Multicore Platform Efficiency Across Remote Sensing Applications

An Exploration of Execution Times and Energy Consumption

Ekaterina Tyutlyaeva, Alexander Moskovsky, Igor Odintsov, Sergey Konyukhov, Alexey Poyda, Mikhail Zhizhin

RSC Technologies

Introduction

A wide range of modern system architectures and platforms targeted for different algorithms and application areas is now available.
Even general-purpose systems have advantages in some computation areas and bottlenecks in another. Scientific applications on specific areas, on the other hand, have different requirements for CPU performance, scalability and power consumption.
The best practice now is algorithm/architecture co-exploration approach, where scientific problem requirements influence the hardware configuration; on the other hand, algorithm implementation is re-factored and optimized in accordance with the platform architectural features.
In this research, two typical modules used for multispectral nighttime satellite image processing are studied:

- *measurement of local perceived sharpness in visible band using Fourier transform;*
- *cross-correlation in moving window between visible and infrared bands.*

Both modules are optimized and studied on wide range of up-to-date testbeds, based on different architectures. Our testbeds include computational nodes based on Intel® Xeon® E5-2697A v4, Intel® Xeon® Gold 6148, Intel® Xeon Phi®, Texas Instruments Sitara AM5728 dual-core ARM Cortex-A15, and NVIDIA® JETSON™ TX2.

The study includes performance testing and energy consumption measurements. The results achieved can be used for assessing serviceability for multispectral nighttime satellite image processing by two key parameters: execution time and energy consumption.

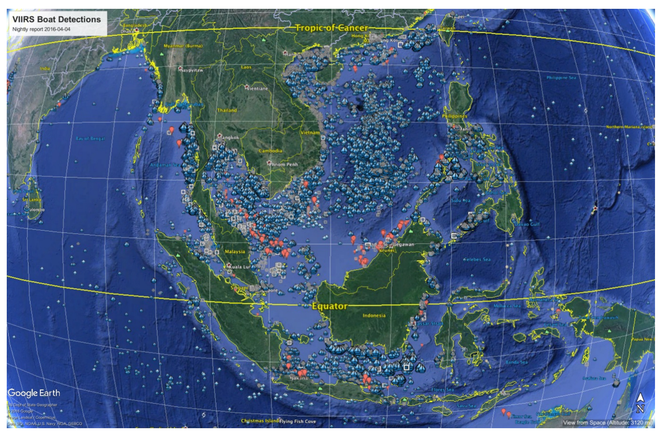

Figure 1: Lights from fishing boats detected from space in South Asia

Hardware

In the Table 1 codenames and specifications of the studied testbeds are listed.

Table 1: Testbeds: Technical Specifications

Codename	Processors (CPU + GPU)	Num Cores	Memory
Broadwell	2 x Intel® Xeon® E5-2697 v4	2 x 16	128 GB: 8 x Samsung 16GB DDR4/2133MHz
Skylake	2 x Intel® Xeon® Gold 6148	2 x 20	192 GB: 12 x 16 GiB DRAM DDR4/2400MHz
KNL	1 x Intel® Xeon Phi® 7250	68	Intel®16 GB MCDRAM + 6 x Micron 32 GB DDR4/2133 MHz
ARM	Used: Texas Instruments Sitara AM5728 dual-core ARM Cortex-A15 Not used: 2x PowerVR SGX544 3D GPU cores; Vivante GC320 2D GPU core	2	DDR3, 2GB
Jetson	ARM Cortex-A57 (quad-core)/2GHz + NVIDIA Denver2 (dual-core) + 256-core Pascal/1300MHz	4 + 2 + 256	8GB 128-bit LPDDR4/1866Mhz

The Sharpness Module

The studied module is the most computationally intensive part of an automatic system for detecting fishing boat lights from nighttime images of the VIIRS multispectral radiometer [1].

The system detects isolated bright spikes that are sharply visible on the sea's night surface. In the moonlight, the interference from clouds and lunar glint are taken into account as well.

In the current work, a module based on Direct Fourier Transformation in the moving window has been re-factored and studied.

Flowchart of the algorithm (see Fig. 2) includes:

- **Input** stage, where Day/Night Band (DNB) observations data image from the Suomi National Polar-orbiting Partnership (Suomi-NPP) Visible Infrared Imaging Radiometer Suite (VIIRS) are used. The data is stored locally in HDF5 format.
- **Stretch**, i.e. Logarithmic transformation of the brightness histogram
- Applying the Adaptive **Wiener** filter [2]
- Computing the **Spike Median Index (SMI)**
- Computing the **Sharpness Index (SI)** [3] in moving window $size_{blk}$ x $size_{blk}$.This routine is the most computationally intensive stage of the module. It repeatedly performs Direct Fourier Transforms and solves overdetermined real linear systems.
- **Output** stage, where the result gets stored locally in the ENVI format.

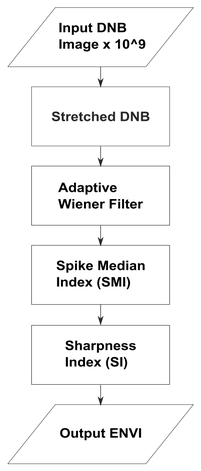

Figure 2: Studied module Flowchart

The Cross-Correlation Module

The Cross-correlation module calculates correlations between two spectral bands, visible and infrared. The main idea of the algorithm is validation of the detected sources in different spectral bands in moonlit conditions.

The validation is carried out through synchronously computing the linear Pearson's correlation between the corresponding moving windows in two spectral images. If the visible and infrared images are locally well correlated, it means that the signal in the visible comes from the moonlit clouds. If the local correlation is weak, it means that the visible signal comes from the sea surface.

Measurements

The testing procedure consisted of the energy consumption measurement and execution time measurement:

- Testing procedure included a series of 10 executions for each combination of input data set and input feature sets;
- Appropriate preparatory steps has been done prior to each execution, especially removing the results of previous computations and cleaning up the caches and swap.

The aggregate result is calculated as median value of measured results.

Energy measurements were conducted using equipment listed below:

- Running Average Power Limit (RAPL) [4] energy sensors available in recent Intel CPUs has been used for measuring energy consumption for Intel testbeds (Broadwell, Skylake, and KNL).
- Hantek DSO2000 Series USB Oscilloscope has been used for power measurements for ARM and NVidia testbeds. Electric current is measured in ampers every second of testing. Voltage is measured before execution of test series.

Execution time measured using PAPI Library (*PAPI_get_real_nsec()* function) on all testbeds.

Intel Testbed Features

The algorithm uses simple parallelization scheme:

- Number of MPI ranks is equal to number of physical cores;
- Each MPI rank processes a specific image.

This simple scheme allows us to use all available cores effectively while processing large archives of multispectral images.

For Intel architectures, the hybrid parallelization scheme is used. Number of OpenMP threads per MPI rank is equal to number of hyperthreads on studied testbed (2 for Broadwell and Skylake; 4 for KNL).

Such scheme reflects the architectural design of studied Intel testbeds and allows to avoid idle time of the processor resources during processing.

In order to achieve the highest performance on Intel testbeds, the vectorization features has been used (data alignment, pragmas, Intel MKL Library and processor-specific compilation options (-ipo -O3 -xMIC-AVX512 (for KNL)/-xCORE-AVX2 (for Broadwell) /-xCORE-AVX512 (for Skylake)).

The Figure 3 below demonstrated application scalability within the cores. According to the implementation features, number of used cores if equal to the number of processed images (1 core load – 1 image processed, 2 core load – 2 images processed etc...). So the ideal timeline case should be straight horizontal line (if all cores works independently and not used shared resources). But the collective work with memory and I/O slightly increases execution time with increased number of used cores and processed pictures. SizeBlk parameter shows the size of moving window for Sharpness Index calculation.

At the Figure 4 the energy consumption is shown. It can be noted, that energy consumption increases with number of used cores. Moreover, there is slight bump up in energy consumption for 2-socket testbeds (after 16 cores for Broadwell testbed and after 20 cores for Skylake testbed). The reason is that first physical CPU is fully loaded to this moment, and further MPI processes are pinned to second CPU.

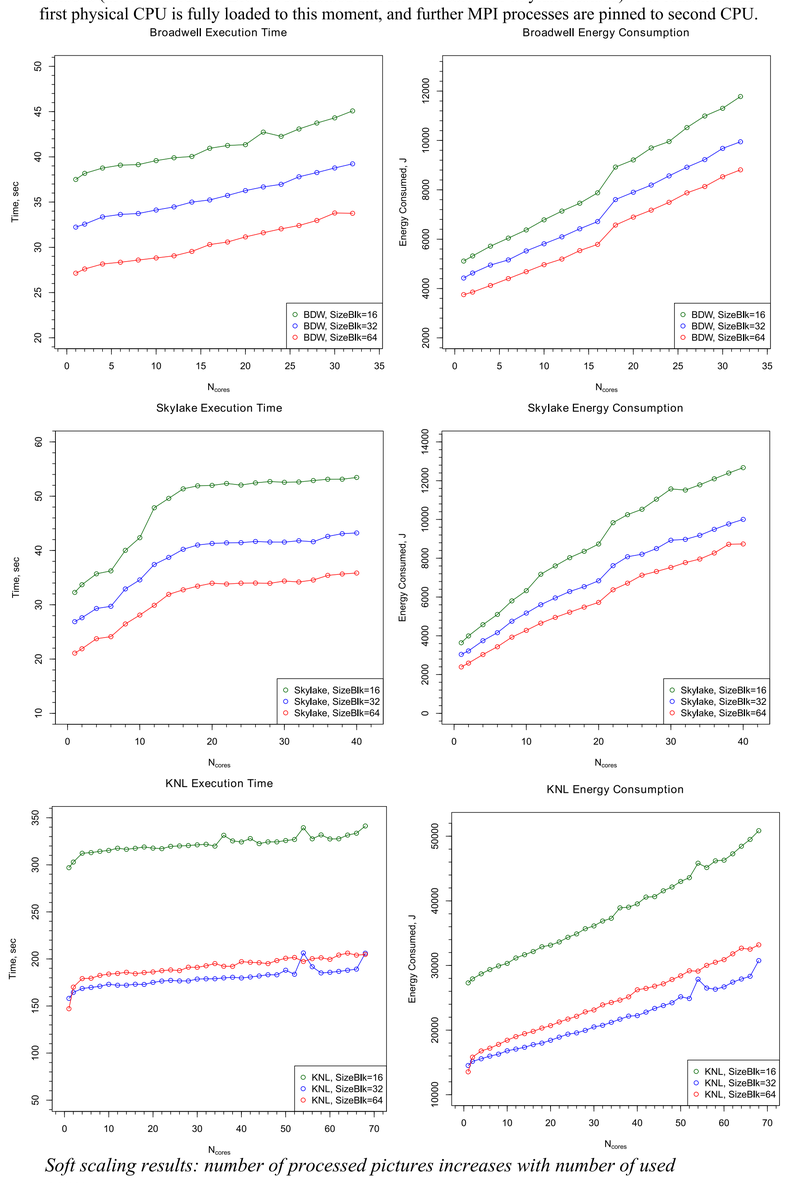

Soft scaling results: number of processed pictures increases with number of used cores.

Figure 3: Execution times

Figure 4: Energy Consumption

Parallel Efficiency

The experimental results of processing times and energy consumption measurements of the studied modules are listed at the Tables 2.1 (the Sharpness module testing results) and 3.1 (the Cross-Correlation module testing results).
The studied module implemented using C++ and MPI (except for Jetson testbed, where CUDA threads are used for parallelization). The numbers of MPI processes and CUDA threads are carefully adjusted according to the available number of cores and implementation for each architecture.

Table 2.1: The Sharpness Module Execution Statistics

Characterization	Broadwell	Skylake	KNL	ARM	Jetson
Images processed	32	40	68	2	1
Execution Time, sec	39.2	44.2	206.1	355.5	33
Energy Consumed, J	9 947	10 002	30 757	3 511	258

Table 2.2: The Sharpness Module Estimated Time To Process 1-year Archive

Characterization	Broadwell	Skylake	KNL	ARM	Jetson
Execution Time, **hours**	17.9	16.1	44.3	2 595	481.8
Energy Consumed, **kJ**	16 337	13 142	23 773	92 269	13 560

Thus, main testing results, such as execution time, energy consumed and number of pictures processed, varies considerably among studied testbeds. One feasible solution to compare results is an assessment of the testbed performance and energy consumption in the archive processing mode.
In is important to outline in this context that rapid technological progress in multispectral imaging area stimulates appearance of new methods and challenges in the analysis and interpretation of multispectral data sets. This, in turn, leads to re-processing of collected data for the last year(s). So the re-processing procedure is maintained systematically.
According to the current data, one Visible Infrared Imaging Radiometer Suite (VIIRS) day/night band (DNB) image corresponds to 5 min observations data.Therefore, the observation data archive per 1 year contains approximately 52560 DNB images. Tables 2.2 and 3.2 shows the estimated time and energy to process 1-year archive using the Sharpness and Cross-Correlation modules in conformity with the experimental results mentioned above.

Table 3.1: The Cross-Correlation Module Execution Statistics

Characterization	Broadwell	Skylake	KNL	ARM	Jetson
Images processed	32 (x 2)	40 (x 2)	68 (x 2)	2 (x 2)	1 (x 2)
Execution Time, sec	25.7	11.6	39	20	5
Energy Consumed, J	4 382	2 846	5 107	179	27

Table 3.2: The Cross-Correlation Module Estimated Time To Process 1-year Archive

Characterization	Broadwell	Skylake	KNL	ARM	Jetson
Execution Time, **hours**	11.7	4.2	8.3	146	73
Energy Consumed, **kJ**	7 197	3 739	3 947	4 704	1 419

NVidia Testbeds Features

Generally, the main difference of the modules implementation for Jetson testbed from others is usage of CUDA (V8.0.62) threads instead of MPI. CUDA threads utilizes GPU resources effectively, so MPI layer not used in this implementation.

The preparation steps, such as data input, output, applying the Wiener filter, stretch and computing the Spike Median Index (SMI) are performed on CPU. The most computationally intensive steps in both algorithms (computing the Sharpness Index and Correlation) implemented using CUDA, and performed using GPU cores only.

The different NVPModels have been studied, which defines maximum frequency of the CPU and GPU cores being used (see Fig 5-6).

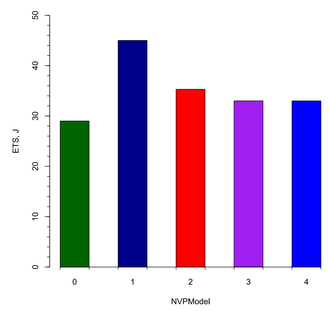

Figure 5: Execution times

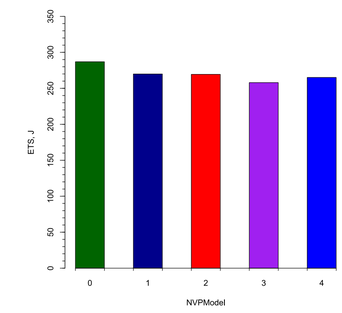

Figure 6: Energy Consumption

The best result in terms of energy consumption has been obtained using the NVPModel number 3, so-called Max-P ARM. Max-P ARM model provides 1.12 Ghz GPU Frequency and uses 4 cores of ARM A57 with 2.0 GHz frequency. Denver 2 cores is not used on this model. The best result in terms of general performance has been obtained using NVPModel 0, Max-N. This model uses all available resources with highest frequency (2 cores of Denver 2 with 2.0 GHz, 4 cores of ARM a57 with 2.0 GHz and all GPU cores with 1.30 GHz). But this model shows worst energy consumption.

Conclusion

This poster presents a performance and energy consumption results while running multispectral nighttime satellite image processing modules on different architectures.

The different architecture-dependent implementation of the same modules has been developed *for the study*. Necessary optimizations for Intel architecture, including compiler-assisted vectorization and hybrid (MPI+OpenMP) parallelization mode have been done. Alternative implementation with CUDA threads has been developed for effectively utilizing Jetson testbed resources. Additionally, the NVPModels on Jetson testbed impact on time and energy consumption is shown.

According to the preliminary investigation, the Intel® Xeon® Gold 6148 –based server (Skylake testbed) demonstrated the best result *r*egarding computation performance and energy consumption criteria.This conclusion is based on the following main observations:

- The best result for the Sharpness module achieved on a Skylake testbed; According to the obtained results, Skylake could process 40 pictures per 44 sec with 10 kJ energy consumed. It is the best result among studied testbeds.
- The fastest result for the Cross-Correlation module also achieved on a Skylake testbed. It could compute cross-correlation between 40 pairs of images in 11.6 seconds. But the most energy efficient result for this module achieved in Jetson testbed. It requires only 27 J per pair of images processing.

As future work, we plan to analyze in details the performance and energy consumption of other types of architectures, including Russian VLIW Elbrus processors. Furthermore, we are going to continue work with our ARM testbed to accelerate the results by using GPU resources that are idle in the current implementation.

References:

1. Elvidge, Christopher D. and Zhizhin, Mikhail and Baugh, Kimberly and Hsu, Feng-Chi: Automatic Boat Identification System for VIIRS Low Light Imaging Data In: Remote Sensing Journal, Vol.7, Num. 3, pp. 3020--3036 ISSN 2072-4292, doi: 10.3390/rs70303020 (2015)

2. Lim, Jae S., Two-Dimensional Signal and Image Processing, Englewood Cliffs, NJ, Prentice Hall, 1990, p. 548, equations 9.26, 9.27, and 9.29

3. Vu, C.T.; Phan, T.D.; Chandler, D.M. S3: A spectral and spatial measure of local perceived sharpness in natural images. IEEE Trans. Image Process. 2012, 21, 934–945

4. Rotem, E., Naveh, A., Ananthakrishnan, A., Weissmann, E., Rajwan, D.: Power-Management Architecture of the Intel Microarchitecture Code-Named Sandy Bridge. IEEE Micro, vol. 32, no. 2, pp. 20-27, March-April 2012. doi: 10.1109/MM.2012.12

Acknowledgments

This research was supported by the grant from the Russian Ministry of Education and Science (The government contract 14.607.21.0165, unique identifier RFMEFI60716X0165).