Correlation Between State and Control Signals in Out-of-Order Issue Logic

Ken Barr, Ken Conley and Serhii Zhak
Checkpoint I: October 19, 2000

Modifications to Project Goals

- More data mining, less implementation
  - Abstracted from particular hardware architecture
  - Correlation statistics
    - What control signals occur at moment instruction is fetched?
    - How consistent is instruction's path through pipeline?
  - Register renaming statistics
    - Quantitative measure of how quickly register resources freed
    - Use as metric in correlation statistics

Preliminary Results

- Correlation statistic <u>does not</u> include register renaming, yet

- Initial statistics show 50-70% rate of consistency for instruction path

  - Run with SpecInt95 benchmark suite

  - Measurement based on mathematical mode

  - Bimodal distribution giving lower-than-expected results

Register Renaming Issues

- Register renaming or large register file necessary for OOO performance

- Should/does register renaming need to be included in prediction?
  - Prediction for register renaming depends on consistent renaming
  - Perfectly consistent renaming equivalent to no renaming (ignoring ISA benefits)

- SimpleScalar register renaming not as versatile as we had hoped

Examining Renaming

- Simplescalar RUU scheme ties renaming to all aspects of OOO execution.

- Hard to isolate renamer
  - We extend Simplescalar to manage a "free list" of physical registers
  - Also, add per-instruction renaming statistics
    - Min/max/avg length of "renamedness"
    - Number of "remaps"
    - Re-execution counter

- Following slides: 500K instructions of a Perl benchmark with 16 RUU_Stations.

Reservation Length

- Every instruction needs physical register for ≥ 3 cycles.

- 80% relinquish their mapping within 7 cycles.

- This gives hope that a scheme could quickly reassign a mapping

- We need to compare this with average CPI
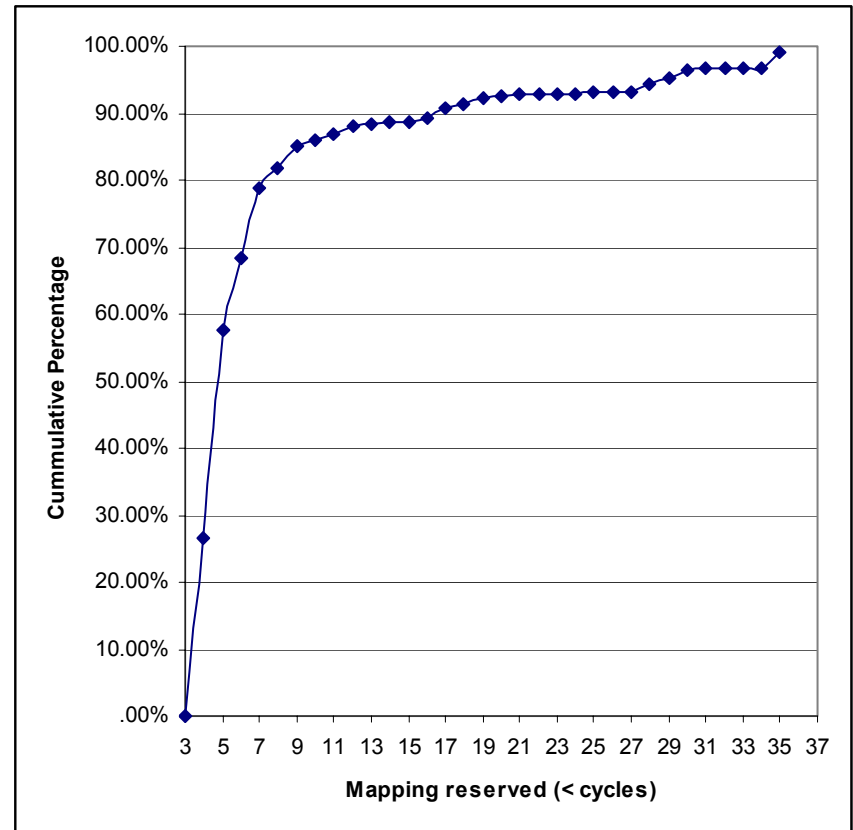
Remap Rate

- remap rate = samename / repetitions

- Discouraging chart!

- But there's some good news:
  - "Samename" counts only mappings that match the first mapping.

    Perhaps two-bit-counter or confidence mechanism could enhance this.
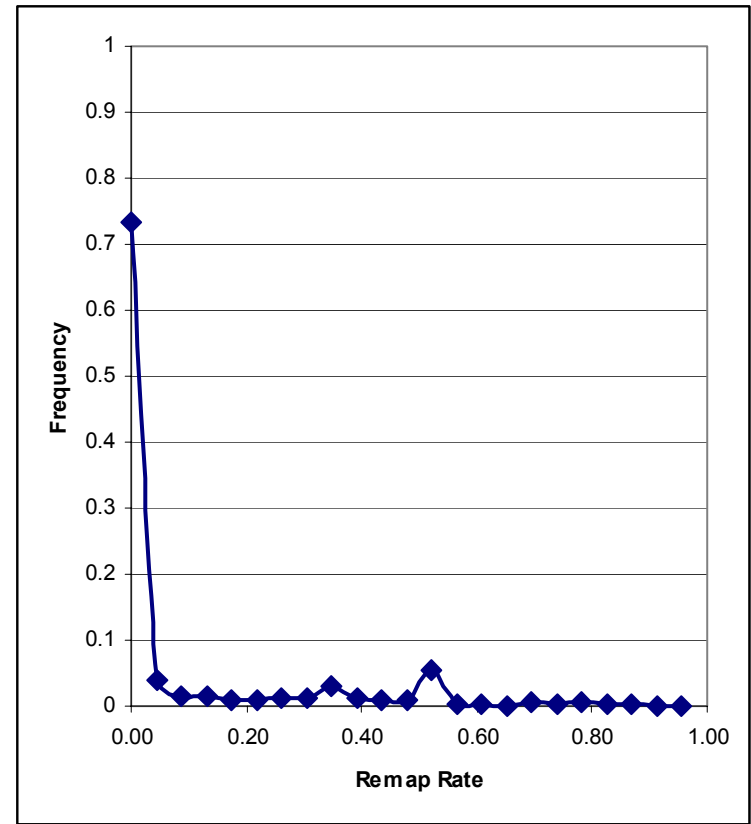
Progress to Checkpoint 2 and Beyond!

- Final version of data mining engines

- Investigate bimodal distributions

- Build custom test patterns

- Run more, longer simulations

- Continue work on SimpleScalar for collection of statistics