

Collective Knowledge Project: a community-driven approach to performance optimization

Grigori Fursin, France



Designing novel computer systems and optimizing their software is becoming too tedious, ad hoc, time consuming and error prone due to enormous number of available design and optimization choices. Empirical autotuning combined with run-time adaptation and machine learning has been demonstrating some potential to address above challenges for several decades but is still far from the widespread production. The main reasons include unbearably long exploration and training times, ever changing tools and their interfaces, lack of a common experimental methodology, lack of diverse and representative benchmarks, and lack of unified mechanisms for knowledge building and exchange apart from publications where reproducibility and reusability of results is often not even considered.

I will present our community-driven solution to above problems based on our open-source Collective Knowledge technology (CK) that can gradually organize, exchange and reuse knowledge and experience in computer engineering. CK helps share various artifacts (benchmarks, data sets, libraries, tools) as unified, reusable and Python-based components with JSON meta description via GITHUB. Researchers can then quickly prototype and crowdsource various experimental workflows such as performance and energy autotuning, design space exploration and run-time adaptation. At the same time, CK continuously analyzes and extrapolates all collected knowledge using powerful data science techniques to automatically model computer systems' behavior, predict better optimizations or hardware configurations, and eventually enable faster, more power efficient, reliable and self-tuning software and hardware. Furthermore, CK can record any unexpected behavior in a reproducible way and expose it to an interdisciplinary community to find missing features and improve models. Live demo of our approach is available at <http://cknowledge.org/repo>.

Grigori Fursin is a chief scientist of the non-profit cTuning foundation (France) and a CTO of a startup dividiti (UK). In the past, he was a senior research scientist at INRIA and a co-founder of the Intel Exascale Lab in France. Grigori has an interdisciplinary background in computer engineering, physics, electronics, machine learning and mathematics. He obtained PhD in Computer Science from the University of Edinburgh in 2004. Grigori pioneered crowdsourcing of machine-learning based program autotuning and hardware co-design to deliver faster, smaller and more power efficient systems. In 2008, he established one of the first public repositories of optimization knowledge (cTuning.org) connected to self-tuning GCC compiler to continuously collect and predict optimizations based on program, hardware and data set features. Since then, cTuning technology has been used and extended in multiple industrial projects together with IBM, ARM, Intel, STMicroelectronics and Synopsys. In 2012, Grigori received INRIA award and 4-year fellowship for "making an outstanding contribution to research". He also recently initiated Artifact Evaluation initiative for CGO and PPOPP conferences to promote sharing and community-driven validation of benchmarks, datasets, tools and experimental results.