Performing DVC at the Voxel Scale

François Hild, Hugo Leclerc, Stéphane Roux

► To cite this version:

François Hild, Hugo Leclerc, Stéphane Roux. Performing DVC at the Voxel Scale. 2013 SEM Annual Conference & Exposition on Experimental & Applied Mechanics, Jun 2013, Lombard, United States. hal-01385681

HAL Id: hal-01385681

<https://hal.science/hal-01385681>

Submitted on 21 Oct 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Performing DVC at the Voxel Scale

F. Hild,^a H. Leclerc^a and S. Roux^a

^aLMT-Cachan, ENS Cachan / CNRS / Université Paris 6 / PRES UniverSud Paris
61 avenue du Président Wilson, F-94235 Cachan Cedex, France
Email: hild@lmt.ens-cachan.fr

ABSTRACT

To analyze the displacement field in volumes imaged by X-ray tomography at several deformation states, a new approach is proposed whereby the displacement is measured down to the voxel scale and determined from a mechanically regularized system using the equilibrium gap method and additional boundary regularizations. It is then possible to compute a displacement vector for each voxel, inducing lower residuals (in terms of experimental data). As representative reconstructed volumes lead to large numbers of degrees of freedom, a dedicated GPU computational strategy has been developed and implemented. A set of volumes of size $100 \times 170 \times 256$ voxels (i.e., more than 13 million kinematic unknowns), which corresponds to a part of a sample made of nodular graphite cast iron and tested in tension, is analyzed.

Keywords: Digital volume correlation, GPU, in situ test, regularization, X-ray tomography.

INTRODUCTION

When displacement fields are measured, one key aspect is to get local fluctuations in the material to understand, for instance, how they can be correlated to the underlying microstructure. This requires the measurements to be as accurate as possible on the smallest possible scale. When dealing with volumes obtained via X-ray tomography, a voxel-scale analysis seems to be one of the ultimate scales of description. The aim of the present paper is to show that this goal can be reached via an adapted digital volume correlation (DVC) technique. Because the number of degrees of freedom increases substantially with such an approach, numerical issues need also to be addressed to meet the challenge of analyzing very large amount of data.

In the early developments of digital volume correlation, local approaches were implemented in which a volume of interest (*i.e.* a small part of the reconstructed volume) is chosen in the reference configuration and registered with a volume in the deformed configuration [1-3]. An alternative approach is to resort to global analyses in which the whole region of interest is analyzed by using, say, 8-noded cube elements to form a mesh (*i.e.* so called C8-DVC [4]). In this approach, a spatial regularization is enforced, namely, the displacement field is assumed *a priori* to be continuous. One way to achieve even smaller spatial resolutions is to regularize with mechanical requirements (*e.g.* static equilibrium [5]). It was shown that calculations could be performed at the voxel scale when coupling DVC with an elastic calculation. This result means that the spatial resolution was enhanced by at least one order of magnitude when compared to standard global and local approaches.

To be applicable to analyze real data (*i.e.* reconstructed volumes whose size can easily reach Gbytes), voxel-scale calculations are computationally very demanding. In the present paper, graphics processing units (GPUs) are used to perform these large scale calculations. With the development of GPUs, there are potential gains to be expected in terms of computation time and memory requirements when processing a series of pictures or volumes. Digital Image Correlation (DIC) has benefitted from the use of GPUs (*e.g.*, FFT-DIC [6] and global DIC [7-9]). This is also the case for local and global DVC [5, 10, 11].

VOXEL-SCALE DIGITAL VOLUME CORRELATION (V-DVC)

Let us consider f the volume in the reference configuration, and g the volume in the deformed configuration. Those two volumes capture the random texture that is passively advected during the experiment. The gray level conservation at any voxel location $\mathbf{x}$ then reads

$$f(\mathbf{x}) = g(\mathbf{x} + \mathbf{u}(\mathbf{x})) \quad (1)$$

where $\mathbf{u}$ is the displacement vector. Because the conservation law is never strictly satisfied due to acquisition noise, reconstruction artifacts and the number of unknowns exceeds the number of equations, it is not possible to determine the displacement vector $\mathbf{u}$ from the *sole* knowledge of f and g . The correlation procedure is generally written on a given domain that contains more than 1 voxel. Thus it aims, for instance, at minimizing the sum of squared differences

$$T_{DVC} = \int_{\Omega} (f(\mathbf{x}) - g[\mathbf{x} + \mathbf{u}(\mathbf{x})])^2 d\mathbf{x} \quad (2)$$

over the considered region of interest Ω in which the displacement field is interpolated as

$$\mathbf{u}(\mathbf{x}) = \sum_n u_n \mathbf{\Psi}_n(\mathbf{x}) \quad (3)$$

where $\mathbf{\Psi}_n$ are (chosen) vector fields, and u_n the associated degrees of freedom gathered in a global vector $\{\mathbf{u}\}$. The measurement problem then consists in minimizing T with respect to the unknown vector $\{\mathbf{u}\}$.

If a voxel-scale determination of the displacement field is sought, additional information is needed to regularize the previous measurement problem. This is achieved by using, say, the equilibrium gap method [12]. To enforce mechanical admissibility in an FE sense, the equilibrium gap is first introduced. If linear elasticity applies, the equilibrium equations read

$$[\mathbf{K}]\{\mathbf{u}\} = \{\mathbf{f}\} \quad (4)$$

where $[\mathbf{K}]$ is the stiffness matrix, and $\{\mathbf{f}\}$ the vector of nodal forces. When the displacement vector $\{\mathbf{u}\}$ is prescribed and if the (unknown) stiffness matrix is not the true one, load residuals $\{\mathbf{f}_r\}$ will arise

$$\{\mathbf{f}_r\} = [\mathbf{K}]\{\mathbf{u}\} - \{\mathbf{f}\} \quad (5)$$

In the absence of body forces, interior nodes of Ω are free from any external load. Consequently, the minimization of the equilibrium gap consists in minimizing the following quantity

$$T_{EG} = \{\mathbf{u}\}' [\bar{\mathbf{K}}] [\bar{\mathbf{K}}] \{\mathbf{u}\} \quad (6)$$

where $'$ is the transposition operator, $[\bar{\mathbf{K}}]$ the rectangular stiffness matrix associated with the inner nodes and those belonging to traction-free boundaries. Similarly, boundary regularization can be considered when the ROI boundary is not a free edge, and an *ad hoc* functional is considered T_{BR} [13]. Because this functional should be invariant under a solid body motion, it has to be based on a homogeneous differential operator (or its discrete form) of second order whose quadratic norm is to be minimized. Therefore, in terms of scaling properties, the boundary regularization behaves in a comparable way to the bulk equilibrium gap term. Yet, these three functionals are not of the same dimensional nature. Consequently, normalized quantities are constructed by normalizing by $T_{DVC}(\mathbf{v})$, $T_{EG}(\mathbf{v})$ and $T_{BR}(\mathbf{v})$, where $\mathbf{v}$ is a chosen trial displacement field (for which mechanical admissibility is not mandatory). One possible choice, among many, is a plane wave $\mathbf{v} = \mathbf{v}_0 \exp(2\pi i \mathbf{n} \cdot \mathbf{x} / \lambda)$, whose wavevector is aligned with $\mathbf{n}$ and where λ is the wave length whose value should be much smaller than the size of the ROI (*i.e.* the cubic root of the ROI volume) and yet much larger than the voxel size. The total functional T to be minimized then reads

$$T(1 + w_e + w_b) = \tilde{T}_{DVC} + w_e \tilde{T}_{EG} + w_b \tilde{T}_{BR} \quad \text{with} \quad \tilde{T}_{\bullet} = \frac{T_{\bullet}(\mathbf{u})}{T_{\bullet}(\mathbf{v})} \quad (7)$$

where the weights w_e and w_b are related to the regularization lengths ℓ_e and ℓ_b

$$w_e = \left(\frac{\ell_e}{\lambda} \right)^4 \quad \text{and} \quad w_b = \left(\frac{\ell_b}{\lambda} \right)^4 \quad (8)$$

Note that the scaling of the weights with respect to the wave length λ (exponent 4) results from the fact that the regularization kernels are the squares ($\times 2$) of second order ($\times 2$) differential operators. The larger the weights, the more the constraints defined by functionals T_{EG} and T_{BR} are enforced. This setting can be viewed as a filtering that enforces mechanical admissibility below a cut-off size equal to ℓ_e or ℓ_b . These length scales can be adjusted with the weight levels w_e and w_b .

Let us also emphasize that the minimization of T leads to a linear system for which a Green function (*i.e.* displacement field solution for a second member is equal to unity on a single voxel) can be defined in a mean-field sense. The latter has an algebraic (power-law) decay away from the point source and vanishes exponentially above a characteristic scale

of order ℓ_e . The key property is that the solution to the linear regularized problem has a complexity that can be compared to that of an elastic problem up to size ℓ_e . As it is expected to seek mode values of this length scale (say of the order of 10 voxels at most), convergence is not expected to be an issue (unless complicated crack geometries are considered). This guides towards an implementation scheme where the challenge is the efficient handling of memory, but not really the efficient convergence. The nature of the problem and hence the implementation strategy would be different if large ℓ_e were to be used.

IMPLEMENTATION

The mechanical regularization gives access to local and detailed kinematic features since a voxel-scale DVC (or V-DVC) is considered. Consequently, it leads to costly systems (*e.g.* about 2×10^8 DOFs for a 400^3 -voxel volume). Fortunately, the long-range coupling between these degrees of freedom is essentially controlled by the correlation kernel, whereas the mechanical regularization only controls short range couplings (at distances smaller than ℓ_e). Thus, a conjugate gradient with a block Jacobi preconditioner is expected to give a sufficient convergence rate to compete against direct or iterative solvers with more complex preconditioners. Jacobi preconditioners generally lead to poor convergence properties but can be highly parallelized, which permits to take full advantage of the current GPU hardware with an SIMD or SIMT programming model. However, if avoiding the need for synchronization is one of the first prerequisite to get correct execution speeds, memory access has also to be aligned within thread blocks as much as possible. Furthermore, and more specifically for GPUs, one has to take into account the fact that the cache memory is currently very small, even on recent models.

To solve these issues, a specific storage scheme has been developed for the considered volumes. Data are stored as tiny blocks (usually $4 \times 4 \times 4$ voxels in each block) that are assembled by chunks of size greater than the maximum warp size for the involved kernels (*i.e.* usually 256). Each block is then associated with a ‘thread,’ meaning that within a block, offsets are typically equal to $\{0 \times 256, 1 \times 256, \dots, 4^3 \times 256\}$. Thus, for a tiny block of size n^d , memory access to the neighboring data, which is mandatory for the mechanical operators) are aligned in $100 n^d / (n + 1)^d$ percent of the cases (≈ 0.51 for 4^3 blocks, ≈ 0.64 for 4^2 blocks for the two dimensional case). However, the main point is that unrolling loops for the processing of a block enables for the use of static aligned offsets, besides the possibility to use registers and to let the compiler find a way to make static and explicit caching of data.

ANALYSIS OF TENSILE TEST

A nodular graphite cast iron specimen is loaded *in situ* in a testing machine installed in ID 19 beam-line at the European Synchrotron Radiation Facility (ESRF) in Grenoble (France) [14]. Two reconstructed volumes are analyzed, namely, one

Fig. 1 -a-Region of interest in the reference configuration. Its size is $100 \times 170 \times 256$ voxels.
 -b-Transverse displacement along the longer edge direction expressed in voxels when $\ell_e = 32$ voxels. Warp factor is equal to 10. The ROI contains the outer (lateral) boundaries. The physical size of one voxel is equal to $5.1 \mu\text{m}$

in the reference configuration (applied load: 22 N), and a second for a 273-N tensile load. For each load level, a set of 600 radiographs was recorded during a 180-degree rotation on a charge coupled device (CCD) camera. This device was coupled with a fluorescent screen via optical lenses. The white beam coming from the synchrotron ring was rendered monochromatic by a multilayered monochromator. The energy of the beam was set to 60 keV. Reconstruction of the tomographic data was performed with a filtered back-projection algorithm developed at ESRF that provided 8-bit gray-scale 3D images with an isotropic voxel size of 5.1 μm . The gauge volume is parallelepipedic, $1.6 \times 0.8 \text{ mm}^2$ in cross-section and 10 mm in length. As can be seen in Figure 1a, the graphite nodules are well dispersed in the matrix and with a rather high volume fraction (14 %). They provide a very good texture for DVC. The characteristic diameter of nodules is of the order of 50 μm and the mean distance between them is 50 μm , or 10 voxels. Furthermore, the X-ray attenuation is very different for graphite (nodules) and iron (ferritic matrix) so that a very good contrast is observed between the graphite nodules and the ferritic matrix.

Figure 1b shows the displacement field obtained with the V-DVC algorithm when the parallelepipedic region of interest includes four free (lateral) faces. Note that the steps visible over the lateral faces corresponds to the voxel-scale discretization of the actual sample roughness (it was as cast with no surface finish). The estimated displacement field is smooth, while staying close with the experimental data (as will be shown in the following).

The mesh used in the previous calculations is shown in Figure 2. Since the size of the analyzed volume is $100 \times 170 \times 256$ voxels, more than 4 million elements are used, and therefore about 13 million degrees of freedom are needed to express the measured displacement field. The use of GPUs makes this type of calculation still tractable.

Fig. 2 Mesh of region of interest in the reference configuration composed of more than 4 million C8 elements.
The colors correspond to the same displacement component as in Figure 1b

An illustration of typical convergence rates is shown in Figure 3. For instance, volumes of size $100 \times 170 \times 128$ voxels are processed in approximately 10 minutes. It is far greater than that needed for a standard C8 global analysis ($\approx 0.2 \text{ s}$ when 16^3 -voxel elements are chosen for the same considered volume) but it is worth remembering that each voxel has 3 degrees of freedom, leading in this case to about 6,528,000 unknowns!

Fig. 3 Root mean square of the displacement corrections per voxel vs. iteration with a fixed $\ell_e = 64$ voxels, and without any multiscale procedure. The analyzed volume contains $100 \times 170 \times 128$ voxels

Figure 4 shows the change of the mean dimensionless correlation residual for the tensile test when convergence is reached. The proposed approach (V-DVC) leads to very low residuals, especially for small regularization lengths for which the displacement fluctuations are better captured. All these trends are confirmed when the present results are compared with a global DVC approach with 8-noded cubes (of edge size ℓ), i.e. C8-DVC. For large ℓ_e values, comparable residuals are observed because in the studied case as the displacement field is mostly affine, and hence no systematic error is made because of the space in which displacements are sought. However, in the general case, the mechanical regularization only removes the non-admissible part of the mechanical displacements below a chosen length scale, while other representations operate also as a high-frequency filter but introduce a somewhat *arbitrary* small scale smoothness that does not do justice to the mechanical behavior. Because the mechanical regularization is more faithful to the very nature of the measured field, the cut-off frequency of the filter can be pushed much higher than for other approaches, and hence the spatial resolution reaches unprecedented levels.

Fig. 4 Dimensionless correlation residual vs. regularization length ℓ_e (in voxels) for V-DVC (red dashed curve) or element size ℓ for C8-DVC (plain green line with square symbols)

CONCLUSION

The presented mechanical regularization allows us to perform measurements at the voxel scale with digital volume correlation. The described GPU implementation (in Cuda language) enabled representative volumes with more than 6 million degrees of freedom to be analyzed in less than 10 minutes on a GPU for a voxel-scale volume correlation. Further optimizations for this type of GPU implementation can still be achieved and will be investigated in the future. Let us also stress that key advantage of the voxel scale discretization is that the description of the geometry (domain boundary, but also inner porosities or cracks) can be described without interference with the mesh geometry as sub-voxel details are not expected to be relevant. In fact, a boundary is not described as such, namely, the voxels that do not belong to the specimen are simply attributed vanishing elastic constants (Figures 1b and 2).

It is to be emphasized that the lower scale discretization (here at the voxel scale) tend not to play any role when the actual displacement field is smooth. The voxel scale was chosen here because it allows for a generic implementation irrespective of the material, its rich or poor texture, and specimen geometry. Thus it enables for efficient implementations. However, the very same strategy can be used with a coarser kinematic discretization if no rapid variations are expected. In that case, regularization is no longer a computationally demanding procedure but it may constitute a very efficient tool to deal with very poor material textures without being too harmful to the description of kinematic fields.

Last, such an approach opens the way of using the mechanical regularization as a tool to identify mechanical parameters, or analyzing cracked samples with complex geometries. By identification, one may imagine elastic properties, either globally or by phase (referenced by their gray levels), or shapes or boundaries (say crack surface) to be precisely defined, or parameters of non-linear constitutive laws.

ACKNOWLEDGMENTS

The support of the French Agence Nationale de la Recherche through ‘RUPXCUBE’ project (ANR-09-BLAN-0009-01) is acknowledged. The tomographic images were obtained at ESRF through a grant for the experiment MA-501 on beamline ID 19. The scans used herein were obtained with the help of Drs. J.-Y. Buffière, A. Gravouil, N. Limodin, W. Ludwig, and J. Rannou.

REFERENCES

- [1] B.K. Bay, T.S. Smith, D.P. Fyhrie, M. Saad, Digital volume correlation: three-dimensional strain mapping using X-ray tomography, *Exp. Mech.*, 39 (1999) 217-226.
- [2] M. Bornert, J.-M. Chaix, P. Doumalin, J.-C. Dupré, T. Fournel, D. Jeulin, E. Maire, M. Moreaud, H. Moulinec, Mesure tridimensionnelle de champs cinématiques par imagerie volumique pour l'analyse des matériaux et des structures, *Inst. Mes. Métrol.*, 4 (2004) 43-88.
- [3] E. Verhulp, B. van Rietbergen, R. Huisjes, A three-dimensional digital image correlation technique for strain measurements in microstructures, *J. Biomech.*, 37 (2004) 1313-1320.
- [4] S. Roux, F. Hild, P. Viot, D. Bernard, Three dimensional image correlation from X-Ray computed tomography of solid foam, *Comp. Part A*, 39 (2008) 1253-1265.
- [5] H. Leclerc, J.-N. Périé, S. Roux, F. Hild, Voxel-scale digital volume correlation, *Exp. Mech.*, 51 (2011) 479-490.
- [6] R.A. Asmara, M. Hariadi, Accelerating Phase Based Motion Estimation with Hierarchical Search Technique Using Parallel Threading in Graphical Processing Unit (GPU), *Int. J. Comput. Sci. Network Security*, 9 (2009) 140-146.
- [7] H. Leclerc, J.-N. Périé, S. Roux, F. Hild, Integrated digital image correlation for the identification of mechanical properties, in: A. Gagalowicz, W. Philips (Eds.) *MIRAGE 2009*, Springer, Berlin, 2009, pp. 161-171.
- [8] J. Marzat, Y. Dumortier, A. Ducrot, Real-Time Dense and Accurate Parallel Optical Flow using CUDA, in: *Proc. WSCG 2009*, 2009, pp. 105-111.
- [9] D. Mitzel, T. Pock, T. Schoenemann, D. Cremers, Video Super Resolution using Duality Based TV-L1 Optical Flow, in: *Lecture Notes in Computer Science*, 2009, pp. 432-441.
- [10] M. Gates, J. Lambros, M. Heath, High-Performance Digital Volume Correlation for Experimental Mechanics Applications, in: *Proc. SEM XII*, Costa Mesa, CA (USA), 2012.
- [11] H. Leclerc, J.-N. Périé, F. Hild, S. Roux, Digital Volume Correlation: What are the limits to the spatial resolution?, *Méc. & Ind.*, (2013) in press, DOI: 10.1051/meca/2012025.
- [12] D. Claire, F. Hild, S. Roux, A finite element formulation to identify damage fields: The equilibrium gap method, *Int. J. Num. Meth. Engng.*, 61 (2004) 189-208.
- [13] Z. Tomičević, F. Hild, S. Roux, Mechanics-Aided Digital Image Correlation, *J. Strain Analysis*, (2013) accepted.
- [14] J. Réthoré, N. Limodin, J.-Y. Buffière, F. Hild, W. Ludwig, S. Roux, Digital volume correlation analyses of synchrotron tomographic images, *J. Strain Analysis*, 46 (2011) 683-695.