Assessing the Performance-Power Potential of Graphics Processors for Spectral Unmixing

Sergio Sánchez, Germán León, Antonio Plaza, and Enrique S. Quintana-Ortí

Abstract

Bla, bla, bla...

Keywords

Graphics processors (GPUs), hyperspectral imaging, energy consumption, high performance computing.

I. Introduction

Dummy entry [1]

II. Hyperspec chain

III. Parallelization on GPUs

IV. Experimental Results

We open this section with a description of the hardware setup (i.e., computational resources and power measurement device) as well as the hyperspectral data testbeds that were employed in the experimental study. The major contribution of this paper, namely the analysis of the performance-power trade-off of current hardware accelerators when applied to process all stages of the two complete spectral unmixing chains, follows next.

A. Hardware configurations

The experimental study was performed on three different platforms, equipped with some of the most recent graphics technology from NVIDIA and state-of-the-art multicore processors from Intel or ARM:

- `Carma`: An NVIDIA Quadro 1000M GPU connected to an ARM Cortex A9 multicore processor (4 cores at 1.3 GHz) with 2 GB of DDR3L RAM.
- `Fermi`: An NVIDIA GeForce GTX 480 ("Fermi") graphics card connected to a single Intel Xeon i7-3770K (ivy-bridge) processor (4 cores at 3.5 GHz) and 16 GB of DDR3 RAM.
- `Kepler`: An NVIDIA Tesla K20c ("Kepler") graphics card connected to a single Intel Xeon i7-3930K (sandy-bridge e) processor (6 cores at 3.2 GHz) and 24 GB of DDR3 RAM.

These three platforms represent two extremes of the spectrum in hardware acceleration using graphics processors. On the one side, there is the NVIDIA Q1000, a 96-core GPU with 2 GB of DDR3 RAM, integrated into a low-power board (Carma development kit) together with a general-purpose processor from ARM. This system has no disk or any other relevant devices attached to it, drawing a mere 12.5 Watts (on average) when idle, i.e., when doing nothing.

On the other side, Fermi and Kepler correspond the last two generations of high-throughput accelerators from NVIDIA and, in both cases, the tested cards integrate a graphics processor (Tesla T10 in the Tesla C1060 and GK110 in the Tesla K20) with a high number of cores (240 and 2,496, respectively) and GDDR5 RAM (4 and 5 GB, respectively). To operate, these boards have to be attached via PCI-e to a server with at least one general-purpose processor in charge of controlling the GPU, playing the same role as the ARM A9 processor in the `Carma` system. However, these are regular servers that contain not only the described components, but also one or more disks and Ethernet ports (the latter embedded into the mainboard). The net result is a much higher idle power consumption: on average, 97.9 and 102.6 Watts, respectively, for the platforms where the Fermi and Kepler boards are attached.

Tuned implementations of the numerical linear algebra operations that appear in the algorithms were obtained from recent releases of Intel MKL (version XX for the two Ivy-Bridge processors) and NVIDIA CUBLAS (version XX for all three GPUs). Currently there exist no tuned implementation of analogous kernels for the ARM, so that we had to rely, in this case, in the legacy implementation of these routines available at netlib[1].

S. Sánchez and A. Plaza are with Hyperspectral Computing Laboratory (HyperComp), Department of Technology of Computers and Communications, University of Extremadura, 10.071–Cáceres, Spain (e-mail: mitisnio@gmail.com,aplaza@unex.es). G. León and E. S. Quintana-Ortí are with Depto. de Ingeniería y Ciencia de Computadores, Universidad Jaume I, 12.071–Castellón, Spain (e-mails: {leon,quintana}@uji.es).

[1] http://www.netlib/org.

Chain #1	VD	PCA	NFINDR	LSU	
System	Time				Total
Carma	1.23	1.35	1.53	0.40	4.51
Fermi	0.18	0.11	0.08	0.05	0.42
Kepler	0.16	0.08	0.07	0.04	0.35
	Avg. power				Avg.
Carma	28.3	19.5	19.6	22.8	22.2
Fermi	248.3	213.9	208.8	218.2	228.2
Kepler	251.5	255.9	204.4	211.4	238.5
	Max. power				Max.
Carma	32.5	22.5	19.6	23.1	32.50
Fermi	322	232	218	236	322.00
Kepler	273	264	208	214	273.00
	Energy				Total
Carma	34.8	26.3	30.0	9.1	100.24
Fermi	44.7	23.5	16.7	10.9	95.84
Kepler	40.2	20.5	14.3	8.5	83.48
	Net Energy				Total
Carma	19.4	9.5	10.9	4.1	43.87
Fermi	27.1	12.8	8.9	6.0	54.70
Kepler	38.2	19.5	13.4	8.0	79.10

TABLE I
CUPRITE. CHAIN 1

In order to measure power, we employed a *WattsUp? Pro .Net* wattmeter. This device is plugged into to the cable that connects the electrical socket to the power supply unit (PSU), and reports external AC for the full platform, with a sampling rate of 1 Hz, an accuracy of $\pm 1\%$ and a resolution of 0.1 Watts. We warmed up the platforms by executing each stage of the chains repeatedly during XXs before the sampling was initiated for that particular stage. Power measures were then continuously recorded while the test (i.e., the stage of the chain) was running during XX more seconds, and power was averaged over this period and multiplied by the execution time of a single execution of the stage to obtain its (total) energy consumption. A fairer metric in order to compare the energy efficiency of the different platforms is the net energy consumption, which was obtained by subtracting the product of idle power and time from the energy consumption. This measure better reflects the energy necessary to perform the work, cancelling the effect of unnecesary components (e.g., the disk) on the power draw. Hereafter, execution time is reported in seconds (s), power in Watts (W) and energy in Joules (J=W·s),

B. Hyperspectral scenes

We leveraged the same two datasets employed for the evaluation in [**?**], so that we can later compare the performance-power ratios of the GPU-equipped systems to those of the multicore architectures that were analyzed in that paper. The first case corresponds to the Airborne Visible Infra-Red Imaging Spectrometer (AVIRIS) `Cuprite` scenario, and is a well-known benchmark for the evaluation of spectral unmixing methods, with online data[2]. The specific image corresponds to an m=350×350-pixel subset of sector `f970619t01p02_02_sc03.a.rfl`, features n=188 spectral bands between 0.4 and 2.5 μm, and requires an execution time below 2.98 s for real-time processing.

The second image was collected over the World Trace Center (`WTC`) after the attacks on September 11, 2011. It consists of m=512×614 pixels and n=224 bands, which corresponds to the standard data cube size recorded by AVIRIS. Compared with the previous scenario, processing this image in real-time requires a longer time, 5.09 s in this case.

C. Power-performance trade-off of the chains

D. Comparative study against multicore architectures

In [**?**], we carried out a complete study of the performance-power balance of 5 different multicore processors using two efficient spectral unmixing methods for identifying the endmembers and estimating their fractional aboundances in hyperspectral images: the Orthogonal Subspace Projection via Gram-Schmidt [2] and ISRA. Three of these systems were low-power processors (Intel Atom, Texas Instruments DSP and ARM Cortex) while the remaining two corresponded to conventional architectures for the desktop and server segments (Intel Xeon and AMD Opteron). That analysis revealed the AMD and the Texas Instruments as the clear winners from the viewpoints, respectively, of performance and energy efficiency. For comparison purposes, we reconsider next the evaluation of ISRA on these two optimal processors and the Intel Xeon:

- `DSP`: A single Texas Instrument C6678 digital signal processor (8 cores at 1.0 GHz) with 512 Mbytes of DDR3 RAM.
- `Xeon`: Two Intel Xeon E5504 processors (4 cores per processor, at 2.0 GHz) with 32 GB of DDR3 RAM.

[2]http://aviris.jpl.nasa.gov/freedata.

Chain #2	HYSIME	SPCA	NFINDR	ISRA	
System	Time				Total
Carma	11.28	1.45	0.91	5.06	18.70
Fermi	0.81	0.12	0.06	0.92	1.91
Kepler	0.63	0.10	0.05	0.87	1.65
	Avg. power				Avg.
Carma	20.6	19.9	20.0	42.6	26.47
Fermi	226.6	220.0	212.9	323.1	272.24
Kepler	215.5	199.3	203.1	261.1	238.19
	Max. power				Max.
Carma	32.8	28	22.5	45.6	45.60
Fermi	281	243	218	336	336.00
Kepler	237	218	205	270	270.00
	Energy				Total
Carma	232.4	28.9	18.2	215.6	495.0
Fermi	183.5	26.4	12.8	297.3	520.0
Kepler	135.8	19.9	10.2	227.2	393.0
	Net Energy				Total
Carma	91.4	10.7	6.8	152.3	261.2
Fermi	104.2	14.6	6.9	207.1	332.9
Kepler	78.1	11.7	5.5	191.8	287.1

TABLE II
CUPRITE. CHAIN 2

Chain #1	VD	PCA	NFINDR	LSU	
System	Time				Total
Carma	3.55	4.05	12.24	1.48	21.32
Fermi	0.44	0.32	0.44	0.18	1.38
Kepler	0.34	0.24	0.33	0.14	1.05
	Avg. power				Avg.
Carma	28.9	20	18.8	23.7	21.0
Fermi	248.8	226.1	211.9	225.8	228.8
Kepler	240.1	234.5	205.9	214.5	224.7
	Max. power				Max.
Carma	42.1	32.5	29.6	26.1	42.10
Fermi	277	273	228	245	277
Kepler	253	251	232	222	253.00
	Energy				Total
Carma	102.6	81.0	230.1	35.1	448.78
Fermi	109.5	72.4	93.2	40.6	315.70
Kepler	81.6	56.3	67.9	30.0	235.89
	Net Energy				Total
Carma	58.2	30.4	77.1	16.6	182.28
Fermi	66.4	41.0	50.1	23.0	180.55
Kepler	46.7	31.6	34.1	15.7	128.1

TABLE III
SUBSETWTC. CHAIN 1

- Opteron: Two AMD Opteron 6128 processors (8 cores per processor, at 2.0 GHz) with 24 Gbytes of DDR3 RAM.

Tuned implementations of BLAS were employed in all cases; see [**?**] for details. The datasets and the power measurement methodology are coherent with those employed for evaluation of the GPU-equipped systems above. **Check the conditions for the experiments with ISRA? Same number of iterations in both cases?**

V. CONCLUDING REMARKS

ACKNOWLEDGMENTS

The author from the Universidad Jaume I was supported by project CICYT TIN2011-23283 and FEDER. Funding from the Spanish Ministry (CEOS-SPAIN project, reference AYA2011-29334-C02-02) is also gratefully acknowledged.

REFERENCES

[1] G. H. Golub and C. F. V. Loan, *Matrix Computations*, 3rd ed. The Johns Hopkins University Press, 1996.
[2] J. C. Harsanyi and C.-I. Chang, "Hyperspectral image classification and dimensionality reduction: An orthogonal subspace projection," *IEEE T. Geoscience and Remote Sensing*, vol. 32, no. 4, pp. 779–785.

Chain #2	HYSIME	SPCA	NFINDR	ISRA	
System	Time				Total
Carma	33.5	4.57	4.74	13.85	56.66
Fermi	2.62	0.37	0.28	2.47	5.74
Kepler	2.62	0.3	0.23	1.61	4.76
	Avg. power				Avg.
Carma	20.6	20.8	19.8	45.5	26.64
Fermi	219.1	224.6	213.5	347.6	274.48
Kepler	198.4	203.9	205.8	280.8	226.97
	Max. power				Max.
Carma	41.7	32.1	21.3	48.2	48.20
Fermi	315	254	247	358	315
Kepler	226	214	210	287	287.00
	Energy				Total
Carma	690.1	95.1	93.9	630.2	1509.2
Fermi	574.0	83.1	59.8	858.6	1575.5
Kepler	519.8	61.2	47.3	452.1	1080.4
	Net Energy				Total
Carma	271.4	37.9	34.6	457.1	800.9
Fermi	317.4	46.9	32.4	616.7	1013.3
Kepler	305.1	36.6	25.5	394.4	761.6

TABLE IV
SUBSETWTC. CHAIN 2

Scenario	Cuprite		WTC	
System	Time	Net Energy	Time	Net Energy
Carma	3.56	109.6	9.98	328.3
Fermi	0.66	143.0	1.73	434.7
Kepler	0.55	89.1	1.12	197.4
DSP	4.94	10.08	12.73	25.46
Xeon	0.69	79.41	2.42	280.86
Opteron	0.34	82.21	1.79	407.85

TABLE V
PERFORMANCE AND NET ENERGY OF THE GPU-EQUIPPED SYSTEMS AND THREE MULTICORE PROCESSORS ON ISRA.