HW/SW Co-designed Processors:
Challenges, Design Choices and a Simulation Infrastructure for Evaluation

José Cano[1], Rakesh Kumar[2], Aleksandar Brankovic[3], Demos Pavlou[4], Kyriakos Stavrou[4], Enric Gibert[5], Alejandro Martínez[6], Antonio González[7]

[1]University of Edinburgh, UK [2]Uppsala University, Sweden [3]Intel [4]11pets [5]Pharmacelera [6]ARM [7]Universitat Politècnica de Catalunya, Spain

The problem

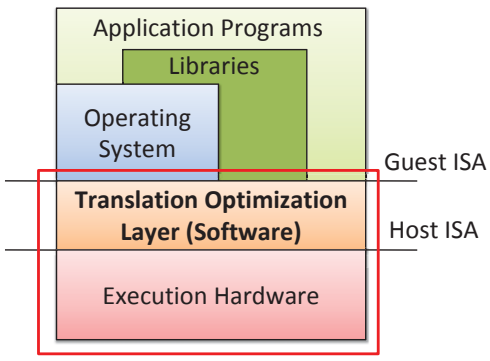

HW/SW co-designed processor

- **HW/SW co-designed processors**
  - –Lot of potential to improve
    - **Performance** (DBTO in SW)
    - **Energy efficiency** (simple HW)
  - –Some projects from **industry**
    - **IBM** DAISY/BOA, **Transmeta** Crusoe/Efficeon, **NVIDIA** Denver
    - But no successful product yet
  - –No major project from **academia**

- These processors need to address some **key challenges** before they can become mainstream
- There are **no simulation infrastructures** for evaluating different **design choices** and **trade-offs** to meet these challenges

Challenges and design choices

- Where to implement (HW or SW) **microarchitectural features**
  - Instruction decoding/reordering, register renaming, memory disambiguation, …
- How to reduce "**startup delay**"
  - One of the major problems of *Transmeta* processors
- When and where to **translate/optimize** the guest binaries
  - As soon as code becomes "hot"?
- How to address **speculative** execution (memory, control)
  - Checkpointing granularity?
- When and how to **profile** the execution
  - Overhead vs opportunity for improvement

DARCO: Simulation infrastructure

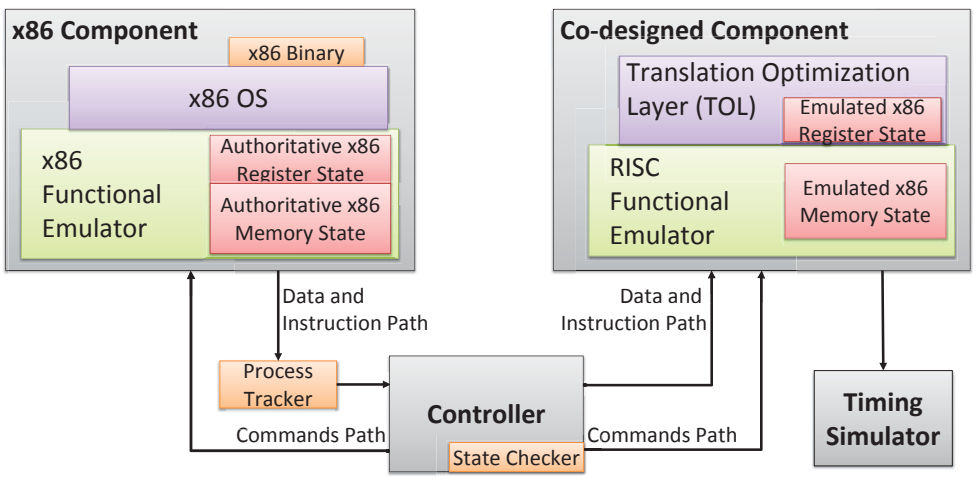

TOL execution modes

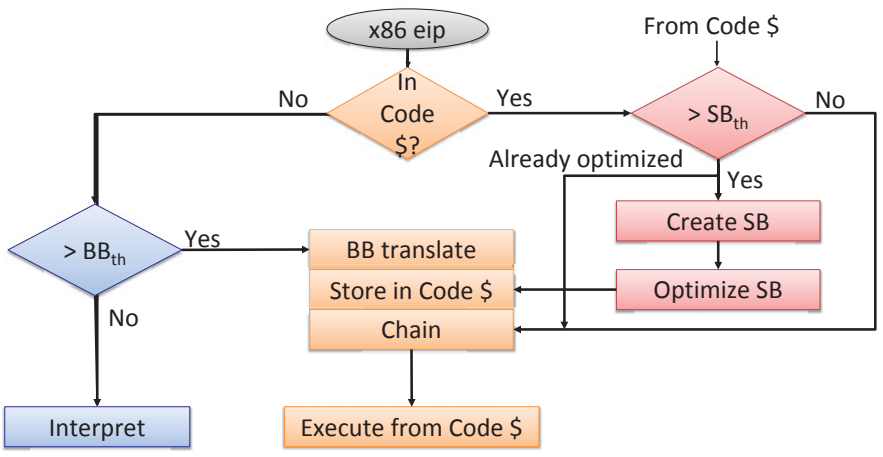

Challenges in building an infrastructure vs DARCO

- **Correctness**
  - It should not change program behavior
  - Arch./memory states compared periodically
- **Minimum TOL overhead**
  - TOL execution time must be small
  - 3-stages translation/optimization, chaining
- **Minimum emulation cost**
  - Host to guest instruction ratio must be low
  - Aggressive/speculative optimizations
- **Support for multiple guest ISAs (front-ends)**
  - User → app → ISA → device
  - Incorporating additional front-ends is simple
- **Plug and play support**
  - Easy to include/evaluate new features
  - Modular design
- **Debugging**
  - Strong debug toolchain
  - Debug mechanism activated if mismatch detected

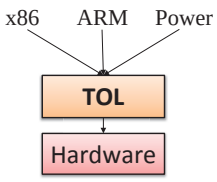

DARCO: Evaluation

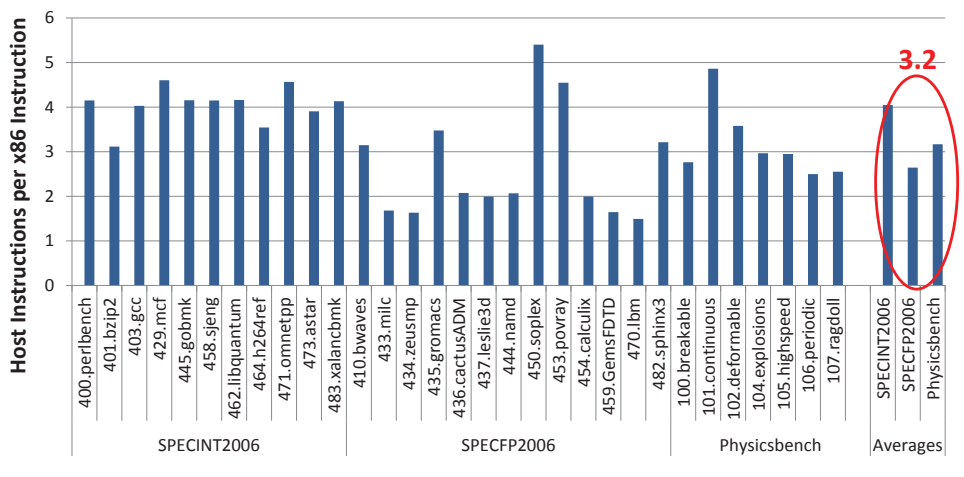

Conclusions

- **HW/SW co-designed processors**
  - Potential to improve energy efficiency and performance
  - Several industrial projects, no major project in academia
- **Challenges**
  - To become mainstream (e.g. startup delay)
  - To build a simulation infrastructure (e.g. software layer overhead)
- **DARCO**
  - May enable academic research in HW/SW co-designed domain
  - Modular infrastructure, easy to add new components/optimizations

THE UNIVERSITY of EDINBURGH informatics

ARM Research Summit, Cambridge, UK - September 11-13, 2017

icsa | Institute for Computing Systems Architecture