

Extra-P



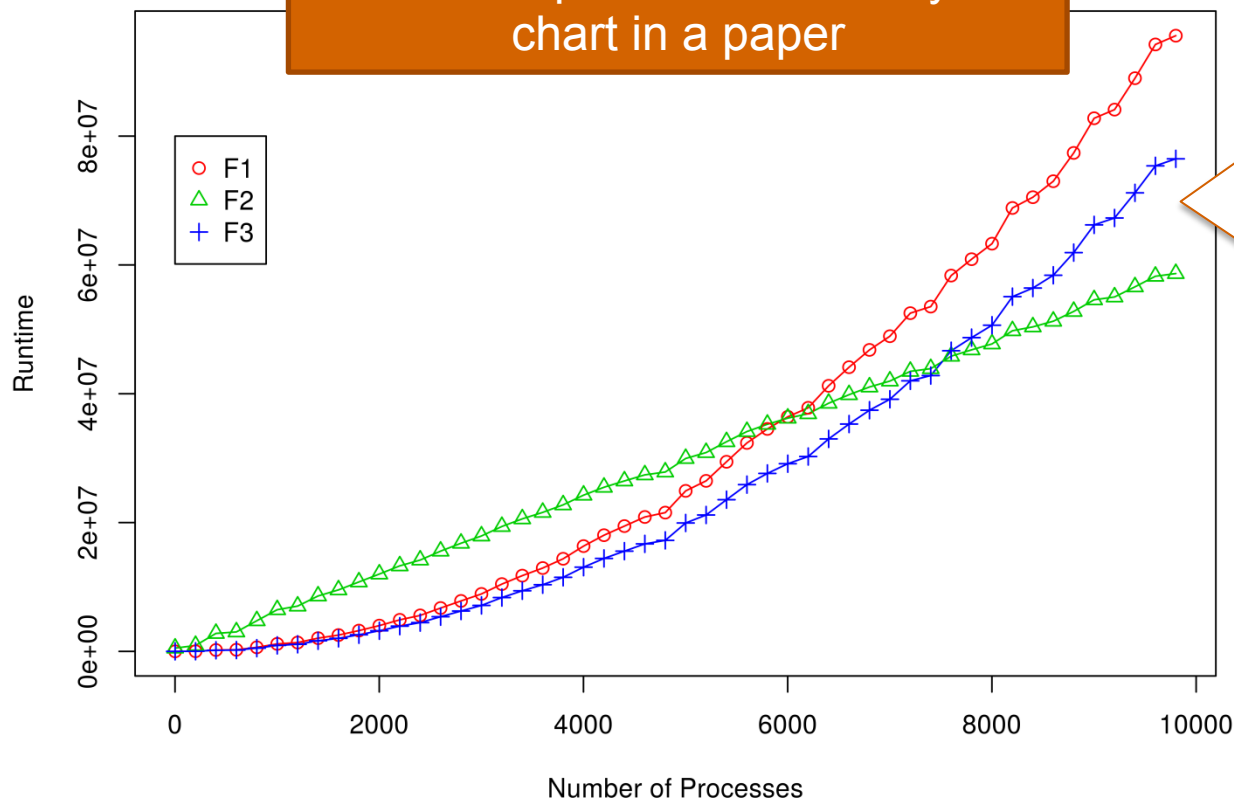
TECHNISCHE
UNIVERSITÄT
DARMSTADT

Follow the Blind Seer – Create Better Performance Models Using Less Information



Why?

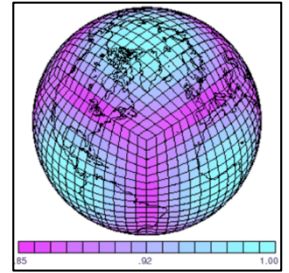
Common performance analysis
chart in a paper



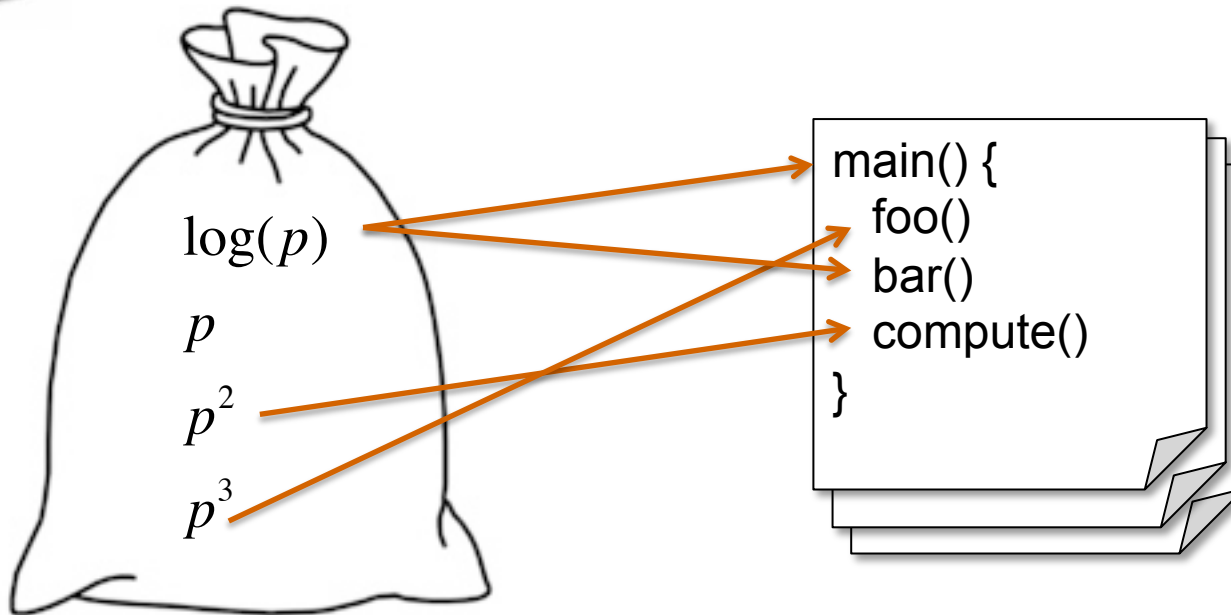
Can't we obtain these
kinds of graphs
automatically?

When?

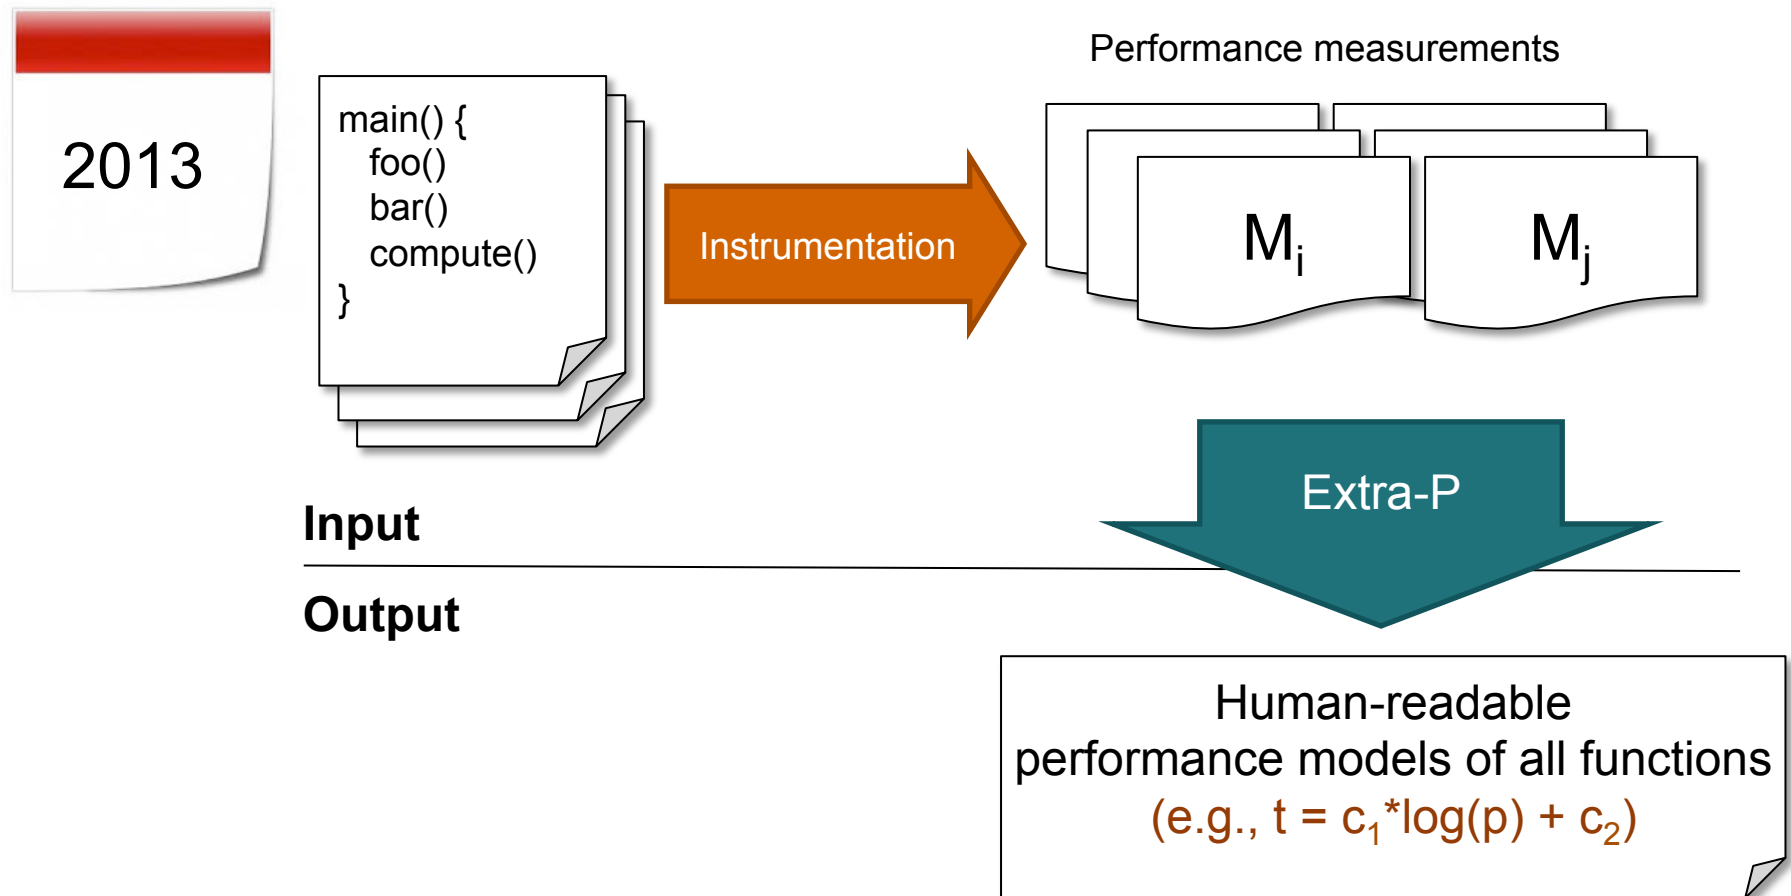
2012



HOMME



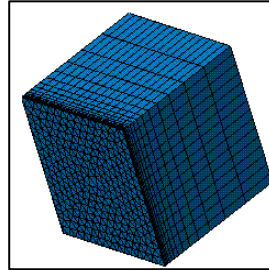
How?



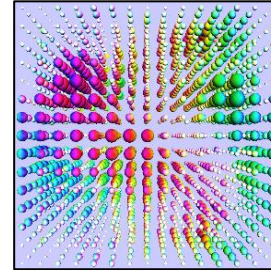
How?



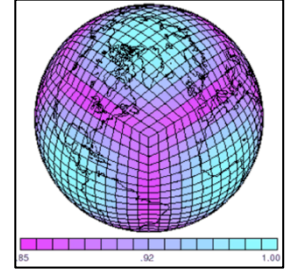
2013



Sweep3d



Milc



HOMME

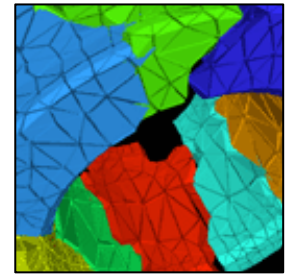
$$f(p) = \sum_{k=1}^n c_k \cdot p^{i_k} \cdot \log_2^{j_k}(p)$$

$$\begin{aligned} n &\in \mathbb{N} \\ i_k &\in I \\ j_k &\in J \\ I, J &\subset \mathbb{Q} \end{aligned}$$

Alexandru Calotoiu, Torsten Hoefler, Marius Poke, Felix Wolf: Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes. In Proc. of the ACM/IEEE Conference on Supercomputing (SC13), Denver, CO, USA, pages 1-12, ACM, November 2013.



What can it do?

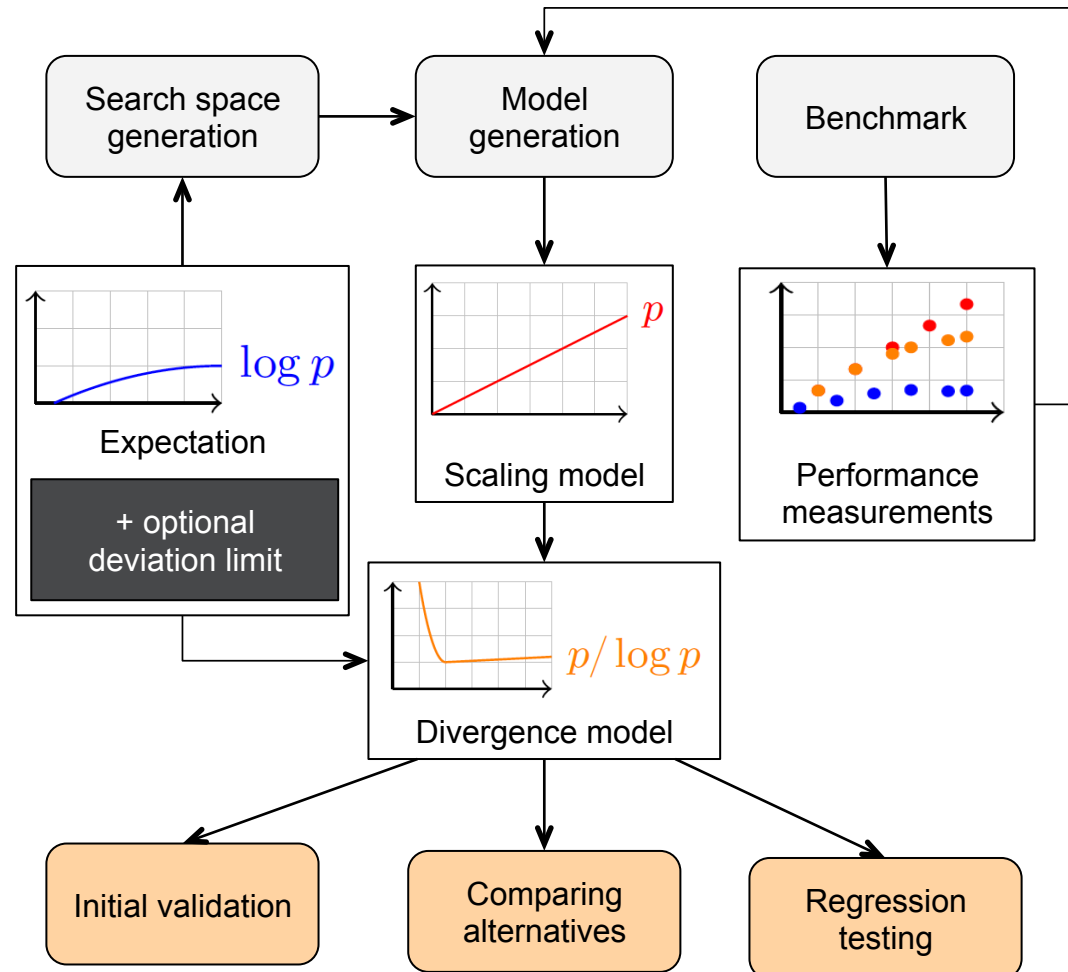


UG4

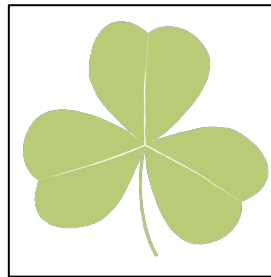
- [1] Sergei Shudler, Alexandru Calotoiu, Torsten Hoefler, Alexandre Strube, Felix Wolf: Exascaling Your Library: Will Your Implementation Meet Your Expectations?. In *Proc. of the International Conference on Supercomputing (ICS), Newport Beach, CA, USA, pages 1-11, ACM, June 2015*
- [2] Andreas Vogel, Alexandru Calotoiu, Alexandre Strube, Sebastian Reiter, Arne Nägel, Felix Wolf, Gabriel Wittum: 10,000 Performance Models per Minute - Scalability of the UG4 Simulation Framework. In *Proc. of the 21st Euro-Par Conference, Vienna, Austria of Lecture Notes in Computer Science, pages 519–531, Springer, August 2015.*
- [3] Christian Iwainsky, Sergei Shudler, Alexandru Calotoiu, Alexandre Strube, Michael Knobloch, Christian Bischof, Felix Wolf: How Many Threads will be too Many? On the Scalability of OpenMP Implementations. In *Proc. of the 21st Euro-Par Conference, Vienna, Austria of Lecture Notes in Computer Science, pages 451–463, Springer, August 2015.*



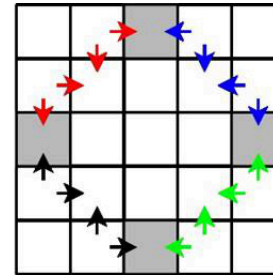
What can it do?



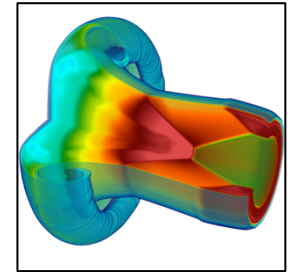
Can EXTRA-P do more?



CloverLeaf



Kripke



BLAST

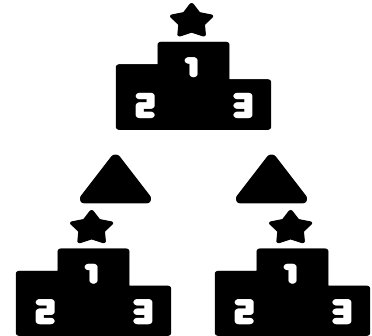
$$f(x_1, \dots, x_m) = \sum_{k=1}^n c_k \prod_{l=1}^m x_l^{i_{kl}} \cdot \log_2^{j_{kl}}(x_l)$$

Alexandru Calotoiu, David Beckingsale, Christopher W. Earl, Torsten Hoefer, Ian Karlin, Martin Schulz, Felix Wolf: Fast Multi-Parameter Performance Modeling. In *Proc. of the 2016 IEEE International Conference on Cluster Computing (CLUSTER)*, Taipei, Taiwan, pages 1-10, IEEE Computer Society, September 2016.



Can EXTRA-P do more?

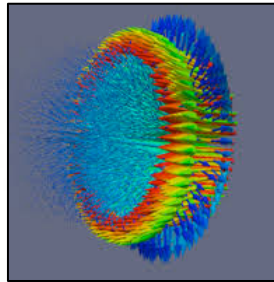
- **Hierarchical search** – Assumes the best multi-parameter model is created out of the combination of the best single parameter hypothesis for each parameter



- **Modified golden section search** – Speeds up the single parameter search by ordering the hypothesis space and then using a variant of binary search to find the model in logarithmic time rather than linear time



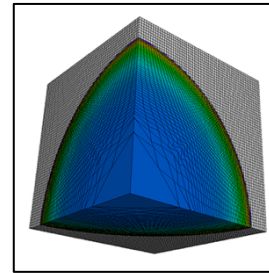
Can Extra-P find more than bugs?



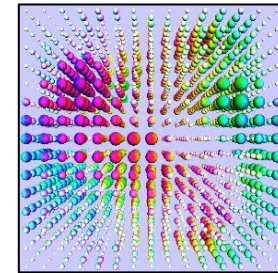
OpenFoam



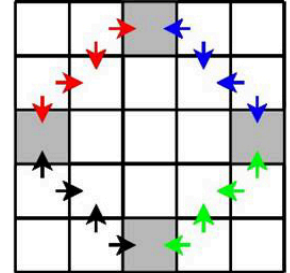
Re-learn



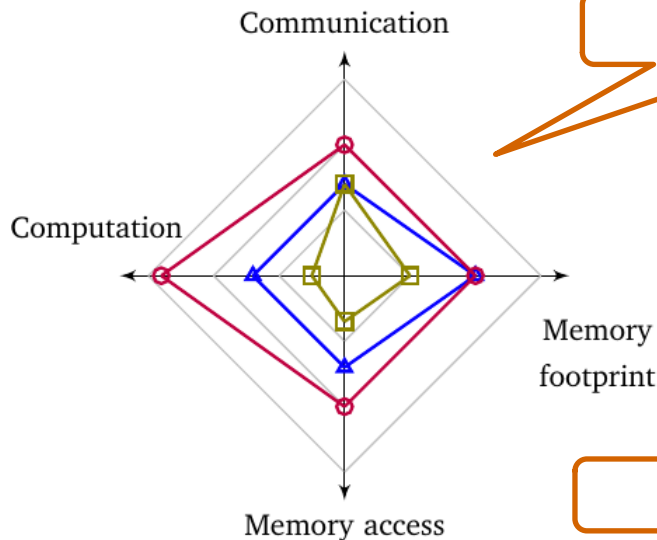
Lulesh



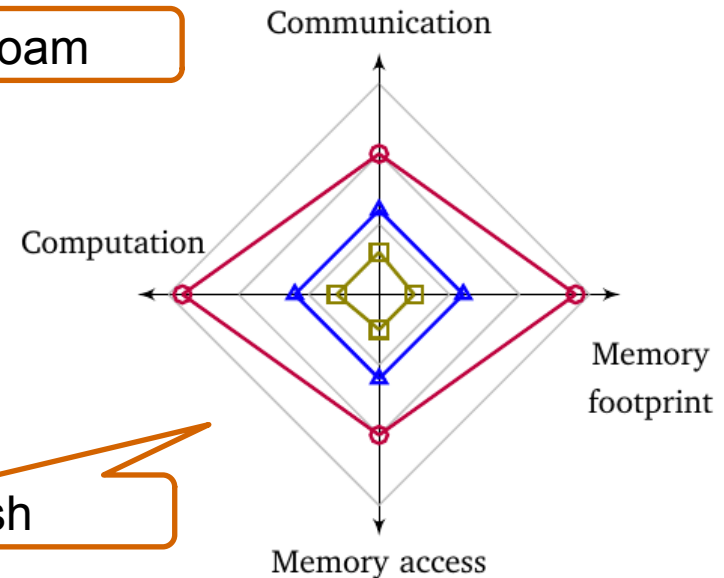
Milc



Kripke



OpenFoam

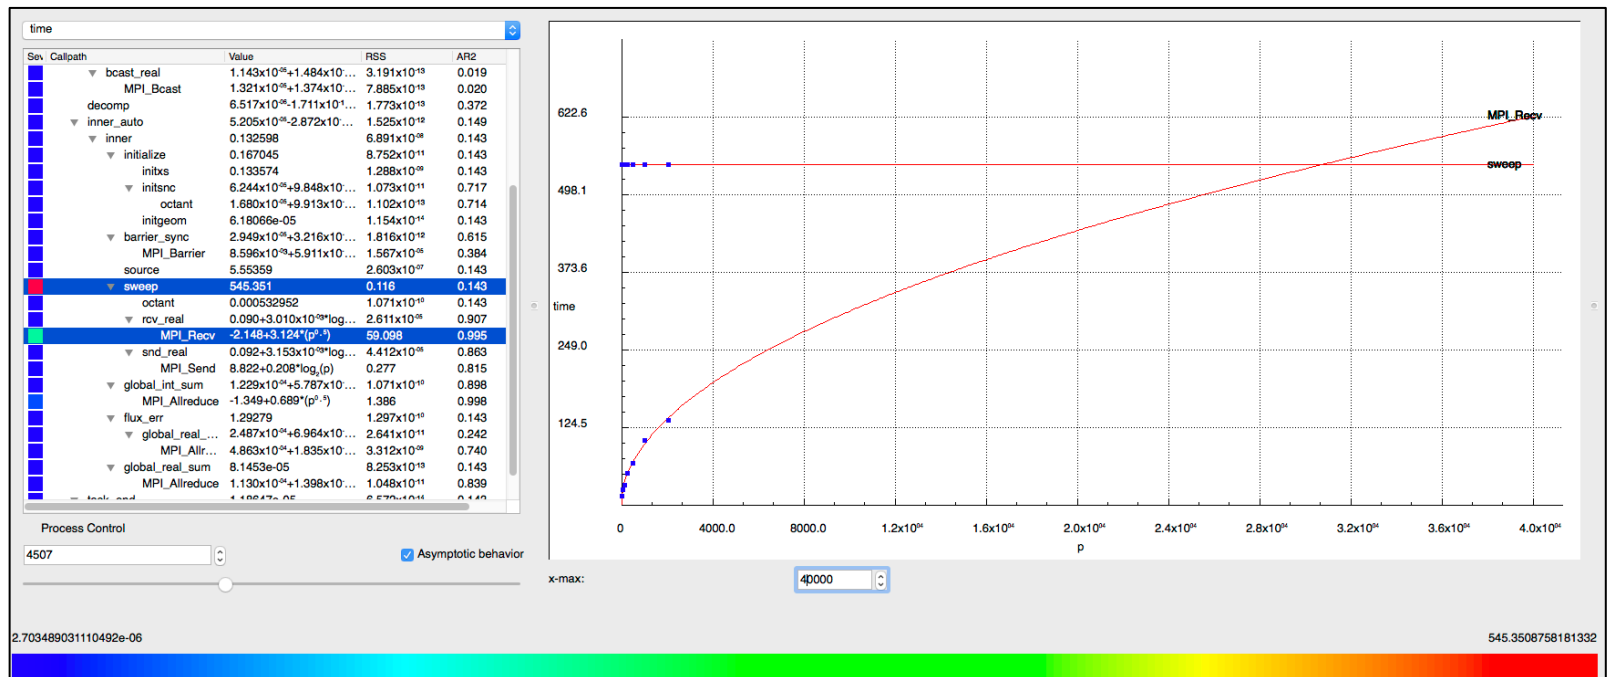


Lulesh

What about now?

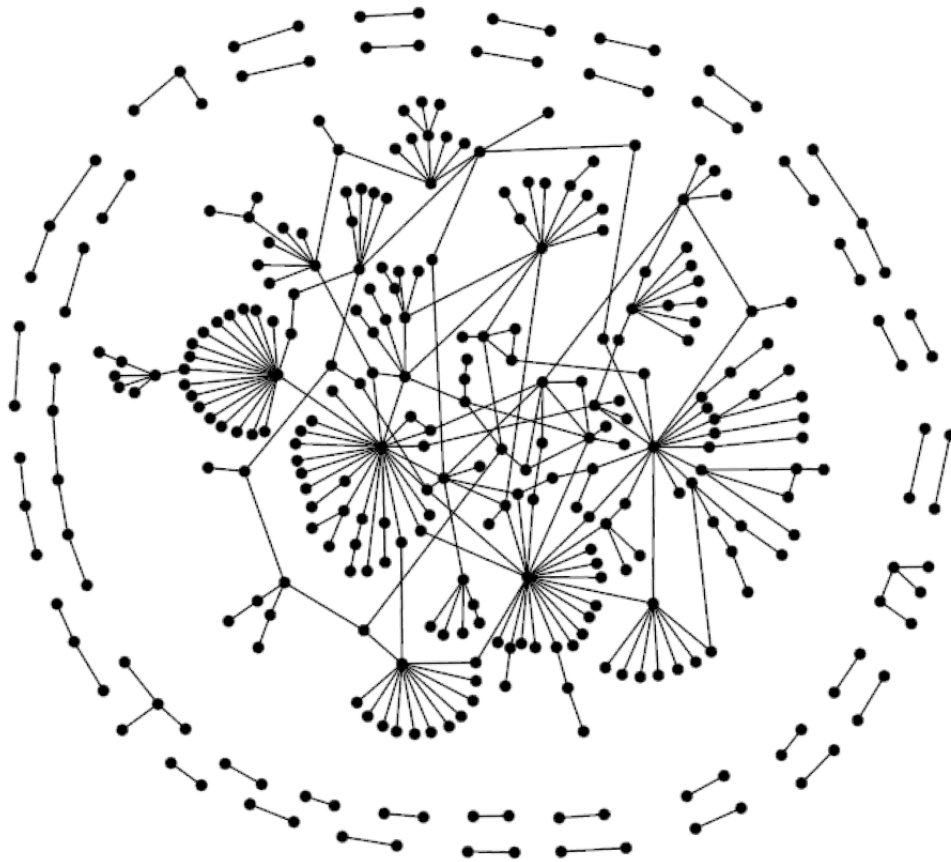
Get it here: <http://www.scalasca.org/software/extra-p/download.html>

Tutorial available on demand – just ask!



Past tutorials: EuroMPI 2015, SC'15, SC'16, 25th. VI-HPS Workshop

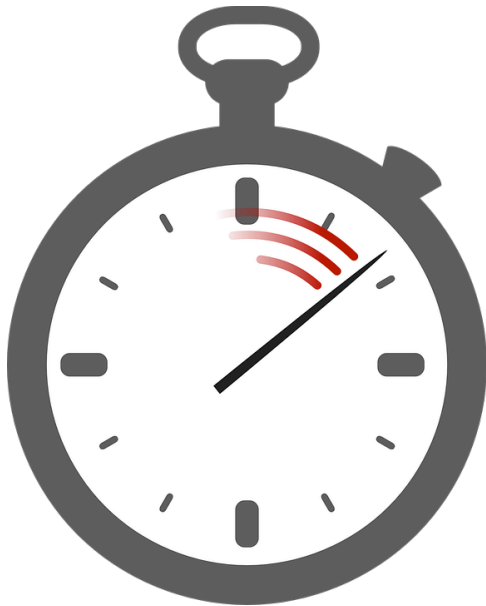
What about tomorrow?



Performance modeling for
graph algorithms

How to identify
metaparameters?

What about tomorrow?



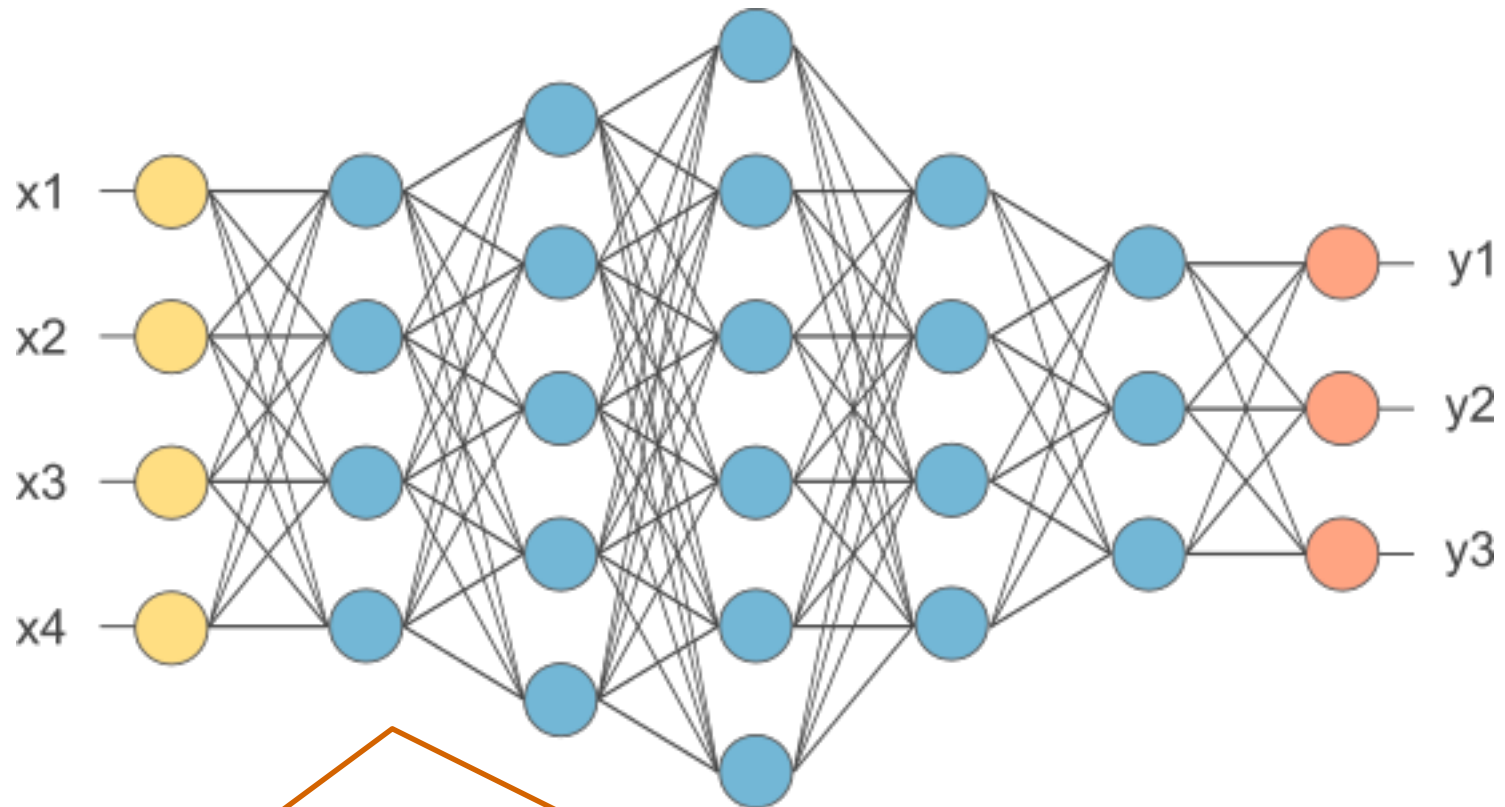
Real-time modeling

Scheduling?

What about tomorrow?



TECHNISCHE
UNIVERSITÄT
DARMSTADT

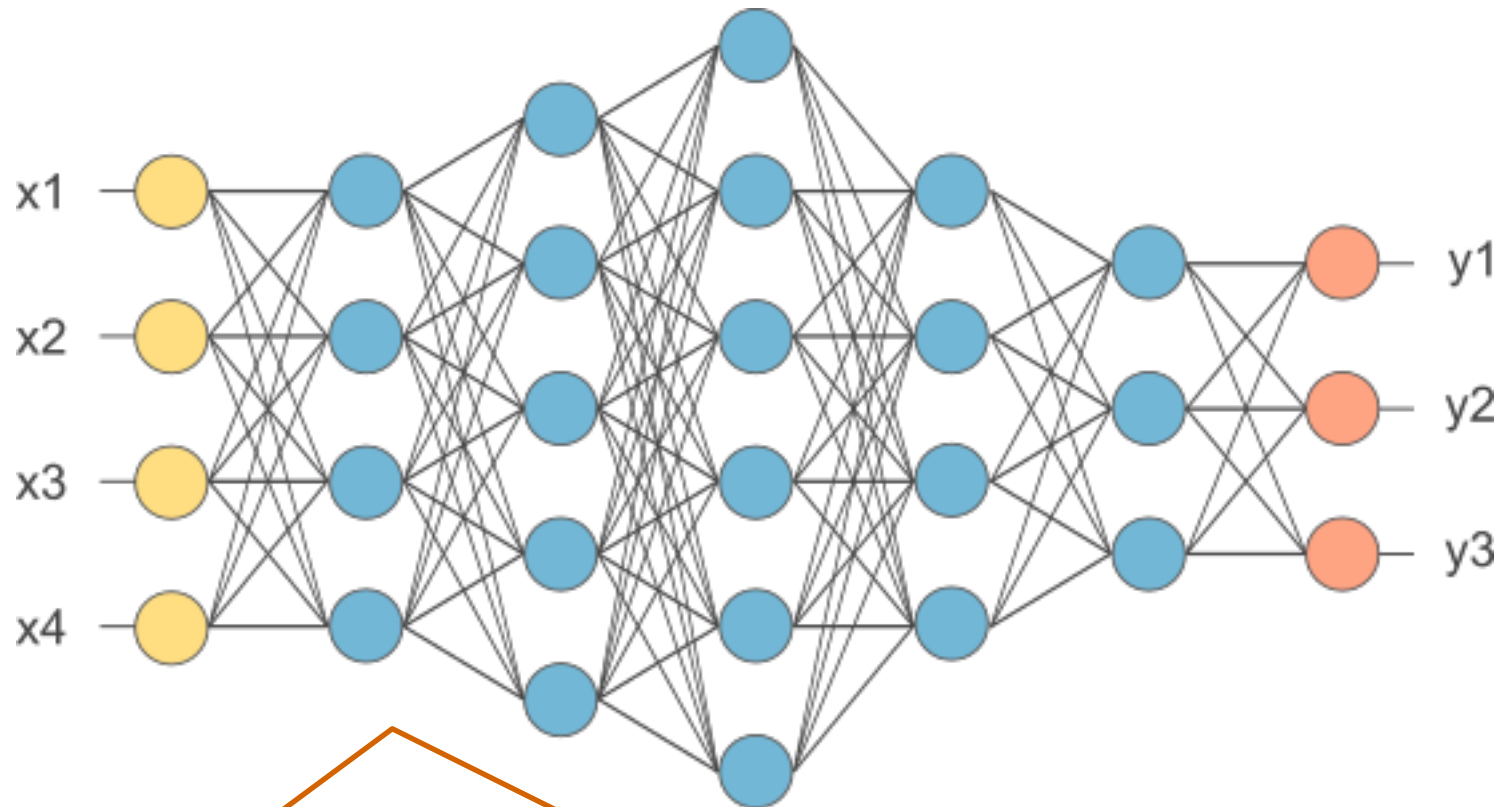


Performance models of DNNs – deployment, characteristics

What about tomorrow?



TECHNISCHE
UNIVERSITÄT
DARMSTADT



Performance models using DNNs – better, worse, or different?