

ELF: Efficient Lightweight Fast Streaming Processing at Scale

Liting Hu

PhD candidate@Georgia Tech

<http://www.istc-cc.cmu.edu/>



Our Goals

Fast

High throughput for batch processing jobs;

Low latency for streaming processing jobs and queries;

Scale

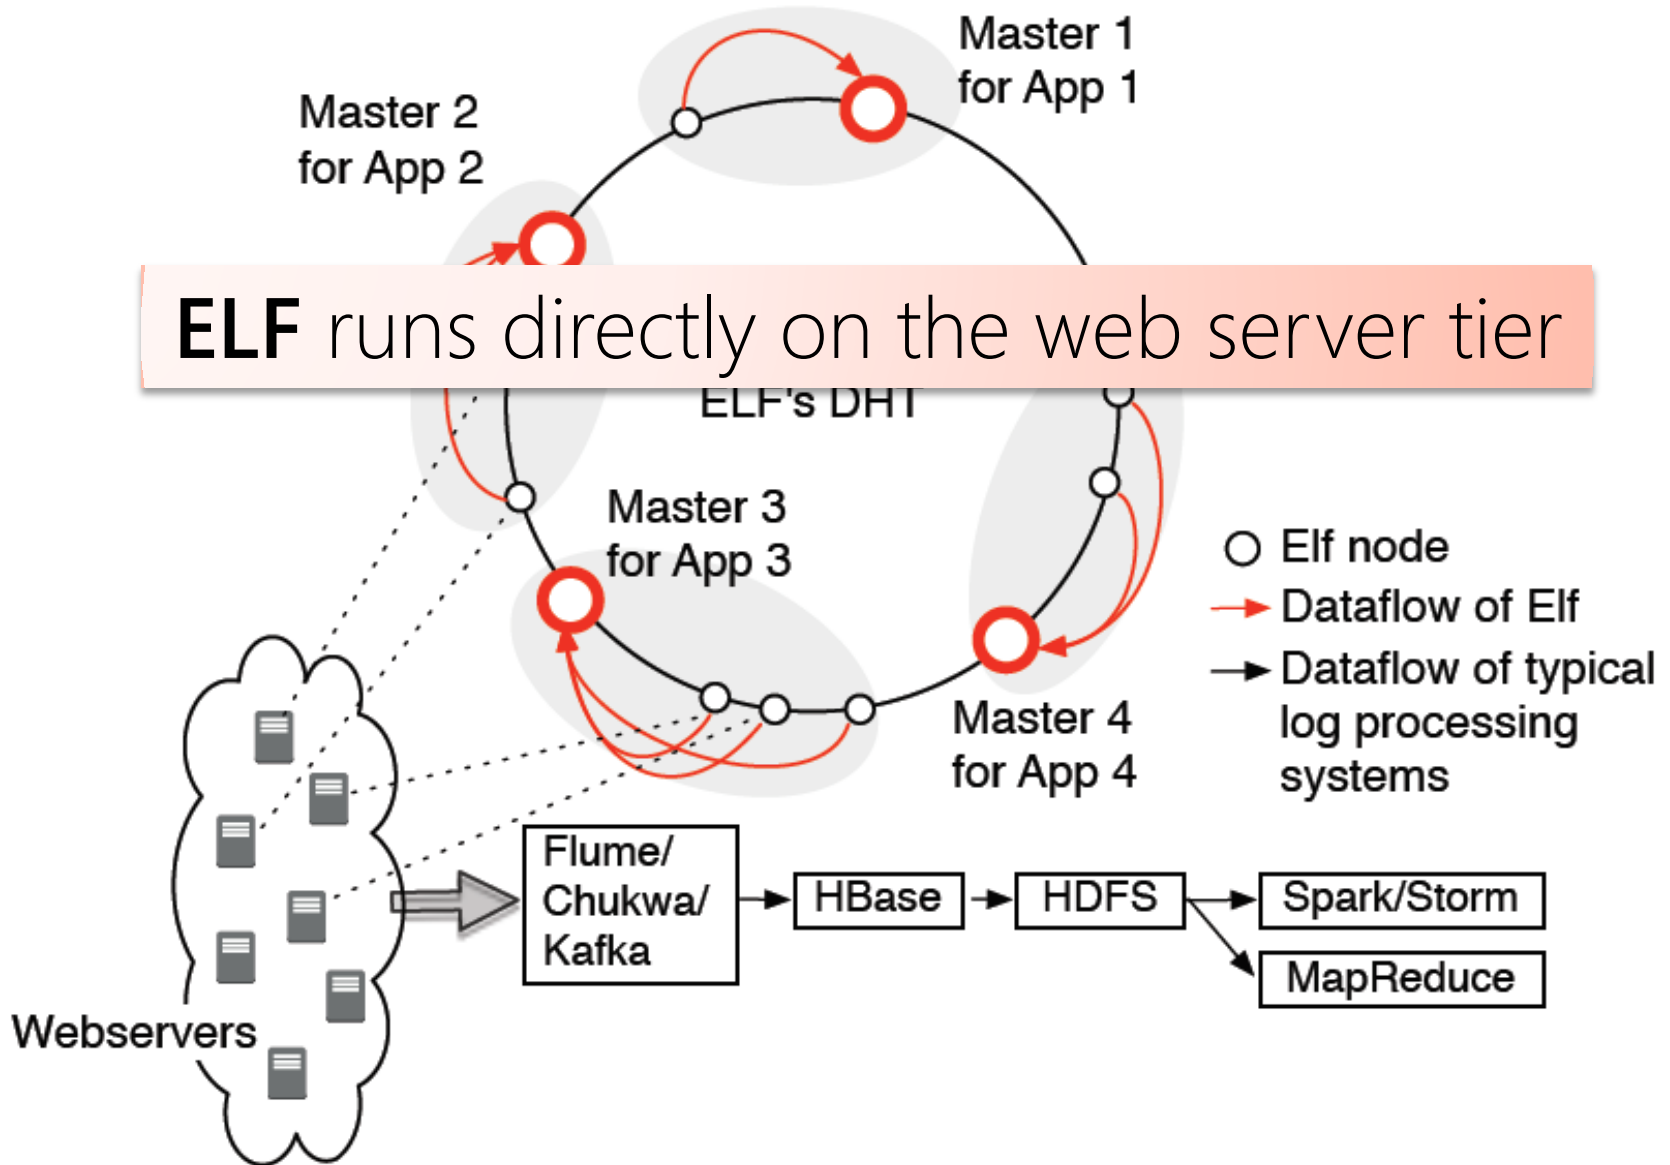
Scale with hundreds of **diverse parallel jobs**, particular for jobs with diverse execution logics (e.g., pipeline or cyclic dataflow), different batch sizes, different sided sliding windows;

Flexible

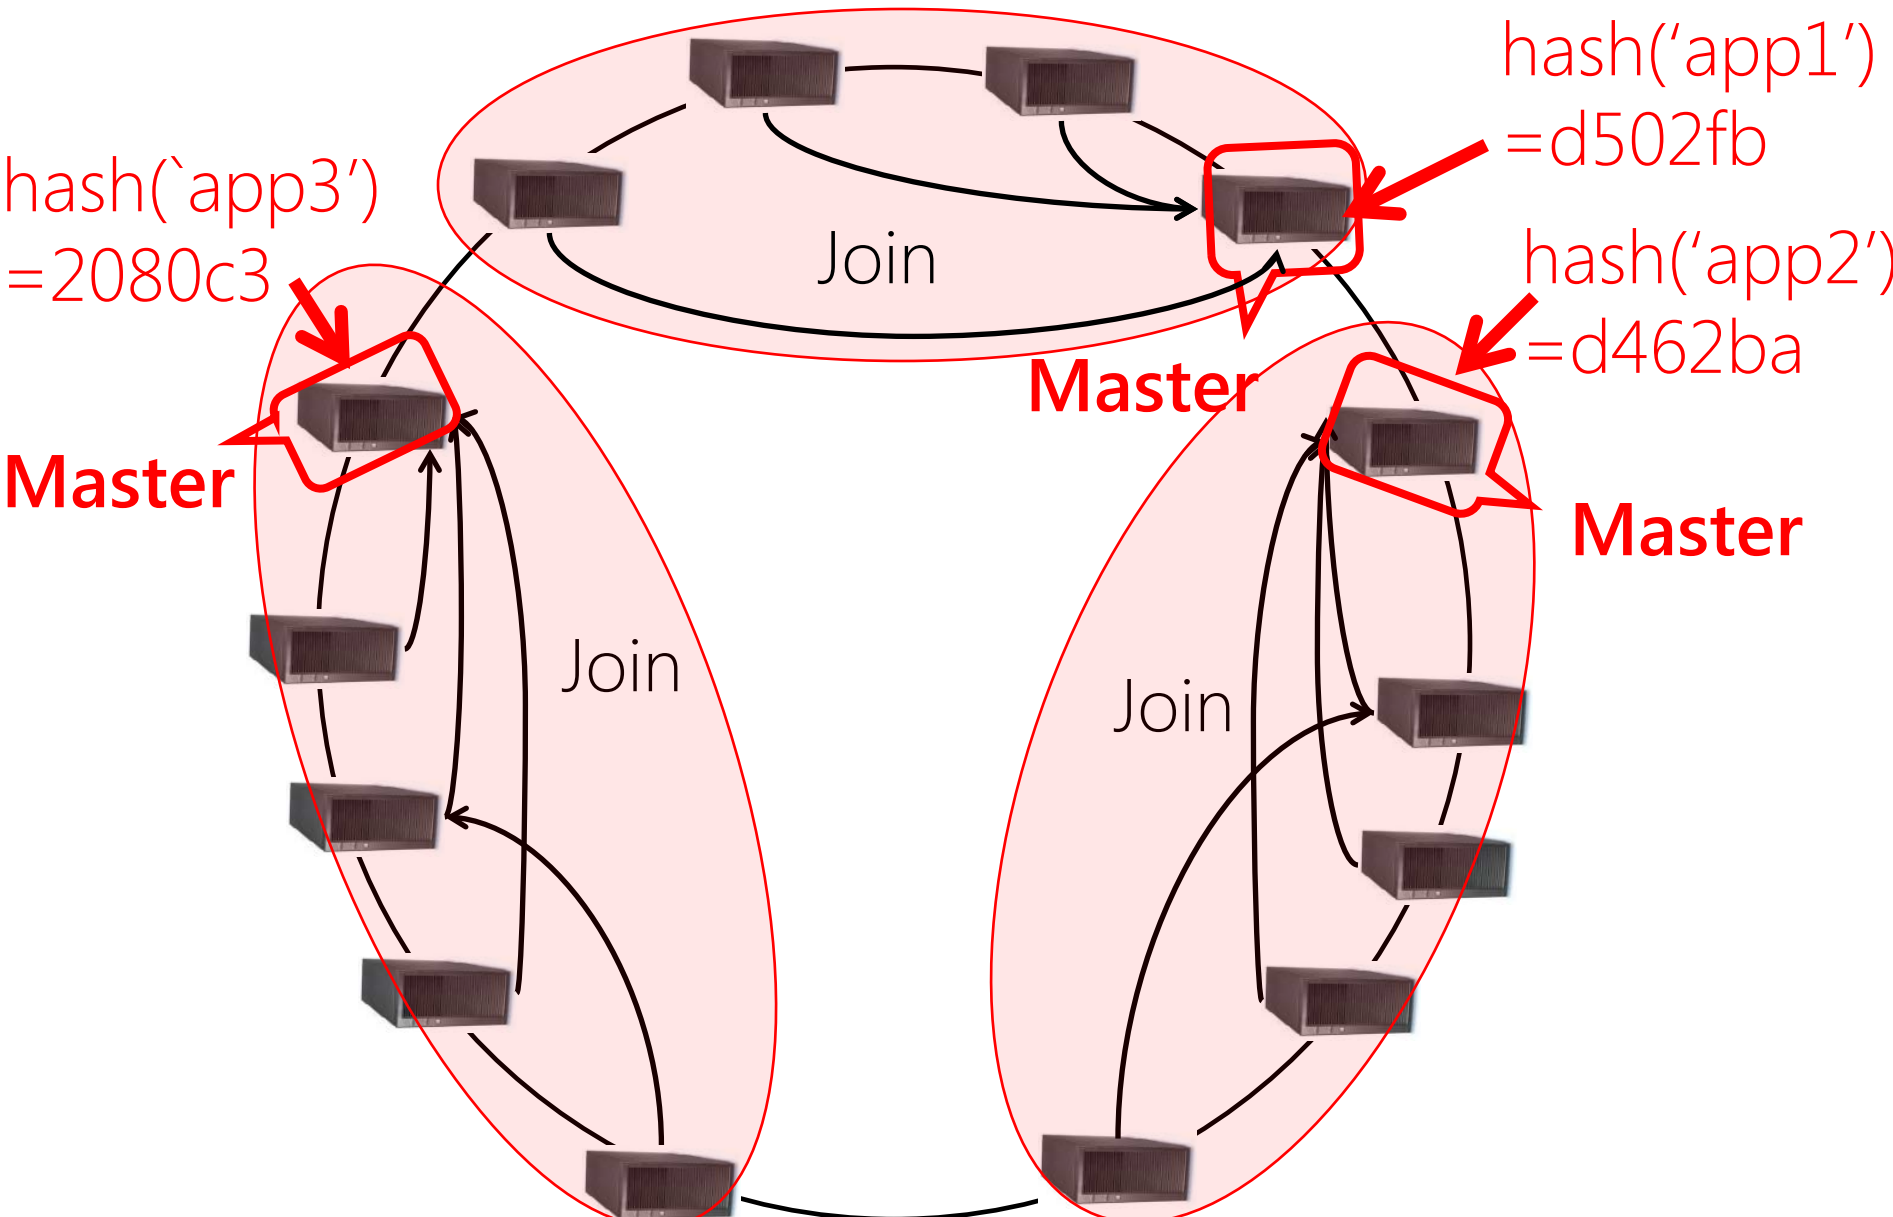
Flexible **coordination** between jobs;

Flexible job **functions changes** at runtime;

Our Design



Many Masters/Many Workers



ELF Software Architecture

