



GEMSCLAIM

Greener Mobile Systems By Cross Layer Integrated Energy Management


Thomas Fahringer, University of Innsbruck, Austria

Consortium:

University of Innsbruck
Queen's University Belfast
RWTH Aachen University
Politenica University of Timisoara


Start: September 1st, 2012

Duration: 3 years



Scientific background

- ▶ Personal computing: desktop computers → mobile systems
- ▶ Mobile systems are becoming more energy demanding
 - ▶ Radio interfaces
 - ▶ High computational load (games, multimedia apps.)
- ▶ “Major challenge for the mobile society in the next decade is that battery capacity is not keeping pace with Moore’s Law.”
by M. Muller, CTO of ARM
- ▶ Trade-off between energy and performance for mobile devices and many other computing systems:
 - ▶ desktop computing, cloud computing, high performance computing



Key challenges and potential impact

Project duration: Sept. 2012 – Aug. 2015

Cross layer energy management and optimization for mobile devices: HW/simulator, OS, compiler

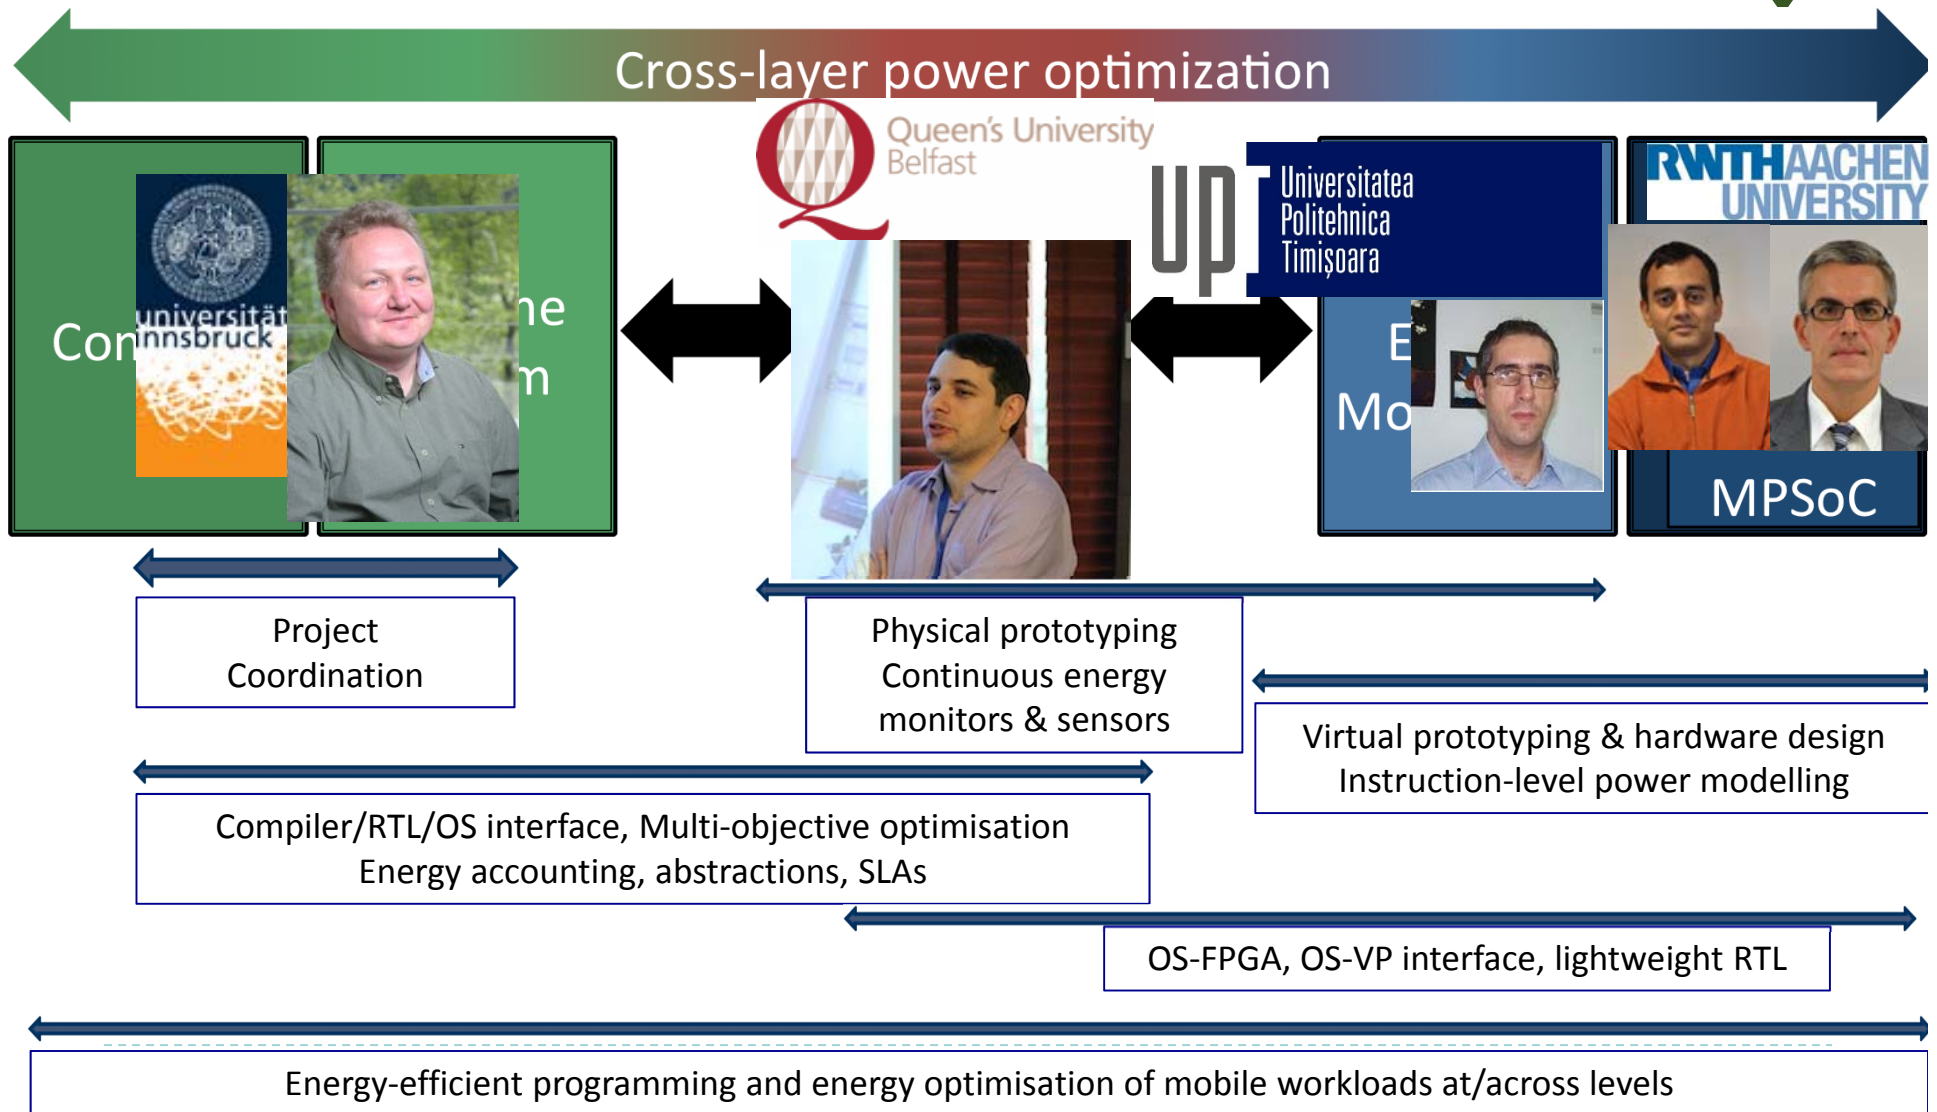
- ▶ Energy-aware optimizing and parallelizing compiler
- ▶ Energy-proportional operating system
- ▶ Customizable HW modeling with energy monitoring facilities
 - ▶ simulator
 - ▶ FPGA


Potential impact

- ▶ control trade-off between energy optimization and performance
- ▶ additional 30 % energy saving for mobile terminals



Consortium and synergies






Work plan and milestones

Milestone	Delivery month	WP involved	Title
M1	6	WP1, WP2	Requirements specification and design of the GEMSCCLAIM HW/SW environment
M2	15	WP1, WP3-WP7	Early prototype of the GEMSCCLAIM HW/SW environment (able to simulate/execute benchmarks applications)
M3	30	WP1, WP3-WP7	Final prototype of the GEMSCCLAIM HW/SW environment (able to perform energy savings benchmarking)
M4	36	WP1, WP7	Benchmarks fully optimized for energy (targeted savings: 30%) on FPGA with GEMSCCLAIM SW; project completion




Project Management

- ▶ **Internal project meetings: every 6 - 8 months**
 - ▶ Timisoara, Feb. 2014, Innsbruck, Sept 2014; Belfast, April 2015
- ▶ **Skype meetings**
 - ▶ once a month
- ▶ **Scientists exchange:**
 - ▶ bilateral meetings (1-2 weeks)
 - ▶ extend project meetings by a few days for technical work
 - ▶ 2 months visit by Phd student
- ▶ **Deliverables and documents**
 - ▶ internal review
 - ▶ Annual reports for chist-era offices and national funding agencies
- ▶ **Financial reporting**
 - ▶ Once a year
- ▶ **document repository**
 - ▶ dropbox
- ▶ **software repository**
 - ▶ gitlab server in Belfast
- ▶ **webpage**
 - ▶ www.gemsclaim.eu



Project Management

- ▶ Quality Assurance
 - ▶ roles and responsibilities for each WP
 - ▶ time plan controlled by project coordinator
 - ▶ redmine
 - ▶ interaction with WP leaders
 - ▶ risk management
 - ▶ regular updates with input from WP leaders
 - ▶ SW Engineering
 - ▶ Agile software development
 - ▶ Test-driven development with focus on requirements and goals



Resources and Funding


- ▶ For projects 1 and 2


N ^o	Partner	Person. months	Total costs in €	Percentage of requested budget
1	UIBK	40,375	143.598,24	43,13
2	QUB	26,4	262.944	58,3
3	UPT	23,314	71.226,00	28,49



Simulation based Power Models

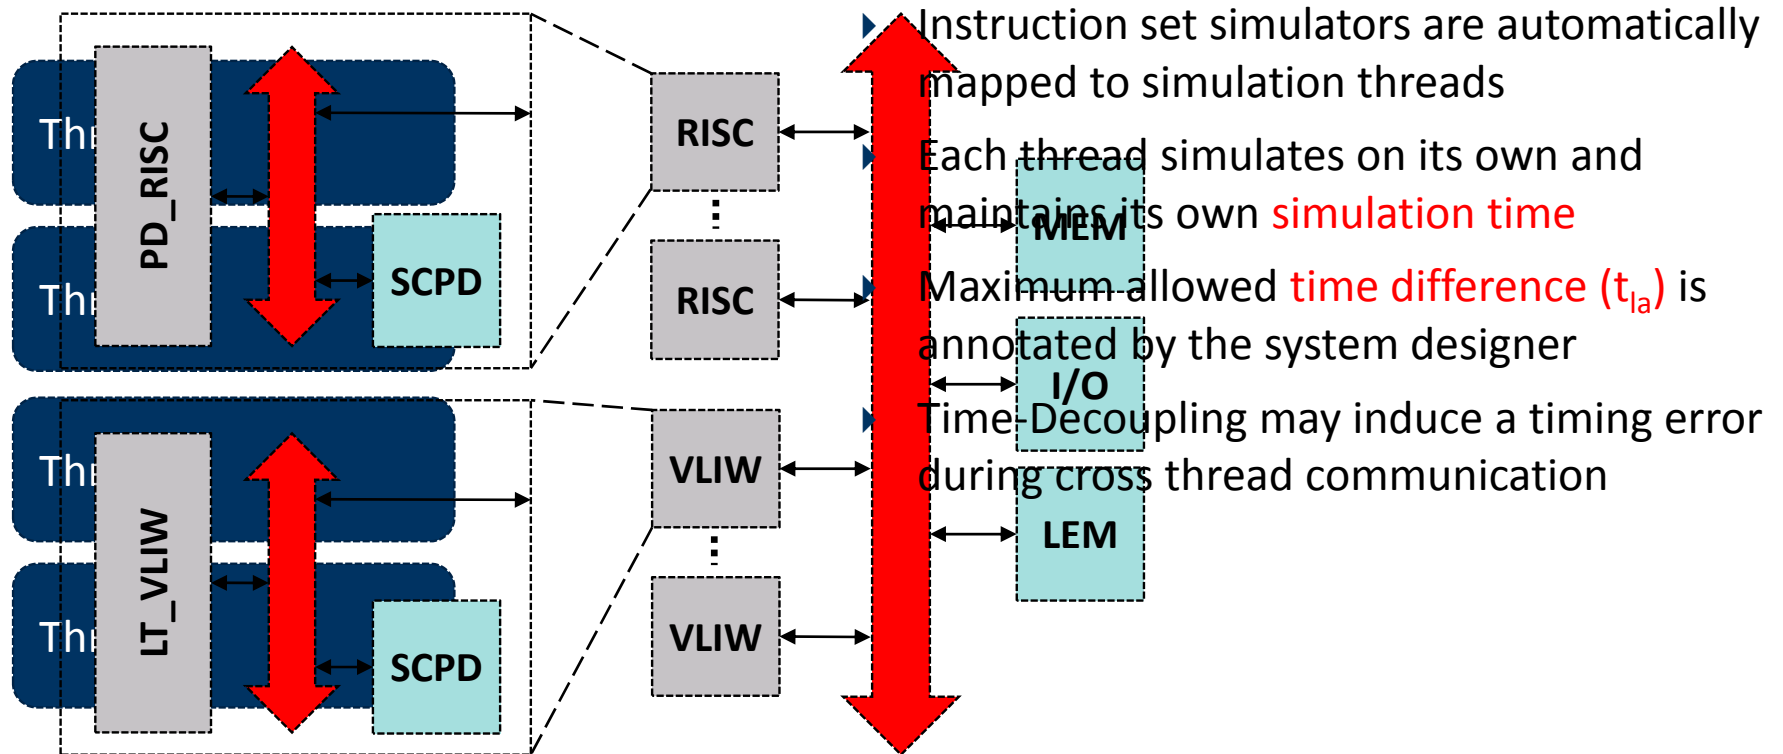
- ▶ Power Models for Buses and Memories estimated from hardware counters
- ▶ Utilization counters (u_{active} , u_{read} , u_{write}) updated live during simulation
- ▶ Power coefficients c_i calibrated offline with real hardware





Parallel SystemC Simulation

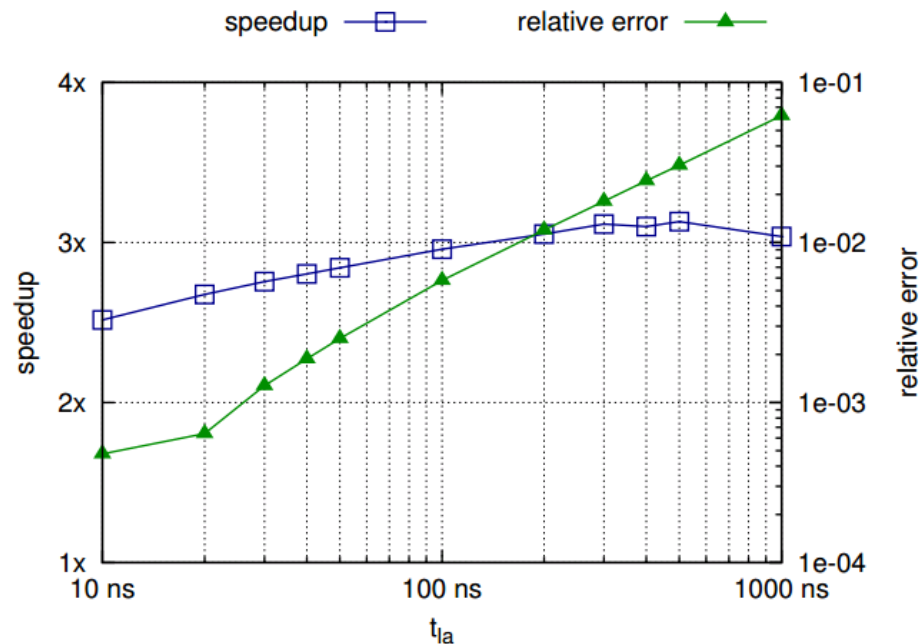
- ▶ Contemporary heterogeneous multi-core system
 - ▶ System parameters (number/type of cores, mem size...) statically configurable
- ▶ GEMSCLAIM simulator is based on SystemC
 - ▶ Exchanging reference SystemC with ICE's SCOPE SystemC kernel enables parallel simulation





Simulation speed vs. accuracy

- ▶ Test setup (parallel simulator uses 4 threads):
 - ▶ 8 RISC + 8 VLIW system
 - ▶ Application: “ocean” from SPLASH-2 benchmark suite



- 3x speedup possible with ~1% error in timing
- Timing error increases linearly with time decoupling

Validation of power models by physical measurements



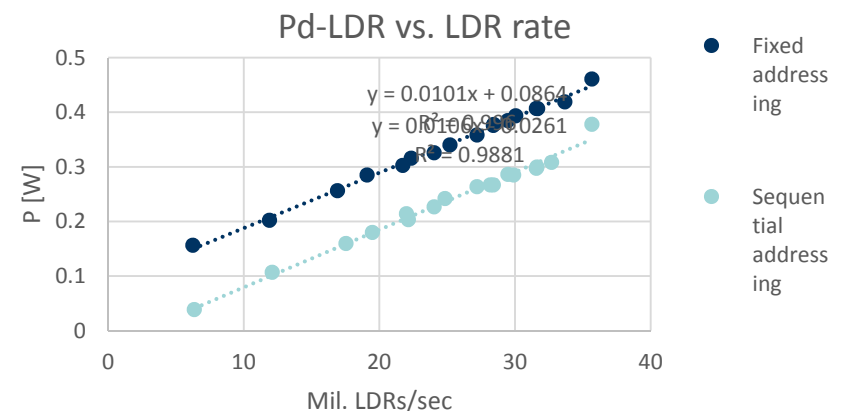
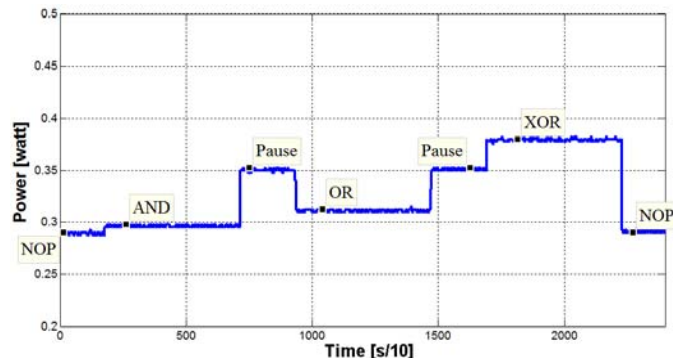
- ▶ Aim

- ▶ Validate and calibrate the power models implemented by GEMSCCLAIM simulator by physical measurements using FPGA prototyping


- ▶ Measurement methodology

- ▶ FPGA components power profiling: LUT, PLL, DSP, BRAM
- ▶ PD_RISC core instruction set power profiling
- ▶ Power profiling of Microblaze cores, DDR3 memory, and AXI interconnects
 - ▶ Instruction level power profiling of MB core
 - ▶ Transaction level power profiling of DDR3 memory and interconnects

- ▶ Measuring samples



Validation of power models by physical measurements



▶ Results

- ▶ Power consumption correlation matrix between simulation values and physical measurements

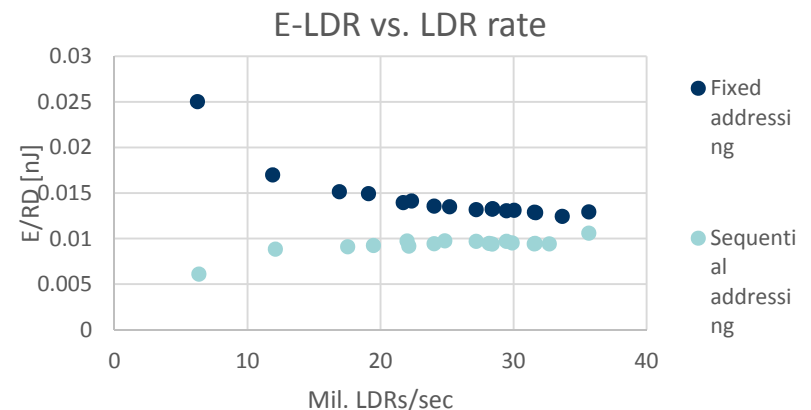
	130nm (faraday)	90nm (faraday)	65nm (faraday)	45nm (Nangate)	45nm (Atlys)
130nm (faraday)	1				
90nm (faraday)	92.11%	1			
65nm (faraday)	98.95%	92.53%	1		
45nm (Nangate)	98.67%	93.16%	98.37%	1	
45nm (Atlys)	85.78%	95.47%	84.34%	86.39%	1

- ▶ Power and energy models for DDR3

- ▶ $P_{\text{DRAM}} = C_1 x u_{\text{RD}} + C_2 x u_{\text{WR}} + P_s$
- ▶ $E_{\text{DRAM}} = N_{\text{RD}} x E_{\text{RD}} + N_{\text{WR}} x E_{\text{WR}} + P_s x t$

▶ Publications


- ▶ 2 joint papers
 - ▶ 1 accepted, 1 submitted
- ▶ 1 single partner paper






Hardware prototyping and energy monitor

- ▶ Load and Energy Monitor (LEM) architecture
 - ▶ Shared bus based infrastructure (adapted from Wishbone bus)
 - ▶ Components
 - ▶ Central LEM
 - ▶ LEM Bus
 - ▶ LEM Sensors





Rethinking the OS



Energy slices (first-class resource)



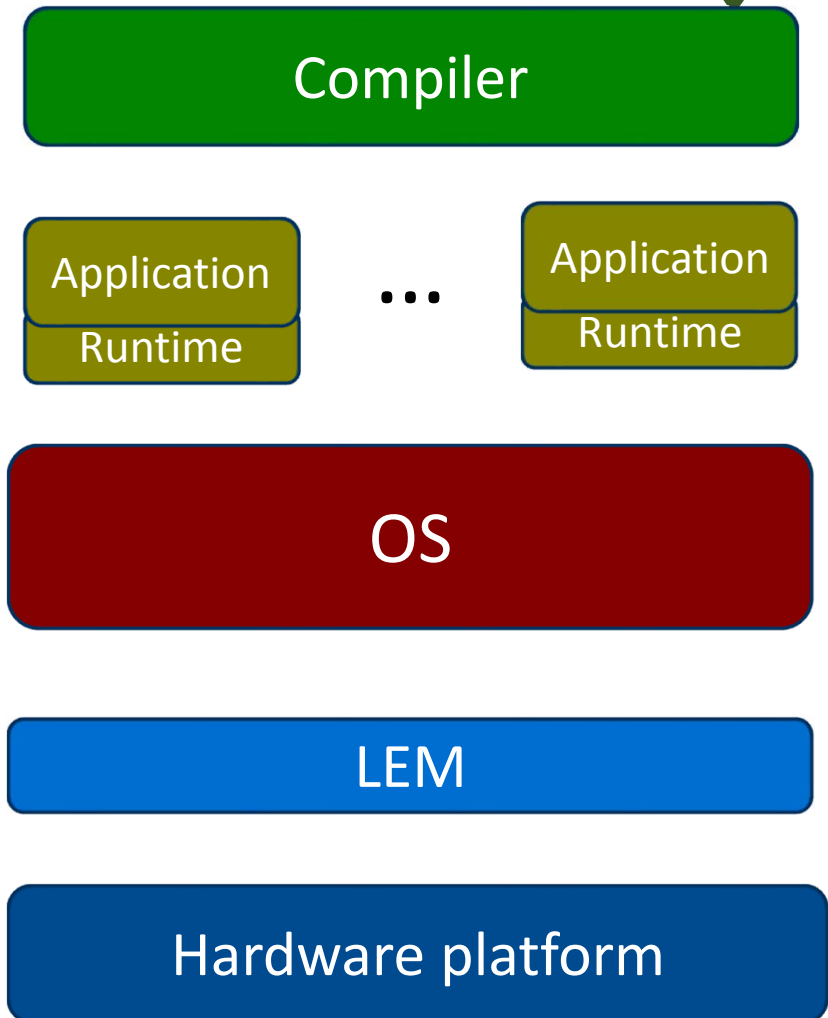
Processor time slices (within energy budget)





GEMSCLAIM OS services


- ▶ POSIX threading interface
 - ▶ Shared Memory abstraction
- ▶ Manage energy consumption of HW components
 - ▶ LEM
- ▶ Export PM interfaces to the application layer
 - ▶ RAPMI (resource allocation and power management interface)
- ▶ Implement energy-aware scheduling policies
 - ▶ Pre-emptive multithreading
 - ▶ Energy Interrupt
 - ▶ Synergy with Application Runtimes
- ▶ Ports
 - ▶ Virtual Platform (simulator)
 - ▶ FPGA Microblaze MPSoC
 - ▶ Legacy platforms (Linux-based)






RAPMI: Component-level power management

- ▶ RAPMI
 - ▶ `rapmi_ctrl(mode, rate, power_level)`
 - ▶ sensor mode
 - ▶ rate
 - ▶ power_level
 - ▶ `rapmi_start()`
 - ▶ Start sensor measuring
 - ▶ `rapmi_stop()`
 - ▶ Stop sensor measuring
 - ▶ `rapmi_read(rapmi_t *pm)`
 - ▶ Read sensor values
 - ▶ `rapmi_set_freq(freq)`
 - ▶ Set core frequency
 - ▶ `rapmi_get_freq()`
 - ▶ Returns core frequency
- ▶ Sensor control
 - ▶ Start, Stop
 - ▶ Read
 - ▶ Set sampling rate
- ▶ HW control
 - ▶ DVFS
 - ▶ Halt cores



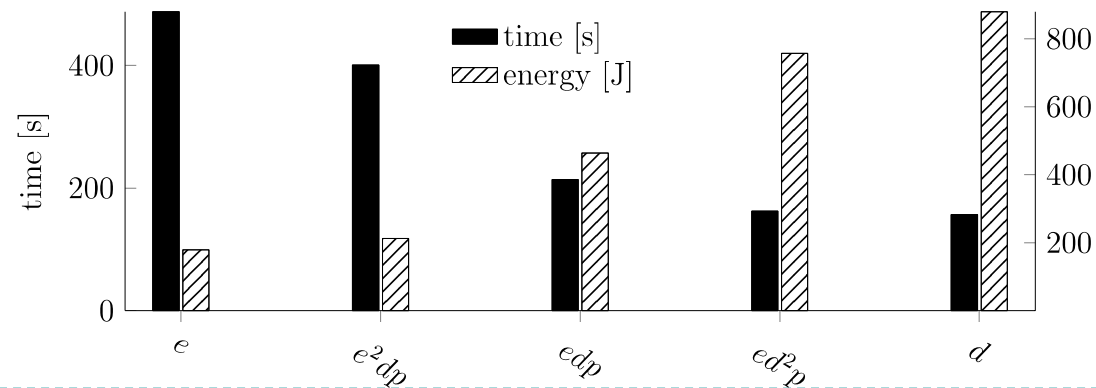
Rethinking OS scheduling


- ▶ Distribute *energy* slices instead of time slices
 - ▶ Energy interrupts (LEM)
- ▶ Priority-based policies
 - ▶ Interactive applications (User assisted)
- ▶ Multi-objective optimization
 - ▶ Energy *and* Performance
 - ▶ Multi-program workloads
- ▶ HW and SW re-configuration
 - ▶ HW: DVFS, Halt cores
 - ▶ SW: DCT (Dynamic Concurrency Throttling) – Synergy between OS and Application Runtimes



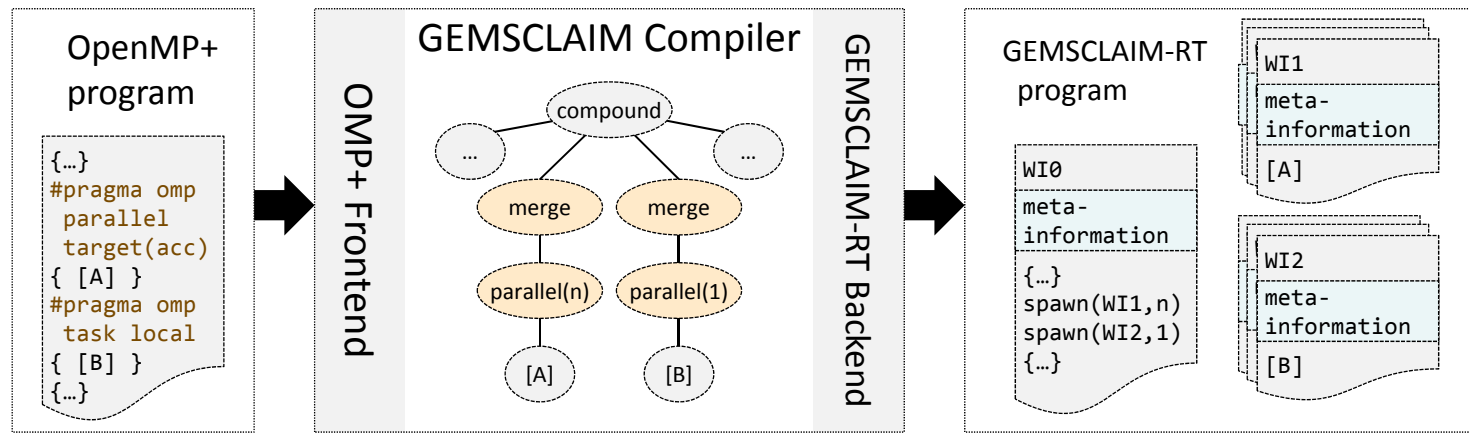
Some results

- ▶ ODRROID-XUE platform
 - ▶ Cores, Memory power sensors
 - ▶ Heterogeneous cores
- ▶ Multi-program workloads
 - ▶ Per-program characterization
 - ▶ Synergistic scheduling (OS runtime and Application Runtime)
- ▶ Different optimizing schedulers
 - ▶ Energy (e)
 - ▶ Delay (d)
 - ▶ Energy-delay products (edp, e^2dp , ed^2p)






Compiler and runtime system




- ▶ The GEMSCCLAIM Compiler is a source-to-source C compiler
 - ▶ supports C programs annotated with OpenMP+
 - ▶ energy-aware extension to OpenMP
- ▶ Programs are analyzed and divided into meaningful code regions
 - ▶ individually tunable and annotated with additional metadata
- ▶ Execution orchestrated by the runtime system
 - ▶ handles dynamic decision making and scheduling
 - ▶ communicates with the GEMSCCLAIM OS.



Compiler runtime system

- ▶ Compiler runtime fully ported to relevant platforms
 - ▶ GEMSCLAIM virtual platform
 - ▶ ARM boards (e.g. ODROID XU+E)
 - ▶ X86/Linux baseline version
- ▶ Supports *upcalls* from OS/low-level layer
 - ▶ E.g. `irt_set_dop(uint32)` allows OS to control per-program degree of parallelism
 - more effective parallelism management than simple thread packing!




OpenMP extensions (OpenMP+)

- ▶ Now fully supported in compiler and runtime system:
 - ▶ Region construct (handling code outside OpenMP regions)
 - ▶ Objective clause (defining multiple optimisation objectives and constraints)
 - ▶ Param clause (define tunable parameters for compiler optimisation)

- ▶ New: *quality of service* constraints

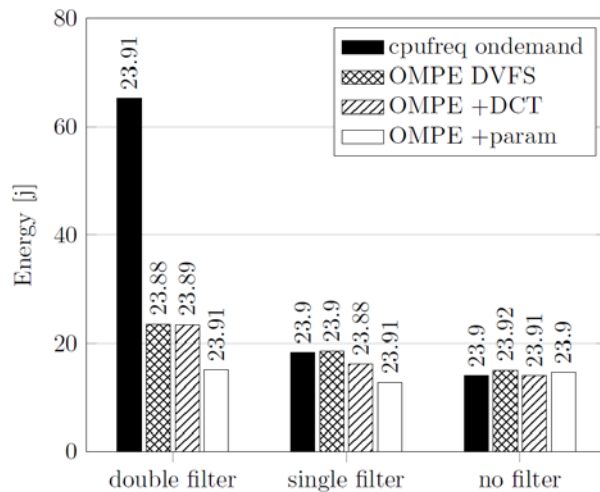
- ▶ Can be mapped to user-defined parameters
- ▶ Example usage in video decoder:

```
1  #pragma omp parallel for schedule(dynamic)
    objective(E : T<1/f_rate; Q<3) param(scaling, range(1:8:1))
2  for (int y=0; y<rows; y+=2*scaling)
3    for (int x=0; x<cols_2; x+=scaling) {
4        ...
5        if(scaling > 1) { ... }
6    }
```

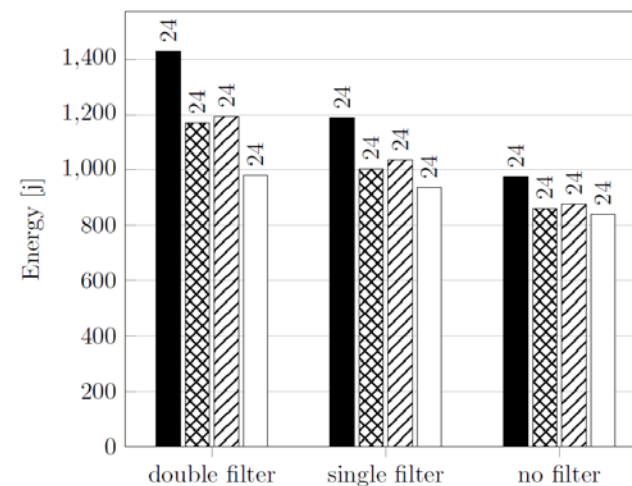


e-optimizer

- ▶ Multi-objective optimization for OpenMP+ programs
 - ▶ Dynamically searches best configuration for given goals and constraints
 - ▶ Combines random sampling (to prevent local minima) with multi-dimensional hill climbing (to quickly converge)
 - ▶ Up to 77% energy savings on mobile and 31% on desktop



Mobile platform, 704x576 resolution



Desktop platform, 1408x1152 resolution



GEMSCCLAIM Publications 2012 and 2013

- ▶ A Multi-Objective Auto-Tuning Framework for Parallel Codes Herbert Jordan, Peter Thoman, Juan J. Durillo, Simone Pellegrini, Philipp Gschwandtner, Thomas Fahringer, and Hans Moritsch . SC '12 November 11 - 15, 2012
- ▶ Fast Dynamic Binary Rewriting to Support Thread Migration in Shared-ISA Asymmetric Multicores: Best Paper Award, Georgakoudis, G., Nikolopoulos, D. & Lalis, S. Feb 2013 Proceedings of the First ACM International Conference on Code Optimisation for Multi- and Many-Cores (COSMIC). ACM New York, NY: ACM, p. 4:1 4 p.
- ▶ Prefetching and Cache Management using Task Lifetimes, Papaefstathiou, V., Katevenis, M. G. H., Nikolopoulos, D. & Pnevmatikatos, D. Jun 2013 27th ACM International Conference on Supercomputing . New York, NY,USA: ACM, p. 325-334 10 p.
- ▶ Adaptive Granularity Control in Task Parallel Programs using Multiversioning. Peter Thoman, Herbert Jordan, Thomas Fahringer. Euro-Par 2013 - Aug 26 Aachen, Germany
- ▶ INSPIRE: The Insieme Parallel Intermediate Representation. Herbert Jordan, Simone Pellegrini, Peter Thoman, Klaus Kofler, Thomas Fahringer. PACT 2013, September 7-11, 2013.
- ▶ BDDT: Block-level Dynamic Dependence Analysis for Task-Based Parallelism (full version), Tzenakis, G., Papatriantafyllou, A., Vandierendonck, H., Pratikakis, P. & Nikolopoulos, D. Sep 2013 Proceedings of the 2013 International Conference on Advanced Parallel Processing Technology. Springer, (Lecture Notes in Computer Science)
- ▶ Programming the Energy Efficiency of High Performance Computing Systems: Keynote Talk: Proceedings of the 4th International Conference on Energy-Aware High Performance Computing (Dresden, Germany) Nikolopoulos, D. Sep 2013 1 p.
- ▶ Deterministic Scale-Free Pipeline Parallelism with Hyperqueues, Vandierendonck, H., Chronaki, K. & Nikolopoulos, D. Nov 2013, Proceedings of Supercomputing'13: International Conference for High Performance Computing, Networking, Storage and Analysis (SC). New York, NY, USA: ACM
- ▶ A Case Study of Automated Testing Implementation in the Automotive Industry, D. Ganea, R. Bogdan, V. Ancusa, M. Popa. 14th IEEE International Symposium on Computational Intelligence and Informatics, CINTI 2013, Budapest, November 19-21, 2013.
- ▶ A. Power Modeling and Estimation during ADL-driven Embedded Processor Design, Wang, Z., Wang, L., Xie, H., and Chattopadhyay, in 4th International Conference on Energy Aware Computing Systems & Applications (ICEAC), Dec 2013, Istanbul, Turkey




GEMSCLAIM Publications 2014

- ▶ M. Marcu, C. Cernazanu, Applications of Smart Metering and Home Appliances' Power Signatures, Proceedings of 2014 IEEE International Instrumentation and Measurement Technology Conference (I2MTC), Montevideo, Uruguay, May 2014
- ▶ Fast Dynamic Binary Rewriting for flexible thread migration on shared-ISA heterogeneous MPSoCs. Georgakoudis, G.; Nikolopoulos, D.S.; Vandierendonck, H.; Lalis, S., Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIV), 2014 International Conference on , vol., no., pp.156,163, 14-17 July 2014
- ▶ R.N. Vasiliu, M. Popa, M. Marcu, Wireless programmable thermostat using Raspberry Pi, Proceedings of 6th International Workshop on Soft Computing Applications (SOFA) Timisoara, Romania, Jul. 2014
- ▶ TProf: An Energy Profiler for Task Parallel Programs. / Manousakis, Ioannis; Zakkak, Foivos S.; Pratikakis, Polyvios; Nikolopoulos, Dimitrios. In Sustainable Computing: Informatics and Systems, 2014.
- ▶ Multi-Objective Auto-Tuning with Insieme: Optimization and Trade-Off Analysis for Time, Energy and Resource Usage. Philipp Gschwandtner, Juan J. Durillo, Thomas Fahringer. Euro-Par 2014, Aug 25-29 Porto, Portugal.
- ▶ On the Potential of Significance-Driven Execution for Energy-Aware HPC. Philipp Gschwandtner, Charalampos Chaliou, Dimitrios S. Nikolopoulos, Hans Vandierendonck, Thomas Fahringer. EnA-HPC 2014 - Sep 1-2 Dresden, Germany.
- ▶ Compiler Multiversioning for Automatic Task Granularity Control. Peter Thoman, Herbert Jordan, Thomas Fahringer. Concurrency and Computation: Practice and Experience, Volume 26, Issue 14, September 25, 2014.
- ▶ Power-Capped DVFS and Thread Allocation with ANN Models on Modern NUMA Systems. Imamura, Satoshi; Sasaki, Hiroshi; Inoue, Koji; Nikolopoulos, Dimitrios. Proceedings of the 32nd IEEE International Conference on Computer Design (ICCD). Institute of Electrical and Electronics Engineers (IEEE), 2014.
- ▶ C. Cernazanu, S. Fedea, A. Amaricai, M. Marcu, Energy Profiling of FPGA Designs, Proceedings of 2014 IEEE International Symposium on Robot and Sensors Environments (ROSE), Timisoara, Romania, Oct. 2014
- ▶ S. Fuicu, A. Avramescu, D. Lascu, R. Padurariu, and M. Marcu, Real-time monitoring using finite state-machine algorithms, HealthyIoT Conference Proceedings (HealthyIoT) si Lecture Notes of ICST, Roma, Italy, Oct. 2014
- ▶ On the viability of microservers for financial analytics. Charles J Gillan, Dimitrios S. Nikolopoulos, Giorgis Georgakoudis, Richard Faloon, George Tzenakis, and Ivor Spence. In Proceedings of the 7th Workshop on High Performance Computational Finance (WHPCF '14)




GEMSCCLAIM Publications 2014 and 2015

- ▶ Power Modeling and Capping for Heterogeneous ARM/FPGA SoCs. Wu, Yun; Nunez-Yanez, Jose; Woods, Roger; Nikolopoulos, Dimitrios. Proceedings of the 2014 International Conference on Field- Programmable Technology (FPT). IEEE Computer Society, 2014
- ▶ Low-Cost Hardware Infrastructure for Runtime Thread Level Energy Accounting, Marius Marcu, Oana Boncalo, Madalin Ghenea, Alexandru Amaricai, Cosmin Cernazanu, Energy Efficiency with Heterogeneous Computing Workshop (EEHCO), Amsterdam, Holland, Jan. 2015
- ▶ Direct FPGA-based Power Profiling for a RISC Processor, Cosmin Cernazanu-Glavan, Marius Marcu, Alexandru Amaricai, Stefan Fedea, Madalin Ghenea, Zheng Wang, Anupam Chattopadhyay, Jan Weinstock, Rainer Leupers, Accepted to 2015 IEEE International Instrumentation and Measurement Technology Conference (I2MTC), Pisa, Italy, May 2015
- ▶ Hardware support for performance measurements and energy estimation of OpenRISC processor, Submitted to 10th IEEE International Symposium on Applied Computational Intelligence and Informatics (SACI), Timisoara, Romania, May 2015
- ▶ Low-Cost Hardware Infrastructure for Runtime Thread Level Energy Accounting, UPT, RWTH and QUB, Invited to IET CDT Journal, Deadline Mar. 2015
- ▶ Low-Cost Hardware Infrastructure for Runtime Shared Memory Thread Level Energy Accounting, Marius Marcu, Oana Boncalo, Madalin Ghenea, Jan Henrik Weinstock; Rainer Leupers, Submitted to 8th ACM International Systems and Storage Conference (SYSTOR), Haifa, Israel, May 2015
- ▶ Parallel SystemC Simulation for ESL Design using Flexible Time Decoupling. Jan Henrik Weinstock; Rainer Leupers; Gert Ascheid. International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation (IC-SAMOS), IEEE Computer Society, 2015 (submitted)
- ▶ OpenMPE: A Language Extension for Application-level Energy Awareness. Ferdinando Alessi, Peter Thoman, Giorgis Georgakoudis, Thomas Fahringer, and Dimitrios S. Nikolopoulos, (submitted)
- ▶ Peter Thoman, Giorgis Georgakoudis, Ferdinando Alessi, Thomas Fahringer, Dimitrios S. Nikolopoulos, "A Unified Runtime Framework To Optimize Energy and Performance On Multi-program Workloads", to be submitted to SC'15




Important GEMSCCLAIM project objectives

- ▶ Cross layer energy management and optimization for mobile devices: HW/simulator, OS, compiler
- ▶ Control trade-off between energy optimization and performance
- ▶ additional 30 % energy saving for mobile terminals
- ▶ **GEMSCCLAIM achieved all of that.**



Conclusions and way forward

- ▶ Cross-layer energy management and optimization requires breaking barriers between layers
 - ▶ Common abstractions, metrics
 - ▶ Synergetic optimization approaches
- ▶ Measuring energy consumption remains challenging
 - ▶ Machine-specific, intrusive, coarse-grain
 - ▶ Hybrid modelling & measurement approaches are the only viable
- ▶ Software needs to evolve to break the energy wall
 - ▶ Energy optimization should become explicit
 - ▶ Sensitivity of energy to software structures is not well understood
- ▶ *GEMSCCLAIM provides promising solutions for these challenges in a holistic approach, demonstrated on a physical HW/SW substrate*



Sustainability

- ▶ Joint FETHPC proposal
- ▶ Part of the SW under Apache 2.0 license (open source)
- ▶ Parts of the GEMSCLAIM OS Runtime are used in current research projects (FP7 NanoStreams) and are a foundation for future research proposals under the Horizon 2020 research programme
- ▶ Parallel simulation technology developed in GEMSCLAIM will be used for upcoming industry collaboration projects.
- ▶ The GEMSCLAIM VP will be used as a basis for future research projects proposed in the context of Horizon-2020 Research Programme of the EU.

www.gemsclaim.eu