Appendix L

General-purpose GPU Radiative Solver

Andrea Tosetto Marco Giardino Matteo Gorlani
(Blue Engineering & Design, Italy)

Abstract

In the scope of the CADET project, a new radiative tool was developed in Blue Engineering. The tool can compute the extended view-factor and extended incident heat fluxes for solar, planetary and albedo contributions using the Monte Carlo Ray Tracing model. The software is implemented using the OpenCL framework, in order to take advantage of the computational capabilities of GPGPU hardware and dedicated computation hardware.

A comparison between tool results and ESATAN TMS results is performed in order to validate the tool.

GPGPU Radiative solver

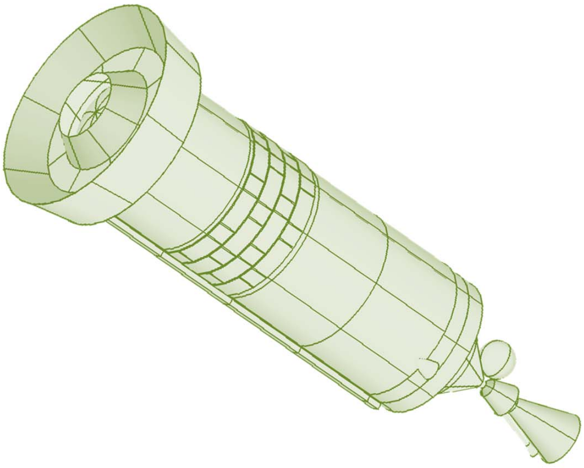

Marco Giardino
m.giardino@blue-group.it
Andrea Tosetto
a.tosetto@blue-group.it
Matteo Gorlani
m.gorlani@blue-group.it

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 1

Contents

- CADET Project
- Radiative heat exchange
- Why GPGPU?
- GPGPU Programming languages
- Software platform structure
- Acceleration structures
- Tool architecture
- Box test
- H10 test
- Tool performances
- Bibliography

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 2

CADET Project

- CApture and DE-orbiting Technologies
- Develop enabling technologies required for Active Debris Removal from LEO orbits:
  - IR and Visible tracking of the target
  - Develop GNC for close rendez-vous, final approach and capture phases
  - technologies, strategies and concepts for target capture and solidarization
- BLUE contribution: onboard IR image processing to get target relative position and motion, with simplified thermal analysis (low power HW)

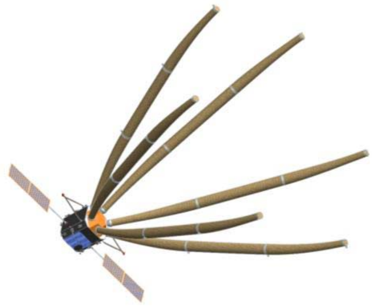

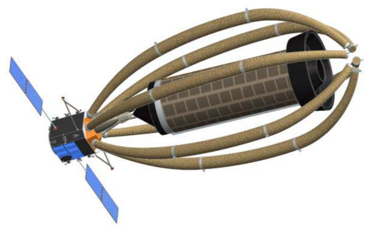

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 3

Radiative heat exchange

Radiation exchange	Heat fluxes
View Factor F_{ij}: fraction of the radiant energy emitted by surface *i* directly intercepted by surface *j*;	**Direct incident heat fluxes**: radiative energy emitted by an external source directly intercepted by a surface;
Radiative exchange factors (REF, also Gebhart factors) B_{ij}: fraction of the radiant energy emitted by surface *i* finally absorbed by surface *j* (including surfaces diffuse reflections);	**Total absorbed heat fluxes**: radiative energy emitted by an external source finally absorbed by a surface (including diffuse reflection component);
Extended view factors F_{ij}: fraction of the radiant energy emitted by surface *i* finally absorbed by surface *j*, directly of through diffuse or specular reflections and through transmissions; → **GR (also REF)**	**Extended incident heat fluxes**: radiative energy emitted by an external source finally absorbed by a surface, directly of through diffuse or specular reflections and through transmissions;

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 4

Why GPGPU?

- Rendering real life pictures is quite similar to evaluate heat fluxes and view factors
- GPU are made for rendering
- Modern GPU are programmable for general purpose calculations using CUDA or OpenCL programming languages (most mature and common)
- GPU are powerful: CPUs ~10GFLOPS, GPUs ~5TFLOPS
- GPU consume less power per GFLOPS (mobile GPUs ~4watts)

General-Purpose computing on Graphics Processing Units (GPGPU)
→ using GPU hardware to perform computations

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 5

GPGPU - history

General-Purpose computing on Graphics Processing Units (GPGPU)
→ using GPU hardware to perform computations

- Dedicated graphic hardware evolved aiming to boost simple operation on huge data sets in order to render more detailed CG scenes
- From 2002 (NVIDIA GeForce 3, ATI Radeon 9700): introduction of programmable rendering pipeline (Direct3D 8.0, OpenGL 1.4)
- Dedicated languages enabled GPU processors to be used for computations:

CUDA, 2007	OpenACC, 2012
OpenCL, 2009	C++ AMP, 2012

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 6

GPGPU – structure (1)

AMD Bulldozer CPU block diagram

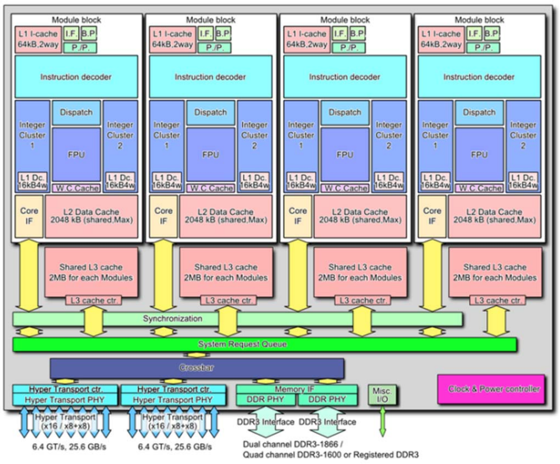

AMD Radeon HD 7790 GPU block diagram

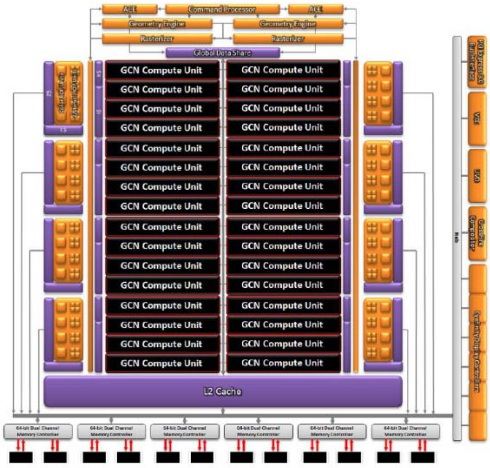

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 7

GPGPU – structure (2)

	CPU	GPU
Architecture	MIMD (Multiple Instruction Multiple Data) → Multi purpose, independent processors/cores	SIMD (Single Instruction Multiple Data) → 1 control unit dispatch commands to multiple ALUs. Texture units
Execution	Branch execution flow, control structures	Optimized for branchless execution (maximize occupancy)
Memory structure	• Low latency, low bandwidth memory (RAM, up to 20 GB/s) • L1, L2, L3 levels cache memory • CPU registers	• High latency, high bandwidth memory (VRAM, up to 300 GB/s) • On chip, block shared memory • GPU registers
Memory management	Automatic (e.g. cache pre-fetch)	User controlled data flow between VRAM and shared memory, registers
Context switch	Software (thread management by OS)	HW (thread switch controlled by the control unit to hide memory latency)
Precision	Single, double, quad, with almost same performances	Preferred single, double available but with huge performance losses (at best, ¼ of single precision on latest hardware)

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 8

GPGPU Programming languages

CUDA	OpenCL
GPGPU language by NVIDIA	Open standard by Khronos Group (OpenGL)
Specific and available only on NVIDIA GPUs	Designed for heterogeneous computing: available on NVIDIA and AMD GPUs, ARM processors, CPUs (Intel & AMD), FPGA vendors
Offline kernel compilation to intermediate language	Runtime kernel compilation of the provided kernel sources by the device specific driver
Mature framework, with libraries (CuFFT, CuBLAS, etc.) and tools (IDE, profiler, debugger, etc.)	Less advanced libraries and tools, generally from open source projects
Work only on NVIDIA hardware	The standard guarantee software portability among different HW (NB execution performances are not guaranteed: they can only be achieved tuning software on each specific hardware architecture)

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 9

GPGPU Platform (OpenCL dialect)

Host (CPU)	Device (GPU, FPGA, etc)
Host code: C, C++, Fortran, C#, Java, Python, ecc.	Device code (kernels): OpenCL C, which is a modified C99 (add vector primitives and vector functions; some restrictions in use of pointers, some standard C99 headers are missing)
Manage the algorithm workflow	Perform computations on the provided data sets
Manage the data transfer between host memory (RAM) and device memory (VRAM); memory can be a shared area between host and device	Manage data flow among device mass memory, device shared memory and processor registers, through kernel's instructions
Manage global execution synchronization	Execution is divided in 1D, 2D or 3D workgroups; block synchronization can be obtained through device shared memory

Host code controls the parallel execution of dedicated kernels on one or more devices; kernels are compiled during initialization

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 10

Computer Graphics Rendering Heritage

- Rasterisation → objects lightning obtained through artifacts (textures)
- Global illumination → direct evaluation of an objects appearance, due to objects and light sources positions, object properties, etc
- Ray Tracing Monte Carlo (RTMC) → Elegant solution to compute global illumination, under development from 1980s [Kay & Kajiya 1986]

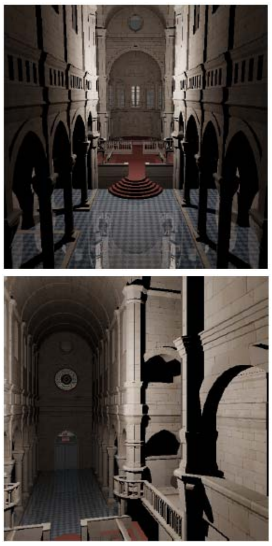

[Fabianowski 2011]

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 11

Acceleration Structures

Organize the geometry elements in order to compute ray-geometry intersection without testing **ALL** the model elements

- Voxels
- Octrees
- K-d trees
- Binary Volume Hierarchy (BVH):
  - Recursive binary tree partition
  - Each element belong to a single tree leaf

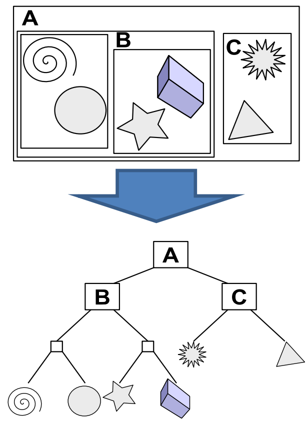

Courtesy Wikipedia

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 12

Tool architecture

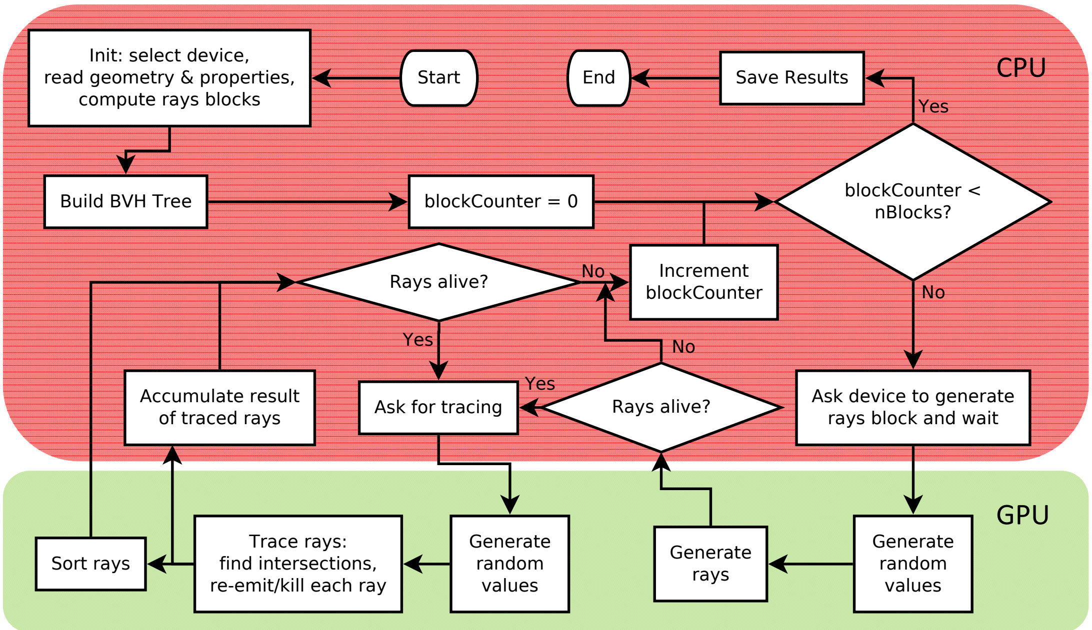

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 13

Open Box test

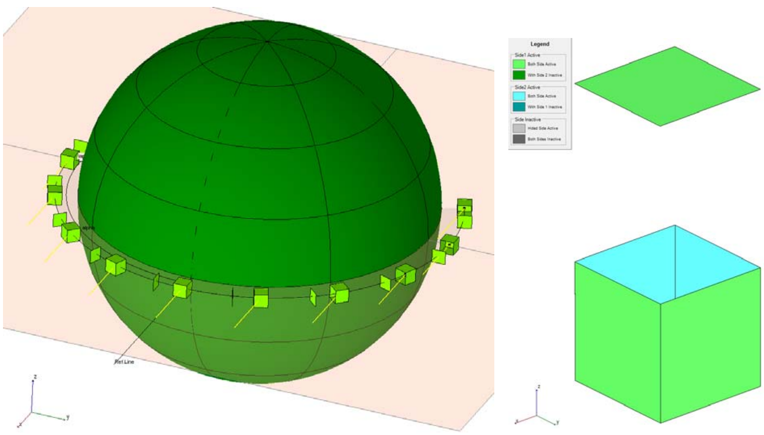

Node	Pos.	Node	Pos.
10	-X,1	11	-X,2
20	-Y,1	21	-Y,2
30	-Z,1	31	-Z,2
40	+X,1	41	+X,2
50	+Y,1	51	+Y,2
60	+Z,1	61	+Z,2

Orbit height	800 km over Earth surface
Sun constant	1000 W/m^2
Ray densities	10k rays/m^2 (Sun), 100k rays/m^2 (Earth, GR),

Side 1 optic	UV: $\alpha = 1$ IR: $\varepsilon = 0.5$, $\tau = 0$, $\rho_{ratio} = 0.5$
Side 2 optic	UV: $\alpha = 0$, $\tau = 0$, $\rho_{ratio} = 0$ IR: $\varepsilon = 0.5$, $\tau = 0$, $\rho_{ratio} = 0.5$

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 14

Open Box test - albedo results

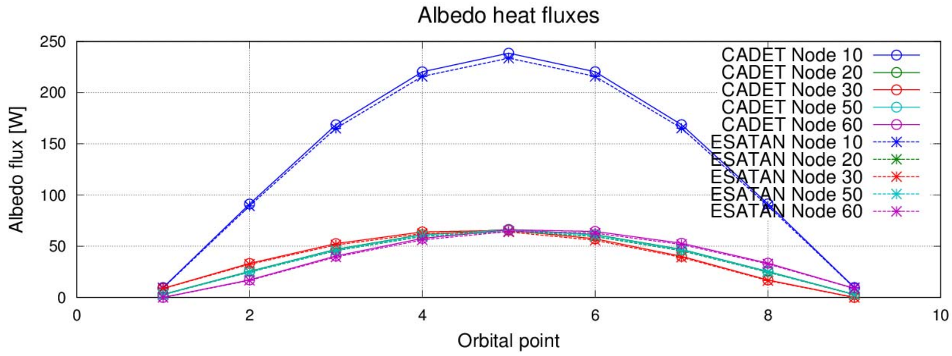

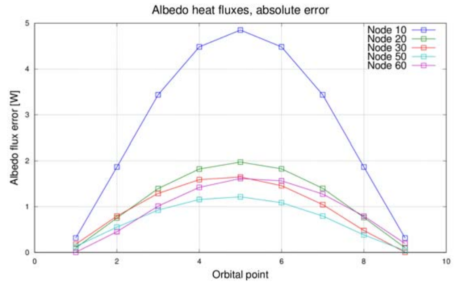

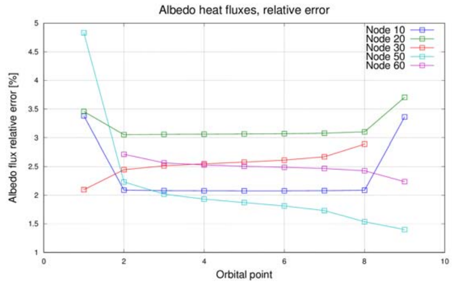

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 15

Open Box test – Earth flux results

Node	10	11	20	21	30	31
CADET	153.7860	6.7952	43.1359	11.6836	42.5930	10.3962
ESATAN	153.6489	6.9998	42.6820	11.4758	42.3568	10.1615
Abs. err.	0.137070	-0.204551	0.453887	0.207798	0.236169	0.234655
Rel. Err. %	0.089209	2.922240	1.063411	1.810749	0.557570	2.309245

Node	40	41	50	51	60	61
CADET	0	34.9459	42.9083	11.7025	42.9015	42.9358
ESATAN	0	35.0538	42.9592	11.4846	42.6880	42.7810
Abs. err.	0	-0.107996	-0.050915	0.217839	0.213462	0.154752
Rel. Err. %	0	0.308085	0.118519	1.896782	0.500051	0.361730

NB due to the selected attitude, Earth fluxes are the same on each point of the orbit

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 16

Open Box test – Sun flux results

		10	11	20	21	30	31	40	41	50	51	60	61	Tot
CADET Sun Fluxes [W]	P1	0.00	0.00	0.00	0.00	1000.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1000.00
	P2	0.00	0.00	0.00	0.18	923.88	0.00	382.68	0.00	0.00	0.09	0.00	0.00	1306.84
	P3	0.00	0.00	0.00	0.00	707.11	0.07	707.11	0.14	0.00	0.00	0.00	0.00	1414.43
	P4	0.00	0.00	0.00	0.00	382.68	0.04	923.88	0.08	0.00	0.00	0.00	0.00	1306.68
	P5	0.00	0.00	0.00	0.00	0.00	0.00	1000.00	0.00	0.00	0.00	0.00	0.00	1000.00
	P6	0.00	0.09	0.00	0.18	0.00	0.09	923.88	0.09	0.00	0.00	382.68	0.18	1307.21
	P7	0.00	0.28	0.00	0.14	0.00	0.42	707.11	0.57	0.00	0.35	707.11	0.07	1416.05
	P8	0.00	0.28	0.00	0.28	0.00	0.26	382.68	0.50	0.00	0.22	923.88	0.09	1308.19
	P9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1000.00	0.00	1000.00

		10	11	20	21	30	31	40	41	50	51	60	61	Tot
Absolute error, CADET – ESATAN [W]	P1	-0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.05
	P2	0.00	0.00	0.00	0.18	-0.02	0.00	0.00	0.00	0.00	0.09	0.00	0.00	0.26
	P3	0.00	0.00	0.00	0.00	-0.02	0.07	0.00	0.14	0.00	0.00	0.00	0.00	0.19
	P4	0.00	0.00	0.00	0.00	-0.02	0.04	0.00	0.08	0.00	0.00	0.00	0.00	0.10
	P5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	P6	0.00	0.09	0.00	0.18	0.00	0.09	0.00	0.09	0.00	0.00	-0.02	0.18	0.63
	P7	0.00	0.28	0.00	0.14	0.00	0.42	0.00	0.57	0.00	0.35	-0.02	0.07	1.81
	P8	0.00	0.28	0.00	0.28	0.00	0.26	0.00	0.50	0.00	0.22	-0.02	0.09	1.61
	P9	-0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.05

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 17

Open Box test – Sun flux results

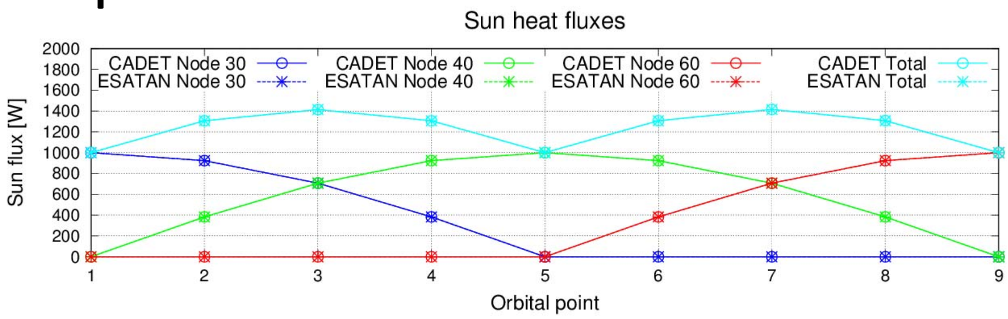

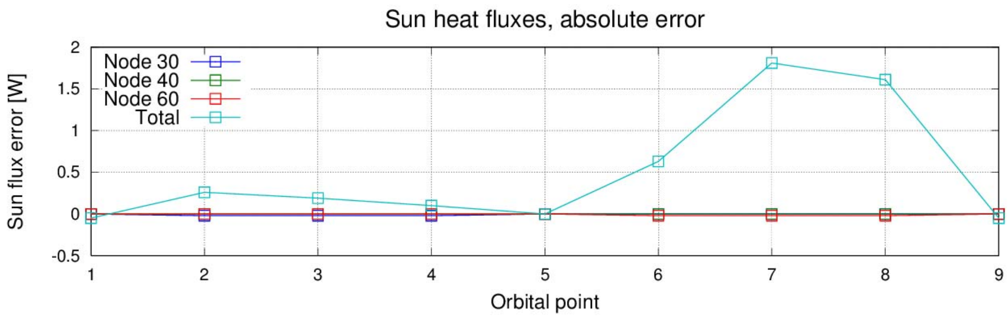

All other nodes: null flux due to attitude and null alpha

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 18

H10 test – Description (1)

- Ariane IV 3rd stage, standard CADET target
- Analyze target thermal behavior to allow IR tracking
- Test orbital parameters:

Radius *r*	7161 km
Eccentricity *e*	0
Inclination *i*	45°
Argument of periapsis *ω*	0°
Ascending node *Ω*	-90°

- Attitude: LOCS, Z body aligned with $-\vec{V}$, Y body aimed to Z Earth
- Optical properties: guess, obtained through aerospace materials data base ($\overline{\alpha}$=0.537, $\overline{\varepsilon}$=0.266)
- Ray density (rays/m^2): 10k (Sun), 100k (Earth, albedo)
- 41 orbital position computed (half orbit, no eclipse)

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 19

H10 test – Description (2)

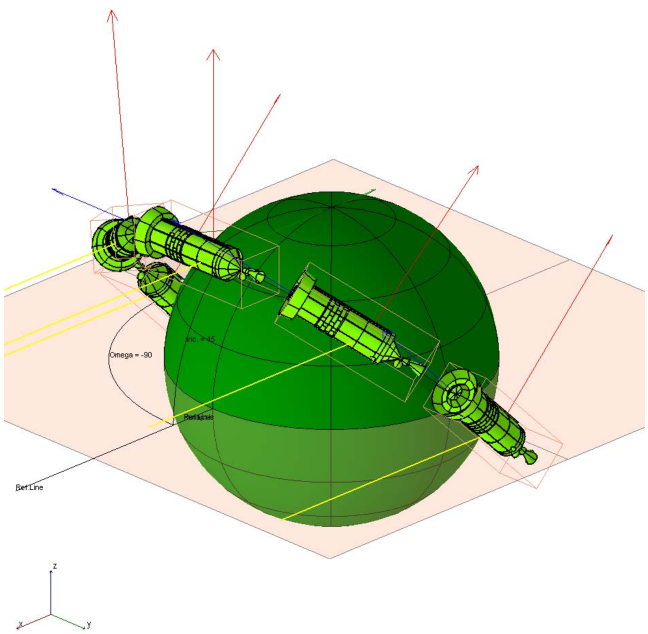

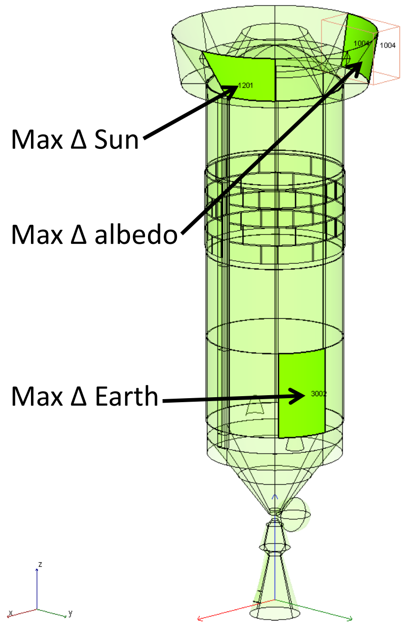

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 20

H10 test – total albedo flux

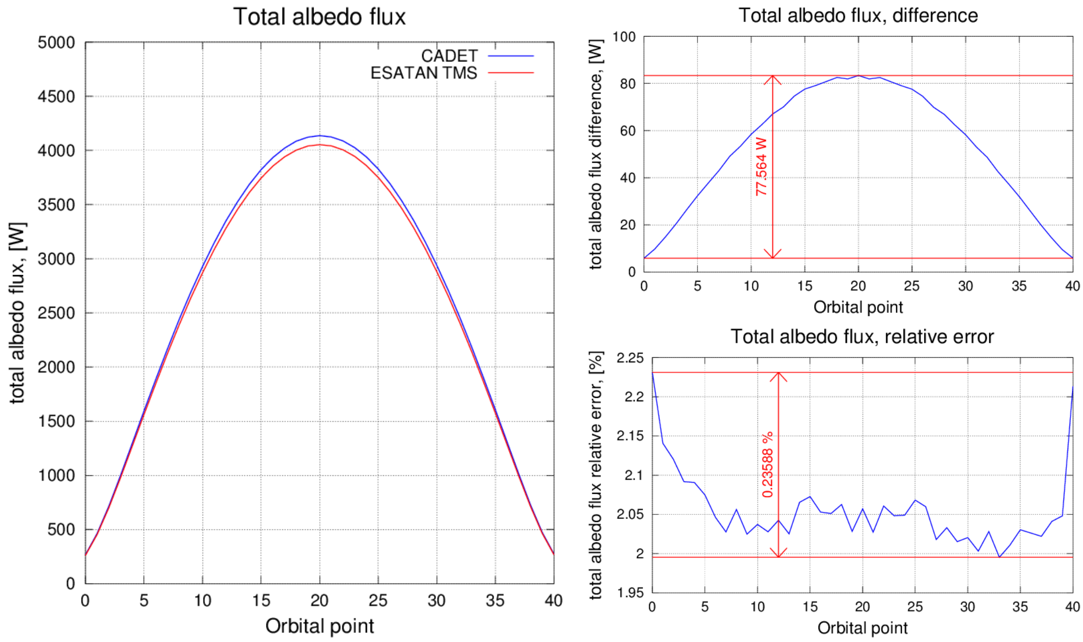

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 21

H10 test – max. albedo flux diff.

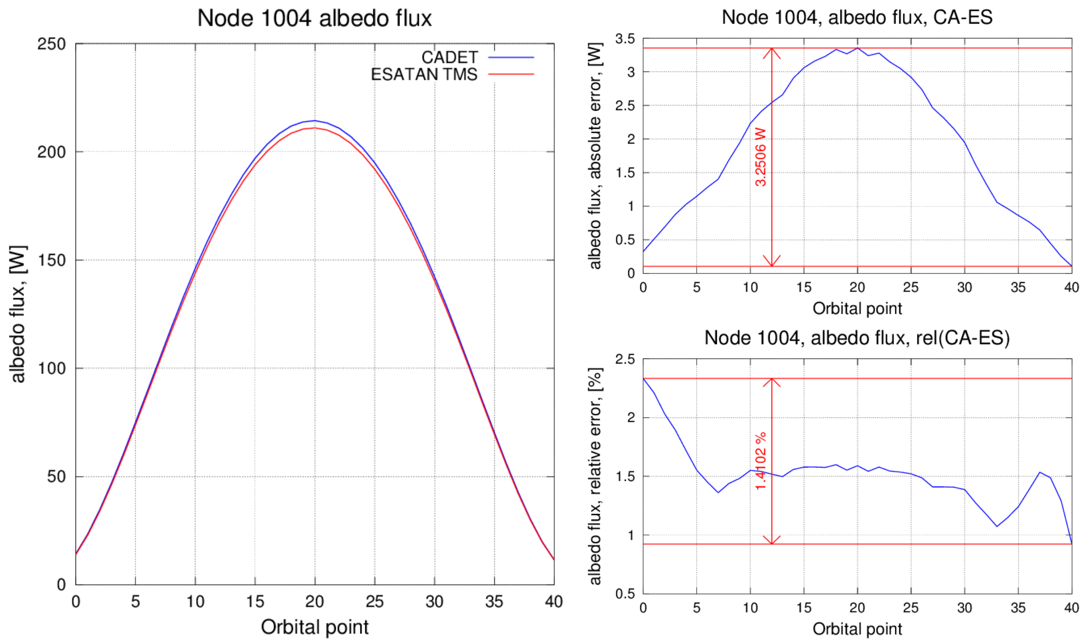

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 22

H10 test – total Earth flux

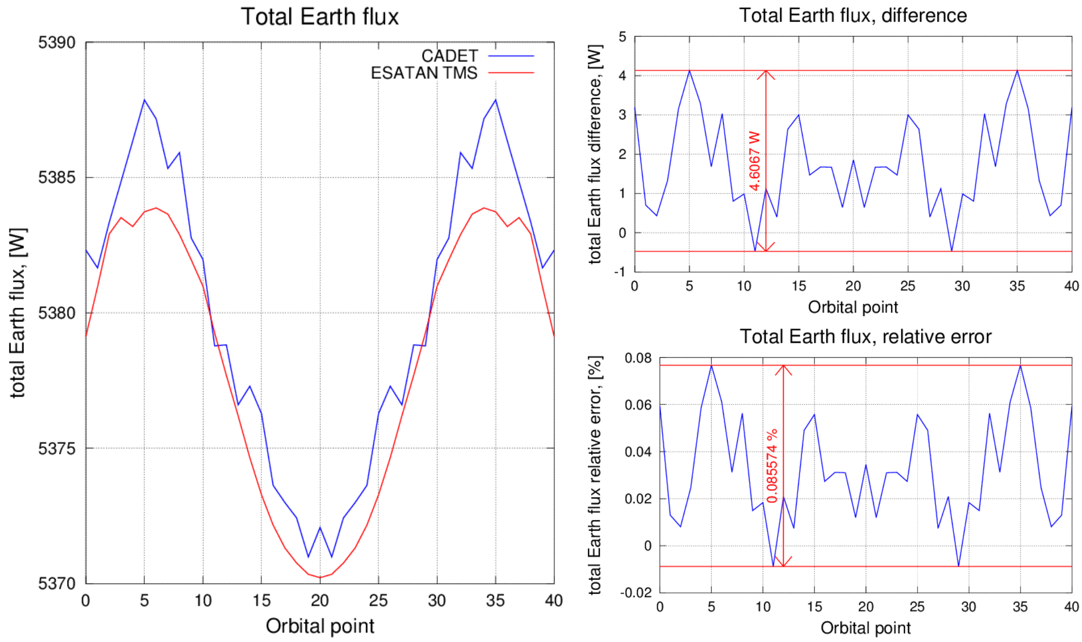

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 23

H10 test – max Earth flux diff.

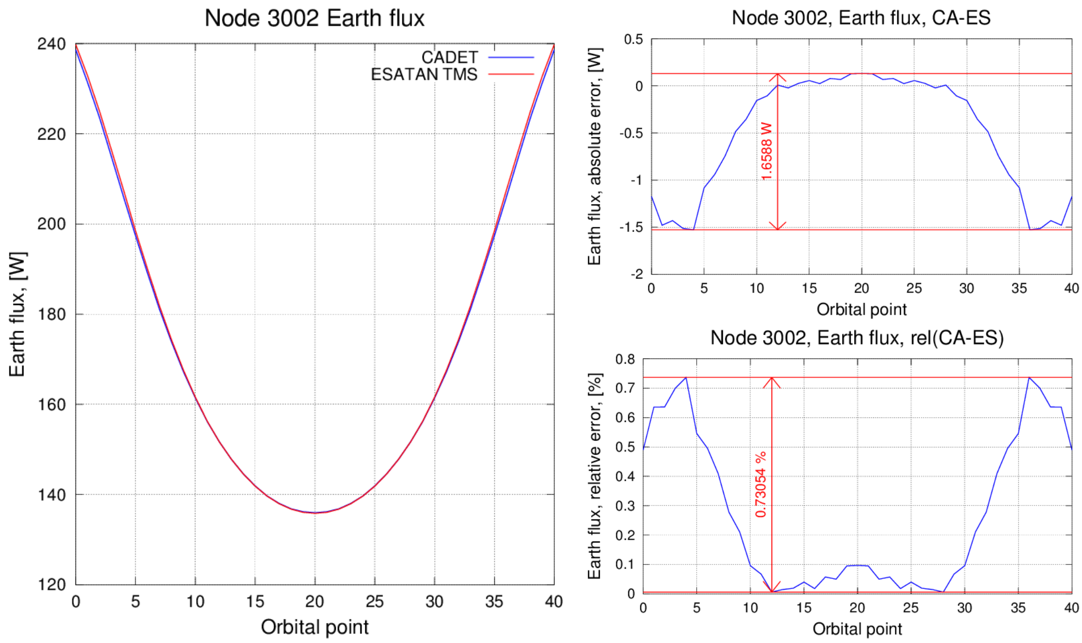

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 24

H10 test – total Sun flux

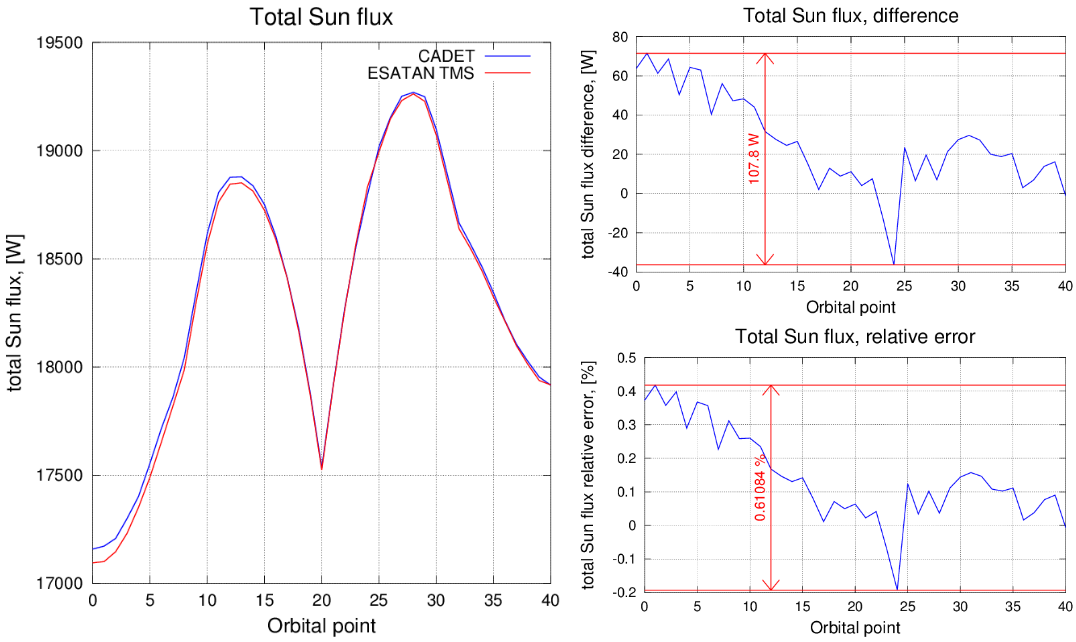

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 25

H10 test – max Sun flux diff.

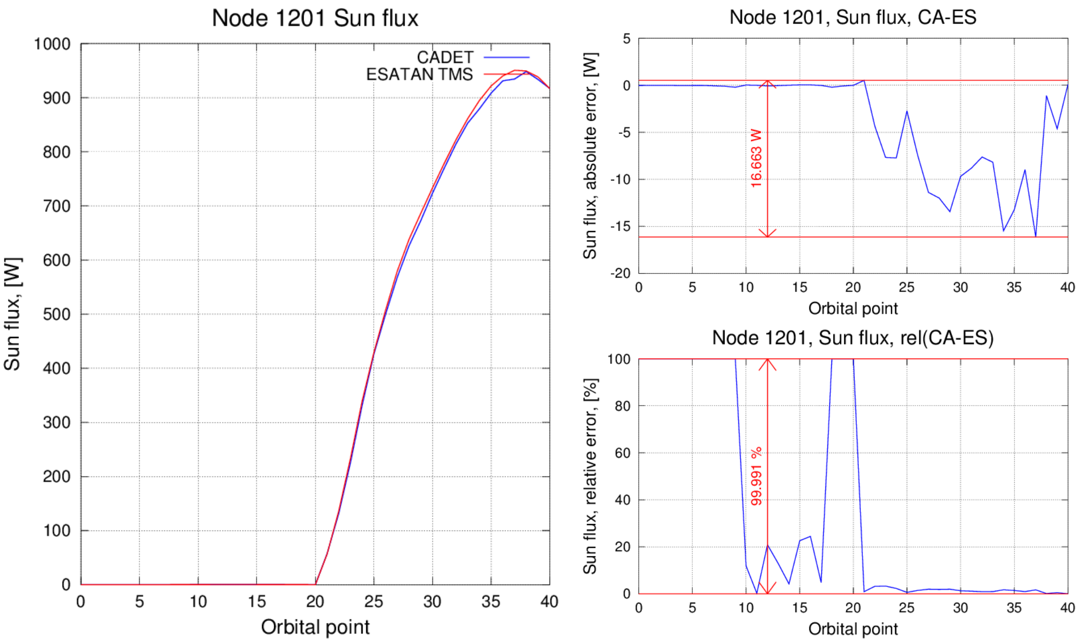

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 26

H10 test – GR

Case	ESATAN TMS	CADET
# GR	9453	9287
Common GR	6495	
Additional GR	2958	2792
$\sum GR_{Uncommon}$	1.2167e-3	57.818e-3
Max $Gr_{Uncommon}$	0.1372e-3	1.079e-3

Case	ESATAN TMS	CADET	Δ X1000	Δ%
$\sum GR$	83.070	83.154	84.240	0.101
$\sum GR_{Model\ to\ space}$	48.304	48.298	-6.476	-0.013
$\sum GR_{Model\ to\ inactive}$	6.595	6.595	0.834	0.012
$\sum GR_{Model\ to\ model}$	28.17	28.26	89.882	0.319
$\sum GR_{Common}$	83.068	83.096	27.859	0.033
Max Δ, GR(3105,3205)	0.31389	0.31152	2.372	0.755

ESATAN and CADET GR cut off value: 1e-6

H10 test – GR

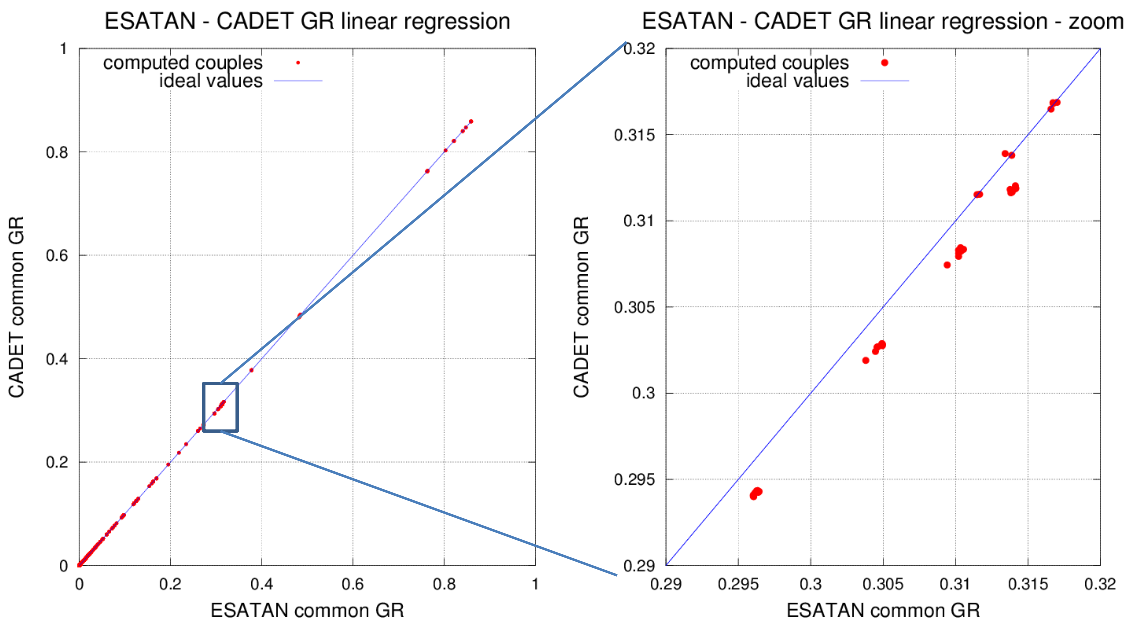

Performance – Test Hardware

Device	Vendor	Type	CU	Clock GHz	Memory Type	Memory GB	Memory Float4 GB/s	Memory Float16 GB/s	Peak performances Int4 GIOPS	Peak performances Float4 GFLOPS	Peak performances Float16 GFLOPS
Core i7 940	Intel	CPU	4+4	2.93	DDR3	18	12.29	12.29	8.90	27.77	25.69
Core i3 2370M	Intel	CPU	2+2	2.40	DDR3	4	9.56	9.06	/	19.12	19.12
Core i3 3220	Intel	CPU	2+2	3.30	DDR3	8	10.96	11.05	/	26.19	26.61
Core i7 4700HQ	Intel	CPU	2.4	2.40	DDR3	8	17.20	17.09	/	35.06	36.71
GT 630	NVIDIA	GPU	2	1.62	GDDR3	1	17.52	4.56	103.07	307.50	280.09
GTX 670M	NVIDIA	GPU	7	1.20	GDDR5	1.5	61.46	15.36	/	786.62	718.22
GT 750M	NVIDIA	GPU	2	1.10	DDR3	4	25.60	11.64	/	740.01	740.00
GTX Titan Black	NVIDIA	GPU	15	0.98	GDDR5	6	298.26	107.37	1010.58	4685.42	4685.42
Quadro 2000	NVIDIA	GPU	4	1.25	GDDR5	1	35.87	8.97	158.94	472.28	459.16
HD 5850	AMD	GPU	18	0.725	GDDR5	1	86.04	34.41	443.27	1738.85	1782.76
HD 7870 GE	AMD	GPU	20	1.00	GDDR5	2	130.31	35.51	505.76	2507.83	437.86
HD 7950 Boost	AMD	GPU	28	1.15	GDDR5	3	278.37	48.54	748.95	3725.15	3650.85
HD 8970M	AMD	APU*	5	0.825	DDR3	2	7.99	2.21	/	251.56	247.00

* Accelerated Processing Unit, AMD definition of a chip which contains a CPU section and a GPU section

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 29

Performance – Results

Device	Tracing, Mr/s Sun	Tracing, Mr/s REF Planet	Execution times, seconds 41 Sun fluxes	Sun, mean	GR + UV REF	41 Earth + Albedo (CPU)	Earth + Albedo, mean	Total	Gain
Core i7 940	2.25	1.61	57.21	1.40	224.28	8.94	0.22	302.66	3.24
Core i3 2370M	1.21	0.65	240.72	5.87	590.10	15.02	0.40	863.11	1.14
Core i3 3220	1.23	1.01	106.41	2.60	378.89	20.14	0.50	524.39	1.87
Core i7 4700HQ	2.25	1.56	57.14	1.39	231.73	5.17	0.13	303.27	3.24
GT 630	2.02	1.33	62.98	1.54	307.10	20.15	0.48	410.16	2.39
GTX 670M	4.29	3.39	29.61	0.72	119.21	12.70	0.31	167.00	5.88
GT 750M	2.73	1.99	46.61	1.14	158.05	5.53	0.19	218.61	4.49
GTX Titan Black	9.42	13.35	13.49	0.33	30.24	7.12	0.17	53.84	18.23
Quadro 2000	1.65	1.29	40.82	1.00	302.21	7.68	0.19	357.95	2.74
HD 5850	3.61	4.34	35.26	0.86	94.15	5.13	0.19	144.74	6.78
HD 7870 GE	6.21	8.52	20.71	0.51	47.60	5.07	0.12	81.58	12.03
HD 7950 Boost	8.18	13.85	15.54	0.38	29.48	6.03	0.15	58.93	16.66
HD 8970M	3.08	2.85	41.23	1.01	143.36	5.16	0.14	197.28	4.98
ESATAN TMS*	/	/	319.00	7.78	650.00	12.50	0.30	981.50	1.00

* Calculation on Intel i7 940 2.93 GHz; single core execution with Intel Turbo Boost, (clock set to 3.2 GHz)
All computations on Windows 7/8 64 machines; tools also works on Linux

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 30

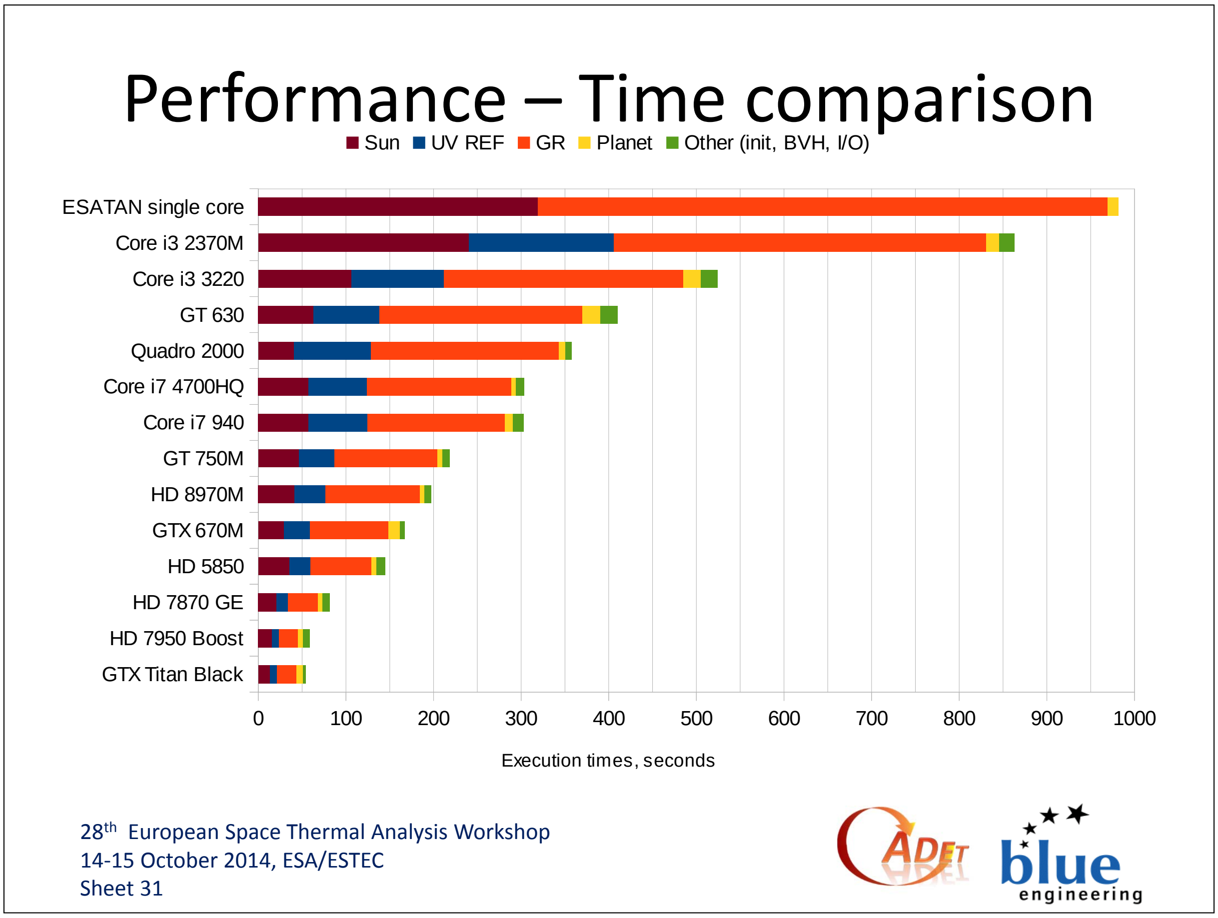
Performance – Time comparison
Sun
UV REF
GR
Planet
Other (init, BVH, I/O)
ESATAN single core
Core i3 2370M
Core i3 3220
GT 630
Quadro 2000
Core i7 4700HQ
Core i7 940
GT 750M
HD 8970M
GTX 670M
HD 5850
HD 7870 GE
HD 7950 Boost
GTX Titan Black
0
100
200
300
400
500
600
700
800
900
1000
Execution times, seconds
28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 31
ADET
blue
engineering

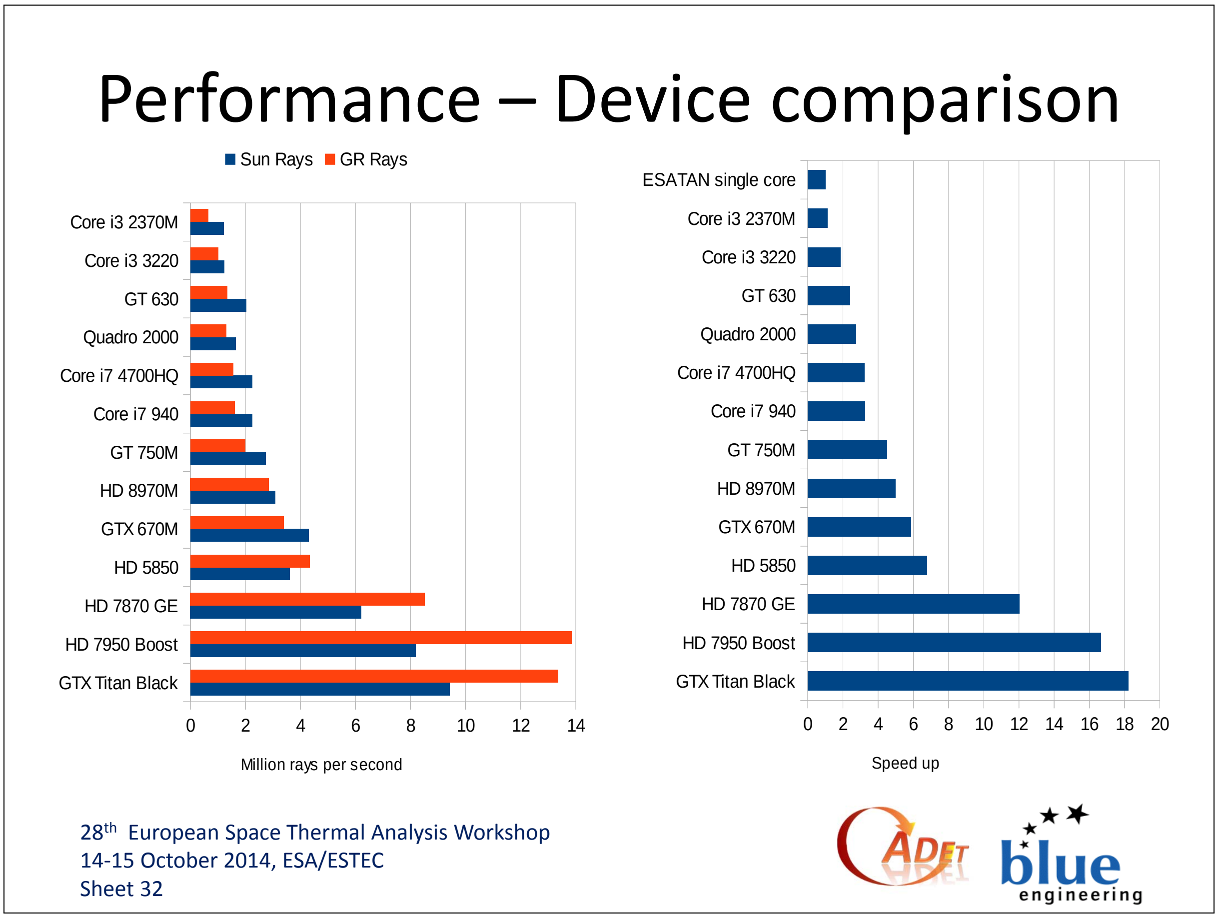
Performance – Device comparison
Sun Rays
GR Rays
Core i3 2370M
Core i3 3220
GT 630
Quadro 2000
Core i7 4700HQ
Core i7 940
GT 750M
HD 8970M
GTX 670M
HD 5850
HD 7870 GE
HD 7950 Boost
GTX Titan Black
0
2
4
6
8
10
12
14
Million rays per second
ESATAN single core
Core i3 2370M
Core i3 3220
GT 630
Quadro 2000
Core i7 4700HQ
Core i7 940
GT 750M
HD 8970M
GTX 670M
HD 5850
HD 7870 GE
HD 7950 Boost
GTX Titan Black
0
2
4
6
8
10
12
14
16
18
20
Speed up
28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 32
ADET
blue
engineering

Performance – CG reference

Reference code [Aila&Laine, 2009]	CADET code
Language: **CUDA**	Language: **OpenCL**
Sibenik Cathedral model (80131 triangles)	Sibenik Cathedral model (80131 triangles)
NVIDIA GT 630	NVIDIA GT 630
Test on diffuse illumination rays	GR/UV REF computation
15M rays/s	**~1.17M rays/s**

Reference code is about 13 times quicker than the developed the code!

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 33

Possible Improvements

- Data flow:
  - Float16 to float4 data items conversion
  - Data caching through constant memory blocks for initialization (BVH data etc.)
  - Data caching through texture units (simpler if OpenCL 1.2 is used)
- Improve BVH structure (Split BVH construction → SBVH)
- Stack-less tracing kernel (reduce register pressure and improve the number of concurrent threads)
- Take advantage of multiple devices (multi GPU, GPU+CPU configurations)
- Evaluate different ray tracing strategies:
  - Offloading part of the ray tracing procedure to the CPU when few rays need to be traced, and leave on device only mass evaluations
  - Use device fission (OpenCL 1.2) to parallelize ray spawning and ray tracing on device, in order to maximize device occupancy

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 34

Conclusions

- CADET radiative solver is implemented and its results are consistent with the ESATAN TMS tool (considering RTMC fluctuations).
- On GPU devices, the tool use the huge computation power of GPU HW and get acceptable performances (up to 18x respect to ESATAN TMS on high end hardware).
- Additional speedup of 10x and more could be achieved improving the tool (compared to NVIDIA computer graphics rendering algorithm).
- New developers' skills are needed to maximize GPGPU performances.

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 35

Bibliography

- Kay, Timothy L., and James T. Kajiya. "*Ray Tracing Complex Scenes*." *SIGGRAPH Comput. Graph.* 20, no. 4 (August 1986): 269–78. doi:10.1145/15886.15916.
- Fabianowski, B. "*Interactive Manycore Photon Mapping.*" Trinity College Dublin, 2011.
- Aila, Timo, and Samuli Laine. "*Understanding the Efficiency of Ray Traversal on GPUs*." In Proceedings of the Conference on High Performance Graphics 2009, 145–49, 2009.
- Renard, Patrice. "*6. Theoretical Background of the Radiative Module.*" In *Thermica User Manual*, 6–1 – 6–28, 1997.
- Munshi, Aaftab. "*The OpenCL Specification, Version 1.1*". Khronos OpenCL Working Group, June 1, 2011.
- Gaster, Benedict R. "*The OpenCL C++ Wrapper API, Version 1.1*". Khronos OpenCL Working Group, June 2010.
- NVIDIA Corporation. "*OpenCL Programming Guide for the CUDA Architecture v4.2*". NVIDIA Corporation, 2012.
- Advanced Micro Devices, Inc. "*AMD OpenCL™ Programming Guide*" 2013.

28th European Space Thermal Analysis Workshop
14-15 October 2014, ESA/ESTEC
Sheet 37