GEMSCLAIM

Greener Mobile Systems By Cross Layer Integrated Energy Management

Thomas Fahringer and Peter Thoman, University of Innsbruck, Austria

Consortium:

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

Start: September 1, 2012

End: August 31, 2015

Chist-era Projects Seminar 2016, Bern, April 28, 2016

Scientific background

Smartphone sales (billions)

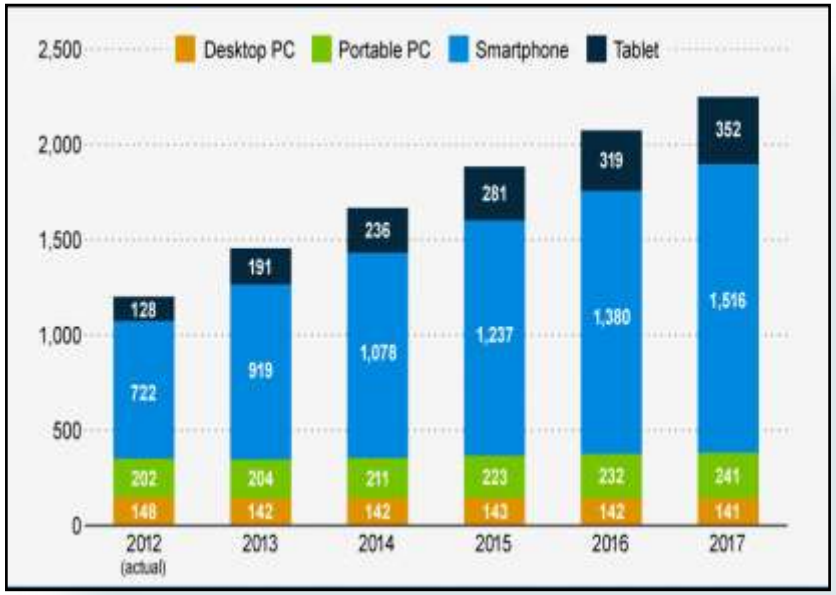

Improvements wanted by mobile device users (%)

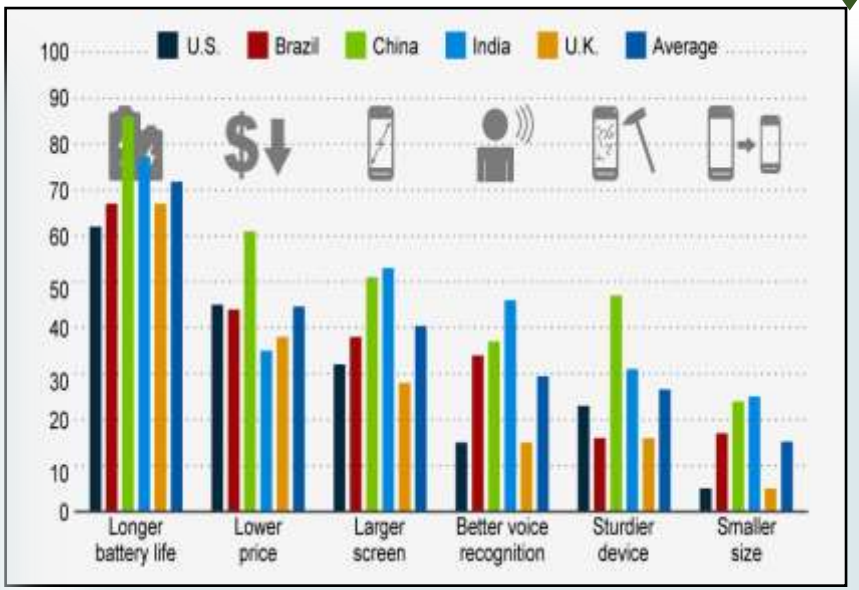

- Mobile systems are becoming more energy demanding
- Longer battery life is the most desired user feature.
- Trade-off between energy and performance for mobile devices and many other computing systems: desktop computing, cloud computing, HPC
- Mobile devices are inherently heterogeneous
  - Risc-like application processor, DSP for baseband processing

chist-era

Key challenges and potential impact

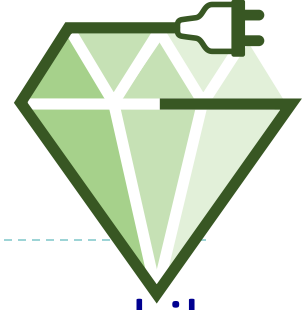

Cross layer energy management and optimization for mobile devices: HW/simulator, OS, compiler

- 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 savings for mobile terminals

Consortium and synergies

Cross-layer power optimization

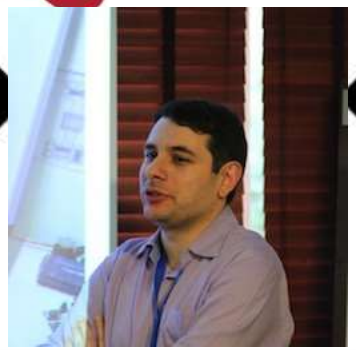

Project Coordination

Physical prototyping
Continuous energy monitors & sensors

Virtual prototyping & hardware design
Instruction-level power modelling

Compiler/RTL/OS interface, Multi-objective optimisation
Energy accounting, abstractions, SLAs

OS-FPGA, OS-VP interface, lightweight RTL

Energy-efficient programming and energy optimisation of mobile workloads at/across levels

GEMSCLAIM Software and Hardware Stack

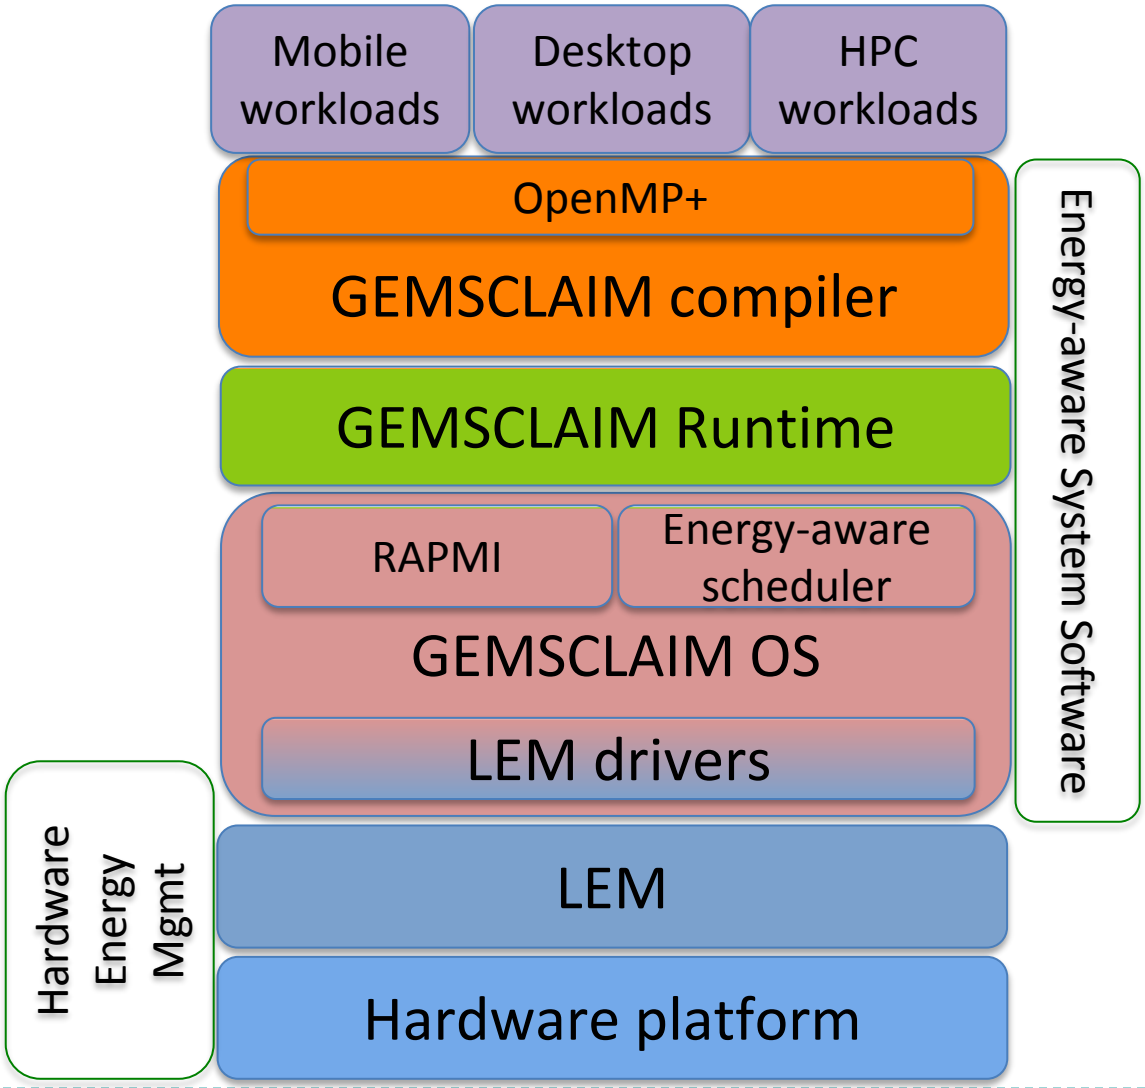

Parallel SystemC Simulation

- 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

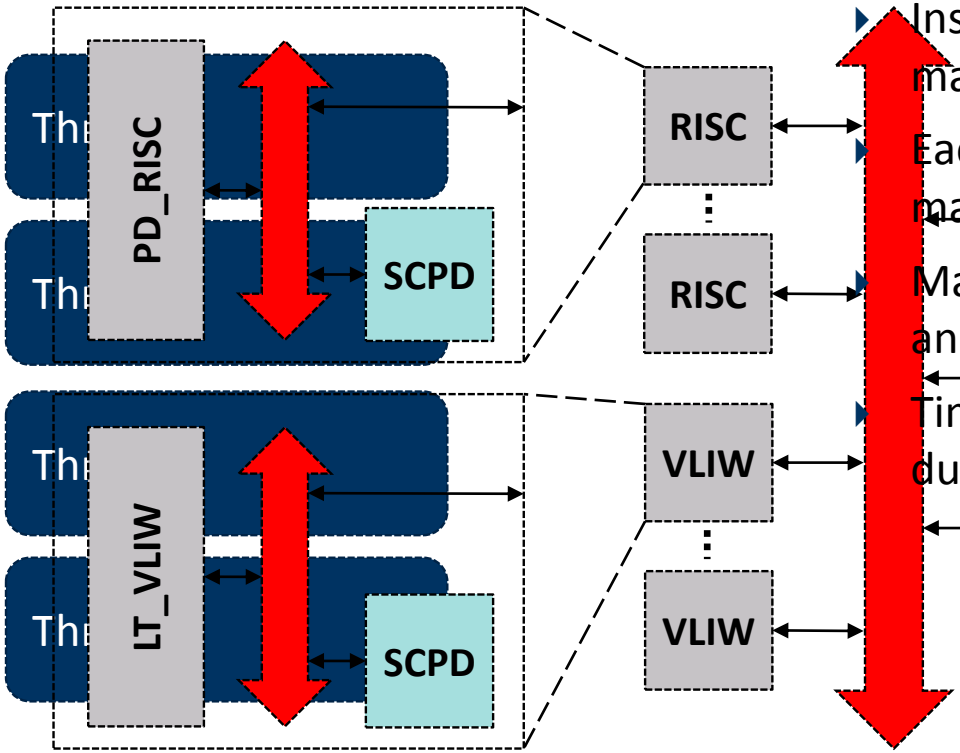

- Instruction set simulators are automatically mapped to simulation threads
- Each thread simulates on its own and maintains its own simulation time
- Maximum allowed time difference (t_{la}) is annotated by the system designer
- Time-Decoupling may induce a timing error during cross thread communication

chist-era

Simulation speed vs. accuracy

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

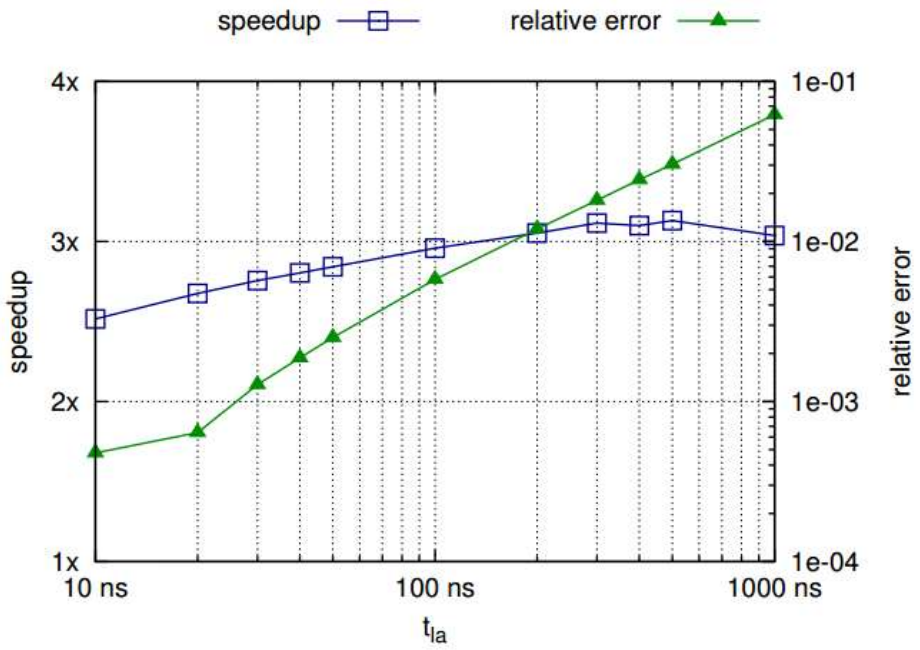

- 3x speedup possible with ~1% error in timing
- Timing error increases linearly with time decoupling

FPGA prototype with PTEA

- Per Thread Energy Accounting (PTEA)
  - Per-core energy accounting (HW)
    - Shared resources
      - Monitor bus transactions
      - Account the energy of the transaction to the core which initiated it
      - Use the interconnect existing support to identify the source of the transaction
    - Local resources
  - Per-thread energy accounting (OS)
    - Energy counters part of thread context

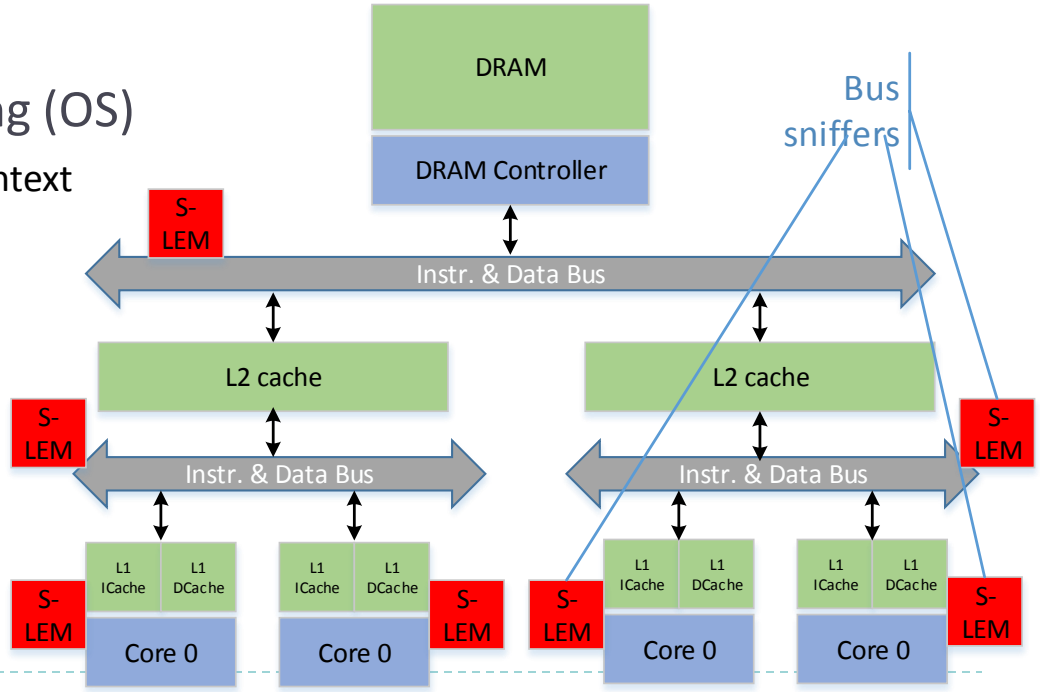

chist-era

FPGA prototype with PTEA

- Platform
  - ZC 702 board with Xilinx Zynq 7020 device
    - PL – programmable logic
    - PS – processing system
- System
  - ARM Cortex A9 Dual Core processing system
    - ARM0 – used to configure and monitor LEM
  - Target reference design block
    - 4 processing cores (Microblaze), with local interrupt controller and local memory
    - LEM sensors: 2 sensors/core, 2 shared sensors/memory
  - Target design
    - 2 clusters of 4 cores blocks
    - 8 Microblaze cores, with local interrupt controller, local timer and local memory
    - 2 LEMs

FPGA prototype with PTEA

- Calibrate and validate FPGA prototype
- WCET benchmarks
  - No OS / bare metal implementation
  - Small enough to be executed from BRAM too
  - Repeatable/ deterministic
  - 2-3 minutes/test x 2/3 tests x 2 boards
  - Normalized energy

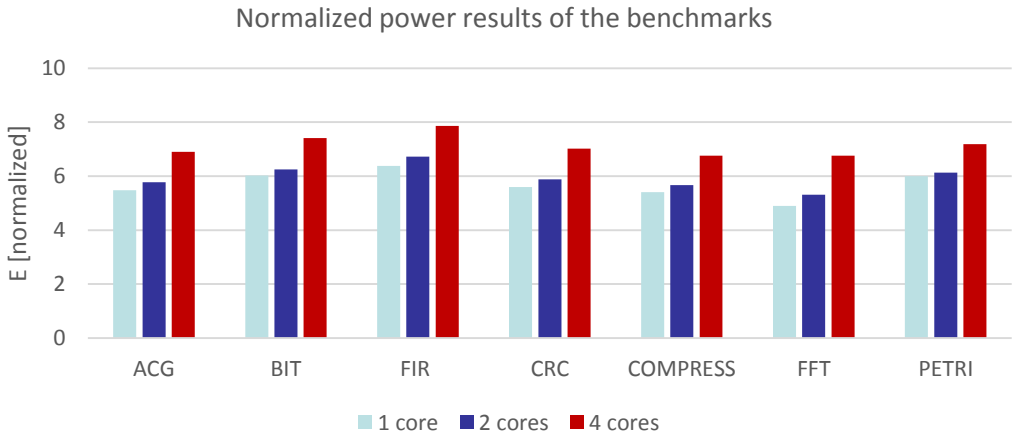

chist-era

FPGA prototype with PTEA

- Correlation results between power estimations of LEM and physical measurements: ~95%

Energy-aware OS prototype

- GEMSCLAIM OS prototype
  - Heterogeneous Multi-processing support
  - SystemC simulation
  - FPGA prototype
- Portable Execution Supervisor
  - Execution controlling daemon on Linux systems
  - ARM big.LITTLE (ODROID-XU3)
    - Exynos 5 Octa
  - x86-64 (Intel servers)
- Energy-awareness
  - Account energy per-thread and on different HW components
  - Control execution for reducing energy consumption

Resource Allocation and Power Management Interface (RAPMI)

- Ported on several platforms
- LEM (SystemC and FPGA)
  - CPU, Memory
  - Sampling period <1us
- ODROID board sensors
  - CPU, Memory, GPU
  - Sampling period ~262us
- Intel RAPL
  - Hardware counters
  - Sampling period ~1ms

Energy-aware scheduling

Energy slices (first-class resource)

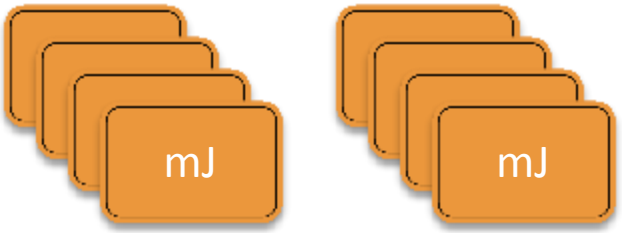

Processor time slices (within energy budget)

Application A

Application B

chist-era

Multi-objective energy-performance optimization

- Target combined energy-performance goals
  - Include SLA information from OpenMP+ extensions
  - Multi-program execution
- Monitor execution of the workload
  - Programs enter and exit
  - Periodically probe programs through the GEMSCLAIM runtime
    - *Scalability* (multi-threading)
      - Guide DCT
    - *Slowdown* (frequency throttling)
      - Decide DVFS

Prototype evaluation

- ODROID-XU3 platform
  - Power sensors
    - CPU and Memory
  - DVFS
  - ARM big.LITTLE heterogeneity
    - A15 and A7 cores
- Multi-program execution
  - Workload generator
    - MediaBench, BOTS, NAS, Rodinia
- Baseline
  - Unmanaged
- Optimizers
  - Energy only (E-only)
  - Performance (execution time) only (P-only)
  - Multi-objective Probing Guided Optimizer (MO)
    - EDP target

Results

- MO
  - vs. Unmanaged
  - ~3x less energy
  - ~2x faster execution
  - vs. P-only
  - ~30% less energy
  - ~90% of performance
  - vs. E-only
  - ~6x faster execution
  - ~30% more energy

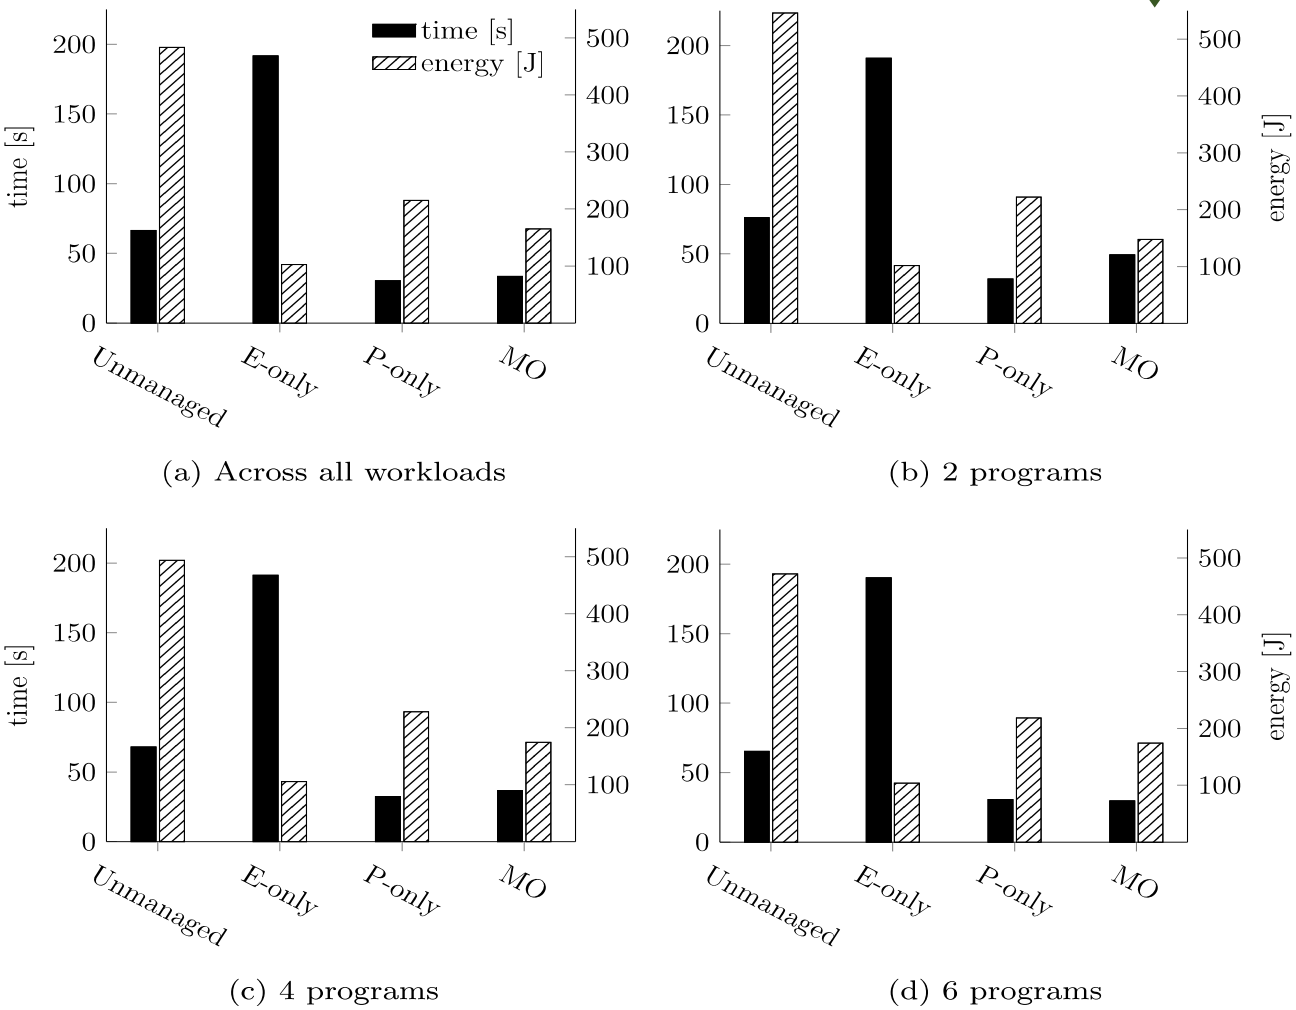

Compiler and runtime system

- The GEMSCLAIM Compiler is a source-to-source C compiler
  - supports C programs annotated with OpenMPE
  - OpenMPE is an energy-aware extension to OpenMP

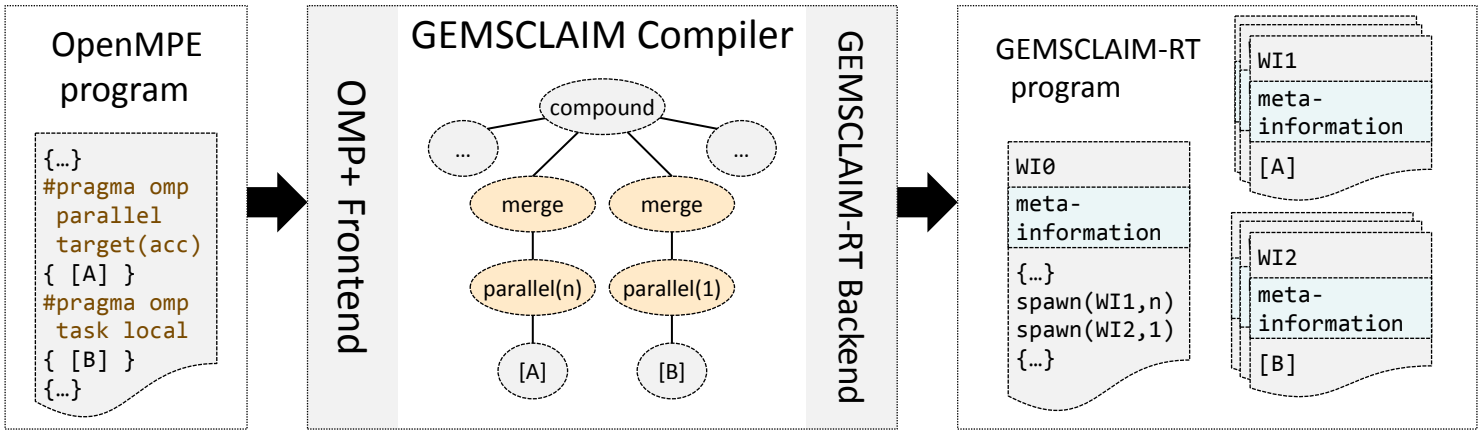

- Programs are divided into meaningful code regions
  - division either by manual annotation or automatically (e.g. based on parallel regions)
  - individually tunable and annotated with additional metadata
  - metadata is derived from analysis and/or user annotations

chist-era

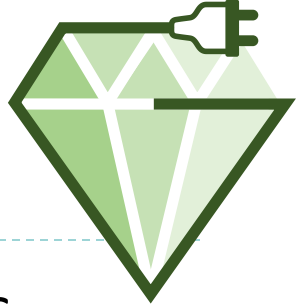

Compiler runtime system

- Compiler runtime available on all relevant platforms
  - GEMSCLAIM virtual platform & ARM boards (e.g. ODROID XU+E) & X86/64-Linux reference 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!
- Performs *downcalls* to OS/low-level layer
  - E.g. for application-level frequency scaling control

  → Can make use of information not available at the OS layer!

OpenMPE – OpenMP for energy

- Set of extensions published at IWOMP 2015 (*Application-level Energy Awareness for OpenMP*)
  - Region construct (handling code outside OpenMP regions)
  - Objective clause (defining multiple optimisation objectives and constraints)
  - Param clause (define tunable parameters for compiler optimisation)

- In addition to energy, power and time: *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 OpenMPE 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

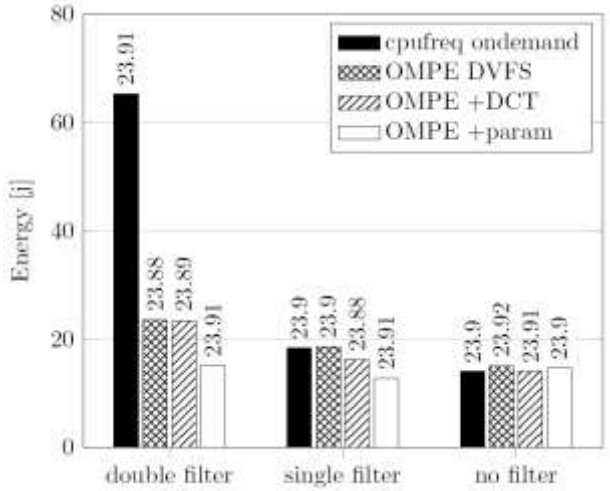

Mobile platform, 704x576 resolution

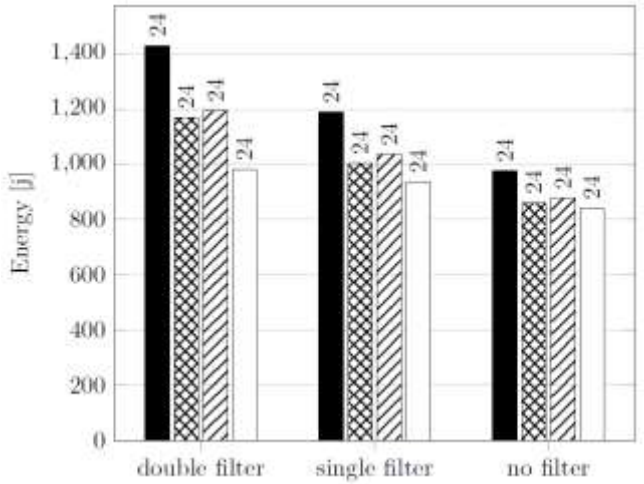

Desktop platform, 1408x1152 resolution

chist-era

GEMSCLAIM 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 Chalios, 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. Fedeac, A. Amaricai, M.Marcu, Energy Profiling of FPGA Designs, Proceedings of 2014 IEEE International Symposium on Robotic 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)
- 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

GEMSCLAIM Publications 2015 (1)

- Hardware support for performance measurements and energy estimation of OpenRISC processor, 10th IEEE International Symposium on Applied Computational Intelligence and Informatics (SACI), Timisoara, Romania, May 2015
- On the Quality of Implementation of the C++11 Thread Support Library. Peter Thoman, Philipp Gschwandtner, Thomas Fahringer. 23rd Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), 2015. Turku, Finland, March 4-6.
- Energy Prediction of OpenMP Applications using Random forest Modeling Approach. S. Benedict, R. R.S., P. Gschwandtner, R. Prodan, T. Fahringer, iWAPT 2015 – May 29, Hyderabad, India.
- Optimizing Task Parallelism with Library-Semantics-Aware Compilation. Peter Thoman, Stefan Moosbrugger, Thomas Fahringer. 21st International Conference on Parallel and Distributed Computing (Euro-Par), 2015. Vienna, Austria, August 24-28.
- Application-level Energy Awareness for OpenMP. Ferdinando Alessi, Peter Thoman, Giorgis Georgakoudis, Thomas Fahringer, and Dimitrios S. Nikolopoulos. Proceedings of the 11th International Workshop on OpenMP (IWOMP 2015). Vol. 9342 Springer International Publishing Switzerland, 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

GEMSCLAIM Publications 2015 (2)

- Evaluating Asymmetric Multicore Systems-on-Chip and the Cost of Fault Tolerance using Iso-Metrics. / Chalios, Charalambos; Nikolopoulos, Dimitrios S.; Catalan, Sandra; Quintana-Orti, Enrique S. In: IET Computers and Digital Techniques, 2015.
- ALEA: Fine-Grain Energy Profiling with Basic Block Sampling. Mukhanov, Lev; Nikolopoulos, Dimitrios S.; de Supinski, Bronis R. In: Proceedings of the 24th International Conference on Parallel Architectures and Compilation Techniques (PACT). 2015.
- Methods and Metrics for Fair Server Assessment under Real-Time Financial Workloads. Georgakoudis, Giorgis; Gillan, Charles J.; Sayed, Ahmed; Spence, Ivor; Faloon, Richard; Nikolopoulos, Dimitrios S. In Press, Concurrency and Computation: Practice and Experience, 2015.
- Iso-Quality of Service: Fairly Ranking Servers for Real-Time Data Analytics. / Georgakoudis, Giorgis; Gillan, Charles J.; Sayed, Ahmed; Spence, Ivor; Faloon, Richard; Nikolopoulos, Dimitrios S. In Press, Parallel Processing Letters, 2015.
- Direct FPGA-based Power Profiling for a RISC Processor. Cosmin Cernazanu, Marius Marcu, Alexandru Amaricai-Boncalo, Stefan Fedeac, Madalin Ghenea, Zheng Wang, Anupam Chattopadhyay, Jan Weinstock, Rainer Leupers. Proceedings of 2015 IEEE International Instrumentation and Measurement Technology Conference (I2MTC2015), May 2015, Pisa, Italy.
- Low-Cost Hardware Infrastructure for Runtime Shared Memory Thread Level Energy Accounting. Marius Marcu, Oana Boncalo, Madalin Ghenea, Alexandru Amaricai, Jan Weinstock, Rainer Leupers, Zheng Wang, Giorgis Georgakoudis, Dimitrios Nikolopoulos, Cosmin Cernazanu-Glavan, Lucian Bara and Marian Ionascu, Architecture of Computing Systems (ARCS), 04-07 April 2016, Nuremberg, Germany

Important GEMSCLAIM project objectives

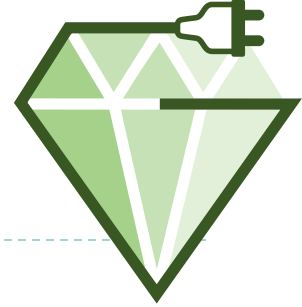

- Cross layer energy management and optimization for mobile devices: HW/simulator, OS, compiler
- Control trade-off between energy optimization and performance
- additional 30 % energy saving for mobile terminals

- GEMSCLAIM 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
- *GEMSCLAIM provides promising solutions for these challenges in a holistic approach, demonstrated on a physical HW/SW substrate*

Sustainability and Valorisation

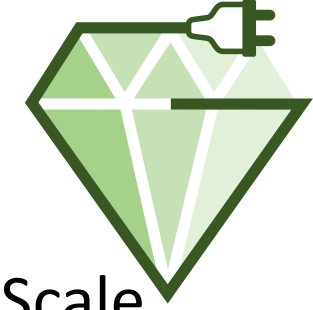

- Joint FETHPC proposal accepted: H2020 FETHPC AllScale
- Part of the SW under Apache 2.0 license (open source)
- Parts of the GEMSCLAIM OS Runtime are used in current research projects (FP7 NanoStreams) and are a foundation for future research proposals under the Horizon 2020 research programme
- Parallel 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