Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 113 (2017) 496–501

Procedia Computer Science

www.elsevier.com/locate/procedia

International Workshop on Data Mining on IoT Systems (DaMIS17)

A virtualized software based on the NVIDIA cuFFT library for image denoising: performance analysis

Ardelio Galletti[a], Livia Marcellino[a,*], Raffaele Montella[a], Vincenzo Santopietro[a], Sokol Kosta[b]

[a]*Department of Science and Technology, University of Naples Parthenope, Centro Direzionale Isola C4, 80143, Naples Italy*

[b]*Aalborg University Copenhagen The Faculty of Engineering and Science The Technical Faculty of IT and Design Department of Electronic Systems Infrastructure, Services and Entrepreneurship Center for Communication, Media and Information Technologies*

Abstract

Generic Virtualization Service (GVirtuS) is a new solution for enabling GPGPU on Virtual Machines or low powered devices. This paper focuses on the performance analysis that can be obtained using a GPGPU virtualized software. Recently, GVirtuS has been extended in order to support CUDA ancillary libraries with good results. Here, our aim is to analyze the applicability of this powerful tool to a real problem, which uses the NVIDIA cuFFT library. As case study we consider a simple denoising algorithm, implementing a virtualized GPU-parallel software based on the convolution theorem in order to perform the noise removal procedure in the frequency domain. We report some preliminary tests in both physical and virtualized environments to study and analyze the potential scalability of such an algorithm.

Peer-review under responsibility of the Conference Program Chairs.

Keywords: Hybrid infrastructures, Virtualization, High Performance Computing, FFT, GPGPU, image denoising;

1. Introduction

Analysis and processing of very large datasets, or big data, poses a significant challenge. Massive datasets are collected, studied and interpolated in numerous domains, from mathematical and engineering sciences [5,15,11,12,13,16] to social networks [19], from biomolecular research, commerce and security [23,31] to IoT applications [3,4,7,8,9,10]. In order to obtain the best performance in analyzing big data, the use of computing machines with high processing power became necessary [20]. Virtualization technologies are currently widely deployed, as their use yields important benefits such as resource sharing, process isolation, and reduced management costs [1,21,26]. Thus, it is straightforward that the usage of virtual machines (VMs) in HPC is an active area of research. From a computational point of view, the hierarchical and heterogeneous high performance computing paradigm, that emerged in the last decade, delivers enough power to process spatial big data leveraging on massive multicore CPUs, general purpose graphic processing units (GPGPUs) and, recently, field-programmable generic arrays (FPGAs) and supported by solid state storage skyrocketing the long

* Corresponding author. Tel.: +39-081-5476676.
E-mail address: livia.marcellino@uniparthenope.it

Peer-review under responsibility of the Conference Program Chairs.
10.1016/j.procs.2017.08.310

term memory access performance[25,30]. In order to mitigate the computational infrastructures total cost of ownership, and accelerate the tools availability to a large users footprint, the use of high performance cloud computing resources pushes to minimize the big data based products time to market. While CPU virtualization and elastic storage is a common practice in public, private and hybrid clouds[22,28], the same techniques are not widely available due to the fact that often the accelerators exploit closed technologies.
In this paper we demonstrate how is possible to use the GPGPU resource leveraging on the RAPID GVirtuS[25,27] GPGPU virtualization service presenting a specific case related to the image denoising process. As is well-known, image denoising algorithms have a very high computational cost, especially when dealing with large-scale images. An effective approach to solve this problem is to use a parallel algorithm. There are many research efforts in this field, using different parallel computing architectures[2,6,17], especially on hybrid environment[14].
Here we present a denoising algorithm, based on a frequency domain convolution leveraging the CUDA environment, which exploits the computational power of the NVIDIA cuFFT library[1]. Our aim is to highlight the benefits in using virtualization without loss of efficiency. We used this simple case study in order to analyze how the overall virtualization cost affects the execution time. To be specific, starting by the Fourier Transform (FT) approach, we use the *convolution theorem* to perform the image denoising process as a pointwise product of FTs on GPU device. Therefore, we report some tests in both physical and virtualized environments and compare them.

The rest of the paper is organized as follows: in section 2 we recall the image denoising problem, the FT decomposition and the convolution theorem; in section 3 we present the cuFFT library of CUDA and how to use it efficiently: moreover a description of the GPU-parallel algorithm and of the virtualization approach with GVirtuS is also provided; the evaluation is carried out in section 4. Finally, in section 5 we conclude the paper.

2. The image denoising in frequency domain

Reconstruction of a signal from a noisy one is a well-known inverse problem. From a mathematical point of view a noisy image can be defined as a piecewise function $f(x, y) = C[u(x, y)]$, where $u(x, y)$ is the image function (non-corrupted by noise) and C is a function that defines the noise we wish to eliminate. Estimation of an unknown signal u from the available noisy data f, is an ill-posed problem that involves to find the noisy-free image, by preserving useful information. Noise in digital image may arise during the acquisition step, caused by poor illumination and/or high temperature. It can be either additive or multiplicative. Some noise models are also called *white*: the noise is spatially uncorrelated and identically distributed. Another often observed noise model is the so-called *salt and pepper* noise, which can be reduced with a median filter. In this work we deal with additive white Gaussian noise, defined as $p_G(u) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(u-\mu)^2}{2\sigma^2}}$ where u represents the image gray level and μ and σ are the Gaussian distribution mean and standard deviation, respectively. Generally, the simplest method to perform image denoising, in the spatial domain, is based on a convolution operation between the function f and the Gaussian function. The related algorithm uses a finite size mask (or kernel) which scannes across the image. The value of each pixel on the output image is the weighted sum of the input pixels within a window where the weights are the values of the mask. This algorithm has a high computational complexity, therefore this naive approach can become very slow for big inputs (requiring too much time for large images). A more efficient approach exploits the computational power of the Fourier Transform (FT) and the related frequency domain operations.

The Fourier transform of a function of time itself is a complex-valued function of frequency, defined as follows:

$$F(\eta) = \int_{-\infty}^{+\infty} f(t) e^{-2\pi i \eta t} dt.$$

The numerical approach for computing FT approximates a non periodic function as sum of a certain number of sine and cosine functions. The component frequencies of the sine a cosine functions that spread across the spectrum are represented as peaks in the frequency domain. Starting by the FT concept, the image denoising process in the frequency domain uses the *Convolution Theorem* for which: *under suitable conditions the FT of a convolution is the*

[1] https://developer.nvidia.com/cufft