Understanding Timing Error Rates For Adder Architectures

Yuanbo Fan, Russ Joseph
Department of Electrical Engineering and Computer Science
Northwestern University

Introduction

Circuit-level timing speculation has been proposed as a technique to improve overall system efficiency by eliminating overheads arising from worst-case design assumptions. Under this design paradigm, the processor runs at voltage, frequency, and thermal operating points which would not guarantee signal setup time constraints for all logic paths. The system is augmented with timing error detection and correction techniques, so that a timing error no longer leads to catastrophic system failure. Instead, we can trade off error rate for energy savings.

To understand the relationship between the gate-level circuit and timing error rates, we use different static CMOS adder implementations and find the timing error rates for each architecture.

ADDER ARCHITECTURES

- Computations in adders can be decomposed into three parts.
- Adders are distinguished by the prefix network.
- Different adder architectures have different characteristics in terms of delay, area and energy consumption

Precomputation: $G_{i:i} = A_i \cdot B_i \quad P_{i:i} = A_i \hat{} B_i$, $G_{0:0} = C_{in} \quad P_{0:0} = 0$

Prefix: $G_{i:j} = G_{i:k} + P_{i:k} \cdot G_{k-1:j}$, $P_{i:j} = P_{i:k} \cdot P_{k-1:j}$

Postcompuation: $S_i = P_i \hat{} G_{i-1:0}$

Ripple Carry

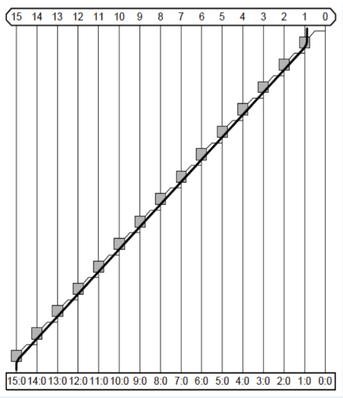

Kogge-Stone

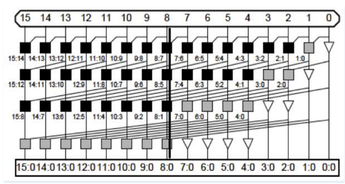

Brent-Kung

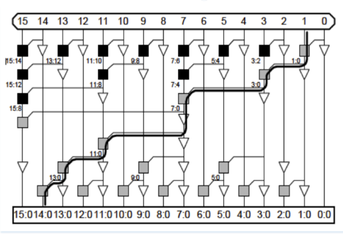

Han-Carlson

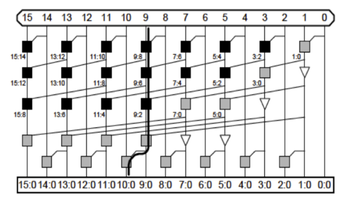

Sklanskey

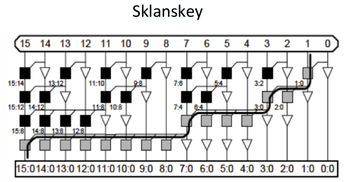

Knowles

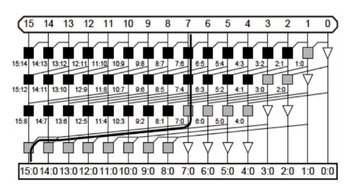

Ladner-Fischer

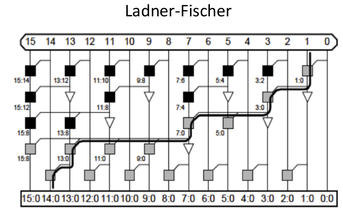

EXPERIMENTS & RESULTS

Experiments:

- Inputs
  Adder-related instructions sampled from SPEC2006 Integer Benchmarks (bzip2, gcc, mcf, perlbench)
- Designs
  32-bit Adders (Ripple Carry, Brent-Kung, Sklansky, Kogge-Stone, Han-Carlson, Knowles, Ladner-Fischer)
- Outputs
  Delay for each data input (Timing-Error Rates calculated at various frequencies)

Observations:

- Comparatively, some benchmarks exhibit lower error rates across all adders.
- Given a fixed frequency, there is some consistency in adder performance across benchmarks.
- At a fixed error rate, different clusters of adders have their own critical operating frequency.
- Different clusters of adders have their own growth rate .

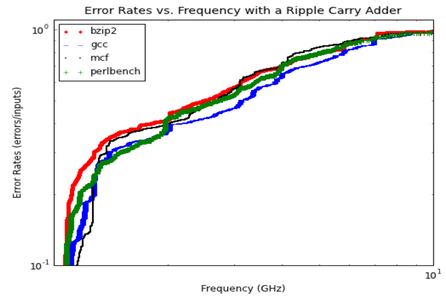

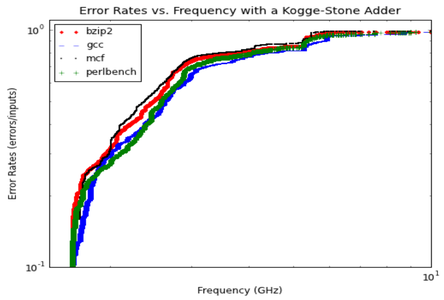

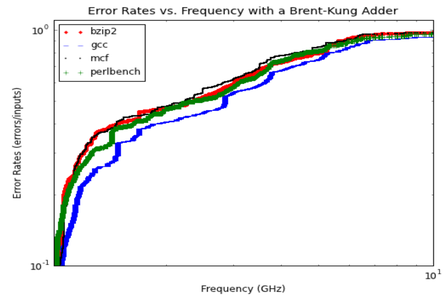

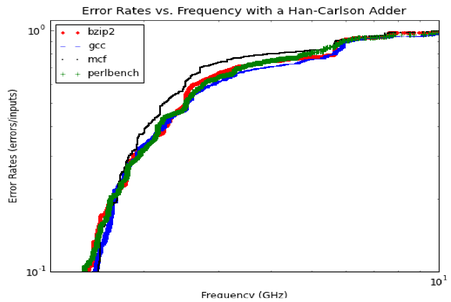

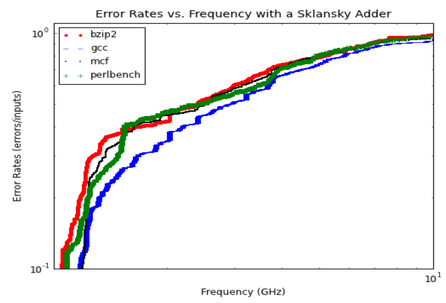

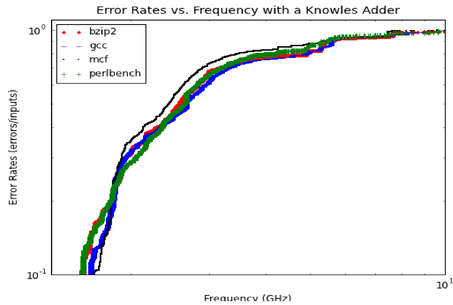

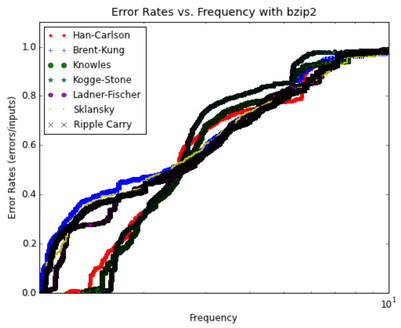

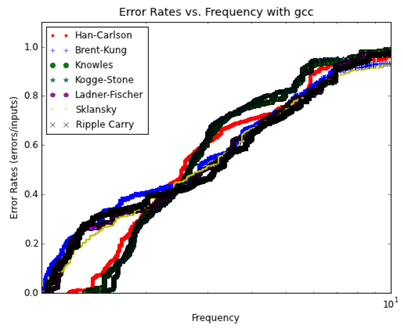

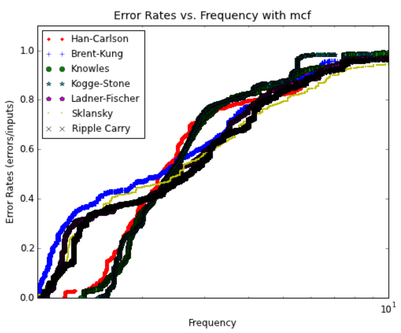

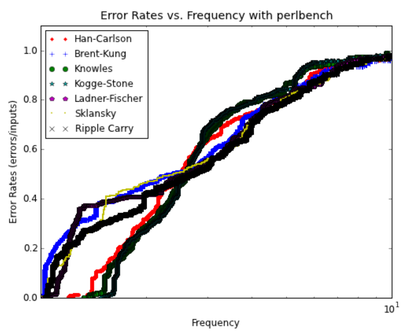

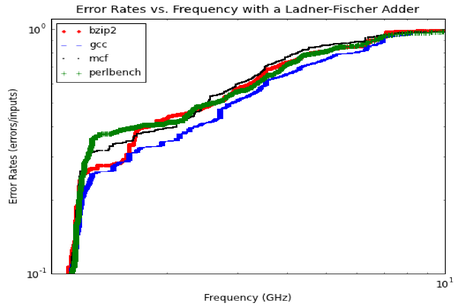

Conclusion & Future Work

- There is substantial variation in timing error rate patterns across adder architectures and workloads.
- We plan to investigate heterogeneous designs that can exploit good pairing between circuit structure and workloads.