# Using Variable Clocking to Reduce Leakage in Synchronous Circuits

Navid Toosizadeh, Safwat G. Zaky and Jianwen Zhu Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada navid.toosizadeh@utoronto.ca, safwat.zaky@utoronto.ca, jzhu@eecg.utoronto.ca

Abstract—There is a growing demand for high-performance, low-power systems, particularly in portable devices. New approaches to design are needed in technologies with feature sizes of 90 nm and below to reduce leakage power and to deal with process variations, which force designers to use increasingly conservative delay estimations. This paper presents a variable clock generator for a conventionally-designed synchronous circuit core. The clock frequency adjusts automatically to interand intra-chip process, voltage and temperature variations, making it possible to design the circuit assuming typical rather than worst-case conditions. The resulting circuit uses much fewer high-speed, low-voltage-threshold cells, and consequently has significantly reduced leakage power. Post-layout test results on a 32-bit microprocessor implemented in 90-nm technology showed 10X less leakage and 19% less dynamic power when operating under typical conditions, compared to a conventional, fixed-frequency implementation. The system is functional under all PVT corners.

#### I. INTRODUCTION

Power reduction is an important objective in the design of today's high-performance systems, particularly in portable devices. Meeting power requirements is becoming increasingly difficult in today's shrinking technologies because of leakage power. As technology scales to smaller feature sizes, leakage power becomes a more substantial portion of the total power, due to two main factors. First, gate length and threshold voltage of transistors are reduced, resulting in a substantial increase in leakage power [1]. Secondly, process variations make the quality of fabricated chips less predictable and hence, more and more conservative delay and clock frequency estimations are used [2]. This results in an undesired over-engineering to ensure that a large percentage of the fabricated chips meet performance and power consumption requirements. In a conventional synchronous design flow using multi-corner static timing analysis, a digital system is designed to deliver the required performance under all Process-Voltage-Temperature (PVT) corners, including worst-case PVT. To this end, many lowvoltage-threshold (LVT) cells are used, resulting in high leakage power consumption.

This paper demonstrates that variable clocking can be used to reduce leakage power while maintaining performance. According to the proposed approach, a digital system is designed, synthesized, optimized and laid out to meet the required clock frequency under typical rather than worstcase PVT conditions. The clock generator adjusts the clock frequency automatically as PVT conditions change, either inter-chip or intra-chip. Since the system adjusts its clock period to PVT conditions, it is referred to as a *PVT-aware*  design. The asynchronous clock generator circuit is based on that previously proposed for VariPipe [3].

A key feature of the proposed approach is that it fits well in a low-power standard-cell ASIC design flow. The synchronous nature of the system is preserved, and it is designed using standard synchronous design techniques. At the same time, the introduction of a clock generator that adjusts its clock frequency to the prevailing PVT conditions brings many of the advantages of asynchronous operation to the synchronous environment. Because the system is designed under typical rather than worst-case PVT conditions, the resulting circuits are smaller and use fewer LVT cells, and thus have significantly reduced leakage. In a case study of a 32-bit DLX microprocessor, leakage power was reduced by a factor of 10 under typical PVT conditions and a factor of 7 under worst-case PVT conditions. Other advantages include a reduction in dynamic power, resilience to PVT variations, and suitability for voltage scaling.

Section II uses an example to demonstrate the potential reduction in leakage power when a system is designed under typical rather than worst-case PVT conditions. The large reduction that is possible provides the main motivation for the proposed design approach. Subsequent sections describe the PVT-aware architecture and the methodology to integrate the proposed architecture in conventional digital design flows. As a case study, the proposed design methodology is demonstrated using a free-license DLX microprocessor, and post-layout results in 90-nm technology are presented. Comparison to related work is presented in Section VIII.

#### II. MOTIVATION

An IC foundry characterizes its technology under different PVT conditions, known as PVT corners. The PVT corners for the 90-nm technology used in this paper are given in Table I for a 1.0 V supply voltage. Unless otherwise stated, the conditions in Table I are those referenced throughout the paper.

TABLE I								
	PVT CORNERS							
PVT corner Process Voltage Temperature								
Best	Fast	1.1	-40°C					
Typical	Typical	1.0	25°C					
Worst	Slow	0.9	125°C					

Multi-threshold libraries in today's technologies feature different implementations for functions, including highvoltage-threshold (HVT), standard-voltage-threshold (SVT) and low-voltage-threshold (LVT) cells, which have different speed and leakage characteristics. Of them, the LVT cells are the fastest and have the highest leakage. They are used by the synthesis and optimization tools in critical paths. The SVT and HVT gates are used in less-critical paths to reduce leakage power.

For worst-case design, the best PVT corner is used for hold time check and the worst PVT corner is used for setup time check. Critical paths are identified under the worst PVT corner. Then, the synthesis and physical design tools are instructed to optimize the design to meet the required performance under the worst-case conditions. Achieving high performance under the worst PVT conditions often requires a large number of high-speed and high-leakage LVT cells.

The test vehicle used in this paper is a Hennessey and Patterson's 32-bit DLX pipeline microprocessor [4] down-loaded from opencores.org [5]. To illustrate the potential for design optimization, the DLX processor was synthesized by Synopsys Design Compiler for a clock frequency of 1 GHz under the worst PVT conditions (Design 1) and also under typical PVT conditions (Design 2). Both designs were constrained for the best area and power optimizations. The design methodology used will be described in Sections IV and V.

The resulting number of cells used in the two designs from each of the three categories of low, standard and highthreshold cells is shown in Table II. The optimization tool had to use a mix of all three cell types in Design 1 to meet the performance constraint. For Design 2, it was able to achieve the desired performance using HVT cells only. As a result, the leakage power of Design 1 is substantially larger than that of Design 2. Also, its dynamic power is higher as it is a larger circuit with more switching capacitances.

It should be noted that the power values in the table were obtained from an initial power analysis to evaluate the two designs. Power analysis based on post-layout simulations will be presented in Section V. The power values given are for the typical PVT corner, which uses a temperature of 25°C. As temperature increases, leakage power increases exponentially and becomes a more substantial portion of the total power.

The lower power consumption of Design 2 demonstrates that there is a substantial room for improvement when a system can be designed to provide the desired performance level under typical PVT conditions, and is equipped with a mechanism that adjusts its run-time speed to accommodate changes in PVT conditions. The architecture to support such a PVT-aware mechanism and the corresponding design flow are presented next.

## III. THE PVT-AWARE ARCHITECTURE

The proposed PVT-aware system uses an on-chip clock generation circuit that adjusts its frequency with PVT conditions. The chip area is divided into multiple regions as shown in Fig. 1, and a PVT-aware completion detection circuit is included in each region. The number of regions depends on many parameters such as the size of the design, the quality of fabrication and the voltage and temperature profiles over the chip.

Each region contains a completion detection circuit. When it receives a pulse from the central clock generator, the circuit sends back a completion signal after an appropriate



Fig. 1. Clock generation circuit,  $CD \equiv Completion Detector, CPG \equiv Clock$ Pulse Generator



Fig. 2. Clock generation circuit schematic

delay. The central clock generator generates the next clock pulse after it has received a response from all regions. This structure is based on VariPipe [3] and Dean's dynamic clocking approach [6].

Fig. 2 shows a more detailed schematic of the clock generation circuit, including the clock detection circuit for one region. The completion circuit consists of a delay element and a toggle. The delay element is selected such that the delay around the clock generation loop meets the requirements of the critical path of the system under typical PVT conditions. During operation, the delay around the loop will depend on the local PVT conditions. Thus, a clock pulse received from the clock generator will emerge from the delay element after a delay that is dependent on the prevailing PVT conditions in the region. The clock pulse is converted to a level by the toggle before it is sent back to the C-element of the clock pulse generator. When all completion signals have been received, the clock pulse generator generates a new clock pulse. Thus, the region introducing the longest delay determines the clock period.

Initially, all toggle elements are reset and so is the output of the C-element. After the reset signal is removed, all toggle elements change state, causing the C-element to toggle its output, thus creating a clock pulse of width CPW at the output of the XOR gate.

In essence, the clock signal generator is a ring oscillator in which the ring delay is equal to the longest of the delays in all regions. Furthermore, the ring delay changes with the prevailing PVT conditions in each region. Thus, the clock frequency changes in response to both intra- and inter-chip variations in PVT.

#### TABLE II

Post-synthesis power breakdown and area of the designs under typical PVT conditions, temperature=25°C

Design	Average leakage power (mW) and <i>Number of cells</i> Average dynamic					Area $(\mu m^2)$
U	HVT	SVT	LVT	Total	power (mW)	
Design 1	0.006	0.026	1.347	1.379	247	115 015
Design I	3,641 cells	2,377 cells	7,133 cells	13,151 cells	34.7	115,215
Design 2	0.023	0	0	0.023	24.0	100.020
Design 2	(12,711 cells)	0	0	(12,711 cells)	54.0	109,930



Fig. 3. Proposed low-power PVT-aware design flow, HDL  $\equiv$  Hardware Description Language, DRC  $\equiv$  Design Rule Check, CTS  $\equiv$  Clock Tree Synthesis, SI  $\equiv$  Signal Integrity, STA  $\equiv$  Static Timing Analysis, ECO  $\equiv$  Engineering Change Order

## IV. DESIGN FLOW

Fig. 3 shows the design flow to integrate the suggested PVT-aware architecture into a conventional standard-cell ASIC design flow [7]. A few extra steps are needed to add the clock generation circuit to the top-level hardware description language (HDL), place the clock generation elements appropriately and tune the delay elements.

The most important deviation from a conventional design flow is that the main core is synthesized and laid out to meet the required clock period under typical PVT conditions. Thus, typical-case rather than worst-case timing libraries are used for setup check and critical path analysis during synthesis and layout.

#### A. Clock generation issues

• The first step to implement the suggested clock generation circuit of Fig. 2 is to create a library of delay elements in the target technology. A delay element can be implemented as a chain of 2n inverters, where n =1,2,...,N. Then, the delay of each delay element is estimated using a static timing analysis (STA) tool. The result will be a table of several delay elements and their corresponding delay values, which can be used in the clock generation circuit.

- Multiple completion detection circuits are implemented to match the delay of the critical path. All of the completion detection circuits are designed to match the same critical path delay. However, they are placed in different regions of the chip as shown in Fig. 1, to be subject to the regional operating PVT conditions and reduce the impact of random dopant fluctuation (RDF).
- The delay elements are adjusted such that the delay of the loop composed of the completion detection circuit and the clock pulse generator is equal to the critical path of the system. The delay of the loop must be tested under different PVT corners to ensure that it matches the critical path under all conditions. When adjusting the delay elements, appropriate margins should be used, because different factors such as crosstalk, inductance, IR drops, noise, etc. may affect the completion detection

circuits and the datapath elements differently.

- During synthesis, it is sufficient to insert delay elements that are approximately 25% longer than the desired clock period. They are trimmed later, during the layout flow.
- The submodules of the clock generation circuit should be pre-placed during floorplanning to avoid a random placement.
- After the layout is completed, the post-layout netlist, the standard delay format (SDF) file and the standard parasitic exchange format (SPEF) file are exported to an STA tool to test the total delay including gate and interconnect delays.
- The clock loop is examined with each completion detection circuit using an STA tool to ensure that the resulting clock period is appropriate.
- Different paths could have different sensitivity to PVT changes and thus, under different PVT conditions, the critical path may change. Using STA analysis, it is not necessary to know what the critical path is. Only the value of the maximum delay under each PVT corner is needed. The critical path producing this delay may be different under different conditions. The delay loops are designed to be always longer than the maximum delay under all available PVT corners.
- If the resulting clock period reported by the STA tool is not correct, the delays inside the completion detection circuits are tuned and a pass of engineering change order (ECO) is performed to fix the layout.
- The clock pulse width determined by CPW in Fig. 2 is tested under all PVT conditions to ensure that the pulse width requirements of sequential elements are not violated.
- The reset signal to the system must be long enough to ensure that all the delay elements are successfully reset and the gates and flip-flops become stable.

V. CASE STUDY: PVT-AWARE DLX MICROPROCESSOR

The DLX processor introduced in Section II was used as a case study. It was implemented both as a PVT-aware system and as a conventional synchronous circuit with a fixed-frequency clock. The PVT-aware design flow of Fig. 3 was implemented in 90-nm technology using the toolset shown in Table III. A low-power design flow similar to that of Fig. 3, using the same tools and optimizations but without the PVT-aware implementation steps, was realized for the conventional synchronous design.

The shortest possible post-layout clock period of the DLX core was found to be 1.244 ns under the worst PVT corner. Hence, the design flow of Fig. 3 was used to implement a PVT-aware DLX processor with the same clock period of 1.244 ns but under typical PVT conditions. The chip was divided into four regions similar to Fig 1, and a completion detection circuit was placed in each quadrant.

# A. Tuning delays

To simplify delay tuning, a library of delay elements in the target technology was implemented as explained in

TABLE III

	TOOLSET	
Task	Tool	Version
Synthesis	Design Compiler	Y-2006.06-SP5
Timing and power analysis	PrimeTime-PX	Y-2006.06-SP3-1
Physical design	SoC Encounter	5.2
Simulation	ModelSim	6.3c

Section IV-A. Delays in the clock generation circuit were chosen to be 25% larger than needed. They were tuned after the place and route step using an engineering change order (ECO) flow. In practice, a margin of 10-15% is used. A 10% margin was used here.

To tune the delays after the place and route, each clock generation loop in Fig. 2 was analyzed by PrimeTime to find the resulting clock period. If the period was longer than needed, the delay element was replaced by a smaller delay element from the delay library, and vice versa. This process was repeated for each completion detection circuit. After tuning the delays, the new netlist was fed back to Encounter to update the layout (ECO).

### B. Implementing the fixed-clock counterpart

The objective in the implementation of the fixed-clock counterpart was to realize a conventional fixed-clock synchronous system that has the same speed of the PVT-aware processor most of the time. The PVT-aware processor has an effective clock period of 1.468 ns under typical conditions. Considering the 10% margin used in the design, the fixed-clock DLX core was constrained to a clock period of 1.34 ns. This clock period was to be met under all three PVT corners, because the clock is fixed. All the optimizations of Fig. 3 were applied to the fixed-clock design.

## VI. EVALUATION

In this section, the PVT-aware and the fixed-clock microprocessors are compared in terms of power consumption and performance. Using several application programs with different power consumptions, it is demonstrated that the leakage of the PVT-aware processor is 10X less under typical conditions and 7X less under worst-case conditions. Other properties of the PVT-aware microprocessor such as resilience to PVT variations and suitability for voltage scaling are also examined. The functionality and power consumption of the PVT-aware DLX processor and its fixedclock counterpart were analyzed using the three benchmark suites given in Table IV. The benchmarks were compiled by DLX GCC [8]. Post-layout simulations of the circuits were performed for each benchmark to record switching activities in the switching activity interchange format (SAIF). These, together with parasitic data (SPEF files), were used by PrimeTime-PX for simulation-based power analysis.

#### A. Power and performance analysis

**Typical PVT:** Average power consumption and execution times under typical PVT conditions are presented in Table V. A typical PVT corner at a temperature of 25°C was used, because this was the only corner provided by the technology supplier to represent typical conditions. Power values in the table do not include memory and IO. The core power is the



BENCHMARKS					
Source	Benchmark				
MiBench [9]	adpcm_coder adpcm_decoder crc32 dijkstra qsort				
PowerStone [10]	bcnt blit compress ucbqsort				
Applications from [11], [12]	Bubble Sort JPEG-DCT MP3-DCT32 MPEG2-Bdist				

power of the processor excluding the clock tree. In the case of the PVT-aware design, the power of the clock generation architecture is partly in the core and partly in the clock tree. The clock generation circuit including the delay lines and the clock pulse generator consumed only 3% of the total power and occupied 0.5% of the total area. Its area is mainly taken by the delay elements.

The PVT-aware processor executes all benchmark programs with the same execution time of the fixed-clock processor under typical conditions. However, it consumes significantly less leakage and dynamic power. To calculate the average power values (highlighted in the table), the energy consumption of every program was calculated and summed, then divided by the sum of the execution times. On average, the leakage power of the PVT-aware processor is 10X less than that of its fixed-clock counterpart and its dynamic power is 19% less for the same performance. The total power of the PVT-aware processor is 21% smaller on average. Since the clock tree power is almost equal for the two designs, only the core power contributes to the power differences.

Table VI shows how the two designs differ in some of the key parameters that affect power consumption. The PVT-aware design is implemented mostly of HVT cells. The fixed-clock system uses a large number of LVT and SVT cells, which have significantly larger leakage power. Also, the fixed-clock system is a bigger circuit with about 14% more area and more switching capacitances. As a result, it consumes more dynamic power than its PVT-aware counterpart.

**Worst-case PVT:** Power requirements, specially leakage power, were next examined under the worst-case PVT corner, at a temperature of 125°C (this was the only available PVT corner with a high temperature). Average power consumption of all benchmarks and their total execution time for this case are presented in Table VII. Leakage power is a substantial portion of the total power for the two designs under these conditions. Leakage power for the PVT-aware processor is 7X lower than for the fixed-frequency processor. The speed of the PVT-aware processor is reduced under worse-than-typical conditions, and thus, its dynamic power is also reduced. The power-delay product of the PVT-aware system is 2.30X smaller than that of its fixed-clock counterpart.



Fig. 4. Performance of PVT-aware and fixed-clock DLX processors under all PVT corners

#### B. Resilience to inter-chip PVT variations

The PVT-aware processor was tested under the three PVT corners. The system executes all benchmarks correctly and tunes itself to suit the prevailing PVT conditions. Fig. 4 shows that the execution time changes as the clock period changes under different PVT conditions.

## C. Resilience to intra-chip PVT variations

The PVT-aware processor was also tested for its resilience to intra-chip PVT variations. An area in the right side of the chip equal to about 1/3 of the total area was selected. Starting with the typical-case SDF file, the delays of all the cells in that area were augmented by 10% using the Design Compiler's derating commands, and a new SDF file was generated. Simulations verified that the system executes all the benchmarks correctly. Table VIII shows the change in clock periods with the increased delay.

### D. Suitability for voltage scaling

The 90-nm technology used in this paper is characterized for two supply voltage levels: 1.0 V and 1.2 V. These characterizations were used to apply voltage scaling to the PVT-aware DLX processor. It was ensured that the pads were compatible with 1.2 V and no hold violation occurred. Then, the PVT-aware processor was tested under typical PVT conditions for both supply voltages. The system automatically adjusted its frequency to changes in supply voltage. The clock frequency of the system with 1.2 V was 1.17 times higher than that for a voltage supply of 1.0 V. This shows that the PVT-aware design is amenable to voltage scaling techniques.

## VII. DISCUSSION

# A. Design space

The proposed PVT-aware approach expands the design space by providing more flexibility in power and performance trade-offs. This can be useful in the implementation of many applications, such as portable systems. The case study presented in this paper shows that a system can be designed to deliver the desired performance under *typical conditions*, which are the conditions that the system is exposed to most of the time. If conditions get worse, the system is still functional as it automatically slows down, and if conditions get better, it will speed up. The advantage of such a system over the fixed-clock design is an overall reduction in power and area, which is a result of using a smaller number of high-leakage high-speed cells to achieve the desired performance.

## TABLE V

POWER AND PERFORMANCE RESULTS FO	R FIXED-CLOCK A	ND PVT-AWARE	DLX PROCESSORS UNDE	R TYPICAL PVT
	ONDITIONS TEM	$PERATURE = 25^{\circ}C$	ч	

Fixed-clock Processor								
D	Average Dynamic power (mW)			Averag	e leakage pow	er (mW)	Total average	Execution time
Program	Core	Clock tree	Total	Core	Clock tree	Total	power (mW)	(µs)
adpcm_coder	17.012	18.794	35.806	1.237	0.006	1.243	37.049	286.135
adpcm_decoder	18.713	18.794	37.507	1.236	0.006	1.242	38.749	1274.348
crc32	17.113	18.794	35.907	1.236	0.006	1.242	37.149	286.851
dijkstra	16.812	18.794	35.606	1.236	0.006	1.242	36.848	143.806
qsort	19.713	18.794	38.507	1.236	0.006	1.242	39.749	729.099
bent	17.812	18.794	36.606	1.236	0.006	1.242	37.848	44.204
blit	17.511	18.794	36.305	1.237	0.006	1.243	37.548	103.950
compress	16.513	18.794	35.307	1.236	0.006	1.242	36.549	1508.450
ucbqsort	19.414	18.794	38.208	1.235	0.006	1.241	39.449	1399.062
Bubble Sort	24.012	18.794	42.806	1.235	0.006	1.241	44.047	12.963
JPEG-DCT	19.813	18.794	38.607	1.236	0.006	1.242	39.849	588.808
MP3-DCT32	20.315	18.894	39.209	1.233	0.006	1.239	40.448	73.323
MPEG2-Bdist	18.613	18.794	37.407	1.235	0.006	1.241	38.648	50.578
Average	18.380	18.795	37.175	1.236	0.006	1.242	38.402	

PVT-aware Processor								
Drogram	Average	Dynamic pow	ver (mW)	Averag	e leakage pow	er (mW)	Total average	Execution time
Flogram	Core	Clock tree	Total	Core	Clock tree	Total	power (mW)	(µs)
adpcm_coder	10.15	18.996	29.146	0.125	0.004	0.129	29.275	286.136
adpcm_decoder	11.25	18.996	30.246	0.125	0.004	0.129	30.375	1274.349
crc32	10.05	18.996	29.046	0.125	0.004	0.129	29.175	286.852
dijkstra	9.75	18.996	28.746	0.125	0.004	0.129	28.875	143.807
qsort	12.45	18.896	31.346	0.125	0.004	0.129	31.475	729.100
bcnt	10.85	18.896	29.746	0.125	0.004	0.129	29.875	44.205
blit	10.35	18.996	29.346	0.125	0.004	0.129	29.475	103.951
compress	9.65	18.996	28.646	0.125	0.004	0.129	28.775	1508.451
ucbqsort	12.15	18.896	31.046	0.125	0.004	0.129	31.175	1399.063
Bubble Sort	15.148	18.796	33.944	0.125	0.004	0.129	34.073	12.964
JPEG-DCT	11.95	18.996	30.946	0.125	0.004	0.129	31.075	588.809
MP3-DCT32	11.85	18.996	30.846	0.125	0.004	0.129	30.975	73.324
MPEG2-Bdist	11.15	18.896	30.046	0.125	0.004	0.129	30.175	50.579
Average	11.133	18.961	30.094	0.125	0.004	0.129	30.223	

TABLE VI

POST-LAYOUT AREA AND LEAKAGE BREAKDOWN UNDER TYPICAL PVT CORNER, TEMPERATURE=25°C

Design	Leakag	Total area of			
C	HVT	SVT	LVT	Total	std. cells ( $\mu$ m <sup>2</sup> )
Elevel al value	5.45	85.96	1150.00	1241.41	151 572
F1Xed-clock	3,821 cells	4,715 cells	5,652 cells	14,188 cells	151,575
DVT outons	14.45	27.12	87.18	128.75	120 629
P v 1-aware	10,680 cells	1,901 cells	662 cells	13,243 cells	129,038

#### TABLE VII

POWER AND PERFORMANCE RESULTS UNDER WORST-CASE PVT, TEMPERATURE=125°C

Design	Av. Leakage	Av. Total	Total Ex.	Power-delay
8	power (mW)	power (mW)	Time (ms)	product (µJ)
Fixed-clock	59.193	87.216	6.502	567.08
PVT-aware	8.092	22.344	11.023	246.30

TABLE VIII

CLOCK PERIOD CHANGES WITH INTRA-CHIP VARIATIONS

PVT	Successive clock periods	Critical path
Chip under typ.	1.367ns , 1.569ns	1.244ns
Chip with augmented delay	1.473ns , 1.600ns	1.378ns

Fig. 5 shows the design space. Using the conventional design approach, the designer of a system may trade off performance for power. Choosing a high clock frequency results in high dynamic power dissipation. To reach higher frequencies, a larger number of high-speed high-leakage LVT

cells are required, which increase leakage power consumption. Similarly, the required area is increased as the target clock frequency increases.

The proposed PVT-aware design approach adds a new curve to the design space, which may be used to achieve the required speed with lower area and power consumption. The distance between the two curves in Fig. 5 increases with the clock frequency. At lower frequencies, both conventional and PVT-aware designs achieve the required speed using a small number of LVT cells. However, as the target clock frequency increases, the number of LVT cells, and hence the total



Fig. 5. Design Space expansion using PVT-aware design

leakage power, rises more rapidly for the conventional design than for the PVT-aware design. Similarly, the difference between the area and dynamic power of the conventional design and those of the PVT-aware design becomes larger as the target clock frequency increases.

Alternatively, a PVT-aware system may be designed to deliver the performance of its fixed-clock counterpart under *worst-case conditions*. When operating under typical conditions, such a system delivers a performance better than that of the fixed-clock counterpart system at the expense of higher power consumption.

Suitability of PVT-aware systems for voltage scaling adds another degree of freedom to the trade-offs available to the designer. PVT-aware systems automatically adjust their frequency to the input voltage. Hence, the input voltage can be reduced using dynamic voltage scaling techniques to conserve power when top performance is not required. At other times, the input voltage may be increased to boost performance when the system is exposed to poor PVT conditions.

In this paper the typical PVT corner was used as the typical operating conditions. Different designs have different typical operating conditions, specifically for voltage and temperature, which should be considered using this approach. Also, analyzing the design under more PVT corners increases the chances of generating a circuit that is functional under a wide range of operating conditions.

### B. Communication with the environment

The PVT-aware architecture results in a variable clock period. Hence, special attention should be paid to how it communicates with its environment to ensure correct data transfers. The problem of transferring data between unsynchronized clock domains already exists in many highspeed systems. Many approaches have been developed to minimize metastability and data loss when different clock domains are connected. They include multi-flop synchronizers, multiplexer recirculation techniques, use of first-in-firstout buffers between different clock domains and handshake techniques [13], [14]. Similar synchronization techniques may be applied for inter-chip and intra-chip data transfers between a PVT-aware system and its environment.

## C. Clock limitations

The clock frequency of the proposed clock generation circuit is limited by the loop delays. In turn, these delays are a function of the size of the regions to which the chip is subdivided. Hence, it may not always be possible for a PVT-aware system to achieve the same clock frequency of a fixed-clock synchronous design, especially for deep pipelines with shallow logic in each stage.

It is possible to apply the PVT-aware clocking approach in a modular way to complex and large systems on chip (SoCs). SoCs are growing in complexity and hence, multiple clock domains are inevitable. Complex SoCs are composed of several modules with different clock domains, which communicate using asynchronous interfaces [15]. In the case of the PVT-aware design, separate clock generation circuits can be used in different modules and similar approaches can be used to connect the modules from different clock domains. Therefore, a large SoC can benefit from the advantages of the proposed PVT-aware design.

#### D. Clock error detection

The clock generator is not self-starting. It is started by the reset signal. Since each clock edge depends on the previous one, a lost clock edge would cause the clock to stop. This is unlikely because the C-element in the clock generation circuit masks transient pulses and glitches of up to a certain width. However, wider glitches may mistakenly be interpreted as a completion signal. Hence, it may be advisable to include an error detection circuit.

## VIII. RELATED WORK

Other studies have been published to design variablespeed pipelines and use typical PVT conditions. The main distinction of the approach presented in this paper is its objective: *reducing leakage power*, whereas previous work mainly focused on increasing speed when the system is subject to favorable PVT conditions. Other differences are discussed below.

The clock generator circuit used in this paper is a simplified version of that introduced in the VariPipe system [3]. It adjusts its speed to the prevailing PVT conditions but not to the current operations in the pipeline. As a result, the design flow is simplified and more suitable for automation.

Dean's STRiP processor [6] opened up the area of dynamic clocking based on PVT conditions. Dean uses tracking cells to match the delay of different functional units. To do so, several functions are selected according to the frequency of their use and partially replicated at the transistor-level. A speed improvement of 2X has been achieved in a 32-bit microprocessor under typical PVT conditions.

The main distinguishing feature of the approach presented in this paper is its simplicity and suitability for being easily incorporated into standard cell design flows. The design of tracking cells in Dean's work is complicated because they should accurately replicate the critical path of the corresponding function. Also, they are implemented at the transistor level, using passive transistors to imitate loads on the transistors of the functional units. Replicas of circuit paths may be large and power consuming compared to the matched delays used in this paper. However, they more accurately model the delay of corresponding paths. The approach presented in this paper can be easily incorporated in standard-cell ASIC design implementation and does not require customized transistor-level design. It employs *static timing analysis* to tune the delays of completion detection circuits. Therefore, the design of completion detection circuits is independent of the function being matched.

The Razor project [16] aims to reduce the voltage margins used in worst-case analysis of synchronous circuits. Power consumption is reduced by lowering the input voltage, but the clock frequency remains unchanged. The design does not guarantee error-free operation. Hence, an error recovery circuit is added to cope with timing errors that may result from the reduced voltage. Errors are detected using the Razor flip-flops, which have a shadow latch operating with a delayed clock to obtain the correct result. An error signal is generated when the contents of the shadow latches are different from those of the main flip-flops. With the error signal, the data in the main flip-flop are replaced with the data in the latch and the pipeline is flushed. Simulation results show a 64% power saving with less than 3% performance penalty in a simplified 64-bit Alpha processor.

Compared to the PVT-aware approach presented in this paper, Razor has the advantage of keeping the clock frequency fixed, which is useful in many applications. However, implementing Razor requires both architectural and circuit changes. The Razor flip-flops and the pipeline flushing mechanism increase the area and complexity of the system. Flushing a pipeline with many stages may result in a significant performance loss. In addition, Razor-based design methodology inherits the disadvantage of traditional synchronous design where circuit is optimized for worstcase conditions. Although such over-design for typical conditions can be partially overcome by lowering the supply voltage, area overhead is still incurred, which in turn leads to increased leakage power. In contrast, the proposed PVTaware mechanism targets typical-case conditions during the synthesis and physical design phases of the implementation. This leads to smaller area and leakage. Meanwhile, clock frequency requirements are met under typical conditions. Under worse conditions, the clock frequency drops, but if required, it can be increased by raising the supply voltage.

The desynchronization method [17], [18] shares some of the objectives presented in this paper. The main difference is that the circuit resulting from desynchronization is asynchronous. The desynchronization method introduces an area overhead of 13.5% in a DLX microprocessor [18]. By comparison, area overhead using the presented PVT-aware architecture is only 0.5% for a similar DLX processor.

## IX. CONCLUSION

This paper proposes a design methodology for synchronous circuits using a clock that adjusts its frequency based on prevailing PVT conditions. The proposed structure enables the synchronous core to be designed and implemented assuming typical rather than worst-case conditions. The resulting circuit achieves the desired performance under typical conditions, but with much reduced leakage power and somewhat reduced dynamic power. The clock frequency drops under worse-than-typical conditions, guaranteeing correct operation. The paper presents a complete design solution for PVT-aware systems, including a system architecture and a design flow using standard-cell ASIC implementation. The suggested methodology adds a new curve to the digital design space that is suitable for many applications with low power and high performance requirements, such as portable devices.

A case study of a DLX microprocessor implemented in 90-nm technology has demonstrated that a PVT-aware system can deliver the same performance as its fixed-clock counterpart under typical PVT conditions, with 10X less leakage and 19% less dynamic power. The clock frequency changes automatically to suit the prevailing PVT conditions. The paper also shows that voltage scaling techniques can be applied to PVT-aware systems, which automatically adjust their speed to the input voltage.

## X. ACKNOWLEDGMENTS

The authors gratefully acknowledge the financial support of Natural Sciences and Engineering Research Council of Canada and of the Government of Ontario.

#### REFERENCES

- W. Kuzmicz, E. Piwowarska, A. Pfitzner, and D. Kasprowicz, "Static power consumption in nano-cmos circuits: Physics and modelling," in *Proc. of the 14th International Conference Mixed Design of Integrated Circuits and Systems*, June 2007, pp. 163–168.
- [2] S. Borkar, T. Karnik, et al., "Parameter variations and impact on circuits and microarchitecture," in Proc. of the ACM/IEEE Design Automation Conf., June 2003, pp. 338–342.
- [3] N. Toosizadeh, S. Zaky, and J. Zhu, "VariPipe: Low-overhead variableclock synchronous pipelines," in *Proc. of IEEE International Conference on Computer Design (ICCD)*, Oct. 2009, pp. 117–124.
- [4] J. Hennessey and D. Patterson, *Computer Architecture: A Quantative Approach.* Morgan Kaufmann, 2007.
- [5] "ASPIDA," in http://www.opencores.org/projects.cgi/web/aspida.
- [6] M. Dean, "STRiP: A self-timed RISC processor," Ph.D. dissertation, Stanford University, 1992.
- [7] "ASIC design flow," in http://www.faradaytech.com/html/products/asic/Design\_Flow.html.
- [8] "DLX GCC," in http://www2.ucsc.edu/courses/cmps111-elm/dlx.
- [9] "MiBench," in http://www.eecs.umich.edu/mibench.
- [10] L. Lee, B. Moyer, and J. Arends, "Instruction fetch energy reduction using loop caches for embedded applications with small tight loops," in *Proc. of ISLPED*, 1999, pp. 267–269.
- [11] B. Gorjiara and D. Gajski, "Automatic architecture refinement techniques for customizing processing elements," in *Proc. of the 45th* ACM/IEEE Design Automation Conf., June 2008, pp. 379–384.
- [12] "NISC technology website," in *http://www.cecs.uci.edu/~nisc*.
- [13] "Clock crossing: domain Closing the loop on clock domain functional implementation problems," in http://w2.cadence.com/whitepapers/cdc\_wp.pdf. Cadence Design Systems, Inc., 2004.
- [14] A. Lines, "Asynchronous interconnect for synchronous SoC design," *IEEE Micro*, vol. 24, no. 1, pp. 32–41, Feb. 2004.
- [15] S. Sirowy, W. Yonghui, S. Lonardi, and F. Vahid, "Clock-frequency assignment for multiple clock domain systems-on-a-chip," in *Proc. of the Design, Automation & Test in Europe Conference & Exhibition*, Apr. 2007, pp. 1–6.
- [16] T. Austin *et al.*, "Making typical silicon matter with Razor," *Computer*, vol. 37, no. 3, pp. 57–65, Mar. 2004.
- [17] J. Cortadella *et al.*, "Desynchronization: Synthesis of asynchronous circuits from synchronous specifications," *IEEE TCAD*, vol. 25, no. 10, pp. 1904–1921, Oct. 2006.
- [18] N. Andrikos, L. Lavango, D. Pandini, and C. Sotiriou, "A fullyautomated desynchronization flow for synchronous circuits," in *Design Automation Conf.*, June 2007, pp. 982–985.