



HAL
open science

Towards Smart Visualization for HPC Simulations

Lokman Rahmani, Gabriel Antoniu, Luc Bougé

► **To cite this version:**

Lokman Rahmani, Gabriel Antoniu, Luc Bougé. Towards Smart Visualization for HPC Simulations. 28th IEEE International Parallel & Distributed Processing Symposium - PhD Forum, May 2014, PHOENIX, United States. , 2014. hal-01290266

HAL Id: hal-01290266

<https://inria.hal.science/hal-01290266>

Submitted on 17 Mar 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Context

- HPC platforms featuring **millions** of cores
- HPC scientific simulation on 100.000+ cores
- **Petabytes** of data to be post-processed and visualized
- Increasing **gap** between computational power and I/O capacity



Figure 1: Blue Waters NCSA supercomputer[1]




Figure 2: CM1 Tornado simulation


Acknowledgements

This work is done within the framework of the Joint Laboratory for Petascale Computing, an international collaboration between INRIA, the University of Illinois at Urbana-Champaign (IL. USA), and Argonne National Laboratory (IL. USA).

References


- [1] <http://www.ncsa.illinois.edu/BlueWaters/>.
 [2] Damaris/viz: <http://damaris.gforge.inria.fr/>.
 [3] VisIt, <https://wci.llnl.gov/codes/visit/>.
 [4] Aladdin grid'5000: <http://www.grid5000.fr>.

Off-line Visualization



- I/O performances issues in the simulation
- I/O performances issues in the visualization tool
- Doesn't scale

In Situ Visualization



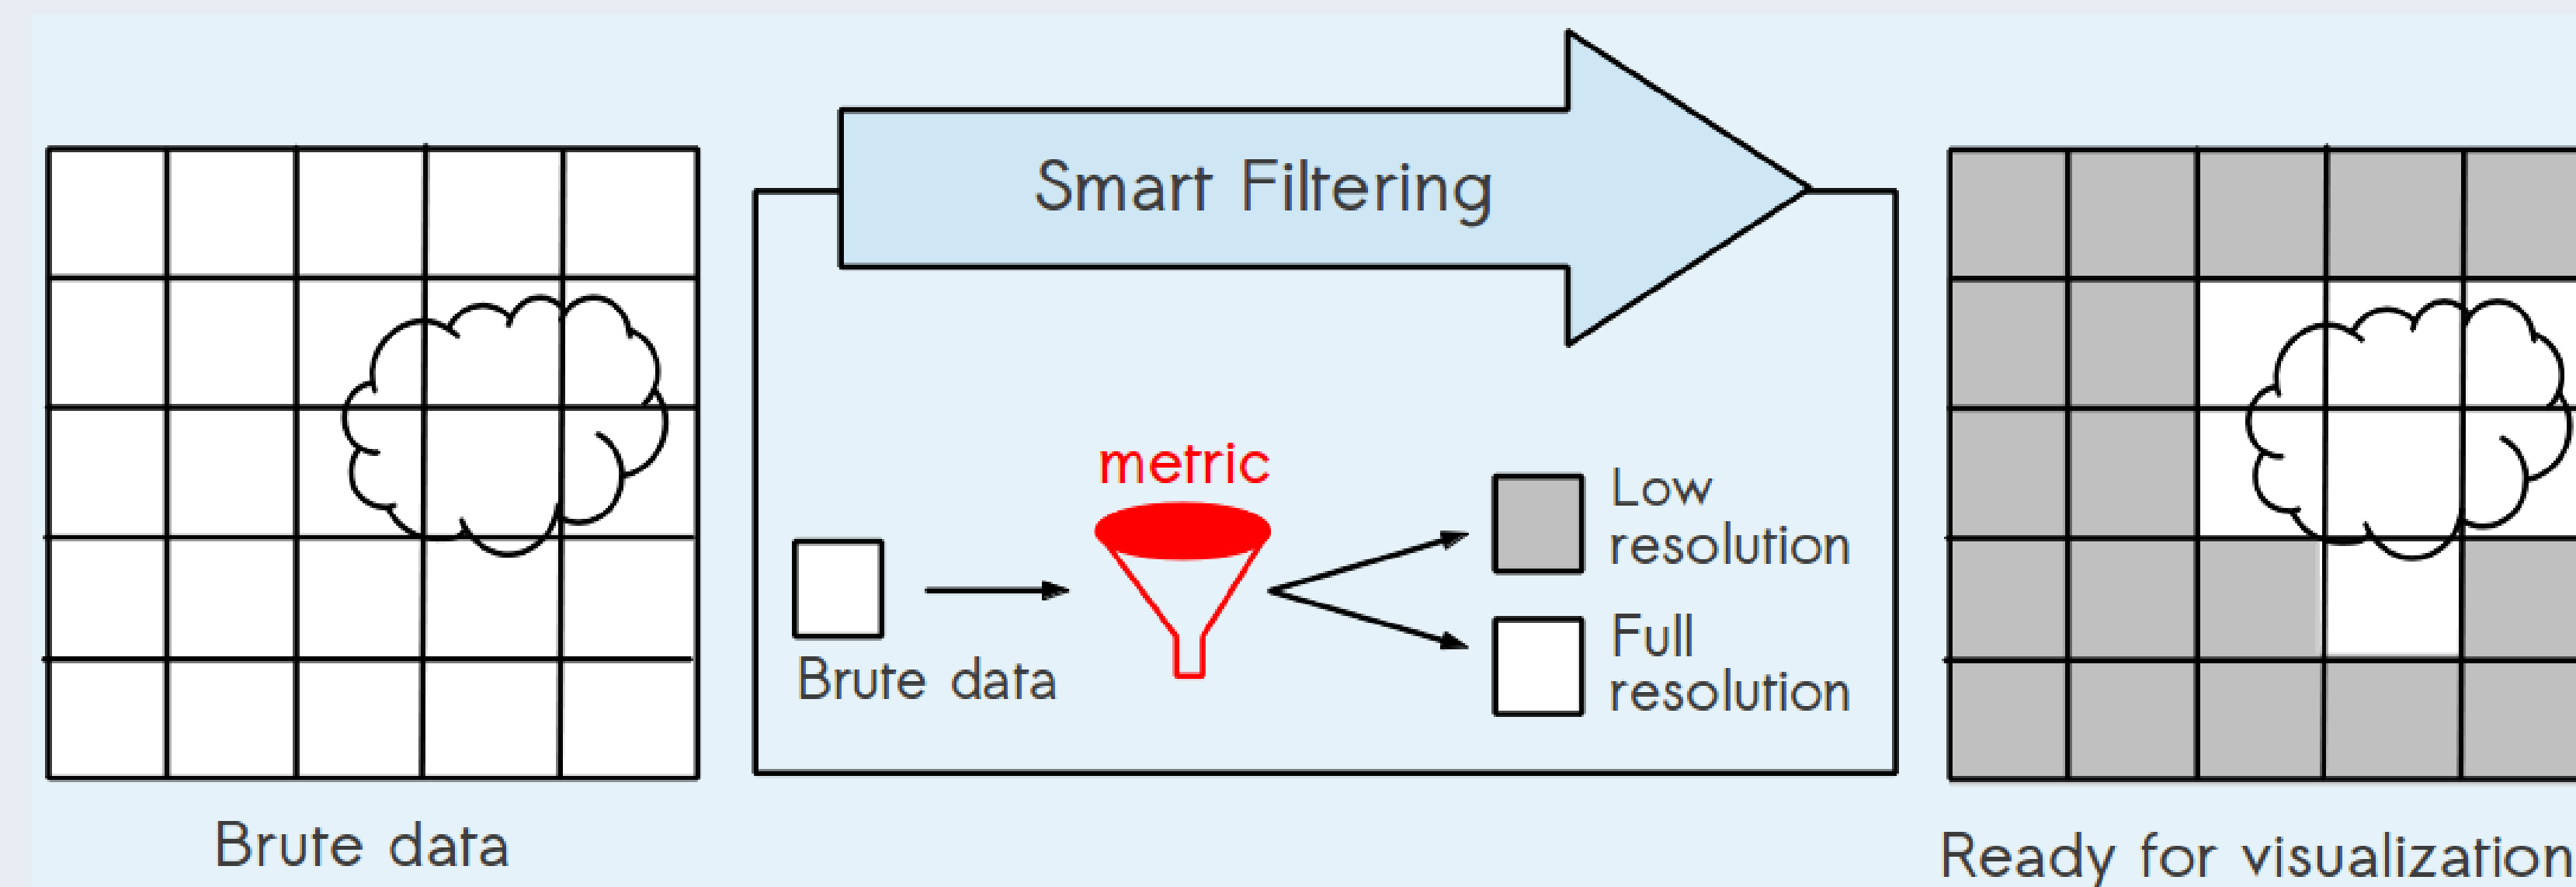
- Direct insight in the simulation
- Bypass the storage system
- Interactive

The Problem : Too much data generated by simulations!

Preparing data for visualization take too much time, which may **block** the simulation. This impacts badly the performances of the whole simulation process.

Our Idea: Smart ISV Framework Architecture

- Define a **semantics** of the data: variation in data
- **Automatically** detect interesting parts of data
- **Adapt** the resolution of the visualization



Implementation and Integration

Integration into existing ISV (Damaris/Viz[2]):

- **Without any code modification** in the simulation
- **No reconfiguration** of the visualization tool (VisIt[3]).

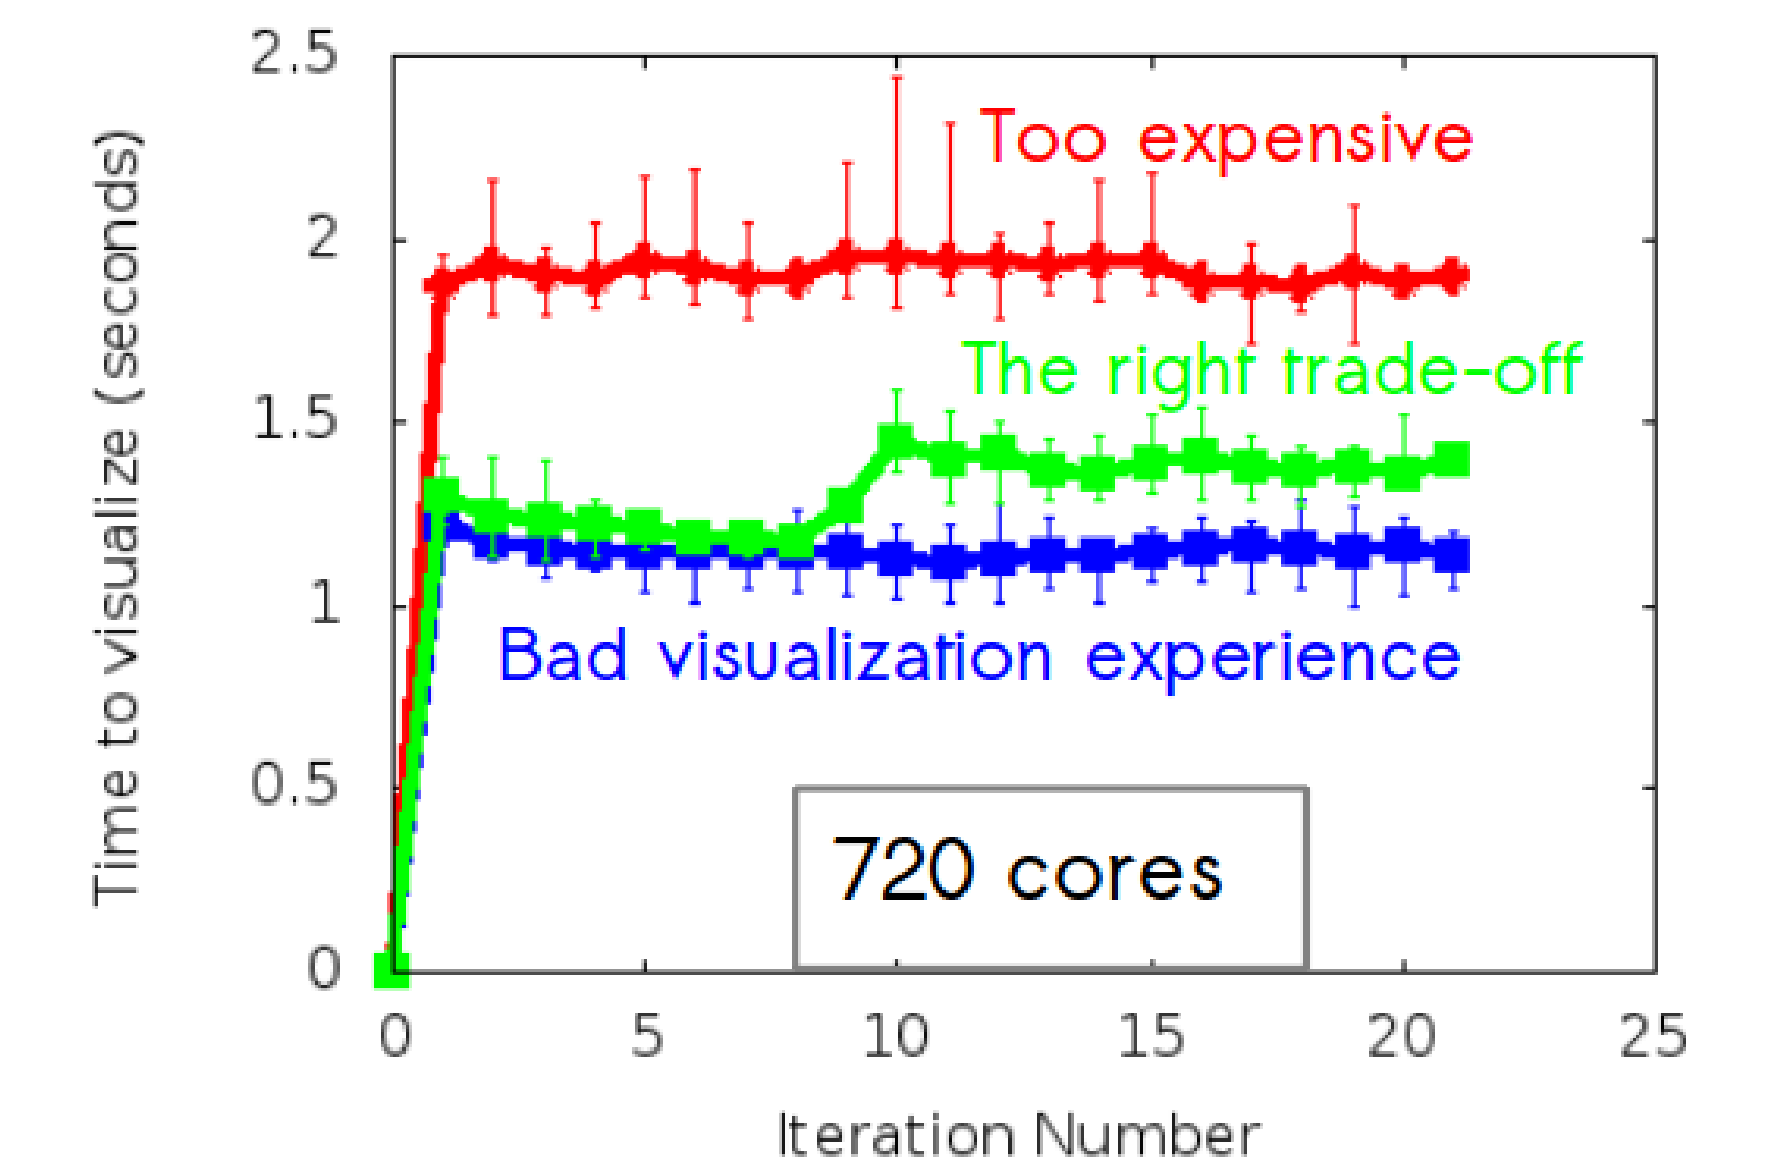
Automatically detect relevant data

Using different **metrics**, with respect to the defined *semantics*:


- Statistics: Coefficient of variation
- Information Theory: Entropy
- Image Processing: Gradient

Results: CM1 application, Grid5000[4] testbed

Impact of Smart ISV: Local Rendering



- Full resolution: Too expensive
- Low resolution: Bad visualization
- Smart ISV: The right trade-off!



Impact of Smart ISV: Scalable Rendering

Conclusion

- 1 Defined data semantics: more **variation** in the data implies interesting data
- 2 **Automatically detect** the relevant parts of the data with respect to the defined semantics.
- 3 Visualizing **only** the relevant parts of the generated data provides up to **40% gain** in visualization time without quality loss.