

# Optimization of Overdrive Signoff

Tuck-Boon Chan<sup>‡</sup>, Andrew B. Kahng<sup>†‡</sup>, Jijia Li<sup>‡</sup> and Siddhartha Nath<sup>†</sup>

<sup>†</sup>CSE and <sup>‡</sup>ECE Departments, UC San Diego, La Jolla, CA 92093  
{tbchan, abk, jil150, sinath}@ucsd.edu

**Abstract**—In modern SOC implementations, multi-mode design is commonly used to achieve better circuit performance and power across voltage-scaling, “turbo” and other operating modes. Although there are many tools for multi-mode circuit implementation, to our knowledge there is no available systematic analysis or methodology for the selection of associated signoff modes. We observe that the selection of signoff modes has significant impact on circuit area, power and performance. For example, incorrect choice of signoff voltages for required overdrive frequencies can result in a netlist with 15% suboptimality in power or 21% in area. In this paper, we propose a concept of *mode dominance* which can be used as a guideline for signoff mode selection. Further, we also propose efficient circuit implementation flows to optimize the selection of signoff modes within several distinct use cases. Our results show that our proposed methodology provides 5-7% improvement in performance compared to the traditional “signoff and scale” method. The signoff modes determined by our methods result in only 0.6% overhead in performance and 8% overhead in power after implementation, compared to the optimal signoff modes.

## I. INTRODUCTION

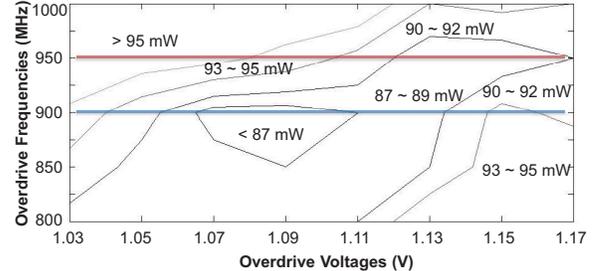
In the era of heterogeneous multi-core SOCs, the performance of single-threaded operations limits the overall speedup of applications. Designers use frequency overdrive at elevated voltages to obtain better performance in consumer electronic devices. An *operating mode* (for simplicity, *mode*) is defined by an (operating frequency, voltage) pair. Devices typically operate at two or three modes, e.g., supply-voltage-scaled (SVS), nominal and turbo (overdrive). The nominal mode corresponds to a low operating voltage and a low frequency whereas the overdrive mode corresponds to a high operating voltage and a high frequency. Due to limited energy budget, laptops and handheld devices operate at nominal or SVS mode for most of their lifetimes. When high performance is needed to boost CPU-intensive tasks, overdrive mode is turned on for a brief period of operation. The average power consumption ( $P_{avg}$ ) for a circuit with both nominal and overdrive modes is

$$P_{avg} = r \times P_{OD} + (1 - r) \times P_{nom} \quad (1)$$

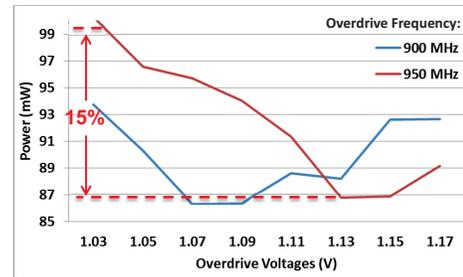
where  $r$  is the total overdrive time normalized to total time when the circuit is turned on ( $0 < r < 1$ ).  $P_{OD}$  and  $P_{nom}$  are the circuit power at overdrive mode and nominal mode, respectively.

We define the *design space* for signoff as the set of combinations of feasible signoff modes. A *point* in such design space indicates  $m$  (frequency, voltage) pairs for  $m$ -mode signoff, where  $m \geq 1$ . Signing off at different points in a design space results in circuits with different area, power and performance. For example, Figure 1(a) shows for a testcase in TSMC 65nm technology that the average power of a circuit can vary by up to 20% depending on the selection of the overdrive modes (the nominal mode is fixed at (500MHz, 0.9V)). Even when the overdrive frequency is also fixed, Figure 1(b) shows for a testcase in TSMC 65nm technology that the average power can vary by up to 15% for different overdrive voltages. Therefore, it is clear that we can reduce design overhead by optimizing the signoff modes. Our work studies the signoff mode optimization problem, which seeks the optimal nominal frequency ( $f_{nom}$ ), nominal voltage ( $V_{nom}$ ), overdrive frequency ( $f_{OD}$ ) and overdrive voltage ( $V_{OD}$ ) with respect to optimization objectives and constraints in terms of circuit area, performance and power.

Traditional multi-corner and multi-mode design is conducted by applying a common constraint during circuit implementation – synthesis, place and route (SP&R) – and verifying every corner and mode at the signoff stage [8]. Other approaches apply additional margins during physical design or implement incremental optimization for all the corners and modes [6]. However, these



(a)  $P_{avg}$  contour plot.



(b)  $P_{avg}$  vs.  $V_{OD}$  for fixed  $f_{OD}$ .

Fig. 1. Circuits signed off at the same nominal frequency (500MHz) and voltage (0.9V) but different overdrive frequencies and voltages. Testcase: *AES*. Technology: TSMC 65nm. (a) Average power ( $P_{avg}$ ) contour plot shows a range of values of up to 20%. (b) Circuit power does not change monotonically with overdrive voltage, but rather tends to have a unimodal behavior.

approaches can introduce poor timing predictability and be very time-consuming [6]. In recent years, EDA tools have offered *Multi-Corner-Multi-Mode (MCMM)* capability [12], [13]. MCMM methodology simultaneously analyzes and optimizes at all corners and modes of operation throughout the SP&R flow, to obtain improved quality of results (QoR). Applying MCMM throughout the entire SP&R flow can result in better timing convergence at the cost of increased runtime. The adaptive MCMM flow introduced in [4] identifies and satisfies constraints only at “dominant” modes, where a mode is said to be dominant if the circuit implementation is mainly constrained by the requirements at that mode. In other words, a circuit that satisfies the constraints at dominant modes should also satisfy design constraints at all other modes. By identifying such modes, the adaptive MCMM flow reduces runtime and memory usage in IC implementation while retaining similar QoR to optimization at all modes and corners. A weakness of the adaptive MCMM technology is that it only focuses on the dominant mode during implementation. Whenever there is a dominant mode, there can be an overdesign at non-dominant modes. For example, our experimental results in Figure 1 show that a circuit implemented to satisfy a dominant mode has up to 15% power consumption overhead for non-dominant modes (i.e., when comparing circuits signed off with overdrive frequency of 950MHz, and overdrive voltages of 1.03V and 1.13V). Circuit power varies with signoff voltage because when signing off at a lower voltage, buffer insertion to meet timing constraints leads to higher power consumption. On the other hand, although circuit area decreases with signoff voltage, power consumption increases with the operating voltage. The optimal signoff voltage should be in between, which leads to the unimodal behavior shown in Figure 1(b). Thus, it is necessary

to define the dominant mode *before* implementation.

In this paper, we propose a method to analyze and identify dominant modes before implementation so that the overdesign resulting from signoff at a dominant mode can be reduced. Moreover, we propose design methodologies to optimize operating mode definitions for multi-mode signoff.

Following are the contributions of our work.

- 1) We propose a methodology to analyze and identify the dominant modes before circuit implementation.
- 2) We show that for signoff optimization, *equivalent dominance* of all modes should be achieved to avoid overdesign.
- 3) Based on the property of equivalent dominance, we reduce the runtime of searching for optimal signoff modes by reducing the solution space for signoff mode selection.
- 4) We develop design methodologies for signoff optimization that show 5-7% improvement in performance compared to traditional “signoff and scale” method. The signoff modes identified by our proposed flow lead to only 0.6% performance and 8% power overheads compared to the optimal result obtained by exhaustive search over all possible combinations of signoff modes.

Our paper is organized as follows. In Section II we define the *design cone*, based on which we give a definition of *equivalent dominance*; we also propose a guideline for multi-mode signoff optimization. Section III formulates three problems for signoff optimization. In Section IV, we propose circuit-implementation flows for signoff optimization. We present our experiments and results in Section V and conclude the paper in Section VI.

The following notation is used in the discussion below.

- Signoff frequencies:  $f_{nom}$  and  $f_{OD}$
- Signoff voltages:  $V_{nom}$  and  $V_{OD}$
- Duty cycle in overdrive mode:  $r$  ( $0 < r < 1$ )
- Power consumption at two modes:  $P_{nom}$  and  $P_{OD}$
- Average power:  $P_{avg} (= (1 - r) \cdot P_{nom} + r \cdot P_{OD})$
- Peak power:  $P_{peak} (= P_{OD})$

## II. DOMINANCE OF MODES

### A. Design Cone

To analyze the dominance of modes, we give definitions of *mode* and *design cone* as follows.

**Definition:** A *mode* is a (frequency, voltage) pair.

**Definition:** Given a mode  $M$ , the *design cone* of mode  $M$  is the union of all the feasible (frequency, voltage) operating modes for circuit implementations that are signed off at mode  $M$ .

For example, Figure 2 illustrates the design cone  $R$  (shaded region) of a nominal mode  $A$ . Since the region of the design cone is determined by frequency vs. voltage tradeoff curves of the circuits signed off at mode  $A$ , the boundary of the design cone is determined by the minimum and maximum circuit frequencies at different voltages.

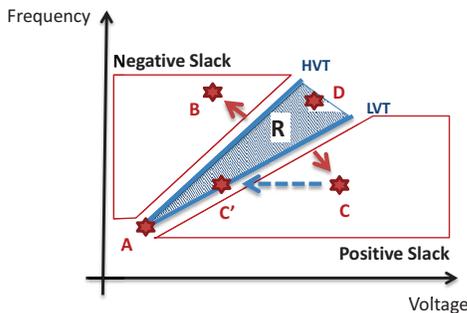


Fig. 2. Design cone and mode slacks. The shaded region is the design cone of mode  $A$ . A circuit signed off with mode  $A$  will have negative (resp. positive) timing slacks when operated at mode  $B$  (resp.  $C$ ).

To study the feasible minimum and maximum frequencies at different voltages, we model the corner cases of timing-critical paths in a digital circuit by simulating chained standard cells with different gate types, threshold voltages ( $V_T$ ) and fanouts. We use standard cells from TSMC 65nm libraries. The simulation results in Figure 3 show that the frequency (reciprocal of path delay) of an inverter chain increases essentially linearly as supply voltage increases. Therefore, we approximate the frequency vs. voltage tradeoff curve of a critical path as a straight line [3], where the boundary of a design cone is determined by the straight lines with maximum and minimum slopes. Even though the frequency vs. voltage tradeoffs may not be exactly straight lines, our approximate tradeoff curves are sufficient for solution space estimation.

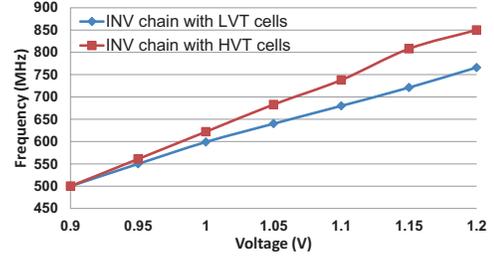


Fig. 3. Frequency vs. voltage tradeoffs for inverter chains. LVT and HVT lines represent different circuits, but both satisfy the timing constraint (500MHz at 0.9V). The HVT line has a steeper slope because HVT cells have higher sensitivity to voltage changes.

Data in Table I show that the slopes of frequency vs. voltage tradeoffs are mainly determined by the threshold voltages of standard cells. Meanwhile, gate type and fanout have little influence on the slope of frequency vs. voltage tradeoffs. Furthermore, changing wire resistance between consecutive inverters from  $0.016\Omega$  ( $0.1\mu\text{m}$  [2]) to  $160\Omega$  (1mm) affects the slope of the frequency vs. voltage tradeoff by less than 2%, while the change due to threshold voltage is  $\sim 30\%$ . We also observe that circuits with high threshold voltage (HVT) cells have steeper tradeoff slopes compared to circuits with low threshold voltage (LVT) cells; this is also observed in [5], where in 45nm CMOS, the slope of the frequency vs. voltage tradeoff is three times larger for HVT than for LVT cells.

TABLE I  
SLOPES OF FREQUENCY VS. VOLTAGE TRADEOFFS FOR DIFFERENT CIRCUITS (CHAINED STANDARD CELLS). DELAY =  $2ns$  AT  $V = 0.9V$ .

$V_T$	Fanout	Slopes (MHz/volt)		
		INV	NAND	NOR
LVT	4	887	800	936
LVT	16	776	787	877
HVT	4	1167	1176	1260
HVT	16	1126	1217	1246

Note that delay and supply voltage of a circuit also affect the frequency vs. voltage slope. However, a design cone is defined at a mode where the delay (reciprocal of frequency) is fixed. Thus, the slopes and the resulting design cone are mainly affected by threshold voltages in the critical paths of a circuit. Since the slopes of frequency vs. voltage tradeoff curves in the design cone are only affected by  $V_T$ , the upper (resp. lower) boundary of the design cone at a given mode can be estimated by synthesizing a circuit at the mode with only HVT (resp. LVT) cells. For example, Figure 3 illustrates that the design cone for mode (500MHz, 0.9V) is bounded by the frequency vs. voltage slopes of HVT and LVT cells. The steeper slope of HVT line is because of HVT cells' high sensitivity to voltage changes, while larger power consumption is due to buffer insertion and larger cells in the HVT-only circuit.

### B. Dominance

When a circuit is operated at a mode which is outside of the design cone corresponding to the signoff mode, positive or negative slacks occur. In Figure 2, point  $A$  indicates the nominal-signoff

mode. When a mode (e.g., mode  $C$ ) is located in the right-bottom region outside of the design cone of the signoff mode (e.g., mode  $A$ ), positive slack is introduced. Such slack can be utilized by either increasing the frequency of mode  $C$  to improve performance, or decreasing the operating voltage to reduce power consumption. We say that the existence of positive timing slacks indicates *overdesign*.

We illustrate the use of positive slack to reduce power without introducing penalties in performance or circuit area, using mode  $A$  and mode  $C$  in Figure 2. We select a mode  $C'$  that is located on the lower boundary of the design cone corresponding to mode  $A$ . Mode  $C'$  has the same frequency as mode  $C$ . By our definition, a design cone represents all circuits that can be signed off at the corresponding mode. Further, the lower boundary of a design cone indicates the circuit with the loosest timing constraints. Thus, any circuit signed off at mode  $A$  satisfies timing constraints at mode  $C'$ , where circuits signed off with mode  $A$  and mode  $C$  can operate at mode  $C'$  without timing violation. Moreover, mode  $C'$  has lower operating voltage than mode  $C$ , which leads to less power consumption, while both have the same performance. Hence, the positive slack can be exploited to reduce power without introducing penalties in performance or area.

On the other hand, when a mode (e.g., mode  $B$ ) is on the left-upper side of the design cone of the signoff mode (e.g., mode  $A$ ), negative timing slack occurs. This is because mode  $B$  has tighter timing constraints. Signing off at mode  $A$  cannot satisfy the timing requirement at mode  $B$ . Such negative slack can be eliminated by increasing the operating voltage at mode  $B$ .

**Definition:** Given two modes  $M_1$  and  $M_2$ , if mode  $M_2$  shows positive slacks with respect to mode  $M_1$ , we define mode  $M_1$  as the *dominant mode*, and mode  $M_2$  as the *dominated mode*.

For example, when considering mode  $A$  and mode  $C$  in Figure 2, mode  $A$  is the dominant mode and mode  $C$  is the dominated mode. The dominant mode has tighter constraints, so when constraints of both modes need to be satisfied, the dominant mode determines the properties of a design. Such properties can be interpreted as area, number of instances, total capacitance, slope of the frequency vs. voltage tradeoff curve, etc. When neither of two modes is dominant with respect to the other, we say that the two modes demonstrate *equivalent dominance*. In other words, their constraints are equivalently strict and the properties of a design are determined by both of the modes. Furthermore, such properties should be similar to those of the design signed off at either of the two modes. In Figure 4, modes  $A$  and  $B$  exhibit equivalent dominance.

**Definition:** Given two modes:  $M_1$  and  $M_2$ , when mode  $M_1$  is in the design cone of mode  $M_2$  and mode  $M_2$  is in the design cone of mode  $M_1$ , we say that mode  $M_1$  and mode  $M_2$  exhibit *equivalent dominance*.

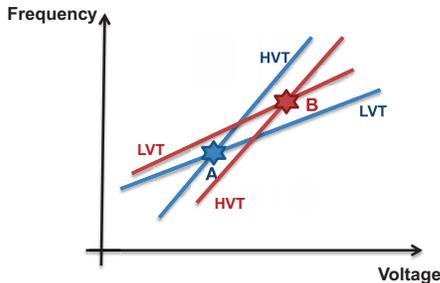


Fig. 4. Modes  $A$  and  $B$  exhibit equivalent dominance, where they are in each other's design cone.

Based on the equivalent dominance concept, we state the following.

**Lemma 1:** If two modes do not exhibit equivalent dominance, then each mode is outside of the design cone of the other mode.

**Proof (by contradiction):** Suppose *Lemma 1* is false (hypothesis), i.e., modes  $M_1$  and  $M_2$  do not exhibit equivalent dominance, but one mode ( $M_1$ ) is located in the design cone of the other

( $M_2$ ). According to the definition of design cone, any point in the design cone of  $M_2$  lies on a frequency vs. voltage tradeoff curve corresponding to a circuit signed off at  $M_2$ . Therefore, there is at least one circuit with a frequency vs. voltage tradeoff curve that passes through both  $M_1$  and  $M_2$ . This means that  $M_2$  is also in the design cone of  $M_1$ . Hence, modes  $M_1$  and  $M_2$  exhibit equivalent dominance, contradicting our initial assumption.  $\square$

**Lemma 2:** Multi-mode signoff at modes which do not exhibit pairwise equivalent dominance leads to overdesign.

**Proof:** If a set of modes does not exhibit pairwise equivalent dominance, then there exist two modes for which equivalent dominant does not hold. According to *Lemma 1*, neither mode is located in the design cone of the other. Then, one of the modes must be dominant, and the other dominated. By definition of a dominated mode, the circuit being implemented at the dominated mode will have positive timing slack. Therefore, at least one mode will be overdesigned if a set of modes does not exhibit pairwise equivalent dominance.  $\square$

Figure 5 shows an example where four modes exhibit equivalent dominance. From the figure, we can expect that as the number of modes increases while still maintaining the mutual equivalent dominance relationship, the feasible design space for signoff defined by the design cones will shrink (eventually to one line).

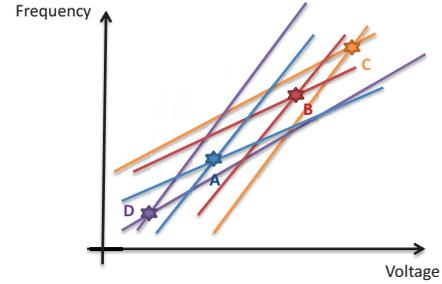


Fig. 5. Four modes exhibit equivalent dominance. The feasible design space is the overlap region of design cones of modes  $A$ ,  $B$ ,  $C$  and  $D$ , which will eventually shrink to line  $D-A-B-C$  as the number of modes increases.

### III. PROBLEM FORMULATION

In order to sign off a circuit that operates at both nominal and overdrive modes, we need to select four parameters:  $f_{nom}$ ,  $V_{nom}$ ,  $f_{OD}$  and  $V_{OD}$ .

**Definition:** We define the problem where  $m$  parameters are given, and  $n$  parameters must be determined, as the  $m + n$  problem. In particular, we are interested in cases where  $m + n = 4$ , and  $m = 0, 1, 2, 3$ .

#### The 3 + 1 Problem

We classify the 3 + 1 problem into two types. (1) The first type, where two frequencies and one voltage are given, is a common scenario in typical IC design flows. This is because  $f_{nom}$  and  $V_{nom}$  are usually defined by the technology node, and  $f_{OD}$  is usually determined by the (market-driven) product specification. Since the performance at both modes is predefined, the objective in this kind of problem can be minimization of power consumption or area. In light of package and reliability requirements, the maximum operating voltage and the peak power consumption are usually set as constraints. (2) In the second type, two voltages and one frequency are given, and we search for the unknown frequency for signoff optimization. Such an optimization can be used to maximize performance under an energy budget.

#### The 2 + 2 Problem

There are four variants of 2 + 2 problems: (1) given two frequencies, search for signoff voltages; (2) given one mode, search for the other mode; (3) given two voltages, search for signoff frequencies; and (4) given a voltage at one mode and a frequency

at the other mode, search for the other two parameters. The third variant is not a use model of interest to real-world product design teams, because designers care mostly about performance or power consumption, neither of which can be determined when only signoff voltages are given [10]. In the fourth variant, the operating voltage at one mode is unrelated to the frequency at the other mode; hence this too does not reflect any practical use model.

In our work, we study the following 2 + 2 problems.<sup>1</sup>

*Definition of the FIND\_OD Problem:*

**Inputs:**  $f_{nom}$ ,  $V_{nom}$  and  $r$

**Objective:** Maximize  $f_{OD}$

**Constraints:**  $P_{peak} \leq C_1$ ;  $P_{avg} \leq C_2$ ;  $V_{OD} \leq C_3$

**Outputs:**  $f_{OD}$  and  $V_{OD}$

*Definition of the FIND\_NOM Problem:*

**Inputs:**  $f_{OD}$ ,  $V_{OD}$  and  $r$

**Objective:** Maximize  $f_{nom}$

**Constraints:**  $P_{peak} \leq C_1$ ;  $P_{avg} \leq C_2$

**Outputs:**  $f_{nom}$  and  $V_{nom}$

*Definition of the FIND\_VOLT Problem:*

**Inputs:**  $f_{nom}$ ,  $f_{OD}$  and  $r$

**Objective:** Minimize  $P_{avg}$

**Constraints:**  $P_{peak} \leq C_1$ ;  $V_{OD} \leq C_2$

**Outputs:**  $V_{nom}$  and  $V_{OD}$

The 2 + 2 problems can always be reduced to 3 + 1 problems by sweeping one unknown parameter. Figure 6 illustrates the reduction relationships. The FIND\_OD problem is reduced to the 3 + 1 problem by sweeping  $V_{OD}$ . A range of  $V_{OD}$  values, together with given  $f_{nom}$  and  $V_{nom}$ , are fed into the 3 + 1 problem solver. Among the output  $f_{OD}$ 's, the one that offers the highest performance is selected as the solution of the FIND\_OD problem. Similarly, the FIND\_NOM problem can be reduced to a 3 + 1 problem by sweeping  $V_{nom}$ . For the FIND\_VOLT problem, where two frequencies are given, one can sweep either  $V_{nom}$  or  $V_{OD}$ . If we sweep  $V_{nom}$ , then from the outputs of the 3 + 1 problem, we select the  $V_{OD}$  and corresponding  $V_{nom}$  that offer minimum power consumption as the output of the FIND\_VOLT problem.

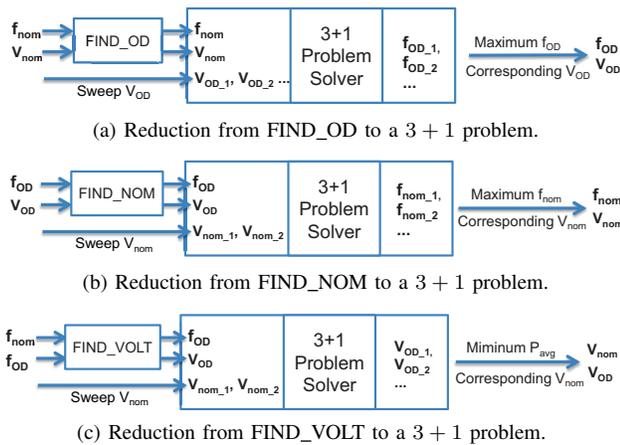


Fig. 6. Reduction from 2 + 2 problems to 3 + 1 problems.

<sup>1</sup>To our knowledge, the 1 + 3 problem would not occur in a real product design context. Moreover, it could be solved by sweeping one parameter at a time and optimally selecting the other two parameters (i.e., reducing to the 2 + 2 problem). The 0 + 4 problem is also not a practically relevant formulation. Therefore, we do not study these problems in this paper.

## IV. METHODOLOGY

**Design space reduction based on equivalent dominance.** In MCMM methodology, all modes need to be analyzed during the implementation. Thus, execution time of MCMM SP&R is significantly slower than that of the conventional single-mode methodology [8]. The design space for signoff increases exponentially with the number of operating modes. Thus exhaustive search for optimal signoff modes (e.g., by implementing circuits with MCMM methodology at many trial combinations of modes in a design space) is infeasible. We propose to reduce the design space for signoff based on the concept of equivalent dominance described in Section II. According to *Lemma 2*, signing off circuits at modes that are not equivalently dominant will lead to overdesigned circuits. *Lemma 2* also tells us that a circuit implementation without overdesign is only feasible when the signoff modes have a mutual equivalent dominance relationship. Since overdesign is equivalent to a lower QoR, we propose to search only the design space for signoff modes in which the equivalent dominance property holds; this is much smaller than the entire feasible design space.

**Design cone approximation.** To identify whether a pair of modes are equivalently dominant, we need to know the design cones of the modes. During the initial step of selecting trial modes when circuits have not yet been implemented, we estimate the design cones using a two-step procedure. First, given the frequency and voltage of a mode, we create HVT-only and LVT-only inverter chains. The number of stages in each inverter chain is selected such that the delays of the inverter chains match the reciprocal of the given frequency at the given voltage. Second, we simulate the HVT and LVT inverter chains at different voltages to obtain the frequency vs. voltage tradeoff curves that define the boundary of a design cone.

**QoR analysis within the design cone.** Within the design cone, we study two extreme cases: circuits dominated by HVT cells, and circuits dominated by LVT cells. In the first case, since HVT cells are slow, buffers are inserted to meet timing constraints. The additional buffers lead to larger area and capacitance. Thus, power consumption increases. In the second case, circuit implementation with a large number of LVT cells increases leakage power overhead. However, LVT cells in non-critical timing paths can be replaced by HVT cells to reduce leakage power. The optimal signoff mode will be located in between these two extreme cases.

### A. 3 + 1 Problems

As mentioned in Section III, there are two types of 3 + 1 problems. Correspondingly, we propose two kinds of methodologies. In the first type of 3 + 1 problem, given two frequencies ( $f_a$ ,  $f_b$ ) and a voltage ( $V_a$ ), we seek to find another voltage ( $V_{var}$ ) that minimizes circuit power. To solve such a problem, we first calculate the approximate design cone for the mode defined by  $f_a$  and  $V_a$  using the design cone approximation method described above. Based on the property of equivalent dominance, we select the range of  $V_{var}$  defined by the intersection of  $f_b$  and the design cone. We then perform a binary search along the feasible range of  $V_{var}$ , i.e., for each candidate  $V_{var}$  value in the binary search, we run an MCMM circuit implementation. Finally, we choose the  $V_{var}$  which results in the circuit with minimum power.

In the second type of 3 + 1 problem, given power constraints, two voltages ( $V_a$ ,  $V_b$ ) and a frequency ( $f_a$ ), we seek to find the maximum frequency ( $f_b$ ). A similar methodology as that for the first type of 3 + 1 problem can be applied. We define the design cone for the mode at  $f_a$  and  $V_a$ . By *Lemma 2*, the intersection of the design cone and  $V_b$  defines a range of  $f_{var}$ . We then perform MCMM circuit implementation with each  $f_{var}$  and given parameters ( $f_a$ ,  $V_a$  and  $V_b$ ). Since power consumption increases monotonically with frequency, binary search can be applied to reduce the overall complexity of solving these problems. We output the maximum frequency obtained under the given power consumption constraints.

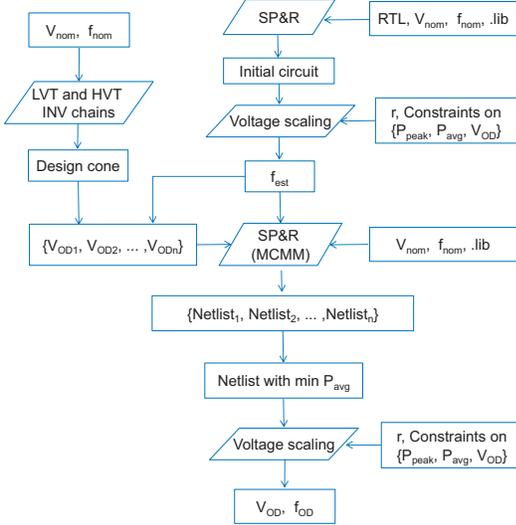


Fig. 7. Flow of our methodology for the 2 + 2 problem where the nominal mode is given.

### B. The FIND\_OD 2+2 Problem

Although we can solve the problem FIND\_OD by running the 3 + 1 solver with trial frequencies within the feasible frequency range, there is a more efficient way to search for the optimal overdrive mode. Specifically, we solve the FIND\_OD problem using the flow depicted in Figure 7. As shown in the Figure, we first implement an *initial gate-level netlist* at the given nominal mode. Then, we estimate the maximum overdrive frequency ( $f_{est}$ ) under power constraints. The  $f_{est}$  is obtained by running timing and power analyses on the initial netlist with increasing  $V_{OD}$ .<sup>2</sup> Second, we obtain the design cone of the nominal mode using the above-described approximation method. The design cone, together with  $f_{est}$ , defines several approximate overdrive modes ( $f_{OD}$  and  $V_{OD}$  pairs) inside the design cone as indicated by the red line in Figure 8. Third, we implement MCMM to sign off circuits at these approximate-optimal overdrive modes and the given nominal mode. Fourth, from the output netlists, we select the one that shows largest positive power slack and implement voltage scaling on that netlist. Under the predefined power and voltage constraints, the output maximum frequency from voltage scaling, along with the corresponding voltage, define our heuristic optimal overdrive mode.

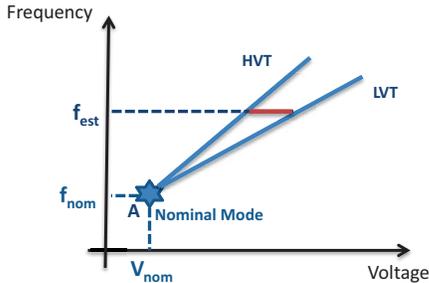


Fig. 8. Estimation of  $f_{est}$  using voltage scaling. The intersection of  $f_{est}$  and the design cone indicates approximate-optimal overdrive modes.

### C. The FIND\_NOM 2+2 Problem

The FIND\_NOM problem is similar to the FIND\_OD problem. The only difference is that we implement the reference circuit and calculate the design cone at the given overdrive mode to obtain a set of approximate-optimal nominal modes. We then apply

<sup>2</sup>We enable timing and power analyses with libraries characterized at different voltages. We increase  $V_{OD}$  by 10mV in each run until circuit power exceeds the pre-defined constraints.

MCMM signoff to obtain a number of circuits with approximate-optimal nominal modes and the fixed overdrive mode. Under power constraints, we then run voltage scaling on the obtained circuits to search for the maximum  $f_{nom}$  and corresponding  $V_{nom}$  which determine the output nominal mode.

### D. The FIND\_VOLT 2+2 Problem

As mentioned in Section III, the FIND\_VOLT problem can be solved by providing a set of voltages to the 3 + 1 solver. We convert the FIND\_VOLT problem to a 3 + 1 problem by providing the minimum voltage of the process as  $V_{nom}$ . We then search for the optimal voltages using exhaustive search with the upper bound for  $V_{nom}$  being the maximum allowed voltage for the process.

## V. EXPERIMENTS AND RESULTS

### A. Implementation Flow and Tools

Our experiments use two RTL designs – AES (~15K instances at 65nm) and JPEG (~40K instances at 65nm) – obtained from the OpenCores website [17]. These designs are implemented using TSMC 65nm HVT, NVT and LVT libraries. We characterize all libraries at operating voltages ranging from 0.8V to 1.2V in steps of 0.01V using *Cadence Library Characterizer vEDI9.1* [12]. The designs are synthesized using *Synopsys Design Compiler vC-2009.06-SP2* [13] and then placed and routed using *Cadence SoC Encounter vEDI10.1* [12]. We further use *Cadence SoC Encounter* for timing and power analysis. *HSPICE* [14] is used for all transistor-level modeling and simulation.

### B. Design of Experiments

We design experiments to solve the 2 + 2 problems using the flows proposed in Section IV. In our experiments, we implement synthesis at nominal mode, and MCMM P&R with both nominal and overdrive modes. To eliminate tool noise, we execute each P&R run three times, perturbing the timing constraints by a small amount in each run [9]. We assume that the duty cycle of nominal mode is 95%.

In the FIND\_OD problem, where the nominal mode is given, we search for the overdrive mode. Three instances of the FIND\_OD problem are studied (Table II). Two of these involve implementation of AES. In both of these cases, nominal mode is defined as  $f_{nom} = 500\text{MHz}$ ,  $V_{nom} = 0.9\text{V}$ . The maximum voltage constraint is 1.2V and the average power constraint is 25mW in both cases, but the peak power constraints are 40mW and 50mW, respectively. The third case is based on the JPEG design, with  $f_{nom} = 600\text{MHz}$ ,  $V_{nom} = 0.9\text{V}$ , maximum voltage constraint of 1.2V, and average and peak power constraints of 60mW and 100mW respectively. Due to the similarity between the FIND\_OD problem and the FIND\_NOM problem, we omit discussion of FIND\_NOM experiments.

In the FIND\_VOLT problem, for which frequencies are given, we implement our flow to search for optimal voltages. Two instances are addressed (Table III); one is based on AES, and the other is based on JPEG. We set  $f_{nom} = 500\text{MHz}$ ,  $f_{OD} = 600\text{MHz}$  for AES and  $f_{nom} = 600\text{MHz}$ ,  $f_{OD} = 720\text{MHz}$  for JPEG. The peak power constraints are 50mW for AES and 100mW for JPEG.

TABLE II  
EXPERIMENTAL SETUP FOR THE FIND\_OD PROBLEM

Case	Design	$f_{nom}$ (MHz)	$V_{nom}$ (V)	$P_{peak}$ (mW)	$P_{avg}$ (mW)	$V_{max}$ (V)
1	AES	500	0.9	50	25	1.2
2	AES	500	0.9	40	25	1.2
3	JPEG	600	0.9	100	60	1.2

TABLE III  
EXPERIMENTAL SETUP FOR THE FIND\_VOLT PROBLEM

Case	Design	$f_{nom}$ (MHz)	$f_{OD}$ (MHz)	$P_{peak}$ (mW)	$V_{max}$ (V)
4	AES	500	600	50	1.2
5	JPEG	600	720	40	1.2

### C. Results

Figure 9 shows frequency vs. voltage tradeoff curves of *AES* signed off at (500MHz, 0.9V) and at (800MHz, 1.1V) with LVT-only cells and with HVT-only cells. Tradeoff curves of inverter chains are also plotted in the same chart. We observe only a slight difference (less than 7%) between the slopes of the inverter chains and real circuits.

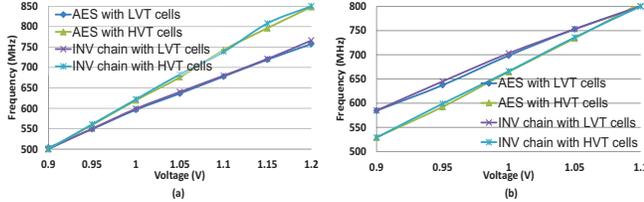


Fig. 9. HSPICE simulation and circuit signoff results for (a) nominal mode and (b) overdrive mode.

Tables IV shows the results of experiments addressing the FIND\_OD problem. Three methods are implemented in our studies: **Signoff&Scale** applies traditional signoff and scale methodology; **Proposed** implements the flow that we propose, and **Reference** uses exhaustive search. The exhaustive search explores the entire feasible solution space for given design parameters, e.g., in *Case 1* we implement MCMM P&R with nominal mode (500MHz, 0.90V) and overdrive modes (670-740MHz, 1.01-1.15V), with step sizes for  $f_{OD}$  and  $V_{OD}$  of 20MHz and 20mV, respectively. The results show that our flow offers 5-7% improvement in overdrive performance compared to the **Signoff&Scale** method while maintaining similar area and power. This is a significant improvement, considering that even 20% improvement in performance per new technology generation is now quite difficult to achieve. The results also show that the overdrive frequency obtained from our proposed method is within 0.6% of that obtained from the **Reference** method. We believe that our proposed method holds promise for determining near-optimal (frequency, voltage) signoff modes in the use cases that we have studied.

Table V gives results for the FIND\_VOLT problem achieved by our proposed method (**Proposed**) and the reference method (**Reference**). The results of the proposed flow show some overheads in power ( $\sim 8\%$ ) compared to the **Reference** method.

TABLE IV

AREA, PERFORMANCE AND RUNTIME OF PROPOSED FLOW FOR FIND\_OD PROBLEM. THE PROPOSED FLOW ACHIEVES 5-7% IMPROVEMENT IN PERFORMANCE COMPARED TO THE SIGNOFF & SCALE FLOW, AND SIMILAR QOR COMPARED TO THE REFERENCE FLOW.

		Signoff&Scale	Proposed	Reference
<i>AES</i> (Case 1)	$f_{OD}$ (MHz)	711	764	768
	$V_{OD}$ (V)	1.14	1.14	1.15
	area ( $\mu\text{m}^2$ )	31029	32016	32020
	$P_{OD}$ (mW)	49.13	49.14	49.76
	$P_{avg}$ (mW)	21.73	20.90	20.24
	#P&R runs	1	7	32
<i>AES</i> (Case 2)	$f_{OD}$ (MHz)	651	688	692
	$V_{OD}$ (V)	1.07	1.08	1.07
	area ( $\mu\text{m}^2$ )	31029	30727	31910
	$P_{OD}$ (mW)	39.51	39.46	39.55
	$P_{avg}$ (mW)	21.54	20.47	20.42
	#P&R runs	1	4	32
<i>JPEG</i> (Case 3)	$f_{OD}$ (MHz)	783	822	825
	$V_{OD}$ (V)	1.08	1.10	1.10
	area ( $\mu\text{m}^2$ )	161250	161366	158938
	$P_{OD}$ (mW)	100.00	97.98	99.75
	$P_{avg}$ (mW)	49.14	48.08	48.12
	#P&R runs	1	5	32

## VI. CONCLUSIONS

We study the multi-mode signoff optimization problem and introduce the concept of *equivalent dominance* among signoff modes. We show that for a multi-mode design, the modes for signoff

TABLE V

AREA, POWER AND RUNTIME OF PROPOSED FLOW FOR THE FIND\_VOLT PROBLEM. THE PROPOSED FLOW ACHIEVES SIMILAR QOR BUT 4X RUNTIME REDUCTION COMPARED TO THE REFERENCE FLOW.

		Proposed	Reference
<i>AES</i> (Case 4)	$V_{nom}$ (V)	0.92	0.91
	$V_{OD}$ (V)	1.02	1.01
	area ( $\mu\text{m}^2$ )	30948	30960
	$P_{avg}$ (mW)	22.28	20.61
	$P_{OD}$ (mW)	41.08	30.38
	#P&R runs	9	33
<i>JPEG</i> (Case 5)	$V_{nom}$ (V)	0.90	0.89
	$V_{OD}$ (V)	0.99	0.99
	area ( $\mu\text{m}^2$ )	164637	168438
	$P_{avg}$ (mW)	55.45	54.46
	$P_{OD}$ (mW)	82.90	90.25
	#P&R runs	9	33

must maintain a mutual equivalent dominance condition to avoid overdesign. Based on the properties of equivalent dominance, we propose guidelines and methodologies to search for the optimal modes for signoff. Our experimental results indicate that the proposed methodologies can identify signoff modes which lead to 5-7% performance improvement compared to the traditional “signoff and scale” methodology. Our experiments further show that circuits signed off with our flow have 0.6% overhead in performance and 8% overhead in average power compared to the essentially optimal results obtained through exhaustive search.

Our future works include (1) developing faster methodologies to search for optimal signoff modes, (2) treating the signoff mode optimization problem according to the theory of efficient global optimization with minimum number of function evaluations, and (3) taking into consideration additional tradeoffs of design metrics such as circuit area, reliability and design time.

### ACKNOWLEDGMENTS

Research supported in part by funding from IMPACT, SRC, NSF, Qualcomm and Samsung.

### REFERENCES

- [1] D. Bull, S. Das, K. Shivashankar, G. S. Dasika, K. Flautner and D. Blaauw, “A Power-Efficient 32 bit ARM Processor Using Timing-Error Detection and Correction for Transient-Error Tolerance and Adaptation to PVT Variation”, *IEEE JSSC* 46(1) (2011), pp. 18-31.
- [2] T.-B. Chan and A. B. Kahng, “Tunable Sensors for Process-Aware Voltage Scaling”, *Proc. ICCAD*, 2012.
- [3] M. Elgebaly and M. Sachdev, “Variation-Aware Adaptive Voltage Scaling System”, *IEEE TVLSI* 15(5) (2007), pp. 560-571.
- [4] F. D. Meersman, “IC Compiler RM: Reference Methodology with Emphasis on Concurrent MCMM & Signoff Driven Design Closure”, [http://www.synopsys.com/news/pubs/snug/sanJose08/wb1\\_mcmm\\_tutorial.pdf](http://www.synopsys.com/news/pubs/snug/sanJose08/wb1_mcmm_tutorial.pdf).
- [5] M. Meijer, B. Liu, R. V. Veen, and J. P. Gyvez, “Post-Silicon Tuning Capabilities of 45nm Low-Power CMOS Digital Circuits”, *Proc. VLSI Circuits Symposium Digest of Technical Papers*, 2009, pp. 110-111.
- [6] A. Mulgaonkar, “Multicorner-Multimode – A Necessary and Manageable Reality of Design”, [http://www.synopsys.com/apps/protected/docs/pdfs/iccpw/icc\\_mcmm\\_wp.pdf](http://www.synopsys.com/apps/protected/docs/pdfs/iccpw/icc_mcmm_wp.pdf).
- [7] C. R. Parthasarathy, M. Denais, V. Huard, G. Ribes, D. Roy, C. Guerin, F. Perrier, E. Vincent and A. Bravaix, “Designing in Reliability in Advanced CMOS Technologies”, *Microelectronics Reliability* 46 (2006), pp. 1464-1471.
- [8] B. M. Riess, “Multi-Corner Multi-Mode Synthesis in Design Compiler – A Must or Just Nice to Have?”, [https://www.synopsys.com/news/pubs/snug/germany2011/a1\\_riess\\_paper.pdf](https://www.synopsys.com/news/pubs/snug/germany2011/a1_riess_paper.pdf)
- [9] K. Jeong and A. B. Kahng, “Methodology From Chaos in IC Implementation”, *Proc. ISQED*, 2010, pp. 885-892.
- [10] S. Dobre, Qualcomm CDMA Technologies, Inc., *personal communication*, July 2012.
- [11] L. Stok and J. Cohn, “There is Life Left in ASICs”, *Proc. ISPD*, 2003, pp. 48-50.
- [12] “Cadence SOC Encounter User Guide.” [http://www.cadence.com/products/di/first\\_encounter/pages/default.aspx](http://www.cadence.com/products/di/first_encounter/pages/default.aspx)
- [13] “Synopsys Design Compiler User Guide.” <http://www.synopsys.com/Tools/Implementation/RTLsynthesis/DCUltra/pages/default.aspx>
- [14] “Synopsys HSPICE User Guide.” <http://www.synopsys.com/Community/Interoperability/HSPICE/Pages/default.aspx>
- [15] “SDC User’s Guide.” [http://www.actel.com/documents/SDC\\_AN.pdf](http://www.actel.com/documents/SDC_AN.pdf)
- [16] “LEF DEF Reference.” <http://www.si2.org/openeda.si2.org/projects/lefdef>
- [17] “OpenCores.” <http://opencores.org>