

## Note: Quasi-real-time analysis of dynamic near field scattering data using a graphics processing unit

G. Cerchiarì,<sup>1,2</sup> F. Croccolo,<sup>1,a)</sup> F. Cardinaux,<sup>1</sup> and F. Scheffold<sup>1,2</sup>

<sup>1</sup>Department of Physics, University of Fribourg, Ch. Du Musée 3, 1700 Fribourg, Switzerland

<sup>2</sup>Fribourg Center for Nanomaterials, University of Fribourg, 1700 Fribourg, Switzerland

We present an implementation of the analysis of dynamic near field scattering (NFS) data using a graphics processing unit. We introduce an optimized data management scheme thereby limiting the number of operations required. Overall, we reduce the processing time from hours to minutes, for typical experimental conditions. Previously the limiting step in such experiments, the processing time is now comparable to the data acquisition time. Our approach is applicable to various dynamic NFS methods, including shadowgraph, Schlieren and differential dynamic microscopy.

Light scattering techniques are powerful tools to characterize the structure and dynamics of simple and complex fluids.<sup>1</sup> Structural information about the system is contained in the angular dependence of the scattered intensity, while its temporal fluctuations are characteristic of the dynamics. Traditional static and dynamic light scattering techniques analyze scattering in the far field, but the same information can also be obtained under near field conditions.<sup>2,3</sup> Near field scattering (NFS) techniques are nowadays available in different layouts using radiation sources ranging from coherent<sup>4</sup> and partially coherent light,<sup>5-7</sup> to white light<sup>8</sup> or x-rays<sup>9</sup> with different experimental constraints. The ease of access to extremely low scattering angles or the possibility to perform spatially resolved experiments are interesting features of NFS that have allowed studies on liquids,<sup>2-5</sup> colloidal dispersions,<sup>10</sup> or biological samples.<sup>11</sup> A drawback of NFS is the heavy computational load involved in the evaluation of data.<sup>4,8,12</sup> Typically, a sequence of thousands of mega-pixel images is necessary to evaluate both the static and dynamic spectra with sufficient accuracy.

In this study, we present an implementation of dynamic NFS using a graphics processing unit (GPU). The highly parallel architecture of GPUs is optimized to execute a certain operation (*kernel*) on multiple data elements (*stream*) simultaneously. It results in an acceleration of processing originally used to handle computer graphics. The use of stream processing for non-graphical applications is expanding due to the availability of high-level programming interfaces<sup>13,14</sup> and has been exploited in applications ranging from computational biology, cryptography, computational chemistry, biomedical imaging, or optics.<sup>13-17</sup> Very recently Lu *et al.* have reported on the application of an algorithm similar to ours for the rapid analysis of fluorescence microscopy images of liquid colloidal suspensions.<sup>17</sup> Here we show that for dynamic NFS a remarkable speedup can be achieved when processing the data using a GPU instead of a central processing unit (CPU). Moreover we introduce an effective method to manage the

graphics random access memory (G-RAM) in order to optimize the number of operations and data transfers in the treatment of NFS data. Overall, we are able to reduce the processing time from hours to minutes, for our typical experimental settings. Previously the limiting step in such experiments, the time required for data processing is now comparable to the actual data acquisition time.

As a representative test experiment, we present NFS results for concentration non-equilibrium fluctuations<sup>18</sup> in a binary fluid subjected to a thermal stress.<sup>4,5,19</sup> The sample is a mixture of tetrahydronaphthalene and n-dodecane ( $c = 50\%$  w/w) contained in a sapphire flat cell with a gap thickness of 1.3 mm. Non-equilibrium fluctuations are driven by a temperature difference of 20 °C ( $T_{avg} = 25$  °C) across the sample. The structure and dynamics of the concentration fluctuations are characterized using a heterodyne NFS instrument having a shadowgraph detection layout.<sup>5,6</sup> A low coherence light source emitting at  $680 \pm 10$  nm (Super Lumen Diodes, Broad Lighter S680) is expanded to a diameter of 20 mm and collimated. The cell is centered in the optical axis with its surface normal to the axis. Finally, a charge-coupled device (Vosskühler, CCD4000) placed at a distance of 26 cm away from the cell acquires intensity-images at rate of 4 Hz. Raw images are composed of a superposition of static contributions (dominated by the primary beam) and a time dependent contribution arising from the scattering due to concentration fluctuations. A sample image is presented in Fig. 1(a) wherein the beam profile as well as residual interferences due to reflections and stray light are visible. In a scattering experiment, the structure and the dynamics of the system under study are characterized by the wave vector dependence of the scattered intensity  $I_s(\mathbf{q})$  and the intermediate scattering function (ISF)  $f(\mathbf{q}, \tau)$ , respectively.<sup>4,5</sup> An elegant way to obtain both quantities from NFS experiments is to use the double-frame differential analysis.<sup>4,5</sup> In a first step, images are normalized by their spatial average values and differences  $\Delta i_m(\mathbf{x}, \tau) = i_m(\mathbf{x}, t) - i_m(\mathbf{x}, t + \tau)$  between normalized images separated by a time delay  $\tau$  are computed. The resulting spatial distributions of intensity show a homogeneous speckle pattern.<sup>18</sup> This is illustrated in Fig. 1 for a time delay of (b)  $\tau = 0.25$  s, (c) and

<sup>a)</sup>Electronic mail: fabrizio.croccolo@unifr.ch.

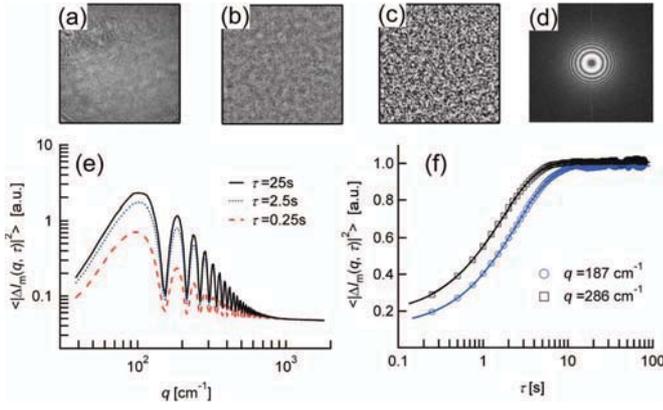


FIG. 1. Upper panels: (a) Shadowgraph image, image differences with time delay of (b) 0.25 s and (c) 25 s, (d) spatial power spectrum of (c). Lower panels: Power spectra as a function of (e) wave vector  $q$  and (f) time delay  $\tau$ .

25 s. The scattered intensity and the ISFs are obtained by analyzing the time delay dependence of power spectra of  $\Delta I_m(\mathbf{x}, \tau)$ :

$$\langle |\Delta I_m(\mathbf{q}, \tau)|^2 \rangle = 2a \{ T(\mathbf{q}) I_s(\mathbf{q}) [1 - f(\mathbf{q}, \tau)] + B(\mathbf{q}) \}, \quad (1)$$

where  $\Delta I_m(\mathbf{q}, \tau)$  is the spatial Fourier transform of  $\Delta I_m(\mathbf{x}, \tau)$ ,  $a$  is a renormalization constant,  $T(\mathbf{q})$  is the transfer function of the imaging optics,<sup>5,6</sup> and  $B(\mathbf{q})$  is the noise background.<sup>4,5</sup> A 2D power spectrum for  $\tau = 25$  s is presented in Fig. 1(d). For the usual case of stationary dynamics, ensemble averaging  $\langle \dots \rangle$  in (1) can be performed over the measurement time. A selection of azimuthally averaged power spectra is presented in Figs. 1(e) and 1(f).

The processing of NFS data using the double-frame differential analysis of Eq. (1) requires performing a limited number of consecutive operations on large data sets. As we will show a remarkable increase in processing speed can be achieved using a parallel processing approach. An equally important aspect of a fast processing is to avoid redundant calculations. For a number  $N$  of images, the direct calculation of the power spectra for all the available  $N - 1$  values of  $\tau$  involves  $N(N - 1)/2$  iterations. However, some redundant calculations can be avoided. In particular, one can take advantage of the linearity of the Fourier transform by evaluating FFT's on single images instead of image differences, thus reducing their number to  $N$ . A further speedup can be obtained by dividing the spatial Fourier transform of each image by its zero-frequency value for image normalization. Finally, only the square moduli of FFT differences are evaluated  $N(N - 1)/2$  times. These processing steps require images to be read from the motherboard random access memory (RAM), transfer it to the G-RAM for the GPU-processing, FFT-transform, calculate differences, and finally calculate and store the average for any given  $\Delta N$  in the memory. For the case of CPU-processing the images are read directly from the RAM. However for the case of GPU computing data has to be transferred from the RAM to the limited G-RAM. The amount of available G-RAM therefore sets a ceiling up to which GPU-processing provides full computational benefits. As a matter of fact, if the number of images to be analyzed

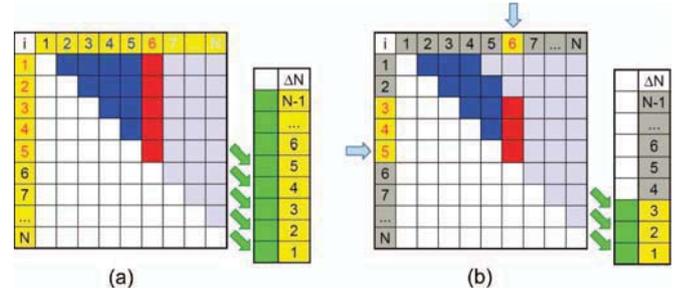


FIG. 2. Organization schemes for the memory in the two distinct cases (a) when all the data can be stored in the G-RAM, and (b) when the data exceeds the G-RAM capacity. The vertical axis denotes image FFT's, the horizontal one denotes the image FFT to compute differences for a given  $\Delta N$ , and the green column contains the calculated power spectra for corresponding time delays  $\tau = \Delta N \times 0.25$  s. Blue elements of the matrix stand for already calculated differences, while red elements are calculated at the current iteration. Cyan arrows indicate images that are uploaded at the current iteration.

exceeds the G-RAM capacity then data has to be transferred multiple times leading to a performance decrease. In the latter case data transfer can be optimized by organizing the memory in a first-in-first-out (FIFO) queue. In order to obtain a “general-purpose” code able to treat any amount of images of any size, no *a priori* optimization has been built into the code. Here we chose to optimize memory allocation prior to starting the calculations. For a memory capacity for only  $N^* < (N - 1)$  images we can write  $N - 1 = \alpha N^* + \beta$ , where  $\alpha, \beta$  are positive integer numbers and  $\beta < N^*$ . The program needs to calculate all the possible FFT differences step by step. To this end  $N^*$  images are uploaded into the FIFO, one image is uploaded into the “current image area” (see Fig. 2), all possible differences are calculated and the results are averaged and stored in the “differences area”. These steps are repeated until the first triangle of matrix elements has been processed. Next one diagonal is filled by subsequently replacing one FIFO image and uploading one “current image”. This procedure is repeated for all the  $\alpha$  diagonals and subsequently the values for the remaining upper-right triangle are computed. For simplicity we discuss only the case  $\beta = 0$  and we find that the number of image transfer operations to the G-RAM is  $N_{FFT} = N \cdot (2\alpha - 3 + 2/\alpha) \rightarrow 2\alpha N$  for  $\alpha \gg 1$  and thus much smaller compared to the total number of differences computed  $\cong N^2/2$ . Our reference motherboard is equipped with an 8-core CPU (INTEL, XEON X3440 at 2.53 GHz) and a GPU-board (NVIDIA, Tesla C2050). The double-frame differential analysis is implemented both in C++/CUDA for the GPU-based code and in C++ for the CPU-based one. The performances of the two versions are compared in order to quantify the speed increase due to the GPU. The total execution time of the program is the key parameter for evaluating its performance. However, we also analyzed separate timings for different operations.<sup>20</sup> With both codes we processed different sets of images with different sizes, spanning from 32 up to 2048 images of  $32 \times 32$  up to  $4 \text{ k} \times 4 \text{ k}$  pixels.

We now focus our attention to the ratio between the CPU and the GPU processing time  $t_{CPU}/t_{GPU}$ . The result of this analysis is shown in Fig. 3 as a contour plot of this ratio (z axis) as a function of image size (x axis) and number

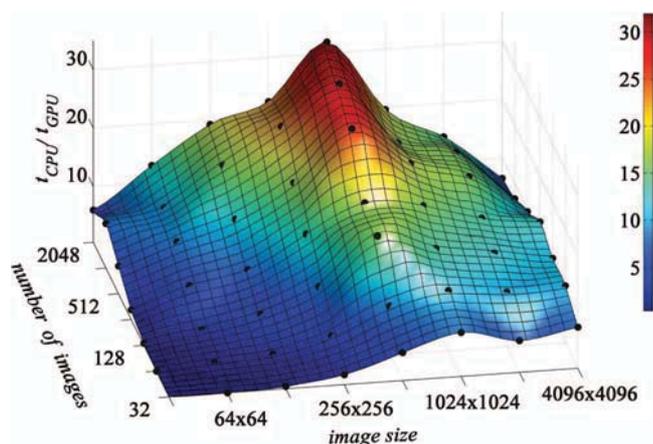


FIG. 3. Ratio between the total processing time of the GPU and the CPU based code as a function of the size of the images and their number.

of images (y axis). For essentially all of our experimental conditions the GPU code provides a decrease of the processing time, the only exceptions being a combination of image size smaller than  $128 \times 128$  and  $N < 128$ . The highest gain in speed reaches  $t_{CPU}/t_{GPU} = 32$  for  $512 \times 512$  sized images and  $N = 2048$  images. Under these conditions, while the CPU code requires about 47 min for analyzing data, the same analysis is performed by the GPU one in less than 2 min. As a general trend we note that increasing the number of images provides an increased speedup ratio up to image sizes of  $512 \times 512$ . This is certainly related to the fact that for 2048 images of  $512 \times 512$  pixels the amount of data slightly exceed the G-RAM capacity, therefore for 2048 images of larger size multiple loading is preventing larger speedups. A more detailed analysis of the timing of the different operations is provided in the supplementary material.<sup>20</sup>

From this we conclude that the performance of the code is limited by the competition between the speed increase due to the GPU parallelization and the slow down due to the additional time needed to transfer data to the GPU. This finding can provide guidelines for the further optimization of the code for specific applications. For example, in our case we chose to analyze all the time delays and to store the entire 2D power spectra, which might be unnecessary in many cases. Relaxing one or more of these requirements would of course improve the performance for large data sets. A widely used

concept would be, for example, to consider only delay times distributed evenly on a logarithmic scale.<sup>21</sup>

We are grateful to S. Biffi for help with MATLAB code, to R. Caravita for his cThread C++ class, to R. Cerbino and L. Chantada for setting up an earlier variant of the NFS experiment, and to R. A. Armenta-Calderón for interesting discussions. G.C., on leave from the University of Milan, would like to thank the Fribourg Center for Nanomaterials for financial support during a summer internship. F. Crocco acknowledges EU financial support: Marie Curie (Contract No. IEF-251131), DyNeFI. Financial support by the Swiss National Science Foundation (Project No. 132736) and the Adolphe Merkle Foundation is gratefully acknowledged.

- <sup>1</sup>B. J. Berne and R. Pecora, *Dynamic Light Scattering* (Wiley, New York, 1976).
- <sup>2</sup>M. Giglio, M. Carpineti, and A. Vailati, *Phys. Rev. Lett.* **85**, 1416 (2000).
- <sup>3</sup>R. Cerbino and A. Vailati, *Curr. Opin. Colloid Interface Sci.* **14**, 416 (2009).
- <sup>4</sup>F. Crocco, D. Brogioli, A. Vailati, M. Giglio, and D. S. Cannell, *Appl. Opt.* **45**, 2166 (2006).
- <sup>5</sup>F. Crocco, D. Brogioli, A. Vailati, M. Giglio, and D. S. Cannell, *Phys. Rev. E* **76**, 41112 (2007).
- <sup>6</sup>F. Crocco and D. Brogioli, *Appl. Opt.* **50**, 3419 (2011).
- <sup>7</sup>D. Brogioli, D. Salerno, F. Crocco, R. Ziano, and F. Mantegazza, *New J. Phys.* **13**, 123007 (2011).
- <sup>8</sup>R. Cerbino and V. Trappe, *Phys. Rev. Lett.* **100**, 188102 (2008).
- <sup>9</sup>R. Cerbino, L. Peverini, M. A. C. Potenza, A. Robert, P. Bösecke, and M. Giglio, *Nat. Phys.* **4**, 238 (2008).
- <sup>10</sup>A. Oprisan, S. Oprisan, and A. Teklu, *Appl. Opt.* **49**, 86 (2010).
- <sup>11</sup>L. G. Wilson, V. A. Martinez, J. Schwarz-Linek, J. Tailleur, G. Bryant, P. N. Pusey, and W. C. K. Poon, *Phys. Rev. Lett.* **106**, 018101 (2011).
- <sup>12</sup>D. Magatti, M. D. Alaimo, M. A. C. Potenza, and F. Ferri, *Appl. Phys. Lett.* **92**, 241101 (2008).
- <sup>13</sup>M. Garland, S. L. Grand, J. Nickolls, J. Anderson, J. Hardwick, S. Morton, E. Philips, Y. Zhang, and V. Volkov, *IEEE Micro* **28**, 13 (2008).
- <sup>14</sup>S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, and K. Skadron, *J. Parallel Distrib. Comput.* **68**, 1370 (2008).
- <sup>15</sup>P. J. Lu, H. Oki, C. A. Frey, G. E. Chamitoff, L. Chiao, E. M. Fincke, C. M. Foale, S. H. Magnus, W. S. McArthur, D. M. Tani, P. A. Whitson, J. N. Williams, W. V. Meyer, R. J. Sicker, B. J. Au, M. Christiansen, A. B. Schofield, and D. A. Weitz, *J. Real-Time Image Proc.* **5**, 179 (2010).
- <sup>16</sup>T. Shimobaba, Y. Sato, J. Miura, M. Takenouchi, and T. Ito, *Opt. Express* **16**, 11776 (2008).
- <sup>17</sup>P. J. Lu, F. Giavazzi, T. E. Angelini, E. Zaccarelli, F. Jargstorff, A. B. Schofield, J. N. Wilking, M. B. Romanowsky, D. A. Weitz, and R. Cerbino, *Phys. Rev. Lett.* **108**, 218103 (2012).
- <sup>18</sup>A. Vailati and M. Giglio, *Nature (London)* **390**, 262 (1997).
- <sup>19</sup>F. Crocco, H. Bataller, and F. Scheffold, "A light scattering study of non equilibrium fluctuations in liquid mixtures to measure the Soret and mass diffusion coefficient" *J. Chem. Phys.* (to be submitted).
- <sup>20</sup>See supplementary material at <http://dx.doi.org/10.1063/1.4755747> for a more detailed analysis of timings for different operations.
- <sup>21</sup>K. Schätzel, M. Drewel, and S. Stimac, *J. Mod. Opt.* **35**, 711 (1988).