Notes 10/7/14

Conte paper- “Benchmark Evaluation for Experimental System Evaluation”

- How do you determine locality?
    - Some would say by cache miss behavior
    - Caches are much bigger now, so this metric is less relevant
- Purpose of characterizing benchmarks
    - Design machines to match workloads
    - Find good subset
    - If the model for a program and a system are coupled together, simulations are not as necessary
- A trace
    - Sequence of occurrences
    - Can be represented as w(t) = ri
        - Where ri is in the set R (e.g. an instruction)
    - The paper examines how long until certain occurrences are repeated with next(w(t)) function.
    - Another measurement is how many things happen between an occurrence repeating with u(w(t)) function
- For a stream of data that doesn't repeat, we will observe only spatial hits, no temporal hits
- Interreference temporal density function shows how many unique references will probably happen between references to a certain item
    - If an item doesn’t repeat, interreference density is infinite
    - This can predict performance for LRU buffers
- Interreference spatial density function is the probability of the a future reference being a certain spatial distance away.
- Can use a stack to keep track of how many unique references are between a certain reference
    - Similar to LRU policy