Working document: A Comparison of Different Quasi-Newton Acceleration Methods for Partitioned Multi-Physics Codes

Rob Haelterman, Alfred Bogaers, Joris Degroote

Abstract In many cases, different physical systems interact, which translates to coupled mathematical models. We only focus on methods to solve (strongly) coupled problems with a partitioned approach, i.e. where each of the physical problems is solved with a specialized code that we consider to be a black box solver.
Running the black boxes one after another, until convergence is reached, is a standard but slow solution technique, known as non-linear Gauss-Seidel iteration. A recent interpretation of this approach as a root-finding problem has opened the door to acceleration techniques based on quasi-Newton methods that can be "strapped onto" the original iteration loop without modification to the underlying code. In this paper, we analyze the performance of different acceleration techniques on different multi-physics problems.

Keywords: Acceleration, Fluid-Structure Interaction, Iterative Method, Partitioned Method, Root-Finding, Quasi-Newton Method.

Disclaimer

For published version see
https://link.springer.com/book/10.1007%2F978-981-10-7488-2

Rob Haelterman
Royal Military Academy, Dept. Mathematics, Renaissancelaan 30, B-1000 Brussels, e-mail: robby.haelterman@rma.ac.be

Alfred Bogaers
Council for Scientific and Industrial Research, Advanced Mathematical Modelling, Modelling and Digital Sciences, Meiring Naudé Road; Brummeria, Pretoria, South Africa, e-mail: abogaers@csir.co.za

Joris Degroote
Ghent University, Dept. Flow, Heat and Combustion Mechanics, Sint-Pietersnieuwstraat 41, 9000 Ghent, Belgium, e-mail: joris.degroote@ugent.be

1 Introduction

Multi-physics problems occur in various disciplines. One example is fluid-structure interaction (FSI), which refers to the coupling between a fluid flow and a moving or deforming structure. This kind of problems can be found in heart valves, flapping flags, flutter of electricity cables, etc. Also problems with more than two fields exist, such as fluid-structure-thermal interaction in sintering processes and gas turbine blades. Mathematically these can be written as the a non-linear system of equations:

$$\begin{cases} f_1(x_1,x_2,\ldots,x_n) = 0 \\ f_2(x_1,x_2,\ldots,x_n) = 0 \\ \qquad\qquad\qquad\vdots \\ f_n(x_1,x_2,\ldots,x_n) = 0 \end{cases} \tag{1}$$

where $f_i : D_F \subset \mathbb{R}^m \to \mathbb{R}^{k_i}$, $x_j \in \mathbb{R}^{m_j}$ $(i,j \in \{1,2,\ldots,n\})$, $\sum_{i=1}^{n} k_i = m$, $\sum_{j=1}^{n} m_j = m$. Each equation describes (the discretized equations of) a physical problem that is spatially or mathematically decomposed. E.g. $f_1(x_1,x_2) = 0$ could give the pressure x_1 on the wall of a flexible tube for a given geometry x_2, while $f_2(x_1,x_2) = 0$, could give the deformed geometry of the wall due to the pressure exerted on it by the fluid. We will assume the problem has the following characteristics [11, 21]:

1. Good solvers exist for each equation of the system. For this reason, we will use a partitioned solution method.
2. The analytic form of f_i $(i = 1,2,\ldots,n)$ is unknown.
3. The problem has a large dimensionality.
4. Evaluating f_i $(i = 1,2,\ldots,n)$ is computationally costly.

2 Fixed-point formulation and quasi-Newton accelaration

A typical solution method for (1) is the fixed-point iteration [22].

Fixed-point iteration

1. Startup: Take initial values x_2^1, x_3^1, $\ldots x_n^1$. Set $s = 1$.
2. Loop until convergence:
   2.1. Solve $f_1(x_1,x_2^s,\ldots,x_n^s) = 0$ for x_1, resulting in x_1^{s+1}.
   2.2. Solve $f_2(x_1^{s+1},x_2,\ldots,x_n^s) = 0$ for x_2, resulting in x_2^{s+1}.
   ...
   2.n. Solve $f_n(x_1^{s+1},x_2^{s+1},\ldots,x_n) = 0$ for x_n, resulting in x_n^{s+1}.
   2.n+1. Set $s = s+1$.

If F' (the Jacobian of $F = (f_1,\ldots,f_n)$) satisfies the condition

$$\forall i < j < n : [F']_{ij} = 0 \tag{2}$$

then we can write the whole process 2.1. to 2.n, as $x_n^{s+1} = H(x_n^s)$ [20]. The problem can now be considered as finding the fixed point of H, or alternatively as finding

the zero of $K(x_n) = H(x_n) - x_n$, where we assume that K has continuous first partial derivatives and a nonsingular Jacobian in a neighborhood of its single zero. It is on this root-finding problem that we apply quasi-Newton (QN) acceleration.

Quasi-Newton acceleration

1. Startup: Take initial values $x_2^1, x_3^1, \ldots x_n^1$. Set $s = 1$.
2. Loop until convergence:
   2.1. Solve $f_1(x_1, x_2^s, \ldots, x_n^s) = 0$ for x_1, resulting in x_1^{s+1}.
   2.2. Solve $f_2(x_1^{s+1}, x_2, \ldots, x_n^s) = 0$ for x_2, resulting in x_2^{s+1}.
   ...
   2.n. Solve $f_n(x_1^{s+1}, x_2^{s+1}, \ldots, x_n) = 0$ for x_n, resulting in $H(x_n^s)$.
   2.n+1. Compute an approximate Jacobian $\hat{K}_s'$ of K (see below)
   2.n+2. $x_n^{s+1} = x_n^s - (\hat{K}_s')^{-1} K(x_n^s)$
   2.n+3. Set $s = s + 1$.

Alternatively a slightly different quasi-Newton step $x_n^{s+1} = x_n^s - \hat{M}_s' K(x_n^s)$ can be used. Here M_s' serves as an approximation to the inverse of the Jacobian at step s, whereas $\hat{K}_s'$ is an approximation of the Jacobian itself. We will designate methods that approximate the Jacobian as Type I methods, and methods that approximate the inverse Jacobian as Type II methods [11].

3 Different choices of quasi-Newton methods

We define $\delta x_s = x_n^{s+1} - x_n^s$, $\delta K_s = K(x_n^{s+1}) - K(x_n^s)$ and $\{\iota_j; j = 1, \ldots, m_n\}$ as the canonical basis for $\mathbb{R}^{m_n}$ and $\langle \cdot, \cdot \rangle$ as the standard Euclidean scalar product.

3.1 Non-Linear Gauss-Seidel (GS)

This method (also called, among others, "Iterative Substructuring Method" or "Picard iteration") is nothing else than the fixed-point iteration described at the beginning of §2. It is seldom considered to be a quasi-Newton method, but can take this form if we set $(\hat{K}_{s+1}')^{-1} = -I$ [22].

3.2 Aitken's δ^2 method ($A\delta^2$)

Aitken's δ^2 method [1] is a relaxation method and as such is again seldom seen as a quasi-Newton method, but it can take its form if we define $(\hat{K}_{s+1}')^{-1} = -\frac{1}{\omega_{s+1}} I$ with

$$\omega_{s+1} = -\omega_s \frac{\langle K(x_n^s), K(x_n^{s+1}) - K(x_n^s) \rangle}{\langle K(x_n^{s+1}) - K(x_n^s), K(x_n^{s+1}) - K(x_n^s) \rangle}. \tag{3}$$

3.3 Broyden's Good Method (BG), Bad Method (BB) and the Switched Broyden method (SB)

1. Broyden's first (or "good") method is a quasi-Newton method that is part of the family of Least Change Secant Update (LCSU) methods [4, 5, 8, 9, 12], where the approximate Jacobian $\hat{K}'_{s+1}$ is chosen as the solution of $\min\{\|\hat{K}' - \hat{K}'_s\|_{Fr}\}$, s.t. $\hat{K}'\delta x_s = \delta K_s$, which leads to the following rank-one update:

$$\hat{K}'_{s+1} = \hat{K}'_s + \frac{(\delta K_s - \hat{K}'_s \delta x_s)\delta x_s^T}{\langle \delta x_s, \delta x_s \rangle} \quad \text{or} \tag{4}$$

$$(\hat{K}'_{s+1})^{-1} = (\hat{K}'_s)^{-1} + \frac{(\delta x_s - (\hat{K}'_s)^{-1}\delta K_s)\delta x_s^T (\hat{K}'_s)^{-1}}{\langle \delta x_s, (\hat{K}'_s)^{-1}\delta K_s \rangle}. \tag{5}$$

   $(\hat{K}'_1)^{-1}$ is typically set to be $-I$, i.e. the first iteration is a GS iteration.
2. Broyden's second (or "bad") method is a quasi-Newton method that uses an approximation $\hat{M}'$ of the inverse Jacobian. It is also part of the family of LCSU methods [4, 9, 12], where $\hat{M}'_{s+1}$ is chosen as the solution of $\min\{\|\hat{M}' - \hat{M}'_s\|_{Fr}\}$, s.t. $\hat{M}'\delta K_s = \delta x_s$,, which leads to the following rank- one update:

$$\hat{M}'_{s+1} = \hat{M}'_s + \frac{(\delta x_s - \hat{M}'_s \delta K_s)\delta K_s^T}{\langle \delta K_s, \delta K_s \rangle} \quad \text{or} \tag{6}$$

$$\left(\hat{M}'_{s+1}\right)^{-1} = \left(\hat{M}'_s\right)^{-1} + \frac{\left(\delta K_s - \left(\hat{M}'_s\right)^{-1}\delta x_s\right)\delta K_s^T \left(\hat{M}'_s\right)^{-1}}{\langle \delta K_s, \left(\hat{M}'_s\right)^{-1}\delta x_s \rangle}. \tag{7}$$

   Again, typically $\hat{M}'_1 = -I$ is chosen.
3. Broyden himself [4] admitted that the "bad" formulation of his algorithm didn't function properly. The reasons for the "good" or "bad" behavior are not well understood, and in some instances the bad method actually outperforms the good method. For this reason we follow an idea suggested in [24] that avoids the need to choose between the two methods and create a switched version of BG/BB (called "SB") based on the following reasoning: as both BG and BB are secant methods we have $\hat{K}'_s \delta x_{s-1} = \delta K_{s-1}$ and $\left(\hat{M}'_s\right)^{-1}\delta x_{s-1} = \delta K_{s-1}$. Using (4) and (7), we get

$$\hat{K}'_{s+1}\delta x_{s-1} - \delta K_{s-1} = \frac{(\delta K_s - \hat{K}'_s \delta x_s)\delta x_s^T}{\langle \delta x_s, \delta x_s \rangle}\delta x_{s-1} \tag{8}$$

$$\left(\hat{M}'_{s+1}\right)^{-1}\delta x_{s-1} - \delta K_{s-1} = \frac{\left(\delta K_s - \left(\hat{M}'_s\right)^{-1}\delta x_s\right)\delta K_s^T \left(\hat{M}'_s\right)^{-1}}{\langle \delta K_s, \left(\hat{M}'_s\right)^{-1}\delta x_s \rangle}\delta x_{s-1}. \tag{9}$$

   Equations (8) and (9) can be considered to be a secant error at the next approximation with respect to the previous iterates. Thus, BG has a smaller error than BB when

$$\frac{\left|\delta x_s^T \delta x_{s-1}\right|}{\langle \delta x_s, \delta x_s \rangle} < \frac{\left|\delta K_s^T \delta K_{s-1}\right|}{\left|\langle \delta K_s, \left(\hat{M}'_s\right)^{-1}\delta x_s \rangle\right|}. \tag{10}$$

The same reasoning can be built-up starting from (5) and (7). We then obtain for the switching condition:

$$\frac{|\delta x_s^T \delta x_{s-1}|}{|\langle \delta x_s, (\hat{K}'_s)^{-1}\delta K_s\rangle|} < \frac{|\delta K_s^T \delta K_{s-1}|}{\langle \delta K_s, \delta K_s\rangle}. \tag{11}$$

To avoid situations where (10) and (11) contradict, we will only use the latter. If (11) is met then the BG-update will be applied, otherwise the BB-update.

3.4 Column-Updating Method (CU), Inverse Column-Updating Method (ICU) and Switched Column-Updating Method (SCU)

1. The Column-Updating method is a quasi-Newton method that was introduced by Martinez [25, 27, 28]. The rank-one update of this method is such that the column of the approximate Jacobian corresponding to the largest coordinate of the latest increment δx_s is replaced in order to satisfy the secant equation $\hat{K}'\delta x_s = \delta K_s$ at each iteration. This results in:
$$(\hat{K}'_{s+1})^{-1} = (\hat{K}'_s)^{-1} + \frac{(\delta x_s - (\hat{K}'_s)^{-1}\delta K_s)\iota_{j_{K,s}}^T (\hat{K}'_s)^{-1}}{\langle \iota_{j_{K,s}}, (\hat{K}'_s)^{-1}\delta K_s\rangle} \tag{12}$$
where $\iota_{j_{K,s}}$ is chosen such that $j_{K,s} = \text{Argmax}\{|\langle \iota_j, \delta x_s\rangle|; j = 1,\ldots,m_n\}$.
$(\hat{K}'_1)^{-1}$ is typically set to be $-I$,
2. The Inverse Column-Updating method (ICU) is a quasi-Newton method that was introduced by Martinez and Zambaldi [23, 26]. It uses a rank-one update such that the column of the approximation of the inverse of the Jacobian corresponding to the largest coordinate of δK_s is replaced in order to satisfy the secant equation $\hat{M}'\delta K_s = \delta x_s$ at each iteration. This results in:
$$\hat{M}'_{s+1} = \hat{M}'_s + \frac{(\delta x_s - \hat{M}'_s \delta K_s)\iota_{j_{M,s}}^T}{\langle \iota_{j_{M,s}}, \delta K_s\rangle}, \tag{13}$$
where $\iota_{j_{M,s}}$ is chosen such that $j_{M,s} = \text{Argmax}\{|\langle \iota_j, \delta K_s\rangle|; j = 1,\ldots,m_n\}$. Again, typically $\hat{M}'_1 = -I$ is chosen.
3. As far as the authors are aware, the idea behind SB has not yet been applied to CU and ICU, despite being straightforward. A similar reasoning as for SB gives the following switching condition:
$$\frac{|\iota_{j_{K,s}}^T \delta x_{s-1}|}{|\langle \iota_{j_{K,s}}, (\hat{K}'_s)^{-1}\delta K_s\rangle|} < \frac{|\iota_{j_{M,s}}^T \delta K_{s-1}|}{\left|\langle \iota_{j_{M,s}}, \delta K_s\rangle\right|}. \tag{14}$$
If this condition is satisfied, then the CU-update is used, otherwise the ICU-update.

3.5 Quasi-Newton Least Squares (QN-LS) and Quasi-Newton Inverse Least Squares (QN-ILS)

1. In the Quasi-Newton Least Squares Method (QN-LS) [14, 16, 17, 18, 30, 31] the approximate Jacobian $\hat{K}'_{s+1}$ is chosen as the solution of $\min\{\|\hat{K}' - \hat{K}'_s\|_{Fr}\}$ s.t. $\hat{K}'\delta x_i = \delta K_i$, $\forall i \in \{1,\ldots,s\}$, which leads to the following rank-one update:

$$(\hat{K}'_{s+1})^{-1} = (\hat{K}'_s)^{-1} + \frac{(\delta x_s - (\hat{K}'_s)^{-1}\delta K_s)\left((I - \mathscr{L}_s\mathscr{L}_s^T)\delta x_s\right)^T (\hat{K}'_s)^{-1}}{\left\langle (I - \mathscr{L}_s\mathscr{L}_s^T)\delta x_s, (\hat{K}'_s)^{-1}\delta K_s \right\rangle}, \quad (15)$$

   with $\mathscr{L}_s = [\bar{L}_1| \ldots |\bar{L}_s]$, where $\bar{L}_k$ is the k^{th} left singular vector of $V_s = [v_1^s \ldots v_{s-1}^s]$, with $v_i^s = x_n^s - x_n^i$ $(i = 1,\ldots,s-1)$. $(\hat{K}'_1)^{-1}$ is typically set to be $-I$. As $\mathscr{L}_1$ doesn't exist we replace it by the zero matrix.
2. The Quasi-Newton Inverse Least Squares Method (QN-ILS) [7, 15] is similar to QN-LS but constructs an approximation to the inverse Jacobian. The approximate Jacobian $\hat{M}'_{s+1}$ is chosen as the solution of $\min\{\|\hat{M}' - \hat{M}'_s\|_{Fr}\}$, s.t. $\hat{M}'\delta K_i = \delta x_i$, $\forall i \in \{1,\ldots,s\}$, which leads to the following rank-one update:

$$\hat{M}'_{s+1} = \hat{M}'_s + \frac{(\delta x_s - \hat{M}'_s\delta K_s)\left((I - \tilde{\mathscr{L}}_s(\tilde{\mathscr{L}}_s)^T)\delta K_s\right)^T}{\langle (I - \tilde{\mathscr{L}}_s(\tilde{\mathscr{L}}_s)^T)\delta K_s, \delta K_s\rangle}, \quad (16)$$

   with $\tilde{\mathscr{L}}_s = [\tilde{\bar{L}}_1|\tilde{\bar{L}}_2| \ldots |\tilde{\bar{L}}_s]$, where $\tilde{\bar{L}}_k$ is the k^{th} left singular vector of $\tilde{V}_s = [\tilde{v}_1^s \ldots \tilde{v}_{s-1}^s]$, with $\tilde{v}_i^s = K(x_n^s) - K(x_n^i)$ $(i = 1,\ldots,s-1)$. $\hat{M}'_1$ is typically set to be $-I$. As $\tilde{\mathscr{L}}_1$ does not exist we replace it by the zero matrix.
3. A similar reasoning as for SB gives the following switching condition:

$$\frac{|\langle\left((I - \mathscr{L}_s\mathscr{L}_s^T)\delta x_s\right)^T \delta x_{s-1}|}{|\langle (I - \mathscr{L}_s\mathscr{L}_s^T)\delta x_s, (\hat{K}'_s)^{-1}\delta K_s\rangle|} < \frac{|\left((I - \tilde{\mathscr{L}}_s(\tilde{\mathscr{L}}_s)^T)\delta K_s\right)^T \delta K_{s-1}|}{\left|\langle (I - \tilde{\mathscr{L}}_s(\tilde{\mathscr{L}}_s)^T)\delta K_s, \delta K_s\rangle\right|} \quad (17)$$

   In this expression we observe that $\left((I - \mathscr{L}_s\mathscr{L}_s^T)\delta x_s\right) \perp V_s$ and $\delta x_{s-1} \in \mathscr{R}(V_s)$. Consequently $\left((I - \mathscr{L}_s\mathscr{L}_s^T)\delta x_s\right)^T \delta x_{s-1} = 0$.
   Similarly $\left((I - \tilde{\mathscr{L}}_s(\tilde{\mathscr{L}}_s)^T)\delta K_s\right)^T \delta K_{s-1} = 0$. This switch will thus never be triggered and for that reason this Switched Quasi-Newton Least Squares method is rejected.

3.6 Non-linear Eirola-Nevanlinna Type I Method (EN1), Type II Method (EN2) and Switched Eirola-Nevanlinna Method (SEN)

1. It is clear from the different update formulas for the approximate (inverse) Jacobian of all the previous methods that they can only be applied starting with $\hat{K}'_2$. In other words, $\hat{K}'_1$ needs to be chosen. Conventionally, this is set to be equal to $-I$. Likewise, for $A\delta^2$, ω_1 needs to be chosen and is set to 1. As a result all of the methods given above will have an identical first iteration, i.e. $x_n^2 = H(x_n^1)$.
   The nonlinear Eirola-Nevanlinna (EN) was proposed by [33] as the nonlinear

counterpart to the linear EN algorithm [10] and is different as it computes $\hat{K}'_1$ based on a virtual choice of $\hat{K}'_0$, set to be equal to $-I$ (which can also be interpreted as setting the initial approximation of the Jacobian of H as zero), which is used to create a first approximation $\hat{K}'_1$. The method is given by

$$(\hat{K}'_{s+1})^{-1} = (\hat{K}'_s)^{-1} + \frac{(p_s - (\hat{K}'_s)^{-1} q_s) p_s^T (\hat{K}'_s)^{-1}}{\langle p_s, (\hat{K}'_{s-1})^{-1} q_s \rangle}, \tag{18}$$

where $p_s = -(\hat{K}'_s)^{-1} K(x_n^{s+1})$ and $q_s = K(x_n^{s+1} + p_s) - K(x_n^{s+1})$. Note that the EN algorithm requires two calls of K (or H) per iteration.

2. Eirola and Nevanlinna did not propose a Type II method, but by generalisation this can be written as [11]:

$$\hat{M}'_{s+1} = \hat{M}'_s + \frac{(p_s - (\hat{K}'_s)^{-1} q_s) q_s^T}{\langle q_s, q_s \rangle}, \tag{19}$$

where p_s and q_s are defined as in the EN1 method.

3. As far as the authors are aware, the idea behind SB has not been applied to EN1 and EN2. The switching condition now becomes

$$\frac{\left| p_s^T \delta x_{s-1} \right|}{\left| \langle p_s, (\hat{K}'_{s-1})^{-1} q_s \rangle \right|} < \frac{\left| q_s^T \delta K_{s-1} \right|}{\langle q_s, q_s \rangle}. \tag{20}$$

When this condition is satisfied, then EN1 (equation (18)) is used, otherwise EN2 (equation (19)).

4 Re-use of previous information

When the problem is time-dependent, and assuming that the changes over one time-step are relatively small, then the approximate (inverse) Jacobian of the previous time-step might be a relatively good initial guess for the approximate (inverse) Jacobian at the next time-step.

One of the primary aims of this investigation is to compare the performance of the various QN methods when the Jacobian, at the start of each new time step, $k+1$, is either reset to $-I$, such that $\left(\hat{K}'_1\right)^{-1}_{k+1} = -I$, (which we will call "Jacobian reset") or set equal to the final approximation from the previous time step, *i.e.* $\left(\hat{K}'_1\right)^{-1}_{k+1} = \left(\hat{K}'_s\right)^{-1}_{k}$, where k indicates the time step counter (which we will call "Jacobian re-use"). A similar procedure is used for quasi-Newton methods using $\hat{M}'_1$ and for Aitken's method.

5 Fluid structure interaction test-cases

Quasi-Newton methods have received significant attention in recent years within the field of partitioned fluid-structure interactions (FSI). In this section we aim to investigate the various QN methods when applied to a number of incompressible, transient, FSI benchmark problems by coupling OpenFOAM for the fluid flow solution and Calculix for the structural deformation. A relaxation factor of $\omega = 0.001$ is used for the first iteration of the first time step, or for the first iteration in every

time step whenever the Jacobian is reset such that

$$x_n^2 = x_n^1 + \omega \delta x_1. \tag{21}$$

This is done to avoid the possibility of an excessively large first displacement guess, which can lead to divergence for strongly coupled problems. x_n^1 at the start of each new time step is set equal to the final converged solution from the previous time step.

5.1 Dynamic piston-channel problem

The piston-channel test problem layout is shown in Figure 1(a), and consists of a 10m long fluid domain which is forced out of the channle by an accelerating unit by unit moving block. The problem is a surprisingly difficult problem to solve when using partitioned solution schemes, and has been investigated in a number of publications (see for example [2, 6]). The coupling strength is sufficiently strong that simple fixed-point iterations, such as Gauss-Seidel iterations, are insufficient to guarantee convergence.

While the test case is intrinsically a one dimensional problem, it is modelled here in three dimensions, with a fluid domain discretised using 10 linear elements and a single linear structural element. The fluid density and viscosity are 1kg/m^3 and 1.0kg/(m s), respectively, where the solid is described by a linear elastic material with a Young's modulus of $E = 10.0$Pa with a zero density and Poisson's ratio. A slip boundary condition is applied to the fluid wall boundaries and the velocity is prescribed as $u(t) = 0.2t$ on the left side of the block (see Figure 1(a)). The simulation is solved here using time step sizes of $\Delta t = 0.02$s for a convergence tolerance of $\varepsilon = \frac{||K(x_n^s)||}{\sqrt{m_n}} = 10^{-8}$. The interface displacement and velocity, along with a 1D-solution is shown in Figure 1(b). The 1D-model is found by simplifying the problem to a 1D mass-spring system, where the solid piston represents a linear spring and the fluid domain a variable mass (see [2] for more information on the simplified 1D-model). The accuracy of the FSI simulation can naturally be improved by increasing the mesh resolution or decreasing the time step size.

The average number of iterations required to reach convergence for each of the QN methods is summarised in Table 1. The results highlight the surprising complexity of the test problem, with both GS and Aitken failing to provide convergent results. The performance of all the other QN methods is virtually identical, with a significant improvement offered by retaining the Jacobian at the start of each new time step.

5.2 Dam break with an elastic obstacle

The dam break problem consists of a collapsing column of water striking an elastic baffle, which has previously been analysed in [3, 32]. The problem layout is shown in Figure 2(a) with a plot of the beam tip displacement shown in Figure 2(b). The FSI simulation is performed using a time step size of $\Delta t = 0.001$s, 3670 linear fluid elements, and 14 quadratic, full integration finite-elements. The mean number

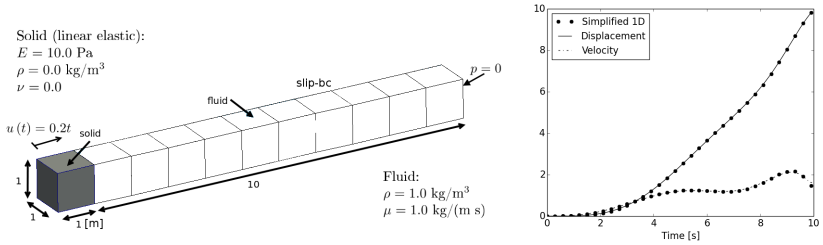

Fig. 1: (a) Piston-channel test problem layout. (b) The interface displacement and velocity for a time step size of $\Delta t = 0.02$s, compared to the simplified 1D mass-spring system.

Table 1: The mean number of iterations required to reach convergence for the piston-channel test problem. Failure to converge is indicated by div() where the time step at which failure occurred is indicated in brackets; top performing method highlighted in bold.

	Jacobian re-use	Jacobian reset
GS	N/A	div(1)
Aδ^2	div(1)	div(1)
BG	3.98	4.73
BB	3.98	**4.72**
SB	3.97	4.73
CU	3.98	div(75)
ICU	3.98	4.73
SCU	3.98	4.73
QN-LS	4.01	div(52)
QN-ILS	4.01	div(438)
EN1	3.91	div(108)
EN2	3.90	div(108)
SEN	**3.89**	div(108)

of iterations to reach convergence for each of the QN methods is summarised in Table 2. The QN-LS family of methods is the top performing family of methods, with significant benefit seen in retaining the Jacobian from the previous time steps.

5.3 Wave propagation in a three dimensional elastic tube

The 3D flexible tube problem was originally proposed in [13], inspired by the type of flow encountered in haemodynamics. The density ratios of the fluid and solid are near unity, which in conjunction with internal incompressible flow results in a very strongly coupled FSI problem that has received much attention in literature [2, 7, 13].

The problem consists of a flexible tube of length $l = 5$cm, with an inner and outer radius of $r_i = 0.5cm$ and $r_0 = 0.6$cm, respectively. The flexible tube is modelled using a St. Venant-Kirchoff material model, with a Young's modulus of $E = 3 \times 10^6$ dynes/cm^2, density $\rho = 1.2$ g/cm^3 and Poisson's ratio of 0.3, where the fluid flow has a density of $\rho = 1.0$ g/cm^3 and a viscosity of $\mu = 0.03$ poise. The problem

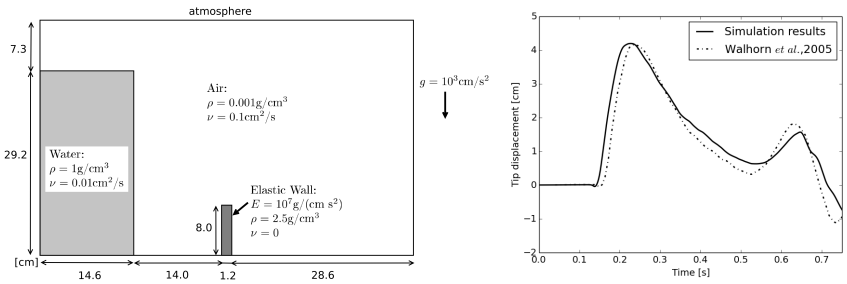

Fig. 2: (a) Dam break with elastic obstacle problem description. (b) Beam tip displacement results for the dam break test problem, compared to the results reproduced from [32].

Table 2: Mean number of iterations required by the QN methods for the dam break test problem. Failure to converge is indicated by div() where the time step at which failure occurred is indicated in brackets; top performing method highlighted in bold.

	Jacobian re-use	Jacobian reset
GS	N/A	23.42
Aδ^2	div(279)	11.45
BG	4.17	7.09
BB	4.32	7.10
SB	4.14	7.07
CU	5.62	7.10
ICU	5.84	div(279)
SCU	5.80	7.29
QN-LS	**3.90**	6.52
QN-ILS	div(279)	**6.49**
EN1	4.37	8.05
EN2	4.58	8.05
SEN	4.58	8.04

is modelled using 600 twenty-noded quadratic solid elements coupled with 6000 linear fluid flow elements resulting in an interface Jacobian size of $m_n = 1880$. The tube walls are fixed on both ends, and a smoothly varying pressure in the form of $p(t) = 1.3332 \times 10^4 \left(\sin\left(2\pi t/0.003 + 1.5\pi\right) + 1\right)/2$ is applied at the inlet over the first 0.003 seconds. The time step size for the simulation is $\Delta t = 0.0001$s with a convergence tolerance $\varepsilon = \frac{||K(x_n^s)||}{\sqrt{m_n}} = 10^{-8}$. The resulting pressure pulse propagation is illustratively shown for different time steps in Figure 3.

The mean number of iterations to reach convergence is summarised in Table 3 for the various QN methods, and the number of coupling iterations required per time step, for each of the family of methods, is shown in Figure 4. The QN-LS method was once again the top performing method.

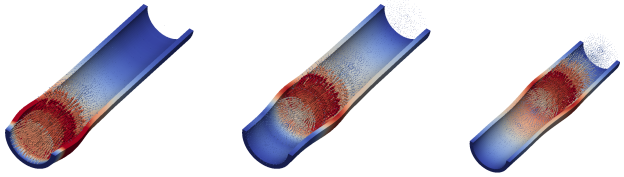

Fig. 3: Pressure pulse propagation at 0.003s, 0.005s and 0.008s (where the wall displacement is magnified by a factor 10).

Table 3: Comparison of the mean number of iterations required to reach convergence for the 3D flexible tube problem. Failure to converge is indicated by div() with the time step at which failure occurred in brackets; top performing method highlighted in bold.

	Jacobian re-use	Jacobian reset
GS	N/A	div(1)
Aδ^2	div(1)	div(1)
BG	6.51	18.02
BB	7.68	div(47)
SB	6.75	18.02
CU	15.63	div(7)
ICU	div(14)	div(12)
SCU	13.89	div(10)
QN-LS	**5.64**	15.94
QN-ILS	5.95	**14.47**
EN1	6.91	div(52)
EN2	10.56	div(11)
SEN	7.37	div(2)

5.4 2D flexible beam

The selected test case is a fluid-structure interaction problem consisting of flow around a fixed cylinder with an attached flexible beam. The beam undergoes large deformations induced by oscillating vortices formed by flow around the circular bluff body. The problem was first proposed by Turek *et al.* [29], and has received substantial numerical verification. The problem layout and material properties are provided in Figure 5(a).

A parabolic inlet boundary condition, with mean flow velocity of $\bar{U} = 1$m/s is slowly ramped up for $t < 0.5$s. A snapshot of the beam tip displacement is illustratively shown in Figure 5(b). The convergence behavior for the various QN methods is summarized in Table 4.

Once again the QN-LS method is the top performing method. Besides for Broyden's method, the switching strategies provided no benefit for any of the other families of QN methods.

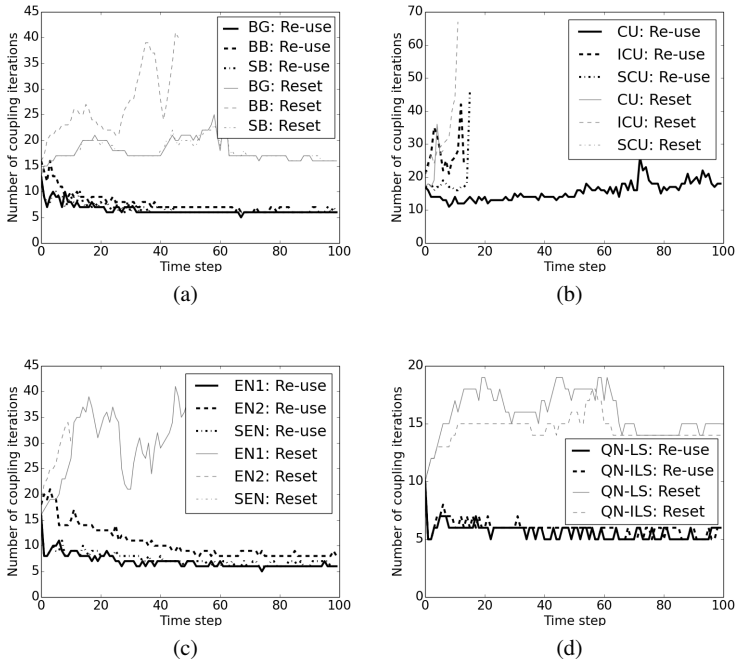

Fig. 4: A comparison of the number of coupling iterations required to reach convergence for the 3D flexible tube test problem.

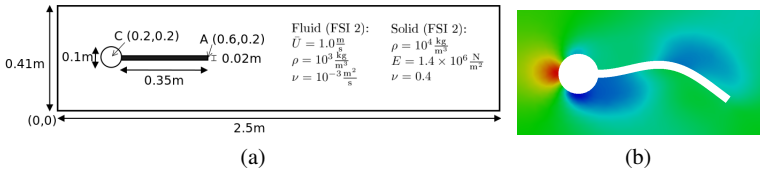

Fig. 5: (a) Flexible beam problem description and (b) snapshot of beam displacement and pressure contours at 8.7 seconds.

6 Other application: Simplified model of plasma heating by RF waves in a plasma

Quasi-Newton acceleration can also be applied to problems outside the field of fluid-structure interaction. The model that we present here is a simplified version of the set of codes commonly used to describe the steady state of plasma heating by radio frequency waves in a tokamak plasma, but still retains all the characteristics needed to validate our ideas as explained in §1. In an abstract form the governing equations can be written as:

Table 4: Comparison of the mean number of iterations required to reach convergence for the 2D flexible beam problem. Failure to converge is indicated by div() with the time step at which failure occurred in brackets; top performing method highlighted in bold.

	Jacobian re-use	Jacobian reset
GS	N/A	div(1)
Aδ^2	10.31	14.00
BG	4.02	7.48
BB	4.25	7.42
SB	3.98	7.48
CU	5.26	7.99
ICU	5.39	7.94
SCU	6.34	8.05
QN-LS	**3.86**	6.69
QN-ILS	3.93	**6.63**
EN1	4.61	8.47
EN2	4.54	8.46
SEN	4.54	8.47

$$\begin{cases} f_1(y_1,y_2,y_3,y_4,y_5,y_6,y_7) & = 0 \\ f_2(y_1,y_2,y_3,y_4,y_6,y_8,y_9,y_{10}) & = 0 \\ f_3(y_1,y_2,y_3,y_5,y_6,y_7,y_8,y_9,y_{10}) & = 0, \end{cases} \tag{22}$$

where f_1 represents a simplified 1-component wave equation, f_2 a simplified Fokker-Planck equation and f_3 a simplified 1D diffusion equation. y_1, y_2 and y_3 are the temperature profile of the species (majority ions, minority ions and electrons); y_4 the effective temperature of the minority ions; y_5, y_6 and y_7 are the power density profiles of the wave damping onto the species; y_8, y_9 and y_{10} are the collisionally redistributed minority power density fraction profiles onto the species. More details can be found in [19]. As the system in (22) does not satisfy condition (2), (y_1,y_2,y_3,y_4) need to be grouped as x_3, as explained in [20]. A convergence criterium of $\frac{\|x_n^{s+1}-x_n^s\|}{\|x_n^s\|} \leq 10^{-7}$ is used. The results are shown in table 5 for different values of launched power P and diffusion coefficient κ. The QN methods clearly outperform G-S and Aitken's method. Of all the QN methods, the Least Squares methods most often give the best results (with a slight edge for QN-ILS over QN-LS), followed by the Broyden methods. While not equivocally so, the switching strategy often (slightly) improves the convergence speed of the underlying methods.

7 Conclusion

We have tested a wide variety of acceleration techniques on different multi-physics problems that are written as a fixed-point problem. While the choice of the best method remains problem dependent, it is clear that the best choice is the class of quasi-Newton methods, of which, more often than not, the Least Squares methods come out on top. Re-using the Jacobian of all the QN methods at the beginning of the iterations of the next time step results in important reductions in the required number of iterations. With a few exceptions, a switching strategy, that hasn't drawn much attention in the past, is shown to offer only a slight boost of performance in

Table 5: Simplified tokamak model. Number of function calls (of H) needed to reach convergence for various values of κ. P=2MW (left) and P=5MW (right). "div" = divergence or no convergence after 100 iterations. The top performing method is highlighted in bold.

P=2MW.										
$\kappa(\cdot 10^{-2})$	3.5	4.0	4.5	5.0	7.5	10	25	50	75	100
G-S	div	div	95	88	67	55	30	19	14	14
$A\delta^2$	63	52	52	49	48	33	24	17	14	15
BG	div	33	30	29	28	23	18	14	**13**	**12**
BB	30	30	29	30	28	23	18	14	**13**	**12**
SB	31	div	28	25	26	21	18	14	**13**	**12**
CU	37	div	div	37	32	22	17	**12**	**13**	**12**
ICU	34	36	33	32	30	24	18	14	14	**12**
SCU	33	38	28	29	30	23	18	14	**13**	**12**
QN-LS	div	**24**	24	23	**21**	20	**16**	14	**13**	**12**
QN-ILS	**25**	25	**23**	**22**	**21**	**19**	**16**	14	**13**	**12**
EN1	36	34	40	30	30	24	22	16	16	16
EN2	34	36	34	34	32	26	22	16	16	16
SEN	35	37	34	34	30	24	22	16	16	16

P=5MW.						
$\kappa(\cdot 10^{-2})$	7.5	10	25	50	75	100
G-S	67	52	28	19	16	15
$A\delta^2$	39	36	23	19	16	16
BG	26	24	19	17	**14**	13
BB	25	24	20	17	**14**	13
SB	24	23	19	17	15	13
CU	36	div	21	17	**14**	**12**
ICU	27	30	21	18	15	14
SCU	26	27	21	17	**14**	14
QN-LS	22	23	**18**	**16**	**14**	13
QN-ILS	**21**	**22**	**18**	**16**	**14**	13
EN1	28	28	24	20	16	16
EN2	28	28	24	20	16	16
SEN	28	28	25	19	16	16

exchange for a negligeable penalty in complexity. The class of Eirola-Nevanlinna methods, which are among the lesser known QN methods, have not shown their worth, and in the authors' opinion do not seem to warrant the complexity that they entail.

References

1. A.C. Aitken, On Bernouilli's numerical solution of algebraic equations. Proc. Roy. Soc. Edinb. **46**, pp. 289–305 (1926)
2. A.E.J. Bogaers, S. Kok, B.D. Reddy, T. Franz, Quasi-Newton methods for implicit black-box FSI coupling *Computer Methods in Applied Mechanics and Engineering*, **279**, pp. 113–132 (2014)
3. A.E.J. Bogaers, S. Kok, B.D. Reddy, T. Franz, An evaluation of quasi-Newton methods for application to FSI problems involving free surface flow and solid body contact *Computers & Structures*, **173**, pp. 71–83 (2016)
4. C.G. Broyden, A class of methods for solving nonlinear simultaneous equations. *Math. Comp.* **19**, pp. 577–593 (1965)
5. C.G. Broyden, Quasi-Newton methods and their applications to function minimization. *Math. Comp.* **21**, pp. 368–381 (1967)
6. J. Degroote, K.-J. Bathe, J. Vierendeels, Performance of a new partitioned procedure versus a monolithic procedure in fluid–structure interaction *Computers & Structures*, **87**/11, pp. 793–801 (2009)
7. J. Degroote, R. Haelterman, S. Annerel, P. Bruggeman, J. Vierendeels, Performance of partitioned procedures in fluid–structure interaction *Computers & structures*, **88**/7, pp. 446–457 (2010)
8. J.E. Dennis, J.J. Moré, Quasi-Newton methods: motivation and theory. *SIAM Rev.* **19**, pp. 46–89 (1977)

9. J.E. Dennis, R.B. Schnabel, Least Change Secant Updates for quasi-Newton methods. *SIAM Rev.* **21**, pp. 443–459 (1979)
10. T. Eirola, O. Nevanlinna, Accelerating with rank-one updates. *Linear Algebra Appl.* **121**, pp. 511–520 (1989)
11. H.-R. Fang, Y. Saad, Two classes of multisecant methods for nonlinear acceleration. *Numerical Linear Algebra with Applications*, **16/**3, pp. 197–221 (2009).
12. A. Friedlander, M.A. Gomes-Ruggiero, D.N. Kozakevich, J.M. Martinez, S.A. dos Santos, Solving nonlinear systems of equations by means of quasi-Newton methods with a nonmonotone strategy. *Optim. Methods Softw.* **8**, pp. 25–51 (1997)
13. J.-F. Gerbeau, M. Vidrascu et al., A Quasi-Newton Algorithm Based on a Reduced Model for Fluid-Structure Interaction Problems in Blood Flows *ESAIM: Mathematical Modelling and Numerical Analysis* **37/**4, pp. 631–647 (2003)
14. R. Haelterman, J. Degroote, D. Van Heule, J. Vierendeels, The Quasi-Newton Least Squares Method: A New and Fast Secant Method Analyzed for Linear Systems, *SIAM J. Numer. Anal.*, **47**(3), pp. 2347-2368 (2009)
15. R. Haelterman, J. Degroote, D. Van Heule, J. Vierendeels, On the Similarities Between the Quasi-Newton Inverse Least Squares Method and GMRes, *SIAM J. Numer. Anal.*, **47**(6), pp. 4660–4679 (2010)
16. R. Haelterman, J. Petit, H. Bruyninckx, J. Vierendeels, On the non-singularity of the quasi-Newton-least squares method, *Journal of Computational and Applied Mathematics*, **257**, pp. 129-131 (2014)
17. R. Haelterman, B. Lauwens, F. Van Utterbeeck, H. Bruyninckx, J. Vierendeels, On the similarities between the quasi-Newton least squares method and GMRes, *Journal of Computational and Applied Mathematics* **273**, pp. 25–28 (2015)
18. R. Haelterman, B. Lauwens, H. Bruyninckx, J. Petit, Equivalence of QNLS and BQNLS for affine problems, *Journal of Computational and Applied Mathematics* **278**, pp 48–51 (2015)
19. R. Haelterman, D. Van Eester, D. Verleyen, Accelerating the solution of a physics model inside a Tokamak using the (Inverse) Column Updating Method, *Journal of Computational and Applied Mathematics* **279**, pp. 133–144 (2015)
20. R. Haelterman, D. Van Eester, S. Cracana, Does Anderson Always Accelerate Picard ?, *14th Copper Mountain Conference on Iterative Methods*, Copper Mountain, USA (2016).
21. R. Haelterman, A. Bogaers, J. Degroote, S. Cracana, Coupling of Partitioned Physics Codes with Quasi-Newton Methods, Lecture Notes in Engineering and Computer Science: Proceedings of The International MultiConference of Engineers and Computer Scientists 2017, 15-17 March, 2017, Hong Kong, pp. 750–755 (2017)
22. C.T. Kelley, Iterative methods for linear and nonlinear equations. *Frontiers Appl. Math.*, SIAM, Philadelphia (1995)
23. V.L.R. Lopes, J.M. Martinez, Convergence properties of the Inverse Column-Updating Method. *Optim. Methods Softw.* **6**, pp. 127–144 (1995)
24. J.M. Martinez, L.S. Ochi, Sobre Dois Metodos de Broyden. *Mat. Apl. Comput.* **1/**2, pp. 135–143 (1982)
25. J.M. Martinez, A quasi-Newton method with modification of one column per iteration. *Computing* **33**, pp. 353–362 (1984)
26. J.M. Martinez, M.C. Zambaldi, An Inverse Column-Updating Method for solving large-scale nonlinear systems of equations. *Optim. Methods Softw.* **1**, pp. 129–140 (1992)
27. J.M. Martinez, On the convergence of the column-updating method. *Comp. Appl. Math.* **12/**2, pp. 83–94 (1993)
28. J.M. Martinez, Practical quasi-Newton method for solving nonlinear systems. *J. Comput. Appl. Math.* **124**, pp. 97–122 (2000)
29. S. Turek, J. Hron, Proposal for Numerical Benchmarking of Fluid-Structure Interaction between an Elastic Object and Laminar Incompressible Flow, In: *Fluid-Structure Interaction*, Ed. H.-J. Bungartz and M. Schäfer, Michael, Series "Modelling, Simulation, Optimisation" Vol. **53**, Springer Berlin Heidelberg, ISSN 1439-7358, pp. 371–385 (2006)

30. J. Vierendeels, Implicit coupling of partitioned fluid-structure interaction solvers using reduced-order models, in *Fluid-Structure Interaction, Modelling, Simulation, Optimization*, H.-J. Bungartz, M. Sch afer, (eds.), *Lect. Notes Comput. Sci. Eng.* **53**, Springer, Berlin, pp. 1-18 (2006)
31. J. Vierendeels, L. Lanoye, J. Degroote, P. Verdonck, Implicit coupling of partitioned fluid-structure interaction problems with reduced order models, *Comput. & Structures*, **85**, pp. 970-976 (2007)
32. E. Walhorn, A. Kölke, B. Hübner, D. Dinkler Fluid–structure coupling within a monolithic model involving free surface flows *Computers & structures*, **83**/25, pp. 2100–2111 (2005)
33. U.M. Yang. A family of preconditioned iterative solvers for sparse linear systems. PhD thesis, Dept. of Computer Science, University of Illinois, Urbana-Champaign, 1995.