

Thermal and Performance Efficient On-chip Surface-wave Communication for Many-core Systems in Dark Silicon Era

AMMAR KARKAR, Department of Electronic and communication engineering, University of Kufa, Iraq

NIZAR DAHIR, College of Information Engineering, Al-Nahrain University, Iraq

TERRENCE MAK, Faculty of Physical Sciences and Engineering, University of Southampton, UK

KIN-FAI TONG, Department of Electrical and Electronic Engineering, University College London, UK

Due to the exceedingly high integration density of VLSI circuits and the resulting high power density, thermal integrity became a major challenge. One way to tackle this problem is Dark silicon. Dark silicon is the amount of circuitry in a chip that is forced to switch off to insure thermal integrity of the system and prevents permanent thermal-related faults. In many-core systems, the presence of Dark Silicon adds new design constraints, in-general, and on the communication fabric of such systems, in particular. This is due to the fact that system-level thermal-management systems tend to increase the distance between high activity cores to insure better thermal balancing and integrity. Consequently, a designing dilemma is created where a compromise has to be made between interconnect performance and power consumption. This study proposes a hybrid wire and surface-wave interconnect (SWI) based Network-on-Chip (NoC) to address the dark-silicon challenge. Through efficient utilization of one hop cross the chip communication SWI links, the proposed architecture is able to offer efficient and scalable communication platform in terms of performance, power and thermal impact. As a result, evaluations of the proposed architecture compared to baseline architecture under dark silicon scenarios show reduction in maximum temperature by 15 °C, average delay up to 73.1%, and energy saving up to ~3X. This study explores the promising potential of the proposed architecture in extending the utilization wall for current and future many-core systems in dark silicon era.

CCS Concepts: • **Hardware** → **Network on chip**; *Very large scale integration design*.

Additional Key Words and Phrases: Networks-on-chip, Dark silicon, Surface wave, Many-core systems, On-chip interconnects, Thermal reliability, Communication efficient, Multicast.

ACM Reference Format:

Ammar Karkar, Nizar Dahir, Terrence Mak, and Kin-Fai Tong. 2021. Thermal and Performance Efficient On-chip Surface-wave Communication for Many-core Systems in Dark Silicon Era. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Even though CMOS technology has enabled the integration of many devices in single chip such as in the case of many-core systems, the supply voltage and the threshold voltage of transistors are not scaling proportionally to their sizes as it has been since the 60's [29]. Therefore, predictions show that power density are expected to exponentially scale with technology due to increase in power leakage and number of devices in the die. These factors have a drastic effect on the integrated devices lifetime and could lead to system failure. Therefore, power became more expensive

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104

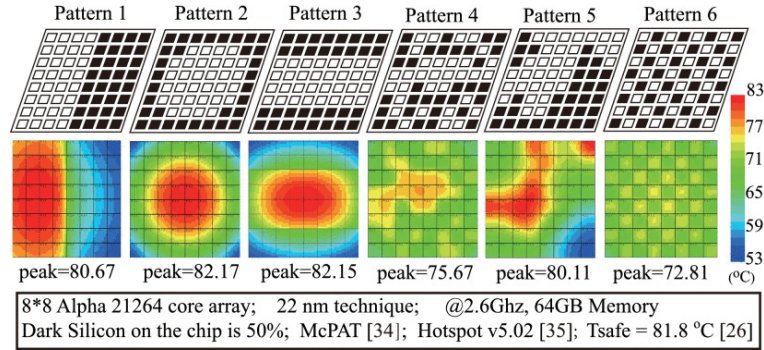


Fig. 1. Various Dark silicon patterns. [21].

than the number of integrated devices in the chip. This has led to dark silicon where some integrated devices have to be turned off to stay within the power and temperature limits [1, 26]. Consequently, the many-core systems have hit the utilization wall due to fact that some of the cores computation power cannot be harvested since they are being deactivated or running with low activity mode [29].

In general, the system components that have to be turned off should be distributed uniformly to achieve homogenous thermal impact and reduce the average and maximum temperatures of the die. Also the components with higher switching activities than the other components are kept apart. For instance, Fig. 1 shows various patterns of mapping 50% of the active processing elements (PE) along with the resulted thermal impact [21]. Obviously, patterns such as pattern 6 where active cores are kept separated have the least peak and average temperature. This is due to the fact that heat dissipation is better when active cores are surrounded by dark cores. Therefore, most mapping techniques that aim to reduce on-chip temperature map tasks with high computations apart. These high computation activity cores usually have high communication activity as well.

In contrast, energy-aware task mapping algorithms that aim to reduce inter-core communication cost tries to place active cores near each other to improve performance and reduce interconnect energy [23]. Otherwise, mapping active cores away from one another in the dark silicon domain increases the average hop counts, causes the performance of interconnects to deteriorate, and escalates power issues [5]. This dilemma is clearly shown in our simulation of the average-delay, see Fig. 2, in which PEs processed the same number of flits (injected and drained) but with one case having all PEs operational and with another case having 50% of PEs turned off in pattern similar to pattern 6. Obviously, the increase of average hope counts caused the network to saturate faster and the average delay increased for dark many-core system as the packet injection rate increased compared to normal operation many-core system. Therefore, the need for an interconnect architecture that is efficient for global and semi-global communication is considered as one of the main challenges for the dark silicon era [21, 27].

This study proposes a hybrid wire and surface wave interconnects (SWI) architecture for the many-core systems that efficiently addresses the dark silicon interconnects dilemma. Such architecture that has been proven to show significant improvements for unicast/multicast global and semi-global communication could be the answer for thermal reliability and performance efficiency [17, 19, 20]. The main contributions of this study are:

- Proposing an efficient hybrid NoC architecture that satisfies dark-silicon era communication requirements using SWI.

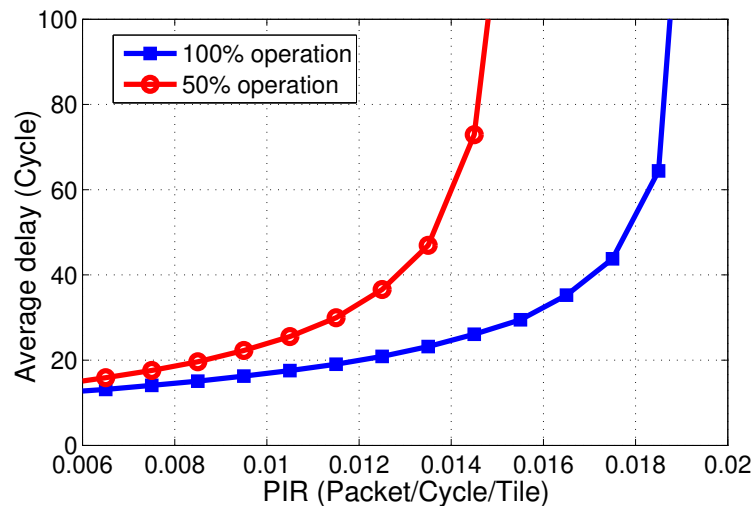


Fig. 2. A 6×4 regular mesh NoC simulations with random traffic that compare the average delay for two cases with same total drained flits: (1) where all the processing elements are active (2) where only 50% of processing elements are active.

- Designing a distance-based weighted-random-robin arbitration technique to improve SWI utilization under dark silicon operation mode.
- Evaluating the proposed architecture in terms of communication latency, thermal stability, power consumptions and area overhead under dark-silicon scenarios.

2 RELATED WORK

Many studies has highlighted the thermal and power scaling issues for systems on chip (SoC) that led to dark silicon scenario and the consequential challenges [26, 29]. One of these main challenges is to efficiently manage the on-chip communication requirements under power and thermal constrains without performance degradation [26]. Therefore, to address this challenge, a thermal-aware routing techniques has been proposed to mitigate the thermal impact of the NoC itself [11]. In addition, other studies has approached this challenge by focusing on reducing the thermal impact of the NoC hardware components by designing them to be more power efficient [9, 33]. However, such designs not just sacrifice the NoC performance, but also does not address the new traffic requirements imposed by the dark silicon domain.

On the other hand, some related work tried to address the conflict between on-chip communication efficiency while guarantee the system temperature reliability [5, 21]. For instance, Liu *et.al.* proposed a thermal-aware task mapping along with reconfigurable NoC, referred to as SMART, to bypass the multi-hop links that is required by separated active cores [21]. However, this NoC is still limited by the underlying costly wire links. Although, the wires are still the cheapest interconnects, they have shown that global and semi-global wires do not scale with the technology in terms of signal latency and dissipate nontrivial energy (J/b) [32]. Even when repeaters are introduced for global communication, they only mitigated the delay issue by making the delay rises linearly with distance. Moreover, the repeaters magnify the power issues of the on-chip interconnects. Consequently, Some studies tried to reduce the repeaters numbers and size with some delay penalty [2]. The drawback of such solutions is neither the power nor latency are optimum. Therefore,

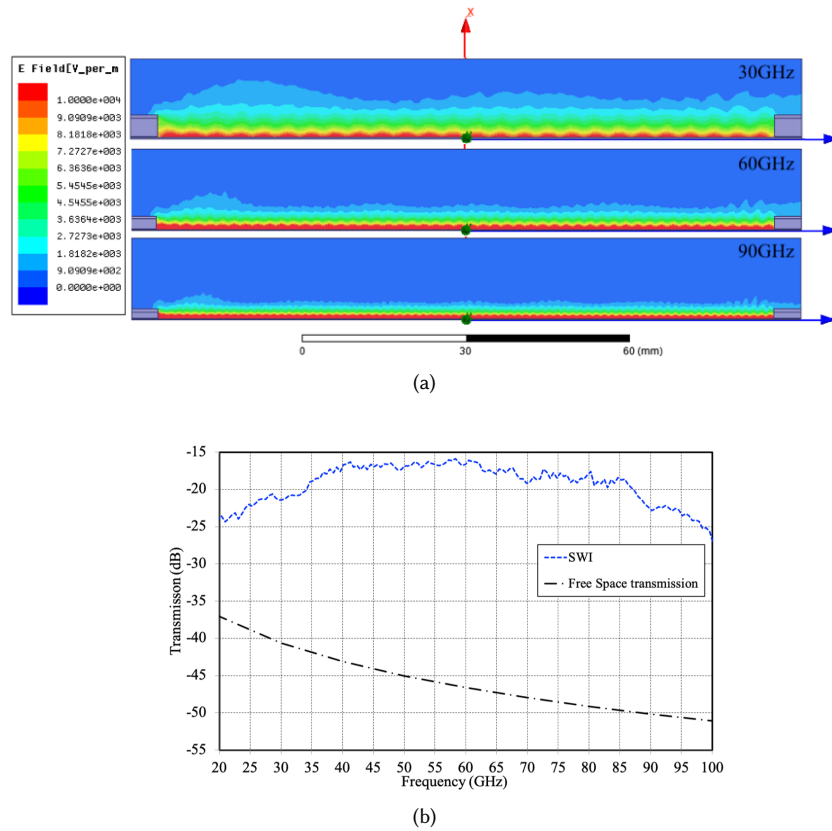


Fig. 3. (a) Trapped surface wave signal decay, (b) transmission (S21) of the signal traveling 140 mm in the millimeter-wave frequency band on the SWI and free space [31]

many emerging interconnects has been proposed to replace or supplement the wire links for future many-core systems [18].

3 PROPOSED SURFACE-WAVE-BASED NOC

3.1 Trapped Surface Wave Interconnect

The trapped surface wave (SW) is a flat electromagnetic (EM) wave that is guided by a surface [31]. This surface is designed as a waveguide; thus, instead of three dimensional free space, the EM is trapped in a two dimensional medium. Consequently, the SW decay rate horizontally along the surface is around $(1/\sqrt{d})$, where d is the distance from the source [14], see Fig. 3. This make the SW one of the most efficient emerging interconnects for multicast and global communication [18].

The the characteristic impedance ($Z_0 = 10 + j300 \Omega$) of the surface is obtained by altering its dimensions and its materials. This technology has already been implemented and tested on PBC for inter-chip communication using switched-beam end-fire planar array with Integrated 2-D butler matrix [3]. In this study, the chosen surface is engineered from a dielectric that is coated with a conductor for low fabrication costs and simple geometry [30, 31]. For instance, an




Fig. 4. Integrated transceiver and integrated transducer (inverted quarter-wavelength monopole) stacked over the designed surface.

integrated waveguide surface for millimeter-wave applications could be manufactured using metal ground sheet of thickness $1\mu\text{m}$. Also, to provide the $10 + j300\ \Omega$, this sheet is coated with thermoplastic substrate layer such as Acrylic resin ($\epsilon_r = 4.5$, thickness = 0.5mm) for better thermal dissipation. Such process is not costly since the coating process is part of conventional CMOS fabrication process. In addition, there is no need for expensive highly polished dielectric wafer ($Ra < 0.01\mu\text{m}$) because the targeted carrier frequency is less than 300 GHz.

In order to have multi-channel based on frequency-division-multiple-access (FDMA) a wide range of frequencies is required. Assuming the use of proper transducer, the SWI frequency bandwidth is only limited by the transceiver capabilities [31]. Thus, since carrier frequencies of integrated transceiver are projected to scale with the CMOS technology switching speed, the required SWI channels frequency bandwidth with the necessary frequency spacing can be achieved. For instance, an integrated transceiver that is proposed and implemented by Chang *et al.* [7] or Carpenter *et al.* [6] could be adopted. In addition, to match the data bandwidth of the baseline architecture wire link, each channel is designed to have 32 sub-channels, which use 16-QAM (quadrature amplitude modulation) to transmit 4 bits. Further details of this communication channel specifications are presented in previous work [19].

In order to launch the modulated signal via the waveguide surface, an integrated transducer is linked to the transceiver [31], as shown in Fig. 4. Such transducer could be designed with a parallel plate waveguide along with monopole or dipole for omni-directional communication [22]. In order to link such transducer layer with transceiver (after it has been fabricated separately), flip-chip bonding and through-silicon-via (TSV) technique could be used. The transceiver and transducer design is beyond the scope of this study.

3.2 Hybrid Wire and Surface-Wave NoC

Designing an efficient interconnect architecture that guarantees thermal reliability should aim to reduce the power cost of the communication. This is one of the significant advantages of SWI since its ability to provide cross the chip communication eliminating the need to go through power hungry links and router. However, like all wireless NoC (WNoC), it is limited by the shared media, which creates congestion over its available frequency bandwidth. Moreover, in terms of wire-based interconnects, local communication seems to scale well with technology scaling unlike global communication [25]. In addition, this type of interconnects has the cheapest implementation cost compared to other interconnect fabrics. Consequently, instead of replacing metal wire NoC completely, the best solution would be to combine both metal and SWI. The resulted multi-layered network architecture will be hybrid interconnects of regular NoC and SWI, which we will refer to in this study as SWI-NoC, as shown in Fig. 5. The chosen first layer is a mesh topology, since it is suitable for a general purpose and local interconnects, its wire links have uniform lengths, and more convenient for chip floor planning. The second layer is the SWI that represents a bus topology. Thus, this

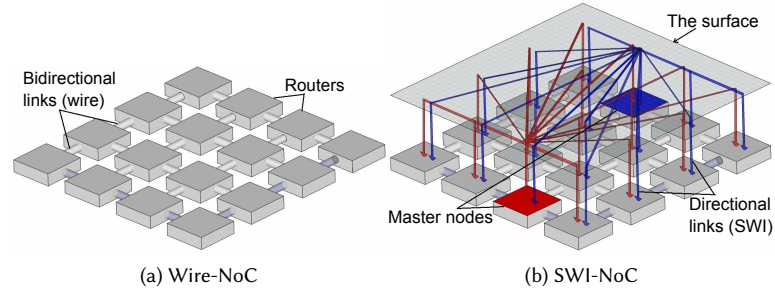


Fig. 5. (a) Wire-NoC: wire-based NoC layer with 4-ary 2-mesh topology; (b) SWI-NoC: hybrid Multilayer-network with two SWI channels and mesh NoC.

architecture would offer one-hop cross the chip communication, which is what non-adjacent cores need to communicate in dark-silicon case.

In this hybrid architecture, in order to link the mesh NoC to the SWI, a sixth port is needed for the routers along with all related control circuits plus the increase of the crossbar size. all the routers in the NoC are linked to a receiver to receive traffic via SWI. However, to reduce contention by dedicating a frequency bandwidth for each transmitter and to reduce area overhead, fewer routers are selected to be linked to transmitters. These routers with TxRx capability are referred to as masters, while the rest with only Rx capability are referred to as slaves. These masters has been distributed with the aim of minimizing average Manhattan distance from all slaves to the nearest master. This master placement should reduce the the average-hop-count of the resulted NoC. Moreover, each master node would be reachable via shortest path of costly routers and wires.

3.3 Utilizing SWI-NoC for global communication

This section presents a Distance-based weighted-random-robin (DW) arbitration algorithm, which aims to improve the utilization of SWI interconnects while avoiding creating a bottleneck. This is an improved arbitration technique from the one proposed in previous study to handle the existence of multicast traffic and dark cores [19]. The basic idea of this technique is: an arbitration between SWI port and wire ports (N, E, W, and S) is carried out based on the calculated weight (W) by eliminating one of these two options. This elimination is done on output options presented by the routing algorithm and before the final forward port of the flit is determined by the arbitration algorithm. Based on distance between the current router (master node) and the destination router, the DW algorithm would select SWI or not. This arbitration by the DW algorithm (see Algorithm.1) is trying to give a higher priority to the SWI as the distance increases. Consequently, the power saving will be increased, the thermal impact would be reduced, and network performance (throughput and delay) will be improved. Moreover, the DW will keep the SWI likely available for the flits need to travel long distances, in terms of hop counts. However, in case of multicast traffic, the arbitration should always chose SWI to utilize its fan-out merits. This algorithm assigns (0) weight (W_0) for a distance of one hop and certain start up weight (W_1) for two hops. This weight increases linearly until it reach (100%) for the maximum possible distance, which in case of mesh NoC is equal to half the network diameter.

Fig.6 shows the implementation of the proposed DW algorithm where set of circular shift register (CSR) are used to store codes that represent the W for each destination node. These weights have been calculated as shown in Algorithm.1

Algorithm 1: Distance-based weighted-random-robin arbitration algorithm (DWA).

Data: X = Network dimension in direction X , Y = Network dimension in direction Y , d = distance to destination by (hop), M = one-to-many traffic flag, P = set of possible output ports linked to wire interconnect.

Result: Chosen output port(s)

```

1 if ( $d > 1$ ) then
2   |  $W = (\frac{W_1}{(X+Y)-2} \times (d - 2)) + (100 - W_1)$ ;
3 else
4   |  $W = W_0$ ;
5 end
6 Circular Shift Right( $W$ );
7 if ( $W[0] = 0$  or  $M = 1$ ) then
8   | return: Surface wave channel;
9 else
10  | return:  $C$ , where  $C \in P$ ;
11 end

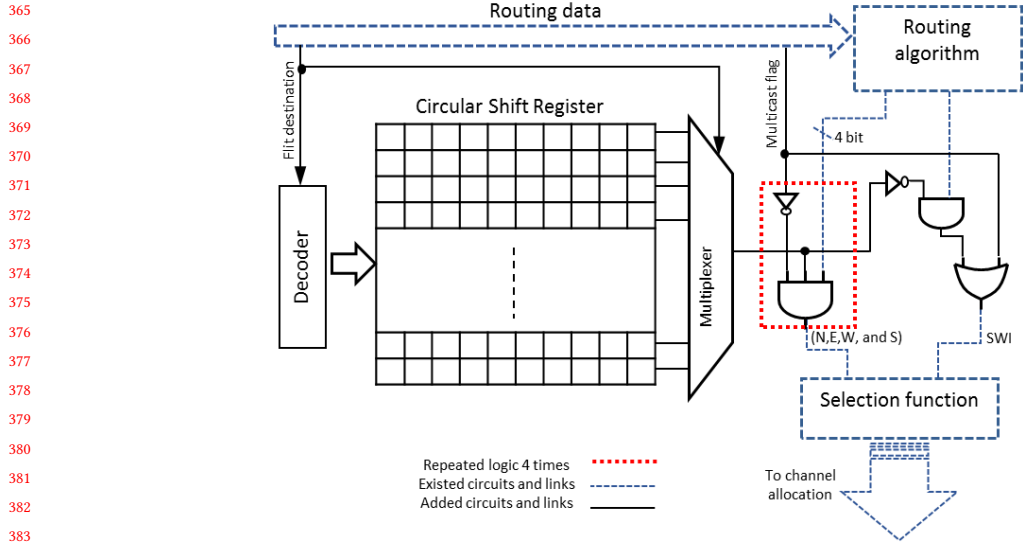
```

(lines: 1-5) and stored in each master router in the design time. For instance, to give SW port a weight of 60% and wire ports a weight of 40%, the code would be: (1001001001) in CSR. As a result, when the CSR shift right after each access and the value of the least significant bit (LSB) is zero, the traffic will be forward via SWI port only. Otherwise, if LSB value is one, it will forward the packet via ports linked to wire links (N, E, S, and W). The right weight word stored in CSR is accessed by decoding the packet destination. Moreover, flit destination is used as switching signal for the multiplexer, as shown in Fig.6). Consequently, the number of CSR is equal to $(N-1)$, where N is the total number of routers in the NoC. On the other hand, the size of CSR depends on the weight precision required. The circuit components and wires drawn in solid line in Fig.6 are the extra circuitry added to the routing unit to implement the proposed DW arbitration. Moreover, this hardware is required only for in a master routers. Thus, the power increase is almost negligible, which crucial in dark silicon domain.

3.4 Utilizing SWI-NoC for multicast communication

Multicast traffic existence is inevitable in many-core systems due to cache coherence protocols. Moreover, in special cases such as artificial spiking neural network (SNN) the multicast ratio reach 100%. In such systems, the multicast destination group members is relatively large, especially if each PE simulated large number of neurons $(N - 1)$, where N is the number of PE. Even with dark silicon where N is getting smaller by reducing the number of active cores, the challenge of one-to-many communication is increased in order to communicate with all the relatively distant PEs. Therefore, the proposed architecture should support high physical channel fan-out to match the high graph degree communication.

The SWI-NoC address the multicast requirement since the SWI offers natural fan-out physical interconnect layer. However, for better utilization of the interconnect fabric, the a multicast routing and contention solving schemes are needed. Thus, an improved tree-based multicast routing is proposed where the one-to many traffic forks only at the nearest master from the source. This single branching point will then utilize the SWI high fan-out feature by transmitting in one hope to all multicast destinations. This is achieved by implementing the following phases; Phase 1: Route all the multicast traffic to the nearest master node using any simple deadlock-free routing algorithm. Phase 2: If



384 Fig. 6. DW arbitration implementation in every master node, where CSR provide the round robin functionality and store the weight code.
385
386

387
388 Table 1. The parameters adopted for the NoC-based multi/many-core tile. This system tile size is $3.6 \times 5.2 \text{ mm}^2$.
389

390

PE components	NoC components	
Two Pentium™ class IA-32 cores	Message passing router	4-port to neighbor routers, 1-port to local cores and 1-port for the SWI. 6 buffers each with 4×128 bit (4 flits)
Two 256 KB private L2 caches	Links	5 bidirectional wire interconnects (16 byte width). 1 surface wave channel (Rx or TxRx), 32 sub-channel with 16-QAM modulation.

391
392
393
394
395
396
397
398
399
400

401 the local PE of the master router is a member of the multicast destination group, the multicast traffic is not allowed to
402 allocate the router local port until it has been forwarded and released the SWI port.
403

404 This routing scheme manages one-to-many communication efficiently and avoid deadlock, for the following reasons:
405 Firstly, by Phase 1, each router handle multicast traffic basically by forwarding such traffic to the nearest master using
406 any simple deadlock-free routing algorithm. For instance, odd-even could be chosen since it is partially adaptive, simple,
407 deadlock free, and offer path diversity [10]. Hence, there is no need for additional complex hardware or complicated
408 algorithms to build the multicast routing tree and to determine the branching points. Moreover, the routing tree would
409 have only one branching point (the nearest master node). In this master node, packets will be replicated only at the
410 destination routers due to the fan-out feature of the SWI. Secondly, Phase 2 avoids multicast dependency between
411 wire-layer and SWI-layer that leads to deadlock. The SWI-layer multicast dependency issues had been already addressed
412 in previous work where a centralized and decentralized arbitration techniques have been suggested [17, 20]. The next
413 section will evaluate the one-to-many communication capability under dark silicon operation mode.
414
415
416

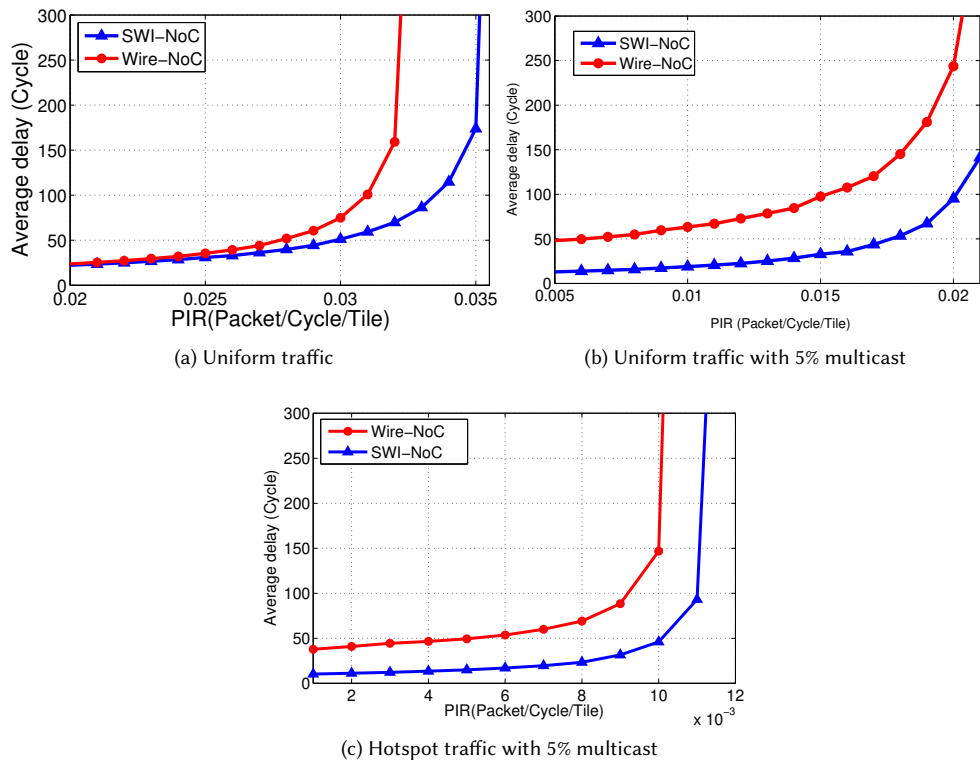


Fig. 7. SWI-NoC improvements over Wire-NoC in terms of average delay under various synthetic traffic and 50% of PEs are dark.

4 EVALUATION

In order to evaluate the merits of the proposed design in conditions of the dark silicon, we used a cycle accurate NoC-simulator that is based on Noxim [12]. This simulator has been integrated with Hotspot 6.0 [34] to measure the thermal impact of the NoC and the PEs. Single chip cloud computing (SCC) tile is adopted as the baseline architecture in this study [24]. Table 1 summarize the tile altered specifications for the SWI-NoC. In this simulator the power consumption, and therefore its thermal impact, of each PE is determined based on bandwidth-based Rent's rule [13]. According to this rule, the bandwidth is proportional to the circuit activity. Since the circuit activity is directly linked to its dynamic power consumption, we calculated the PE power based on PE interface bandwidth and based on the reported power consumption PE ratio compared to the routers in SCC chip [24].

4.1 Performance evaluation

Many-core systems performance becoming more and more dependents on NoC performance to the degree that many-core systems became communication-centric rather than computation-centric [4]. However, as mentioned earlier, due to thermal constrains, the active PEs need to be distant from other active PEs as much as possible. This has increased communication cost and degrade its overall performance. As a result, in this section, average delay of proposed SWI-NoC under dark-silicon circumstances is evaluated against regular Wire-NoC.

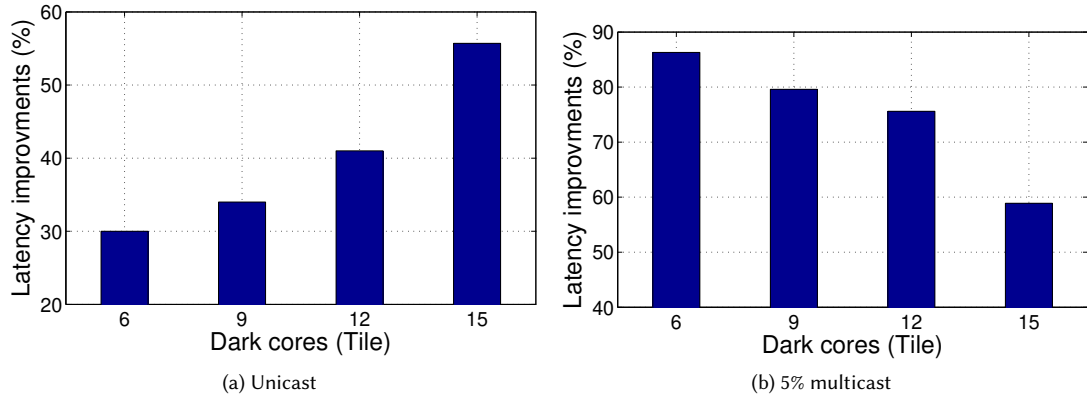


Fig. 8. 6×4 SWI-NoC average delay improvements over Wire-NoC with different number of dark cores under uniform traffic with and without multicast.

Figure 7 compares between SWI-NoC and Wire-NoC under synthetic traffic: uniform, uniform with 5% multicast and hotspot with 5% multicast, where 50% of the cores are dark and distributed uniformly. The multicast percentage is chosen based on cache coherence protocols communication requirements in many-core systems which is found to be between 3.1% to 12.4% multicast traffic [15]. On the other hand, the hotspot traffic places the 4 cores hotspot at the corners to have better thermal dissipation. Obviously, the proposed architecture shows significant improvements in performance since it has been designed to cope with the needs for global and semi-global communication. In contrast, wire-NoC suffer from costly multi-hop communication since the active cores that need to communicate are separated to reduce its thermal impact as mentioned earlier. Thus, for uniform traffic, Wire-NoC saturated much faster than SWI-NoC. In the case of exiting multicast traffic, even for zero-load-latency (ZLL), the average delay of Wire-NoC is approximately four times for uniform and hotspot traffic (47.9 and 37.8 cycle, respectively) that of SWI-NoC (12.9 and 10.2 cycle, respectively). This latency increases exponentially as the packet injection rate (PIR) increases, as shown in Fig. 7. This is due to the fact that the propose architecture achieves double efficiency for multi-hop and multicast communication.

In addition, Fig. 8 shows the performance improvements of SWI-NoC over Wire-NoC under two traffic scenarios: In the first under uniform traffic, the latency improvements of SWI-NoC over Wire-NoC are increased as the dark silicon size increased. In contrast, in the second under uniform traffic with 5% multicast, the SWI-NoC improvements decrease as the dark silicon increased. This is because of the silver lining of dark silicon which decreases the multicast destination groups required for cache consistency [17]; thus reducing source hotspots on the baseline architecture. Nonetheless, communication performance improvements is proportional to the multicast percentage. Consequently, the NoC-SWI optimum not just for general purpose applications but also for SNN under the dark silicon domain.

4.2 Thermal evaluation

Thermal issues are what motivated the need for dark-silicon. Therefore, any proposed interconnect architecture needs to be also thermal efficient. In this subsection the thermal impact is been evaluated for both the baseline architecture (Wire-NoC) and the proposed architecture (SWI-NoC). Firstly, for normal operation where all the processing elements are active, Fig. 10 shows that the spatial thermal impact of SWI-NoC are significantly lower than the wire-NoC under

521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572

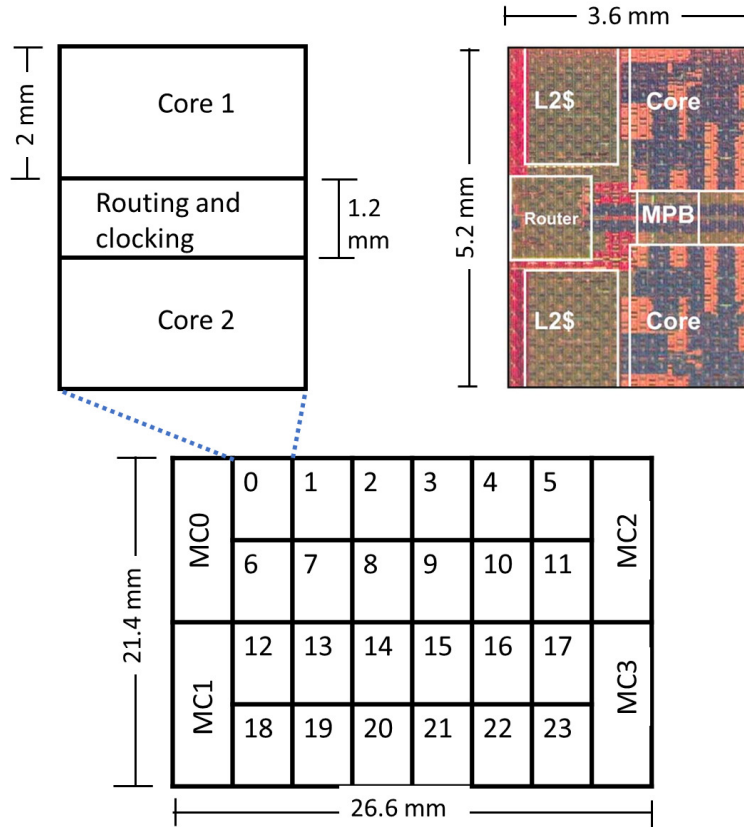


Fig. 9. SCC chip floorplan for the hotspot simulator.

synthetic uniform traffic. Moreover, unlike SWI-NoC, hotspots are formed in Wire-NoC with a temperature above 90 °C (maximum temperature is 94.7 °C). Consequently, the proposed architecture might be able to eliminate the need for dark silicon and thus enable higher computation power under moderate traffic.

On the other hand, if dark silicon is forced on the system, the proposed architecture still exceed the baseline architecture in terms of maximum temperature (Wire-NoC: 87.9°C, SWI-NoC:72.8 °C) and average temperature (Wire-NoC: 79.4°C, SWI-NoC: 67.8 °C), as shown in Fig.11. Although 50 % of the PEs are inactive, the Wire-NoC had manage to reduce the maximum temperature by only ≈ 8°C. In contrast, Fig.11 also shows that thermal impact of SWI-NoC with 50 % dark silicon is still within acceptable limits. Thus the SWI-NoC is more efficient for dark silicon era than Wire-NoC, considering the high impact of dark silicon on Wire-NoC performance, as mentioned earlier, for such low reduction of the maximum temperature.

The gap of maximum temperature of Wire-NoC and SWI-NoC is increasing as the PIR is increased in dark silicon many-core system, see Fig. 12. This is due to the fact that the proposed interconnect architecture are efficiently handling semi-global and global communication that is known in such cases. Moreover, SWI-NoC reduce travel through power-costly routers and wires as will be discussed in the next section. This will reduce dissipated power and therefore the on-chip temperature.

573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624

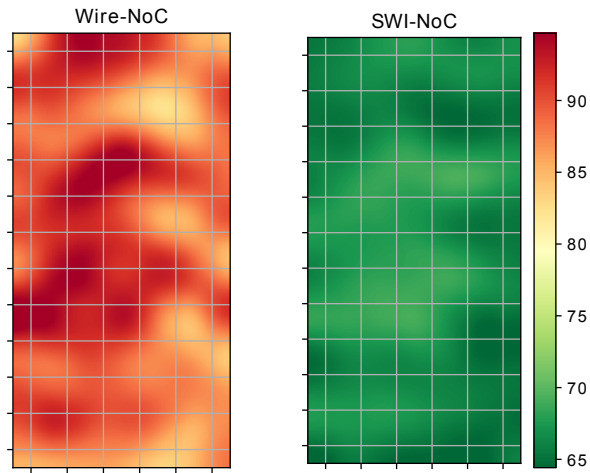


Fig. 10. Thermal spatial impact comparison between Wire-NoC and SWI-NoC with normal operation where all PE are active.




Fig. 11. Thermal spatial impact comparison between Wire-NoC and SWI-NoC with normal operation where 50% PE are operational.

4.3 Power Reduction

As mention earlier, the dark-silicon is a consequence outcome compelled by the increasing power density in SoC. Thus, this section presents an evaluation of the proposed architecture in terms of power consumption.

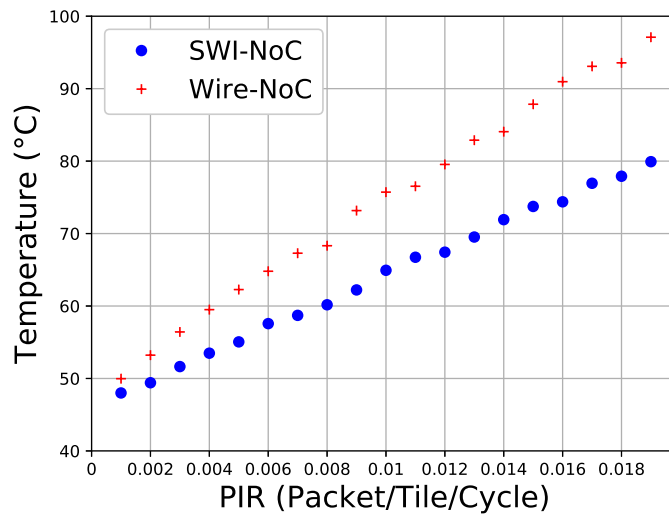


Fig. 12. Comparison of the maximum temperature between Wire-NoC and proposed SWI-NoC architecture as the PIR increased and 50% of the PE are Dark.

4.3.1 *Router and wires power.* The power of the router (static and dynamic) including the extra components for SWI-NoC is calculated using Orion area and power models 2.0 [16] under 45nm technology. With 35% toggling percentage, 2GHz frequency, and 1.1 supply Voltage, router power found to be 549.8 mW. This is match reported power measurements of routers in SCC (baseline architecture) [24]. Moreover, wire links power dissipation is calculated for directions lengths (3.6mm) and (5.2mm) that is according to the baseline architecture.

4.3.2 *transceiver power.* The power consumption of the transceiver (TxRx) is projected by Chang et.al. [7, 8], which is calculated to be 24mW per sub-channel. In addition, the power dissipation of the SWI is calculated based on the analytical model that has been presented previously [19]. All the aforementioned values were used in our adjusted Noxim simulator to calculate the total power consumption of the baseline and the proposed architecture.

4.3.3 *Overall Power Evaluation.* Fig. 13 compares between the total energy of SWI-NoC and that of Wire-NoC when average delay = $2 \times ZLL$ for different NoC sizes under uniform traffics with 5% of multicast and 50% of PEs is turned off. Clearly, SWI-NoC shows significant energy savings compared to wire-NoC. This power saving is exponentially increase as the NoC size increase until the Wire-NoC total communication energy is ($\sim 3x$) the SWI-NoC energy. This is due to the fact that SWI eliminate the need for flits to cross through energy-costly wires and routers especially since the communicating cores are at least 2 hops away. Thus, in Wire-NoC. as the NoC size increase the cost of communication is exponentially increased. In contrast, SWI prioritize global over local communication using DW arbitration technique, see section 3.3, and thus provides one hope cross the chip for such required communication in dark silicon domain. Therefore, the proposed architecture has two advantages in the dark-silicon era. Firstly, it reduces interconnect power budget and scale linearly, which results in mitigating thermal impact of the interconnect itself. Secondly, it is power efficient in terms of satisfying global and semi-global communication requirements of dark silicon operation mode.

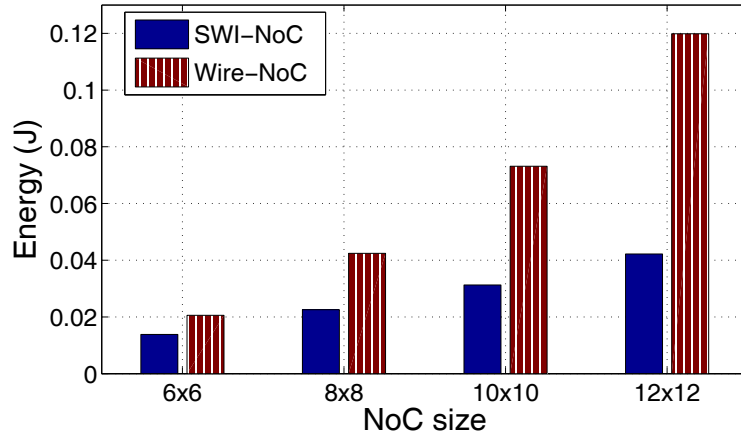


Fig. 13. Interconnect energy comparison between Wire-NoC and SWI-NoC for different NoC sizes where only 50% of PEs are active.

Table 2. Area overhead evaluation for SWI-NoC over Wire-NoC for 45nm technology.

NoC component		Area per item (mm^2)	
Component	No.	Wire-NoC	SWI-NoC
Router	24	1.0853	1.5124
Transmitter	4	-	0.1558
Receiver	24	-	0.0083
Global arbiter	1	-	0.0552
VCTM table	24	-	-
Wire Links	1	13.653	13.653
Total extra area over Wire-NoC			11.13
NoC/SCC-die area (%)		7	8.96

4.4 Area Overhead Evaluation

In this section, the proposed architecture on-chip area overheads is evaluated and compared with baseline architecture. For the required transceivers, the proposed transceiver by Chang *et al.* [8] is adopted. In calculating the transceiver area, the assumption is that the active parts are the only parts that are scaled down when shifting to 45nm technology, where the area is proportional to the square of the scaling factor [28]. On the other hand, since the passive parts are proportional to the channels' operational frequency range, they remain almost the same size. As a result, the calculated area of transmitter and receiver sub-channel are $4870\mu m^2$ and $260\mu m^2$, respectively. Secondly, the baseline router has been modeled using Orion 2.0 [16] and modeling error is 6% less than the reported implemented router area [24], which is acceptable for the purpose of comparison evaluation in this section. In addition, using Orion 2.0, the extra area resulted from adding the extra router port linked to SW channels over the baseline router is calculated to be $0.427mm^2$.

Table 2 compares on-die area of the proposed architecture and the baseline architecture. Clearly, the area overhead of the SWI-NoC is negligible, which is 2% of the total die area. This is especially true in the dark-silicon era in which the SoC is limited by power budget and its thermal impact rather than the number of on-die devices. This fact makes the SWI-NoC optimum for interconnect architecture for many-core systems.

5 CONCLUSION

This study tackled the inter-chip especial communication demands driven by dark silicon many-core systems using hybrid wire and SWI NoC. In order to insure thermal reliability, active cores are forced to be scattered over the chip, which leads that bulk of the communication is power-hungry global and semi-global communication. Given the one-hop cross the chip links of SWI, the proposed architecture has been proven that it is firstly power efficient with low thermal impact compared to the baseline architecture. Secondly, satisfy communication requirements of distant active cores in dark silicon situation. Thus, resolve the dilemma of having to scarifies performance to save power. Future work would include dynamically reconfigurable hybrid architecture that is thermal and traffic contention aware.

REFERENCES

- [1] Arghavan Asad, Ozcan Ozturk, Mahmood Fathy, and Mohammad Reza Jahed-Motlagh. 2017. Optimization-based power and thermal management for dark silicon aware 3D chip multiprocessors using heterogeneous cache hierarchy. *Microprocessors and Microsystems* 51 (2017), 76 – 88. <https://doi.org/10.1016/j.micpro.2017.03.011>
- [2] K. Banerjee and A. Mehrotra. 2002. A power-optimal repeater insertion methodology for global interconnects in nanometer designs. *Electron Devices, IEEE Transactions on* 49, 11 (nov 2002), 2001 – 2007. <https://doi.org/10.1109/TED.2002.804706>
- [3] P. Baniya and K. L. Melde. 2019. Switched-Beam Endfire Planar Array With Integrated 2-D Butler Matrix for 60 GHz Chip-to-Chip Space-Surface Wave Communications. *IEEE Antennas and Wireless Propagation Letters* 18, 2 (2019), 236–240. <https://doi.org/10.1109/LAWP.2018.2887259>
- [4] Tobias Bjerregaard and Shankar Mahadevan. 2006. A Survey of Research and Practices of Network-on-chip. *ACM Comput. Surv.* 38, 1, Article 1 (June 2006). <https://doi.org/10.1145/1132952.1132953>
- [5] Shan Cao, Zoran Salcic, Zhaolin Li, Shaojun Wei, and Yingtao Ding. 2016. Temperature-aware multi-application mapping on network-on-chip based many-core systems. *Microprocessors and Microsystems* 46 (2016), 149 – 160. <https://doi.org/10.1016/j.micpro.2016.03.010>
- [6] A. Carpenter, Jianyun Hu, Jie Xu, M. Huang, Hui Wu, and Peng Liu. 2012. Using Transmission Lines for Global On-Chip Communication. *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on* 2, 2 (June 2012), 183–193. <https://doi.org/10.1109/JETCAS.2012.2193519>
- [7] M.-C.F. Chang, V.P. Roychowdhury, Liyang Zhang, Hyunchol Shin, and Yongxi Qian. 2001. RF/wireless interconnect for inter- and intra-chip communications. *Proc. IEEE* 89, 4 (Apr 2001), 456–466. <https://doi.org/10.1109/5.920578>
- [8] M. C F Chang, J. Cong, A. Kaplan, Chunyue Liu, M. Naik, J. Premkumar, G. Reinman, E. Socher, and Sai-Wang Tam. 2008. Power reduction of CMP communication networks via RF-interconnects. In *Microarchitecture, 2008. MICRO-41. 2008 41st IEEE/ACM International Symposium on*. 376–387. <https://doi.org/10.1109/MICRO.2008.4771806>
- [9] H. Cheng, J. Zhan, Y. Zhao, Y. Xie, J. Sampson, and M. J. Irwin. 2015. Core vs. uncore: The heart of darkness. In *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6. <https://doi.org/10.1145/2744769.2647916>
- [10] Ge-Ming Chiu. 2000. The odd-even turn model for adaptive routing. *Parallel and Distributed Systems, IEEE Transactions on* 11, 7 (jul 2000), 729 –738. <https://doi.org/10.1109/71.877831>
- [11] Nizar Dahir, Ra’ed Al-Dujaily, Terrence Mak, and Alex Yakovlev. 2014. Thermal Optimization in Network-on-Chip-Based 3D Chip Multiprocessors Using Dynamic Programming Networks. *ACM Trans. Embed. Comput. Syst.* 13, 4s, Article 139 (April 2014), 25 pages. <https://doi.org/10.1145/2584668>
- [12] F. Fazzino, M. Palesi, and D. Patti. 2010. Noxim: Network-on-Chip Simulator. <http://noxim.sourceforge.net/>
- [13] Daniel Greenfield, Arnab Banerjee, Jeong-Gun Lee, and Simon Moore. 2007. Implications of Rent’s Rule for NoC Design and Its Fault-Tolerance. In *Proceedings of the First International Symposium on Networks-on-Chip (NOCS ’07)*. IEEE Computer Society, Washington, DC, USA, 283–294. <https://doi.org/10.1109/NOCS.2007.26>
- [14] J. Hendry. 2010. Isolation of the Zenneck surface wave. In *Antennas and Propagation Conference (LAPC), 2010 Loughborough*. 613 –616. <https://doi.org/10.1109/LAPC.2010.5666898>
- [15] N.E. Jerger, Li-Shiuan Peh, and M. Lipasti. 2008. Virtual Circuit Tree Multicasting: A Case for On-Chip Hardware Multicast Support. In *ISCA ’08. 35th International Symposium on Computer Architecture*. 229–240. <https://doi.org/10.1109/ISCA.2008.12>
- [16] A.B. Kahng, Bin Li, Li-Shiuan Peh, and K. Samadi. 2012. ORION 2.0: A Power-Area Simulator for Interconnection Networks. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 20, 1 (Jan 2012), 191–196. <https://doi.org/10.1109/TVLSI.2010.2091686>
- [17] A. Karkar, T. Mak, N. Dahir, R. Al-Dujaily, K. Tong, and A. Yakovlev. 2018. Network-on-Chip Multicast Architectures Using Hybrid Wire and Surface-Wave Interconnects. *IEEE Transactions on Emerging Topics in Computing* 6, 3 (July 2018), 357–369. <https://doi.org/10.1109/TETC.2016.2551043>
- [18] A. Karkar, T. Mak, K. Tong, and A. Yakovlev. 2016. A Survey of Emerging Interconnects for On-Chip Efficient Multicast and Broadcast in Many-Cores. *IEEE Circuits and Systems Magazine* 16, 1 (Firstquarter 2016), 58–72. <https://doi.org/10.1109/MCAS.2015.2510199>
- [19] A.J. Karkar, J.E. Turner, K. Tong, R. Al-Dujaily, T. Mak, A. Yakovlev, and Fei Xia. 2013. Hybrid wire-surface wave interconnects for next-generation networks-on-chip. *Computers Digital Techniques, IET* 7, 6 (November 2013), 294–303. <https://doi.org/10.1049/iet-cdt.2013.0030>
- [20] Ammar Jallawi Mahmood Karkar. 2016. *Interconnects architectures for many-core era using surface-wave communication*. Ph.D. Dissertation. Newcastle University.

- 781 [21] W. Liu, L. Yang, W. Jiang, L. Feng, N. Guan, W. Zhang, and N. Dutt. 2018. Thermal-Aware Task Mapping on Dynamically Reconfigurable Network-
782 on-Chip Based Multiprocessor System-on-Chip. *IEEE Trans. Comput.* 67, 12 (Dec 2018), 1818–1834. <https://doi.org/10.1109/TC.2018.2844365>
- 783 [22] David M Pozar. 2009. *Microwave Engineering*. John Wiley & Sons.
- 784 [23] Pradip Kumar Sahu and Santanu Chattopadhyay. 2013. A survey on application mapping strategies for network-on-chip design. *Journal of systems*
785 *architecture* 59, 1 (2013), 60–76.
- 786 [24] P. Salihundam, S. Jain, T. Jacob, S. Kumar, V. Erraguntla, Y. Hoskote, S. Vangal, G. Ruhl, and N. Borkar. 2011. A 2 Tb/s 6×4 Mesh Network for a
787 Single-Chip Cloud Computer With DVFS in 45 nm CMOS. *Solid-State Circuits, IEEE Journal of* 46, 4 (April 2011), 757–766. [https://doi.org/10.1109/](https://doi.org/10.1109/JSSC.2011.2108121)
788 [JSSC.2011.2108121](https://doi.org/10.1109/JSSC.2011.2108121)
- 789 [25] Semiconductor Industry Association. 2011. ITRS: International Technology Roadmap for Semiconductors . <http://www.itrs.net/reports.html> [online].
- 790 [26] M. Shafique and S. Garg. 2017. Computing in the Dark Silicon Era: Current Trends and Research Challenges. *IEEE Design Test* 34, 2 (April 2017),
791 8–23. <https://doi.org/10.1109/MDAT.2016.2633408>
- 792 [27] M. Shafique, S. Garg, T. Mitra, S. Parameswaran, and J. Henkel. 2014. Dark silicon as a challenge for hardware/software co-design. In *2014 International*
793 *Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*. 1–10. <https://doi.org/10.1145/2656075.2661645>
- 794 [28] J. Srinivasan, S.V. Adve, P. Bose, and J.A. Rivers. 2004. The impact of technology scaling on lifetime reliability. In *Dependable Systems and Networks,*
795 *2004 International Conference on*. 177 – 186. <https://doi.org/10.1109/DSN.2004.1311888>
- 796 [29] M. B. Taylor. 2013. A Landscape of the New Dark Silicon Design Regime. *IEEE Micro* 33, 5 (Sep. 2013), 8–19. <https://doi.org/10.1109/MM.2013.90>
- 797 [30] J.E. Turner, M.S. Jessup, and Kin-Fai Tong. 2012. A Novel Technique Enabling the Realisation of 60 GHz Body Area Networks. In *Wearable and*
798 *Implantable Body Sensor Networks (BSN), 2012 Ninth International Conference on*. 58 –62. <https://doi.org/10.1109/BSN.2012.23>
- 799 [31] J. Wan, K. F. Tong, and C. H. Chan. 2019. Simulation and Experimental Verification for a 52 GHz Wideband Trapped Surface Wave Propagation
800 System. *IEEE Transactions on Antennas and Propagation* 67, 4 (2019), 2158–2166.
- 801 [32] Linda Wilson. 2013. International technology roadmap for semiconductors (ITRS). *Semiconductor Industry Association* 1 (2013).
- 802 [33] J. Zhan, J. Ouyang, F. Ge, J. Zhao, and Y. Xie. 2015. DimNoC: A dim silicon approach towards power-efficient on-chip network. In *2015 52nd*
803 *ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6. <https://doi.org/10.1145/2744769.2744824>
- 804 [34] Runjie Zhang, Mircea R Stan, and Kevin Skadron. 2015. Hotspot 6.0: Validation, acceleration and extension. *University of Virginia, Tech. Rep* (2015).
- 805
- 806
- 807
- 808
- 809
- 810
- 811
- 812
- 813
- 814
- 815
- 816
- 817
- 818
- 819
- 820
- 821
- 822
- 823
- 824
- 825
- 826
- 827
- 828
- 829
- 830
- 831
- 832