

Low-power Networks-on-Chip: Progress and Remaining Challenges

Mark Buckler^{†‡}, Wayne Burleson^{†‡}, Greg Sadowski[†]

AMD Research, Boxborough, Mass.[†]

University of Massachusetts, Amherst, Mass.[‡]

mark.buckler@amd.com, wayne.burleson@amd.com, greg.sadowski@amd.com

Abstract

After a long period of academic and industrial research, networks-on-chips (NoCs) are starting to be incorporated into commercial multi-processor designs. NoCs have proven themselves to scale better than bus-based designs and they are here to stay. It is still important to note, however, that even well-designed NoCs consume a large portion of a given system's power budget. This brief paper and accompanying presentation discuss what options are available to designers who need to reduce NoC power consumption, their benefits, and their limitations. Techniques discussed here include general NoC system design as well as disruptive interconnect mediums and their associated strategies.

Keywords

Networks-on-Chip, Low Power, Cache Coherence, DVFS, Asynchronous, Low Swing, 3D, Wireless, Nanophotonics

1. Introduction

Before the concept of a network-on-chip (NoC) was proposed, system-on-chips (SoCs) relied on complex bus structures to connect processors to memory and I/O. Moore's Law has continued since that time; however, clock rates have stagnated due to power issues. The need for more processing power (without clock increases) and the ability to add more transistors on a chip has led to designers increasing both the number and diversity of processors on chips. The old bus structures were improved to account for these multi-processor system-on-chips (MPSoCs), but eventually the bus designs could not sustain the large degree of interconnect scaling and complexity. Eventually, the NoC emerged as a solution to this problem by "routing packets instead of wires" and has increased in popularity since then [1]. This trend has led to companies like Arteris, Sonics, Blendics, and iNoCs providing this style of interconnect solution. In fact, many companies are beginning to choose pre-designed NoC IP solutions over designing their own NoCs in house.

Due to the long-standing need to reduce SoC power consumption, research in low-power NoCs has existed for at least a decade. Unfortunately, NoCs are not inherently low power. Some examples cite power numbers as high as 35% of total chip power [2]. The restrictions on SoC power usage have only become stronger, influencing the engineers who design NoCs to reduce power whenever possible. Low-power research areas include traffic management, signaling strategies, and interconnect paradigms. Traffic management involves research into topics like cache coherence and compression. Signaling strategies include asynchronous communication, dynamic voltage, and low swing. Interconnect paradigms include 3D, nanophotonic, and wireless interconnects.

2. Traffic Management

In many SoCs, the bulk of NoC traffic is to maintain cache coherence. For this reason, design and management of the cache is critical and must be considered when distributing the cache among CPU and GPU cores [3]. Methods have been developed to reduce the power used in cache-hierarchy management using both data locality and knowledge of the NoC's physical structure [4][5]. Coherence-free systems have been proposed to avoid coherence protocols, but the industry largely favors cache-coherent systems [6]. One successfully demonstrated method to decrease cache coherency power usage combined bus-based snooping coherency and NoC-based directory coherency [7].

Power reduction has also been achieved through efficient use of data compression [8], error detection/correction encoding [9], and heterogeneous interconnect [10]. Other techniques achieved power reduction by differentiating among different kinds of traffic (such as 1-to-many/many-to-1 [11] or request/response [12]) and optimizing for each type. Hardware techniques focus on router designs and microarchitecture [13]. Although bufferless NoC designs have been proposed, their benefits are minimal (1.5% savings) [14].

3. Signaling Strategies

3.1. Asynchronous Communication

Distributing a global clock across an entire NoC continues to be difficult and very power-hungry as technology scaling continues while die area remains the same. For this reason, the globally asynchronous/locally synchronous (GALS) NoC was proposed. Studies have verified that GALS NoCs save both energy and latency by removing the global clock but require overhead in the form of synchronizer circuits and extra router wires for flow control [15]. These extra router wires manifest as a requirement for more space for the NoC, sometimes as high as 25% increased switch area (while still maintaining 21% power reduction, given certain factors) [16]. Those numbers were improved to an impressive 57% power reduction when using the butterfly fat tree (BFT) network topology [17].

Area overhead can be reduced with specialized circuitry for routers and other asynchronous components [18]. Even without using the full GALS approach, gains can be achieved with asynchronous circuitry. One recent paper uses router crossbars with built-in asynchronous repeated link circuits. This technique has achieved single-clock-cycle latency along with a 2.2X power savings [19]. Source-synchronous communication using bundled data also has been proposed. This technique routes the clock (as a pulse) along with the data. Source-synchronous systems reduce power through their removal of the global clock [20].

3.2. Dynamic Voltage and Frequency Scaling

Similar to other parts of the SoC, the NoC does not always need to operate at its maximum possible level of performance. For this reason, dynamic voltage and frequency scaling (DVFS) can optimize dynamic power. Clock- and power-gating can be considered extreme cases of DVFS and make sense to minimize dynamic and static power, respectively. NoCs must also take into account DVFS changes in the chip nodes that modify incoming and outgoing data rates. Recent work has shown that savings as high as 33% can be seen when applying DVFS to the NoC and low-level cache (LLC) when sharing a voltage/frequency domain [21]. Another proposed design includes dynamic reconfigurable NoC interconnect in addition to DVFS, allowing for energy savings and latency reduction [22]. A simplified binary DVFS control using only a high and a low voltage state also has been proposed to be sufficient for NoC switches [23].

3.3. Low Swing

A low-swing signaling attempts to save energy by reducing the voltage potential between high and low states (lowering the swing) on large chip wires. New low-swing techniques have proven to reduce clock power by 66% [24]. With reduced swing comes increased sensitivity to noise, however, requiring special care to ensure reliability [25].

Due to the analog nature of this technique, work until now has focused on differential signaling and both voltage- and current-mode transceiver circuit designs [24][26]. Unfortunately, highly custom circuits pose a problem for modern SoCs, which often are designed using synthesized circuits. For this reason, focus also has been given to creating low-swing solutions that can be easily implemented using mainstream SoC design techniques [27].

4. Interconnect Paradigms

4.1. 3D Interconnect

The long-awaited emergence of 3D VLSI and die-stacking technology has motivated additional work in the corresponding NoCs. 3D promises shorter interconnect and reduced capacitance, as well as excellent inter-layer connections with the use of through-silicon vias (TSVs) [28]. TSVs also make the circuit design of 3D routers and 3D routing schemes significantly different [29]. Unfortunately, the state of technology today prevents more than two logic layers to be stacked in one package due to thermal concerns. Designs have been proposed with more than two layers, suggesting that one layer could be dedicated to the NoC [30]. These thermal concerns have caused researchers to explore the possibility of thermal-aware 3D NoC architectures that can help mitigate thermal issues [31].

4.2. Nanophotonic Interconnect

Although a nanophotonics-based NoC has not yet been developed due to technology limitations, silicon photonics have now been demonstrated in a 90-nm process [32]. This kind of progress has increased interest in nanophotonics as a way to replace traditional metal wires for long-haul connections in NoCs. Full analysis of planned nanophotonic networks has shown significant promise for both increased performance and decreased power consumption using athermal ring resonators and on-chip lasers that enable quick

power-gating [33]. Nanophotonics promises bit rates almost independent of distance, higher bandwidth from frequency-division multiplexing (FDM), and lower power due to dissipation at the endpoint only. These promised benefits allow for the potential to improve performance by 60% and decrease power by 80% [34]. NoC laser energy also can be reduced by 49% using busses controlled by distributed on-chip lasers [35]. While these pure photonic designs are very attractive, the first practical photonic NoC likely will be some combination of photonics and traditional metal wires [36]. Although recent nanophotonics research is very promising, there is more work to be done on the process side before nanophotonic NoCs can be fully realized.

4.3. Wireless Interconnect

Both photonics and wireless NoC designs are part of a trend to integrate formerly off-chip communication techniques into the on-chip network to increase performance and reduce power. Miniature on-chip antennas could be used to transmit and receive information, and the technology already exists to create them on silicon. A wireless NoC would save power and area because small transmitters do not need large capacitive transmission lines and do not require multi-hop connections. Hybrid designs have been proposed with wireless used for long-distance on-chip transmissions [37][38].

Wireless NoCs can use FDM (similar to the concept's use in nanophotonics) and time-division multiplexing (TDM) along with low-power transceivers to achieve 34% power reduction compared to leading NoCs [39]. Another wireless NoC design uses a sub-divided mesh topology to improve the performance of other wireless NoC designs [40]. Wireless systems face unique challenges, however. For now, designers are limited to using existing millimeter-wave antennas using CMOS technology, but future carbon nanotube (CNT) antennas will significantly reduce the overhead [41]. Use of these CNT antennas is not possible yet due to the need for process scaling that has not yet been achieved.

5. Conclusions

The successful application of both high-level design strategies and interconnect paradigms can be very effective in limiting NoC power usage. Well-designed high-level systems manage to combine traffic management and signaling strategies into an efficient whole. Challenges associated with high-level design largely consist of improving these areas of design and their methods of integration.

While traffic management and signaling strategies are also important for NoCs employing low-power interconnect paradigms, they are not the biggest challenge. Process limitations are the greatest factor when considering a new interconnect paradigm. 3D interconnects still require improvements in TSV yield, photonics have only recently been miniaturized to the nanometer level, and wireless NoCs still rely on less efficient CMOS millimeter-wave antennas. For these reasons, the interconnect paradigms described in this paper are not yet ready for mainstream design. Some forms of interconnect such as 3D will be available in the near term, however, showing that there is a range of near- and longer-term solutions to on-chip communication. Low-power NoC designers should be aware of their current limitations while still looking forward to future opportunities.

References

1. Dally, W.J.; et al., "Route packets, not wires: on-chip interconnection networks," DAC, 2001.
2. Kim, J.S.; et al., "Energy characterization of a tiled architecture processor with on-chip networks," ISLPED, 2003.
3. Xu, T.C.; et al., "Explorations of optimal core and cache placements for Chip Multiprocessor," NORCHIP, 2011.
4. Hyungjun, Kim; et al., "Reducing network-on-chip energy consumption through spatial locality speculation," NoCS, 2011.
5. Fensch, C.; et al., "Designing a Physical Locality Aware Coherence Protocol for Chip-Multiprocessors," IEEE Tr. Computers, 2013.
6. Milo, Martin; et al., "Why on-chip cache coherence is here to stay," Communications of the ACM, 2012.
7. Hui, Zhao; et al., "A hybrid NoC design for cache coherence optimization for chip multiprocessors," DAC, 2012.
8. Yuho, Jin; et al., "Adaptive data compression for high-performance low-power on-chip networks," MICRO, 2008.
9. Po-Tsang, Huang; et al., "Low Power and Reliable Interconnection with Self-Corrected Green Coding Scheme for Network-on-Chip," NoCS, 2008.
10. Flores, A.; et al., "Heterogeneous Interconnects for Energy-Efficient Message Management in CMPs," IEEE Tr. Computers, 2010.
11. Tushar, Krishna; et al., "Towards the ideal on-chip fabric for 1-to-many and many-to-1 communications," MICRO, 2011.
12. Volos, S.; et al., "CCNoC: Specializing On-Chip Interconnects for Energy Efficiency in Cache-Coherent Servers," NoCS, 2012.
13. Kim, J., "Low-cost router microarchitecture for on-chip networks," MICRO, 2009.
14. Michelogiannakis, G.; Sanchez, D.; Dally, W.J.; Kozyrakis, C., "Evaluating Bufferless Flow Control for On-chip Networks," NOCS, 2010.
15. Gebhardt, D.; et al., "Comparing Energy and Latency of Asynchronous and Synchronous NoCs for Embedded SoCs," NOCS, 2010.
16. El Ghany, M.A.A.; et al., "Power analysis for Asynchronous CLICHE Network-on-Chip," SOCC, 2010.
17. Rashed, M.; et al., "Power characteristics of Asynchronous Networks-on-Chip," SOCC, 2011.
18. Horak, M.N.; et al., "A Low-Overhead Asynchronous Interconnection Network for GALS Chip Multiprocessors," IEEE Tr. Computer-Aided Design of Integrated Circuits and Systems, 2011.
19. Chen, Chia-Hsin Owen; et al., "SMART: A single-cycle reconfigurable NoC for SoC applications," DATE, 2013.
20. Yoon Seok, Yang; et al., "WaveSync: A low-latency source synchronous bypass network-on-chip architecture," ICCD, 2012.
21. Xi, Chen; et al., "In-network Monitoring and Control Policy for DVFS of CMP Networks-on-Chip and Last Level Caches," NoCS, 2012.
22. Liang, Guang; et al., "Run-time communication bypassing for energy-efficient, low-latency per-core DVFS on Network-on-Chip," SOCC, 2010.
23. Kumar, Yadav, et al., "A Simple DVFS Controller for a NoC Switch," PRIME, 2012.
24. Yi, Liu; et al., "A novel low-swing transceiver for interconnection between NoC routers," IDCTA, 2011.
25. Ejlali, A.; et al., "Performability/Energy Tradeoff in Error-Control Schemes for On-Chip Networks," IEEE Tr. VLSI Systems, 2010.
26. Mensink, E.; et al., "Power Efficient Gigabit Communication Over Capacitively Driven RC-Limited On-Chip Interconnects," IEEE Jou. Solid-State Circuits, 2010.
27. Postman, J.; et al., "SWIFT: A Low-Power Network-On-Chip Implementing the Token Flow Control Router Architecture With Swing-Reduced Interconnects," IEEE Tr. VLSI Systems, 2013.
28. Xu, T.C.; et al., "Optimal number and placement of Through Silicon Vias in 3D Network-on-Chip," DDECS, 2011.
29. Ahmed, A.B.; et al., "Low-overhead Routing Algorithm for 3D Network-on-Chip," ICNC, 2012.
30. Nandakumar, V.S.; et al., "Low power, high throughput network-on-chip fabric for 3D multicore processors," ICCD, 2011.
31. Rahmani, A.-M; et al., "Design and management of high-performance, reliable and thermal-aware 3D networks-on-chip," Circuits, Devices & Systems, IET, 2012.
32. Assefa, S.; et al., "A 90nm CMOS integrated Nanophotonics technology for 25Gbps WDM optical communications applications," IEDM, 2012.
33. Kurian, G.; et al., "Cross-layer Energy and Performance Evaluation of a Nanophotonic Manycore Processor System Using Real Application Workloads," IPDPS, 2012.
34. Morris, R.; et al., "Extending the Performance and Energy-Efficiency of Shared Memory Multicores with Nanophotonic Technology," IEEE Tr. Parallel and Distributed Systems, 2013.
35. Chao, Chen; et al., "Runtime Management of Laser Power in Silicon-Photonic Multibus NoC Architecture," IEEE Jou. Selected Topics in Quantum Electronics, 2013.
36. Adi, C.A.D.; et al., "An Efficient Path Setup for a Photonic Network-on-Chip," ICNC, 2010.
37. Abd El Ghany, M.A.; et al., "Hybrid Mesh-Ring wireless Network on Chip for multi-core system," ISOCC, 2012.
38. Ganguly, A.; et al., "Scalable Hybrid Wireless Network-on-Chip Architectures for Multicore Systems," IEEE Tran. Computers, 2011.
39. DiTomaso, D.; et al., "Energy efficient modulation for a wireless network-on-chip architecture," NEWCAS, 2012.
40. Ling, Wang; et al., "A hybrid chip interconnection architecture with a global wireless network overlaid on top of a wired network-on-chip," SoC, 2012.
41. Carloni, L.P.; et al., "Networks-on-chip in emerging interconnect paradigms: Advantages and challenges," NoCS, 2009.