

# CURRICULUM VITAE

**Name:** Grigori Fursin, Ph.D.  
**Current job:** Tenured Research Scientist, INRIA Saclay, France  
**Address:** INRIA Saclay, Parc Orsay Universite  
4, rue Jacques Monod, 91893, Orsay cedex, France  
**E-mail** [grigori.fursin@inria.fr](mailto:grigori.fursin@inria.fr)  
**Website:** <http://fursin.net/research>  
**Birthday:** 07/1977, Moscow, Russia  
**Languages:** English, Russian, French (beginner)



## SUMMARY

---

From an early age I have been fascinated by computers and by extraordinary capabilities of the brain for learning, resiliency and adaptation to ever changing environment. Gradually, it became my lifelong dream to understand how brain works and to use this knowledge to open new horizons in science and technology.

Therefore, since 1993, I started independent, self-motivated and long-term research endeavor to enable adaptive self-tuning computer systems based on my interdisciplinary background in physics, electronics, computer science, mathematics, statistics and machine learning.

Nowadays, I enjoy leading challenging, innovative, collaborative and interdisciplinary research on automating and systematizing design and optimization of computer systems trying to go far beyond the state of the art. I attempt to revisit current over-complicated, under-performing and power-hungry computer technology with self-learning and self-tuning, adaptive technology combined with collaborative participation of users and computers through the internet within community-driven Collective Tuning Center (cTuning.org). I hope that this technology on automating code, compiler and architecture tuning will considerably improve performance, power consumption and reliability of new computing systems particularly for Exascale era, reduce their development cost and time to market by predicting better optimizations, architecture designs and run-time adaptation. I hope that this in turn will have a global positive impact on society by boosting innovation and research in science and technology as well as improving reproducibility of academic and industrial research.

I am very happy that most of my concepts, techniques and tools are now used and extended in academia and industry (IBM, Synopsys, Intel, Google and others) and included in real production compilers including GCC. They have been included in the HiPEAC 2012-2020 research roadmap, referenced in patents and conferences/journals such as PLDI, MICRO, CGO, CASES, HiPEAC, IJPP, TACO, and in world-wide press-releases from IBM. In 2010-2011, I have been invited to help establish new Intel Exascale Lab in France based on my past research. I am also honored to be regularly invited to give talks about my research at companies and universities including Intel, IBM, UIUC, INRIA, University of Edinburgh, Paris South University, and ICT.

In my spare time, I enjoy thinking about emerging technologies, intelligent systems, and understanding brain functionality. I am also 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 learning French language.

### Current research focus:

- Consolidate all my past research and developments
- Systematize and automate code and architecture co-design and optimization within Collective Tuning Center based on multidisciplinary approaches including empirical tuning, machine learning, statistical analysis, and multi-user collective participation
- Develop adaptive self-tuning computer systems to automatically balance performance, size, power and reliability, reduce time-to-market, development cost and improve return on investment particularly for the future Exascale systems
- Transfer my knowledge to industry through consulting or creation of collaborative laboratories

## Major research achievements:

- 2010 - 2011: helped to establish **new Intel Exascale Lab in France** based on my past research and developments to enable Exascale systems by 2018
- 2009 - 2010: participated in the preparation of the **EU HiPEAC 2012-2020 research vision** on future computer systems that also included most of my past research directions
- 2006 - 2009: designed **world's first machine learning enabled, self-tuning compiler (MILEPOST GCC / cTuning CC)** that can automatically adapt its optimization heuristic to any underlying reconfigurable architecture, and developed **community-based Collective Tuning Center (cTuning.org)** with unified collaborative optimization repository, web-services, tools, benchmarks and datasets. When enough optimization information is collected in the repository, it can be extrapolated to predict better optimizations for new programs and architectures using statistical, collective and machine learning techniques (*released as MILEPOST GCC and cTuning CC; available in mainline GCC and is used in academia and industry including IBM, Synopsys, CAPS; reference publications: FKMP2011, FT2010, MCFP2010, YYLP2010, DJBP2009, Fur2009, TOFP2009, CFAP2007, FC2007, FCOP2007, ABCP2006, FOTP2005, FOK2002*)
- 2005 - 2010: developed a concept to considerably speed up program optimization and compiler tuning through **collective optimization** where exploration and learning is transparently distributed among multiple users (*released within Collective Tuning Center and MILEPOST GCC; reference publications: FT2010, Fur2009, FCOP2005*); developed a concept to characterize programs and architectures through reactions to optimizations to predict better optimizations or hardware designs (*reference publications: FT2010, DCFP2007*);
- 2005 - 2010: developed **many practical techniques** to automate program and architecture co-design and multi-objective optimization; compiler design and tuning for reconfigurable architectures; run-time adaptation; predictive scheduling for heterogeneous CPU/GPU architectures using empirical, statistical, polyhedral, collective and machine learning techniques (predictive modeling, active learning, decision trees) based on my interdisciplinary background (*reused in MILEPOST project and cTuning.org*)
- 2004 - 2006: developed a concept to **enable run-time adaptation for statically compiled programs** through code cloning and reaction to dynamic program, environment and architecture behavior (*released as UNIDAPT framework, partially implemented in GCC and supported by Google; reference publications: FCOP2005, FT2010*)
- 2003 - 2006: developed a concept of the **Interactive Compilation Interface** to open up production compilers through light-weight event-based plugin framework to transform them into powerful interactive research toolsets (*available in mainline GCC; reference publications: FKMP2011, FC2007, FCOP2005*)
- 1999 - 2004: developed an unconventional approach to **treat computer programs, compiler, run-time systems and architectures as physical systems** to be able to apply empirical, statistical and machine learning techniques to automatically analyze and optimize their performance (*now mainstream technology; reference publications: FOK2002, Fur2004, FCOP2005, FKMP2011*); developed a concept for program and architecture characterization through reactions to semantically non-equivalent changes (*used and extended in industry; reference publications: FOTP2001, FOPT2004, Fur2004*)
- 1998 - 1999: developed a concept to access parallel high-performance computer systems through Internet web-services (*extended these ideas in my research on collective optimization*)
- 1993 - 1997: started parallelizing and optimizing my modeling software using random search strategies to better adapt to a given architecture (*extended these ideas in my PhD proposal and later research on empirical program optimization using statistical techniques and machine learning*)

## Major developments:

*I am quite sad about current situation in academic and industrial research where it is often impossible to reproduce results of publications. Hence, I spend considerable amount of time to implement and release all my techniques in open-source and production tools:*

- 2006 - 2009: designed and released world's first (according to IBM) machine learning enabled, self-tuning, adaptive compiler (**MILEPOST GCC / cTuning CC**), (*announced in world-wide IBM's press-release; <http://cTuning.org/ctuning-cc>*); developed **Collective Tuning Center** to collect optimization cases from the community, train MILEPOST GCC, predict better optimizations or designs and improve reproducibility of research (*announced in world-wide IBM's press-release; <http://cTuning.org>*)
- 2004 - 2006: developed and released **UNIDAPT framework** to enable dynamic adaptation for statically compiled programs (*partially available in mainline GCC, <http://cTuning.org/unidapt>*)
- 2005 - 2009: developed collaborative benchmarks and datasets (**cBench**) for systematic program and architecture characterization and optimization as well as training for MILEPOST GCC
- 2004 - 2009: developed **Continuous Collective Compilation Framework (CCC)** for automatic code and architecture tuning and sharing of optimization data in cTuning repository to train MILEPOST GCC
- 2003 - 2006: developed and released **Interactive Compilation Interface (ICI)** for Open64, PathScale and GCC to transform production compilers into interactive research toolsets (*available in mainline GCC, <http://cTuning.org/ici>*)
- 1999 - 2003: developed **Edinburgh Optimizing Software** to treat programs, compilers and architectures as physical systems for empirical optimization and analysis through reactions to transformations
- 1998 - 1999: developed a unified interface to access parallel high-performance computer systems through Internet to distribute and speed up my neural network modeling software
- 1993 - 1997: developed semiconductor neural network combined with software modeling to improve character recognition

## Major publications:

Since most of my approaches to code and architecture co-design and optimization are mostly non-traditional and interdisciplinary based on physics, biology, statistics, machine learning and collective participation of users, it usually takes a while before this technology is accepted by academia and industry. Nevertheless, I am very happy that most of my concepts, techniques and tools have been eventually accepted and referenced in conferences/journals such as PLDI, MICRO, CGO, CASES, HiPEAC, IJPP, TACO, and in world-wide press-releases from IBM.

Main references: FKMP2011, FT2010, MCFP2010, YYLP2010, FT2009, DJBP2009, Fur2009, JGVP2009, LF2009, LFF2007, DCFP2007, CFAP2007, FC2007, FCOP2007, ABCP2006, FCOP2005, FOTP2005, FOTP2004, FOK2002

## Diplomas:

- 07/2004: Ph.D. in computer science from the University of Edinburgh, UK with ORS award (*title: "Iterative Compilation and Performance Prediction for Numerical Applications", advisor: Prof. Michael O'Boyle*)
- 06/1999: M.S. in computer engineering from Moscow Institute of Physics and Technology, Russia with highest honours (*title: "Unifying remote access to high-performance computing systems through the Web"*)
- 06/1997: B.S. in physics and electronics from Moscow Institute of Physics and Technology, Russia with highest honors (*interdisciplinary courses in physics, electronics, mathematics, statistics, machine learning, biology and computer science*)

**Professional experience:**

- 09/2011 - cur.: Tenured research scientist at INRIA Saclay, France
- 03/2010 - 08/2011: Head of the application characterization group and director of research at Intel Exascale Lab, France (on industrial sabbatical from INRIA invited to help establish new Intel Lab in Europe with novel research directions)
- 09/2007 - 02/2010: Tenured research scientist at INRIA Saclay, France
- 12/2005 - 08/2007: Postdoctoral researcher at INRIA Futurs and Paris South University, France
- 10/2004 - 11/2005: Visiting researcher at INRIA Futurs and Paris South University, France
- 01/2002 - 11/2005: Research associate at the University of Edinburgh, UK
- 02/2000 - 03/2000: Visiting researcher at Paris South University, France
- 02/1999 - 12/2001: Research assistant at the University of Edinburgh, UK
- 09/1997 - 01/1999: Research assistant at the Institute for High-Performance Computing of the Russian Academy of Sciences, Russia
- 01/1994 - 06/1994: Research assistant at Moscow Institute of Physics and Technology, Russia

**Institution building:**

- 2010 - 2011: Helped to establish new Intel Exascale Lab (Intel/CEA/GENCI/UVSQ) in France
- 2008 - cur.: Founded public, community-driven Collective Tuning Center (cTuning.org)

**Awards and distinctions:**

- 2009: IBM world-wide press release about world's first machine learning enabled self-tuning compiler MILEPOST GCC and about Collective Tuning Center (cTuning.org)
- 2005 - cur.: Several best paper awards
- 2005 - 2006: EU HiPEAC postdoctoral fellowship
- 2000 - 2001: International Overseas Research Student Award at the University of Edinburgh, UK
- 1999: M.S. with highest honours and golden medal from Moscow Institute of Physics and Technology, Russia
- 1997: B.S. with highest honours from Moscow Institute of Physics and Technology, Russia
- 1996 - 1998: International George Soros Award and Fellowships.

**Major funding:**

I am very grateful to European Commission, HiPEAC, Intel, Google, University of Edinburgh, UVSQ and INRIA for generously funding my research since undergraduate years:

- 2010 - 2011: Funding from Intel Exascale Lab
- 2009: Funding from Google Summer of Code Program and from INRIA
- 2009: Funding for visiting students from ICT, China
- 2006 - 2009: Funding from EU FP6 MILEPOST project
- 2005 - 2006: HiPEAC postdoctoral fellowship
- 2001 - 2004: Funding from the University of Edinburgh
- 1999 - 2000: Funding from EU FP5 MHAOTEU project
- 1999 - 2000: UK ORS fellowship
- 1996 - 1998: Funding from International Soros Science Education Program

**Invited talks:**

Since 2003, I am honoured to be regularly invited to give talks about my research on self-tuning computer systems at various companies, universities, laboratories and organizations including:

- IBM TJ Watson Lab, USA; IBM Toronto Lab, Canada; IBM Haifa Lab, Israel;
- Intel USA; Intel Russia; Reservoir Lab, USA;
- University of Illinois at Urbana Champagne, USA; Imperial College, UK;
- Institute of Computing Technology of Chinese Academy of Sciences;
- Institute for System Programming of Russian Academy of Sciences;
- Paris South University, INRIA, CNRS, UVSQ, France;
- European Commission (Information and Brokerage Conference on Information and Communication Technologies presenting advances in research on machine learning for code and architecture optimization)

### **Industrial collaborations:**

- Intel, IBM, CAPS Enterprise, Google, ARC (now Synopsys), AMD, STMicroelectronics, NXP, CEA, GENCI

### **Academic collaborations:**

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

### **Technology transfer:**

Though I primarily focus on a very long-term research for the future computer systems, I am also interested in transferring my ideas to industry. However, since my approaches to code and architecture co-design and optimization are mostly non-traditional and interdisciplinary based on physics, biology, statistics, machine learning and collective participation of users, it usually takes a while before this technology is accepted by industry. Nevertheless, I truly believe in my long-term research vision and very determined, I am never discouraged and spend considerable effort to popularize my techniques through publications, open-source tools and invited talks. I am very happy that now most of my techniques and tools are actively used and extended in companies including IBM, Intel, ARC (now Synopsys), Google.

- 2010 - 2011: Invited to establish new Intel Exascale Lab in France and create a group on application characterization based on my past research and developments.
- 2005 - 2009: Moved my research on code and architecture co-design and optimization based on empirical tuning for multiple objectives (performance, size, compilation time, etc) combined with statistical, collective and machine learning techniques to production compilers including Open64, PathScale and GCC within MILEPOST project and within community-driven Collective Tuning Center (<http://cTuning.org>). This technology is used by various companies and universities world-wide including IBM, CAPS Enterprise and Synopsys.
- 2009: Moved Interactive Compilation Interface to mainline GCC (collaboration with IBM, Google and HiPEAC); Partially moved UNIDAPT framework to GCC (collaboration with ICT, China and funded by Google Summer of Code program).

### **Organizing events:**

- 2011 - cur: International ACM EXADAPT workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era ([exadapt.org](http://exadapt.org)), co-located with FCRC/PLDI 2011, ASPLOS 2012 (collaboration with Google and University of Virginia)
- 2011: BOF on collaboratively mining rich information to prepare the Exascale challenges, co-located with SuperComputing 2011 (collaboration with Intel)
- 2010: Speedup-Test tutorial on statistical methodology to evaluate program speedups and their optimization techniques to improve quality and reproducibility of academic research (collaboration with UVSQ)
- 2010: International Workshop on GCC Research Opportunities, co-located with HiPEAC 2010 (collaboration with IBM)
- 2009: Tutorial on cTuning tools and techniques at HiPEAC Computing Week at Infineon
- 2007: International GCC tutorial, co-located with HiPEAC 2007 (collaboration with IBM, STMicroelectronics, RedHat)
- 2007 - cur: International SMART workshop series on Statistical and Machine Learning Approaches applied to Architectures and Compilation, co-located with HiPEAC 2007-2010, CGO 2011
- 2001: Co-organizer of the CPC 2001

**Chairing events:**

- 2011: Session chair on dynamic optimization at CGO 2011
- 2011: EXADAPT 2011 workshop co-chair at FCRC/PLDI 2011
- 2010: GROW 2010 workshop co-chair at HiPEAC 2010
- 2007: SMART 2007 workshop co-chair at HiPEAC 2007

**Participating in program committees:**

HiPEAC/ACM TACO 2012; CGO 2011; EXADAPT 2011; ICPADS 2009; iWAPT 2009; GROW 2009-2011; IPDPS 2008; CASES 2007; SMART 2007-2011

**Reviewing:**

- 2012: ACM TACO/HiPEAC 2012 (member of the board of distinguished reviewers)
- 2011: 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 (invited international expert)
- 2010: EU HiPEAC 2012-2020 research vision
- 2009: Springer-Verlag book on auto-tuning techniques
- 2000 - cur.: ACM TACO, CGO, HiPEAC, PACT, TPDS, CASES, ICS, DATES, CPE, MASCOTS, CPC, IJHPSA, ICPADS, JPDC, CLSS, iWAPT, Open64, SMART, GROW

**Teaching:**

I do not have to teach at INRIA officially, but I enjoy it. Therefore, from 2005 I started lecturing at Paris South University about my research on code and architecture co-design and optimization using empirical, statistical, machine learning and adaptive techniques. Eventually, in 2008-2009, I was appointed as an organizer of the MS course on "Future Computing Systems" at Paris South University, France with a focus on future adaptive self-tuning computing systems. I temporarily stopped it during my industrial sabbatical at Intel Exascale Lab in 2010-2011, but plan to resume it soon.

- 2005-2010: Part-time lecturer and organizer of the M2R course on "Future Computing Systems" at Paris South University, France

**Current students:**

- 2009 - cur.: Abdul Wahid Memon, PhD student at UVSQ, France; government funded
- 2009 - cur.: Yuriy Kashnikov, PhD student at UVSQ, France; funded by Intel Exascale Lab, France

**Professional memberships:**

- HiPEAC, ACM, IEEE

**General interests and expertise:**

- systematization and automation of code and architecture design and optimization
- non-traditional interdisciplinary approaches to design of computer systems
- self-tuning adaptive computer systems
- adaptive compilers, programs, architectures and run-time systems
- automatic tuning of computer systems to improve performance, reliability while reducing power consumption, size and contentions
- automatic improvement in development cost, time to market and return of investment for new computer systems
- cloud computing, data centers, embedded architectures, supercomputers
- collaborative repositories and tools for design and optimization of computer systems
- reproducible and collaborative research
- machine learning and artificial intelligence
- automatic classification, predictive modeling, active learning
- brain understanding and modeling
- bio-inspired design of computer systems
- emerging information technologies
- knowledge transfer to industry, consulting and startups

# EXTENDED CURRICULUM VITAE

## PROFESSIONAL EXPERIENCE

---

- 09/2011-  
cur. - Tenured research scientist and UNIDAPT Group founder (<http://unidapt.org>) at INRIA Saclay, France. I continue developing foundations of adaptive, self-tuning computing systems particularly for exascale era.
- 03/2010-  
08/2011 - Director of research and head of application characterization group at Intel Exascale Lab in France (<http://www.exascale-computing.eu>) (on personal industrial sabbatical from INRIA invited to help to establish new Intel Lab in France from scratch). I prepared long-term R&D and build a team on application characterization and optimization from scratch. My project grew from 2 to 8 members by the end of my sabbatical and we successfully passed the 1st year review while securing future funding.
- 09/2007-  
02/2010 - Tenured research scientist and UNIDAPT Group founder (<http://unidapt.org>) at INRIA Saclay, France. 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 / Paris South University, 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, MICRO 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.
- 02/2000-03/2000 - Visiting researcher at Paris-Sud XI University (Paris South 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 Wen-mei 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; examiners: Prof. Francois Bodin and Dr. Marcelo Cintra.

In this thesis I introduced a novel and simple approach to quickly detect realistic lower bound of the execution time of memory intensive applications by modifying individual assembler instructions including converting 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. I evaluated this technique on in-order Alpha and out-of-order Intel processors (I wrote my own tools for assembler analysis, instrumentation and modification for a range of platforms) and a variety of small kernels and large SPEC applications, and showed that this technique can be effectively used to detect parts of a program which can benefit from memory optimizations (and avoid wasting time for long iterative or manual optimizations on a code that do not have any potential for further optimizations). I also showed that focusing on code optimizations that reduce cache misses on out-of-order architecture using hardware counters can be totally misleading and my approach provides more realistic upper bound.

I also wrote my own source-to-source Fortran and C transformation tool and showed how to automate program optimization using empirical iterative compilation on program, function and loop-level based on array padding, loop tiling and unrolling combined with random and hill-climbing search strategies that beat state-of-the-art commercial COMPAQ/HP and Intel compilers on various SPEC2000 benchmarks. All these techniques and tools have been actively used and extended in academia and industry as well as motivating development of various R&D tools.

I implemented and released all my techniques in an open-source Edinburgh Optimization Software (EOS). 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).

[*FOTP2001, FOK2002, FOTP2004, FUR2004*].

## EDUCATION

---

- 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. 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.

## ORGANIZING EVENTS

---

- 2012 - **EXADAPT 2012** (at **ASPLOS 2012**): Organizer together with Robert Hundt from Google, Yuriy Kashnikov from Intel Exascale Lab and Jason Mars from the University of Virginia (2nd International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era)  
<http://exadapt.org>
- 2011 - **BOF** (at **SC 2011**): Organizer together with Marie-Christine Sawley from Intel Exascale Lab (collaborative design and tuning of Exascale computer systems).  
**EXADAPT 2011** (at **PLDI 2011/FCRC 2011**): Organizer together with Robert Hundt from Google, Yuriy Kashnikov from Intel Exascale Lab and Jason Mars from the University of Virginia (ACM SIGPLAN 1st International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era, Keynote on "Autotuning in the Exascale Era": Prof. Katherine Yelick (LBNL and UC Berkeley, USA) )  
<http://exadapt.org>
- **SMART 2011** (at **CGO 2011**): Organizer (5<sup>th</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation, Chair: Prof. Francois Bodin (CAPS Entreprise, France), Keynote on "Automatic Performance Tuning and Machine Learning": Prof. Markus Püschel (ETH Zürich, Switzerland) )  
<http://cTuning.org/workshop-smart2011>
- 2010 - **GROW 2010** (at **HiPEAC 2010**): co-organizer and editor of informal proceedings (2<sup>nd</sup> International Workshop on GCC Research Opportunities)  
<http://cTuning.org/workshop-grow10>
- **SMART 2010** (at **HiPEAC 2010**): co-organizer and editor of informal proceedings (4<sup>th</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation, Chair: Prof. David Whalley, Keynote on "Moving adaptation into individual optimizations": Prof. Keith Cooper (Rice University, USA) )  
<http://cTuning.org/workshop-smart10>
- **Speedup-Test 2010** tutorial (at **HiPEAC 2010**): co-organizer (Statistical Methodology to Evaluate Program Speedups and their Optimization Techniques)  
<http://www.hipeac.net/conference/pisa/speedup>
- 2009 - **SMART 2009** (at **HiPEAC 2009**): Organizer and PC member (3<sup>rd</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation, Chair: Prof. David Padua, Panel: Francois Bodin, Marcelo Cintra, Bilha Mendelson, Lawrence Rauchwerger, Per Stenstrom)  
<http://www.hipeac.net/smart-workshop.html>
- **cTuning tools tutorial** (at **HiPEAC Computing Week at Infineon**): organizer and presenter  
[http://www.hipeac.net/computing\\_systems\\_week\\_munich](http://www.hipeac.net/computing_systems_week_munich)
- 2008 - **SMART 2008** (at **HiPEAC 2008**): Organizer (2<sup>nd</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation, Chair: Prof. Michael O'Boyle)  
<http://www.hipeac.net/smart-workshop-08.html>

- 2007 - **2<sup>nd</sup> GCC Tutorial** (at **HiPEAC 2007**): organizer and presenter ("Interactive Compilation Interface and plugin system to transform GCC into research toolset")  
<http://www.hipeac.net/gcc-tutorial>
- **SMART 2007** (at **HiPEAC 2007**): 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-07.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 2001**: local organizer (9<sup>th</sup> Workshop on Compilers for Parallel Computers)

## CHAIRING EVENTS

---

- 2011 - **CGO 2011**: Session chair on dynamic compilation of the International Symposium on Code Generation and Optimization, session chair on dynamic compilation  
<http://www.cgo.org/cgo2011>
- **EXADAPT 2011** (at **PLDI 2011/FCRC 2011**): Co-chair of the ACM SIGPLAN 1st International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era  
<http://exadapt.org>
- 2010 - **GROW 2010** (at **HiPEAC 2010**): Co-chair of the 2<sup>nd</sup> International Workshop on GCC Research Opportunities  
<http://cTuning.org/workshop-grow10>
- 2007 - **SMART 2007** (at **HiPEAC 2007**): Co-chair of the 1<sup>st</sup> Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation  
<http://www.hipeac.net/smart-workshop-07.html>

## PARTICIPATING IN PROGRAM COMMITTEES

---

- 2011 - **CGO 2011**: International Symposium on Code Generation and Optimization  
<http://www.cgo.org/cgo2011>
- **GROW 2011** (at **CGO 2011**): 3<sup>rd</sup> International Workshop on GCC Research Opportunities  
<http://grow2011.inria.fr>
- 2010 - **GROW 2010** (at **HiPEAC 2010**): 2<sup>nd</sup> International Workshop on GCC Research Opportunities  
<http://cTuning.org/workshop-grow10>
- 2009 - **iWAPT 2009**: International Workshop on Automatic Performance Tuning  
<http://www.iwapt.org/2009>
- **ICPADS 2009**: International Conference on Parallel and Distributed Systems, multi-core architectures track  
<http://www.comp.polyu.edu.hk/conference/icpads09>
- **Open64 Workshop** (at **CGO 2009**):  
<http://www.capsl.udel.edu/conferences/open64/2009>
- **SMART 2009** (at **HiPEAC 2009**): 3<sup>rd</sup> International Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation  
<http://www.hipeac.net/smart-workshop.html>
- **GROW 2009** (at **HiPEAC 2009**): 1<sup>st</sup> International Workshop on GCC Research Opportunities  
<http://www.doc.ic.ac.uk/~phjk/GROW09>
- 2008 - **IPDPS 2008**: IEEE International Parallel and Distributed Processing Symposium  
[http://www.ipdps.org/ipdps2008/2008\\_cfp.html](http://www.ipdps.org/ipdps2008/2008_cfp.html)
- **SMART 2008** (at **HiPEAC 2008**): 2<sup>nd</sup> International Workshop on Statistical and Machine Learning Approaches applied to Architectures and Compilation  
<http://www.hipeac.net/smart-workshop-08.html>
- 2007 - **CASES 2007**: International Conference on Compilers, Architecture, and Synthesis for Embedded Systems  
<http://www.irit.fr/recherches/ARCHI/MARCH/CASES2007>

## REVIEWING

---

- 2012 - HiPEAC 2012/ACM TACO (member of the board of distinguished reviewers)  
<http://www.hipeac.net/conference/paris/>  
<http://www.hipeac.net/conference/paris/reviewers>
- 2011 - Research proposals for the open grant competition of the Russian Federation to attract leading scientists to Russian universities:  
<http://eng.mon.gov.ru/pro/ved/uch/4311>
- 2010 - EU HiPEAC 2012-2020 research vision:  
<http://www.hipeac.net/roadmap>
- 2009 - Springer-Verlag book on auto-tuning techniques
- 2000 - cur. - ACM TACO, CGO, HiPEAC, PACT, TPDS, CASES, ICS, DATES, CPE, MASCOTS, HiPEAC, CPC, IJHPSA, ICPADS, JPDC, CLSS, iWAPT, Open64, SMART, GROW

## TEACHING

---

*I do not have to teach at INRIA officially, but I enjoy it. Therefore, from 2005 I started lecturing at Paris South University about my research on code and architecture co-design and optimization using empirical, statistical, machine learning and adaptive techniques. Eventually, in 2008-2009, I was appointed as an organizer of the MS course on "[Future Computing Systems](#)" at Paris South University, France with a focus on future adaptive self-tuning computing systems. I temporarily stopped it during my industrial sabbatical at Intel Exascale Lab in 2010-2011, but I plan to resume it soon. You can find more info about my teaching activities [at this page](#).*

- 2008-2009 - Organizer of the M2R course on "Future Computing Systems" in Paris South University, France with a focus on my research on adaptive self-tuning systems; part-time lecturer
- 2007-2008 - Part-time lecturer, M2R course in Paris South University, France with a focus on my research on continuous adaptive program optimization combined with machine learning techniques
- 2006-2007 - Part-time lecturer of the M2R course in Paris South University, France with a focus on my research on continuous adaptive program optimization combined with machine learning techniques
- 2005-2006 - Part-time lecturer of the M2R course in Paris South University, France with a focus on my research on adaptive and feedback driven compilation

## ADVISING/COLLABORATING

---

- 2011
  - Dr. Pablo Oliveira (postdoctoral researcher at Université de Versailles-Saint-Quentin-en-Yvelines, France).
- 2010 - cur.
  - Yuriy Kashnikov and Abdul Wahid Memon (PhD students at Université de Versailles-Saint-Quentin-en-Yvelines, France).
- 2009
  - 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 [FT2010, HPWP2010, FT2009, FCOP2005]  
(<http://socqhop.appspot.com/org/home/google/gsoc2009/gcc>).
  - 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>).
  - Dr. Cosmin Oancea (postdoctoral researcher from Cambridge University) to develop software TLS tools and techniques to improve automatic program parallelization.
- 2008
  - 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].
  - 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
  - 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
  - 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
  - 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
  - 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"

## RESEARCH AND DEVELOPMENT 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, FT2010, FKMP2011, 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:* FKMP2011, FMTP2008, MCFP2010, 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, FT2010

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, MCFP2010, HPWP2010, FKMP2011

*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 statically enable run-time optimizations and self-tuning binaries through cloning of program hot spots, applying various aggressive optimizations to clones for different optimization cases (that may improve performance/power/fault-tolerance, etc), statically integrating low-overhead program/system behaviour monitoring routines (using hardware counters) and selecting appropriate versions at run-time as a reaction to different program behavior, architectural changes or contentions.

For the first time, I utilized my Interactive Compilation Interface for PathScale compiler with loop vectorization, tiling, unrolling, interchange, fission/fusion, pipelining, prefetching and array padding to make static self-tuning binaries that can automatically learn from the past experience and adapt/react to various environments, run-time behavior and contentions that is important to improve efficiency and cost of both embedded systems and HPC data centers (cloud computing).

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. For example, we used it later to support collective optimization [FT2010,FT2009,FMPP2007], predictive scheduling for heterogeneous multicore systems [JGVP2009], automatically create adaptive multiversioning libraries and programs [LCWP2009], etc. based on continuous transparent run-time program optimization and adaptation as a reaction to dynamic changes in program behavior and environment. Since 2007 it is being actively extended by Google Inc. for data centers and cloud computing.

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

*Publications:* FCOP2005, FMPP2007, LFF2007, LCWP2009, JGVP2009, Fur2009, FT2009, HPWP2010, FT2010

- **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, MCFP2010, FKMP2011

*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 proposed and developed portable distributed client/server framework to enable unified remote access to heterogeneous high-performance computers as Internet web-services at Russian Academy of Sciences (some similar ideas are now used in GRID and cloud computing).

1996-1998 - I proposed and 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 was actively 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 AND DISTINCTIONS

---

- 2010 - HiPEAC paper award for the publication “Evaluating Iterative Optimization across 1000 Data Sets” (Y. Chen et al) at the ACM SIGPLAN 2010 Conference on Programming Language Design and Implementation (PLDI 2010), Toronto, Canada, January 2010.
- 2009 - IBM made a world-wide press release about MILEPOST GCC and Collective Tuning Center which I designed to show a possibility for collective optimization of computer systems through collaborative sharing of optimization data and tuning tools combined with statistical and machine learning techniques.  
<http://www-03.ibm.com/press/us/en/pressrelease/27874.wss>
- HiPEAC paper award for the publication “Portable Compiler Optimization Across Embedded Programs and Microarchitectures using Machine Learning” (C. Dubach et al) at the 42<sup>nd</sup> IEEE/ACM International Symposium on Microarchitecture (MICRO), New York, USA, December 2009.
- 2005 - My publication “A Practical Method For Quickly Evaluating Program Optimizations” was ranked highest at the conference on High Performance Embedded Architectures & Compilers (HiPEAC 2005). This research has laid the foundations for realistic run-time adaptation of statically compiled programs and has been actively used and extended in academia and industry since then.
- 1999-2000 - ORS award from the University of Edinburgh to fund my PhD program in the UK (UK Scholarship for International Research Students)  
<http://www.orsas.ac.uk/applications>
- 1999 - Golden medal for B.S. and 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”
- 1993 - Medal for studies during secondary school

## FUNDING

---

*I am very grateful to the following organizations and companies who generously sponsored my independent interdisciplinary research on developing self-tuning adaptive computer systems from the early undergraduate years and up to now: European Commission FP5 and FP6 programs, HiPEAC network of excellence, University of Edinburgh, INRIA, UVSQ, Google and Intel.*

- 2010-2011 - Funding from Intel Exascale Lab (**amount is under NDA**). I was invited to help establish new Intel Exascale Lab in France (Intel, CEA, GENCI, UVSQ) as a project leader and scientific advisor. I took personal sabbatical from INRIA and helped to prepare research directions for Exascale application characterization, build a team (that grew from 2 to 8 members by the end of my sabbatical), pass successfully the review and secure the rest of the funding for the Lab.
- 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>
- 2009 - 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 and machine learning:  
<http://socghop.appspot.com/org/home/google/gsoc2009/gcc>
- 2008 - EU HiPEAC collaborative grant for the project "Context-aware optimization and run-time adaptation of sequential libraries for multi-core systems". I have been the principal investigator and collaborated with my colleagues from Imperial College and ARM to develop static libraries adaptable at run-time to different program and architecture behaviour (program phases, shared resources contentions, switching to another architecture during virtualization, etc) based on my previously developed technique on enabling dynamic adaptation of statically compiled programs.
- 2006-2009 - Funding from the EU FP6 MILEPOST project for INRIA (**~650,000 euros**) (<http://cTuning.org/project-milepost>) to develop the first machine-learning enabled compiler. It was a collaborative project involving IBM, University of Edinburgh, CAPS Enterprise, ARC and INRIA. I have been a technical leader and the architect of the MILEPOST GCC and cTuning collaborative optimization framework. I have been also responsible for recruiting engineers, students, postdocs and interns at INRIA for this project.
- 2006-2007 - EU HiPEAC collaborative grant for the project "Exploring optimization techniques and runtime code selection mechanisms for heterogeneous systems". I was principle investigator and collaborated with my colleagues from the UPC (Spain) to develop predictive scheduling for heterogeneous architectures based on my previously developed technique on enabling dynamic adaptation of statically compiled programs.

- 2005-2006 - HiPEAC postdoctoral research grant (**~75,000 euros**) to collaborate with the Alchemy group at INRIA Futurs, France. I proposed to make statically compiled programs adaptable at run-time through static cloning of hot functions with different optimizations that are dynamically selected based on reaction to program run-time behavior and phases. This research has laid foundations for developing practical self-tuning programs.
- 2001-2004 - Funding from the University of Edinburgh to support my PhD research. -
- 1999-2000 - Funding from the EU FP5 MHAOTEU project to support my PhD research on empirical program optimization.
- 1999-2000 - ORS award (UK Scholarship for International Research Students) to support my PhD research at the University of Edinburgh on empirical program optimization. I proposed to optimize programs as physical systems through empirical analysis and reaction to code modifications based on my interdisciplinary background. I also proposed to use neural networks from my previous research to tune compiler heuristic and optimize program but it took us 5 more years before we could apply systematic machine learning to program optimization in the MILEPOST project.
- 1996-1998 - Research grants from International Soros Science Education Program to support my independent undergraduate research on modelling and designing semiconductor neural networks.

## 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 (invited talk)*

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

*Presented at the University of Illinois at Urbana-Champaign and at Intel Illinois, USA, April 2009 (invited talk)*

- "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 (invited talk by European Commission).*

- "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 (invited talk)*

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

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

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

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

*Presented at IBM Haifa Lab (compiler group), Israel, April 2007 (invited talk)*

- "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 (invited talk)*

- “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 (invited talk)*

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

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

*Presented at the UPC, Barcelona, Spain, January 2007 (invited talk)*

- “Towards continuous collective compilation”

*Presented at the ICSA seminar, University of Edinburgh, UK, December 2006 (invited talk)*

- “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 (invited talk)*

- “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 (presentation of MILEPOST proposal for EU FP6 funding)*

- “A practical method for quickly evaluating program optimizations”

*Presented at the ICSA seminar, University of Edinburgh, UK, December 2005 (invited talk)*

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

*Presented at the LRI, Paris-Sud XI University, France, April 2003 (invited talk)*

[FUR2004]

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

**Based on FTOP2001, FOK2002, FOTP2004.**

*In this thesis I introduced a novel and simple approach to quickly detect realistic lower bound of the execution time of memory intensive applications by modifying individual assembler instructions including converting 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. I evaluated this technique on in-order Alpha and out-of-order Intel processors (I wrote my own tools for assembler analysis, instrumentation and modification for a range of platforms) and a variety of small kernels and large SPEC applications, and showed that this technique can be effectively used to detect parts of a program which can benefit from memory optimizations (and avoid wasting time for long iterative or manual optimizations on a code that do not have any potential for further optimizations). I also showed that focusing on code optimizations that reduce cache misses on out-of-order architecture using hardware counters can be totally misleading and my approach provides more realistic upper bound.*

*I also wrote my own source-to-source Fortran and C transformation tool and showed how to automate program optimization using empirical iterative compilation on program, function and loop-level based on array padding, loop tiling and unrolling combined with random and hill-climbing search strategies that beat state-of-the-art commercial COMPAQ/HP and Intel compilers on various SPEC2000 benchmarks. All these techniques and tools have been actively used and extended in academia and industry as well as motivating development of various R&D tools.*

*I implemented and released all my techniques in an open-source Edinburgh Optimization Software (EOS).*

## PUBLICATIONS (REFERENCES)

---

Microsoft Research: <http://journalogy.com/Author/968902.aspx?query=grigori%20fursin>

Google Scholar: <http://scholar.google.com/scholar?q=grigori+fursin>

DBLP: <http://www.informatik.uni-trier.de/~ley/db/indices/a-tree/f/Fursin:Grigori.html>

---

## PUBLICATIONS (INTERNATIONAL JOURNALS)

---

[FKMP2011]

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, Chris Williams, Michael O'Boyle. **MILEPOST GCC: machine learning enabled self-tuning compiler.** *International Journal of Parallel Programming (IJPP)*, June 2011, Volume 39, Issue 3, pages 296-327

- *Concept is included in the HiPEAC 2012-2020 research roadmap*

[FT2010]

Grigori Fursin and Olivier Temam. **Collective Optimization: A Practical Collaborative Approach.** *ACM Transactions on Architecture and Code Optimization (TACO)*, December 2010, Volume 7, Number 4, pages 20-49

- *Concept is included in the HiPEAC 2012-2020 research roadmap*

[LF2009]

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

[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

[FOK2005]

Grigori Fursin, Michael O'Boyle and Peter Knijnenburg. **Evaluating Iterative Compilation.** *Lecture Notes in Computer Science*, Volume 2481, pages 362-376, 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

## PUBLICATIONS (INTERNATIONAL CONFERENCES)

---

[MCFP2010]

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, Architecture, And Synthesis For Embedded Systems (CASES 2010)*, October 2010, Scottsdale, AZ, USA

[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)

- **HiPEAC paper award**

[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 is now published in ACM TACO (FT2010)**

[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.**

[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)

- **I obtained a HiPEAC collaborative grant to develop this technique with Victor Jimenez.**
- **Similar approaches for gluing/adapting applications for heterogeneous architectures are used in Intel's Qilin and in CAPS Enterprise's HMPP.**

[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

---

[DCFP2007]

Christophe Dubach, John Cavazos, Björn Franke, Grigori Fursin, Michael O'Boyle and Olivier 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)

[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**

[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**
- **Continuation of research presented in FOK2002 and FUR2004**
- **Research extended in FKMP2011 (MILEPOST GCC)**

---

[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**
- **Extensions for transparent collective optimization is available in our ACM TACO publication (FT2010)**
- **Concept is included in the HiPEAC 2012-2020 research roadmap**

*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 behaviour monitoring routines. It has been referenced in patents and considerably extended in academia and industry. For the first time, we utilized Interactive Compilation Interface for PathScale compiler with loop vectorization, tiling, unrolling, interchange, fission/fusion, pipelining, prefetching and array padding to make static self-tuning binaries that can automatically learn from the past experience and adapt/react to various environments, run-time behavior and contentions that is important to improve efficiency and cost of various HPC systems. Since 2007 it is being actively extended by Google for data centers (cloud computing).*

[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)

[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

## PUBLICATIONS (INTERNATIONAL WORKSHOPS)

---

### [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)

### [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

*This paper introduces collective tuning infrastructure (<http://cTuning.org>) and repository (<http://cTuning.org/cdatabase>) to start continuous parameterization of all computing systems and to automate, simplify and systematize code and architecture design, characterization and optimization. Collecting enough data about various architectures, compilers, programs, benchmarks, kernels and datasets will help to quickly predict better program optimizations or architecture designs using machine learning techniques thus considerably reducing time to market and enabling self-tuning, adaptive computing systems.*

### [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)

### [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

- *Extended version is now published in IJPP ([FKMP2011](#)).*

### [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

---

[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

[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)

- **More info is now available in the extended version published in IJPP ([FKMP2011](#)).**

[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

[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

*This paper introduces a concept of empirical optimization (iterative compilation) of large applications to automatically adapt them to a given hardware using several basic search strategies. Our 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 the foundations for further research on focused optimizations using statistical techniques, machine learning and run-time adaptation.*

[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

## PUBLICATIONS (TECHNICAL REPORTS)

---

[FOTP2001]

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

[ABBP2001]

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

[AFGP2001]

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

## PUBLICATIONS (NATIONAL CONFERENCES)

---

[FUR1997A]

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

[FUR1997B]

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

[FUR1995]

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

## PUBLICATIONS (MISCELLANEOUS)

---

[FHMP2011]

Grigori Fursin, Robert Hundt, Jason Mars, Yuriy Kashnikov. **Introducing ACM SIGPLAN International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era (<http://exadapt.org>).** *ACM International Conference Proceeding Series*, co-located with PLDI, June 2011, San Jose, USA

## REFERENCES

---

- **Prof. Michael O'Boyle,**  
*Director of ICOSA*  
School of Informatics, University of Edinburgh,  
Mayfield Road, Edinburgh, EH9 3JZ, Scotland, UK
- **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  
91893-ORSAY Cedex, France
- **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