The Effects of Blocked Requests on Simulated Run Time for Ouroborus ONoC Architecture

Bryanna Davison
University of California Berkeley, EECS Department

Alex Berlanga
University of California Santa Barbara, College of Engineering

Lei Zhang
University of Maryland Eastern Shore, Department of Engineering and Aviation Sciences

Motivation

Advances in computation have slowed as traditional electronic signal transmission speeds have plateaued. Further electronic signals are susceptible to interference from neighboring wires. Optical signals can be transmitted at a much faster speed and don't suffer from as much interference as electrical signals. Optical signals also generally use less power and generate less heat. However even with these advantages, optical chip architecture is a much less mature field than it's electrical counterpart. Optical Network on Chip (ONoC) architecture is an auspicious design for optical computing applications. It features low latency, power consumption, and high bandwidth.

We focus on the Ouroborus Network ONoC architecture for our research. In this model, processing cores sit on a ring as seen in Figure 1. By changing the ordering of these cores, improvements can be realized in simulated runtime and latency per packet. Changing the ordering of these cores is known as reconfiguring the network and any set of these 16 nodes in a certain order is known as a reconfiguration. Using splash-II benchmarks, we can evaluate differences between different reconfigurations.

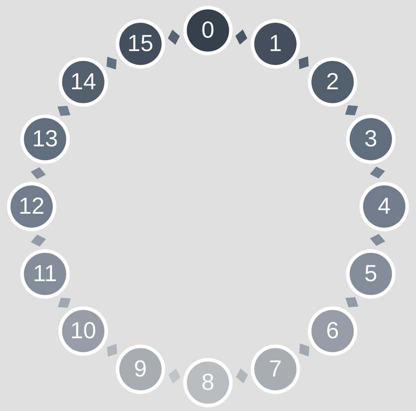

Fig. 1 Visualization of Ourobors Network

We propose that the simulated runtime of a ONoC is correlated more strongly to the total number of times any request is blocked than the communication cost of how traffic is assigned.

Approach

We created a simulation of the Ouroborus Network in python. It simulates running Splash-II benchmarks on the Ouroborus network. Benchmarks are read into the simulation as log files with volumes, source and destination core locations, and traffic start times. This simulation was used to find all of the data presented.

Results

To discuss our results we'll be using data collected from running the vips benchmark with five reconfigurations :

- Numerical (e.g. nodes arranged 0 to 15)
- Three randomly ordered
- Nearest Neighbor sorted (heuristic algorithm)

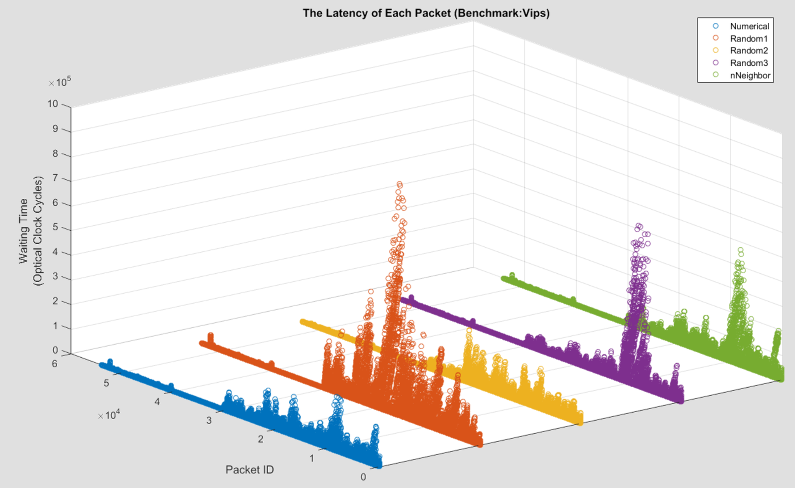

Fig 2. The latency experienced by each individual packet as it is processed by the ONoC. Example of benchmark vips.

Fig 3. Simulator results from running vips benchmark

Vips Benchmark Data		
Simulation Run Time (OCC)	Times Blocked	Com. Cost
33907460	3755630939	159493120
41935780	9692131935	169289728
34270620	2670864164	160273920
32197520	1632682983	170331648
34379180	3445467009	166157824

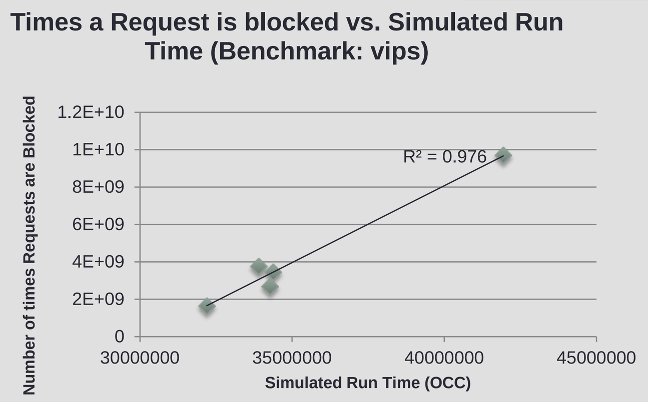

Fig. 4 Data obtained from simulator

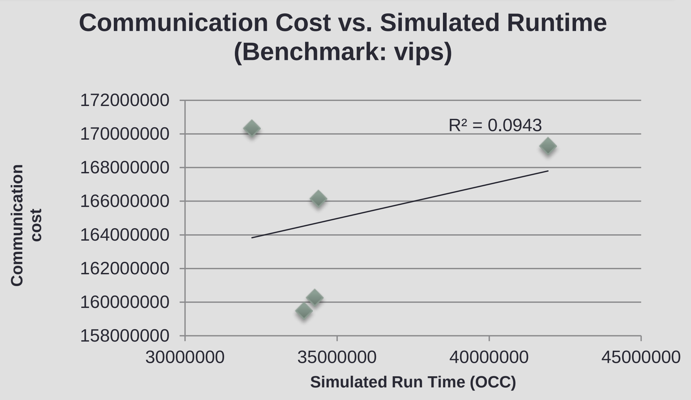

Fig. 5 Data obtained from simulator

In figure 2, latency caused by each configuration can be seen visually. The Random1 reconfiguration shows much higher latency on many packets than the other reconfigurations. This is also reflected in its run time, which was the longest of the 5 tested reconfigurations (fig. 3). This indicates some reconfigurations are more optimized for processing requests quicker and that it is related to their run times. We compared the simulator results from each configuration to the calculated communication cost and the number of times requests were blocked. Communication cost was calculated using this equation:

$$Comunication\ Cost\ = \sum Cores\ Occupied\ by\ Request \times \sum DataVolume$$

We found that the number of blocked requests had a much higher correlation to the simulated run time when compared using a linear regression fit. Figures 4 and 5 show this result graphically with the vips benchmark data. The correlations of the data sets of other reconfigurations can be seen below in figure 6.

Benchmark	r^2 Num Requests Blocked vs. Simulated Run Time	r^2 Communication cost vs. Simulated Run Time
vips	0.976	0.0943
swaptions	0.8393	0.7847
streamcluster	0.8597	0.1428
raytrace	0.9666	0.03
freqmine	0.776	0.2803
fluidanimate	0.4597	0.0697
dedup	0.7961	0.2582
Average:	0.8104	0.2371

Fig. 6 An analysis of simulation data. Table of correlation of determinations using a linear fit performed in Excel.

We found that for the benchmarks we ran that the average correlation of determination was 0.81 for comparing the number of blocked requests to simulated runtime while the average for comparing communication cost to simulated run time was 0.24.

Conclusions

By minimizing the number of hops between high traffic nodes it was assumed that performance could by reducing the communication cost. However from the results, it can be understood that a reduction in communication cost does not necessarily lead to an increase in performance, but when reducing the amount of delay for each request, a linear correlation can be seen. Therefore, the criteria for creating a reconfiguration algorithm should aim to reduce this delay and consider the architecture of the Ouroboros optical Network-on-Chip.

Future Directions

After this preliminary conclusion, we would like to test more configurations and obtain more data to support our proposal that packet delay is what we should focus on optimizing. With this data, we would like to begin to explore algorithms to reconfigure the Ouroboros network more efficiently. Also we plan to add additional functionality to our simulator, such as adding more channels for communication between cores and simulating power consumption.

Salisbury UNIVERSITY

Acknowledgements

Dr. Enyue Lu for her dedication to the EXERCISE REU
The Nationaal Science Foundation
Salisbury University

References

[1] Rand Morris, Avinash Karanth Kodi, and Ahmed Louri, "Dynamic Reconfiguration of 3D Photonic Networks-on-Chip for Maximizing Performance and Improving Fault Tolerance," presented at the 2012 IEEE/ACM 45th Annual International Symposium of Microarchitecture, Vancouver, BC, 1-5 Dec 2012.
[2] Lei Zhang, Xianfang Tan, Mei Yang, and Yingtao Jiang, "A Centralized Optical Network-on-Chip Architecture with Space-Division Multiple Access," presented at the 2014 IEEE Optical Interconnects Conference, San Diego, CA, USA, 4-7 May 2014.