Collective TUNING
Using big data and collective intelligence to design efficient adaptive computer systems

Log in

Making computer engineering a science; Systematizing program and system analysis and optimization using auto-tuning, machine learning and crowdsourcing; Enabling self-tuning computer systems

If you are interested in preview, tutorial or presentation about cTuning/Collective Mind crowd-tuning and machine learning technology, its current industrial usages or new publication model, do not hesitate to get in touch with Grigori Fursin!

cTuning-related recent and upcoming events and news:

- 2013.Summer: We are very busy validating new Collective Mind framework and preparing it for release...
- 2013.06.05: Grigori gave a keynote about cTuning/Collective Mind technology at iWAPT'13 @ ICCS'13
- 2013.05.02: cTuning/cMind thematic session on "making computer engineering a science" was held at HiPEAC computing systems week/ACM ECRC in Paris - thanks to your feedback we plan to continue this collaborative activity!
- 2013.03.27: Grigori gave a keynote on statistical crowdsourcing of auto-tuning and tutorial on Collective Mind Framework at the Conference on Advanced Topics and Auto Tuning in High Performance Scientific Computing in NTU, Taiwan
- 2013.02.27: Grigori was an external PhD examiner at Politecnico di Milano for PhD students: Michele Tartara, Ettore Speziale, Simone Corbetta and Paolo Grassi
- 2013.01.22: We organized ADAPT workshop at HiPEAC 2013 sponsored by NVidia and Microsoft Research

Home

Continuing innovation in science and technology is vital for our society and requires ever increasing computational resources. However, delivering such resources became intolerably complex, ad-hoc, costly and error prone due to an enormous number of available design and optimization choices combined with the complex interactions between all software and hardware components, and a large number of incompatible analysis and optimization tools. As a result, understanding and modeling of the overall relationship between end-user algorithms, applications, compiler optimizations, hardware designs, data sets and run-time behavior, essential to provide better solutions and computational resources, became simply infeasible as confirmed by many recent long-term international research visions about future computer systems.

Currently, we see 2 main paths for computer engineering:

- continue developing various ad-hoc techniques and publish non-reproducible experimental results
- join long-term, collaborative initiative to systematize R&D methodology in computer engineering and change publication model to favor reproducible and truely collaborative research

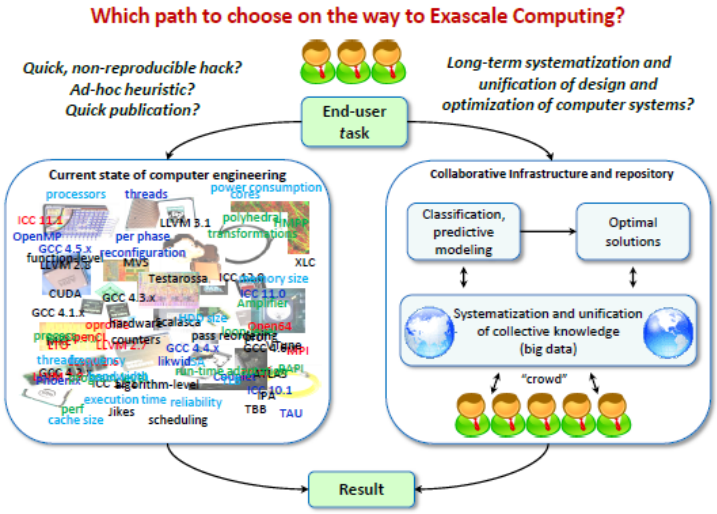

Since 1996, we are working on gradual systematization of knowledge about design and optimization of computer systems based on our background in physics and AI. We are developing collaborative experimental methodology, repository, tools and publication model for computer engineering (cTuning initiative) that favors collaborative discovery, sharing and reuse of knowledge. This technology allows users to:

- share experimental data, tools, models and interfaces (in workgroups through private cTuning repository or with anyone through public cTuning repository)
- collaboratively explore large optimization spaces
- apply classification and predictive models to existing data to explain complex behavior of existing computer systems
- validate data and models by the community (similar to physics, biology and other sciences)
- collaboratively identify anomalies in behavior and suggest how to improve it (statically or dynamically)
- extrapolate this knowledge to build more efficient computer system in terms of performance, power, size, reliability and other characteristics

	Having common collaborative R&D repository and infrastructure allows users to focus their effort on novel approaches

combined with with data mining, classification and predictive modeling rather than spending considerable effort on building new tools with already existing functionality or using some ad-hoc tuning heuristics. It also allows conferences and journals to favor publications that can be collaboratively validated by the community.

Since 2007, we have released all experimental data and tools for our cTuning-related publications [Fur2009], [FCOP2007], [FKMP2011], [FT2010] which allowed further collaborative validation, reproducibility and extension of this technology together with IBM, CAPS, ARC, Intel/CEA Exascale Lab, Google, University of Edinburgh, ICT, UPC, NCAR!

We believe that cTuning may finally complete the puzzle about how to consolidate various ad-hoc techniques and tools together to build efficient self-tuning computer systems. Therefore we strongly advocate for further collaborative R&D and new publication model. If you are interested, join us at upcoming events or use our mailing list and collaborative wiki to participate in discussions and collaborative developments.

We would like to thank all our colleagues for interesting and sometimes tough discussions, feedback, collaborations and support during development of cTuning technology.

You can find more details about history and implementation of cTuning here and in the following reference publications [Fur2009], [FKMP2011], [FT2010], [Fur2004].