

CURRICULUM VITAE - GRIGORI FURSIN, PHD

Current position	<ul style="list-style-type: none">• CTO, dividiti, UK (since 2015)• Chief Scientist and Technologist, non-profit cTuning foundation, France (since 2014)
Experience	Leading successful industrial and academic R&D projects for more than 20 years.
Languages	English - fluent (British citizen); Russian - native; French (spoken) - intermediate
Address	At the moment, I live in Paris suburbs. However, I regularly commute to UK/USA where I have my current and main industrial and academic projects.
Education	<ul style="list-style-type: none">• PhD in computer science with ORS award from the University of Edinburgh, UK (2004)
Website	fursin.net/research
Prof. memberships	ACM, HiPEAC, IEEE
LinkedIn	fr.linkedin.com/in/grigorifursin
Personal e-mail	gfursin@gmail.com



I have an interdisciplinary background in computer engineering, electronics, physics, mathematics, machine learning and neural networks with more than 20 years of experience in leading successful and influential research and development projects. I hope that my techniques and open source technology will eventually enable efficient computing via knowledge sharing and adaptation while improving the quality and reproducibility of our experimentation, and boosting innovation in science and technology. Therefore, I am always happy to collaborate with companies and universities where I can continue implementing my ideas and long-term vision.

PROFESSIONAL CAREER

- 2015-cur.: CTO, dividiti, UK.
- 2014-cur.: Chief Scientist and Technologist, non-profit cTuning foundation, France.
- 2014-cur.: Consultant and scientific advisor (knowledge management, brain-inspired computing, machine learning based performance/energy/size/cost autotuning, run-time adaptation, SW/HW co-design).
- 2012-2014: Tenured Research Scientist (associate professor) at INRIA, France.
- 2010-2011: Head of application characterization and optimization group at Intel Exascale Lab, France.
- 2008-cur.: Founder of cTuning.org to crowdsource machine-learning based autotuning as an open research.
- 2007-2010: Guest lecturer at the University of Paris-Sud, France.
- 2007-2010: Tenured Research Scientist (assistant professor) at INRIA, France.
- 2000-2006: Research Associate at the University of Edinburgh, UK.
- 1999-cur.: Evangelist of a collaborative and reproducible research and experimentation in computer engineering.
- 1993-1999: Research Assistant at MIPT, Russia.
- 1992-1993: Founder and CTO of a software startup in Moscow.

EDUCATION

- 2004: PhD in computer science with ORS award from the University of Edinburgh, UK.
- 1999: MS in computer engineering with golden medal (summa cum laude) from Moscow Institute of Physics and Technology, Russia.
- 1997: BS in electronics, mathematics and machine learning (summa cum laude) from Moscow Institute of Physics and Technology, Russia.

ACADEMIC PARTNERS

Edinburgh University (UK), Cambridge University (UK), University of Copenhagen (Denmark), Manchester University (UK), Imperial College London (UK), University of Pittsburgh (USA), UCAR (USA), INRIA (France), ENS Paris (France).

ASSORTED TECHNICAL KNOWLEDGE

Main technical knowledge (continuously acquire new ones): Linux, Windows, Android, Python, scikit, neural networks, decision trees, SVM, agile development, large-scale project management, APIs, GCC, LLVM, polyhedral optimizations, ARM compilers, Intel Compilers, Intel VTUNE, C, C++, Java, Fortran, Basic, GPU, OpenCL, CUDA, MPI, OpenMP, PHP, R, MySQL, FPGAs, Elasticsearch, Hadoop, Jenkins, html, apache2, mediawiki, drupal, OpenOffice, Eclipse, SVN, GIT, GIMP2, Adobe Photoshop, Visual Studio, Microsoft Office, Android Studio

HOBBIES

Traveling, discovering new cultures, gardening, active sport (football, skiing, swimming, snorkeling, climbing, jogging, ...), photography, reading.

Contents

1 RESEARCH SUMMARY	3
2 CURRENT RESEARCH FOCUS	3
3 BIOGRAPHY	4
4 EDUCATION (Z)	4
5 PROFESSIONAL EXPERIENCE (J)	5
6 INSTITUTION BUILDING (I)	8
7 EDITOR (Ed)	10
8 KEYNOTES (K)	10
9 SOCIAL ACTIVITIES (O)	10
10 STARTUPS (C)	11
11 EXAMINER	11
12 EXPERT SERVICE (E)	11
13 MAJOR RESEARCH ACHIEVEMENTS (M)	12
14 PUBLIC OR IN-HOUSE REPOSITORIES OF KNOWLEDGE (R)	15
15 AWARDS, PRIZES AND FELLOWSHIPS (A)	17
16 MAJOR FUNDING (F)	18
17 MAJOR SOFTWARE AND DATASETS (S)	19
18 HARDWARE (H)	25
19 TALKS (T)	25
20 PARTICIPATING IN PROGRAM COMMITTEES AND REVIEWING	28
21 TEACHING AND ORGANIZING COURSES (L)	29
22 ADVISING/COLLABORATING (Q)	29
23 ORGANIZING/CHAIRING EVENTS (E)	31
24 PUBLICATIONS AND OTHER DISSEMINATION MATERIAL (P)	37
24.1 THESES	37
24.2 INTERNATIONAL JOURNALS	38
24.3 INTERNATIONAL CONFERENCES	39
24.4 INTERNATIONAL WORKSHOPS	42
24.5 NATIONAL CONFERENCES AND WORKSHOPS	45
24.6 POSTERS	46
24.7 TECHNICAL REPORTS, NEWSLETTERS, EXTENDED ABSTRACTS, INTRODUCTIONS	46
24.8 PRESS RELEASES	48

1 RESEARCH SUMMARY

I have an interdisciplinary background in computer engineering, electronics, physics, mathematics and machine learning with more than 20 years of experience in setting up and leading successful and influential research and development projects based on my long-term vision and interests [M1].

To be able to continue my original research on designing semiconductor neural network accelerators for bio-inspired self-adaptive computers and AI systems (1994-1997), I desperately needed faster, cheaper, more power efficient and reliable computer systems as well as unified mechanisms for sharing knowledge in a reproducible way across colleagues.

In 1998, I switched to computer engineering with an ambitious goal to radically change ad-hoc, error prone and time consuming benchmarking, optimization and co-design of computer systems across all their software and hardware layers (heterogeneous multi-core architectures, compilers, run-time libraries, applications) into a unified physics-based "big data" problem.

I then pioneered a new and complete solution by combining statistical analysis, machine learning based autotuning (classification and predictive analytics), data mining (finding missing features), multi-versioning combined with run-time adaptation and online learning as a dynamic reaction to behavior changes, my public cTuning.org technology and Collective Mind repository of knowledge/Collective Knowledge infrastructure and repository, infrastructure for experiment crowdsourcing using commodity mobile phones, and collective intelligence [Pub25, P24, Pub3, M3, Pub26, Pub5, Pub6, Pub50].

My technology combined with collective knowledge can help automatically and intelligently explore large design and optimization spaces, detect some hardware and software bugs, automatically build and improve models of program behavior on various hardware (performance, energy, faults), detect scalability and other issues (particularly in OpenCL, CUDA and MPI codes), enable run-time adaptation for statically compiled programs across numerous data sets and architecture designs, and much more [Pub25, Pub3].

All these techniques, my open source research SDK and shared artifacts were used and extended in various academic and industrial projects together with IBM, Intel, ARC(Synopsys), STMicroelectronics, ARM, and other companies to improve efficiency and utilization of computer systems while dramatically reducing their development costs and time to market, included in mainline GCC (plugin framework) [S16], referenced in international press-releases by IBM [Ref] and Fujitsu [Ref], published in major conferences and journals including PLDI, MICRO, CGO, TACO, HiPEAC, CASES [Pub25], received several awards [A2, A1], used in guest lectures [L1], and included in the EU HiPEAC long-term research vision (2012-2020).

At the same time, I spent considerable effort since 1999 to promote sharing of all research artifacts in a reusable and reproducible way along with publications, and, more importantly, to develop a supporting technology (rather than only writing papers and manifestos). My practical cTuning.org technology and experience gradually helped initiate new publication model in computer engineering where experimental results and all research artifacts are continuously shared, discussed, reproduced and improved by the community [Pub26, E3, E1, E2, E4, E5]. I hope that my techniques and open source technology will eventually enable efficient computing via knowledge sharing and adaptation while improving the quality and reproducibility of our experimentation, and boosting innovation in science and technology. Therefore, I am always happy to collaborate with companies and universities where I can continue implementing my ideas and long-term vision.

2 CURRENT RESEARCH FOCUS

- Systematize and automate software and hardware analysis, benchmarking, validation, optimization and co-design to enable faster, smaller, cheaper, more power efficient and reliable computer systems using public cTuning/Collective Mind Infrastructure and Repository (plugin-based knowledge management system) based on multidisciplinary approaches including empirical tuning, machine learning, statistical analysis, and crowdsourcing using commodity mobile phones [S2,R2].
- Continue preparing and promoting new publication model in computer engineering to validate experimental results by the community [M2,E12, Pub50, S2, R2].
- Transfer knowledge to industry through consulting (development of fast and power-efficient, intelligent and self-tuning computer systems; machine learning based design and optimization space modeling and exploration; brain modeling, etc) and creation of interdisciplinary laboratories.

3 BIOGRAPHY

I am a computer scientist, engineer and manager with an interdisciplinary background in computer engineering, electronics, physics, mathematics and machine learning, and with more than 20 years of research and development experience. I am always interested to lead highly challenging, innovative and interdisciplinary research projects particularly related to a development of faster, smaller, cheaper, more power efficient and reliable self-tuning computer systems; collaborative discovery, systematization, sharing and reuse of knowledge; machine learning; big data analytics; artificial brain and collective intelligence.

In 1993, at the age of 16, I joined my first official R&D project as a Research Assistant designing and optimizing semiconductor neural network accelerators for a possible brain-inspired supercomputer computer [M10, M9]. After tedious attempts to optimize and parallelize my neural network modeling software for several supercomputers, I decided to switch to computer engineering to automate and systematize this process using my interdisciplinary knowledge. Eventually, I was one of the first researchers to radically change ad-hoc, error prone, time consuming and costly process of designing, benchmarking and optimizing the next generation of computer systems across all software and hardware layers into a unified big data problem [Pub3, Pub3, M3, P24, Pub5, Pub6, Pub50]. I then started tackling, systematizing and speeding it up using my open source plugin-based cTuning and Collective Mind multi-dimensional and multi-objective auto-tuning infrastructure and repository, statistical analysis, machine learning (classification and predictive analytics), data mining (finding missing features), run-time adaptation with static multi-versioning, adaptive exploration of large optimization spaces, online tuning, differential analysis, crowdsourcing using commodity mobile phones and collective intelligence.

Besides publishing in major conferences and journals including PLDI, MICRO, CGO, TACO, IJPP, CASES and HIPEAC, I spend considerable effort to release all my code and data including tools, benchmarks, data sets and predictive models along with my articles at cTuning.org and c-mind.org/repo to ensure reproducibility. As a side effect, this approach initiated a new open publication model in computer engineering where experimental results and all research artifacts are continuously shared along with articles to be validated and improved by the community [Pub26, E4,E5, Pub26, M2,E12, Pub50]. Eventually, most of my techniques became a mainstream, have been included into mainline GCC, validated and extended by industry with IBM, ARC (Synopsys), Intel, Google, STMicroelectronics and ARM. These techniques enabled practical open-source machine learning based self-tuning compiler (MILEPOST GCC) [M4,R4,R2] considered by IBM to be the first in the world [Pub59]. They also dramatically reduced development costs and time to market of the new embedded reconfigurable devices from ARC (Synopsys) while improving their performance, power consumption, time, size and ROI [X9]. Finally, my techniques were included in the European Union's long-term IT research vision for 2012-2020 [X8, X3] In 2008, I established international nonprofit cTuning.org foundation for community-driven systematization, automation and acceleration of the design, benchmarking and optimization of all existing computer systems across the whole software and hardware stack using my repository of knowledge, open source tools and big data analytics. My foundation and community-driven approach has been referenced in press-releases from IBM in 2008-2009 and from Fujitsu in 2014.

I delivered more than 60 regular and invited talks, lectures and keynotes in the major international companies and universities in Europe, USA, China, Canada and Russia; founded SMART and ADAPT workshops that ran consecutively for 8 years sponsored by Google, NVidia, Intel and Microsoft; prepared and taught advanced MS course on future self-tuning computer systems in the Paris South University. In 2010-2011, I was on industrial leave invited to help establish Intel Exascale Lab in France while preparing long-term research directions and serving as the head of program optimization and characterization group [I2]. In 2012, I rejoined INRIA and received a personal award and 4-year fellowship for "making an outstanding contribution to research" [A2]. I hope my research will help boost innovation in science and technology.

4 EDUCATION (Z)

#	Year	Degree/course
[Z1]	2013-10	Participating in Dagstuhl seminar on "Automatic Application Tuning for HPC Architectures" in Germany Website: http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=13401
[Z2]	2008-07	Attending 4th International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems (ACACES 2008) organized by HIPEAC in L'Aquila, Italy <ul style="list-style-type: none">• Associated poster [Pub46] Website: http://www.hipeac.net/acaces2008 Courses: http://www.hipeac.net/acaces2008/index.php?page=courses

[Z3] 2007-07 Attending 3rd International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems (ACACES 2007) organized by HiPEAC in L'Aquila, Italy

- *Associated poster [Pub47]*

Website: <http://www.hipeac.net/acaces2007>
Courses: <http://www.hipeac.net/acaces2007/index.php?page=courses>

[Z4] 2006-07 Attending 2nd International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems (ACACES 2006) organized by HiPEAC in L'Aquila, Italy

Website: <http://www.hipeac.net/acaces2006>
Courses: <http://www.hipeac.net/acaces2006/index.php?page=courses>

[Z5] 2004-05 PhD in computer science from the University of Edinburgh, UK

- *Awards: ORS award*
- *Advisor: Prof. Michael O'Boyle (University of Edinburgh, UK)*
- *Examination board: Francois Bodin (CAPS Entreprise, France), Marcelo Cintra (University of Edinburgh, UK)*
- *Associated thesis [Pub1]*

[Z6] 1999-06 MS summa cum laude in computer engineering from Moscow Institute of Physics and Technology and Institute of High-Performance Computing Systems of Russian Academy of Sciences, Russia

- *Awards: golden medal*
- *GPA=4.0/4.0; TOEFL=593; GRE=x/800/780*
- *Associated thesis [Pub2]*

[Z7] 1999-06 MS summa cum laude in computer engineering from Moscow Institute of Physics and Technology and Institute of High-Performance Computing Systems of Russian Academy of Sciences, Russia

- *Awards: golden medal*
- *GPA=4.0/4.0; TOEFL=593; GRE=x/800/780*
- *Associated thesis [Pub2]*

[Z8] 1997-06 BS summa cum laude in physics and electronics from Moscow Institute of Physics and Technology, Russia

- *Interdisciplinary courses in physics, electronics, mathematics, statistics, machine learning, brain simulation and computer science*
- *Associated publications [Pub41, Pub40, Pub39]*

[Z9] 1993-06 Graduated from Moscow Secondary School No249 with medal

[Z10] 1993-06 Graduated from Moscow College of Physics and Technology, Russia

- *GPA=4.0/4.0*

5 PROFESSIONAL EXPERIENCE (J)

#	Year	Job
[J1]	2015-07 cur.	- CTO at dividiti (UK)

[J2] 2014-11 - Chief Scientist and Technologist *at the cTuning foundation (France)* developing open-source infrastructure and repository for systematic, collaborative, automatic and reproducible benchmarking, optimization and co-design of computer systems based on crowdsourcing, predictive analytics and collective intelligence

- *Funded by EU FP7 TETRACOM project [F1]*
- *Developed Collective Knowledge software and repository [S1, R1]*
- *Received HiPEAC award [A1]*
- *Publications [Pub25]*

[J3] 2011-11 - Advisory board member *of Exascale LLC (USA)*
cur.

[J4] 2011-09 - Tenured research scientist (on sabbatical) *in INRIA Saclay (France)* directing R&D of the novel Collective Mind concept for collaborative, systematic and reproducible benchmarking, optimization and co-design of computer systems using public repository of knowledge, plugin-based auto-tuning, run-time adaptation, big data, predictive analytics, machine learning, data mining, statistical analysis, feature selection, crowdsourcing and collective intelligence

- *Funded by INRIA 4 year fellowship [A2]*
- *Initiated new publication model supported by HiPEAC where experimental results (tools, data, models) are continuously shared and validated by the community [E12, E14]*
- *Associated software [S2, S9, S3, S4, S5]*
- *Associated publications [Pub50, Pub29, Pub51, Pub52, Pub42, Pub53, Pub43, Pub4, Pub54, Pub44]*
- *Associated events [E10, E12, E13, E15, E16, E17]*
- *Associated public repository of knowledge [R2]*

[J5] 2010-03 - Director of research, head of the application characterization and optimization group, and one of co-founders *in Intel Exascale Lab (France)* preparing long-term R&D vision based on cTuning technology while building from scratch and leading team of 8 researchers, engineers and students

- *More info about this activity [I2]*

[J6] 2008-01 - Visiting scientist *in Institute of Computing Technology of Chinese Academy Of Sciences* preparing collaboration on extending cTuning and MILEPOST technology

- *Funded by ICT (China) [F6]*
- *Associated publications [Pub4, Pub10, Pub28, Pub30]*

[J7] 2007-09 - Tenured research scientist *in INRIA Saclay (France)* leading R&D in EU MILEPOST project on building practical machine learning based research compiler (MILEPOST GCC) and public plugin-based auto-tuning framework and repository of knowledge (cTuning.org)

- *Funded by EU MILEPOST project [F7]*
- *Associated job [J9]*
- *Associated public repository of knowledge [R4]*

- [J8]** 2007-09 - Adjunct professor at *University of Paris-Sud (France)* preparing and teaching my novel approach on collaborative, systematic and reproducible benchmarking, optimization and co-design of computer systems using public repository of knowledge, plugin-based auto-tuning, run-time adaptation, big data, predictive analytics, machine learning, data mining, statistical analysis, feature selection, crowdsourcing and collective intelligence
- *Promoting my new research and publication model where experimental results (tools, data, models) are continuously shared and validated by the community [Pub26, E12, E14]*
 - *Sharing all my related code and data at cTuning.org and later Collective Mind repository [R2]*
 - *Regularly giving invited lectures in the UK, USA, Canada, Russia and Asia*
- Sharing all notes online through cTuning foundation: <http://cTuning.org/lab/education>*
Enabling interactive publications: <http://c-mind.org/online-graph-for-interactive-articles>
-
- [J9]** 2006-07 - Technical leader of the *EU FP6 035307 MILEPOST project* directing development of the world's first practical machine learning based self-tuning compiler and of the first public repository of optimization knowledge to crowdsource optimization and co-design of computer systems
- *Associated publications [Pub5, Pub6, Pub10, Pub11, Pub28, Pub12, P24, Pub13, Pub30, Pub14, Pub15, Pub46]*
 - *Associated public repository of knowledge [R4]*
 - *IBM press-release [Pub59]*
-
- [J10]** 2005-12 - Postdoctoral researcher and principal investigator in *INRIA Saclay (France)* preparing foundations of crowdsourcing auto-tuning combined with machine learning and public repositories of knowledge
- *Prepared INRIA part in the EU MILEPOST project proposal on crowdsourcing auto-tuning*
 - *Associated publications [Pub32, Pub47, Pub33, Pub16, Pub17, Pub34, Pub18, P39, Pub19, Pub20, Pub21, Pub8, Pub22]*
-
- [J11]** 2004-10 - Visiting scientist in *INRIA Saclay (France)* leading R&D of a novel concept of statically enabling dynamic optimizations using multi-versioning and run-time adaptation
- *Funded by EU HiPEAC fellowship [A6]*
 - *Prepared base for crowdtuning*
 - *Associated publications [Pub23]*
-
- [J12]** 2002-01 - Research associate in the *University of Edinburgh (UK)* leading research and developing framework and repository for program online auto-tuning, polyhedral optimization, and machine learning-based software/hardware co-design and co-optimization
- *Associated publications [Pub21, Pub36, Pub1, Pub9, Pub48, Pub37, Pub56, Pub38]*
-
- [J13]** 2000-02 - Visiting scientist in *Paris South University and INRIA Saclay (France)* leading development of novel memory/CPU characterization technique via semantically non-equivalent binary patching
- *Associated publications [Pub1, Pub9, Pub56, Pub38, Pub24]*
 - *Extensions of this work [Pub50, Pub45]*
-
- [J14]** 1999-02 - Research assistant (research associate since 2000) in the *University of Edinburgh (UK)* leading R&D in 2 workpackages in the EU MHAOTEU project on program behavior analysis and auto-tuning for HPC systems
- *Associated publications [Pub1, Pub9, Pub56, Pub38, Pub57, Pub24, Pub58]*
-

[J15]	1999-02 cur.	-	<p>Evangelist of a collaborative and reproducible research and experimentation in computer engineering</p> <ul style="list-style-type: none"> • Developing public repositories of knowledge and common research and development tools [R7, R6, R5, R4, R2] • Enabling new publication model where experimental results (tools, data, models) are continuously shared and validated by the community [Pub26, E12, E14] • • Established international not-for-profit cTuning foundation [O3] <p>cTuning foundation activities on collaborative and reproducible research, development and experimentation in computer engineering: http://cTuning.org/reproducibility Public wiki on reproducible research: http://c-mind.org/reproducibility My motivation, history and manifesto: http://cTuning.org/history</p>
[J16]	1998-09 1999-01	-	<p>Programmer in the Laboratory for Computer Technologies in Teaching in Moscow Insitute of Physics and Technology (Russia) leading development of educational web-based software for undegraduate courses in computer engineering and machine learning</p>
[J17]	1997-09 1999-02	-	<p>Research assistant and principal investigator in the Institute for High-Performance Computing of the Russian Academy of Sciences (Russia) directing research on unifying remote access to high-performance computing systems as a web service and speeding up own neural network modelling software</p> <ul style="list-style-type: none"> • Associated thesis [Pub2]
[J18]	1994-01 1994-06	-	<p>Research assistant in Moscow Insitute of Physics and Technology (Russia) leading development of simulation and visualization software in a project "Computer simulation of non-linear wave processes in gaseous streams"</p>
[J19]	1994-01 1997-06	-	<p>Research assistant and principal investigator in Moscow Institute of Physics and Technology (Russia) directing research on developing a public research, development and experimentation toolset to design, model and optimize semiconductor neural networks as a practical step towards collaborative reverse engineering of a brain and development of a brain-inspired neuro-computer</p> <ul style="list-style-type: none"> • Project description [M10] • Associated publications [Pub41, Pub40, Pub39]
[J20]	1992-02 1993-06	-	<p>CTO in own Moscow-based startup developing and selling software for automation of financial reporting in companies</p>

6 INSTITUTION BUILDING (I)

#	Year	Description
---	------	-------------

- [I1] 2014-07 - **Established non-profit cTuning foundation** to continue community-driven development of the open-source cTuning technology (Collective Knowledge infrastructure and repository) as well as methodology for collaborative and reproducible research and experimentation in computer engineering
- *Continue supporting my open-source tools and repository for systematic, collaborative and reproducible research and experimentation in computer engineering ("big data" predictive analytics, machine learning, data mining, statistical analysis, autotuning, run-time adaptation)*
 - *Continue validating new publication model in computer engineering where all research artifacts (tools, benchmarks, datasets, models) are continuously shared and validated by the community [Pub26,E4,E5,E3,E1,E2]*
 - *Awards [A1]*
 - *Guest lectures [L1]*
 - *Funded by [F1]*
 - *Associated job [J2]*
 - *Associated publications [Pub25]*
 - *Associated events [E4,E5,E3,E1,E2]*
 - *Associated software [S1]*
 - *Associated public repository for collaborative and reproducible experimentation with interactive graphs and articles [R1]*

Website: <http://cTuning.org>

- [I2] 2010-03 - **Helped establish Intel Exascale Lab in France while preparing long term research and development agenda based on cTuning technology**
- 2011-08
- *On industrial leave from INRIA invited by Intel to help establish new Exascale Lab in France preparing long-term R&D vision based on cTuning technology [I3] while building from scratch and leading team of 8 researchers, engineers and students*
 - *Funded by [F3]*
 - *Associated job [J5]*
 - *My first team members [Q6, Q4, Q3]*
 - *Associated publications [Pub53, Pub43, Pub55, Pub45]*
 - *Associated events [E17, E18, E19]*
 - *Associated in-house software and repository [S6, S10]*

- [I3] 2008-01 - **Founded public, community-driven cTuning.org portal** to start collaborative systematization of analysis, benchmarking, optimization and co-design of computer systems using extensible public repositories of knowledge, plugin-based auto-tuning, run-time adaptation, crowdsourcing, big data, predictive analytics (machine learning, data mining, statistical analysis, feature detection) and collective intelligence while serving as a president and CTO
- cur.
- *Proposed and implemented idea of crowdsourcing program analysis, optimization and run-time adaptation using plugin-based repository and infrastructure for the EU MILEPOST project [F7] that was effectively used to train machine-learning based self-tuning compiler MILEPOST GCC [S9] and meta-optimizer cTuning CC [S8]*
 - *Initiated new publication model in computer engineering where all research artifacts (tools, benchmarks, datasets, models) are continuously shared and validated by the community [E12]*
 - *Awards [A2]*
 - *Keynotes [K2, K1]*
 - *Funded by [F7, F6, F5, F4, F2]*
 - *Associated job [J9, J4]*
 - *Associated publications [Pub50, Pub29, Pub51, Pub42, Pub53, Pub43, Pub4, Pub44, Pub5, Pub6, Pub45, Pub10, Pub11, Pub28, Pub12, P24, Pub13, Pub30, Pub14, Pub15, Pub7, Pub46, Pub31]*
 - *Associated events [E10, E12, E12, E16, E23, E25]*
 - *Associated software [S10, S2, S5, S9, S7, S8, S13, S11, S4, S12, S3]*
 - *Associated public repository of optimization knowledge [R2]*

Website: <http://cTuning.org>

7 EDITOR (ED)

#	Year	Type	Description
[Ed1]	2015-04	Special journal issue	Guest editors: Alex Jones (<i>University of Pittsburgh, USA</i>) and Grigori Fursin (<i>INRIA, France</i>). Special Issue on Reproducible Research Methodologies, IEEE Transactions on Emerging Topics in Computing (TETC) <i>IEEE TETC Website: http://www.computer.org/portal/web/tetc</i> <i>Flyer and CFP: http://www.computer.org/cms/Computer.org/transactions/cfps/cfp_tetcsi_rrm.pdf</i> <i>Related OCCAM project: http://www.occamportal.org</i> <i>Related Collective Mind project: http://c-mind.org</i> <i>Related Collective Mind repository: http://c-mind.org/repo</i>

8 KEYNOTES (K)

#	Year	Name
[K1]	2013-06	Keynote at iWAPT 2013 / ICCS 2013 in <i>Barcelona, Spain</i> ("Crowdsourcing autotuning: challenges and possible solutions") <i>Website: http://iwapt.org/2013</i>
[K2]	2013-03	Keynote at HPCS 2013 in <i>NTU, Taipei, Taiwan</i> ("Systematizing tuning of computer systems using crowdsourcing and statistics") <i>Website: http://goo.gl/emTlc</i>

9 SOCIAL ACTIVITIES (O)

#	Year	Description
[O1]	2014-11 - cur.	Initiator and co-chair of Artifact Evaluation for CGO and PPOPP conferences (together with Bruce Childers from the University of Pittsburgh, USA) <ul style="list-style-type: none">• <i>Since 2015, this initiative is supported by ACM</i> <i>AE PPOPP'16: http://cTuning.org/event/ae-ppopp2016</i> <i>AE CGO'16: http://cTuning.org/event/ae-cgo2016</i> <i>AE PPOPP'15: http://cTuning.org/event/ae-ppopp2015</i> <i>AE CGO'15: http://cTuning.org/event/ae-cgo2015</i>
[O2]	2009-09 - cur.	Supporter of Doctors Without Borders organization (MSF) <i>Website: http://www.msf.org</i>

[O3] 2008-01 - **Founder of the non-profit cTuning foundation**
2014-10

- *Non-profit cTuning foundation enables and promotes collaborative, reproducible and systematic research, experimentation and development in computer engineering. It develops and supports public repository of knowledge and related infrastructure. cTuning foundation also helps academic and industrial partners systematize, automate and speed up optimization, benchmarking and co-design of computer systems across all software and hardware layers (applications, compilers, run-time libraries, heterogeneous multi-core architectures) for minimal execution time, power consumption, failures and other costs while reducing time to market for new solutions. We were among the first researchers to convert this complex task into a unified big data problem and tackle it using open source Collective Mind/Collective Knowledge public repository of knowledge, common plugin-based auto-tuning infrastructure, run-time adaptation, machine learning, data mining, feature selection, crowdsourcing and collective intelligence*
- *Started in 2008 as an outcome of the EU MILEPOST project (2006-2009) [J9]*
- *Officially registered in France in 2014*

Website: <http://cTuning.org>

10 STARTUPS (C)

#	Year	Description
[C1]	2015-07 cur.	- CTO of dividiti, UK reusing my knowledge and experience in "big data" predictive analytics, data mining, autotuning and run-time adaptation
[C2]	1992-02 1993-06	- Founder and CTO of a small startup developing and selling software for automation of financial reporting in companies (used profits to fund my own research project on brain-inspired computing during undergraduate studies)

11 EXAMINER

- 2015-11 - **PhD examiner** for Luka Stanic (INRIA Grenoble, France) (expected)
- 2013-02 - **PhD examiner** for Ettore Speziale (Politecnico di Milano, Italy)
- 2013-02 - **PhD examiner** for Michele Tartara (Politecnico di Milano, Italy)
- 2013-02 - **PhD examiner** for Paolo Grassi (Politecnico di Milano, Italy)
- 2013-02 - **PhD examiner** for Simone Corbetta (Politecnico di Milano, Italy)

12 EXPERT SERVICE (E)

#	Year	Description
[X1]	2015	Invited as an international expert to unify and improve Artifact Evaluation across ACM conferences (based on my experience with public repositories of knowledge and Artifact Evaluation for PPOPP and CGO)
[X2]	2013	Invited as an international expert to review research proposals for the open grant competition of the Russian Federation to attract leading scientists to Russian universities with a total budget of around 200,000,000 euros
[X3]	2013	Invited to contribute to EU HiPEAC roadmap on advanced computing (2013 - 2020) <i>Online document: http://hipeac.net/roadmap</i>
[X4]	2012	Invited as an international expert to review research proposals for the open grant competition of the Russian Federation to attract leading scientists to Russian universities with a total budget of around 200,000,000 euros

[X5]	2012 - cur.	Consulting several major international IT companies (names are currently under NDA) to design faster and more power efficient production computer systems (software and hardware)
[X6]	2011	Invited as an international expert to prepare common EU-Russia IT call (related to GPGPU, programming models, performance and power tuning)
[X7]	2011	Invited as an international expert to review research proposals for the open grant competition of the Russian Federation to attract leading scientists to Russian universities with a total budget of around 300,000,000 euros
[X8]	2009	Invited to contribute to EU HiPEAC roadmap on advanced computing (2009 - 2020) long-term ideas extending cTuning and MILEPOST technology <i>Online document: http://hipeac.net/roadmap</i>
[X9]	2009	Consulting ARC (Synopsys) to apply cTuning and MILEPOST technology for multi-objective tuning (performance/code size/power) of customers' programs

13 MAJOR RESEARCH ACHIEVEMENTS (M)

#	Year	Description
[M1]	2014 - cur.	<p>Initiated Collective Knowledge Project aggregating all my past open-source cTuning research and developments for big data driven and cost aware computer engineering as a natural science</p> <ul style="list-style-type: none"> • <i>Continued new publication model in computer engineering [M3] where all research artifacts (tools, benchmarks, datasets, models) are continuously shared and validated by the community, and extended in new cM Lego-style R&D scenarios</i> • <i>Awards [A1]</i> • <i>Guest lectures [L1]</i> • <i>Partially funded by [F1]</i> • <i>Associated job [J2]</i> • <i>Associated publications [Pub25]</i> • <i>Associated events [E4,E5,E3,E1,E2]</i> • <i>Open-source Collective Knowledge Infrastructure [S1]</i> • <i>Public Collective Knowledge repository for collaborative and reproducible experimentation with interactive graphs and articles [R1]</i> • On-going, community-driven effort <p><i>Collective Knowledge live repository with interlinked code, data, experimental results, predictive models, interactive graphs and articles, etc.: http://cknowledge.org/repo</i> <i>Collective Knowledge Framework for collaborative, reproducible and cost-aware computer engineering: http://github.com/ctuning/ck</i></p>

[M2] 2011 - cur. **Developed theoretical Collective Mind foundations and supporting plugin-based infrastructure with public web-services to enable collaborative, systematic and reproducible analysis, design and optimization of adaptive computer systems based on extensible public repositories of knowledge, crowdsourcing, online tuning and machine learning, and to initiate new publication model with reproducible results where all research artifacts are continuously shared, validated and extended by the community**

- *Aggregates all my past ideas and concepts from [M10, M9, M8, M7, M6, M5, M4]*
- *Partially funded by INRIA 4 year fellowship [A2]*
- *Continued new publication model in computer engineering [M3] where all research artifacts (tools, benchmarks, datasets, models) are continuously shared and validated by the community, and extended in new cM Lego-style R&D scenarios*
- *Publications [Pub50, Pub29, Pub51, Pub42, Pub53, Pub43]*
- *Associated events [E12, E10, E14, E17]*
- *Collective Mind Framework (plugin-based knowledge management system) [S2]*
- *Collective Mind Public repository [R2]*
- **Collective Mind framework and repository discontinued for a much smaller and faster Collective Knowledge Infrastructure [M1]**

Collective Mind live repository: <http://c-mind.org/repo>

Online advice web service to predict optimizations based on features: http://c-mind.org/cm/predict_opts

Universal auto-tuning and learning pipeline for top-down multi-objective optimization: http://c-mind.org/cm/autotuning_pipeline

[M3] 2007 - cur. **Evangelized and pushed reproducible research in computer engineering to masses and initiate new publication model with reproducible results where all research artifacts (experimental pipelines, benchmarks, codelets, data sets, tools, models) are continuously shared, validated and extended by the community**

- *Since it was extremely difficult to persuade community which mainly focuses on publications to start sharing their research artifacts, in 2007, I developed and opened a cTuning.org web portal with a public repository, and shared all my past experimental results and artifacts to set up an example. Since then this approach was gradually picked up by the community and now being evaluated in major computer science conferences including CGO and PPOPP [E4, E5].*
- *Partially funded by [A2,F1]*
- *Awards [A1]*
- *Publications [Pub25, Pub50, Pub29, Pub51, Pub42, Pub53, Pub43]*
- *Associated events [E4,E5,E3,E1,E2,E12, E10, E14, E17]*
- *Open-source Collective Knowledge Infrastructure [S1]*
- *Public Collective Knowledge repository for collaborative and reproducible experimentation with interactive graphs and articles [R1]*
- *Open-source Collective Mind Framework (plugin-based knowledge management system) [S2]*
- *Public Collective Mind Public repository [R2]*
- **Our vision paper on community-driven reviewing of papers and artifacts [Pub26]**
- **On-going, community-driven effort backed up by ACM since 2015 ...**

CrowdTuning portal: <http://cTuning.org>

Collective Mind live repository: <http://c-mind.org/repo>

Online advice web service to predict optimizations based on features: http://c-mind.org/cm/predict_opts

Universal auto-tuning and learning pipeline for top-down multi-objective optimization: http://c-mind.org/cm/autotuning_pipeline

[M4] 2006 - cur. **Prepared theoretical foundations and led development of the first practical machine learning based self-tuning compiler (MILEPOST GCC and cTuning CC) and plugin-based multi-objective auto-tuning framework (execution time, code size, compilation time, power consumption or any other user defined metric) combined with collective participation of multiple users (cTuning.org)**

- *Aggregates my past ideas and concepts from [M10, M9, M8, M7, M6, M5]*
- *Funded by [F7, F4]*
- *Collaboration with IBM (Israel), University of Edinburgh (UK), ARC (now Synopsys, UK), CAPS Enterprise (France), and ICT (China)*
- *Initiated new publication model in computer engineering [M3] where all research artifacts (tools, benchmarks, datasets, models) are continuously shared and validated by the community within cTuning plugin-based framework [M2]*
- *Publications [Pub21, Pub20, Pub17, Pub16, Pub31, Pub46, Pub30, Pub13, Pub12, Pub11, Pub5]*
- *Associated events [E28, E26, E25, E24, E23, E21, E19]*
- *MILEPOST GCC software [S9]; most of technology is now available in mainline GCC and is being added to major commercial compilers*
- *cTuning CC software [S8]*
- *cTuning repository [R4]*
- **Considered by IBM to be the first practical machine-learning based compiler in the world (IBM press-release [Pub59])**
- **All benchmarks, datasets, tools, models and experimental results have been released to public for collaborative validation and extension!**

cTuning collaborative portal: <http://cTuning.org>

Online advice web service to predict optimizations based on features: <http://cTuning.org/cpredict>

[M5] 2004 - cur. **Developed concept of statistical collaborative optimization (collective optimization) that dramatically speeded up analysis and multi-objective optimization of computer systems by transparently distributing them among multiple users and combining with statistical analysis**

- *Includes ideas and concepts from [M9, M8, M7, M6]*
 - *Funded by [A6, F7, F4]*
 - *Publications [Pub2, Pub22, Pub8, Pub18, Pub14, P24, Pub10, Pub6]*
 - *Software [S16, S14, S13, S12, S11, S10]*
 - *Public repository [R4]*
-

[M6] 2004 - cur. **Developed concept of split-compilation to statically enable dynamic optimization and adaptation using code multi-versioning and low-overhead run-time adaptation as a reaction to program behavior, architecture changes, and dataset parameters**

- *Associated software [S13, S3]*
 - *Funded by [A6, F4, F5]*
 - *Supports predictive scheduling for heterogeneous architectures [Pub15] and adaptive libraries combined with automatically built run-time decision trees [Pub30]*
 - *Publications [Pub23, Pub34, Pub33, Pub14, P24, Pub28, Pub6, Pub50]*
 - *Highest ranked paper introducing technique [A5, Pub23]*
 - *Prepared for mainline GCC during Google Summer of Code program in 2009 [F4]*
 - *Finalized in mainline GCC 4.8 in 2012*
 - **Used and extended by Google**
-

[M7] 2004 - cur. **Developed concept of making hardwired tools and applications interactive using simple, event-based plugin framework**

- Supports GCC, LLVM, Open64, PathScale compilers; enables self-tuning applications [M6, Pub15, Pub50]
- Funded by [A6, A2, F7, F4, F2]
- Associated software [S16, S3, S4]
- Publications [Pub23, Pub8, P39, Pub15, Pub30, Pub28, Pub11, Pub5, Pub6, Pub50]
- **Now available in mainline GCC and extensively used by industry and academia world wide**

[M8] 1999 - cur. **Developed multiple unconventional interdisciplinary approaches for empirical program and architecture analysis, optimization and co-design through reactions to (possibly semantically non-equivalent) modifications (source or binary instruction, loop or thread level patching)**

- Used to quickly detect application CPU/memory bounds and performance bottlenecks without the need for slow simulators or possibly misleading hardware counters, or characterize programs and architectures similarity through a vector of reactions to canonical transformations to predict most profitable optimizations
- Funded by [F8, F7, F3]
- Software [S18, S4, S10, S2]
- Publications [Pub38, Pub56, Pub9, Pub1, Pub16, Pub14, Pub6]
- **Used in EU MHAOTEU project [J14], my PhD R&D [Z5], and in all later projects on program and architecture co-design, crowd-tuning and machine learning**
- **Adopted and extended in Intel Exascale Lab within DECAN tool (decremental analysis via binary patching) [Pub45]**

[M9] 1997 - 1999 **Developed concept of unified access to HPC resources for non-specialists as a simple web service**

- Used to simplify deployment of my parallel (MPI-based) neural network modelling software on distributed and varying HPC resources
- Partially funded by [A9, J20]
- Software [S19, S20]
- Publications [Pub2, Pub39, Pub40, Pub41]
- Repository [R7]
- **Used in EU MHAOTEU project [J14], my PhD R&D [Z5] and later work on cTuning.org and crowd-tuning**

[M10] 1993 - 1999 **Developed prototype of a public research, development and experimentation toolset to design, model and optimize semiconductor neural networks as a practical step towards collaborative reverse engineering of a brain and development of a brain-inspired neuro-computer**

- Partially funded by [A9, J20]
- Hardware: [H1] (own ADC/DAC board for automatic measurement of characteristics of semiconductor (neural) devices)
- Software [S20]
- Publications [Pub39, Pub40, Pub41]
- Repository [R7]
- **Since modeling was already too slow, unreliable and power hungry, it forced me to switch most of the following R&D effort to make program and architecture optimization practical and combine it with statistical analysis and machine learning using interdisciplinary background**
- **Ideas from this project was reused in all my next interdisciplinary projects on program and architecture characterization, optimization and co-design using machine learning and crowd-sourcing**

14 PUBLIC OR IN-HOUSE REPOSITORIES OF KNOWLEDGE (R)

#	Year	Description
---	------	-------------

[R1] 2014-11 - **Collective Knowledge public repository (CK aka cTuning4)** to continue improving whole experimental setup sharing (code, data, dependencies, experimental results, models) along with interactive articles

- *Developed by the non-profit cTuning foundation*
- *Opened to public in 2015*
- *Received award [A1]*
- *Included all past and current semantically connected research artifacts from my research, development and experimentation (hundreds of codelets and benchmarks; thousands of datasets; GCC, LLVM, Open64, PathScale, ICC compiler optimization description; tools and scripts; online tuning plugins; machine learning plugins; adaptive exploration plugins; graph plotting plugins; data mining plugins; machine learning based meta compiler; MILEPOST GCC, etc*
- *Supports our initiatives on Artifact Evaluation and new publication models where results and papers are validated and improved by the community [M3]*
- *Partially funded by [F1]*
- *Powered by Collective Knowledge Framework [S1]*

Collective Knowledge live repository: <http://cknowledge.org/repo>

Examples of interactive graphs: <http://cknowledge.org/repo/web.php?wcid=graph>:

Examples of interactive reports: <http://cknowledge.org/repo/web.php?wcid=report>:

[R2] 2011-09 - **Collective Mind public repository (cM aka cTuning3)** to start collaborative systematization of analysis, design and optimization of computer systems based on extensible public repositories of knowledge, crowdsourcing, online tuning and machine learning, and to initiate new publication model where all research artifacts are continuously shared, validated and extended by the community

- *Opened to public in 2013*
- *Included all past and current semantically connected research artifacts from my research, development and experimentation (hundreds of codelets and benchmarks; thousands of datasets; GCC, LLVM, Open64, PathScale, ICC compiler optimization description; tools and scripts; online tuning plugins; machine learning plugins; adaptive exploration plugins; graph plotting plugins; data mining plugins; machine learning based meta compiler; MILEPOST GCC, etc*
- *Connected with Android Collective Mind Node [S5] to crowdsource program and architecture characterization and multi-objective autotuning (execution time, code size, compilation time, power consumption) using any available Android-based mobile phone, tablet or laptop*
- *Used for the new publication model [E12]*
- *Funded by [A2, F2]*
- *Powered by Collective Mind Framework [S2]*

Collective Mind live repository: <http://c-mind.org/repo>

Online advice web service to predict optimizations based on features: http://c-mind.org/cm/predict_opts

Universal autotuning and learning pipeline for top-down multi-objective optimization: http://c-mind.org/cm/autotuning_pipeline

[R3] 2010-03 - **In-house Codelet Tuning Repository for Intel Exascale Lab (aka cTuning2)** to decompose large applications into codelets for continuous characterization and tuning

- *Developed with my team [Q6, Q4, Q3] as customizable repository for Intel Exascale Lab, CEA, GENCI and UVSQ (France)*
 - *Funded by [F3]*
 - *Powered by Codelet Tuning Infrastructure [S6]*
 - ***Discontinued for [R2]***
-

- [R4] 2006-01 - **cTuning.org public repository (aka cTuning1)** to start collaborative systematization of analysis, design and optimization of computer systems based on extensible public repositories of knowledge, crowdsourcing, online tuning and machine learning
- *Opened to public in 2008*
 - *Included past research experimental results on program and architecture multi-objective tuning (execution time, code size, compilation time, power consumption) for reproducibility and collaborative extensions*
 - *Connected with MILEPOST GCC [S9] for continuous and online training and improvement of the prediction models*
 - *Funded by [F7, F6, F5, F4, F2]*
 - *Powered by cTuning Framework [S10]*
 - **Gradually being discontinued for [R2]**
- cTuning live repository: <http://ctuning.org/wiki/index.php/CDatabase>
 Online advice web service to predict optimizations based on features: <http://cTuning.org/cpredict>*

- [R5] 2004-06 - **In-house collaborative optimization repository** for research on multi-objective program and architecture autotuning and co-design combined with machine learning
- *Powered by FCO framework [S14]*
 - *Funded by [A6, F6]*
 - **Discontinued for [R4]**

- [R6] 1999-02 - **In-house collaborative optimization repository** for research on multi-objective program and architecture characterization, optimization and co-design with first experiments on predictive modeling
- *Powered by EOS framework [S18]*
 - *Developed and used in the EU FP5 MHAOTEU project [J14]*
 - *Funded by [F8, A6]*
 - **Discontinued for [R5]**

- [R7] 1993-03 - **In-house Experimental Repository** for research, development and experimentation on novel, semi-conductor neural networks, and on providing unified access to HPC resources as a web service
- *Powered by SCS framework [S19]*
 - *Partially funded by [A9, J20]*
 - **Discontinued for [R6]**

15 AWARDS, PRIZES AND FELLOWSHIPS (A)

#	Year	Description
[A1]	2014-12	HiPEAC award for transferring cTuning technology to ARM to systematize benchmarking and combining it with "big data" predictive analytics
[A2]	2012-04 - 2014-10	INRIA award and fellowship for "making an outstanding contribution to research" for making program and architecture optimization and co-design more practical, systematic and reproducible by combining cTuning plugin-based autotuning technology with statistical analysis, machine learning and community-driven curation
[A3]	2010-11	HiPEAC award for paper "Evaluating Iterative Optimization across 1000 Data Sets" [Pub10], PLDI 2010, Canada
[A4]	2009-12	HiPEAC award for paper "Portable Compiler Optimization Across Embedded Programs and Microarchitectures using Machine Learning" [Pub12], MICRO 2009, NY, USA

[A5]	2005-01	Highest ranked paper "A practical method for quickly evaluating program optimizations" [Pub23] at HiPEAC 2005, Barcelona, Spain
[A6]	2004-10 2005-11	- EU HiPEAC Fellowship to collaborate with INRIA Saclay (France)
[A7]	2000-01 2001-12	- International Overseas Research Student Award (fellowship) for PhD research from the UK government
[A8]	1999-06	Golden Medal for MS studies and thesis from Moscow Institute of Physics and Technology (Russia)
[A9]	1996-08 1998-08	- International Soros Science Education Program fellowship "in recognition and appreciation of outstanding achievements in the study of science at the university level" (300 students in Russia)

16 MAJOR FUNDING (F)

#	Year	Description
[F1]	2014-11 2015-04	- EU FP7 609491 TETRACOM funding to validate open-source cTuning technology in ARM (open-source infrastructure and repository for collaborative and reproducible autotuning and "big data" predictive analytics) <i>Developed Collective Knowledge Infrastructure and Repository (by non-profit cTuning foundation):</i> http://github.com/ctuning/ck
[F2]	2013-07 2013-10	- HiPEAC industrial internship funding for Abdul Memon to validate PhD results in STMicroelectronics (France) working on "Auto Tuning Optimization System Acceleration for Embedded Linux Stacks" <i>HiPEAC description:</i> http://www.hipeac.net/content/auto-tuning-optimization-system-acceleration
[F3]	2010-04 2011-08	- Funding from Intel and CEA to invited to help establish new Intel Exascale Lab in France based on cTuning technology, serve as a head of application characterization and optimization group, and direct research and development
[F4]	2009-06 2009-08	- Funding from Google (GSOC program) to move cTuning and MILEPOST technology to mainline GCC (Interactive Compilation Interface and code multiversioning to make statically compiled programs adaptable at run-time) for 2 students from ICT (China) <i>GSOC description page for adding and extending plugin-based framework (ICI) to mainline GCC:</i> http://cTuning.org/gsoc2009-yjhuang <i>GSOC description page for statically enabling dynamic optimizations in GCC through code cloning:</i> http://cTuning.org/gsoc2009-lpeng
[F5]	2008-03 2008-05	- EU HiPEAC funding for Victor Jimenez (PhD student from UPC, Spain) to visit My group in INRIA Saclay (France) <ul style="list-style-type: none"> • <i>Extending my previous work on statically enabling dynamic optimization [Pub23] with predictive scheduling for heterogeneous architectures [Pub15]</i>
[F6]	2008-01 2008-04	- Funding from ICT (China) for students and faculty exchange to extend cTuning and MILEPOST technology
[F7]	2006-09 2009-09	- Funding from EU FP6 MILEPOST project to develop practical machine learning based self-tuning compiler <ul style="list-style-type: none"> • <i>I was one of the initiators of this project responsible for INRIA part - development of a repository of knowledge, common plugin-based auto-tuning infrastructure, statistical analysis, predictive modeling, GCC plugin-based framework</i>

[F8] 1999-02 - **Funding from EU FP5 MHAOTEU project** to develop tools for memory hierarchy optimization for High-
2001-12 Performance Computer Systems

- *I was responsible for University of Edinburgh part developing polyhedral source-to-source compiler, plugin-based distributed analysis and auto-tuning infrastructure and a repository of knowledge while reusing some of my M.S. developments for unifying access to supercomputers through the web*

17 MAJOR SOFTWARE AND DATASETS (S)

- Description of current cTuning tools developments: <http://cTuning.org/ctools>
- Older collaborative Wiki for cTuning tools: <http://cTuning.org/wiki/index.php/CTools>
- Older detailed description of some past software developments: http://fursin.net/research_dev
- Description of some hobby software developments: http://fursin.net/hobbies_research_soft

#	Year	Description
[S1]	2014-11 - cur.	<p>Collective Knowledge Framework and Repository (CK aka cTuning4, BSD-license) - Collective Knowledge (CK) is a light-weight, portable, modular and python-based framework, repository, web service and SDK to organize, describe, cross-link and share user code, data, experimental setups and meta information as unified and reusable components with JSON API via standard Git services (such as GITHUB or BitBucket).</p> <ul style="list-style-type: none">• <i>Opened to public in 2015</i>• <i>Pre-released in May, 2015 (V1.2, permissive and simplified BSD license)</i>• <i>Partially funded by EU FP7 TETRACOM 6-months grant [A2]</i>• <i>Supports our new publication model in computer engineering where all research artifacts (tools, benchmarks, datasets, models) are continuously shared and validated by the community [Pub26,E4,E5,E3,E1,E2]</i>• <i>Awards [A1]</i>• <i>Guest lectures [L1]</i>• <i>Associated publications [Pub25]</i>• <i>Associated events [E4,E5,E3,E1,E2]</i>• <i>Associated live repository [R1]</i> <p>Website: http://github.com/ctuning/ck Wiki: http://github.com/ctuning/ck/wiki cM Google Group discussions: http://groups.google.com/group/collective-mind</p>

[S2] 2011-09 - **Collective Mind Framework and Repository (cM aka cTuning3)** - plugin-based knowledge management system to preserve, systematize and share all research, development and experimentation artifacts using private or in-house web and plugin-based, customizable, schema-free, NoSQL repository of knowledge combined with crowdsourcing and machine learning; collaborative and agile implementation and systematization of experimental scenarios combined with statistical analysis and data mining; plugin-based program and architecture autotuning and co-design combined with crowdsourcing, machine learning and run-time adaptation

- Released in 2013 (V1.0beta, standard BSD license)
- Partially funded by INRIA 4 year fellowship [A2]
- Includes software [S3, S4, S5]
- Aggregated and unified all my past research and development ideas and prototypes within new Collective Mind Framework and Repository to systematize collaborative research, development and experimentation
- Continued new publication model in computer engineering where all research artifacts (tools, benchmarks, datasets, models) are continuously shared and validated by the community, and extended in new cM Lego-style R&D scenarios
- Technology: easily customizable for any end-user R&D scenario through plugins; agile development methodology; NoSQL databases with JSON representation; ElasticSearch indexing; unified python plugins for web-services, autotuning, data mining and machine learning; OpenME interface to connect applications and tools written in C, C++, Fortran, PHP to cM; support for practically any Unix and Windows-based desktops, laptops, supercomputers, cloud servers, and even tablets and mobile phones with ARM, Intel, ST, Loongson, AMD, NVidia and other chips; powerful graph capabilities
- Publications [Pub50, Pub29, Pub51, Pub42, Pub53, Pub43]
- Associated events [E12, E10, E14, E17]
- Associated live repository and online advice service [R2]

Website: <http://cTuning.org/tools/cm>

cM SVN download: <http://sourceforge.net/p/c-mind>

cM wiki: <http://ctuning.org/cm/wiki>

New cM development tasks: <https://sourceforge.net/p/c-mind/tickets>

Old cM development tasks: <http://unidapt.org/redmine/projects/cm>

cM Google Group discussions: <http://groups.google.com/group/collective-mind>

[S3] 2011-09 - **OpenME event-based plugin framework and unified interface** to making rigid, hardwired applications and tools interactive, tunable and adaptive

- Released in 2013
- Technology: C, C++, Fortran, PHP, python event-based plugin framework
- Included plugins to open up latest GCC and LLVM for fine-grain program analysis and autotuning, and to open up HPC applications for online tuning and adaptation on heterogeneous architectures [Pub50, Pub15]
- Current version: 1.0beta (LGPL v2 license)
- Associated publications [Pub50, Pub29, Pub51, Pub42, Pub53, Pub43]

Website: <http://cTuning.org/tools/cm>

[S4] 2011-09 - **Alchemist plugin** for fine-grain program feature extraction, decremental analysis, and optimization in GCC

- Collaboration with STMicroelectronics [Q7, F2]
- Pre-released in 2013
- Current version: 1.0beta (GPL v2 license)
- Technology: GCC dynamic plugin written in C
- Associated publications [Pub50, Pub29, Pub51, Pub42, Pub53, Pub43]

Website: <http://cTuning.org/tools/cm>

[S5] 2011-09 - **Android Collective Mind Node** to crowdsource program and architecture characterization and multi-objective autotuning (execution time, code size, compilation time, power consumption) using any available Android-based mobile phone, tablet or laptop

- Released in 2013
- Connected to public Collective Mind Repository [R2]
- Associated publications [Pub50, Pub6, P24, Pub30]

Google Play Website: https://play.google.com/store/apps/details?id=com.collective_mind.node

[S6] 2010-03 - **Colaborative Codelet Tuning Infrastructure (CTI aka cTuning2)** to decompose large applications into codelets for continuous characterization and tuning

- Collaboration with Intel, CEA, GENCI and UVSQ
- Developed the concept based on my cTuning technology and developed first prototype with my team [Q6, Q4, Q3] as Intel Exascale Lab in-house autotuning infrastructure and repository extending cTuning1 framework and repository
- Funded by [F3]
- Associated job [J5]
- More info about this activity [I2]
- Technology: mixed MySQL and NoSQL database; customizable python, PHP and C plugin and web-based infrastructure
- Availability: private use - after 2012 some parts developed by my team members became available under GPL v3 license
- **Discontinued for [S2]**

[S7] 2010-01 - **KDatasets** to Multiple datasets for cBench [S11] (1000 per benchmark)

- Released in 2010
- Last version: 1.0 (GPL v2 license)
- Funded by [F6]
- **Partially merged with Collective Mind repository [R2]**

Website: <http://kdatasets.appspot.com>

[S8] 2009-01 - **cTuning CC - machine learning based self-tuning meta-optimizer for any compiler including GCC, LLVM, ICC, etc.**

- Released in 2010
- Last version: 2.5 (GPL v2 license)
- **Moved to Collective Mind Framework as a ctuning.compiler plugin [S2]**
- IBM press-release [Pub59]

Website: <http://cTuning.org/ctuning-cc>

- [S9] 2006-07 - **MILEPOST GCC - machine learning based self-tuning compiler**
2009-06
- Released in 2009
 - Last version: 2.5 (GPL v2 license)
 - **Merged with cTuning CC [S8]**
 - Funded by [F7]
 - Collaboration with IBM (Israel), University of Edinburgh (UK), ARC (now Synopsys, UK), CAPS Enterprise (France), and ICT (China)
 - Developed within EU FP6 MILEPOST project [J9]
 - More info about this activity [I3]
 - Associated live repository and online advice service [R4]
 - **Considered by IBM to be the first practical machine-learning based compiler in the world (IBM press-release [Pub59])**

Website: <http://cTuning.org/milepost-gcc>

- [S10] 2006-01 - **cTuning and CCC (Continuous Collective Compilation) frameworks** for collaborative user-defined program and architecture characterization, multi-objective optimization (execution time, code size, compilation time, power consumption) and co-design of computer systems using public repositories of knowledge, plugin-based autotuning, machine learning and crowdsourcing
2010-03
- Released in 2008
 - Current version: 2.5 (GPL v2 license)
 - Collaboration with IBM (Israel), University of Edinburgh (UK), ARC (now Synopsys, UK), CAPS Enterprise (France), and ICT (China)
 - Developed within EU FP6 MILEPOST project [J9]
 - More info about this activity [I3]
 - Technology: mixed MySQL and NoSQL database; customizable PHP,C,C++ plugin and web-based infrastructure
 - **Discontinued for [S2]**
 - IBM press-release [Pub59]

Development website: <https://sourceforge.net/projects/cccpf/files/CCC/V2.5-beta>
Tools: <http://cTuning.org/tools>

- [S11] 2006-01 - **cBench (Collective Benchmark)** to unify and systematize benchmarking with multiple datasets for realistic and collaborative program and architecture autotuning and co-design combined with machine learning; unify training and tuning of MILEPOST GCC []; enable research on split compilation []
2010-03
- Released in 2008
 - Last version: 1.1 (GPL v2 license)
 - **Discontinued and merged with Collective Mind repository [R2]**

Website: <http://ctuning.org/wiki/index.php/CTools:CBench>

- [S12] 2006-01 - **cDatasets (Collective Datasets)** to multiple datasets for cBench [S11] (20..100 per benchmark)
2010-03
- Released in 2008
 - Last version: 1.1 (GPL v2 license)
 - **Discontinued and merged with Collective Mind repository [R2]**

Website: <https://sourceforge.net/projects/cbenchmark/files/cDatasets>

- [S13] 2004-10 - **UNIDAPT - universal plugin-based program run-time adaptation framework**
2010-03
- *Software support for my technique for split compilation (to dynamic optimization and adaptation for statically compiled programs using multi-versioning and light run-time adaptation mechanisms) [Pub23, Pub33]*
 - *Supports predictive scheduling for heterogeneous architectures [Pub15] and adaptive libraries combined with automatically built run-time decision trees [Pub30]*
 - *Funded by [A6, F4]*
 - *Publications [Pub23, Pub15, Pub30, Pub34, Pub33, Pub14, P24, Pub28, Pub6, Pub50]*
 - *Finalized in mainline GCC 4.8 in 2012*
 - *Prepared for mainline GCC during Google Summer of Code program in 2009 [F4]*
 - **Redesigned and unified in [S3]**
- Website: <http://sourceforge.net/projects/unidapt>
GCC multi-versioning description: <http://gcc.gnu.org/wiki/FunctionMultiVersioning>
-
- [S14] 2004-06 - **Framework for Continuous Optimization (FCO)**
2006-06
- *Released in 2006*
 - *Licence: GPL v2*
 - *Collaboration with Institute of Computing Technology (China) to tune applications and compilers for LoongSon and Godson processors [F6, J6]*
 - *Discontinued for S10*
- Old description: <http://fursin.net/wiki/index.php5?title=Research:Developments:FCO>
Framework with ICI for Open64 compiler: <http://open64-ici.cvs.sourceforge.net/open64-ici>
Framework with ICI for PathScale compiler: <http://pathscale-ici.cvs.sourceforge.net/pathscale-ici>
-
- [S15] 2004-06 - **Code, data and experiment sharing tool** to decentralize collection of huge amount of experiments
2004-12 during tuning GCC optimization heuristic
- *I started implementing a plugin framework and a plugin for GCC to transparently tune optimization heuristic and embed special function to collect run-time info about behavior of real programs and data sets in real environments from multiple users. Originally, I connected my framework to MySQL database, but it could not cope with huge amount of data. Therefore, I tried to implement another solution - sharing best optimizations/speedups via P2P networks such as overnet, bittorrent and edonkey. Unfortunately, the solution became rather complex and unstable due to a lack of stable and universally acceptable P2P tools. So, I had to move back to MySQL database in the EU FP6 MILEPOST project [S9]. However, I later implemented my own P2P sharing mechanism with noSQL Hadoop-based repository in [S1, S2] while investigating capabilities of new third-party P2P tools.*
-
- [S16] 2004-06 - **Interactive Compilation Interface (ICI)** to open up production compilers (GCC, Open64, PathScale, etc.) through light-weight event-based plugin framework and transform them into powerful interactive
2009-06 research toolsets
- *Publications [Pub23, Pub8, P39, Pub15, Pub30, Pub28, Pub11, Pub5, Pub6, Pub50]*
 - *Released in 2006; Merged with mainline GCC in 2009 [F4]*
 - *Last version: 2.5 (GPL v2 license)*
 - **Redesigned and unified in [S3]**
- Website: <http://cTuning.org/ici>
-

- [S17] 1999-02 - **Source-to-source polyhedral transformation server** - source-to-source C and Fortran polyhedral transformation server based on MARS compiler
2006-06
- *First released in 2001*
 - *Developed during EU FP5 MHAOTEU project [F4] and PhD studies [Z5]*
 - *Last version: V1.15i (GPL v2 license)*
 - *Used in first experiments to predict complex optimizations (not just 1 transformation) using machine learning and program semantic and dynamic features (hardware counters)*
 - *Discontinued for compilers with Interactive Compilation Interface or OpenME [S14, S16, S3]*
- Description: <http://fursin.net/wiki/index.php5?title=Research:Developments:EOS>*
-
- [S18] 1999-02 - **Edinburgh Optimizing Software (EOS)** - plugin-based client-server program and architecture characterization and autotuning framework
2006-06
- *Released in 2001*
 - *Uses source-to-source C and Fortran polyhedral transformation server based on MARS compiler [S17]*
 - *Developed during EU FP5 MHAOTEU project [F4] and PhD studies [Z5]*
 - *Last version: V2.2 (GPL v2 license)*
 - *Technology: NoSQL based database; java and C plugins; socket communication between modules; java based GUI*
 - *Included plugins for program memory/CPU characterization through semantically non-equivalent assembler/binary patching [Pub38, Pub56, Pub9, Pub1, Pub45]; own source-to-source compiler; fine-grain autotuning plugins (unrolling, array padding) with partial polyhedral optimization support (tiling, fusion/fission, vectorization)*
 - *Used in first experiments to predict complex optimizations (not just 1 transformation) using machine learning and program semantic and dynamic features (hardware counters)*
 - *Discontinued for S14 and later for S10*
- Description: <http://fursin.net/wiki/index.php5?title=Research:Developments:EOS>*
-
- [S19] 1997-06 - **SuperComputer Service (SCS)** - framework to provide and unify remote access to high-performance computing systems for non-specialists as a simple web service
1999-02
- *Released in 1999*
 - *Last version: V1.3 (GPL license)*
 - *Technology: MySQL based database; java, perl, C, Visual Basic, Visual C modules; standard http/ftp communication; web-based GUI*
 - *Used to simplify execution of my neural network modelling software on distributed and varying HPC resources*
 - *Publications [Pub2, Pub39, Pub40, Pub41]*
 - *Repository [R7]*
 - *Partially funded by [A9, J20]*
 - *Discontinued for S18*
-

- [S20] 1993-02 - **Semiconductor Brain** - semiconductor and modelled neural networks with my own ADC/DAC PC board and analysis software
- Released in 1997
 - Last version: V2.1 (GPL license)
 - Works with special hardware (ADC/DAC board): [H1]
 - Validated by improving recognition and restoration of characters by neural network in noised environments
 - Technology: Visual Basic and assembler; PSpice analog circuit and digital logic simulation software; MPI for HPC
 - Partially funded by [A9, J20]
 - Publications [Pub39, Pub40, Pub41]
 - Repository [R7]
 - Due to very slow simulation and limitation of semiconductor technology decided to switch to program and architecture optimization to eventually enable fast and low-power neural networks and machine learning
-

- [S21] 1991-08 - **Productivity tool** to automatically pack files to remote and removable disks with limited space using fast and random strategy instead of greedy one
- My first autotuning experience - at that time, we had removable disks with only 360Kb, and it was extremely challenging to archive files on multiple disks. As a practical solution, I created a tool that checked free space on all disks and size of all files to randomly find the best packing strategy (within amount of time specified by a user). Interestingly, it worked as good as a greedy algorithm but about an order of magnitude faster which was important on very slow computers of that time. This was my first successful autotuning experience that I reused in all my further research.
-

18 HARDWARE (H)

#	Year	Description
[H1]	1993-02 - 1999-02	ADC/DAC board for personal computer - automating measurement of characteristics of semiconductor devices <ul style="list-style-type: none"> • Released to MIPT colleagues in 1997 • Related DOS/Windows software: [S20] • Used in various undergraduate MIPT laboratories • Partially funded by [A9, J20] • Publications [Pub39, Pub40, Pub41]

19 TALKS (T)

- [T1] 2015-07 - **ACM Workshop on Artifact Evaluation, New York, USA** (Our Artifact Evaluation Experience with PPOPP and CGO)
- [T2] 2015-06 - **Facebook, New York, USA** (Collective Knowledge Project: enabling more efficient computing via social networking and machine learning)
- [T3] 2015-06 - **PyData'15, Bloomberg, London, UK** (Collective Knowledge: python and scikit-learn based open research SDK for collaborative data management and exchange)
- [T4] 2015-03 - **University of Copenhagen, Denmark** (Collective Knowledge Project: Towards Collaborative and Reproducible Computer Engineering)
- [T5] 2015-02 - **CGO/PPOPP'15, San Francisco, CA, USA** (our Artifact Evaluation initiative, experience and outcome)
SlideShare: <http://www.slideshare.net/GrigoriFursin/presentation-fursin-aecgoppopp2015>
AE CGO'15: <http://cTuning.org/event/ae-cgo2015>
AE PPOPP'15: <http://cTuning.org/event/ae-ppopp2015>

- **[T6]** 2015-01 - **TETRACOM workshop @ HiPEAC'15, Amsterdam, Netherlands** (our TETRACOM Collective Knowledge project with ARM)
- **[T7]** 2015-01 - **CPC'15, London, UK** (Collective Mind usage scenarios: holistic and machine-learning based approach to auto-tuning ad run-time adaptation)
- **[T8]** 2015-01 - **ADAPT'15, Amsterdam, Netherlands** (our new publication model proposal via Reddit and Slashdot)
- **[T9]** 2014-06 - **ACM SIGPLAN TRUST'14 panel @ PLDI'14, Edinburgh, UK** (new open publication model proposal)
SlideShare: <http://www.slideshare.net/GrigoriFursin/panel-at-acmsigplantrust2014>
- **[T10]** 2014-06 - **Joint UIUC-ANL-INRIA workshop, Nice, France** (Collective Mind: bringing reproducible research to the masses)
- **[T11]** 2014-05 - **University of Tennessee, USA** (Collective Mind: community-driven systematization and automation of program optimization)
- **[T12]** 2014-04 - **INRIA Grenoble, France** (Collective Mind Technology: towards collaborative, systematic and reproducible auto-tuning)
- **[T13]** 2014-04 - **SEA 2014, UCAR, Boulder, USA** (Collective Mind: a collaborative curation tool for program optimization)
SlideShare: <http://www.slideshare.net/GrigoriFursin/presentation-fursin>
- **[T14]** 2014-03 - **Collaborations Workshop 2014, Oxford, UK** (Collective Mind Framework: systematizing and crowdsourcing multi-objective auto-tuning)
- **[T15]** 2013-12 - **INRIA / Paris South University TAO group, Orsay, France** (Collective Mind Framework: systematizing and crowdsourcing multi-objective auto-tuning)
- **[T16]** 2013-11 - **10th ANR-INRIA-UIUC workshop, Urbana, IL, USA** (Crowdtuning and public Collective Mind repository of knowledge)
- **[T17]** 2013-11 - **Google, Paris, France** (Collective Mind Framework: systematizing and crowdsourcing multi-objective auto-tuning)
- **[T18]** 2013-10 - **Dagstuhl seminar on auto-tuning for HPC, Germany** (Crowdsourcing analysis, tuning and learning of computer systems)
- **[T19]** 2013-09 - **INRIA Lille (ADAM group), France** (Collective Mind: plugin based autotuning and machine learning; collaborative and reproducible research and development)
- **[T20]** 2013-09 - **Minisymposium on Auto-tuning at PARCO 2013, Munich, Germany** (Crowdtuning: systematizing auto-tuning using predictive modeling and crowdsourcing)
- **[T21]** 2013-06 - **iWAPT 2013 @ ICCS 2013, Barcelona, Spain** (Crowdsourcing auto-tuning: challenges and possible solutions)
- **[T22]** 2013-05 - **ACM ECRC 2013/HiPEAC computing week 2013, Paris, France** (Thematic session: making computer engineering a science - preserving whole experimental pipeline and all research artifacts for reproducible research and development)
- **[T23]** 2013-03 - **HPSC 2013, Taipei, Taiwan** (Collective Mind: novel methodology, framework and repository to crowd-source auto-tuning)
SlideShare 1: <http://www.slideshare.net/GrigoriFursin/presentation-fursin-hpsc2013fursin1>
SlideShare 2: <http://www.slideshare.net/GrigoriFursin/presentation-fursin-hpsc2013fursin2>
- **[T24]** 2012-12 - **STMICROELECTRONICS, Grenoble, France** (Collective Mind (cTuning3) infrastructure and repository: universal, customizable, collaborative, and systematic experimentation on design and optimization of computer systems)
- **[T25]** 2012-10 - **IMEC, Ghent, Belgium** (Collective Tuning: systematization of analysis and optimization of computer systems)
- **[T26]** 2012-07 - **ARM, Cambridge, UK** (Collective Mind (cTuning3) infrastructure and repository: universal, customizable, collaborative, and systematic experimentation on design and optimization of computer systems)
- **[T27]** 2012-05 - **PARALLELS, Moscow, Russia** (Collective Mind Initiative - plugin-based framework and repository to systematize, automate and speed up optimization and co-design of computer systems)
- **[T28]** 2012-04 - **HiPEAC computing week, Goteborg, Sweden** (Collaborative and reproducible characterization, optimization and design of computer systems)
- **[T29]** 2012-04 - **ASEArch 2012, Oxford, UK** (Auto-tuning tools: collective approach)
- **[T30]** 2012-03 - **INRIA review commission, Paris, France** (Collective Mind Initiative - plugin-based framework and repository to systematize, automate and speed up optimization and co-design of computer systems)
- **[T31]** 2012-03 - **EPCC, Edinburgh, UK** (Collective Mind Initiative - plugin-based framework and repository to systematize, automate and speed up optimization and co-design of computer systems)
- **[T32]** 2012-03 - **ARM, Cambridge, UK** (Foundations of collective tuning)
- **[T33]** 2012-03 - **EXADAPT panel @ ASPLOS 2012, London, UK** (Collaborative research methodology for large-scale computer systems)
- **[T34]** 2012-02 - **SEA 2012, UCAR, Boulder, CO, USA** (Collaborative application characterization and optimization)

- **[T35]** 2012-02 - **STMicroelectronics, Grenoble, France** (Collective Mind Initiative - plugin-based framework and repository to systematize, automate and speed up optimization and co-design of computer systems)
- **[T36]** 2011-10 - **Visit of US Department of Energy delegation (Dan Hitchcock) to INRIA, Paris, France** (Collective Mind Initiative - plugin-based framework and repository to systematize, automate and accelerate optimization and co-design of computing systems)
- **[T37]** 2010-12 - **Intel Exascale Lab, Versailles, France** (Codelet Tuning Infrastructure aka cTuning2; plugin-based auto-tuning; collaborative program characterization and optimization combined with predictive modeling)
- **[T38]** 2010-01 - **HiPEAC speedup tutorial, Pisa, Italy** (cTuning.org:: Collective Tuning Initiative and public repository of optimization knowledge for reproducible research)
- **[T39]** 2009-06 - **GCC Summit 2009, Montreal, Canada** (cTuning.org:: Collective Tuning Initiative - plugin-based framework and repository to systematize, automate and accelerate optimization of computing systems)
- **[T40]** 2009-06 - **Infion, Munich, Germany** (cTuning.org:: Collective Tuning Initiative - plugin-based framework and repository to systematize, automate and accelerate optimization of computing systems)
- **[T41]** 2009-06 - **EU Commission, Munich, Germany** (cTuning.org and MILEPOST GCC for HiPEAC1 network of excellence final review)
- **[T42]** 2009-05 - **UVSQ, France** (cTuning.org public repository of optimization knowledge; interactive compilation; machine-learning based compilers; decremental CPU/memory characterization; split compilation through static multi-versioning and run-time adaptation)
- **[T43]** 2009-04 - **University of Illinois at Urbana Champaign and Intel Illinois, IL, USA** (cTuning.org public repository of optimization knowledge; interactive compilation; machine-learning based compilers; decremental CPU/memory characterization; split compilation through static multi-versioning and run-time adaptation)
- **[T44]** 2009-01 - **HiPEAC 2009, Paphos, Cyprus** (Collective optimization)
- **[T45]** 2009-01 - **SMART 2009, Paphos, Cyprus** (Finding representative sets of optimizations for adaptive multiversioning applications)
- **[T46]** 2008-11 - **HiPEAC computing week, Paris, France** (MILEPOST GCC; Interactive Compilation Interface; static multi-versioning to enable dynamic optimizations)
- **[T47]** 2008-10 - **EU-Russia Information and Brokerage Event, Moscow, Russia** (EU FP7 MILEPOST project results: cTuning.org public repository of knowledge and machine learning based self-tuning MILEPOST GCC compiler)
- **[T48]** 2008-09 - **ScalPerf 2008, Bertinoro, Italy** (Enabling Dynamic Optimization and Adaptation for Statically Compiled Programs Using Function Multi-Versioning)
- **[T49]** 2008-08 - **Reservoir Labs, NY, USA** (cTuning.org public plugin-based repository of optimization knowledge and machine-learning based self tuning compiler)
- **[T50]** 2008-06 - **IBM, Toronto, Canada** (cTuning.org public plugin-based repository of optimization knowledge and machine-learning based self tuning compiler)
- **[T51]** 2008-06 - **GCC Summit 2008, Toronto, Canada** (Enabling dynamic and interactive selection of optimal optimization passes in machine learning based MILEPOST GCC)
- **[T52]** 2008-02 - **Imperial College, London, UK** (Continuous collective program optimization and run-time adaptation)
- **[T53]** 2008-01 - **HiPEAC adaptive cluster, Goteborg, Sweden** (Continuous program optimization and adaptation)
- **[T54]** 2008-01 - **Institute of Computing Technology, Beijing, China** (plugin-based public repository and auto-tuning infrastructure; collective optimization; statically enabling dynamic optimizations)
- **[T55]** 2007-10 - **EU Commission, University of Edinburgh, UK** (EU MILEPOST project review: plugin-based public repository and auto-tuning infrastructure; collective optimization; statically enabling dynamic optimizations)
- **[T56]** 2007-07 - **IBM, Toronto, Canada** (continuous program optimization and machine learning)
- **[T57]** 2007-07 - **GCC Summit 2007, Ottawa, Canada** (procedure cloning in GCC to statically enable dynamic optimization and run-time adaptation)
- **[T58]** 2007-04 - **INRIA evaluation commission, Paris, France** (overview of past R&D for Alchemy group review)
- **[T59]** 2007-04 - **IBM, Haifa, Israel** (statistical collaborative optimization; machine learning based optimization; plugin-based auto-tuning framework, and public repository of knowledge)
- **[T60]** 2007-02 - **Intel, Moscow, Russia** (machine learning based optimization)
- **[T61]** 2007-02 - **Institute of System Programming of Russian Academy of Sciences, Moscow, Russia** (machine learning based optimization)
- **[T62]** 2007-01 - **HiPEAC 2007, Ghent, Belgium** (MiDataSets: Creating the Conditions for More Realistic Evaluation of Iterative Optimization)
- **[T63]** 2007-01 - **SMART 2007 @ HiPEAC 2007, Ghent, Belgium** (Building practical interactive compiler)
- **[T64]** 2007-01 - **UPC, Barcelona, Spain** (Continuous run-time adaptation and optimization for statically compiled programs)
- **[T65]** 2006-12 - **University of Edinburgh, UK** (Collaborative statistical optimization)
- **[T66]** 2006-11 - **EU Commission, Ghent, Belgium** (HiPEAC review: GCC with plugins as common EU R&D platform)
- **[T67]** 2006-10 - **National Seoul University, Korea** (Overview of my R&D)

- **[T68]** 2006-03 - **IBM TJ Watson, NY, USA** (A Practical Method for Quickly Evaluating Program Optimizations; Using Machine Learning to Focus Iterative Optimization)
- **[T69]** 2006-03 - **EU Commission, Brussels, Belgium** (Presentation of MILEPOST project proposal - continuous collective optimization; public repositories of knowledge; machine learning)
- **[T70]** 2005-12 - **University of Edinburgh, UK** (Practical run-time evaluation of optimizations)
- **[T71]** 2005-11 - **HiPEAC 2012, Barcelona, Spain** (Practical run-time evaluation of optimizations)
- **[T72]** 2005-06 - **HP, Paris, France** (Collective online tuning)
- **[T73]** 2003-04 - **Paris South University, France** (Iterative compilation and performance prediction)
- **[T74]** 2001-06 - **CPC 2012, Edinburgh, UK** (Fast and Accurate Evaluation of Memory Performance Upper-Bound)
- **[T75]** 2001-05 - **EU Commission, UPC, Spain, Barcelona** (Results of 2 workpackages for the final EU MHAOTEU project review)
- **[T76]** -

20 PARTICIPATING IN PROGRAM COMMITTEES AND REVIEWING

- **CCGrid 2015** (PC and reviewer)
- **IEEE TPDS journal 2015** (reviewer)
- **ACM TACO journal 2015** (reviewer)
- **ACM TACO journal 2014** (reviewer)
- **PARCO journal 2014** (reviewer)
- **CGO 2014** (external reviewer)
- **IPDPS 2014** (PC and reviewer)
- **IEEE MICRO 2013** (reviewer)
- **SBAC-PAD 2013** (PC and reviewer)
- **MuCoCoS 2013** (PC and reviewer)
- **PACT 2013** (reviewer)
- **CC 2013** (PC and reviewer)
- **IPDPS 2013** (PC and reviewer)
- **MuCoCoS 2012** (PC and reviewer)
- **PhD Thesis, Michele Tartara, Politecnico di Milano, Italy, 2012** (reviewer)
- **HiPEAC/ACM TACO 2012** (PC and reviewer)
- **MASCOTS 2011** (reviewer)
- **CPE 2011** (reviewer)
- **ACM OSR** (reviewer)
- **CGO 2011** (PC and reviewer)
- **GROW 2011** (PC and reviewer)
- **IEEE TEVC 2010** (reviewer)
- **Springer-Verlag book on auto-tuning techniques 2010** (reviewer)
- **CLSS 2009** (reviewer)
- **CGO 2010** (reviewer)
- **TPDS 2009** (reviewer)
- **GROW 2010** (PC and reviewer)
- **ICPADS 2009** (PC and reviewer)
- **iWAPT 2009** (PC and reviewer)
- **Open64 2009** (PC and reviewer)
- **JPDC 2009** (reviewer)
- **GROW 2009** (reviewer)
- **JPDC 2008** (reviewer)
- **CGO 2008** (reviewer)
- **IPDPS 2008** (PC and reviewer)
- **IJHPSA 2008** (reviewer)
- **DATES 2008** (reviewer)
- **HiPEAC 2007** (reviewer)
- **ACM ICS 2007** (reviewer)
- **CASES 2007** (PC and reviewer)
- **Transactions on HiPEAC 2006** (reviewer)
- **DATES 2006** (reviewer)
- **PACT 2005** (reviewer)

21 TEACHING AND ORGANIZING COURSES (L)

#	Year	Type	Name
[L1]	2015-03	Lecturer	Guest lectures on "systematic, collaborative and reproducible experimentation in computer engineering via Collective Knowledge Framework and Repository", University of Copenhagen (Denmark)
[L2]	2013-03	Lecturer	Guest lectures on "systematizing tuning of computer systems using crowd-sourcing and statistics", National Taiwan University (Taipei, Taiwan) Lecture 1 slides: http://www.slideshare.net/GrigoriFursin/presentation-fursin-hpsc2013fursin1 Lecture 2 slides: http://www.slideshare.net/GrigoriFursin/presentation-fursin-hpsc2013fursin2
[L3]	2008-09 - 2008-11	Lecturer	Organizing and teaching MS course on "Future computing systems", University Paris Sud (France)
[L4]	2007-10	Lecturer	Preparing and teaching MS course on "Adaptive and Feedback Driven Compilation and Optimization; Machine Learning", University Paris Sud (France)
[L5]	2005-10	Lecturer	Preparing and teaching MS course on "Adaptive and Feedback Driven Compilation and Optimization; Machine Learning", University Paris Sud (France)

22 ADVISING/COLLABORATING (Q)

#	Year	Type	Name
[Q1]	2013-07 - 2013-09	Intern	Vincent Grevendonk <i>in ARM (UK)</i>
[Q2]	2012-09 - 2013-05	MS student	Michael Pankov <i>in Bauman Moscow State Technical University (France)</i> <ul style="list-style-type: none"> • <i>Discussing model-driven optimization search space exploration for compilers</i>
[Q3]	2011-01 - 2011-08	Postdoctoral researcher	Pablo Oliveira <i>in Intel/CEA Exascale Lab (France)</i> <ul style="list-style-type: none"> • <i>Collaborating on program characterization and adaptive auto-tuning combined with machine learning extending cTuning and MILEPOST technology</i> • <i>Most of this work was under NDA until 2012</i> • <i>Presentation of our activities [E17]</i>
[Q4]	2010-04 - 2011-08	MS student and expert engineer	Frank Talbart <i>in Intel/CEA Exascale Lab (France)</i> <ul style="list-style-type: none"> • <i>Helping to develop Collaborative Codelet Tuning Infrastructure (cTuning2 aka CTI)</i> • <i>Most of this work was under NDA until 2012</i> • <i>Presentation of our activities [E17]</i>

[Q5]	2010-04 2011-08	-	PhD student	Souad Koliai <i>in UVSQ (France)</i> funded by Intel/CEA Exascale Lab (France) <ul style="list-style-type: none"> • Collaborating to extend my program behavior characterization technique via semantically non-equivalent binary patching [Pub1, Pub9, Pub56, Pub38] within new DECAN framework • Most of this work was under NDA until 2012 • Associated publication [Pub45]
[Q6]	2009-09 2013-07	-	PhD student	Yuriy Kashikov <i>in UVSQ (France)</i> funded by INRIA and Intel/CEA Exascale Lab (France) <ul style="list-style-type: none"> • Thesis "A holistic approach to predict effective compiler optimizations using machine learning" extending cTuning and MILEPOST technology • Prof. William Jalby kindly agreed to be Yuriy's co-advisor while I was preparing HDR in France • Associated publications [Pub54, Pub5, Pub55, Pub28] • Associated events [E15, E18]
[Q7]	2009-09 cur.	-	PhD student	Abdul Wahid Memon <i>in UVSQ (France)</i> funded by national government, INRIA and HiPEAC/STMicroelectronics internship <ul style="list-style-type: none"> • Thesis is extending cTuning and MILEPOST technology • Associated publications [Pub29, Pub5]
[Q8]	2009-06 2009-08	-	MS student	Yuanjie Huang <i>in Institute of Computing Technology (China)</i> funded by Google Summer of Code <ul style="list-style-type: none"> • Extending MILEPOST technology within GCC to enable fine-grain program optimization • Associated publications [Pub4, Pub10, Pub28] <p>Google Summer of Code description page: http://cTuning.org/gsoc2009-yjhuang</p>
[Q9]	2009-06 2009-08	-	MS student	Liang Peng <i>in Institute of Computing Technology (China)</i> funded by Google Summer of Code <ul style="list-style-type: none"> • Adding static multi-versioning capabilities to GCC and per-function fine-grain optimization for run-time adaptation • Associated publications [Pub10, Pub28] <p>Google Summer of Code description page: http://cTuning.org/gsoc2009-lpeng</p>
[Q10]	2009-04 2009-08	-	Postdoctoral researcher	Cosmin Oancea <i>in INRIA Saclay (France)</i> <ul style="list-style-type: none"> • Exploratory research on program memory behavior characterization and run-time adaptation
[Q11]	2009-01 2009-06	-	Expert engineer	Zbigniew Chamski <i>in INRIA Saclay (France)</i> funded by EU MILEPOST project <ul style="list-style-type: none"> • Moving Interactive Compilation Interface (event-based plugin framework) to mainline GCC • Publication [Pub5]

[Q12]	2008-03 2008-08	-	Postdoctoral researcher	Abid Muslim Malik in <i>INRIA Saclay (France)</i> funded by EU MILEPOST project <ul style="list-style-type: none"> • <i>Exploratory research on semantic program feature analysis to improve optimization predictions in MILEPOST GCC</i>
[Q13]	2008-01 2008-02	-	MS student	Lianjie Luo in <i>Institute of Computing Technology (China)</i> <ul style="list-style-type: none"> • <i>Extending my previous work on statically enabling dynamic optimization [Pub23] with minimal representative sets of versions</i> • <i>Associated publication [Pub30]</i>
[Q14]	2007-03 2007-04	-	PhD student	Victor Jimenez in <i>UPC (Spain) visiting INRIA Saclay (France)</i> funded by EU HiPEAC fellowship <ul style="list-style-type: none"> • <i>Extending my previous work on statically enabling dynamic optimization [Pub23] with predictive scheduling for heterogeneous architectures [Pub15]</i> • <i>Associated publication [Pub15]</i>
[Q15]	2006-09 2008-08	-	PhD student	Cupertino Miranda in <i>INRIA Saclay (France)</i> <ul style="list-style-type: none"> • <i>Helping to develop Interactive Compilation Interface for MILEPOST GCC</i> • <i>Associated publications [Pub46, Pub31, Pub47, Pub33]</i>
[Q16]	2006-09 2007-08	-	PhD student	Piotr Lesnicki in <i>INRIA Saclay (France)</i> <ul style="list-style-type: none"> • <i>Co-advising with Albert Cohen and Olivier Temam to extending my previous work on statically enabling dynamic optimization [Pub23] for split-compilation</i> • <i>Associated publication [Pub34]</i>
[Q17]	2003-06 2003-08	-	MS student	Edwin Bonilla in <i>the University of Edinburgh (UK)</i> <ul style="list-style-type: none"> • <i>Exploring possibilities to use machine learning for fine-grain program optimization prediction (in GCC)</i> • <i>Publication [Pub21]</i>

23 ORGANIZING/CHAIRING EVENTS (E)

#	Year	Type	Event
[E1]	2016-03	Artifact evaluation	Grigori Fursin (<i>dividiti, UK/cTuning foundation, France</i>) and Bruce Childers (<i>University of Pittsburgh, USA</i>). Artifact evaluation for CGO'16 , Barcelona, Spain, March 2016 <i>AE website: http://cTuning.org/event/ae-cgo2016</i> <i>CGO'16 website: http://cgo.org/cgo2016</i>
[E2]	2016-03	Artifact evaluation	Grigori Fursin (<i>dividiti, UK/cTuning foundation, France</i>) and Bruce Childers (<i>University of Pittsburgh, USA</i>). Artifact evaluation for PPOPP'16 , Barcelona, Spain, March 2016 <i>AE website: http://cTuning.org/event/ae-ppopp2016</i> <i>PPOPP'16 website: http://conf.researchr.org/home/PPoPP-2016</i>

[E3]	2015-11	Artifact evaluation	Bruce Childers (<i>University of Pittsburgh, USA</i>), Grigori Fursin (<i>dividiti, UK / cTuning foundation, France</i>), Shriram Krishnamurthi (<i>Brown University, USA</i>) and Andreas Zeller (<i>Saarland University, Germany</i>). Dagstuhl Perspectives Workshop: Artifact Evaluation for Publications , Dagstuhl, Germany, November 2015
[E4]	2015-02	Artifact evaluation	Grigori Fursin (<i>INRIA / cTuning foundation, France</i>) and Bruce Childers (<i>University of Pittsburgh, USA</i>). Artifact evaluation for PPOPP'15 , San Francisco Bay Area, CA, USA, February 2015 AE website: http://cTuning.org/event/ae-ppopp2015 PPOPP'15 website: http://ppopp15.soe.ucsc.edu
[E5]	2015-02	Artifact evaluation	Grigori Fursin (<i>INRIA / cTuning foundation, France</i>) and Bruce Childers (<i>University of Pittsburgh, USA</i>). Artifact evaluation for CGO'15 , San Francisco Bay Area, CA, USA, February 2015 AE website: http://cTuning.org/event/ae-cgo2015 CGO'15 website: http://cgo.org/cgo2015
[E6]	2015-01	Special journal issue	Alex Jones (<i>University of Pittsburgh, USA</i>) and Grigori Fursin (<i>INRIA, France</i>). Special Issue on Reproducible Research Methodologies, IEEE Transactions on Emerging Topics in Computing (TETC) <ul style="list-style-type: none"> • <i>Submission Deadline: Sept 1, 2014</i> • <i>Reviews Completed: Nov 1, 2014</i> • <i>Major Revisions Due (if Needed): Dec 15, 2014</i> • <i>Reviews of Revisions Completed (if Needed): Jan 3, 2015</i> • <i>Minor Revisions Due (if Needed): Jan 20, 2015</i> • <i>Notification of Final Acceptance: Jan 31, 2015</i> • <i>Publication Materials for Final Manuscripts Due: Feb 15, 2015</i> • <i>Publication date: 2nd Issue of 2015</i> <p>IEEE TETC Website: http://www.computer.org/portal/web/tetc Submission website: https://mc.manuscriptcentral.com/tetc-cs Flyer and CFP: http://www.computer.org/cms/Computer.org/transactions/cfps/cfp_tetcsi_rrm.pdf Related OCCAM project: http://www.occamportal.org Related Collective Mind project: http://c-mind.org Related Collective Mind repository: http://c-mind.org/repo</p>
[E7]	2014-06	Workshop	Grigori Fursin (<i>INRIA, France</i>), Alex Jones (<i>University of Pittsburgh, USA</i>), Daniel Mosse (<i>University of Pittsburgh, USA</i>) and Bruce Childers (<i>University of Pittsburgh, USA</i>). 1st ACM SIGPLAN International Workshop on Reproducible Research Methodologies and New Publication Models (TRUST) co-located with PLDI 2014 , Edinburgh, UK, June 2014 <ul style="list-style-type: none"> • <i>Workshop Date: June 12, 2014</i> • <i>Abstract submission deadline: March 7, 2014 (Anywhere on Earth)</i> • <i>Paper submission deadline: March 14, 2014 (Anywhere on Earth)</i> • <i>Notification date: April 14, 2014</i> • <i>Final version deadline: May 2, 2014</i> • <i>Related ADAPT'14 panel on reproducible research methodologies and new publication models in computer engineering [E11]</i> <p>Website: http://c-mind.org/events/trust2014 ACM DL: http://dl.acm.org/citation.cfm?id=2618137</p>

[E8]	2014-05	Panel	<p>Marisa Gil (<i>BSC/UPC, Spain</i>), Chris Fensch (<i>University of Edinburgh, UK</i>) and Grigori Fursin (<i>INRIA, France</i>). Panel "Is Current Research on Heterogeneous HPC Platforms inline with Real-world Application needs?" co-located with HiPEAC Spring Computing Systems Week 2014, BSC, Barcelona, Spain, May 2014</p> <ul style="list-style-type: none"> • <i>Participants: Paolo Faraboschi (HP Labs, Spain), Ana Lucia Varbanescu (University of Amsterdam, Netherlands), Mats Brorsson (KTH, Sweden), Pooyan Dadvand (CIMNE/UPC, Spain) and Paul Keir (CodePlay, UK);</i> <p><i>Website and slides: http://www.hipeac.net/add/res/102/481 CSW Website: http://www.hipeac.net/csw/2014/barcelona</i></p>
[E9]	2014-02	Workshop	<p>Alex Jones (<i>University of Pittsburgh, USA</i>), Grigori Fursin (<i>INRIA, France</i>), Daniel Mosse (<i>University of Pittsburgh, USA</i>) and Bruce Childers (<i>University of Pittsburgh, USA</i>). Workshop on Reproducible Research Methodologies (REPRODUCE) co-located with HPCA 2014, Orlando, Florida, USA, February 2014</p> <ul style="list-style-type: none"> • <i>Submission Deadline: January 8, 2014</i> • <i>Notification of Acceptance: January 23, 2014</i> • <i>Final (Camera-ready) Manuscripts Due: January 30, 2014</i> • <i>Workshop Date: February 15, 2014</i> <p><i>Website: http://www.occamportal.org/reproduce</i></p>
[E10]	2014-01	Workshop	<p>Christophe Dubach (<i>University of Edinburgh, UK</i>) and Grigori Fursin (<i>INRIA, France</i>). 4th International workshop on Adaptive Self-Tuning Computing Systems (ADAPT) with a special focus on reproducibility co-located with HiPEAC 2014, Vienna, Austria, January 2014</p> <ul style="list-style-type: none"> • <i>Keynote: "Towards Resource Management in Parallel Architectures Under the Hood" by Prof. Per Stenstrom (Chalmers University of Technology, Sweden)</i> • <i>Associated panel on reproducible research methodologies and new publication models in computer engineering [E11]</i> • <i>Shared research material and experimental results for 2 papers have been validated by our volunteers: Alberto Magni from the University of Edinburgh, UK and Sascha Hunold from Vienna University of Technology, Austria</i> • <i>Sponsored by Nvidia, USA</i> <p><i>Website: http://adapt-workshop.org Final program: http://adapt-workshop.org/program.htm Shared artifacts (paper 1): http://c-mind.org/repo/?view_cid=876bf92b1409cb46:26177154df5a20ef Shared artifacts (paper 2): http://c-mind.org/repo/?view_cid=876bf92b1409cb46:ac4b5d8f1a252bd8</i></p>

[E11]	2014-01	Panel	<p>Grigori Fursin (INRIA, France), Alex Jones (University of Pittsburgh, USA), Daniel Mosse (University of Pittsburgh, USA) and Bruce Childers (University of Pittsburgh, USA). Panel on reproducible research methodologies and new publication models in computer engineering co-located with ADAPT 2014 , Vienna, Austria, January 2014</p> <ul style="list-style-type: none"> • <i>Participants: Jack Davidson (University of Virginia / Co-Chair of ACM's Publication Board, USA); Lieven Eeckhout (Ghent University / Intel ExaScience Lab, Belgium); Sascha Hunold, Jesper Larsson Traff (Vienna University of Technology, Austria); Anton Lohmotov (ARM, UK); Alex K.Jones, Daniel Mosse, Bruce Childers (University of Pittsburgh, USA); Grigori Fursin (INRIA, France);</i> • <i>Associated workshop [E10]</i> <p>Website and slides: http://adapt-workshop.org/program.htm Brief blog note: http://collective-mind.blogspot.fr/2014/02/adapt14-panel-on-reproducible-research.html</p>
[E12]	2013-05	Thematic session	<p>Grigori Fursin (INRIA, France). Thematic session - making computer engineering a science co-located with ACM ECRC 2013 / HiPEAC computing week 2013 , Paris, France, May 2013</p> <ul style="list-style-type: none"> • <i>Keynote on "OCCAM: Open Curation for Computer Architecture Modeling" by Prof. Bruce Childers and Prof. Alex Jones (University of Pittsburgh, USA)</i> • <i>Talks by Vittorio Zaccaria (Politecnico di Milano, Italy), Christophe Guillon and Christian Bertin (STMicroelectronics, France), Christoph Reichenbach (Johann-Wolfgang Goethe Universitat Frankfurt, Germany)</i> • <i>Newsletter in [Pub51]</i> • <i>Sponsored by HiPEAC, EU</i> <p>Website: http://www.hipeac.net/thematic-session/making-computer-engineering-science Backup: http://ctuning.org/making-computer-engineering-a-science-2013</p>
[E13]	2013-01	Workshop	<p>Christophe Dubach (University of Edinburgh, UK) and Grigori Fursin (INRIA, France). 3rd International workshop on Adaptive Self-Tuning Computing Systems (ADAPT) co-located with HiPEAC 2013 , Berlin, Germany, January 2013</p> <ul style="list-style-type: none"> • <i>Keynote on "Autotuning Recursive Functions" by Prof. Markus Pueschel (ETHZ, Switzerland) sponsored by Microsoft Research</i> • <i>Introduction in [Pub52]</i> • <i>Sponsored by Microsoft Research, UK and Nvidia, USA</i> <p>Website: http://homepages.inf.ed.ac.uk/cdubach/adapt2013 ACM DL: http://dl.acm.org/citation.cfm?id=2484904</p>
[E14]	2012-04	Thematic session	<p>Grigori Fursin (INRIA, France). Thematic session - "Collective characterization, optimization and design of computer systems" co-located with HiPEAC spring computing week 2012 , Goteborg, Sweden, April 2012</p> <ul style="list-style-type: none"> • <i>Talks by Marisa Gil (UPC, Spain), Lasse Natvig (NTNU, Norway), David Whalley (Florida State University, USA), Cristina Silvano (Politecnico di Milano, Italy)</i> • <i>Sponsored by HiPEAC, EU</i> <p>Website: http://www.hipeac.net/thematic-session/collective-characterization-optimization-and-design-computer-systems Backup: http://ctuning.org/events/ctuning-hipeac-theme-2012</p>

[E15]	2012-03	Workshop	<p>Grigori Fursin (<i>Intel/CEA Exascale Lab, France</i>), Jason Mars (<i>University of Virginia, USA</i>), Yuriy Kashnikov (<i>Intel/CEA Exascale Lab, France</i>) and Robert Hundt (<i>Google, USA</i>). 2nd International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era (EXADAPT) co-located with ASPLOS 2012, London, UK, March 2012</p> <ul style="list-style-type: none"> • <i>Keynote on "Self-Tuning Bio-Inspired Massively-Parallel Computing" by Prof. Steve Furber (Manchester University, UK)</i> • <i>Introduction in [Pub54]</i> • <i>Associated panel [E16]</i> • <i>Sponsored by Google, USA</i> <p>Website: http://exadapt.org/2012/index.html ACM DL: http://dl.acm.org/citation.cfm?id=2185475</p>
[E16]	2012-03	Panel	<p>Grigori Fursin (<i>INRIA, France</i>). Panel - Joint EXADAPT 2012 / GPGPU 2012 round table on "Leveraging GPUs and Self-Tuning Systems on the Road to Exascale" co-located with EXADAPT 2012, London, UK, March 2012</p> <ul style="list-style-type: none"> • <i>Participants: Steve Furber (University of Manchester, UK), Anton Lokhmetov (ARM, UK), Paul Kelly (Imperial College London, UK)</i> • <i>Associated workshop [E15]</i> <p>Website: http://exadapt.org/2012/joint_exadapt_gpgpu_round_table.html</p>
[E17]	2011-11	BOF	<p>Grigori Fursin (<i>INRIA, France</i>) and Marie-Christine Sawley (<i>Intel/CEA Exascale Lab, France</i>). BOF - Collaboratively mining rich information to prepare the Exascale challenges @ SC 2011 co-located with SuperComputing 2011, Seattle, WA, USA, November 2011</p> <ul style="list-style-type: none"> • <i>Presenting and discussing long-term R&D of my team at Intel/CEA Exascale Lab, France</i> • <i>Continuation in [Pub50, Pub29]</i> <p>Website: http://www.sc11.supercomputing.org/schedule/event_detail.php?evId=bof164 Backup: http://ctuning.org/sc2011-bof-characterization</p>
[E18]	2011-05	Workshop	<p>Grigori Fursin (<i>Intel/CEA Exascale Lab, France</i>), Jason Mars (<i>University of Virginia, USA</i>), Yuriy Kashnikov (<i>Intel/CEA Exascale Lab, France</i>) and Robert Hundt (<i>Google, USA</i>). 1st International ACM Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era (EXADAPT) co-located with PLDI 2011 / FCRC 2011, San Jose, USA, June 2011</p> <ul style="list-style-type: none"> • <i>Keynote on "Autotuning in the Exascale Era" by Prof. Katherine Yelick (LBNL and UC Berkeley, USA)</i> • <i>Introduction in [Pub55]</i> • <i>Sponsored by Google, USA and ACM, USA</i> <p>Website: http://exadapt.org/2011 ACM DL: http://dl.acm.org/citation.cfm?id=2000417</p>

[E19]	2011-04	Workshop	<p>Grigori Fursin (<i>Intel/CEA Exascale Lab, France</i>) and John Cavazos (<i>University of Delaware, USA</i>). 5th International Workshop Statistical and Machine learning approaches to ARchitecture and compilaTion (SMART) co-located with CGO 2011 , Chamonix, France, April 2011</p> <ul style="list-style-type: none"> • <i>Keynote on "Automatic Performance Tuning and Machine Learning" by Prof. Markus Pueschel (ETHZ, Switzerland)</i> • <i>Chair: Prof. Francois Bodin (CAPS Entreprise, France)</i> • <i>Sponsored by CAPS Entreprise, France</i> <p>Website: http://ctuning.org/workshop-smart2011 Final program: http://ctuning.org/workshop-smart2011-program</p>
[E20]	2010-01	Tutorial	<p>Sid Touati (<i>UVSQ, France</i>) and Grigori Fursin (<i>INRIA, France</i>). Tutorial - "Speedup-Test: Statistical Methodology to Evaluate Program Speedups and their Optimisation Techniques" co-located with HiPEAC 2010 , Pisa, Italy, January 2010</p> <p>Website: http://www.hipeac.net/conference/pisa/speedup</p>
[E21]	2010-01	Workshop	<p>Grigori Fursin (<i>INRIA, France</i>) and John Cavazos (<i>University of Delaware, USA</i>). 4th International Workshop Statistical and Machine learning approaches to ARchitecture and compilaTion (SMART) co-located with HiPEAC 2010 , Pisa, Italy, January 2010</p> <ul style="list-style-type: none"> • <i>Keynote on "Moving adaptation into individual optimizations" by Prof. Keith Cooper (Rice University, USA)</i> • <i>Chair: Prof. David Whalley (Florida State University, USA)</i> • <i>Sponsored by Intel/CEA Exascale Lab, France</i> <p>Website: http://ctuning.org/workshop-smart10 Final program: http://ctuning.org/workshop-smart10-program</p>
[E22]	2010-01	Panel	<p>Dorit Nuzman (<i>IBM Haifa, Israel</i>) and Grigori Fursin (<i>INRIA, France</i>). 2nd International Workshop on GCC Research Opportunities (GROW) co-located with HiPEAC 2010 , Pisa, Italy, January 2010</p> <ul style="list-style-type: none"> • <i>Keynote on "Using GCC as a toolbox for research: GCC plugins and whole-program compilation" by Diego Novillo (Google, Canada)</i> <p>Website: http://ctuning.org/workshop-grow10</p>
[E23]	2009-06	Tutorial	<p>Grigori Fursin (<i>INRIA, France</i>). cTuning tools tutorial on collaborative and reproducible program and architecture characterization and autotuning co-located with HiPEAC computing systems week , Infineon, Munich, Germany, June 2009</p>
[E24]	2009-01	Workshop	<p>Grigori Fursin (<i>INRIA, France</i>) and John Cavazos (<i>University of Delaware, USA</i>). 3rd International Workshop Statistical and Machine learning approaches to ARchitecture and compilaTion (SMART) co-located with HiPEAC 2009 , Paphos, Cyprus, January 2009</p> <ul style="list-style-type: none"> • <i>Associated panel [E25]</i> • <i>Chair: Prof. David Padua (UIUC, USA)</i> <p>Website: http://www.hipeac.net/smart-workshop.html Final program: http://unidapt.org/index.php/Dissemination:SMART09</p>

[E25]	2009-01	Panel	<p>Grigori Fursin (<i>INRIA, France</i>) and John Cavazos (<i>University of Delaware, USA</i>). Panel - Can machine learning help to solve the multicore code generation issues? co-located with HiPEAC 2009 , Paphos, Cyprus, January 2009</p> <ul style="list-style-type: none"> • <i>Chair: Francois Bodin (CAPS-Enterprise, France); participants: Marcelo Cintra (University of Edinburgh, UK), Bilha Mendelson (IBM Haifa, Israel), Lawrence Rauchwerger (Texas A&M University, USA), Per Stenstrom (Chalmers University of Technology, Sweden)</i> • <i>Associated workshop [E24]</i> <p><i>Website: http://www.hipeac.net/smart-workshop.html</i> <i>Video: http://cTuning.org/smart2009-panel-video</i></p>
[E26]	2008-01	Workshop	<p>Grigori Fursin (<i>INRIA, France</i>) and John Cavazos (<i>University of Delaware, USA</i>). 2nd International Workshop Statistical and Machine learning approaches to ARchitecture and compilaTion (SMART) co-located with HiPEAC 2008 , Goteborg, Sweden, January 2008</p> <ul style="list-style-type: none"> • <i>Chair: Prof. Michael O'Boyle (University of Edinburgh, UK)</i> <p><i>Website: http://www.hipeac.net/smart-workshop-08.html</i></p>
[E27]	2007-01	Tutorial	<p>Albert Cohen (<i>INRIA, France</i>), Ayal Zaks (<i>IBM Haifa, Israel</i>), Dorit Nuzman (<i>IBM Haifa, Israel</i>), Diego Novillo (<i>Red Hat, USA</i>), Roberto Costa (<i>STMicroelectronics, Italy</i>), Grigori Fursin (<i>INRIA, France</i>) and Sebastian Pop (<i>Ecole des Mines de Paris, France</i>). 2nd HiPEAC GCC Tutorial: How To and Return on Experience co-located with HiPEAC 2007 , Ghent, Belgium, January 2007</p> <p><i>Website: http://www.hipeac.net/gcc-tutorial</i></p>
[E28]	2007-01	Workshop	<p>Grigori Fursin (<i>INRIA, France</i>) and John Cavazos (<i>University of Delaware, USA</i>). 1st International Workshop Statistical and Machine learning approaches to ARchitecture and compilaTion (SMART) co-located with HiPEAC 2007 , Ghent, Belgium, January 2007</p> <p><i>Website: http://www.hipeac.net/smart-workshop-07.html</i> <i>Final program: http://www.hipeac.net/smart-workshop-07-program.html</i></p>
[E29]	2001-05	Workshop	<p>Michael O'Boyle (<i>University of Edinburgh, UK</i>) and Grigori Fursin (<i>University of Edinburgh, UK</i>). 9th Workshop on Compilers for Parallel Computers (CPC) , Edinburgh, UK, June 2001</p> <ul style="list-style-type: none"> • <i>Local organizers: Michael O'Boyle (University of Edinburgh, UK), Grigori Fursin (University of Edinburgh, UK)</i> <p><i>Website: http://www.icsa.informatics.ed.ac.uk/cpc2001</i></p>

24 PUBLICATIONS AND OTHER DISSEMINATION MATERIAL (P)

Google Scholar page with citations and H-index: <http://scholar.google.com/citations?user=IwcnpkwAAAAJ>

24.1 THESES

[Pub1] Grigori Fursin. **Iterative Compilation and Performance Prediction for Numerical Applications.**
PhD thesis, University of Edinburgh, UK, May 2004 (cID=29db2248aba45e59:976aac68c6834af6)



- Based on publications [Pub38, Pub37, Pub9]
- PhD degree [Z5]
- Concept [M9]
- Presented ideas are now being used and extended in Intel Exascale Lab (France) [I2, Pub45]
- Associated public software [S19]

[Pub2] Grigori Fursin. **Unifying remote access to high-performance computing systems (HPC) as a web service.**

MS Thesis, Moscow Institute of Physics and Technology and Institute of High-Performance Computing Systems of Russian Academy of Sciences, Moscow, Russia, May 1999 (cID=29db2248aba45e59:c3a327ce9649509d)



- Concept [M9]
- MS degree [Z7]
- Associated public software [S19]

24.2 INTERNATIONAL JOURNALS

[Pub3] Grigori Fursin, Renato Miceli, Anton Lokhmotov, Michael Gerndt, Marc Baboulin, Allen D. Malony, Zbigniew Chamski, Diego Novillo and Davide Del Vento. **Collective Mind: towards practical and collaborative autotuning.**

Special issue on Automatic Performance Tuning for HPC Architectures, Scientific Programming Journal, IOS Press, Edinburgh, UK, August 2014 (cID=29db2248aba45e59:6f40bc99c4f7df58)



- This work summarizes our long-term vision on making autotuning practical and reproducible using public repository of knowledge, common plugin-based autotuning infrastructure, predictive analytics (machine learning, data mining, statistical analysis, feature detection), crowd-sourcing and collective intelligence
- **This work extends previous article [P24]**
- Related Collective Mind infrastructure and repository [S2]
- **This work supports our initiative on open research and publication model where all experimental results and related material is continuously shared, validated and improved by the community [Pub26]. To set up an example, we continue sharing all benchmarks, datasets, tools, models and experimental results in Collective Mind repository (c-mind.org/repo) and in a new version: Collective Knowledge Repository (cknowledge.org/repo)**

CK-powered interactive article with all shared artifacts for reproducibility at cknowledge.org/repo: <http://cknowledge.org/repo/web.php?wcid=29db2248aba45e59:6f40bc99c4f7df58>

[Pub4] Yang Chen, Shuangde Fang, Yuanjie Huang, Lieven Eeckhout, Grigori Fursin, Olivier Temam and Chengyong Wu. **Deconstructing iterative optimization.**

ACM Transactions on Architecture and Code Optimization (TACO), Volume 9, Number 3, pages 21:1-21:30, September 2012 (cID=29db2248aba45e59:77927d4cd49b4dc6)



- [Pub5] Grigori Fursin, Yuriy Kashnikov, Abdul Wahid Memon, Zbigniew Chamski, Olivier Temam, Mircea Namolaru, Elad Yom-Tov, Bilha Mendelson, Ayal Zaks, Eric Courtois, Francois Bodin, Phil Barnard, Elton Ashton, Edwin Bonilla, John Thomson, Christopher Williams and Michael O'Boyle. **Milepost GCC: Machine Learning Enabled Self-tuning Compiler.**
International Journal of Parallel Programming, pages 296-327, June 2011
([CID=29db2248aba45e59:a31e374796869125](#))



- *Implements concept [M4]*
- *Funded by [F7, F4]*
- *Collaboration with IBM (Israel), University of Edinburgh (UK), ARC (now Synopsys, UK), CAPS Entreprise (France), and ICT (China)*
- *MILEPOST GCC software [S9]; most of technology is now available in mainline GCC and is being added to major commercial compilers*
- *cTuning CC software [S8]*
- *cTuning repository [R4]*
- *Award [A2]*
- ***Considered by IBM to be the first practical machine-learning based compiler in the world (IBM press-release [Pub59])***
- ***All benchmarks, datasets, tools, models and experimental results have been released to public for collaborative validation and extension!***

- [Pub6] Grigori Fursin and Olivier Temam. **Collective optimization: A practical collaborative approach.**
ACM Transactions on Architecture and Code Optimization (TACO), Volume 7, Number 4, pages 20:1-20:29, December 2010 ([CID=29db2248aba45e59:530e5f456ea259de](#))



- *Implements concept [M5]*
- *Award [A2]*
- *Extended paper [Pub14]*
- ***All benchmarks, datasets, tools, models and experimental results have been released in cTuning public optimization repository of knowledge [R4] for collaborative validation and extension!***
- ***Continues in [M2]***

- [Pub7] Shun Long and Grigori Fursin. **Systematic search within an optimisation space based on Unified Transformation Framework.**
International Journal of Computational Science and Engineering (IJCSE), Volume 4, Number 2, pages 102-111, July 2009 ([CID=29db2248aba45e59:c3ea0b108239a1da](#))



- [Pub8] Grigori Fursin, Albert Cohen, Michael O'Boyle and Olivier Temam. **Quick and Practical Run-Time Evaluation of Multiple Program Optimizations.**
Transactions on High-Performance Embedded Architectures and Compilers, Volume 1, pages 13-31, 2006 ([CID=29db2248aba45e59:fd899287e64b4999](#))



- ***Extended version of [Pub23]***

- [Pub9] Grigori Fursin, Michael O'Boyle, Olivier Temam and Gregory Watts. **A fast and accurate method for determining a lower bound on execution time.**
Concurrency: Practice and Experience, Volume 16, Number 2-3, pages 271-292, January 2004
([CID=29db2248aba45e59:b254c18c8794ba29](#))



- *Concept [M8]*
- *Extended version of publication [Pub38]*
- *Associated public software [S18]*
- *Extended in PhD thesis [Pub1] and DECAN framework [Pub45]*

24.3 INTERNATIONAL CONFERENCES

[**Pub10**] Yang Chen, Yuanjie Huang, Lieven Eeckhout, Grigori Fursin, Liang Peng, Olivier Temam and Chengyong Wu. **Evaluating iterative optimization across 1000 datasets.**
Proceedings of the 2010 ACM SIGPLAN conference on Programming language design and implementation (PLDI), pages 448-459, Toronto, Canada, June 2010 (acceptance rate: 20% (41/204))
(*clD=29db2248aba45e59:dabab2e9ff26c745*)



- *Extension of [Pub18]*
- *Being added to public repository of knowledge [R2]*

[**Pub11**] Mircea Namolaru, Albert Cohen, Grigori Fursin, Ayal Zaks and Ari Freund. **Practical aggregation of semantical program properties for machine learning based optimization.**
Proceedings of the International Conference on Compilers, Architectures and Synthesis for Embedded Systems (CASES'10), pages 197-206, Scottsdale, Arizona, USA, October 2010 (acceptance rate=29%)
(*clD=29db2248aba45e59:1dc5655fcaa4c4ca*)



[**Pub12**] Christophe Dubach, Timothy M. Jones, Edwin Bonilla, Grigori Fursin and Michael O'Boyle. **Portable compiler optimisation across embedded programs and microarchitectures using machine learning.**
Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), pages 78-88, New York, NY, USA, December 2009 (acceptance rate: 25% (52/209))
(*clD=29db2248aba45e59:2fb0f364a428b631*)



- *HiPEAC paper award [A4]*
- *Includes concept [M8]*
- *Christophe Dubach received BCS/CPHC Distinguished Dissertation Award'09 for his related thesis "Using Machine-Learning to Efficiently Explore the Architecture/Compiler Co-Design Space" supervised by Prof. Michael O'Boyle.*

[**Pub13**] John Thomson, Michael O'Boyle, Grigori Fursin and Björn Franke. **Reducing training time in a one-shot machine learning-based compiler.**
Proceedings of the 22nd international conference on Languages and Compilers for Parallel Computing (LCPC), Newark, DE, USA, October 2009 (*clD=29db2248aba45e59:84bd982340c913b0*)



[**Pub14**] Grigori Fursin and Olivier Temam. **Collective Optimization.**
Proceedings of the 4th International Conference on High Performance Embedded Architectures and Compilers (HiPEAC), pages 34-49, Paphos, Cyprus, January 2009 (acceptance rate: 28% (27/97))
(*clD=29db2248aba45e59:2df284a66bdf5468*)



- *Implements concept [M5]*
- *Considerably extended in journal version [Pub6]*

[**Pub15**] Víctor J. Jiménez, Lluís Vilanova, Isaac Gelado, Marisa Gil, Grigori Fursin and Nacho Navarro. **Predictive Runtime Code Scheduling for Heterogeneous Architectures.**
Proceedings of the 4th International Conference on High Performance Embedded Architectures and Compilers (HiPEAC), pages 19-33, Paphos, Cyprus, January 2009 (acceptance rate=28% (27/97))
(*clD=29db2248aba45e59:3b9c0386749be8ad*)



- *I initiated collaborative HiPEAC project and obtained funding [F5] to extend [Pub23, M6]*
- *Similar approaches for gluing/adapting applications for heterogeneous architectures are now used in Intel's Qilin and in CAPS Enterprise's HMPP frameworks*
- *Being extended in [Pub50, M2]*

- [Pub16] Christophe Dubach, John Cavazos, Björn Franke, Grigori Fursin, Michael O’Boyle and Olivier Temam. **Fast compiler optimisation evaluation using code-feature based performance prediction.** *Proceedings of the 4th international conference on Computing Frontiers*, pages 131-142, Ischia, Italy, May 2007 (acceptance rate=50% (28/56)) (cid=29db2248aba45e59:5725dd49f40f29dc)



- [Pub17] John Cavazos, Grigori Fursin, Felix Agakov, Edwin Bonilla, Michael O’Boyle and Olivier Temam. **Rapidly Selecting Good Compiler Optimizations using Performance Counters.** *Proceedings of the International Symposium on Code Generation and Optimization (CGO)*, pages 185-197, San Jose, USA, March 2007 (cid=29db2248aba45e59:24fbaabd00976640)



- We added hardware counters to [S14, S10, R5] and used PCA to improve machine-learning based optimization prediction

- [Pub18] Grigori Fursin, John Cavazos, Michael O’Boyle and Olivier Temam. **MiDataSets: creating the conditions for a more realistic evaluation of Iterative optimization.** *Proceedings of the 2nd international conference on High performance embedded architectures and compilers (HiPEAC)*, pages 245-260, Ghent, Belgium, January 2007 (acceptance rate=29%) (cid=29db2248aba45e59:404c91542f38bd58)



- Preparing R&D on collective optimization [M5, M4, M2]
- Released in public repositories of knowledge [R4, R2]

- [Pub19] Shun Long, Grigori Fursin and Björn Franke. **A cost-aware parallel workload allocation approach based on machine learning techniques.** *Proceedings of the 2007 IFIP international conference on Network and Parallel Computing (NPC)*, pages 506-515, Dalian, China, September 2007 (cid=29db2248aba45e59:e7f43b0b29e66f09)



- [Pub20] John Cavazos, Christophe Dubach, Felix Agakov, Edwin Bonilla, Michael O’Boyle, Grigori Fursin and Olivier Temam. **Automatic performance model construction for the fast software exploration of new hardware designs.** *Proceedings of the International Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES)*, pages 24-34, Seoul, Korea, October 2006 (acceptance rate=41% (41/100)) (cid=29db2248aba45e59:5d3c4fe024f79be8)



- **Finalist best paper award**

- [Pub21] Felix Agakov, Edwin Bonilla, John Cavazos, Björn Franke, Grigori Fursin, Michael O’Boyle, John Thomson, Mark Toussaint and Christopher K. I. Williams. **Using Machine Learning to Focus Iterative Optimization.** *Proceedings of the International Symposium on Code Generation and Optimization (CGO)*, pages 295-305, New York, NY, USA, March 2006 (acceptance rate=36% (29/80)) (cid=29db2248aba45e59:3854f3510e3e6c91)



- **Best presentation award**
- Concept [M4]
- **Extended journal version in [Pub5]**

- [Pub22] Björn Franke, Michael O’Boyle, John Thomson and Grigori Fursin. **Probabilistic source-level optimization of embedded programs.** *Proceedings of the 2005 ACM SIGPLAN/SIGBED conference on Languages, Compilers, and Tools for Embedded Systems (LCTES)*, pages 78-86, Chicago, IL, USA, June 2005 (acceptance rate=26% (25/95)) (cid=29db2248aba45e59:c83fc3a0393b4339)



- **Introducing probabilistic approach for adaptive exploration of large optimization spaces (effectively reducing tuning dimensions)**

- [Pub23] Grigori Fursin, Albert Cohen, Michael O'Boyle and Olivier Temam. **A practical method for quickly evaluating program optimizations.**
Proceedings of the First International Conference on High Performance Embedded Architectures and Compilers (HiPEAC), pages 29-46, Barcelona, Spain, November 2005 (acceptance rate=20% (17/84))
(cID=29db2248aba45e59:5e16c10fa43f8bb4)



- **Highest ranked paper [A5]**
- *In this paper I introduced a novel concept to statically enable run-time optimizations and self-tuning binaries through code cloning and integrated low-overhead program/system behaviour monitoring plugin that can evaluate or select multiple optimization online using program phases [M6]. It has been referenced in patents and extended in academia and industry. We utilized Interactive Compilation Interface for PathScale compiler with loop vectorization, tiling, unrolling, interchange, fission/fusion, pipelining, prefetching and array padding to make static binaries adaptable and reactive to various environments and run-time behaviour to improve execution time and power consumption.*
- *Presented ideas are now being used and extended by Google for data centers (cloud computing) to build adaptive applications and save energy [F4]*
- *Funded by [A6, F4]*
- *Associated public software [S13, S14, S16]*
- *Multi-versioning support for adaptive applications is now available in mainline GCC since version 4.8 [F4]*

- [Pub24] Jaume Abella, Sid Touati, A Anderson, C Ciuraneta, J M Codina M Dai, Christine Eisenbeis, Grigori Fursin, Antonio Gonzalez, Joseph Llosa, Michael O'Boyle, A Randrianatoavina, J Sanchez, Olivier Temam, Xavier Vera and Gregory Watts. **The MHAOTEU Toolset for Memory Hierarchy Management.**
16th IMACS World Congress on Scientific Computation, Applied Mathematics and Simulation, Lausanne, Switzerland, August 2000 (cID=29db2248aba45e59:4e8d86d598edbe3b)



- *Part of EU FP5 MHAOTEU project [J14]*
- *Associated public software [S18]*

24.4 INTERNATIONAL WORKSHOPS

- [Pub25] Grigori Fursin, Abdul Wahid Memon, Christophe Guillon and Anton Lokhmotov. **Collective Mind, Part II: Towards Performance- and Cost-Aware Software Engineering as a Natural Science.**
18th International Workshop on Compilers for Parallel Computing (CPC'15), London, UK, London, UK, January 2015 (cID=29db2248aba45e59:cd11e3a188574d80)



- *Partially funded by [F1]*
- *This work summarizes our long-term vision on enabling cost-aware computer engineering autotuning using methodology from physics, biology and other natural sciences*
- **This work extends our previous article [Pub3]**
- *Related Collective Knowledge infrastructure and repository (CK) [S1]*
- *Related Collective Mind infrastructure and repository (deprecated for CK) [S2]*
- **This work supports our initiative on open research and publication model where all experimental results and related material is continuously shared, validated and improved by the community [Pub26]. To set up an example, we continue sharing all benchmarks, datasets, tools, models and experimental results in Collective Mind repository (c-mind.org/repo) and in a new version: Collective Knowledge Repository (cknowledge.org/repo)**

Some reddit discussions: https://www.reddit.com/r/programming/comments/3bh9v0/so_all_software_engineers_need_is_a_collective
CK-powered interactive article with all shared artifacts for reproducibility at cknowledge.org/repo:
<http://cknowledge.org/repo/web.php?wcid=29db2248aba45e59:cd11e3a188574d80>

[Pub26] Grigori Fursin and Christophe Dubach. **Experience report: community-driven reviewing and validation of publications.**

Proceedings of the 1st Workshop on Reproducible Research Methodologies and New Publication Models in Computer Engineering (TRUST 2014) co-located with PLDI 2014, Edinburgh, UK, June 2014 (cID=29db2248aba45e59:40a4b58adfb594a8)

- *We propose a new and open publication model for reproducible research where articles, experiments and artifacts are reviewed by the community based on our practical related experience during MILEPOST, cTuning and Collective Mind projects since 2008*



[Pub27] Grigori Fursin. **Collective Tuning Initiative: automating and accelerating development and optimization of computing systems.**

Proceedings of the GCC Developers' Summit, Montreal, Canada, June 2009 (cID=29db2248aba45e59:0c44d9a2db3de3c9)

- *Introduced concept of crowd-tuning using public repositories of knowledge, autotuning, machine learning and crowd-sourcing, and new publication model where results are continuously validated and extended by the community [M5]*
- *Award [A2]*
- *MILEPOST GCC software [S9]; most of technology is now available in mainline GCC and is being added to major commercial compilers*
- *cTuning CC software [S8]*
- *cTuning public optimization repository of knowledge [R4]*
- ***All benchmarks, datasets, tools, models and experimental results have been released to public for collaborative validation and extension! See our new research and publication proposal [Pub26]!***
- ***Continues in [M2, Pub50, Pub3]***



[Pub28] Yuanjie Huang, Liang Peng, Chengyong Wu, Yuriy Kashnikov, Joern Renneke and Grigori Fursin. **Transforming GCC into a research-friendly environment: plugins for optimization tuning and re-ordering, function cloning and program instrumentation.**

2nd International Workshop on GCC Research Opportunities (GROW), co-located with HiPEAC'10 conference, Pisa, Italy, January 2010 (acceptance rate: 57% (8/14)) (cID=29db2248aba45e59:e94906b2eacf7d3f)



[Pub29] Grigori Fursin. **Collective Tuning Initiative: automating and accelerating development and optimization of computing systems.**

Proceedings of the GCC Developers' Summit, Montreal, Canada, June 2009 (cID=29db2248aba45e59:0c44d9a2db3de3c9)

- *Introduced concept of crowd-tuning using public repositories of knowledge, autotuning, machine learning and crowd-sourcing, and new publication model where results are continuously validated and extended by the community [M5]*
- *Award [A2]*
- *MILEPOST GCC software [S9]; most of technology is now available in mainline GCC and is being added to major commercial compilers*
- *cTuning CC software [S8]*
- *cTuning public optimization repository of knowledge [R4]*
- ***All benchmarks, datasets, tools, models and experimental results have been released to public for collaborative validation and extension! See our new research and publication proposal [Pub26]!***
- ***Continues in [M2, Pub50, Pub3]***



- [Pub30] Lianjie Luo, Yang Chen, Chengyong Wu, Shun Long and Grigori Fursin. **Finding representative sets of optimizations for adaptive multiversioning applications.**
3rd Workshop on Statistical and Machine Learning Approaches Applied to Architectures and Compilation (SMART'09), co-located with HiPEAC'09 conference, Paphos, Cyprus, 2009 (acceptance rate=62% (8/13)) (CID=29db2248aba45e59:199e2c0b9d4abc9)



- Extends [Pub23, M6]
- Funded by [F4, F6]

- [Pub31] Grigori Fursin, Cupertino Miranda, Olivier Temam, Mircea Namolaru, Elad Yom-Tov, Ayal Zaks, Bilha Mendelson, Phil Barnard, Elton Ashton, Eric Courtois, Francois Bodin, Edwin Bonilla, John Thomson, Hugh Leather, Chris Williams and Michael O'Boyle. **MILEPOST GCC: machine learning based research compiler.**
Proceedings of the GCC Developers' Summit, Ottawa, Canada, June 2008 (CID=29db2248aba45e59:a43178b20901ba15)



- Implements concept [M4]
- Funded by EU MILEPOST project [J9, F7]
- **Considerably extended in journal version [Pub5]**

- [Pub32] Veerle Desmet, Grigori Fursin, Sylvain Girbal and Olivier Temam. **Leveraging Modular Simulation for Systematic Design Space Exploration.**
4th HiPEAC Industrial Workshop on Compilers and Architectures organized by ARM Ltd., Cambridge, UK, November 2007 (CID=29db2248aba45e59:1ab772451f5f1c5d)



- [Pub33] Grigori Fursin, Cupertino Miranda, Sebastian Pop, Albert Cohen and Olivier Temam. **Practical Run-time Adaptation with Procedure Cloning to Enable Continuous Collective Compilation.**
Proceedings of the GCC Developers' Summit, Ottawa, Canada, June 2008 (CID=29db2248aba45e59:d60d0d44ec4a042e)



- Moving concept of statically enabling dynamic optimizations [M6] to mainline GCC

- [Pub34] Piotr Lesnicki, Albert Cohen, Grigori Fursin, Marco Cornero, Andrea Ornstein and Erven Rohou. **Split Compilation: an Application to Just-in-Time Vectorization.**
International Workshop on GCC for Research in Embedded and Parallel Systems (GREPS'07) co-located with PACT'07, Brasov, Romania, September 2007 (CID=29db2248aba45e59:63c8a602a044dba9)



- [Pub35] Grigori Fursin and Albert Cohen. **Building a Practical Iterative Interactive Compiler.**
1st Workshop on Statistical and Machine Learning Approaches Applied to Architectures and Compilation (SMART'07), colocated with HiPEAC 2007 conference, Ghent, Belgium, January 2007 (acceptance rate=58% (7/12)) (CID=29db2248aba45e59:92b413c92fd1bb6d)



- In this paper I introduced a novel concept to convert hardwired blackboxed production compilers into interactive research toolsets [M7]
- Extended journal version in [Pub5]
- Associated public software [S16, S9]
- **Discontinued for OpenME interface [S3, Pub50]**
- **Now available in mainline GCC since version 4.6 [F4]**

- [Pub36] Shun Long and Grigori Fursin. **A Heuristic Search Algorithm Based on Unified Transformation Framework.**
Proceedings of the 7th International Workshop on High Performance Scientific and Engineering Computing (HPSEC), pages 137-144, Oslo, Norway, June 2005 (CID=29db2248aba45e59:366084436b9f2df5)



- Our first experiments on program autotuning using polyhedral transformations (iterative optimization using polyhedral model)

[Pub37] Grigori Fursin, Michael O'Boyle and Peter Knijnenburg. **Evaluating Iterative Compilation.**
Proceedings of the 15th Workshop on Languages and Compilers for Parallel Computing (LCPC), pages 305-315, College Park, MD, USA, 2002 (CID=29db2248aba45e59:71496986d1f8072c)



- *In this paper I introduced a concept of empirical optimization for large applications to automatically adapt them to a given hardware using several basic search strategies including random and hill-climbing. This approach considerably outperformed state-of-art compilers on Intel, Alpha and several other popular architectures for several large SPEC applications. This technique has also laid foundations for further research on systematic program and architecture optimization and co-design using statistical analysis, machine learning and run-time adaptation [M5, M4, M2]*
- *Associated public software [S18]*

[Pub38] Grigori Fursin, Michael O'Boyle, Olivier Temam and Gregory Watts. **Fast and Accurate Evaluation of Memory Performance Upper-Bound.**
Proceedings of the 9th Workshop on Compilers for Parallel Computers (CPC), pages 163-172, Edinburgh, UK, 2001 (CID=29db2248aba45e59:91c6126884641b59)



- *In this paper I introduced a novel, simple and fast approach to detect program performance anomalies or CPU/memory bounds via semantically non-equivalent assembler patching [M8]. We add or remove various assembler instructions to convert array accesses to scalars in various ways without preserving the semantics of the code while avoiding code crashing to be able to directly compare original and transformed programs. This technique does not need any slow simulation and proved to be realistic particularly on out-of-order processors where hardware counters can be totally misleading. This technique also advise how to optimize code, i.e. if code is CPU bound, we should focus on ILP optimizations; while if the code is memory bound, we should focus on polyhedral transformations or reduce processor frequency to save power.*
- *Associated public software [S18]*
- *Extended in journal version [Pub9], PhD thesis [Pub1] and DECAN framework [Pub45]*

24.5 NATIONAL CONFERENCES AND WORKSHOPS

[Pub39] Grigori Fursin. **Measurement of characteristics of neural elements with the aid of personal computer.**
Proceedings of the national conference on physical processes in devices of electronic and laser engineering at Moscow Institute of Physics and Technology, Moscow, Russia, 1997 (CID=29db2248aba45e59:154e2c842ea60cb5)



- *Concept [M10, M9]*
- *Associated public software [S20, S19]*
- *Associated hardware [H1]*

[Pub40] Grigori Fursin. **Modeling of processes of learning and recognition in modified neural network.**
Proceedings of the national conference on physical processes in devices of electronic and laser engineering at Moscow Institute of Physics and Technology, Moscow, Russia, 1997 (CID=29db2248aba45e59:5fe68265adaf4b24)



- *Concept [M10, M9]*
- *Associated public software [S20, S19]*

[Pub41] Grigori Fursin. **Restoration of symbols with noise by neural network.**
Proceedings of the national conference on physical processes in devices of electronic and laser engineering at Moscow Institute of Physics and Technology, pages 112-117, Moscow, Russia, 1995 (CID=29db2248aba45e59:ecb835592497469e)



- *Concept [M10, M9]*
- *Associated public software [S20, S19]*

24.6 POSTERS

- [Pub42] Grigori Fursin. **Poster: cTuning.org: novel collaborative methodology, publication model, framework and repository to crowdsource autotuning.**
HiPEAC conference poster session, Berlin, Germany, January 2013
(*cid=29db2248aba45e59:db899893d4b71ba7*)



- [Pub43] Grigori Fursin. **cTuning.org: novel extensible methodology, framework and public repository to collaboratively address Exascale challenges.**
Poster at SuperComputing Companion (SC), pages 1403, Salt Lake City, Utah, USA, 2012
(*cid=29db2248aba45e59:54ac7ceace451a83*)



- [Pub44] Grigori Fursin. **Poster: cTuning.org: Collaborative initiative to create open-source repository and tools to share and reuse knowledge about designs and optimizations of computer systems.**
HiPEAC conference poster session, Paris, France, January 2012 (*cid=29db2248aba45e59:60ac26514a4499ed*)



- [Pub45] Souad Koliai, Grigori Fursin, Tipp Moseley and William Jalby. **DECAN: Decremental Performance Analysis Tool via Binary Patching.**
Poster at the Workshop on Languages and Compilers for Parallel Computing (LCPC), USA, 2010
(*cid=29db2248aba45e59:2a45819c829656cc*)



- *Extension of publication [Pub9]*
- *Extension of concept [M8]*
- *Discontinued for Alchemist plugin within Collective Mind Framework [S2, S4]*

- [Pub46] Grigori Fursin, Cupertino Miranda, Olivier Temam, Mircea Namolaru, Elad Yom-Tov, Ayal Zaks, Bilha Mendelson, Phil Barnard, Elton Ashton, Eric Courtois, Francois Bodin, Edwin Bonilla, John Thomson, Hugh Leather, Chris Williams and Michael O'Boyle. **MILEPOST GCC: machine learning based research compiler.**
Poster at ACACES summer school, Italy, 2008 (*cid=29db2248aba45e59:4f7549e2a7c1c8a3*)



- [Pub47] Grigori Fursin, Cupertino Miranda, Sebastian Pop and Albert Cohen. **Enabling Interactivity in GCC for Fine-Grain Optimizations.**
Poster at ACACES summer school, Italy, 2007 (*cid=29db2248aba45e59:0684ae3aa9f423de*)



- [Pub48] Grigori Fursin. **Poster: Iterative Compilation and Performance Prediction.**
Division of Informatics, University of Edinburgh, UK, 2003 (*cid=29db2248aba45e59:4b4079a332016220*)



24.7 TECHNICAL REPORTS, NEWSLETTERS, EXTENDED ABSTRACTS, INTRODUCTIONS

- [Pub49] Grigori Fursin. **Crowdsourcing autotuning: challenges and possible solutions.**
Extended abstract at Dagstuhl Seminar 13401, Dagstuhl, Germany, January 2014
(*cid=29db2248aba45e59:cd75d4dd0a5957aa*)



[Pub50] Grigori Fursin. **Collective Mind: cleaning up the research and experimentation mess in computer engineering using crowdsourcing, big data and machine learning.**

INRIA technical report HAL-00850880, France, 2013 (CID=29db2248aba45e59:1943b3f46fabae4)



- **Extended journal version: [Pub3]**
- *This work summarizes my long-term vision on collaborative, systematic and reproducible benchmarking, optimization and co-design of computer systems across all software and hardware layers using public Collective Mind repository of knowledge, common plugin-based autotuning framework, big data, predictive analytics (machine learning, data mining, statistical analysis, feature detection), crowdsourcing and collective intelligence*
- **This work extends my previous article [P24]**
- *Should be publicly available at some point in autumn, 2014*
- *Related Collective Mind infrastructure and repository [S2]*
- **This work supports my initiative on open research and publication model where all experimental results and related material is continuously shared, validated and improved by the community [Pub26]. To set up an example, I continue sharing all benchmarks, datasets, tools, models and experimental results in Collective Mind repository (c-mind.org/repo)**

[Pub51] Grigori Fursin. **HiPEAC Thematic Session on "Making Computer Engineering a science": cleaning up the mess.**

HiPEAC newsletter 35, 2013 (CID=29db2248aba45e59:bfb3134d4c078236)



- *Introducing open-source Collective Mind Framework (plugin-based knowledge management system) [S2] and public repository [R2] to start collaborative systematization of computer engineering and initiate new publication model where all research artifacts are shared, validated and extended by the community [M2]*

[Pub52] Christophe Dubach and Grigori Fursin. **Introducing 3rd International Workshop on Adaptive Self-Tuning Computing Systems.**

ACM Digital Library, 2013 (CID=29db2248aba45e59:c574c168b737dd92)



[Pub53] Grigori Fursin. **cTuning.org: Novel Extensible Methodology, Framework and Public Repository to Collaboratively Address Exascale Challenges.**

Extended abstract at SuperComputing Companion (SC), pages 1401-1402, Salt Lake City, Utah, USA, November 2012 (CID=29db2248aba45e59:241959a66d0e6895)



[Pub54] Grigori Fursin, Yuriy Kashnikov, Jason Mars and Robert Hundt. **Introducing 2nd International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era (EXADAPT).**

ACM Digital Library, 2012 (CID=29db2248aba45e59:72986f50f464ba31)



[Pub55] Grigori Fursin, Yuriy Kashnikov, Jason Mars and Robert Hundt. **Introducing ACM SIGPLAN International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era (EXADAPT).**

ACM Digital Library, 2011 (CID=29db2248aba45e59:f8c881aed904c178)



[Pub56] Grigori Fursin, Michael O'Boyle, Olivier Temam and Gregory Watts. **A Fast and Accurate Evaluation of Memory Performance Upper-Bound.**

ESPRIT project No 24942 technical report, 2001 (CID=29db2248aba45e59:48d8a274c1c9dea9)



- *Part of EU FP5 MHAOTEU project [J14]*
- *Associated public software [S18]*

[**Pub57**] Jaume Abella, Grigori Fursin, Antonio Gonzalez, Joseph Llosa, Michael O'Boyle, Abhishek Prabhat, Olivier Temam, Sid Touati, Xavier Vera and Gregory Watts. **Advanced Performance Analysis.**
MHAOTEU ESPRIT project No 24942 technical report M3.D2, 2001
(*cid=29db2248aba45e59:47759c6d5adb7b84*)



- *Part of EU FP5 MHAOTEU project [J14]*
- *Associated public software [S18]*

[**Pub58**] Jaume Abella, Cedric Bastoul, Jean-Luc Bechenec, Nathalie Drach, Christine Eisenbeis, Paul Feautrier, Björn Franke, Grigori Fursin, Antonio Gonzalez, Toru Kisku, Peter Knijnenburg, Joseph Llosa, Michael O'Boyle, Julien Sebot and Xavier Vera. **Guided Transformations.**
MHAOTEU ESPRIT project No 24942 technical report M3.D2, 2001
(*cid=29db2248aba45e59:a9874a698e3c2d58*)



- *Part of EU FP5 MHAOTEU project [J14]*
- *Associated public software [S18]*

24.8 PRESS RELEASES

[**Pub59**] **World's First Intelligent, Open Source Compiler Provides Automated Advice on Software Code Optimization (IBM Research and European Union Provide Software Developers with Performance Gains and Faster Time-To-Market).**
IBM MILEPOST project press release, Haifa, Israel and Armonk, NY, USA, June 2009
(*cid=29db2248aba45e59:065ef420f809ff7b*)



Powered by Collective Knowledge