



Publication Year	2019
Acceptance in OA	2020-12-15T14:09:54Z
Title	Astrophysical Code Migration into Exascale Era
Authors	Goz, David, BERTOCCO, SARA, TORNATORE, Luca, TAFFONI, Giuliano
Handle	http://hdl.handle.net/20.500.12386/28858
Serie	ASTRONOMICAL SOCIETY OF THE PACIFIC CONFERENCE SERIES
Volume	523

Astrophysical Code Migration into Exascale Era

D. Goz, S. Bertocco, L. Tornatore, and G. Taffoni

INAF - Osservatorio Astronomico di Trieste - via Tiepolo 11, 34131 Trieste
Italy; david.goz@inaf.it

Abstract. The ExaNeSt and EuroExa H2020 EU-funded projects aim to design and develop an exascale ready computing platform prototype based on low-energy-consumption ARM64 cores and FPGA accelerators. We participate in the application-driven design of the hardware solutions and prototype validation. To carry on this work we are using, among others, *Hy-Nbody*, a state-of-the-art direct N -body code. Core algorithms of *Hy-Nbody* have been improved in such a way to increasingly fit them to the exascale target platform. Waiting for the ExaNest prototype release, we are performing tests and code tuning operations on an ARM64 SoC facility: a SLURM managed HPC cluster based on 64-bit ARMv8 Cortex-A72/Cortex-A53 core design and powered by a Mali-T864 embedded GPU. In parallel, we are porting a kernel of *Hy-Nbody* on FPGA aiming to test and compare the performance-per-watt of our algorithms on different platforms. In this paper we describe how we re-engineered the application and we show first results on ARM SoC.

1. Introduction

The current market offers low-power micro-processor hardware solutions integrating enough transistors to include an on-chip floating-point unit capable of running typical HPC (High Performance Computing) applications. They are less expensive and more power-efficient than standard HPC devices. For this reason SoC (System on Chip) solutions are a possible approach to actually reduce the costs of HPC in terms of time and power consumption and this becomes extremely important when designing the new generation of HPC supercomputer, the exascale platforms. The ExaNest H2020 project (Katevenis et al. 2016) aims at the design and development of an exascale-class prototype computing system built upon power-efficient hardware able to execute real-world applications coming from a wide range of scientific and industrial domains, including also HPC for astrophysics (Ammendola et al. 2017). The ExaNeSt basic compute unit consists of low-energy-consumption ARM CPUs, FPGAs and low-latency interconnects (Katevenis et al. 2018).

Programmers will have to re-engineer their applications in order to fully exploit this new exascale platform based on heterogeneous hardware. We studied whether a direct N -body code for real scientific production may benefit from embedded GPUs given that the powerful high-end GPUs already have demonstrated to provide tremendous performance benefit for N -body code. To the best of our knowledge, this is the first work to implement such algorithm on embedded GPUs and to compare results with multi-core solutions on a SoC implementation.

2. Code implementation

Hy-*N*body is a direct N -body code that relies on the Hermite 6th order time integrator and that has been conceived to exploit hybrid hardware. The code is derived from HiGPUs (Capuzzo-Dolcetta et al. 2013; Capuzzo-Dolcetta & Spera 2013; Spera 2014), which has been widely used for simulations of star clusters with up to ~ 8 million bodies (Spera & Capuzzo-Dolcetta 2015; Spera et al. 2015), and of galaxy mergers (Bortolas et al. 2016). The kernels of Hy-*N*body have been developed with OpenCL in order to write efficient code for hybrid (CPU/GPU/FPGA) architecture. Kernels have been optimized using (i) *vectorization*, to increase the number of operations per cycle, and exploiting the (ii) *local memory* of the device, to reduce the latency of data transactions. The OpenCL host code is parallelized with hybrid MPI+OpenMP programming. A one-to-one correspondence between MPI processes and computational nodes is established and each MPI process manages all the OpenCL-compliant devices of the same type available per node. Inside of each shared-memory computational node, parallelization is achieved by means of OpenMP environment.

The Hermite 6th order integration schema requires double precision (DP) arithmetic in the evaluation of inter-particles distance and acceleration in order to minimize the round-off error. Full IEEE-compliant DP-arithmetic is efficient in available CPUs and GPGPUs, but it is still extremely resource-eager and performance-poor in other accelerators like embedded GPUs or FPGAs. The extended-precision (EX) numeric type is a valuable alternative in porting our application on devices not specifically designed for scientific calculations, such as embedded GPUs or FPGAs. We implemented in Hy-*N*body the EX-arithmetic as proposed by Thall (2006).

On SoC the memory is shared between CPU and GPU, so using local memory as a cache with associated barrier synchronization can waste both performance and power. For this reason, we implemented a specific embedded-GPU-optimized version of all kernels of Hy-*N*body.

3. Testbed description

We deployed a cluster based on heterogeneous hardware (CPU+GPU) to validate and test the Hy-*N*body code. Each computational node is a Rockchip Firefly-RK3399 single board computer. It is a six core 64-bit High-Performance Platform, based on SoC with the ARM big.LITTLE architecture. The main characteristics of this cluster, named INCAS¹, are listed in Table 1, while full details are in Bertocco et al. (2018).

4. Performance results

We just focused on the most computationally demanding kernel of the Hermite 6th order algorithm (with N bodies the kernel has $O(N^2)$ computational cost) and compared the performances on ARM CPUs. The left panel of Figure 1 shows the ratio of the best running time achieved by the CPUs as a function of the number of particles for both arithmetics. ARM Cortex-A72 with two cores is faster than Cortex-A53 with four cores by approximately a factor of two.

¹INTensive Clustered Arm-Soc

Table 1. The main characteristics of our cluster used to test the Hy-Mbody code.

Cluster name	INCAS
Nodes available	8
SoC	Rockchip RK3399 (28nm HKMG Process)
CPU	Six-Core ARM 64-bit processor (Dual-Core Cortex-A72 and Quad-Core Cortex-A53)
GPU	ARM Mali-T864 MP4 Quad-Core GPU
RAM	4GB Dual-Channel DDR3 (per node)
Network	1000Mbps Ethernet
Power	DC12V - 2A (per node)
Operating System	Ubuntu version 16.04 LTS
Compiler	gcc version 7.3.0
MPI	OpenMPI version 3.0.1
OpenCL	OpenCL 2.2
Job scheduler	SLURM version 17.11

High-end GPGPUs have already proved to speedup the solution of the direct N -body problem. In this work we aim to evaluate the performance of low-power embedded ARM GPU. We studied the best running time on ARM Cortex-A72x2 as the ratio over the best execution time taken by our ARM-optimized GPU implementation, as shown in the right panel of Figure 1. The ARM-optimized implementation is as fast as the dual-core implementation on the ARM Cortex-A72x2 using DP-arithmetic, as long as the GPU is kept fed with enough particles, while is almost three times faster using EX-precision.

5. Future development

ExaNeSt project is facing, among others, the challenge of the sustainable power consumption focusing on efficient hardware acceleration. For this reason, we are planning also to quantitatively measure the impact of our algorithms on energy consumption on SoC, shedding some light on their suitability for exascale applications. The findings from this research activity on ARM SoC are fundamental in order to also enhance our capabilities to exploit FPGAs for HPC, which in comparison to both CPUs and GPUs provide higher throughput-per-watt.

6. Conclusions

In light of our findings, embedded GPUs appear to be attractive from a performance perspective as soon as their double-precision compute capability increases. However, we demonstrated that the extended-precision approach can be a solution to supply enough power to execute scientific computation and benefit at maximum of the SoC devices.

SoC technology will play a fundamental role on future exascale heterogeneous platforms that will involve millions of specialized parallel compute units. Programmers will have to re-design their codes in order to fully exploit embedded accelerators, because of restricted hardware features compared to high-end GPGPUs.




Figure 1. Left: Running time comparison between ARM Cortex-A53x4 and Cortex-A72x2 CPUs for both DP-arithmetic (solid line) and EX-arithmetic (dashed line) as a function of the number of particles. Right: comparison of the time to solution between ARM Cortex-A72x2 CPU and Mali-T864 GPU for both DP-arithmetic (solid line) and EX-arithmetic (dashed line) as a function of the number of particles.

Acknowledgments. This work was carried out within the ExaNeSt (FET-HPC) project (grant no. 671553) and the ASTERICS project (grant no. 653477), funded by the European Union’s Horizon 2020 research and innovation program.

References

- Ammendola, R., Biagioni, A., Cretaro, P., Frezza, O., Cicero, F. L., Lonardo, A., Martinelli, M., Paolucci, P. S., Pastorelli, E., Simula, F., Vicini, P., Taffoni, G., Pascual, J. A., Navaridas, J., Lujan, M., Goodacre, J., Chrysos, N., & Katevenis, M. 2017, in 2017 Euromicro Conference on Digital System Design (DSD), 510
- Bertocco, S., Goz, D., Tornatore, L., & Taffoni, G. 2018, in INAF-OATs technical report, 222
- Bortolas, E., Gualandris, A., Dotti, M., Spera, M., & Mapelli, M. 2016, MNRAS, 461, 1023. 1606.06728
- Capuzzo-Dolcetta, R., & Spera, M. 2013, Computer Physics Communications, 184, 2528. 1304.1966
- Capuzzo-Dolcetta, R., Spera, M., & Punzo, D. 2013, J. Comput. Phys., 236, 580. 1207.2367
- Katevenis, M., Ammendola, R., Biagioni, A., Cretaro, P., Frezza, O., Lo Cicero, F., Lonardo, A., Martinelli, M., Paolucci, P., Pastorelli, E., Simula, F., Vicini, P., Taffoni, G., Pascual, J., Navaridas, J., Luján, M., Goodacre, J., Lietzow, B., Mouzakitidis, A., Chrysos, N., Marazakis, M., Gorlani, P., Cozzini, S., Brandino, G., Koutsourakis, P., Ruth, J., Zhang, Y., & Kersten, M. 2018, Microprocessors and Microsystems, 61, 58
- Katevenis, M., Chrysos, N., Marazakis, M., Mavroidis, I., Chaix, F., Kallimanis, N., Navaridas, J., Goodacre, J., Vicini, P., Biagioni, A., Paolucci, P. S., Lonardo, A., Pastorelli, E., Cicero, F. L., Ammendola, R., Hopson, P., Coates, P., Taffoni, G., Cozzini, S., Kersten, M., Zhang, Y., Sahuquillo, J., Lechago, S., Pinto, C., Lietzow, B., Everett, D., & Perna, G. 2016, in 2016 Euromicro Conference on Digital System Design (DSD), 60
- Spera, M. 2014, ArXiv e-prints. 1411.5234
- Spera, M., & Capuzzo-Dolcetta, R. 2015, ArXiv e-prints. 1501.01040
- Spera, M., Mapelli, M., & Bressan, A. 2015, MNRAS, 451, 4086. 1505.05201
- Thall, A. 2006, in ACM SIGGRAPH 2006 Research Posters (New York, NY, USA: ACM), SIGGRAPH '06. URL <http://doi.acm.org/10.1145/1179622.1179682>