# Performance, Power and Cooling Trade-Offs with NCFET-based Many-Cores

Martin Rapp*, Sami Salamin*, Hussam Amrouch*, Girish Pahwa†, Yogesh Chauhan†, Jörg Henkel*

*Department of Computer Science, Karlsruhe Institute of Technology, Karlsruhe, Germany
†Electrical Engineering Department, Indian Institute of Technology Kanpur, Kanpur, India
{martin.rapp, sami.salamin, amrouch, henkel}@kit.edu, {girish, chauhan}@iitk.ac.in

*Abstract*—**Negative Capacitance Field-Effect Transistor (NCFET) is an emerging technology that incorporates a ferroelectric layer within the transistor gate stack to overcome the fundamental limit of sub-threshold swing in transistors. Even though physics-based NCFET models have been recently proposed, system-level NCFET models do not exist and research is still in its infancy. In this work, we are the first to investigate the impact of NCFET on performance, energy and cooling costs in many-core processors. Our proposed methodology starts from physics, where the effects of ferroelectricity do occur, all the way up to the system level, where the performance and power of a many-core is ultimately affected. Our methodology allows for the first time the exploration of the new trade-offs between performance gains and power losses that NCFET brings to system-level designers. It demonstrates that an optimal ferroelectric thickness does exist. In addition, it reveals that current power management techniques fail when NCFET (with a thick ferroelectric layer) comes into play.**

*Index Terms*—**Negative Capacitance, NCFET, FinFET, Many-Cores, Processor Performance, Emerging Technology**

## I. Introduction

NCFET incorporates a ferroelectric material within the gate stack of the transistor that manifests itself as a negative capacitance (NC) in which a voltage amplification for the internal gate of transistor is provided. NC bypasses the fundamental thermionic limit of the sub-threshold swing of a transistor (60 mV/decade at room temperature), which is imposed by the so-called "Boltzman tyranny" [1], [2]. Hence, pMOS and nMOS transistors achieve the same ON currents but at a lower operating voltage ($V_{dd}$). Even though the concept of NC in ferroelectric materials was postulated since decades, it was until 2008 when the first direct measurement of NC was demonstrated [2]. After which, many experimental works [3] have proved the great potential that NCFET has akin to its outstanding electrical characteristics compared to conventional FETs. Importantly, GlobalFoundries has recently demonstrated the compatibility of NCFET with the existing standard CMOS fabrication process, which is essential for any emerging technology to be used in commercial applications. They fabricated the first NCFET-based circuit using industrial 14nm FinFET technology [3]. Such a compatibility would not be possible without the discovery of ferroelectricity in $HfO_2$ material [4], which is the standard high-k material for transistor's dielectric.

NCFET technology impacts the efficiency of circuits in two primary ways: 1) It allows circuits to be clocked at higher frequencies compared to conventional FET, while $V_{dd}$ remains the same. 2) It allows circuits to operate at a lower $V_{dd}$, while still achieving the same frequency as conventional FET. While there is a limited amount of work that has explored the impact



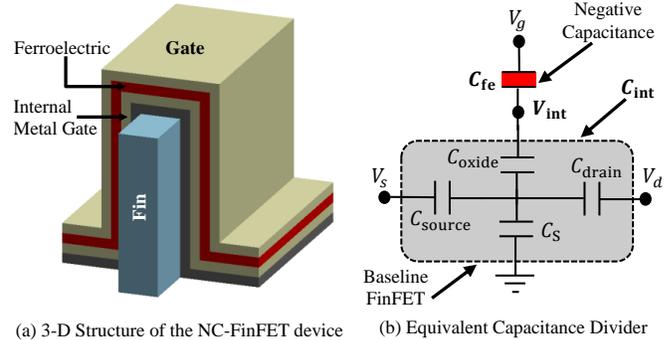(a) 3-D Structure of the NC-FinFET device    (b) Equivalent Capacitance Divider

Fig. 1. (a) shows the structure of an NC-FinFET device with metal-ferroelectric-metal-insulator-semiconductor (MFMIS) configuration. (b) presents the equivalent capacitance divider consisting of the baseline FinFET and the negative capacitance (NC) introduced by the ferroelectric layer.

of NCFET on the speed and energy of circuits [5], [6], this work is the first to investigate the NCFET technology in the context of many-core processors.

**In this work, we are the first to:**

- Propose a methodology to model NCFET at the system level and therefore bridge the gap between existing physics-based models and the system level.
- Investigate the impact of NCFET on the performance, power, energy and cooling trade-offs of a many-core.
- Demonstrate that existing power management techniques fail in NCFETs with thick ferroelectric layers.

**At the system level, we answer the following key questions:**
(1) How far does NCFET increase the thermally-sustainable utilization (number of active cores) at the peak V/F-levels?
(2) How far does NCFET increase the thermally-sustainable CPU frequencies of a many-core under full utilization? How does this impact the performance and energy consumption?
(3) How far does NCFET lead to savings in cooling costs, while still fulfilling the required thermal constraint?

## II. Negative Capacitance FETs

The principle idea of NCFET is to incorporate a thin layer of a ferroelectric material within the gate stack of a transistor. The presence of NC under some conditions due the ferroelectricity provides a voltage amplification that pushes the sub-threshold swing of a transistor to go beyond its fundamental limit. In Fig. 1, we demonstrate our used configuration in FinFET along with the equivalent capacitance divider circuit. The voltage amplification ($A$) provided by NC depends on the applied voltage on the transistor gate ($V_g$), the ferroelectric capacitance ($C_{fe}$) and the internal baseline FET capacitance ($C_{int}$). The

internal voltage $V_{int}$, which the transistor perceives at the end under the effects of NC, can be expressed as follows:

$$V_{int} = AV_g \; ; \; A \approx \frac{|C_{\text{fe}}|}{|C_{\text{fe}}| - C_{\text{int}}} \; ; \; |C_{\text{fe}}| > C_{\text{int}} \Rightarrow A > 1 \quad (1)$$

The thicker the ferroelectric layer, the higher is $A$, resulting in larger ON currents in NCFETs. The maximum frequency of a circuit ($f$) is determined by the maximum delay of all paths of a circuit. The latter depends on the delay of the pMOS and nMOS transistors that form the critical path ($CP$). The delay of a transistor is reversely proportional to its ON current. Hence, increases in the ON current directly result in a boost in the circuit's frequency, without the need to increase $V_{dd}$.

$$f = \frac{1}{t_{CP}} \; ; \; t_{CP} = \sum_{i=1}^{n} t_d(i) : i \in CP \; ; \; t_d \propto \frac{1}{I_{ON}} \quad (2)$$

To ensure a hysteresis-free operation, which is a prerequisite for CMOS logic, the condition $|C_{\text{fe}}| > C_{\text{int}}$ must be met. However, $C_{\text{fe}}$ exhibits a negative value due to the ferroelectricity effect. Hence, the total gate capacitance $C_T$ of NCFET is *always larger* than the gate capacitance of the baseline FET ($C_{int}$), as Eq. 3 demonstrates.

$$C_T = \frac{C_{\text{fe}} \cdot C_{\text{int}}}{C_{\text{fe}} + C_{\text{int}}} \; ; \; |C_{\text{fe}}| > C_{\text{int}} \wedge C_{\text{fe}} < 0 \Rightarrow C_T > C_{\text{int}} \quad (3)$$

Due to the increase in the gate capacitance, NCFET-based circuits will, at the same voltage, consume larger dynamic power compared to the conventional FET ($P_{dyn} \approx C \cdot V_{dd}^2 \cdot f$). However, because NC provides a boost in the circuit's frequency, NCFET achieves the same baseline $f$ at a lower $V_{dd}$. Reduction in $V_{dd}$, in turn, results in quadratic savings in the dynamic power. Therefore, at the same baseline $f$, NCFET-based circuits exhibit considerable power savings in the dynamic power. Increasing the thickness of the ferroelectric layer leads to a higher amplification. Hence, the baseline $f$ can be achieved at even lower $V_{dd}$ and dynamic power.

## III. NCFET PROCESSOR-LEVEL INVESTIGATIONS

Fig. 2 illustrates our proposed methodology to model NCFET at the system level. We start from physics, where the effects of ferroelectricity occur, and traverse up to the final GDSII chip-level (i.e. final level for chip fabrication), where the frequency and power of a processor will be ultimately affected. This allows us to accurately investigate how different NCFET technologies (represented by different ferroelectric layer thicknesses) impact the efficiency of a processor. Then, we develop abstract, yet accurate models for the system level that capture how the power and frequency of a processor behave under voltage scaling for both NCFET and baseline (i.e. conventional FinFET). Later, these models are integrated within a system-level many-core simulator, in which we investigate how the figure of merit (i.e. frequency, power, energy, execution time, temperature, and cooling cost) will be affected by the NCFET technology.

We start from physics, where we model the voltage amplification due to the ferroelectric layer based on the Landau-Khalatnikov (L-K) theory [2], [14]. The model is implemented within the the industry standard BSIM-CMG model [7], [15] for FinFET devices, which takes short-channel and quantum-mechanical effects into account. Then, we employ the 7nm
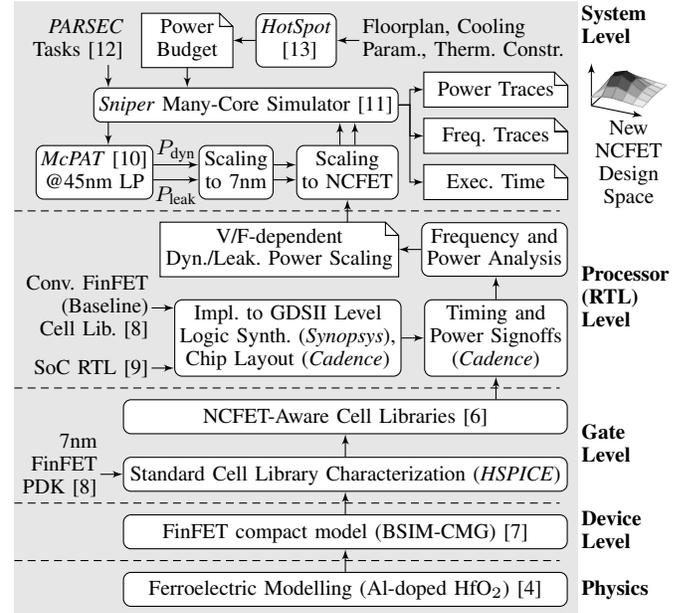


Fig. 2. General overview of our proposed methodology in which we link physics to system level to investigate NCFET at the system level. Starting from our physics-based model for ferroelectricity effects, we create NCFET-aware cell libraries for 7nm FinFET. Libraries are fully compatible with standard tool flows of chip design. Therefore, we use them within commercial signoff tools for accurate frequency and power analysis of a full SoC design. Based on this analysis for the baseline and different NCFET technologies across the entire voltage range (0.2 V to 0.7 V), high-level power scaling models are developed and employed within a system-level simulator towards analyzing how NCFET impacts the system-level design space, i.e. the trade-off new between performance, energy and cooling of a many-core that NCFET brings.

FinFET Process Design Kit (PDK) [8] to get the required FinFET parameters, as well as the post-layout SPICE netlist of standard cells. Afterwards, we characterize, using the *HSPICE* circuit simulator, the delay and power of standard cells available within the 7nm PDK. The cell library characterization is done for the baseline (conventional FinFET) and NCFET starting from the nominal $V_{dd}$ (0.7 V as defined in the PDK) down to 0.2 V [6]. This provides us, for the entire voltage range, with baseline (NCFET-unaware) and NCFET-aware standard cell libraries. All libraries are *fully compatible* with existing commercial EDA tool flows like Synopsys and Cadence. Therefore, the standard design flow for chip implementation (i.e. from logic synthesis all the way to the final GDSII level) can be fully applied. To explore different NCFET technologies, we created the NCFET-aware cell libraries for varied ferroelectric layer thicknesses ($t_{fe}$): 1 nm, 2 nm, 3 nm, and 4 nm. These NCFET technologies are referred to as TFE1, TFE2, TFE3 and TFE4, respectively. This enables us to explore the new trade-offs between frequency and power that NCFET brings. Note that thicker ferroelectric layers (i.e. $t_{fe} \geq 5$ nm) leads to hysteresis, which must be avoided in CMOS logic [6]. Therefore, our exploration is up to $t_{fe} = 4$ nm.

To study the impact of NCFET at the processor level, we target a full System-on-Chip (SoC) design. To this end, we employ the state-of-the-art *OpenPiton* SoC, which is an open-source RTL SoC [9] based on the *OpenSPARC T1* core [16]. *OpenPiton* is designed for many-core processors. In our implementation, we employ a full tile in which a CPU core, a Floating Point Unit (FPU), caches, and Network-
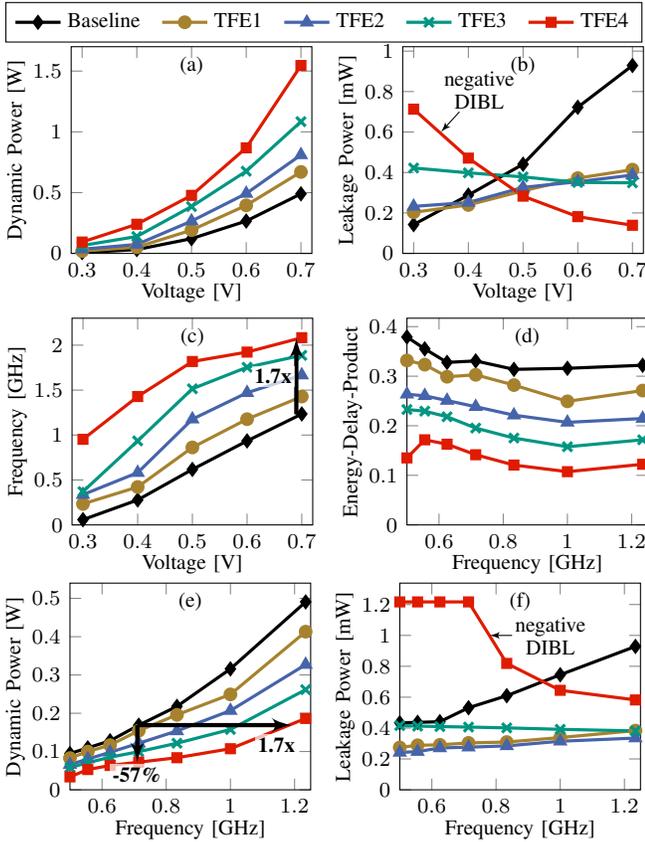
Fig. 3. Figure-of-merit analysis for the studied SoC. All the frequency and power analysis is performed at the final GDSII chip-level using standard commercial signoff tool flows. (a and b) demonstrate the impact of NCFET on the dynamic and leakage power at varied $V_{dd}$. (c) shows the impact of NCFET on boosting the frequency, compared to the baseline (i.e. conventional FinFET) at the same $V_{dd}$. (d) demonstrate the relation between Energy-Delay-Product (EDP) and frequency. (e) summarizes the relation between the dynamic power and frequency. For the same power budget, NCFET (TFE4) boosts the frequency by around 68% and for the same frequency level, NCFET (TFE4) provides 57% dynamic power saving. (f) shows how NCFET impacts the leakage power over varied frequency levels.

on-Chip (NoC) routers are all implemented from RTL to the GDSII level. Implementing such a complex SoC design provides us with very representative and accurate analysis. This is unlike [6], which considered merely a CPU core and performed the analysis only at the logic-synthesis level (i.e. without any chip layout). Such an analysis suffers from inaccuracy as it does not take the impact of parasitics into account. On the other hand, [5] studied the impact of a single NCFET configuration (i.e. single ferroelectric thickness) considering the parasitics effects in some circuits like AES and multiplier but not for a full SoC and for different NCFET configurations, as done in our work. We implement the SoC only once with the baseline cell library and then perform power and timing analysis also with the NCFET-aware cell libraries, which are fully compatible. *This allows us to pinpoint the impact of the negative capacitance technology alone and thus exclude any potential effects from different layouts.*

**Figure-Of-Merit Analysis:** Fig. 3 summarizes the frequency and power analysis for the studied SoC. All results are obtained using standard commercial signoff tool flows for accurate power and delay analysis. Fig. 3a demonstrates how

the dynamic power is affected by NCFET compared to the baseline (i.e. conventional FinFET without any ferroelectric layer). As shown, for the same $V_{dd}$ level, NCFET always consumes higher dynamic power and as the thickness ($t_{fe}$) increases, the increase in dynamic power becomes larger. This is because of the increase in the total chip's capacitance due to the increase in the transistor gate's capacitance caused by the negative capacitance (i.e. integration of a ferroelectric layer) within the transistor's gate stack (details in Section I). For TFE4 case, the increase in dynamic power reaches around 3x at the nominal $V_{dd}$ (0.7 V). Fig. 3b shows how the leakage power is affected by NCFET over different $V_{dd}$. In the baseline, $V_{dd}$ reductions lead to lower leakage power. However, the dependency of leakage on $V_{dd}$ becomes weaker in NCFET for the case of TFE1 and TFE2. In TFE3, the leakage power becomes almost independent of $V_{dd}$. *Importantly, the dependency of leakage and $V_{dd}$ becomes reverse in the case of TFE4. This is due to the negative Drain-Induced Barrier Lowering (DIBL) effect that appears in NCFET at thicker ferroelectric layers [17]. Revealing such effects is only possible due to our proposed methodology that traverses all layers from physics to processor level.* Fig. 3c shows the impact of NCFET on boosting the frequency. As can be seen, at the same $V_{dd}$ level, NCFET always increases the frequency. The gain in frequency becomes larger at thicker ferroelectric layers. It can also be noticed from the analysis that the baseline frequency can be always achieved at lower $V_{dd}$ in NCFET. Fig. 3d demonstrates the relation between the Energy-Delay-Product (EDP) and frequency, in which NCFET always exhibits lower EDP. Finally, we summarize in Fig. 3e and 3f the relation between frequency and dynamic / leakage power. As shown in Fig. 3e, when having a baseline frequency of 0.7 GHz, as an example, TFE4 results in around 57% saving in dynamic power, which is due to the capability of TFE4 to meet that frequency level at a lower $V_{dd}$ (0.3 V) compared to the baseline, which needs to operate at the nominal $V_{dd}$ of 0.7 V to meet the same frequency. On the other hand, when having a power budget of 0.2 W, as an example, TFE4 provides 68% increase in the frequency compared to the baseline.

## IV. SYSTEM-LEVEL EVALUATION: EXPERIMENTAL SETUP

We apply the processor-level frequency and power analysis to the system level to investigate the design-space of a many-core built upon NCFET technology. We study a many-core that operates under a thermal constraint. We explore two different cooling scenarios: passive air, as in mobile CPUs, i.e. static air without a fan [18], and forced-convection air, as in desktop CPUs. We use the *HotSniper* toolchain [19] that combines the *Sniper* many-core simulator [11] with *McPAT* [10] for area and power estimations and *HotSpot* [13] for thermal simulations. We simulate a many-core with 5×5 out-of-order cores that are all connected by a NoC. The private per-core L1-D and L1-I caches have a size of 32 KB, each. The L2 caches are also private and have a size of 512 KB, each, 12.5 MB in total. We run tasks from the *PARSEC* benchmark suite [12].

Fig. 2 shows the system-level experimental setup. Because *McPAT* does not supports 7nm FinFET, which is our target, we perform simulations at 45nm using low-power devices, which
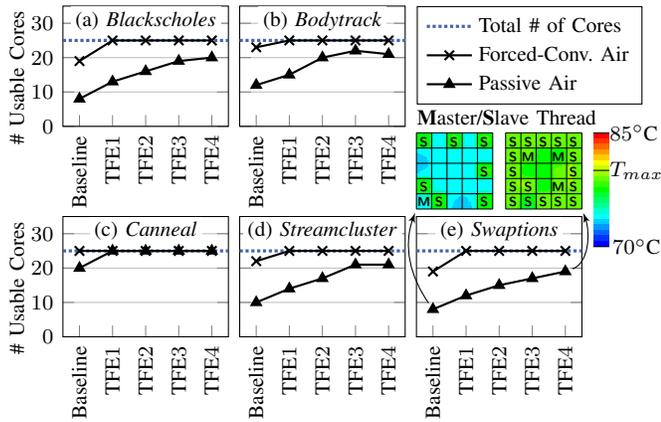
Fig. 4. NCFET technologies increase the number of usable cores without violating the temperature constraint compared to the baseline (conventional FinFETs). The highest benefit is observed with TFE3 or TFE4 (by up to 2.5x).

are closest to the devices used in the analysis in Section III. In order to scale the area and power from 45nm to 7nm, we employ the same approach as described in Section III in which we implement the *OpenPiton* SoC using both a 45nm Bulk CMOS [20] and 7nm conventional FinFET standard cell library. We then compare the designs to develop scaling factors for area, dynamic and leakage power. To simulate an implementation with NCFET, we further scale the dynamic and leakage power according to the frequency and power analysis described in Section III. We extract frequency-dependent scaling factors from Fig. 3e and Fig. 3f by comparing the power consumption of different technologies at iso-frequency. Since the peak frequency of the simulated many-core is 2.4 GHz, whereas the peak frequency of the *OpenPiton* SoC is 1.25 GHz with the baseline (conventional FinFET), we normalize the frequency to properly extract the scaling factors.

We use *McPAT* to estimate the area of the modeled processor. The scaled area of each core at 7nm is $11.4\,\text{mm}^2$ including caches and routers. This results in a total chip area of $286\,\text{mm}^2$ ($17\,\text{mm} \times 17\,\text{mm}$). Dynamic Voltage and Frequency Scaling (DVFS) sets the frequency from a minimum of $1.0\,\text{GHz}$ up to $2.4\,\text{GHz}$ in steps of $100\,\text{MHz}$. Inactive cores are power-gated.

We use *HotSpot* to model the cooling system. The ambient and maximum temperature are set at $45\,°\text{C}$ and $80\,°\text{C}$, respectively. The cooling scenario determines the thermal resistance between the heatsink and ambiance, which is set at $0.6\,\frac{\text{K}}{\text{W}}$ for passive air and $0.25\,\frac{\text{K}}{\text{W}}$ for forced-convection air [18]. We use *HotSpot*'s default configuration for all remaining parameters. This results in 50 W and 100 W Thermal Design Power (TDP) for both cooling scenarios, respectively.

## V. EFFICIENCY IMPROVEMENTS DUE TO NCFET

We consider two classes of tasks and study how NCFET increases their performance: tasks with high Thread Level Parallelism (TLP) and tasks with high Instruction Level Parallelism (ILP).

### A. Number of Usable Cores at the Peak V/f-Levels

Tasks with high ILP benefit more from high voltage/frequency (V/f)-levels than from high number of threads. Hence, cores running such tasks should be operated at the peak V/f-levels. Due to high power densities and limited cooling

capabilities, it is not always possible in modern technology nodes to simultaneously activate all cores without violating the thermal constraint. We investigate how NCFET affects the number of simultaneously usable cores.

The number of simultaneously usable cores at the peak V/f-levels depends on three factors: (i) the cooling system, (ii) the power consumptions of the threads, and (iii) the mapping of threads to cores. We explore the two cooling scenarios described in Section IV. The power consumption depends on the tasks' characteristics. We run *PARSEC* tasks with up to eight slave threads. Their characteristics range from highly memory-bound (e.g., *canneal*) to highly compute-bound (e.g., *swaptions*). The mapping of cores to tasks influences the peak temperature – and therefore the number of usable cores – because it influences the formation of hotspots. We use an Integer Linear Program to obtain the thermally optimal mapping. We employ a simple DVFS policy that sets the peak V/f-levels to cores that execute an active thread and sets the minimum V/f-levels to cores with an idle thread. Cores with no thread mapped to it are power-gated. This policy is used in all experiments to enable fair comparisons.

Fig. 4 shows the number of usable cores for the different tasks during execution of the Region of Interest (RoI) with different NCFET technologies and cooling scenarios. NCFET always allows to power on more cores than the baseline while still meeting the thermal constraint. With forced-convection air, power reductions with TFE1 already allow all benchmarks to turn on all cores. Further insights can be gained with passive air cooling. Increasing $t_{fe}$ increases the number of usable cores. TFE3 results in the highest number for *bodytrack*, whereas TFE4 results in the highest number for all other benchmarks. This is because NCFET affects dynamic power and leakage power differently (see Fig. 3e and 3f). TFE4 results in the lowest dynamic power among all technologies, but also has the highest leakage for low V/f-levels due to the negative DIBL effect. Power consumption of active threads, which are executed at the peak V/f-levels, is dominated by the dynamic power. Contrarily, idle threads, which are executed at the minimum V/f-levels, mainly dissipate leakage power. With *bodytrack*, two threads are idle during the RoI, whereas only one thread (the master thread) is idle with the other benchmarks. Therefore, *bodytrack* suffers more from the leakage power increase of TFE4, and TFE3 results in the highest number of cores. *It is noteworthy that the increase in the number of usable cores does not come at the cost of higher peak temperatures, reduced frequencies or higher cooling cost.*

### B. Increasing the Maximum Thermally-Sustainable Frequency

Tasks with high TLP benefit more from high number of threads than from high V/f-levels. These tasks should be executed with as many threads as there are available cores. To avoid thermal violations, cores cannot be operated at the peak V/f-levels any more, but DVFS needs to be employed to select thermally-sustainable V/f-levels. This section presents an investigation how NCFET increases the maximum sustainable frequency *under full utilization* and how this affects the performance and energy consumption. We enforce the thermal constraint by per-core power budgets. Since all cores
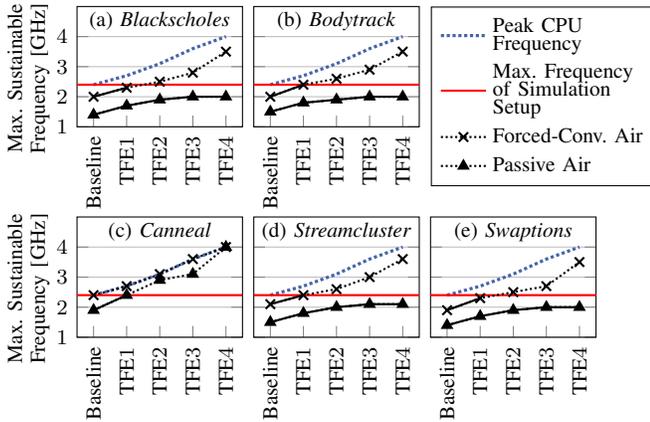
Fig. 5. NCFET increases the maximum thermally sustainable frequency of the slave threads during the RoI. Highest benefits are seen with TFE3 or TFE4. Frequencies above 2.4 GHz cannot be simulated by our setup and therefore are obtained by employing the power and frequency scaling from Section III.
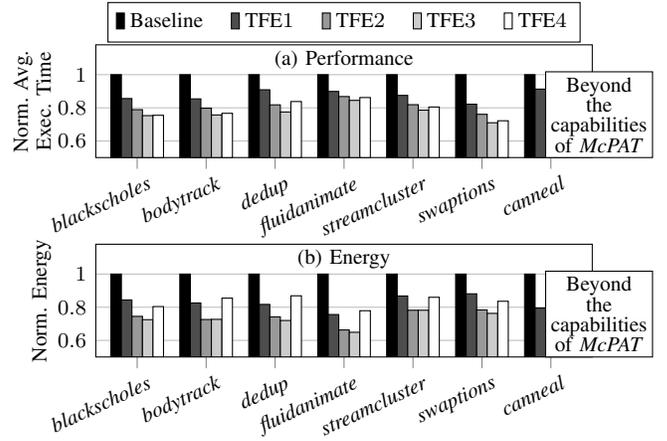


Fig. 6. NCFET decreases the average execution time under a given thermal budget (passive air cooling). TFE3 gives the highest improvements of up to 29 % (23 % on average for TFE3). This is the main reason for an energy decrease by up to 35 % (27 % on average). It is beyond the capabilities of *McPAT* to perform simulations where the maximum sustainable frequency is larger than the peak CPU frequency of the baseline design (conventional FinFETs).

are active, we divide the global power budget (TDP) evenly among all cores to determine the per-core power budget. This budget is enforced by per-core DVFS control loops that set the V/f-level to the highest level that does not lead to a power budget violation and therefore is thermally sustainable.

**Sustainable Frequencies:** We extract the maximum sustainable frequencies directly from simulation runs. Since the *McPAT* tool employed in our setup is unaware of the NCFET technology, the maximum frequency in the simulations is limited to 2.4 GHz, which is the peak CPU frequency of the baseline design for the selected micro-architectural parameters. However, Fig. 3c shows that NCFET raises the peak CPU frequency by up to 1.7x compared to the baseline, leading to a peak CPU frequency of 4.0 GHz for TFE4. Therefore, there is a gap between the peak CPU frequency and the maximum simulatable frequency. For runs that reached the maximum simulatable frequency without fully utilizing the power budget, we use the (normalized) frequency and power trends from Fig. 3e and 3f to extract the actual maximum sustainable frequency. However, this approach does not allow us to report performance and energy consumption for these runs.

Fig. 5 shows the maximum sustainable frequencies of the slave threads during the RoI. NCFET raises the sustainable frequency by up to 2x for *canneal*, which is almost unaffected by the thermal constraint. Two different trends are seen with the other benchmarks. For passive air cooling, TFE4 shows almost no improvement over TFE3, whereas with forced-convection air cooling, TFE4 raises the maximum sustainable frequency significantly over TFE3. This is because absolute frequencies are lower with TFE3 and the negative DIBL effect increases leakage more strongly, which sets a lower limit to the dynamic power and therefore limits the frequency. This experiment shows that trends observed at the processor level (peak frequency) do not directly translate to a thermally constrained many-core (sustainable frequency). Unlike previous work [6], which explored only a single core, this work investigates an NCFET-based many-core and is able to observe such trends.

**Performance and Energy:** As described earlier, our setup cannot simulate frequencies higher than 2.4 GHz. We limit this investigation to scenarios where the maximum sustainable

frequency is lower than this value. Therefore, we focus on the passive air cooling scenario. For *canneal*, the sustainable frequencies exceed 2.4 GHz even under passive air cooling for TFE2, TFE3 and TFE4. We exclude these combinations.

Fig. 6a shows the performance (execution time) for different benchmarks and NCFET technologies. The increase in the maximum sustainable frequencies directly improves the execution time. TFE3, which resulted in the highest sustainable frequencies, also results in the lowest execution times among all NCFET technologies, closely followed by TFE4. Fig. 6b shows the improvements in the energy consumption. All NCFET technologies share the same power budget, which is determined by the thermal constraint. Therefore, the power consumption is mostly the same among all NCFET technologies. The reduction in the energy consumption stems mainly from the reduction in the execution time. Idle threads are the second contributor to energy consumption. The power consumption of these threads, which is mostly leakage power, depends on the NCFET technology. This is the reason why TFE4 results in significantly less energy savings than TFE3 even though the execution time is similar. TFE3 results in the highest savings in the execution time and energy consumption. *Again, it is noteworthy that the gained performance along with energy savings does not come at the cost of higher peak temperatures or increase in cooling costs.*

## VI. IMPACT ON THE COOLING SYSTEM

NCFET reshapes the existing trade-off between cooling costs and achievable performance, where higher performance comes at the cost of higher power dissipation, and therefore higher cooling costs. We demonstrate how, for the same performance, NCFET decreases the cooling costs. We fully utilize the many-core in this investigation and use the same DVFS policy as in Section V-A, which uses the peak V/f-levels for active threads and the minimum levels for idle threads. We modify the thermal resistance of the heat sink $R_{th}$, while keeping the remaining cooling system unchanged. This corresponds to changing the air convection.
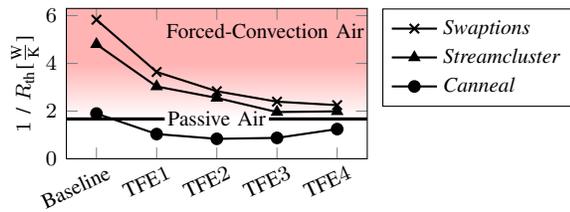
Fig. 7. NCFET decreases the required cooling capabilities by allowing a higher thermal resistance of the heat sink ($R_{\text{th}}$) while maintaining the same maximum temperature of 80°C.

Fig. 7 shows the required $R_{th}$ depending on the technology for three representative tasks. *Canneal* consumes little power. Only very low air convection is required to maintain thermal safety in the baseline. NCFET allows to reduce the required cooling capabilities below the value for passive air by reducing the dynamic power. This headroom could be exploited by downsizing the heatsink itself. For higher $t_{fe}$ (TFE3 and TFE4), the increase in the leakage power outweighs the gains in the dynamic power. *Swaptions*, which exhibits the highest power values, requires a lower $R_{th}$ of $\frac{1}{5.8}$=0.17 $\frac{\text{K}}{\text{W}}$ with conventional FinFETs. Such a value is only possible with forced-convection air. NCFET allows to decrease this almost to passive air. This results in 90 % savings in the cooling costs [21]. The optimal $t_{fe}$ depends on the task. TFE2, TFE3 and TFE4 result in the lowest cooling requirements for *canneal*, *streamcluster* and *swaptions*, respectively.

## VII. IMPACT OF NCFET ON DVFS POLICIES

Reducing V/f-levels of idle cores is a commonly employed strategy in all current technologies because it decreases both dynamic and leakage power. However, this does not hold true anymore with NCFET with high $t_{fe}$ (e.g. TFE4). Here, due to the negative DIBL effect (see Fig. 3f), decreasing the V/f-levels to manage on-chip power/temperature increases the leakage. *This is, in fact, a paradigm shift for existing power management techniques because the underlying assumption that lowering the voltage suppresses leakage breaks down.* This observation has so far not been exploited and forms an opportunity for future NCFET-aware DVFS policies. *Modeling such effects is only possible due to our proposed methodology that traverses all layers from physics to processor level.*

## VIII. CONCLUSION

This work presented a methodology to model the emerging technology NCFET at the system level and an investigation of the impact of NCFET on a many-core. Starting from a physics-based model, we traversed to the processor level, where we extracted power- and frequency trade-offs for different thicknesses of the ferroelectric layer $t_{fe}$. NCFET increases the maximum CPU frequency while voltage remains the same, or decreases the dynamic power consumption at a given frequency. For large $t_{fe}$, the negative DIBL effect causes leakage to increase with *lower* voltages. We applied our analysis to the system-level to investigate performance and energy of a thermally-constrained many-core. We showed that NCFET allows to power on more cores at the peak V/f-levels, or to increase the thermally-sustainable frequency at full system utilization. This increases the performance while

simultaneously reducing the energy consumption. We showed that NCFET reshapes the trade-offs between performance and cooling by reducing the required cooling costs while meeting the same performance. We demonstrated that further scaling (i.e., increasing $t_{fe}$) does not necessarily increase the gains, but instead, there is always a trade-off that determines the optimal $t_{fe}$. *The implications of the shown analysis of NCFET at the system level has a wide-reaching effect on known system-level policies, as they must be re-thought.*

## REFERENCES

[1] M. Hoffmann, F. P. Fengler, M. Herzig *et al.*, "Unveiling the Double-Well Energy Landscape in a Ferroelectric Layer," *Nature*, p. 1, 2019.
[2] S. Salahuddin and S. Datta, "Use of Negative Capacitance to Provide Voltage Amplification for Low Power Nanoscale Devices," *Nano Letters*, vol. 8, no. 2, pp. 405–410, 2008.
[3] Z. Krivokapic, U. Rana1, R. Galatage *et al.*, "14nm Ferroelectric FinFET Technology with Steep Subthreshold Slope for Ultra Low Power Applications," in *IEEE Int. Electron Devices Meeting (IEDM)*, 2017.
[4] J. Müller, T. S. Boscke, U. Schröder *et al.*, "Ferroelectricity in Simple Binary ZrO$_2$ and HfO$_2$," *Nano Letters*, pp. 4318–4323, 2012.
[5] S. K. Samal, S. Khandelwal, A. I. Khan *et al.*, "Full Chip Power Benefits with Negative Capacitance FETs," in *Int. Symp. on Low Power Electronics and Design (ISLPED)*, July 2017.
[6] H. Amrouch, G. Pahwa, A. D. Gaidhane, J. Henkel, and Y. S. Chauhan, "Negative Capacitance Transistor to Address the Fundamental Limitations in Technology Scaling: Processor Performance," *IEEE Access*, vol. 6, 2018.
[7] "BSIM-CMG Model," http://bsim.berkeley.edu/models/bsimcmg.
[8] L. T. Clark, V. Vashishtha, L. Shifren *et al.*, "ASAP7: A 7-nm FinFET Predictive Process Design Kit," *Microelectronics Journal*, vol. 53, pp. 105–115, 2016.
[9] 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)*, 2016, pp. 217–232.
[10] S. Li, J. H. Ahn, R. D. Strong *et al.*, "The McPAT Framework for Multicore and Manycore Architectures: Simultaneously Modeling Power, Area, and Timing," *Transactions on Architecture and Code Optimization (TACO)*, vol. 10, no. 1, p. 5, 2013.
[11] T. E. Carlson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the Level of Abstraction for Scalable and Accurate Parallel Multi-Core Simulation," in *High Performance Computing, Networking, Storage and Analysis (SC)*. ACM, 2011, p. 52.
[12] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC Benchmark Suite: Characterization and Architectural Implications," in *Parallel Architectures and Compilation Techniques (PACT)*, 2008, pp. 72–81.
[13] W. Huang, S. Ghosh, S. Velusamy *et al.*, "HotSpot: A Compact Thermal Modeling Methodology for Early-Stage VLSI Design," *Transactions on Very Large Scale Integration (VLSI) Systems*, pp. 501–513, 2006.
[14] L. D. Landau and I. M. Khalatnikov, "On the Anomalous Absorption of Sound near a Second Order Phase Transition Point," *Dokladii Academii Nauk*, vol. 96, pp. 469–472, 1954.
[15] G. Pahwa, T. Dutta, A. Agarwal *et al.*, "Analysis and Compact Modeling of Negative Capacitance Transistor with High ON-Current and Negative Output Differential ResistancePart II: Model Validation," *Transactions on Electron Devices (TED)*, vol. 63, no. 12, pp. 4986–4992, 2016.
[16] *OpenSPARC T1 Microarchitecture Spec.*, Sun Microsystems, 2008.
[17] G. Pahwa, T. Dutta, A. Agarwal, and Y. S. Chauhan, "Designing Energy Efficient and Hysteresis Free Negative Capacitance FinFET with Negative DIBL and 3.5 XI ON Using Compact Modeling Approach," in *European Solid-State Circuits Conference (ESSCIRC)*, 2016, pp. 49–54.
[18] *Heatsink 180AB*, ABL Aluminium Components, 2014.
[19] A. Pathania and J. Henkel, "HotSniper: Sniper-Based Toolchain for Many-Core Thermal Simulations in Open Systems," *Embedded Systems Letters (ESL)*, 2018.
[20] NanGate, "Open Cell Library," http://www.nangate.com.
[21] D. Shin, S. W. Chung, E.-Y. Chung, and N. Chang, "Energy-Optimal Dynamic Thermal Management: Computation and Cooling Power Co-Optimization," *Transactions on Industrial Informatics (TII)*, 2010.