3D Fabric for Multicore Resilient Systems

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

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 using 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.

1 OUR IDEA

We are using 3D circuits 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).

2 PROGRESS AND RESULTS

We have implemented both 2D-StageNet architecture and simulated the 3D design. We have done Synthesis and Place and Route for both the designs. For the 2D design, we have the 4 core design with crossbars between each stage. In order to approximately simulate the 3D design, we have used 1-core simulation results clubbed with the TSVs delay obtained from [2].

Initially, the previous design adapted from EECS 470, had a very long critical path because the complete multiplier was in the execute stage rather than being pipelined. So we were not seeing any major improvements in our timing in 3D design vs 2D design. For instance, for a tclk =13ns, 2D crossbars added 1ns delay and 3D TSV added 0.2ns delay, which was not appreciable improvement compared to tclk. So for the time being, we have removed the multiplier from our design in order to have a better representation of a simple core that have almost equal distribution in all the stages. To improve about this, we might later change the multiplier to a pipelined 8 or 16-stage multiplier. Though, this is something that is not critical to our results.

Table 1. basically shows the working frequency for each design. The result for 2D design has been obtained from the layout for 4-core design. We tried to increase the frequency to obtain the maximum frequency that design would work without any negative slack. For 3D design, as the inter-core connection is happening through TSVs, we added the TSV delay to clock period of each individual core. First, we created the layout for one core and it corresponding crossbars and found the maximum clock frequency in which the design would work. Then, we added the TSV delay [2] to the critical path delay.

As we are assuming that that in no fault condition, cores would not use outside crossbars and pipeline stages would connect internally, the frequency for both of them is the same.

	2D design	3D design
No Fault	**3 ns (333 MHz)**	**3 ns (333 MHz)**
Fault	**7 ns (143 MHz)**	**3.45 ns (290 MHz)**

Table 1. Simulation Results for 2D and 3D design in case of faulty and fault-free scenario

But, in case of faults there is more than 2x improvement in Tclk for the 3D design as compared to baseline 2D design. Our results seem to be promising and we hope that we would be able to improve them with further simulations. More detailed simulation results can be found in the Appendix section.

3 FUTURE STEPS

Initially, we had started with our baseline 2D model and the 3D model with crossbar switches and TSVs respectively. Later on further analysis we have decided to make improvements to the baseline design in 2D itself. We will add pipelines to the crossbars switches so that we could increase the frequency of operation with a tradeoff of increased CPI. Also, the simulation will for various fault injections. We injected single and multiple faults manually in different stages of the pipeline.

4 APPENDIX

Table 2. shows our detailed simulation results for 3 designs - 1 core with switches, 1 core with switches and 4 cores with switches. We came up with the timing for the 3 designs after synthesis and place and route both. The 4 core design timing indicates the timing for the 2D baseline. And the 1 core timing along with TSV timing was used to estimate the 3D cores timing and frequency. The ones in yellow are the results without the multiplier and are more appropriate. We were able to reduce the Tclk for the designs on subsequent runs as can be seen in the table. While doing these simulations we tried to keep the cell density same, in order to have a fair comparison between the designs.

Design	Clock	Slack (syn)	Slack (apr)	Area	Density
One Core with switches	10ns	met	violated(-0.004)	600*600	0.705
	3ns	met	violated(-0.018)	600*600	0.86
	13ns	met	met	600*600	0.7
	5ns	met	met	600*600	0.7
	4ns	violated			
	4.5	met			

4 Cores with switches (flatten)	**10ns**	**met**		**1200*1200**	**0.713**
	8ns	**met**		**1200*1200**	**0.713**
	7ns	**met**			
One Core baseline (without switches)	**3ns**	**met**	**violated(-0.004)**	**600*600**	**0.78**
	6ns	**met**	**violated(-0.004)**	**600*600**	**0.67**
	8ns				
	3ns	**met**	**met**	**600*600**	**0.67**

Table 2. (Our simulations are running. These are intermediate results)

5 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.

[2] You, Jhih-Wei, et al. "Performance characterization of TSV in 3D IC via sensitivity analysis." Test Symposium (ATS), 2010 19th IEEE Asian. IEEE, 2010.