Correlating Sub-Phenomena in Performance Data in the Frequency Domain

Tom Vierjahn[1,4], Marc-André Hermanns[3,4], Bernd Mohr[3,4], Matthias S. Müller[2,4], Torsten W. Kuhlen[1,4], Bernd Hentschel[1,4]
[1] Visual Computing Institute, RWTH Aachen University, Germany
[2] Chair for High-Performance Computing, RWTH Aachen University, Germany
[3] Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH, Germany
[4] JARA-HPC, Aachen, Germany

1 Introduction

Finding and understanding correlated performance behaviour in high-performance computing applications is
- a time-consuming task, but
- key in understanding and optimization.

Thus, we propose to use
- automatic correlation analysis in the frequency domain
- allowing for filtering-out known sub-phenomena
- in order to detect new, unknown phenomena.

2 Performance Data

- Performance profiles store data (e.g., execution time) per system resource (e.g., cores, threads, …).
- These can be arranged in a Cartesian space, the system topology (e.g., cores x threads)
- Performance data thus constitutes a space domain signal
  $v(\mathbf{x})$ $\mathbf{x}$: location in the system topology
- for which we can compute a spectrum
  $V(\mathbf{k})$ $\mathbf{k}$: frequency

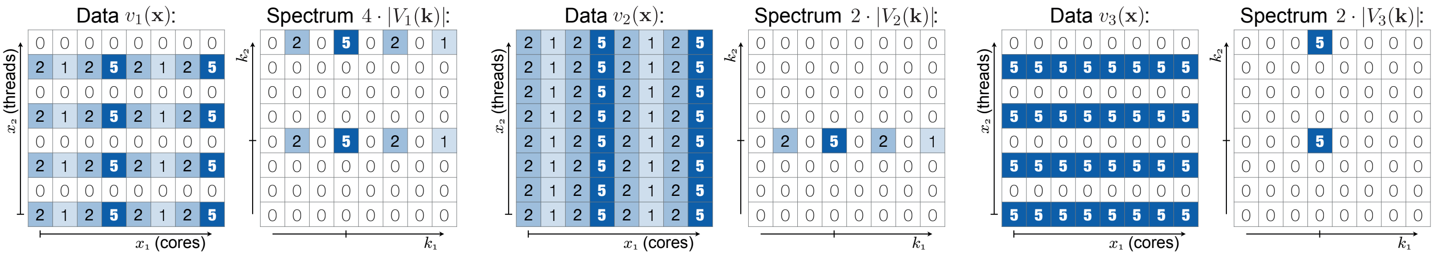

3 Automatic Correlation Analysis

Pearson correlation
- is sensitive to dominant sub-phenomena,
- weights all data equally,
- obscures new patterns.

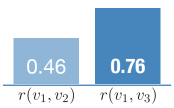

Filtered correlation
- has selective sensitivity,
- filters-out known sub-phenomena,
- provides new insights.

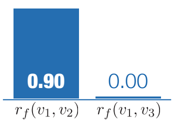

Computing filtered correlation (cosine-weighting, cross-correlation):

$$W_i^2(\mathbf{k}) = \frac{k_i^2}{\mathbf{k}^\top \mathbf{k}} \quad , \quad W_i^2(\mathbf{0}) = 0$$

$$g_{y,z} = \mathcal{F}^{-1}\left[\sum_i f_i \cdot W_i^2(\mathbf{k}) \cdot V_y^*(\mathbf{k}) \cdot V_z(\mathbf{k})\right](\mathbf{0})$$

$$r_f = \frac{g_{a,b}}{\sqrt{g_{a,a} \cdot g_{b,b}}}$$

4 Interactive Visualization, Results

Interactive analysis tool inspired by Cube and ParaProf:

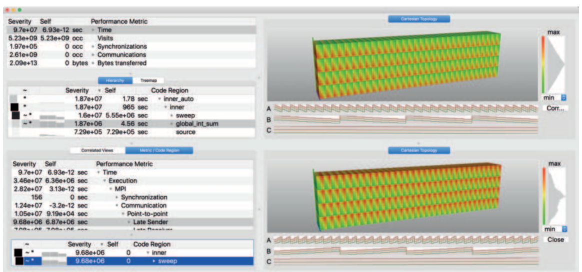

Efficient analysis with little memory overhead:

Code	Threads	Size	Over-head	Views ≠ 0	Correl. Analysis
NEKBONE	1,835,008	8.5 GB	0.6 MB	624	52.3 s
psOpen	65,384	1.0 GB	2.0 MB	2,031	1.8 s
Sweep3D	65,384	0.4 GB	53.1 kB	850	0.6 s

Acknowledgements

This work has been partially funded by the German Federal Ministry of Research and Education (BMBF) under grant number 01IH13001D (Score-E), and by the Excellence Initiative of the German federal and state governments.

['pɪvət]
pvt performance visualization toolkit

Visual Computing Institute
Virtual Reality & Immersive Visualization
Prof. Dr. Torsten W. Kuhlen

JARA|HPC

RWTH AACHEN UNIVERSITY