Considerations for the use of cellular electrophysiology models within cardiac tissue simulations

Jonathan Cooper[a,*], Alberto Corrias[a,1], David Gavaghan[a], Denis Noble[b]

[a] *Oxford University Computing Laboratory, University of Oxford, Oxford, UK*
[b] *Department of Physiology, Anatomy & Genetics, University of Oxford, Oxford, UK*

Abstract

The use of mathematical models to study cardiac electrophysiology has a long history, and numerous cellular scale models are now available, covering a range of species and cell types. Their use to study emergent properties in tissue is also widespread, typically using the monodomain or bidomain equations coupled to one or more cell models. Despite the relative maturity of this field, little has been written looking in detail at the interface between the cellular and tissue level models. Mathematically this is relatively straightforward and well-defined. There are however many details and potential inconsistencies that need to be addressed, in order to ensure correct operation of a cellular model within a tissue simulation. This paper will describe these issues and how to address them.

Simply having models available in a common format such as CellML is still of limited utility, with significant manual effort being required to integrate these models within a tissue simulation. We will thus also discuss the facilities available for automating this in a consistent fashion within Chaste, our robust and high-performance cardiac electrophysiology simulator.

It will be seen that a common theme arising is the need to go beyond a representation of the model mathematics in a standard language, to include additional semantic information required in determining the model's

*Corresponding author. Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford, UK, OX1 3QD.

Email addresses: jonathan.cooper@comlab.ox.ac.uk (Jonathan Cooper), alberto@nus.edu.sg (Alberto Corrias), david.gavaghan@comlab.ox.ac.uk (David Gavaghan), denis.noble@dpag.ox.ac.uk (Denis Noble)

[1] Present address: Division of Bioengineering, National University of Singapore, Singapore

interface, and hence to enhance interoperability. Such information can be added as metadata, but agreement is needed on the terms to use, including development of appropriate ontologies, if reliable automated use of CellML models is to become common.

Keywords: cardiac electrophysiology, model interfaces, CellML, Chaste

1. Introduction

Computational modelling of cardiac electrophysiology has developed extensively over the last 50 years at all spatial scales. Numerous models at the cellular scale are now available, covering a range of species, cell types, and experimental conditions. As well as being used to study cellular scale phenomena, these models are frequently also embedded within a framework that describes cardiac tissue, in order to investigate the propagation of electrical activity that gives rise to the heartbeat. Such emergent behaviour at the tissue level can be very effectively modelled (Clayton et al., 2010).

With its long history, this field is relatively mature. Many of the cellular level models have been curated and are available in a standard computer-readable format from the CellML model repository (Lloyd et al., 2008), providing easy access to checked encodings of the model equations. Several tissue-level simulation environments are also available (Niederer et al., submitted). However despite this, little has been written looking in detail at the interface between these levels. Perhaps this is because, as we shall describe shortly, this interface is relatively straightforward and well-defined from a mathematical point of view. However, in order to ensure correct operation of a cellular level model within a tissue simulation, there are many details and potential inconsistencies that must be taken into account and addressed.

Tissue-level cardiac electrophysiology is usually modelled using the bidomain equations (or the monodomain simplification thereof). These consist of two partial differential equations (PDEs) describing the intracellular and extracellular potential fields (ϕ_i and ϕ_e) through the cardiac tissue as a reaction-diffusion system, coupled at each point in space with a system of ordinary differential equations (ODEs) representing the concentrations of ions and other variables at the cellular level (see, for example, Keener and Sneyd, 1998). Let Ω denote the region occupied by the cardiac tissue. In the parabolic–elliptic formulation, which describes ϕ_e and the transmembrane voltage $V_m = \phi_i - \phi_e$,

the bidomain equations are then

$$\chi\left(\mathcal{C}_m\frac{\partial V_m}{\partial t}+I_{\text{ion}}(\boldsymbol{u},V_m)\right)-\nabla\cdot\left(\sigma_i\nabla\left(V_m+\phi_e\right)\right)=I_{\text{i}}^{(\text{vol})}, \tag{1a}$$

$$\nabla\cdot\left(\left(\sigma_i+\sigma_e\right)\nabla\phi_e+\sigma_i\nabla V_m\right)=-I_{\text{total}}^{(\text{vol})}, \tag{1b}$$

$$\text{and for each point in } \Omega,\quad \frac{\partial \boldsymbol{u}}{\partial t}=\boldsymbol{f}\left(\boldsymbol{u},V_m\right), \tag{1c}$$

where (1a) and (1b) describe the spatial and temporal evolution of the electical potentials, and (1c) describes the remaining cellular level behaviour. Here σ_i is the intracellular conductivity tensor, σ_e is the extracellular conductivity tensor, χ is the surface area to volume ratio, $\mathcal{C}_m$ is the membrane capacitance per unit area, $\boldsymbol{u}$ is a set of cell-level variables (such as gating variables, ion concentrations, etc.), and $I_{\text{ion}} \equiv I_{\text{ion}}(\boldsymbol{u}, V_m)$ is the transmembrane ionic current per surface area. $I_{\text{total}}^{(\text{vol})} = I_{\text{i}}^{(\text{vol})} + I_{\text{e}}^{(\text{vol})}$, where the source terms $I_{\text{i}}^{(\text{vol})}$ and $I_{\text{e}}^{(\text{vol})}$ are the intra- and extra-cellular stimuli per unit volume. Appropriate boundary and initial conditions must also be applied; see (Pathmanathan et al., 2010a) for details.

In this paper, we examine three classes of issues concerning the interoperability of CellML models of cellular electrophysiology within cardiac tissue simulations, and indicate possible strategies to address each of them.

The first class of issues arises from the fact that the cellular models available in the CellML model repository are not formulated as components of a tissue model, but represent entities somewhere between an isolated cell and a patch of cell membrane. They thus do not provide straightforward definitions of the terms I_{ion} and $\boldsymbol{f}$ as they appear in (1). Rather, their equations appear in the form

$$\frac{\mathrm{d}V_m}{\mathrm{d}t}=-\frac{I_{\text{ion}}(\boldsymbol{u},V_m)+I_{\text{stim}}}{\mathcal{C}_m}, \tag{2a}$$

$$\frac{\mathrm{d}\boldsymbol{u}}{\mathrm{d}t}=\boldsymbol{f}(\boldsymbol{u},V_m), \tag{2b}$$

or a variation thereof. The first challenge is therefore to identify the relevant variables within the cell model for coupling to the tissue model, and extract the necessary equations. This is the topic of Section 2.1.

The next class of issues, treated in Section 2.2, is that of inconsistencies between the models being connected. This may arise through differing use of units, from variations in how the models are structured, or from differences

in parameter values. The final class of issues we address, in Section 2.3, are those faced by software that is intended to be generic enough to be able to deal with different cell types. A sample simulation, illustrating the importance of addressing these issues correctly, is presented in Section 3.

These issues are of particular importance to those seeking to exploit the full potential of having cellular models available in a standard language such as CellML, by being able to reuse these models within tissue simulations without significant manual effort. For each issue, we therefore also describe the support available in Chaste (Pitt-Francis et al., 2009), via the PyCml software (Cooper, 2009; Garny et al., 2008), for automatically addressing it when processing a suitably annotated cell model. This allows the transparent inclusion of any cell model represented in CellML within a tissue simulation. The cardiac portion of Chaste is a powerful, efficient, parallel and well-engineered monodomain/bidomain solver. Chaste is open-source and available for download at `http://web.comlab.ox.ac.uk/chaste/`. All of the functionality described in this paper is present in the version 2.2 release.

We conclude in Section 4 by identifying common threads arising from these issues, and discuss some of the wider implications for the computational modelling field.

2. Issues to consider

The main focus of this paper is on the interfaces between cellular and tissue level models as systems of equations. Firstly however, it must be noted that the choice of numerical scheme for solving the equations also has an impact on this interface. This question has been investigated in some detail elsewhere and so will not be addressed here. For example, where and how in the spatial domain the cellular models are evaluated can have a significant impact on the accuracy of certain features of tissue simulations, such as the conduction velocity (Pathmanathan et al., 2011). The choice of stimulus currents is also important, and incompatibilities can arise from mixing formulations (Pathmanathan et al., 2010a).

2.1. Identification of model components

The first class of issues is concerned with the need to identify those portions of the cellular model which form the interface to the tissue model. This requires being able to handle variations in model variable naming conventions (different authors use different names for key quantities) and also variations

in how the model equations are structured in the CellML encoding of the model. These only cause difficulty when automating the reuse of models—if coding up the cellular model 'by hand' within a tissue simulator, the relevant portions are readily identified by eye.

The cellular model variables in (2) that form the interface to (1) are V_m, I_{ion}, and I_{stim}. Since this is a small set of quantities, one might first consider simply testing for all names known to be used. Besides being prone to failure on new models that might use a novel nomenclature, this approach also requires testing for a surprising number of alternatives, and may for some models identify incorrect variables, particularly in the case of I_{ion}. The multiplicity of options derives in part from the structure of CellML, in which models are decomposed into multiple components as an aid to reuse (Lloyd et al., 2004). There may thus be variation in both component names (typically cell or membrane are used for the external interface of the model) and variable names. Common cases for the latter include V, E, or Vm for V_m, and i_Stim, I_st, or i_pulse for I_{stim}, with variations in capitalisation also, as can be seen in Table 1. Identification of I_{ion} is further complicated by variations in the structure of the model equations, described in more detail in Section 2.2. In particular, there is often not a single variable for I_{ion} in the model; rather it is implicit as a sum of component currents within the equation for $\mathrm{d}V_m/\mathrm{d}t$. Hence analysis of the equation is required for robust automatic identification of this term, since matching purely on variable names could potentially include currents which are internal to the cell, rather than only the transmembrane currents comprising I_{ion}.

Some manual configuration for each model is thus required for a truly robust solution. While this could be done using a separate configuration file (as was done in earlier Chaste releases) a better and more extensible approach is to annotate model variables using terms from a suitable ontology, building on the CellML metadata framework (Beard et al., 2009). The ontology provides standardised names for important quantities, eliminating the variation in naming. Chaste currently uses its own mini-ontology for naming, which consists of a short list of identifiers built in to the software.[2] However, it is intended to work towards a community standard, so that models available from the CellML repository are pre-annotated and ready to use directly.

[2]See `https://chaste.comlab.ox.ac.uk/chaste/tutorials/release_2.2/ChasteGuides/CodeGenerationFromCellML.html` for more details.

Model	**Component name**	V_m	I_{stim}	$\mathcal{C}_m$	**Units of currents**
DiFrancesco and Noble 1985	membrane	V	i_pulse	C	nA
Luo and Rudy 1991	membrane	V	I_stim	C	μA/cm^2
Dokos et al. 1996	membrane	E	N/A	C	pA
Noble et al. 1998	membrane	V	i_Stim	Cm	nA
Jafri et al. 1998	membrane	V	I_stim	Cm	μA/mm^2
Viswanathan et al. 1999	membrane	V	I_st	Cm	μA/μF
Matsuoka et al. 2003	membrane	Vm	i_ext	Cm	pA
Hund and Rudy 2004	cell	V	i_Stim	Acap	μA/μF
Mahajan et al. 2008	cell	V	i_Stim	N/A	μA/μF
Stewart et al. 2009	membrane	V	N/A	Cm	pA/pF

Table 1: Names used for key variables, and units used for ionic currents, in a small sample of cellular models

However, annotating every variable does introduce an additional task, whether for the model author, curator, or user. To reduce this burden, PyCml is capable of analysing the model mathematics to identify most of the interface. Only V_m and I_{stim} are required to be explicitly identified. From these, we can determine I_{ion} as described in the next section. Finding the derivatives defined in the model is straightforward, and so extraction of $\mathrm{d}\boldsymbol{u}/\mathrm{d}t$ simply requires ignoring $\mathrm{d}V_m/\mathrm{d}t$. The subsidiary equations in the model can then be analysed to determine which are required in computing $\mathrm{d}\boldsymbol{u}/\mathrm{d}t$ and I_{ion}, and suitable C++ code generated for use in Chaste (Cooper, 2009).

2.2. Inconsistencies when connecting models

A common problem arising in any model coupling exercise is the existence of incompatibilities between the component models, due to differing conventions or methodologies followed in developing the constituent pieces. Three categories of such incompatibilities were defined by Terkildsen et al. (2008) as unit, structural, and parameter inconsistencies. These can usefully be applied in our scenario, although frequently multiple types of inconsistency occur together.

The simplest case to resolve is that of *unit inconsistencies*. These occur

when the cellular and tissue level models use different physical units for the same quantity. A common example is the representation of time, which is generally measured in either milliseconds (as in Chaste) or in seconds. In the latter case a conversion is required; it is a straightforward scaling in this case since the dimensions match. Since all quantities in a CellML model must be explicitly associated with their units, it is possible to check automatically that the cellular model is internally consistent and add suitable conversions at component interfaces (Cooper and McKeever, 2008). (If the cellular model is not internally consistent, no conversions can be performed reliably, and the model must be corrected to ensure correct behaviour.) The interface to the tissue model can be addressed using the same functionality, by adding a new component to the CellML model containing variables in the units required by the tissue model, and connecting these to the corresponding variables in the cellular model. The tissue model then interacts only with this new interface component.

Other quantities at the interface must also be given in consistent units. For V_m we know of no model that does not use millivolts (mV), but for ionic currents the situation is more complex due to the existence of *structural inconsistencies*. These refer to differences in how the biological system is represented by the mathematical equations. There are three different conventions used in cellular models for representing transmembrane ionic currents. The first is to use whole-cell current, which is expressed in (multiples of) Amperes (A) (e.g. Noble et al., 1998). Alternatively the current may be normalised, either by the cell membrane surface area (e.g. Luo and Rudy, 1991), or by its capacitance (e.g. Stewart et al., 2009), leading to units dimensions of $\mathrm{A/m^2}$ or A/F respectively. If the cellular and tissue level models use different conventions, a simple scaling conversion is impossible. Furthermore, where normalisation is employed, *parameter inconsistencies* may also play a role, with the cellular and tissue models using different estimates for the same biological quantity (whether cell surface area or membrane capacitance).

Fortunately, we can make use of the parameter inconsistency to perform a suitable conversion for ionic currents into the units expected by Chaste, $\mu\mathrm{A/cm^2}$, automatically. The three cases are as follows.

1. Cellular model uses current per unit area ($\mathrm{A/m^2}$). These are the same dimensions as used by Chaste for the tissue model, and so a simple scaling is sufficient.
2. Cellular model uses current per unit capacitance (A/F). In this case we

note that the units of the ionic current are dimensionally equivalent to those for $\mathrm{d}V_m/\mathrm{d}t$, and so the equation for $\mathrm{d}V_m/\mathrm{d}t$ should not include a scaling by $\mathcal{C}_m$ as in (2)—this is already incorporated into the currents. We thus need only to change the normalisation by scaling using Chaste's estimate for the membrane capacitance, measured in $\mu\mathrm{F/cm}^2$, to obtain currents in the expected dimensions.

3. Cellular model uses whole-cell current (A). Ideally, the conversion should be done by dividing by the (electrically active) cellular surface area defined by the cellular model. However, this information is not always available. Instead, we estimate it using the membrane capacitance, which is given in pure capacitance units (F) within the cellular model, but in $\mu\mathrm{F/cm}^2$ by the tissue model. Since these are conceptually supposed to represent the same quantity, the ratio of the two yields an estimate of the electrically active area. Hence the only additional manual effort required in this case is annotation of $\mathcal{C}_m$ in the same way as for V_m. Note that this also assumes that the cellular model represents a single cell, and hence does not apply to models such as that of DiFrancesco and Noble (1985) which consider a multicellular preparation. In such cases, specific configuration is required.

If the tissue model uses a different convention from that used in Chaste, similar conversions can be applied.[3]

There are further structural inconsistencies involved in the connection of ionic currents between cellular and tissue level. The first concerns the intracellular stimulus current, given by $I_{\mathrm{i}}^{(\mathrm{vol})}$ in the tissue model and I_{stim} in the cellular model. Within Chaste, the former is measured in $\mu\mathrm{A/cm}^3$ and the latter in $\mu\mathrm{A/cm}^2$, as is I_{ion}. When the cellular model is being used in a tissue simulation, $\mathrm{d}V_m/\mathrm{d}t$ is not evaluated, and so I_{stim} is not required there. However, for models such as that of Hund and Rudy (2004) which assign the stimulus current to particular ionic species in order to conserve charge, I_{stim} is still used in computing $\mathrm{d}\boldsymbol{u}/\mathrm{d}t$. It must therefore be calculated as

[3]This approach to units conversions is conceptually similar to the process of non-dimensionalisation, in which quantities intrinsic to a system are used to scale the equations, yielding a new system using only dimensionless variables. While this technique is common in mathematical biology, it has not generally been applied in physiological modelling; perhaps due both to the complexity of the models and a desire to keep parameters in physical units.

$I_{\text{stim}} = I_{\text{i}}^{(\text{vol})}/\chi$, where χ is given in cm^{-1}.

A particular advantage we have observed of performing these conversions is that it reduces the amount of per-model 'tweaking' required to achieve a successful tissue simulation. For example, the same stimulus can now be applied by Chaste to produce an action potential in each of a sample of 32 cellular models. This is in stark contrast to the default stimuli coded in the models, which vary significantly.

Secondly, as we have alluded to already, the form of Equation (2a) can vary significantly. The convention followed by the majority of cell models, and expected by Chaste, is for (2a) to have the form

$$\frac{\mathrm{d}V_m}{\mathrm{d}t} = -\frac{I_1 + \ldots + I_n + I_{\text{stim}}}{\mathcal{C}_m}, \tag{3}$$

whence the input to (1), I_{ion}, is implicit and given by $I_{\text{ion}} = I_1 + \ldots + I_n$. The CellML versions of some models (e.g. Priebe and Beuckelmann, 1998) introduce intermediate equations, for example $\mathrm{d}V_m/\mathrm{d}t = \text{dVdt}$, $\text{dVdt} = -I_{\text{tot}}$, $I_{\text{tot}} = I_1 + \ldots + I_n + I_{\text{stim}}$. Other models (e.g. Mahajan et al., 2008) use the opposite sign for the currents comprising I_{ion}, whereas others (e.g. Jafri et al., 1998) instead invert just I_{stim}. Each of these variations can be automatically accounted for by PyCml, by analysing the model equations in the following steps.

1. Firstly, the currents forming I_{ion} must be identified. This is done by examining the variables appearing in (2a) and selecting those with suitable units (excluding I_{stim}). There are a few subtleties to this. If I_{stim} is defined in the model (i.e. it is not self-excitatory) then 'suitable units' means dimensionally equivalent to the units of I_{stim}. Otherwise, the dimensions could be any of the three options for ionic currents discussed above. Once I_{ion} has been identified, the convention actually used in the model can be determined.

   The equations defining the model are thus systematically searched, starting from the definition of $\mathrm{d}V_m/\mathrm{d}t$, for variables with suitable units. If any are found, the sum of these (with the exception of I_{stim}) is considered to be I_{ion}. If none are found at first, the definitions of variables occurring in this equation are then searched; if no suitable variables are found, the definitions of variables occurring in those definitions are searched, and so on. The search is performed in this fashion in order to ensure that currents are not double-counted, by including both a

variable representing a sum of currents and the variables representing the individual currents included in the sum.

One complication is that A/F is dimensionally equivalent to mV/ms, and so if intermediate equations are used, a variable such as dVdt could be mistakenly identified as an ionic current, despite the model using (say) the whole-cell current convention. If I_{stim} is not present (and so the convention used is unknown), the search is thus performed first with A/F excluded from the units options, stopping after two levels of definitions have been examined. This assumes that ionic currents will occur either in the equation defining $\mathrm{d}V_m/\mathrm{d}t$, in one of the equations defining variables occurring therein, or in the definitions of variables occurring in those equations. A model would need to be structured very unusually to break this assumption, which is made partly for efficiency, but primarily to guard against the possibility of features located deeper within the equations being misidentified as currents, e.g. if an ionic current was defined by dividing something measured in amps by a capacitance. No existing models do so as far as we are aware. If no suitable currents are found, the search is restarted with A/F included, and no early termination.

A further complication is that I_{stim} may occur within the definition of currents identified as part of I_{ion} (this can occur if intermediate equations are used, or with conservative cell models). A check for this case is made when generating the code for I_{ion}, and if I_{stim} is found, it is replaced by zero.

2. Secondly, the sign of I_{ion} must be checked. Since PyCml already includes the ability to evaluate portions of a CellML model (Cooper, 2009), we can utilise this, faking the values assigned to variables, to compute the sign of I_{ion} by evaluating the right-hand side of (2a). Each of the currents identified in step 1 is (temporarily) assigned the value 1, as are other time-varying variables.[4] If $\mathrm{d}V_m/\mathrm{d}t$ evaluates strictly positive, I_{ion} is then considered to be negated with respect to (3).
3. Finally, if it is present, the sign of I_{stim} must be determined. Again this can be deduced by fake evaluation of (2a), but with different values assigned to the variables. In this case, I_{stim} is set to 1, other ionic

[4]Constant variables retain their values given in the model, in order to account for pathological cases, not seen in any models thus far, such as $\mathrm{d}V_m/\mathrm{d}t = X \cdot I_{\text{ion}}/\mathcal{C}_m$ with $X = -1$.

currents are set to 0, and all other time-varying variables are set to 1. The stimulus is then negated (as compared with (3)) if $\mathrm{d}V_m/\mathrm{d}t$ evaluates strictly positive.

2.3. Issues for generic simulation software

Other issues are faced by software such as Chaste which is intended to be generic, dealing with models of different cell types, and providing facilities for both single-cell and tissue level simulation. The primary example concerns identification of the intracellular stimulus I_{stim}. This will appear in models of ventricular cells, but generally not in self-excitatory cells such as from the sino-atrial node. If no stimulus is found, and the model represents a ventricular cell, processing software should report an error since this indicates that the stimulus has not been suitably annotated (as described in Section 2.1). Otherwise, the stimulus current would be considered as part of I_{ion}, and so an additional stimulus would be applied to the cellular model, leading to incorrect results.

If, however, the cellular model represents a self-excitatory cell, it may not contain a stimulus current, and so no error should be reported in this case. (It should be noted that it is *not* an error if a self-excitatory cell *does* contain a stimulus—a Purkinje cell may exhibit both spontaneous and paced behaviour in vivo, and some models (e.g. Sampson et al., 2010) therefore consider both activities.) Hence annotation of the model itself is required (or some alternative method of configuration) to indicate the type of cell being modelled, and hence whether to expect a stimulus current.

Various other technical issues, not central to the thrust of this paper, may also be considered. For example, we also allow metadata annotations that specify minimum and maximum possible values for variables (e.g. to specify a probability that must lie between zero and one, or a concentration that must be positive). Including these enables Chaste to check during a simulation whether a variable has gone out of range, and terminate the simulation early with an error if this occurs. Such a situation is typically due to parameters chosen for the numerical method being unsuitable.

3. Example simulation

To demonstrate the functionality available in Chaste for addressing the issues described above, proof-of-concept simulation results are presented here

(input files and a movie from a sample simulation are present in the supplementary material). The simulation consists of a 1 cm fibre with the first 3 mm comprising self-excitatory Purkinje cells (Stewart et al., 2009), and the remainder ventricular cells (Noble et al., 1998). An action potential is initiated by the Purkinje cells, and propagates along the fibre. The cellular models are chosen to illustrate both the handling of multiple cell types described in this section, and the units conversions for ionic currents described in Section 2.2—the models use different conventions, both differing from Chaste's, as can be seen from Table 1. The Purkinje model also illustrates a challenging case for identification of I_{ion}, since it does not include a stimulus current, and has currents given in dimensions of A/F.

To demonstrate the importance of interfacing models properly, six variations of the simulation are presented, half with conversions and half omitting conversion of the ionic currents. Note that simulation time must be units converted (the cellular models use different units) or the numerical method fails. Three different values for the tissue model membrane capacitance $\mathcal{C}_m$ are also used, to show the effect it has on the coupling. This effect is difficult to distinguish precisely, since $\mathcal{C}_m$ also has an effect on action potential propagation within the tissue model.

Figure 1 shows the action potential at two nodes, one a Purkinje node and one a ventricular node, with and without an appropriate interface conversion of units for the ionic currents. In Purkinje cells differences are small because the ionic current conversion for the Purkinje cell model involves merely scaling by $\mathcal{C}_m$. For ventricular cells, on the other hand, where the difference between the ionic current units in the model (nA) and the ones expected by Chaste (μA/cm^2) is significant, the errors become extremely large.

Figure 2 shows the conduction velocity along the cable with and without appropriate units conversion for the ionic currents. Results for three different values of membrane capacitance are shown. For Purkinje nodes, where the units conversion is a linear scaling by the value of capacitance, differences are small and inversely related to the variations of the cell capacitance, i.e. failing to convert the units of the ionic currents caused an underestimation of conduction velocity when the cell capacitance is increased, while it caused an overestimation of conduction velocity when the cell capacitance is decreased. For the ventricular cells, on the other hand, the I_{ion} term is also divided by the model's value for the membrane capacitance, and since this value (in nF) is less than 1 the I_{ion} term is smaller when conversions are not used, leading to reduced conduction velocity for all cases.

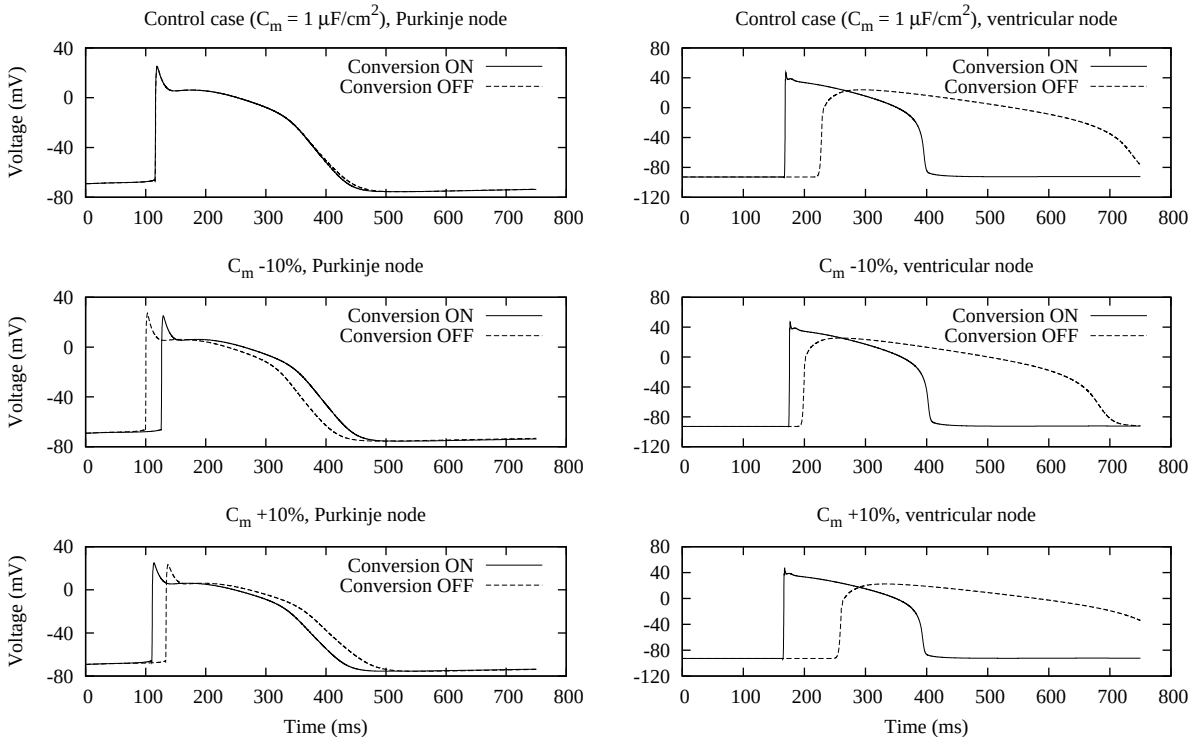

Figure 1: Effects of omission of appropriate units conversion for the ionic currents on the action potential profile at two nodes along the cable: node number 10 (a Purkinje node, left panels) and a node at the end of the cable (a ventricular node, right panels). Action potentials are plotted for the case of appropriate units conversion ("Conversion ON", solid lines) and without appropriate units conversion ("Conversion OFF", dashed lines) for three different values of membrane capacitance: the top panels show the control case ($\mathcal{C}_m = 1\,\mu\mathrm{F/cm^2}$), the middle panels relate to the case of capacitance reduced by 10%, while the bottom panels correspond to the case of capacitance increased by 10%.

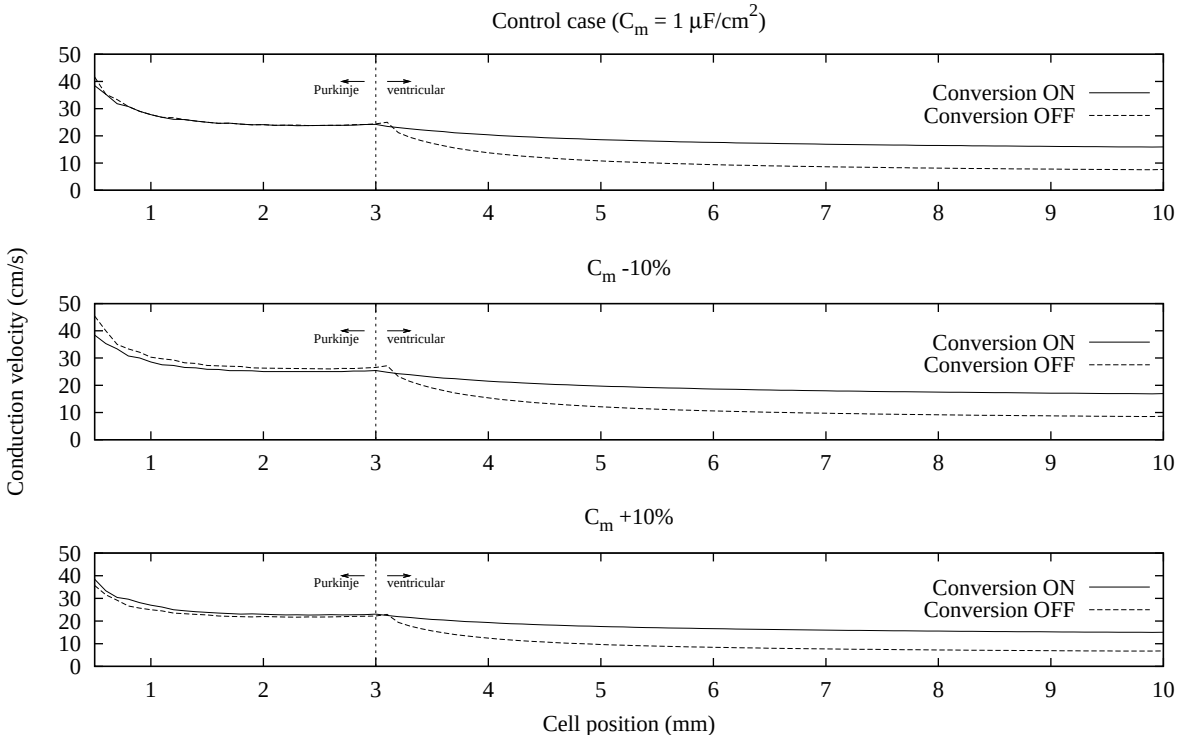

Figure 2: Effects of omission of appropriate units conversion for the ionic currents on conduction velocity. The abscissa represents the position of the cell along the 1D cable (the first five cells are not shown, as end effects would distort the y-axis scale). The vertical dashed line marks the junction bewteen Purkinje cells and ventricular cells. Conduction velocities are plotted for the case of appropriate units conversion ("Conversion ON", solid lines) and without appropriate units conversion ("Conversion OFF", dashed lines) for three different values of membrane capacitance: control case ($\mathcal{C}_m = 1\,\mu\mathrm{F/cm}^2$, top panel), reducing capacitance by 10% (middle panel), and increasing capacitance by 10% (bottom panel).

4. Discussion and conclusions

It is clear from the issues described above that, despite having easy access to mathematical models in standard formats, there are still significant limitations to easy model reuse if these model descriptions contain merely the mathematical equations making up the model. This applies even if the equations themselves have been curated and verified as a correct representation, and arranged logically into components. Additional semantic information is required in order to determine the interface of a model, and hence to enable interoperability.

We have described approaches to resolving many of the issues that arise in the context of incorporating cellular electrophysiology models within a monodomain or bidomain framework. Similar issues are likely to arise in other contexts also. They may involve merely identifying which variables represent the same biological concept, as discussed in Section 2.1, a requirement for any model coupling exercise. More complex logic may also be required, especially when coupling models that use different mathematical formulations (such as coupling ordinary and partial differential equations, as in this case). A model's equations may require re-arrangement in order to be used within the other formulation, as in Section 2.2.

A closely related context is that of multiphysics models of the heart, in particular coupling mechanics to the electrophysiology. As discussed by Terkildsen et al. (2008), similar issues indeed arise purely at the cellular level. In that case the parameter inconsistencies were considered and addressed manually, but one could envisage a more automated approach being possible were this to be generalised to more cell models. Support for electromechanical coupling within Chaste is currently being investigated (see e.g. Pathmanathan et al., 2010b) and the framework described in this paper will be extended accordingly. This framework can already be used to incorporate models of fibroblasts, and a publication on this topic is in preparation.

The situation is particularly challenging when dealing with a modelling language that is general in scope, such as CellML. This can require further effort on the part of tool developers in order to analyse the mathematics of models and determine the underlying biological structure, as we have done. Languages targeted at a particular domain have the potential to be designed so as to make this knowledge explicit, and hence will be easier to process. The disadvantage of this approach is that, naturally, a specific language is needed for each biological domain, potentially resulting in reduced uptake by

the modelling community.

The use of metadata encoded using the standard Resource Description Framework (RDF, `http://www.w3.org/RDF`) according to the CellML metadata framework does provide a generic approach to adding such semantic information to model descriptions. Using a standard format for such annotations has several benefits, including allowing them to be stored with the 'canonical' version of the model in the CellML repository (even though it may not 'understand' the semantic meaning of specific annotations used). However, to realise the full potential of this approach, further agreement is required on how to annotate models—going beyond the framework to the specific terms to use. This necessitates the development of ontologies by the modelling community, codifying knowledge of how biological systems, and the models representing them, are structured. Only through such standardisation can the utility of the models available from the CellML repository be enhanced, and reliable automated use of CellML models within tissue-level simulations become common.

Within our framework at present, the equations of the tissue model, and hence the specification of its desired interface, are hardcoded within the Chaste source code. It would be desirable to have this portion of the coupled model also available encoded in a markup language, and work on FieldML (Christie et al., 2009) is progressing in this direction. Interfacing between CellML and FieldML models will require careful consideration of issues such as those we have discussed.

Finally, while the work described above makes it possible to use any cellular model within a tissue simulation automatically, this does not imply that doing so for a particular scenario is in any way biologically realistic. A particular cell model may be unsuited to use in a tissue context, or it may have been developed to represent particular experimental conditions, with parameter values and initial conditions specified accordingly, and give unexpected results when used outside that regime. Further work is thus required on describing the scientific questions which the model was developed to address, and hence assisting users in determining its suitability for use in their study. Some initial work on this topic is presented by Cooper et al. (2011, this issue).

Acknowledgements

The authors would like to thank all other members of the Chaste development team for fruitful discussions on the topics considered in this manuscript.

Role of the funding source

JC is partially supported by the European Commission DG-INFSO under grant numbers 223920 (VPH-NoE) and 224381 (preDiCT). AC was also partially supported under grant number 224381 (preDiCT). The funding bodies had no other role in this work.

Editors' note

Please see also related communications in this issue by Quinn et al. (2011) and Bradley et al. (2011).

References

Beard, D.A., Britten, R., Cooling, M.T., Garny, A., Halstead, M.D.B., Hunter, P.J., Lawson, J., Lloyd, C.M., Marsh, J., Miller, A., Nickerson, D.P., Nielsen, P.M.F., Nomura, T., Subramanium, S., Wimalaratne, S.M., Yu, T., 2009. CellML metadata standards, associated tools and repositories. Physical and Engineering Sciences 367, 1845–1867.

Bradley, C., Bowery, A., Britten, R., Budelmann, V., Camara, O., 2011. OpenCMISS: A multi-physics & multi-scale computational infrastructure for the VPH/Physiome project. Progress in Biophysics and Molecular Biology NUMBER, [note to publisher: please update before print].

Christie, G.R., Nielsen, P.M.F., Blackett, S.A., Bradley, C.P., Hunter, P.J., 2009. FieldML: concepts and implementation. Phil Trans R Soc A 367, 1869–1884.

Clayton, R.H., Bernus, O., Cherry, E.M., Dierckx, H., Fenton, F.H., Mirabella, L., Panfilov, A.V., Sachse, F.B., Seemann, G., Zhang, H., 2010. Models of cardiac tissue electrophysiology: progress, challenges and open questions. Progress in Biophysics and Molecular Biology 104, 22–48.

Cooper, J., 2009. Automatic validation and optimisation of biological models. Ph.D. thesis. University of Oxford.

Cooper, J., McKeever, S., 2008. A model-driven approach to automatic conversion of physical units. Softw Pract Exper 38, 337–359.

Cooper, J., Mirams, G., Niederer, S., 2011, this issue. High throughput functional curation of cellular models. Progress in Biophysics and Molecular Biology NUMBER, [note to publisher: please update before print].

DiFrancesco, D., Noble, D., 1985. A model of cardiac electrical activity incorporating ionic pumps and concentration changes. Phil Trans Roy Soc B 307, 353–398.

Dokos, S., Celler, B., Lovell, N., 1996. Ion currents underlying sinoatrial node pacemaker activity: a new single cell mathematical model. Journal of Theoretical Biology 181, 245–272.

Garny, A., Nickerson, D., Cooper, J., dos Santos, R.W., McKeever, S., Nielsen, P., Hunter, P., 2008. CellML and associated tools and techniques. Phil Trans R Soc A 366, 3017–3043.

Hund, T.J., Rudy, Y., 2004. Rate dependence and regulation of action potential and calcium transient in a canine cardiac ventricular cell model. Circulation 110, 3168–3174.

Jafri, M.S., Rice, J.J., Winslow, R.L., 1998. Cardiac Ca2+ dynamics: the roles of ryanodine receptor adaptation and sarcoplasmic reticulum load. Biophys J 74, 1149–1168.

Keener, J., Sneyd, J., 1998. Mathematical physiology. volume 8 of *Interdisciplinary Applied Mathematics*. Springer.

Lloyd, C.M., Halstead, M.D., Nielsen, P.F., 2004. CellML: its future, present and past. Progress in Biophysics and Molecular Biology 85, 433–450.

Lloyd, C.M., Lawson, J.R., Hunter, P.J., Nielsen, P.F., 2008. The CellML model repository. Bioinformatics 24, 2122–2123.

Luo, C., Rudy, Y., 1991. A model of the ventricular cardiac action potential: depolarization, repolarization, and their interaction. Circ Res 68, 1501–1526.

Mahajan, A., Shiferaw, Y., Sato, D., Baher, A., Olcese, R., Xie, L.H., Yang, M.J., Chen, P.S., Restrepo, J.G., Karma, A., Garfinkel, A., Qu, Z., Weiss, J.N., 2008. A rabbit ventricular action potential model replicating cardiac dynamics at rapid heart rates. Biophysical Journal 94, 392–410.

Matsuoka, S., Sarai, N., Kuratomi, S., Ono, K., Noma, A., 2003. Role of individual ionic current systems in ventricular cells hypothesized by a model study. The Japanese Journal of Physiology 53, 105–123.

Niederer, S.A., Kerfoot, E., Benson, A., Bernabeu, M.O., Bernus, O., Bradley, C., Cherry, E.M., Clayton, R., Fenton, F.H., Garny, A., Heidenreich, E., Land, S., Maleckar, M., Pathmanathan, P., Plank, G., Rodríguez, J.F., Roy, I., Sachse, F.B., Seemann, G., Skavhaug, O., Smith, N.P., submitted. N-version benchmark evaluation of cardiac tissue electrophysiology simulators. Phil Trans R Soc A .

Noble, D., Varghese, A., Kohl, P., Noble, P., 1998. Improved guinea-pig ventricular cell model incorporating a diadic space, i_{K_r} and i_{K_s}, length- and tension-dependent processes. Canadian Journal of Cardiology 14, 123–134.

Pathmanathan, P., Bernabeu, M.O., Bordas, R., Cooper, J., Garny, A., Pitt-Francis, J.M., Whiteley, J.P., Gavaghan, D.J., 2010a. A numerical guide to the solution of the bidomain equations of cardiac electrophysiology. Progress in Biophysics and Molecular Biology 102, 136–155.

Pathmanathan, P., Chapman, S.J., Gavaghan, D.J., Whiteley, J.P., 2010b. Cardiac electromechanics: the effect of contraction model on the mathematical problem and accuracy of the numerical scheme. The Quarterly Journal of Mechanics and Applied Mathematics 63, 375–399.

Pathmanathan, P., Mirams, G., Southern, J., Whiteley, J., 2011. The significant effect of the choice of ionic current integration method in cardiac electro-physiological simulations. International Journal for Numerical Methods in Biomedical Engineering .

Pitt-Francis, J., Pathmanathan, P., Bernabeu, M.O., Bordas, R., Cooper, J., Fletcher, A.G., Mirams, G.R., Murray, P., Osbourne, J.M., Walter, A., Chapman, J., Garny, A., van Leeuwen, I.M.M., Maini, P.K., Rodriguez, B., Waters, S.L., Whiteley, J.P., Byrne, H.M., Gavaghan, D.J., 2009. Chaste:

a test-driven approach to software development for biological modelling. Computer Physics Communications 180, 2452–2471.

Priebe, L., Beuckelmann, D.J., 1998. Simulation study of cellular electric properties in heart failure. Circulation Research 82, 1206–1223.

Quinn, T., Antzelevitch, C., Bollensdorf, C., Bub, G., RAB., B., Chen, P., 2011. Minimum Information about a Cardiac Electrophysiology Experiment (MICEE): Standardised reporting for model reproducibility, interoperability, and data sharing. Progress in Biophysics and Molecular Biology NUMBER, [note to publisher: please update before print].

Sampson, K.J., Iyer, V., Marks, A.R., Kass, R.S., 2010. A computational model of Purkinje fibre single cell electrophysiology: implications for the long QT syndrome. The Journal of Physiology 588, 2643–2655.

Stewart, P., Aslanidi, O.V., Noble, D., Noble, P.J., Boyett, M.R., Zhang, H., 2009. Mathematical models of the electrical action potential of Purkinje fibre cells. Phil Trans R Soc A 367, 2225–2255.

Terkildsen, J.R., Niederer, S., Crampin, E.J., Hunter, P., Smith, N.P., 2008. Using Physiome standards to couple cellular functions for rat cardiac excitation–contraction. Experimental Physiology 93, 919–929.

Viswanathan, P.C., Shaw, R.M., Rudy, Y., 1999. Effects of IKr and IKs heterogeneity on action potential duration and its rate dependence: a simulation study. Circulation 99, 2466–2474.