

# Nixing Scale Bugs in HPC

Saeed Taheri and Ganesh Gopalakrishnan (University of Utah, Salt Lake City, UT)  
Sindhu Devale and Martin Burtscher (Texas State University, San Marcos, TX)  
[www.cs.utah.edu/fv/HybridDebugger](http://www.cs.utah.edu/fv/HybridDebugger)



XPS Exploratory Award CCF 1438963, 1439002

## Objectives

### Challenges of HPC Software Debugging Include

- Heterogeneity: hardware (CPUs, GPUs, ...)
- Extreme scale: number of threads/cores
- Rapid evolution: new CPUs/GPUs/libraries
- Reality: High manual effort to annotate code
- Reality: No tools that collect enough debugging information per large-scale run with low overhead (Service Units or Core Hours get exhausted, precluding “second run”)

### Wish List

A tool for gathering information from HPC software that maintains a history of events and their causal relationships such that, upon failure detection, one can query and navigate the history to narrow down the likely cause of the bug.

### Challenges in Realizing These Goals

- **Scalability**
  - Detecting memory corruption (e.g., out of bounds access) and data races requires heavy-weight instrumentation
- **Solution**
  - For now, focus on synchronization/control bugs (deadlocks, livelocks)
  - Higher degrees of scrutiny on relatively newer pieces of code

- **Handling Heterogeneity**
  - Tracking control-flow / happens-before across different types of execution hardware has not been addressed before
- **Solution**
  - Develop synchronization action collection methods across CPUs, GPUs, and code written under different concurrency models

## Prior Work (exemplars)

### Commercial Debuggers

- Very good at detailed trace collection
- Good at minutely examining execution state
- Poor at handling scale
- Little help toward identifying the root cause
- Little help bridging concurrency models

### Research Lab Tools

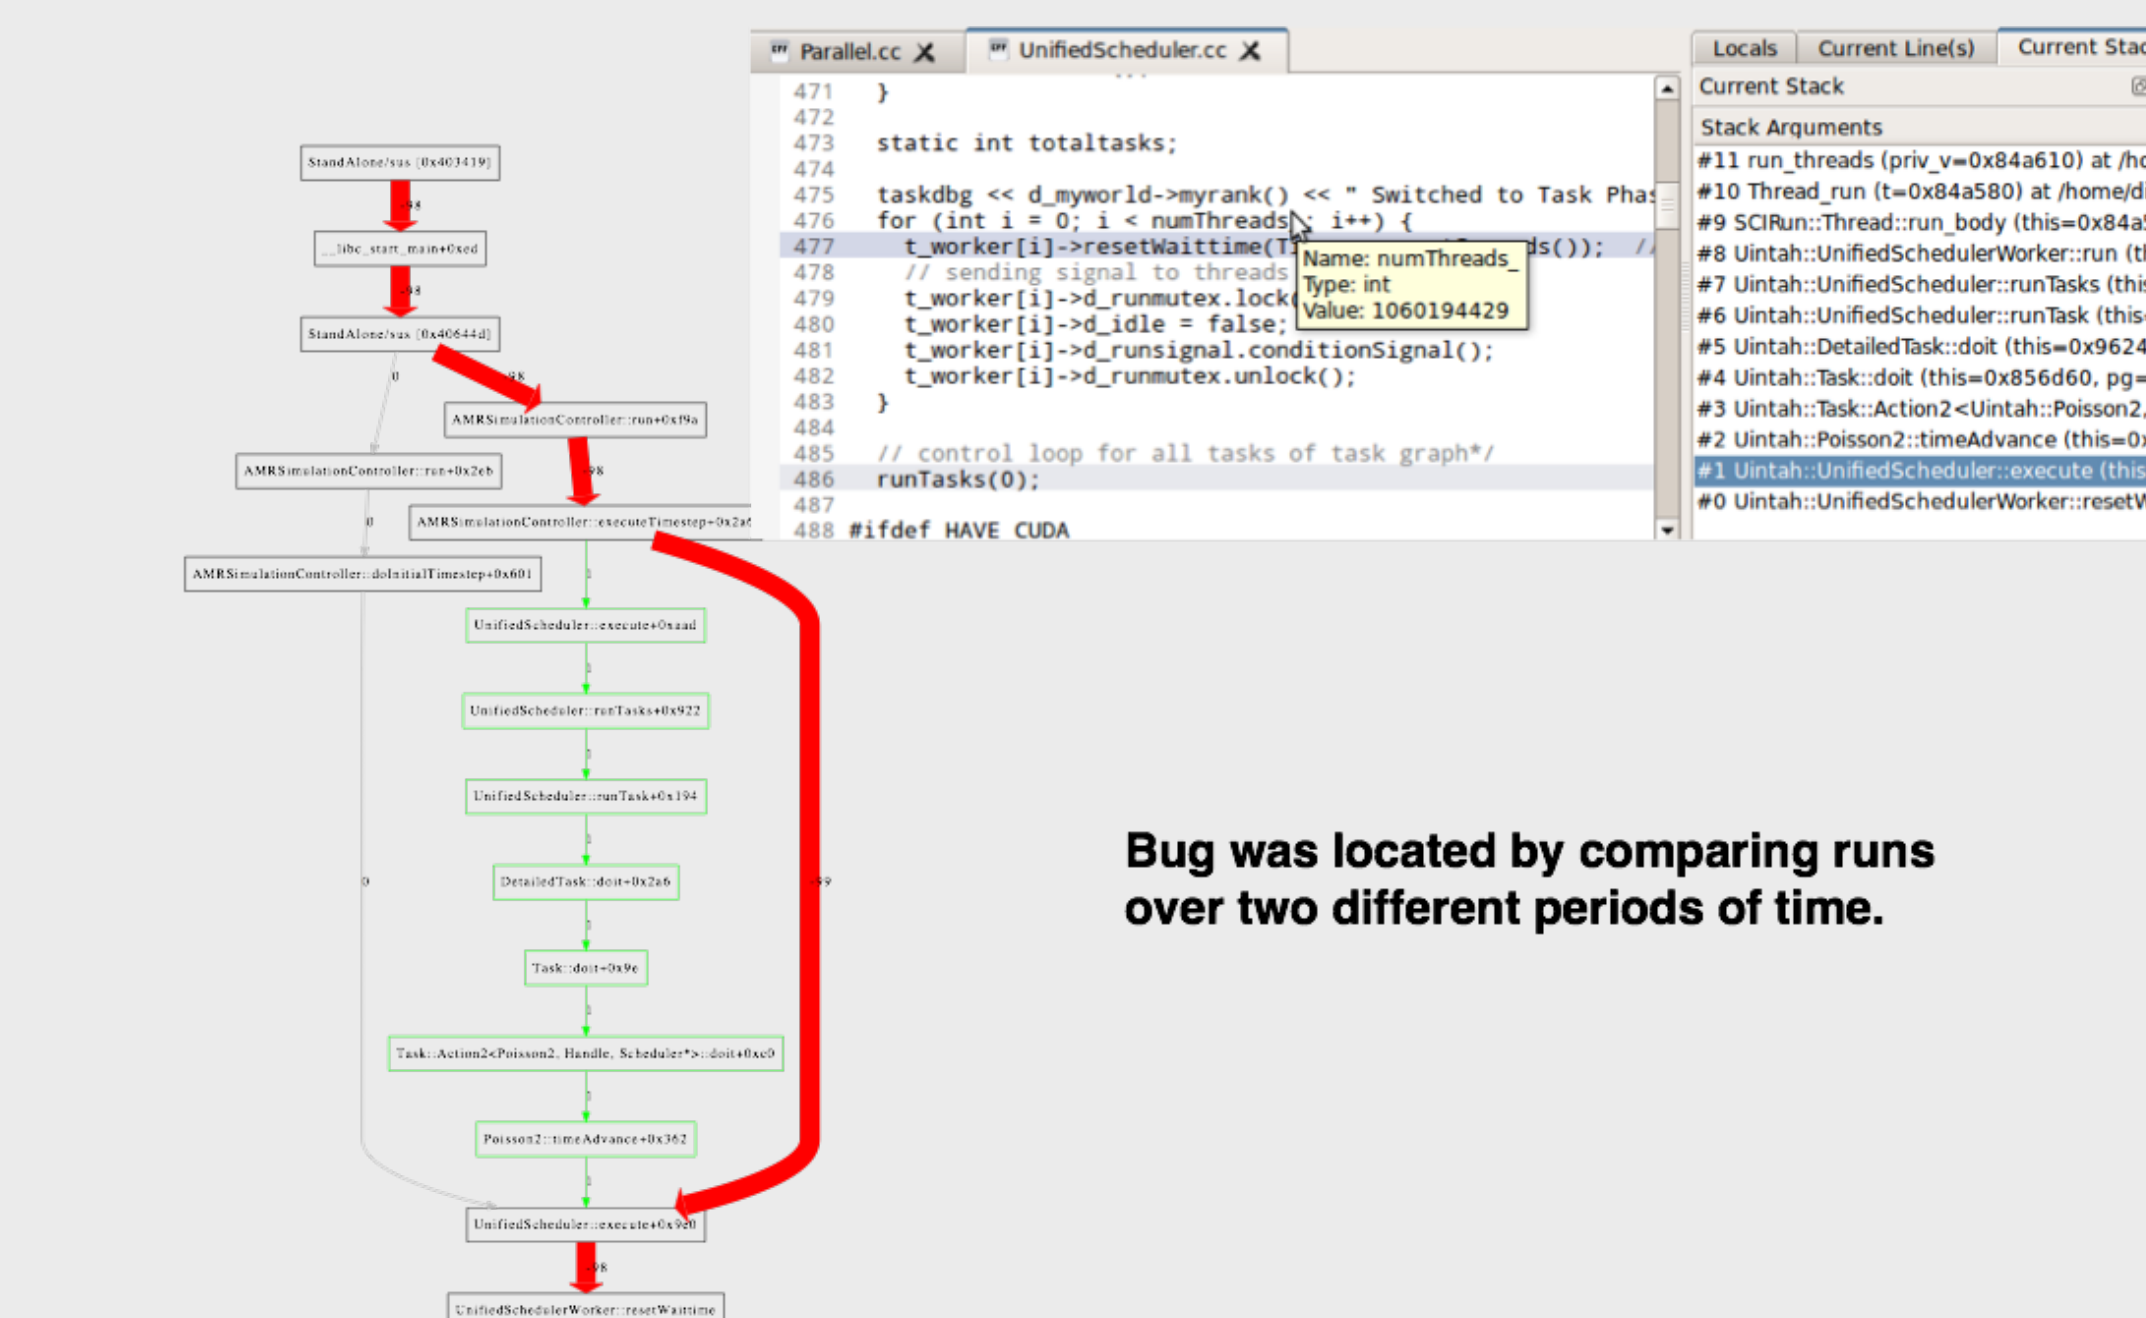
- Stack trace collection based
  - Often for MPI
- Progress tracking
  - Often based on loop progress order
- Example : STAT, AutomaDeD, Protometer (LLNL), Dynoptic (UW)
- Do not exploit behavioral differences

### Use of Coalesced Stack Trace Graphs

- Proven useful in large code base
- Summarizes stack nests
- Does not handle heterogeneity
- While overhead is low, it was not focused toward synchronizations across multiple concurrency models

Example use of Coalesced Stack Trace Graphs in detecting uninitialized variables in Uintah code base (University of Utah; see LCPC 2014)

### Nondeterminism due to Uninitialized Variable (Poisson2 Example)



## Ongoing Work

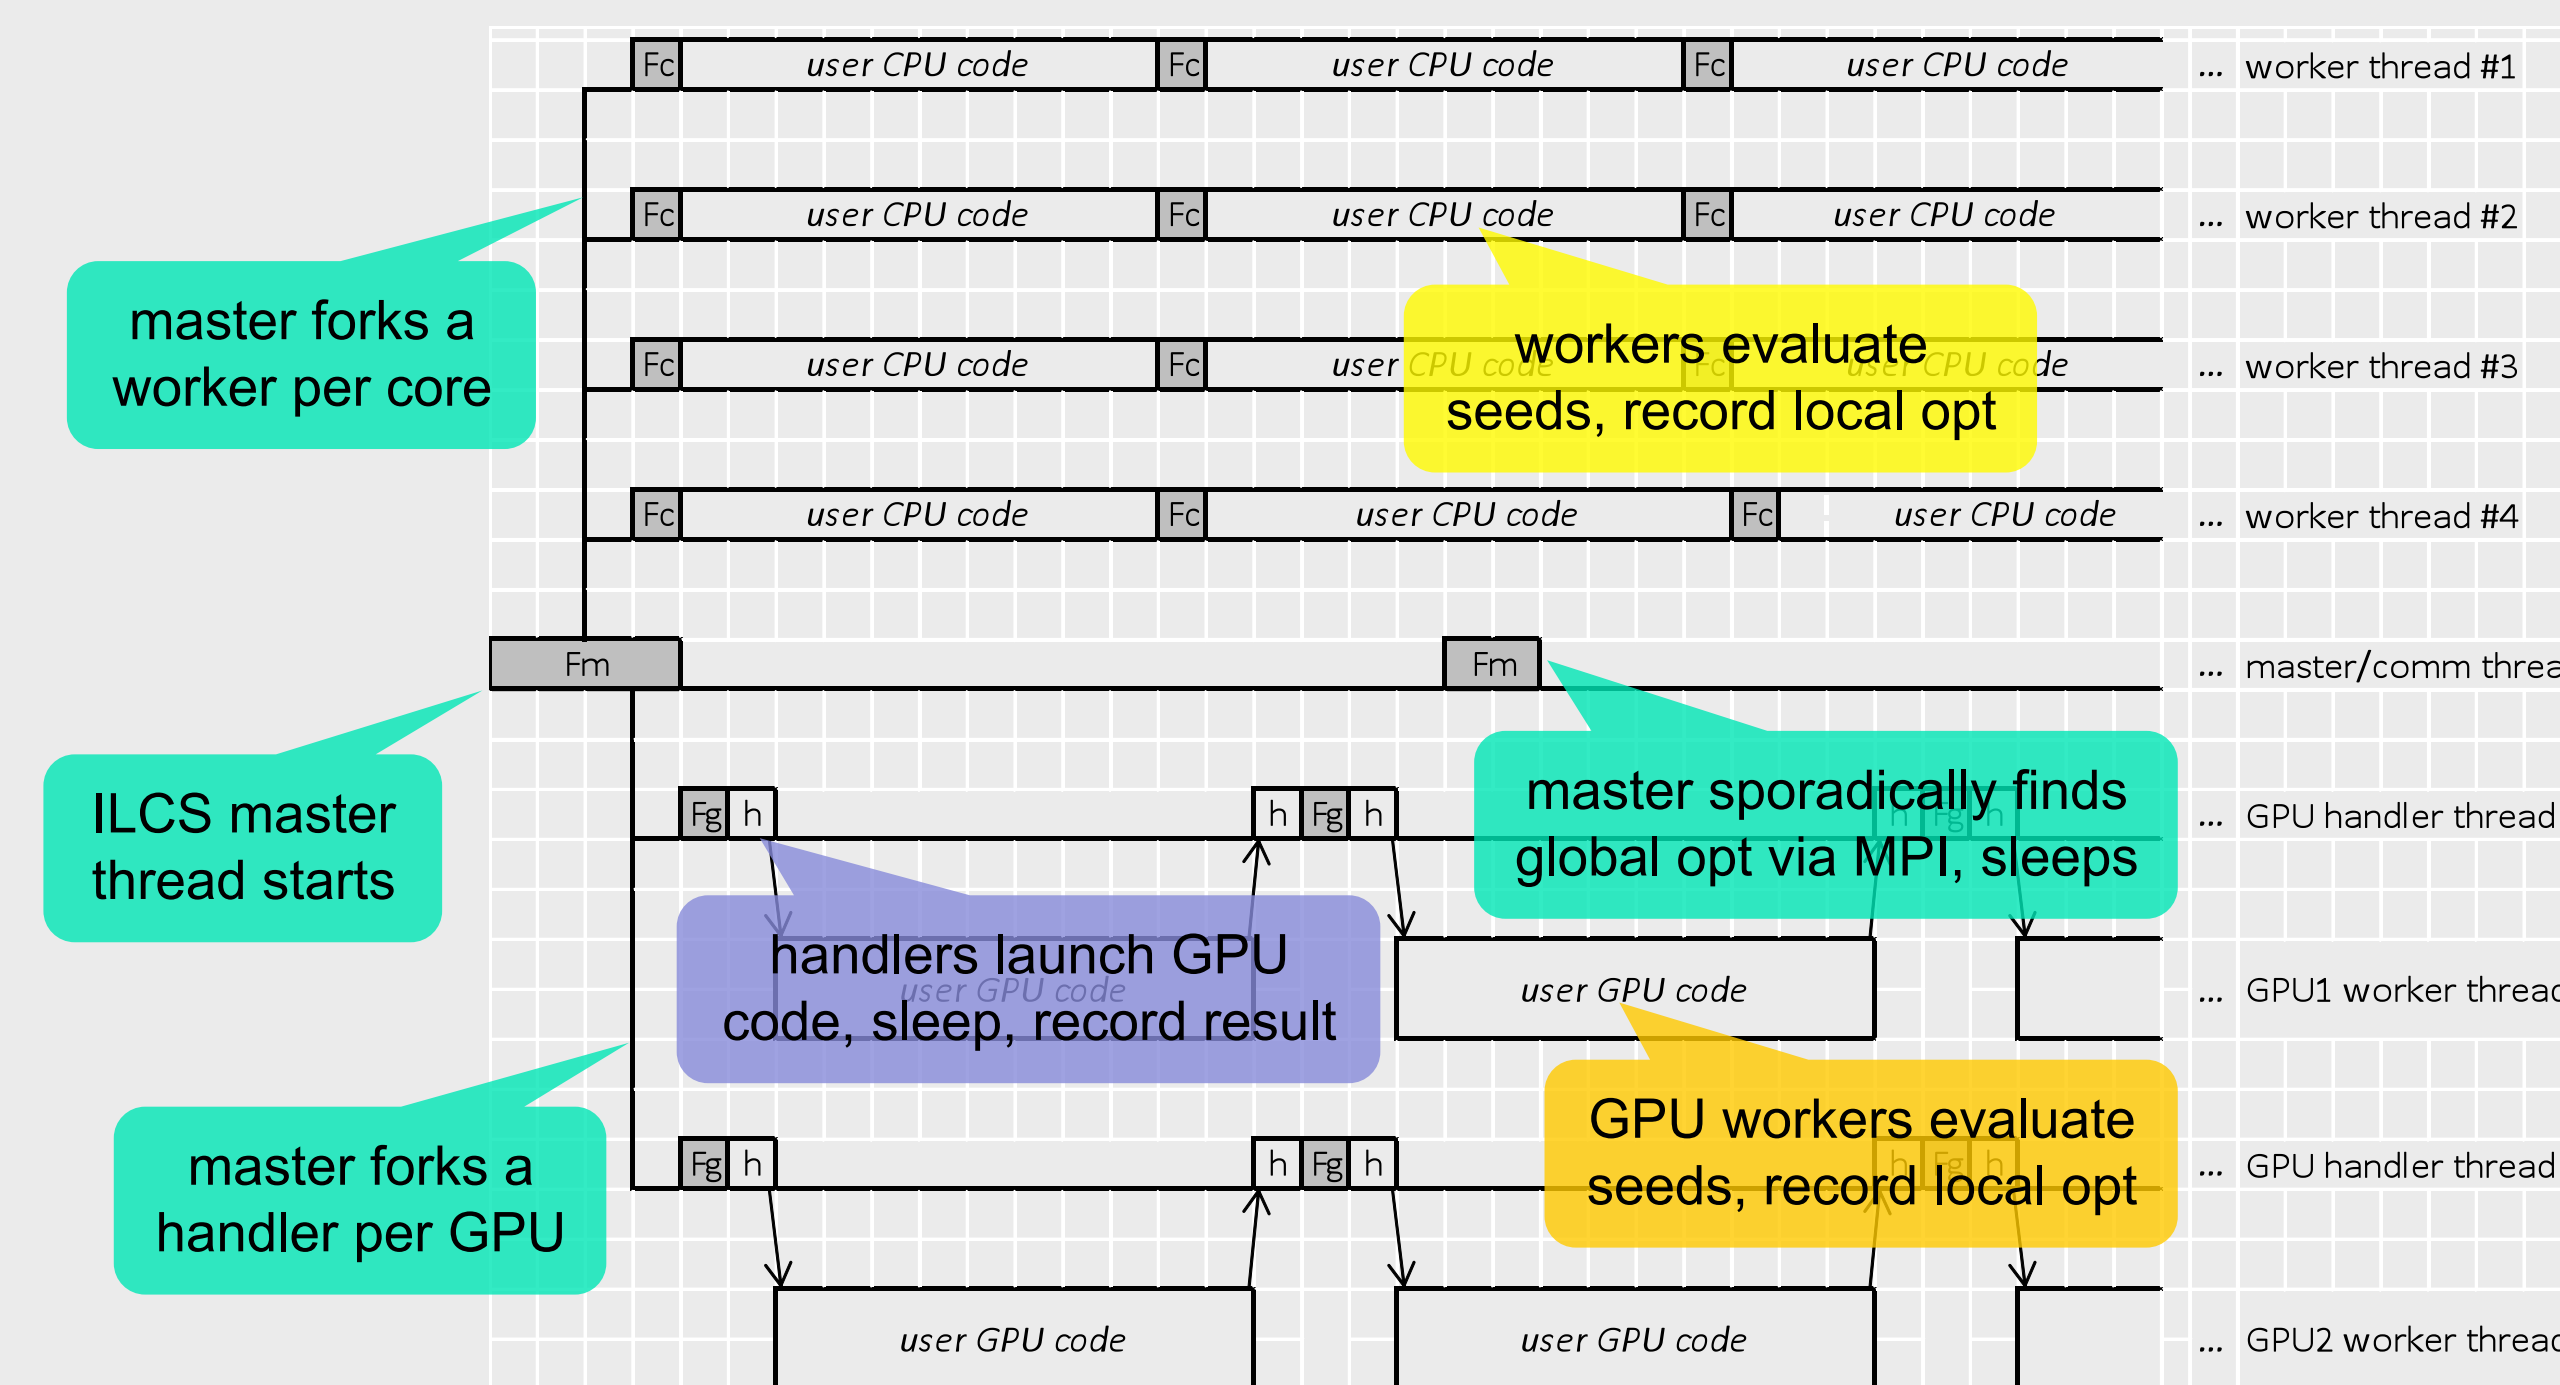
### Case Study: ILCS

- A heterogeneous concurrent program called the Iterative Local Champion Search (ILCS) has been chosen


### ILCS combines three flavors of concurrency

- MPI
- OpenMP
- GPU

### Execution Flow of ILCS



A Scalable Heterogeneous Parallelization Framework for Iterative Local Searches



### Largest Scale Tested

| system    | compute nodes | total CPUs | total GPUs | total CPU cores | total GPU cores |
|-----------|---------------|------------|------------|-----------------|-----------------|
| Keeneland | 128           | 256        | 384        | 2048            | 196,608         |
| Ranger    | 2048          | 8192       | 0          | 32768           | 0               |
| Stampede  | 1024          | 2048       | 0          | 16384           | 0               |


### Other Case Studies Planned

- A Rigorous Global Optimizer for Floating-Point Precision Estimation
- Parallelized versions of a GPU Data Race Checker
- Using parallel execution frameworks to parallelize verification


## Merit, Impact, Milestones, Students

### Work in progress

#### 1) CPU ↔ GPU control tracing methods



#### 2) CPU ↔ XeonPhi control tracing methods



#### 3) Binary Instrumentation Methods for Heterogeneous Trace Collection

### Intellectual Advances, Broader Impact

- Ways to track “happens before” at scale
- Ways to mine the tracked information for finding the root causes of bugs
- Enabling science at scale
  - Achieving Extreme Scale requires the use of powerful debugging tools
- An understanding of what information to collect, how collection scales, and how it facilitates debugging
- Effective utilization of HPC resources
- System heterogeneity is only bound to increase — impactful beyond HPC

### Anticipated Tools and Utilities

- Control tracing utilities for heterogeneous concurrency
- A graphical tool for querying collected control dependencies for debugging

### Student Training

- Saeed Taheri, PhD
- Sindhu Devale, MS
- Kurstie Lenear (BS, graduated)