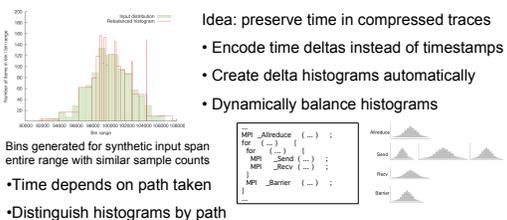


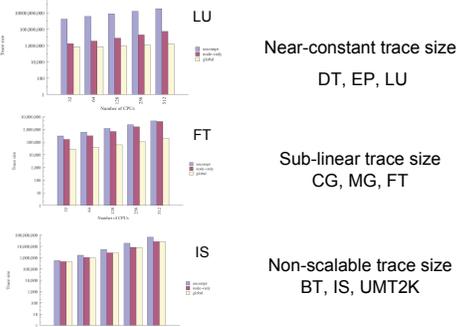
An Open Framework for Scalable, Reconfigurable Performance Analysis

Todd Gamblin[‡], Prasun Ratn[†], Bronis R. de Supinski^{*}, Martin Schulz^{*}, Frank Mueller[†], Robert J. Fowler[‡], Daniel Reed[‡]
^{*}Lawrence Livermore National Laboratory, [†]North Carolina State University, [‡]Renaissance Computing Institute

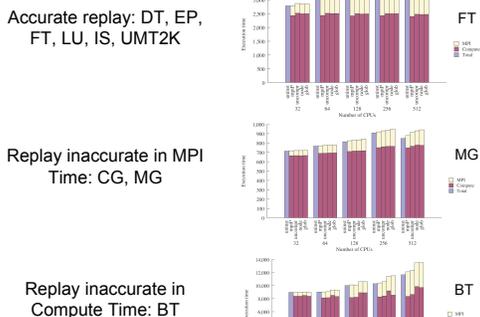
ScalaReplay: Replay Using Histogram Timing Annotations



Trace sizes (NAS Benchmarks and UMT2K)



Replay accuracy (NAS Benchmarks and UMT2K)

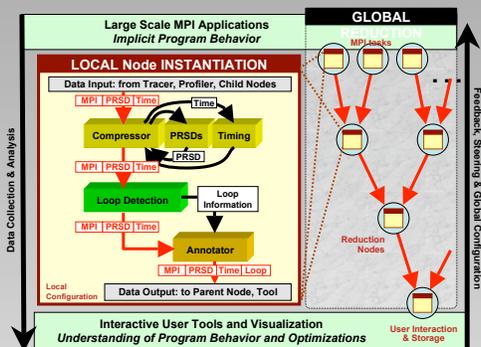


Description of problem

- Size of machines is rapidly increasing (130,000+ processors)
- Tools will be overwhelmed with data
- Need scalable, online measurement and analysis

ScalaTrace: Reconfigurable, Scalable Performance Analysis

Trace representation using Power Regular Section Descriptors (PRSDs)



- Near-constant low-overhead MPI traces
- Annotate with additional reconfigurable data, e.g.
 - Time using adaptive histograms (left)
 - Progress rates, load imbalance (right)

Future Directions

- Flexible framework for application-specific tools
 - Ability to select compression schemes for different fields, data types
 - Foster combination, interaction between data collection, analysis mechanisms
 - Adaptive, self-tuning runtime systems
- Near term:
 - Adaptive wavelet transform topology, equivalence class detection
 - Full time-sequence compression

Evolutionary Load-Balance Analysis with Scalable Data Collection

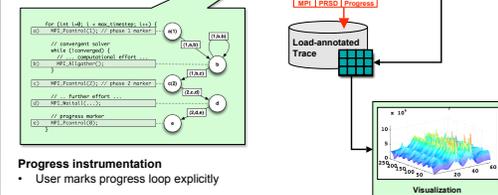
Idea: Normalize measurements and models based on application semantics.

Progress loops

- Typically outer loops in SPMD codes
- Indicate absolute progress towards some domain-specific goal
- Basis for comparison of load over time

Effort loops

- Variable-time loops, represent load
- Data-dependent execution



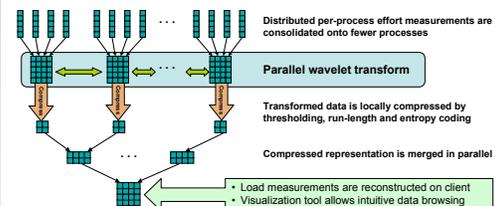
Progress instrumentation

- User marks progress loop explicitly

Effort modeled with code regions

- Dynamically detected at runtime
- Split MPI-op trace at collectives and wait operations
- User can further divide code into phases with instrumentation

Scalable Compression Using Wavelet Transform



Load Balance in ParaDIS

Models dislocation dynamics in crystals

- Dislocations discretized as nodes & arms
- Recursive spatial domain decomposition
- Balancer subdivides nodes/arms along x, then y, then z

