Voltage Island Management in Near Threshold Manycore Architectures to Mitigate Dark Silicon

Cristina Silvano[1]   Gianluca Palermo[1]   Sotirios Xydis[2]   Ioannis Stamelakos[1]

[1]Politecnico di Milano - Dipartimento di Elettronica, Informazione e Bioingegneria
{cristina.silvano, gianluca.palermo, ioannis.stamelakos}@polimi.it
[2]Institute of Communication and Computer Systems - National Technical University of Athens
sxydis@microlab.ntua.gr

Abstract—**The power-wall problem driven by the stagnation of supply voltages in deep-submicron technology nodes, is now the major scaling barrier for moving towards the manycore era. Although the technology scaling enables extreme volumes of computational power, power budget violations will permit only a limited portion to be actually exploited, leading to the so called *dark silicon*. Near-Threshold voltage Computing (NTC) has emerged as a promising approach to overcome the manycore power-wall, at the expenses of reduced performance values and higher sensitivity to process variations. Given that several application domains operate over specific performance constraints, the performance sustainability is considered a major issue for the wide adoption of NTC. Thus, in this paper, we investigate how performance guarantees can be ensured when moving towards NTC manycores through variability-aware voltage and frequency allocation schemes. We propose three aggressive NTC voltage tuning and allocation strategies, showing that STC performance can be efficiently sustained or even optimized at the NTC regime. Finally, we show that NTC highly depends on the underlying workload characteristics, delivering average power gains of 65% for thread-parallel workloads and up to 90% for process-parallel workloads, while offering an extensive analysis on the effects of different voltage tuning/allocation strategies and voltage regulator configurations. [1]**

I. INTRODUCTION

The end of Dennard's scaling [1] poses designers in front of the so called power/utilization wall. Projections show that the gap between the number of cores integrated on a chip and the number of cores that can be utilized will continue to widen on future technology nodes [2]. As a result, *dark silicon* - transistor count under-utilization due to power budget - has been recently emerged as a major design challenge that jeopardizes the well-established core count scaling path in current and future chip generations.

To address the dark silicon problem, researchers have proposed techniques at the micro-architectural level [3], [4], [5] down to physical and device level [6], [7]. Near-Threshold Voltage Computing (NTC) [8] represents a promising technique to mitigate the effects of dark silicon, allowing a large number of cores to operate simultaneously under a given manycore power envelope. Thus, NTC has emerged as a key enabler for extreme-scale computing platforms [9]. In comparison with the conventional Super-Threshold Voltage Computing (STC), computations at NTC regime are performed in a very energy efficient manner, unfortunately at the expense of reduced performance and high susceptibility to parametric process variations.

In this paper, we investigate the power efficiency potential of manycore architectures at the NTC regime, considering process variation as well as power delivery architectures supporting multiple V_{dd} domains, under strict performance constraints originated from multicore architectures at the STC regime. Unlike previous works on variation-aware voltage allocation that target the STC regime [10], [11], [12], we propose the formation of voltage islands (VIs) for the minimization of the impact of within-die variations, which are more evident at NTC, in both performance and power. Then, we show how process variations can be efficiently exploited for further boosting the performance of an NTC manycore. To support the aforementioned research objectives, an exploration framework for manycore architectures operating at NTC has been developed to investigate the power efficiency under different workloads, while sustaining the performance when moving from the ST to the NT region.

Evaluation results on both thread-parallel (parallel-application view - high synchronization) and process-parallel (cloud-based application view - low synchronization) workloads show the high dependence of NTC efficiency to the workload's characteristics. Moving to NT regime for a 128-core architecture, while sustaining performance values obtained by a 16-core architecture at STC, average power gains >90% are delivered for process-parallel workloads, while 65% power gains for the thread-parallel workload set. We also show that given a best-effort V_{dd} tuning scenario (i.e. let NTC manycore to run faster than the requested STC constraint), a performance improvement of 27% can be achieved at the expense of 45% NTC power overhead. However, even with 45% power overhead, the maximum power dissipated by the NTC manycore is around 10W. Finally, analyzing the V_{dd} distributions at NTC, we demonstrate that the utilization of multiple VIs together with efficient integrated regulators can be considered a feasible option at NTC to efficiently deal with the process variability.

II. STATE OF THE ART

Near-threshold voltage operation relies on the aggressive tuning of the V_{dd} very close to the transistors' threshold voltage V_{th}, to a region where still $V_{dd} > V_{th}$. This decrement of the supply voltage increases the potential for energy efficient computation, e.g. by reducing V_{dd} from the nominal 1.1 V to 500 mV, energy gains of $10\times$ are reported [8]. NTC is the

---

region that delivers interesting trade-offs regarding energy efficiency and transistor delay, since super-threshold V_{dd} quickly reduces energy efficiency while sub-threshold V_{dd} leads to slower transistors. However, NTC comes together with two major drawbacks: *(i)* reduced performance and *(ii)* increased sensitivity to process variations.

Performance reduction at NTC is exposed through the limited maximum achievable clock frequency. This is an implicit effect due to the reduction of the $V_{dd} - V_{th}$ difference, applied when moving to the NTC region. Performance degradation can be compensated by exploiting trade-off points corresponding to higher task parallelism at lower clock frequencies. Thus, an important open question for NTC to be investigated is the following: *Is the inherent parallelism of applications enough to retain the performance levels of super-threshold design with lower power consumption, thus making it worth going to near-threshold operation?* Pinckey et al. [13] studied the limits of voltage scaling together with task parallelization knobs to address the performance degradation at NTC by considering a clustered micro-architectural template with cores sharing the local cache memory. They proved that under realistic application/architecture/technology features (i.e. parallelization efficiency, inter-core communication, V_{th} selection, etc.) the theoretical energy optimum point ($\frac{dEnergy}{dV_{dd}} = 0$) moves from the sub-threshold to the near-threshold region. Considering a single supply voltage per die, the energy optimum point can be found within an interval of 200 mV higher V_{th}, thus implicitly defining the upper limits of the NTC region.

The second important challenge for manycore architectures operating at NTC regime is their increased sensitivity to process variations. The transistor delay is heavily affected by the variation of V_{th} at NT voltages compared to the one in super-threshold voltages [14], [15]. In addition, failure rate of conventional SRAM cells is increased in low voltage operation [16], [17]. As a consequence, the operating frequency of the cores varies considerably, reducing the yield. In addition, variation's effects on the total power of the chip have to be carefully considered, due to the exponential dependency of leakage current upon V_{th}.

We focus our study on the NTC design space defined by [8] and [18]. Specifically, we target power efficient NTC manycore architectures that sustain STC performance levels by considering their increased sensitivity to process variation. Performance sustainability is a critical issue for the adoption of the NTC, since best effort approaches are more suitable for managing performance fluctuations due to process variability. In comparison to previous work [8], [18] where only a single system-wide power domain is considered, we differentiate our approach by exploring multiple voltage domain NTC architectures through variation-aware voltage island (VI) formation techniques.

III. PROPOSED METHODOLOGY FOR SUSTAINING PERFORMANCE AT NTC

Voltage island formation combined with V_{dd} and frequency tuning have been proved very efficient for mitigating core-to-core frequency and leakage variations [11]. There are four power management schemes supporting voltage/frequency islands: Single-Voltage/Single-Frequency
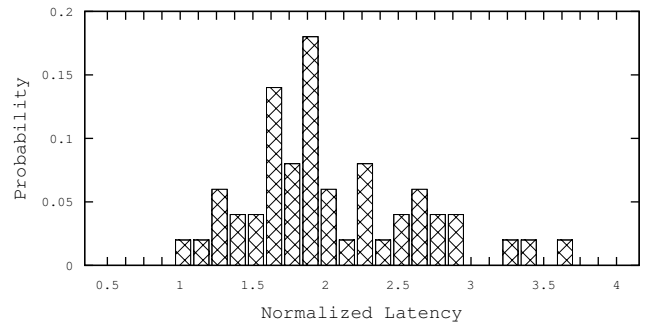

Fig. 1. Performance distribution on a 128-core NTC manycore implementing the EnergySmart [18] approach.

(SVSF) for all cores, Single-Voltage/Multiple-Frequencies (SVMF), Multiple-Voltages/Single-Frequency (MVSF) and Multiple-Voltages/Multiple-Frequencies (MVMF). While the SVSF scheme usually leads to overdesigned power management decisions, the SVMF, MVSF and MVMF schemes provide a larger set of tuning knobs for mitigating process variations. The tuning of these knobs considering only variability mitigation scenarios [18] provides no guarantees regarding the performance of the NTC manycore. In order to exploit the energy efficiency potential of NTC architectures for realistic workloads, applications running at NTC mode should ideally sustain their STC performance figures. Moving to NTC considering only the case of targeting a best-effort application domain, will limit NTC's applicability since the notion of service level agreements (SLAs), used in current data-center infrastructures and emerging cloud-based workloads, would not be efficiently supported. To further motivate the aforementioned claim, Figure 1 shows the performance distribution for a 128-core NTC manycore that implements the best-effort EnergySmart power management SVMF approach [18]. The results are obtained for the executions of the *BARNES* application over 100 different variation maps. The normalized performance value of 1 corresponds to the *nominal* performance of the application. As shown, the performance of NTC manycore platforms are not controllable and spread out over a wide range of normalized values (from 1 to 3.7) due to the underlying process variability. Thus, the adoption of NTC for applications, exhibiting specific performance and/or throughput constraints, requires careful selection and tuning of the power management scheme. In the following sections, we explore several variation-aware power management tuning strategies that will enable performance sustainability at NTC.

Figure 2(a) shows an abstract view of the target tile-based manycore architecture as well as the intra-tile organization. Although in this paper we limit the analysis to a 4 core per tile, the discussion is general and can be extended to other cluster organizations such as those proposed in [19], [18] and [20] exploring more coarse/fine-grained clusters. The intra-tile architecture is composed of 4 cores per tile and a last level cache (LL$) shared among all the cores in the tile. Each core owns a private instruction and data cache (P$). The Intel Nehalem processor [21] configuration for the core and the P$ has been adopted as reference.
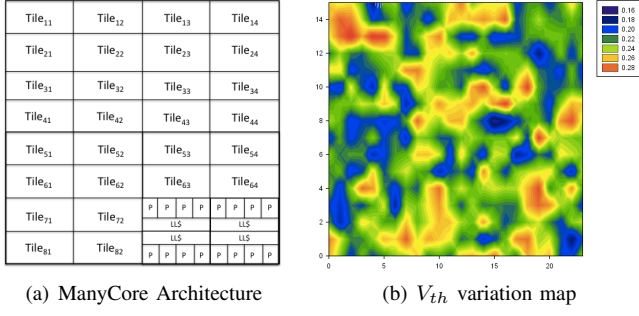
(a) ManyCore Architecture   (b) V_{th} variation map

Fig. 2. Tile-based manycore architecture (a) and corresponding V_{th} variation map (b).

A. Workload Dependent NTC Frequency for Sustained Performance

So far, application workloads have been originally developed and characterized for the STC regime. In order to sustain STC performance figures (i.e. latency or throughput) when moving to the NTC regime, the inherent parallelism of the applications should be exploited [13] to alleviate the impact of the reduced clock frequencies at NTC. Assuming a minimum allowed latency L_{min} and maximum core count constraint, C_{max} for the NTC manycore, we first calculate the clock frequency of the platform at NTC regime, f_{NTC}, that satisfies the performance constraint. Let $L_{C_{max}}$ be the performance, in terms of latency, at the STC regime of a manycore architecture with C_{max} number of cores, running at f_{STC}. At STC region, $L_{min} - L_{C_{max}} > 0$ is the available latency slack due to the higher degree of parallelism of the architecture, that can be exploited to run the application at lower frequency. Utilizing this positive slack, the f_{NTC} is calculated as follows:

$$f_{NTC} = \frac{L_{C_{max}}}{L_{min}} \times f_{STC} \qquad (1)$$

The calculated f_{NTC} refers to the target clock frequency of each core at NTC for sustaining STC performance, without considering the spatial effects of process variations. Assuming B as the set of component blocks in the floorplan and D the set of dies, we define $V_{th}^{(i,j)}$, $i \in B$, $j \in D$ that corresponds to the V_{th} of the architecture's component i in sample die j. Once extracted, $V_{th}^{(i,j)}$ is used for allocating to each component the lowest possible $V_{dd}^{(i,j)}$ for sustaining the f_{NTC} frequency constraint given that:

$$f_{NTC} \propto \frac{(V_{dd}^{(i,j)} - V_{th}^{(i,j)})^{\beta}}{V_{dd}^{(i,j)}} \qquad (2)$$

where β is a technology-dependent constant (≈ 1.5). The extraction of the f_{NTC} and the per component $V_{dd}^{(i,j)}$, enables the adoption of different power management schemes for NTC operation with guaranteed performance sustainability.

B. Going as Fast as STC: VI Formation and Variability Aware V_{dd} Allocation at NTC

Given this NTC scenario, the f_{NTC} and the $V_{dd}^{(i,j)}$ values are used by a MVSF power management scheme to form the voltage island domains and allocate their NTC voltages. The adoption of the MVSF scheme mitigates variability effects, while at the same time it derives an iso-frequency view of the manycore platform. The iso-frequency view of the platform facilitates the application development and porting, because it enables a symmetric platform from the performance point of view. Once the VIs have been defined, we compute the per island V_{dd} assignment that satisfies the f_{NTC} constraint.

More specifically, for the j^{th} die, $j \in D$, each VI, $k \in VI$, operates in its own $V_{dd}^{(k,j)}$, tuned for the $VI_{k,j}$ group of processors and memories. In $VI_{k,j}$, the core with the highest $V_{th}^{(i,j)}$, $i \in B$, $j \in D$ determines the V_{dd} for the specific voltage island, to satisfy the VI_k's critical path timing. Analyzing the trade-off by moving towards coarse grained VI granularities, we reduce area cost since less voltage regulation logic is allocated at the expenses of degrading the power efficiency of the manycore with respect to the finest possible granularity. For B_k, $k \in VI$, the set of resources found in VI_k and from Eq. 2, we calculate $V_{dd}^{(k,j)}$ according to the following relation:

$$V_{dd}^{(k,j)} = \max_{i \in B_k, j \in D} \left[V_{dd}^{(i,j)} \right] \qquad (3)$$

C. Going even Faster: Variability-aware V_{dd} Allocation Combined with Best-effort Frequency Assignment under Minimum Performance Requirements

The MVSF approach presented in the previous section guarantees the performance at NTC by allocating in a variability-aware manner the V_{dd} to each VI, in order to enable each VI to run at f_{NTC} (i.e. the minimum clock frequency requested to sustain STC performance without timing violations). However, as shown in Figure 1, the effects of process variability are not monolithic: process variation might generate on-chip regions with higher V_{th} values that reduce the achievable clock frequency as well as regions with lower V_{th} values that enable clock frequencies higher than the f_{NTC} to be allocated. The existence of positive frequency slacks at specific regions of the manycore platform can be exploited by moving from the previous MVSF approach to a MVMF power management scheme to further push system performance. The adoption of a MVMF scheme enables multiple frequencies to be allocated within a single VI tailored to the performance capabilities of the VI's components, i.e. the underlying tile architecture. However, it is worth noting that MVMF will not impact the V_{dd} allocation of the VIs, which depends on the maximum V_{th} found within each VI, thus performance guarantees continue to be valid. Thus, under the MVMF scenario, the NTC manycore is becoming heterogeneous, by including tiles of processing cores that run at least as fast as f_{NTC} or even faster, implying that the performance is not only sustained, but even optimized with respect to the STC reference configuration.

The frequency allocation within each VI is performed by applying locally the EnergySmart approach [18], since each VI can be considered as an SVMF configuration. Since the $V_{dd}^{(k,j)}$, $k \in VI$, $j \in D$, is allocated according to Eq. 3, it implies that the maximum achievable frequency, $f_{tile}^{(k,j)}$, of each tile within VI_k is bounded as follows:

$$f_{NTC} \leq f_{tile}^{(k,j)} \leq f_{MAX}^{k,j} \qquad (4)$$

TABLE I.   EXPERIMENTAL SETUP: PLATFORM PARAMETERS

Parameters	Value
Process Technology	22nm
STC Frequency	3.2GHz
STC Supply Voltage	1.05V
Nominal $V_{th}/\sigma_{V_{th}}$	0.23V/0.025
Number of Cores/Core Area	$128/6mm^2$
Tile/VI Size	4cores/4tiles
Private Cache Size/Area	$320KB/4.14mm^2$
Last Level Cache Size – Area	$8\ MB\ /\ 15.52mm^2$

where $f_{MAX}^{k,j}$ corresponds to the maximum frequency supported by $V_{dd}^{(k,j)}$ and f_{NTC} is the minimum frequency to sustain the performance. Given the NTC voltage allocation, the power overheads of allowing higher clock frequencies than f_{NTC} to be assigned, is expected to be limited due to the linear but upper bounded frequency increment. We foresee the proposed MVMF scheme to be proved very advantageous for multi-process workloads exhibiting efficient scalability due to limited synchronization, where performance boost of a single core leads to direct throughput improvements.

D. Fine-grained VI Formation by Decoupling Cores from Cache Hierarchies

The two aforementioned VI formation strategies consider the tile as the finest granularity. However, the coarser the granularity, the smaller the optimization impact of the tuning procedure, because the average or worst case effects are becoming the dominant coefficients. Providing voltage and frequency knobs at the finest granularity, the tuning procedure is becoming more complex, but also more aggressive, thus offering further optimization potentials. Given the tile-based NTC manycore architectural template considered so far, we identify the finest possible granularity by decoupling within each tile the V_{dd} of the cores from the V_{dd} allocated to the cache memory hierarchy. Recent advances in memory design have shown that extreme voltage and frequency scaling of SRAM modules close to NTC regime with sufficient resilience regarding memory content flipping hazards is now available [22]. The core-cache decoupling will enable each tile component to be tailored according to its own process variability features. Performance guarantees could be satisfied with less emphasis on the platform's components, thus leading to extra power efficiency. The basic core-cache decoupling presents a power reduction due to the reduced granularity of the VI that we measured around 3%. However this decoupling can open a research path towards the exploitation of more specific cache optimization approaches (such as [17]) to get further power savings.

So far, a major barrier to such fine-grained tuning is the low efficiency of on-chip voltage regulators, showing 10%-15% efficiency loss. However, recent advancements in fully-integrated voltage regulators like Intel's FIVR technology [23], or the low-drop out (LDO) voltage regulator scheme proposed in [24], show that cost- and power-effective on-chip voltage regulation at fine-grained does not represent anymore a visionary scenario.

IV.   EXPERIMENTAL RESULTS

In this section, we present the experimental evaluation of the proposed methodology to sustain performance in the Near

Threshold region.

A. Experimental Setup

The Sniper multicore simulator [25] and the McPAT power modeling framework [26] have been used for the performance and power characterization respectively, while the Various-NTV microarchitectural model [27] has been employed to capture the process variation at the NT regime. A summary of the experimental setup used to evaluate the methodology is presented in Table I. Core and caches types, sizes and area are taken from the Intel Nehalem architecture. The target platform is a 128 core many-core chip at NTC (at 22nm technology node) composed of 32 tiles, each one including 4 cores and a shared last level cache (LL$) of 8MB and 8 voltage islands (4 tiles each). Although in this paper we are going to present the results obtained by considering single values for the tile size and VI granularity, the approach can be easily generalized to other architectural topologies [20]. Maximum V_{dd} has been set to 1.05V and the frequency to 3.2 GHz for the STC regime, according to parameter values derived from [28] for conservative technology scaling. By assuming a maximum power budget of 80W at STC, the performance to be sustained at NTC (L_{min}) corresponds to a 16 core architecture in the STC regime. From Various-NTV, we extracted 100 different variation maps by using a 24x16 grid based on the core/cache granularity.

Finally, the target applications have been derived from the SPLASH-2 benchmark suite [29], where the "large dataset" workload, provided within Sniper [25], has been adopted. The target applications have been used for the validation in two different scenarios. The first scenario consists of the single application multiple threads (SAMT) approach, where we supposed to run a single application on the entire platform by using its internal parallelism at thread level (128 threads). The second scenario consists of multiple applications multiple threads (MAMT), where multiple instances of the same application are running (one per tile) and the internal parallelism at the thread-level is used within each tile (4 threads). This second version gives a sort of "cloud-oriented" view of the platform. The applications considered in the SAMT version exhibit different behaviors by scaling from 16 to 128 cores: close to ideal (*RADIOSITY*), medium (*BARNES, WATER-NSQ*) and limited scaling (*RAYTRACE, WATER-SP*). Additionally, we examined an *AVERAGE* case workload, that aggregates in a single execution sequence the five applications, treating them as a single benchmark. In that way, we manage to see what happens in an *average* case, where there is a combination of benchmarks that scale well and others that don't scale well. On the opposite, all the applications in the MAMT version present an almost ideal scaling passing from 16 cores (2 application instances over 2 tiles) to 128 cores (32 application instances).

B. Power Gains: NTC vs STC

Figure 3 shows the power consumption comparison when passing from 16 cores at STC to 128 cores at NTC for each benchmark in both SAMT and MAMT versions. The power values for the same benchmark on SAMT and MAMT versions are not comparable because the application performance are different in the two cases. All the MAMT versions of the applications and the RADIOSITY-SAMT deliver large power
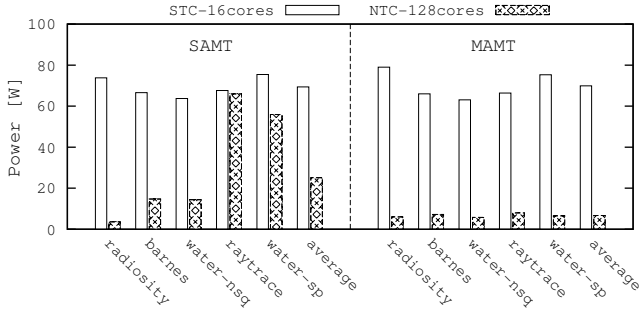
Fig. 3. Power reduction: 16-core STC chip versus 128-core NTC for both SAMT and MAMT versions of the target applications
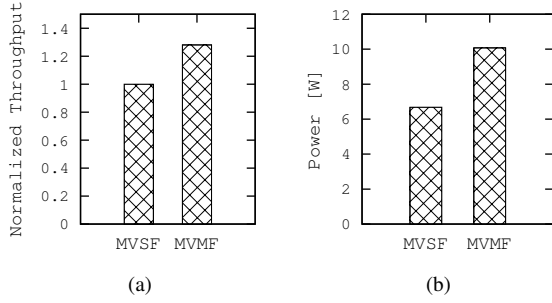

Fig. 4. Impact of MVMF vs MVSF in terms of (a) Throughput and (b) Power
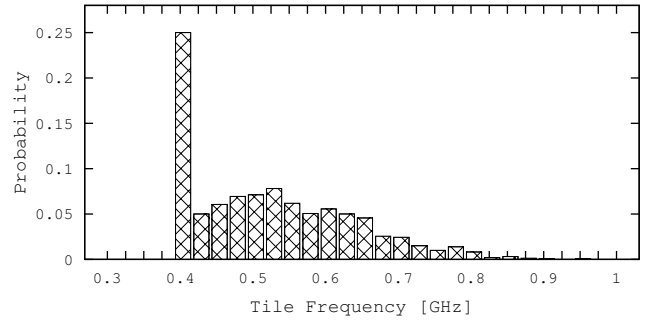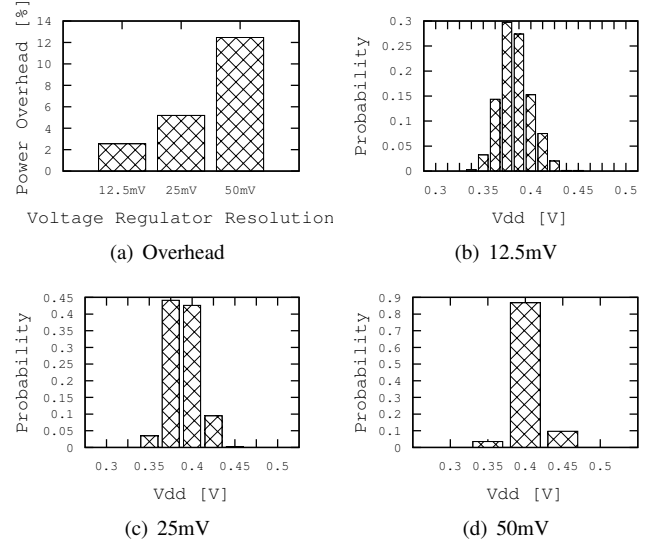

Fig. 5. Tile frequency distribution in MVMF mode

Fig. 6. Voltage regulator analysis: Power overhead (a) and V_{dd} probability distribution (b-d) for three voltage regulator resolutions

gains ($> 90\%$) due to the almost ideal performance scaling as the number of cores increases. The rest of the applications in SAMT version present a power gain that depends on the scaling capability, since it impacts the minimum frequency to be sustained and thus the minimum V_{dd} to be deployed to the voltage islands. For the remaining applications, Figure 3 shows a 75% decrement in power for BARNES and WATER-NSQ, around 25% for WATER-SP and an almost identical power for RAYTRACE. The AVERAGE-SAMT workload (composed of a sequential mix of all applications) delivers a power gain of 65%.

C. Relaxing the Isofrequency Constraint

Figure 4 shows the power/performance impact of the relaxation on the isofrequency constraint. To better evaluate this scenario, we present the experimental data considering only the MAMT version of the AVERAGE case. As stated in the previous section, while the MVMF has ideally an advantage due to the increment of the tile frequency, this can be really exploited only when the application is aware of this performance asymmetry. This is not the case of the SAMT version of our target applications. To have a clear view of the performance improvement we adopted the application throughput concept as the rate of jobs (application instances) completed within a time interval. As expected, the MVMF approach offers a performance speedup due to the frequency increment in the tiles not affected by the critical V_{th}. However, the performance improvement ($\approx 27\%$) is balanced by an increased power overhead ($\approx 45\%$).

Additionally, Figure 5 shows the tile frequency distribution across the 100 variation maps by using the MVMF mode. The

minimum frequency is $400MHz$ to guarantee the application performance in terms of throughput. As expected, the minimum value is the most probable since there is at least 1 tile per VI (the one that limits the V_{dd} scaling) running at that frequency. Regarding the other values, we can notice that the distribution shows a long tail meaning that there is a large margin that can be used for further speedups.

D. Voltage Regulators Analysis

The analysis conducted so far considers an ideal scenario where we can deliver all the requested on-chip voltage levels. According to state-of-the-art power supply architectures, we want to start including realistic constraints to the results, so in this section we analyze the impact of the on-chip voltage regulator resolution on power efficiency. We analyzed three different voltage regulator resolutions, delivering voltage with a precision of *(i)* 12.5mV, *(ii)* 25mV and *(iii)* 50mV. Figure 6 presents: the average power overhead for each voltage regulator precision in Figure 6(a) and the V_{dd} distribution according to each regulator resolution in Figures 6(b) - 6(d). The power overhead and the V_{dd} distributions have been calculated across the 100 variation maps considering a target frequency of $400MHz$ to be sustained.

In Figure 6(a) we refer to power overhead as the normalized

average difference between the power consumed in the ideal case (voltage regulator delivering arbitrary V_{dd} values) and the power corresponding to specific values of voltage precision. As expected, the higher is the resolution the smaller is the overhead since we are closer to the ideal case, passing from a 12% at 50mV to less than 3% at 12.5mV. This limited overhead value is interesting also considering the results shown in Figures 6.b-d, where it can be noticed that the V_{dd} distribution is very concentrated, which makes the use of the cost-efficient LDO on-chip regulation [24] schemes feasible to the NTC regime.

V. CONCLUSION

This paper focuses on the emerging NTC paradigm as a key enabler for the power-efficient scaling of manycore architectures. While power efficiency is guaranteed by definition at the NTC regime, performance guarantee is still an open challenge. Sustaining STC performance figures during NTC operation is a critical issue for the wider adoption of the NTC paradigm. Towards this direction, we presented a set of techniques for variability-aware voltage island formation and voltage/frequency tuning that enable moving to NTC regime while sustaining STC performance guarantees. Extensive experimentation showed the optimization potentials of moving towards near-threshold voltage computing, outlining its high dependency on both workload characteristics and voltage tuning strategy.

REFERENCES

[1] R.H. Dennard, F.H. Gaensslen, V.L. Rideout, E. Bassous, and A.R. LeBlanc. Design of ion-implanted MOSFET's with very small physical dimensions. *Solid-State Circuits, IEEE Journal of*, 9(5):256–268, 1974.

[2] Hadi Esmaeilzadeh, Emily Blem, Renee St. Amant, Karthikeyan Sankaralingam, and Doug Burger. Dark silicon and the end of multicore scaling. In *Proceedings of the 38th annual international symposium on Computer architecture*, ISCA '11, pages 365–376, 2011.

[3] N. Goulding-Hotta, J. Sampson, G. Venkatesh, S. Garcia, J. Auricchio, P. Huang, M. Arora, S. Nath, V. Bhatt, J. Babb, S. Swanson, and M.B. Taylor. The GreenDroid mobile application processor: An architecture for silicon's dark future. *Micro, IEEE*, 31(2):86–95, 2011.

[4] V. Govindaraju, Chen-Han Ho, and K. Sankaralingam. Dynamically specialized datapaths for energy efficient computing. In *High Performance Computer Architecture (HPCA), 2011 IEEE 17th International Symposium on*, pages 503–514, 2011.

[5] Yatish Turakhia, Bharathwaj Raghunathan, Siddharth Garg, and Diana Marculescu. HaDeS: architectural synthesis for heterogeneous dark silicon chip multi-processors. In *DAC*, pages 173–178. ACM, 2013.

[6] Arun Raghavan, Yixin Luo, Anuj Chandawalla, Marios C. Papaefthymiou, Kevin P. Pipe, Thomas F. Wenisch, and Milo M. K. Martin. Computational sprinting. In *HPCA*, pages 249–260. IEEE, 2012.

[7] Francesco Paterna and Sherief Reda. Mitigating dark-silicon problems using superlattice-based thermoelectric coolers. In *Proceedings of the Conference on Design, Automation and Test in Europe*, DATE '13, pages 1391–1394, San Jose, CA, USA, 2013. EDA Consortium.

[8] Ronald G. Dreslinski, Michael Wieckowski, David Blaauw, Dennis Sylvester, and Trevor N. Mudge. Near-threshold computing: Reclaiming Moore's law through energy efficient integrated circuits. *Proceedings of the IEEE*, 98(2):253–266, 2010.

[9] J. Torrellas. Extreme-scale computer architecture: Energy efficiency from the ground up. In *Proceedings of the Conference on Design, Automation and Test in Europe*, DATE '14, 2014.

[10] A. Das, S. Ozdemir, G. Memik, and A. Choudhary. Evaluating voltage islands in CMPs under process variations. In *Computer Design, 2007. ICCD 2007. 25th International Conference on*, pages 129–136, 2007.

[11] Sohaib S. Majzoub, Resve A. Saleh, Steven J. E. Wilton, and Rabab K. Ward. Energy optimization for many-core platforms: communication and PVT aware voltage-island formation and voltage selection algorithm. *Trans. Comp.-Aided Des. Integ. Cir. Sys.*, 29(5):816–829, May 2010.

[12] Sebastian Herbert, Siddharth Garg, and Diana Marculescu. Exploiting process variability in voltage/frequency control. *IEEE Trans. VLSI Syst.*, 20(8):1392–1404, 2012.

[13] N. Pinckney, K. Sewell, R. G. Dreslinski, D. Fick, T. Mudge, D. Sylvester, and D. Blaauw. Assessing the performance limits of parallelized near-threshold computing. In *Proceedings of the 49th Design Automation Conference*, pages 1147–1152, 2012.

[14] M. Eisele, J. Berthold, D. Schmitt-Landsiedel, and R. Mahnkopf. The impact of intra-die device parameter variations on path delays and on the design for yield of low voltage digital circuits. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 5(4):360–368, 1997.

[15] D. Markovic, C.C. Wang, L.P. Alarcon, Tsung-Te Liu, and J.M. Rabaey. Ultralow-power design in near-threshold region. *Proceedings of the IEEE*, 98(2):237–252, 2010.

[16] L. Chang, R.K. Montoye, Y. Nakamura, K.A. Batson, R.J. Eickemeyer, R.H. Dennard, W. Haensch, and D. Jamsek. An 8T-SRAM for variability tolerance and low-voltage operation in high-performance caches. *Solid-State Circuits, IEEE Journal of*, 43(4):956–963, 2008.

[17] A. Sasan, H. Homayoun, A. M. Eltawil, and F. J. Kurdahi. Inquisitive defect cache: A means of combating manufacturing induced process variation. *IEEE Trans. VLSI Syst.*, 19(9):1597–1609, 2011.

[18] Ulya R. Karpuzcu, Abhishek A. Sinkar, Nam Sung Kim, and Josep Torrellas. EnergySmart: Toward energy-efficient manycores for near-threshold computing. In *HPCA*, pages 542–553, 2013.

[19] Ronald G. Dreslinski, Bo Zhai, Trevor N. Mudge, David Blaauw, and Dennis Sylvester. An energy efficient parallel architecture using near threshold operation. In *PACT*, pages 175–188, 2007.

[20] Ioannis Stamelakos, Sotirios Xydis, Gianluca Palermo, and Cristina Silvano. Variation aware voltage island formation for power efficient near-threshold manycore architectures. In *Proceedings of the ASP-DAC*, ASP-DAC '14, 2014.

[21] D. Kanter. Inside Nehalem: Intel's future processor and system. http://www.realworldtech.com, 2008.

[22] Tobias Gemmeke, Mohamed M. Sabry, Jan Stuijt, Praveen Raghavan, Francky Catthoor, and David Atienza. Resolving the memory bottleneck for single supply near-threshold computing. In *Proceedings of the Conference on Design, Automation and Test in Europe*, DATE '14, 2014.

[23] Intel's fourth generation Core CPU Haswell. FIVR Fully Integrated Voltage Regulator, http://www.intel.com, 2013.

[24] Hamid Reza Ghasemi, Abhishek A. Sinkar, Michael J. Schulte, and Nam Sung Kim. Cost-effective power delivery to support per-core voltage domains for power-constrained processors. In *Proceedings of the 49th DAC*, pages 56–61, 2012.

[25] Trevor E. Carlson, Wim Heirman, and Lieven Eeckhout. Sniper: Exploring the level of abstraction for scalable and accurate parallel multicore simulations. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2011.

[26] Sheng Li, Jung Ho Ahn, Richard D. Strong, Jay B. Brockman, Dean M. Tullsen, and Norman P. Jouppi. McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures. In *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO 42, pages 469–480, 2009.

[27] Ulya R. Karpuzcu, Krishna B. Kolluru, Nam Sung Kim, and Josep Torrellas. VARIUS-NTV: A microarchitectural model to capture the increased sensitivity of manycores to process variations at near-threshold voltages. In *IEEE/IFIP International Conference on Dependable Systems and Networks, DSN*, pages 1–11, 2012.

[28] S. Borkar. The exascale challenge. In *VLSI Design Automation and Test (VLSI-DAT), 2010 International Symposium on*, pages 2–3, 2010.

[29] Steven Cameron Woo, Moriyoshi Ohara, Evan Torrie, Jaswinder Pal Singh, and Anoop Gupta. The SPLASH-2 programs: characterization and methodological considerations. *SIGARCH Comput. Archit. News*, 23(2):24–36, May 1995.