



Information Society
Technologies

HiPEAC *info*¹²

COMPILATION ARCHITECTURE

appears quarterly | October 2007

Network of Excellence on High Performance Embedded Architectures and Compilers

- 2 Message from the HiPEAC coordinator
- 2 Message from the project officer
- 3 HiPEAC²
- 3 HiPEAC Journal
- 4 HiPEAC 2008 Conference
- 4 Community News
- 5 In the spotlight:
 - Dilworth Decomposition
 - 2nd CEDI in Zaragoza (Spain)
 - International workshop on GCC for Research in Embedded and Parallel Systems (GREPS)
- 6 PhD News
- 8 Upcoming events



**The Future of Computing
Thank you, Stamatis**

www.HiPEAC.net


HiPEAC 2008 Conference:
Göteborg, Sweden, January 27-29, 2008

Message from the HiPEAC coordinator

Mateo Valero
Coordinator
UPC Barcelona
mateo@ac.upc.edu

Dear Colleagues,

Last September, the HiPEAC Network of Excellence entered its final year. We are keen to tackle this last period and bring it to a close successfully, particularly as we know that this is only the first phase of a bigger project. I have the pleasure to announce to you that the new HiPEAC proposal submitted to the first FP7 call was approved. This means that HiPEAC will continue within FP7, with Koen de Bosschere from Ghent University as the coordinator of this new phase. We are working to create an improved network by capitalising on the lessons learnt from the present HiPEAC. I would like to thank project officer, Mercè Griera i Fisa for her continuous availability, support, and dedication to the HiPEAC Network, and I would like to welcome Panos Tsarchopoulos as the new HiPEAC EC officer.



But as mentioned above, we still have one more year of HiPEAC ahead of us. Thus, I would like to encourage all of you, who are interested in joining our community, to apply for membership to our current HiPEAC Network. Visit our website to find out how to apply for membership.

HiPEAC is gaining visibility, a fact that is being contributed to by its continuity. It is, therefore, in the best interest of our community to continue to offer high quality events and results.

A new issue of the HiPEAC Journal was produced in September, 2007 and it

is currently available on our website. We would like our journal to become the main HiPEAC community publication; therefore, I encourage all members to submit their best papers there. Remember that it just cannot happen without you! You will find the submission procedures on our website.

Message from the project officer

Panos Tsarchopoulos
Panagiotis.Tsarchopoulos@ec.europa.eu

Last month, a new Project Officer was put in charge of HiPEAC. After devoting much time and effort to ensuring HiPEAC remained a successful Network of Excellence, Mercè Griera i Fisa moved to new professional responsibilities within the European Commission. Apart from having been the Project Officer for HiPEAC, Mercè was also involved in the development of the "Computing Systems" research objective that formed part of the 1st Call for Proposals of the Information and Communication Technologies in FP7.

As you know, the Call's deadline was in May 2007 and the Commission is currently negotiating the contracts with the proposals that received the highest scores. We are particularly happy with the exceptional quality of the proposals we received as witnessed by the fact that three out of every four proposals received were marked above the threshold.

The application areas addressed by the proposals cover various sectors, for example, consumer, medical, communications, automotive, avionics, etc. A main focus

of the proposals received was embedded computing (in most cases, heterogeneous) for multi-processor systems on a chip and new programming paradigms, either to target the design of reconfigurable logic elements uniformly as part of a heterogeneous array, or to address the programming of data-driven and stream-based architectures. Programmable logic was a more general focus, including new approaches to compilation and synthesis. High-performance embedded computing and the issues that arise in hard real-time multi-processor systems were also addressed. Other areas covered by the proposals received include general purpose computing (addressing mainly parallel programming models) and high-performance computing (mainly software tools for programming multi-core systems and accelerators based on programmable logic).

This wealth of high-quality proposals addressing the full spectrum from hard real-time embedded computing up to high-performance computing shows that we have a vibrant European research community in the area of Computing Systems.

Over the next couple of months, our Cordis website will have more details about the selected projects.

With the 1st Call having already passed, we are now preparing the next Call in the area of Computing Systems that is scheduled for 2009. In this context, the HiPEAC roadmap is an important source of input. Additional input will be gathered from the workshop and session on Computing Systems organized during the "High Level Conference on Nanotechnologies" in Braga, Portugal, 19–21 November. You can find more information about this conference at <http://www.fct.mctes.pt/hlc/default.htm>.

I look forward to working with you as your new Project Officer.

Panos Tsarchopoulos
Project Officer



ACACES 2007, our Third International Summer School on Advanced Computer Architecture and Compilation, took place last July in L'Aquila, Italy. 203 students attended, of which 23 were female. The course offering was rated 'very good' by the participants. In parallel, HiPEAC industry members and other industry participants were invited to an industrial brainstorming meeting, which was very much appreciated. In addition, a woman's meeting took place, in which female participants had the opportunity to discuss gender issues with Barbara Ryder and Kathryn McKinley, both of whom are very active in the gender committees in the US. The week was concluded with a party, featuring a live rock & roll band, to the satisfaction of many; unfortunately, I was unable to attend.

The HiPEAC Winter Conference 2008 is scheduled for January 27–29, 2008 in Göteborg. Last July, the authors of

selected papers were already notified. In parallel to the Conference, there will be workshops and tutorials, which have proven very successful in previous years. As in past events, HiPEAC will provide grants to HiPEAC members' students to attend the Conference. Please visit the Conference website for further information.

A new Call for collaborations opened in early August and closed in mid-September. The Call included clusters and fellows, internships, 3-month travel grants and publications in top conferences. Proposals with participation of HiPEAC industry members are always welcome and encouraged. Besides clusters, HiPEAC continues to foster industry-academia relationships on high-performance by means of internships and industrial workshops.

The fourth HiPEAC Industrial Workshop

is being held in November 2007 in Cambridge, UK. As usual, representatives from companies and universities are expected to attend a week of events focusing on compiler and architecture technologies at the ARM premises. The HiPEAC General Cluster Meeting will be collocated with the HiPEAC Industrial Workshop. Take a look at our website for information on how to register for these events!

Finally, many members of the HiPEAC community have been present at the Symposium on the future of computing in Delft in memory of Stamatis Vassiliadis. Our participation is intended to be a small tribute to a very much-missed colleague and friend.

Mateo Valero
HiPEAC Coordinator

HiPEAC²

We are proud to announce that our community managed to have a HiPEAC² network of excellence approved by the commission in the Computing Systems research objective of the 1st Call for Proposals of the Information and Communication technologies in FP7. This means that HiPEAC will continue until the end of 2011 and we will be able to continue to further strengthen our community.

Although the scope has widened, the name of the new network stays the same, but now HiPEAC stands for "High Performance **and** Embedded Architecture and Compilation Network of Excellence". HiPEAC² will start on February 1, 2008. During a transition period of 7 months in which HiPEAC¹ and HiPEAC² will co-exist, the HiPEAC¹ activities will gradually become HiPEAC² activities.

The most important difference between HiPEAC¹ and HiPEAC² is that the bottom-up organization for the clusters has been replaced by a top-down organization. The changes are (i) that the number of

clusters is now fixed to 9; (ii) that topics of the clusters are predefined (based on the HiPEAC roadmap and the topics called for in the 1st call); and (iii) that each cluster is supported by a senior researcher.

The clusters will, however, have the same freedom as before to organize themselves the way they want, and the total budget spent on (collaborative) research has even increased. The clusters will be open to anybody who is contributing to the research and networking objectives of the clusters. The cluster performance will be reviewed during an annual review meeting by the steering committee. Our hope is that the nine clusters will grow into vibrant research communities that will further strengthen our community.

The HiPEAC² network will be officially launched during a kick-off meeting on January 30, in Göteborg, Sweden. I hope to be able to welcome many of you there.

Koen De Bosschere
HiPEAC² Coordinator



HiPEAC Journal

Transactions on High-performance Embedded Architectures and Compilers

The following papers were accepted this summer:

M. Ahn, Y. Paek.
Fast Code Generation for Embedded Processors with Aliased Heterogeneous Registers.
Transactions on High-Performance Embedded Architectures and Compilers, 2(2):40-59, 2007.

Aneesh Aggarwal.
Complexity Effective Bypass Networks.
Transactions on High-Performance Embedded Architectures and Compilers, 2(3):89-108, 2007.

HiPEAC 2008 Conference

The first two editions of the HiPEAC conference series in Barcelona (2005) and Ghent (2007) have really demonstrated that the topics covered by HiPEAC attract a lot of interest. This year, about 200 people attended the conference and its satellite events. It is a great privilege, but also a challenge, for me to host the third HiPEAC conference which takes place in Göteborg, Sweden on January 27–29, 2008.

The offerings of the next edition are rich and diverse. Prof. Mats Brorsson – the workshop and tutorial chairman – has collected a set of four workshops on topics that are all central to the HiPEAC roadmap: multi-cores, compiler optimizations, reconfigurable computing, and interconnection networks. Additionally, a tutorial on the Sunflower Toolsuite is also being offered.

The conference will be as rich, in terms of activities and presentations, as in previous years. 78 papers were submitted to the conference. The program chairs – Prof. Manolis Katevenis and Prof. Rajiv Gupta – together with their program committee have selected 25 regular papers to be presented at the conference and printed



in the proceedings. These papers span many important and timely topics, such as multi-core processors, reconfigurable systems, compiler optimization, power-aware techniques, and much more. The conference will also offer a keynote speech and several social events to promote interaction and exchange of ideas in informal settings such as a tour at the Universeum – a science exhibition center and aquarium – which is also where the banquet will take place.

After the main conference, several HiPEAC events will take place. Presumably, the most important one is the kick-off meeting for the new HiPEAC Network of Excellence under FP7 that will launch several new initiatives, with the existing HiPEAC as a platform. This kick-off meeting will take place on January 30, followed by HiPEAC research cluster meetings that will be run on January 30–31.



In conclusion, it will be a very full week of various events that are central to the activities in HiPEAC, and there are many reasons for you to attend. Please make sure that you visit the conference web page (<http://www.hipeac.net/conference>) and register for these great events in time.

See you in Göteborg in January next year!

Per Stenström
Chalmers University of Technology

Community news

With Peter Knijnenburg's passing away earlier this year, we lost a great scientist, colleague and friend. Peter made significant contributions in many areas and is probably best known for his work on iterative compilation where he showed creative insight many years before the subject became mainstream.

Peter was unusual in that he was highly versatile in his thinking and approach. This is probably best illustrated by his ability to move from physics to the "theory of theories" - category theory and then on to linear algebraic pro-

gram frameworks, iterative compilation and SMT quality of service. He was a highly prolific researcher publishing many papers in his short academic life and working with over 90 different co-authors. He possessed exceptional analytical skill and was able to succinctly communicate his research contributions. He was also generous with his time and ideas.

For many Peter Knijnenburg will be remembered for his warm hospitality when hosting the CPC workshop in Amsterdam 2003. Peter's passing is a



great loss to the research community. He is sadly missed and will always be remembered with affection by his colleagues.

Dilworth Decomposition



Do you know about a fundamental algorithm in lattices and order theory called Dilworth Decomposition? And what is its relationship with HiPEAC? Let me explain...

When you deal with code optimization and automatic parallelization for high performance or embedded computing, you may be asked to maximize the parallelism extraction in order to maximize the speedup. Such parallelism is constrained by the inherent data dependences present in the code. At any application granu-

larity (coarse grain: functions and process; medium grain: loop nests; fine grain: instructions), data dependences are represented by a graph where nodes represent tasks (functions, threads, instructions, etc.) and edges represent data dependences. If such a graph is acyclic, then Dilworth decomposition provides a polynomial time algorithm to partition such a directed acyclic graph into a minimal set of chains (paths) and, equivalently, to construct a maximal anti-chain (a maximal set of parallel nodes). Thanks to this algorithm, we

are able to analyze the maximal parallelism between tasks. It is also used to compute the maximal memory requirement (in terms of registers and scalar) of any valid schedule. In order to help the community to use such an algorithm, we are pleased to announce the release of a C++ implementation based on the LEDA graph library. Our implementation is dedicated to Vincent Bouchitté, who taught us well in graph theory, lattices, and orders.

<http://www.prism.uvsq.fr/~touati/sw/DD>

2nd CEDI in Zaragoza (Spain)

For the second time, most of the Spanish conferences on computer-related subjects, which are normally organized separately every year, have joined under the common umbrella of a big event called CEDI (Congreso Español de Informática). Following on from the first event in Granada, this 2nd CEDI was held in Zaragoza from September 11th to 14th. It consisted of 25 conferences, 23 invited talks, 9 workshops, and more than 1400 Spanish computer scientists and engineers, along with business partners and representatives. Fifteen scientific societies were also involved in the organization. One of the most active organizing groups was the Computer Architecture Group at Zaragoza

University (gaZ), which is an active member of HiPEAC.

Although the CEDI is a Spanish event, a large number of attendees were HiPEAC members from all over Europe. Three of the conferences in the CEDI were closely related to HiPEAC subjects: the XVIII Jornadas de Paralelismo (JJPAR 2007), the VII Jornadas de Computación Reconfigurable y Aplicaciones (JCRA 2007), the 2nd Simposio de Sistemas Tiempo Real, and the I Workshop en Robótica (ROBOT 2007). The staging of these computer architecture and compiler-related conferences brought together more than 300 researchers.

Prof. Yale Patt (Univ. of Texas) kicked-



off the plenary opening session with a conference entitled The Microprocessor Ten Years from Now: Why it is Relevant to all Informatics. A. Moshovos (Univ. of Toronto) was the invited speaker for the opening session of the JJPAR, where he talked about Region-Centric Memory Design. Robot demonstrations, product presentations, and round tables were also held.

Further information about CEDI can be found at <http://www.congresocedi.es>

International workshop on GCC for Research in Embedded and Parallel Systems (GREPS)

Researchers from INRIA Futurs, France and IBM Haifa Research Lab, Israel recently organized an international workshop on GCC for Research in Embedded and Parallel Systems (GREPS). The workshop took place on September 16 in conjunction with PACT 2007 in Brasov, Romania. This full-day workshop featured 11 presentations related to GCC by leading researchers and two distinguished keynote speeches by Paul Kelly from Imperial College, London entitled

"GCC in Software Performance Research: Just Plug In" and by Benoît Dupont de Dinechin from STMicroelectronics, Grenoble, France entitled "GCC for Embedded VLIW Processors: Why Not?" About 30 people registered and attended the workshop, which provided extensive, fruitful discussions about research with GCC. The overall feedback was excellent, with a unanimous interest in attending a second workshop. Attendees commented: "I would encourage my students to



participate"; "great opportunity to meet other researchers working in GCC". The support from HiPEAC was vital for the success of this workshop. For more information, please visit the workshop webpage at <http://sysrun.haifa.il.ibm.com/hr/greps2007>.

Systematic Analysis of the Cache Behavior of Irregular Codes

By **D. Andrade (dcanosa@udc.es)**
Prof. B.B. Fraguera, Prof. R. Doallo,
University of A Coruña, Spain
April 11, 2007

The performance of memory hierarchies, in which caches play an essential role, is critical in today's general-purpose and embedded computing systems because of the growing memory-bottleneck problem.

This thesis develops an automated framework for the fast and accurate prediction of the cache behaviour of codes with irregular access patterns. Namely, we propose an analytical modelling based on PME (Probabilistic Miss Equations) to predict the cache misses generated by irregular accesses originated by data-dependent conditional structures and indirections.

Experiments using the model predictions for different computer architectures and input data sets have proven its ability to successfully guide the optimizations of a compiler.

Fuzzy Computation

By **Carlos Álvarez Martínez**
(calvarez@ac.upc.edu)
Prof. Mateo Valero Cortés,
Dr. Jesús Corbal San Adrián
Universitat Politècnica de Catalunya
(UPC), Spain
May 16, 2007

This thesis presents the new paradigm of Fuzzy Computation for Multimedia that stands for trading computation precision by processor's performance and power

consumption. A lot of media applications trade precision by storage (JPEG, MP3), but this thesis' proposal not only introduces a new dimension (power/performance) but it also allows trade precision by the three factors at the same time (storage at application level, and power and performance at processor level).

Additionally, this thesis proposes and evaluates different techniques that implement fuzzy computation as Tolerant Reuse (instruction and region) or Tolerant

Arithmetic Units. Tolerant Reuse allows a classical reuse schema to make a hit when similar (not identical) inputs arrive. Tolerant Arithmetic, on the other hand, uses a simpler technique than the usual hardware to achieve a nearly exact computational result. These techniques increase the processor's performance, while also reducing its energy consumption at the cost of introducing errors undetected by human senses (i.e. maintaining high signal-to-noise ratios).

Optimizing VLIW Architectures for Multimedia Applications

By **Esther Salami San Juan**
(esalami@ac.upc.edu)
Prof. Mateo Valero Cortés
Universitat Politècnica de Catalunya
(UPC), Barcelona, Spain
June 1, 2007

This thesis concentrates on improving VLIW architectures in the context of a

multimedia workload. As part of this thesis, we analyze the problem of memory disambiguation in multimedia codes, and propose a dynamic memory-disambiguation technique specially targeted at multimedia loops or any other applications with similar memory access patterns. The proposal is fully described, implemented and evaluated, as well as compared against

advanced interprocedural pointer analysis. Second, we propose and evaluate enhancing a μ SIMD-VLIW core with conventional vector-processing capabilities to exploit the Data Level Parallelism (DLP) of multimedia kernels in a cost-effective way. Hardware simplicity and power efficiency are features of both vector and VLIW architectures, thus allowing for lower clock rates and lower voltages.

Low-power Issue Queue Design for Out-of-Order Processors

By **Marco A. Ramirez Salinas**
(mramirez@ac.upc.edu)
Prof. Mateo Valero (DAC-UPC),
Prof. Alex V. Veidenbaum (UCI),
Dr. Adrian Cristal (BSC)
and Dr. Luis Villa (CIC-IPN)
Universitat Politècnica de Catalunya
(UPC), Spain
June 7, 2007

We studied the issue logic complexities on its three internal operations: allocation,

wakeup and selection. Analysis concludes that after the clock-distribution network, the instruction issue logic is one of the major consumers of power. To measure the impact on the micro-architecture, three types of IQs were designed using SPICE CMOS and wires models. Results show that issue logic is a hot-spot in the chip; mainly due to the fact that wakeup is based on CAM. We propose and evaluate: 1) A mapping-block mechanism for partitioned RAM-CAM-based IQ design;

only blocks with a successor are explicitly enabled for wakeup comparison. 2) A direct instruction wakeup mechanism for a totally RAM-based IQ design, using a simple method for mapping pointers to first consumer, and handling instructions with multiple dependents by a second level of pointers. Our proposals save around 80% of IQ-power with minimal additional hardware and negligible loss of performance.

Data Assignment and Scheduling in Multi-tasked Embedded Systems

By José Ignacio Gómez Pérez
(jigomez@dacya.ucm.es)

Prof. Luis Piñuel, Prof. Manuel Prieto, Prof. Francky Catthoor
Universidad Complutense de Madrid, Spain

July 2007

Modern, embedded, multimedia applications require an enormous computa-

tional performance at low power consumption. Given their data-dominated behaviour, the memory system plays a crucial role in the whole optimization process. This work explores data allocation and scheduling techniques for software-controlled memory hierarchy, automatically generating a set of feasible mapping solutions that trade-off performance and energy consumption.

At every layer of the heterogeneous memory hierarchy we try to find the most energy-efficient data assignment, while taking into account task interactions, thus optimizing the usage of the available bandwidth. We finally propose a technique to integrate access scheduling with data placement in order to minimize the number of unexpected memory access conflicts at run time.

Application Security through Program Obfuscation

By Matias Madou
(matias.madou@elis.ugent.be)

Prof. Koen De Bosschere,
Ghent University, Belgium
September 7, 2007

Program obfuscation is a software technique used to protect the inner workings of a program by transforming its structure at the level of the data representa-

tion, the data flow, the control flow, and the individual instructions. This thesis proposes a framework consisting of three complexity metrics to evaluate the quality of obfuscating transformations. This framework is then applied to three structural obfuscating transformations. Several de-obfuscating transformations are then proposed, and it is shown that these de-obfuscating transforma-

tions effectively reduce the complexity of the obfuscated program back to the original complexity, proving that existing static obfuscation techniques can easily be broken. Finally, a new technique called trace obfuscation is presented. This technique proves to be much harder to break than static-program obfuscation techniques.

Memory Footprint Reduction for Operating-system Kernels

By Dominique Chanet
(Dominique.Chanet@elis.UGent.be),
Prof. Koen De Bosschere,
Ghent University, Belgium
September 20, 2007

In this thesis, we describe a number of techniques for substantially reducing the memory footprint of an operating-system kernel. As a first step, existing

link-time compaction techniques are adapted for use on the OS kernel. If the hardware and software configuration of the system are known and fixed, it is possible to go further by specializing the OS kernel for this known configuration. Superfluous functionality is removed from the kernel, while the remaining functionality is adapted for the specific usage patterns that can be derived from

the analysis of the system's hardware and software. As a last optimization, we introduce techniques that allow for the removal of rarely executed kernel code from memory. This code is then loaded on demand, only if and when it is needed. For a case study performed on the Linux 2.4 kernel, the combined techniques result in a memory footprint reduction of up to 48%.

Enhanced Applicability of Loop Transformations

By Martin Palkovic
(palkovic@imec.be)
Prof. Henk Corporaal,
Prof. Francky Catthoor
University of Technology Eindhoven,
The Netherlands
September 24, 2007

Data transfers and storage of large arrays in background memories are dominating contributors to the chip area and power consumption of embedded systems. High-

level memory optimizations contribute to their cost-efficient realization. In these optimizations, an important step involves loop transformations across the global program scope. These transformations can be performed on a geometrical model extracted from the program.

Due to the limitations of current geometrical models, the applicability of the loop transformations is limited. In this dissertation, we propose several applicability-

enhancing techniques for loop transformations. Automatic, hierarchical rewriting separates and encapsulates the details of the application into functions and hides undesired constructs. Our system scenario approach uses combined compile and run-time techniques to enlarge the exploration space for optimizations. Trade-off-oriented loop transformations show that targeting the best solution for locality is not necessarily optimal and that systematic trade-off guidelines are crucial.

Upcoming events

4th HiPEAC Industrial Workshop on Compilers and Architectures,

ARM Cambridge, UK, November 26, 2007, http://www.hipeac.net/industry_workshop4



Conference on Nanotechnologies

Braga, Portugal, November 20–21, 2007, <http://www.fct.mctes.pt/hlc/default.htm>

MICRO-40, The 40th Annual IEEE/ACM International Symposium on Microarchitecture.

Chicago, Illinois, USA, December, 1–5, 2007, <http://www.microarch.org/micro40/>

ISSPIT 2007, The 7th IEEE International Symposium on Signal Processing and Information Technology,

Cairo, Egypt, December 15–18, 2007, <http://www.isspit.org/isspit/2007/>

HiPEAC Workshops and Tutorials

Göteborg, Sweden, January 27, 2008

- Workshop on Statistical and Machine Learning approaches applied to Architectures and Compilation (SMART) <http://www.hipeac.net/smart-workshop.html>
- Workshop on Interconnection Network Architectures: On-Chip, Multi-Chip <http://www.disca.upv.es/jflich/wina/wina.html>
- Workshop on Reconfigurable Computing http://ce.et.tudelft.nl/HiPEACRC_WS/HiPEACRC_WS_2008/
- Workshop on Programmability Issues for Multi-Core Computers (MULTIPROG) <http://www.hipeac.net/smart-workshop.html>
- Tutorial on The Sunflower Toolsuite <http://www.es.ele.tue.nl/~pstanley/hipeac08-tutorial/>



HiPEAC Conference 2008

Göteborg, Sweden, January 28–29, 2008, <http://www.hipeac.net/hipeac2008>

HiPEAC² Kick-off meeting

Göteborg, Sweden, January 30, 2008

HPCA-14, 14th International Symposium on High-Performance Computer Architecture,

Salt Lake City, Utah, February 16–20, 2008, <http://www.ece.arizona.edu/hpca08/>

PPoPP'08 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming,

Salt Lake City, Utah, February 20–23, 2008, <http://research.ihost.com/ppopp08/>

SAC 2008, 23rd ACM Symposium on Applied Computing,

Fortaleza, Brazil, March 16–20, 2008, <http://www.acm.org/conferences/sac/sac2008>

ACM International Conference on Computing Frontiers,

Ischia, Italy, May 5–7, 2008, <http://www.computingfrontiers.org/>



The 5th International Conference on Ubiquitous Intelligence and Computing (UIC-08),

Oslo, Norway, June 23–25, 2008, <http://www.iu.hio.no/uic08/>

ACACES 2008, fourth HiPEAC Summer School

L'Aquila, Italy, July 13–19, 2008

Contributions

If you are a HiPEAC member and you want to contribute to this newsletter, please contact Sylvie Detournay at sylvie.detournay@elis.ugent.be