Networking Challenges and Prospective Impact of Broadcast-Oriented Wireless Networks-on-Chip

Sergi Abadal
NaNoNetworking Center in Catalonia
UPC - BarcelonaTech
Barcelona, Spain
abadal@ac.upc.edu

Mario Nemirovsky
ICREA Professor
Barcelona Supercomputing Center (BSC)
Barcelona, Spain

Eduard Alarcón
NaNoNetworking Center in Catalonia
UPC - BarcelonaTech
Barcelona, Spain

Albert Cabellos-Aparicio
NaNoNetworking Center in Catalonia
UPC - BarcelonaTech
Barcelona, Spain

ABSTRACT

The cost of broadcast has been constraining the design of manycore processors and of the algorithms that run upon them. However, as on-chip RF technologies allow the design of small-footprint and high-bandwidth antennas and transceivers, native low-latency (a few clock cycles) and low-power (a few pJ/bit) broadcast support through wireless communication can be envisaged. In this paper, we analyze the main networking design aspects and challenges of Broadcast-oriented Wireless Network-on-Chip (BoWNoC), which are basically reduced to the development of Medium Access Control (MAC) protocols able to handle hundreds of cores. We evaluate the broadcast performance and scalability of different MAC designs, to then discuss the impact that the proposed paradigm could exert on the performance, scalability and programmability of future manycore architectures, programming models and parallel algorithms.

Categories and Subject Descriptors

B.4.3 [**Hardware**]: Interconnections (subsystems)—*Topology (e.g. bus, point-to-point)*; C.0 [**Computer Systems Organization**]: General—*System Architectures*

Keywords

Network-on-Chip, Wireless On-Chip Communication, Multicast, Broadcast, MAC Protocols, Manycore Processors

1. INTRODUCTION

In the ever-changing world of microprocessor design, multicore architectures are currently the dominant trend for both conventional and high-performance computing. These architectures consist of the interconnection of several independent processors or cores, as well as of a multilevel cache to improve overall performance. Communication between cores and the memory hierarchy is not only a requirement to ensure the correct operation of a multiprocessor, but also a main determinant of its performance [14].

Networks-on-Chip (NoCs) are currently the paradigm of choice to meet the communication demands of multiprocessors beyond a handful of cores [7]. The packet-switched and point-to-point nature of NoCs provides better scalability in terms of throughput than buses and improved energy-delay profile as long as most of the traffic is local. In light of this, multiprocessing architectures and algorithms generally take base on the well-known *locality of reference* properties and attempt to reduce the amount of global and broadcast communication, which incur into non-local traffic and are thus detrimental to most switched NoCs. Performance reductions are more severe in scaled processors, a scenario wherein communication is of critical importance [38].

The limited scalability of conventional NoCs in the presence of certain types of traffic is a factor that sets important restrictions upon the design of architectures and algorithms for manycore processors. For instance, the prohibitive cost of broadcast communication has important implications upon the performance, complexity, and programmability of manycore systems. In shared memory, having ordered broadcast capabilities could simplify the design of coherence protocols and increase their performance by reducing the number of race conditions [41] and the amount of indirection in certain situations [9], respectively. In message passing, a number of algorithms would highly benefit from better one-to-all and all-to-all routines [43]. Further, efficient hardware synchronization through broadcast would ease the complexity of parallel programming [33].

Recently, a plethora of works have tried to address these issues by proposing ways to improve the performance and scalability of NoCs. Topologies with higher connectivity [32, 38], sophisticated router pipeline designs [20, 35], or the use of high-radix switches [5] have consistently delivered better performance for most types of traffic. However, the inherently point-to-point nature of these networks prevents

NOCS '15 September 28 - 30, 2015, Vancouver, BC, Canada
 ISBN 978-1-4503-3396-2/15/09 $15.00
DOI: http://dx.doi.org/10.1145/2786572.2788710.

the design of truly scalable broadcast schemes without significantly affecting other traffic flows.

The advent of novel interconnect technologies has opened a set of new design opportunities to complement conventional NoCs [6,10,19]. Vertical stacking enables new topologies not feasible in planar environments and brings nodes closer both physically and logically, thereby improving performance. The use of nanophotonic components has been also proposed as it promises great energy efficiency, as well as outstanding performance via speed-of-light propagation and virtually unlimited bandwidth. Similarly, RF communication via transmission lines offers improved energy efficiency and *long-range* communication support. However and as indicated in Section 2.1, all these technologies have their caveats, which may reduce their applicability.

Last but not least, the concept of Wireless Network-on-Chip (WNoC) has recently garnered considerable attention due to its unique latency, reconfigurability and broadcast capabilities [1, 10]. In a WNoC, a transmitting antenna radiates RF signals at a given frequency, which propagate throughout the chip and may be received by any other antenna tuned to the same frequency. In other words, all antennas tuned at the same frequency channel share the medium similarly to in a bus. Most works make the case for the use of multiple channels, approach that reduces the probability of two or more antennas trying to access to the same channel simultaneously. These advanced designs have reported outstanding performance and power improvements when augmenting a conventional NoC by reducing the cost critical transmissions, alleviating certain network bottlenecks, or adapting to the traffic of different applications [11,13,25,31].

In this paper, we present the concept of Broadcast-oriented Wireless Network-on-Chip (BoWNoC, see Fig. 1), a WNoC design that basically differs from most proposals in that all cores share a single broadband channel. This approach reduces the attainable throughput of the WNoC, but instead provides a simple and effective vehicle for the ordered delivery of multicast and broadcast traffic in a single hop. As such, the BoWNoC will only transport broadcast and latency-critical messages, thereby complementing a conventional NoC that will deal with the rest of communication flows. Although this strategy is not new [27, 34], this is the first work that discusses the use of a wireless plane for broadcast purposes as a complement of a unicast and throughput-oriented network. The motivation of this research resides in the limitations cast upon current manycore systems, which we expect to largely eliminate with the introduction of BoWNoC. Besides presenting BoWNoC, the contributions of this work are:

- An outline of its enabling technologies, design aspects and networking challenges.
- An evaluation of its attainable broadcast performance, which is later compared to that of conventional NoCs with state-of-the-art broadcast support.
- A qualitative analysis of the impact of this paradigm upon the scalability, performance and programmability of future manycore systems.

The remainder of the paper is as follows. In Section 2, we present the concept of BoWNoC and outline their main characteristics and differences with respect to other designs. Section 3 summarizes the methodology used for the modeling and simulation of different BoWNoC schemes. Section 4 presents the results of a performance evaluation that compares the performance of BoWNoC with that of other advanced NoCs. Then, in Section 5, we discuss the prospective impact that the improved capabilities of BoWNoC could have on the design of future architectures and algorithms. Section 6 concludes the paper.

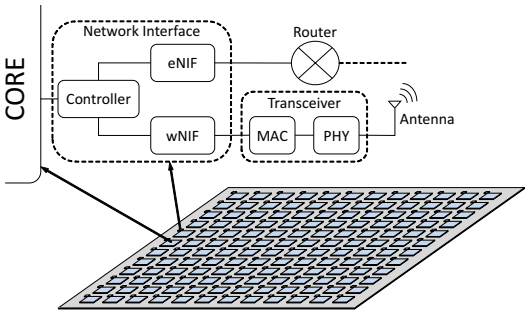

Figure 1: Schematic diagram of a Broadcast-oriented Wireless Network-on-Chip.

2. BROADCAST-ORIENTED WIRELESS NETWORK-ON-CHIP

2.1 Motivation and Related Work

Ever since NoCs started replacing buses for scalability reasons, several works have acknowledged the lack of proper multicast/broadcast support. Given the unicast nature of NoCs, multicast packets need to be replicated at some point. If the destination set is large, this process causes a severe performance drop in the whole network [16]. This situation is expected to worsen as the number of cores grows, and may lead to a remarkable slowdown in the execution speed of multiprocessors [21].

To remedy this, improved router designs have sought to minimize latency by reducing the pipeline depth [35] or even allowing multiple hops within the same clock cycle [22], and to maximize throughput by balancing the load [21]. Prototypes have shown impressive performance gains and even demonstrated that ordered delivery of broadcasts is possible in switched NoCs, which typically do not guarantee ordering [9]. However, these works have recognized important latency and throughput scalability concerns beyond 100 cores: messages still need to traverse a potentially large number of routers to reach the furthest nodes, and contention will still be significant for large destination sets.

Another set of proposals broke away from the switched paradigm and, instead, have proposed to use globally shared media to carry broadcast communication flows. Manevich *et al* proposed to overlay buses on top of a mesh NoC [27]. Oh *et al* presented a shared transmission line ring that supports the broadcast of RF signals [34]. Such design uses very similar principles than IBM's nanophotonic platform Corona [42], which has a shared waveguide exclusively for broadcast signals. Kurian *et al* also leveraged nanophotonics to create a set of independent logical rings, each of which broadcasts signals from a given transmitter [23]. However, all cases entail scalability concerns that may limit their usefulness: the design complexity of buses and transmission

(a) Upstream (b) Downstream

Figure 2: 4-way concentration in BoWNoC.

lines does not scale well and worsens when broadcast capabilities are required, whereas the logarithmic nature of losses in the nanophotonic case complicates the design of scalable networks within strict laser power budgets [2].

2.2 Design Aspects and Benefit

BoWNoC adopts wireless communication technologies to create a globally shared medium and uses it as a broadcast plane. This is accomplished by resorting to the inherently broadcast nature of WNoC and tuning all the antennas to a unique frequency channel. The use of a single channel forces BoWNoC to manage access to the shared medium through arbitration, but also guarantees one-hop delivery with all nodes having a consistent view of the order of delivery. In other WNoC designs, nodes are divided in different subsets and assigned one channel per subset. Thus, broadcast messages still need to resort to the switched NoC, thus requiring multiple hops and losing ordering guarantees.

Figure 1 depicts the scheme that bridges the processors with the network. Ideally, one small-footprint antenna and transceiver will be integrated within each computing tile to provide broadcast capabilities at the core level. If it is necessary to reduce the number of on-chip antennas as per area or power reasons, k-way concentration can be leveraged with an asymmetric design. As shown in Fig. 2, a switch can arbiter access of k controllers to the antenna, whereas received signals can be amplified and delivered to all k controllers. For small k values, the latency overhead would be one clock cycle per direction at most.

The shared channel should support very high data rates in order to ensure that the transmission of a single message does not take more than a few clock cycles. Even so, it is expected that the attainable throughput will be low as compared to in conventional NoCs. Consequently, the use of BoWNoC will be limited to broadcast transmissions and, in rare occasions, latency-critical messages (this basically concerns to the design of the controller). Overlaying such a network over a throughput-oriented NoC would not only provide scalable, uniform, and low-latency support for broadcast, but also dramatically reduce contention effects in the conventional NoC plane by preventing it from having to deal with such type of traffic. To quantify the potential improvements achievable with this scheme, we evaluated the performance of **E-MESH** depicted in Sec. 3.1 with and without an idealized instance of BoWNoC. Figure 3 shows that the low-load latency is cut by a factor proportional to both the system size and the percentage of broadcast transmissions, whereas the saturation throughput increases significantly for a wide range of broadcast percentages.

2.3 Main Enabling Technologies

To maximize the benefits of the BoWNoC approach, it is required that antennas and transceivers (i) be commensurate in size with future computing tiles, (ii) offer data rates in the order of one flit per clock cycle, and (iii) consume low power. Fortunately, this three requisites could be met by increasing the frequency at which the wireless network operates, as this increases its capacity and reduces both the size and energy per bit of the antennas and transceivers [2]. Assuming a clock frequency of 1 GHz, an operation frequency of a few hundreds of GHz seems an appropriate target. Antennas would have a lateral size of a few tens of micrometers and potential to deliver data rates of several tens of Gbps.

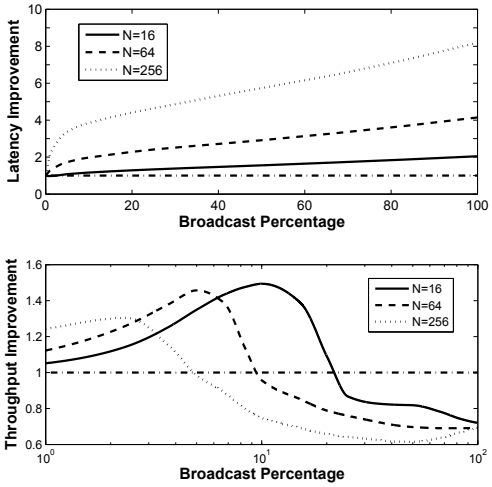

Figure 3: Improvement in terms of (a) latency and (b) throughput delivered by BoWNoC as a function of the broadcast percentage and system size.

Reaching such frequency bands becomes possible as CMOS technology evolves and advanced devices such as FinFETs and III-V on silicon are implemented [1, 24, 39, 45]. A growing number of publications also addresses mmWave and THz on-chip antennae that are suitable for silicon-based transceivers [15, 29]. Additionally, novel technologies such as graphene could introduce further improvements at the chip scale: graphene-based planar antennas are expected to radiate in the THz band while being two orders of magnitude smaller than their metallic counterparts [1], whereas graphene-based transceivers can be envisaged as the predicted operating frequency of graphene devices lies within the same band [44]. Another technology that may confirm the feasbility of BoWNoC is surface wave technology [17]. Instead of propagating in all directions, radiated signals are bound to the surface and propagate along it as in a transmission line. This shrinks the spreading loss and, if feasible, opens the door to higher throughput schemes without the design complexity of transmission lines.

2.4 MAC: The Grand Challenge

The design of a WNoC implies addressing important challenges at the different levels of the protocol stack [1,10]. For instance, finding lightweight coding and modulation schemes able to maintain a graceful balance between speed, power, and transceiver complexity is one of such challenges. But, given that the medium is shared by a potentially large number of nodes, resource management and access control are arguably the most critical design challenges. On top of that, there is a need to revise addressing and routing mechanisms

in those WNoCs where packets may go through both the wired and wireless planes to reach its destination.

The BoWNoC paradigm has substantial differences with respect to other WNoC designs, subtly affecting the design challenges. Since the wired and wireless planes are isolated in BoWNoC (i.e., a message will only go through one of the planes to reach its destinations), a revision of the addressing and routing mechanisms may not be needed. Instead, even more emphasis must be put on the criticality of the MAC protocol design as a given channel will be shared by a noticeably larger number of nodes.

To take full advantage of the broadcast capabilities of BoWNoC, the MAC protocol must be:

1. **Scalable:** the zero-load broadcast latency of BoWNoC is scalable as it remains constant when varying the number of cores. Therefore, the MAC protocol must also be scalable (which is a challenge in itself) in order to retain such latency advantage. Due to this, distributed mechanisms may be preferred.
2. **Performance and cost effective:** the overhead introduced by the protocol is another important aspect. Lengthy or power-hungry arbitration mechanisms cannot be applied in this scenario due to its stringent performance and power requirements. The main challenge here is to design a protocol that, while keeping the overhead low, performs well even in the expected presence of bursty traffic [3].
3. **Fair:** fairness generally implies providing service with a rather uniform latency, with independence on the source of the message. This is a difficult task when accounting with few resources and using a distributed protocol. The expected presence of hotspot traffic [3] may further complicate this goal.

In order to achieve this goals, the chip scenario offers unique opportunities not available in conventional wireless networks. The possibility of implementing a dedicated lightweight wired network to help with arbitration, or of exploiting spatiotemporal traffic correlations to predict the source of broadcast transmissions, represent two clear examples.

The WNoC scenario admits both collision-free schemes (any combination of multiplexing and centralized or distributed arbitration) and contention-based schemes (where cores contend to access the medium) as long as the area and power overheads are low. First designs heavily relied on frequency-multiplexing [25], which were later enhanced with basic time-multiplexing or token-passing arbitration schemes to scale to a larger number of cores and still avoid collisions [13,31]. Another collision-free mechanism has been recently proposed in [12], consisting of distributed protocol where nodes request access using tightly-synchronized code-multiplexed packets. Contention-based mechanisms, which can virtually eliminate the overheads associated to accessing the channel, have received less attention due to their poor performance when high loads cause recurrent collisions. To address this issue, authors of [28] propose an adaptive scheme that switches between a contention-based protocol and a token-passing scheme depending on the level of contention.

3. SIMULATION METHODOLOGY

In this work, we evaluate the performance of a set of representative BoWNoC designs and benchmark them against a state-of-the-art switched NoC. The simulator employed to this end is PhoenixSim [8], a cycle-accurate NoC simulator. It includes a complete set of methods and primitives for the evaluation of switched NoCs, on top of which we implemented the necessary modules for the simulation of WNoCs.

Table 1: Simulation Parameters

System	
Die Size	400 mm^2
Number of Cores	16 - 1024
Operation Frequency	2 GHz
Flit Size	128 bits
Wireline Network-on-Chip	
Topology	Mesh
Router Pipeline	1-cycle (bypass)
Flow Control	Credit-based
Wireless Network-on-Chip	
Number of Channels	1
Channel Capacity	256 Gbps[1]
Propagation Delay	<0.25 cycles

To model the system, we used the scheme depicted in Figure 1 with the parameters shown in Table 1. Cores simply inject traffic following the directives of a centralized generator, which models flows with a given burstiness and distributes them according to a given spatial distribution [40]. All packets are broadcast and, for simplicity, their length equals to the flit size. We stress the network with a variable number of injected packets between 10K and 100K, with appropriately sized warm-up and cool-down periods.

The wireless network is modeled as a shared medium in which any node can transmit following the policies set by the MAC protocol. During the propagation time, receiving nodes may still not be aware that the medium is being used. In contention-based schemes, this can lead to collisions even if the transmitters check the medium before occupying the channel. Nodes that are not transmitting will automatically detect collisions whenever two incoming transmissions overlap in time. Both messages are internally discarded and the collision is generally notified to the MAC module.

3.1 Investigated Network Architectures

We aim to cover a large fraction of the solution space by considering the following schemes. Unless noted, acknowledging is implicit since negligible bit error rates can be assumed with proper coding and power allocation. Wireless capabilities are given at the core level.

Multihop Token Passing [W-MTKN] - where only the core that possesses the token is able to transmit. We consider a asynchronous wired ring for token passing purposes, as represented in Fig. 4. Each MAC module is connected to the ring through a register controlled by a *seize* bit, which is enabled when the node wants to transmit or it is "multiple of M" hops away from the previous transmitter (whose ID is known by all). This way, the token will traverse M MAC modules within the same clock cycle unless someone along the path needs to trasmit. M is limited by technology and

[1]The wireless capacity value is chosen to be roughly the rate at which cores can inject data (i.e. flit size × frequency). Capacity requirements will be greatly relaxed if lower clock speeds are used, which is highly probable [18], or shorter messages are transmitted.

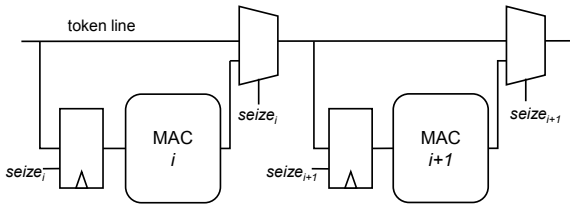

Figure 4: Multihop ring for token passing.

energy consumption [20]: we choose $M = 4$ as a conservative value. Token passing overlaps with part of the transmission.

Carrier Sensing [W-CSMA] **-** with which we aim to represent contention-based protocols. We model a simple non-persistent MAC protocol based on collision avoidance: when a node is ready to send data, it listens to the medium and only transmits if the medium is idle. Otherwise, it waits during a backoff period and checks again whether the medium is idle. We adopt a negative acknowledgment (NACK) strategy to reduce the control overhead: NACKs are sent through the same channel than data, seeking to create a burst of NACKs that will be interpreted by the source as an erroneous transmission. A packet will be forwarded to an alternative network plane in the unlikely case that it exceeds the maximum number of retries (8).

Centralized Arbitration [W-CARB] **-** where contention is progressively resolved using a lightweight wired network consisting of a hierarchical set of 4:1 switches. Transmitters send their source ID to the buffer through such arbitration network, the last layer of which is connected to a FIFO buffer. This buffer grants access to the node whose request is in the head of the buffer using, to this end, a dedicated 1-bit wire. We assume that each switch stage and the access granting take one clock cycle, and that the latter partly overlaps with the wireless transmission. The main reason for evaluating such potentially unrealistic scheme is to quantify the improvement margin of the protocols mentioned above in terms of throughput.

Routed Mesh [E-MESH] **-** as baseline, we consider an aggressive mesh design where router and link traversals take one clock cycle each when there is no contention, which is achievable with bypass strategies [35]. Multicast is tree-based with deterministic routing. We consider multiport switch allocation and flit forking at the crossbar [21], which allows to keep the router traversal time to one clock cycle even for multicast packets.

4. PERFORMANCE EVALUATION

The networks discussed above are evaluated in terms of latency for low loads and of throughput for high loads. For the latter, we obtain the throughput for a given latency limit to be considered the maximum admissible for the on-chip scenario. Here, this limit is set to 250 clock cycles.

Figure 5 shows how the performance of the different options scales with the number of cores. In this case, we modeled traffic as broadcast with exponential arrivals and uniform source distribution. It is observed that **W-CSMA** offers the best performance in terms of latency, which is of a few clock cycles and is maintained constant for all system sizes for low loads, with a reasonable throughput for high loads. Even though the latency of **W-CARB** slightly increases with the system size due to the growing number of

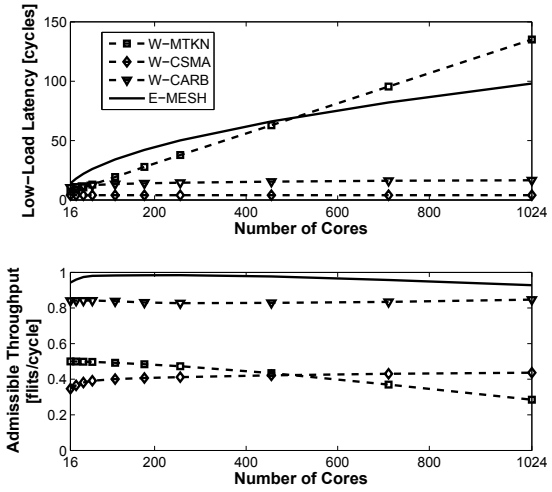

Figure 5: Low-load latency and admissible throughput (at 250 cycles) as a function of the system size.

switch stages that messages need to traverse to reach the buffer, the results are extremely competitive and yielding the best throughput among the wireless options. The propagation of signals to all cores, which takes ∼0.2 clock cycles in this configuration, prevents the throughput from being larger; however, this could be improved by increasing the wireless capacity a bit further. In **W-MTKN**, the latency grows linearly with the number of cores due to the increasing diameter of the token ring, to the point of limiting the admissible throughput for high core counts. Both metrics are improvable, though, by using adaptive designs such as the one proposed in [28] or by letting tokens do more hops in the same clock cycle. Another possible enhancement would consist of partially or completely overlapping the wireless transmission and the passing of the token. Finally, **E-MESH** shows poor latency scalability but the best throughput of all options, accomplished by virtue to the aggressive router design and the huge bisection bandwidth of the network. Note, though, that such high broadcast throughput implies significant in-network contention for the rest of traffic flows.

4.1 Sensitivity Study

Several studies have revealed that on-chip communication in general, and multicast in particular, is potentially bursty and may be concentrated around a reduced number of cores [3, 40]. Burstiness can be modeled through the Hurst exponent H, which takes values between 0.5 (exponential) and 1 (extremely bursty). We generate bursty traffic by alternating ON/OFF periods, the length of which follows Pareto distributions defined by H [26]. Spatial concentration is modeled through a gaussian parameter σ, which takes values between 0 (concentrated) and ∞ (spread out) and describes the percentage of the load to be randomly assigned to each node. Previous results assumed $H = 0.5$ and $\sigma/N = 100$.

Due to the random nature of bursts, we performed 15 runs for each $\{H, \sigma, \lambda\}$ tuple, where λ is the aggregated load, and calculated the geometric mean. Figure 6 the results of the evaluation in a 256-core system for different levels of burstiness and concentration. On the one hand, all networks see their performance significantly reduced as the burstiness of

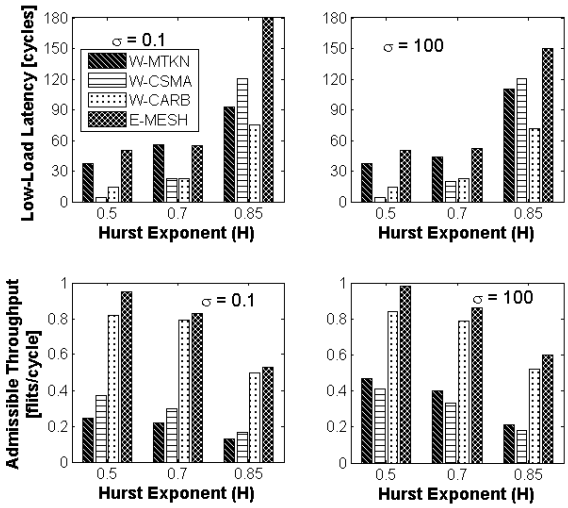

Figure 6: Low-load latency and admissible throughput (at 250 cycles) in a 256-core system as a function of the temporal and spatial characteristics of traffic.

traffic increases: **W-CSMA** and **E-MESH** show particularly dramatic growths in terms of latency, whereas in all cases the throughput is approximately cut to half. On the other hand, the impact of the traffic concentration on performance is generally minor with notable exceptions in **W-MTKN** and **E-MESH**, where the throughput is slightly higher if the load is spread out.

4.2 Energy

In a conventional mesh, broadcasts invariably imply $N-1$ link traversals and buffer writes/reads, as well as a total of $2(N-1)$ flits being output by routers (either towards the next router or when being ejected). Assuming 45nm technology and a 128-bit datapath, low-swing links would approximately consume 40 fJ/bit/hop [35], while the router considered in this work consumes 117 fJ per buffered bit and between 65 and 221 fJ per each bit that performs a crossbar traversal [21]. With all this, the dynamic cost of a broadcast in such an optimized scheme sits between 0.29 and 0.6 pJ/bit/core.

In BoWNoC, the cost will highly depend on the modulation used. Recent implementations of simple On-Off Keying (OOK) working at 60GHz attain speeds around 16Gbps while consuming 17mW at the transmitter and 15mW at the receiver [45]. With these figures and given that in a broadcast we have one transmitter and $(N-1)$ receivers, the total energy consumption would be of around 0.94 pJ/bit/core, assuming a negligible MAC overhead. This is between 2X and 4X larger than in an aggressive switched NoC design, difference that could be compensated by the use of concentration as outlined in Fig. 1. We also envisage a strong reduction of the energy consumption in BoWNoC by virtue of (1) custom designs increasing the power at the transmitter in order to relax the consumption at the receiver, and (2) the use of higher frequencies allowing the design of transceivers with lower energy per bit (see [2]).

5. IMPACT ON FUTURE MANYCORE SYSTEMS

The capabilities and limitations of the on-chip interconnect have traditionally guided the design of parallel architectures and of the algorithms that run over them. For a few cores, buses with ordered broadcast capabilities are feasible and broadcast-based architectures yield better performance. As processors scale and the interconnect design shifts to switched NoCs, though, broadcast becomes a costly feature. This encourages the design of systems that actively avoid it, impacting upon programmability and performance.

Having an effective broadcast plane with uniform latency and total ordering would therefore be invaluable as it potentially relaxes a large number of constraints cast upon architects and parallel programmers. As we discuss next, the potential advantages are manyfold and we speculate that they will be thoroughly investigated once the feasibility of the BoWNoC approach is demonstrated.

5.1 Synchronization Primitives

Synchronization among a large set of cores is one of the functionalities that may benefit most from BoWNoC. Notification and release stages of most methods (e.g. barriers) can be easily implemented with broadcast. Moreover, unique properties of the transmission of RF signals through a shared medium allow the design of alternative schemes for such synchronization methods [4, 33]. By reducing the cost of global synchronization, programmability is improved as maintaining sequential consistency becomes easier and less expensive. Moreover, the bottleneck of synchronization-intensive kernels and applications (e.g. after many OpenMP loops there is an implicit barrier) would be alleviated.

5.2 Architectures

Shared memory systems normally heavily rely on an architecture that, among other aspects, ensures the coherency and consistency of shared data and, by doing so, generates on-chip communication. With the advent of NoCs, broadcast-based snoopy methods for coherence have progressively given way to directory-based protocols that only use multicast to invalidate shared data [37]. However, this comes at the cost of significant area, energy, and performance overheads [9], as well as of higher design complexity.

The availability of inexpensive broadcast may not represent the return of snoopy coherence, but could help increase performance and reduce the complexity of existing protocols. Variants of token coherence [30], which decouple performance and correctness (thereby reducing complexity), generate considerable broadcast. Their use may be extended beyond a few tens of cores. Using renewed synchronization methods, a reduction of the race conditions of coherence protocols could be also envisaged as proposed in [41]. Additionally, explicit communication could be enforced to help coherence protocols in performance-degrading situations (e.g. automatic updates on each write to a producer address in producer-consumer access patterns).

5.3 Algorithms

In systems such as message passing, communication is explicit and the programmer needs to orchestrate it in order to optimize performance and ensure correctness. One-to-all and all-to-all communication routines such as `MPI_Bcast` or

`MPI_Allgather` are available to the programmer, yet their use may be limited in manycore processors due to their increasing cost and reduced performance. The main issue is that certain widespread applications exhibit bottlenecks due to these collective operations [43]. With BoWNoC, collective communication routines could be redefined seeking greater performance and lower cost. This would increase the performance of unmodified versions of the applications; however, larger benefits will be reaped if algorithms are re-defined taking into consideration the improvement of collective communication support. In extreme cases, programmers may want to go back to the original problem and re-implement a solution without global communication constraints.

5.4 Programming Models

The work in [36] discussed several programming models that aim to combine the flexibility of shared medium and the hand-tuned performance of message passing, and that would benefit from an improved broadcast mechanism. One example is *Consumer Tagging*, where the programmer tags addresses susceptible of being producers and the consumers associated to them. Writes to the tagged addresses will be automatically updated without the intervention of the coherence protocol. With broadcast, the programmer does not even need to specify the consumer set. Other examples are *Adaptive Constraint-Based Programming* and *Application Heartbeats*, both of which periodically broadcast information. We also speculate that an efficient broadcast plane could be also used to improve the performance and redundancy properties of the well-known MapReduce model.

6. CONCLUSIONS

As advancements in RF technologies push the operating frequencies towards the mmWave and terahertz bands, the implementation of antennas and transceivers comparable in size with the processing cores becomes feasible. This enables the creation of wireless on-chip networks with unprecedented broadcast capabilities even in manycore settings. We have shown that, when overlaid to a conventional NoC, such networks will provide dramatic performance improvements for a wide range of cases. To harness such potential, it is indispensable to account for scalable, lightweight and effective MAC schemes capable of coping with the stringent communication requirements of manycore chips. Through performance evaluations, we have observed that simple MAC implementations could consistently provide latencies several times lower and throughputs comparable to those of conventional NoCs. As discussed, the availability of such a broadcast medium will have multiple implications on future systems from the architectural and algorithmic perspectives.

Acknowledgments

This work was supported by SAMSUNG under the GRO Program, INTEL under the Doctoral Student Honor Program, and the Technical University of Catalunya (UPC).

7. REFERENCES

[1] S. Abadal, E. Alarcón, M. C. Lemme, M. Nemirovsky, and A. Cabellos-Aparicio. Graphene-enabled Wireless Communication for Massive Multicore Architectures. *IEEE Communications Magazine*, 51(11):137–143, 2013.

[2] S. Abadal, M. Iannazzo, M. Nemirovsky, A. Cabellos-aparicio, and E. Alarcón. On the Area and Energy Scalability of Wireless Network-on-Chip: A Model-based Benchmarked Design Space Exploration. *IEEE /ACM Transactions on Networking*, 23(5):1, 2015.

[3] S. Abadal, R. Martínez, E. Alarcón, and A. Cabellos-Aparicio. Multicast On-Chip Traffic Analysis Targeting Manycore NoC Design. In *Proceedings of the PDP '14*, pages 370–378, 2014.

[4] J. L. Abellán, J. Fernández, and M. E. Acacio. Efficient Hardware Barrier Synchronization in Many-Core CMPs. *IEEE Transactions on Parallel and Distributed Systems*, 23(8):1453–1466, 2012.

[5] N. Abeyratne, R. Das, Q. Li, K. Sewell, B. Giridhar, R. G. Dreslinski, D. Blaauw, and T. Mudge. Scaling towards kilo-core processors with asymmetric high-radix topologies. In *Proceedings of the HPCA '13*, pages 496–507. Ieee, Feb. 2013.

[6] C. Batten, A. Joshi, V. Stojanovic, and K. Asanovic. Designing Chip-Level Nanophotonic Interconnection Networks. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(2):137–153, 2012.

[7] D. Bertozzi, G. Dimitrakopoulos, J. Flich, and S. Sonntag. The fast evolving landscape of on-chip communication. *Design Automation for Embedded Systems*, Apr. 2014.

[8] J. Chan, G. Hendry, A. Biberman, K. Bergman, and L. P. Carloni. PhoenixSim: A Simulator for Physical-Layer Analysis of Chip-Scale Photonic Interconnection Networks. In *Proceedings of DATE '10*, pages 691–696, 2010.

[9] B. Daya, C.-H. O. Chen, S. Subramanian, W.-C. Kwon, S. Park, T. Krishna, J. Holt, A. P. Chandrakasan, and L.-S. Peh. SCORPIO: a 36-core research chip demonstrating snoopy coherence on a scalable mesh NoC with in-network ordering. In *Proceedings of the ISCA-41*, pages 25–36, 2014.

[10] S. Deb, A. Ganguly, P. P. Pande, B. Belzer, and D. Heo. Wireless NoC as Interconnection Backbone for Multicore Chips: Promises and Challenges. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(2):228–239, 2012.

[11] D. DiTomaso, A. Kodi, D. Matolak, S. Kaya, S. Laha, and W. Rayess. A-WiNoC : Adaptive Wireless Network-on-Chip Architecture for Chip Multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, PP(99), 2014.

[12] K. Duraisamy, R. G. Kim, and P. P. Pande. Enhancing Performance of Wireless NoCs with Distributed MAC Protocols. In *Proceedings of the ISQED '15*, pages 406 – 411, 2015.

[13] A. Ganguly, K. Chang, S. Deb, P. P. Pande, B. Belzer, and C. Teuscher. Scalable Hybrid Wireless Network-on-Chip Architectures for Multi-Core Systems. *IEEE Transactions on Computers*, 60(10):1485–1502, 2010.

[14] J. Hennessy and D. Patterson. *Computer architecture: a quantitative approach*. Morgan Kaufmann, 2012.

[15] D. Hou, Y. Z. Xiong, W. Hong, W. L. Goh, and J. Chen. Silicon-based on-chip antenna design for millimeter-wave/THz applications. In *Proceedings of*

the EDAPS '11, pages 1–4, 2011.

[16] N. E. Jerger, L.-S. Peh, and M. Lipasti. Virtual Circuit Tree Multicasting: A Case for On-Chip Hardware Multicast Support. In *Proceedings of the ISCA-35*, pages 229–240. Ieee, June 2008.

[17] A. J. Karkar, J. E. Turner, K. Tong, R. Al-Dujaily, T. Mak, A. Yakovlev, and F. Xia. Hybrid wire-surface wave interconnects for next-generation networks-on-chip. *IET Computers & Digital Techniques*, 7(6):294–303, Nov. 2013.

[18] U. R. Karpuzcu, A. Sinkar, N. S. Kim, and J. Torrellas. EnergySmart: Toward energy-efficient manycores for Near-Threshold Computing. In *Proceedings of the ISCA '13*, pages 542–553, 2013.

[19] J. Kim and K. Choi. Exploiting New Interconnect Technologies in On-Chip Communication. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(2):124–136, 2012.

[20] T. Krishna, C. Chen, W. Kwon, and L. Peh. Smart: Single-Cycle Multihop Traversals over a Shared Network on Chip. *IEEE Micro*, 34(3):43–56, 2014.

[21] T. Krishna, L. Peh, B. Beckmann, and S. K. Reinhardt. Towards the ideal on-chip fabric for 1-to-many and many-to-1 communication. In *Proceedings of the MICRO-44*, pages 71–82, 2011.

[22] T. Krishna and L.-S. Peh. Single-Cycle Collective Communication Over A Shared Network Fabric. In *Proceedings of the NoCS '14*, pages 1–8, 2014.

[23] G. Kurian, J. Miller, J. Psota, J. Eastep, J. Liu, J. Michel, L. Kimerling, and A. Agarwal. ATAC: A 1000-Core Cache-Coherent Processor with On-Chip Optical Network. In *Proceedings of the PACT*, pages 477–488. ACM, 2010.

[24] S. Laha, S. Kaya, D. W. Matolak, W. Rayess, D. DiTomaso, and A. Kodi. A New Frontier in Ultralow Power Wireless Links: Network-on-Chip and Chip-to-Chip Interconnects. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34(2):186–198, 2015.

[25] S.-B. Lee, S.-W. Tam, I. Pefkianakis, S. Lu, M.-C. F. Chang, C. Guo, G. Reinman, C. Peng, M. Naik, L. Zhang, and J. Cong. A scalable micro wireless interconnect structure for CMPs. In *Proceedings of the MOBICOM '09*, page 217, 2009.

[26] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, 1994.

[27] R. Manevich, I. Walter, I. Cidon, and A. Kolodny. Best of both worlds: A bus enhanced NoC (BENoC). In *Proceedings of the NoCS '09*, pages 173–182, 2009.

[28] N. Mansoor and A. Ganguly. Reconfigurable Wireless Network-on-Chip with a Dynamic Medium Access Mechanism. In *Proceedings of the NoCS '15*, 2015.

[29] O. Markish, B. Sheinman, O. Katz., D. Corcos, and D. Elad. On-chip mmWave Antennas and Transceivers. In *Proceedings of the NoCS '15*, 2015.

[30] M. Martin. Token Coherence: decoupling performance and correctness. In *Proceedings of the ISCA-30*, pages 182–193, 2003.

[31] D. Matolak, A. Kodi, S. Kaya, D. DiTomaso, S. Laha, and W. Rayess. Wireless networks-on-chips: architecture, wireless channel, and devices. *IEEE Wireless Communications*, 19(5), 2012.

[32] U. Ogras and R. Marculescu. "It's a small world after all": NoC performance optimization via long-range link insertion. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 14(7):693–706, July 2006.

[33] J. Oh, M. Prvulovic, and A. Zajic. TLSync: support for multiple fast barriers using on-chip transmission lines. In *Proceedings of ISCA-38*, pages 105–115, 2011.

[34] J. Oh, A. Zajic, and M. Prvulovic. Traffic steering between a low-latency unswitched TL ring and a high-throughput switched on-chip interconnect. In *Proceedings of the PACT '13*, pages 309–318, 2013.

[35] S. Park, T. Krishna, C.-H. Chen, B. Daya, A. Chandrakasan, and L.-S. Peh. Approaching the theoretical limits of a mesh NoC with a 16-node chip prototype in 45nm SOI. In *Proceedings of the DAC 2012*, page 398. ACM Press, 2012.

[36] J. Psota, J. Miller, G. Kurian, H. Hoffman, N. Beckmann, J. Eastep, and A. Agarwal. ATAC : Improving Performance and Programmability with On-Chip Optical Networks. In *Proceedings of the ISCAS '10*, pages 3325–3328, 2010.

[37] A. Ros, M. E. Acacio, and J. M. García. Cache Coherence Protocols for Many-Core CMPs. In *Parallel and Distributing Computing*, pages 93–118. 2010.

[38] D. Sanchez, G. Michelogiannakis, and C. Kozyrakis. An Analysis of On-Chip Interconnection Networks for Large-Scale Chip Multiprocessors. *ACM Transactions on Architecture and Code Optimization*, 7(1), 2010.

[39] E. Seok, D. Shim, C. Mao, R. Han, S. Sankaran, C. Cao, W. Knap, and K. K. O. Progress and challenges towards terahertz CMOS integrated circuits. *IEEE Journal of Solid-State Circuits*, 45(8):1554–1564, 2010.

[40] V. Soteriou, H. Wang, and L. Peh. A Statistical Traffic Model for On-Chip Interconnection Networks. In *Proceedings of MASCOTS '06*, pages 104–116, 2006.

[41] D. Vantrease, M. H. Lipasti, and N. Binkert. Atomic Coherence: Leveraging nanophotonics to build race-free cache coherence protocols. In *Proceedings of the HPCA '11*, pages 132–143. Ieee, Feb. 2011.

[42] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. Beausoleil, and J. Ahn. Corona: System implications of emerging nanophotonic technology. *ACM SIGARCH Computer Architecture News*, 36(3):153–164, 2008.

[43] J. Vetter and A. Yoo. An Empirical Performance Evaluation of Scalable Scientific Applications. In *Proceedings of the SC '02*, pages 1–16, 2002.

[44] Y. Wu, K. a. Jenkins, A. Valdes-Garcia, D. B. Farmer, Y. Zhu, A. a. Bol, C. Dimitrakopoulos, W. Zhu, F. Xia, P. Avouris, and Y.-M. Lin. State-of-the-art graphene high-frequency electronics. *Nano letters*, 12(6):3062–7, June 2012.

[45] X. Yu, J. Baylon, P. Wettin, D. Heo, P. Pande, and S. Mirabbasi. Architecture and Design of Multi-Channel Millimeter-Wave Wireless Network-on-Chip. *IEEE Design & Test*, 31(6):19–28, 2014.