

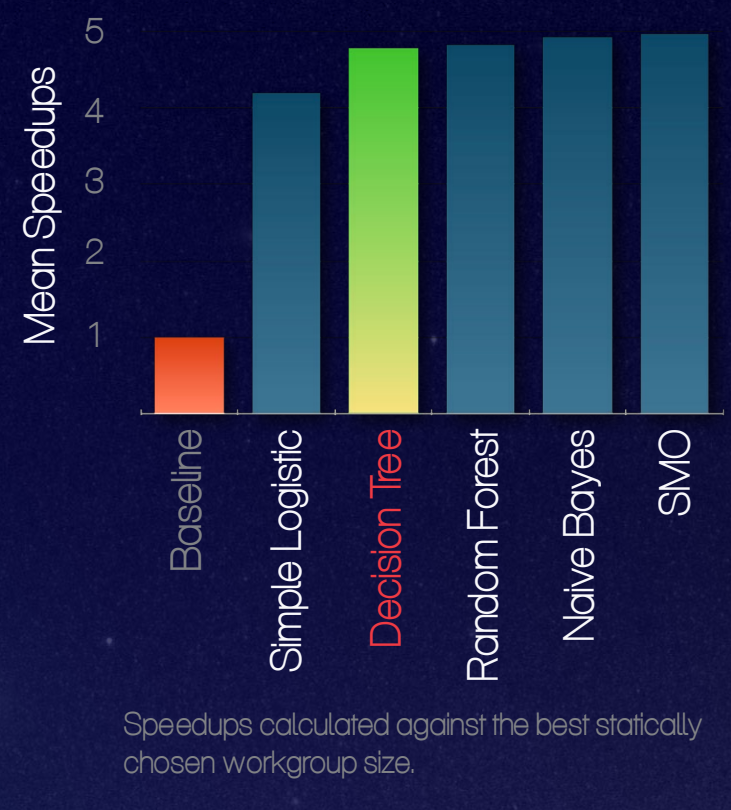


Autotuning OpenCL Workgroup Sizes

Collaborative autotuning for Stencils outperforms human experts

Chris Cummins Pavlos Petoumenos Michel Steuwer Hugh Leather
c.cummins@ed.ac.uk

5.0x
speedup!
Predicting OpenCL workgroup sizes of 429 stencil programs, execution devices, and datasets.

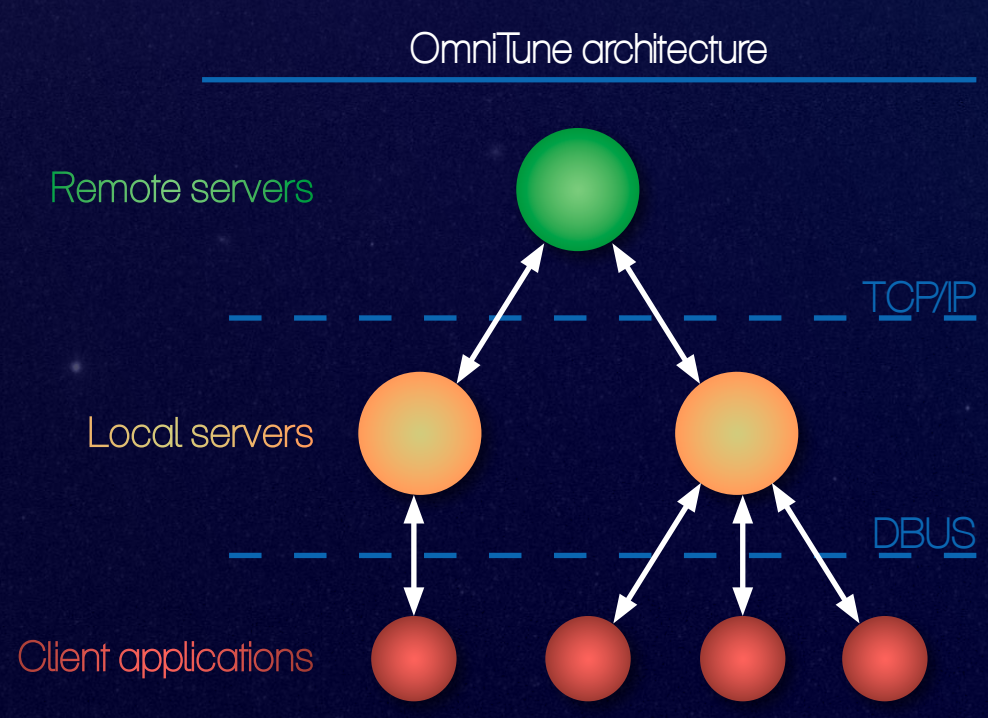
Hand tuning programs is **expensive** and time consuming



We **automate** this task using collaborative machine learning

Introducing OmniTune ...

- OmniTune generates synthetic benchmark programs to use for empirical testing
- OmniTune collaboratively gathers performance data by testing different parameter values
- OmniTune uses machine learning to predict parameters for unseen programs at runtime



Machine Learning features

- Source code**
 - Instruction densities
 - Num basic blocks
 - Stencil region size
 - etc.
- Device**
 - Compute devices
 - Memory sizes
 - Cache types
 - etc.
- Dataset**
 - Container width
 - Container height
 - Input data type
 - etc.

Publications

Chris Cummins, Pavlos Petoumenos, Michel Steuwer and Hugh Leather: "Autotuning OpenCL Workgroup Size for Stencil Patterns" ADAPT 2016.

Chris Cummins, Pavlos Petoumenos, Michel Steuwer and Hugh Leather: "Towards Collaborative Performance Tuning of Algorithmic Skeletons" HPLGPU 2016.

Read more ...
<http://chriscummins.cc>

