

Non-symmetric Algebraic Multigrid Preconditioners for the Bidomain Reaction–Diffusion system

Micol Pennacchio and Valeria Simoncini

Abstract We deal with the efficient solution of the so-called *bidomain* system which is possibly the most complete model for the cardiac bioelectric activity. We study the performance of a non-symmetric structured algebraic multigrid preconditioner on the formulation generally used of the bidomain model, i.e. the one characterized by a parabolic equation coupled with an elliptic one. Our numerical results show that, for this formulation, the non-symmetric preconditioner provides the best overall performance.

1 The Bidomain Model

The excitation process in the myocardium is a complex phenomenon characterized by rapid ionic fluxes through the cellular membrane separating the intracellular and the interstitial fluid in the myocardium [7]. The *bidomain* is the most complete model for the cardiac bioelectric activity, and it consists of a non-linear Reaction-Diffusion (R-D) system of equations for the intra- and extracellular potential u_i and u_e , coupled through the transmembrane potential $v := u_i - u_e$ [9]. The nonlinearity arises through the current-voltage relationship across the membrane which is described by a set of nonlinear ODEs (see [7]). The anisotropic properties of the media are modeled by the intra- and extracellular conductivity tensors $M_i = M_i(\mathbf{x})$ and $M_e = M_e(\mathbf{x})$ that satisfy a uniform ellipticity condition, see [10].

Micol Pennacchio
Istituto di Matematica Applicata e Tecnologie Informatiche del CNR, via Ferrata, 1,
I-27100 Pavia, Italy, e-mail: micol@imati.cnr.it

Valeria Simoncini
Dipartimento di Matematica, Università di Bologna, Piazza di Porta S. Donato, 5,
I-40127 Bologna, Italy; CIRSA, Ravenna and IMATI-CNR, Pavia, Italy e-mail: vale-
ria@dm.unibo.it

1.1 (u_e, v) formulation

The R-D system governing the cardiac electric activity may be written in various forms involving different combinations of the variables u_i, u_e, v ; see, e.g., [12]. Here we deal with the formulation generally used for the numerical simulations, i.e. with a parabolic equation for the transmembrane potential v coupled with an elliptic equation for the extracellular potential u_e :

find $(v(\mathbf{x}, t), u_e(\mathbf{x}, t))$, $\mathbf{x} \in \Omega$, $t \in [0, T]$ such that

$$\begin{cases} c_m \partial_t v - \operatorname{div} M_i \nabla v + I_{ion} = \operatorname{div} M_i \nabla u_e + I_{app} & \text{in } \Omega \times]0, T[\\ -\operatorname{div} M \nabla u_e = \operatorname{div} M_i \nabla v & \text{in } \Omega \times]0, T[\\ \mathbf{n}^T M_i \nabla v = 0, \quad \mathbf{n}^T M \nabla u_e = 0 & \text{on } \Gamma \times]0, T[\\ v(\mathbf{x}, 0) = 0 & \text{in } \Omega. \end{cases} \quad (1)$$

with $M = M_i + M_e$ bulk conductivity tensor. Due to the presence of different time and space scales, the numerical solution of the bidomain system represents a very intensive computational task: realistic three dimensional simulations typically yield discrete problems with millions of unknowns, and time steps of the order of 10^{-2} milliseconds or less. To reduce the computational cost, different numerical techniques have been developed and are currently under investigation [4, 5, 6, 8, 11, 17]. Here we employ a semi-implicit method in time, that only requires the solution of linear systems at each time step and allows performing larger time steps than explicit schemes. By using a finite element discretization in space and a semi-implicit scheme in time, we get:

$$\mathcal{B} \boldsymbol{\xi}^{k+1} = \mathbf{b} \quad \text{with} \quad \mathcal{B} = \begin{bmatrix} C_t + A_i & A_i \\ A_i & (A_i + A_e) \end{bmatrix}, \quad (2)$$

with $\mathbf{b} = [C_t \mathbf{v}^k - I_{ion}^h(\mathbf{v}^k) + I_{app}^h; \mathbf{0}]$, $\mathbf{v}^k = \mathbf{u}_i^k - \mathbf{u}_e^k$, $\boldsymbol{\xi}^{k+1} = [\mathbf{v}^{k+1}; \mathbf{u}_e^{k+1}]$.

Whatever the method chosen for discretizing the problem, a huge computational effort is required to solve the associated linear system in (2) at each time step, whose conditioning considerably worsens as the problem dimension increases, resulting in an unacceptable increase in the computational costs of the whole simulation. Preconditioning is therefore mandatory. Attempts in the recent literature have employed diagonal preconditioners, Symmetric Successive Over Relaxation [10], Block Jacobi preconditioners with incomplete LU factorization (ILU) [18], multigrid [1, 13]; see [19] for an overview.

Classically, the system with \mathcal{B} was solved by means of a nested iteration that can be explicitly stated as a block Gauss–Seidel method involving the two diagonal blocks. More recent publications have demonstrated that appropriately preconditioned Krylov subspace methods can largely outperform these classical methods; see [11, 12] and references therein. In this paper we study a nonsymmetric block triangular preconditioner for the system in (2).

2 Block Preconditioners

In the coefficient matrix \mathcal{B} in (2), the (1,1) block is positive definite while the (2,2) block is only positive semi-definite. Moreover, all matrices are square and symmetric. It is therefore natural to derive preconditioners that exploit this structure, as is the case in various multilevel methods as well as in saddle point problems [2]. In [12] we analyzed symmetric structured preconditioners and in particular a block diagonal preconditioner \mathcal{P}_d and a block factorized preconditioner \mathcal{P}_f :

$$\mathcal{P}_d = \text{blockdiag}(K, D), \quad \mathcal{P}_f = \begin{bmatrix} I & O \\ A_i K^{-1} & I \end{bmatrix} \begin{bmatrix} K & A_i \\ O & D \end{bmatrix}, \quad (3)$$

where K is a symmetric and positive definite approximation to the (1,1) block, while D is a symmetric and positive definite approximation either to the Schur complement $(A_i + A_e) - A_i(C_t + A_i)^{-1}A_i$, or to the (2,2) block $A_i + A_e$. Note that for $K = C_t + A_i$ and $D = (A_i + A_e) - A_i(C_t + A_i)^{-1}A_i$ then $\mathcal{P}_f = \mathcal{B}$. In [12] we also experimentally verified that more general structured preconditioners may also be appealing. For instance, the following ‘‘one-sided’’ version of \mathcal{P}_f has been used for symmetric (indefinite) saddle point problems (see, e.g., [2]):

$$\mathcal{P}_M = \begin{bmatrix} K & A_i \\ O & D \end{bmatrix}.$$

If K and D exactly reproduce the (1,1) block and the Schur complement, respectively, the nonsymmetric preconditioned matrix would be the matrix

$$\mathcal{B}\mathcal{P}_M^{-1} = \begin{bmatrix} I & O \\ A_i(C_t + A_i)^{-1} & I \end{bmatrix},$$

whose spectrum consists of the single unit eigenvalue, so that a minimal residual method such as GMRES ([15]) would converge in at most two iterations. In general, the behavior of the approximate versions of K and D is less predictable; moreover, a good approximation of the Schur complement may be very expensive to obtain. The performance of \mathcal{P}_M within the indefinite saddle point context highly overcomes its nonsymmetric nature. The situation is considerably different in our context, where the original matrix is positive (semi)definite. Applying \mathcal{P}_M destroys symmetry. Remarkably, however, the use of \mathcal{P}_M in our 2D problem provides some interesting numerical results, as already shown in [12]. Here we would like to complete the analytical analysis of this preconditioner for the other ideal case $D = A_i + A_e$, or when K and D are spectrally equivalent approximations to the ideal cases. In the following we assume that A_i^{-1} stands for the pseudo-inverse whenever the matrix is singular. Singularity does not effect the analysis as all vectors are assumed to lie in the range of the considered matrices. For $K = A_i + C_t$, $D = A_i + A_e$, it can be easily verified that

$$\mathcal{BP}_M^{-1} = \begin{bmatrix} I & O \\ A_i(A_i + C_t)^{-1} & I - \mathcal{S} \end{bmatrix}, \quad \mathcal{S} = A_i(A_i + C_t)^{-1}A_i(A_i + A_e)^{-1}. \quad (4)$$

The following result shows that the spectrum of \mathcal{BP}_M^{-1} is bounded independently of the mesh parameter for judiciously chosen D .

Theorem 1. *With the previous notation, let $K = A_i + C_t$ and let \mathcal{X} be an eigenvector matrix of \mathcal{BP}_M^{-1} . If $D = A_i + A_e$, then*

$$\lambda_{\min}(\mathcal{BP}_M^{-1}) = 1 - \mu \quad \lambda_{\max}(\mathcal{BP}_M^{-1}) = 1, \quad \mathcal{X} = \begin{bmatrix} I & O \\ (A_i + A_e)A_i^{-1} & Y \end{bmatrix},$$

with $\mu \leq (1 + \lambda_{\min}(A_e, A_i))^{-1}$, μ constant independent of h , and Y eigenvector matrix of $I - \mathcal{S}$.

If instead D is such that there exist positive constants α_1, α_2 such that $\alpha_1 \mathbf{x}^T D \mathbf{x} \leq \mathbf{x}^T (A_i + A_e) \mathbf{x} \leq \alpha_2 \mathbf{x}^T D \mathbf{x}$ for all \mathbf{x} in the range of $A_i + A_e$, then either $\lambda(\mathcal{BP}_M^{-1}) = 1$ or $\alpha_1(1 - \mu) \leq \lambda(\mathcal{BP}_M^{-1}) \leq \alpha_2$, with μ defined above. Moreover,

$$\mathcal{X} = \begin{bmatrix} I & O \\ M & Y \end{bmatrix},$$

with $M = -(G - I)^{-1}A_i(A_i + C_t)^{-1}$, $G = (A_i + A_e - A_i(C_t + A_i)^{-1}A_i)D^{-1}$ and Y eigenvector matrix of G .

Proof. We shall see that the eigenvalues of \mathcal{S} are real and non-negative. From the structure of the matrix \mathcal{BP}_M^{-1} it thus follows that $\lambda_{\min}(\mathcal{BP}_M^{-1}) = 1 - \lambda_{\max}(\mathcal{S})$ and $\lambda_{\max}(\mathcal{BP}_M^{-1}) = 1$. To analyze the eigenvalues λ of \mathcal{S} we consider the eigenvalue problem

$$A_i(A_i + C_t)^{-1}A_i(A_i + A_e)^{-1}\mathbf{x} = \lambda\mathbf{x} \quad \Leftrightarrow \quad A_i(A_i + C_t)^{-1}A_i\mathbf{u} = \lambda(A_i + A_e)\mathbf{u}.$$

Clearly, $\lambda = 0$ for $\mathbf{u} \in N(A_i) = N(A_e) = N(A_i + A_e)$. Moreover, since the matrices on both sides are symmetric and positive definite in the range of A_i, A_e , $\lambda \geq 0$. We can write

$$\lambda = \frac{\mathbf{u}^T A_i(A_i + C_t)^{-1}A_i\mathbf{u}}{\mathbf{u}^T (A_i + A_e)\mathbf{u}}, \quad \mathbf{u} \notin N(A_i).$$

Thanks to [12, Lemma 4.1] we obtain that $\lambda \leq (1 + \lambda_{\min}(A_e, A_i))^{-1} < 1$. Using the conductivity coefficients defined in Section 5 of [12], the two stiffness matrices are related as $c_1 \mathbf{v}^T A_e \mathbf{v} \leq \mathbf{v}^T A_i \mathbf{v} \leq c_2 \mathbf{v}^T A_e \mathbf{v}$ independently of the mesh. Thus, $\lambda_{\min}(A_e, A_i)$ is bounded by a quantity that only depends on the conductivity tensors of the two stiffness matrices, and not on the grid. This completes the proof for $D = A_i + A_e$. One can readily verify that \mathcal{X} satisfies $\mathcal{BP}_M^{-1}\mathcal{X} = \mathcal{X}\text{blockdiag}(I, I - \Lambda)$, where Λ is the eigenvalues matrix of $I - \mathcal{S}$.

For general D we have

$$\mathcal{BP}_M^{-1} = \begin{bmatrix} I & O \\ A_i(A_i + C_t)^{-1} & I \end{bmatrix} \begin{bmatrix} I & O \\ O & (A_i + A_e - A_i(C_t + A_i)^{-1}A_i)D^{-1} \end{bmatrix}.$$

The eigenvalues θ 's of the (2,2) block in the second factor satisfy

$$\begin{aligned} \theta &= \frac{\mathbf{x}^T(A_i + A_e - A_i(C_t + A_i)^{-1}A_i)\mathbf{x}}{\mathbf{x}^T D \mathbf{x}} \\ &= \frac{\mathbf{x}^T(A_i + A_e)\mathbf{x}}{\mathbf{x}^T D \mathbf{x}} \left(1 - \frac{\mathbf{x}^T A_i(C_t + A_i)^{-1}A_i\mathbf{x}}{\mathbf{x}^T(A_i + A_e)\mathbf{x}} \right) = \gamma_1 \gamma_2. \end{aligned}$$

Using the spectral equivalence of D , we have $\alpha_1 \leq \gamma_1 \leq \alpha_2$. Moreover, using the definition of μ above, $(1 - \mu) \leq \gamma_2 \leq 1$, from which the result follows. The fact that the given \mathcal{X} is an eigenvector matrix can be readily verified.

We observe that due to the matrix structure, we expect $\text{cond}(\mathcal{X})$ to be mesh independent in the exact case ($K = A_i + C_t$, $D = A_i + A_e$). A spectral analysis when K is an approximation to the (1,1) block, that is it is not exact, is much more involved. With a convenient splitting of \mathcal{B} , we write

$$\begin{aligned} \mathcal{BP}_M^{-1} &= \left(\begin{bmatrix} K & A_i \\ A_i & D \end{bmatrix} + \begin{bmatrix} C_t + A_i - K & O \\ O & (A_i + A_e) - D \end{bmatrix} \right) \mathcal{P}_M^{-1} \\ &= \begin{bmatrix} I & O \\ A_i K^{-1} & I - A_i K^{-1} A_i D^{-1} \end{bmatrix} + \\ &+ \begin{bmatrix} (C_t + A_i - K)K^{-1} & -(C_t + A_i - K)K^{-1} A_i D^{-1} \\ O & (A_i + A_e - D)D^{-1} \end{bmatrix} \equiv R + E. \end{aligned}$$

If K , D are spectrally equivalent to $A_i + C_t$ and $A_i + A_e$ respectively, the spectrum of R is also spectrally equivalent to that of the exactly preconditioned matrix (the one obtained for $K = A_i + C_t$, $D = A_i + A_e$). The matrix E represents a perturbation to the ideal case, and its size depends on the accuracy of the preconditioning blocks.

Assume that all relevant¹ eigenvalues of the (2,2) block of R are less than one, and let X be an eigenvector matrix of R ; in fact, it is possible to derive a more explicit structure for X , but such a description is beyond the scope of this paper. Then we have (cf, e.g., [16])

$$|\lambda(\mathcal{BP}_M^{-1}) - \lambda(R)| \leq \|X^{-1}EX\| \leq \text{cond}(X)\|E\|,$$

where $\text{cond}(X)$ is the spectral condition number of X and $\|\cdot\|$ is the matrix norm induced by the Euclidean vector norm. Therefore, if K and D are good approximations to the corresponding blocks, then we expect the spectrum of \mathcal{BP}_M^{-1} not to deviate significantly from that of R , unless the eigenvector matrix X is very ill conditioned. If the condition number of the eigenvector matrix of \mathcal{BP}_M^{-1} is moderate (cf. Th. 1), we also expect that a nonsymmetric

¹ That is, those associated to eigenvectors in the range of the given matrices.

solver like GMRES will converge in approximately the same number of iterations as for the exact case, whenever K , D are appropriately chosen. Spectral mesh independence can also be expected in this case. Our numerical results fully confirm all these considerations, as reported in Table 1.

3 Numerical results

In this section we report on our experiments with the exact and “inexact” versions of the block triangular preconditioner \mathcal{P}_M . We consider a square domain $\Omega = [0, 1]^2$ modeling a block of myocardium with cardiac fibers parallel to a diagonal of the square and the conductivity coefficients defined as in [12]. The meshes considered on Ω were built by using a Delaunay triangulation algorithm. The number of mesh nodes for each refined grid was $2n$ with $n \in \{2705, 10657, 42305, 168577, 673025\}$, whereas the time step τ was chosen to be equal to $4 \cdot 10^{-2}$ msec. All experiments correspond to a typical temporal instant in the time step evolution, so that the right-hand side includes information generated during the previous time steps. All computations were performed with Matlab 7.4.0 (R2007a) on a iMac Intel Core 2 Duo 2 GbRAM 2.66GHz and 6Mb L2 cache.

In the approximate (inexact) case, the matrices K and D are implicitly defined by applying an Algebraic MultiGrid (AMG) method to approximate the corresponding blocks. As in [12], we reorder each block matrix of \mathcal{B} by using the matlab function `symrcm`. We employ the AMG code available in the HSL library, the `HSL_MI20` routine, equipped with a Matlab interface [3]. This function implements the classical (Ruge-Stüben) AMG method, as described in [14]. The code was used as a black box: Gauss-Seidel smoothing was used in all instances. The multilevel method is often built on originally singular matrices. To increase the robustness of the preconditioning strategy, in some cases we generated the preconditioner by using a shifted (nonsingular) matrix, with a shift equal to $\varepsilon^{1/2}$ and $\varepsilon \approx 10^{-16}$ the Matlab machine precision.

Table 1 CPU time and number of iterations (in parenthesis) for: \mathcal{P}_f with AMG; \mathcal{P}_M with exact and AMG-based blocks, when using the nonsymmetric solver GMRES, FOM and the symmetric solver CG (with regularization). Here $K = \text{AMG}(C_t + A_i)$ and $D = \text{AMG}(A_i + A_e)$.

n	\mathcal{P}_f	exact \mathcal{P}_M		AMG-based \mathcal{P}_M		
		CG	GMRES	GMRES	FOM	CG
2705	0.41 (6)	1.1 (21)	0.38 (6)	0.15 (6)	0.27 (6)	0.32 (22)
10657	0.88 (7)	2.98 (13)	1.95 (7)	0.5 (7)	0.52 (7)	0.64 (12)
42305	2.82 (8)	11.85 (11)	8.92 (7)	2.15 (8)	2.17 (8)	2.07 (10)
168577	9.92 (8)	58.11 (11)	44.83 (7)	8.99 (8)	9.06 (8)	8.27 (10)
673025	47.47 (10)	315.49 (10)	249.95 (7)	41.02 (9)	41.32 (9)	36.21 (11)

All our results are displayed in Table 1 for the considered discretization meshes: we report CPU times (in seconds) and in parenthesis number of iterations. In the second column we recall the numbers to obtain a reduction of the residual norm by a factor of 10^{-6} , with the best performing preconditioner analyzed in detail in [12], namely P_f and AMG for computing K and D . The subsequent two columns show the performance of the exact preconditioner \mathcal{P}_M , namely with $K = A_i + C_t$ and $D = A_i + A_e$, when using either the nonsymmetric solver GMRES or CG (see below for further comments on the latter method). The last columns show the performance of \mathcal{P}_M when AMG is used to build K and D as approximations to $A_i + C_t$ and $A_i + A_e$, respectively. Note that a sparse direct solver employs 26.15 secs, after a proper reordering, to solve the whole system in (2) for $n=168577$ (an “out of memory” results for $n=673025$). The reported timings clearly confirm the competitiveness of the AMG-based preconditioner compared to the exact version and the sparse direct solver. In particular, results are reported for the minimal residual method GMRES, and for both FOM (Full Orthogonalization Method) and CG [15]. If the preconditioned problem were symmetric and positive definite, then FOM and CG would be mathematically equivalent. Since $\mathcal{B}\mathcal{P}_M^{-1}$ is nonsymmetric, we expect CG to behave more poorly than FOM, which is a Galerkin-type method devised for nonsymmetric problems. It is however quite surprising that CG converges very quickly in spite of the full nonsymmetry of the problem. In Figure 1 we display the convergence history of the methods, in terms of residual norms: the CG curve deviates from the expected one, represented by FOM, as soon as nonsymmetry is detected. However, the spectral properties are so favourable that nonsymmetry does not prevent the method from converging in just a few more iterations. In fact, CG may be viewed in this case as a (highly) truncated full orthogonalization procedure; see, e.g. [12] and references therein. Due to the very cheap short-term recurrence, the CG timings are also very competitive; cf. the last column of Table 1.

In the exact case, our theory predicts mesh independence, and this is confirmed in the table. In addition, the use of AMG preconditioning maintains mesh independence in the inexact case, with in general a number of iterations only slightly higher than in the exact case.

References

1. T. M. Austin, M. L. Trew, and A. Pullan. Solving the cardiac bidomain equations for discontinuous coefficients. *IEEE Trans. Biomed. Eng.*, **53**:1265–1272 (2006)
2. M. Benzi, G. H. Golub, and J. Liesen. Numerical solution of saddle point problems. *Acta Numerica* **14**–137 (2005)
3. J. Boyle, M. D. Mihajlović and J.A. Scott. HSL_MI20: an efficient AMG preconditioner. *RAL Tech.Rep. RAL-TR-2007-021* (2007)
4. E. Cherry, H. Greenside, C. S. Henriquez. Efficient simulation of three-dimensional anisotropic cardiac tissue using an adaptive mesh refinement

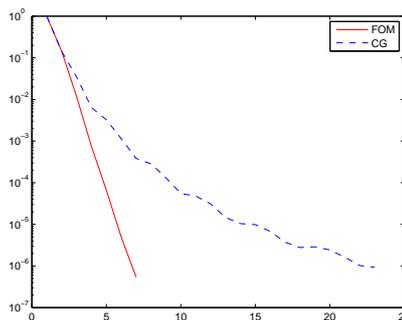


Fig. 1 Convergence history of FOM and CG on the nonsymmetric matrix \mathcal{BP}_M^{-1} , $n = 2705$.

- method. *Chaos* **13**, 853-865 (2003)
5. P. Colli Franzone, P. Deuffhard, B. Erdmann, J. Lang, L. F. Pavarino. Adaptivity in space and time for reaction-diffusion systems in electrocardiology. *SIAM J. Sci. Comput.* **28**, 942-962 (2006)
 6. P. Colli Franzone, M. Pennacchio, L. Guerri. Accurate computation of electrograms in the left ventricular wall. *Math. Mod. Meth. Appl. Sci.*, **10** 507-538 (2000)
 7. J. Keener and J. Sneyd. *Mathematical Physiology. Vol. I-II. Second Edition* Springer, New York, (2009)
 8. M. Pennacchio. The mortar finite element method for the cardiac "bidomain" model of extracellular potential. *J. Sci. Comput.*, **20**, 191-210 (2004)
 9. M. Pennacchio, G. Savaré, P. Colli Franzone. Multiscale modeling for the bioelectric activity of the heart. *SIAM J. Math. Anal.* **37**, 1333-1370 (2006)
 10. M. Pennacchio, V. Simoncini. Efficient algebraic solution of reaction-diffusion systems for the cardiac excitation process. *J. Comput. Appl. Math.* **145**, 49-70 (2002)
 11. M. Pennacchio, V. Simoncini. Substructuring preconditioners for mortar discretization of a degenerate evolution problem. *J. Sci. Comput.* **36** 391-419 (2008)
 12. M. Pennacchio, V. Simoncini. Algebraic Multigrid Preconditioners for the Bidomain Reaction-Diffusion system. *Appl. Numer. Math.* **59**, 3033-3050 (2009)
 13. G. Plank, M. Liebmann, R. Weber dos Santos, E. J. Vigmond, and G. Haase. Algebraic multigrid preconditioner for the cardiac bidomain model. *IEEE Trans. Biomed. Eng.* **54**, 585-596 (2007)
 14. J.W. Ruge, K. Stüben. Algebraic multigrid. In S.F. McCormick, editor, *Multigrid Methods*, v.3 of *Frontiers in Applied Mathematics*, pp. 73-130, SIAM (1987)
 15. Y. Saad. *Iterative methods for sparse linear systems*. The PWS Publishing Company (1996)
 16. G.W. Stewart, J-G. Sun. *Matrix Perturbation Theory*. Academic Press (1990)
 17. J. Trangenstein and C. Kim. Operator splitting and adaptive mesh refinement for the Luo-Rudy I model. *J. Comput. Phys.* **196**, 645-679 (2004)
 18. E. J. Vigmond, F. Aguel, and N. A. Trayanova. Computational techniques for solving the bidomain equations in three dimensions. *IEEE Trans. Biomed. Eng.* **49**, 1260-1269 (2002)
 19. E. J. Vigmond, R. Weber dos Santos, A. J. Prassl, M. Deo, G. Plank. Solvers for the cardiac bidomain equations. *Progress in Biophysics & Molecular Biology* **96**, 3-18 (2008)