Scalability-Oriented Multicast Traffic Characterization

Sergi Abadal*, Raúl Martínez† Eduard Alarcón*, and Albert Cabellos-Aparicio*

*NaNoNetworking Center in Catalonia (N3Cat)
Universitat Politècnica de Catalunya, Barcelona, Spain
Corresponding Email: abadal@ac.upc.edu
†INTEL Barcelona Research Center, Barcelona, Spain

***Abstract*—Multicast on-chip communications are expected to become an important concern as the number of cores grows and we reach the manycore era. The increasing importance such traffic flows directly contrasts with the diminishing multicast performance of current Network-on-Chip (NoC) designs, and has lead to a surge of research works that seek to improve on-chip multicast support. Within this context, one-to-many traffic models may become useful for the early-stage design and evaluation of these proposals. However, existing models do not distinguish between unicast and multicast flows and often do not consider different multiprocessor sizes. To bridge this gap, a multicast scalability analysis is presented, aiming to provide tools for the modeling of multicast communications for NoC design and evaluation purposes.**

***Index Terms*—Multicore Processors, Multicast, Broadcast, On-Chip Traffic Analysis, Network-on-Chip, Scalability**

I. Introduction

Networks-on-Chip (NoCs) are emerging as the way to interconnect the components of a multiprocessor. As recent years have seen a rapid increase in the core density, it is crucial to guarantee the scalability of NoCs to avoid communication to become the performance bottleneck in next-generation multiprocessors. Among other issues, NoCs suffer a significant performance drop in the presence of multicast (one-to-many) or reduction (many-to-one) flows [1].

In the particular case of multicast, messages are generally broken down into multiple unicast packets and served independently. Such approach is not only highly power-inefficient, but also the cause of the aforementioned network performance drop that typically implies a reduction of the multiprocessor performance. It is expected that this effect will exacerbate in denser networks since, as it is shown in this work, multicast flows grow in intensity and number of destinations with the core count. Also, multicast is intensively used in new architectural innovations for many-core processors [2].

For all this, significant research efforts have been recently devoted to improving on-chip multicast support (see [1]–[4] and references therein). Perhaps due to the lack of realistic multicast models, these approaches have been generally tested using synthetic traffic and considering a fixed network size. As a result, their impact upon the network performance is imprecise and their scalability remains largely unknown. Given that one-to-many traffic models may become useful for the early-stage design and evaluation of NoCs, in this paper we characterize multicast traffic from a scalability perspective as a first step towards complete traffic models. Unlike existing characterization efforts [5]–[7], our approach differentiates between unicast and multicast flows and is not bound to a given network size.

TABLE I
Simulation Parameters

Number of cores	4, 8, 16, 32, 64
L1 Cache (I & D)	32 KB, 2-way, 2 cycles
L2 Cache	512 KB, 8-way, 10 cycles
Coherency	MESI, HyperTransport (HT)
Main Memory Latency	30 ns
Network-On-Chip	2D-Mesh, 1-cycle link, 5-cycle router

II. Framework

To perform a scalability-oriented multicast traffic analysis, we simulate different multiprocessors running SPLASH-2 and PARSEC (*simmedium* input set) in order to obtain a set of traces. All multiprocessors share the same basic architecture (see Table I) and use the number of cores and coherency mechanism as parameters. Simulations are carried out with GEM5 [8], which has been slightly modified so that the network interfaces would register the time of arrival, origin, destinations, type and size of each multicast message that needs to be injected into the NoC. Since GEM5 only admits up to 64 cores thus far, scalability trends obtained with such methodology could be used to extrapolate a first approximation of the multicast traffic requirements of manycore processors.

III. Results

Multicast Traffic Intensity: The number of multicast messages per instruction is a NoC-agnostic measure of the multicast intensity. Figure 1 (left) plots the number of multicast messages per one million instructions for both MESI and HyperTransport (HT) coherence schemes in the test configuration. It is observed that HT has multicast requirements one order of magnitude larger than that of MESI. More importantly, it is shown that most applications become more multicast intensive as the number of cores grows. Although such increase is application-dependent and does not follow a common scaling

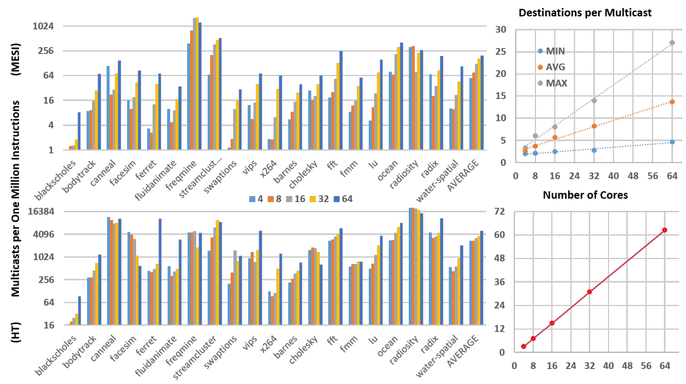

Fig. 1. Number of multicast messages per 10^6 instructions (left) and number of destinations per multicast (right) as a function of the number of cores, assuming MESI (top) or HT coherence (bottom).

trend, fitting methods on the average values yield a logarithmic relation between multicast intensity and number of cores. Application scalability limitations may explain such tendency.

Number of Destinations: This is an important metric given that the performance of conventional NoCs is inversely proportional to the number of destinations per message. Figure 1 (right) shows the how the number of destinations per multicast (averaged over all the applications) scale with the number of cores N. The number of destinations scales as $\mathrm{O}(\sqrt{N})$ and as $\mathrm{O}(N)$ when assuming MESI and HT coherence, respectively. In the former case the metric is application-dependent; whereas, in the latter case, the trend is application-independent since the coherence protocol issues a broadcast for each coherence operation.

Spatial Distribution: The study of the injection spatial distribution may be useful for the identification of hot spots, which are especially concerning in the case of multicast. To express this traffic characteristic, we calculate the coefficient of variation (COV) of the number of injected multicast per node as $c_v = \sigma/\mu$, where σ and μ are the standard deviation and mean of the multicasts injected by each node. A higher COV means a higher concentration of the multicast injection over given cores. Figure 2 shows how the COV (averaged over all the applications) scales with the number of cores in MESI and HT. In both cases, the concentration increases with the core count.

Temporal Distribution: Besides enabling the identification of periodic multicast-intensive phases, studying the temporal distribution of multicast message injection provides knowledge on the burstiness of such traffic. Related works have shown that on-chip traffic is bursty (self-similar) in general [5] and, provided that multicast traffic is a subset of the on-chip traffic, it is reasonable to deduce that multicasts will also exhibit self-similarity. We calculated the Hurst exponent H ($0.5 < H \leq 1$, a value close to 1 denotes strong self-similarity) applying the RS plot method [5] to the full-system traces and the averaging it over all the applications. In light of the results of Figure 3, it can be concluded that multicast traffic is self-similar and that burstiness increases with the core count.

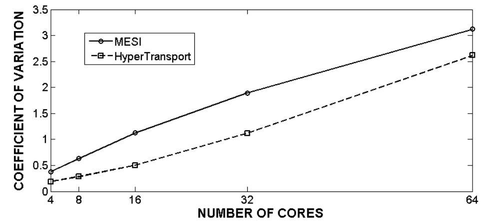

Fig. 2. Averaged coefficient of variation of the spatial injection distribution for MESI and HT as a function of the number of nodes.

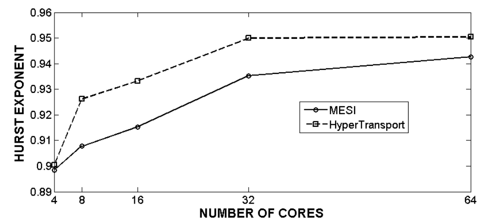

Fig. 3. Averaged Hurst exponent for MESI and HT as a function of the number of nodes.

IV. Conclusions

Our multicast traffic characterization shows that the one-to-many communication requirements of MESI and HT grow with the number of cores due to an increase of both the number of multicasts per instruction and the destinations per multicast. Further, spatial injection unbalance and temporal burstiness worsen with the core count. In light of this, the need for multicast-efficient NoCs becomes patent as we reach the manycore era.

Acknowledgments

The authors gratefully acknowledge support from INTEL through the Doctoral Student Honor Program.

References

[1] 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 *44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-44)*, vol. 2, 2011, pp. 71–82.

[2] N. E. Jerger, L.-S. Peh, and M. Lipasti, "Virtual Circuit Tree Multicasting: A Case for On-Chip Hardware Multicast Support," in *2008 International Symposium on Computer Architecture*. Ieee, Jun. 2008, pp. 229–240.

[3] J. Kim and K. Choi, "Exploiting New Interconnect Technologies in On-Chip Communication," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems (JETCAS)*, vol. 2, no. 2, pp. 124–136, 2012.

[4] S. Abadal, E. Alarcón, M. C. Lemme, M. Nemirovsky, and A. Cabellos-Aparicio, "Graphene-enabled Wireless Communication for Massive Multicore Architectures," *IEEE Communications Magazine*, vol. 51, no. 11, pp. 137–143, 2013.

[5] V. Soteriou, H. Wang, and L. Peh, "A Statistical Traffic Model for On-Chip Interconnection Networks," in *14th IEEE International Symposium on Modeling, Analysis, and Simulation*. Ieee, 2006, pp. 104–116.

[6] J. Bahn and N. Bagherzadeh, "A generic traffic model for on-chip interconnection networks," in *Proceedings of the 1st International Workshop on Network on Chip Architectures (NoCArc)*, 2008, pp. 22–29.

[7] N. Barrow-Williams, C. Fensch, and S. Moore, "A communication characterisation of Splash-2 and Parsec," in *IEEE International Symposium on Workload Characterization (IISWC)*, 2009, pp. 86–97.

[8] N. Binkert, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, D. a. Wood, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, and T. Krishna, "The gem5 simulator," *ACM SIGARCH Computer Architecture News*, vol. 39, no. 2, 2011.