A study of NoC Exit Strategies

Mikael Millberg & Axel Jantsch
KTH - Royal Institute of Technology, Sweden
{micke, axel}@imit.kth.se

The throughput of a network is limited due to several interacting components. Analysing simulation results made it clear that the component that was worth attacking was the exit bandwidth between the network and the connected resources. The obvious approach is to increase this bandwidth; the benefit is a higher throughput of the network and a significant lowering of the buffer requirements at the entry points of the network; this because worst case scenarios now happens at a higher injection rate. The result we present shows significant differences in throughput as well as in average and worst case latency.

Offering services with best effort performance, naturally, gives no hard guarantees due to the dynamic behaviour of any general purpose system. In order to make use of such services statistical performance measures are instead utilised. This leads to that the traffic often has to be below a certain threshold for which the desired statistical properties can be given. These properties can be derived, either, from a rigid reasoning based on the implementation or from an analysis of simulations. From this analysis the desired properties can be given with a safety margin.

Given that we are bound to offer services with statistical characteristics on performance, how do we do this to a low cost? The cost in this context is the required buffers needed to guarantee no packet losses together with a safety margin in terms of injection rate to "guarantee" a worst, and average, case latency.

The approach that we have chosen within, Nostrum [1], for giving the service of *Best effort*, is by utilising deflective routing, with no explicit buffering, to keep the size of the switches small[2]. Through simulations with uniform random traffic patterns, we observed that there seems to exist a "hard" limitation on the network. On a 4x4 mesh this limit is reached for an injection rate of 0.63 packets/node/cycle for the best performing routing strategy tried out. Once this limit is reached packets start queuing up at the entry points of the network and the worst case latencies grows exponentially.

Analysing simulation results made it obvious that an increased bandwidth between the network and the connected resources would make the network perform better. The price for this solution is that packets now need to be buffered at the exits of the network, but from an overall perspective the total buffers required is lowered for a moderately to heavily loaded network. At the limit injection rate the average latency was reduced from 25 to 15 clock cycles and the observed worst case latency from 280 to 180 clock cycles. This will give better margins before the network saturates or a higher throughput with the previous margin kept. All this assumes that there no single node or bisection cut of the network are exposed to a static over-utilisation.

The validity of the chosen approach is not restricted to uniformly random traffic patterns on meshes but also applicable to "any" topology where the traffic pattern involves potential network exit congestions due to multiple sources having the same destination and/or multiple routing paths are possible.

The network exit strategy has not received as much attention as other parts of network design. Most work that in detail analyse cost and performance of a router and the network as a whole, e.g. [3, 4, 5] assume an ideal packet ejection model, which means that packets are absorbed by the receiving node as soon as they are delivered by the network. In [6] an ejection policy is studied that reduces the cost and complexity of the router while minimizing the impact on performance. However, to our knowledge no study about the trade-offs involved in increasing the network exit bandwidth has been reported.

[1] M Millberg et al. The Nostrum backbone - a communication protocol stack for networks on chip. VLSI Design Conference, January 2004.
[2] E. Nilsson et al. Load distribution with the proximity congestion awareness in a network on chip. DATE 2003.
[3] A. A. Chien. A cost and speed model for k-ary n-cube wormhole routers. IEEE Transactions on Parallel and Distributed Systems, 9(2):150-162, Feb. 1998.
[4] L. S. Peh and W. J. Dally. A delay model for router micro-architectures. IEEE Micro, 2001.
[5] E. Rijpkema, et al. Trade offs in the design of a router with both guaranteed and best-effort services for networks on chip. DATE, 2003.
[6] Z. Lu and A. Jantsch. Flit ejection in on-chip wormhole-switched networks with virtual channels. NorChip, 2004.