

PARSEC vs. SPLASH-2: A Quantitative Comparison of Two Multithreaded Benchmark Suites

Christian Bienia (Princeton University),
Sanjeev Kumar (Intel),
Kai Li (Princeton University)

Outline



- Overview
 - What is PARSEC?
 - Why a new benchmark suite?
- Objectives of PARSEC
 - Technology Trend 1: Proliferation of CMPs
 - Technology Trend 2: Change of Technology Constraints
 - Technology Trend 3: Growth of World Data
- Characteristics Analysis
 - Methodology
 - Results
- Conclusions



What is PARSEC?

- Princeton Application Repository for Shared-Memory Computers
- Benchmark Suite for Chip-Multiprocessors
- Started as Joint-Venture between Intel and Princeton University
- Freely available at:

<http://parsec.cs.princeton.edu/>

- You can use it for your research

But what distinguishes PARSEC from SPLASH-2?

Requirements for a Benchmark Suite for CMPs



- **Multithreaded Applications**
Future programs must run on multiprocessors
- **Emerging Workloads**
Increasing CPU performance enables new applications
- **Diverse**
Multiprocessors are being used for more and more tasks
- **State-of-Art Techniques**
Algorithms and programming techniques evolve rapidly
- **Support Research**
Our goal is insight, not numbers

Assessment of Situation



	<i>Multithreaded</i>	<i>Emerging Workloads</i>	<i>Diverse</i>	<i>Not HPC-Focused</i>	<i>Research</i>
<i>SPEC CPU2006</i>	No	No	Yes	No	No
<i>SPEC OMP2001</i>	Yes	No	Yes	No	No
<i>SPLASH-2</i>	Yes	No	Yes	No	Yes
<i>ALPBench</i>	Yes	Yes	No	Yes	Yes
<i>BioBench</i>	No	No	No	No	Yes
<i>BioParallel</i>	Yes	No	No	No	Yes
<i>MediaBench II</i>	No	No	No	Yes	Yes
<i>MineBench 2.0</i>	Yes	No	No	Yes	Yes
<i>PhysicsBench</i>	Yes	Yes	No	Yes	Yes

This is why we created PARSEC

Workloads



<i>Program</i>	<i>Application Domain</i>	<i>Parallelization</i>
<i>Blackscholes</i>	Financial Analysis	Data-parallel
<i>Bodytrack</i>	Computer Vision	Data-parallel
<i>Canneal</i>	Engineering	Unstructured
<i>Dedup</i>	Enterprise Storage	Pipeline
<i>Facesim</i>	Animation	Data-parallel
<i>Ferret</i>	Similarity Search	Pipeline
<i>Fluidanimate</i>	Animation	Data-parallel
<i>Freqmine</i>	Data Mining	Data-parallel
<i>Streamcluster</i>	Data Mining	Data-parallel
<i>Swaptions</i>	Financial Analysis	Data-parallel
<i>Vips</i>	Media Processing	Data-parallel
<i>X264</i>	Media Processing	Pipeline

PARSEC is substantially different from SPLASH-2

Outline



- Overview
 - What is PARSEC?
 - Why a new benchmark suite?
- Objectives of PARSEC
 - Technology Trend 1: Proliferation of CMPs
 - Technology Trend 2: Change of Technology Constraints
 - Technology Trend 3: Growth of World Data
- Characteristics Analysis
 - Methodology
 - Results
- Conclusions

Objectives of PARSEC



- PARSEC was designed to capture recent technology trends:
 - Proliferation of CMPs
 - Multiprocessors are used in more and more areas
 - Change of Technology Constraints
 - Different software optimizations required for CMPs
 - Growth of World Data
 - Huge increase of stored data which must be processed
- These trends are changing programs

Impact of Technology Trends



- Proliferation of CMPs:
 - New application areas (e.g. video games)
 - New parallelization models (e.g. pipelining)
- Change of Technology Constraints:
 - Constrained off-chip bandwidth
 - Shared caches
- Growth of World Data:
 - Huge increase of input data
 - Higher importance of linear algorithms

We show that these trends affect program characteristics

Outline



- Overview
 - What is PARSEC?
 - Why a new benchmark suite?
- Objectives of PARSEC
 - Technology Trend 1: Proliferation of CMPs
 - Technology Trend 2: Change of Technology Constraints
 - Technology Trend 3: Growth of World Data
- Characteristics Analysis
 - Methodology
 - Results
- Conclusions

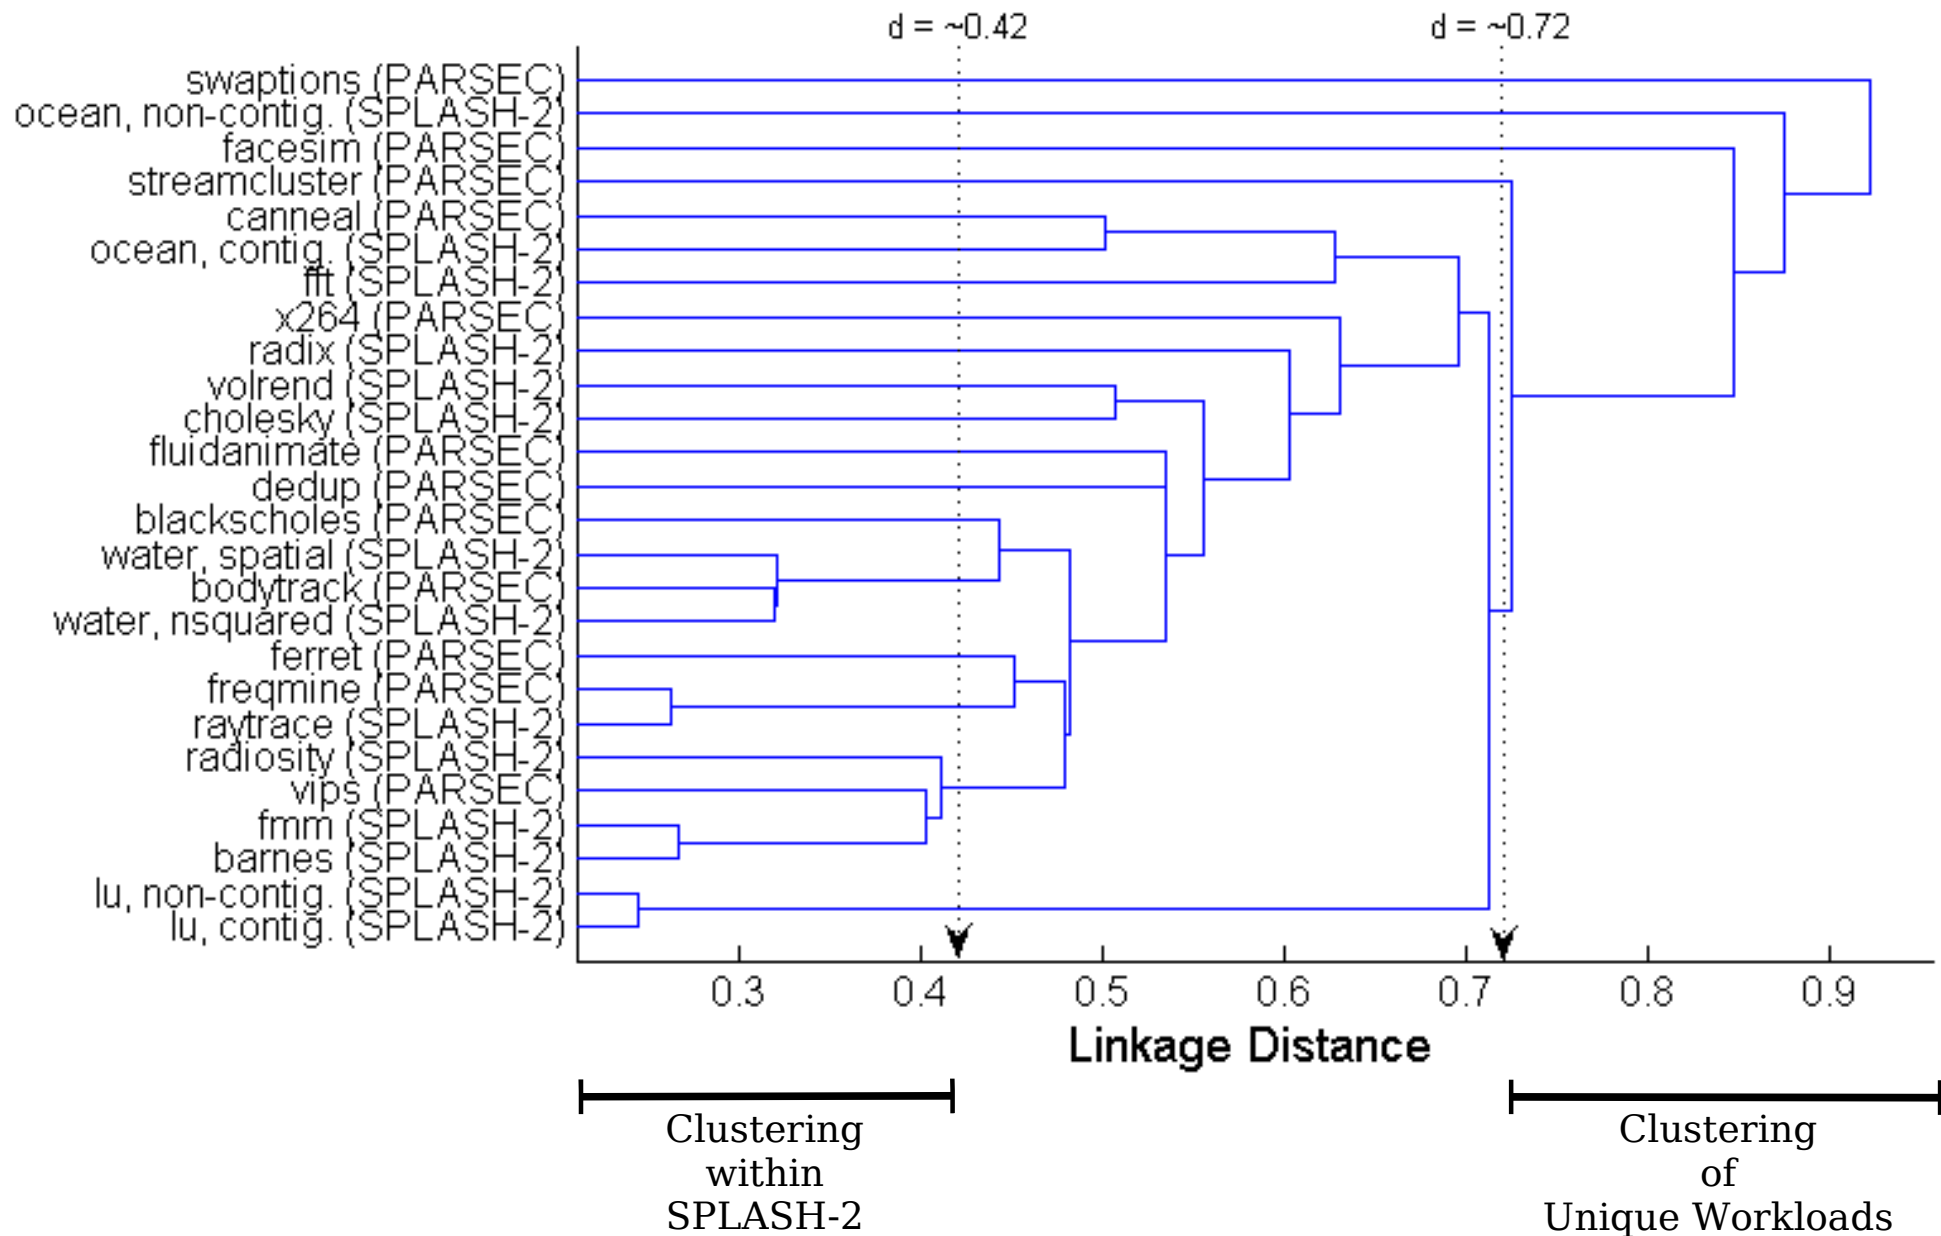
Methodology



- Simulate abstract cache hierarchy with CMP\$im
- Preprocess chosen characteristics with Principal Component Analysis (PCA) to eliminate correlation
- Compute similarity with hierarchical clustering
- Visualize results with dendrograms and scatter plots
- 44 characteristics chosen:
 - Instruction mix (4 characteristics)
 - Working set (8 characteristics)
 - Sharing (32 characteristics)




Redundancy & Similarity






Redundancy & Similarity






Redundancy & Similarity



PARSEC is more diverse than SPLASH-2

Systematic Differences



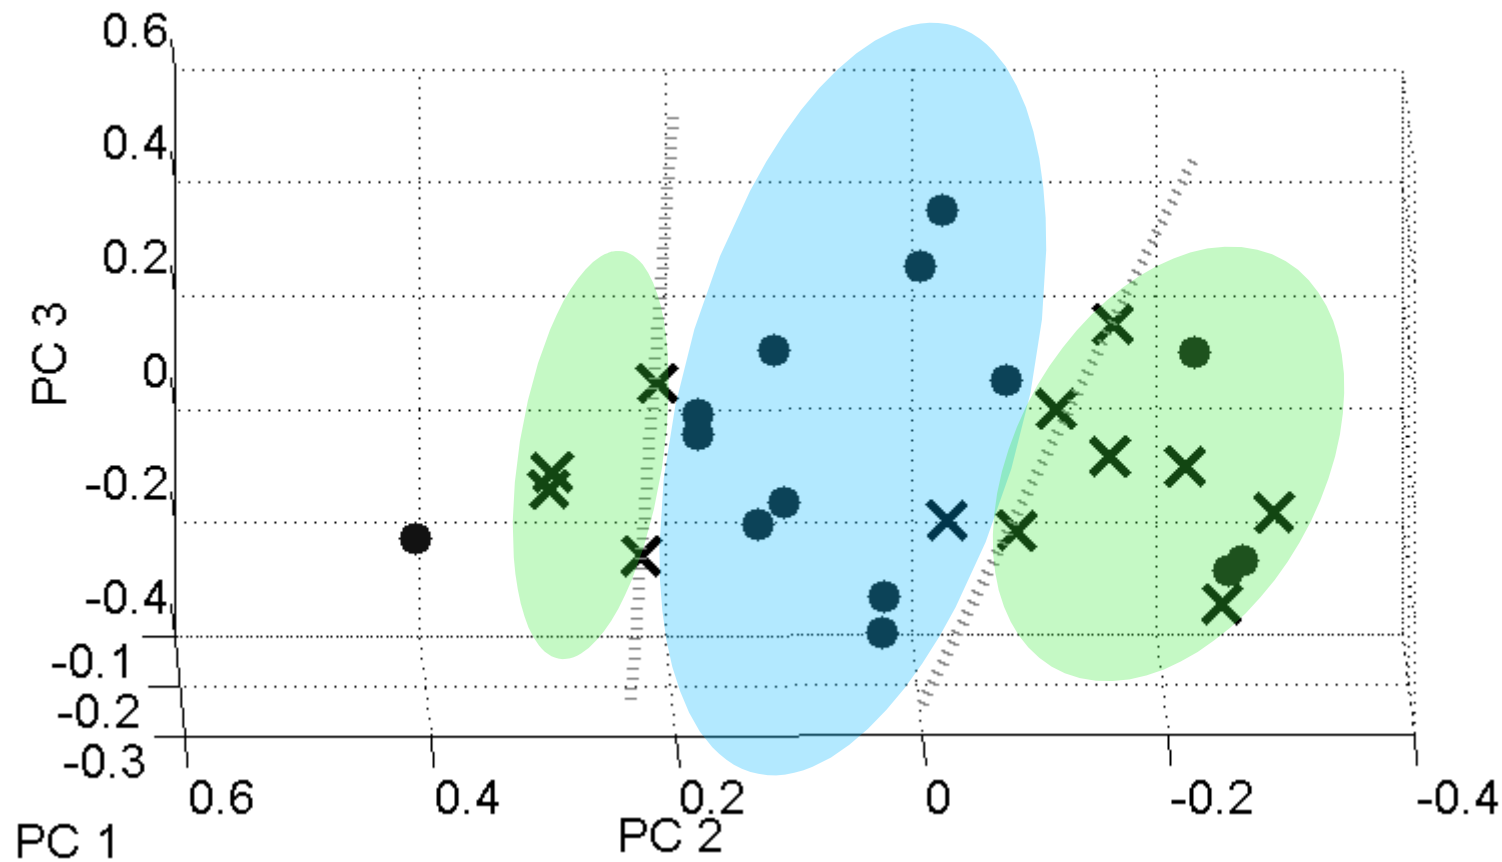
× PARSEC

● SPLASH-2

Benchmark suites
cluster in different
areas, little overlap

PARSEC and SPLASH-2 have little in common

Instruction Mix Differences




× PARSEC

● SPLASH-2

PARSEC workloads use cores differently

Working Set Differences

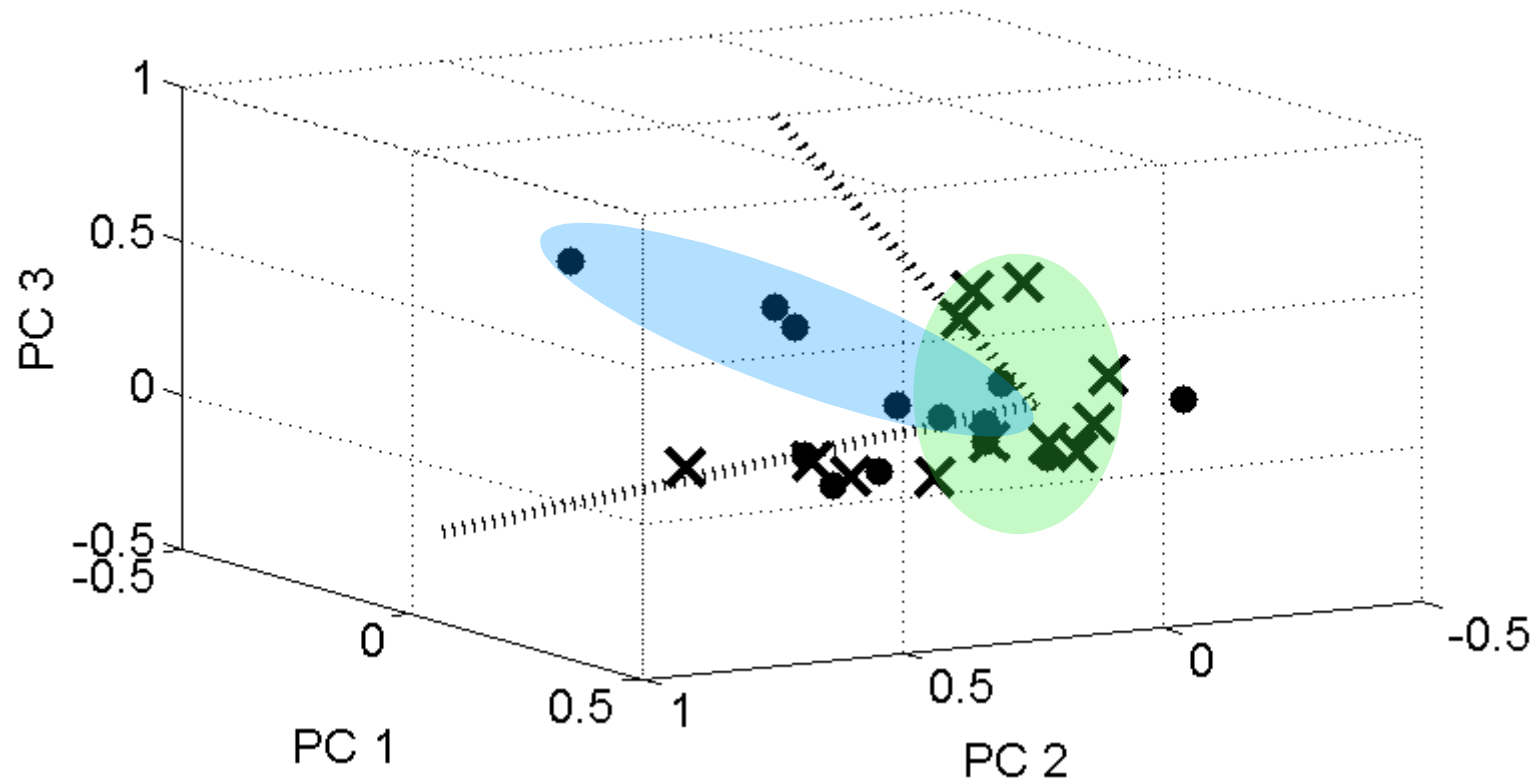


× PARSEC

● SPLASH-2

Some PARSEC workloads use memory differently

Sharing Behavior Differences



× PARSEC

● SPLASH-2

PARSEC workloads communicate differently

Outline



- Overview
 - What is PARSEC?
 - Why a new benchmark suite?
- Objectives of PARSEC
 - Technology Trend 1: Proliferation of CMPs
 - Technology Trend 2: Change of Technology Constraints
 - Technology Trend 3: Growth of World Data
- Redundancy & Similarity
 - Methodology
 - Results
- Conclusions

Conclusions



- PARSEC and SPLASH-2 are substantially different benchmark suites
- PARSEC is more diverse
- No single reason for differences

You should expect different results

Thank you!

Questions?