

Circuits and Architectures for Field Programmable Gate Array With Configurable Supply Voltage

Yan Lin, *Student Member, IEEE*, Fei Li, and Lei He, *Member, IEEE*

Abstract—Field programmable gate arrays (FPGAs) with supply voltage (Vdd) programmability have been proposed recently to reduce FPGA power, where the Vdd-level can be customized for FPGA circuit elements and unused circuit elements can be power-gated. In this paper, we first design novel Vdd-programmable and Vdd-gateable interconnect switches with minimal number of configuration SRAM cells. We then evaluate Vdd-programmable FPGA architectures using the new switches. The best architecture in our study uses Vdd-programmable logic blocks and Vdd-gateable interconnects. Compared to the baseline architecture similar to the leading commercial architecture, our best architecture reduces the minimal energy-delay product by 54.39% with 17% more area and 3% more configuration SRAM cells. Our evaluation results also show that LUT size 4 gives the lowest energy consumption, and LUT size 7 leads to the highest performance, both for all evaluated architectures.

Index Terms—Digital integrated circuits, field programmable gate arrays (FPGAs), power supplies.

I. INTRODUCTION

FIELD programmable gate arrays (FPGAs) provide an attractive design platform with low nonrecurring engineering (NRE) cost and a short time-to-market. Due to the large number of transistors required for field programmability and the low utilization rate of FPGA resources, existing FPGAs consume more power compared to ASICs [1]. As the process technology advances to nanometer technology and low-energy embedded applications are explored for FPGAs, power consumption becomes a crucial design constraint for FPGAs.

Several recent papers have studied FPGA power modeling and optimization. [2]–[4] presented power evaluation frameworks for generic parameterized FPGA architectures and showed that both interconnect and leakage power are significant for FPGAs in nanometer technologies. [5] quantified the leakage power of a commercial FPGA architecture in 90 nm technology. FPGA power optimization by computer-aided design (CAD) algorithms and novel circuits/architectures have also been studied. [6] proposed active leakage power reduction by reconfiguring the input vectors of multiplexers. [7] studied the interaction between a suite of power-aware FPGA CAD algorithms without changing the existing FPGA circuits and architectures, and showed that technology mapping and clustering algorithms are most effective in reducing power. The

following work focused on designing power-efficient FPGA circuits and architectures. [8] studied region-based power-gating and placement to reduce leakage power of unused FPGA logic blocks. [9] applied fine-grained power-gating to FPGA interconnects. [10], [11] proposed dual-Vdd and Vdd-programmable FPGA logic blocks as well as simple, yet effective CAD algorithms to leverage such logic blocks. Field programmability of supply voltage was shown to be necessary to reduce FPGA power using the dual-Vdd technique. The concept of programmable-Vdd introduced in [11] was further extended to the FPGA interconnects in [12]–[14].

Previously, conventional FPGA architecture evaluation has been performed using the metrics of area, delay and energy. [15] showed that LUT size 4 achieves the smallest area, and [16] showed that LUT size 5 or 6 leads to the best performance in nonclustered FPGAs. [17] evaluated cluster-based island style FPGAs using the metric of area-delay product in 0.35- μm technology, and showed that the range of LUT sizes from 4 to 6 and cluster sizes between 4 and 10 can produce the best area-delay product. The following work further extended FPGA architecture evaluation considering energy. [2] showed that LUT size 3 consumes the smallest energy in 0.35 μm technology. [4] showed that LUT size 4 consumes the smallest energy and LUT size 7 leads to the best performance in 100-nm technology.

However, the emerging power-efficient circuits and architectures may lead to different FPGA power characteristics, therefore, calling for an architecture evaluation considering these power optimization techniques. In this paper, we study Vdd-programmable FPGAs that were originally proposed in [10], [11]. We first design a new set of Vdd-programmable circuits and develop several new architecture classes for Vdd-programmable FPGAs with different levels of Vdd programmability for interconnects. We then study the effect of cluster and LUT sizes on FPGA area, energy and delay. Using the power evaluation framework from [4], [18], we evaluate the energy reduction by our new Vdd-programmable architecture classes compared to those with fixed Vdd-level.

The rest of the paper is organized as follows. Section II first describes the background related to FPGA architecture and evaluation methodology and then evaluates the baseline architecture class. Section III presents the novel circuit designs for Vdd-programmable and Vdd-gateable interconnect switches with a minimal number of configuration SRAM cells. Sections IV and V propose three Vdd-programmable architecture classes and evaluate their energy, delay and area with comparison to the baseline case. We conclude this paper in Section VI. An extended abstract of this paper was presented in [18].

Manuscript received October 20, 2004; revised April 7, 2005. This work is supported in part by the National Science Foundation under NSF CAREER Award CCR-0093273 and under NSF Grant CCR-0306682.

The authors are with the Department of Electrical Engineering, University of California, Los Angeles, CA 90034 USA (e-mail: lhe@ee.ucla.edu).

Digital Object Identifier 10.1109/TVLSI.2005.857180

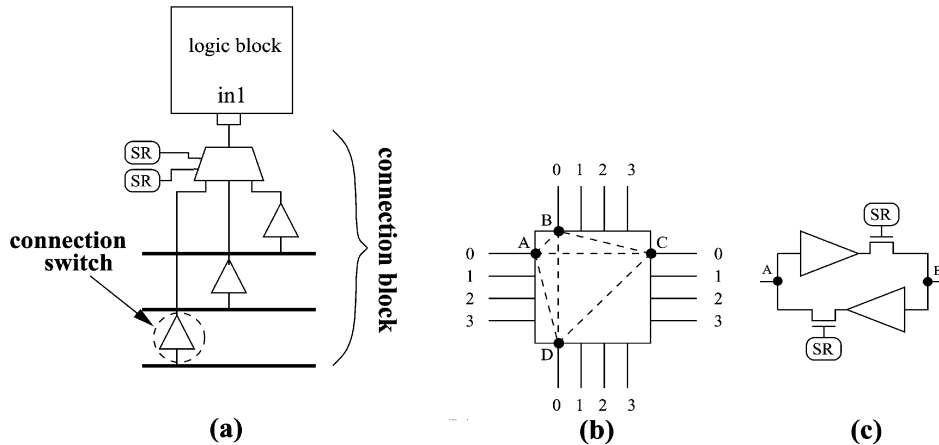


Fig. 1. (a) Connection block. (b) Switch block. (c) Routing switches.

II. BACKGROUND

A. Cluster-Based Island Style FPGAs

We assume cluster-based island style FPGA architecture [3], [19] for all classes of FPGAs studied in this paper. Each logic block includes N fully connected basic logic elements (BLEs). Each BLE includes one k -input lookup table (LUT) and one flip-flop (DFF). The combination of cluster size N and LUT size k is the architectural setting we evaluate in this paper. The logic blocks are surrounded by routing channels consisting of wire segments. The input and output pins of a logic block can be connected to the wire segments in routing channels via a *connection block* [see Fig. 1(a)]. A routing *switch block* is located at the intersection of a horizontal channel and a vertical channel. Fig. 1(b) shows a subset switch block [20], where the incoming track can be connected to the outgoing tracks with the same track number¹. We implement routing switches [see Fig. 1(c)] by tri-state buffers and use two tri-state buffers for each connection so that it can be programmed independently for either direction. We define an *interconnect segment* as a wire segment driven by a tri-state buffer or a buffer². We use the smallest square FPGA array for each benchmark circuit, and decide the routing channel width CW in the same way as the architecture study in [19], [21], i.e., $CW = 1.2 CW_{\min}$ where CW_{\min} is the minimum channel width required to route the given circuit successfully. The channel width CW represents a “low-stress” routing situation that usually occurs in commercial FPGAs for “average” circuits.

B. Evaluation Framework

This paper uses fpgaEVA-LP2 [4], [18] as the evaluation framework. fpgaEVA-LP2 includes a *BC-netlist* generator and a cycle-accurate power simulator *Psim*. The BC-netlist generator takes placement and routing results by VPR [19] and generates the Basic Circuit netlist (BC-netlist) annotated with post-layout capacitance and delay. *Psim* then performs cycle-accurate simulation on the BC-netlist to obtain FPGA

TABLE I
DEVICE AND INTERCONNECT MODEL AT 100 NM TECHNOLOGY

Device model			
normal-Vt	Vdd (V)	NMOS-Vt (V)	PMOS-Vt (V)
high-Vt	1.3	0.2607	-0.3030
	1.3	0.4693	-0.5454
Interconnect model			
wire width	wire spacing	wire thickness	dielectric const.
0.56 μ m	0.52 μ m	1.08 μ m	2.7

power consumption. There are three types of power sources in FPGAs, switching power, short-circuit power and static power. The first two types of power contribute to the dynamic power and can only occur when a signal transition happens at the gate output. There are two types of signal transitions. *Functional transition* is the necessary signal transition to perform the required logic functions between two consecutive clock ticks. *Spurious transition* or *glitch* is the unnecessary signal transition due to the imbalanced path delays to the inputs of a gate. Glitch power can be a significant portion of the dynamic power. The short-circuit power is consumed when both PMOS and NMOS transistors are turned on in a gate. The third type of power, static (leakage) power, is the power consumed when there is no signal transition for a gate or a circuit element.

In fpgaEVA-LP2, the power components of switching power, short-circuit power and static power for logic blocks are pre-calculated per switch or per unit time by SPICE simulation, and so is the leakage power for interconnects. The interconnect switching power is calculated by a switch-level model with extracted parasitics, and its short-circuit power is calculated as a portion of switching power. This portion can be pre-calculated by SPICE simulation for a variety of input signal transition time and load capacitances.

In this paper, we use the Berkeley predictive device model [22] and ITRS predictive interconnect model [23] at the 100-nm technology node. Table I summarizes the values of the key model parameters used throughout the rest of the paper. It has been shown in [4], [18] that fpgaEVA-LP2 achieves high fidelity as well as high accuracy. The average absolute error is 8.26% compared to SPICE simulation.

¹Without loss of generality, we assume subset switch block in this paper.

²We interchangeably use the terms of switch and buffer/tri-state buffer.

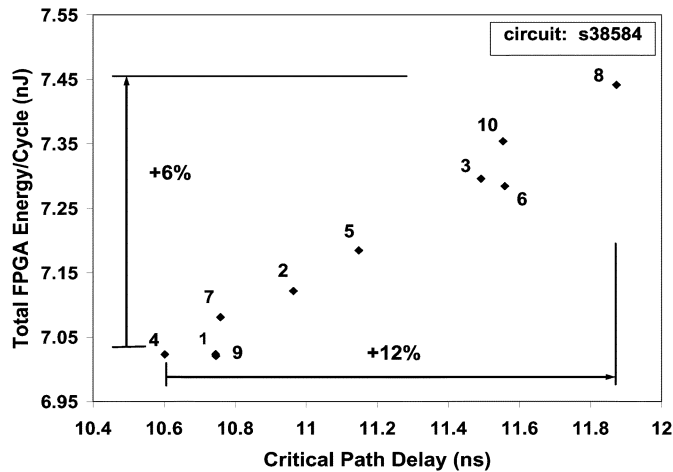


Fig. 2. Impact of random seed on FPGA energy and delay.

 TABLE II
 STATISTICS OF MCNC BENCHMARK CIRCUITS ($N = 10, k = 4$)

circuit	# of nets	# of logic blocks	# of I/O blocks
alu4	782	162	22
apex2	1246	213	41
apex4	849	134	28
bigkey	1542	294	426
clma	7995	1358	144
des	1325	218	501
diffeq	1291	195	103
dsip	1139	588	426
elliptic	2617	666	245
ex1010	3033	513	20
ex5p	834	194	71
frisc	3240	731	136
misex3	828	181	28
pdc	2933	624	56
s298	908	66	10
s38417	5426	982	135
s38584	4502	1046	342
seq	1138	274	76
spla	2091	461	122
tseng	918	305	174

C. Evaluation Methodology and Results for Baseline Architecture Class

Our architecture evaluation starts with the placement and routing results by VPR. For a given FPGA architecture and benchmark circuit, VPR can generate different placement and routing results by using different seeds in its placement algorithm. Fig. 2 shows the FPGA energy and delay using ten different VPR seeds for the same circuit s38584. We label the seed value beside each data point. The delay variation is 12% and the energy variation is 5%. This variation due to VPR seeds may affect our architecture evaluation. Because the delay variation is more sensitive to the VPR seeds than the energy variation, we decide to use the min-delay solution among all VPR seeds for every benchmark circuit. Note that the min-delay solution often consumes low energy. For the architecture evaluation in this paper, Energy (E), Delay (D), Energy-Delay Product (ED) and Area (A) are always the geometric means of those values over the 20 largest MCNC benchmark circuits [24] in Table II.

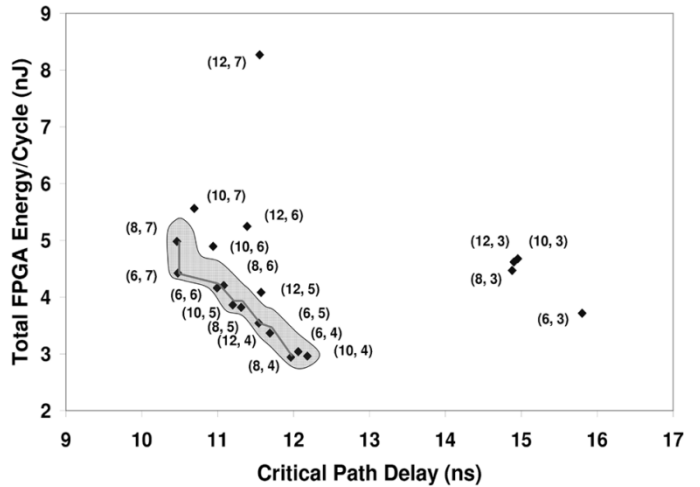


Fig. 3. Energy-delay tradeoff for single-Vdd dual-Vt FPGA class (Class0). The polyline represents the strictly dominant architectures and the enclosed area covers the relaxed dominant architectures.

Using the above methodology, we perform an architecture evaluation for the single-Vdd dual-Vt FPGA architectures from [10], defined as FPGA architecture *Class0* in this paper. The entire FPGA uses the uniform supply voltage of 1.3 V, but high-Vt is applied to all FPGA configuration SRAM cells to reduce SRAM leakage power. The high-Vt configuration cells do not incur runtime performance degradation because they are constantly in read status after an FPGA is configured, and their read and write operations are irrelevant to the runtime performance. This high-Vt SRAM technique has already been used in commercial FPGAs and is also applied to all FPGA architectures in this paper.

Fig. 3 presents the evaluation results for single-Vdd dual-Vt FPGA Class0. Each data point in the figure is an FPGA architecture represented by a tuple (N, k) , where $N \in \{6, 8, 10, 12\}$ is the cluster size and $k \in \{3, 7\}$ is the LUT size. If one architecture (N_1, k_1) has smaller delay and less energy consumption than another architecture (N_2, k_2) , we say that architecture (N_1, k_1) is superior to (N_2, k_2) . We define *strictly energy-delay dominant architectures* as the set of superior data points in the entire energy-delay tradeoff space. Those architectures are highlighted by the polyline in Fig. 3. Our results also show that some of the architectures may have fairly similar energy and delay values such as architectures $(N = 8, k = 4)$, $(N = 6, k = 4)$ and $(N = 10, k = 4)$, and all of of them can be considered as valid solutions. To avoid pruning out architectures with slightly worse energy and delay, we further define *relaxed energy-delay dominant architectures*. If architectures (N_1, k_1) and (N_2, k_2) have both an energy and delay difference of less than r (*relaxation parameter*), then neither of them can dominate the other one. With $r = 2\%$ in this paper, the relaxed dominant architectures are data points inside the highlighted area in Fig. 3. Min-delay and min-energy architectures are the two extreme cases among those energy-delay dominant architectures. The min-delay architecture is $(N = 8, k = 7)$ and the min-energy architecture is $(N = 8, k = 4)$ for the FPGA Class0 in Fig. 3, and the energy and delay differences between the two extreme cases are 57% and 14%, respectively. It shows that a significant tradeoff between energy and delay can be obtained

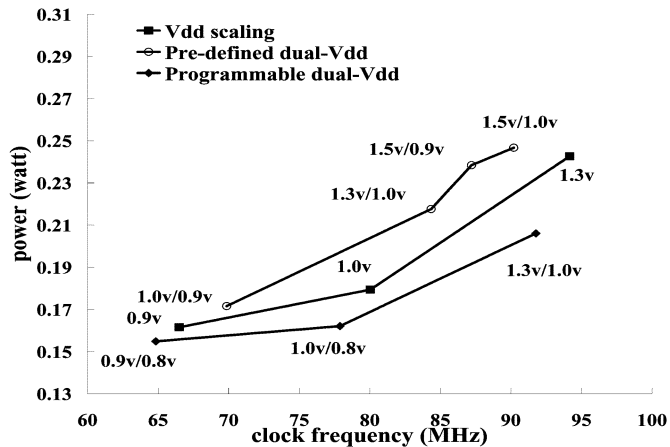


Fig. 4. Comparison of three power reduction solutions for benchmark s38584.

by varying the cluster size and LUT size. Note that our min-energy architecture ($N = 8, k = 4$) is also the min-area architecture found by [17]. Commercial FPGAs such as Xilinx Virtex-II [25] coincidentally use a cluster size of 8 and an LUT size of 4, indicating that their architectures may have used a min-area solution, which doubles to be a min-energy architecture in the single-Vdd architecture class.

D. Field Programmability of Vdd Supply Versus Pre-Determined Vdd Pattern

A higher supply voltage leads to a higher performance but a larger power consumption. Leveraging this, Vdd scaling lowers the supply voltage of an entire design or a large circuit module for power reduction. Alternatively, dual-Vdd applies high supply voltage (VddH) to the logic on the critical path and low supply voltage (VddL) to the logic not on the critical path. Under given performance constraints, dual-Vdd or multi-Vdd is able to reduce more power than Vdd scaling does for ASICs [26]–[29]. The success of dual-Vdd is due to the fact that ASIC designers are able to customize Vdd patterns for different applications. However, such flexibility does not exist for FPGAs using a pre-defined dual-Vdd. Hence, pre-defined dual-Vdd FPGA fabric may have limited power reduction compared to single-Vdd scaling.

Assuming generic cluster-based FPGA architecture and MCNC benchmark circuit s38584, the power and performance curves for both Vdd scaling and fixed dual-Vdd patterns [10] are shown in Fig. 4. The Vdd-level is decided uniformly for all logic blocks in Vdd scaling. Furthermore, each logic block has a pre-determined Vdd-level in dual-Vdd, and various dual-Vdd patterns are tried to obtain the best result. It is easy to see from this figure that the fixed dual-Vdd pattern consumes more power than Vdd scaling for a given frequency. This is due to the fact that the pre-determined Vdd pattern imposes extra placement constraints to match Vdd-level while increasing interconnect delay (and power) [11]. In contrast, the FPGA fabric using programmable dual-Vdd logic blocks [11] reduces power significantly compared to Vdd scaling. It is clear that field programmability of the power supply is required to achieve FPGA power reduction via dual-Vdd. Therefore, we use field programmable dual-Vdd for logic blocks and interconnects in this paper.

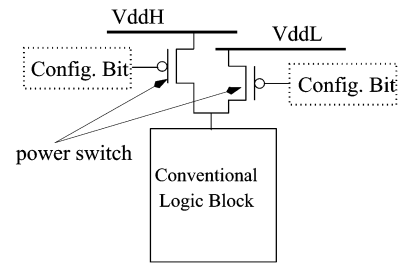


Fig. 5. Vdd-programmable logic block.

III. FPGA CIRCUITS FOR VDD PROGRAMMABILITY

A. Previous Work and Section Overview

Programmable dual-Vdd has been introduced in [10], [11] and applied to logic blocks to reduce FPGA power. We define Vdd programmability as the flexibility to select Vdd-levels for used circuit elements and the capability to power-gate unused circuit elements. Fig. 5 shows the Vdd-programmable logic block. Two extra PMOS transistors, called *power switches* or *power transistors*³, are inserted between the conventional logic block and the dual-Vdd power rails for Vdd selection and power-gating. We use normal-Vt power transistors with gate-boosting same as those in [11] to reduce area overhead and achieve effective leakage reduction. 210X minimum width PMOS power transistors [11] are used for a logic block containing 10 4-LUTs with delay overhead bounded by 5% when considering the worst case simultaneous switching current.

In this section, we further apply Vdd programmability to interconnect switches. Normal-Vt power transistors with gate-boosting are used for interconnect switches. We design two types of interconnect switches, a Vdd-programmable switch and a Vdd-gateable switch. A Vdd-programmable switch provides three power states which include VddH, VddL and power-gating. Different from a Vdd-programmable switch, a Vdd-gateable switch only provides two power states between a pre-determined Vdd and power-gating, but it has significantly fewer configuration SRAM cells for Vdd programmability. The detailed circuit designs of Vdd-programmable and Vdd-gateable switches are discussed below.

B. Vdd-Programmable Interconnect Switch

Fig. 6 shows the design of Vdd-programmable interconnect switches (both routing switch and connection switch). A Vdd-level converter is needed whenever a VddL interconnect switch drives a VddH interconnect switch. In other cases, the level converter can be bypassed. As shown in Fig. 6(a), a pass transistor M1 and a MUX together with a configuration SRAM cell can be used to implement a configurable level conversion. The transistor M1 is used to prevent signal transitions from propagating through the level converter when it is bypassed, therefore eliminating the dynamic power of an unused level converter. Only one configuration bit is needed to realize the level converter selection and signal gating for an unused level converter. The same asynchronous level converter circuit in [10], [11] is used

³The terms power switch and power transistor are used interchangeably in this paper.

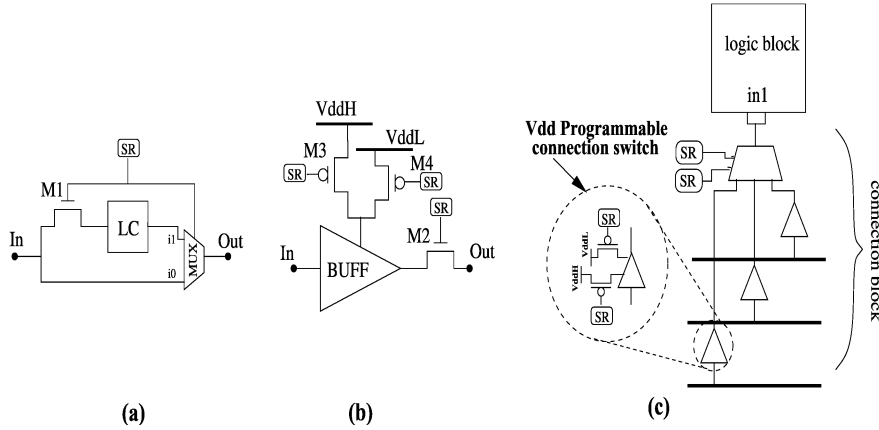


Fig. 6. (a) Configurable level conversion; (b) Vdd-programmable routing switch; (c) Vdd-programmable connection block. (SR stands for SRAM cell and LC stands for level converter).

TABLE III

DELAY AND POWER OF A VDD-PROGRAMMABLE ROUTING SWITCH. WE USE $7 \times$ MINIMUM WIDTH TRI-STATE BUFFERS FOR ROUTING SWITCHES AND $4 \times$ MINIMUM WIDTH PMOS TRANSISTORS FOR POWER TRANSISTORS

Vdd	routing switch delay (ns)		energy per switch (Joule)	
	w/o Vdd programmability	w/ Vdd programmability (increase %)	w/o Vdd programmability	w/ Vdd programmability
1.3v	5.90E-11	6.86E-11 (+16.27%)	3.3049E-14	3.2501E-14
1.0v	6.45E-11	7.55E-11 (+17.05%)	1.6320E-14	1.6589E-14

and sized to achieve a bounded delay with minimum power consumption.

For the Vdd-programmable routing switch in Fig. 6(b), two PMOS power transistors M3 and M4 are inserted between the tri-state buffer and VddH, VddL power rails, respectively. Turning off one of the power transistors can select a Vdd-level for the routing switch. By turning off both power transistors, an unused routing switch can be power-gated. The pass transistor M2 must be kept to prevent a sneak path [30], i.e., a current path that flows from Vdd to ground through a set of “on” transistors which belong to different gates. SPICE simulation shows that power-gating the routing switch can reduce leakage power of an unused routing switch by a factor of over 1000. There are power and delay overheads associated with the insertion of power transistors. As shown in Table III, the dynamic power overhead is almost negligible. This is because the power transistors stay either ON or OFF after configuration and there is no charging or discharging at their source/drain capacitors. The delay overhead associated with the power transistor insertion can be bounded when the power transistor is properly sized. Another type of routing resources is the connection block in Fig. 6(c). The multiplexer-based implementation chooses only one track in the routing channel and connects it to the logic block input pin. The buffers between the routing tracks and the multiplexer are connection switches. Similar to the routing switch, programmable-Vdd is also applied to the connection switch. The multiplexer must be kept to prevent the sneak path.

There are three SRAM cells for each Vdd-programmable routing switch in Fig. 6(b). For a connection block containing N Vdd-programmable connection switches in Fig. 6(c), there are $2N + \lceil \log_2 N \rceil$ configuration SRAM cells, among which

$\lceil \log_2 N \rceil$ SRAM cells are for multiplexer and the other $2N$ extra SRAM cells are for N Vdd-programmable connection switches. We can use combinational logic such as a decoder to reduce the number of extra SRAM cells introduced by Vdd programmability. As shown in Fig. 7(a), we first define a Vdd-programmable switch module with three signal ports, $VddH_En$, $VddL_En$ and $Pass_En$. By setting these three control signals, we can set the Vdd-programmable switch for Vdd selection and power-gating.

We design a new Vdd-programmable routing switch in Fig. 7(b). $Pass_En$ can be generated by $VddH_En$ and $VddL_En$ with a NAND2 gate. Table IV summarizes the configurations for the Vdd-programmable routing switch and the truth table of the relevant control signals. Similarly, Fig. 7(c) presents a new design of Vdd-programmable connection block with reduced configuration SRAM cells. For a connection block containing N connection switches, we use a $\lceil \log_2 N \rceil : N$ decoder and $2N$ NAND2 gates as the control logic. There is a disable signal $Dec_Disable$ for the decoder. Each decoder output is connected to $Pass_En$ of one connection switch. Setting $Pass_En$ of a connection switch to “0” can power-gate this switch by setting both $VddH_En$ and $VddL_En$ to “1” with NAND2 gates. When the whole connection block is not used, all N outputs of the decoder are set to “0” to power-gate all the connection switches by asserting $Dec_Disable$. When the connection block is in use, $Dec_Disable$ is not asserted. By using $\lceil \log_2 N \rceil$ configuration bits for the decoder, only one $Pass_En$ is set to “1” and others are set to “0”, i.e., only one connection switch inside the connection block is selected, connecting one track to the logic block input, while other unused connection switches are power-gated. Another configuration bit Vdd_Sel is used to select the Vdd-level for the selected connection switch. Table V summarizes configurations for Vdd-programmable connection switch and the truth table of relevant control signals.

By replacing the conventional connection switch with the new Vdd-programmable switch in Fig. 7(a), the pass transistor in the Vdd-programmable switch can now prevent sneak paths. Therefore, the multiplexer implemented by the NMOS pass transistor tree can be removed from the new Vdd-programmable connection block. Table VI shows the delay and power of a new

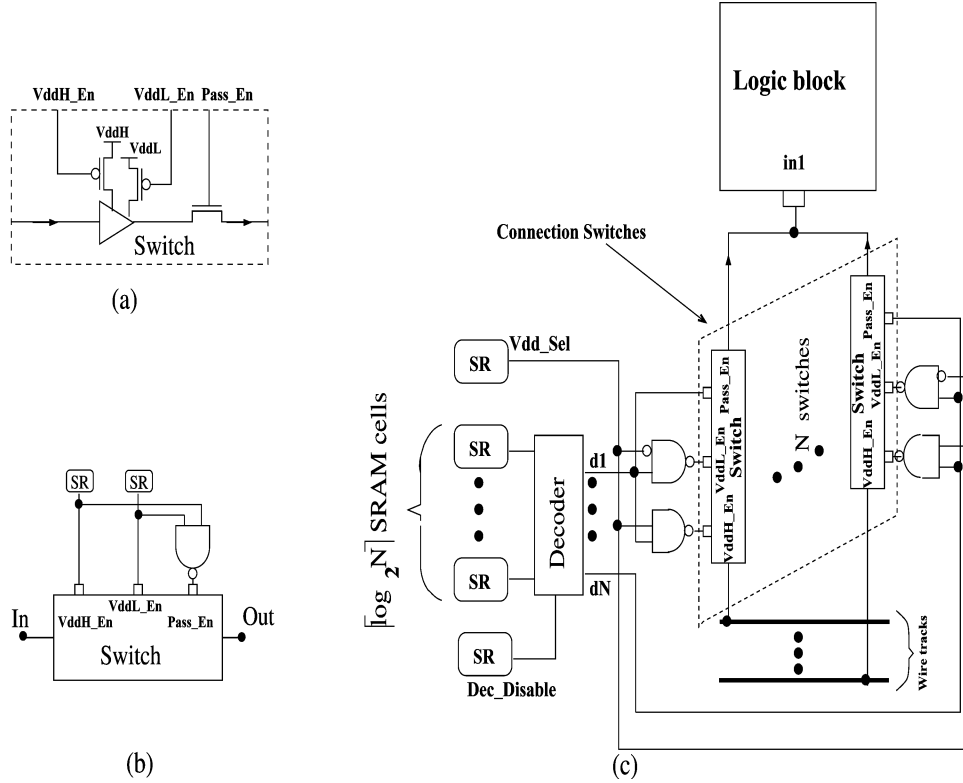


Fig. 7. (a) Vdd-programmable switch. (b) SRAM-efficient Vdd-programmable routing switch. (c) SRAM-efficient Vdd-programmable connection block.

TABLE IV
CONFIGURATIONS FOR A VDD-PROGRAMMABLE ROUTING SWITCH

state	$VddH_En$	$VddL_En$	$Pass_En$
VddH	0	1	1
VddL	1	0	1
power-gated	1	1	0

Vdd-programmable connection block. The delay and dynamic energy per signal transition are reduced by 28% and 19%, respectively, when Vdd-level is 1.3 V. The delay and power reduction is due to removing the multiplexer.

For a connection block containing N connection switches, only $\lceil \log_2 N \rceil + 2$ configuration SRAM cells are needed to provide Vdd selection and power-gating capability for each individual connection switch inside the connection block. Compared to a conventional connection block, only two extra configuration SRAM cells are introduced. Similar to the configuration SRAM cell, we use high- Vt transistors for control logic to reduce leakage overhead since the delay of control logic will not affect system runtime performance. We also use minimum width transistors for control logic to reduce area overhead. In this paper, we use the same area model from [19], in which the area is counted by the number of minimum width transistor areas while considering the parallel diffusions technique for large transistors. Given a transistor with channel width W , the transistor area measured by the minimum width transistor with channel width W_{min} is

$$\text{Area}(W) = 0.5 + \frac{W}{2 \cdot W_{min}}. \quad (1)$$

Table VII compares the number of configuration SRAM cells, leakage and area between the Vdd-programmable routing switch/connection blocks in Figs. 6 and 7, respectively. The Vdd-programmable routing switch and connection block in Fig. 7, called *SRAM-efficient switches*, have smaller area and less leakage, and will be used in the rest of the paper.

C. Vdd-Gateable Interconnect Switch

Compared to the Vdd-programmable switch, a Vdd-gateable interconnect switch only provides two power states between a pre-determined Vdd-level and power-gating, but it can dramatically reduce the number of extra SRAM cells for Vdd programmability. Fig. 8(a) shows the circuit design for a Vdd-gateable switch. Based on a conventional tri-state buffer, we insert a PMOS transistor M2 between the power rail and the tri-state buffer to provide the power-gating capability⁴. When a switch is not used, transistor M1 is turned off by the configuration cell SR. At the same time, we can turn off M2 to power-gate the unused switch. Similarly, both M1 and M2 are turned on by the configuration cell SR when the switch is used. Thus, we do not need to introduce an extra SRAM cell for power-gating capability. Fig. 8(b) presents Vdd-gateable routing switches. We can reduce leakage power by a factor of over 1000 for an unused switch when it is power-gated. Similar to the Vdd-programmable switch, the pass transistor M1 must be kept to prevent the sneak path. However, there is a delay overhead associated with the M2 insertion. We properly size M2 for the tri-state buffer to achieve a delay increase by 16%.

⁴Transistor M1 can be placed between the ground and the NMOS transistor of the buffer.

TABLE V
CONFIGURATIONS FOR A VDD-PROGRAMMABLE CONNECTION SWITCH

state	<i>Dec_Disable</i>	<i>Vdd_Sel</i>	<i>Pass_En</i>	<i>VddH_En</i>	<i>VddL_En</i>
power-gated	1	-	0	-	-
power-gated	0	-	0	-	-
VddH	0	1	1	0	1
VddL	0	0	1	1	0

TABLE VI
DELAY AND POWER OF NEW VDD-PROGRAMMABLE CONNECTION BLOCK. WE USE $4 \times$ MINIMUM WIDTH TRI-STATE BUFFER FOR CONNECTION SWITCHES AND $1 \times$ MINIMUM WIDTH PMOS TRANSISTOR FOR POWER TRANSISTORS

Vdd	connection switch delay (ns)		energy per switch (Joule)	
	w/o Vdd programmability	w/ Vdd programmability (increase %)	w/o Vdd programmability	w/ Vdd programmability (increase %)
1.3v	2.93E-10	2.10E-10 (-28.33%)	3.84E-14	3.11E-14 (-19.01%)
1.0v	3.70E-10	2.22E-10 (-40.00%)	3.09E-14	2.04E-14 (-33.98%)

Similar to the Vdd-programmable switch, the dynamic power overhead associated with the insertion of a PMOS M2 is almost negligible because transistor M2 is always ON when the routing switch is used and there is no charging or discharging at its source–drain capacitors.

The design of the Vdd-gateable connection block is shown in Fig. 8(c). We only need $\lceil \log_2 N \rceil$ configuration SRAM cells to control N connection switches in a connection block via a decoder with complementary outputs to achieve the power-gating capability for each connection switch at the same time. We use another configuration bit, *Dec_Disable*, to disable the decoder when we apply power-gating to the whole connection block. Similar to the SRAM-efficient design of the Vdd-programmable connection block, we use high- V_t and minimum width transistors for the decoder to reduce leakage and area overhead. Alternatively, N configuration SRAM cells can be used to control the same number of connection switches without using the decoder. Table VIII compares the number of SRAM cells, leakage and area for a nondecoder based and decoder based connection blocks containing 32 connection switches. The decoder based Vdd-gateable connection block consumes less area and leakage power, and will be used in the rest of this paper. Similar to the SRAM-efficient Vdd-programmable connection block, the decoder based Vdd-gateable connection block reduces the delay and energy per signal transition by 28% and 19% when compared to the conventional connection block due to removing the multiplexer.

IV. ARCHITECTURE EVALUATION FOR VDD-PROGRAMMABLE FPGAS

In this section, we first evaluate architecture classes *Class1* and *Class2* for Vdd-programmable FPGAs. Both apply programmable dual-Vdd to each logic block. In addition, *Class1* applies programmable dual-Vdd to each interconnect segment, and inserts a configurable level conversion circuit in front of each routing/connection switch as well as at the inputs/outputs of the logic blocks. *Class2* uses Vdd-gateable routing/connection switches in FPGA interconnects. Therefore, the interconnect switches in architecture *Class2* only have two configurable states, high-Vdd (VddH) and power-gating. As we use VddH for

interconnects in architecture *Class2*, level converters are only needed at the logic block outputs, but not at the logic block inputs nor in routing channels. All these architecture classes (with *Class3* to be presented in Section V) are summarized in Table IX.

We apply a simple yet practical design flow from [11]. Starting with a single-Vdd gate level netlist, we apply technology mapping and timing-driven packing [19] to obtain the single-Vdd cluster-level netlist. We then perform single-Vdd timing-driven placement and routing by VPR [19] to generate the basic circuit netlist (BC-netlist). We calculate the power sensitivity $\Delta P / \Delta V_{dd}$, which is the power reduction by changing VddH to VddL, for each circuit element. The total power P includes both switching power P_{sw} and leakage power P_{lkg} . For each node i , we have switching power $P_{sw}(i) = 0.5 f_{clk} \cdot E_i \cdot C_i \cdot V_{dd}^2$, where E_i and C_i are transition density and load capacitance, and leakage power $P_{lkg} = I_{lkg}(V_{dd}) \cdot V_{dd}$. We pre-characterize I_{lkg} and device delay at each Vdd level using SPICE simulation. We assume that the transition density for each circuit element will not change when some circuit elements are assigned to VddL, and therefore we only need to calculate the power sensitivity for each circuit element once. A greedy algorithm is performed for Vdd assignment considering iteratively updated timing slack. Assuming that all the circuit elements are initially assigned to VddH, we iteratively perform the following steps. Timing analysis is performed to obtain the circuit elements on the path with the largest timing slack. We then assign VddL to the element with the largest power sensitivity. The configurable level converter can be enabled as needed. After updating the circuit timing, we accept the assignment if the critical path delay does not increase. Otherwise, we reject the assignment and restore the circuit element supply voltage to VddH. In either case, the circuit element will be marked as 'tried' and will not be re-visited in subsequent iterations. After the dual-Vdd assignment, we obtain a dual-Vdd BC-netlist without degrading system performance. For FPGA *Class1*, the Vdd assignment unit is a logic block or an interconnect switch. For FPGA *Class2*, the Vdd assignment unit is a logic block. For both *Class1* and *Class2*, power-gating is applied to all unused logic blocks and interconnect switches. Finally, we perform the energy and delay evaluation for the dual-Vdd design.

Fig. 9 presents the energy-delay tradeoff in terms of different architectures, i.e., different combinations of cluster size N and LUT size k , for three FPGA classes: *Class0*, *Class1* and *Class2*. Considering the VddL/VddH ratio between 0.6 and 0.7 suggested in [31], we use 1.3 v for VddH and 0.8 v for VddL in our experiments. We only show the relaxed dominant architectures in the figure and the polylines represent the strictly dominant architectures. Similar to the baseline FPGA *Class0*, the min-delay

TABLE VII

COMPARISON OF THE NUMBER OF SRAM CELLS, LEAKAGE AND AREA BETWEEN A VDD-PROGRAMMABLE ROUTING SWITCH/CONNECTION BLOCK AND AN SRAM-EFFICIENT VDD-PROGRAMMABLE ROUTING SWITCH/CONNECTION BLOCK. WE USE 32:1 CONNECTION BLOCK AND THE CONTROL LOGIC FOR SRAM-EFFICIENT DESIGN CONTAINS A STANDARD 5:32 DECODER AND 64 NAND2 GATES. AREA IS PRESENTED IN MINIMUM WIDTH TRANSISTOR AREAS

Vdd-programmable routing switch										
Figure 6 (b)			SRAM-efficient design in Figure 7 (b)				Figure 7 (b) compared to Figure 6 (b)			
SRAM cells			SRAM cells			NAND2				
number	leakage (watt)	area	number	leakage (watt)	area	leakage (watt)	area	Δ number of SRAM cells	Δ leakage (watt)	Δ area
3	2.32E-8	21.87	2	1.55E-8	14.58	3.49E-10	2.50	-1	-7.38E-9 (-33%)	-4.79 (-22%)
32:1 Vdd-programmable connection block										
Figure 6 (c)			SRAM-efficient design in Figure 7 (c)				Figure 7 (c) compared to Figure 6 (c)			
SRAM cells			SRAM cells			control logic				
number	leakage (watt)	area	number	leakage (watt)	area	leakage (watt)	area	Δ number of SRAM cells	Δ leakage (watt)	Δ area
69	5.32E-7	503.01	7	5.42E-8	43.74	3.30E-8	311	-62	-4.56E-7 (-86%)	-148.27 (-29%)

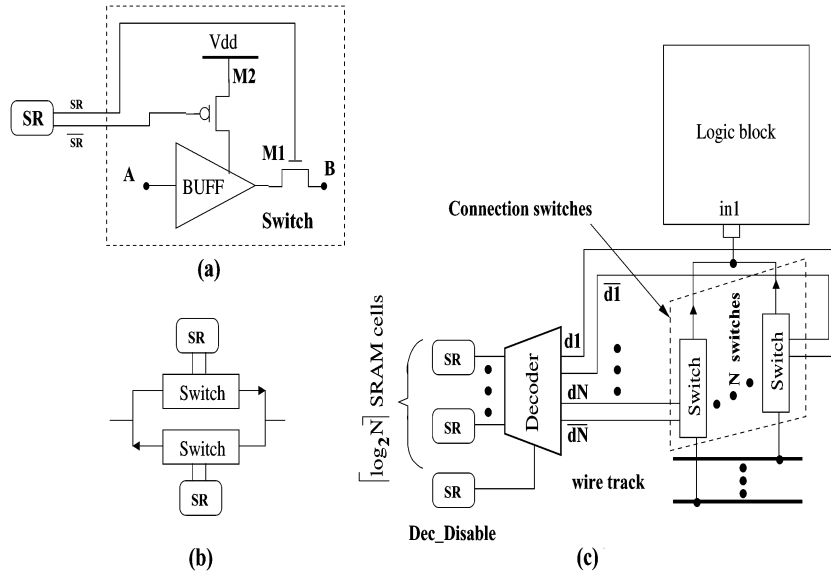


Fig. 8. (a) Vdd-gateable switch. (b) Vdd-gateable routing switches. (c) Vdd-gateable connection switches. (SR stands for SRAM cell).

TABLE VIII

COMPARISON OF THE NUMBER OF SRAM CELLS, LEAKAGE AND AREA BETWEEN A 32:1 NON-DECODER BASED VDD-GATEABLE CONNECTION BLOCK AND A 32:1 DECODER BASED VDD-GATEABLE CONNECTION BLOCK. FOR THE DECODER BASED VDD-GATEABLE CONNECTION BLOCK, WE USE A 5:32 DECODER WITH COMPLEMENTARY OUTPUT. AREA IS PRESENTED IN MINIMUM WIDTH TRANSISTOR AREAS

32:1 Vdd-gateable connection block										
non-decoder based connection block			decoder based connection block				compared to baseline: w/o decoder			
SRAM cells			SRAM cells			5:32 decoder				
number	leakage (watt)	area	number	leakage (watt)	area	leakage (watt)	area	Δ number of SRAM cells	Δ leakage (watt)	Δ area
32	2.47E-7	233.28	6	4.63E-8	43.74	2.00E-8	94.25	-26	-1.81E-7 (-73%)	-95.29 (-41%)

TABLE IX

SUMMARY OF BASELINE ARCHITECTURE CLASS AND VDD-PROGRAMMABLE ARCHITECTURE CLASSES

Architecture Class	Logic block	Interconnect
Class0 (baseline)	single Vdd	single Vdd
Class1	programmable dual-Vdd	programmable dual-Vdd w/ level converters
Class2	programmable dual-Vdd	VddH and Vdd-gateable
Class3	programmable dual-Vdd	programmable dual-Vdd w/o level converters

architecture is ($N = 8, k = 7$) for both Class1 and Class2. The min-energy architecture is ($N = 8, k = 4$) for Class1 and

($N = 12, k = 4$) for Class2. This shows that LUT size 7 gives the best performance and LUT size 4 leads to the lowest energy consumption for these Vdd-programmable FPGAs.

We then use the metrics of energy E , delay D and energy-delay product ED to compare the two classes of Vdd-programmable FPGAs (Class1 and Class2) and the baseline FPGA (Class0). We first compare the min-energy and min-delay architectures within each FPGA architecture class. As shown in Table X, in terms of min-energy architecture, Class1 and Class2 reduce energy by 28.57% and 54.08% respectively when compared to the baseline architecture class. In terms of min-delay architecture, the delay increase due to Vdd programmability is




Fig. 9. Energy and delay tradeoff for the baseline single-Vdd dual-Vt FPGA (Class0) and the two classes of Vdd-programmable FPGAs (Class1 and Class2). The figure only shows relaxed energy-delay dominant solutions and the strictly dominant solutions are represented by polylines.

TABLE X
COMPARISON BETWEEN VDD-PROGRAMMABLE FPGAs (CLASS1, CLASS2 AND CLASS3) AND THE BASELINE FPGA (CLASS0) USING ENERGY E , DELAY D AND ENERGY-DELAY PRODUCT (ED)

Arch. Class \rightarrow	Class0 (baseline)	Class1	Class2	Class3
min- E arch. (N,k)	(8,4)	(8,4)	(12,4)	(12,4)
energy (nJ/cycle)	2.94	2.10	1.35	1.18
energy saving (%)	-	28.57%	54.08%	59.86%
min- D arch. (N,k)	(8,7)	(8,7)	(8,7)	(8,7)
delay (ns)	10.46	10.82	10.82	10.82
delay increase (%)	-	3%	3%	3%
min- ED arch. (N,k)	(8,4)	(8,4)	(12,4)	(12,4)
ED product (nJ · ns)	35.19	26.05	16.05	14.03
ED reduction	-	25.97%	54.39%	60.13%

only 3% for both Class1 and Class2. We also use the min-ED (i.e., the minimum energy-delay product) architecture within each architecture class and obtain the ED product reduction. FPGA Class1 and Class2 reduce the ED product by 25.97% and 54.39%, respectively.

V. IMPROVED FPGA ARCHITECTURES

A. FPGA Architectures and Related CAD Algorithm

Vdd-programmable interconnects can reduce interconnect dynamic energy, which is not available in Vdd-gateable interconnects. However, as presented in Section IV, fully Vdd-programmable FPGA architecture Class1 consumes more energy than FPGA architecture Class2 which uses Vdd-gateable interconnects. This is because of the leakage overhead of the large number of Vdd-level converters in routing channels, which provides Vdd programmability for each individual interconnect switch. To achieve a better energy-delay tradeoff, we design an improved fully Vdd-programmable FPGA architecture Class3. It uses the same SRAM-efficient interconnect switches as FPGA architecture Class1, but inserts level converters only at logic block inputs and outputs. Since there is no level converter in routing channels, we need a CAD algorithm to guarantee

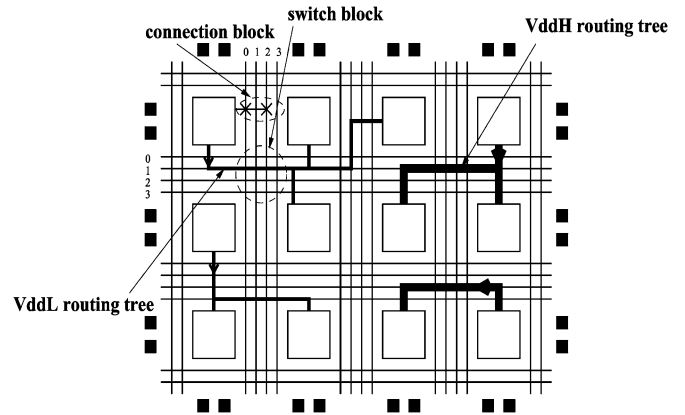


Fig. 10. Improved fully Vdd-programmable FPGA architecture Class3. No level converter is inserted in routing tracks.




Fig. 11. Energy and delay tradeoff for all FPGA architecture classes. The figure only shows relaxed energy-delay dominant solutions and the strictly dominant solutions are represented by polylines.

that no VddL interconnect switch drives VddH interconnect switches. We tackle the problem by choosing a routing tree as the Vdd assignment unit. Similar to FPGA Class1, the same design flow and the sensitivity-based Vdd-level assignment algorithm is used to decide the Vdd-level for each interconnect routing tree. The only difference is that we use an interconnect routing tree as the assignment unit for FPGA Class3 while an interconnect switch is used as the assignment unit for Class1. Since two routing trees will not intersect with each other in routing channels, we do not need level converters in routing channels⁵. Fig. 10 illustrates the situation where a VddH routing tree and VddL routing tree can share the same routing track without level converters in routing channels.

B. Energy and Delay Evaluation

In this section, we evaluate the improved fully Vdd-programmable architecture Class3. Fig. 11 shows the energy-

⁵As we perform Vdd-level assignment after routing, there is no impact on routability using a routing tree as the assignment unit.

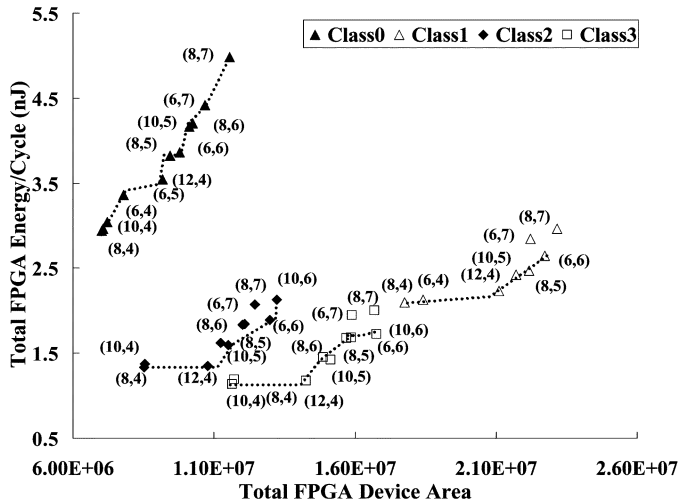


Fig. 12. Energy versus area curve for all architecture classes. This figure only shows relaxed energy-delay dominant solutions and min-area solution within each FPGA class. The polylines represent the lowest energy-area envelop. Area is measured in minimum width transistor areas.

delay evaluation for the improved architecture Class3 compared to the evaluation results for architecture Class0, Class1 and Class2. As shown in Fig. 11, we can see that the improved architecture Class3 achieves better energy-delay tradeoff than architecture Class1, and is even better than Class2. This is because FPGA Class3 removes the level converters in routing channels, but still can reduce interconnect dynamic energy. This reduction is not available in architecture Class2 which uses Vdd-gateable interconnect switches.

Similar to Class0, the min-delay architecture is ($N = 8$, $k = 7$) for Class3. The min-energy architecture is ($N = 12$, $k = 4$) for Class3 ($N = 12$, $k = 4$) also gives the minimum energy-delay product ED in architecture Class3. We can see that for our improved FPGA architecture Class3, again, LUT size 7 gives the best performance and LUT size 4 leads to the lowest energy consumption. Compared to the min-energy (min-delay) architecture within baseline architecture Class0, the min-energy architecture in Class3 reduces energy by 59.86%, and the min-delay architecture in Class3 has a 3% delay overhead due to Vdd programmability. The min-ED architecture in FPGA Class3 reduces energy-delay product ED by 60.13%. As shown in Table X, FPGA Class3 gives the lowest energy as well as the lowest ED .

C. Energy Versus Area

In this paper, area is measured by the total device area because FPGA area is dominated by devices [19]. Fig. 12 presents the energy-area curve for all FPGA architecture classes. The total device area includes both the logic block and interconnect device area. The area overhead of extra configuration SRAM cells, power transistors, control logics, and Vdd-level converters are included. In this figure, we show the relaxed ED-dominant architectures as well as the min-area architecture in each FPGA class. The min-area architecture is ($N = 8$, $k = 4$) for Class0 and Class1, ($N = 10$, $k = 4$) for Class2 and Class3. We can see that LUT size 4 not only gives the lowest energy consumption, but also gives the minimum area. In FPGA Class2 and Class3,

the min-area architecture ($N = 10$, $k = 4$) consumes similar energy as the min-energy architecture ($N = 12$, $k = 4$) but it has a much smaller area.

Vdd programmability increases the total number of SRAM cells required to store those extra configuration bits. However, SRAM cells consume extra area and are vulnerable to soft errors and the total number of SRAM cells should be minimized. Table XI presents the increase in SRAM cell number and the total area overhead due to Vdd programmability. The SRAM cells include those used in LUTs. Only dominant architectures are presented in the table. Vdd-programmable FPGA Class1 increases the SRAM cell number by 132%. This shows that a large number of extra SRAM cells are needed to provide fine-grained Vdd programmability for interconnects. FPGA Class2 only increases the SRAM cell number by 3% because only two power states (VddH and power-gating) are provided for FPGA interconnect switches and the original SRAM cells for interconnection programmability can be shared for Vdd programmability. Compared to FPGA Class1, the improved FPGA Class3 uses the same Vdd-programmable switches but only increases the SRAM cell number by 28%. This is because FPGA Class3 removes the configurable level converters in routing channels. On average, FPGA Class1 has 118% area overhead, FPGA Class2 has 17% area overhead and FPGA Class3 has 52% area overhead. Both Class2 and Class3 introduce less SRAM and area overhead while reducing more energy compared to Class1. Compared to FPGA Class3, Class2 reduces a comparable amount of energy while it gives much smaller SRAM and area overhead. Therefore, Class2 is the best architecture class considering performance, energy and area.

D. Breakdown of Energy, and Area Overhead

Fig. 13 presents the energy breakdown of architecture ($N = 12$, $k = 4$), the min-energy or close to min-energy architecture for all FPGA architecture classes. The logic energy is consumed by LUTs, flip-flops and MUXes in logic blocks. The local interconnect energy is consumed by internal routing wires and buffers within logic blocks. Routing wires outside logic blocks, programmable interconnect switches in routing channels and their configuration SRAM cells contribute to global interconnect energy. The energy reduction of both FPGA Class2 and Class3 are mainly due to global interconnect leakage energy reduction. By power-gating interconnect switches with an intrinsically low utilization rate for field programmability (about 3% on average for architecture ($N = 12$, $k = 4$) as shown in Table XII), both FPGA Class2 and Class3 can dramatically reduce global interconnect leakage. Class1 fails to do so due to large leakage overhead of Vdd-level converters in routing channels. The figure also shows that global interconnect dynamic energy, 59.24% of total FPGA energy for Class2 and 52.34% for Class3, becomes dominant after applying the programmable-Vdd technique.

Fig. 14 presents the area overhead breakdown of architecture ($N = 12$, $k = 4$) for FPGA architecture Class2 and Class3. The area overhead of routing switches and connection blocks is introduced by power transistors, extra configuration SRAM

TABLE XI

TOTAL NUMBER OF CONFIGURATION SRAM CELLS AND DEVICE AREA OVERHEAD FOR DIFFERENT VDD-PROGRAMMABLE FPGAS. SRAM CELLS INCLUDE THOSE USED IN LUTS AND TOTAL DEVICE AREA INCLUDES BOTH LOGIC BLOCK AND INTERCONNECT AREA. THE DEVICE AREA IS IN MINIMUM WIDTH TRANSISTOR AREA

Dominant Arch. (N,k)	total # of SRAM cells on chip				total device area			
	Class0 baseline	Class1 (% overhead)	Class2 (% overhead)	Class3 (% overhead)	Class0 baseline	Class1 (% overhead)	Class2 (% overhead)	Class3 (% overhead)
(8,7)	649218	88%	2%	17%	11541440	100%	15%	44%
(6,7)	621929	89%	2%	20%	10689783	108%	16%	49%
(6,6)	469504	128%	3%	31%	10114162	125%	19%	57%
(10,5)	374174	164%	3.4%	33%	9793576	126%	17%	55%
(12,4)	317391	190%	4%	40%	9173613	130%	17%	55%
Average	-	132%	3%	28%	-	118%	17%	52%

TABLE XII

INTERCONNECT SWITCH UTILIZATION RATE OF FPGA ARCHITECTURE ($N = 12, k = 4$)

circuit	routing switch			connection switch			interconnect switch		
	total #	used #	utilization	total #	used #	utilization	total #	used #	utilization
alu4	22843	2446	10.71%	99216	2421	2.44%	122059	4867	3.99%
apex2	48721	4835	9.92%	211484	3810	1.80%	260205	8645	3.32%
apex4	28598	3256	11.39%	124267	2293	1.85%	152865	5549	3.63%
bigkey	52464	3735	7.12%	227448	2475	1.09%	279912	6210	2.22%
clma	416040	36919	8.87%	1803360	20438	1.13%	2219400	57357	2.58%
des	92130	5291	5.74%	399360	3004	0.75%	491490	8295	1.69%
diffeq	25300	2953	11.67%	109850	2506	2.28%	135150	5459	4.04%
dsip	76510	3729	4.87%	331695	1995	0.60%	408205	5724	1.40%
elliptic	90888	8892	9.78%	394212	5635	1.43%	485100	14527	2.99%
ex1010	126912	12779	10.07%	550368	8587	1.56%	677280	21366	3.15%
ex5p	30046	3099	10.31%	130559	2089	1.60%	160605	5188	3.23%
frisc	158600	13886	8.76%	687700	8509	1.24%	846300	22395	2.65%
misex3	26722	3100	11.60%	116064	2503	2.16%	142786	5603	3.92%
pdv	181375	17304	9.54%	786500	9989	1.27%	967875	27293	2.82%
s298	40440	3553	8.79%	175500	3804	2.17%	215940	7357	3.41%
s38417	148648	14140	9.51%	644436	10282	1.60%	793084	24422	3.08%
s38584	127432	11910	9.35%	552500	8113	1.47%	679932	20023	2.94%
seq	36938	4278	11.58%	160381	3244	2.02%	197319	7522	3.81%
spla	109282	10329	9.45%	473993	6789	1.43%	583275	17118	2.93%
tseng	17738	2361	13.31%	77077	1847	2.40%	94815	4208	4.44%
Avg.	-	-	9.62%	-	-	1.61%	-	-	3.11%

TABLE XI

TOTAL NUMBER OF CONFIGURATION SRAM CELLS AND DEVICE AREA OVERHEAD FOR DIFFERENT Vdd-PROGRAMMABLE FPGAS. SRAM CELLS INCLUDE THOSE USED IN LUTS AND TOTAL DEVICE AREA INCLUDES BOTH LOGIC BLOCK AND INTERCONNECT AREA. THE DEVICE AREA IS IN MINIMUM WIDTH TRANSISTOR AREA

Dominant Arch. (N,k)	total # of SRAM cells on chip				total device area			
	Class0 baseline	Class1 (% overhead)	Class2 (% overhead)	Class3 (% overhead)	Class0 baseline	Class1 (% overhead)	Class2 (% overhead)	Class3 (% overhead)
(8,7)	649218	88%	2%	17%	11541440	100%	15%	44%
(6,7)	621929	89%	2%	20%	10689783	108%	16%	49%
(6,6)	469504	128%	3%	31%	10114162	125%	19%	57%
(10,5)	374174	164%	3.4%	33%	9793576	126%	17%	55%
(12,4)	317391	190%	4%	40%	9173613	130%	17%	55%
Average	-	132%	3%	28%	-	118%	17%	52%

cells and control logics. The area overhead of logic blocks is introduced by Vdd-level converters at logic block inputs/outputs, power transistors and associated configuration SRAM cells. The area overhead of FPGA Class2 due to routing switches, connection blocks and logic blocks are 3.87%, 11.31%, and 1.95%, respectively. The area overhead of FPGA Class3 due to routing switches, connection blocks and logic blocks are 16.93%, 34.22%, and 3.19%, respectively. The area overhead due to connection blocks is dominant for both FPGA Class2 and Class3.

From another point of view, the area overhead of FPGA Class2 due to power transistors and control logics are 10.22% and 5.75%, respectively. The area overhead due to extra configuration SRAM cells is less than 1% for FPGA Class2. For FPGA Class3, the area overhead due to power transistors, control logics and extra configuration SRAM cells are 19.05%, 25.47%, and 8.02%, respectively. Power transistors introduce the largest area overhead for FPGA Class2 while control logics introduce the largest area overhead for FPGA Class3.

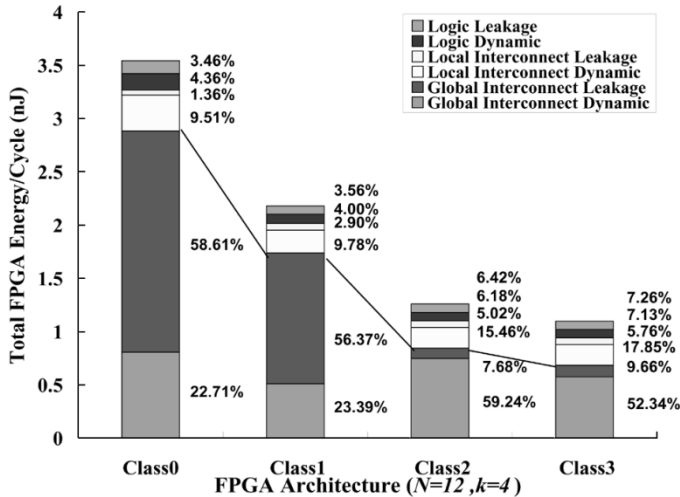


Fig. 13. Energy breakdown of architecture ($N = 12, k = 4$) for all classes.

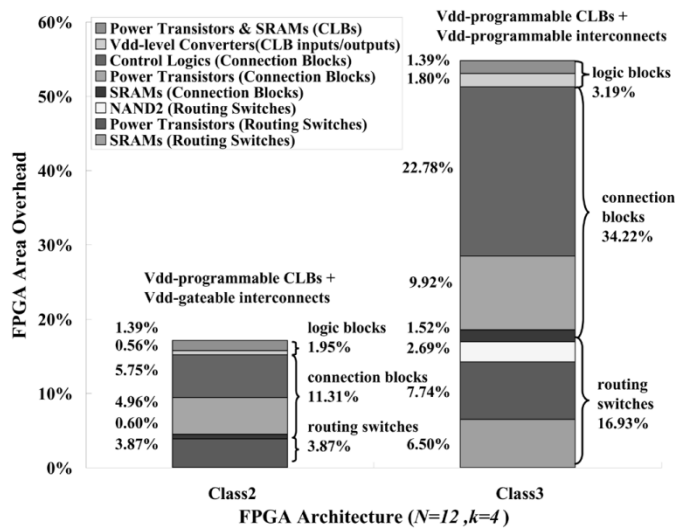


Fig. 14. Area overhead breakdown of architecture ($N = 12, k = 4$) for FPGA architecture Class2 and Class3.

VI. CONCLUSIONS AND DISCUSSIONS

In this paper, we have designed novel Vdd-programmable and Vdd-gateable interconnect switches with minimal number of configuration SRAM cells. Using the new switches, we have proposed three new classes of Vdd-programmable FPGA architectures and conducted FPGA architecture evaluation for these architecture classes. We conclude that the best architecture class in our study uses Vdd-programmable logic blocks and Vdd-gateable interconnects considering area, power and performance tradeoff. It reduces the energy-delay product over the MCNC benchmark set by 54.39% with 17% area increase and 3% more configuration SRAM cells compared to the baseline architecture class using high-Vdd for both logic blocks and interconnects. Our evaluation results also show that within each architecture class, LUT size 4 gives the lowest energy consumption as well as the smallest total device area while LUT size 7 leads to the highest performance.

Increased area due to Vdd programmability makes wire segments longer and wire capacitance per segment larger, which results in larger energy consumption. We do not consider longer wire segments due to Vdd programmability in this paper. As the load capacitance of a routing switch is usually dominated by its fanout routing switches in FPGAs, energy reduction may become slightly smaller for all FPGA classes considering this factor. The FPGA class using Vdd-programmable logic blocks and Vdd-gateable interconnects has the smallest area overhead of 17% within all Vdd-programmable classes, which leads to only 8% longer wire segments (as $1.08 \times 1.08 \approx 1.17$). The impact of longer wire segments due to larger area on this class should be much smaller than those on other classes due to the fact that it has the smallest area overhead. Therefore, we believe that the FPGA class using Vdd-programmable logic blocks and Vdd-gateable interconnects is still the best architecture class considering this factor.

There are a few alternative architecture classes that can be studied in the future. One may apply single-Vdd with power-gating to both logic blocks and interconnects with one configuration SRAM cell and one power transistor for each logic block or interconnect segment. As the area overhead due to Vdd-programmable logic blocks is small, about 3% for architecture ($N = 12, k = 4$), the new architecture class may have a slightly smaller area compared to the FPGA class using Vdd-programmable logic blocks and Vdd-gateable interconnects, and may be effective when the utilization rate of logic blocks is low. Moreover, we assume the smallest FPGA array for each benchmark circuit in this paper, but our evaluation methodology can be easily extended to the new class with the practical utilization rate. Furthermore, one configuration SRAM cell and two power transistors for field Vdd-level selection without power-gating. In this case, VddL can be applied to the unused circuit elements to reduce leakage. We speculate that Vdd-gating is able to reduce more energy than pure Vdd-selection does, but we have not studied this architecture class in this paper.

Our recent work [32] studied the device and FPGA architecture co-optimization for higher power reduction compared to this paper. In addition, exploring the solution space containing power transistor size, supply voltage, threshold voltage, and FPGA architecture virtually removes the area overhead of power transistors.

The state-of-the-art commercial FPGAs have applied uni-directional routing switches in routing architecture and used depopulated local interconnects inside logic blocks [21], [33]. As these interconnect features may have a great impact on power and performance, in the future we will conduct architecture evaluation considering these features with Vdd programmability.

ACKNOWLEDGEMENT

The authors like to thank L. Cheng and Ms. H.-Y. P. Wong from the Electrical Engineering Department, University of California, Los Angeles, for generating circuit models for a variety of basic FPGA circuits and for helpful discussions. The authors used computers donated by Intel.

REFERENCES

- [1] E. Kusse and J. Rabaey, "Low-energy embedded FPGA structures," in *Proc. Int. Symp. Low Power Electronics and Design*, Aug. 1998, pp. 155–160.
- [2] K. Poon, A. Yan, and S. Wilton, "A flexible power model for FPGAs," in *Proc. 12th Int. Conf. on Field-Programmable Logic and Applications*, Sep. 2002.
- [3] F. Li, D. Chen, L. He, and J. Cong, "Architecture evaluation for power-efficient FPGAs," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, Feb. 2003.
- [4] F. Li and L. He, "Power modeling and characteristics of field programmable gate arrays," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, Oct. 2005.
- [5] T. Tuan and B. Lai, "Leakage power analysis of a 90 nm FPGA," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2003.
- [6] J. H. Anderson, F. N. Najm, and T. Tuan, "Active leakage power optimization for FPGAs," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, Feb. 2004.
- [7] J. Lamoureux and S. J. Wilton, "On the interaction between power-aware FPGA CAD algorithms," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2003, pp. 701–708.
- [8] A. Gayasen, Y. Tsai, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, and T. Tuan, "Reducing leakage energy in FPGA's using region-constrained placement," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, Feb. 2004.
- [9] Y. Lin, F. Li, and L. He, "Routing track duplication with fine-grained power-gating for FPGA interconnect power reduction," in *Proc. Asia South Pacific Design Automation Conf.*, Jan. 2005.
- [10] F. Li, Y. Lin, L. He, and J. Cong, "Low-power FPGA using pre-defined dual-V_{dd}/dual-V_t fabrics," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, Feb. 2004.
- [11] F. Li, Y. Lin, and L. He, "FPGA power reduction using configurable dual-v_{dd}," in *Proc. Design Automation Conf.*, Jun. 2004.
- [12] —, "V_{dd} programmability to reduce fpga interconnect power," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2004.
- [13] J. H. Anderson and F. N. Najm, "Low-power programmable routing circuitry for FPGAs," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2004.
- [14] A. Gayasen, K. Lee, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, and T. Tuan, "A dual-v_{dd} low power FPGA architecture," in *Proc. Int. Conf. Field-Programmable Logic and its Application*, Aug. 2004.
- [15] J. Rose, R. Francis, D. Lewis, and P. Chow, "Architecture of field-programmable gate arrays: The effect of logic functionality on area efficiency," *IEEE J. Solid-State Circuits*, 1990.
- [16] S. Singh, J. Rose, P. Chow, and D. Lewis, "The effect of logic block architecture on FPGA performance," *IEEE J. Solid-State Circuits*, 1992.
- [17] E. Ahmed and J. Rose, "The effect of LUT and cluster size on deep-submicron FPGA performance and density," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, Feb. 2000, pp. 3–12.
- [18] Y. Lin, F. Li, and L. He, "Power modeling and architecture evaluation for FPGA with novel circuits for v_{dd} programmability," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, Feb. 2005.
- [19] V. Betz, J. Rose, and A. Marquardt, *Architecture and CAD for Deep-Submicron FPGAs*: Kluwer Academic Publishers, 1999.
- [20] G. G. Lemieux and S. D. Brown, "A detailed router for allocating wire segments in field-programmable gate arrays," in *Proc. ACM Physical Design Workshop*, Apr. 1993.
- [21] D. Lewis *et al.*, "The stratix routing and logic architecture," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, Feb. 2003.
- [22] (2002) Predictive Technology Model. Univ. of Berkeley Device Group. [Online]. Available: <http://www.device.eecs.berkeley.edu/ptm/mosfet.html>
- [23] (2002) International Technology Roadmap for Semiconductor. [Online]. Available: <http://public.itrs.net/>
- [24] S. Yang, "Logic Synthesis and Optimization Benchmarks, Version 3.0," Microelectronics Center of North Carolina (MCNC), 1991.
- [25] "Virtex-II 1.5 v Platform FPGA Complete Data Sheet," Xilinx Corporation, 2002.
- [26] D. E. Lackey *et al.*, "Managing power and performance for system-on-chip designs using voltage islands," in *Proc. Int. Conf. Computer-Aided Design*, 2002, pp. 195–202.
- [27] K. Usami and M. Horowitz, "Clustered voltage scaling techniques for low-power design," in *Proc. Int. Symp. Low Power Electronics and Design*, 1995.
- [28] C. Chen, A. Srivastava, and M. Sarrafzadeh, "On gate level power optimization using dual-supply voltages," *IEEE Trans. on VLSI Systems*, vol. 9, pp. 616–629, Oct. 2001.
- [29] M. Hamada, Y. Ootaguro, and T. Kuroda, "Utilizing surplus timing for power reduction," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2001, pp. 89–92.
- [30] B. Calhoun, F. Honore, and A. Chandrakasan, "Design methodology for fine-grained leakage control in MTCMOS," in *Proc. Int. Symp. Low Power Electronics and Design*, Aug. 2003.
- [31] M. Hamada *et al.*, "A top-down low power design technique using clustered voltage scaling with variable supply-voltage scheme," in *Proc. IEEE Custom Integrated Circuits Conf.*, 1998, pp. 495–498.
- [32] L. Cheng, P. Wong, F. Li, Y. Lin, and L. He, "Device and architecture co-optimization for FPGA power reduction," in *Proc. Design Automation Conf.*, Jun. 2005.
- [33] D. Lewis *et al.*, "The stratix II routing and logic architecture," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, Feb. 2005.



Yan Lin (S'05) received the B.E. degree in automation from Tsinghua University, Beijing, China, in 2002 and the M.S. degree in electrical engineering from University of California, Los Angeles (UCLA) in 2004. He is currently working toward the Ph.D. degree in electrical engineering department at UCLA.

His research interests include computer-aided design of VLSI circuits and systems, programmable logic fabrics and low-power design.



Fei Li received the B.S. and M.S. degrees in electrical engineering from Fudan University, Shanghai, China, in 1997 and 2000, respectively, the M.S. degree in computer engineering from University of Wisconsin, Madison in 2002, and the Ph.D. degree in electrical engineering from University of California, Los Angeles (UCLA) in 2005.

His research interests include computer-aided design of VLSI circuits and systems, programmable device architecture and low-power design.



Lei He (S'94–M'99) received the Ph.D. degree in computer science from University of California, Los Angeles (UCLA), in 1999.

He was a faculty member at electrical and computer engineering, University of Wisconsin, Madison between 1999 and 2001, and joined the Faculty of Electrical Engineering, UCLA, in 2002. His research interests include: 1) modeling and design considering inductance effects and process variations for integrated circuits and packages; 2) programmable logic fabrics and reconfigurable systems; and 3) power-efficient circuits and systems. He has published over 100 technical papers and has served as TPC members of Design Automation Conference, International Symposium on Low Power Electronics and Design, and International Symposium on Field Programmable Gate Array.

Dr. He was granted the NSF CAREER award in 2000, UCLA Chancellor's faculty career development award (highest class) in 2003, IBM faculty partner award in 2003, and Northrop Grumman Excellence in Teaching Award in 2005.