

# Fine-Grained Architecture in Dark Silicon Era for SRAM-Based Reconfigurable Devices

Sadegh Yazdanshenas and Hossein Asadi, *Member, IEEE*

**Abstract**—In this paper, we present a fine-grained *dark silicon* architecture to facilitate further integration of transistors in *SRAM-Based Reconfigurable Devices* (SRDs). In the proposed architecture, we present a technique to power off inactive configuration cells in non-utilized or underutilized *Logic Blocks*. We also propose a routing circuitry capable of turning off the configuration cells of *Connection Blocks* (CBs) and *Switch Boxes* (SBs) in the routing fabric. Experimental results carried out on MCNC benchmark show that power consumption in configuration cells of lookup tables, CBs, and SBs can, on average, be reduced by 27%, 75%, and 4%, respectively.

**Index Terms**—SRAM-Based Reconfigurable Devices, Dark Silicon, Power Consumption, Routing Fabric, Dependability.

## I. INTRODUCTION

IN the past decade, *SRAM-based Reconfigurable Devices* (SRDs) have gained much popularity in wide range of applications due to short design and implementation time, inexpensive design update, and opportunity to reconfigure device with workload variation. The most commonly adapted SRD architecture in industry is the island style architecture, where *Logic Blocks* (LBs) are surrounded by a sea of routing fabric [1]. The routing fabric consists of vertical and horizontal channels, *Connection Blocks* (CBs), and *Switch Boxes* (SBs). While CBs provide connectivity between LBs and routing channels, SBs are employed at the intersection of vertical and horizontal routing channels to provide routing flexibility. LBs in industrial SRDs range from a small set of *Look-Up Table* (LUT) and hard logic to complex *Digital Signal Processing* (DSP) and processor blocks.

With aggressive transistor downscaling, the number of transistors in SRDs has already passed six billions transistors per a single chip [2]. Such aggressive scaling, however, has faced major challenges such as power and reliability. The dominant power consumption in SRDs is static power [3], which is mainly attributed to configuration cells used to program different resources available in SRDs. It is projected that the leakage power per SRAM cell increases drastically for each upcoming technology generations, creating a *power wall* for further scaling of transistor feature size [4]. One possible solution to overcome the power wall is to selectively power off the inactive regions, called *dark silicon* [5]. This concept urges that some parts of the design should be inactive in order to avoid *power wall*.

In addition to power limitation of SRDs, the susceptibility of SRAM configuration bits to energetic particles along with

the enormous number of sensitive configuration bits results in an unacceptable error rate for enterprise and safety-critical applications [6]. This potentially creates a *reliability wall* to further integration of transistors in a single chip unless heavy redundancies are employed [4].

In this paper, we present a fine-grained architecture to turn off configuration cells in different resources of SRDs within inactive regions. In the proposed architecture, typical resources of SRDs such as LUTs, CBs, and SBs are examined to find possibility of turning off a portion of the resource configuration bits. In addition to the significant power reduction, by turning off the unused configuration bits in the proposed architecture, the number of sensitive configuration bits to particle strikes is also significantly reduced, enhancing the circuit reliability.

Experimental results show that the static power of configuration bits in LUTs is, on average, reduced by 27% while imposing less than 6% area overhead. In CBs, the power consumption of configuration cells is, on average, reduced by 75% at the cost of 20% increased CB area. In SBs, the power consumption of configuration cells is, on average, reduced by 4% while the SB area is increased by 19%. In addition, our results demonstrate that the number of susceptible configuration bits to soft errors in SBs and CBs are also reduced by 77% and 5%, respectively, allowing further integration beyond the reliability wall.

The rest of this paper is organized as follows. Sec. II reviews the previous work. Sec. III presents the proposed architecture. Sec. IV details the experimental setup and then reports the results. Sec. V discusses limitations of the proposed architecture. Finally, Sec. VI concludes the paper.

## II. RELATED WORK

Aggressive transistor scaling has been hindered by emergence of a phenomena known as the power wall [4]. Power wall has put a limit on the number of transistors that can simultaneously be active on a single chip. Hence, in order to utilize the vast amount of silicon area on a chip, a fraction of the chip area has to be turned off or under clocked [5]. This phenomena, known as *Dark Silicon* allows further scaling of multi-core systems [7].

Previous work on overcoming the power wall in SRDs has aimed at reducing either dynamic or static power consumption. Since the latter is dominant in SRDs [3], we focus on the previous work aiming at static power consumption. When it comes to leakage power consumption, configuration bits of SRDs play a major role. One approach to reduce power consumption in configuration memory is using more complicated, expensive, low-power fabrication processes such as use of triple oxide, multiple-Vt, and variable transistor gate length [8]. [9] has proposed the use of a sleep transistor in SRD silicon area to turn off power in unused regions. They also further boost

Manuscript received January 18, 2014; revised May 5, 2014; accepted June 13, 2014. All authors are with the Department of Computer Engineering, Sharif University of Technology, Tehran, Iran. Emails: syazdanshenas@ce.sharif.edu, asadi@sharif.edu

Copyright (c) 2014 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending an email to pubs-permissions@ieee.org.

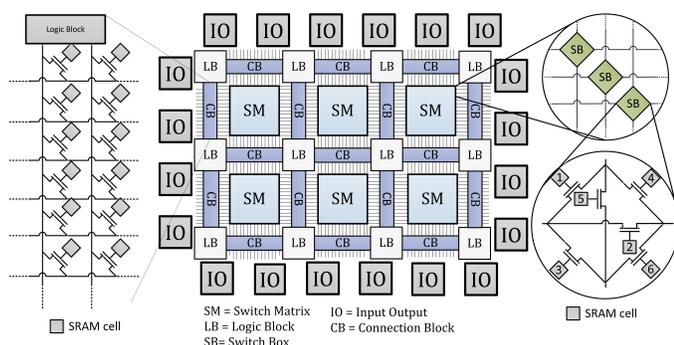


Fig. 1: General view of an island style architecture

their power reduction scheme by *Region Constraint Placement* (RCP) to create more opportunities for turning off a greater portion of silicon area. This constraint on placement, however, limits CAD refinements carried out during the placement phase. In addition, RCP targets coarse-grained blocks while fine-grained granularity could provide more opportunities to turn off unused resources without compromising CAD objectives. [10] has proposed *Dynamically Controlled Power Gating* (DCPG) scheme which uses a controller to turn off different regions of SRDs using a centralized controller. Such controller requires complex routing to different blocks, resulting in more expensive chips. DCPG works at a relatively coarse-grained scheme, providing few opportunities for power saving when device utilization is high. In addition, limited forms of power gating already exist in state-of-the-art industrial SRDs [11].

There are also other works that try to reduce power consumption in SRDs without turning off device regions. These works include using multi-voltage sources [12], use of heterogeneous routing resources [13], and power-aware CAD algorithms [14]. The power reduction achieved by these techniques, however, is not expected to be as high as those that completely turn off unused resources. For example, [14] has used efficient algorithms for power saving in the entire CAD flow and has been able to obtain 12.6%, 7.6%, 3.0%, and 2.6% improvement for clustering, technology-mapping, placement, and routing algorithms, respectively. Another important point is that by employing such aggressive power-aware algorithms, other design objectives such as timing and area which are the primary concerns of CAD tools, are compromised. Nevertheless, these works are valuable in that they can be additive to the techniques aiming at turning off unused silicon area.

### III. PROPOSED ARCHITECTURE

The main aim of the proposed architecture is to effectively turn off unused logic resources and routing configuration cells in an island style SRDs. The typical resources of an island style architecture including LBs, CBs, and SBs have been depicted in Fig. 1. All of these resources are programmed using configuration SRAM cells. The most fundamental element used in the proposed architecture is a SRAM cell, called *Cu-SRAM*, that can be cut off using a controlling SRAM cell, referred to as *CSRAM*. As shown in Fig. 2, a CSRAM cell is a regular SRAM cell that is used to cut off a single or a group of Cu-SRAM cells. CSRAMs are programmed during system

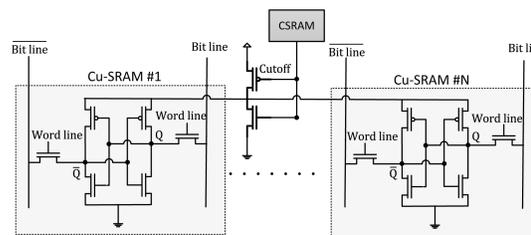


Fig. 2: SRAM cell with cut-off ability

reconfiguration to control the power supply of a group of Cu-SRAM cells, which are all either *on* or *off*. This brings an opportunity to turn off a group of Cu-SRAMs. By employing CSRAM and Cu-SRAM cells, we examine different parts of SRDs to effectively turn off unused or under utilized parts in order to save power and enhance dependability.

#### A. Cut-off SRAM

In the proposed architecture, we propose to replace all SRAM configuration cells with Cu-SRAM cells, as shown in Fig. 2. The main advantage of a Cu-SRAM over a regular SRAM cell is that it can be turned off by activating the cut-off signal when the cell is not utilized in the design. A group of Cu-SRAM cells controlled by a CSRAM, however, impose two extra transistors and also more power in the cell active mode as compared to group of regular SRAM cells. Our Hspice simulations reveal that the leakage power of a group of 64 Cu-SRAM cells, shown in Fig. 2, using 45nm technology is about 2.86E-06 watts during the active mode and 7.45E-08 watts during cut-off while the leakage power of a normal minimum-sized SRAM cell in 45nm is 5.19E-08 watts.

One major limitation of the cut-off transistor is that it reduces the switching speed of Cu-SRAM cells. The switching speed of configuration bits does not, however, affect the circuit performance in the normal operation of SRDs since the Cu-SRAM cells are used to hold configuration bits and are not part of circuit datapath. We have used minimum-sized SRAMs for both Cu-SRAM and CSRAMs. The reduced switching speed of Cu-SRAMs, however, slows down the programming rate of configuration bits. Additionally, since the power supply of Cu-SRAMs is derived by CSRAMs, the CSRAM cells should be programmed before Cu-SRAMs during circuit reconfiguration.

#### B. Logic Blocks

The main logic resource employed in LBs is LUT. In general, once a design is mapped to an SRD device, it is expected that a significant fraction of LUTs remain unused or underused. The unused or underused LUTs bring the opportunity of turning of the corresponding configuration bits and saving significant power. In case, a LUT is not used in a design, all corresponding configuration bits can be turned off using one CSRAM cell, as shown in Fig. 3a. In this figure, a LUT-6 consists of 64 SRAM cells is turned off by employing one CSRAM cell. Here, a single CSRAM cell is shared among all cut-off transistors of LUT-6 configuration bits.

Despite a significant power saving achieved by turning off unused LUT-6 cells, our study reveals that a majority of LUTs are underused rather than being completely unused in a design. We have conducted a study investigating the percentage of unused and underused LUTs for MCNC benchmark circuits mapped to a SRD device employing LUT-6. The results of

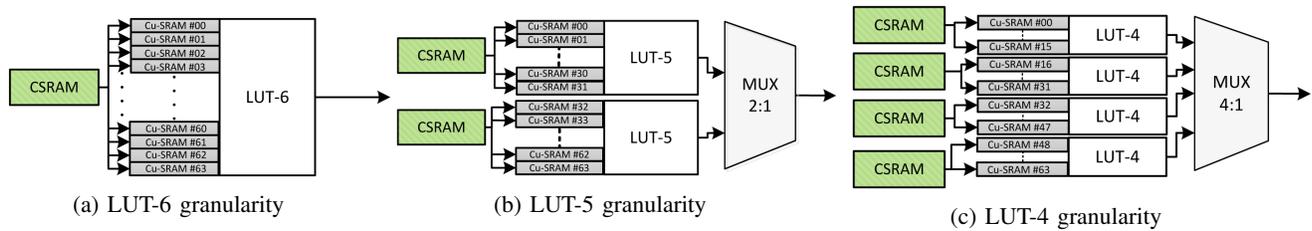


Fig. 3: Proposed dark silicon architecture to turn off configuration cells in LUTs with different granularities

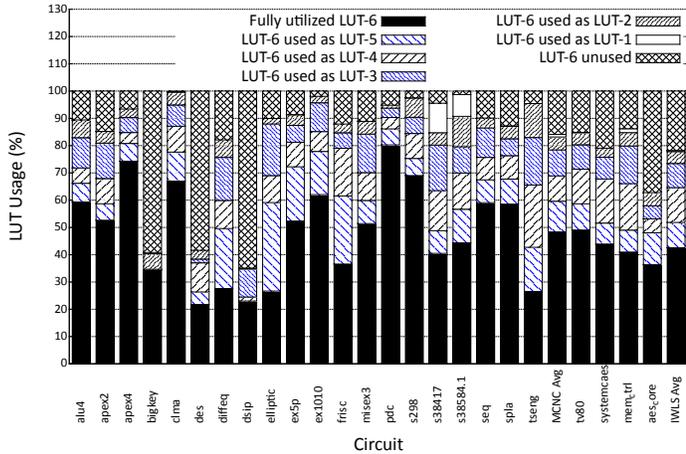


Fig. 4: Percentage of unused or underused LUT-6 elements

this study, reported in Fig. 4, reveal that unused LUTs, on average, contribute to less than 16% of the total LUTs while underused LUTs contribute to more than 35% of the total LUTs. In particular, the results show that, on average, 11%, 9%, and 9% of LUT-6 cells are used as LUT-5, LUT-4, and LUT-3, respectively. In order to demonstrate that larger circuits also exhibits such degree of unused and underused LUTs, we have carried out the same experiment on four large circuits from IWLS-05 benchmark suite [15]. The results, as reported in Fig. 4, demonstrate that the larger circuits also show similar behaviour to MCNC benchmark suite.

In order to turn off the unused configuration bits in underused LUTs, we propose to use a CSRAM cell for smaller granularities of LUTs, such as LUT-5 and LUT-4, as shown in Fig. 3b and Fig. 3c, respectively. In Fig. 3c, as an example, by employing a CSRAM cell for each LUT-4, a three-fourth of a LUT-6 cell can be turned off in case it is used to implement a 4-input function. As can be seen in Fig. 5, a more fine-grained architecture brings an opportunity to turn off a greater number of unused SRAM cells. This comes at the cost of greater number of CSRAM cells, which, in turn, imposes more area and power overheads.

### C. Connection Blocks

CBs are either used to connect the routing channel to input pins of LBs or to connect output pins of LBs to the routing channel. We call the former input CBs and the latter output CBs. Due to the nature of LBs which receive several inputs and generate only one output, the number of input CBs is much greater than the number of output CBs. Hence, power saving in input CBs will significantly improve the overall power efficiency of CBs. Here, we focus on input CBs rather than output CBs.

In input CBs, the proposed scheme is motivated by the fact

that at most one of the lines for each LB input is activated at a time. In other words, if the input line is used, the line has one source and only one of the corresponding configuration cells is activated. The remaining configuration cells are deactivated and can be turned off. In order to be able to turn off as many of these inactive cells as possible, we propose to group the configuration bits of an input CB into few sets, where each set is controlled by a CSRAM cell. This allows us to have at most one active set while the other sets are turned off using the corresponding CSRAM cells. In general, the configuration cells in an input CB in a device with channel width of  $n$  can be divided into  $k$  sets controlled by  $k$  CSRAM cells, where the channel width of each set is equal to  $n/k$ . This creates an opportunity to turn off at least  $k-1$  out of  $k$  sets. For smaller values of  $k$ , fewer CSRAM cells are required but fewer Cu-SRAMs are turned off. For larger values of  $k$ , however, more opportunity to turn off configuration bits is provided at the cost of greater number of CSRAM cells.

### D. Switch Boxes

SBs provide connectivity between horizontal and vertical routing channels. A SB pattern is typically represented by a sequence of zeros and ones, which represent *off* and *on* configuration bits, respectively. The order of configuration bits in a symbolic notation of a SB is shown in Fig. 1.

In order to explore possibility of turning off configuration bits in SBs, we have investigated the distribution of different SB patterns. Our study over MCNC benchmarks shows that the distribution of SB patterns is not uniform. While some patterns such as "000000" and "010010" are highly frequent, some other patterns such as "000011" and "000110" are less frequent in designs mapped to SRDs. After characterizing SB patterns, we further investigate the *on* and *off* frequency of SB configuration bits in different SB patterns and explore possibility of turning off a group of SB configuration bits.

By profiling the SB patterns and the *on* and *off* frequency of SB configuration bits, we categorize the configuration bits of a typical SB into a single group. This group is controlled by one CSRAM. The proposed scheme allows turning off 23% of SB configuration cells with only one CSRAM cell, as it will be detailed in Sec. IV. Note that grouping the configuration bits into two or three sets brings more opportunity to turn off unused switches. However, such scheme imposes higher number of CSRAM cells per SB. Our results demonstrate that the power penalty imposed by two or three CSRAM cells fades away the power reduction achieved in such schemes. As such, we use one set of grouping in SBs in the proposed architecture.

## IV. EXPERIMENTAL RESULTS

In order to evaluate the efficiency of the proposed architecture, we have implemented MCNC benchmark circuits using

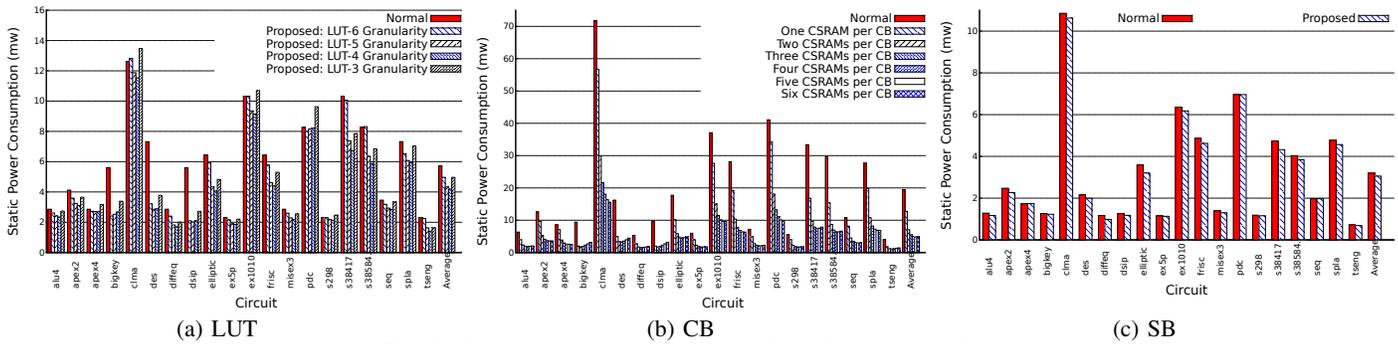


Fig. 5: Static power consumption of configuration cells

VPR 6.0 [1] toolset. For this purpose, we first used ABC toolset [16] to map MCNC benchmarks to LUT-6 elements. Then, VPR 6.0 is used to perform placement and routing of target designs. The optimization objective in the experiments is area. As such, all designs are placed on a minimum-grid FPGA size. HSpice toolset is also used to compute power usage of CSRAM and Cu-SRAM cells. For HSpice simulations, we have used 45nm *Predictive Transistor Model* (PTM) [17].

#### A. Power Efficiency Results

We have extracted the power reduction achieved by the proposed architecture in LUTs for different granularities (LUT-6, LUT-5, LUT-4, and LUT-3). As the size of Cu-SRAM cell sets decreases, more power saving opportunities are obtained. However, there exists a threshold in which the power consumption of CSRAM cells exceed the power saving obtained by turning off the Cu-SRAM cells. As compared to the baseline SRD architecture, the proposed architecture reduces the power by 13%, 25%, 27%, and 14% when LUT-6 cells are power controlled in LUT-6, LUT-5, LUT-4, and LUT-3 granularities, respectively. Hence, the best power efficiency is achieved when a CSRAM is used to control LUTs in LUT-4 granularity. As can be seen in Fig. 5a, the power consumption can, on average, be reduced for LUT-4 groups by 27% and up to 62% while imposing only 6% area overhead. For smaller or larger groups, this power saving declines.

Fig. 5b reports power saving achieved when the proposed architecture is applied to CBs at different granularities. In this experiment, we have reported the power saving when a CB is divided into one, two, three, four, five, and six sections. It is shown that the best power saving is obtained when routing tracks are divided into five sections. At this point, 77% of original configuration cells in CBs can be completely turned off, resulting in enhanced power efficiency. The average area overhead for CBs in MCNC benchmark is about 20%.

Fig. 5c reports the power saving results gained by the proposed SB architecture. Although 23% of original SBs can be turned off, the power consumption is reduced by only 4% as shown in Fig. 5c. The insignificant power reduction in SBs is due to power overhead of CSRAM cells in utilized SBs which contribute to 77% of the total SBs in a SRD device. The area overhead of the proposed architecture for SBs is about 19%. In overall, the proposed architecture, on average, reduces the static power consumption of configuration cells by 57%.

#### B. Dependability Enhancement

Previous study has demonstrated that short faults are one of the major threats to the dependability of SRDs [18]. A short

fault typically occurs when two nets are erroneously connected together by turning on an unused configuration bit. The unused configuration bit gets *on* by hitting an energetic particle strike. This can short two different nets in the design.

The proposed architecture enhances the dependability of SRD routing fabric by reducing the number of susceptible nets to soft errors. This is achieved by the fact that an unused Cu-SRAM is not susceptible to particle strike since diffusion areas of SRAM transistors are inactive in case of power outage. Our results reported in Fig. 6 show that the proposed architecture eliminates 77% and 5% of susceptible configuration bits to soft errors in CBs and SBs, respectively. We have used the number of active short sensitive cells as a measure of dependability.

It is noteworthy to mention that if unused (or OFF) CSRAM is turned on due to a particle strike, its corresponding Cu-SRAM cells will be turned on. Such Cu-SRAMs do not affect the system reliability while they increase the power consumption of the system. In case a particle strike hits a used (or ON) CSRAM and turns it off, all the corresponding Cu-SRAMs will be unwantedly turned off. This will can definitely change the circuit functionality and can affect the system reliability. To overcome this issue, we employ asymmetric SRAM cell proposed in [19] that makes cells immune to soft errors when they have a specific logical value. For this purpose, we use one-optimized asymmetric cell so that the CSRAMs become immune to one-to-zero bit-flips.

#### C. Comparison With Related Work

The most relevant past research to our proposed architecture is [10] which uses a centralized controller to dynamically power off logic clusters. There are several differences between this work and our proposed architecture. First, the proposed architecture is a fine-grained approach which provides more power optimization opportunities than a coarse grained approach [20]. Second, the proposed architecture improves the dependability of the device by turning off short sensitive cells in unused parts of switch box and connection blocks while a coarse-grained approach is unable to improve the circuit dependability in unused parts of switch box and connection blocks. Lastly, we avoid using a centralized controller to improve the scalability of the proposed scheme. It is also worth mentioning that previous power gating techniques use a single sleep transistor to turn off unused resources [10]. Using a single sleep transistor to cut off the voltage node is not applicable when applying a fine-grained scheme. This is due to the fact that the voltage node of SRAM cells has to be

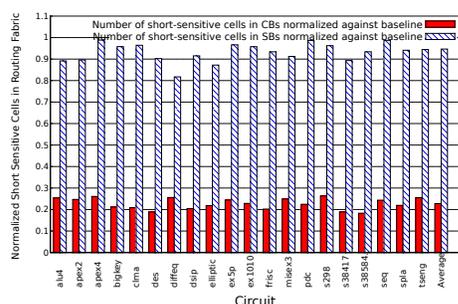


Fig. 6: Percentage of susceptible cells to soft errors in the proposed architecture for SBs and CBs normalized against the baseline architecture

actively grounded to avoid float nodes and therefore to avoid unwanted short faults.

### V. LIMITATIONS OF THE PROPOSED ARCHITECTURE

The major shortcoming of the proposed architecture is the increased buffer size in the routing fabric due to area overhead in the proposed architecture. The contribution of the power consumption of routing buffers to the total power consumption is different for various devices and might limit the number of CSRAMs when maximizing power efficiency. Moreover, Cu-SRAMs can negatively affect the configuration time when uploading a new configuration bitstream. Our Hspice simulations reveal that considering a group of 16 Cu-SRAM cells controlled by an SRAM cell, the configuration time of SRAM cells is increased between 1% to 23% depending on the number of SRAM cells that are bit flipped. Our Hspice analysis shows that power consumption is unaffected by the proposed architecture since the rate of reconfiguration is orders of magnitude less than the frequency of circuits operation.

In addition, one may argue that if a complex fabrication process such as the use of triple oxide, multiple-Vt, high K metal gate, FinFeT transistors, and variable transistor gate length are employed in SRDs, SRAM power consumption will become insignificant in the total power consumption. Despite the merits of device-level techniques, architectural schemes are still additive to device level schemes and can further help reduce power consumption. While a device level technique alone may allow few generations of scaling, device-level techniques together with architectural-level schemes can lead to few more generations of scaling. Nonetheless, it should be mentioned that device-level schemes need design refinements and generally are not straightforward for future technology generations while architecture-level solutions are typically applicable to a wide range of technology nodes and emerging technologies.

Another limitation of this work, not taken into consideration, is the adaptability of LUTs in commercial SRDs that allows several small functions to be employed instead of a single LUT-6. This limitation was imposed by our technology map, placement, and routing tools which do not target commercial SRDs.

### VI. CONCLUSION

In this paper, we presented a fine-grained dark silicon architecture for commonly used island style SRDs. The proposed architecture significantly reduces the number of active

configuration cells in LBs, CBs, and SBs. The reduced number of active configuration bits results in power savings up to 62%, 78%, and 15% in LBs, CBs, and SBs, respectively. The average power saving for these resources is 27%, 75%, and 4% while the worst case power saving is 1%, 67% and 0%, respectively. The reduced number of active configuration bits also results in fewer number of susceptible configuration bits to soft errors, resulting in improved circuit error rate. Taking area overheads into consideration, CBs and LUTs are the most appealing resources of SRDs for the proposed scheme while SBs provide less opportunity for power saving.

### REFERENCES

- [1] J. Rose, J. Luu, C. W. Yu, O. Densmore, J. Goeders, A. Somerville, K. B. Kent, P. Jamieson, and J. Anderson, "The vtr project: architecture and cad for fpgas from verilog to routing," in *International Symposium on Field Programmable Gate Arrays (FPGA)*, 2012, pp. 77–86.
- [2] Xilinx-Corporation, "Virtex-7 fpga family," Sep 2013. [Online]. Available: <http://www.xilinx.com/products/silicon-devices/fpga/virtex-7/>
- [3] F. Li, Y. Lin, L. He, D. Chen, and J. Cong, "Power modeling and characteristics of field programmable gate arrays," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 24, no. 11, pp. 1712–1724, 2005.
- [4] S. Borkar, N. P. Jouppi, and P. Stenstrom, "Microprocessors in the era of terascale integration," in *Design, Automation and Test in Europe (DATE)*, 2007, pp. 237–242.
- [5] M. B. Taylor, "Is dark silicon useful?: harnessing the four horsemen of the coming dark silicon apocalypse," in *Design Automation Conference (DAC)*, 2012, pp. 1131–1136.
- [6] H. Asadi, M. Tahoori, B. Mullins, D. Kaeli, and K. Granlund, "Soft error susceptibility analysis of sram-based fpgas in high-performance information systems," *IEEE Transactions on Nuclear Science*, vol. 54, no. 6, pp. 2714–2726, Dec. 2007.
- [7] H. Esmailzadeh, E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in *38th Annual International Symposium on Computer Architecture (ISCA)*, 2011, pp. 365–376.
- [8] "Power consumption at 40 and 45 nm," White Paper, Xilinx, April 2009.
- [9] A. Gayasen, Y. Tsai, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, and T. Tuan, "Reducing leakage energy in fpgas using region-constrained placement," in *International Symposium on Field Programmable Gate Arrays (FPGA)*, 2004, pp. 51–58.
- [10] A. A. Bsoul and S. J. Wilton, "An fpga architecture supporting dynamically controlled power gating," in *International Conference on Field-Programmable Technology (FPT)*, 2010, pp. 1–8.
- [11] Xilinx-Corporation, "Virtex-7 2000t faq," 2011.
- [12] A. Gayasen, K. Lee, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, and T. Tuan, "A dual-v dd low power fpga architecture," in *Field Programmable Logic and Application (FPL)*, 2004, pp. 145–157.
- [13] A. Rahman, S. Das, T. Tuan, and A. Rahut, "Heterogeneous routing architecture for low-power fpga fabric," in *IEEE Custom Integrated Circuits Conference*, 2005, pp. 183–186.
- [14] J. Lamoureux and S. J. Wilton, "On the interaction between power-aware fpga cad algorithms," in *IEEE/ACM International conference on Computer-aided design (ICCAD)*, 2003, p. 701.
- [15] K. McElvain, "Iwls93 benchmark set: Version 4.0," in *MCNC International Workshop on Logic Synthesis*, vol. 93, 1993.
- [16] R. Brayton and A. Mishchenko, "Abc: An academic industrial-strength verification tool," in *Computer Aided Verification*, 2010, pp. 24–40.
- [17] (2013) Predictive technology model (ptm). [Online]. Available: <http://ptm.asu.edu/>
- [18] M. A. Abdul-Aziz and M. B. Tahoori, "Soft error reliability aware placement and routing for fpgas," in *IEEE International Test Conference (ITC)*, 2010, pp. 1–9.
- [19] B. Gill, C. Papachristou, and F. Wolff, "A new asymmetric sram cell to reduce soft errors and leakage power in fpga," in *Design, Automation and Test in Europe (DATE)*, April 2007, pp. 1–6.
- [20] A. Rahman, S. Das, T. Tuan, and S. Trimmerger, "Determination of power gating granularity for fpga fabric," in *IEEE Custom Integrated Circuits Conference (CICC)*, 2006, pp. 9–12.