Chapter 1
A Perspective on Dark Silicon

Anil Kanduri, Amir M. Rahmani, Pasi Liljeberg, Ahmed Hemani, Axel Jantsch, and Hannu Tenhunen

1.1 Introduction

The possibilities to increase single core performance has ended due to limited instruction level parallelism and a high penalty when increasing frequency. This prompted designers to move towards multi-core paradigms [1], largely supported by transistor scaling [2]. Scaling down transistor gate length makes it possible to switch them faster at a lower power, as they have a low capacitance. In this context, an important consideration is *power density* - the power dissipated per unit area. Dennard's scaling establishes that reducing physical parameters of transistors allows operating them at lower voltage and thus at lower power, because power consumption is proportional to the square of the applied voltage, keeping power density constant [3]. Dennard's estimation of scaling effects and constant power density is shown in Table 1.1. Theoretically, scaling down further should result in more computational capacity per unit area. However, scaling is reaching its physical limits to an extent that voltage cannot be scaled down as much as transistor gate length leading to failure of Dennardian trend. This along with a rise in leakage current results in increased power density, rather than a constant power density. Higher power density

Anil Kanduri
University of Turku, Finland, e-mail: spakan@utu.fi

Amir M. Rahmani
University of Turku, Finland, e-mail: amirah@utu.fi

Pasi Lijeberg
University of Turku, Finland, e-mail: pasi.liljeberg@utu.fi

Ahmed Hemani
KTH Royal Institute of Technology, Sweden, e-mail: hemani@kth.se

Axel Jantsch
TU Wien, Vienna, Austria, e-mail: axel.jantsch@tuwien.ac.at

Hannu Tenhunen
KTH Royal Institute of Technology, Sweden, e-mail: hannu@kth.se

implies more heat generated in a unit area and hence higher chip temperatures which has to be dissipated through cooling solutions, as increase in temperature beyond a certain level results in unreliable functionality, faster ageing and even permanent failure of the chip. To ensure a safe operation, it is essential for the chip to perform within a fixed power budget [4]. In order to avoid too high power dissipation, a certain part of the chip needs to remain inactive, the inactive part is termed Dark Silicon [5]. Hence, we have to operate working cores in a multi-core system at less than their full capacity, limiting the performance, resource utilization and efficiency of the system.

1.2 The Dark Silicon Phenomenon

Too high power density is the chief contributor to dark silicon. It is to be noted that trends in computer design that are technology node centric have a significant impact on area, voltage, frequency, power, performance, energy and reliability issues. Most of these parameters are cyclically dependent on one another in a way that improving one aspect will deteriorate the other. Many-core systems face a new set of challenges at the verge of dark silicon to continue providing the expected performance and efficiency. Computational intensity of future applications such as deep machine learning, virtual reality, big data, etc., demands further technology scaling, leading to further rise in power densities and dark silicon issue. Increase in power density leads to thermal issues [6], hampering the chip's performance and functionality. Subsequently, issues of reliability and ageing have come up, in addition to limitation in performance. ITRS projections have predicted that by 2020, designers would face up to 90% of dark silicon, meaning that only 10% of the chip's hardware resources are useful at any given time when high operating frequencies are aplied [7]. Figure 1.1 shows the amount of usable logic on a chip as per projections of [5] [7]. Increasing dark silicon directly reflects on performance to a point that many-core scaling provides zero gain [8]. Thermal awareness and dark silicon sensitive resource allocation are key techniques that can address these challenges.

1.2.1 Power Density

Working components of a chip consume power and generate heat on the chip's surface area that increases chip's temperature. Higher chip temperatures adversely effect the chip's functionality, induce unreliability and quicken the ageing process. Heat accumulated on the die has to be dissipated through cooling solutions in order to maintain a safe chip temperature. The amount of heat generated on the chip's surface area depends on its power consumption. Most of the computer systems must function within a given power envelope, since heat dissipated in a given area is restricted by cooling solutions. The power budget of a chip is one of the principal

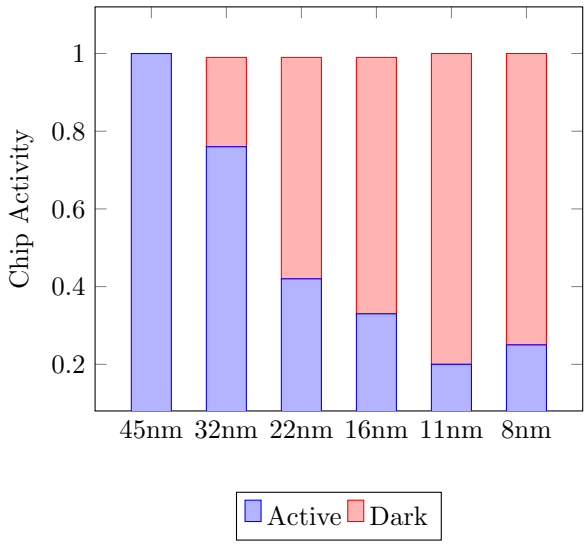

Fig. 1.1: Projections of Dark Silicon [7] [5]

design parameters of a modern computer system. With every technology node generation, power densities have been steadily increasing, while the areas and cooling solutions are stale. Increase in power density that comes with integration capacity is shown in Figure 1.2, which is estimated based on 2013 ITRS projections [9] and conservative projections by Borkar [10]. We see an exponential rise in power density, while the ideal scaling principle of Dennard states that it would remain constant. Increased power density is currently handled by unwanted dark silicon.

The scaling conclusions from Dennard's principle [3] are summarized in Table 1.1. Scaling transistor's physical parameters viz., length (L), width (W) and oxide thickness (t_{ox}) translates into reduced voltage and increased frequency. This means that scaling physical parameters by a constant factor k would also scale down voltage, capacitance and scale up frequency by k. This reduces power consumption and chip area by the same factor of k^2, keeping the power density constant. This argument held only until a reaching certain point, called the utilization wall [11]. The major causes for increase in power density are discussed in following sections.

1.2.2 Power Consumption in CMOS Chips

Power consumption in CMOS based chips mainly consists of *dynamic power* - the power consumed when transistors are switching between states and *static power* - the power drawn by leakage currents, even at an idle state, as expressed in Equation 1.1.

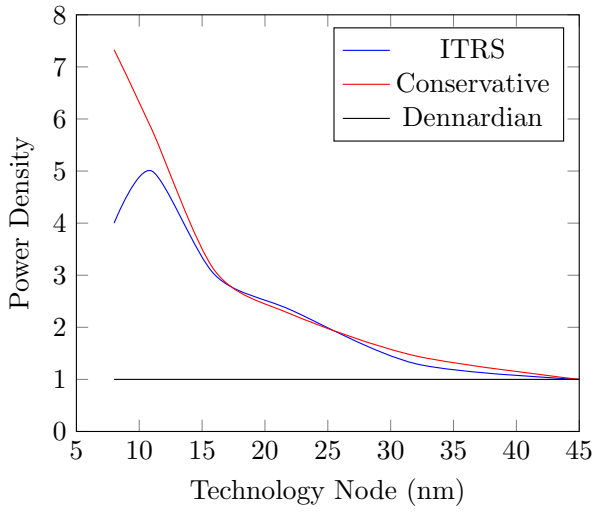

Fig. 1.2: Increase in Power Density with scaling

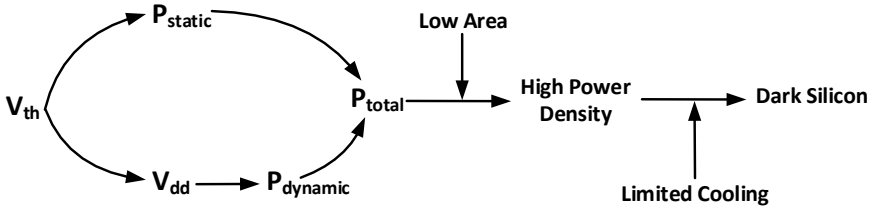

Fig. 1.3: Chief contributors to Dark Silicon

$$P_{total} = P_{dynamic} + P_{static} \tag{1.1}$$

The dynamic power component, as in Equation 1.2, is the power consumed to charge or discharge capacitive loads across the gate. Dynamic power is proportional to the capacitance C, the frequency F at which gates are switching, squared supply voltage V, and activity A of the chip (the number of gates that are switching). In addition, a short circuit current I_{short} flows between the supply and ground terminals for a brief period of time τ, whenever the transistor switches states [12]. Dynamic power consumption can be minimized by reducing any of these parameters.

$$P_{dynamic} = ACV^2F + \tau AVI_{short} \tag{1.2}$$

With transistor scaling, the voltage and frequency of the chip can also be scaled, as proposed in [3] and shown in Table 1.1. Frequency varies with supply and threshold voltages as,

$$F \propto \frac{(V - V_{th})^2}{V} \tag{1.3}$$

Reducing physical parameters reduces voltage, capacitance and hence the *RC* delay of a transistor [13]. Noticeable aspect here is that frequency varies almost as voltage, assuming a smaller threshold voltage. With higher level of integration, reduced frequency can be compensated with parallel multi-processor architectures which yield better performance [14]. Voltage scaling can reduce power by factor of a square, encouraging transistor scaling. Unlike a perfect switch, a CMOS transistor draws some power even when it is idle (i.e., not switching between the states) [15]. This can be attributed to the leakage current arising at lower threshold voltages and gate oxide thickness. More details on leakage power are presented in Section 1.2.4.

1.2.3 Slack Voltage Scaling

As predicted by Gordon Moore, semi-conductor industry has been successfully improving transistors of increase speed with decrease in gate length for every technology node generation [2]. Moore's prediction has obtained validation through Dennard's scaling law [3] which establishes an empirical relation between technology node scaling and power density such that power density remains constant even with aggressive integration of transistors in a given area. The reason behind validity of Dennard's scaling is that reduction in gate length of transistors also lowers operating voltage and load capacitance - which in turn reduces dynamic power consumption significantly [14]. Dennardian scaling held as long as voltage and transistor gate length were scaled down at a similar rate. However, voltage scaling is reaching its limits and it is increasingly challenging to scale down further. Figure 1.4 shows projections of voltage and gate length scaling from ITRS reports of 2011 [7], 2013 [9] and conservative approach [10] (Data of [7] and [10] retrieved from [5]). All the projections clearly indicate that voltage scaling is almost stuck in comparison with integration capacity, which continues to scale gracefully. Increasingly wide gap between voltage scaling and transistor scaling results in higher power density and dark silicon. This can also be considered a corollary on failure of Dennard's scaling, where it is expected that voltage and technology nodes could be scaled at a reasonably similar rate. Scaling the supply voltage is pegged with scaling the threshold voltage. The requirements on performance and reliability determines minimum and maximum limits of the supply voltage. Along with them, the risk of scaling down threshold voltage is the reason for limitation in voltage scaling [16]. Since the voltage no longer scales with the transistor size, ideal Dennardian scaling will no longer hold. Instead, the power density starts increasing with every new technology node generation, as indicated in [11].

Table 1.1 shows realistic scaling in the post-Dennardian era. Assuming that voltage scales as a factor s instead of transistor scaling factor k, the power density becomes k^2/s^2, a non-constant and greater than unity quantity. This indicates that aggressive technology scaling (high k) and/or slack voltage scaling (low s) put together

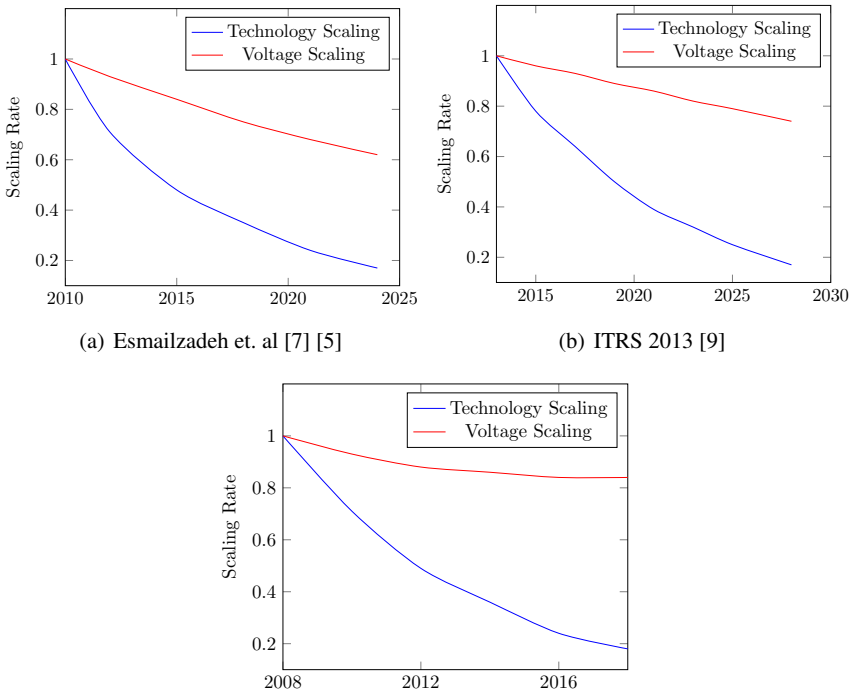

(a) Esmailzadeh et. al [7] [5]

(b) ITRS 2013 [9]

(c) Borkar et. al [10] [5]

Fig. 1.4: Technology Node versus Voltage Scaling

results in a high k^2/s^2 value, which is nothing but higher power density. Safer limit

Table 1.1: Practical scaling in post-Dennardian Era

Parameter	Symbol	Ideal Scaling	Practical Scaling
Transistor dimensions	L, W, t_{ox}	1/k	1/k
Area	A = $1/L \times W$	$1/k^2$	$1/k^2$
Capacitance	C	1/k	1/k
Gate Delay	RC	1/k	1/k
Frequency	F = 1/RC	k	k
Supply Voltage	V_{dd}	1/k	1/s
Power	P = CV^2F	$1/k^2$	$1/s^2$
Power Density	D = P/A	1	k^2/s^2

on power consumption can be achieved by operating only a section of the entire chip's resources, resulting in dark silicon [14]. From this view point, accumulated heat has to be dissipated by controlling power consumption, which practically is not possible, unless there is dark silicon.

1.2.4 Leakage Power

Power has to be viewed as summation of static and dynamic components, as in Equation 1.1. Reducing voltage and frequency would obviously reduce dynamic power consumption and also total power, as long as the static component is negligible. However, an increasingly critical component is the static power that is drawn by the leakage current I_{leak}, which is increasing with scaling [17]. Leakage current is split into two components - sub-threshold leakage and gate oxide leakage, as in Equation 1.4.

$$P_{static} = VI_{leak} = V \times (I_{sub} + I_{ox}) \tag{1.4}$$

Leakage current is the static power drawn even at an off state (not switching), meaning that a chip can still consume appreciable power without being active [14]. Off-state leakage current is experimentally driven in [18], and is represented as in Equation 1.5.

$$I_{off} = K_1 We^{\frac{-V_{th}}{nV_\theta}} (1 - e^{\frac{-V}{V_\theta}}) \tag{1.5}$$

where K_1 and n are experimentally derived constants , W is transistor width, V_θ is 25 mV. With V_θ being much smaller than V, this is further simplified in [12] as in Equation 1.6.

$$I_{off} \propto e^{\frac{-qV_{th}}{kT}} \tag{1.6}$$

where k is a constant and T is temperature. This indicates that leakage current is low when the exponential component is small, meaning that threshold voltage has to be high. In other words, leakage current increases exponentially with a decrease in threshold voltage (negative exponential component). Reduction in transistor physical parameters reduces voltage and thus also reduces threshold voltage. This results in higher leakage current and thus higher static power being drawn. It can be said that with every technology node generation, leakage current increases exponentially with a decrease in threshold voltage [13]. Also, leakage power is a cumulative of all transistors' leakage, while the number of transistors more or less doubles for every technology node generation, contributing to increasing leakage power [14]. Reduction in transistor physical dimensions also reduces oxide thickness, to an extent that electrons start to leak also through the oxide resulting in leakage current [15]. Oxide thickness is reaching around 1.5 nm, results in significant leakage current (I_{ox}) [19]. Gate oxide leakage current is shown in Equation 1.7.

$$I_{ox} = K_2 W (\frac{V}{T_{ox}})^2 e^{\frac{-\alpha T_{ox}}{V}} \tag{1.7}$$

where K_2 and α are experimental constants. Analogous to off-state leakage, gate oxide leakage depends exponentially on oxide thickness. The squared denominator and negative exponential component of Equation 1.7 shows that gate oxide leakage increases exponentially with decrease in oxide thickness. Since oxide thickness depends on other transistor parameters, it has to be scaled down with technology node scaling, thus increasing I_{ox}. While dynamic power was a lone significant contributor to total power, static power is alarmingly increasing and would exceed dynamic component with further technology scaling [12]. In addition, leakage increases exponentially at higher temperatures, implying that a chip would have far more leakage if it operates to its full potential [20]. Power density would not have been constant with scaling, however it would have been in a relatively controllable reach, had there been no static power.

1.2.5 Thermal Issues

CMOS based chips are limited by battery technology and cooling solutions that can dissipate chip's heat. Therefore, a safe peak operating temperature and the corresponding power that can be dissipatable within given volume are used as design time parameters. The upper limit on power consumption, and thus temperature, is popularly known as Thermal Design Power (TDP) [21]. Safe and reliable operation of a chip is guaranteed as long as power consumption and heat dissipation stays within TDP guidelines [21]. However, TDP is not the maximum power that can be consumed by a chip, but it is only a safe upper bound. TDP is calculated assuming a worst case application running at worst case voltage and frequencies of components. Although, a typical application might consume power that is less than TDP, yet runs under TDP constraint. Performance has to be sacrificed through voltage and frequency reduction, clock and power gating in order to stay within TDP. This is acceptable for applications that are power hungry, however other applications that may never exceed TDP would still end up running at lower than full compute capacity of the chip. In this sense, current systems are rather conservative. Such conservative estimates of the upper bound can be met through inactivating part of he chip, contributing to dark silicon. Thus far, most common practice of ensuring safe chip functionality has been by estimating TDP [21] in a conservative manner. Recent upgrades on AMD and Intel Corporations' CPUs have the option of a configurable TDP, however they are limited to a maximum of 3 modes, without any fine grained control [22][23].

1.3 Challenges and Consequences of Dark Silicon

Increase in dark silicon reduces the number of simultaneously active components on a chip and hampers performance, energy efficiency, reliability, ageing and effective resource utilization of computer systems [5]. With the dark silicon phenomenon being a hardware and device level issue, other layers of the computing stack are un-aware of the sources of inefficiency. This makes computer design challenging, considering the lack of support from workload characteristics, programming languages through compilers. The key challenges and consequences of dark silicon are summarized in the following sections.

1.3.1 Performance

With a section of the chip being inactive, expected performance from many-core systems can never be realized in practice. In the quest for higher performance, designers opt for denser chips, with the expectation of increased performance. Transistor scaling has paved the way to integrate more cores at a lower power consumption, effectively increasing compute capacity per area [2]. Inherent parallelism, if present, in workloads could take advantage of multi and many-core computers to show improved performance to a certain extent [24]. Dark silicon changes this consensus, since we cannot power up all the available resources at any given time [25]. Figure 1.5 shows cyclic dependency of performance requirements and dark silicon, where integration - the solution to high performance becomes the subsequent cause for low performance. Consequently, there is no significant performance gain from technology node scaling and moving into the future, there is no clear direction on attacking this problem [26].

1.3.2 Energy Efficiency

Extensive work has been done towards energy efficiency in computer systems [27], [28], [29], [30], although they are restricted to dynamic power component. Deterioration in performance due to dark silicon also affects energy efficiency of the system [31]. With static power dominating the total power consumption, achievable performance in a fixed energy budget i.e., *performance per watt*, is declining. Around half of the energy budget is wasted towards static power, literally implying spending energy for doing nothing. Assuming that transistor scales down by 30% in physical dimensions, voltage and area being scaled the same way, energy savings should be as much as 65% [14]. However, voltages and subsequently power densities are no longer scaling better, thus also limiting the possible energy gains.

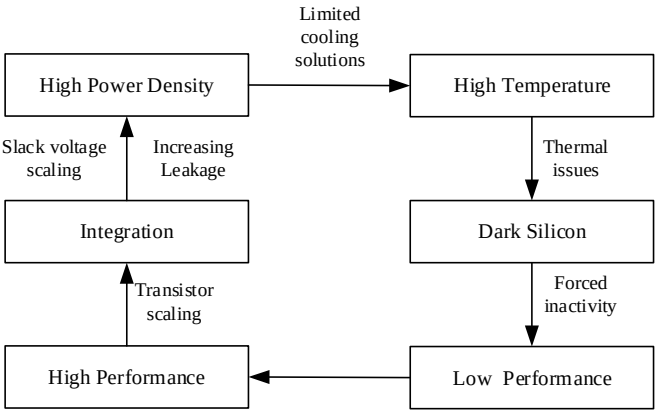

Fig. 1.5: Performance degradation due to dark silicon

1.3.3 Resource Allocation and Utilization

Applications that are designed at a higher level (e.g., algorithms, programming languages, compilers) and run-time and operating systems that allocate resources to these applications are blind sided since dark silicon is only visible at a lower level. Run-time resource allocation can influence performance and energy efficiency of many-core systems significantly [32]. A naive allocation can result in unbalanced load across the chip, which potentially might lead to hotspots. State-of-the-art run-time resource management techniques for many-core systems target benefits in terms of network performance, minimizing congestion, system throughput, power optimization, etc. [33][34][35], without considering the dark silicon scenario at all. Subsequently, schedulers allocates tasks to hardware expecting them to finish efficiently, while inefficiency is inevitable. As a result, most of the system resources go under-utilized and poorly allocated, limiting both performance and energy gains.

1.3.4 Thermal Management

High power density and accumulation of heat results in an increase of average and peak temperatures of the chip. This opens the window for manifesting soft errors and unreliability in computation [36], in addition to decrease in life time of the chip [37]. This ages the chip faster and puts an additional penalty on manufacturing cost per reliable chip. Another important aspect is the exponential dependence of gate oxide leakage on temperature. A chip working at full throttle can dissipate more heat and thus high temperature, which in turn increases leakage too [20]. Again, the omnipresent cyclic dependency of temperature, leakage, performance and energy

efficiency surface here. Further, dissipating heat requires a heat sink and a fan that blows air through it, adding to manufacturing costs. For a data centric level compute platform, more sophisticated cooling solutions have to be used, further increasing maintenance costs.

1.4 Solutions for Dark Silicon

Architectural and run-time management techniques can mitigate dark silicon to a large extent and improve energy efficiency. In this book, some of the potential directions and solutions to attack, mitigate and exploit dark silicon are discussed from an architectural and run-time management perspectives. Figure 1.6 summarizes a typical heterogeneous multi-core/ many-core computing platform implementing design time and run-time optimization techniques tailored for dark silicon era. A NoC-based many-core system is chosen to cater wider implementation scenarios with future application characteristics in view, while traditional bus-based multi-core system is used to better represent heterogeneity. The *Architecture and Design* phase is a design time approach, based on architectural asymmetry and heterogeneity of cores, task specific accelerators that are expected to run energy efficiently. All the three phases augment each other in avoiding, mitigating or minimizing dark silicon.

1.4.1 Architecture and Implementation Perspective

Operating cores with same instruction set architecture (ISA) but different micro-architectures, possibly at different voltage and frequency levels results in asymmetric cores with different power-performance characteristics. Such an asymmetric multi-core system offers the choice of executing specific application on a core with suitable power-performance characteristics that can improve energy efficiency.

As of now, multi-core systems have been using more common conventional methods such as DVFS, clock and power gating for power management. As an advancement, near threshold computing (NTC), where the chip's supply voltage is scaled down to as low as threshold voltage offers ultra low power consumption, at the loss of performance. The cores that operate at low power and performance have penalty on per-core performance, yet the overall throughput of the chip is better than in case of dark cores that are simply idle. Near threshold computing leads to better per-chip-throughput cores, also termed Dim Silicon. Energy efficiency can be improved by selectively down scaling the voltage, inducing Dim Silicon over dark silicon. *Chapter 2* focuses on aspects of near threshold computing and dim silicon for mitigating dark silicon. It provides *Lumos*, a framework for design space exploration of future many-core systems quantifying dark silicon effect. The framework analytically models technology node scaling effects at near threshold operation.

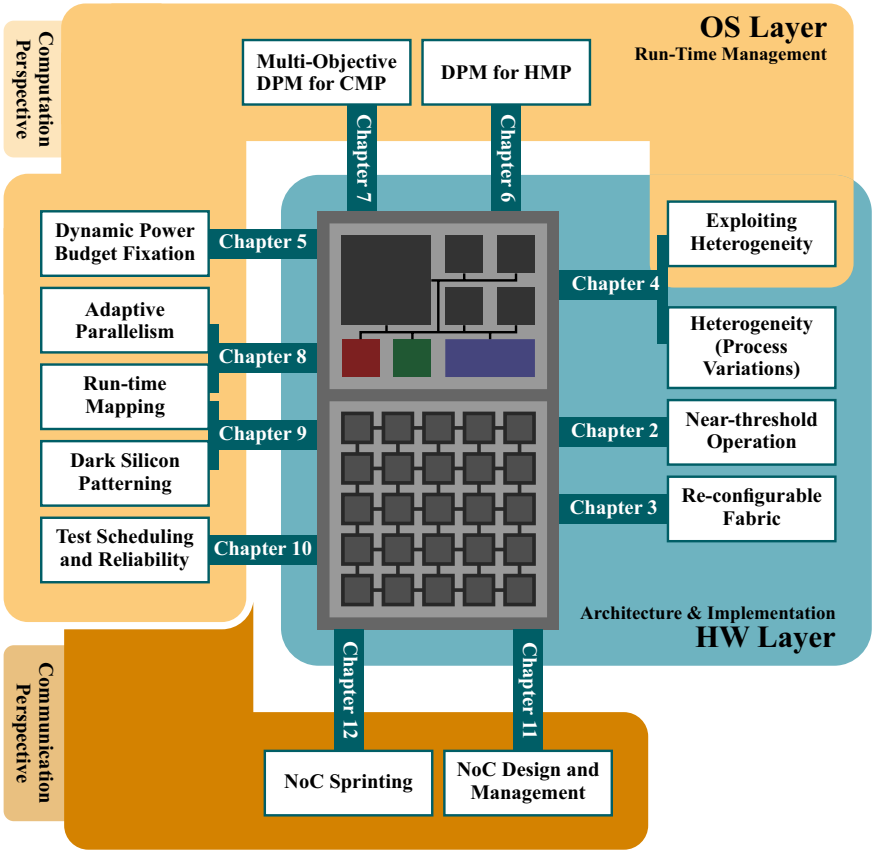

Fig. 1.6: Design and run-time solutions to Dark Silicon

Current compute platforms are still largely homogeneous, calling for innovations from an architectural stand point. Designing cores that are different at micro-architectural as well as instruction set architectural (ISA) levels provides core level heterogeneity. Heterogeneous systems offer highly customized processing units that can run specific tasks and applications at high performance and low energy consumption. Emerging accelerators like neural networks, fuzzy, quantum and approximation paradigms are expected to be dominate future many-core systems. Both asymmetric and its super set, heterogeneous cores, provide customized hardware that can run certain tasks or applications more efficiently at the loss of generality. Yet, appropriate design and choice of asymmetric and/or heterogeneous cores can mitigate dark silicon significantly, over than conventional homogeneous systems that use same cores for even different tasks. Further, variable quality of service (QoS) and throughput requirements of certain applications presents a chance

for coarse and fine-grained (i.e., per-cluster and per-core) dynamic voltage and frequency scaling to enhance overall energy efficiency. *Chapter 3* presents a special case of heterogeneous design leveraging process variations induced by technology scaling. Further, they provide architecturally synthesized heterogeneous platform for minimizing dark silicon, analyzing the parameters of area and power budgets, workloads to be executed and a library of cores.

Design of re-configurable logic at different levels of granularity lends adaptability to the architecture. The compute platform can be re-configured on the fly subject to power consumption and workload characteristics, improving performance and energy efficiency along with area. Further, domain specific custom accelerators implemented on re-configurable fabric can be tightly coupled with traditional cores, making it a special case of heterogeneity. Dark silicon can be minimized by adjusting the granularity and the extent of re-configurability. Further, dynamic compilation techniques that control degree of parallelism and choice of specialized hardware makes the decision of re-configuring hardware more comprehensive. *Chapter 4* proposes use of a coarse grained re-configurable fabric as an alternative to traditional core design for mitigating dark silicon. Hardware - architecture, storage and interconnects are re-configured dynamically in view of workload and resource availability through dynamic compilation and run-time decision making.

1.4.2 Run-time Resource Management - Computational Perspective

Run-time management operates in an observe-decide-act loop to enhance the energy efficiency of underlying hardware. Observe phase involves monitoring critical parameters of chip's performance including instantaneous power consumption, active and dark cores, network congestion, temperature accumulation, etc. The Decide phase uses a pro-active strategy, relying on application mapping, resource allocation and scheduling to reduce potential dark silicon. Act phase employs a reactive strategy that applies different actuators such as clock and power gating, voltage and frequency scaling, hardware re-configuration, task migration, to mitigate dark silicon.

Current computer systems use a conservative estimate of upper bound on power budget, the Thermal Design Power (TDP). However, a fixed TDP value for variable application scenarios results in poor utilization of available power budget. Run-time estimation of power budget subject to the set of applications currently executing on the chip and their spatial alignment can provide a thermally safe upper bound on power budget, variable at run-time. This variable upper bound, Thermal Safe Power (TSP), improves utilization of available power budget and minimizes dark silicon that otherwise manifests from conservative design-time estimate of TDP. *Chapter 5* introduces the concept of TSP, the variable yet safe upper bound on power budget, evaluated dynamically based on current set of applications being executed. It also provides a light weight C instrumentation library that can compute the safe power

budget, Thermal Safe Power (TSP), on the fly as a function of spatial alignment of working cores.

Individual actuators like voltage and frequency scaling including near threshold operation, clock and power gating, task allocation and migration etc., can manage power consumption, occasionally by sacrificing performance. However, a suitable combination of more than one of these knobs will guarantee an acceptable level of performance within a given energy budget. Dynamic power management (DPM) techniques that use the right combination of these knobs can monitor and manage power consumption with a fine grained control. Observation, decision making and control of power by considering all possible knobs and actuators would yield a comprehensive power management platform tailored for the dark silicon era. *Chapter 6* present dynamic power management (DPM) framework for asymmetric and heterogeneous multi-processor systems (HMPs), employing per-cluster dynamic voltage and frequency scaling. Subject to TDP and QoS demands of applications, it further enhances energy efficiency using run-time task migration and task allocation techniques.

Traditional power management techniques were dominantly triggered by instantaneous power consumption, and actuating voltage and frequency. With dark silicon scenario, combination of several critical parameters should be used as knobs to monitor and similarly, combination of several actuators are needed for fine-grained power management. Knobs that are needed to be monitored in addition to power consumption are: upper bound on power budget (TDP/TSP), critical temperature, thermal profile, network congestion, applications' characteristics - arrival, QoS demands, criticality, etc. Adaptive actuators like adjustable degree of parallelism (DoP), computational and interconnect sprinting, application mapping - mapping/queuing the application, contiguity and dispersion, coarse and fine-grained reconfigurability, can be used in addition to traditional actuators like voltage and frequency scaling, clock and power gating. *Chapter 7* propose a comprehensive power management framework for chip multi-processors (CMPs) for the dark silicon era that considers workload characteristics, network congestion, QoS demands, per-core and per-chip power requirements. They use actuators like voltage and frequency scaling, clock and power gating and mapping decisions, within available power budget, maximizing utilization. It is to be noted that resource allocation though run-time application mapping lends a fine grained control on several critical chip parameters and impacts performance and energy efficiency.

Dynamic application mapping involves allocating task-per-core at run-time, given an unpredictable sequence of incoming applications. An efficient mapping effects network parameters like congestion, dispersion, spatial availability, power consumption, etc. With varied limitations of parallelism and sequential bottleneck, different applications run efficiently at a different degrees of parallelism. Also considering heterogeneous systems, it is advantageous to consider task-per-core from the degree of parallelism stand point. *Chapter 8* presents a run-time mapping approach controlling the degree of parallelism of applications to minimize average energy consumption. In addition, they provide reliability and soft error awareness leveraging process variations to find optimal set of cores for mapping an application.

Dark silicon agnostic techniques greedily strive for performance via contiguous mappings. Given variable workload characteristics and application scenarios, aforementioned TSP can be maximized with appropriate dark silicon aware application mappings. Mapping applications sparsely can avoid potential hot-spots and distribute heat accumulation evenly across the chip area, in turn increasing the safer limit of power budget (TSP). Gain in power budget may be used to activate more cores improving performance and resource utilization, over compensating the loss from dispersed mapping. *Chapter 9* proposes dark silicon patterning, an efficient run-time mapping scheme that maximizes power budget and utilization. It employs sparsity among mapped applications to balance heat distribution across the chip and thus maximize utilizable power budget. Inevitable dark cores are patterned alongside active cores providing the necessary cooling effect.

Despite the loss in performance and energy efficiency caused by dark silicon, an alternative approach to attack dark silicon is to exploit it for chip maintenance tasks like testing, reliability and balancing the life time. Testing a core consumes power resources and interrupts applications' execution. Test procedures can be scheduled on specific set of cores based on chip activity such that available power budget is efficiently utilized for both computation and testing, which otherwise is wasted as dark cores. *Chapter 10* exploits dark silicon to schedule test procedures on idle cores. With new application to be mapped, they leverage the possibility of few cores being left idle due to limited power budget, and schedule test procedures on such cores. They propose a mapping technique that selects working cores to evenly balance the idle periods for testing.

1.4.3 Design and Management - Communication Perspective

Interconnection network forms the other major component of a computing platform, accounting for performance latency and power consumption. Power management actuators, like power gating and voltage and frequency scaling, can also be applied to interconnect components at both design time and run-time, giving fine grained control. An optimized network with suitable energy efficient and reliable interconnect components can increase overall energy efficiency of the system. *Chapter 11* proposes *darkNoC*, an adaptive design time energy efficient interconnect for exploiting dark silicon, over conventional voltage and frequency scaling. They also present *SuperNet*, a combinatorial run-time optimizable interconnection approach for executing in either energy efficiency, performance or reliability modes, through voltage and frequency scaling.

Operating a section of the chip at highest possible voltage and frequency for a brief period of execution time is known as computational sprinting. Sprinting can accelerate performance with an increase in power consumption, possibly violating the safe upper bound. However, the sprinting time is kept low enough so that increase in power does not reflect directly in temperature accumulation. Since thermal safety is maintained, relatively high performance and utilization can be achieved despite dark

silicon. Sprinting can be employed to specific tasks or applications particularly on requirements and demands. *Chapter 12* proposes fine-grained interconnect sprinting that activates only a selected number of routers for short burst of heavy communication, instead of activating all the communication resources. Specific communication channels are chosen for this purpose based on application and energy demands.

1.5 Summary

Dark Silicon has become a growing concern for computer systems ranging from mobile devices up to data centers. Architecture, design and management strategies need re-thinking when moving into the future exascale computing. Although dark silicon scenario is a challenge, it still presents opportunities to innovate and pursue new research directions. An overview of dark silicon regime has b een presented in this chapter. State-of-the-art solutions to minimize, mitigate and attack dark silicon are organized into the subsequent chapters of this book. Several interesting discussions and key insights leading to open research directions addressing energy efficiency are covered.

References

1. H. Sutter, "The free lunch is over: A fundamental turn toward concurrency in software," *Dr. Dobb's journal*, vol. 30, no. 3, pp. 202–210, 2005.
2. G. E. Moore, "Cramming more components onto integrated circuits," *Readings in computer architecture*, p. 56, 2000.
3. R. H. Dennard, V. Rideout, E. Bassous, and A. Leblanc, "Design of ion-implanted mosfet's with very small physical dimensions," *IEEE Journal of Solid-State Circuits*, vol. 9, no. 5, pp. 256–268, 1974.
4. M.-H. Haghbayan, A.-M. Rahmani, A. Y. Weldezion, P. Liljeberg, J. Plosila, A. Jantsch, and H. Tenhunen, "Dark silicon aware power management for manycore systems under dynamic workloads," in *Proc. of IEEE International Conference on Computer Design*, 2014, pp. 509–512.
5. H. Esmaeilzadeh, E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in *Proc. of IEEE International Symposium on Computer Architecture*, 2011, pp. 365–376.
6. J. Kong, S. W. Chung, and K. Skadron, "Recent thermal management techniques for microprocessors," *In ACM Computing Surveys*, pp. 13:1–13:42, 2012.
7. Semiconductor Industry Association, "International technology roadmap for semiconductors (ITRS), 2011 edition," 2011.
8. H. Esmaeilzadeh, E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," *IEEE Micro*, no. 3, pp. 122–134, 2012.
9. Semiconductor Industry Association, "International technology roadmap for semiconductors (ITRS), 2013 edition," 2013.
10. S. Borkar, "The exascale challenge," *In Proc. of VLSI Design Automation and Test*, 2010.
11. G. Venkatesh, J. Sampson, N. Goulding, S. Garcia, V. Bryksin, J. Lugo-Martinez, S. Swanson, and M. B. Taylor, "Conservation cores: reducing the energy of mature computations," in *ACM SIGARCH Computer Architecture News*, vol. 38, no. 1, 2010, pp. 205–218.
12. T. Mudge, "Power: A first-class architectural design constraint," *Computer*, no. 4, pp. 52–58, 2001.
13. S. Borkar, "Getting gigascale chips: Challenges and opportunities in continuing moore's law," *Queue*, vol. 1, no. 7, p. 26, 2003.
14. S. Borkar and A. A. Chien, "The future of microprocessors," *Communications of the ACM*, vol. 54, no. 5, pp. 67–77, 2011.
15. H. Sasaki, M. Ono, T. Yoshitomi, T. Ohguro, S.-i. Nakamura, M. Saito, and H. Iwai, "1.5 nm direct-tunneling gate oxide si mosfet's," *IEEE Transactions on Electron Devices*, vol. 43, no. 8, pp. 1233–1242, 1996.
16. S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi, and V. De, "Parameter variations and impact on circuits and microarchitecture," in *Proc. of ACM Design Automation Conference*, 2003, pp. 338–342.
17. N. S. Kim, T. Austin, D. Baauw, T. Mudge, K. Flautner, J. S. Hu, M. J. Irwin, M. Kandemir, and V. Narayanan, "Leakage current: Moore's law meets static power," *IEEE Computer*, vol. 36, no. 12, pp. 68–75, 2003.
18. A. P. Chandrakasan, W. J. Bowhill, and F. Fox, *Design of high-performance microprocessor circuits*. Wiley-IEEE press, 2000.
19. C.-H. Choi, K.-Y. Nam, Z. Yu, and R. W. Dutton, "Impact of gate direct tunneling current on circuit performance: A simulation study," *IEEE Transactions on Electron Devices*, vol. 48, no. 12, pp. 2823–2829, 2001.
20. G. Sery, S. Borkar, and V. De, "Life is cmos: Why chase the life after?" in *Proc. of ACM Design Automation Conference*, 2002, pp. 78–83.
21. "Intel Corporation. Intel Xeon Processor - Measuring Processor Power, revision 1.1," in *White paper, Intel Corporation*, April, 2011.
22. "Intel Corporation. Fourth Generation Mobile Processor Family Data Sheet," in *White paper, Intel Corporation*, July, 2014.

23. "AMD. AMD Kaveri APU A10-7800," accessed: 2015-02-28. [Online]. Available: http://www.phoronix.com/scan.php?page=article&item=amd_a_45watt&num=1
24. G. M. Amdahl, "Validity of the single processor approach to achieving large scale computing capabilities," in *Proc. of ACM Spring Joint Computer Conference*, 1967, pp. 483–485.
25. G. Venkatesh, J. Sampson, N. Goulding-Hotta, S. K. Venkata, M. B. Taylor, and S. Swanson, "Qscores: Trading dark silicon for scalable energy efficiency with quasi-specific cores," in *Proc. of IEEE/ACM International Symposium on Microarchitecture*, 2011, pp. 163–174.
26. H. Esmaeilzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger, "Power limitations and dark silicon challenge the future of multicore," *ACM Transactions on Computer Systems*, vol. 30, no. 3, pp. 11:1–11:27, 2012.
27. Y. Almog, R. Rosner, N. Schwartz, and A. Schmorak, "Specialized dynamic optimizations for high-performance energy-efficient microarchitecture," in *Proc. of International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization*, 2004, pp. 137–148.
28. K. Khubaib, M. A. Suleman, M. Hashemi, C. Wilkerson, and Y. N. Patt, "Morphcore: An energy-efficient microarchitecture for high performance ilp and high throughput tlp," in *Proc. of IEEE/ACM International Symposium on Microarchitecture, 2012*, pp. 305–316.
29. G. Semeraro, G. Magklis, R. Balasubramonian, D. H. Albonesi, S. Dwarkadas, and M. L. Scott, "Energy-efficient processor design using multiple clock domains with dynamic voltage and frequency scaling," in *Proc. of IEEE International Symposium on High-Performance Computer Architecture, 2002*, pp. 29–40.
30. D. M. Brooks, P. Bose, S. E. Schuster, H. Jacobson, P. N. Kudva, A. Buyuktosunoglu, J.-D. Wellman, V. Zyuban, M. Gupta, and P. W. Cook, "Power-aware microarchitecture: Design and modeling challenges for next-generation microprocessors," *IEEE Micro*, vol. 20, no. 6, pp. 26–44, 2000.
31. M. B. Taylor, "Is dark silicon useful?: Harnessing the four horsemen of the coming dark silicon apocalypse," in *Proc. of ACM Design Automation Conference*, 2012, pp. 1131–1136.
32. E. L. de Souza Carvalho, N. L. V. Calazans, and F. G. Moraes, "Dynamic task mapping for mpsocs," *IEEE Design & Test of Computers*, vol. 27, no. 5, pp. 26–35, 2010.
33. C.-L. Chou and R. Marculescu, "Contention-aware application mapping for network-on-chip communication architectures," in *Proc. of IEEE International Conference in Computer Design*, 2008, pp. 164–169.
34. M. Fattah, M. Daneshtalab, P. Liljeberg, and J. Plosila, "Smart hill climbing for agile dynamic mapping in many-core systems," in *Proc. of IEEE/ACM Design Automation Conference*, 2013.
35. M. Fattah, P. Liljeberg, J. Plosila, and H. Tenhunen, "Adjustable contiguity of run-time task allocation in networked many-core systems," in *Proc. of IEEE Asia and South Pacific Design Automation Conference*, 2014, pp. 349–354.
36. J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, "The case for lifetime reliability-aware microprocessors," in *Proc. of IEEE International Symposium on Computer Architecture*, 2004, pp. 276–287.
37. J. Srinivasan, S. Adve, P. Bose, and J. Rivers, "The impact of scaling on processor lifetime reliability," in *Proc. of International Conference on Dependable Systems and Networks*, 2004.