Virtualization of Deeply Pipelined Magnetologic

Michael J. Hall, Roger D. Chamberlain
Department of Computer Science and Engineering

Washington University in St. Louis
School of Engineering & Applied Science

Supported by
NSF Grants CNS-0931693, CNS-1527510,
Exegy, Inc., and VelociData, Inc.

Motivation

SpinFET Magnetic Tunnel Junction

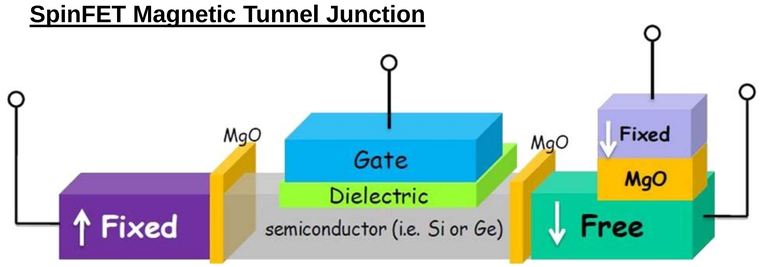

Nikonov and Young, "Overview of beyond-CMOS devices and a uniform methodology for their benchmarking," *Proceedings of the IEEE*, vol. 101, no. 12, Dec. 2013.

Magnetologic gates

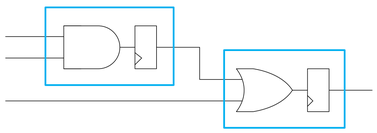

Conventional gates propagate signals combinationally.

Magnetologic gates have state, meaning that each gate may act as a pipeline stage.

For a large circuit, this can become a deeply pipelined circuit.

Magnetologic circuit with feedback

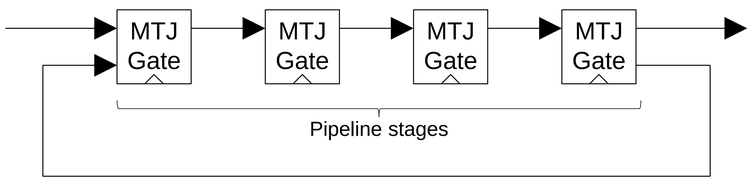

Queueing Model

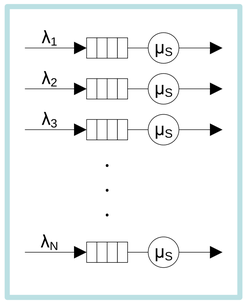

Model definition:

Tput, *Latency*, *Occupancy* = *f* (*Circuit*, *Tech*, *C*, *N*, *S*, R_S, λ)

Variable	Definition
Circuit	Logical circuit description (e.g. written in Verilog or VHDL)
Tech	Target technology (e.g. Magnetologic, FPGA, or ASIC)
C	Pipeline depth (fine-grain contexts)
N	Total number of contexts (requires secondary memory if *N*>*C*)
S	Cost of a context switch (for secondary memory)
R_S	Scheduling period (number of rounds of *C* contexts that execute before context-switching to secondary memory)
λ	Arrival rate (e.g. data elements per second)
OL	Offered load (the ratio of the aggregate arrival rate (of all streams) to the peak service rate of the system (i.e., when *S* = 0))

M/G/1 modeling expressions

Service rate:
$$\mu_s = \frac{R_S}{(R_S N + SN/C)\cdot t_{CLK}} \ \text{elements/s}$$

Total achievable throughput:
$$T_{TOT} = N\cdot\mu_s = \frac{R_S}{(R_S + S/C)\cdot t_{CLK}} \ \text{elements/s}$$

Waiting time expressions: (Latency)
$$W = C\cdot t_{CLK} + \left[\frac{(1-p_s)\,T_V t_{CLK}}{2}\right] + \left[\frac{p_s C\cdot t_{CLK}}{2}\right] + \left[\frac{\lambda C^2 t_{CLK}^2}{2(1-\rho)}\right] + \left[\frac{\rho T_V t_{CLK}}{(1-\rho)\,R_S}\right] \ \text{seconds}$$

$$T_V = R_S(N-C) + SN/C \qquad p_s = \frac{R_S C}{R_S N + SN/C}$$

$$\rho = \lambda/\mu_s \qquad OL = N\cdot\lambda\cdot t_{CLK}$$

Hall and Chamberlain, "Using M/G/1 queueing models with vacations to analyze virtualized logic computations," in *2015 33rd IEEE International Conference on Computer Design (ICCD)*, Oct 2015, pp. 78–85.

Secure Hash Algorithm (SHA-2) Results

Secure Hash Algorithm (SHA-2)

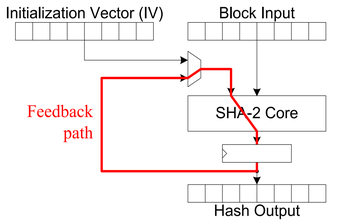

Total achievable throughput plot

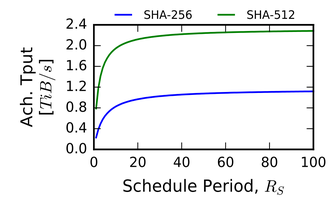

Given:
- *Tech*=Magnetologic, *N*=2*C*, *S*=2,000, *OL* varies
- *Circuit*=SHA-256 / 512

Design Params: R_S

Latency and schedule period optimization plots

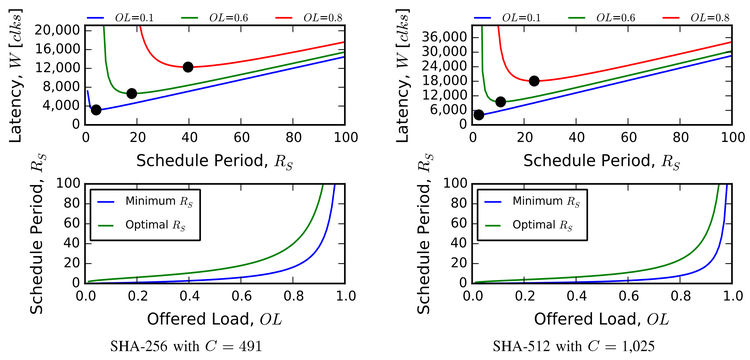

SHA-256 with $C = 491$ SHA-512 with $C = 1{,}025$

Virtualized Logic Computation

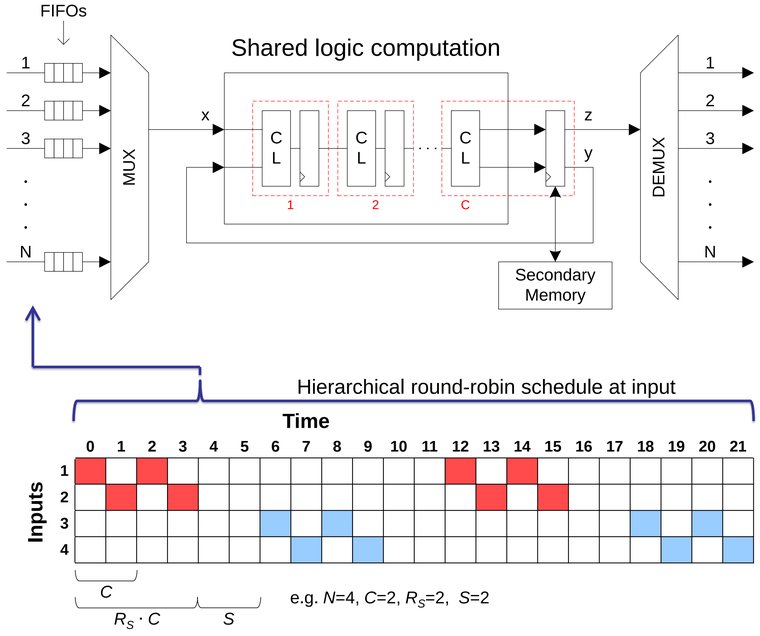

Linear Congruential Generator (LCG) Results

Linear Congruential Generator (LCG)

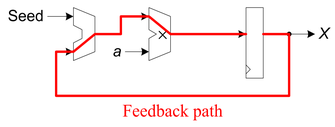

$$X_{i+1} = a\cdot X_i$$

where *a* = 65,539.

- Estimate a pipeline depth at *C* = 100 (the ratio of 32-bit adder delay to fanout-4 (FO4) inverter delay is approximately 50, and LCG requires two adders).
- 20 GHz clock rate (t_{CLK} = 50 ps) (for a SpinFET MTJ in Nikonov and Young 2013)

Results

Tech=Magnetologic, *N*=*C*, *S*=0

	Random Numbers	Bytes
Single-stream throughput	200 million/s	760 MiB/s
Total achievable throughput	20 billion/s	75 GiB/s

Conclusions and Future Work

Conclusions

- Using the M/G/1 model, we show the relationship between circuit parameters, offered load, throughput, and latency.
- We show dramatic aggregate throughput gains when multiple "virtual" copies of a (deeply pipelined) magnetologic circuit are exploited.
- Virtualized logic computations are able to utilize deeply pipelined circuits.

Future Work

- Generalize the arrival process assumptions (Markovian) to incorporate other distributions. A real system may act quite differently (e.g. buffering up data and sending in bursts).
- Extend the model to support additional scheduling algorithms. The current scheduling algorithm (hierarchical round-robin) is not work-conserving (an empty queue will still get scheduled).