TAMPERE UNIVERSITY OF TECHNOLOGY

Open Source Toolset for HW/SW Co-Design of High Performance Low Power Co-Processors

Pekka Jääskeläinen
Customized Parallel Computing group
http://cpc.cs.tut.fi
Tampere University of Technology,
Finland

In OpenSuco 3 Workshop (ISC18)
Frankfurt, June 28 2018

Outline

- TTA-Based Co-Design Environment (TCE)
  - TTA: Transport Triggered Architecture
- Examples of recent TCE processor co-designs
- *If time allows: Screencast demo of the tools*

TTA-BASED CO-DESIGN ENVIRONMENT (TCE)

TCE Design Flow Overview

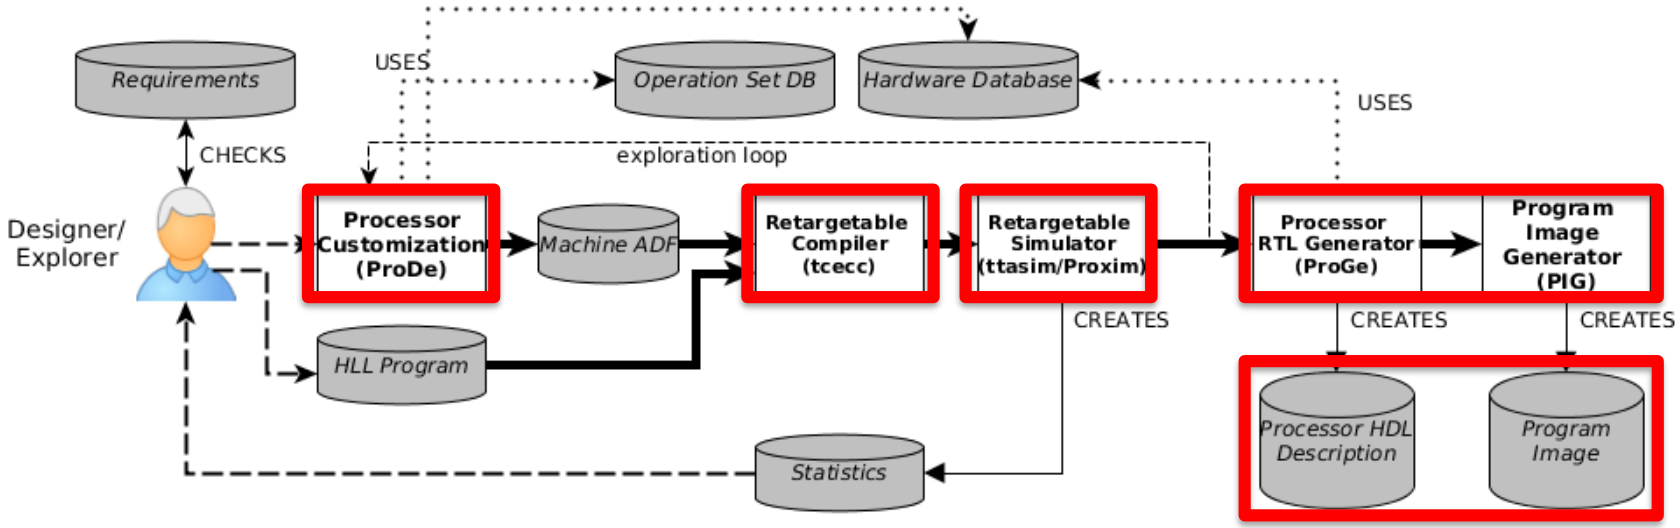

- Core tool is **tcecc: a retargetable Clang/LLVM-based compiler**
- MIT-licensed open source project available at http://github.com/cpc/tce
- Developed in TUT since 2003 in various research projects
- Mature toolkit that has been used for several custom processor design cases
  - Video compression, computer vision, software defined radio, audio, CNN inference...
  - Mostly academic case studies, but also known commercial designs

Transport Triggered Architectures

- Used as the base architecture template or ”paradigm” in TCE
- **Speciality: Software programmed data path transfers**
- Similar/related ideas have been proposed since 1970s
  - Any TTA-related patents surely expired by now
- Studied in TU Delft since late 1980s
  - For **more scalable instruction-level parallel designs:** improved VLIW
- TTA research continued since early 2000s in our group
  - Using TTAs for rapidly customized low power high performance programmable accelerators
  - FP16, wide SIMD, TLP (multicore, SIMT), energy efficient instruction streams…
  - Design case studies (DSP, SDR, DNN, graphics: http://tut.fi/vga …) from ultra low power cores up to wide-SIMD multicores

EXAMPLES OF RECENT TCE PROCESSOR CASE STUDIES

Example TCEMC design case by Tampere Univ. of Technology:

AivoTTA: Custom DSP for CNN-Based Object Recognition

- Co-designed for traffic sign and face detection networks
- Potential use cases in e.g. nanodrones, smart cameras
- Special integer vector MADD unit for convolutions
- Four register files with minimal ports:
  Only 1 rd/wr port each!
  - 8x1024b, 8x256b, 8x32b, 16x32b
- Loop buffer to reduce imem accesses + improve ifetch power
- C and OpenCL C supported
- 28 nm FDSOI (after place'n route):
  - 11 mW 16 GOPS @ 400 MHz
  - 116 mW 57 GOPS @ 1.4 GHz
- **World class power-performance for a compiler-programmed DNN inference accelerator**

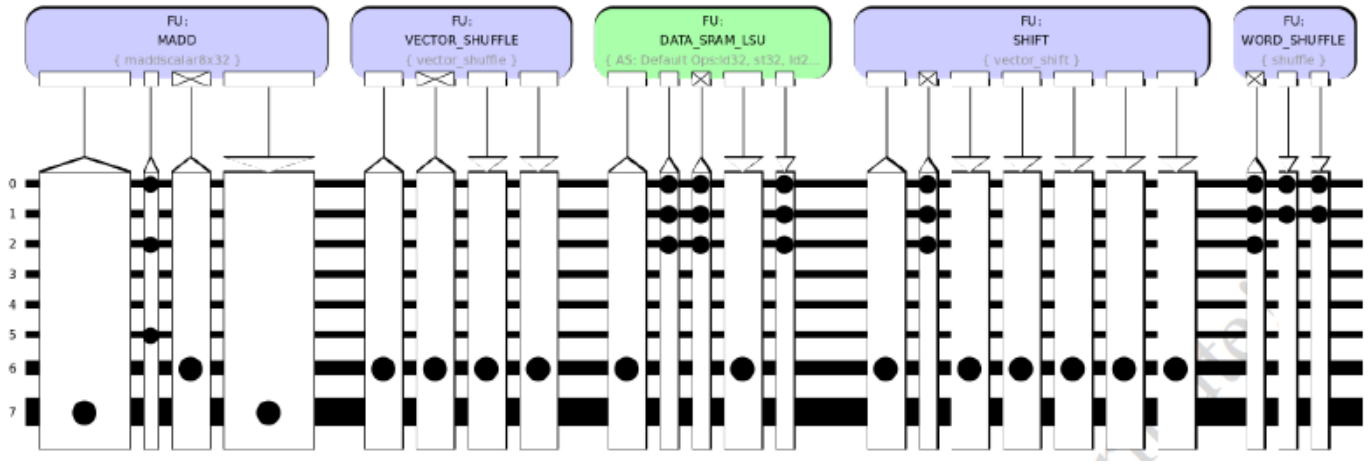

Example TCE design case by University of Turku:

A 5.3 pJ/op Approximate TTA VLIW Tailored for Machine Learning

- Minimum energy point 0.35 V near threshold operating voltage for **ultra low power** execution
- Features for approximate computing
  - Detect errors in computation, replace with safe values
- Manufactured on 28 nm FDSOI
- **About 320 µW (incl. memories) on ML workloads**
- Published in Elsevier Microelectronics Journal 61 (2017) 106–113

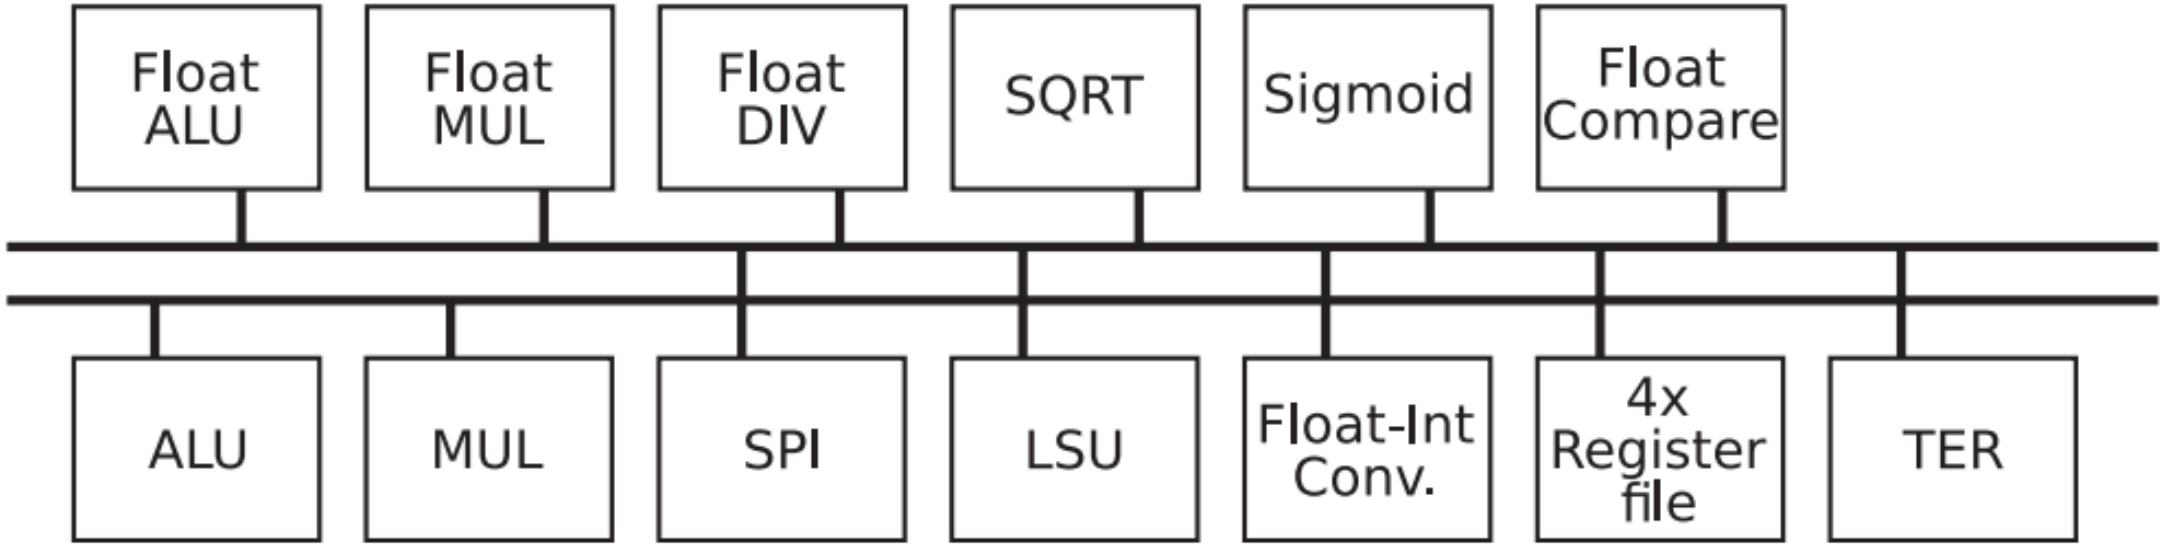

Example TCEMC design case by Leibniz Universität Hannover:

Customized High Performance Low Power Processor for Binaural Speaker Localization

- Custom DSP targeted to battery driven hearing aid devices with support for compute heavy algorithms
  - Very low power consumption, high computational performance, small form factor
- 32 x int32 SIMD (1024b) datapath
- Synthesized on 28 nm FDSOI
- 12 mW at 50 MHz, 1V
- **2-split SIMD RF, 1 write port each**
  - **Only 10.5% of total power thanks to software bypassing**
- Published in 2016 IEEE International Conference on Electronics, Circuits and Systems (ICECS)

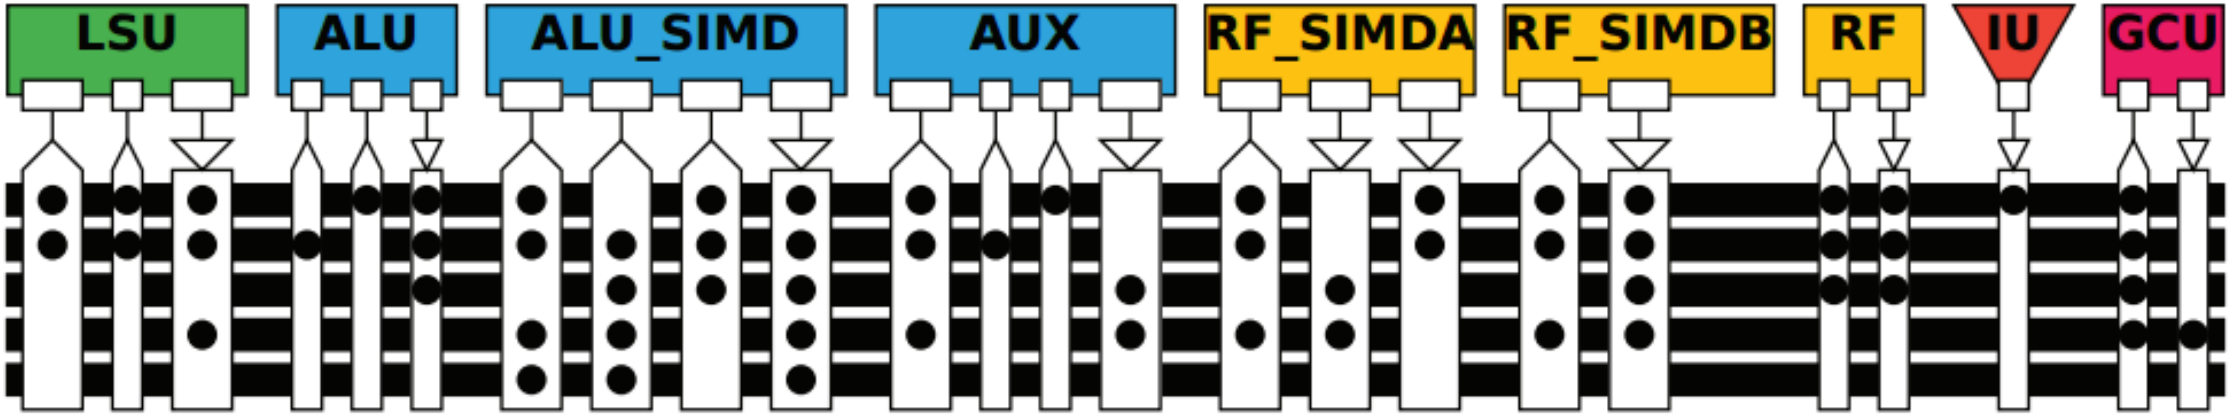

Thanks!
Questions?

Pekka Jääskeläinen
pekka.jaaskelainen@tut.fi

TCE source code (MIT licensed): http://github.com/cpc/tce

Customized Parallel Computing group: http://cpc.cs.tut.fi