

# Development of a modular and fully-digital PCIe-based interface to Real-Time Digital Simulator

Master Thesis Defense Presentation

Steffen Vogel

ACS | Automation of Complex Power Systems



# Agenda

- Objectives
- Architecture
  - VILLASfpga
  - VILLASnode

#### Linux

- Real-time Optimizations
- Driver API

# FPGA

- Datamovers / Synchronization
- RTDS Interface
- Example Applications
- Future work







High-level Architecture



# **Motivation**

#### Flexible framework

- Extensible
- Reuse of existing intefaces
- Integration into VILLASnode concept = Hard / Soft real-time
- No VHDL knowledge
   C / C++, Simulink models



# **Objectives**

- Framework for power system co-simulation on FPGAs
  - Extensible
  - Hard Real Time

Leverage flexibility of Linux and performance of specialized DRTS

- Example: Interface to RTDS
  - Synchronized data exchange with sub 10 uS time-steps
- Get familiar with latest FPGA design tools & methodologies
  - ≡ High Level Synthesis (C / C++)
  - System Generator for DSP (Simulink + Matlab)
  - Electronic System Level Design (ESL)

Optimize Linux for hard real-time simulation



# Applications

- Real-time co-simulation
  - One subsystem running on Linux
- Hybrid simulation: EMT + dynamic phasors
  - DFT implementation on FPGA
  - One subsystem is running in phasor domain on Linux
- Interface algorithms for geographically distributed simulation
  - Prediction and compensation of communication delay
  - DFT / wavelet transforms
- Custom communication protocols for RTDS: EtherCat...
- User interaction & monitoring
  - WebSockets, FIWARE
- Data logging & replay

6

■ Big Data (we could capture every timestep)



# Architecture: VILLASnode





# Architecture: VILLASnode + VILLASfpga



Full picture of possible VILLASnode topology.



# Hardware: Xilinx VC707 Evaluation Board





LC Multi-mode Fiber

SFP Module





# Architecture: VILLASfpga





R\\\

10

### **Architecture: IP Integrator**



Vivado IP Integrator Block Diagram



# **Architecture: IP Integrator (Model)**



Vivado IP Integrator Block Diagram of Model Hierarchy



A PCIe-based Co-simulation Interface | Steffen Vogel | Thesis Defense | 29.04.16

12

# **Architecture: IP Integrator (PCIe)**





E.ON Energy Research Center

13

# Interconnect

- Virtex 5 (ML507)
  - Memory Mapped: Processor Local Bus (PLB)
  - Streaming: Local Link (LL)
  - Influenced by IBM PowerPC architecture

- Virtex 6/7 (ML605 / VC707)
  - Memory Mapped: AXI4 & AXI4-Lite
  - Streaming: AXI4-Stream
  - Advanced Microcontroller Bus Architecture (AMBA) introduces by ARM Inc.



Interconnect Standards from ARM



# AXI4-Stream

- AXI4-Stream enables **pipelining** 
  - FIFOs instead of Addressable Mem
- Cut-through instead of Store-and-forward
  - Reduced latency
- Examples
  - ≡ Filters
  - Recursive / Short-term DFT





#### Implicit synchronization





# Xilinx ML507 Evaluation Board

- Integrated PPC440 CPU
- 5th generation Virtex FPGA
- Supported by ISE IDE only

Two boards are used at ACS



© Xilinx Inc.



# Xilinx ML605 Evaluation Board

- 6th generation Virtex FPGA
- Supported by ISE IDE only

We have one in OPAL-RT OP5600



© Xilinx Inc.



# Xilinx VC707 Evaluation Board

- 7th generation Virtex FPGA
- Supported by Vivado IDE
- We bought one



© Xilinx Inc.



# **Existing RTDS Interfaces: GTNET cards**

- Supports a variety of protocols: IEC61850 Goose & SV, IEC60870-5-104, DNP, UDP, TCP
- GTNET-SV: IEC 61850-9-2 Sampled Values
   Up to 256 samples / cycle (12.8 kHz for f=50 Hz)
- GTNETv2-SKT
  - Up to 300 values at 1 kHz
  - Less then 60 values at 5 kHz
  - ≡ GTNETx1 limited to 100 Hz



#### © RTDS Technologies



# **Existing RTDS Interfaces: GTFPGA / MMC support unit**

- First digital RTDS interface was demonstrated by CAPS / ABB in 2008
  - 16-bit LVDS digital interface between to GTFPGA cards
- Closed source IP (compiled NGC netlist)
- Up to 64 values per direction and time step (FP or Integer)



© CAPS, Michael Sloderbeck



# Linux, Drivers & Optimization

ACS I Automation of Complex Power Systems





# **Linux Real-time Optimizations**

#### Linux

- Tickless kernel!
- Pinned threads and IRQs (isolcpus)
- Preemptible Kernel: PREEMPT\_RT patchset

### System

- Avoid System Management Interrupts
- Reduce PCIe bus contention and auxiliary IRQs
- Lock-less datastructures: queues, lists, stacks





# SW Driver

- Linux provides APIs to implement drivers in userspace
  - Userspace IO (UIO)
  - Virtual Function IO (VFIO)
  - ≡ /sys/bus/pci/[bdf]/resources
  - = /dev/mem + /proc/[pid]/pagemaps
- Ideally we would implement the driver as a custom kernel module
  - Disadvantages:
    - = Context-switches become critical for small timestep simulations
    - = Maintenance of out-of-tree module
    - = Compilation of kernel / module





# Comparison

# Userspace IO (UIO)

- No bus-mastering
  - = No DMA
  - = No MSI interrupts
- /dev/mem + sysfs
  - Bus-mastering possible with dirty hacks
  - No device isolation by IOMMU
  - DMA continuous allocation only with hugepages
  - No MSI support
  - Requires superuser privileges
- Virtual Function IO (VFIO)
  - "Successor" of UIO
  - Mostly used by KVM and DPDK
  - Safe bus-mastering by isolating devices in IOMMU containers.
  - No superuser privileges required



# **Driver Comparison**

	UIO Userspace IO	<b>VFIO</b> Virtual Function IO	<b>Sysfs</b> /sys/bus/pci + /proc/pagemap	<b>LKM</b> Loadable Kernel Module
Bus-mastering		<b>v</b>	ugly	<b>V</b>
Interrupts	Legacy only	Legacy + MSI		<b>v</b>
DMA		<ul> <li>✓</li> </ul>	ugly	<ul> <li>✓</li> </ul>
IOMMU Isolation		<ul> <li>✓</li> </ul>		( )
Req. permissions	Superuser	flexible	Superuser	flexible
Users	DPDK, Closed Source Drivers	KVM, DPDK	Nobody	Upstream Drivers
Userspace	<ul> <li>✓</li> </ul>	v	<ul> <li>✓</li> </ul>	( )
IPC	mmap()	mmap() ioctl()	read() write() mmap()	flexible



# Datamovers: CPU <-> FPGA

- Transfer data between FPGA and CPU
- Convert from memory-mapped to packet streaming bus
- Usually two channels:
  - Device to Host: Stream to Memory-mapped (S2MM)
  - Host to Device: Memory-mapped to Stream (MM2S)

- Direct Memory Access (DMA)
  - Data transfer is delegated to DMA controller on FPGA
  - CPU is not occupied
  - PCI Bus Mastering: Device writes to host memory
- Programmed / Memory Mapped IO (MMIO / PIO)
  - Direct register access
  - ≡ FIFO



#### **PCIe Interface: Datamovers**





# **Address Spaces & Methods**



Address spaces and Access methods



#### **PCIe Interface: Datamovers**





# **Synchronization Methods**

# Interrupts

- Legacy (INTX)
  - = Out-of-band: Data-races possible
  - = Only 4 physical interrupt lines
- Message signalled (MSI)
  - = In-band: ordering with data preserved
  - = Lower latency
  - = Supports up to 32 vectors
- Polling
  - Host RAM
    - = Directly poll on DMA buffer descriptors
  - PCI mapped memory (BARs)
    - = Poll Interrupts status registers
  - Network Interface Quibbles (NIQ)
    - = long polling
    - = delayed PCI read-completion TLPs (up to 50ms)



# **PCIe Interface: Latency**

### Response time to interrupt generated on FPGA

- ≡ Polling: ~ 1 uS
- ≡ Interrupts: > 2 uS





# **PCle Interface: Jitter**

- Periodic timer IRQ every 50 uS ~ 170k clks
- Similiar results for polling on real-time and non real-time kernel
- Big differences for IRQs
  - Differences in PREEMPT\_RT interrupt handling



Comparison of interrupt jitter (Linux 4.5.7 vs 4.0.4-rt x86\_64 SMP)



# **Synchronization**

- Events require synchronization
  - Simulation started / stopped
  - New timestep started
  - Deadline missed / overrun
  - Data transfer completed



- "Hard" synchronized
  - Reference to global time
  - Attached HIL devices
  - Multiple HIL devices

"Soft" synchronized
 Only logical order is preserved



# Synchronization: "Hierarchy"



#### **FPGA** implementation

- 8 clocks
- 4 clock domains



**Clock Domain Synchronization** 



# Timing





👉 Update

36



# Timing 2

- Different communication schemes are possible
- Limitations imposed by capabilities of RTDS\_InterfaceModule
  - Sends data only at beginning of TS
  - Sends data only once per TS



A PCIe-based Co-simulation Interface Steffen Vogel | Thesis Defense | 29.04.16

Computation

Update

# PCIe / RTDS Interface: Overruns

- Biggest matrix dimension which can be solved by LU decomposition?
  - Net-lib's LAPACK
- Detecting overruns by abnormal RTT
- Parallel communication pattern, time step: 50 us





# **Round-trip Time: RTDS-Linux-RTDS**







A PCIe-based Co-simulation Interface | Steffen Vogel | Thesis Defense | 29.04.16

**RTDS-Linux-RTDS** 





40



# **Example Applications**

ACS I Automation of Complex Power Systems





Original network

# To be decoupled into two subsystems (SS1 + SS2)





Decoupled network by Ideal Transformer Model (ITM)





- SS2 is replaced by CBuilder component
- CBuilder component includes solver
- Added ∆dt delay for control signals Vss2A & Iss1A



- With GTFPGA interface
- CBuilder code is cross-compiled on Linux
  - Some C and header files are used with some tricks



# **Results: Simple Circuit**

- $f_0 = 1$  kHz, Δt = 25 μs
- 1 time step latency
- 1 additional time step due to coupling of control and power system solution (T2 vs T0 zone)





# Example: Sliding DFT

- Simple Recursive implementation
- Window over fundamental period:  $z^{-N}$
- 3 Additions, 2 Multiplications per frequency and sample

$$S_k(n) = e^{j2\pi k/N} [S_k(n-1) + x(n) - x(n-N)]$$



Performance: 68 clk cycles + 2 clk cycles per harmonic (pipelined)

≡ IEEE 754 Single Precission FP

E. Jacobsen and R. Lyons, "The sliding DFT," IEEE Signal Processing Magazine, vol. 20, no. 2, pp. 74-80, Mar. 2003.

47



# **High-level Synthesis**

#### IP Integrator





```
void hls dft(hls::stream<axis> &input, hls::stream<axis> &output,
            float fharmonics[MAX HARMONICS], ap int<8> num harmonics, ap int<8> decimation) {
        #pragma HLS INTERFACE s axilite port=return,fharmonics,num harmonics,decimation bundle=ctrl
        #pragma HLS INTERFACE axis port=input,output
        #pragma HLS STREAM depth=64 variable=input,output
        #pragma HLS DATAFLOW
       /** Previous coefficients for incremental update */
        static std::complex<float> coeffs[MAX HARMONICS];
       /* Time */
       static float t;
       static ap int<32> decimation cnt;
       /** Sliding window of samples */
       static ap shift reg<float,NSAMPLES> windows[MAX VALUES];
       /** AXI Stream signals */
       axis real, imag, refph;
       for (int index = 0; index < MAX VALUES; index++) {</pre>
               /* Read real-valued time-domain data from AXI Stream interface */
               axis in = input.read();
               /* Shift and get data from SLR */
               float newest = in.data;
               float oldest = windows[index].shift(newest, NSAMPLES-1);
               for (int i = 0; i < num harmonics; i++) {</pre>
                       #pragma HLS PIPELINE II=2
```

49



# **Results: Dynamic Phasor – EMT transformation**

- Alignment of reference phase can be problematic
  - Simulations must be started simulateously



![](_page_49_Picture_4.jpeg)

# **Results: Dynamic Phasor – EMT transformation**

#### DFT32 block on RTDS is subject to distortions in case of fast transients

![](_page_50_Figure_2.jpeg)

**E.ON Energy Research Center** 

# Conclusion

Real-time co-simulation interfaces with latencies below 1 time-step are feasible

- ≡ 2 time step round-trip (parallel, 1 time step per direction)
- 1 time step round-trip (serial, RTDS->Linux->RTDS)
- Time-step periods **over 25 uS** are pretty stable
- Results align with first tests made between OPAL-RT and RTDS
- Presented architecture is flexible and extensible
  - Comes at the cost of increased complexity and step learning curve
- Interaction of control system and power system solution of RTDS may add an additional time-step latency
- Hard real-time on Linux is possible but fragile
  - Already small changes, different kernel versions or improper coding style will cause issues

![](_page_51_Picture_11.jpeg)

# **Future work: Applications**

Use the framework: It's there.

- Additional interfaces
  - Industrial ethernet: EtherCAT, PowerLink, ...
    - Connection to PGS, CWD, EBS labs
  - PHIL amplifiers:
    - = OPAL RT-Link ORION, Aurora
- FPGA based models: C/C++, Simulink
- Real-time co-simulation of power and communication systems
  - Linux runs NS3 as a communication emulator
- Multi-physics co-simulation
  - RTDS: electrical
  - Linux: thermal
  - Multi-rate interface is required
- Flexible control algorithms running on Linux

![](_page_52_Picture_15.jpeg)

### Future work: Framework

- Add FMC extension module for more fiber optic connections
  - FPGA-to-FPGA
  - FPGA-to-RTDS
- Sync RTDS as Slave
- Small time-step mode (< 10 uS)</p>
- Test other targets / OS's
  - RTOS: VxWorks
  - AMP: HermitCore? Baremetal?
  - Architectures: NUMA, ARM (Zynq)
- Simplify usage by adding more automation (ESL)

![](_page_53_Picture_11.jpeg)

FMC Quad-SFP Extension Module

![](_page_53_Picture_13.jpeg)

E.ON Energy Research Center

#### **Acknowledgments**

Thanks to my supervisors Marija Stevic & Stefan Lankes for their advise and support.

Thanks to RTDS Technologies for their support and access to the GTFPGA interface.

Thanks to OPAL-RT Technologies for testing and support of the RTDS-OPAL co-simulation interface.

Thanks to Xilinx Inc. for sponsoring software licenses within the XUP and their excellent tools which made this work possible.

![](_page_54_Picture_5.jpeg)

**E.ON Energy Research Center** 

![](_page_54_Picture_7.jpeg)

![](_page_54_Picture_8.jpeg)

![](_page_54_Picture_9.jpeg)

![](_page_54_Picture_10.jpeg)

![](_page_55_Picture_0.jpeg)

#### Contact

E.ON Energy Research Center Mathieustraße 10 52074 Aachen Germany Steffen Vogel stvogel@eonerc.rwth-aachen.de http://www.eonerc.rwth-aachen.de

http://www.steffenvogel.de

ACS I Automation of Complex Power Systems

![](_page_55_Picture_6.jpeg)

![](_page_55_Picture_7.jpeg)