

# CURRICULUM VITAE

**Name:** Grigori Fursin, Ph.D.  
**Current job:** Project Leader, EXATEC Lab (Intel / CEA / GENCI / UVSQ)  
**Address:** Batiment Descartes, UVSQ  
45 avenue des Etats-Unis, 78000 Versailles, France  
**E-mail** [Grigori.Fursin@unidapt.org](mailto:Grigori.Fursin@unidapt.org)  
**Website:** <http://unidapt.org/people/gfursin>  
**Birthday:** 1977/07  
**Languages:** English, Russian, French (beginner)



## SUMMARY

---

I enjoy leading challenging, innovative, collaborative and interdisciplinary R&D projects (more than 15 years of experience). My long term goal is to revisit current over-complicated, under-performing and power-hungry computing technology, and develop self-tuning, adaptive computing systems based on interdisciplinary research techniques including empirical automatic code/ architecture optimization and statistical, collective and machine learning techniques. Automating code, compiler and architecture design and optimization will help to improve performance, power consumption and reliability of computing systems, reduce their development cost and time to market. I hope that this in turn will boost innovation and research in bio-informatics, medicine, physics, aeronautics, robotics, chemistry, finances, image/video processing and other important areas.

I started moving empirical automatic performance tuning, iterative optimization (compilation), statistical and machine learning techniques to real production compilers and computing systems, and started the whole system parameterization to enable future self-tuning, adaptive computing systems. I designed Collective Tuning Center (<http://cTuning.org>) and led the development of the first publicly available machine learning enabled self-tuning compiler based on automatic correlation between program features, run-time behavior and optimizations (MILEPOST GCC, <http://cTuning.org/milepost-gcc> and cTuning CC, <http://cTuning.org/ctuning-cc>) that currently improves program execution time, code size, compilation time, power consumption and enabling run-time adaptation. Though it is very basic, it opens up many novel research opportunities for the community.

I developed multiple practical techniques based on empirical optimizations, statistical analysis and machine learning to enable fast dynamic iterative feedback-directed compilation, self-tuning adaptive applications, predictive parallelization, run-time adaptation and scheduling for heterogeneous multi-core architectures, collective optimization to leverage experience of multiple users, systematic architecture and compiler design space exploration, simple program and architecture characterization based on reaction to optimizations.

These techniques and tools have been used and extended in multiple international research projects supported by generous grants, used in practice in industry and included in real production compilers including GCC. They have been included in the HiPEAC 2012-2020 research roadmap (I have been reviewing and writing parts on automatic self-adaptation, split-compilation and collective optimization), referenced in patents and conferences/journals such as MICRO, PLDI, ASPLOS, CGO, PACT, and in press-releases from IBM. From time to time, I give invited talks about my techniques and tools at companies and universities including Intel, IBM, UIUC, ICT, etc.

In my spare time, I enjoy thinking about emerging technologies and various intelligent systems, and interested in the knowledge transfer, consulting and startups. In my leisure time, I enjoy traveling, playing football, learning to play guitar, participating in community activities or struggling to learn French.

**Jobs:**

- 2010 - cur: Project leader, EXATEC Lab (Intel/CEA/GENCI/UVSQ), France.
- 2007 - 2010: Tenured research scientist at INRIA Saclay (France), UNIDAPT Group founder (<http://unidapt.org>) and Collective Tuning Center founder / R&D coordinator (<http://cTuning.org>) (*currently on sabbatical*)
- 2006-2009: R&D coordinator of the MILEPOST project (<http://cTuning.org/project-milepost>, <http://cTuning.org/milepost-gcc>)
- 2000-2005: Research associate at the University of Edinburgh (UK)

**Education:**

- 2005 - 2007: Postdoctoral researcher at INRIA Futurs (France)
- 1999 - 2004: Ph.D. from the University of Edinburgh with ORS award (UK)
- 1993 - 1999: B.S & M.S. (physics, electronics, computer engineering) with highest honors from MIPT (Russia)

**Main R&D:**

- Collective Tuning Center (cTuning.org) - community-driven collaborative center to share optimization cases and develop common R&D tools with common APIs for the future smart self-tuning adaptive systems (including programs, compilers, run-time systems, architectures) based on empirical compilation and performance evaluation, statistical collective optimization and machine learning.
- cTuning CC - cTuning compiler collection to enable transparent collective optimization.
- MILEPOST GCC - first open source machine learning based research compiler that revolutionized code and architecture design and optimization.
- Interactive Compilation Interface - high-level research API and event-driven plugin system to “hijack” production compilers and transform them into extensible interactive research optimization toolsets. We participated in discussions on GCC low-level plugin framework and added some of the ICI to the mainline GCC 4.5.
- CCC and UNIDAPT framework - extensible framework to automate iterative compilation and enable run-time adaptation for statically compiled programs.
- cBench/MiDataSets - collection of open source programs and datasets for realistic research on program optimization and performance evaluation.

**Recent teaching activities:**

- 2008/2009: M2R Course organizer (Future Computing Systems) at LRI, Paris South University, France and teaching part of the course on adaptive self-tuning systems.

**Service:**

- **CGO'11** PC member, **GROW'10** co-organizer, **SMART'10** co-organizer
- **ICPADS'09** PC member, **iWART'09** PC member, **GROW'09** PC member
- **SMART'09,08,07** organizer and PC member, **IPDPS'08** PC member
- **CASES'07** PC member, **HiPEAC GCC Tutorial'07** co-organizer

**Collaboration:**

- Intel, GENCI, CEA, CAPS Enterprise
- UVSQ (France), University of Edinburgh (UK), Imperial College (UK), UPC (Spain), ICT (China), UIUC (USA), Ghent University (Belgium)

**More information about my research activities is available at:**

- <http://fursin.net/research> , <http://unidapt.org> , <http://cTuning.org>

## PROFESSIONAL EXPERIENCE

---

- 03/2010-  
cur. - Project leader at the new EXATEC Lab in France (Intel Labs Europe/CEA/GENCI/UVSQ) (*currently, on sabbatical from INRIA*). Further information will be available through the official communication channels of the Center.
- 09/2007-  
02/2010 - Tenured research scientist and UNIDAPT Group founder (<http://unidapt.org>) at INRIA Saclay, France (*currently on sabbatical*). I continue research to automate code, compiler and architecture design and optimization based on collective optimization, iterative compilation, run-time program adaptation, machine learning (predictive modeling and reinforcement learning), auto-parallelization, architecture design space exploration, hardware/software co-design, performance prediction and low power optimization techniques. All tools and techniques are now publicly available at <http://cTuning.org> in enable collaborate R&D for the future smart self-tuning computing systems, particularly in the presence of rapidly evolving multi-core heterogeneous architectures.
- 12/2005-  
08/2007 - Postdoctoral researcher in the Alchemy group at INRIA Futurs, France. I had been working with Prof. Olivier Temam to develop a new concept to *enable continuous run-time optimization and adaptation for statically compiled programs* and to *speed up iterative compilation* by 3 orders of magnitude using a *run-time low-overhead program phase detection scheme and function versioning*. At the same time, I started developing an *Interactive Compilation Interface (ICI)* for Open64/PathScale compilers and GCC to create *self-tuning intelligent compilers and systems*.
- We also developed a novel technique to *characterize programs or architectures using program reaction to optimizations* (transformations). I had been collaborating with my colleagues from the University of Edinburgh, UK to *introduce this technique as well as statistical search and machine learning* to enable *optimization knowledge reuse* among different programs and architectures using static and dynamic program and architecture features.
- These techniques are now used in multiple international research projects and are referenced in CGO, PLDI, PACT and other publications and IBM patent application.
- 01/2002-  
11/2005 - Research associate at the Institute for Computing Systems Architecture, University of Edinburgh, UK. I had been working with Prof. Michael O'Boyle to *introduce iterative compilation at a fine-grain level* (function, loop or instruction) to *automatically find best optimization settings for large applications* (rather than kernels) on *rapidly evolving architectures that beat state-of-the art optimizing compilers*. I had also been working with Prof. Olivier Temam to develop a *fast and accurate technique* to determine *lower bound of the execution time* of memory intensive applications by replacing all array accesses with scalars to have a stopping criterion for iterative compilation.

## PROFESSIONAL EXPERIENCE

---

- 02/2000-03/2000 - Visiting researcher at Paris-Sud XI University, France. I had been working in the Alchemy group with Prof. Olivier Temam to develop a fast and accurate technique to predict lower bound of the execution time of memory intensive applications by replacing all array accesses with scalars to have a stopping criterion for iterative compilation.
- 02/1999-12/2001 - Research assistant at the University of Edinburgh, UK in the European Project "MHAOTEU" (Memory Hierarchy Analysis and Optimization Tools for the End-User). I developed source-to-source compiler with several popular transformations, program instrumentation tool (similar to ATOM/PIN), performance prediction tool and iterative compilation framework (Edinburgh Optimizing Software) to automatically optimize memory intensive applications at fine-grain level within EU project MHAOTEU (Memory Hierarchy Analysis and Optimization Tools for the End-User). This software had been used and extended in several M.S. and Ph.D. projects at the University of Edinburgh (UK).
- 09/1998-01/1999 - Programmer at the Laboratory for Computer Technologies in Teaching at Moscow Institute of Physics and Technology, Russia. I developed a software/hardware infrastructure to measure characteristics of semiconductor devices that involved designing and implementing a special DSP board, developing of novel communication protocol and client software for data collection, analysis and visualization. This system is still used to teach undergraduate and postgraduate students during electronics labs at Moscow Institute of Physics and Technology (Russia).
- 02/1998-01/1999 - Research assistant at the Institute for High-Performance Computing in the Russian Academy of Sciences in the project "Remote access to high-performance computer systems through the Internet". I suggested to access heterogeneous high-performance computers through unified Internet services and dynamically adapt/manage/balance their resources based on user requirements and applications. I implemented the first prototype but due to lack of interest from my scientific advisers, I decided to peruse PhD at the University of Edinburgh. Interestingly, some similar ideas can now be found in cloud computing technology.
- 01/1994-06/1994 - Research assistant at Moscow Institute of Physics and Technology, Russia in the project "Computer simulation of non-linear wave processes in gaseous streams". I developed and optimized simulation and visualization framework in the research project "Computer simulation of non-linear wave processes in gaseous streams" at Moscow Institute of Physics and Technology (Russia).

## EDUCATION

---

- 2008 - Attended ACACES'2008 - International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems (with courses by Josep Torrellas, Dean Tulsen, Babak Falsafi and Christos Kozyrakis).
- 2006 - Attended ACACES'2006 - International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems (with courses by Wenmei Hwu, David Padua, David Whalley and Michael Hind).
- 2005 - Attended ACACES'2005 - First International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems (with courses by Josh Fisher, Ayal Zaks, Trevor Mudge and Rajiv Gupta).
- 1999-2004 - Received Ph.D. degree (computer science) from the University of Edinburgh, UK, advisor: Prof. Michael O'Boyle, thesis title: "Iterative Compilation and Performance Prediction for Numerical Applications", first advisor (before moving to industry): Prof. Nigel Topham.

We introduced *iterative compilation (empirical iterative feedback-directed optimization)* at a *fine-grain level* (function, loop or instruction) to automatically find best optimization settings for *large applications* (rather than kernels) on rapidly evolving architectures that beat state-of-the art optimizing compilers. We also developed a fast *and accurate technique* to determine *lower bound of the execution time* of memory intensive applications by replacing all array accesses with scalars to have a stopping criterion for iterative compilation [FOTP2001, FOK2002, FOTP2004, FUR2004].

These techniques have been used and extend in multiple international projects to automate code and architecture design and optimization based on machine learning and statistical techniques (in collaboration with the University of Edinburgh, INRIA, Ghent University, ICT, IBM, ARC, CAPS Enterprise, NXP and STMicro among others).

- 1997-1999 - Received M.S. degree with medal from Moscow Institute of Physics & Technology, Russia (High-Performance Computing, GPA=4.00/4.00). I suggested to access heterogeneous high-performance computers through unified Internet services and dynamically adapt/manage/balance their resources based on user requirements and applications. I implemented the first prototype but due to lack of interest from my scientific advisers, I decided to peruse PhD at the University of Edinburgh. Interestingly, some similar ideas can now be found in cloud computing technology.
- 1996-1999 - Had research practice at the Institute for High-Performance Computer Systems of Russian Academy of Sciences.
- 1993-1997 - Received B.S. degree with highest honors (summa cum laude) from Moscow Institute of Physics & Technology, Russia (Department of Physical & Quantum Electronics, GPA=3.98/4.00).
- 1990-1993 - Finished Moscow Physical & Technical College, Russia (GPA=4.00/4.00).
- 1983-1993 - Finished Moscow Secondary School, Russia with medal.

## PROFESSIONAL ACTIVITIES (PC MEMBER; EVENT ORGANIZER)

---

- 2010
  - GROW'10: co-organizer and editor of informal proceedings (2<sup>nd</sup> International Workshop on GCC Research Opportunities)  
<http://cTuning.org/workshop-grow10>
  - SMART'10: co-organizer and editor of informal proceedings (4<sup>th</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation)  
<http://cTuning.org/workshop-smart10>
  - Speedup-Test'10 tutorial at HiPEAC'10: co-organizer (Statistical Methodology to Evaluate Program Speedups and their Optimization Techniques)  
<http://www.hipeac.net/conference/pisa/speedup>
- 2009
  - iWAPT'09: PC member (International Workshop on Automatic Performance Tuning)
  - ICPADS'09: PC member (International Conference on Parallel and Distributed Systems), multi-core architectures track
  - Open64 Workshop at CGO'09: PC member
  - SMART'09: Organizer and PC member (3<sup>rd</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation)
  - GROW'09: PC member (1<sup>st</sup> International Workshop on GCC Research Opportunities)
- 2008
  - IPDPS'08: PC member (IEEE International Parallel and Distributed Processing Symposium)
  - SMART'08: PC member (2<sup>nd</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation)
- 2007
  - CASES'07: PC member (International Conference on Compilers, Architecture, and Synthesis for Embedded Systems)
  - 2<sup>nd</sup> HiPEAC GCC Tutorial: organizer (<http://www.hipeac.net/gcc-tutorial>)
  - SMART'07: chair and organizer (1<sup>st</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation  
<http://www.hipeac.net/smart-workshop.html>).

*In recent years, we have shown that machine learning and statistical search techniques can be effectively used to create and optimize production compilers (such as MILEPOST GCC) and architectures. Therefore, together with my colleagues, I decided to organize this workshop to promote new ideas and to present recent developments in compiler and architecture design using machine learning, statistical approaches, and search in order to enhance their performance, scalability, and adaptability.*
- 2001
  - CPC'01: local organizer (9<sup>th</sup> Workshop on Compilers for Parallel Computers)

## PROFESSIONAL ACTIVITIES (TEACHING)

---

- 2008/2009 - Future Computing Systems - M2R Course organizer and teaching part of the course on adaptive self-tuning systems (LRI, Paris South University, France)
- 2007/2008 - Continuous adaptive iterative compilation and machine learning techniques (part of the postgraduate course at LRI, Paris South University, France)
- 2006/2007 - Continuous adaptive iterative compilation and machine learning techniques (part of the postgraduate course at LRI, Paris South University, France)
- 2005/2006 - Adaptive and feedback driven compilation (part of the postgraduate course at LRI, Paris South University, France)

## PROFESSIONAL ACTIVITIES (REVIEWING)

---

### Conferences, workshops and books:

- HiPEAC research vision (writing and reviewing): <http://www.hipeac.net/roadmap>
- Springer-Verlag book on autotuning techniques (on-going)
- CGO, PACT, TPDS, CASES, ICS, DATES, HiPEAC, CPC, IJHPSA, ICPADS, JPDC, CLSS, iWAPT, Open64, SMART, GROW

## PROFESSIONAL ACTIVITIES (MEMBERSHIPS)

---

- HiPEAC member
- ACM member (SIGARCH - Special Interest Group on Computer Architecture)
- IEEE member (Computer Society)

# PROFESSIONAL ACTIVITIES (COLLABORATIONS)

---

## Industry:

- Intel
- CEA
- GENCI
- IBM
- CAPS Enterprise
- STMicroelectronics
- Thales
- NXP
- AMD
- ARC

## Academia:

- Université de Versailles-Saint-Quentin-en-Yvelines, France
- University of Edinburgh, UK
- Imperial College, UK
- INRIA, France
- UPC, Spain
- UIUC, USA
- Ghent University, Belgium
- ICT, China

## PROFESSIONAL ACTIVITIES (STUDENTS, POSTDOCS)

---

- 2010 - Advising 2 PhD students from Université de Versailles-Saint-Quentin-en-Yvelines, France: Yuriy Kashnikov and Abdul Wahid Memon.
- 2009 - Mentoring 2 Google Summer of Code'2009 students from ICT, China (Yuanjie Huang and Liang Peng) to develop run-time adaptation techniques, iterative compilation and machine learning at fine-grain level and enable collective optimization in production compilers such as GCC [HPWP2010, FT2009, FCOP2005] (<http://socqhop.appspot.com/org/home/google/qsoc2009/qcc>).
- Advising Abdul Wahid Memon and Menjato Rakoto (M2R students from Paris South University) during their internship at my group - we are extending cTuning tools and techniques to automate program optimization using machine learning (<http://cTuning.org>).
- Collaborating with Dr. Cosmin Oancea (postdoctoral researcher from Cambridge University) to develop software TLS tools and techniques to improve automatic program parallelization.
- 2008 - Advised Lianjie Luo (M.S. student at ICT, China) to develop self-tuning applications and libraries based on machine learning techniques using Open64 with ICI and CCC framework [LCWP2009].
- Advised Abid Malik (postdoc at INRIA, France) to improve GCC ICI and investigate abilities for pass reordering and fine-grain optimizations using machine learning and statistical techniques.
- 2007 - Advised Victor Jimenez (Ph.D. student at UPC, Barcelona) to develop predictive run-time adaptation and scheduling techniques for heterogeneous computing systems with multiple ISA [JGVP2009].
- 2006 – 2007 - Advised Piotr Lesnicki (Ph.D. student at Paris-South University, France) to develop automatic adaptive split-compilation techniques (Java and .NET) on multi-core embedded systems [LCFP2007].
- 2006 - Advised Hamid Daoud (M.S. student at Paris-South University, France) to tune GCC loop unrolling optimization heuristic using ICI combined with machine learning techniques available in WEKA.
- Advised Cupertino Miranda (engineer at INRIA, France) to implement new version of Interactive Compilation Interface (GCC high-level plugin framework) to enable research on fine-grain program optimizations and to start implementing run-time adaptation technique for statically compiled programs with varying context [FMPP2007].
- 2003-2004 - Advised Edwin Bonilla (M.S. student at the University of Edinburgh, UK) to combine iterative compilation and machine learning at loop level for his M.S. thesis "Predicting Good Compiler Transformations Using Machine Learning"

## MAJOR R&D PROJECTS

---

2008-cur - **Collective Tuning Center (cTuning) / Collective Optimization Infrastructure:**  
Since 1997, I developed multiple research techniques and tools to automate code and architecture design and optimization using empirical iterative feedback-directed compilation and novel techniques based on statistical collective optimization and machine learning (thanks to many fruitful collaborations). Since we obtained very promising results, I decided to move all my developments to public domain at [cTuning.org](http://cTuning.org) to enable further collaborative community-based research and development.

The core of cTuning infrastructure is a Collective Optimization Database that allows easy collection, sharing, characterization and reuse of a large number of optimization cases from the community. The infrastructure also includes collaborative R&D tools with common API (CCC framework, MILEPOST GCC with ICI and static feature extractor, cBench and UNIDAPT framework) to automate optimization, produce adaptive applications and enable realistic benchmarking. I developed several tools and open web-services to substitute default compiler optimization heuristic and predict good optimizations for a given program, dataset and architecture based on static and dynamic program features and standard machine learning techniques.

cTuning infrastructure provides a novel fully integrated, collaborative, "one button" approach to improve existing underperforming computing systems ranging from embedded architectures to high-performance servers based on systematic iterative compilation, statistical collective optimization and machine learning. Our experimental results show that it is possible to reduce execution time (and code size) of some programs from SPEC2006 and EEMBC among others by more than a factor of 2 automatically. It can also reduce development and testing time considerably. Together with the first production quality machine learning enabled interactive research compiler (MILEPOST GCC) this infrastructure opens up many research opportunities to study and develop future realistic self-tuning and self-organizing adaptive intelligent computing systems based on systematic statistical performance evaluation and benchmarking. Finally, using common optimization repository is intended to improve the quality and reproducibility of the research on architecture and code optimization of my research and developments,

*Website(s):*

- Main: <http://cTuning.org>
- COD: <http://cTuning.org/cdatabase>
- Optimization prediction web service: <http://cTuning.org/cpredict>

*Publications:* Fur2009, FT2009

*IBM press-release:* <http://www-03.ibm.com/press/us/en/pressrelease/27874.wss>

2010-cur - **cTuning Compiler Collection (cTuning CC):** cTuning CC is a free, open source compiler collection that combines multiple tools and techniques (MILEPOST GCC, ICI, CCC framework, cTuning web-services and cBench) as the first practical step toward self-tuning, adaptive computing systems based on industrial tools, empirical techniques, transparent collective optimization, statistical analysis and machine learning.

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

*Publications:* under preparation

2008-cur - **Collective Benchmark (cBench):** Since there was not enough programs/datasets available in MiBench/MiDataSets, I opened a repository of open-source sequential and parallel programs with multiple datasets assembled by the community to enable realistic benchmarking and research on program and architecture optimization. The source code of individual programs is simplified to ease portability. All the benchmarks now include scripts to be used with CCC Framework to perform automatic optimizations (iterative compilation) using GCC, LLVM, GCC4CIL, Open64, PathScale, Intel and other compilers on a wide range of architectures.

*Website:* <http://cTuning.org/cbench>

*Publications:* FCOP2007, Fur2009, YYLP2010

2006-cur - **MILEPOST GCC:** I have been coordinating the development of the first machine learning enabled open-source self-tuning research compiler that can adapt to any architecture using iterative feedback-directed compilation, machine learning and collective optimization (in collaboration with the MILEPOST consortium). It combines the strength of the production quality GCC that supports more than 30 families of architectures and can compile real, large applications including Linux, and the flexibility of the Interactive Compilation Interface that transforms GCC into a research compiler. It is currently based on predictive modeling using program and machine-specific features, execution time, hardware counters and off-line training. MILEPOST GCC includes static program feature extractor developed by IBM Haifa. It automatically adjusts its optimization heuristics to improve execution time, code size, or compilation time of a given program on a given architecture. MILEPOST/cTuning technology is orthogonal to GCC and can be used in any future adaptive self-tuning compiler using common Interactive Compilation Interface. It opens up many research possibilities and is a part of the community-driven cTuning infrastructure.

*Website:* <http://cTuning.org/wiki/index.php/CTools:MilepostGCC>

*Publications:* FMTP2008, Fur2009, ABCP2006

*IBM press-release:* <http://www-03.ibm.com/press/us/en/pressrelease/27874.wss>

**MiDataSets for MiBench:** Together with Prof. Olivier Temam, we prepared multiple datasets for embedded benchmarks to enable systematic and realistic research on empirical program optimization, run-time adaptation and benchmarking.

*Website:* <http://cTuning.org/cbench>

*Publications:* FCOP2007, Fur2009, FT2009, YYLP2010

2004-cur - **Interactive Compilation Interface (ICI):** To support my PhD and postdoctoral research, I started developing an Interactive Compilation Interface (ICI) and an event-driven plug-in system to transform production quality compilers into open interactive collaborative research toolsets. I needed such system to avoid developing optimization infrastructure from scratch and to enable rigorous, systematic and statistical performance evaluation of iterative feedback-directed compilation and collective optimization. This tool supports my research to automate code and architecture design and optimization and to develop smart self-tuning adaptive computing systems based on empirical iterative compilation, statistical analysis and machine learning.

Originally I developed ICI for Open64/PathScale compilers but later moved to GCC since it is a unique production quality compiler that supports more than 30 families of architectures and multiple languages. GCC with ICI is now used in multiple international research projects and supported by the HiPEAC network of excellence. Since 2007, we participated in negotiations to implement similar plugin system in mainline GCC. It resulted in a collaborative development of the plugin framework for GCC 4.5. Some of the ICI-related events will also be in the mainline GCC 4.5.

*Website:* <http://cTuning.org/ici>

*Publications:* FCOP2005, FC2007, FMTP2008, LCWP2009, Fur2009, HPWP2010

*Older Open64/PathScale ICI developments:* <http://open64-ici.sourceforge.net>

**UNiversal aDAPTation Framework (UNIDAPT):** During my postdoctoral research in the Alchemy group at INRIA Futurs (France) I developed a new concept to enable collective optimization [FT2009] and speed up iterative feedback-directed compilation by several orders of magnitude [FCOP2005] based on continuous transparent run-time program optimization and adaptation. This technique enabled simple dynamic optimizations for statically compiled programs by cloning hot functions, applying aggressive optimizations for different optimization cases (performance/power/fault-tolerance, etc), using run-time low-overhead monitoring of program behavior (hardware counters) and learning reaction of different clones to different program behavior continuously. This technique opened up many research possibilities and has been used in multiple research projects in collaboration with UPC, ICT, IBM, CAPS Enterprise, STMicro and has been supported by MILEPOST, HiPEAC and Google Summer of Code grants.

*Website:* <http://cTuning.org/unidapt>

*Publications:* FCOP2005, FMPP2007, LFF2007, FT2009, LCWP2009, JGVP2009, Fur2009, FT2009, HPWP2010

- **Continuous Collective Compilation Framework (CCC):** To support my research on systematic program and architecture design and optimization, I started developing a collaborative modular plugin-enabled R&D infrastructure to automate iterative feedback-directed compilation and architecture design space exploration, distribute optimization among multiple users and gather static and dynamic optimization profile data in the Collective Optimization Database. It supports optimizations at global/coarse-grain and finer grain level if compiler supports ICI. It can help end-users optimize their programs, libraries and whole OS automatically (improve execution time/code size, etc) or test compilers. This framework has been used in multiple international research projects including MILEPOST project to automate generation of training sets to substitute default compiler optimization heuristic with a predictive optimization plugin.

*Website:* <http://cTuning.org/ccc>

*Publications:* FCOP2005, FCOP2007, CFAP2007, FMTP2008, LCWP2009, Fur2009,YYLP2010

*Older developments:*

- EOS: <http://fursin.net/wiki/index.php5?title=Research:Developments:EOS>
- FCO: <http://fursin.net/wiki/index.php5?title=Research:Developments:FCO>

- 1999-2004 - I developed source-to-source compiler with several popular transformations, program instrumentation tool (similar to ATOM/PIN), performance prediction tool and iterative compilation framework (Edinburgh Optimizing Software) to automatically optimize memory intensive applications at fine-grain level within EU project MHAOTEU (Memory Hierarchy Analysis and Optimization Tools for the End-User). This software had been used and extended in several M.S. and Ph.D. projects at the University of Edinburgh (UK).

*Publications:* FOK2002, FCOP2004, FUR2004

- 1998-1999 - I developed portable distributed client/server framework to enable simple remote access to heterogeneous high-performance computers as Internet web-service at Russian Academy of Sciences (similar technology is now used in GRID and cloud computing).

- 1996-1998 - I developed a software/hardware infrastructure to measure characteristics of semiconductor devices that involved designing and implementing a special DSP board, developing of novel communication protocol and client software for data collection, analysis and visualization. This system is still used to teach undergraduate and postgraduate students during electronics labs at Moscow Institute of Physics and Technology (Russia).

- 1994 - I developed and optimized simulation and visualization framework in the research project "Computer simulation of non-linear wave processes in gaseous streams" at Moscow Institute of Physics and Technology (Russia).

## AWARDS, PRESS-RELEASES AND GRANTS

---

- 2010
  - Funding from UVSQ (France) to sponsor Prof. Keith Cooper's keynote speech at SMART'10 workshop co-located with HiPEAC'10 in Pisa, Italy  
<http://cTuning.org/workshop-smart10>
  - HiPEAC paper award for the publication "Portable Compiler Optimization Across Embedded Programs and Microarchitectures using Machine Learning" presented at the 42<sup>nd</sup> IEEE/ACM International Symposium on Microarchitecture (MICRO), New York, USA, December 2009.
- 2009
  - IBM press release about MILEPOST GCC and Collective Tuning Center  
<http://www-03.ibm.com/press/us/en/pressrelease/27874.wss>
  - Google Summer of Code'09 program grants for 2 students (I have been mentoring Yuanjie Huang and Liang Peng from ICT, China to extend GCC ICI plugin framework and enable fine-grain optimization selection and function run-time adaptation based on collective optimization [FT2009] and machine learning [FMTP2008]):  
<http://socghop.appspot.com/org/home/google/gsoc2009/gcc>
- 2008
  - HiPEAC grant for the project "Context-aware optimization and run-time adaptation of sequential libraries for multi-core systems" (I have been the principal investigator)
  - HiPEAC grant for the project "Context-aware optimization and run-time adaptation of sequential libraries for multi-core systems" (I have been the principal investigator)
- 2006-2009
  - EU funding through the SARC project to develop novel scalable and adaptive heterogeneous multi-core computer architecture
- 2006-2007
  - HiPEAC grant for the project "Exploring optimization techniques and runtime code selection mechanisms for heterogeneous systems" (I have been the principal investigator)
- 2005-2006
  - HiPEAC Postdoctoral research grant to collaborate with the Alchemy group at INRIA Futurs, France
- 1999,2000
  - ORS award (UK Scholarship for International Research Students):  
<http://www.orsas.ac.uk/applications>
- 1999
  - Golden medal for M.S. studies from Moscow Institute of Physics & Technology (Russia)
- 1997
  - International George Soros award "In recognition and appreciation of outstanding achievements in the study of science at the university level"
- 1996-1998
  - Research grants from International Soros Science Education Program
- 1993
  - Medal for studies during secondary school

## MAJOR FUNDING FOR MY RESEARCH GROUP

---

2010-2012 - *TBA*

2006-2009 - **650,000 euros** - Funding from EU FP6 program / INRIA for the MILEPOST project (<http://cTuning.org/project-milepost>) to develop the first machine-learning enabled compiler. I have been R&D coordinator and the architect of the MILEPOST GCC / MILEPOST optimization framework. I have been also responsible for recruiting engineers, students, postdocs and interns.

## PUBLICATIONS

---

[YYLP2010]

Yang Chen, Yuanjie Huang, Lieven Eeckhout, Grigori Fursin, Liang Peng, Olivier Temam, Chengyong Wu. **Evaluating Iterative Optimization across 1000 Data Sets.** *Proceedings of the ACM SIGPLAN 2010 Conference on Programming Language Design and Implementation (PLDI 2010)*, June 2010, Toronto, Canada (acceptance rate: 20%, 41/204)

[HPWP2010]

Yuanjie Huang, Liang Peng, Chengyong Wu, Yuriy Kashnikov, Jörn Renneke, and Grigori Fursin. **Transforming GCC into a research-friendly environment: plugins for optimization tuning and reordering, function cloning and program instrumentation.** *2<sup>nd</sup> International Workshop on GCC Research Opportunities (GROW'10) co-located with HiPEAC'10*, Pisa, Italy, January 2010 (acceptance rate: 57%, 8/14)

[FT2009]

Grigori Fursin and Olivier Temam. **Collective optimization.** *Proceedings of the International Conference on High Performance Embedded Architectures & Compilers (HiPEAC 2009)*, Paphos, Cyprus, January 2009 (acceptance rate: 28%, 27/97)

- **Extended version has been accepted for ACM Transactions on Architecture and Code Optimization (TACO).**
- **Concept is included in the HiPEAC 2012-2020 research roadmap.**

[DJBP2009]

Christophe Dubach, Timothy M. Jones, Edwin V. Bonilla, Grigori Fursin, and Michael F.P. O'Boyle. **Portable Compiler Optimization Across Embedded Programs and Microarchitectures using Machine Learning.** *Proceedings of the 42<sup>nd</sup> IEEE/ACM International Symposium on Microarchitecture (MICRO)*, New York, USA, December 2009 (acceptance rate: 25%, 52/209)

- **HiPEAC paper award.**
- **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.**

[TOFP2009]

John Thomson, Michael O'Boyle, Grigori Fursin and Björn Franke. **Reducing Training Time in a One-shot Machine Learning-based Compiler.** *Proceedings of the 22<sup>nd</sup> International Workshop on Languages and Compilers for Parallel Computing (LCPC'09)*, Newark, Delaware, USA, October 2009

[Fur2009]

Grigori Fursin. **Collective Tuning Initiative: automating and accelerating development and optimization of computing systems.** *Proceedings of the GCC Summit'09*, Montreal, Canada, June 2009

[LCWP2009]

Lianjie Luo, Yang Chen, Chengyong Wu, Shun Long and Grigori Fursin. **Finding representative sets of optimizations for adaptive multiversioning applications.** *3<sup>rd</sup> International Workshop on Statistical and Machine Learning Approaches Applied to Architectures and Compilation (SMART'09) co-located with HiPEAC'09*, Paphos, Cyprus, January 2009 (acceptance rate=62%, 8/13)

[JGVP2009]

Victor Jimenez, Isaac Gelado, Lluís Vilanova, Marisa Gil, Grigori Fursin and Nacho Navarro. **Predictive runtime code scheduling for heterogeneous architectures.** *Proceedings of the International Conference on High Performance Embedded Architectures & Compilers (HiPEAC 2009)*, Paphos, Cyprus, January 2009 (acceptance rate: 28%, 27/97)

[LF2009]

Shun Long and Grigori Fursin. **Systematic search within an optimisation space based on Unified Transformation Framework.** *International Journal of Computational Science and Engineering (IJCSSE)*, Vol.4, No.2, pages 102-111, 2009 (submitted in 2005)

[FMTP2008]

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, Michael O'Boyle. **MILEPOST GCC: machine learning based research compiler.** *Proceedings of the GCC Developers' Summit*, Ottawa, Canada, June 2008

[DFGP2007]

Veerle Desmet, Grigori Fursin, Sylvain Girbal and Olivier Temam. **Leveraging Modular Simulation for Systematic Design Space Exploration.** *4<sup>th</sup> HiPEAC Industrial Workshop on Compilers and Architectures organized by ARM Ltd.*, Cambridge, UK, November 2007

[LCFP2007]

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) in conjunction with PACT'07*, Brasov, Romania, September 2007

[LFF2007]

Shun Long, Grigori Fursin, Björn Franke. **A Cost-Aware Parallel Workload Allocation Approach based on Machine Learning Techniques.** *Proceedings of the IFIP International Conference on Network and Parallel Computing (NPC 2007)*, LNCS-4672, pages 506-515, Dalian, China, September 2007

[FMPP2007]

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, July 2007

[DCFP2007]

Christophe Dubach, John Cavazos, Björn Franke, Grigori Fursin, Michael O'Boyle and Oliver Temam. **Enabling fast compiler optimization evaluation via code-features based performance predictor.** *Proceedings of the ACM International Conference on Computing Frontiers*, Ischia, Italy, May 2007 (acceptance rate=50%,28/56)

[CFAP2007]

John Cavazos, Grigori Fursin, Felix Agakov, Edwin Bonilla, Michael F.P.O'Boyle and Olivier Temam. **Rapidly Selecting Good Compiler Optimizations using Performance Counters.** *Proceedings of the 5<sup>th</sup> Annual International Symposium on Code Generation and Optimization (CGO)*, San Jose, USA, March 2007 (acceptance rate=32%, 27/84)

[FC2007]

Grigori Fursin and Albert Cohen. **Building a Practical Iterative Interactive Compiler.** *1<sup>st</sup> International Workshop on Statistical and Machine Learning Approaches Applied to Architectures and Compilation (SMART'07) co-located with HiPEAC'*, Ghent, Belgium, January 2007 (acceptance rate=58%, 7/12)

[FCOP2007]

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 International Conference on High Performance Embedded Architectures & Compilers (HiPEAC 2007)*, Ghent, Belgium, January 2007 (acceptance rate=29%)

[CDAP2006]

John Cavazos, Christophe Dubach, Felix Agakov, Edwin Bonilla, Michael F.P. 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 2006)*, Seoul, Korea, October 2006 (acceptance rate=41%, 41/100)

**(finalist best paper award)**

[FCOP2006]

Grigori Fursin, Albert Cohen, Michael O'Boyle and Oliver Temam. **Quick and practical run-time evaluation of multiple program optimizations.** *Transactions on High-Performance Embedded Architectures and Compilers*, 1(1), pages 13-31, 2006

[ABCP2006]

F. Agakov, E. Bonilla, J. Cavazos, B. Franke, G. Fursin, M.F.P. O'Boyle, J. Thomson, M. Toussaint and C.K.I. Williams. **Using Machine Learning to Focus Iterative Optimization.** *Proceedings of the 4<sup>th</sup> Annual International Symposium on Code Generation and Optimization (CGO)*, New York, NY, USA, March 2006 (acceptance rate=36%, 29/80)

**(best presentation award)**

[FCOP2005]

Grigori Fursin, Albert Cohen, Michael O'Boyle and Oliver Temam. **A Practical Method For Quickly Evaluating Program Optimizations.** *Proceedings of the 1<sup>st</sup> International Conference on High Performance Embedded Architectures & Compilers (HiPEAC 2005)*, number 3793 in LNCS, pages 29-46, Barcelona, Spain, November 2005

**(highest ranked paper, acceptance rate=20%,17/84)**

*This paper presents a novel concept to statically enable run-time optimizations and self-tuning binaries through function cloning and integrated low-overhead program/system behavior monitoring routines. It has been referenced in IBM patent application, in a number of PLDI, CGO and PACT publications and has been extended in a number of international research projects.*

[FOTP2005]

B. Franke, M. O'Boyle, J. Thomson and G. Fursin. **Probabilistic Source-Level Optimisation of Embedded Systems Software**. *Proceedings of the Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'05)*, pages 78-86, Chicago, IL, USA, June 2005 (acceptance rate=26%,25/95)

[LF2005]

Shun Long and Grigori Fursin. **A heuristic search algorithm based on Unified Transformation Framework**. *Proceedings of the 7<sup>th</sup> International Workshop on High Performance Scientific and Engineering Computing (HPSEC-05)*, pages 137-144, Oslo, Norway, June 2005

[FOTP2004]

Grigori Fursin, Mike O'Boyle, Olivier Temam, and Gregory Watts. **Fast and Accurate Method for Determining a Lower Bound on Execution Time**. *Concurrency Practice and Experience*, 16(2-3), pages 271-292, 2004

[FOK2002]

G.G.Fursin, M.F.P.O'Boyle, and P.M.W. Knijnenburg. **Evaluating Iterative Compilation**. *Proceedings of the 15<sup>th</sup> Workshop on Languages and Compilers for Parallel Computing (LCPC'02)*, College Park, MD, USA, pages 305-315, 2002

[FOTP2001]

Grigori Fursin, Mike O'Boyle, Olivier Temam, and Gregory Watts. **Fast and Accurate Evaluation of Memory Performance Upper-Bound**. *Proceedings of the 9<sup>th</sup> Workshop on Compilers for Parallel Computers (CPC'2001)*, pages 163-172, Edinburgh, UK, 2001

[ATAP2000]

Abella, J., S. A. Ali Touati, A. Anderson, C. Ciuraneta, J. M. Codina, Min Dai, C. Eisenbeis, G. Fursin, A. Gonzalez, J. Llosa, M. O'Boyle, A. Randrianatoavina, J. Sanchez, O. Temam, X. Vera, and G. Watts. **MHAOTEU Tools for Memory Hierarchy Management**. *IMACS'2000, 16<sup>th</sup> IMACS World Congress on Scientific Computation, Applied Mathematics and Simulation*, Lausanne, Switzerland, August 2000

## PH.D. THESIS

---

[FUR2004]

Grigori Fursin. **Iterative Compilation and Performance Prediction for Numerical Applications**. *Ph.D. thesis*, University of Edinburgh, Edinburgh, UK, January 2004

## PRESENTATIONS

---

- “Collective Tuning Initiative”

*Presented at the GCC Summit’09, Montreal, Canada, June 2009*

*Presented at the HiPEAC industrial workshop and HiPEAC clusters, Infineon, Munich, Germany, June 2009*

*Presented at the University of Versailles, France, May 2009*

- “Collective Tuning Initiative: collective optimization, run-time adaptation and machine learning”

*Presented at the University of Illinois at Urbana-Champaign, USA, April 2009*

- “Collective Optimization”

*Presented at HiPEAC’09, Cyprus, January 2009*

- “Enabling Dynamic Optimization and Adaptation for Statically Compiled Programs Using Function Multi-versioning”

*Presented at SMART’09, Cyprus, January 2009*

- “MILEPOST project - using machine learning to automate and speed up program optimization for reconfigurable processors”

*Presented at the Information and Brokerage Conference on Information and Communication Technologies in the EU’s 7<sup>th</sup> Framework, Moscow, Russia, October 2008.*

- “Enabling Dynamic Optimization and Adaptation for Statically Compiled Programs Using Function Multi-Versioning”

*Presented at ScalPerf’08 (Scalable Approaches to High Performance and High Productivity Computing), Bertinoro, Italy, September 2008*

- “Continuous adaptive program optimizations”

*Presented at Reservoir Labs and IBM TJ Watson Research Center, New York, USA, August 2008*

*Presented at Imperial College (Software Performance Engineering Laboratory), London, UK, February 2008*

*Presented at the Institute of Computing Technology (Chinese Academy of Sciences), Beijing, China, January 2008*

- “Program iterative continuous optimizations, run-time adaptation and machine learning”

*Presented at IBM Toronto Lab (compiler group), Canada, July 2007*

- “Machine learning techniques for iterative program optimizations and run-time adaptation”

*Presented for TAO group (machine learning group), LRI, Paris-Sud XI University, INRIA & CNRS, France, June 2007*

- “Overview of current activities: Interactive Compilation Interface for fine-grain program optimizations, dataset sensitivity, machine learning to speed up optimizations and DSE, run-time program adaptation, optimizations for heterogeneous computing systems, continuous

collective optimizations, HiPEAC activities”

*Presented at Intel (compiler group), Moscow, Russia, February 2007*

*Presented at the ISP RAS (Institute for System Programming, Russian Academy of Sciences), Moscow, Russia, February 2007*

- “Continuous run-time adaptation and optimization of statically compiled programs”

*Presented at the UPC, Barcelona, Spain, January 2007*

- “Towards continuous collective compilation”

*Presented at the ICSA seminar, University of Edinburgh, UK, December 2006*

- “Continuous adaptive run-time optimizations for scientific applications” and

- “Using machine learning for compiler optimizations”

*Presented at the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, USA, March 2006*

- “Continuous collective compilation for the MilePost project (Machine Learning Techniques for Adaptive Optimization)”

*Presented at the MilePost EC negotiation meeting, European Commission, Brussels, March 2006*

- “A practical method for quickly evaluating program optimizations”

*Presented at the ICSA seminar, University of Edinburgh, UK, December 2005*

- “Advanced iterative compilation and performance prediction for scientific applications”

*Presented at the LRI, Paris-Sud XI University, France, April 2003*

## TECHNICAL REPORTS AND MISCELLANEOUS

---

- Grigori Fursin, Mike O'Boyle, Olivier Temam, and Gregory Watts. **A Fast and Accurate Evaluation of a Memory Performance Upper-Bound.** *Report for the MHAOTEU ESPRIT project No 24942*, February, 2001
- Jaume Abella, Cédric Bastoul, Jean-Luc Béchenec, Nathalie Drach, Christine Eisenbeis, Paul Feautrier, Björn Franke, Grigori Fursin, Antonio Gonzalez, Toru Kisku, Peter Knijnenburg, Josep Llosa, Michael O'Boyle, Julien Sébot, and Xavier Vera. **Guided Transformations.** *Report M3.D2 for the MHAOTEU ESPRIT project No 24942*, February 2001
- Jaume Abella, Grigori Fursin, Antonio Gonzalez, Josep Llosa, Michael O'Boyle, Abhishek Prabhat, Olivier Temam, Sid Ahmed Ali Touati, Xavier Vera, and Gregory Watts. **Advanced Performance Analysis.** *Report M3.D2 for the MHAOTEU ESPRIT project No 24942*, February, 2001
- Grigori Fursin. **Simulation 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*, Moscow Institute of Physics & Technology, pages 102-111, Moscow, Russia, 1997
- Grigori Fursin. **Measurement of characteristics of neural elements with the aid of personal computer.** *Proceedings of the national conference on devices of electronic and laser engineering*, Moscow Institute of Physics & Technology, pages 20-28, Moscow, Russia, 1997
- 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*, Moscow Institute of Physics & Technology, pages 112-117, Moscow, Russia, 1995

## REFERENCES

---

- **Prof. François Bodin**  
*CTO CAPS Entreprise, Professor at Université de Rennes 1*  
CAPS, Immeuble CAP NORD, Bât A - 2ème étage  
4 Allée Marie Berhaut, 35000 Rennes, France
- **Prof. Olivier Temam,**  
INRIA Saclay, Parc Club Orsay Université  
ZAC des vignes, 3, rue Jacques Monod – Bât G  
91893-ORSAY Cedex, France
- **Prof. Michael O'Boyle,**  
ICSA, School of Informatics, University of Edinburgh,  
Mayfield Road, Edinburgh, EH9 3JZ, Scotland, UK
- **Dr. Bilha Mendelson**  
*Senior manager (Code Optimization and Quality Technologies)*  
COT department, IBM, Haifa University Campus,  
Carmel Mountain Range, Haifa, 31904, Israel
- **Prof. David Padua (IEEE Fellow)**  
University of Illinois at Urbana-Champaign,  
201 N. Goodwin Avenue, Urbana,  
IL 61801-2302, USA