Dynamic Power and Energy Management for NCFET-based Processors

Sami Salamin, *Student, IEEE*, Martin Rapp, *Student, IEEE*, Jörg Henkel, *Fellow, IEEE*, Andreas Gerstlauer, *Senior Member, IEEE*, Hussam Amrouch, *Member, IEEE*

***Abstract*—Power and energy consumption are the key optimization goals in all modern processors. Negative Capacitance Field-Effect Transistors (NCFETs) are a leading emerging technology that promises outstanding performance in addition to better energy efficiency. The thickness of the added ferroelectric layer as well as frequency and voltage are the key parameters that impact the power and energy of NCFET-based processors in addition to the characteristics of runtime workloads.**

Unlike existing CMOS technologies, operating NCFET-based processors at a higher frequency than the required minimum can result in power/energy minimization. The optimal operating point, however, strongly depends on dynamic workload characteristics and technology parameters. In this work, we propose and implement the first NCFET-aware power and energy management approach that minimizes the processor's power and energy through optimal voltage/frequency selection under different runtime scenarios. Such an NCFET-aware approach does not result in any trade-off between power/energy and performance. Instead, it can achieve higher performance while minimizing energy. A comprehensive, simulation-based evaluation of our runtime management under realistic workloads demonstrates up to 58% energy saving with 2.1x higher performance, and 46% power saving compared to conventional NCFET-unaware management techniques, over the total execution of a benchmark. Compared to state-of-the-art NCFET-aware management techniques, our technique provides up to 49% energy saving and 32% power saving.

***Index Terms*—Negative Capacitance (NC), Negative Capacitance Field-Effect Transistors (NCFET), Power Management Technique, Energy Management Technique.**

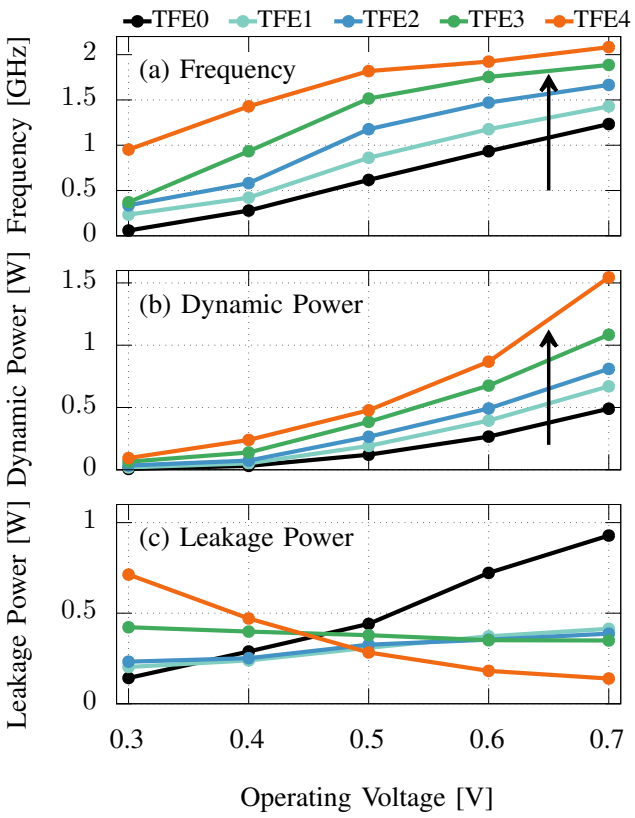

Fig. 1: (a) NCFET boosts the maximum frequency of the processor at a given voltage. Gains increase with a thicker ferroelectric layer. (b) NCFET increases the dynamic power due to the increase in the frequency and total gate capacitance of the transistor. However, NCFET reduces the dynamic power at a fixed frequency as it allows to operate at a lower voltage. (c) NCFET with a thin ferroelectric layer weakens the dependency of leakage on voltage. At higher thicknesses, the dependency is reversed due to the negative DIBL.

I. Introduction

Power and energy minimization of a processor is the primary concern in almost all applications [1], especially in battery-powered devices [2]. The power consumption of any processor depends on its operating frequency (f) and operating voltage (V). Energy consumption is determined by integrating the power trace over the total execution time of the workload.

In conventional FinFET, both total power and total energy are minimized by operating at the minimum voltage (V_{min}) that is required to sustain the minimum required frequency (f_{min}) under a given performance constraint (e.g., fixed deadline). Reducing the supply voltage of a processor reduces both dynamic and leakage power super-linearly. Similarly, despite the increase in execution time, due to the approximately linear relationship between frequency and voltage in the normal super-threshold region, both dynamic and leakage energy are minimized at reduced voltage and frequency [2], [3]. Therefore, traditional power/energy management techniques operate by selecting the minimum frequency f_{min} and the

Manuscript received April 18, 2020; revised June 12, 2020; accepted July 6, 2020. This article was presented in the International Conference on Compilers, Architecture, and Synthesis for Embedded Systems 2020 and appears as part of the ESWEEK-TCAD special issue.

Corresponding author: Sami Salamin

Sami Salamin, Martin Rapp, Jörg Henkel, and Hussam Amrouch are with the Chair for Embedded Systems, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany (e-mail: {sami.salamin, martin.rapp, henkel, amrouch}@kit.edu).

Andreas Gerstlauer is with Department of Electrical and Computer Engineering, University of Texas, Austin, U.S.A. (e-mail: gerstl@ece.utexas.edu).

corresponding minimum voltage (V_{min}) (i.e., V/f pair) to exploit these dependencies [4].

Because these dependencies vary among different technologies, power and energy management techniques must be revisited and investigated when a new technology is introduced. This holds even more when it comes to emerging technologies in which the underlying physics fundamentally differ from conventional FinFET technologies.

Negative Capacitance Field-Effect Transistor (NCFET) technology is one of the promising emerging technologies that provides a considerable improvement in a circuit's performance over conventional FinFETs. This improvement is due to an additional ferroelectric (FE) layer within the gate stack of the transistor, which manifests itself as a Negative Capacitance (NC). NC results in a voltage amplification at the internal gate of the transistor, which boosts the electric field. This, in turn, has two key implications [5]: (1) NCFET-based circuits can operate at a higher frequency at the same operating voltage (V), (2) NCFET-based circuits can achieve the same frequency at lower operating voltage leading to considerable power savings.

Power and performance of NCFET-based processors: Prior works [5], [6] have shown that NCFET-based processors exhibit an observable performance enhancement compared to conventional FinFETs due to the voltage amplification provided by the added FE layer at the transistor's gate. Fig. 1(a)[1] shows how the maximum frequency (performance) of a processor (details in Section VI) at given V increases when a thicker FE layer is employed. The FE layer thickness is referred to as TFEx, where x is the layer thickness in nanometers. TFE0, in this context, refers to conventional FinFETs technology in which no FE layer is in use. The FE layer thickness is limited to less than 4nm to ensure hysteresis-free operation, which is essential for logic circuits [7].

NC increases the total gate capacitance of FinFETs, which, together with an increased frequency, results in a higher dynamic power at the same operating voltage (Fig. 1(b)). Importantly, however, increasing the thickness of the FE layer inverses the dependency of leakage power on V due to the negative drain-induced barrier lowering effect (DIBL) [8], as shown in Fig. 1(c). DIBL is a short-channel effect referring originally to a reduction of the threshold voltage of the transistor (V_{th}) and thus an increase in the leakage current when the voltage decreases [8]. Therefore, reducing V at a high thickness *increases* the leakage power, instead of decreasing it as in conventional FinFETs technologies.

Due to negative DIBL, the minimum *total power* consumption under any fixed frequency can be achieved at an optimal voltage (V_{opt}) that is higher than V_{min} in NCFET-based processors. This is because of the novel trade-off between leakage and dynamic power [9]. However, V_{opt} depends on the dynamic workload characteristics, which need to be considered to minimize power. Furthermore, a higher voltage will also allow the processor to be operated at a higher than needed frequency. Under a fixed performance (i.e., throughput) constraint, a power-optimal strategy may therefore be to alternate between higher-frequency and idle modes rather than operating at a stable V/f point.

Unique NCFET characteristics also demand different approaches for energy vs. power optimization. The work in [10] showed the importance of specific voltage and frequency selection for energy minimization considering the novel trade-offs in NCFET. Results showed that total energy in NCFET technology will be minimized by operating a processor at a frequency and voltage that is *higher* than the minimum, but not necessarily the same as when minimizing power. Similar to power, the precise operating point is, however, strongly dependent on workload characteristics. The investigation in [10] was performed using an analytical model to select optimal V/f pairs for synthetic benchmarks that ignored dynamic runtime characteristics of workloads and only provided an upper bound on the achievable gains.

In this work, we provide a detailed and complete runtime management technique for NCFET-based processors that covers both energy and power optimization for realistic workload scenarios. We extend prior energy management to account for dynamic workload behavior, and we further expand our approach to dynamic power management, where we determine optimal policies for periodically operating a processor between a higher frequency and idle mode to minimize power under average performance constraints.

Our novel contributions within this paper are as follows:
(1) We extend the previously published analytical energy model of NCFET-based processors into a comprehensive power and energy model. The model allows designers to explore the simultaneous impacts of voltage, frequency, workload characteristics and ferroelectric layer thickness. We use the model to study two use-cases: minimizing power under a performance/throughput constraint, and minimizing energy for a fixed amount of work.
(2) We present novel NCFET-aware power and energy minimization algorithms that select optimal V/f pairs at runtime. This algorithm can provide the core of dynamic power/energy management techniques for future NCFET-based processors.
(3) We explore the dependency of the V/f pair that is selected by our algorithm and the optimal power/energy on workload and technology parameters. We also quantify the gains based on an accurate circuit-level simulation of real workloads in addition to an analytical exploration.

Our results show that an NCFET processor at the highest ferroelectric thickness can achieve minimal power/energy not at the lowest voltage and frequency, but at a significantly higher operating point. This differs from conventional FinFET for which the optimal energy point is at a near-threshold voltage where the processor's performance is severely degraded. Furthermore, the optimal voltage and frequency strongly depend on the workload characteristics. This is due to the reverse leakage dependency, which ultimately requires novel NCFET-aware power/energy management algorithms.

II. Background

In this section, we cover the background of this work. This includes voltage selection for power minimization and

[1] Section VI describes our methodology to model NCFET-based processors.

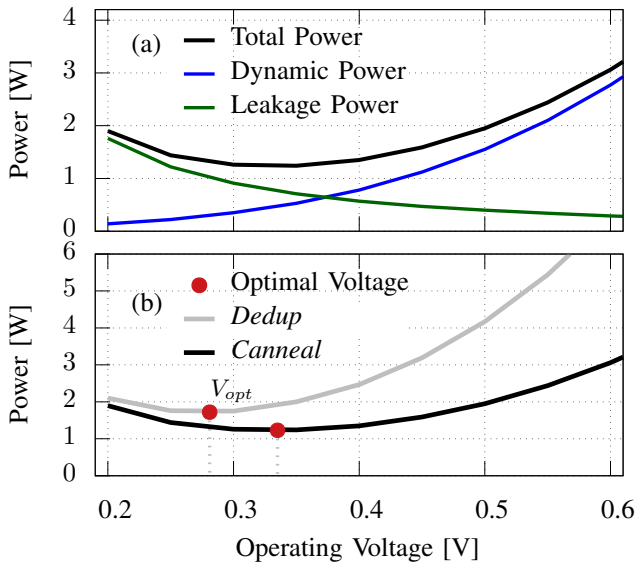

Fig. 2: (a) Total power and its components (i.e., leakage and dynamic) of the master thread of PARSEC *canneal* running on a multi-core system designed in TFE4 at a fixed frequency (1 GHz) depend on the operating voltage. The total power decreases as voltage increases until it reaches an inflection point where it starts to increase again. (b) Total power consumption of two different workloads *canneal* and *dedup* running at the same frequency (1 GHz) over voltage. The total power of different workloads is minimized at different V_{opt}. Note that the ability of NCFET to operate at such low V (0.2 V) is due to the inherent voltage amplification provided by the integrated negative capacitance and does not result in near- or even sub-threshold computing.

its workload dependence, frequency selection for energy minimization, and FE layer thickness dependence in NCFET.

Voltage selection for power minimization: Fig. 2(a)[2] shows the total power consumption and its components (i.e., leakage and dynamic power) of a multi-core system designed in TFE4 running the master thread of PARSEC [11] *canneal* at 1 GHz. The minimum voltage required to sustain this frequency is 0.2 V. Increasing the operating voltage up to 0.6 V *while keeping the frequency constant*, as shown in the figure, increases the dynamic power but decreases the leakage power. As a result, the total power consumption at first reduces until an inflection point is reached after which it starts to increase again as dynamic power becomes dominant. This demonstrates that optimal voltage V_{opt} selection is a must, as the power is minimized at a higher voltage (V_{opt}≈0.35 V in this example) than the minimum voltage in NCFET [9].

Workload dependency: Different workloads have different runtime activities and hence power components differently contribute to the total power. Accordingly, the proportions of the leakage and dynamic power to the total power change dynamically with workload behavior at runtime. As V_{opt} depends on the trade-off between dynamic and leakage power, V_{opt} is expected to change with the workload being executed.

[2] Generated using our methodology replicating the approach from [9].

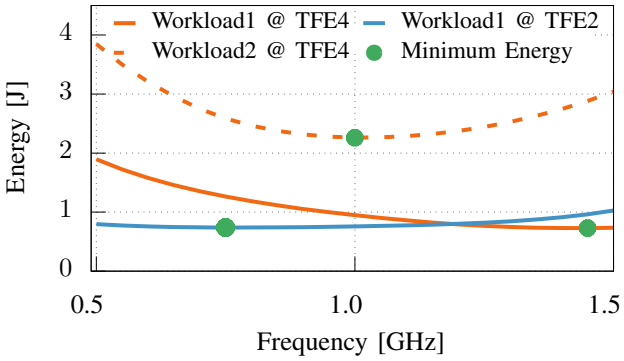

Fig. 3: Energy consumption over frequency (and corresponding voltage) of two synthetic workloads with different dynamic power consumption running on a processor designed in TFE4. Energy is not minimized at the minimum operating frequency but at a higher frequency. As the two workloads have different dynamic power consumption their energy-minimizing frequencies differ. The figure also shows that the energy of the same workload can be optimized at different frequencies for two different FE layers (TFE2 and TFE4), showing the importance of selecting the optimal thickness at design time.

Fig. 2(b)[2] shows the total power across a wide range of voltages for two different workloads running on a multi-core system designed in TFE4. Similarly, sweeping V from V_{min} = 0.2 V to 0.6 V while keeping the frequency constant at 1 GHz in order to explore V_{opt} shows that V_{opt} differs for the two workloads. This demonstrates that the optimal voltage selection should follow the dynamics of workloads. Ignoring such attributes makes any power/energy management technique suboptimal w.r.t. NCFETs technology.

Frequency and voltage selection for energy minimization: Operating a processor at a higher voltage allows the frequency to be increased accordingly. As in [10], operating an NCFET-based processor at a higher frequency than f_{min} could result in a lower energy. Increasing the operating frequency, and the corresponding voltage, increases the dynamic energy, but decreases the leakage energy more strongly. Hence, the total energy decreases. This continues until an inflection point is reached where the dynamic energy becomes prominent and, therefore, increasing the frequency starts to increase the total energy. To show such trends, Fig. 3 plots the energy of two synthetic workloads with different dynamic power consumption in a TFE4 processor as a function of frequency. The figure shows that energy is minimized at a higher frequency than f_{min}=0.5 GHz. More importantly, it also shows how the energy of two different workloads is minimized at different frequencies. However, *the dynamic characteristics of the workload change at runtime*, which is not considered in this example.

Thickness dependency: Fig. 3 plots the energy of a workload using two different FE layer thicknesses. As can be seen, energy depends on the frequency and thickness of the FE layer. Different thicknesses are optimal (minimum energy) at different frequencies under the same workload, showing the importance of selecting the optimal thickness at design time.

In summary, (1) while a processor's power is always minimized at the lowest possible voltage that sustains the required frequency in conventional FinFETs, this does not hold in NCFET at higher FE layer thicknesses [9]. Instead, operating at a higher voltage is beneficial. However, this, in turn, also allows operating at a higher frequency. (2) while a processor's energy is always minimized at the lowest voltage/frequency that satisfies the performance constraint in conventional FinFETs, this does not hold anymore in NCFET [10]. Instead, operating at a higher voltage/frequency is beneficial. The optimal voltage/frequency is, however, different for power vs. energy optimization. (3) thickness dependence in addition to workload characteristics must be considered when selecting voltage/frequency pairs for power and energy minimization. Therefore, developing new NCFET-aware power and energy optimization techniques is indispensable, which is precisely the goal of this work.

III. Related Work

The related work can be clustered in two groups: prior work on power and energy minimization for FinFETs and prior work on NCFETs.

Traditional power and energy minimization for FinFETs: Many techniques have been proposed to minimize a processor's power and energy. We only focus on power/energy management that is based on V/f pair selection, which is the core of DVFS (dynamic voltage and frequency scaling) techniques. DVFS is used in almost all modern processors to optimize power/energy while meeting performance requirements [12]. Conventional DVFS selects the minimum frequency and voltage required to satisfy the performance constraint, and thereby minimizes power/energy in conventional FinFETs [13], [14]. However, because NCFET with a thick FE layer inverses the dependency of leakage power on voltage, existing DVFS techniques become sub-optimal as has been discussed earlier.

When it comes to the optimal energy point in conventional FinFETs, many studies (e.g., [3], [15]) showed that operating processors at near-threshold voltage achieves such a goal. However, operating at such a low voltage leads to severe performance degradation, which limits the applicability of near-threshold computing. We, therefore, focus on the normal super-threshold regime in this paper.

NCFET: Recently, several works have explored NCFET processor design and optimization. [5], [7] presented a comparison between conventional FinFET and NCFET processors under different configurations (i.e., ferroelectric thicknesses). The study in [5] showed how NCFETs impact the performance, power and temperature of a processor. Moreover, recently, a small number of works demonstrated voltage and frequency scaling for NCFETs [9], [10]. In [9], a dynamic voltage scaling (DVS) technique has been proposed to optimize the power consumption of NCFET many-core systems under fixed performance constraints. The work assumes a constant frequency and hence it only explores voltage selection. Furthermore, the work focused solely on power (not energy) minimization, and it studied only a single FE layer thickness, despite the key role that such a parameter plays in NCFET as Fig. 1 demonstrates.

In prior work [10], we performed an analytical exploration of energy optimization in NCFETs by selecting the optimal operating point of frequency and voltage. Analytical results showed that a higher frequency can result in lower energy providing the maximum possible gain. Our previous work, however, did not address dynamic workload dependencies and solely considered energy but not power optimization for synthetic workloads using analytical best-case models.

IV. NCFET-Aware Power and Energy Models

In the following, we present the application, power, frequency and energy models that are needed for the *analytical design-space exploration* in this work. The results of this exploration will later be used to develop runtime V/f pair selection algorithms that will be evaluated using detailed circuit-level simulations. Power, energy and frequency models are developed individually for each thickness and we refer to thickness by (x).

1) Application Model: In the context of this paper, the optimal frequency (f_{opt}) is the operating frequency at which the processor's power/energy is minimized. $V_{min}(f_{opt})$ is the minimum voltage required to sustain f_{opt}. Note that, due to the negative DIBL that impacts NCFET transistors, the minimum power/energy could be achieved at a higher voltage unlike conventional transistors. Therefore, $V_{opt}(f_{opt})$ is the optimal voltage for operating at f_{opt} [9].

To simplify the application model within the design-space exploration, we assume that the performance is linearly affected by frequency. This assumption, even though it is not always true (e.g., due to different clock domains for CPU cores and memory, etc.), greatly reduces the design-space corners. In Section VI, a realistic evaluation will later be presented based on a complex circuit-level simulation for a processor executing actual workloads.

To explore the design-space, we represent a workload by its *ratio of dynamic to total power* on a processor using the highest FE layer thickness running at the common highest frequency ($\hat{f}$) among all thicknesses (i.e., TFE4 at 1.2 GHz). By sweeping this ratio, we can explore a large variety of workload domains from memory-bound to compute-bound applications.

2) Optimization Cases: In this work, we consider two optimization cases:

(a) Power minimization under performance constraints: Minimize average power for executing a given amount of work (W) under a fixed deadline (T) in a given period. The frequency can not be lower than a threshold (f_d) to fulfill T. If the processor is operated at a higher frequency, its execution comprises two parts: First, the processor is active for T_{active}. Then, the processor is idle for $T - T_{active}$. During the idle phase, the processor is assumed to be clock-gated, i.e., its dynamic power is suppressed, but leakage continues. However, in the idle state, the processor can operate at a different voltage (V_{leak}) than in the active state to minimize leakage. Based on dependencies shown in Fig. 1(c), the highest voltage ($V_{leak} = 0.7\,\text{V}$) will be used with TFE3-4 in idle state and the lowest voltage ($V_{leak} = 0.3\,\text{V}$) with TFE0-2.

(b) Energy minimization under fixed work (W): Minimize total energy for executing a fixed amount of work (i.e., non-periodic work). Energy consumption is only incurred during the execution time of W and therefore no idle time exists.

In all scenarios, we assume a single thread is being executed on a single core. *For simplicity, we refer to case (a) as power minimization and case (b) as energy minimization in the remaining parts of this work.*

3) Power and Frequency Models: To develop the dynamic power, leakage power and frequency models, we follow the same methodology as in [5] (see Section VI). We examine dynamic and leakage powers and frequency of the processor for the whole voltage range from 0.2V to 0.7 V, with 50mV steps. To accurately estimate the resulting powers and frequencies, industrial standard signoff tools (Cadence Voltus and Tempus) are used. The process is repeated for each FE layer thickness individually (1-4 nm including FinFET). Finally, and similar to [9], we fitted the results into mathematical equations to use them within our technique.

The minimum voltage $V_{min}^{(x)}(f)$ at thickness x required to sustain a frequency f, and inversely, the maximum sustainable frequency $f_{sus}^{(x)}(V)$ at thickness x and V are given by:

$$V_{min}^{(x)}(f) = \left(\frac{\frac{1}{f} - c_{freq}^{(x)}}{a_{freq}^{(x)}} \right)^{\frac{1}{b_{freq}^{(x)}}} \tag{1}$$

$$f_{sus}^{(x)}(V) = \frac{1}{a_{freq}^{(x)} \cdot V^{b_{freq}^{(x)}} + c_{freq}^{(x)}}, \tag{2}$$

where $a_{freq}^{(x)}$, $b_{freq}^{(x)}$, $c_{freq}^{(x)}$ are constant fitting parameters.

Leakage and peak dynamic power when operating at $f_{sus}^{(x)}(V)$ are:

$$P_{leak}^{(x)}(V) = a_{leak}^{(x)} \cdot V^{b_{leak}^{(x)}} \tag{3}$$

$$P_{dyn,peak}^{(x)}(V) = a_{dyn}^{(x)} \cdot V^{b_{dyn}^{(x)}} + c_{dyn}^{(x)} \tag{4}$$

Here, $a_{dyn}^{(x)}$, $b_{dyn}^{(x)}$, $c_{dyn}^{(x)}$, $a_{leak}^{(x)}$, $b_{leak}^{(x)}$ are constant fitting parameters.

When operating at a frequency lower than $f_{sus}^{(x)}$ and the corresponding V, dynamic power is scaled linearly, and hence:

$$P_{dyn,max}^{(x)}(V,f) = \frac{f}{f_{sus}^{(x)}(V)} \cdot P_{dyn,peak}^{(x)}(V) \tag{5}$$

4) Workload-Dependence of Power and Energy: Dynamic power consumption $P_{dyn}^{(x)}(V,f)$ is affected by the running workload, which induces some switching activity on the processor. The dynamic power consumption is scaled by a factor $r_{dyn} \geq 0$ from the peak dynamic power $P_{dyn,max}^{(x)}(V,f)$:

$$P_{dyn}^{(x)}(V,f) = r_{dyn} \cdot P_{dyn,max}^{(x)}(V,f) \tag{6}$$

The total power consumption $P_{total}^{(x)}(V,f)$ is the sum of dynamic and leakage power:

$$P_{total}^{(x)}(V,f) = P_{dyn}^{(x)}(V,f) + P_{leak}^{(x)}(V) \tag{7}$$

r_{dyn} is not constant since it represents the current workload activity that depends on the dynamic/total power ratio as a variable. We define the dynamic/total power ratio as the r_{dyn} observed at the highest thickness $\hat{x}$=4 nm, the highest common frequency $\hat{f}$ and the minimum required voltage $\hat{V} = V_{min}^{(4)}(\hat{f})$:

$$\begin{aligned} dyn/tot &= \frac{P_{dyn}^{(\hat{x})}(\hat{V},\hat{f})}{P_{total}^{(\hat{x})}(\hat{V},\hat{f})} \\ &= \frac{r_{dyn} \cdot P_{dyn,max}^{(\hat{x})}(\hat{V},\hat{f})}{r_{dyn} \cdot P_{dyn,max}^{(\hat{x})}(\hat{V},\hat{f}) + P_{leak}^{(\hat{x})}(\hat{V})} \end{aligned} \tag{8}$$

r_{dyn} can be calculated from a given dyn/tot as follows:

$$r_{dyn} = \frac{dyn/tot \cdot P_{leak}^{(\hat{x})}(\hat{V})}{P_{dyn,max}^{(\hat{x})}(\hat{V},\hat{f}) \cdot (1 - dyn/tot)} \tag{9}$$

(a) Average power under performance constraint: The average total power consumption under a performance constraint f_{min} (i.e., Optimization Scenario (a)) $P_{avg}^{(x)}(V,f)$ is the sum of power in active and idle reigns, where V_{leak} is selected at design time as described earlier:

$$\begin{aligned} P_{avg}^{(x)}(V,f) = (P_{dyn}^{(x)}(V,f) + P_{leak}^{(x)}(V)) \cdot \frac{W}{f} \\ + P_{leak}^{(x)}(V_{leak}) \cdot (T - \frac{W}{f}) \end{aligned} \tag{10}$$

(b) Total energy: Total energy for Scenario (b) is:

$$E_{total}^{(x)}(V,f) = (P_{dyn}^{(x)}(V,f) + P_{leak}^{(x)}(V)) \cdot \frac{W}{f} \tag{11}$$

5) Optimal Frequency and Voltage Selection: V_{opt} and f_{opt} that minimize total power can be obtained from the power model in the form of a minimization problem:

$$V_{opt}(f, r_{dyn}) = \underset{V_{min}^{(x)}(f) \leq V \leq V_{max}^{(x)}}{\arg\min} P_{avg}^{(x)}(V,f) \tag{12}$$

$$f_{opt}(r_{dyn}) = \underset{f_{min}^{(x)} \leq f \leq f_{max}^{(x)}}{\arg\min} P_{avg}^{(x)}(V_{opt}(f, r_{dyn}), f) \tag{13}$$

V_{opt} and f_{opt} that minimize total energy can be similarly obtained from the energy model:

$$V_{opt}(f, r_{dyn}) = \underset{V_{min}^{(x)}(f) \leq V \leq V_{max}^{(x)}}{\arg\min} E_{total}^{(x)}(V,f) \tag{14}$$

$$f_{opt}(r_{dyn}) = \underset{f_{min}^{(x)} \leq f \leq f_{max}^{(x)}}{\arg\min} E_{total}^{(x)}(V_{opt}(f, r_{dyn}), f) \tag{15}$$

Power/energy management (i.e., DVFS selection) is, therefore, an optimization problem that can be solved by exploring the design space of $P_{total}^{(x)}(V,f)$ or $E_{total}^{(x)}(V,f)$ over all possible voltages and frequencies.

V. Exploration and Optimization

In the following, we present our NCFET-aware power and energy management technique for power/energy minimization. We then perform an analytical design space exploration to determine the impact of FE layer thickness on optimal power and energy as a function of workload parameters.

1) Frequency and Voltage Selection: Power/energy management techniques (e.g., DVFS) operate under some constraints, e.g., to maximize performance given a power/energy budget or

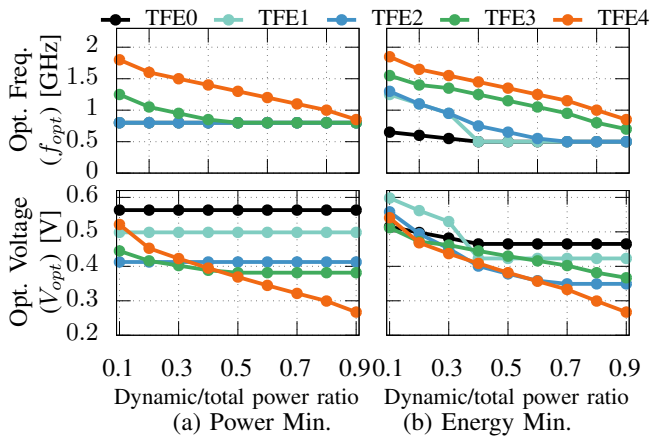

Fig. 4: Optimal frequency and voltage selected by our technique over dynamic/total ratio for thicknesses TFEx targeting (a) power minimization using W=10^6 and (b) energy minimization with f_d=0.8GHz, W=10^6 and T=20ms.

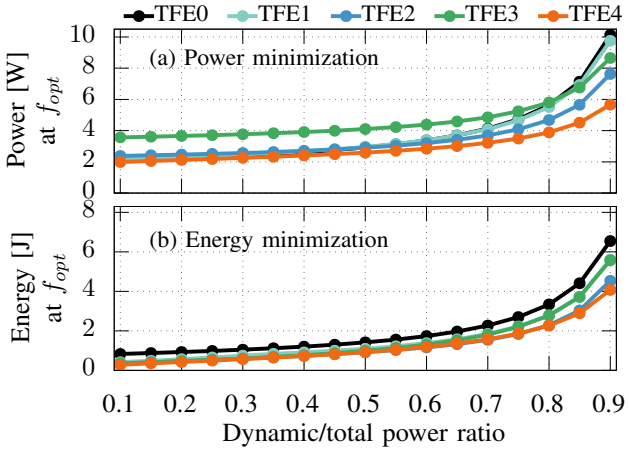

Fig. 5: Optimal power and energy over dynamic/total power ratios for different TFEx operating at optimal frequency f_{opt} using W=10^6 and T=$10W$. (a) Power minimization. TFE4 has the minimum power all the time. (b) Energy minimization. TFE4 has the minimum energy most of the time. For a small region 0.6-0.75, TFE2 has slightly better energy.

to minimize power/energy given a performance goal. The constraints are fulfilled by selecting the proper V/f pairs. V_{opt}/f_{opt} selection following Eq. (13) and Eq. (15) is an optimization problem that can be solved using a search algorithm. Our algorithm to perform the required optimization is summarized in Algorithm 1. It performs a search by sweeping across all possible frequency and voltage steps. Since the number of discrete frequency/voltage settings is limited by the hardware and closed-form analytical models can be evaluated quickly, the search is fast. In addition, it can be applied either online (i.e., executing Algorithm 1 at runtime) or offline (i.e., pre-characterizing a processor at design time and selecting predefined V/f pairs at runtime). The offline technique works by characterizing the processor operating point pairs V_{opt}/f_{opt} as a function of workload ratios and performance at design time.

Algorithm 1 NCFET-aware power/energy management to select the optimal frequency (f_{opt}) and the corresponding optimal voltage (V_{opt}) that minimizes power or energy.

Input: Power models for TFEx: $P^{(x)}_{total}$, Energy models for TFEx: $E^{(x)}_{total}$, voltage range $[V_{min}, V_{max}]$, voltage step ϵ, frequency array (f_r), dyn/total power ratio, Work W, Time T, Deadline frequency f_d

Output: Optimal frequency f_{opt} at optimal voltage V_{opt}

$f_{opt} \leftarrow f$
$rdyn$ at given $dyn/total$ ratio ▷ Eq. (9)
for each f in f_r **do**
  $V_{opt1} \leftarrow V^{(x)}_{min}(f)$ ▷ Eq. (1)
  repeat
    $energy = E^{(x)}_{total}$ for W ▷ Eq. (11)
    $power = P^{(x)}_{avg}$ for W, T, $f \geq f_d$ ▷ Eq. (10)
    if $Min.>$ {energy OR power} **then** ▷ Opt. goal
      $f_{opt} \leftarrow f$, $V_{opt} \leftarrow V_{opt1}$
    $V_{opt1} \leftarrow V_{opt1} + \epsilon$ ▷ iterative update
  until $V_{opt1} = V_{max}$ ▷ Termination criteria
return f_{opt}, V_{opt}

Given actual measured or derived workload characteristics and performance goals in the current power/energy management epoch, operating points V_{opt}/f_{opt} can then be selected by the operating system or hardware at runtime.

We further examine how the optimal frequency and the corresponding optimal voltage that minimize power or energy using our technique depends on possible workload characteristics. To cover a wide range of workloads, we examine dynamic/total power ratios in the range of 0.1-0.9 for W=10^6 cycles, T=20ms (time to finish $10W = 10^7$ cycles at f_{min}=0.5GHz), and a performance constraint of f_d=0.8GHz to meet T. The studied range is similar to what has been presented in [9] for many real applications. Our experimenters (see Section VII) also show that the ratio for evaluated applications does not exceed 0.7 at TFE4. Please note, the dynamic to total power ratios are characterized at TFE4 and $\hat{f} = 1.2$ GHz. As TFE4 reduces dynamic power and increases leakage power at a given frequency [5], the equivalent ratios for the same workload are expected to be larger at TFE0, i.e. in traditional FinFET. Solving Equations 9 and 8, the equivalent range of 0.1-0.9 at TFE4 is 0.19-0.95 at TFE0.

The optimal frequencies and voltages for power and energy minimization cases are shown in Fig. 4(a) and (b), respectively. For power minimization, frequencies and voltages are shown for the active phase. The voltage is set to V_{leak} in the idle phase (see Section IV(2)). The figure shows that different frequencies and voltages are selected based on the optimization case (i.e., power or energy). Moreover, for all cases, TFE4 exhibits the best performance over all thicknesses.

2) Thickness Exploration: In the following, we show the impact of FE layer thickness on the minimum power and energy of the processor. Using the optimal V/f pairs shown in Fig. 4, we examine the resulting power and energy over the dynamic/total power ratios.

Minimum power and energy results for different thicknesses over dynamic/total power ratios are shown in Fig. 5. In

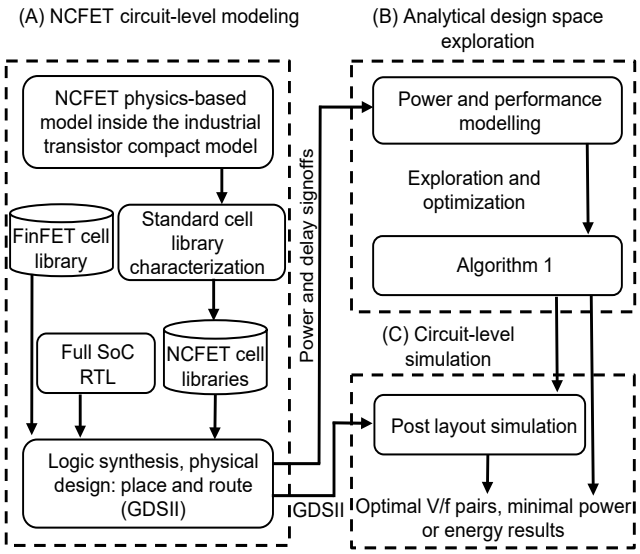

Fig. 6: Methodology for NCFET-based processor modeling besides analytical and experimental NCFET-aware power/energy management evaluation. Circuit-level modeling (A) is a cross-layer implementation that links physics- and circuit-level aspects that are involved in the design of an NCFET-based processor. Afterwards, in (B), the NCFET-aware power/energy optimization is applied to select the optimal V/f pairs that minimize power or energy. Finally, in (c), resulting power and energy savings are validated through circuit-level simulations.

Fig. 5(a), for the power minimization case, the minimal power is at TFE4 for all workloads. TFE3, on the other hand, has the highest power for most of the examined workload ratios. This is due to its highest leakage power when operating at V_{leak} compared to all other thicknesses (see Fig. 1). By contrast, the minimal energy (Fig. 5(b)) is again at TFE4 for most of the examined workloads. TFE2 and TFE4 are very close, where TFE2 is optimal with small improvements over TFE4 for a small region of dynamic/total power ratios between 0.6-0.75. However, TFE4 achieves by far the best performance (i.e., higher f_{opt}, see Fig. 4(a)). As a result, in both scenarios, *TFE*4 shows the overall optimal power/energy and performance tradeoff among all FE layer thicknesses. The preference, therefore, is for *TFE*4 as it minimizes energy and power at high performance.

VI. Validation Methodology

Fig. 6 summarizes our experimental phases and the evaluation methodology that was used to develop and examine our NCFET-aware power and energy management technique. We validate our NCFET-aware power/energy management technique both *analytically and experimentally*. Our methodology, in Fig. 6, consists of three parts described in the following.

A. NCFET Circuit-Level Modeling

As NCFET is an emerging technology and still in its infancy, a real processor chip has not been fabricated yet. Therefore, processor modeling is performed through circuit-level simulation based on mature physics-based NCFET transistor models.

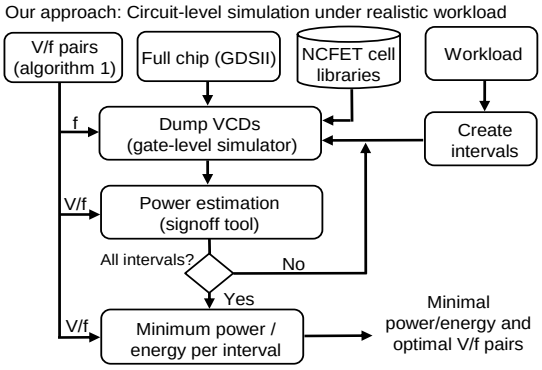

Fig. 7: Methodology for simulation of our NCFET-aware power/energy management technique at circuit-level (i.e., gate-level). Since we do not know the dynamic/total power ratio a priori, the process works by dividing the execution of a workload into intervals and examining the power/energy of each interval for all V_{opt}/f_{opt} pairs obtained analytically.

A physics-based NCFET model [16] was integrated within industrial compact models of FinFETs (BSIM-CMG) [17]. Then, NCFET-aware cell libraries were characterized as in [6] based on the open-source 7nm FinFET PDK [18]. Our libraries are fully compatible with existing EDA tool flows. With such compatibility, we can directly deploy them to perform a full chip design starting from logic synthesis all the way to implementation of a full layout. Our cell libraries cover a wide range of voltages (0.2 to 0.7V) in order to accurately estimate how the power and performance of the processor are affected by NCFET as a function of voltage. The analysis is applied at room temperature of $25°C$. The thickness of the employed FE layer is limited to 4nm to ensure hysteresis-free operation [6].

In this work, we focus on studying the impact of NCFET on power, energy and performance of a CPU core itself. We consider a full tile of the state-of-the-art OpenPiton SoC, which consists of a core, an L1.5 cache, an L2 cache, a floating-point unit (FPU), a CPU-Cache Crossbar (CCX) arbiter, a Memory Inter-arrival Time Traffic Shaper (MITTS), and three P-Mesh NoC routers [19]. We have configured the OpenPiton as summarized in Table I. We modified the default configuration to have relatively small caches. This is important for our study to avoid caches from biasing results as explained below. Within the tile itself, we consider the impact of all memory modules (e.g., caches) and internal interconnects (e.g., inter-chip network routers) for power and energy evaluation. We do not consider the memory and interconnect outside of the tile.

To reveal the trends of power and performance of a processor at different FE layers, we first synthesized the tile's circuit-level description into a gate-level implementation using standard logic synthesis with NCFET cell libraries. The gate-

TABLE I: OpenPiton configuration.

Component	Configuration	Component	Configuration
Cores	1	L1 I-cache	4KB
Threads	1	L1 D-cache	4KB
FPU	enabled	L1.5 & 2 caches	8KB, 64KB
TLBs	16	Reg. file WB	2w, 3r

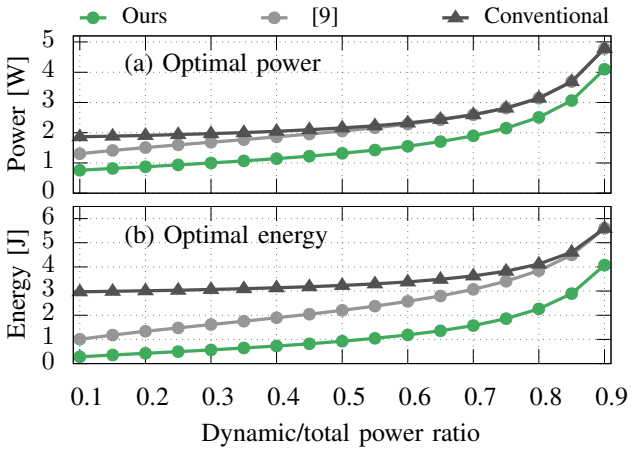

Fig. 8: Analytical results of the optimal power and energy consumption with TFE4 over the dynamic/total power ratio for the three used scenarios targeting (a) power and (b) energy minimization. Our scenario, operating at V_{opt}/f_{opt} shows the minimal power and energy among all scenarios. The conventional technique, where V_{min}/f_{min} is selected, is the worst scenario. State-of-the-art [9], where V_{opt}/f_{min} is selected to achieve a trade-off between leakage and dynamic power, is sub-optimal.

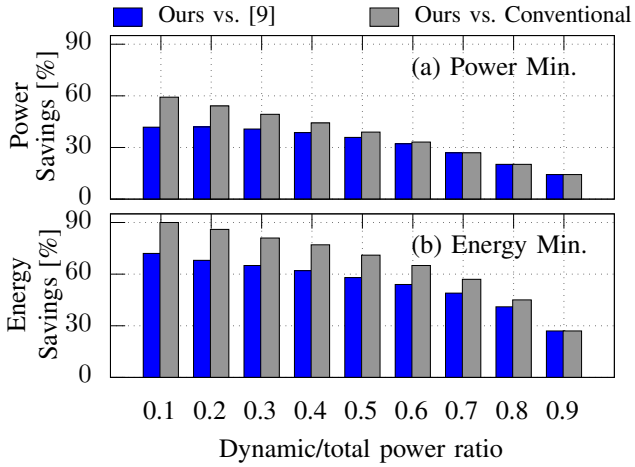

Fig. 9: (a) Power savings and (b) Energy savings by operating at V_{opt}/f_{opt} selected by our scenario in comparison with state-of-the-art [9] and conventional technique. Savings for power minimization are up to 60 % and up to 41 % compared to conventional and state-of-the-art [9], respectively. By contrast, savings for energy minimization are up to 90 % and 72 %.

level design is then further realized using physical design processes (i.e., full circuit layout) as the standard method to examine the final chip before fabrication. Our physical implementation of memory structures such as caches is logic-based as neither SRAM modules nor memory compilers are available for NCFET yet. Therefore, a flip-flop (FF) based SRAM implementation is used in both conventional FinFET (TFE0) and NCFET (TFE1-4) to enable fair comparisons. We selected our configuration (Table I) to ensure that caches do not dominate the power, area or delay in our implementation. With our configuration, the power contribution of caches is on average $< 21\%$ of the overall OpenPiton's power. This in line with reported results in [19] [20].

Finally, using signoff tools, we apply static time analysis (STA) to obtain the power and frequency of the processor across the whole voltage range and for all FE thicknesses. Resulting power and frequency are shown in Fig. 1

B. Analytical evaluation

We first perform an analytical validation based on optimal configurations to evaluate the efficiency of our technique in comparison with other techniques. This requires a design space exploration to expose optimal configurations (i.e., V/f pairs for every FE layer thickness) that minimize power/energy. This exploration covers all possible operating corners under given constraints using synthetic applications.

Based on the power and frequency characterization of the processor in Section VI-A (see Fig. 1), we perform fitting and calibrations of the power and frequency models described in Section IV(3). Afterwards, we apply Algorithm 1 under given runtime constraints (e.g., T and W) to determine the optimal configurations for all possible workloads and FE thicknesses.

For fair comparisons, we also implement state-of-the-art power and energy management techniques to analytically compare our technique against their power and energy consumption. Finally, we report power and energy savings that our technique achieves compared to other techniques.

C. Circuit-level Simulations for Realistic Workloads

The synthetic workloads that are used for the analytical evaluation (see Section IV and Section VI-B) are assumed to have certain properties that may differ from real applications. These properties were selected to allow a comprehensive analytical exploration. Real workloads may show execution phases with different power at runtime. Therefore, to minimize power/energy of such a workload, it will be required to operate at different frequencies and voltages at runtime. Furthermore, the performance might not scale linearly with frequency as assumed in the analytical exploration. Therefore, V/f pair selection should follow dynamic workload characteristics. To consider the dynamic characteristics of workloads, we evaluate our technique for real applications using circuit-level simulation. However, as power/energy management techniques, such as DVFS, are complex algorithms, it is still infeasible to simulate the full technique at the circuit-level. This is due to the large simulation complexity and computational cost. Therefore, we apply our technique manually by simulating runtime selections.

In practice, power/energy management always consist of two parts: (1) predicting future workload characteristics, and (2) selecting (V/f) pairs based on the prediction. To make our simulation independent of any DVFS algorithm and to provide an upper bound on the possible gains, we employ an oracle approach that assumes workload characteristics are known in advance. This allows us to focus on reporting the benefits from V/f selection as studied in this paper in isolation while ignoring

the impact of potentially wrong decisions made by heuristic or predictive techniques.

Note that our NCFET-aware power and energy management can be employed for both pre-characterized and unknown applications. As the optimal V/f pair selection depends on the dynamic power consumption of the workload, it needs to be known or estimated when determining the V/f pair for the next execution phase. In case of a pre-characterized workload, this is trivial. In case of unknown applications, our V/f selection approach needs to be combined with a heuristic or predictive technique for workload characterization.

Fig. 7 summarizes the circuit-level (i.e., gate-level) simulation that is used in this work to evaluate V/f pair selection at runtime under realistic workloads. This simulation is complex, computationally expensive and time consuming. However, it is still important to provide an accurate evaluation that accurately considers the actual workload characteristics. We mimic the functionality of the power/energy management technique using the optimal configurations obtained from the analytical exploration (i.e., V_{opt}/f_{opt} pairs). We first split the execution time into equal intervals. Next, we estimate the power and energy of every V_{opt}/f_{opt} pair for each interval. This process is a gate-level (i.e., post-layout) simulation, where VCD files are generated for every f_{opt}. Then, power is estimated using the power signoff tool for every interval by operating the circuit at V_{opt}. Finally, powers/energies are compared for similar intervals to find V_{opt}/f_{opt} with the minimal power/energy.

VII. Experimental Results

Here, we evaluate the efficacy of our technique. We present the achievable power and energy savings using our NCFET-aware power and energy management in comparison with conventional techniques and state-of-the-art. Comparisons are categorized into analytical and simulation results.

As shown previously, TFE4 shows the minimum power and energy over all thicknesses at f_{opt}. TFE4 also shows the highest frequency over all thicknesses (i.e., best performance). Therefore, we examine our technique only for TFE4. We examine the power and energy for three different scenarios: (1) *NCFET-aware voltage and frequency selection (ours)*: the processor operates at f_{opt} and $V_{opt}(f_{opt})$. (2) NCFET-aware voltage selection (state of the art) following [9]: the processor operates at the minimum frequency f_{min} that is required to meet the performance goal and the optimal voltage $V_{opt}(f_{min})$ following Eq. (12) and Eq. (14). The comparison of this scenario to the first scenario shows the impact of NCFET-aware frequency selection. (3) Conventional (NCFET-unaware) technique where the processor operates at minimum frequency f_{min} required to meet a performance goal and the minimum voltage V_{min} required to sustain that frequency. All scenarios are summarized in Table II.

We use Synopsys Design Compiler to synthesize the OpenPiton tile. Using Cadence Innovus, we further design the floorplan and power delivery network (PDN), and we perform place and route including clock tree synthesis. We employ Cadence Tempus Timing as signoff tool for delay analysis and Voltus IC Power Integrity as signoff tool for power analysis. The synthesized, placed and routed gate-level netlist is simulated using Synopsys VCS together with the OpenPiton emulator, which provides a testbench that emulates full system functionality of all off-chip components including I/O, DRAM controllers, chip bridge, P-Mesh chipset crossbar, and P-Mesh inter-chip network routers.

A. Analytical results

We implemented the scenarios in Table II and ran the same configurations with each scenario for both power and energy cases. We use the same set of parameters as in Section V, i.e., $W = 10^6$ cycles, $T = 20$ms. We then sweep the ratio of dynamic to total power from 0.1 to 0.9 to test all techniques. For each scenario, we explore the whole design-space under all the possible configurations, searching for the minimal power/energy.

The results of the three scenarios are demonstrated in Fig. 8(a) for power minimization and Fig. 8(b) for energy minimization. Results, in both cases, show that our scenario achieves the minimal power and energy in all cases.

Power and energy savings of the analytical results are summarized in Fig. 9. The conventional scenario shows the highest power and energy consumption among all cases for all dynamic/total power ratios as it is completely NCFET-unaware. This highlights, again, that existing power management techniques cannot be used for future NCFET-based processors. Instead, new NCFET-aware techniques need to be developed, which we present in this work. The savings using our approach compared to a conventional technique reach up to 60% in power, and up to 90% in energy. The state-of-the-art scenario achieves lower power/energy than a conventional technique, as the state-of-the-art is NCFET-aware albeit for voltage selection only. Our technique, which also optimizes the frequency, outperforms the state-of-the-art by up to 41% and 72% for power and energy minimization cases, respectively.

B. Simulation results

The analytical results are an upper-bound on the possible power and energy gains under synthetic workloads. Actual gains with real workloads might be different. Therefore, we further examine realistic workloads under the same previously-mentioned three scenarios for fair comparisons.

We evaluate our NCFET-aware power/energy management at the circuit-level as described in Section VI-C following the runtime scenarios that describe the different optimization cases (see Section IV(2)). We select three representative micro-benchmarks from [21] for simulations: Matrix-Matrix multiplication (MM), Memory Test (MT), and Fast Fourier Transform (FFT). Because we perform gate-level simulations, which are computationally very expensive, we are forced to select short-running benchmarks. We load the application directly into the DRAM emulator within the OpenPiton platform.

TABLE II: Scenarios for comparison in the evaluation.

Scenario	Frequency	Voltage
(1) NCFET-aware voltage and frequency selection (ours)	f_{opt}	V_{opt}
(2) NCFET-aware voltage selection (state-of-the-art) [9]	f_{min}	V_{opt}
(3) Conventional technique (NCFET-unaware)	f_{min}	V_{min}

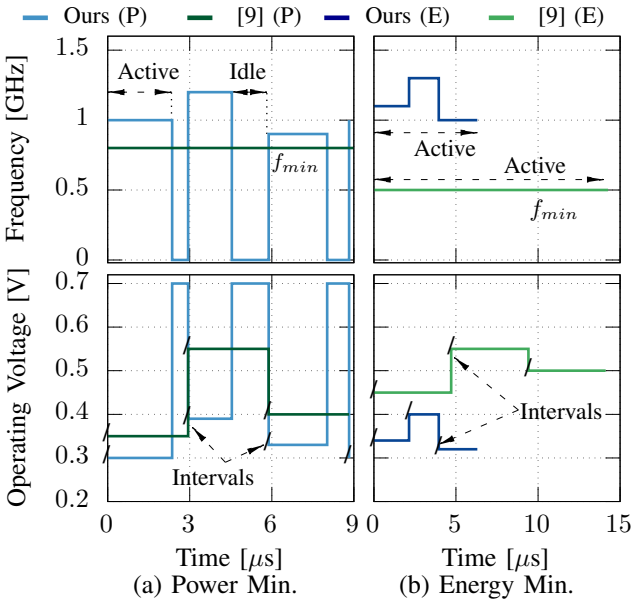

Fig. 10: Samples of three intervals while running FFT benchmark showing the selected frequencies and voltages by our technique and state-of-the-art [9] for (a) power minimization (P) and (b) energy minimization (E). State-of-the-art selects f_{min} for both cases. f_{min}=f_d=0.8 GHz and f_{min}=0.5 GHz for power and energy minimization, respectively. By contrast, our technique always selects higher frequencies as shown in the top row for (a) and (b). For power minimization, our technique switches between idle and active regions. *Different V_{opt} are selected for the same interval based on the optimization goal*. For the power minimization case, two different voltages are selected in the active and idle region. The voltage is set to V_{leak}=0.7 V to reduce leakage power in the idle region.

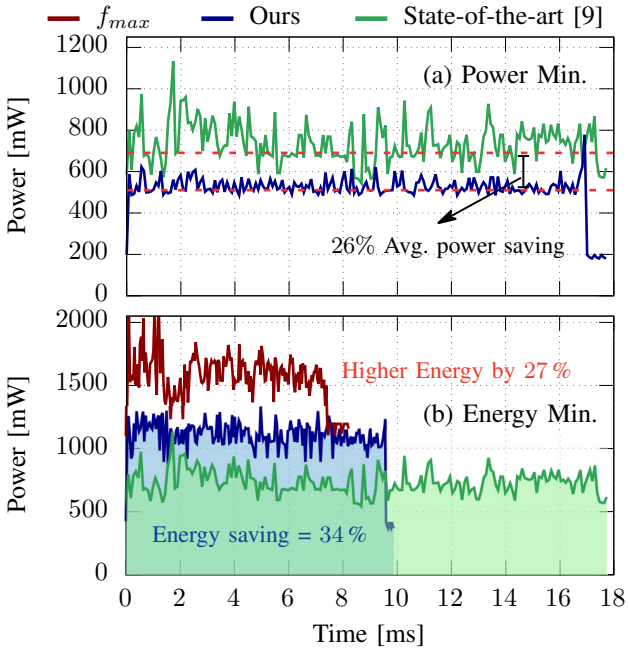

Fig. 11: (a) Power traces of FFT benchmark for power minimization case using our technique and state-of-the-art. Both scenarios finish at the same time. Our technique results in lower average power. The average power is reduced by 26% in this example. (b) Power traces of FFT benchmark for energy minimization case using our technique and state-of-the-art. In our technique, the processor rushes to finish the execution early and shows 1.8x better performance in addition to lower energy (i.e., highlighted area under the curve). Energy savings in this example are up to 34%. Additionally, we compare our technique against a simple rush-to-completion strategy at f_{max} and V_{opt} that provides fastest execution but has 27% higher energy compared to our technique.

All workloads are otherwise cache contained during gate-level simulation. The average dynamic/total power ratios for FFT, MM, and MT benchmarks at TFE4, $\hat{f}$ = 1.2GHz and nominal voltage $\hat{V} = V_{min}(\hat{f})$ (matching the definition of dynamic/total power ratios used throughout the paper) are 0.49, 0.61, and 0.68, respectively. Note that, in practice, the examined workload ratios will depend on techniques that are applied to reduce leakage power/energy (e.g., low-leakage power gating approaches for caches). However, even with additional circuit optimizations, the ratios are well within the range of possible ratios studied throughout the paper.

We split workload executions into intervals to select V/f pair per interval. Note that optimal frequency and voltage selection does not depend on the interval size, but is only a function of average workload behavior in each interval. Due to limitations on the length of gate-level simulations, we select the interval size to be 20−30 instructions. In reality, a larger interval would be selected based on trade-offs between fast reaction time vs. decision-making (i.e., estimating or predicting f/V pair) and frequency/voltage switching overheads.

Assuming an oracle approach that can predict the operating point before the interval start, and using V_{opt}/f_{opt} pairs that resulted from the analytical optimization, we estimate the power and energy for every interval on all pairs. After that, we compare the power/energy consumption of all similar intervals to select V_{opt}/f_{opt} pairs that optimally minimize power/energy. We apply the same previously used constraints, $f_{min} = f_d =$ 0.8 GHz for power minimization and f_{min} = 0.5 GHz for energy minimization.

To highlight the fundamental differences for runtime frequency and voltage selection, and for visualization, we plot samples of frequencies and voltages for three intervals when running the FFT benchmark. Fig. 10(a) and (b) show the selected frequencies and voltages by our technique in comparison with state-of-the-art [9] for power and energy minimization cases, respectively.

As shown for frequency selection, the state-of-the-art technique uses f_{min} all the time. For power minimization, state-of-the-art selects $f_{min} = f_d$ = 0.8 GHz, while for energy minimization it selects f_{min} = 0.5 GHz. Importantly, our technique always selects higher frequencies than f_{min}. Frequencies, as expected, vary with intervals as the activity varies, which causes a varying dynamic power consumption. However, average power minimization results in lower frequencies to be selected compared to the energy minimization case. This is due to the idle phase, where the processor has to wait for the deadline T (i.e., leaks more) and, therefore, a slower execution

is more beneficial to balance between powers and times of the active and idle regions.

As shown for voltage selection, for the two optimization cases, different V_{opt} are selected by our scenario for each interval. For the power minimization case, the voltage in the active region is relatively lower than the voltage in the active region of the energy minimization case as a lower frequency is in use. Moreover, in the idle region, the voltage is set to $V_{leak} = 0.7\,\text{V}$ to reduce leakage power (see Section IV). The state-of-the-art selects higher voltages in both cases. Importantly, *the figure shows that different V_{opt}/f_{opt} are selected for the same interval based on the optimization goal.*

The total runtime power/energy is the combinations of all intervals operating at the optimal V/f pairs. Fig. 11(a) shows the power trace over time when running the FFT benchmark at the optimal V/f pairs for the power minimization case using our scenario and state-of-the-art. As shown, our scenario results in lower average power compared to state-of-the-art. The power savings for this example, are on average 26% compared to state-of-the-art. Fig. 11(b) shows the runtime power consumption for the energy minimization case when running the FFT benchmark at the optimal V/f pairs. As shown, in our scenario the processor rushes to finish earlier. Hence, the power consumption is higher than in the state-of-the-art scenario, but with shorter execution time. The overall energy (i.e., area under the curve) is, however, reduced. This is because the reduction in execution time is larger than the increases in power (i.e., $E = P{\cdot}D$). The energy savings for this example are up to 34% compared to state-of-the-art. Moreover, performance is improved as $f_{opt}{>}f_{min}$. In this example, performance is enhanced by 1.8x. This again still comes with lower energy, as no energy-performance trade-off has to be made.

In addition, we compare energy consumption against a simple rush-to-completion strategy that runs at maximum frequency f_{max} and then goes to idle [22]. We apply f_{max} at the optimal voltage V_{opt} (similar to the state-of-the-art for fair comparison). As shown in Fig. 11(b), such an approach results in 27% higher energy compared to our technique. While operating at f_{max} achieves the fastest execution time, the associated increase in dynamic energy at higher voltages outweighs the gains from a larger idle time in TFE4.

The power and energy savings using our scenario in comparison with the state-of-the-art [9] and a conventional technique are summarized in Fig. 12. The average power savings are up to 32% compared to state-of-the-art scenario, and up to 46% compared to the conventional scenario. The energy savings are up to 42% compared to state-of-the-art, and up to 58% compared to a conventional scenario. Moreover, the performance improvements can reach up to 2.1x for the MT.

As a reference, using optimal V/f pairs selected by our technique for the energy minimzation case, the actual dynamic to total power ratios of FFT, MM, and MT benchmarks are reduced to on average 0.44, 0.53, and 0.64, respectively.

Again, as shown before, the state-of-the-art scenario results in higher savings than the conventional technique, as state-of-the-art is NCFET-aware, albeit for voltage selection only. Crucially, however, our results show that, depending on the workload, minimal power/energy is achieved at a higher frequency than the performance constraint would require. In other words, even optimal power management may necessitate more complex frequency optimization than investigated in [9].

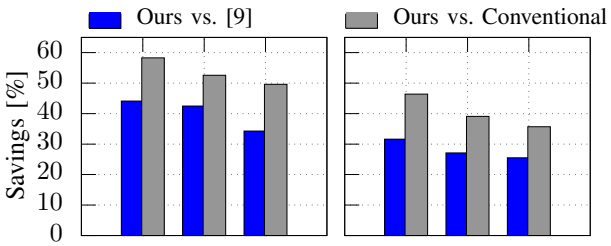

(a) Energy savings (a) Power savings

Fig. 12: Circuit-level simulation (a) energy savings, (b) power savings with our technique in comparison to state-of-the-art and conventional techniques using three benchmarks. Savings are up to 58% and up to 49% compared to conventional and state-of-the-art [9] scenarios, respectively, for the energy minimization case. Fore power minimization case, savings are up to 46% and 32% compared to conventional and state-of-the-art [9] scenarios, respectively

VIII. Conclusions

NCFET is a promising emerging technology that provides outstanding performance in addition to better power optimization compared to conventional FinFET technology. As conventional power and energy minimization techniques are unaware of the inverse dependency that leakage power exhibits in NCFET, they can not be used with NCFET-based processors. In this work, we presented the first NCFET-aware management technique to optimize the power/energy of NCFET-based processors. We introduced an approach that selects the optimal frequency and voltage offline or at runtime. Our analysis further demonstrated an analytical design space exploration in addition to circuit-level simulation for selecting the optimal operating frequency f_{opt} and voltage V_{opt} based on thickness and application characteristics. The optimal frequency to achieve minimal power/energy is larger than the minimum frequency. The largest ferroelectric thickness provides both the best power/energy and performance. Finally, applying our technique at circuit-level to characterize a real workload behavior, our technique results in up to 58% energy saving compared to conventional NCFET-unaware and up to 49% compared to state-of-the-art NCFET-aware management techniques. Moreover, our technique results in up to 46% average power saving compared to conventional and up to 32% compared to state-of-the-art techniques.

References

[1] J. Lee, Y. Zhang, Q. Dong *et al.*, "19.2 a 6.4pj/cycle self-tuning cortex-m0 iot processor based on leakage-ratio measurement for energy-optimal operation across wide-range pvt variation," in *2019 IEEE International Solid- State Circuits Conference - (ISSCC)*, Feb 2019.

[2] S. Salamin, H. Amrouch, and J. Henkel, "Selecting the optimal energy point in near-threshold computing," in *2019 Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2019, pp. 1691–1696.

[3] S. Goswami, B. Chowdhury, and M. Chanda, "Analytical modelling of power dissipation and voltage swing of cmos logic circuit for near-threshold computing," in *Devices for Integrated Circuit (DevIC)*, 2017.

[4] P. Meinerzhagen, C. Tokunaga, A. Malavasi *et al.*, "An energy-efficient graphics processor featuring fine-grain dvfs with integrated voltage regulators, execution-unit turbo, and retentive sleep in 14nm tri-gate cmos," in *Solid - State Circuits Conference - (ISSCC)*, Feb 2018.

[5] M. Rapp, S. Salamin, H. Amrouch *et al.*, "Performance, Power and Cooling Trade-Offs with NCFET-based Many-Cores," *Design Automation Conference (DAC)*, 2019.

[6] H. Amrouch, G. Pahwa, A. D. Gaidhane *et al.*, "Negative Capacitance Transistor to Address the Fundamental Limitations in Technology Scaling: Processor Performance," *IEEE Access*, vol. 6, 2018.

[7] S. K. Samal, S. Khandelwal, A. I. Khan *et al.*, "Full chip power benefits with negative capacitance FETs," in *2017 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*, July 2017.

[8] G. Pahwa, T. Dutta, A. Agarwal *et al.*, "Designing energy efficient and hysteresis free negative capacitance FinFET with negative DIBL and 3.5X IONusing compact modeling approach," in *2016 46th European Solid-State Device Research Conference (ESSDERC)*, Sep. 2016, DOI:10.1109/ESSDERC.2016.7599584.

[9] S. Salamin, M. Rapp, H. Amrouch *et al.*, "NCFET-Aware Voltage Scaling," *The International Symposium on Low Power Electronics and Design (ISLPED)*, 2019.

[10] S. Salamin and M. Rapp and H. Amrouch and A. Gerstlauer and J. Henkel, "Energy optimization in ncfet-based processors," in *Design, Automation Test in Europe Conference Exhibition (DATE)*, 2020.

[11] C. Bienia, S. Kumar, J. P. Singh *et al.*, "The PARSEC Benchmark Suite: Characterization and Architectural Implications," in *Parallel Architectures and Compilation Techniques (PACT)*, 2008.

[12] E. Calore, A. Gabbana, S. Schifano *et al.*, "Software and dvfs tuning for performance and energy-efficiency on intel knl processors," *Journal of Low Power Electronics and Applications*, vol. 8, p. 18, 06 2018.

[13] L. Mo, A. Kritikakou, and O. Sentieys, "Energy-quality-time optimized task mapping on dvfs-enabled multicores," vol. 37, no. 11, Nov 2018.

[14] S. Dighe, S. R. Vangal, P. Aseron *et al.*, "Within-die variation-aware dynamic-voltage-frequency-scaling with optimal core allocation and thread hopping for the 80-core teraflops processor," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 1, pp. 184–193, Jan 2011.

[15] H. Kaul, M. A. Anders, S. K. Mathew *et al.*, "A 320 mv 56 μw 411 gops/watt ultra-low voltage motion estimation accelerator in 65 nm cmos," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 1, 2009.

[16] G. Pahwa, T. Dutta, A. Agarwal *et al.*, "Analysis and Compact Modeling of Negative Capacitance Transistor with High ON-Current and Negative Output Differential Resistance—Part II: Model Validation," *Transactions on Electron Devices (TED)*, vol. 63, no. 12, 2016.

[17] "BSIM-CMG Technical Manual," October 2019, http://www-device.eecs.berkeley.edu/bsim/?page=BSIMCMG.

[18] L. T. Clark, V. Vashishtha, L. Shifren *et al.*, "ASAP7: A 7-nm FinFET predictive process design kit," *Microelectronics Journal*, vol. 53, 2016.

[19] J. Balkind, M. McKeown, Y. Fu *et al.*, "OpenPiton: An Open Source Manycore Research Framework," in *Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, ser. ASPLOS, 2016.

[20] M. McKeown, A. Lavrov, M. Shahrad *et al.*, "Power and energy characterization of an open source 25-core manycore processor," in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2018, pp. 762–775.

[21] "OpenSPARC T1," February 2019. [Online]. Available: https://www.oracle.com/technetwork/systems/opensparc/opensparc-t1-page-1444609.html

[22] J. Haj-Yahya, A. Mendelson, Y. Ben-asher *et al.*, *Energy Efficient High Performance Processors Recent Approaches for Designing Green High Performance Computing*, 2018, ISBN 978-981-10-8554-3.

Sami Salamin received his B.Sc. degree in computer systems engineering and M.Sc. degree (first rank) from Palestine Polytechnic University, Hebron, Palestine in 2005 and 2012, respectively. Since 2016, he is pursuing his Ph.D. degree with the Chair of Embedded Systems (CES), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany, under the supervision of Prof. Jörg Henkel. He holds HiPEAC Paper Award. He is a student member of IEEE. ORCID 0000-0002-1044-7231

Martin Rapp joined the Chair for Embedded Systems in 2017 as a research assistant. He received a B.Sc. degree – with distinction – and M.Sc. degree – with distinction – in Computer Science from the Karlsruhe Institute of Technology in 2014 and 2016, respectively. Currently, he is pursuing his PhD under the supervision of Prof. Dr. Jörg Henkel. His current research interests include NCFET and resource management for many-core architectures with a focus on thermal management. He is a student member of IEEE. ORCID 0000-0002-5989-2950

Jörg Henkel (M'95-SM'01-F'15) is the Chair Professor for Embedded Systems at Karlsruhe Institute of Technology. Before that he was a research staff member at NEC Laboratories in Princeton, NJ. He received his diploma and Ph.D. (Summa cum laude) from the Technical University of Braunschweig. His research work is focused on co-design for embedded hardware/software systems with respect to power, thermal and reliability aspects. He has received six best paper awards throughout his career from, among others, ICCAD, ESWeek and DATE. For two consecutive terms he served as the Editor-in-Chief for the ACM Transactions on Embedded Computing Systems. He is currently the Editor-in-Chief of the IEEE Design&Test Magazine and is/has been an Associate Editor for major ACM and IEEE Journals. He has led several conferences as a General Chair incl. ICCAD, ESWeek and serves as a Steering Committee chair/member for leading conferences and journals for embedded and cyber-physical systems. Prof. Henkel coordinates the DFG program SPP 1500 "Dependable Embedded Systems" and is a site coordinator of the DFG TR89 collaborative research center on "Invasive Computing". He is the chairman of the IEEE Computer Society, Germany Chapter, and a Fellow of the IEEE.

Andreas Gerstlauer (S'97-M'04-SM'11) is an Associate Professor with the Electrical and Computer Engineering Department at The University of Texas at Austin. He received the Ph.D. degree in Information and Computer Science from the University of California, Irvine (UCI) in 2004. Prior to joining UT Austin in 2008, he was an Assistant Researcher with the Center for Embedded Computer Systems (CECS) at UCI. His research interests include system-level design automation, system modeling, design languages and methodologies, and embedded hardware and software synthesis. He has co-authored 3 books and more than 100 refereed conference and journal publications. His work has received several best paper nominations from, among others, DAC, DATE and HOST, and best paper awards at DAC'16 and SAMOS'15. He is the recipient of a 2016-2017 Humboldt Research Fellowship. He has been General and Program Chair for conferences such as MEMOCODE and CODES+ISSS, and he currently serves as Associate Editor for ACM TODAES and ACM TECS journals. He is a senior member of the IEEE. ORCID 0000-0002-6748-2054.

Hussam Amrouch (S'11-M'15) is a Research Group Leader at the Chair for Embedded Systems (CES), Karlsruhe Institute of Technology (KIT), Germany. He is leading of the Dependable Hardware research group. He received his Ph.D. degree (Summa cum laude) degree from KIT in 2015. His main research interests are emerging technologies, design for reliability from physics to system level and machine learning. He holds seven HiPEAC Paper Awards. He has three best paper nominations at DAC'16, DAC'17 and DATE'17 for his work on reliability. He currently serves as Associate Editor at Integration, the VLSI Journal. He is a member of the IEEE. ORCID 0000-0002-5649-3102