On the effect of linear algebra implementations in real-time multibody system dynamics

**Manuel González · Francisco González ·
Daniel Dopico · Alberto Luaces**

Abstract   This paper compares the efficiency of multibody system (MBS) dynamic simulation codes that rely on different implementations of linear algebra operations. The dynamics of an N-loop four-bar mechanism has been solved with an index-3 augmented Lagrangian formulation combined with the trapezoidal rule as numerical integrator. Different implementations for this method, both dense and sparse, have been developed, using a number of linear algebra software libraries (including sparse linear equation solvers) and optimized sparse matrix computation strategies. Numerical experiments have been performed in order to measure their performance, as a function of problem size and matrix filling. Results show that optimal implementations can increase the simulation efficiency in a factor of 2–3, compared with our starting classical implementations, and in some topics they disagree with widespread beliefs in MBS dynamics. Finally, advices are provided to select the implementation which delivers the best performance for a certain MBS dynamic simulation.

1 Introduction

Dynamic simulation of multibody systems (MBS) is of great interest for dynamics of machinery, road and rail vehicle design, robotics and biomechanics. Computer simulations performed by MBS simulation tools lead to more reliable, optimized designs and significant reductions in cost and time of the product development cycle. The computational efficiency of these tools is a key issue for two reasons. First, there are some applications, like hardware-in-the-loop settings or human-in-the-loop devices, which cannot be developed unless MBS simulation is performed in real-time. And second, when MBS simulation is used in virtual prototyping, faster simulations allow the design engineer to perform what-if-analyses and optimizations in shorter times, increasing productivity and interaction with the model. Therefore, computational efficiency is an active area of research in MBS, and it holds a relevant position in MBS-related scientific conferences and journals.

A great variety of methods to improve simulation speed have been proposed during the last years [1–3]. Most of these methods base their efficiency improvements on the development of new dynamic formulations. However, although implementation aspects can also play a key factor in the performance of numeric simulations, their effect on real-time multibody system dynamics has not been studied in detail. Some recent contributions have investigated the possibilities of parallel implementations [4], but comprehensive comparisons about serial implementations in MBS dynamics have not been published yet.

M. González (✉) · F. González · D. Dopico · A. Luaces
Escuela Politécnica Superior, Universidad de A Coruña,
Mendizábal s/n, 15403 Ferrol, Spain
e-mail: lolo@cdf.udc.es
URL: http://lim.ii.udc.es

F. González
e-mail: fgonzalez@udc.es

D. Dopico
e-mail: ddopico@udc.es

A. Luaces
e-mail: aluaces@udc.es

Multibody dynamics codes make an intensive use of linear algebra operations. This is especially true in global methods, which use a relative big number of coordinates and constraint equations to define the position of the system; these methods usually lead to $O(N^3)$ algorithms, where N is the number of bodies, and spend around 80% of the CPU time in matrix computations. Topological methods lead to $O(N)$ algorithms due to the reduced size of the involved matrices, and therefore the weight of matrix computations is also reduced. However, if flexible bodies are considered, matrix computations take a significant percentage of simulation time even for topological methods.

As a result, the implementation of linear algebra operations is critical to the efficiency of MBS dynamic simulations. These operations can be grouped into two categories: (a) operations between scalars, vectors and matrices, and (b) solution of linear systems of equations; two additional orthogonal categories can be established based on the data storage format: dense storage or sparse storage. Many efficient implementations for these routines have been made freely available in the last decade. Their performance has been compared in previous works, both in an application-independent context [5–7] and under the perspective of a particular application like Finite Element Analysis [8] or Computational Chemistry [9]. But, as it will be explained in this paper, these studies do not fit the particular features of MBS dynamics, and therefore their conclusions cannot be extrapolated to this field.

The goal of this paper is to compare the efficiency of different implementations of linear algebra operations, and study their effect in the context of MBS dynamic simulation. Results will provide guidelines about which numerical libraries and implementation techniques are more convenient in each case. This information will be very helpful to researchers developing high-performance or real-time multibody simulation codes.

The remainder of the paper is organized as follows: Sect. 2 describes the test problem and the dynamic formulation used in the numerical experiments to compare the efficiency of different implementations; Sects. 3 and 4 present efficient implementations for dense and sparse linear algebra, respectively; Sect. 5 compares the results obtained in Sects. 3 and 4 and extrapolates them to other dynamic formulations; finally, Sect. 6 provides conclusions, advices for efficient implementations and areas of future work.

2 Benchmark setup

In order to study the effect of linear algebra implementations in MBS dynamic simulations, a test problem will be solved with a particular dynamic formulation using different software implementations. A starting implementation will also

be described, since its efficiency will serve as a reference to measure performance improvements.

2.1 Test problem

The selected test problem (Fig. 1) is a 2D one degree-of-freedom assembly of four-bar linkages with N loops, composed by thin rods of 1 m length with a uniformly distributed mass of 1 kg, moving under gravity effects. Initially, the system is in the position shown in Fig. 1, and the velocity of the x-coordinate of point B_0 is +1 m/s. The simulation time is 20 s. This mechanism has been previously used as a benchmark problem for multibody system dynamics [3,10].

2.2 Dynamic formulation

The N-four-bar mechanism has been modeled using planar natural coordinates (global and dependent) [11], leading to $2N+2$ variables (the x and y coordinates of the B points), and $2N+1$ constraints, associated with the constant length condition of the rods. The equations of motion of the whole multibody system are given by the well-known index-3 augmented Lagrangian formulation in the form:

$$\mathbf{M\ddot{q}} + \mathbf{\Phi_q^T}\alpha\mathbf{\Phi} + \mathbf{\Phi_q^T}\lambda^* = \mathbf{Q}$$
$$\lambda_{i+1}^* = \lambda_i^* + \alpha\mathbf{\Phi}_{i+1}, \quad i = 0, 1, 2, \ldots \tag{1}$$

where $\mathbf{M}$ is the mass matrix (constant for the proposed test problem), $\mathbf{\ddot{q}}$ are the accelerations, $\mathbf{\Phi_q}$ the Jacobian matrix of the constraint equations, α the penalty factor, $\mathbf{\Phi}$ the constraints vector, λ^* the Lagrange multipliers and $\mathbf{Q}$ the vector of applied and velocity dependent inertia forces. The Lagrange multipliers for each time-step are obtained from an iteration process, where the value of λ_0^* is equal to the λ^* obtained in the previous time-step.

As integration scheme, the implicit single-step trapezoidal rule has been adopted. The corresponding difference
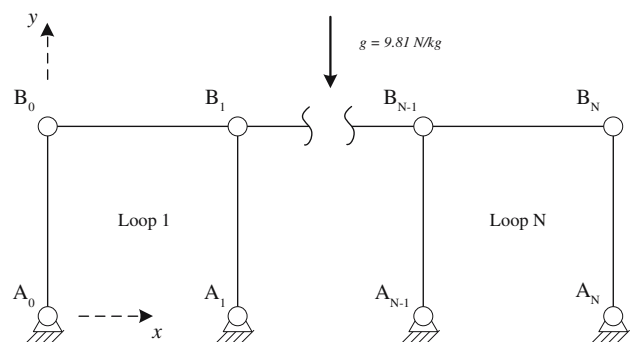

Fig. 1 N-four-bar mechanism

equations in velocities and accelerations are:

$$\dot{\mathbf{q}}_{n+1} = \frac{2}{\Delta t}\mathbf{q}_{n+1} + \hat{\dot{\mathbf{q}}}_n; \quad \hat{\dot{\mathbf{q}}}_n = -\left(\frac{2}{\Delta t}\mathbf{q}_n + \dot{\mathbf{q}}_n\right) \tag{2}$$

$$\ddot{\mathbf{q}}_{n+1} = \frac{4}{\Delta t^2}\mathbf{q}_{n+1} + \hat{\ddot{\mathbf{q}}}_n; \quad \hat{\ddot{\mathbf{q}}}_n = -\left(\frac{4}{\Delta t^2}\mathbf{q}_n + \frac{4}{\Delta t}\dot{\mathbf{q}}_n + \ddot{\mathbf{q}}_n\right)$$

Dynamic equilibrium can be established at time-step $n+1$ by introducing the difference Equation (2) into the equations of motion (1), leading to a nonlinear algebraic system of equations with the dependent positions as unknowns:

$$\mathbf{f}(\mathbf{q}) = \mathbf{M}\mathbf{q_{n+1}} + \frac{\Delta t^2}{4}\mathbf{\Phi}_{\mathbf{q}_{n+1}}^T(\alpha\mathbf{\Phi}_{n+1} + \lambda_{n+1})$$

$$- \frac{\Delta t^2}{4}\mathbf{Q}_{n+1} + \frac{\Delta t^2}{4}\mathbf{M}\hat{\ddot{\mathbf{q}}}_{\mathbf{n}} = 0 \tag{3}$$

Such system, whose size is the number of variables in the model, is solved through the Newton–Raphson iteration

$$\left[\frac{\partial\mathbf{f}(\mathbf{q})}{\partial\mathbf{q}}\right]_i \Delta\mathbf{q_{i+1}} = -\left[\mathbf{f}(\mathbf{q})\right]_i \tag{4}$$

using the approximate tangent matrix (symmetric and positive-definite)

$$\left[\frac{\partial\mathbf{f}(\mathbf{q})}{\partial\mathbf{q}}\right] \cong \mathbf{M} + \frac{\Delta t}{2}\mathbf{C} + \frac{\Delta t^2}{4}\left(\mathbf{\Phi}_{\mathbf{q}}^T\alpha\mathbf{\Phi}_{\mathbf{q}} + \mathbf{K}\right) \tag{5}$$

where $\mathbf{C}$ and $\mathbf{K}$ represent the contribution of damping and elastic forces of the system (which are zero for the test problem). Once convergence is attained into the time-step, the obtained positions $\mathbf{q}_{n+1}$ satisfy the equations of motion (1) and the constraint conditions $\mathbf{\Phi} = 0$, but the corresponding sets of velocities $\dot{\mathbf{q}}^*$ and accelerations $\ddot{\mathbf{q}}^*$ may not satisfy $\dot{\mathbf{\Phi}} = 0$ and $\ddot{\mathbf{\Phi}} = 0$. To achieve this, cleaned velocities $\dot{\mathbf{q}}$ and accelerations $\ddot{\mathbf{q}}$ are obtained by means of mass-damping-stiffness orthogonal projections, reusing the factorization of the tangent matrix:

$$\left[\frac{\partial\mathbf{f}(\mathbf{q})}{\partial\mathbf{q}}\right]\dot{\mathbf{q}} = \left[\mathbf{M} + \frac{\Delta t}{2}\mathbf{C} + \frac{\Delta t^2}{4}\mathbf{K}\right]\dot{\mathbf{q}}^* - \frac{\Delta t^2}{4}\mathbf{\Phi}_{\mathbf{q}}^T\alpha\mathbf{\Phi}_{\mathbf{q}}$$

$$\left[\frac{\partial\mathbf{f}(\mathbf{q})}{\partial\mathbf{q}}\right]\ddot{\mathbf{q}} = \left[\mathbf{M} + \frac{\Delta t}{2}\mathbf{C} + \frac{\Delta t^2}{4}\mathbf{K}\right]\ddot{\mathbf{q}}^* \tag{6}$$

$$- \frac{\Delta t^2}{4}\mathbf{\Phi}_{\mathbf{q}}^T\alpha\left(\dot{\mathbf{\Phi}}_{\mathbf{q}}\dot{\mathbf{q}} + \dot{\mathbf{\Phi}}_t\right)$$

This method, described in detail in [12], has proved to be a robust and efficient global formulation [13,14]. All numerical experiments will be performed using a time-step Δt of 1.25×10^{-3} s and a penalty factor α of 10^8 s.

2.3 Starting implementation

In our starting implementation, the simulation algorithm was implemented using Fortran 90 and the Compaq Visual Fortran compiler. Two versions were developed: (a) a dense matrix storage version, using Fortran 90 matrix manipulation

Table 1 Percentage of the total CPU time required by each algorithm phase in the starting implementation for typical problem sizes: dense version in small problems (10 loops, 22 variables) and sparse version in medium-size problems (40 loops, 82 variables)

Stage	Dense (%)	Sparse (%)
Evaluation of residual and tangent matrix, Eqs. (3) and (5)	48	15
Evaluation of right-term in orthogonal projections, Eq. (6)	4	13
Tangent matrix factorizations and back-substitutions, Eqs. (4 and 6)	44	51
Other	4	21

capabilities and the linear equation solver included with this compiler (IMSL Fortran Library, from Visual Numerics), and (b) a sparse matrix storage version, using the MA27 sparse linear equation solver from the Harwell Subroutine Library. These two implementations, typical in the multibody community, have been tuned and improved by our group during the last years, and they have proved to be faster than commercial codes [13,14]. Its efficiency will serve as a reference to measure the performance improvements achieved with the new implementations proposed in this paper.

Table 1 shows the results of a CPU usage profiling in our starting implementations, for both dense and sparse versions, applied to representative problem sizes. As stated in the introduction, matrix computations consume most of the CPU time.

In order to test alternative implementations, the authors have developed a new MBS simulation software, implemented in C++, which can be easily configured to use different matrix storage formats and linear algebra algorithms and implementations. Numerical experiments have been performed on an AMD Athlon64 CPU. After testing different operating systems and compilers, results show that their effect on the performance is an order of magnitude lower than the effect of linear algebra implementations. Final CPU times have been measured using the GNU gcc compiler and the Linux O.S., without loss of generality.

3 Efficient dense matrix implementations

Global formulations applied to reduced rigid models (e.g. an industrial robot), or topological formulations applied to medium-sized rigid models (e.g. a complete road vehicle), lead to algorithms that operate with small-sized matrices of dimensions less than 50×50. In these cases, dense linear algebra is frequently used in MBS dynamics, since it is *supposed* to provide equal or higher performance than sparse implementations. Achieving real-time in the simulation of these small problems can be a challenge in hardware-in-the

loop settings (e.g. advanced Electronic Stability Control systems for automobiles), due to the low computing power of embedded microprocessors, the small time-steps required for hardware synchronization and the added control logic.

A straightforward way to increase the performance of dense matrix computations is by using an efficient implementation of Basic Linear Algebra Subprograms (BLAS). BLAS [15] is a standardized interface that defines routines to perform low level operations between scalars, dense vectors and dense matrices. A Fortran 77 reference implementation is available, and more efficient implementations have been developed by hardware vendors and researchers. These optimized BLAS versions exploit hardware features of modern computer architectures to get the best computational efficiency. In addition to the reference Fortran 77 implementation, three optimized BLAS implementations have been tested:

- Automatically Tuned Linear Algebra Software (ATLAS), which employs empirical techniques to generate an optimal implementation for any hardware architecture [7].
- GotoBLAS, based on optimized assembler kernels, handwritten for the most popular hardware architectures [16].
- ACML, developed by the microprocessor manufacturer AMD for its CPUs [17]. Other hardware vendors also provide their own implementations (MKL from Intel, SCSL from SGI, etc.).

Dynamic simulations can also make a profit of these optimized BLAS implementations in the solution of dense linear equation systems, provided the LAPACK library is used [18], since its linear equation solvers are based on low-level BLAS operations. In addition to the reference LAPACK implementation, written in Fortran 77, some optimized BLAS implementations like ATLAS and ACML supply their own optimized versions of the LAPACK linear solvers.

The proposed test problem, with a number of loops ranging from 1 to 20 (i.e. number of variables ranging from 4 to 42), was solved using different BLAS and LAPACK implementations to perform all matrix computations. Since the tangent matrix in the proposed dynamic formulation is symmetric and positive-definite (SPD), only the lower triangular part of the matrix is computed, and the LAPACK routines DPOTRF and DPOTRS have been used as linear equation solver. Performance results are shown in Fig. 2, where the legend text is encoded in the form "BLAS implementation + LAPACK implementation" (except for the starting implementation), and the combinations are ordered by increasing efficiency.

Results in Fig. 2 clearly show the advantage of using BLAS and LAPACK, which speed-up the simulation in a factor between 2 and 5, depending on the problem size, compared with our previous starting implementation. The low
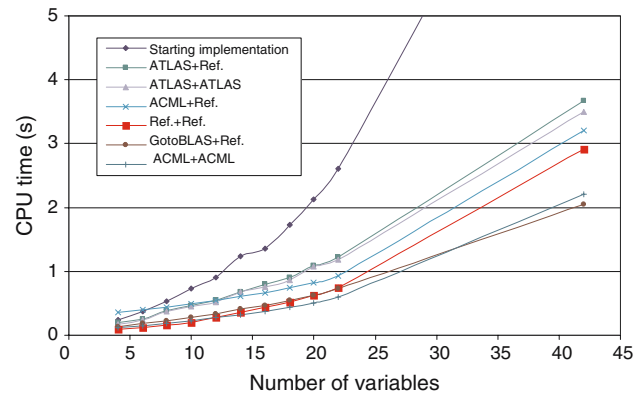

Fig. 2 Performance of different dense BLAS and LAPACK implementations

performance of the ATLAS implementation, compared to the BLAS reference implementation, can be explained by its high sensitiveness to the development environment (e.g. compiler version) and its current unstable state (it is under strong development). The vendor implementation (ACML) and GotoBLAS deliver the best results except for very small problems (up to 10 variables). The implementation named "Ref. + Ref." delivers the best performance for very small problems, and 70–80% of the performance of the best implementations for medium-size problems (three times more efficient than our starting implementation); in addition, it has a very good portability (it is written in plain Fortran 77) and usability: the installation process is straightforward, which is not always true for other implementations.

Since some MBS dynamic formulations lead to a non-symmetric tangent matrix [19], the same numerical experiment has been executed using general algorithms (not SPD-specific) to compute all matrix operations; CPU times are about 15% higher, but the efficiency ranking of Fig. 2 is maintained.

4 Efficient sparse matrix implementations

In MBS dynamics, sparse matrix techniques are used in global formulations applied to medium- or big-sized rigid models; as an example, a global model of an automobile leads to matrices of dimension about 200×200 [14]. If flexible bodies are considered, the matrix size increases, making sparse techniques profitable even if topological formulations are used: a topological model of the same automobile, with some of its bodies characterized as flexible elements (described by component mode synthesis), leads to matrices of dimension about 100×100. In any case, MBS models developed with real-time formulations hardly ever lead to matrix sizes bigger than $1{,}000 \times 1{,}000$, significantly smaller than

the typical sizes in other applications like Finite Element Analysis (FEA) or Computer Fluid Dynamics (CFD).

Regarding the sparsity, the proposed test problem and MBS dynamic formulation lead to a tangent matrix of size $2N+2$ and $12N+4$ structural non-zeros. For matrices of size 50×50, 100×100 and 500×500, the corresponding number of non-zeros is 12, 6 and 1%. These are representative values for MBS simulations, and they are quite higher than typical values in other applications that require sparse matrix technology (FEA, CFD).

Hence, MBS dynamics has two characteristics which make its sparse matrix computations different from other applications:

(a) Matrix computations are very repetitive, and the sparse patterns remain constant during the simulation. Therefore, symbolical preprocessing can be applied to almost all matrix expressions at the beginning of the simulation, in order to accelerate the numerical evaluations during the simulation. Section 4.1 presents some tips to exploit this feature.
(b) The involved sparse matrices are relatively small and dense, compared with the typical values in sparse matrix technology. Section 4.2 evaluates how sparse linear equation solvers perform in these circumstances.

4.1 Optimized sparse matrix computations

Several numerical libraries are available nowadays to support sparse matrix computations: MTL, MV++, Blitz++, SparseKIT, etc. For our new implementations, we have chosen uBLAS, a C++ template class library that provides BLAS functionality for sparse matrices [20]. Its design and implementation unify mathematical notation via operator overloading and efficient code generation via expression templates. Even though, the performance of some matrix operations can be further improved if some special algorithms are used. Results of CPU usage profiling (similar to Table 1) guided us to optimize three operations.

The first optimized operation is the rank-k update of symmetric matrix, $\Phi_q^T \alpha \Phi_q$, computed in Eq. (5). Since the sparse structure of the Jacobian matrix Φ_q is constant, a symbolic analysis is performed in order to pre-calculate the sparse pattern of the result matrix and to create a data structure that holds the operations needed to evaluate it during the simulation. In our starting sparse implementation, a similar approach was taken, but the Jacobian matrix was stored as dense, to simplify the operations at the cost of a higher memory usage.

The second optimized operation is the matrix addition computed in Eq. (5). Our starting sparse implementation used the Harwell MA27 routine as linear equation solver, which requires the sparse matrix to be stored in coordinate format
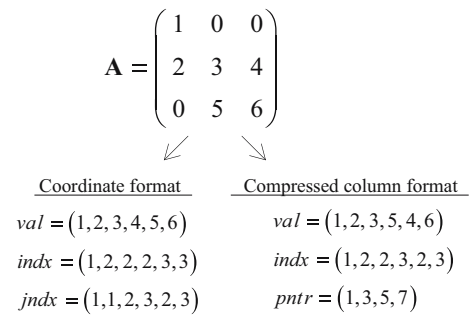
$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 3 & 4 \\ 0 & 5 & 6 \end{pmatrix}$$

Coordinate format

$val = (1, 2, 3, 4, 5, 6)$

$indx = (1, 2, 2, 2, 3, 3)$

$jndx = (1, 1, 2, 3, 2, 3)$

Compressed column format

$val = (1, 2, 3, 5, 4, 6)$

$indx = (1, 2, 2, 3, 2, 3)$

$pntr = (1, 3, 5, 7)$

Fig. 3 Sparse storage formats used in out implementations

(Fig. 3), and allows duplicated entries in the matrix structure. Therefore, the matrix addition is not actually computed, since the different terms are appended as duplicated entries in the tangent matrix. Our new implementation uses the compressed column storage format (Fig. 3), since it is required by the sparse linear equation solvers tested in Sect. 4.2. With this storage format, matrix additions require complex data traversing that slows down the performance. The following approach was taken in order to optimize the operation:

$$\mathbf{B} = t_1 \mathbf{A_1} + t_2 \mathbf{A_2} \tag{7}$$

In the preprocessing stage, the sparse pattern of $\mathbf{B}$ is calculated as the union of $\mathbf{A_1}$ and $\mathbf{A_2}$ sparse patterns, and the resulting pattern is added to $\mathbf{A_1}$ and $\mathbf{A_2}$. In this way, $\mathbf{A_1}$, $\mathbf{A_2}$ and $\mathbf{B}$ share the same sparse pattern (same *indx* and *pntr* arrays in the compressed column storage format shown in Fig. 3), and therefore, the matrix addition can be computed as a vector addition of the *val* arrays:

$$\mathbf{val_B} = t_1 \mathbf{val_{A_1}} + t_2 \mathbf{val_{A_2}} \tag{8}$$

This technique increases the number of non-zeros (NNZ) of the addend matrices. In the proposed MBS dynamic formulation, the NNZ of the mass matrix $\mathbf{M}$ is increased in a 10% approximately, which slows down the matrix–vector multiplications needed in the right terms of Eqs. (3) and (6). However, the simulation timings show that this slowdown is negligible compared with the gains derived from the fast matrix addition.

Finally, the third optimized operation concerns sparse matrix access. The write operation $\mathbf{A}(i, j) = a_{ij}$, straightforward in dense storage, needs additional position lookup when the compressed column storage is used. In the proposed formulation, the update of the Jacobian matrix Φ_q in each iteration takes 10–15% of the CPU time. The involved operations are rather simple, and most of this time is spent in matrix access. In order to optimize this procedure, a preprocessing stage evaluates the Jacobian matrix and registers the order in which entries $\Phi_q(i, j)$ are written in the *val* array of the compressed column format, creating a vector that holds indices to these positions, in the same order of evaluation.

Table 2 Efficiency of the optimized sparse matrix operations

Sparse operation	CPU time (ms)		Ratio
	Not optimized	Optimized	
(1) Rank-k update of symmetric matrix	2528.2	9.4	269
(2) Matrix addition	140.9	1.9	74
(3) Jacobian matrix evaluation	11.6	3.8	3

Latter, in the simulation stage, access to the Jacobian matrix is performed using this index vector, without the need to map (i, j) indices to memory addresses for each write operation.

Table 2 summarizes the performance gains delivered by the proposed optimizations, compared with the performance delivered by the uBLAS default algorithms (which are similar to other generic sparse matrix libraries). The numerical experiment used the matrix terms derived from an N-four-bar mechanism with $N = 40$ loops, which leads to a tangent matrix of size 82×82. Results show the importance of optimizing rank-k updates and matrix additions, since the performance delivered by off-the-shelf sparse matrix libraries is not satisfactory for these repetitive operations.

4.2 Evaluation of sparse linear equation solvers

Data in Table 1 shows that, in our starting sparse implementation, about 50% of the total CPU time is spent in tangent matrix factorizations and back-substitutions, Eqs. (4) and (6). Thus, the main performance improvements in MBS dynamic simulation can be achieved by using a more efficient sparse linear solver. During the last decade, sparse solvers have significantly improved the state of the art of the solution of general sparse linear equation systems, and more than 30 sparse solver libraries are freely available in the World Wide Web [21].

The efficiency of sparse solvers varies greatly depending on the matrix size, structure, number of non-zeros, etc. In addition, solving a sparse linear equation system usually involves three stages: preprocessing (ordering, symbolic factorization), numerical factorization and back substitution; some solvers are very fast in the first stage, while others perform better in the second or third stage. The performance of sparse solvers has been compared in previous works [5,6], but the conditions of these studies (in particular, matrix sizes and percentage of non-zeros) do not fit the above-mentioned particular features of MBS dynamics, and therefore their conclusions cannot be extrapolated to this field. As a result, it is almost impossible to determine, without numerical experiments, which sparse solver will deliver the best performance in an MBS dynamic simulation.

Given the large number of existing sparse solvers, a selection process is required in order to narrow the scope. Solvers for shared memory or distributed memory parallel machines have been discarded, since the small matrix sizes in MBS real-time dynamics (almost fit in the CPU cache memory) makes them unprofitable. The same argument applies to iterative solvers and out-of-core solvers, designed for very big linear equation systems. From the remaining solvers, those that performed best in previous comparative studies have been selected:

- Cholmod, a left-looking supernodal symmetric positive definite solver [22].
- KLU, a solver specially designed for circuit simulation matrices [23].
- SuperLU (serial version), an unsymmetric general purpose solver [24].
- Umfpack, an unsymmetric multifrontal solver [25].
- WSMP, a symmetric indefinite solver [26].

Despite the coefficient matrix is symmetric positive-definite in the proposed dynamic formulation, we have included in the numerical experiments some general, non-symmetric solvers (KLU, SuperLU, Umfpack), since other dynamic formulations lead to a non-symmetric coefficient matrix [19]. In these cases, the whole coefficient matrix (upper and lower parts) is computed, while with symmetric solvers only half matrix is used in the formulation equations. Each solver supports its own set of reordering strategies; all of them have been tested to select the best one in each simulation. In addition, all the optimizations described in the previous Section were applied to our new sparse implementation.

The proposed test problem, with a number of loops ranging from 10 to 500 (i.e. number of variables ranging from 22 to 1,002), was solved using different sparse solvers. Performance results are shown in Fig. 4 for a number of variables up to 160, since the trends are preserved for higher number
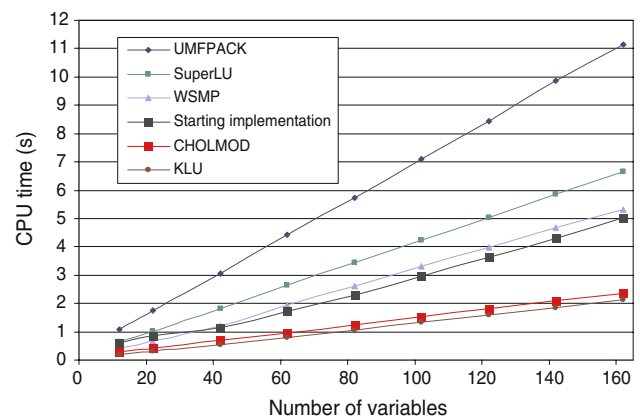

Fig. 4 Performance of different sparse linear equation solvers as a function of problem size

of variables. The legend text shows the name of the sparse solvers, ordered by increasing efficiency.

Surprisingly, KLU is the fastest solver, despite being a general solver that does not exploit the symmetric positive definite condition of the coefficient matrix; in addition, it has been designed for circuit simulation problems, which lead to very sparse matrices, the opposite case of MBS dynamics. However, these results have been obtained by using the KLU *refactor* routine for numerical factorizations, which reuses the pivoting strategy generated in the preprocessing stage. In multibody problems where the elements of the tangent matrix of Eq. (5) may significantly change their relative values during the simulation (e.g. due to violent impacts), the initial pivoting strategy may become invalid and the *refactor* routine would probably accumulate high numerical errors. To avoid this, the KLU solver can recalculate the pivoting strategy in each numerical factorization, but this method increases the CPU times in a 50%. On the other hand, Cholmod, a symmetric positive definite solver, performs at 85% of KLU, despite recalculating the pivoting strategy in each numerical factorization. Our best new sparse implementations (using KLU or Cholmod) perform faster that our starting implementation, in a factor from 2 (small problems) to 3 (big problems of 1,000 variables).

4.3 Effect of dense BLAS implementation

Some sparse solvers, like Cholmod, rely on dense BLAS to increase their performance. In addition, some sparse matrix operations (e.g. the optimized matrix addition described in Sect. 4.1) are actually computed as dense vector operations using BLAS routines. Results shown in Fig. 4 have been generated using the reference BLAS implementation. The same numerical experiment has been executed using the faster, optimized GotoBLAS and ACML implementations, and CPU times have decreased only in a 2–3%. Hence, the reference BLAS implementation is recommended for MBS dynamics in sparse implementations, since it provides the best compromise between performance and usability.

5 Sparse versus dense implementations

As stated previously, dense linear algebra is frequently used in MBS dynamics for small problems (dimension of the coefficient matrix lower than 50), since it is supposed to provide higher performance than sparse implementations [27]. Our starting sparse implementation, which already employs some of the optimizations described in Sect. 4.1, disagrees with this assumption, and this fact is reinforced with the performance of our new optimized implementations: sparse versions perform always faster than dense versions even for

small problems, in a factor which ranges from 1.5 (problems of 10 variables) to 5 (problems of 50 variables).

However, this conclusion has been obtained for the proposed test problem and dynamic formulation, and it could be argued that it cannot be generalized to other situations that lead to a coefficient matrix with a higher percentage of non-zeros, as in the case of highly constrained mechanism or topological formulations. The objection could be made to the efficiency ranking shown in Fig. 4. In order to get insight about this subject, the numerical experiments used to generate Fig. 4 were repeated, but in this case artificial non-zeros were introduced in the mass matrix **M**, in order to generate a tangent matrix with a variable percentage of non-zeros. Figure 5 shows the CPU times for a mechanism of 48 loops (100 variables), as a function of matrix filling. Results show that two sparse implementations, based on the Cholmod and WSMP sparse solvers, are always faster than the best dense implementation, even with 100% of non-zeros in the tangent matrix. This surprising fact can be explained by two factors: (a) Cholmod and WSMP rely on dense BLAS routines to perform the factorization, and therefore they start to operate as dense solvers as the matrix filling increases; (b) the percentage of non-zeros is always lower in the Jacobian matrix than in the tangent matrix, hence optimized sparse implementations achieve significant time savings in Jacobian operations, in comparison with dense implementations.

Results for other problem sizes are synthesized in Fig. 6: the different regions represent the points (*problem size, matrix filling*) where each implementation delivers the best performance. For most MBS problems and dynamic formulations, a sparse implementation based on the KLU solver will be the frontrunner. However, topological formulations with a symmetric tangent matrix will benefit from a sparse implementation based on the WSMP solver, specially when they are applied to rigid models, which result in a higher matrix filling.
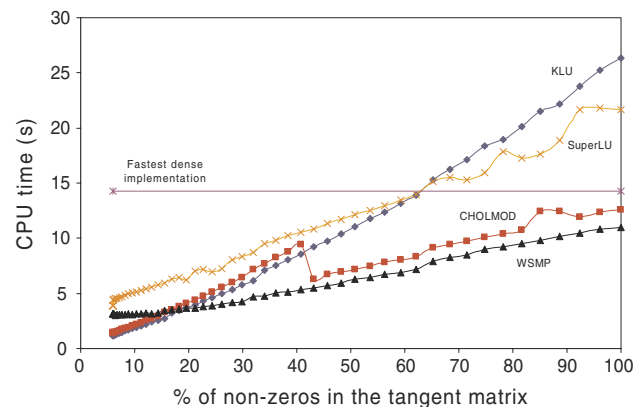

Fig. 5 Performance of different sparse linear equation solvers as a function of tangent matrix filling, for a problem size of 100 variables
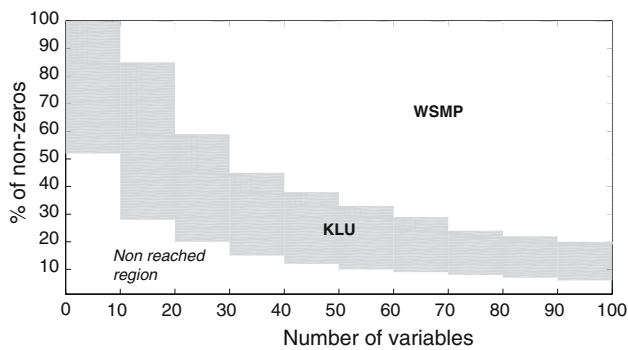
Fig. 6 Best implementation, as a function of problem size and percentage of non-zeros in the tangent matrix
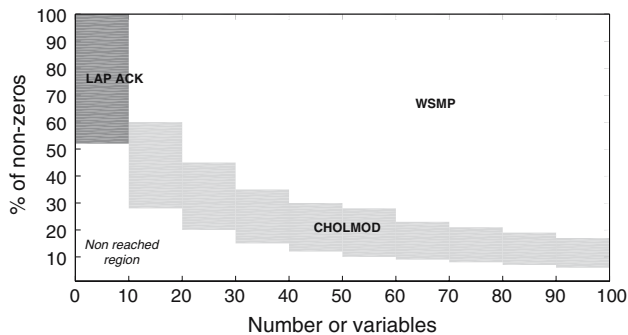

Fig. 7 Best implementation, as a function of problem size and percentage of non-zeros in the tangent matrix (*refactor* routine of KLU is not used)

Figure 6 has been obtained by using the KLU *refactor* routine for numerical factorizations. As explained in Sect. 4.2, this may cause trouble in problems where the entries of the tangent matrix change their relative values significantly during the simulation. If the refactor routine is not used, Fig. 7 is obtained. In this case KLU is replaced by Cholmod, WSMP increases its influence area, and the dense implementation based on LAPACK emerges for very small problems (less than 10 variables), but with a very small advantage. Conversely, two exceptions can be mentioned: (a) for dynamic formulations with symmetric indefinite tangent matrices, WSMP would be the frontrunner for almost all the situations, since Cholmod does not support them; (b) for dynamic formulations with unsymmetric tangent matrices, KLU would be the again frontrunner for almost all the situations (even if the *refactor* routine is avoided), since WSMP does not support them.

6 Conclusions

Regarding the implementation aspects of MBS dynamic simulations, the following conclusions can be established:

- Efficient linear algebra implementations can speed up the efficiency in a factor of 2–3, compared with traditional implementations. In other words, problems of double or triple size can be solved with the same resources.
- The proposed optimizations based on symbolic preprocessing of the sparse matrix computations can deliver huge speedups, since off-the-shelf sparse matrix libraries do not take advantage of the constant sparse pattern of operations during the dynamic simulation.
- Optimized sparse implementations are recommended since they perform better than optimized dense implementations, even for small-sized problems or relatively dense matrices. This disagrees with the widespread belief in MBS dynamics.
- Concerning sparse linear equation solvers, it has been found that KLU, an unfamiliar solver designed for circuit simulation, performs very well with many of the linear equation systems resulting from MBS dynamics. In addition, it was found that the reference BLAS implementation provides the best compromise between performance and usability for sparse implementations.

The results from numerical experiments are summarized in Table 3, which provides a simple decision rule to select the best linear equation solver for MBS dynamics, based on matrix type, size and percentage of non-zeros. Efficient implementations of global MBS dynamic formulations can be easily achieved, provided the above recommendations are followed. All the recommended software libraries are freely available, and the proposed optimization techniques are not bounded to any programming language.

As a consequence of the abovementioned conclusions, the limit for problem size where global formulations perform better than topological formulations, established in the order of 40 variables [14], should be revised. This limit was obtained using dense implementations, and it might get higher if the proposed optimized sparse implementations were used, since their effects on the efficiency are higher in global formulations than in topological formulations. In addition, further work must be carried out in order to determine if the proposed recommendations are still valid for other formulations,

Table 3 Decision rules for selecting the best sparse solver for MBS dynamics, based on matrix type, size and percentage of non-zeros

Type of tangent matrix	(No. of variables) × (% of non-zeros − 10)	
	<900	>900
Symmetric positive definite	KLU (smooth problems)	WSMP
	Cholmod (rough problems)	
Symmetric	KLU	WSMP
Unsymmetric	KLU	KLU

since all the numerical experiments have been performed using a particular global formulation.

References

1. Cuadrado J, Cardenal J, Morer P (1997) Modeling and solution methods for efficient real-time simulation of multibody dynamics. Multibody Syst Dyn 1:259–280

2. Bae DS, Lee JK, Cho HJ, Yae H (2000) An explicit integration method for realtime simulation of multibody vehicle models. Comput Meth Appl Mech Eng 187:337–350

3. Anderson KS, Critchley JH (2003) Improved 'Order-N' performance algorithm for the simulation of constrained multi-rigid-body dynamic systems. Multibody Syst Dyn 9:185–212

4. Anderson KS, Mukherjee R, Critchley JH, Ziegler J, Lipton S (2007) POEMS: parallelizable open-source efficient multibody software. Eng Comput 23:11–23

5. Gupta A (2002) Recent advances in direct methods for solving unsymmetric sparse systems of linear equations. ACM Trans Math Softw 28:301–324

6. Scott JA, Hu YF, Gould NIM (2006) An evaluation of sparse direct symmetric solvers: an introduction and preliminary findings. Appl Parallel Comput State Art Sci Comput 3732:818–827

7. Whaley RC, Petitet A, Dongarra JJ (2001) Automated empirical optimizations of software and the ATLAS project. Parallel Comput 27:3–35

8. Turek S, Becker C, Runge A (2001) The FEAST indices. Realistic evaluation of modern software components and processor technologies. Comput Math Appl 41(10–11):1431–1464

9. Yu JSK, Yu CH (2002) Recent advances in PC-Linux systems for electronic structure computations by optimized compilers and numerical libraries. J Chem Inform Comput Sci 42:673–681

10. Gonzalez M, Dopico D, Lugrís U, Cuadrado J (2006) A benchmarking system for MBS simulation software: problem standardization and performance measurement. Multibody Syst Dyn 16:179–190

11. García de Jalón J, Bayo E (1994) Kinematic and dynamic simulation of multibody systems—the real-time challenge. Springer, New York

12. Bayo E, Ledesma R (1996) Augmented Lagrangian and mass-orthogonal projection methods for constrained multibody dynamics. Nonlinear Dyn 9:113–130

13. Cuadrado J, Gutierrez R, Naya MA, Morer P (2001) A comparison in terms of accuracy and efficiency between a MBS dynamic formulation with stress analysis and a non-linear FEA Code. Int J Numer Meth Eng 51:1033–1052

14. Cuadrado J, Dopico D, González M, Naya M (2004) A combined penalty and recursive real-time formulation for multibody dynamics. J Mech Des 126:602–608

15. NIST (2006) Basic linear algebra subprograms. http://www.netlib.org/blas/

16. Goto K (2006) GotoBLAS. http://www.tacc.utexas.edu/resources/software/

17. AMD (2007) AMD Core Math Library. http://developer.amd.com/acml.jsp

18. NETLIB (2007) LAPACK. http://www.netlib.org/lapack/

19. Dopico D, Lugrís U, González M, Cuadrado J (2006) Two implementations of IRK integrators for real-time multibody dynamics. Int J Num Meth Eng 65:2091–2111

20. Walter J, Kock M (2006) UBLAS. http://www.boost.org/libs/numeric/

21. Dongarra JJ (2004) Freely available software for linear algebra on the web. http://www.netlib.org/utk/people/JackDongarra/la-sw.html

22. Chen Y, Davis TA, Hager WW, Rajamanickam S (2006) Algorithm 8xx: CHOLMOD, supernodal sparse cholesky factorization and update/downdate. http://www.cise.ufl.edu/~davis/techreports/cholmod/tr06-005.pdf

23. Davis TA, Stanley K (2004) KLU: a Clark Kent Sparse LU Factorization Algorithm for Circuit Matrices. http://www.cise.ufl.edu/~davis/techreports/KLU/pp04.pdf

24. Demmel JW, Eisenstat SC, Gilbert JR, Li XYS, Liu JWH (1999) A supernodal approach to sparse partial pivoting. Siam J Matrix Anal Appl 20:720–755

25. Davis TA (2004) Algorithm 832: UMFPACK V4.3—an unsymmetric-pattern multifrontal method. ACM Trans Math Softw 30:196–199

26. Gupta A, Joshi M, Kumar V (1998) WSSMP: a high-performance serial and parallel symmetric sparse linear solver. Appl Parallel Comput 1541:182–194

27. Cuadrado J, Dopico D (2004) A combined penalty and semi-recursive formulation for closed-loops in MBS. In: Proceedings of the eleventh world congress in mechanism and machine science, vols 1–5, pp 637–641