Adv. Radio Sci., 5, 291–295, 2007
www.adv-radio-sci.net/5/291/2007/

Advances in Radio Science

Design issues of arithmetic structures in adiabatic logic

Ph. Teichmann, J. Fischer, F. Chouard, and D. Schmitt-Landsiedel

Institute for Technical Electronics, Technical University Munich, Munich, Germany

Abstract. Since adiabatic logic uses a supply that incorporates both supply voltage and clock signal in one line, adiabatic logic systems have a built-in micro-pipelined architecture. Considering this fact, different design constraints have to be observed compared to static CMOS designs. Complex arithmetic building blocks, like multipliers, mainly consist of adders. Therefore, a comparison of adder structures is performed. Based on these results, multipliers and complex systems can be built. A Discrete Cosine Transformation (DCT) is taken as example for an arithmetic system. Comparing an adiabatic logic implementation of a DCT to its static CMOS counterpart, a significant saving factor of more than 10 can be achieved with the adiabatic system.

1 Introduction

Adiabatic logic circuits are known for offering an energy dissipation less than the $E=\frac{1}{2}CV_{DD}^2$ limit in static CMOS circuits. The Positive Feedback Adiabatic Logic (PFAL) (Vetuli et al., 1996) family has proven to be a reasonable choice for todays VLSI integration (Amirante, 2004; Fischer, 2006) as it is ultra low-power and robust against parameter variations. A chain of PFAL gates is operated via a four-phase clock which acts as power supply and clocking line for the gates, and it is therefore called power-clock. As consequence, cascaded gates form a micropipeline which is an inherent property of adiabatic logic. This fact has to be observed if complex systems are investigated.

Arithmetic structures are basic building blocks in a variety of digital signal processing tasks. E. g. adders are used to build more complex arithmetic functions like multipliers, filters, CORDICs etc. Many proposals for adders focusing on different design goals like power, area and performance have been presented for static CMOS circuits (Koren, 2002; Zimmermann, 1997; Beaumont-Smith and Lim, 2001). But PFAL's inherent properties make it desirable to analyze the known arithmetic structures with respect to their applicability in adiabatic logic.

This work will first introduce the idea of adiabatic logic and the micropipeline in Sect. 2 followed by the results gained from the considerations on adders and multipliers in Sect. 3 followed by a case study of a DCT in Sect. 4. The results are concluded in Sect. 5.

Correspondence to: Ph. Teichmann
(teichmann@tum.de)

2 Brief introduction of adiabatic logic

In static CMOS the fundamental limit for the energy dissipation per switching event of $E=\frac{1}{2}CV_{DD}^2$ can only be lowered by reducing the capacitance C or by scaling of the supply voltage V_{DD}. To overcome this limit adiabatic circuits have been proposed, that trade frequency for energy. By lowering the operating frequency, theoretically an asymptotic convergence to zero dissipation can be reached. The energy dissipation in a PFAL adiabatic logic gate is described by $E=\frac{RC}{T}CV_{DD}^2$, where R is the path resistance, C is the capacitance at the output, T is the rise/fall time of the power-clock signal. A low threshold voltage V_{th} on the one hand reduces the path resistance R due to an improved gate overdrive voltage, on the other hand it increases the leakage currents which limit the energy dissipation in the mid- and low-frequency range. PFAL circuits show a minimum in energy dissipation at the crossing point of the adiabatic losses ($\propto f$) and the leakage losses ($\propto \frac{1}{f}$). For a 130 nm CMOS process this minimum is retrieved around 100 MHz.

The PFAL logic family uses a four-phase power-clock signal, that acts as the supply voltage and the clocking signal (see Fig. 1), and inherently imposes the circuits beeing operated in a pipelined fashion. The pipeline style gives a first conception how a system in adiabatic logic should be implemented, as overhead due to synchronization degrades the savings gained through the application of adiabatic logic, furthermore pipelining leads to a rise in latency.

Published by Copernicus Publications on behalf of the URSI Landesausschuss in der Bundesrepublik Deutschland e.V.

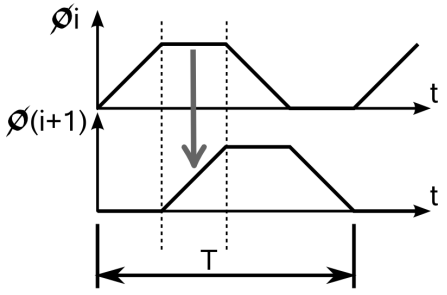

Fig. 1. Two preceeding phases Φi and $\Phi(i+1)$ of an adiabatic four-phase power-clock are shown. When Φi is in its stable high phase, the data is evaluated in the succeeding stage, indicated by the arrow.

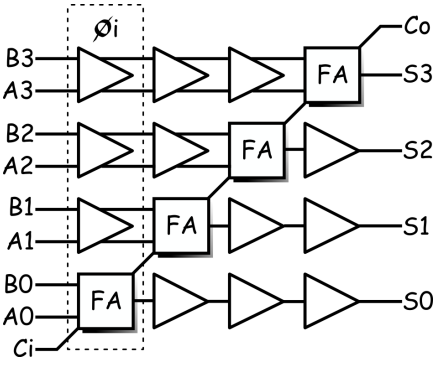

Fig. 2. A N=4 ripple-carry adder in adiabatic logic. Notice the overhead due to the synchronization buffers. Each input bit of A and B has to be buffered, leading to two buffers per bit position. The dashed box shows a clock domain of the power clock Φi.

3 Binary adder structures

The inherent properties of adiabatic logic must be taken into account in the design of adiabatic arithmetic structures. Large overheads in structures result in suboptimal designs.

One advantage of PFAL is its dual-rail signal representation, so a signal is inverted without additional gates. Subtractors are easily obtained from adders, by swapping the signal rails of the subtrahend and feeding a logic one into the carry-in signal of the adder. The results for the adders gained in this section are valid for the subtractors as well.

Basically two groups of adders exist, namely the carry-propagate and the parallel-prefix adders (Sklansky, 1960), (Kogge and Stone, 1973), (Ladner and Fischer, 1980), (Brent and Kung, 1982) and (Han and Carlson, 1987). The ripple-carry adder (RCA) is the simplest implementation of a carry-propagate adder, using N full-adder (FA) cells in static CMOS, where N is the input width of the data words. But for adiabatic logic, the rippling of the carry signal leads to an overhead of synchronizing buffer

stages of $\mathcal{O}(N^2)$, making the ripple-carry structure in adiabatic logic (Fig. 2) an inproper choice if we talk about high bit width N. Additionally, the adiabatic ripple-carry adder (RCA) utilizes $N/4$ clock cycles, leading to a relatively fast rising latency. Only for several suqsequent addition operations, the RCA can be used in a nested way, as pictured in Fig. 3, can be applied efficiently, as the absolute overhead of synchronizing buffers per arithmetic operations remains almost constant. Especially for butterfly structures, as used in the Discrete Cosine Transformation (DCT) in Sect. 4, the nested RCA is advantageous. Other approaches try to reduce the critical path, i.e. the path of the rippling carry, to gain speed in static CMOS designs, respectively to reduce the latency in adiabatic logic. As we are talking about synchronous, pipelined designs, approaches like the carry-bypass adder cannot be applied to adiabatic logic. The carry-select scheme splits the input words into smaller groups, i.e. 2 of size $N/2$ each and pre-calculates the sums

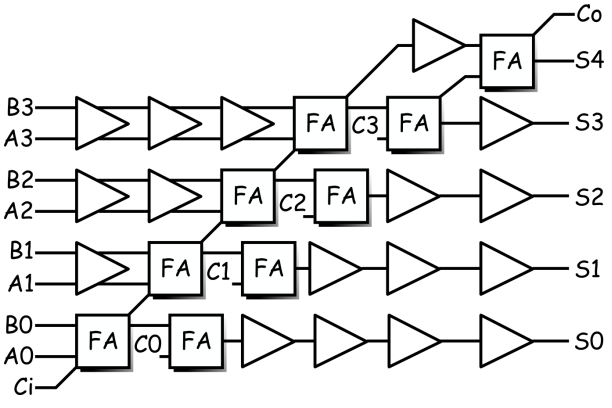

Fig. 3. For a nested RCA (here with two suqsequent addition operations) structure, the relative overhead due to synchronization is reduced. Please note, that inputs C also have to be delayed via input buffers.

for both input carry alternatives. The incoming carry from the preceeding group selects the appropriate output via a multiplexer. Such a design trades area, respectiveley energy, against speed in static CMOS and latency in adiabatic logic respectively. Additionally an overhead from the multiplexers arises for such an adder. If we take the RCA from Fig. 2 and split the adder in two groups of 2 bits, the design uses 3 blocks consisting of 2 full-adder cells and three buffers each, one 6:3 multiplexer and buffers that synchronize the output bits of the lower block to the outputs of the multiplexer. In Fig. 4 for the arrangement with a multiplexer of logic depth equal to 1 it can be seen, that this design uses more full-adders than the RCA in Fig. 2, but less buffers. If the RCA and the carry-select adders are compared for high bit width N it can be seen that the RCA suffers from a higher energy consumption. This estimation does not account for the energy

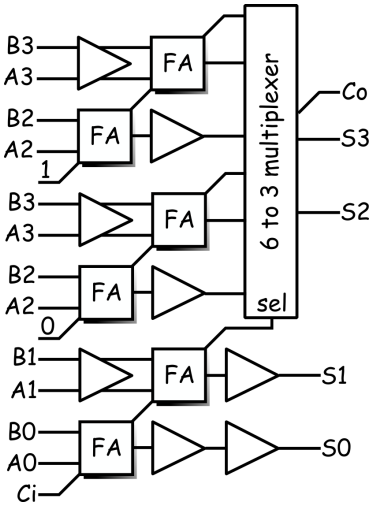

Fig. 4. A 4 bit RCA carry-select adder using a multiplexer with logic depth of 1.

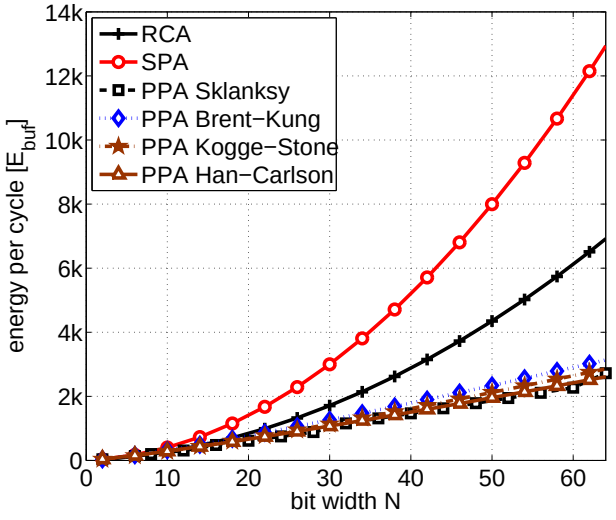

Fig. 5. Comparison of various adder structures up to 64bits. The parallel-prefix adders show the lowest energy consumption.

dissipation caused in the multiplexer. But as only $N/2$ XOR gates are used for the multiplexer, the energy consumed by the multiplexer is negligible. To avoid the synchronizing overhead observed for the RCA, parallel-prefix schemes are investigated. The parallel-prefix adders (PPA) allow to calculate the output of the m-th bit via logical combination of the inputs 0 to $mu-1$. Therefore, theoretically each output of the adder can be calculated using one complex gate. Because of practical reasons in stacking transistor devices, only a maximum of 3 or 4 inputs is combined at a time. In literature different schemes for the parallel prefix algorithm are reported, differing in logical depth and fanout. Here a variety of PPA structures (Sklansky, Brent-Kung, Kogge-Stone, Han-Carlson) are estimated to see, which adder structure is desirable in adiabatic logic. E.g. the Sklansky structure has the lowest logical depth, but the layout is supposed to be irregular. The Han-Carlson structure applies a more regular layout for the price of a higher logical depth. Additionally, the RCA structure and a serial prefix adder (SPA) are estimated in respect on their energy consumption. Estimations are performed by counting the number X_i of each class of gates i (e.g. INV, FA) and multiplying it with its corresponding mean energy value E_i. The whole system's energy dissipation E is then calculated by

$$E = \sum_i X_i E_i. \quad (1)$$

Gates are characterized by means of SPICE simulation; industrial 130nm CMOS process parameters of low-V_{th} devices are used. The gates are simulated at a frequency of 100MHz and a supply voltage of V_{DD}=1.2V. In Fig. 5 the investigated adder structures are presented. The RCA rises in energy consumption quadratically with the bit width N. The parallel-prefix adders present the lowest energy consumption. The serial-prefix adder consumes the largest amount of energy and thus is no alternative. In the detail view in Fig. 6 only the RCA and the most promising PPA structures are presented. Up to 8bit input width, the RCA is the best choice.

For the decision between Sklansky and Han-Carlson different properties have to be taken into consideration. As said before, the Han-Carlson structure has a more regular design, thus saving effort in layouting. Additionally the Han-Carlson adder has a fixed maximum fan-out of 2. For the Sklansky structure, the effort for layout will be higher, and the fanout rises with $N/2$. So the decision, which design is a suitable low-power adder structure will depend on the post-layout extraction, taking into account the rise of the energy consumption due to parasitic capacitances.

4 Discrete cosine transformation: a case study

In image processing the Discrete Cosine Transformation (DCT) is used to compress pictures. A DCT with minimal hardware amount is presented in (Heyne et al., 2006). As the DCT is a strongly parallel algorithm, that can be efficiently implemented using the adiabatic micropipeline, it is a good demonstrator to show the energy savings gained by using adiabatic logic in comparison to static CMOS. The DCT is built of adders, subtractors and buffers as can be seen in Fig. 7.

Different static CMOS implementations have been compared. Flip-flops are power-consuming components in a synchronous design in static CMOS. Different levels of pipelining have been investigated and been opposed to the adiabatic logic DCT implementation. The adiabatic reference design

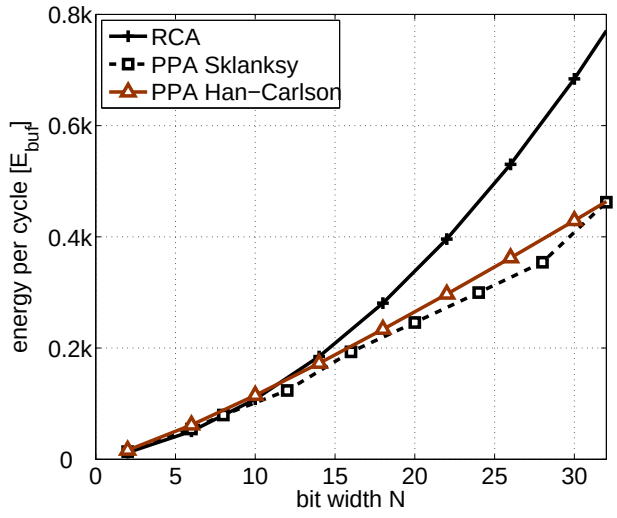

Fig. 6. Comparison of RCA and most promising PPA adders, Han-Carlson and Sklansky PPA, up to 32bits. The RCA exhibits the lowest energy consumption up to 8bit.

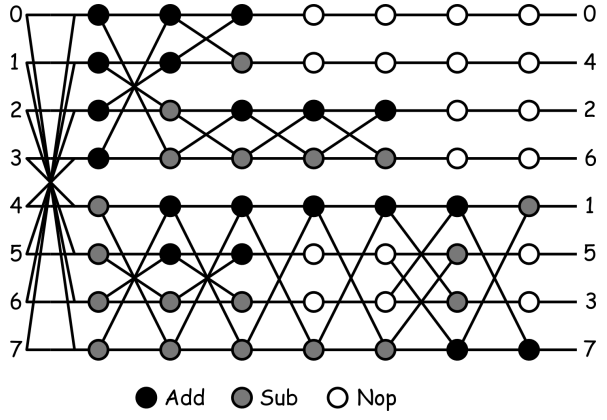

Fig. 7. Scheme of the DCT, where a Nop (No Operation) is a synchronizing buffer stage.

is implemented as nested RCA structure, thus a minimum number of FA cells are used and the overhead due to synchronization is reduced due to nesting. For static CMOS a 2D-pipelined carry-propagate adder (CPA) has been estimated first. To reduce the number of flip-flops, in **V1** to **V3** a CSA has been used to avoid skewing and deskewing at the inputs, respectively at the outputs. In **V1** a flip-flop is introduced after each CSA stage, in **V2** after every second CSA stage, and in **V3** no pipelining was introduced at all. Results are calculated according to

$$E_{AL} = X_{FA,AL} * E_{FA,AL} + X_{BUF,AL} * E_{BUF,AL} \quad (2)$$

$$E_{CMOS} = X_{FA,CMOS} * E_{FA,CMOS} \quad (3)$$

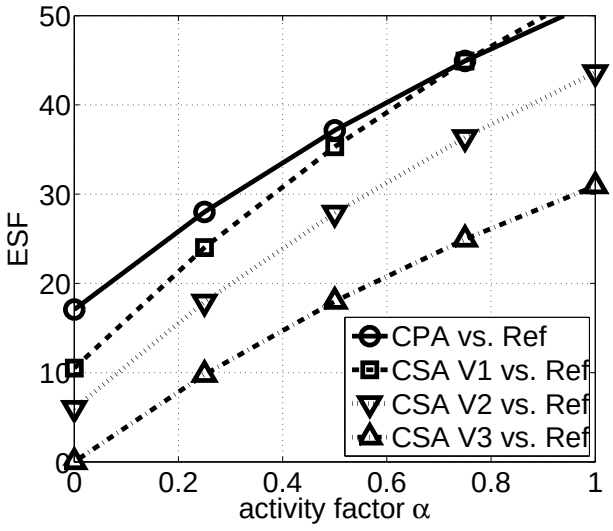

Fig. 8. Comparing the differing static CMOS DCT implementations to the adiabatic DCT. For static CMOS, first a 2D-pipelined carry-propagate adder (CPA) has been implemented and **V1** to **V3** show carry-save array (CSA) structures with different degrees of pipelining, where **V3** has no pipelining at all.

$$+ X_{INV,CMOS} * E_{INV,CMOS}$$
$$+ X_{FF,CMOS} * E_{FF,CMOS}$$

where the index AL stands for adiabatic logic, FA for full-adder, BUF for a buffer, INV for an inverter, and FF for a flip-flop.

Especially for static CMOS an important parameter for energy dissipation is the activity α, whereas in PFAL activity has a minor impact on the dissipation. To allow a fair comparison, activity has to be considered in the estimations via

$$E(\alpha) = E_{\alpha=0} + \alpha\,(E_{\alpha=1} - E_{\alpha=0})\,, \quad (4)$$

where $E_{\alpha=0}$ and $E_{\alpha=1}$ are the respective mean energy dissipation values for activity α=0 and α=1.

The results of the comparison are given in Fig. 8. All estimations are based on the simulations in an industrial 130nm CMOS process, at 100 MHz and with a V_{DD}=1.2V. The mean energy dissipation values for the used gates are summarized in Table 1 for α=0 and α=1.

In Fig. 8 we see that the best implementation for static CMOS is **V3**, that uses no pipelining. There the energy saving factor

$$\mathrm{ESF} = \frac{E_{\mathrm{diss,CMOS}}}{E_{\mathrm{diss,AL}}} \quad (5)$$

reaches its minimum value. Depending on the activity the ESF reaches up to a factor of 30 for the highest activity of the system. Applying a realistic activity value of α=30% (Ye and Roy, 2001) an ESF of factor 10 can be gained. The

dissipated energy per cycle [J]					
family	AL		CMOS		
gate	FA	BUF	FA	INV	FF
α=0	2.10f	0.24f	3.48a	1.59a	6.46f
α=1	2.67f	0.33f	37.6f	2.66f	19.0f

Table 1. Simulated energy per cycle dissipation values for adiabatic and static CMOS gates. All estimations are performed for α=0 and α=1.

generation of the four-phase clock itself is lossy, the ESF must take into account the efficiency of the oscillator. Oscillators with an efficiency of 50% have been reported. There is still room for further improvements of the oscillator efficiency. For a system level comparison the losses on the clock net in an static CMOS system have to be considered here as well. This requires detailed information on the topologie of the clock net. Predictions for high-performance systems show that almost as much energy in the clock distribution is used as in the logic itself. At least flip-flops are used at the inputs of the structure and the outputs of the structure, leading to a slight increase of the dissipation even for the adder **V3** on system level. Thus the factor of 10 presented here should also be a good indication for the savings in a complete system achievable with adiabatic logic.

5 Conclusions

In this paper the inherent properties of adiabatic logic have been analyzed for arithmetic building blocks. Adder structures were investigated and compared for the first time showing that due to the synchronization overhead of $\mathcal{O}(N^2)$ the ripple-carry adder is only suitable for bit sizes $N \leq 8$. Parallel-prefix adders are a good choice, as the reduced depth of the adder also reduces the area and the synchronization overhead and thus the power consumption. The Han-Carlson and the Sklansky architecture show the lowest energy consumption of all parallel-prefix schemes. The Discrete Cosine Transformation is an arithmetic structure making use of adiabatic logic's inherent pipelining structure. The nested ripple-carry adder scheme allows major savings against comparable designs in static CMOS. For an activity factor of 30%, savings of around factor 10 can be achieved.

Adiabatic logic is therefore a circuit family that allows the implementation of ultra low-power digital signal-processing tasks.

References

Amirante, E.: Adiabatic Logic in Sub-Quartermicron CMOS Technologies, vol. 13 of Selected Topics of Electronics and Micromechatronics, Shaker, 2004.

Beaumont-Smith, A. and Lim, C.-C.: Parallel prefix adder design, in: 15th IEEE Symposium on Computer Arithmetic, 218–225, 2001.

Brent, R. P. and Kung, H. T.: A regular layout for parallel adders, IEEE Transactions on Computers, C-31, 260–264, 1982.

Fischer, J.: Adiabatische Schaltungen und Systeme in Deep-Submicron-CMOS-Technologien, vol. 24 of Selected Topics of Electronics and Micromechatronics, Shaker, 2006.

Han, T. and Carlson, D. A.: Fast area-efficient VLSI adders, in: Proceedings of the 8th Symp. Comp. Arithmetic, 49–56, 1987.

Heyne, B., Sun, C.-C., Goetze, J., and Ruan, S.-J.: A Computationally Efficient High-Quality CORDIC based Loefller DCT, in: 14th European Signal Processing Conference (EUSIPCO), 2006.

Kogge, P. M. and Stone, H. S.: A parallel algorithm for the efficient solution of a general class of recurrence equations, IEEE Transactions on Computers, C-22, 786–793, 1973.

Koren, I.: Computer Arithmetic Algorithms, A K Peters, second edn., 2002.

Ladner, R. E. and Fischer, M. J.: Parallel Prefix Computation, J. Assoc. Computing Machinery, 27, 831–838, 1980.

Sklansky, J.: Conditional-sum addition logic, IRE Transactions, EC-9, 226–231, 1960.

Vetuli, A., Pascoli, S. D., and Reyneri, L. M.: Positive feedback in adiabatic logic, Electronics Lett., 32, 1867–1869, 1996.

Ye, Y. and Roy, K.: QSERL: Quasi-Static Energy Recovery Logic, IEEE Journal of Solid-State Circuits, 36, 239–248, 2001.

Zimmermann, R.: Binary Adder Architectures for Cell-Based VLSI and their Synthesis, Ph.D. thesis, ETH Zurich, 1997.