

High Throughput and Low Power Dissipation in QCA Pipelines using Bennett Clocking

Marco Ottavi, *Member, IEEE*, Salvatore Pontarelli, Erik DeBenedictis, *Member, IEEE*, Adelio Salsano, *Member, IEEE*, Peter Kogge, *Fellow, IEEE*, and Fabrizio Lombardi *Fellow, IEEE*,

Abstract—This paper presents a detailed analysis of an architectural pipeline scheme for Quantum-dot Cellular Automata (QCA); this scheme utilizes the so-called Bennett clocking for attaining high throughput and low power dissipation. In this arrangement, computation stages (utilizing Bennett clocking) and memory stages combine the low power dissipation of reversible computing with the high throughput feature of a pipeline. An example of the application of the proposed scheme to an XOR tree circuit (parity generator) is presented; a detailed analysis of throughput and power consumption is provided to show the effectiveness of the proposed architectural solution for QCA.

I. INTRODUCTION

Among so-called emerging technologies that have been proposed to overcome the limitations of CMOS at the “end of the technology roadmap”, Quantum-dot Cellular Automata (QCA) shows features that are very promising to achieve both high computational throughput and low power dissipation. The QCA computational paradigm [1] [2] [3] is readily suited to pipelined architectures with high speed (in the order of THz), while radically departing from the traditional nature of switch-based operation of CMOS, i.e. avoiding the movement of charge from V_{dd} to ground and the resulting energy dissipation. An operating single cell [4] and a functional logic gate have been demonstrated [5] using metal dot implementations at cryogenic temperatures. Recent advances show promising results for manufacturing atomic silicon quantum dots [6]; moreover, molecular scale QCA may make fabrication of QCA cells possible at nanometer dimensions for room temperature operation [7] [8].

In addition to great promise due to its small size and high computational speed, it has been shown that QCA has great potential for low power operation. The reversible computational paradigm is particularly well suited to QCA because Timler has shown that in a clocked, information preserving system, the energy dissipation of a QCA circuit can be significantly lower than $k_B T \ln 2$ [9]. *Reversible computation* is drawing increasing interest as a low power computation paradigm because it may overcome the fundamental power limitation

of today’s irreversible approaches based on bulk CMOS. It is foreseen [10] that in few decades, one of the main obstacle to further integrate computing will be posed by thermal considerations. As bit energies approach the absolute thermodynamic lower bound of $k_B T \ln 2$, [11], then only the energy associated to the physical information of a bit (physical entropy) may be used to encode each logical bit. This assumption relies on the relationship between information and computation, i.e. when destroying the information of a computed bit, its energy needs to be irreversibly “thermalized”, or converted to a thermal energy $Q = T\Delta S = T k_B \ln 2$ where T is the temperature, k_B is Boltzmann’s constant and ΔS is the increase in entropy related to the loss of free energy of one bit.

However, information does not need to be destroyed during computation as shown in [12]. In the Bennett scheme, intermediate results of computation are stored rather than destroyed. Once the output is computed and saved, computation can be reversed to decompute the intermediate results, i.e. instead of destroying the intermediate results, they are transformed back into the original input value. This approach allows power consumption to be reduced to an arbitrarily low level; however, it incurs in a computation cost (either in time or space [13]). This approach has inspired the introduction of the so-called “Bennett clocking” scheme for QCA [14]. In such a scheme, the intermediate values of computation are saved “in place” by locking the QCA cells that computed the intermediate results. The Bennett scheme has very low power consumption and no space overhead because it requires no modification to the QCA circuit; however due to its fine-grained pipeline structure, it introduces a significant time overhead to the commonly employed clocking scheme (often referred to as Landauer clocking) for QCA. A QCA circuit using only a Bennett scheme would require an unacceptably long time to generate outputs. A hybrid solution is therefore needed to reduce power dissipation with no significant loss of throughput. Such hybrid solution should combine the advantages of pipelining and reversibility to achieve high throughput and low power consumption. This paper proposes a hybrid design approach for QCA circuits; it combines regions of Bennett clocked logic with memory regions (or *stages*) to facilitate pipelining. The proposed approach is evaluated in terms of throughput and power consumption due to bit erasures under different pipeline granularity. It should be noticed that an analysis of the overall power consumption would also include power consumption due to the clocking layer, however a detailed analysis of this component which is highly material dependent, is beyond the scope of this work, therefore it is not dealt in this manuscript.

M. Ottavi, S. Pontarelli and A. Salsano are with the University of Rome “Tor Vergata” ITALY

E. DeBenedictis is with Sandia National Laboratories, Albuquerque NM.

P. Kogge is with University of Notre Dame, South Bend IN

F. Lombardi is with Northeastern University, Boston MA

This research was partially funded by the Italian Ministry for University and Research; Program “Incentivazione alla mobilità di studiosi stranieri e italiani residenti all’estero”, D.M. n.96, 23.04.2001

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy’s National Nuclear Security Administration under Contract DE-AC04-94AL85000.




Fig. 1. a) Bistable QCA cell . b) QCA binary wire
Inverter: reversible power dissipation $\ll k_B T \ln 2$

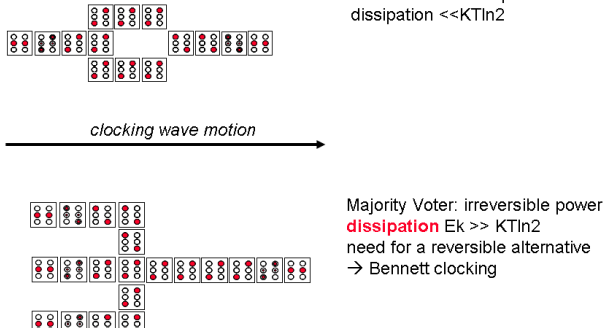


Fig. 2. QCA gates

This paper is organized as follows: section 2 introduces molecular QCA, section 3 presents the proposed approach and discusses in detail the stage organization (for computation and memory). Section 4 discusses performance evaluation. Section 5 presents as a case study, the performance evaluation of an XOR based parity checker. Section 6 provides the conclusion to this manuscript.

II. MOLECULAR QCA

Quantum-dot Cellular Automata (QCA) is a computation paradigm based on a cell made of four or six quantum dots (depending on the implementation technology) and two extra charges. The charges can tunnel between the dots of the cell but they cannot tunnel outside the cell. Coulombic repulsion between the extra charges leads to two stable states in which the charges are in antipodal locations along one of the two diagonals (figure 1 a)). A logic zero corresponds to the configuration in which a line through the extra charges has a negative slope. A logic one corresponds to a positive slope. The center dots are used to facilitate clocking by means of an electric field perpendicular to the plane of the QCA cell. Using the clocking field, the extra charges can be drawn into the center dots (thus making the cell neutral) or pushed onto the corner dots (forcing the cell to assume a logic value).

Through Coulombic interaction, the information of a single cell is propagated to other cells to form a binary wire (figure 1 b)). In QCA the basic logic gate is the majority voter, i.e. the output cell assumes the polarity of the majority of the inputs (figure 2). Along with the QCA inverter, this forms a functionally complete logic set. Figures 1 b) and 2 also show that the propagation of a signal in QCA is carried out through a sequence of four *clock phases* denoted by switch, hold, release, and relax. These clock phases are generated by a traveling wave of the E field perpendicular to the QCA plane. In the switch phase, a cell takes a new configuration when the charges move from the center to the corner dots. In the hold phase, a cell has a definite configuration and can drive the value of neighboring cells. In the release phase, the cell loses its configuration as the extra charges are drawn into

the center dots. Finally, when a cell is in the relax phase, it cannot influence the configuration of neighboring cells. In a six dot cell, this corresponds to the configuration when the extra charges are in the center dots. In figure 1 b), the cells on the far left and right of the shown QCA wire segment are in the relax phase. The second cell (from the left) is releasing its value. The third, fourth, and fifth cells have a definite value and are in the hold phase. These cells drive the sixth cell that is in the switch phase, thus assuming a new value. As the seventh cell (far right) is in the relax phase, then it has no value and cannot influence the new configuration being assumed by the sixth cell. The clocking scheme in which this pattern ripples down the wire in one direction is the traditional QCA clocking scheme, commonly referred to as Landauer clocking. Its advantage is that it allows information to be pipelined at a very fine-grained level. However, the wire and the inverter are logically reversible functions of low power consumption ($\ll k_B T \ln 2$ as they do not destroy information), the majority voter is logically irreversible because the information associated with the minority input (if present) is erased. The energy of the minority input is then thermalized, thus increasing the entropy by $\Delta S = k_B \ln 2$ and dissipating at least $k_B T \ln 2$ (approximately the kink energy E_k).

To take advantage of the low power computation potential of QCA, information cannot be destroyed as previously described. Bennett clocking provides a solution to this problem. If the intermediate results (in this case the inputs to the majority gate) are saved, then the majority function can be decomputed after its output has been latched. In the context of QCA, this can be done with no space overhead by employing Bennett clocking for the circuit [14]. This approach forms the foundation of the hybrid solution presented in this manuscript. While Bennett clocking does not incur in a space overhead, it does however entail a performance degradation in terms of throughput. The hybrid approach discussed below leverages the strengths of the Bennett scheme as well as the strengths of pipelining for QCA circuit operation.

III. PROPOSED APPROACH

This section describes the proposed approach to attain high performance for both power consumption and throughput. The QCA design is divided into computational and memory stages, the computational stages are clocked by the Bennett scheme and do not dissipate power. The memory stages are used to increase the throughput of the pipeline. In this paper, the QCA circuit is partitioned into M stages, each stage has i_j inputs and o_j outputs (obviously $i_j = o_{(j-1)}$).

A. Computation Stages

This section describes the clocking scheme for the combinational parts of the QCA circuit. Prior to describing the proposed clocking scheme a review of clock distribution techniques and clocking schemes for QCA is presented. In this paper, the distribution mechanism introduced by [15] is utilized; an E field generated on a layer of metallic wires above (or below) the QCA layer controls the tunneling within individual QCA cells. The cells are not directly connected to the clocking circuitry, this provides a substantial advantage at




Fig. 3. Proposed pipelined approach: top view

molecular scale. Moreover the continuous transition of the E field on the leading edge of the wave reduces the possibility of occurrence of kink, or error on a QCA wire.

The traveling E field is generated by providing each of the wires with a voltage (phase shifted from the neighbor by $\frac{\pi}{2}$) and a conducting ground layer on the other side of the QCA layer. Hennessy has shown that the E field generated with such a circuit can assume a sinusoidal shape, allowing for Landauer clocking. The z component of the vector \vec{E} acting as clock signal can be described by the wave equation [16], i.e. $E_z(x, t) = E_0 \cos(\kappa x - \omega t)$. Computation and switching of cells occur only on the leading edge of the wave, thus providing directionality to the QCA circuit and virtually eliminating the probability of kink. This is a space continuous implementation of the classic four phases-four zones clocking scheme introduced in [2]. This clocking strategy has been called a “traveling wave”, a “computational wave”[8], and “Landauer”[9] clocking. Here, in the context of reversibility, “Landauer” clocking is used to describe this clocking approach.

The highest (maximum) performance in terms of speed is related to the maximum applicable clock speed and is derived from the tunneling phenomena between quantum dots. To maintain the adiabatic nature and solution of the Schrödinger equation, the switching time t^* of the E field on a QCA cell must be greater than the tunneling speed between quantum dots [2]. Consequently, the fastest applicable clock period on a cell is $T_l = 2t^*$ and therefore $\omega \leq \omega_0 = \frac{2\pi}{2t^*}$. The constraint on the maximum applicable period is used in a later section to assess the throughput of a Landauer clocking schemes; in general for o outputs, $Tr = \frac{o}{2t^*}$.

A Bennett scheme has two steps: computation and de-computation. In the first step, it performs the computation on the inputs and propagates to the outputs without deleting intermediate results. In the second step, the intermediate results are decomputed by the clock “backing off”. So, the release of the cells starts from the outputs and is traced back to the inputs, eventually releasing the whole circuit. This process does not erase (delete) any information because every cell that is released, can “copy” its contents to the still locked cell that originally produced the information in the cell being released. This process prevents information from being thermalized [9].

Hence, circuits implemented with Bennett clocking do not dissipate energy over the course of a computation/decomputation cycle; at the end of a Bennett region’s computation/de-computation cycle, both the original inputs and outputs are stored. The speed of computation is function




Fig. 4. Clocking wave for the Bennett scheme. T represents the time period of the pipeline stage.

of the time required for the clocking signals to propagate back and forth across the region. Furthermore, it is not necessary to make any modifications to an irreversible QCA circuit to make it reversible. In this case, reversibility is accomplished through clocking rather than by circuit, avoiding any circuit overhead required for accomplishing reversibility by Landauer clocking. A circuit clocked with the Bennett scheme has also the advantage that it requires no modification to the layout to avoid deleting the information at the inputs (as it would happen if the inputs were propagated to the outputs). The power dissipated when losing a bit of information is almost equal to the kink energy, $E_{diss} \simeq E_k \gg k_B T \ln 2$. The value of the dissipated energy is obtained from the non equilibrium equation i.e. a set of first-order differential equations for the coherence vector of the QCA cells in contact with the thermal environment [16].

For the computation stages, a Bennett scheme can be implemented using Hennessy’s clocking implementation strategy [15] by applying suitable signals, $\Phi_1 \dots \Phi_n$, (figure 5 a) to the buried wires. The signals needed for Bennett style clocking are very different from the signals needed Landauer style clocking. For the Bennett clock, once the QCA cells have been locked, they must remain locked throughout the remaining part of the computation phase and be released in the de-computation phase as described earlier (figure 4). The pattern of waveforms present on each wire to produce this effect is shown in figure 5 a); Φ_i remains high at V_{max} until Φ_{i+1} reaches V_{min} . With this clocking scheme, data is provided as output from a stage at every period $t = T$, i.e the time required for the clock to sweep forward, latch the output and then retract back the de-computation of all intermediate solutions. As discussed above, to preserve adiabaticity, the switch time on a cell must be at least t^* . For d as the lateral size of a QCA cell, and $N = \lambda_c/d$ as the width of the Bennett-clocked region in number of cells, then the period is given by $T_b = \frac{2\lambda}{d} t^* = 2Nt^*$.

B. Memory Stages

Each memory stage is a single buffer (register) that is used to separate the different stages of the pipeline. A memory stage provides the inputs to a Bennett clocked zone and latch the outputs. The memory stage implementation is straightforward; it can be implemented with a single vertical row of QCA cells, or the minimum number of cells related to the achievable pitch




Fig. 5. a) Clock signal to the buried wires, b) Clocking signal for the memory zones

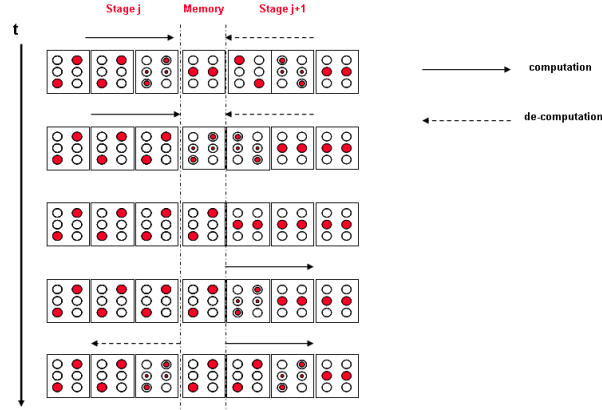


Fig. 6. Asymmetric interaction on memory cell

of the clocking wires.

In the simplest design, the contents of the latch are overwritten on each cycle when a new input is stored. This results in a dissipation given by the number of bits (as stored in each latch) multiplied by the number of latches. However, the properties of QCA cells and the clock can be exploited to reduce this dissipation. As shown in figure 5 b) the clocking signal is sinusoidal with the same period T as for the Bennett clocking scheme. Rather than using a traditional QCA circuit design in which completely locked cells drive the value of neighboring switching cells (with fully relaxed cells on the other side), an asymmetric interaction is used. In this case, the cells that would normally be in the relax phase (with no value) are instead in the process of releasing their values at the same time as the latch is assuming its own values. The directionality of the circuit is preserved because the signal from the driver cell is still stronger than the releasing cell. However, if the data being released is the same as the new data being latched, then that bit will not be dissipated. Instead, it will be “copied” into the new bit being stored.

The signals applied to each buried clocking wire for this asymmetric interaction are shown in figure 5 a). Phase Φ_1 of stage $j+1$ releases its information, while the memory stage is switching; phase Φ_n of stage j is in the hold phase. This allows the new value to propagate correctly to stage j while avoiding the deletion of the information in stage $j+1$ when the value is the same. The propagation in the two opposite directions is illustrated in figure 6; two opposite values interact on the memory cells located in the center. Since the cell on the left locks its value (hold phase) earlier than the one on the right, then the Coulombic interaction (quadrupole moment) on the memory cell is stronger and therefore, it causes the memory cell to assume its value.

IV. PERFORMANCE EVALUATION

The performance of the proposed solution is evaluated in terms of both throughput and power consumption. In general these figures of merit could be in conflict; an increase in the number of pipeline stages leads to a higher throughput. However, an increase in the number of pipeline stages also increases the possible discarded bits of information, resulting in a higher power consumption. Therefore the number of stages must be carefully selected such that computational performance and power consumption can be best assessed as per application requirements.




Fig. 7. Pipelined stages with Bennett clocking

Consider a pipelined circuit with no feedback loop; this circuit is assumed to have M stages and o outputs. As the period of a Bennett stage is given by $T_b = 2Nt^*$, then the throughput is $T_r = \frac{o}{T_b} = \frac{o}{2Nt^*}$ and the initial latency L_b is proportional to $T_b/2$ (figure 7) as $L_b = \frac{MT_b}{2}$

For this M staged pipeline, the power consumption $P(t)$ as function of time is given by

$$P(t) = E_{diss} \cdot \sum_{j=0}^{\infty} \sum_{i=0}^{M-1} K_i(t) \delta(t - \frac{jT_b}{2}) \quad (1)$$

where $K_i(t)$ is the number of inputs of stage i that change value at time t , E_{diss} is the energy dissipated (thermalized) when a bit is deleted on the stage latches.

The time varying value of $K_i(t)$ accounts for the random time variability of the data in the pipeline on the memory stage i . On average, it is likely to be nearly half of the bits stored in memory. The power dissipation of a circuit is therefore spatially localized on the memory stages and is a time varying function composed of a train of pulses. It accounts for the dissipation occurring at the discrete time instants $t = jT/2$ (where j is an integer) on the memory stages. As shown in figure 7 at each $t = jT/2$ power dissipation occurs only on $\lfloor M/2 \rfloor$ i.e. deletion of data occurs only in half of the memory stages in which the waves for computing and decomputing meet. So, at time $t = nT/2$ the number of coefficients $K_i(t) \neq 0$ is $\lfloor M/2 \rfloor$.

V. CASE STUDY: XOR TREE PARITY CHECKER

The size of the Bennett clocked zones can vary from a minimum of two QCA cells (the single cell case is a pathological case as equivalent to a Landauer scheme with a clock as a traveling wave) to the whole circuit (i.e. a purely

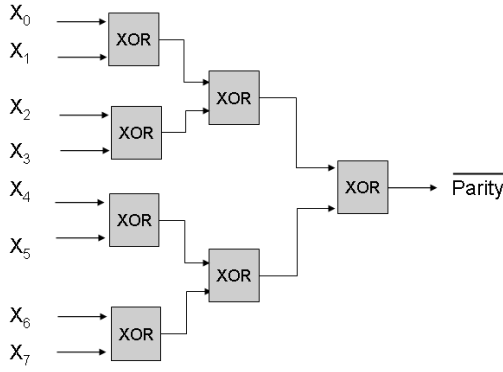


Fig. 8. Case Study: XOR tree parity checker

Bennett clocked circuit). As stated previously, it is expected that increasing the zone size, throughput and power consumption would both decrease, i.e. degradation in computational performance but improvement in power dissipation

Power dissipation depends also on the circuit functionality: a circuit made of only wires and inverters (so reversible by definition), has the best performance with Landauer clocking (no information is deleted). However at the input/output, a circuit made of majority voters requires the introduction of a Bennett scheme to reduce dissipation due to deletion of information. For low power dissipation a Bennett clocked stage must have a number of MVs such that the number of bits of information to be deleted in that stage using Landauer clocking is significantly higher than the number of inputs deleted in the Bennett stage (the number of bits deleted in a stage is not necessarily equal to the MVs as shown next).

An example of the proposed scheme is analyzed in detail; an M stage tree made of XOR gates generates the parity bit for $w = 2^M$ inputs. A worst-case analysis of throughput and power dissipation is calculated for the XOR based parity bit generator by using the previously introduced analysis. For the same XOR tree, different clocking schemes are employed. Landauer clocking is used to provide an irreversible reference for comparison. For the Landauer clocked case, the throughput is $Tr_l = \frac{1}{T_l}$, where T_l is the period of the Landauer clocking wave. A single result is generated as output on each cycle after the pipeline is full.

With the Bennett scheme, throughput and power consumption depend on the period of the Bennett clocked regions, i.e. the period depends on the width of these regions. So, let $T_b = 2Nt^*$, where N is the width (in number of XOR gates) for the region under consideration. Since the same circuit is being compared, then there is again one output per clock period, i.e. $Tr_b = \frac{1}{T_b} = \frac{1}{2Nt^*}$. For both the Landauer and Bennett clocked schemes, the worst case dissipation for an XOR gate is given by $2E_{diss}$. For the Landauer case, consider the XOR function implemented as shown in figure 9. At most, the combination of inputs leads to a dissipation of $2E_{diss}$. The Bennett case is simpler as there are two inputs to each XOR gate. No dissipation will occur within the XOR gate, but the inputs may be written over on the next cycle. This, then, also leads to a worst case dissipation of $2E_{diss}$.

To compare the power dissipation of a M stages XOR tree




Fig. 9. Dissipation of the XOR gate

clocked with the Landauer and Bennett schemes, the following assumptions and definitions are used:

- 1) the dissipated energy of a thermalized bit of information is considered equal to the kink energy, i.e. $E_{diss} \simeq E_k$
- 2) the kink energy value is $E_k = 3.14577 \cdot 10^{-20}$ Joule obtained for a molecular squared cell of lateral size $l = 1.5\text{nm}$ [17] and relative permittivity $\epsilon_r = 1$ (no dielectric material between cells)
- 3) the number of stages of the XOR tree is k ;
- 4) the number of stages of the pipeline is M ;
- 5) the number of stages of the XOR tree per pipeline stage is $c = \frac{k}{M}$
- 6) the values of dissipated energy are calculated over the respective period of computation for each scheme; then the corresponding power values are considered averaged on the same period
- 7) by considering the worst case scenario, the value of $K_i(t)$ from equation 1 is not time dependent; therefore the deleted information is always equal to the number of inputs of stage i

Based on the previous assumptions and analysis, the energy dissipated in a period T_b for a Bennett clocked scheme in the XOR tree is

$$E_B = \int_0^{T_b} P_B(t)dt = 2E_k \sum_{i=0}^{M-1} 2^{ic} = 2E_k \frac{2^{cM} - 1}{2^c - 1} \quad (2)$$

where the sum of a geometric progression of ratio 2^c is used. So, the energy dissipated in a Landauer clocked tree is the sum of the energy dissipated by the whole tree, i.e.

$$E_L = \int_0^{T_l} P_L(t)dt = 2E_k \sum_{i=0}^{M-1} 2^i = 2E_k(2^M - 1) \quad (3)$$

Figures 10 through 12 show the effect of keeping the circuit size (k) constant while varying the number of pipeline stages (M) and consequently the number of XOR gates per pipeline stage (c). Although only the integer values of c are meaningful, the intermediate values are also shown for better understanding of the overall trend of this plot. As c applies only to Bennett clocking, the results for Landauer clocking are constant. Figure 10 shows a comparison of the throughput for the Bennett and Landauer schemes. As expected, the Landauer scheme shows a higher throughput and the gap in performance increases with an increase in c , i.e. as the pipeline stages become wider and the depth of the pipeline decreases, then he throughput decreases.

Figure 11 shows the advantage of the Bennett scheme for dissipated energy per period of computation. As the pipeline

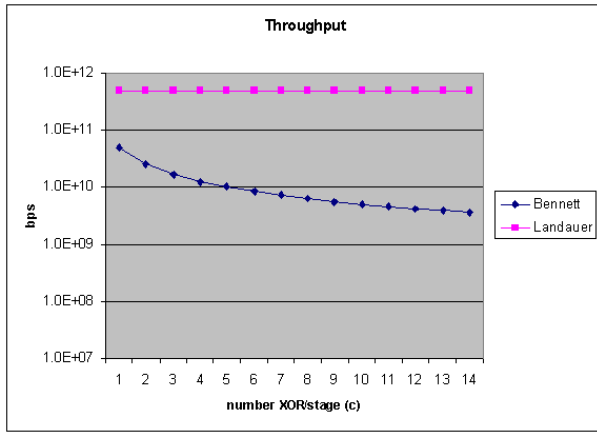


Fig. 10. Throughput

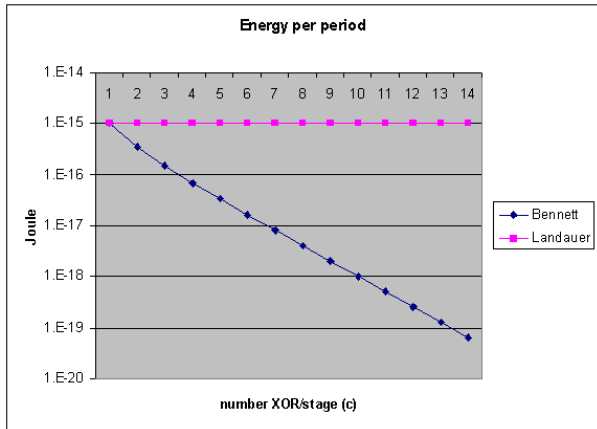


Fig. 11. Energy dissipation per period

depth decreases (c increases), the power dissipated per computation period improves because there are fewer latches (whose contents are dissipated). Even when the entire circuit is in a single Bennett stage, the dissipation does not drop to zero because the original inputs are still being deleted every T_b . Finally figure 12 plots the amount of operations (output bits) per Joule per second, i.e. introducing time too. The results show that there exists a point (the intersection of the curves) for c to attain better performance ($c \leq 6$) for Landauer clocking (for $c > 6$ Bennett clocking is better), i.e. for low values of c the throughput of a pipelined approach with the Bennett scheme is not sufficient to overcome the penalty in terms of power dissipation, while by increasing the size of the pipeline stages (higher c) the power dissipation has a higher impact with respect to a reduction in performance.

VI. CONCLUSION

This paper has introduced a pipelined architecture for low power QCA circuits using a Bennett clocking scheme; this is a clocking scheme that allows intermediate results to be de-computed rather than erased, thus avoiding power dissipation due to loss of information. This architecture allows designers to adjust the level of reversibility in a QCA circuit based on throughput and/or power dissipation. The utilization of Bennett clocking and pipelining enhances performance of the design space to assess the often conflicting figures of merit of power

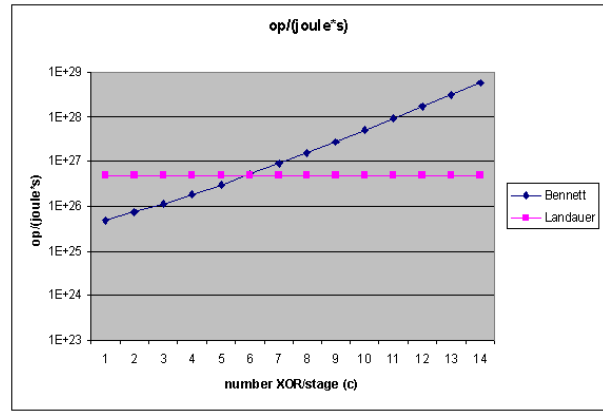


Fig. 12. Comparison of operations per joule per second

and throughput. An example of a QCA circuit that utilizes the proposed scheme, has been analyzed in detail; this circuit is the parity generator (XOR tree). It has been shown that a Bennett clocked pipeline can provide substantial power saving over a Landauer clocked circuit operation.

REFERENCES

- [1] C.S. Lent, P.D. Tougaw, W. Porod, and G.H. Bernstein, "Quantum cellular automata," *Nanotechnology*, vol. 4, no. 1, pp. 49–57, 1993.
- [2] C.S. Lent and P.D. Tougaw, "A device architecture for computing with quantum dots," in *Proc. of the IEEE*, Mar 1997, vol. 85, pp. 541–557.
- [3] P.D. Tougaw and C.S. Lent, "Logical devices implemented using quantum cellular automata," in *Journal of Applied Physics*, 1994, vol. 75(3), pp. 1818–1825.
- [4] A.O. Orlov, I. Amlani, G.H. Bernstein, C.S. Lent, and G.L. Snider, "Realization of a functional cell for quantum-dot cellular automata," in *Science*, 1997, vol. 277, pp. 928–931.
- [5] I. Amlani, A.O. Orlov, G. Toth, G.H. Bernstein, C.S. Lent, and G.L. Snider, "Digital Logic Gate Using Quantum-Dot Cellular Automata," *Science*, vol. 284, no. 5412, pp. 289–291, 1999.
- [6] M Baseer Haider, Jason L Pitters, Gino A. DiLabio, Lucian Livadaru, Josh Y Mutus, and Robert A. Wolkow, "Controlled coupling and occupation of silicon atomic quantum dots at room temperature," *Physical Review Letters* 102, 2009.
- [7] M. Lieberman, S. Chellamma, B. Varughese, Y.L. Wang, C.S. Lent, G.H. Bernstein, G. Snider, and F.C. Peiris, "Quantum-dot cellular automata at a molecular scale," in *Molecular Electronics II, Annals of the New York Academy of Sciences*, 2002, vol. 960, pp. 225–239.
- [8] EP Blair and CS Lent, "An architecture for molecular computing using quantum-dot cellular automata," *Nanotechnology*, 2003. *IEEE-NANO 2003. 2003 Third IEEE Conference on*, vol. 1, 2003.
- [9] J. Timler and C. Lent, "Maxwell's demon and quantum-dot cellular automata," *Journal of Applied Physics*, vol. 94, no. 2, pp. 1050, 2003.
- [10] MP Frank, "Introduction to reversible computing: motivation, progress, and challenges," in *CF '05: Proc. 2nd conference on Computing frontiers*, New York, NY, USA, 2005, pp. 385–390, ACM Press.
- [11] MP Frank, "The physical limits of computing," *Computing in Science & Engineering [see also IEEE Computational Science and Engineering]*, vol. 4, no. 3, pp. 16–26, 2002.
- [12] C.H. Bennett, "Logical reversibility of computation," *IBM J. Res. Develop.*, vol. 17, no. 6, pp. 525–532, 1973.
- [13] C.H. Bennett, "Time/space trade-offs for reversible computation," *SIAM Journal on Computing*, vol. 18, no. 4, pp. 766–776, 1989.
- [14] C.S. Lent, S.E. Frost, and P.M. Kogge, "Reversible computation with quantum-dot cellular automata (QCA)," *Proceedings of the 2nd conference on Computing frontiers*, pp. 403–403, 2005.
- [15] K. Hennessy and C.S. Lent, "Clocking of molecular quantum-dot cellular automata," in *JOURNAL OF VACUUM SCIENCE and TECHNOLOGY B*, 2001, vol. 19(5), pp. 1752–1755.
- [16] CS Lent, Liu M., and Lu Y., "Bennett clocking of quantum-dot cellular automata and the limits to binary logic scaling," *Nanotechnology*, vol. 17, pp. 4240–4251, 2006.
- [17] SE Frost, TJ Dysart, PM Kogge, and CS Lent, "Carbon nanotubes for quantum-dot cellular automata clocking," *Nanotechnology*, 2004. *4th IEEE Conference on*, pp. 171–173, 2004.