institut für astrophysik
UNIVERSITÄTSSTERNWARTE WIEN
universität wien

State of the University of Vienna's flight operating system on the GR740

Armin Luntzer, Roland Ottensamer

Department of Astrophysics, University of Vienna

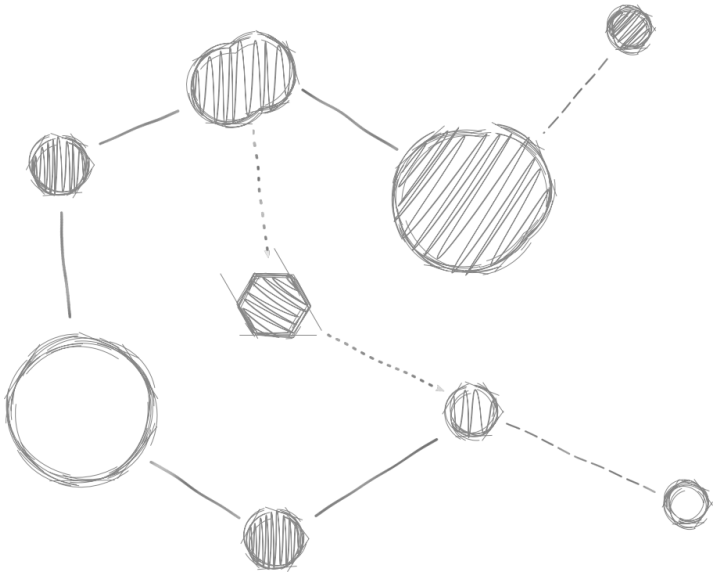

What we do

On-Board Software

- UVIE DoA is a provider of ECSS payload instrument flight software
- system tasks: procedures, communication
- data processing: reduction and compression
- software for half of ESA‘s space telescopes:

  Herschel/PACS (decomissioned)

  CHEOPS (about to launch)

  in development:

  SMILE-SXI, ARIEL-FGS, ATHENA-WFI, PLATO (FPGA Data Compressor)

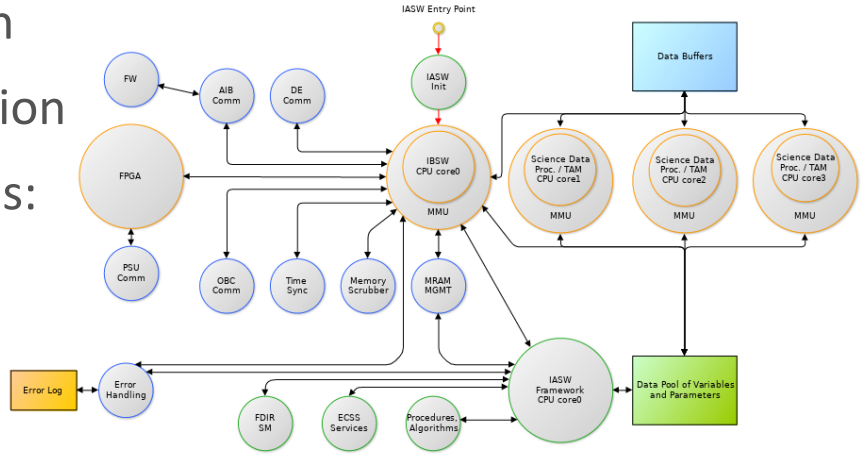

Use Case: ARIEL FGS

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Instrument Data Processing Tasks

- target acquisition (pattern recognition)
- centroiding for AOCS (payload in the loop)
- on-board science data reduction and compression

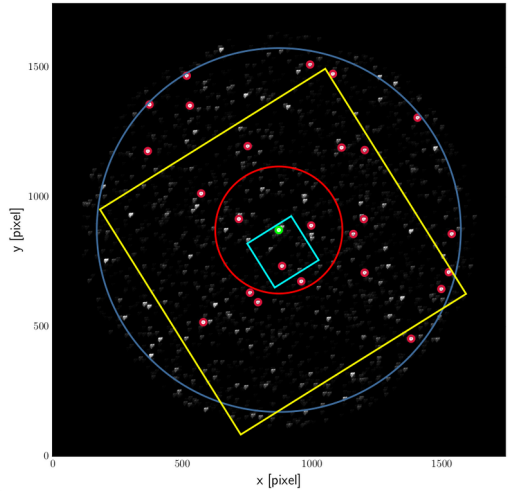

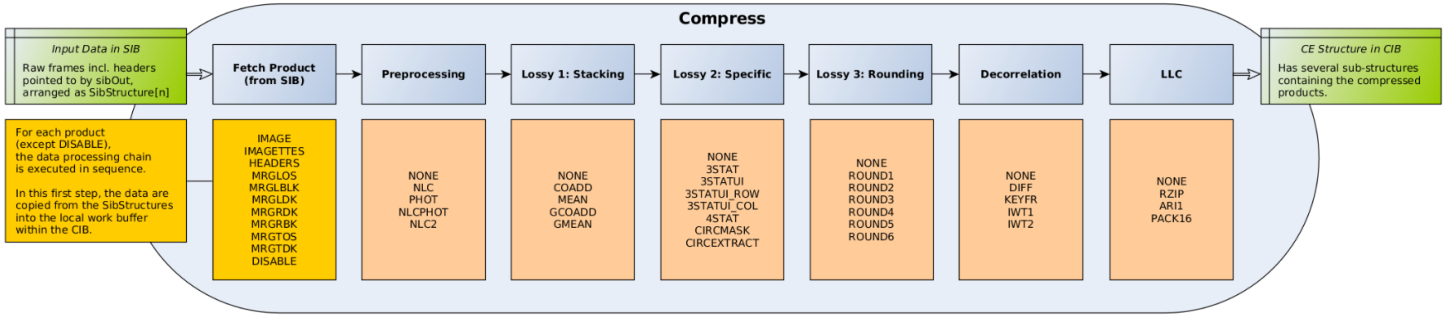

Use Case: ARIEL FGS

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Throughput vs Complexity

- (image) data processing typically memory-dominated
  - image frames can be large (or high frequency)
  - each sample is touched multiple times in different stages
- control tasks typically CPU-dominated:
  - centroid frame frequency up to 10 Hz (31x31 px)
  - „age“ of centroid information is critical (< 20 ms processing time)

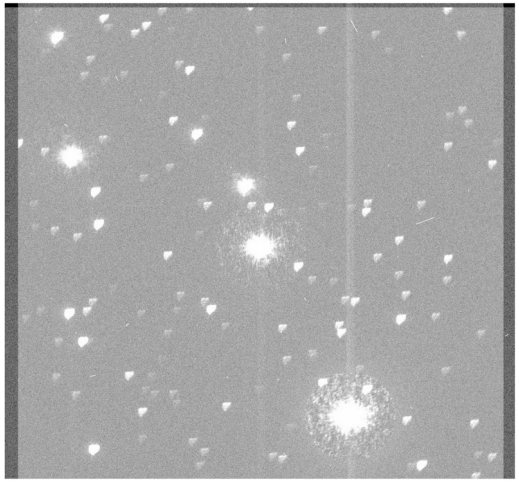

full frame with margins

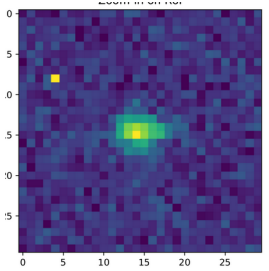

ROI imagette
for centroiding

Benchmarks

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Data Throughput is Critical

- CPUs are as fast as they are
- computational time is determined by algorithm
- memory access is shared

  typically 1 magnitude slower than (load) instruction latency
- Test 1: synthetic, copying of buffers
- Test 2: reduction and compression pipeline

Benchmarks

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Test 1: Synthetic Buffer Copy

- simple 32-bit word copy in a loop
- 4 individual 1 MiB buffers, 1 per CPU
- each CPU copies the buffer of another CPU
- 4 EDF threads (1 per CPU, 95% utilisation)
- 1 RR thread for status printout, no CPU affinity, uses free time slices as available
- modified test for less than 4 CPUs

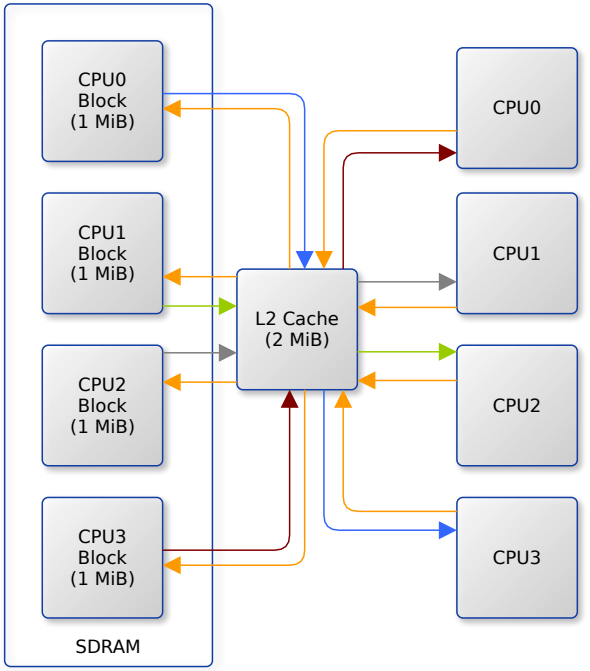

Benchmarks

Test 2: Reduction/Compression

- basic processing network
- 4 steps from CHEOPS pipeline:
  - CCD nonlinearity correction
  - lossy 2-bit rounding
  - 1D differencing decorrelation
  - arithmetic compression
- up to 4 EDF threads (4x pipeline)
- 1 per CPU, 95% utilisation

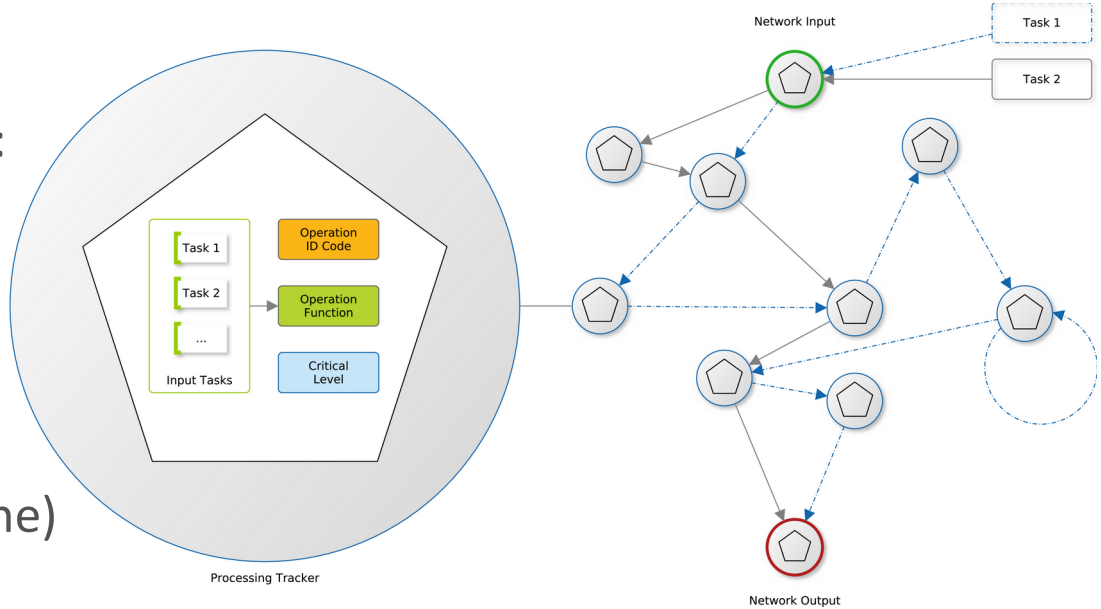

Test Results

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

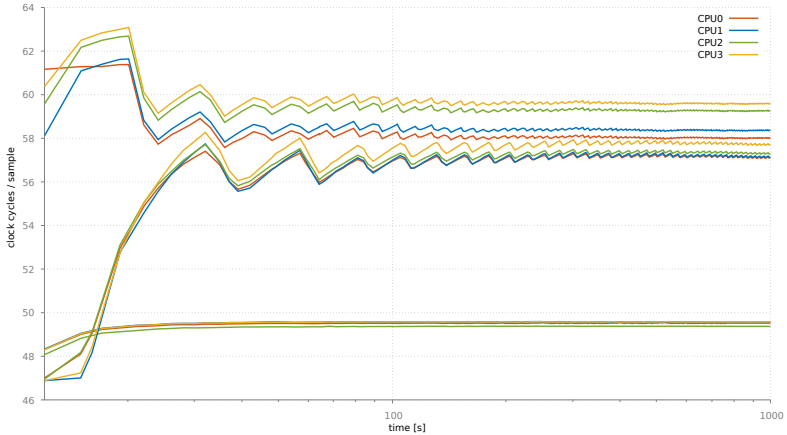

Test 1: Synthetic Buffer Copy

- worst case: 60 cycles/sample
- best case: 49 cycles/sample
- 3 patterns emerge, equally likely
- probably due to (random) shift in thread wakeup
- implies: worst case is real envelope, best case is optimal distribution of 5% down-time between CPUs (allocated EDF runtime is 95% per period)

Test Results

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Test 1: Synthetic Buffer Copy

- comparative tests to see effect of L2 cache

  4 CPUs: 60 c/s, 3 CPUs: 42 c/s

  2 CPUs: 21 c/s, 1 CPU: 18 c/s

- on GR712 eval board using SDRAM:

  2 CPUs: 27 c/s, 1 CPU: 16 c/s

- conclusion: with high memory loads, noticeable stalls occur for 3 or more involved CPUs because of L2 cache size

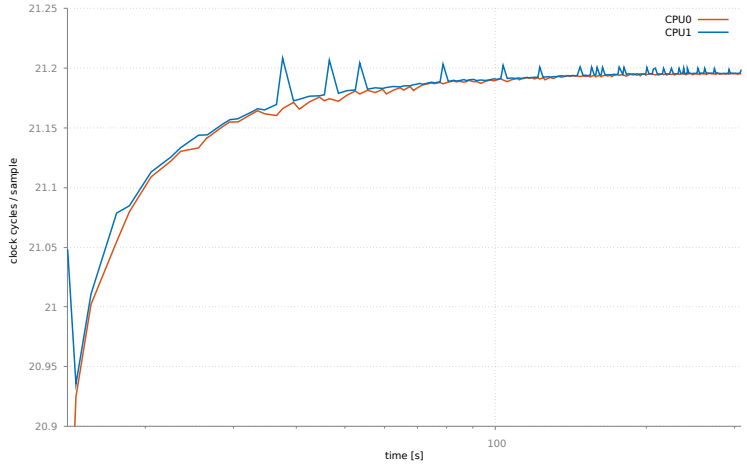

Test Results

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Test 2: Reduction/Compression

- execution time varies significantly for more than one CPU
- can be explained by randomly distributed memory accesses during runtime (initial thread wakeups)
- overall performance scales very well:

  4 CPU case is less than 2x slower compared to single CPU

GR740 eval	
CPUs	c/s
4	825 ± 275
3	700 ± 275
2	650 ± 225
1	475 ± 0.0

CHEOPS DPU (GR712)	
CPUs	c/s
2	635 ± 162
1	542 ± 0.0

Test Results

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Test 2: Reduction/Compression

- results establish a baseline for ATHENA-WFI with GR740
- behaviour is a good indicator on how to design and distribute the science data processing between cores
- tests were enabled by our flight operating system

UVIE Flight OS

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Features

- configurable number of CPUs in SMP mode
- threads, real-time scheduling available (Earliest Deadline First)
- all scheduler clocks run in tickless mode
- SRMMU support, paging and virtual memory
- loadable module and user-supplied executable support (ELF files)
- run-time configuration interface
- compiles with bcc2 gcc or clang

UVIE Flight OS

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Real-Time perfomance (@250 MHz)

- typical OS boot time: 12 ms (3 ms per CPU)
- typical achievable tick interval: 13-22 µs

  appears to depend on CPU id and selected timer, not investigated in detail
- typical task creation time (incl. mem alloc): 9-13 µs
- non-periodic EDF thread exec start after wakeup: 51-70 µs

  this includes an extensive schedulability test, which may be forcibly skipped
- best case absolute deadline: t_{create} + 86 µs

  (deadline - wcet) must be at least the tick interval; assumes schedulability test applied

UVIE Flight OS

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Outlook

- OS to be qualified and used with SMILE (ESA-S, joint ESA/CAS) Soft X-ray Imager (SXI) instrument's DPU with GR712 (launch 2022)
- to be used in ARIEL (ESA-M) FGS Control Electronics (FCE) with GR712 (launch 2028)
- to be used in ATHENA (ESA-L) Wide Field Imager (WFI) Instrument Control and Power Unit (ICPU) with GR740 (launch 2031-2032)

UVIE Flight OS

institut für astrophysik UNIVERSITÄTSSTERNWARTE WIEN universität wien

Summary

- we have selected the GR740 as platform for ATHENA-WFI ICPU
- the SMP system performance assessment was enabled by our OS
- many thanks to ESA/ESTEC for the loan of the GR740 eval board

institut für astrophysik
UNIVERSITÄTSSTERNWARTE WIEN
universität wien

Questions?

universität wien

Serving your computational needs.

Since 1365.