Contech: High Performance Parallel Program Instrumentation

Brian P. Railing

Problem: How to efficiently record all of a parallel program's execution

Solution: Open source, compiler-based instrumentation - http://bprail.github.io/contech/

CONTECH

Parallel Programs

Language

C C++ Fortran

Runtime API

PThreads OpenMP MPI Cilk

ISA

X86 ARM

Contech Instrumentation Overhead

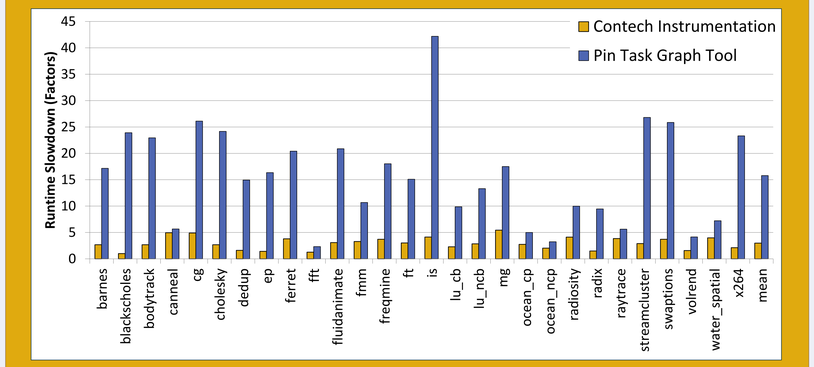

Parallel Components Map to Tasks

Task Graph Legend

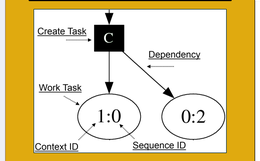

Graphs also record execution within each task

Cilk Example

```
fib(2);

int fib(int n) {
  if (n < 2)
    return n;
  int a = cilk_spawn
                fib(n-1);
  int b = fib(n-2);
  cilk_sync;
  return a + b;
}
```

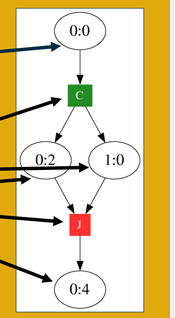

3x Average Slowdown

LLVM + Contech

Binary

Task Graph

Compiler Inlines High Performance Instrumentation

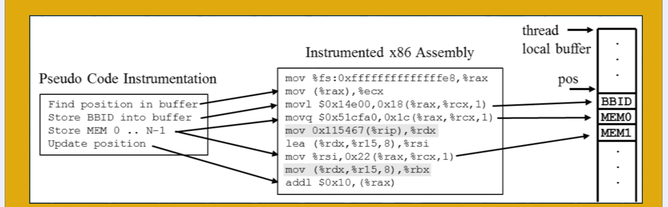

Example Analysis Loop

```
For each task in task graph
  For each basic block in task
    For each memory access in basic block
      Do analysis()
```

Parallel Program Analysis

Dynamic CFG

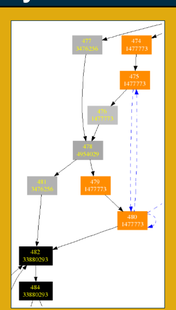

Identify Resource Constraints by Task

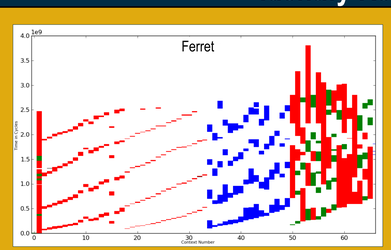

Constraint
Floating Point
Integer
Memory

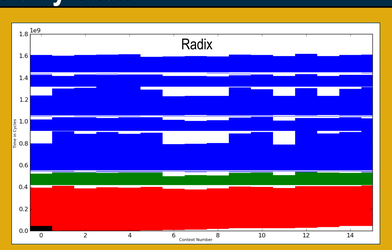

Programs show resource diversity across both threads and time.

Cache Simulation

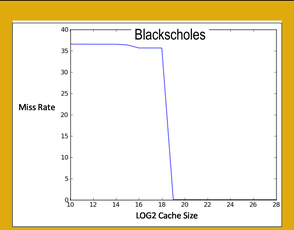

Data Race Detection

Does a happens before relation exist between every access to a given address?