Geometry-Aware Visualization of Performance Data

Tom Vierjahn, Bernd Hentschel, Torsten W. Kuhlen
Visual Computing Institute, RWTH Aachen University, Germany
JARA – High-Performance Computing, Germany

1 Introduction

Visualizing performance data [1] in intuitive domains [2] helps analysts optimize massively parallel applications.
Thus, we propose a tool featuring linked views that

- guides analysts towards *important* parts of performance data,
- visualizes performance data in its spatial context:

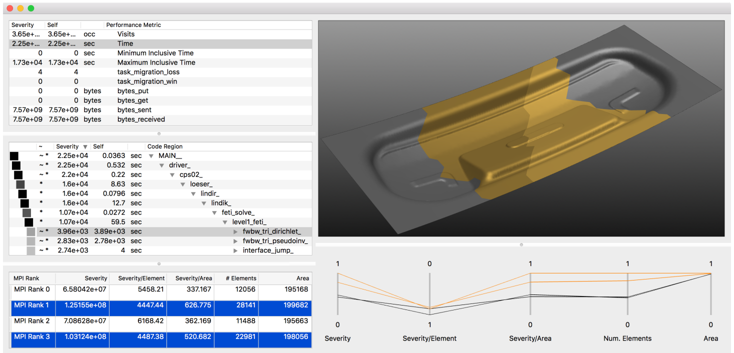

2 Performance Profiles

Profiles summarize performance data according to

- *performance metrics* $m \in \mathcal{M}$
- for the *call paths* $c \in \mathcal{C}$
- per *system resource* $s \in \mathcal{S}$.

Let $v_{m,c} : \mathcal{S} \mapsto \mathbb{R}$ denote a *severity view*, then

- $v_{m,c}(s')$ yields the *severity* of, e.g., execution time spent in function c, for an MPI rank $s' \in \mathcal{S}_{MPI} \subseteq \mathcal{S}$.

With a mapping to the simulated geometry

- $v_{m,c}(s')$ yields the severity for the individual geometry parts.

3 Detecting Variation

For severity views

- with little variation a single number represents performance,
- with large variation a detailed analysis is required.

Large-variation views are detected via the variation coefficient

$$q_{m,c} = \sigma_{m,c} \cdot \mu_{m,c}^{-1}$$

$\mu_{m,c}$: mean severity, $\sigma_{m,c}$: standard dev. in $v_{m,c}$

4 Visualization

A severity view gets selected in two tree-view widgets:

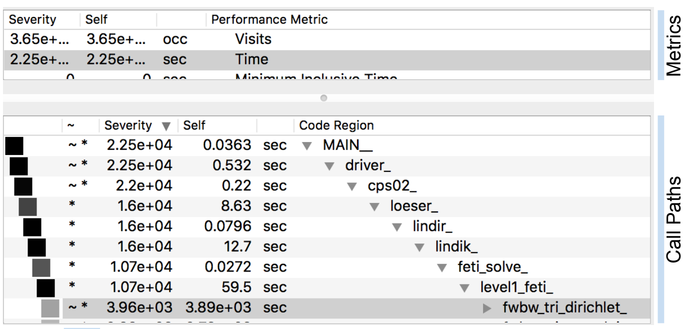

Large-variation indicator: ~ in respective view, * in descendants

Performance data is visualized in its spatial context:

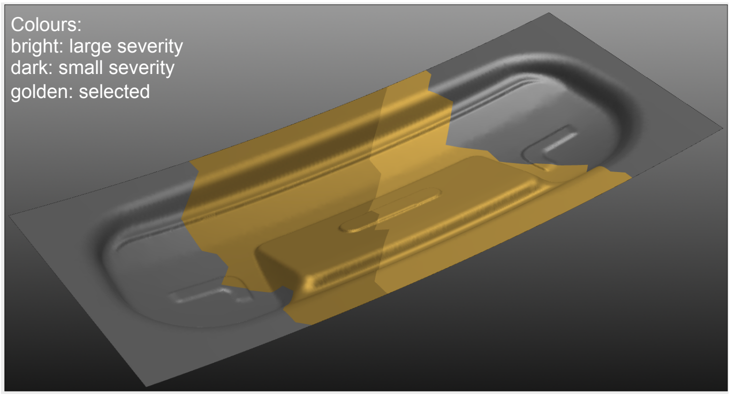

The data is related to the properties of the simulated geometry:

MPI Rank	Severity	Severity/Element	Severity/Area	# Elements	Area
MPI Rank 0	6.58042e+07	5458.21	337.167	12056	195168
MPI Rank 1	1.25155e+08	4447.44	626.775	28141	199682
MPI Rank 2	7.08628e+07	6168.42	362.169	11488	195663
MPI Rank 3	1.03124e+08	4487.38	520.682	22981	198056

A parallel-coordinates plot summarizes the above data:

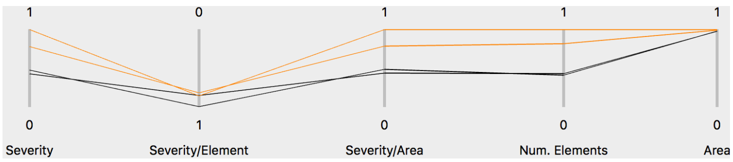

5 Results

Preliminarily evaluation (4 thin nodes of SuperMUC, Phase 1):

- search-space reduction by 29 %;
- data forms two almost separate classes;
- MPI ranks 1 and 3 required most CPU-time;
- MPI ranks 1 and 3 are computing high-detail geometry.

6 Conclusion

- Our tool greatly reduces the search space,
- quickly guides analysts towards *important* severity views,
- relates performance phenomena to the simulation domain,
- thus helps simulation experts understand performance data.

Acknowledgements

Partially funded by the German Federal Ministry of Research and Education (BMBF) and by the Excellence Initiative of the German federal and state governments through the Jülich Aachen Research Alliance – High-Performance Computing.

References

[1] Isaacs K.E., Giménez A., Jusufi I., Gamblin T., Bhatele A., Schulz M., Hamann B., Bremer P.-T.: State of the Art of Performance Visualization. In EuroVis - STARs, 2014.
[2] Schulz M., Levine J.A., Bremer P.-T., Gamblin T., Pascucci V.: Interpreting Performance Data across Intuitive Domains. In Proc. 40th Int. Conf. Parallel Process., 2011.

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

RWTH AACHEN UNIVERSITY