

Digital volume correlation: what are the limits to the spatial resolution?

HUGO LECLERC¹, JEAN-NOËL PÉRIÉ², FRANÇOIS HILD^{1,a} AND STÉPHANE ROUX¹

¹ LMT-Cachan, ENS Cachan/CNRS/UPMC/UniverSud Paris, 61 avenue du Président Wilson, 94235 Cachan Cedex, France

² Université de Toulouse ; INSA, UPS, Mines Albi, ISAE, ICA (Institut Clément Ader), 133 avenue de Rangueil, 31077 Toulouse, France

Received 14 December 2011, Accepted 19 September 2012

Abstract – Most of the norms used in the field of digital image (and volume) correlation to register two images (or volumes) lead to ill-posed problems. One of the frequent solutions is to enforce a restricted kinematics requiring a compromise between the richness of the solution (i.e., the spatial resolution) and the measurement uncertainty. An alternative route is to use a displacement norm that permits to alleviate this compromise by the means of a mechanical regularization used when the gray levels do not give enough information. It is then possible to compute a displacement vector *for each pixel or voxel*, inducing lower residuals (in terms of experimental data) while decreasing the noise sensitivity. The resolution performance of these different approaches is discussed, and compared for the analysis of a tensile test on a cast iron specimen based on a pair of tomographic images. As representative reconstructed volumes lead to a large number of degrees of freedom, a dedicated GPU computational strategy has been developed and implemented.

Key words: Global digital image correlation / GPU / regularization

Résumé – **Corrélation d'images volumiques : quelles limites à la résolution spatiale ?** La plupart des mesures de distance entre paire d'images ou de volumes conduisent, pour identifier les champs de déplacement correspondants, à des problèmes mal posés. Il s'en suit un compromis entre richesse des descriptions (c'est-à-dire la résolution spatiale) et incertitudes de mesure, étant entendu que la plupart des solutions adoptées conduisent à imposer des cinématiques restreintes. Une alternative est d'introduire une métrique du champ de déplacement qui permette, par le biais des équations de la mécanique, d'exploiter une connaissance a priori de l'objet d'étude pour pallier le manque d'information des images aux petites échelles. Ce faisant, il est possible d'obtenir un champ de déplacement dont la discrétisation spatiale est *identique à celle de l'image elle-même* (pixel ou voxel). La résolution de ces différentes approches est discutée et illustrée par l'analyse d'un essai de traction sur une éprouvette de fonte tomographiée in situ. L'implémentation GPU mise en œuvre permet de traiter des volumes à grand nombre de degrés de liberté.

Mots clés : Corrélation d'images / GPU / régularisation

1 Introduction

Computed tomography (CT) is a well-known imaging technique in the medical field. The applications of CT systems in human medicine rely on the different absorptions of biological tissues to image bones, muscles, and organs. Low X-ray radiations are used to make their dose as small as possible for the patient. This type of system has been one of the revolutions in medical imaging [1].

It is also possible to measure and inspect externally and internally components [2] and materials [3] using industrial tomographs or synchrotron radiation facilities. One of the main applications is to establish a complete metrology of parts, even though most of it is invisible to the eye. The full 3D geometry can then be compared to the CAD model of the observed part, or CAD data can be reconstructed from CT data. This type of technique is becoming very popular in the field of non destructive evaluation, and some new industrial applications use fast CT systems online. Similarly, the microstructure of

^a Corresponding author: hild@lmt.ens-cachan.fr

different materials can be visualized in a nondestructive manner.

Another application in the field of mechanics of materials is not only to image the material, but also to perform a mechanical test in situ and then analyze the reconstructed 3D volumes to measure 3D displacement fields in the bulk of the sample. This is achieved by resorting to digital volume correlation [4–6]. The technique consists in measuring displacement fields by registering interrogation volumes (i.e., whose size defines the spatial resolution). This type of technique is currently developed in research labs to monitor in situ mechanical tests [4–12], or to identify and validate constitutive models [13–15]. The same type of revolution as experienced in the field of medicine is now spreading in the area of material science and mechanics of materials. It is a safe bet to forecast similar developments in high-tech industries in the very near future.

As for any new measurement technique, metrological issues (i.e., resolution, uncertainty) are very important to address. The resolution of a measuring system is the “smallest change in a quantity being measured that causes a perceptible change in the corresponding indication” [16]. When dealing with pictures in general, reconstructed volumes as will be discussed in the sequel, the displacement resolution is not independent of the size of the interrogation window, or interrogation volume. There is always a compromise between the standard displacement resolution and the spatial resolution (e.g., the size of the interrogation volume), namely, the larger the interrogation volume, the smaller the displacement resolution. This result is very general to many optical techniques utilized to measure kinematic fields, in particular digital image correlation [17, 18] and digital volume correlation [19].

The displacement resolution can be seen as a frequency resolution, that is, the highest (or Nyquist) frequency controls the displacement resolution. To enhance the displacement resolution, the interrogation volume has to be increased. However, this increase is at the cost of capturing less displacement changes at small scales. This compromise is a very general property pertaining to different fields such as quantum mechanics (i.e., Heisenberg’s principle [20, 21]) to signal processing [22]. The introduction of wavelets [23] was one way of addressing this issue, offering a continuous compromise between spatial and frequency localization. The question addressed herein could then be formulated as: “is there a way of breaking Heisenberg’s limit?”, namely, what are the spatial resolution limits in the context of digital image (and volume) correlation?

When displacement fields are measured, one 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. With the previously described limitation, a voxel-scale analysis seems impossible. The aim of the present paper is to show that, on the contrary, this goal can be reached via an adapted DVC technique. Because the number of degrees

of freedom increases substantially with such an approach, numerical issues are also addressed to meet the challenge of analyzing very large amounts of data.

In the following, the discussion will be focused on digital volume correlation. However, the same trends and results are to be expected when resorting to digital image correlation. 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 its counter-part in the deformed configuration [4–7]. In computer vision other routes, for instance based upon variational formulations, have been followed for a long time. In the work of Horn and Schunck, a spatial regularization was introduced to reduce the displacement fluctuations [24]. When discontinuities arise, the previous approach is not appropriate per se [25]. “Smoother” penalizations based on robust statistics are implemented [26, 27]. An alternative approach is to resort to a global analysis in which the whole region of interest is analyzed by using, say, 8-noded cube elements to form a mesh (i.e., C8-DVC [11]). In this approach, a spatial regularization is enforced, namely, the displacement field is assumed a priori to be continuous. It will be shown herein that this regularization allows for a decrease of the displacement resolution when compared to a local approach for an equivalent spatial resolution. One way to achieve even smaller spatial resolutions is to regularize with mechanical requirements (e.g., static equilibrium [28]). It was shown that DVC 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. This is the case, for example, for accelerating particle image velocimetry (PIV) techniques [29–32] in fluid mechanics. This is also possible when dealing with digital image correlation, its counter-part in solid mechanics. FFT-DIC [33] and global DIC [34–36] have benefitted from the use of GPUs. Applications in radiotherapy for online registrations [37], or analyses of images from the International Space Station [38] are possible with very short computation time. Significant gains have also been achieved for some DVC implementations [39, 40].

Section 2 introduces the main elements to measure displacement fields with local and global approaches and its corresponding noise sensitivity. In Section 3, mechanically regularized DVC is introduced. The displacement resolution is assessed for the introduced approaches and shows that a mechanical regularization outperforms alternative approaches. However, this procedure, allowing for a large

increase in the number of kinematic degrees of freedom, challenges the computer implementation. To meet this goal, an original GPU implementation of DVC coupled with mechanical regularization is proposed in Section 4. The discussion on the spatial resolution will be illustrated in Section 5 by the analysis of an in situ tensile test on a nodular graphite cast iron sample, imaged by tomography at the European Synchrotron Radiation Facility (Grenoble, France).

2 Local and global digital volume correlation

To compare the performance of local and global approaches to DVC, the same displacement interpolations will be considered, namely, a trilinear interpolation based on 8-noded cube elements. In the first case, the continuity requirement will not be enforced so that a local approach will be implemented. When a global approach is chosen, the continuity of nodal displacements is enforced. It will be further regularized in Section 3 by requiring the measured displacement to be (at least locally) the solution to an elastic problem.

2.1 Local DVC

The registration of two gray level volumes in the reference (f) and deformed (g) configurations is based upon the conservation of the gray levels

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

where \mathbf{u} is the unknown displacement field. It consists in finding the displacement field by minimizing the sum of squared differences φ_c^2

$$\varphi_c(\mathbf{x}) = |g(\mathbf{x} + \mathbf{u}(\mathbf{x})) - f(\mathbf{x})| \quad (2)$$

where $\varphi_c(\mathbf{x})$ defines the field of correlation residuals. The minimization of φ_c^2 is a non-linear and ill-posed problem. In particular, if no additional information is available, it is *impossible* to determine the displacement for each voxel independently since there are three unknowns for a given (scalar) gray level difference. Consequently, a weak formulation is chosen. After integration over a Volume Of Interest (VOI), the correlation residual ϕ_c^2 reads

$$\phi_c^2 = \int_{\text{VOI}} [g(\mathbf{x} + \mathbf{u}(\mathbf{x})) - f(\mathbf{x})]^2 d\mathbf{x} \quad (3)$$

in which the displacement field is expressed in a (chosen) basis

$$\mathbf{u}(\mathbf{x}) = \sum_n u_n \boldsymbol{\psi}_n(\mathbf{x}) \quad (4)$$

where $\boldsymbol{\psi}_n$ are (chosen) vector functions, and u_n the associated degrees of freedom. The measurement problem then consists in minimizing ϕ_c^2 with respect to the unknowns u_n . A Newton iterative procedure is followed to circumvent the nonlinear aspect of the minimization problem.

Let \mathbf{u}^i denote the displacement at iteration i , and $\{\mathbf{u}\}^i$ the vector containing all the unknown degrees of freedom. By assuming small increments $d\mathbf{u} = \mathbf{u}^{i+1} - \mathbf{u}^i$ of the sought displacement field, a first order Taylor expansion is carried out $g(\mathbf{x} + d\mathbf{u}(\mathbf{x})) \approx g(\mathbf{x}) + d\mathbf{u}(\mathbf{x}) \cdot \nabla g(\mathbf{x}) \approx g(\mathbf{x}) + d\mathbf{u}(\mathbf{x}) \cdot \nabla f(\mathbf{x})$. Consequently, the linearized correlation residual reads

$$\phi_{lin}^2 = \frac{1}{2} \int_{\text{VOI}} [g^i(\mathbf{x}) + \sum_n du_n \boldsymbol{\psi}_n(\mathbf{x}) \cdot \nabla f(\mathbf{x}) - f(\mathbf{x})]^2 d\mathbf{x} \quad (5)$$

with

$$g^i(\mathbf{x}) = g(\mathbf{x} + \mathbf{u}^i(\mathbf{x})) \quad (6)$$

where du_n denotes the correction to the degree of freedom u_n . The minimization of ϕ_{lin}^2 with respect to all displacement corrections $\{d\mathbf{u}\}$ leads to the following linear system

$$[\mathbf{M}]\{d\mathbf{u}\} - \{\mathbf{b}\}^i = \{\mathbf{0}\} \quad (7)$$

with

$$M_{mn} = \int_{\text{VOI}} [\nabla f(\mathbf{x}) \cdot \boldsymbol{\psi}_m(\mathbf{x})][\nabla f(\mathbf{x}) \cdot \boldsymbol{\psi}_n(\mathbf{x})] d\mathbf{x} \quad (8)$$

$$b_m^i = \int_{\text{VOI}} [f(\mathbf{x}) - g^i(\mathbf{x})][\nabla f(\mathbf{x}) \cdot \boldsymbol{\psi}_m(\mathbf{x})] d\mathbf{x} \quad (9)$$

The iterations stop when the displacement corrections $\{d\mathbf{u}\}$ reach a small level chosen by the user.

In the following, 8-noded VOIs for which the displacement field is interpolated with trilinear shape functions are considered. Contrary to classical procedures (i.e., only the mean displacement is assigned to the center of each analyzed VOI [4,5]), the displacements of the eight nodes of each VOI are considered. The size ℓ of each VOI is defined by the length of any of its edge. Equation (8) shows that ℓ should be chosen so that $[\mathbf{M}^e]$ still remains definite and possibly well-conditioned. Consequently, ℓ should be larger than the correlation length of the underlying texture, but not too large so that displacement fluctuations are still captured. For any point \mathbf{x} within a VOI, the displacement along any direction reads

$$u(\mathbf{x}) = \sum N_i(\mathbf{x}) u_i^e \quad (10)$$

where $N_i(\mathbf{x})$ are the shape functions, here equal to

$$N_i(x, y, z) = \left(\frac{1}{2} \pm \frac{x}{\ell}\right) \left(\frac{1}{2} \pm \frac{y}{\ell}\right) \left(\frac{1}{2} \pm \frac{z}{\ell}\right) \quad (11)$$

It is worth mentioning that the local approaches clearly have the advantage of enabling for easy and efficient implementations. Assuming that in DVC computations the number of VOIs is significantly higher than the number of “cores” in a GPU, the straightforward and probably the best approach would be to map each VOI to a GPU thread. In other words, as the algorithm requires no synchronization, a Single Instruction Multiple Data (i.e., SIMD) or Single Instruction Multiple Threads (i.e., SIMT) implementation would be direct [41].

To estimate the displacement resolution, the following a priori analysis is performed. The un-noised reference being unknown, the noise is assigned to the deformed volume with a variance $2\sigma^2$, where σ^2 is the variance of (Gaussian) white noise. For each degree of freedom u_i^e (i.e., any component of the nodal displacements), the standard deviation $\sigma_{u_i^e}$ is expressed as [28]

$$\sigma_{u_i^e} = \sqrt{2}\sigma p \sqrt{K_{ii}^e} \quad (12)$$

with $[\mathbf{K}^e] = [\mathbf{M}^e]^{-1}$, where $[\mathbf{M}^e]$ is the elementary matrix associated with a C8 element, and p the physical size of one voxel. If the correlation length of the texture is larger than a few voxels and less than the VOI size ℓ , a “mean-field” hypothesis can be made. The squared image gradient appearing in matrix $[\mathbf{M}^e]$ (Eq. (8)) is extracted from the integral and changed into its expectation value

$$\langle \partial_\alpha f \partial_\beta f \rangle = \Gamma_f^2 \delta_{\alpha\beta} \quad (13)$$

where $\Gamma_f^2 = (1/2)\langle \|\nabla f\|^2 \rangle$ is half the mean square of the image gradient, and δ Kronecker’s delta. With this hypothesis, each element i, j of sub-matrix $[\mathbf{M}_{\alpha\beta}^e]$ is the integral over the VOI of the product of interpolation functions $N_i N_j$. Therefore sub-matrix $[\mathbf{M}_{\alpha\beta}^e]$ is given by

$$[\mathbf{M}_{\alpha\beta}^e] \propto \frac{\Gamma_f^2 \ell^3 p}{216} \delta_{\alpha\beta} \times \begin{bmatrix} 8 & 4 & 4 & 2 & 4 & 2 & 2 & 1 \\ 4 & 8 & 2 & 4 & 2 & 4 & 1 & 2 \\ 4 & 2 & 8 & 4 & 2 & 1 & 4 & 2 \\ 2 & 4 & 4 & 8 & 1 & 2 & 2 & 4 \\ 4 & 2 & 2 & 1 & 8 & 4 & 4 & 2 \\ 2 & 4 & 1 & 2 & 4 & 8 & 2 & 4 \\ 2 & 1 & 4 & 2 & 4 & 2 & 8 & 4 \\ 1 & 2 & 2 & 4 & 2 & 4 & 4 & 8 \end{bmatrix} \quad (14)$$

The covariance matrix associated with the measured degrees of freedom, which is proportional to the inverse of $[\mathbf{M}_{\alpha\beta}^e]$ [28]), is thus expressed as

$$[\mathbf{C}_{\alpha\beta}] \propto \frac{16\sigma^2 p^2}{\Gamma_f^2 \ell^3} \delta_{\alpha\beta} \times \begin{bmatrix} 8 & -4 & -4 & 2 & -4 & 2 & 2 & -1 \\ -4 & 8 & 2 & -4 & 2 & -4 & -1 & 2 \\ -4 & 2 & 8 & -4 & 2 & -1 & -4 & 2 \\ 2 & -4 & -4 & 8 & -1 & 2 & 2 & -4 \\ -4 & 2 & 2 & -1 & 8 & -4 & -4 & 2 \\ 2 & -4 & -1 & 2 & -4 & 8 & 2 & -4 \\ 2 & -1 & -4 & 2 & -4 & 2 & 8 & -4 \\ -1 & 2 & 2 & -4 & 2 & -4 & -4 & 8 \end{bmatrix} \quad (15)$$

Consequently, the standard displacement resolution of each degree of freedom becomes

$$\sigma_{u_i^e} = \frac{8\sqrt{2}\sigma p}{\Gamma_f \ell^{3/2}} \quad (16)$$

Equation (16) illustrates the compromise discussed in the introduction section, namely, the smaller the VOI size ℓ (i.e., the spatial resolution), the larger the displacement resolution. For very small sizes, matrix $[\mathbf{M}^e]$ may no longer be invertible so that minimum VOI sizes of 10 to 20 voxels are generally chosen. The previous result is therefore the manifestation of Heisenberg’s limit in the context of a local approach to DVC.

For any quantity of interest linearly related to the measured nodal displacements $\varphi = \sum_i \kappa_i u_i^e$, or $\varphi = \{\boldsymbol{\kappa}\}^t \{\mathbf{u}^e\}$, its standard uncertainty reads

$$\sigma_\varphi = (\{\boldsymbol{\kappa}\}^t [\mathbf{C}] \{\boldsymbol{\kappa}\})^{1/2} \quad (17)$$

where κ_i are known coefficients collected in vector $\{\boldsymbol{\kappa}\}$. For any point \mathbf{x} within a C8 VOI, the displacement along any direction is expressed by Equation (10). Consequently, Equation (17) can be utilized to map out the uncertainty $\sigma_u(\mathbf{x})$. The shape of the uncertainty function has a minimum at the center and a maximum value at the mesh nodes. If one single information is to be kept from the local analysis then the most reliable one is the displacement at its center

$$\bar{u}^e = \frac{1}{8} \sum_{i=1,8} u_i^e \quad (18)$$

and the corresponding standard deviation becomes

$$\sigma_{\bar{u}^e} = \frac{1}{8} \sqrt{\sum_{m=1,8} \sum_{n=1,8} C_{\alpha m \alpha n}^e} = \frac{\sigma_{u_i^e}}{8} \quad (19)$$

Almost an order of magnitude is gained in terms of displacement resolution when the displacement at the center is considered compared to that of the nodal displacements.

2.2 Global DVC

In a global approach, the correlation residual Φ_c^2 is computed over the whole Region of Interest (ROI)

$$\Phi_c^2 = \int_{\text{ROI}} [g(\mathbf{x} + \mathbf{u}(\mathbf{x})) - f(\mathbf{x})]^2 d\mathbf{x} \quad (20)$$

in which the displacement field is continuous by construction for a finite element discretization. A Newton iterative procedure is followed to circumvent the non-linear aspect of the minimization problem [11]. In order to compare with the local approach, 8-noded elements will also be used. Consequently, the region of interest is discretized with C8 finite elements for which the displacement field is interpolated with trilinear shape functions. The size of each element is defined by the length ℓ of any of its edge. A global matrix $[\mathbf{M}]$ is assembled by considering elementary matrices $[\mathbf{M}^e]$ whose components are defined in Equation (8) and vector $\{\mathbf{b}\}$ is assembled by calculating elementary vectors $\{\mathbf{b}^e\}$ expressed in Equation (9). A single global problem is solved iteratively

$$[\mathbf{M}]\{d\mathbf{u}\} = \{\mathbf{b}\}^i \quad (21)$$

where vector $\{\mathbf{b}\}^i$ is updated thanks to the picture corrections g^i , and vector $\{\mathbf{u}\}$ now collects all sought degrees of freedom.

GPU implementations of global DVC algorithms appear to be less obvious than those dedicated to local DVC due to the need for synchronization and linear algebra kernels. The effects at large wavelengths are relatively small, giving the opportunity to use, for instance, a simple block Jacobi preconditioner in a global non linear conjugate gradient [42]. This algorithm is suited to massively parallel architectures and its main parts are very close in terms of implementation to those dedicated to local DVC with the exception of some gathering operations. Currently, one iteration takes approximately 300 ms on an NVIDIA GTX 280 board for $400 \times 400 \times 400$ -voxel volumes with C8 elements of size $\ell = 16$ pixels.

The resolution analysis is applied to the finite element implementation [28]. Each inner node belongs to eight elements so that the shape function associated with a given inner node spans over 8 elements. In that case, the spatial resolution is equal to 2ℓ in each direction. Consequently, the variance is divided by 8, and the standard displacement resolution σ_{u_i} is divided by $2\sqrt{2}$ when compared with the value of a single C8 element (i.e., in a local approach). However, this result does not account for the additional requirement of continuity in the displacement field (i.e., the global approach deals with systems that are written on a level of the ROI). To estimate the additional gain associated with the continuity constraint, a global matrix is formed by assembling elementary matrices $[\mathbf{M}_{\alpha\beta}^e]$. A mesh made of 10^3 elements is considered and it was checked that the following numerical values do not vary more than 0.0001 when finer discretizations are used. For inner nodes, it is found that the standard displacement resolution

$$\sigma_{u_i} \approx 2.28 \frac{\sqrt{2}\sigma p}{\Gamma_f \ell^{3/2}} \quad (22)$$

When the same spatial resolution is considered (i.e., ℓ for a local approach, and 2ℓ for inner nodes of a global approach), the multiplicative factor induced by continuity is thus equal to $2.28 \times 2^{3/2}/8 \approx 0.8$. Consequently, there is a 20% gain provided by the continuity associated with global approaches. The same result is true when nodes belonging to an outer face (i.e., the spatial resolution is then $4^{1/3}\ell$), an outer edge (i.e., the spatial resolution is then $2^{1/3}\ell$) or they are an outer corner (i.e., the spatial resolution is then ℓ) nodes are analyzed. For the latter,

$$\sigma_{u_i} \approx 6.45 \frac{\sqrt{2}\sigma p}{\Gamma_f \ell^{3/2}} \quad (23)$$

to be compared with the displacement resolution of C8 nodes with a local approach given in Equation (16). The enforcement of continuity of the measured displacement field leads to a 20% decrease of the displacement resolution for C8 elements. In 2D, for 4-noded quadrilaterals (i.e., Q4 elements), it was shown by following the same

type of derivation that a 13% decrease is obtained [45], and in 1D, for 2-noded elements, it corresponds to 7%. Even though there is a gain when a global approach is used, the same trend is observed when the displacement resolution is related to the spatial resolution. Except for a pre-factor, it is found that $\sigma_{u_i} \propto \ell^{-3/2}$ in both cases. In a log-log plot, this would mean that the displacement resolution would be lower, but with the same dependence with the spatial resolution.

Because the number of connectivities of each node of the mesh is different (i.e., from 1 to 8) and the corresponding spatial resolution (i.e., from ℓ to 2ℓ), it leads to a displacement resolution σ_{u_i} that is not uniform over the whole ROI. This is also the case of $\sigma_{\bar{u}}$. For inner nodes, by using the same approach as above, the following approximations are obtained for $\sigma_{\bar{u}}$

$$\sigma_{\bar{u}} \approx 0.50 \frac{\sqrt{2}\sigma p}{\Gamma_f \ell^{3/2}} \quad (24)$$

A factor 2 is found when compared with a local approach since (see Eq. (19))

$$\sigma_{\bar{u}^e} \approx \frac{\sqrt{2}\sigma p}{\Gamma_f \ell^{3/2}} \quad (25)$$

Had Q4 elements been used in 2D, a factor 1.59 would have been found instead [45], and 1.2 for linear elements in 1D.

As a first approximation, the results obtained in this section show that the element size ℓ in a global approach can be reduced at least by a factor of 2 to achieve the same displacement resolution when compared to a local approach. The discretization can therefore be refined at a constant resolution when a global approach is used instead of a local one. However, the limitation by the correlation length of the texture is still an issue (i.e., the global matrix \mathbf{M} associated with Newton's scheme still needs to be definite so that system (21) is invertible).

3 Regularized digital volume correlation

A voxel-based digital image correlation (V-DVC) procedure is introduced [28]. If a voxel-scale determination of the displacement field is sought, additional information is needed to regularize the measurement problem. This will be achieved by using the equilibrium gap method [43] for inner nodes.

3.1 Equilibrium gap

To enforce mechanical admissibility in an FE sense, the equilibrium gap is first introduced. If linear elasticity applies, as will be assumed hereafter, the equilibrium equations read

$$[\mathbf{K}]\{\mathbf{u}\} = \{\mathbf{f}\} \quad (26)$$

where $[\mathbf{K}]$ denotes 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, force residuals $\{\mathbf{f}_r\}$ will arise

$$\{\mathbf{f}_r\} = [\mathbf{K}]\{\mathbf{u}\} - \{\mathbf{f}\} \quad (27)$$

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

$$\Phi_m^2 = \{\mathbf{u}\}^t [\mathbf{K}]^t [\mathbf{K}] \{\mathbf{u}\} \quad (28)$$

where t is the transposition operator, and $2\Phi_m^2$ corresponds to the sum of the squared norm of all equilibrium gaps at each interior node.

3.2 Boundary conditions

The equilibrium gap functional considers the residual body forces at *inner* nodes. Along the boundary of the ROI neither static nor kinematic information is available. Any displacement field prescribed on the boundary gives rise to a displacement field for which $\Phi_m = 0$. Hence, without boundary conditions, the mechanical regularization vanishes on the former. Similarly, close to the edges, the DVC algorithm gives rise to a higher level of uncertainty than in the bulk, because of the fewer number of neighboring elements [44].

In the following, no boundary regularization is introduced except in the special case where the boundary of the ROI meets a free surface for which a complete static information (i.e., zero tractions $\boldsymbol{\sigma} \cdot \mathbf{n} = \mathbf{0}$) is available and should be used instead of the above form. In that case, these nodes are considered as an “interior” entity, and their Young’s modulus becomes vanishingly small.

3.3 Correlation procedure

The two functionals introduced previously are considered jointly. First, a normalization procedure is introduced to make them dimensionless. A reference displacement field is first chosen in the form of a plane wave, $\mathbf{v}(\mathbf{x}) = e^{2i\pi\mathbf{x}/\lambda}$ where λ is its wavelength. This field will be used as a “gauge” for any displacement field $\mathbf{u}(\mathbf{x})$. Thus, the normalized functionals are defined as

$$\tilde{\Phi}_c^2 = \frac{\Phi_c^2(\mathbf{u})}{\Phi_c^2(\mathbf{v})}, \quad \tilde{\Phi}_m^2 = \frac{\Phi_m^2(\mathbf{u})}{\Phi_m^2(\mathbf{v})} \quad (29)$$

It is worth noting that the Young’s modulus, which was present in Φ_m , no longer appears in $\tilde{\Phi}_m$, and only Poisson’s ratio remains influential.

To solve the coupled minimization of correlation residuals (Φ_c^2) and equilibrium gap (Φ_m^2), a weighted sum of normalized residuals is considered

$$\Phi_t^2 = \frac{1}{1 + w_m} \left(\tilde{\Phi}_c^2 + w_m \tilde{\Phi}_m^2 \right) \quad (30)$$

where w_m is a dimensionless weight attached to $\tilde{\Phi}_m^2$. It is to be stressed that it introduces a characteristic length.

The wavelength dependence of the latter functional is of fourth order whereas $\tilde{\Phi}_c^2$ is wavelength-independent. Thus, the introduction of the equilibrium gap kernel can be seen as a fourth order low-pass filter, damping out deviations from mechanical admissibility below a given wavelength. Parameter w_m is then expressed as

$$w_m = \left(\frac{\ell_m}{\lambda} \right)^4 \quad (31)$$

where ℓ_m is the characteristic length for $\tilde{\Phi}_m^2$. The wavelength λ used for computing the regularization functional is counted in units of voxels. Thus, the above expression can be seen as a definition of the regularization length scale ℓ_m (also expressed in voxels). This length scale acts in a similar manner as the element size, ℓ , in the case where the global approach is used with no regularization. However, this comparison is solely based on dimensionality arguments, and hence a difficulty is to set a common meaningful gauge to compare both regularization lengths ℓ and ℓ_m . It is proposed to match the cut-off lengths, defined by the reduction of the power of the residual signal by a factor of 2 with respect to the original signal, of the low-pass filters associated respectively with the projection on a (tri-)linear finite element basis and the $\tilde{\Phi}_m^2$ functional. This leads to a cut-off wavelength of 2ℓ to be compared with ℓ_m . This factor of 2 will be used in the following as a genuine way of comparing the actual cut-off wavelengths of the different regularization schemes.

When the material parameters are known, the minimization of Φ_t^2 with respect to the unknown degrees of freedom can be performed [28]. It corresponds to an *integrated* way of measuring displacement fields since the displacement field will satisfy equilibrium equations, the constitutive law in addition to gray level conservation. The above mechanical regularization procedure can be used for an arbitrary supporting mesh. In the following, the elements will be considered to be 8-node cube (C8) elements of *voxel size* so that the underlying mesh is identical to that of the image *itself*. This is an extreme limit that cannot be handled using local or global approaches as the number of degrees of freedom is much higher than the available information (one gray level difference per voxel).

The influence of noise can be tackled in the same way as previously, studying the standard uncertainty of the displacement due to image noise, which is considered to be Gaussian and white. Ignoring edge effects, a displacement Green function, $[\mathbf{G}] = G_{i\alpha j\beta}$, is introduced giving the displacement component α at node i due to a unit source term localized at node j along the component β , and obeying

$$N_{k\gamma i\alpha} G_{i\alpha j\beta} = \frac{1}{\Gamma_f p^2} \Delta_{k\gamma j\beta} \quad (32)$$

with

$$[\mathbf{N}] = [\mathbf{M}^0] + A \left(\frac{\ell_m}{\lambda} \right)^4 [\mathbf{K}]^t [\mathbf{K}] \quad (33)$$

and

$$A = \frac{\{\mathbf{v}\}^t [\mathbf{M}]^0 \{\mathbf{v}\}}{\{\mathbf{v}\}^t [\mathbf{K}]^t [\mathbf{K}] \{\mathbf{v}\}} \quad (34)$$

where $[\mathbf{M}]^0$ is the assembly of the dimensionless matrices $1/(\Gamma_f^2 \ell^3 p) [\mathbf{M}^e]$ defined in Equation (14) with $\ell = 1$ voxel since single voxels are used as basic elements. The interest of introducing this Green function is that the resulting perturbative displacement field $\delta \mathbf{u}$ induced by a white noise $\eta(\mathbf{x})$ on the image is written as a mere convolution and dimensional factors that have been singled out, $\delta \mathbf{u} = \mathbf{G} \star \eta$. From the latter equation, the full covariance matrix in the displacement reads

$$C_{i\alpha j\beta} = 2\sigma^2 p^3 \left(\sum_{k,\gamma} G_{i\alpha k\gamma} G_{k\gamma j\beta} \right) \quad (35)$$

and, in particular, the variance of the nodal displacement assumes the following simple expression

$$\sigma_{u_r}^2 = C_{i\alpha i\alpha} = 2\sigma^2 p^3 \|\mathbf{G}\|^2 \quad (36)$$

where $\|\dots\|^2$ denotes the L2 norm of the Green function. This norm can equivalently (and conveniently) be computed in Fourier space, as Equation (33) that stands for a differential equation in a discretized form is turned into a simple algebraic equation. Noting that the operator $[\mathbf{K}]^t [\mathbf{K}]$ is in Fourier space homogeneous to $\|\mathbf{k}\|^4$, where wavevectors are denoted as \mathbf{k} , the Green function is essentially a constant $1/(\Gamma_f p^2)$ for low frequencies, cut off at $\|\mathbf{k}\|^* \propto \lambda/(A^{1/4} \ell_m)$. The L2 norm can thus be written as

$$\|\mathbf{G}\|^2 \propto \frac{1}{\Gamma_f^2 p} \left(\frac{\lambda}{A^{1/4} \ell_m} \right)^3 \quad (37)$$

so that

$$\sigma_{u_r} \propto \frac{\lambda^{3/2}}{A^{3/8}} \frac{\sigma p}{\Gamma_f \ell_m^{3/2}} \quad (38)$$

The second term has a similar scaling dependence as the one obtained in un-regularized DIC (see Sect. 2), with a substitution from ℓ to ℓ_m . Thus the regularization length plays the same role as the element size or the ZOI size in the previous sections. The prefactor comes from the specific choice of the reference or gauge displacement field \mathbf{v} .

As will be shown in the sequel, the limitation in terms of spatial resolution with local and global approaches in comparison with the correlation length of the underlying texture can be by-passed thanks to the proposed regularization (see Eq. (33)) because matrix $[\mathbf{N}]$ is still definite even though matrix $[\mathbf{M}]$ may non longer fulfil this property. The choice of ℓ_m still remains. The larger ℓ_m , the more the measured displacement field will satisfy mechanical admissibility. This is a choice that is left to the user. It is to be emphasized that this choice is the recurrent compromise between resolution and spatial resolution that cannot be decided a priori. However, the correlation residual field is extremely useful to decide on the choice of acceptable parameters. The latter shows how

good or bad the registration is spatially, and possible ill-behaviors appear generally very clearly [11, 13–15, 46, 47]. The interpretation in terms of too low and especially of too high regularization length scales becomes soon very intuitive with little practice.

4 Implementation

The mechanical regularization gives access to local and detailed kinematic features since a voxel-scale DVC is considered. Consequently, it leads to costly systems (e.g., 192×10^6 DOFs for a 400^3 -voxel tomography). 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 coupling (at distances smaller than ℓ_m). Thus, as in the C8 case, 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 CPU and 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 (i.e., mandatory for the mechanical operators) are aligned in $100 \frac{n^d}{(n+1)^d}$ 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.

An illustration of typical convergence rates is shown in Figure 1. A set of volumes of size $100 \times 170 \times 128$ voxels is processed in approximatively 10 min. It is far greater than that needed for a “standard” C8 global analysis (≈ 0.2 s when 16^3 -voxel elements are chosen) but one has to recall that each voxel has 3 degrees of freedom, leading in this case to about 6 528 000 unknowns!

5 Analysis of tomographic images

To illustrate the previous discussion, let us consider an experimental test case. A nodular graphite cast iron

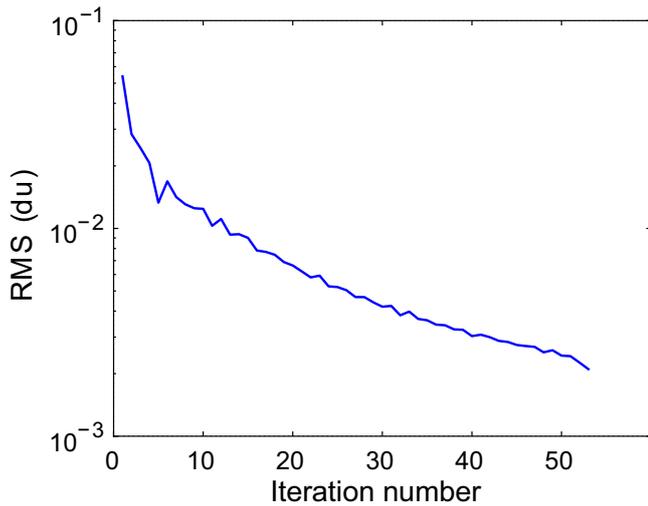


Fig. 1. Rms of the displacement updates per voxel vs. iteration number for the sample presented in Section 5 with a fixed $\ell_m = 64$ voxels, and without a multiscale approach.

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) [46]. Two tomographic images are analyzed, namely, one in the reference state (applied load: 22 N), and a second under a 273-N tensile load. The testing machine used in the experiment is shown in Figure 2a. For each load level, a set of 600 radiographs was recorded during a 180-degree rotation on a charge coupled device (CCD) camera with a square array of 2048×2048 pixels (exposure time: 3 s). This detector 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. The gauge volume is parallelepipedic, $1.6 \times 0.8 \text{ mm}^2$ in cross-section and 10 mm in length. The voxel size is equal to $5.1 \mu\text{m}$. The graphite nodules are well dispersed in the matrix and with a rather high volume fraction (i.e., 14%). They provide a very good texture for DVC. The characteristic diameter of nodules is of the order of $50 \mu\text{m}$ and the mean distance between nodules $50 \mu\text{m}$, or 10 voxels). Furthermore, the X-ray attenuation is very different for graphite and iron so that a very good contrast is observed between the graphite nodules and the ferritic matrix (Fig. 2b). Therefore analyses dealing with VOIs or elements whose size is greater than 8 voxels do not need any regularization.

Figure 3 is a plot of the standard deviation of the displacements obtained by correlation of the reference volume and the same volume with additional noise ($\sigma = \sqrt{2}$ gray level). This procedure was preferred in comparison to the standard artificial rigid body translation because the latter usually underestimates the displacement resolution associated with reconstructed volume since it does not

account for noise [19]. The abscissa represents respectively the element size ℓ for C8 approaches and $\ell_m/2$ for V-DVC. The results are compared using these lengths to obtain the same cut-off wavelengths for the global and regularized cases (see Sect. 3.3). For the local approach, the $-3/2$ power law dependence predicted by Equation (16) is found for large VOIs. The fact that the resolution departs from this slope for as smaller VOI sizes is an indication of gradual divergence. For global and regularized approaches, the predicted $-3/2$ power law dependence is found for small element sizes or regularization lengths since boundary effects are less pronounced. As expected from the theoretical results, the global C8-approach yields lower displacement uncertainties than the local C8-DVC. For small element lengths, a ratio of $0.3 \approx 0.8/2\sqrt{2}$ is found between the two displacement resolutions, in accordance with the results of Sections 2.1 and 2.2. In this last situation, the boundary effects are negligible. Last, the regularized version (V-DVC) out-performs both C8 approaches. In the present case, the displacement fluctuations induced by a random white noise are not mechanically admissible, and are therefore filtered out by the mechanical regularization.

Figure 4 shows the displacement field that is obtained with the V-DVC algorithm either for a parallelepipedic region of interest strictly contained within the specimen, or including four free (lateral) faces. In the latter case, slightly higher residuals are obtained, but better results are obtained in terms of noise sensitivity (i.e., free borders can be seen as an additional and very robust constraint). It leads to smooth results, while staying close with the experimental data (as will be shown in the following). Similar displacement fields (with however a much coarser spatial resolution) were obtained based on local and global C8 approaches.

Figure 5 shows the change of the dimensionless residual for the tensile test. The local approach tends to give lower residuals for *medium sized* elements, compared to *equivalent* global ones since the former contains more degrees of freedom. Moreover, V-DVC leads to far better residuals, especially for small regularization lengths for which the displacement fluctuations are better captured. For large ℓ_m or ℓ all correlation techniques converge to a comparable residual, because in the studied case, the displacement field is mostly affine, and hence no systematic error is made because of the choice of the space in which the displacement is sought. However, in the general case, the mechanical regularization only removes the non admissible part of the mechanical displacements, while the C8 representation operates as a somewhat *arbitrary* high-frequency filter that does not do justice to the mechanical behavior.

6 Conclusion

The presented mechanical regularization allows us to significantly alleviate the compromise between spatial resolution and measurement uncertainties in digital volume

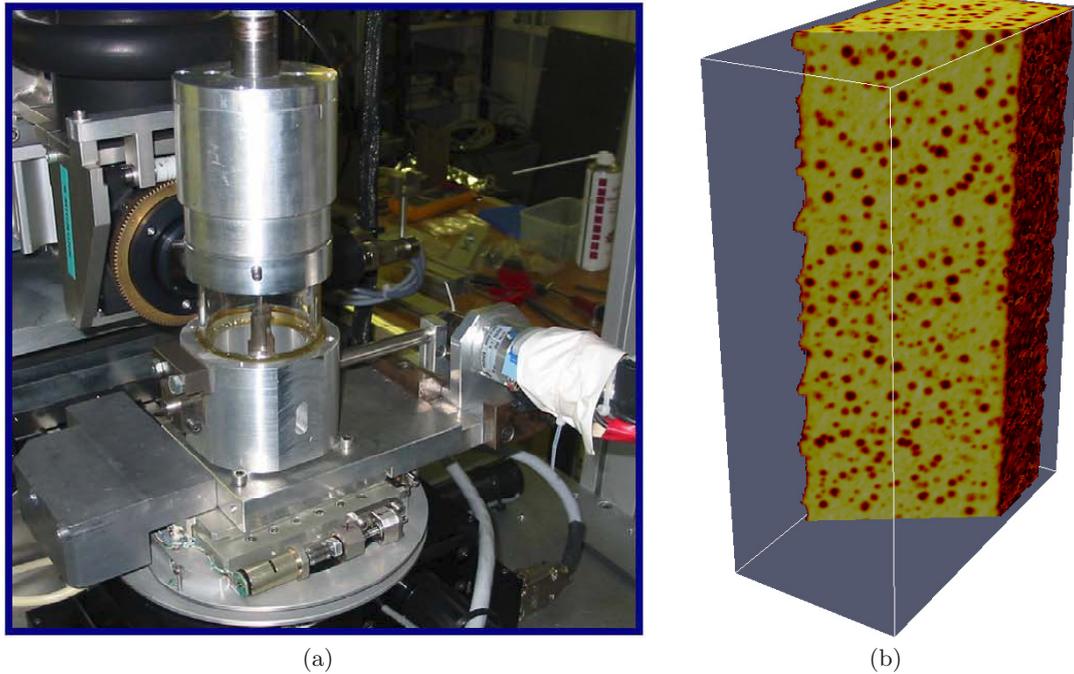


Fig. 2. In situ testing machine used to perform the experiment reported herein (a). Region of interest in the reference configuration (b). Its size is $100 \times 170 \times 256$ voxels.

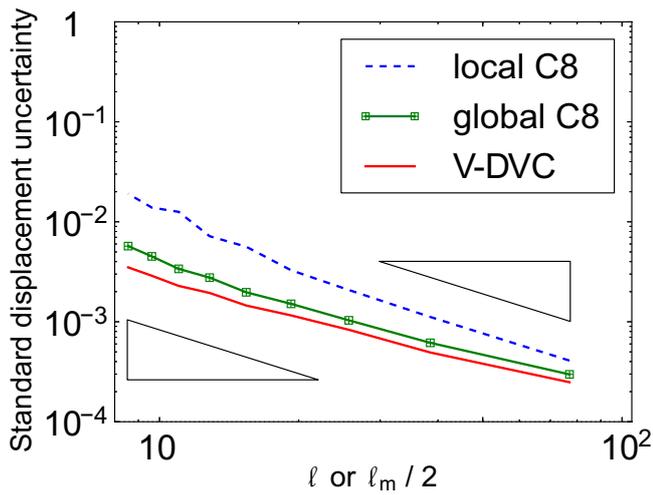


Fig. 3. Standard uncertainty of the displacement vs. element size l or regularization length $l_m/2$ (in voxels) for a Gaussian white noise ($\sigma = \sqrt{2}$ gray level). The $-3/2$ power law dependence is depicted by black triangles.

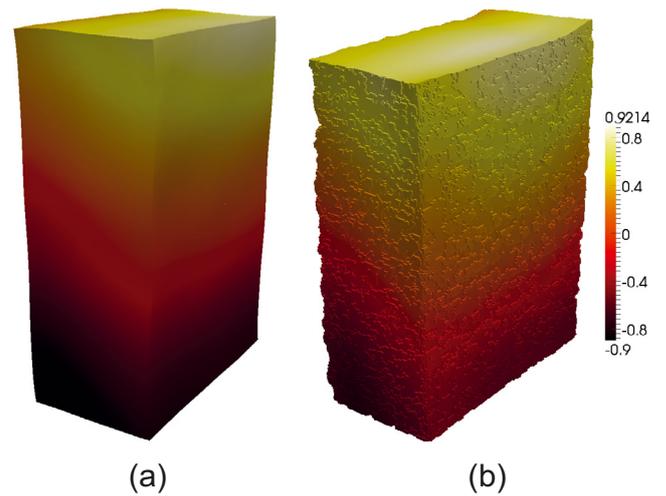


Fig. 4. Transverse displacement along the longer edge direction expressed in voxels of the sample described in Section 5 with $l_m = 32$ voxels. Warp factor is equal to 10. The ROI lies inside the sample (a) or contains the outer (lateral) boundaries (b). The size of the analyzed volume is $100 \times 170 \times 256$ voxels.

correlation. The illustrations in this paper have been obtained using three degrees of freedom per voxel, which can be seen as a “maximum” spatial resolution. Refining further the displacement discretization can be done within the presented theoretical framework, however the image discretization itself will limit the possible expected gain. It has been shown that while considering this spatial resolution, it is possible to decrease the uncertainty level and the distance to the experimental data, with the

additional benefit of obtaining a mechanically admissible field over a *tunable* scale. The latter is provided by the regularization length introduced herein. Depending on the values of the regularization length, the displacement fluctuations whose wavelengths are greater than the regularization length and that are not mechanically admissible are filtered out.

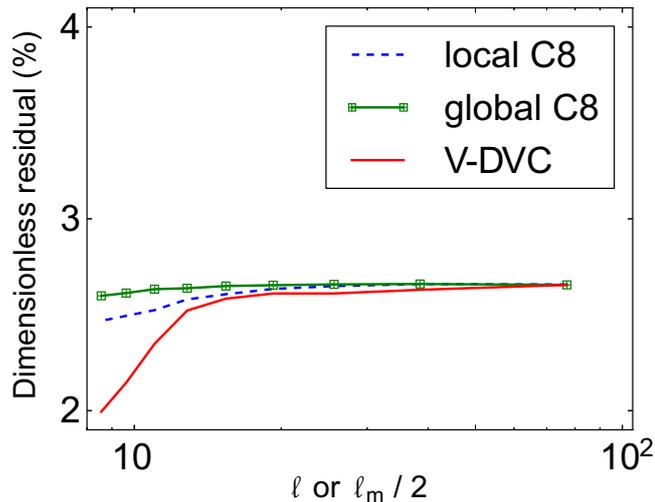


Fig. 5. Dimensionless correlation residual vs. element size ℓ or regularization length $\ell_m/2$ (in voxels).

The described GPU implementation allowed us to analyze representative volumes with more than 6 million degrees of freedom in less than 10 min on a GPU (Cuda language) for a voxel-scale volume correlation. This analysis out-performed local and global DVC analyses with C8 elements. For the latter, it was shown that a global approach itself out-performs a local approach with the same number of elements since less degrees of freedom are needed in the former compared to the latter. Moreover, the continuity of the displacement field should make the global approach more stable than the local one for fine discretizations.

Further optimizations for this type of GPU implementation can still be achieved and will be investigated in the future. In parallel, 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.

Acknowledgements. 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.A. Mae, The revolution in medical imaging rosen Pub. Group, New York, NY (USA), 2003
- [2] Industrial CT scanning, http://en.wikipedia.org/wiki/Industrial_CT_scanning, accessed in February, 2012
- [3] J. Baruchel, J.-Y. Buffière, E. Maire, P. Merle, G. Peix, X-Ray Tomography in Material Sciences, Hermes Science, Paris France, 2000
- [4] 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
- [5] T.S. Smith, B.K. Bay, M.M. Rashid, Digital volume correlation including rotational degrees of freedom during minimization, *Exp. Mech.* 42 (2002) 272–278
- [6] 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
- [7] E. Verhulp, B. van Rietbergen, R. Huiskes, A three-dimensional digital image correlation technique for strain measurements in microstructures, *J. Biomech.* 37 (2004) 1313–1320
- [8] A. Germaneau, P. Doumalin, J.C. Dupré, Full 3D measurement of strain field by scattered light for analysis of structures, *Exp. Mech.* 47 (2007) 523–532
- [9] N. Lenoir, M. Bornert, J. Desrues, P. Bésuelle, G. Viggiani, Volumetric digital image correlation applied to X-ray microtomography images from triaxial compression tests on argillaceous rock, *Strain* 43 (2007) 193–205
- [10] F. Forsberg, R. Mooser, M. Arnold, E. Hack, P. Wyss, 3D micro-scale deformations of wood in bending: Synchrotron radiation CT data analyzed with digital volume correlation, *J. Struct. Biol.* 164 (2008) 255–262
- [11] 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
- [12] F. Hild, E. Maire, S. Roux, J.-F. Witz, Three dimensional analysis of a compression test on stone wool, *Acta Mat.* 57 (2009) 3310–3320
- [13] N. Limodin, J. Réthoré, J.-Y. Buffière, A. Gravouil, F. Hild, S. Roux, Crack closure and stress intensity factor measurements in nodular graphite cast iron using 3D correlation of laboratory X ray microtomography images, *Acta Mat.* 57 (2009) 4090–4101
- [14] N. Limodin, J. Réthoré, J.-Y. Buffière, F. Hild, S. Roux, W. Ludwig, J. Rannou, A. Gravouil, Influence of closure on the 3D propagation of fatigue cracks in a nodular cast iron investigated by X-ray tomography and 3D Volume Correlation, *Acta Mat.* 58 (2010) 2957–2967
- [15] J. Rannou, N. Limodin, J. Réthoré, A. Gravouil, W. Ludwig, M.-C. Baïetto-Dubourg, J.-Y. Buffière, A. Combescure, F. Hild, S. Roux, Three dimensional experimental and numerical multiscale analysis of a fatigue crack, *Comp. Meth. Appl. Mech. Eng.* 199 (2010) 1307–1325
- [16] ISO/IEC guide 99-12:2007, International Vocabulary of Metrology – Basic and General Concepts and Associated Terms, VIM, International Organization for Standardization, Geneva (Switzerland), 2007
- [17] S. Bergonnier, F. Hild, S. Roux, Digital image correlation used for mechanical tests on crimped glass wool samples, *J. Strain Analysis* 40 (2005) 185–197
- [18] G. Besnard, F. Hild, S. Roux, Finite-element displacement fields analysis from digital images: Application to Portevin-Le Châtelier bands, *Exp. Mech.* 46 (2006) 789–803
- [19] N. Limodin, J. Réthoré, J. Adrien, J.-Y. Buffière, F. Hild, S. Roux, Analysis and artifact correction for volume correlation measurements using tomographic images from a laboratory X-ray source, *Exp. Mech.* 51 (2011) 959–970

- [20] W. Heisenberg, Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik, *Z. Phys.* 43 (1927) 172–198
- [21] W. Heisenberg, *Physikalische Prinzipien der Quantentheorie*, Hirzel, Leipzig (Germany), 1930
- [22] D. Gabor, Theory of communication, *J. Ins. Electl. Eng.* 93 (1946) 429–457
- [23] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, San Diego (California), USA, 1998
- [24] B.K.P. Horn, B.G. Schunck, Determining optical flow, *Artificial Intelligence* 17 (1981) 185–203
- [25] A. Mitiche, P. Bouthemy, Computation and analysis of image motion: A synopsis of current problems and methods, *Int. J. Comp. Vision.* 19 (1996) 29–55
- [26] M. Black, *Robust Incremental Optical Flow*, Ph.D. dissertation, Yale University, 1992
- [27] J.-M. Odobez, P. Bouthemy, Robust multiresolution estimation of parametric motion models, *J. Visual Comm. Image Repres.* 6 (1995) 348–365
- [28] H. Leclerc, J.-N. Périé, S. Roux, F. Hild, Voxel-scale digital volume correlation, *Exp. Mech.* 51 (2011) 479–490
- [29] T. Schiwietz, R. Westermann, GPU-PIV, edited by B. Girod, M.A. Magnor, H.-P. Seidel, in: *Proceedings Vision, Modeling and Visualization Conference*, Aka GmnH, 2004, pp. 151–158
- [30] P. Kondratieva, J. Georgii, R. Westermann, H. Petermeier, W. Kowalczyk, A. Delgado, A real-time model-based approach for the reconstruction fluid flows induced by microorganisms, *Exp. Fluids* 45 (2008) 203–222
- [31] F. Champagnat, A. Plyer, G.L. Besnerais, Y.L. Sant, How to calculate dense PIV vector fields at video rate, *Proceedings 8th International Symposium of Particle Image Velocimetry*, 2009, PIV09-0091
- [32] S. Tarashima, M. Tange, S. Someya, K. Okamoto, GPU accelerated direct cross-correlation PIV with window deformation, *Proc. 15th Int Symp on Applications of Laser Techniques to Fluid Mechanics*, 2010, p. 8
- [33] 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
- [34] H. Leclerc, J.-N. Périé, S. Roux, F. Hild, Integrated digital image correlation for the identification of mechanical properties, in: *MIRAGE 2009*, A. Gagalowicz, W. Philips, eds., (Springer, Berlin, 2009), LNCS 5496 161–171
- [35] J. Marzat, Y. Dumortier, A. Ducrot, Real-Time Dense and Accurate Parallel Optical Flow using CUDA, *Proc. WSCG 2009*, pp. 105–111
- [36] D. Mitzel, T. Pock, T. Schoenemann, D. Cremers, Video Super Resolution using Duality Based TV-L1 Optical Flow, in: *Lecture Notes in Computer Science*, 5748 (2009) 432–441
- [37] K.Ø. Noe, B.D. De Senneville, U.V. Elstrøm, K. Tanderup, T.S. Sørensen, Acceleration and validation of optical flow based deformable registration for image-guided radiotherapy, *Acta Oncologica* 47 (2008) 1286–1293
- [38] 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 Jr., D.M. Tani, P.A. Whitson, J.N. Williams, W.V. Meyer, R.J. Sicker, B.J. Au, M. Christiansen, A.B. Schofield, D.A. Weitz, Orders-of-magnitude performance increases in GPU-accelerated correlation of images from the International Space Station, *J. Real-Time Image Proc.* 5 (2010) 179–193
- [39] D. Ruijters, B.M. Ter Haar Romeny, P. Suetens, GPU-accelerated elastic 3D image registration for intra-surgical applications, *Comput. Meth. Programs Biomed.* 103 (2011) 104–112
- [40] J.A. Shackelford, N. Kandasamy, G.C. Sharp, On developing B-spline registration algorithms for multi-core processors, *Phys. Med.* 55 (2010) 6329–6351
- [41] NVIDIA CUDA C Programming Guide, NVIDIA Corporation, (http://developer.download.nvidia.com/compute/DevZone/docs/html/C/doc/CUDA_C_Programming_Guide.pdf, accessed in December 2011)
- [42] Y. Saad, H.A. van der Vorst, Iterative Solution of Linear Systems in the 20th Century, *J. Comput. Appl. Math.* 123 (2000) 1–33
- [43] D. Claire, F. Hild, S. Roux, A finite element formulation to identify damage fields: The equilibrium gap method, *Int. J. Num. Meth. Eng.* 61 (2004) 189–208
- [44] J. Réthoré, G. Besnard, G. Vivier, F. Hild, S. Roux, Experimental investigation of localized phenomena using Digital Image Correlation, *Phil. Mag.* 88 (2008) 3339–3355
- [45] F. Hild, S. Roux, Comparison of local and global approaches to digital image correlation, *Exp. Mech.* 2012 [in press, DOI 10.1007/s11340-012-9603-7]
- [46] 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
- [47] F. Hild, A. Fanget, J. Adrien, E. Maire, S. Roux, Three dimensional analysis of a tensile test on a propellant with digital volume correlation, *Arch. Mech.* 63 (2011) 1–20