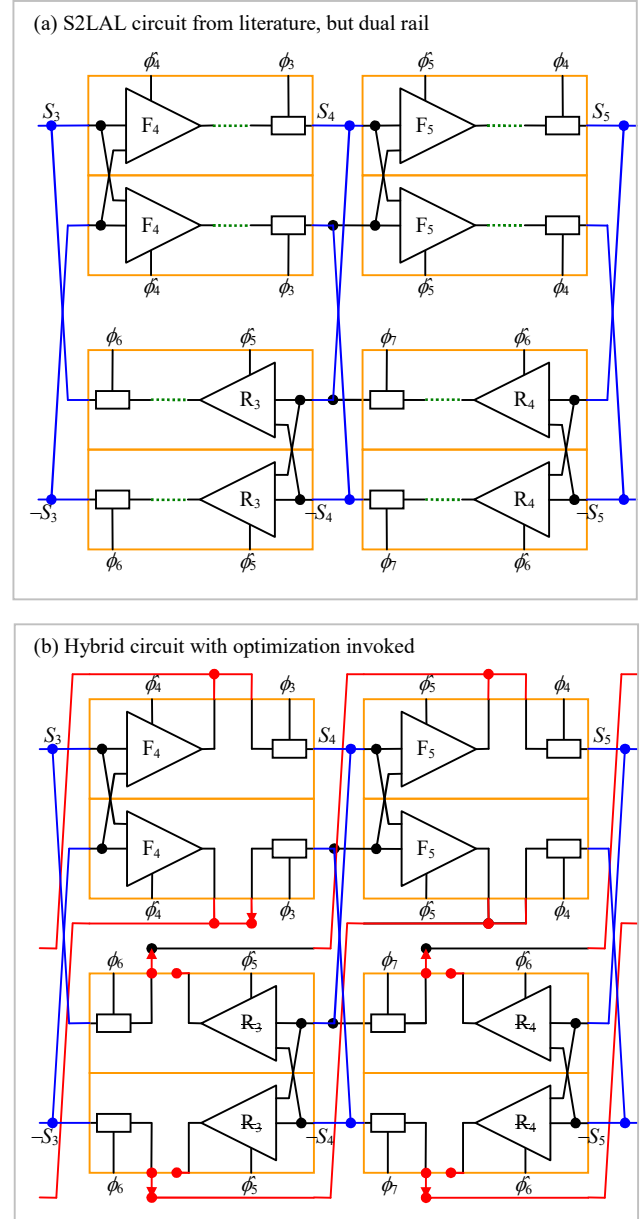


Fig. 2. (a) S2LAL from literature; assumes green dotted lines are wires. (b) Hybrid with/without merging. **Green** dotted lines have been deleted and **red** lines added; ~~adiabatic amplifiers~~ can be deleted. Switch currently positioned for merging, but flipping switches results in circuit in the literature. See text for spice structure.

was divided into 1 μ s startup period (deemed unreliable) and 9 μ s of data collection. Comparing the latter 9 μ s of the simulations, the alternate circuit had 81.9% as much dissipation. The ~~adiabatic amplifiers~~ with strikethrough were not present for this simulation.

VI. LAYOUT TEST

This section includes general discussion about layout issues based on a test Q2LAL layout using the Sky130 PDK. (Note that this section is about Q2LAL where the previous one was S2LAL.)

To assure that power supply loading is independent of whether the system is processing 0s or 1s, the two rails, i. e. S and $-S$, should have the same layout. This design objective can be satisfied if the floor plan has two copies of the circuit in Fig. 1c positioned vertically, but flipped vertically as shown in Fig. 3a. This approach leads to symmetric layout by design—except for a crossover.

The crossover is needed because an adiabatic amplifier requires both S and $-S$ input signals but only produces one output Q' . The strategy is for either the S or $-S$ signal to enter on the left and then pass the other signal through the common side. Thus, each rail gets both signals, but from different sides. This requires the crossover block shown. Such a crossover cannot be created with a vertical reflection and has to be hand edited into the layout when the two-rail cell is created.

Thus, compared to Fig. 1b, Fig. 3b illustrates two rails but only one stage. Compared to the circuits in Fig. 2, Fig. 3b illustrates just one stage.

VII. ADVANTAGE

The reduction in the number of transistors is enumerated in this section. There is also a reduction in wiring congestion, but this is not easily summarized.

An S2LAL stage [2] comprises two adiabatic amplifiers of 3 transistors each plus two transmission gates of 2 transistors each, for a total of 10 transistors. The simplification reduces the circuit to 7 transistors, a reduction of 30%

An Q2LAL stage [3] comprises two adiabatic amplifiers of 4 transistors each plus two transmission gates of 2 transistors each, for a total of 12 transistors. The simplification reduces the circuit to 8 transistors, a reduction of 1/3.

An SCRL stage [1] comprises two split-rail invertors of two transistors each plus two transmission gates of 2 transistors each, for a total of 8 transistors. This simplification reduces the circuit to 6 transistors, or a reduction of 25%. The SCRL circuits known to the author would need to have their clocking changed to accommodate this circuit simplification, so the change is less transparent than for the other families.


ACKNOWLEDGMENT

Michael P. Frank developed S2LAL and a consistent terminology [2], both of which became a starting point for this work. This document uses Mike's terminology, including diagrams, with his permission.

REFERENCES

- [1] Saed G. Younis. *Asymptotically Zero Energy Computing Using Split Level Charge Recovery Logic*. No. AI-TR-1500. Massachusetts Institute of Technology Artificial Intelligence Laboratory, 1994.
- [2] Frank, Michael P., et al. "Reversible Computing with Fast, Fully Static, Fully Adiabatic CMOS," *2020 IEEE International Conference on*

(a) Layout floor plan



(b) Test layout

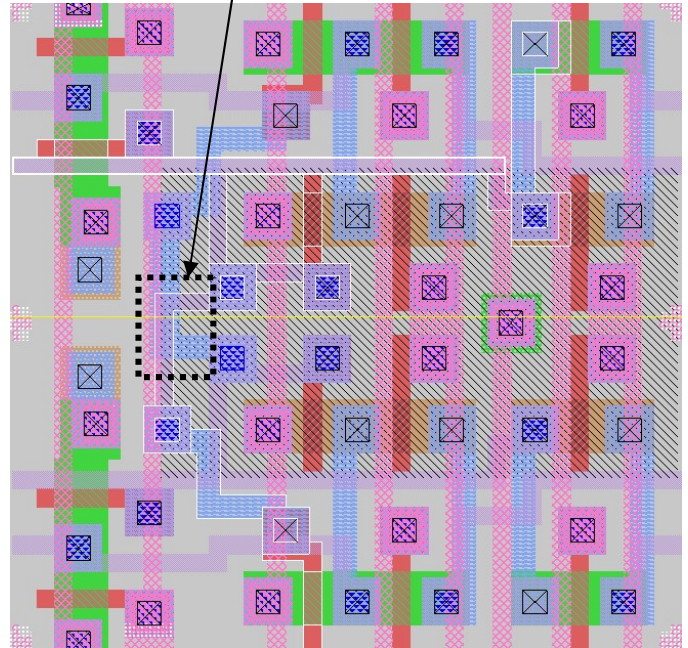


Fig. 3. (a) Test layout floor plan for Q2LAL with vertical reflection except for a crossover. (b) Two copies of layout. Information flow is left-right and there is a vertical overlay of (pink) level-2 metal containing (from left-to-right) GND, Clamp, ϕ_0 , ϕ_1 , ϕ_2 , V_p , ϕ_3 , ϕ_4 .

Rebooting Computing (ICRC), Atlanta, GA, USA, 2020, pp. 1-8, doi: 10.1109/ICRC2020.2020.00014.

- [3] Erik P. DeBenedictis, *Quiet 2-Level Adiabatic Logic*. Zettaflops, LLC technical report ZF009, online at <https://debenedictis.org/erik/CATC/Q2LAL.pdf>
- [4] Athas, William C., et al. "Low-power digital systems based on adiabatic-switching principles." *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 2.4 (1994): 398-407.