institut für astrophysik
UNIVERSITÄTSSTERNWARTE WIEN

universität wien

A Lightweight Operating System for the SSDP

Armin Luntzer

Franz Kerschbaum, Roland Ottensamer, Christian Reimers

Department of Astrophysics, University of Vienna

Background

institut für astrophysik
UNIVERSITÄTSSTERNWARTE WIEN
universität wien

NGAPP (Next Generation Astronomy Processing Platform)

- RUAG Space Austria (RSA), ESA contract number: 40000107815/13/NL/EL/fk
- Collaboration with University of Vienna, Department of Astrophysics (UVIE)
- Evaluation of MPPB (Massively Parallel Processor Breadboarding) prototype
  - **Software benchmarking of MPPB (UVIE)**
  - **Analysis of MPPB features w.r.t. DPU design and applications (RSA)**

Background

institut für astrophysik
UNIVERSITÄTSSTERNWARTE WIEN
universität wien

NGAPP (Next Generation Astronomy Processing Platform)

– Outcome

- **Inputs to SSDP (Scalable Sensor Data Processor) ASIC specification**
- **Analysis of MPPB features w.r.t. DPU design and applications (RSA)**
- **Software performance library (UVIE)**
- **Custom operating system prototype (UVIE)**

Background

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

MPPB Recap

- designed by RECORE Systems
- 1 LEON2 processor (Controller)
- 2 Xentium VLIW DSP cores
- Network-on-Chip (NoC)
- High-speed interfaces (SpW, ADC/DAC)
- 50 MHz system clock
- current version 2.0 with SSDP-like feature set and characteristics

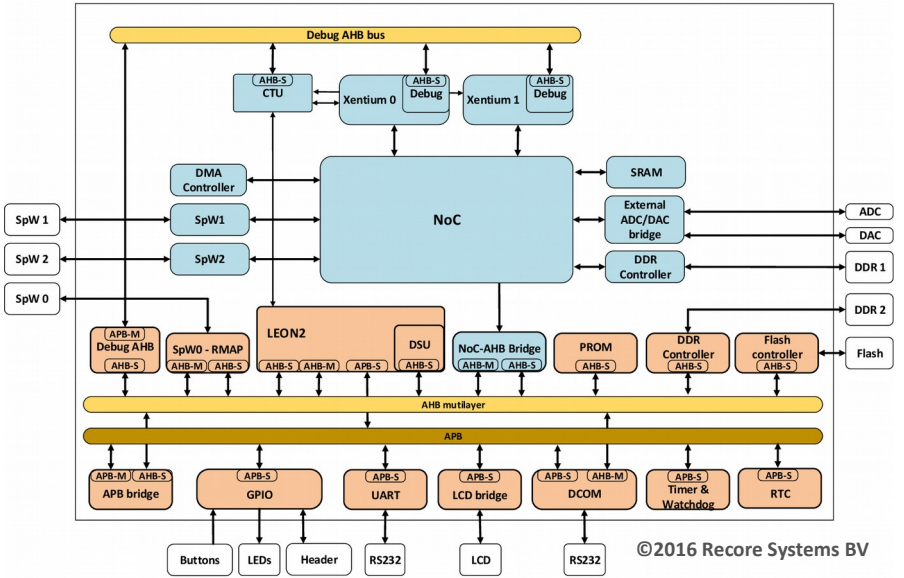

Background

MPPB v2 Characteristics

- DMA
  - **packet-like data transfers between NoC devices**
  - **2D stride support**
  - **parallel channels**
- Xentium DSP subsystem
  - **10 parallel execution units**
  - **Local tightly-coupled memory (TCM): 4 x 8 kiB banks**
  - **I-cache: 16 kiB**

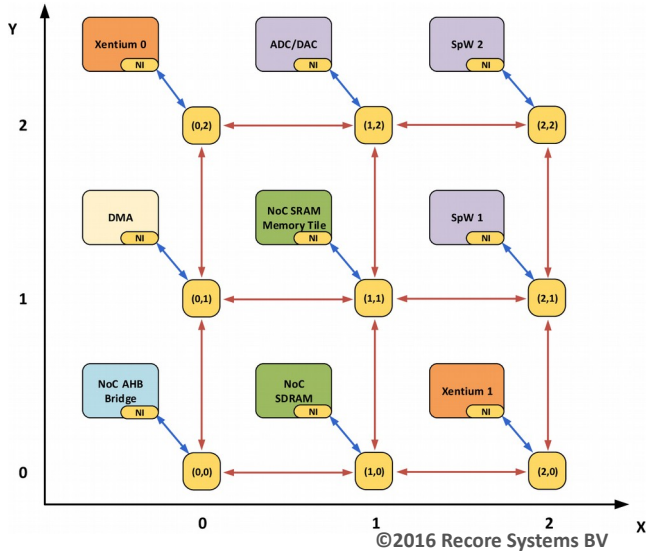

Background

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Xentium Programs: Design Decisions

- classic approach:
  - **single monolithic program running on DSP**
  - **data processing chain is implemented as a series of function calls/operations**
- downsides:
  - **code size of complex processing tasks can exceed size of i-cache quickly**
    - code must be re-fetched via the NoC, causing stalls
  - **even small changes require full program re-validation**
    - high maintenance effort!

Background

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Xentium Programs: Design Decisions

- alternative approach developed by UVIE during NGAPP:
  - **multiple tiny programs (kernels), one for each functional step**
  - **kernels can (and should) be dumb, i.e. perform only one task, without any knowledge of their environment**
- upsides:
  - **code can be arbitrarily small, thus will always fit the i-cache (break up into sub-kernels)**
  - **Processing chain can be created from independent building blocks**
  - **changes affect isolated components only**
  - **simpler units generally have less execution paths and are easier to trace for WCET**

Background

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Xentium Programs: Connecting the Dots

- data must be passed between kernels for processing
- using a single buffer isn't a good idea
  - **kernels can't be dumb**
  - **no flexible scheduling**

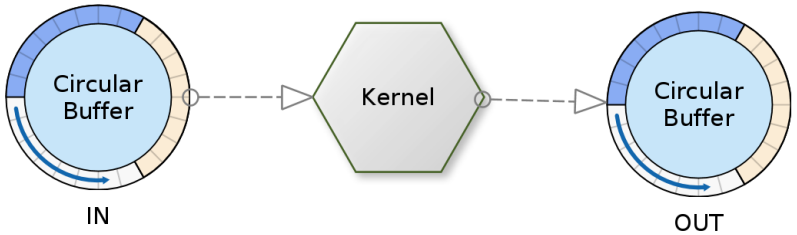

the solution: connect kernels via circular buffers

Background

Xentium Programs: Forging the Chain

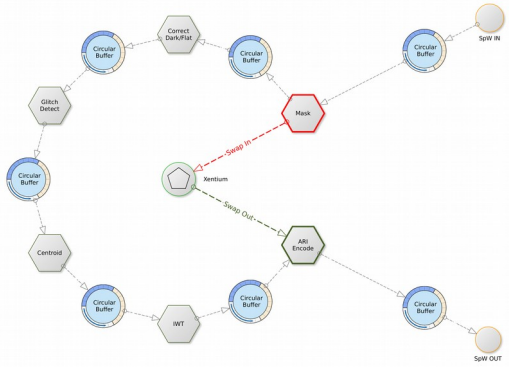

- buffers transport metadata references of data products
- easy creation of a pipeline: just assign buffers
- one kernel's output buffer is another one's input
- metadata defines processing requirements interpreted by kernels
- the circular buffers are inherently accessible from multiple DSPs
  - **"split" or "funnel" the data products into or from separate streams**
  - **easy data-parallel execution on multiple Xentiums for throughput-critical sections**

Background

Xentium Programs: Scheduling

- kernels must be scheduled by some metric
- circular buffer fill levels are most suitable:
  - **define a critical input buffer fill level per kernel**
  - **if buffer becomes critical, swap kernel executing in Xentium**
  - **also serves as DSP load metric**

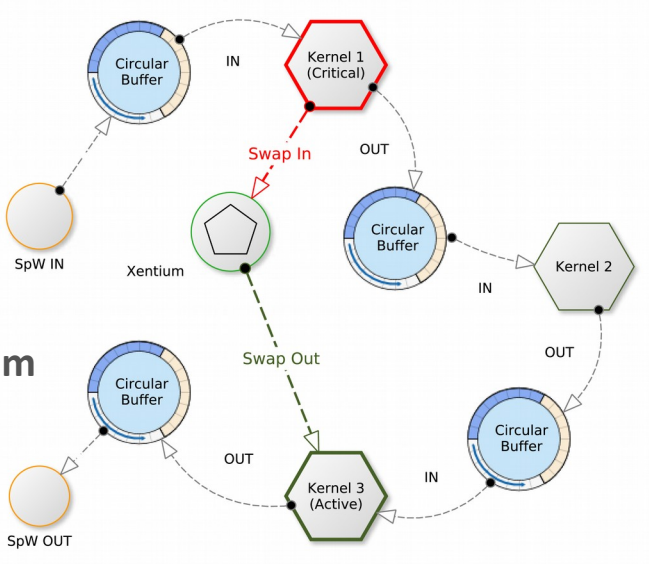

Background

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Xentium Programs: Buffer-Level-Based Scheduling

- complex interaction of NoC nodes and DMA is condensed into a single value
- self-balancing system for optimum resource usage
  - **can further use auto-tuning of buffer criticality levels for self-optimisation**
  - **easy iterative search metric: minimise kernel swaps per time unit**
- fast kernels can accumulate a lot of data at their input
  - **smaller number of kernel switches lead to greater efficiencies**
  - **more processing time is automatically assigned to computational-intensive kernels**

Major difference to a fixed scheme: data accumulates and forward-propagates in the pipeline driven by required computational time of a processing stage

Background

MPPB Experience Summary

- NoC DMA feature was very comfortable to use
- high interrupt rates: typically several 10k, > 100k peak
  - **complex transfers to Xentiums, high-throughput kernels**
  - **still very much manageable, just needs snappy ISR code**

IMPORTANT: MPPB v2 added notifications to Xentiums on DMA transfer completion, drastically reducing IRQ load on the LEON

- Xentiums with hand-coded ASM in chain concept are extremely powerful
  - **Xentium assembly coding is easy!**
  - **nice mnemonics and highly effective pipeline „loop“ feature for parallelisation**

Use Case: Space Astronomy

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Payload Instrument Data Processing

- raw science data >> TM budget
  - **typically 5-10x**
  - **100k – 1M samples per second**
  - **16 bit data types**
- SNR is low
- new science is buried in the noise
  - **noise must be preserved**

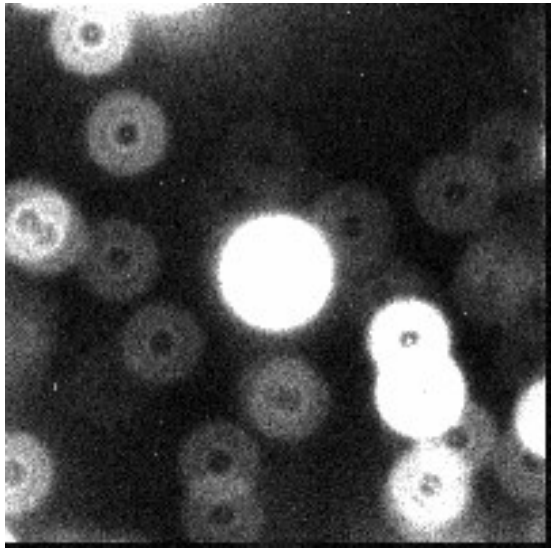

CHEOPS Simulated Observation of ultra-high precision photometry

Use Case: Space Astronomy

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Payload Instrument Data Processing

- general on-board tasks
  - **data deglitching**
  - **calibration**
  - **ancillary data processing, e.g centroiding for high-precision tracking**
- compression
  - **lossy preprocessing steps: temporal or spatial resampling, rounding, transforms**
  - **reversible decorrelation or prediction**
  - **entropy coding**

Use Case: Space Astronomy

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

On-Board Data Processing Chain

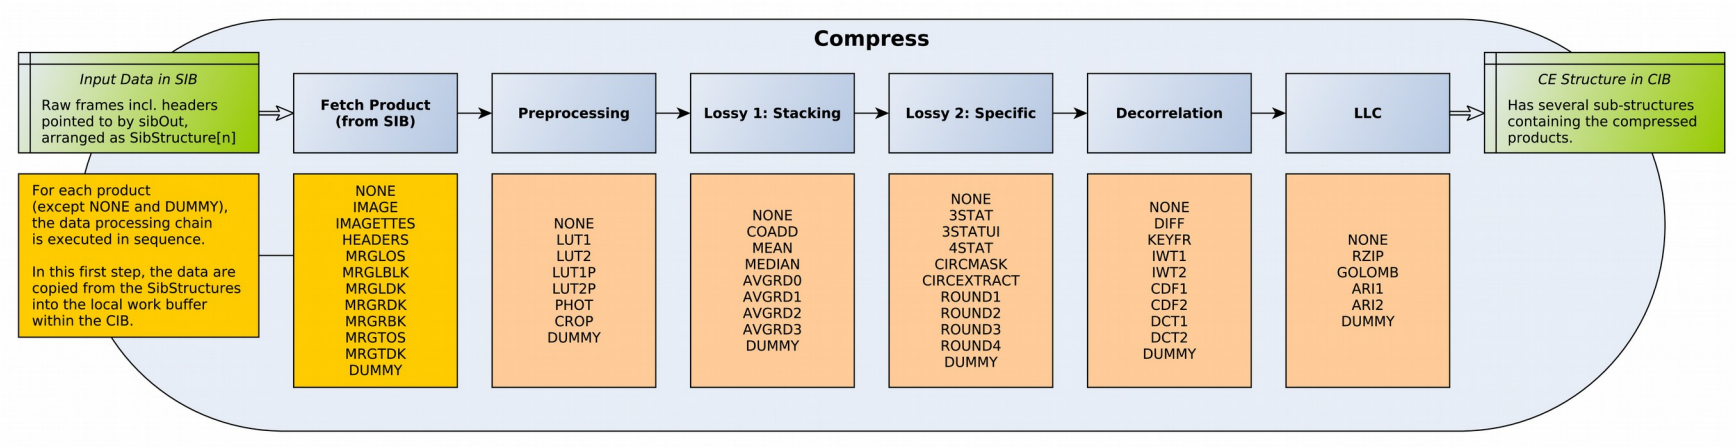

CHEOPS BEE compression stage

An OS for the SSDP

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Why a Custom Operating System?

- particular focus on driving the Xentiums
  - **ensure available computational resources are used efficiently**
- make sensible use of the platform's resources:
  - **Xentium: "kernel" concept, computational and support code should fit the i-caches**
  - **LEON: all performance-critical driver and system code should fit its cache as well**
  - **be scalable to accommodate more than two Xentium nodes**
  - **support synchronisation/operation of multiple SSDPs via high-speed interconnect**

An OS for the SSDP

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

LeanOS – a Lean Operating System

- nationally funded by the FFG under project number 847987, 4 FTEs
- official project start Q1/2015
- actual start when MPPB v2 arrived Q1/2016
- development up to qualifiable level TRL 6-7
- tailored to the Network-On-Chip/DSP concept of the SSDP
- delivery expected Q4/2017
- final product released under an Open Source license

An OS for the SSDP

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Key User Requirements

- it shall be lean and efficient
- it shall support a scalable number of Xentium DSP cores
- it shall support Fault Detection, Isolation and Recovery (FDIR)
- it shall be easy to use
- it shall be designed with applicable S/W standards to be space qualifyable
- it shall come with support documentation and demo applications

An OS for the SSDP

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Key Functional Requirements

- trap/interrupt handling
- SMP readiness
- threads
- “tickless” timing and real-time schedulers (FIFO, RR, EDF)
- run-time configuration interface
- NoC DMA and I/O interface drivers (SpW, ADC/DAC)
- Xentium kernel support infrastructure and scheduler

An OS for the SSDP

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Key Feature: Xentium Processing Chain

- built by linking circular buffers to kernels
- buffers transport metadata packets
- metadata describes data product
  - **“fingerprints” of completed and pending processing operations**
  - **kernels route through metadata packets if operation does not apply**
  - **different types of products can pass through same pipeline**
  - **optional: fingerprint-based “routing table“ to different output buffers for use with multiple processing chains**

An OS for the SSDP

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Next Steps

- release of Software Requirements document
- Architectural Design (update of NGAPP design)
- Alpha release Q4/2016
- Beta release Q2/2017
- Final release Q4/2017

An OS for the SSDP

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Outlook and Plans

- run time configuration interface already in use with CHEOPS DPU IFSW
- planned follow-up: Xentium DSP kernel library
- OS core (threads etc.) use in Solar wind Magnetosphere Ionosphere Link Exporer (SMILE) Soft X-ray Imager (SXI) instrument LEON3 DPU
- use in Advanced Telescope for High ENergy Astrophysics (ATHENA) Wide Field Imager (WFI) Instrument Control and Power Unit (ICPU) with SSDP

institut für astrophysik
UNIVERSITÄTSSTERNWARTE WIEN
universität wien

Questions?

universität wien

Serving your computational needs.

Since 1365.