3D Fabric for Multicore Resilient Systems

The Master and PhDs
Javad Bagherzadeh, Sugandha Gupta, Byoungchan Oh

1 INTRODUCTION

Technological trends into the nanometer regime have led to significantly higher failure rates. Consequently, high reliability and fault tolerance are now getting more emphasis. We are attempting to solve these issues of reliability and tolerance on a simple pipeline, generally used in many-core designs and GPUs. We have based our work on StageNet [1], which uses a reconfigurable and adaptable network of replicated and isolated processor pipeline stages to maximize the useful lifetime of a chip. The network is formed by replacing the direct connections at each pipeline stage boundary by a crossbar switch interconnection.

2 OUR IDEA

We propose to use 3D structure to deal with these issues. By taking advantage of the third dimension, we can place more units in shorter distance with each other in different layers and make the communication latency between them minimum by connecting them vertically using Through Silicon Vias (TSVs).

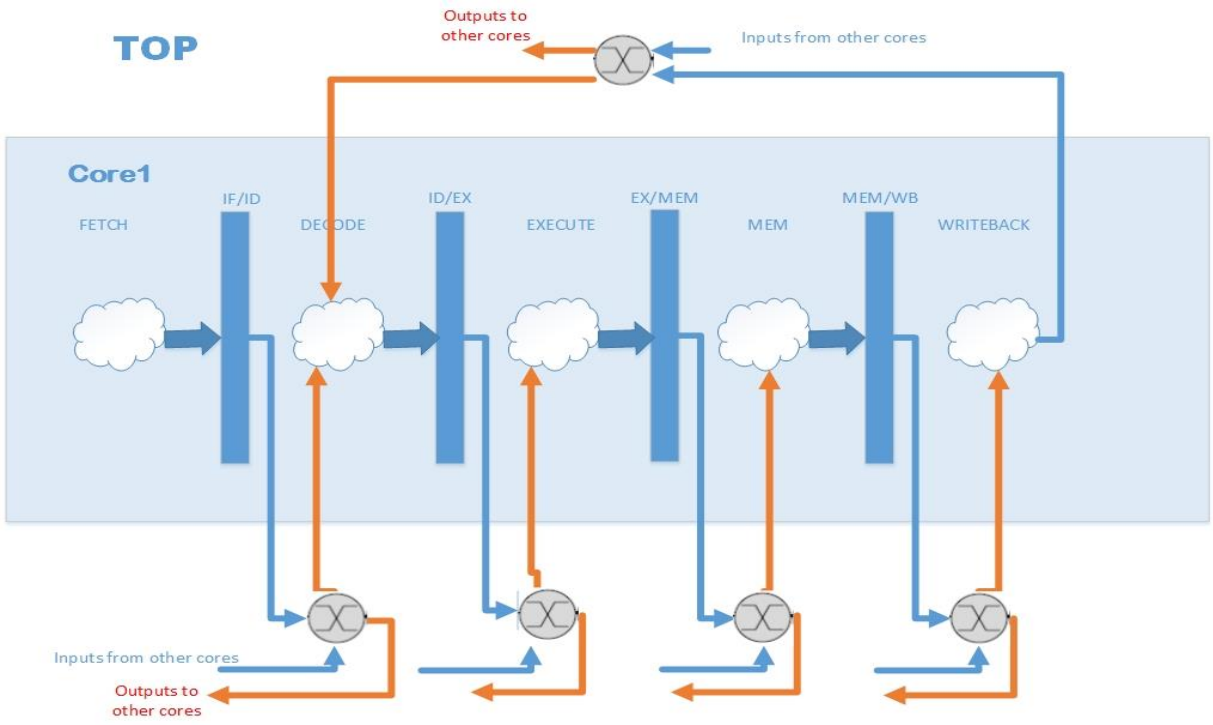

Figure 1: A 5-stage pipeline structure with 5 units and switches between stages

3 PROGRESS

We have implemented the baseline 2D-StageNet architecture in system verilog. It has 4 cores, each with 5 pipeline stages. Hence, there are 5 cross-bar switches for each of the pipeline stages and are connected to each of the cores. Every signal goes through the switch to go to next stage of every core as shown in Figure 1. For simplicity, we have shown only one core here and only the inputs and outputs of the rest of the cores.
We have developed our test benches and verified the design for all the cores. The cross-bar switch functionality has been verified too. All of the 2D design has been simulated and verified at the gate-level also.

4 CHALLENGES AND SOLUTIONS

This section describes certain challenges we faced till now in the implementation and how we resolved them.

In our design, in case of a fault in a particular stage, the stage gets bypassed because of the cross-bar switch. But in case of a fault-free circuit, we had to decide how the internal pipelines stages of a particular core would be connected. One option was to connect them directly with each other and the other option was to connect them through the cross-bar switch every time. The latter is ideally not required for the circuit to function as going through the crossbar would increase the critical path and hence reduce performance. But we chose the second one, as we thought it was unwise to have two different paths during execution in a faulty scenario.

Another design decision we had to make was regarding the layout of the cores. First option is to have each pipeline stage on one 3D layer. For example, the fetch stage will be in layer 1, decode in layer 2, and so on. Second option is to have one complete core with all the pipeline stages in one single 3D layer. A disadvantage of using the second floor plan design is that it would create hot-spots due to big units like execute stacked on the top of each other. This would not happen in the former case, and it would have better heat dissipation. Another disadvantage of the second case is that, core 0 (in the bottom layer) might need to talk to core 3 (in the top layer) which would result in a delay equivalent to 3 TSVs. While, in the first one, only adjacent layers need to communicate with each other and hence have will only 1 TSV delay. We are yet to decide which option to choose.

5 FUTURE STEPS

Next, we plan to run SPICE simulations for the estimation of latencies of the TSVs. Also we plan to make the layout for our 4 cores, to be used for both 2D and 3D design and simulations soon.

6 REFERENCES

[1] Gupta, Shantanu, et al. "The stagenet fabric for constructing resilient multicore systems." Proceedings of the 41st annual IEEE/ACM International Symposium on Microarchitecture. IEEE Computer Society, 2008.