Automatic Detection of Performance Anomalies in Task-Parallel Programs Work in progress on the Aftermath trace analysis tool

## Andi Drebes

Université Pierre et Marie Curie Laboratoire d'Informatique de Paris VI andi.drebes@lip6.fr

Joint work with: Antoniu Pop, Karine Heydemann Albert Cohen, Nathalie Drach

RACING'14, May 30th, 2014







## Context

### Hardware and software environment

- ► Multi-core / many-core systems
- How to exploit the hardware efficiently?
- ► Task-parallel languages based on fine-grained tasks

## Context

#### Hardware and software environment

- Multi-core / many-core systems
- How to exploit the hardware efficiently?
- Task-parallel languages based on fine-grained tasks

## Performance debugging

- Requires analysis of complex interactions at execution time: Application / Run-time / Machine
- ► Possible solution: Record dynamic events to a trace file
- ► Post-mortem (offline) analysis

## Context

### Hardware and software environment

- Multi-core / many-core systems
- How to exploit the hardware efficiently?
- Task-parallel languages based on fine-grained tasks

## Performance debugging

- Requires analysis of complex interactions at execution time: Application / Run-time / Machine
- ► Possible solution: Record dynamic events to a trace file
- ► Post-mortem (offline) analysis

## Aftermath

- ► Trace visualization and support for manual analysis
- ► Originally developed for OpenStream language & run-time
- ► Work in progress: Automate repetitive tasks

## 1. Aftermath

- 2. Insufficient parallelism and its causes
- 3. Performance anomalies during task execution
- 4. Work in progress: Status
- 5. Summary & Questions



# Aftermath











#### Task execution (dark blue)



#### Task execution (dark blue)



### Task creation (white)



## Searching for work (light blue)



#### Basic statistics for run-time states



#### Heatmap indicating task duration (white: fast, red: slow)



# NUMA heatmap indicating locality of memory accesses (blue: local, pink: remote)

Andi Drebes - Automatic Detection of Performance Anomalies in Task-Parallel Programs



#### Navigating through a trace

- Huge amounts of high-dimensional data
- Lots of features  $\rightarrow$  lots of possibilities for analysis
- ► Where to look? What to look for?
- Expertise & lots of time required



#### Navigating through a trace

- Huge amounts of high-dimensional data
- Lots of features  $\rightarrow$  lots of possibilities for analysis
- Where to look? What to look for?
- Expertise & lots of time required

## Guide the user through performance analysis

- ► Refine analysis in several steps
  - Start by analyzing parallelism
  - Analyze what happens inside tasks
- Automate repetitive tasks

# Detecting insufficient parallelism



## Ideal situation

All CPUs are in task execution state without any interruption

# Detecting insufficient parallelism



#### 

## Ideal situation

All CPUs are in task execution state without any interruption

## **Realistic scenario** Task creation, memory allocation, idle time, over-synchronization, ...

# Detecting insufficient parallelism



#### P<sub>0</sub> P<sub>1</sub> P<sub>2</sub> : : P<sub>n-1</sub> Tak escution Other stats

## Ideal situation

All CPUs are in task execution state without any interruption

## **Realistic scenario** Task creation, memory allocation, idle time, over-synchronization, ...

## Detect insufficient parallelism / high overhead automatically





d: Duration of the interval



- d: Duration of the interval
- $d_{e,i}$ : Time that processor *i* spends in task execution state



- d: Duration of the interval
- $d_{e,i}$ : Time that processor *i* spends in task execution state



- d: Duration of the interval
- $d_{e,i}$ : Time that processor *i* spends in task execution state



- d: Duration of the interval
- $d_{e,i}$ : Time that processor *i* spends in task execution state



- d: Duration of the interval
- $d_{e,i}$ : Time that processor *i* spends in task execution state
- $t_e$ : Threshold for task execution, e.g.  $t_e = 0.95$



- d: Duration of the interval
- $d_{e,i}$ : Time that processor *i* spends in task execution state
- $t_e$ : Threshold for task execution, e.g.  $t_e = 0.95$

Consider that there is sufficient parallelism if inequation holds:

$$\sum_{i=1}^{n} d_{e,i} > t_e \cdot n \cdot d$$

# Detecting the cause of insufficient parallelism

## Multiple stages during analysis

- ► If inequation does not hold, find out why
- Possible causes: task creation overhead, memory allocation, not enough tasks available for execution, ...
- Use thresholds for associated states:
  - $t_c$  (task creation),  $t_i$  (idle time)

# Detecting the cause of insufficient parallelism

## Multiple stages during analysis

- ► If inequation does not hold, find out why
- Possible causes: task creation overhead, memory allocation, not enough tasks available for execution, ...
- Use thresholds for associated states:
  - $t_c$  (task creation),  $t_i$  (idle time)



## Interval selection

- ► Multiple intervals: initialization, termination, etc.
- Repeat analysis for different intervals

Choose Interval






#### During task execution

 Performance anomaly possible even at 100% task execution (ineffective use of caches, remote memory accesses, branch misprediction)

# Detecting performance anomalies during task execution

### During task execution

 Performance anomaly possible even at 100% task execution (ineffective use of caches, remote memory accesses, branch misprediction)



#### Impact on the distribution of task duration

- Slowdown of all tasks
- Different groups / peaks

Andi Drebes - Automatic Detection of Performance Anomalies in Task-Parallel Programs

#### Hardware performance counters

- ► Implemented in hardware, no slowdown of the application
- ► Low tracing overhead if sampled at beginning / end of a task
- Dozens of hardware events can be monitored

#### Hardware performance counters

- ► Implemented in hardware, no slowdown of the application
- Low tracing overhead if sampled at beginning / end of a task
- Dozens of hardware events can be monitored

#### Automatic analysis of performance counters

- ► Which hardware events are relevant?
- Manual testing tedious & time consuming



#### Per-CPU performance counter

- Absolute value v(c, i, t), monotonically increasing
  - ► c: Counter (e.g. cache misses)
  - ► *i*: Processor identifier
  - ► t: Timestamp



#### Per-CPU performance counter

- Absolute value v(c, i, t), monotonically increasing
  - ► c: Counter (e.g. cache misses)
  - ► *i*: Processor identifier
  - ► t: Timestamp



#### Per-CPU performance counter

- Absolute value v(c, i, t), monotonically increasing
  - ► c: Counter (e.g. cache misses)
  - ► *i*: Processor identifier
  - ► t: Timestamp
- Sampled at the beginning and end of a task



Per-CPU performance counter

- Absolute value v(c, i, t), monotonically increasing
  - ► c: Counter (e.g. cache misses)
  - ► *i*: Processor identifier
  - ► t: Timestamp
- Sampled at the beginning and end of a task

#### Break down counter evolution to task instances

► Increase of c by task T:  $N_{c,T} = v(c, i, e) - v(c, i, s)$ 



#### Perform linear regression

- ► Assume linear model: d<sub>T<sub>j</sub></sub> = α · N<sub>c,T<sub>j</sub></sub> + β (α and β constant)
- ► Compare coefficient of determination with threshold



#### Perform linear regression

- ► Assume linear model: d<sub>T<sub>j</sub></sub> = α · N<sub>c,T<sub>j</sub></sub> + β (α and β constant)
- ► Compare coefficient of determination with threshold



#### Perform linear regression

- ► Assume linear model: d<sub>T<sub>j</sub></sub> = α · N<sub>c,T<sub>j</sub></sub> + β (α and β constant)
- ► Compare coefficient of determination with threshold

# Shortcuts & refinement

### Variation of task duration

- Determine coefficient of variation for task duration
- Only perform analysis if significant

# Shortcuts & refinement

### Variation of task duration

- Determine coefficient of variation for task duration
- Only perform analysis if significant

### Task types

- Different task types in an application
  - Auxiliary tasks: initialization, termination
  - ► Work tasks: matrix multiplication, decomposition, etc.
- ► Performance anomaly not necessarily present in *all* types

# Shortcuts & refinement

### Variation of task duration

- Determine coefficient of variation for task duration
- Only perform analysis if significant

#### Task types

- Different task types in an application
  - Auxiliary tasks: initialization, termination
  - ► Work tasks: matrix multiplication, decomposition, etc.
- ► Performance anomaly not necessarily present in *all* types

#### Topology of the machine

- Anomaly only present on subset of processors
- Example: Memory accesses local for one NUMA node, remote on another

Choose task type



















Andi Drebes - Automatic Detection of Performance Anomalies in Task-Parallel Programs





CPU 0 CPU 1		
CPU 2 CPU 3		
CPU S CPU 6		
CPU 7 CPU 8		
CPU 10 CPU 11		
CPU 12 CPU 13		
CPU 14 CPU 15		
CPU 17 CPU 18		
CPU 19 CPU 20		
CPU 21 CPU 22		
CPU 24 CPU 25		
CPU 26 CPU 27		
CPU 29		
CPU 31 CPU 32		
CPU 33 CPU 34		
CPU 36 CPU 37		
CPU 38 CPU 39		
CPU 40 CPU 41		
CPU 43 CPU 44		
CPU 45 CPU 46		
CPU 48		
CPU 50 CPU 51		
CPU 52 CPU 53		
CPU 55 CPU 55		
CPU 57 CPU 58		
CPU 59 CPU 60		
CPU 62 CPU 63		
		9004 1004 2004 2004 5004 5004









#### Analysis of parallelism

- ▶ Per-interval analysis of time spent on task execution
- ▶ Per-interval analysis of time spent in run-time states

### Analysis of parallelism

- ▶ Per-interval analysis of time spent on task execution
- ▶ Per-interval analysis of time spent in run-time states
- Support for thresholds
- Loop performing analysis on set of intervals

#### Analysis of parallelism

- ▶ Per-interval analysis of time spent on task execution
- Per-interval analysis of time spent in run-time states
- Support for thresholds
- Loop performing analysis on set of intervals

#### **Correlation of performance indicators**

- Support for performance counters
- Task duration histogram

#### Analysis of parallelism

- ▶ Per-interval analysis of time spent on task execution
- Per-interval analysis of time spent in run-time states
- Support for thresholds
- Loop performing analysis on set of intervals

#### **Correlation of performance indicators**

- Support for performance counters
- Task duration histogram
- Analysis of the variation of task durations
- Breaking down performance counter values to task instances
- Linear regression

## Summary

#### Aftermath

- ► Tool for trace-based analysis of task-parallel programs
- Currently provides only support for manual analysis
- Available at http://openstream.info/aftermath
# Summary

## Aftermath

- ► Tool for trace-based analysis of task-parallel programs
- Currently provides only support for manual analysis
- Available at http://openstream.info/aftermath

### Automatic analysis of parallelism based on thresholds

- Amount of time spent on task execution sufficiently high?
- If not, perform subsequent threshold-based analysis for states associated with overhead of the run-time system

# Summary

## Aftermath

- Tool for trace-based analysis of task-parallel programs
- Currently provides only support for manual analysis
- Available at http://openstream.info/aftermath

#### Automatic analysis of parallelism based on thresholds

- Amount of time spent on task execution sufficiently high?
- If not, perform subsequent threshold-based analysis for states associated with overhead of the run-time system

#### Automatic correlation of performance indicators

- Indicate which events are relevant
- Break down counter evolution to task instances
- Correlate with task duration