A Multi-Pronged Approach to Benchmark Characterization

Nikola Puzovic[1], Sally A. McKee[2], Revital Eres[3], Ayal Zaks[3], Paolo Gai[4], Stephan Wong[5], and Roberto Giorgi[1]

Department of Information Engineering,[1]
University of Siena,
Siena, Italy
{puzovic, giorgi}@dii.unisi.it

Department of Computer Science and Engineering,[2]
Chalmers University of Technology,
Göteborg, Sweden
mckee@chalmers.se

IBM Haifa Labs,[3]
Haifa University Campus,
Haifa, Israel
{eres, zaks}@il.ibm.com

Evidence srl,[4]
Pisa, Italy
pj@evicence.eu.com

Computer Engineering Laboratory,[5]
Delft University of Technology
Delft, The Netherlands
J.S.S.M.Wong@tudelft.nl

Intro

Behavior of workloads

- ✓ Easy to determine when characteristics of the target are known.
  - ✓ Execute on the simulator/HW and gather statistics.
- ✓ What to do when modeling tools do not exist?
  - ✓ Use multiple sources to collect as much information as possible.

Multi-Pronged Approach

- ✓ System is composed of two parts:
  - ✓ Fixed part (service core) + accelerators (to be designed)
- ✓ Collect statistics from different sources
  - ✓ Execution on existing HW, emulation, stats from the compiler (ILP).
  - ✓ Combination of these is driving the design of accelerators.

Case Study: ERA (Embedded Reconfigurable Architectures project)

Architecture

- ✓ A Control processor and a number of reconfigurable accelerators.
- ✓ Network on Chip (NoC) and SW/HW scheduler.
- ✓ Xilinx ML605 used to prototype and test the architecture.

Accelerators

- ✓ VLIW accelerators execute compute-intensive portions of benchmarks

Reconfiguration

- ✓ System can be reconfigured to exploit different requirements:
  - ✓ Save power when possible, provide computational resources when needed.

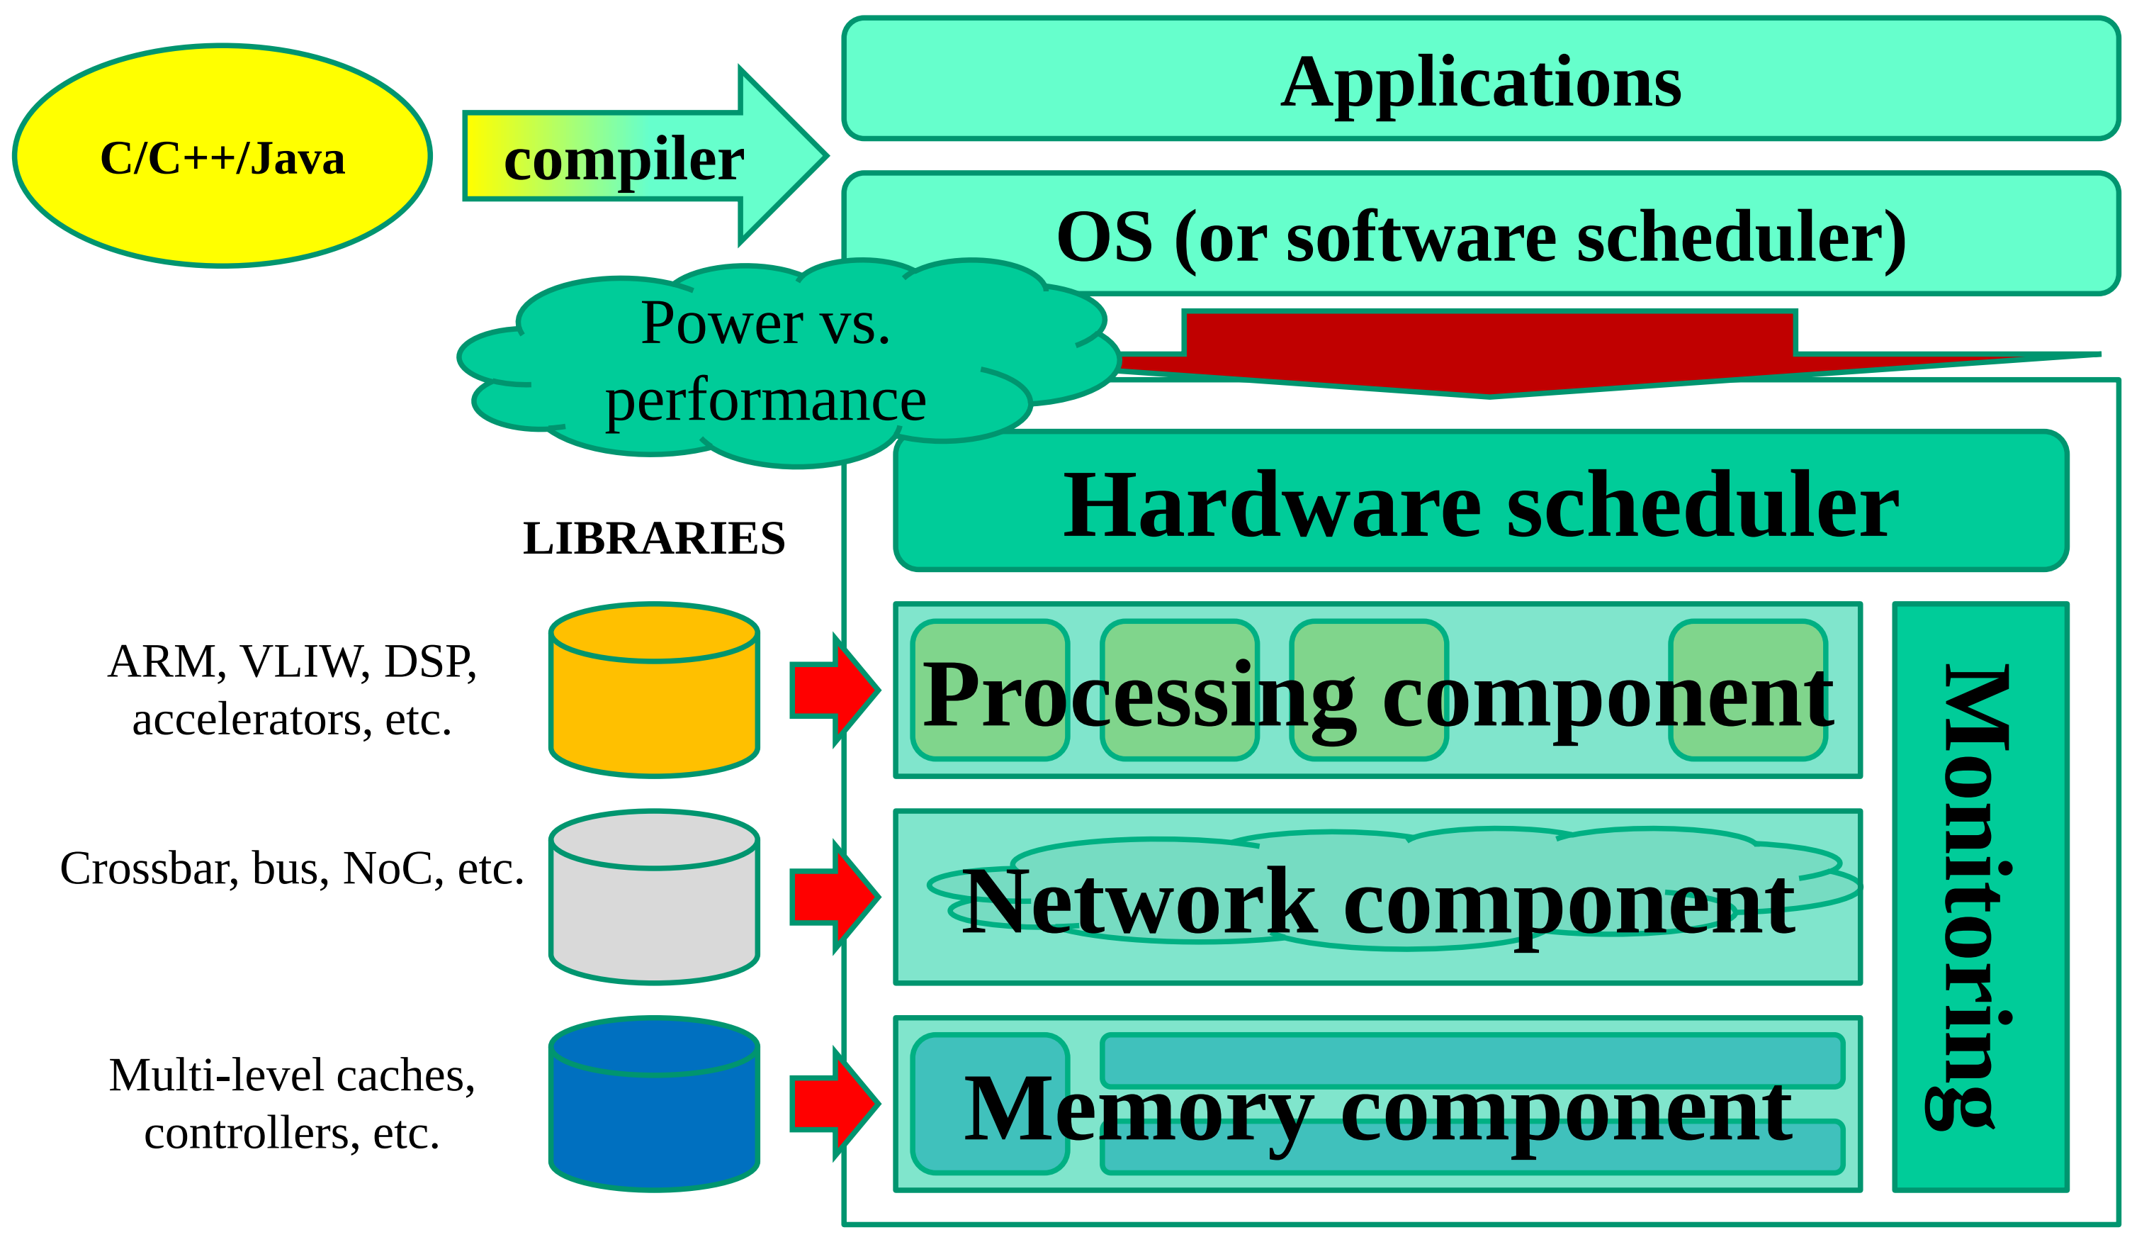

Applications

JPEG

- ✓ JPEG compression/decompression
  - ✓ Commonly used in embedded systems
- ✓ Implementation from MiBench

H.264

- ✓ MPEG4 Advanced Video Coding
- ✓ Improved video quality at lower bitrates (wrt other compression standards).
- ✓ Implementation from PARSEC

MPEG2

- ✓ MPEG2 audio/video compression
- ✓ Widely used in digital television broadcast (over the air or via cable).
- ✓ Custom parallelized implementation.

YUV2RGB

- ✓ Video conversion application: Reads a YUV video stream from the camera and converts it into RGB.
- ✓ Commonly used in mobile phones.

Tesseract

- ✓ Open source OCR engine
  - ✓ Ranked in top 3 OCR engines (1995)
- ✓ Adapted to omit GUI and to measure execution time

AC3

- ✓ JPEG compression/decompression
  - ✓ Commonly used in embedded systems
- ✓ Implementation from MiBench

ECDS

- ✓ Elliptic Curve Digital Signature
  - ✓ Security level similar to RSA with keys 5 times shorter.
- ✓ Implementation from Basicrypt suite.

SUSAN

- ✓ Recognizes corners and edges in Magnetic Resonance Images of the brain.
- ✓ Implementation from MiBench parallelized with OpenMP directives

Multi-Pronged Approach

Statistics

- ✓ Whole-program statistics: to characterize benchmarks
  - ✓ e.g. memory intensive vs. computing intensive
- ✓ Interval-based statistics: to detect phases in benchmarks
  - ✓ See how behavior changes during the execution.

Metrics

- ✓ Microarchitecture independent metrics.
- ✓ Memory behavior: stats on reuse distance, access strides, working set size).
- ✓ Processor: stats on registers (dependency distance, degree of use), on instruction mix.

Sources

- ✓ QEMU for MicroBlaze
  - ✓ Uses Dynamic Binary Instrumentation to extract statistics.
  - ✓ Gathers both whole-program and interval-based statistics
- ✓ Traces (hardware)
  - ✓ Usage of real HW to strengthen the results (Lauterbach tools for Xilinx ML605)
  - ✓ Extracts traces that are later combined with statistics from QEMU.
- ✓ Compiler (ILP)
  - ✓ First, locate hot spots in the code by executing it on x86 or PowerPC
  - ✓ Then, use gcc to estimate the ILP available.

Conclusions

Status

- ✓ Implemented gathering of statistics in QEMU and gcc.
  - ✓ Fine-tuning still needs to be done for both.
- ✓ In progress: combining stats from QEMU and real HW.

Use in general purpose systems

- ✓ Method is not bound to embedded systems only.
  - ✓ Can be ported to general purpose platforms.

Acknowledgements

This work is supported by the European Commission FP7 collaborative project ERA #249059.

ERA

COOPERATION

European Commission
Information Society and Media