ParSim: a Tool for Workload Modeling and Reproduction of Parallel Applications

Andrea Rosà*, Walter Binder*, Lydia Y. Chen†, Marco Gribaudo‡ and Giuseppe Serazzi‡

*Università della Svizzera Italiana, Faculty of Informatics, Lugano, Switzerland; †IBM Research Lab Zurich, Rüschlikon, Switzerland;
‡Politecnico di Milano, Dipartimento di Elettronica, Informazione e Bioingegneria, Milano, Italy

Introduction

- ParSim is a **software tool** developed for:
  - **Modeling** the traffic of an application
  - **Reproducing** an arbitrary workload
  - **Analyzing** its behavior
- To evaluate performance indices:
  - **Response time**
  - **Throughput**
  - ...
- Working in **multicore** and **virtualized** environments:
  - **Parallel** applications
  - **Pseudo-distributed** services

Motivations

- Lack of **configurable** benchmarks for parallel environments:
  - **Specific** workload and **limited configurability**
  - **Difficult** to customize
- Need for **specific metrics**
  - Suitable for parallel environments

Workflow

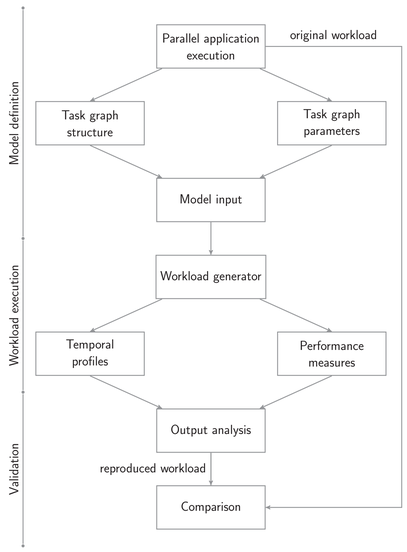

- **Model definition**:
  - Derive the original execution model
    - Task graph **structure**
  - Iterative fit workload **parameters**
- **Workload execution**:
  - Reproduce the **modeled** workload
    - As specified by the task graph
  - Profile the workload during the reproduction:
    - Collect real-time performance and system metrics
  - Obtain **performance measures** and graphical **temporal profiles** about:
    - The **execution**
    - The **processes**
    - The **phases**
- **Validation**
  - Compare original and reproduced workload
    - Similar behavior

Tool architecture

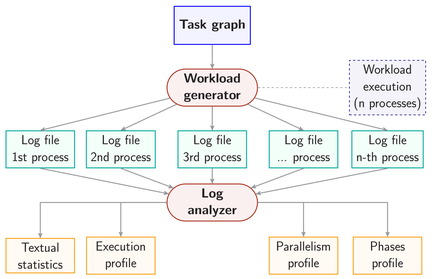

- Three main components:
  - The **task graph**
    - Model definition
  - The **workload generator**
    - Workload execution
  - The **log analyzer**
    - Workload analysis and validation

Task graph

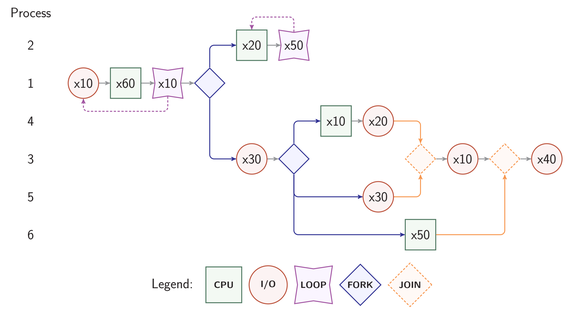

- Models the **workload** to reproduce
- Composed by **phases**:
  - The atomic elements of a task graph
  - Include the desired **number of operations**
  - Each phase can be of a single **type**:
    - CPU-bound; I/O-bound; LOOP; FORK; JOIN
- ParSim can reproduce complex task graphs:
  - Alternated CPU-I/O loads
  - Presence of synchronous and asynchronous processes
  - Partial synchronizations

Analysis

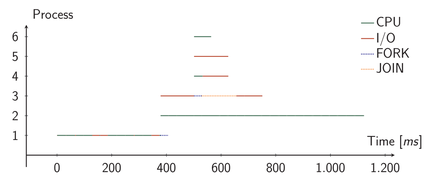

Execution profile

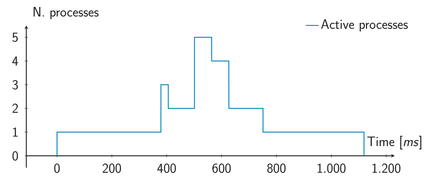

Parallelism profile

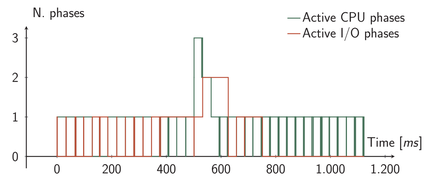

CPU-I/O phases profile

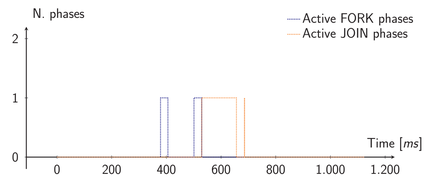

FORK-JOIN phases profile

- Three graphical **temporal profiles**:
  - **Execution profile**
    - Workload evolution over time
    - State of all processes at each time instant
    - Impact of concurrent processes on the application behavior
  - **Parallelism profile**
    - Number of active processes at each time instant
    - Real parallelism degree of the application over time
  - **Phases profile**
    - Number of active phases at each time instant
    - Alternation of **CPU-I/O** and **FORK-JOIN** phases
    - Coexistence of different phases over time
    - Impact of synchronizations between processes
- Detailed textual **performance measures**
  - Several metrics collected for each phase:
    - **Starting** and **ending** time
    - Utilized **CPUs**
    - **Average duration** of each operation
  - In-depth workload analysis

Case study and validation

1. Model definition
   - Monitoring the execution through log files
   - Task graph construction
   - Parameters fitting
2. Workload execution and analysis
3. Validation
   - Comparison between original and reproduced workload

- Of a real and publicly available parallel application
  - Embedded in the **Hadoop MapReduce** framework
  - Standard *pi* benchmark with 10 mappers and 1 reducer
  - Set up in a **pseudo-distributed** environment

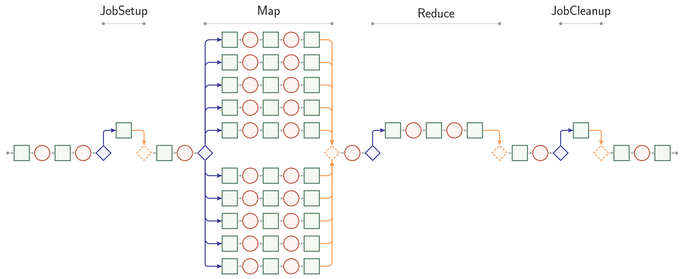

Constructed task graph

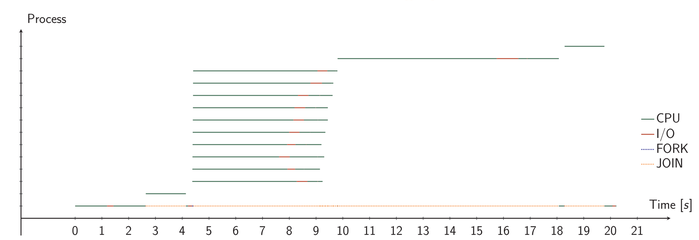

Original workload

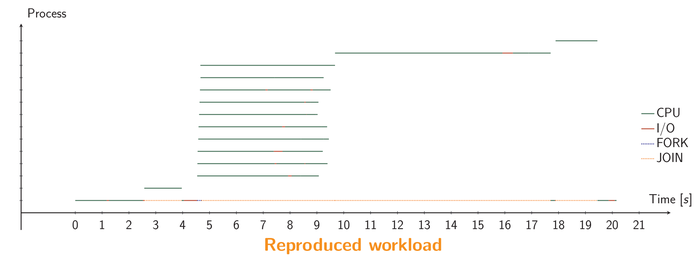

Reproduced workload

Contacts & Info

Andrea Rosà
andrea.rosa@usi.ch
+41 58 666 4325
http://www.inf.usi.ch/phd/rosaa
Università della Svizzera Italiana (USI)
Dynamic Analysis Group
Faculty of Informatics
Via Giuseppe Buffi 13
CH-6904 Lugano
Switzerland

To download ParSim and obtain more information, please visit our website at
http://www.inf.usi.ch/phd/rosaa/parsim

Acknowledgments

This work has been supported by the Swiss National Science Foundation (project 200021_141002)