

Markov chain models of coupled calcium channels: Kronecker representations and iterative solution methods

Hilary DeRemigio^{1,4}, M Drew LaMar^{1,4}, Peter Kemper²
and Gregory D Smith^{1,3}

¹ Department of Applied Science, The College of William and Mary, Williamsburg, VA 23187, USA

² Department of Computer Science, The College of William and Mary, Williamsburg, VA 23187, USA

³ Mathematical Biosciences Institute, The Ohio State University, Columbus, OH 43210, USA

E-mail: greg@as.wm.edu

Received 2 April 2008

Accepted for publication 17 June 2008

Published 14 July 2008

Online at stacks.iop.org/PhysBio/5/036003

Abstract

Mathematical models of calcium release sites derived from Markov chain models of intracellular calcium channels exhibit collective gating reminiscent of the experimentally observed phenomenon of stochastic calcium excitability (i.e., calcium puffs and sparks). Calcium release site models are stochastic automata networks that involve many functional transitions, that is, the transition probabilities of each channel depend on the local calcium concentration and thus the state of the other channels. We present a Kronecker-structured representation for calcium release site models and perform benchmark stationary distribution calculations using both exact and approximate iterative numerical solution techniques that leverage this structure. When it is possible to obtain an exact solution, response measures such as the number of channels in a particular state converge more quickly using the iterative numerical methods than occupation measures calculated via Monte Carlo simulation. In particular, multi-level methods provide excellent convergence with modest additional memory requirements for the Kronecker representation of calcium release site models. When an exact solution is not feasible, iterative approximate methods based on the power method may be used, with performance similar to Monte Carlo estimates. This suggests approximate methods with multi-level iterative engines as a promising avenue of future research for large-scale calcium release site models.

 This article has associated online supplementary data files

1. Introduction

The stochastic gating of voltage- and ligand-gated ion channels in biological membranes observed by single-channel recording techniques is often modeled using continuous-time discrete-state Markov chains (CTMCs) [1, 2]. While the scientific literature developing stochastic models for the behavior of ion channels is largely focused on single channels or populations of independent channels, the application and extension of these techniques to the collective gating of interacting ion

channels is an important topic of current research. For example, interacting aggregated CTMCs have been used by Ball and colleagues to simulate and analyze membrane patches containing several ion channels [3, 4]. A second example is the simulation of plasma membrane receptor arrays involved in bacterial chemotaxis where conformational energies (and thus transition rates) depend on the state of neighboring receptors [5–7]. A third example, the subject of this paper, are simulations of clusters of intracellular Ca^{2+} -regulated Ca^{2+} channels—1,4,5-trisphosphate receptors (IP_3Rs) and ryanodine receptors (RyRs) located on the surface of the endoplasmic reticulum or sarcoplasmic reticulum

⁴ These authors contributed equally to this work.

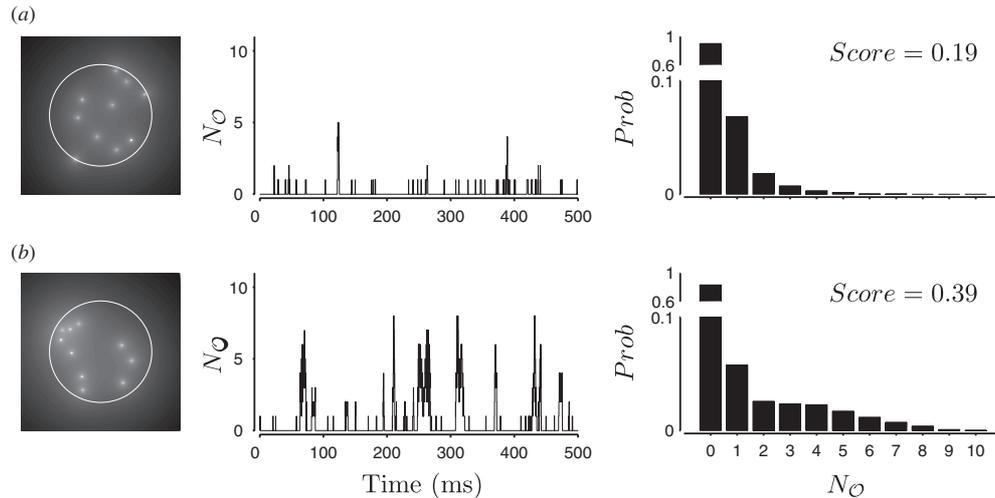


Figure 1. (a) Left: $[Ca^{2+}]$ near $3 \times 3 \mu m^2$ endoplasmic reticulum membrane with 12 Ca^{2+} -regulated Ca^{2+} channels modeled as three-state Markov chains (see figure 2(a)) with positions randomly chosen from a uniform distribution on a disc of radius $2 \mu m$ (source amplitude $0.05 \mu A$). Buffered Ca^{2+} diffusion is modeled as in [13] (see appendix A). Middle: stochastic dynamics of the number of open channels at the release site (N_O) that does not include robust puffs/sparks. Right: probability distribution of the number of open channels leading to a low puff/spark *Score* of 0.19. (b) Different random channel positions result in a release site that exhibits robust Ca^{2+} puff/sparks (middle) and an elevated *Score* of 0.39 (right).

membrane—that give rise to localized intracellular $[Ca^{2+}]$ elevations known as Ca^{2+} puffs and sparks [8–12].

When Markov chain models of Ca^{2+} -regulated Ca^{2+} channels are coupled via a mathematical representation of buffered diffusion of intracellular Ca^{2+} , simulated Ca^{2+} release sites may exhibit the phenomenon of ‘stochastic Ca^{2+} excitability’ where the IP₃R or RyRs open and close in a concerted fashion [14, 13] (see figure 1 for representative simulations). Such models are stochastic automata networks (SANs) that involve a large number of functional transitions, that is, the transition probabilities of one automata (i.e., an individual channel) may depend on the local $[Ca^{2+}]$ and thus the state of the other channels. Because the number of channels in the open class of states, $N_O(t)$, can in principle be back-calculated from microfluorometric measurements of elevated local $[Ca^{2+}]$, our simulations and analysis focus on the stochastic dynamics of $N_O(t)$ (see figures 1(a) and (b), middle).

While the relationship between single-channel kinetics of Ca^{2+} -regulated channels and the collective phenomenon of Ca^{2+} puffs and sparks is not fully understood, several groups have presented mathematical modeling studies of Ca^{2+} release sites that provide insight into the emergent properties of stochastic Ca^{2+} excitability [13–23]. For example, it has been shown that allosteric interactions between intracellular Ca^{2+} channels may lead to synchronous gating [17, 24], but such direct coupling is not required [13, 14, 23]. Rather, Ca^{2+} puffs and sparks can readily be observed when the coupling between single-channel models is mediated entirely via the buffered diffusion of intracellular Ca^{2+} simulated through numerical solution of a system of nonlinear reaction–diffusion equations. In the above mentioned studies, the specific single-channel model chosen, the release site geometry, and the description of the cytosolic milieu all contribute to the measured statistics of simulated puffs and sparks such as amplitude, duration

and inter-event interval. The IP₃R or RyR models used often include transitions representing fast Ca^{2+} activation and slower Ca^{2+} inactivation, two phenomena that have been repeatedly (but not uniformly) observed in single-channel recordings from planar lipid bilayer and nuclear patch experiments [25–31].

Many of the studies mentioned above assume mean-field Ca^{2+} coupling where the channels have no explicit spatial positions at the release site. These reduced simulations assume that the local $[Ca^{2+}]$ experienced by each channel depends on the number of open channels at the Ca^{2+} release site, as though the channels were indistinguishable. In prior work we have shown that mean-field simulations are often in reasonable agreement with results obtained using spatially explicit release site models [32]. However, in some cases channel position can significantly influence release site dynamics. For example, figures 1(a) and (b) show two representative Ca^{2+} release site simulations that are identical except for channel positions (left). These differences in channel positions lead to minimal release site activity in figure 1(a) (middle) but robust release site activity in figure 1(b) (middle) that is reminiscent of the phenomenon of Ca^{2+} puffs or sparks. Note that the presence or absence of puff/sparks in Ca^{2+} release site simulations can be determined from the steady-state distribution of the number of open channels at the release site (figure 1, right) using a response measure dubbed the puff/spark *Score* [13],

$$Score = \frac{\text{Var}[f_O]}{\text{E}[f_O]} = \frac{1}{N} \frac{\text{Var}[N_O]}{\text{E}[N_O]}, \quad (1)$$

where $f_O = N_O/N$ is the fraction of open channels. The puff/spark *Score* takes values between 0 and 1, and a *Score* of greater than approximately 0.3 indicates the presence of robust stochastic Ca^{2+} excitability (as in figure 1(b)).

While response measures such as the puff/spark *Score* and the probability distribution of N_O can be estimated via

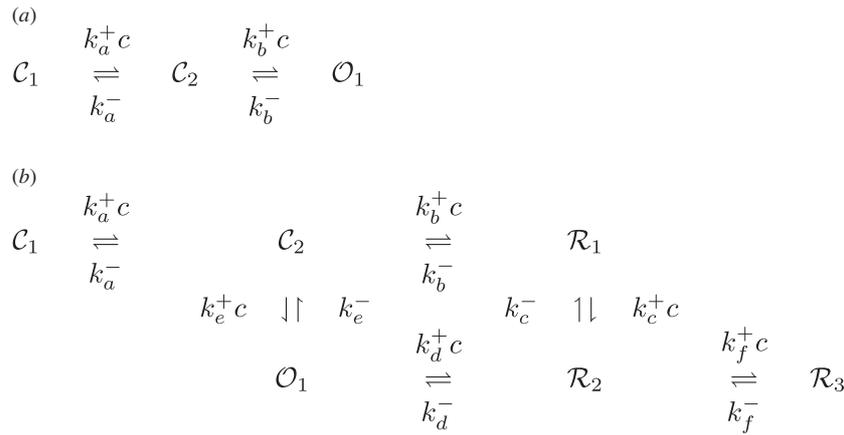


Figure 2. (a) Three-state single-channel model with Ca^{2+} -mediated activation that has two closed ($\mathcal{C}_1, \mathcal{C}_2$) and one open (\mathcal{O}_1) state. Parameters in $\mu\text{M}^{-1} \text{ms}^{-1}$: $k_a^+ = 1.5$, $k_b^+ = 150$; in ms^{-1} : $k_a^- = 50$, $k_b^- = 1.5$. (b) Six-state single-channel model with Ca^{2+} -mediated activation and inactivation. Parameters in $\mu\text{M}^{-1} \text{ms}^{-1}$: $k_a^+ = 1.5$, $k_b^+ = k_d^+ = 0.015$, $k_c^+ = k_e^+ = 300$, $k_f^+ = 3.0$; in ms^{-1} : $k_a^- = 49.5$, $k_b^- = k_d^- = 0.2475$, $k_c^- = k_e^- = 6.0$, $k_f^- = 0.03$.

Monte Carlo simulation, these quantities can also be directly calculated without simulation. In this case the stationary distribution of the generator matrix for a Ca^{2+} release site model is obtained using numerical linear algebra techniques, and response measures such as the puff/spark *Score* and the distribution of $N_{\mathcal{O}}$ are subsequently calculated from the stationary distribution, i.e., the steady-state probability of each release site state (see section 3). Because prior work indicates that the direct approach is computationally more efficient than Monte Carlo simulation [13], we aimed to apply advanced solution methods for Markov chains to models of coupled Ca^{2+} -regulated Ca^{2+} channels. Using a Kronecker-structured representation of the generator matrix for a Ca^{2+} release site model and memory-efficient algorithms applicable to large-scale Markov chains, this paper implements and benchmarks exact and approximate stationary distribution calculations for release sites with explicit channel positions and up to 1.6 million distinct states.

The remainder of this paper is organized as follows. In section 2 we briefly review our model formulation for coupling Ca^{2+} -regulated Ca^{2+} channels via the buffered diffusion of intracellular Ca^{2+} , and we present a Kronecker-structured representation for such Ca^{2+} release site models. In section 3 we review a variety of exact and approximate solution methods for the stationary analysis of CTMCs. In sections 4.1–4.3 we perform benchmark calculations using various exact solvers and analyze the performance of four numerical methods as a function of problem size, e.g., the time required to accurately calculate response measures such as the puff/spark *Score*. In sections 4.4 and 4.5 we present results using an approximate solution technique and examine its reliability by comparison to exact solutions. In section 4.6 we analyze and discuss the extent to which the currently available approximate methods are applicable given the state-space explosion of physiologically realistic Ca^{2+} release site models. In section 5 we make a specific proposal for future research in iterative numerical solution methods for Markov chain models of coupled Ca^{2+} -regulated Ca^{2+} channels.

2. Modeling the coupled gating of Ca^{2+} -regulated Ca^{2+} channels

The stochastic dynamics of single-channel gating has been successfully modeled using continuous-time discrete-state Markov chains (CTMCs) [1, 2]. In this paper we consider two Ca^{2+} -regulated Ca^{2+} channel models: a three-state channel that is activated by Ca^{2+} (figure 2(a)) and a six-state model that includes both fast Ca^{2+} activation and slow Ca^{2+} inactivation (figure 2(b)). In the state-transition diagrams shown in figure 2, $k_i^+ c$ and k_i^- with $i \in \{a, \dots, f\}$ are transition rates with units of reciprocal time, k_i^+ is an association rate constant with units of $\text{conc}^{-1} \text{time}^{-1}$, and c is the local $[\text{Ca}^{2+}]$ experienced by the Ca^{2+} -regulatory site of the channel. If this local $[\text{Ca}^{2+}]$ is specified, the transition-state diagrams shown in figures 2(a) and (b) define CTMCs that take on values in their respective state-spaces: $\mathcal{S} = \{\mathcal{C}_1, \mathcal{C}_2, \mathcal{O}_1\}$ and $\mathcal{S} = \{\mathcal{C}_1, \mathcal{C}_2, \mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3, \mathcal{O}_1\}$, respectively.

In the six-state model, Ca^{2+} -mediated transitions out of the open state can be accelerated due to the increase in local $[\text{Ca}^{2+}]$ when a Ca^{2+} -regulated Ca^{2+} channel is open [33, 34]. Assuming the formation and collapse of Ca^{2+} microdomains is fast compared to channel gating (see appendix B), we can denote the background and domain $[\text{Ca}^{2+}]$ experienced by the channel when closed and open as c_{∞} and c_d , respectively. With this assumption the generator matrices for both the three- and six-state models take the form

$$Q = K_- + (c_{\infty} I + c_d I_{\mathcal{O}}) K_+, \quad (2)$$

where K_- and K_+ are $M \times M$ matrices that collect the unimolecular (k_i^-) and bimolecular (k_i^+) transition rates, I is the $M \times M$ identity matrix, $I_{\mathcal{O}} = \text{diag}\{e_{\mathcal{O}}\}$ and $e_{\mathcal{O}}$ is an $M \times 1$ vector indicating open states of the single-channel model [13]. For example, for the three-state model of figure 2(a) we have

$$K_- = \begin{pmatrix} 0 & 0 & 0 \\ k_a^- & -k_a^- & 0 \\ 0 & k_b^- & -k_b^- \end{pmatrix}, \quad K_+ = \begin{pmatrix} -k_a^+ & k_a^+ & 0 \\ 0 & -k_b^+ & k_b^+ \\ 0 & 0 & 0 \end{pmatrix},$$

$e_O = (0, 0, 1)$, and because the product $I_O K_+$ is a zero matrix, the generator matrix Q is simply

$$Q = \begin{pmatrix} -k_a^+ c_\infty & k_a^+ c_\infty & 0 \\ k_a^- & -k_a^- - k_b^+ c_\infty & k_b^+ c_\infty \\ 0 & k_b^- & -k_b^- \end{pmatrix}. \quad (3)$$

While this matrix can be read off from figure 2(a) with the replacement of c_∞ for c , in the case of the six-state model the Ca^{2+} -mediated transition out of state \mathcal{O}_1 leads to nonzero $I_O K_+$ and an $\mathcal{O}_1 \rightarrow \mathcal{R}_2$ transition rate of $k_a^+(c_\infty + c_d)$. All other Ca^{2+} -dependent transition rates are given by setting $c = c_\infty$ in figure 2(b).

2.1. SAN descriptor for two Ca^{2+} -regulated Ca^{2+} channels

In our model formulation, the interaction between channels located at the same release site is mediated through the buffered diffusion of intracellular Ca^{2+} (see [13] for a complete description). Briefly, the N channels at the Ca^{2+} release site have positions chosen from a two-dimensional uniform distribution on a disc of radius 0.1–2.0 μm (see figure 1, left). When in the open state, each channel contributes to the landscape of $[\text{Ca}^{2+}]$ throughout the Ca^{2+} release site—the so-called Ca^{2+} microdomain—and influences the local $[\text{Ca}^{2+}]$ experienced by other channels. For simplicity we assume that the formation and collapse of individual peaks within the Ca^{2+} microdomain occurs quickly compared to channel gating. We also assume the presence of a single high-concentration Ca^{2+} buffer and the validity of superposing local $[\text{Ca}^{2+}]$ increases due to each of the N channels [35, 36]. Thus, channel interactions can be summarized by an $N \times N$ ‘coupling matrix’ $C = (c_{ij})$ that gives the increase over c_∞ experienced by channel j when channel i is open. The diagonal elements of C represent the quantity denoted above as ‘domain $[\text{Ca}^{2+}]$ ’ (c_d), the increase in $[\text{Ca}^{2+}]$ above background that an open channel contributes to its own Ca^{2+} regulatory site [13]. For example, in the case of two identical channels the Ca^{2+} coupling matrix takes the form

$$C = \begin{pmatrix} c_d & c_{12} \\ c_{21} & c_d \end{pmatrix},$$

where $0 < c_{12} = c_{21} < c_d$. See appendix A for further discussion of how numerical values for these concentrations are specified.

Note that the expanded generator matrix for two coupled Ca^{2+} -regulated Ca^{2+} channels has the Kronecker representation $Q^{(2)} = Q_-^{(2)} + Q_+^{(2)}$, where

$$Q_-^{(2)} = K_- \oplus K_- = K_- \otimes I + I \otimes K_- \quad (4)$$

collects the unimolecular transition rates and \oplus and \otimes denote the Kronecker sum and product, respectively (see chapter 9 in [37]). The transition rates involving Ca^{2+} take the form

$$Q_+^{(2)} = D_1^{(2)} (K_+ \otimes I) + D_2^{(2)} (I \otimes K_+), \quad (5)$$

where each term in the sum represents Ca^{2+} -mediated transitions for each channel. The diagonal matrices $D_1^{(2)}$ and $D_2^{(2)}$ give the $[\text{Ca}^{2+}]$ experienced by channels 1 and 2,

respectively, in every configuration of the release site. For example, for channel 1,

$$\begin{aligned} D_1^{(2)} &= \text{diag}\{c_\infty(e \otimes e) + c_d(e_O \otimes e) + c_{21}(e \otimes e_O)\} \\ &= c_\infty(I \otimes I) + c_d(I_O \otimes I) + c_{21}(I \otimes I_O), \end{aligned}$$

where e is an $M \times 1$ vector of ones. For example, for two coupled three-state channels (figure 2(a)), the transition rate associated with a $\mathcal{C}_2 \mathcal{O}_1 \rightarrow \mathcal{O}_1 \mathcal{O}_1$ transition is the (6, 9) entry of $Q^{(2)}$ given by $k_b^+(c_\infty + c_{21})$ when the release site states are ordered lexicographically, $\mathcal{S}^{(2)} = \{\mathcal{C}_1 \mathcal{C}_1, \mathcal{C}_1 \mathcal{C}_2, \dots, \mathcal{O}_1 \mathcal{C}_2, \mathcal{O}_1 \mathcal{O}_1\}$. This Ca^{2+} -dependent transition rate of channel 1 undergoing a $\mathcal{C}_2 \rightarrow \mathcal{O}_1$ transition (recall figure 2(a)) involves the background $[\text{Ca}^{2+}]$ (c_∞) as well as the concentration above background experienced by channel 1 when channel 2 is open (c_{21}). For two coupled six-state channels (figure 2(b)), the $\mathcal{O}_1 \mathcal{O}_1 \rightarrow \mathcal{O}_1 \mathcal{R}_2$ transition rate is the (15,17) entry of $Q^{(2)}$ and is given by $k_d^+(c_\infty + c_{12} + c_d)$.

Using the Kronecker identities such as $(I \otimes I_O)(I \otimes K_+) = I \otimes I_O K_+$, equation (4) and (5) can be combined and rearranged as

$$\begin{aligned} Q^{(2)} &= X_\infty \oplus X_\infty + c_d(I_O K_+ \otimes I) + c_{12}(I_O \otimes K_+) \\ &\quad + c_{21}(K_+ \otimes I_O) + c_d(I \otimes I_O K_+), \end{aligned} \quad (6)$$

where $X_\infty = K_- + c_\infty K_+$. The final four terms of this expression are arranged so that the column in which a term appears corresponds to the channel that is changing state (and thus the placement of K_+ in the left or right side of the Kronecker product corresponding to channel 1 or 2, respectively). Similarly, within a given column, there are two terms in which the matrix I_O takes every possible position (determined by the row in which the term appears). Each of these terms in a given column is contributing an increase in the local $[\text{Ca}^{2+}]$ consistent with the context (open and closed channels) in which the given channel is changing state. Such compact Kronecker-structured representation is an example of a stochastic automata network descriptor for two coupled Ca^{2+} -regulated Ca^{2+} channels [38].

2.2. SAN descriptor for N Ca^{2+} -regulated Ca^{2+} channels

In the case of N channels coupled at the Ca^{2+} release site, the expanded generator matrix—i.e., the SAN descriptor—is given by

$$Q^{(N)} = \bigoplus_{n=1}^N X_\infty + \sum_{i,j=1}^N \bigotimes_{n=1}^N X_{ij}^n, \quad (7)$$

where $X_\infty = K_- + c_\infty K_+$ as in equation (6) and

$$X_{ij}^n = \begin{cases} I_O & \text{for } i \neq j, i = n \\ c_{ij} K_+ & \text{for } i \neq j, j = n \\ c_d I_O K_+ & \text{for } i = j = n \\ I & \text{otherwise.} \end{cases} \quad (8)$$

Note that all states of the expanded Markov chain $Q^{(N)}$ are reachable, the matrices I, I_O and X_{ij}^n are all $M \times M$, and $2N^2 - N$ of the N^3 matrices denoted by X_{ij}^n are *not* identity matrices. The iterative solution methods discussed in the following section utilize the SAN descriptor for N coupled Ca^{2+} -regulated Ca^{2+} channels given by equations (7) and (8).

Table 1. State-space explosion for a Ca^{2+} release site with NM -state channels, M^N configurations, $L^{(N)} = NM^{N-1}L^{(1)}$ transitions between configurations and $Z^{(N)} = NM^{N-1}L^{(1)} + M^N$ nonzero entries in $Q^{(N)}$, where $L^{(1)}$ is the number of transitions in the single-channel model.

N	Three-state channel			Six-state channel		
	M^N	$L^{(N)}$	$Z^{(N)}$	M^N	$L^{(N)}$	$Z^{(N)}$
1	3^1	4	7	6^1	12	18
3	3^3	108	135	6^3	1296	1512
7	3^7	20 412	22 599	6^7	3919 104	4199 040
19	3^{19}	2.94×10^{10}	3.06×10^{10}	6^{19}	2.32×10^{16}	2.38×10^{16}

3. Solution methods for large-scale Markov chains

The limiting probability distribution of a finite irreducible CTMC is the unique stationary distribution $\pi^{(N)}$ satisfying global balance [37], that is,

$$\pi^{(N)} Q^{(N)} = \mathbf{0} \quad \text{subject to} \quad \pi^{(N)} \mathbf{e}^{(N)} = 1, \quad (9)$$

where $Q^{(N)}$ is the Ca^{2+} release site SAN descriptor for N coupled channels (equations (7) and (8)) and $\mathbf{e}^{(N)}$ is an $M^N \times 1$ column vector of ones. Although Monte Carlo simulation techniques, such as Gillespie's method [39], can be implemented to estimate response measures, such as the distribution of the number of open channels ($N_{\mathcal{O}}$) and the puff/spark *Score* (recall figure 1), this is an inefficient approach when the convergence of the occupation measures to the limiting probability distribution is slow. This problem is compounded by the state-space explosion that occurs when the number of channels (N) or number of states per channel (M) is large (i.e., physiologically realistic). Table 1 illustrates the state-space explosion for release sites composed of the three- and six-state single-channel models of figure 2. Because the model formulation (section 2) accounts for release site ultrastructure (i.e., the spatial location of each channel), the number of configurations of the Ca^{2+} release site grows by a factor of M each time a channel is added, and when N is large the storage requirements of explicitly forming the expanded generator matrix $Q^{(N)}$ are excessive. Furthermore, the occupation measures for the limiting probability distribution are slow to converge and interpreting simulation results involving M^N release site configurations is difficult for large N . Fortunately, both space requirements and quality of results can be addressed using the Kronecker-structured representation of equations (7) and (8)—i.e., the Ca^{2+} release site SAN descriptor—in combination with various iterative numerical methods that leverage its Kronecker structure to solve for $\pi^{(N)}$.

3.1. Exact numerical methods

Many methods are available to solve equation (9) with different ranges of applicability (see [37] for review). For larger models, a variety of iterative methods are applicable including the power method (POWER) and the methods of Jacobi and Gauss–Seidel, along with variants that use relaxation, e.g., Jacobi with relaxation (JOR) and Gauss–Seidel with relaxation (SOR). Such methods require space for iteration vectors and $Q^{(N)}$

but usually converge quickly. More sophisticated projection methods, such as the generalized minimum residual method (GMRES) and the method of Arnoldi (ARNOLDI), have better convergence properties but require more space. While the best method for a particular Markov chain is unclear in general, several options are available for exploration including the iterative methods described above, which can also be enhanced by preconditioning, aggregation–disaggregation (AD) or Kronecker-specific multi-level (ML) methods [40, 41].

Due to multi-level (ML) method's superior performance in this context (see section 4) and their ability to leverage the block structure that is naturally present in the Kronecker representation [42], we will describe them in more detail here. ML methods were inspired by multigrid methods used to solve partial differential equations, as well as aggregation/disaggregation techniques for the reduction of Markov chains [43]. ML methods are iterative algorithms defined on multiple levels of increasing coarseness through which the solution process proceeds in cycles until a given termination criterion is met. The levels are given by the nested block structure of the generator matrix, with the blocks defined implicitly by the Kronecker structure or explicitly through partitioning (see figure 3). One moves from fine-to-coarse and coarse-to-fine representations of the Markov chain via aggregation and disaggregation, respectively, where at each level of refinement an iterative method is implemented (i.e., the 'smoother'). One ML cycle consists of the recursive traversal of these levels from the finest to the coarsest and back. There are multiple implementations of ML methods, with various places for variability: the type of cycle (V, W or F), the type of smoother (e.g. JOR or SOR), the number of iterations of the smoother at each level and the method of selection for which automata (i.e., channels) to aggregate at each stage, e.g., fixed (FIX), cyclic (CYC) or dynamic (DYN). For comparisons of these implementations on a specific set of examples, see below, [40] and the supplementary material stacks.iop.org/PhysBio/5/036003.

Partitioning induces a block structure of the generator matrix that can be useful in the implementation of both exact and approximate methods. For exact methods, such as the ML methods described above, this block structure often suggests aggregation/disaggregation strategies that can improve convergence times. Partitioning can also be chosen consistent with response measures of interest, thereby facilitating the use of these measures as convergence criteria. For example, if we partition the three-state single-channel model of figure 2(a) into closed and open states, the resulting block structure facilitates efficient computation of the performance measures dependent on the number of open channels, such as the distribution of $N_{\mathcal{O}}$ and the puff/spark *Score*. As illustrated in figure 3(a), we can partition the states of the three-state single-channel model $\mathcal{S} = \{\mathcal{C}_1, \mathcal{C}_2, \mathcal{O}_1\}$ using the closed and open aggregate classes and write $\mathcal{P} = \{\mathcal{C}, \mathcal{O}\}$ where $\mathcal{C} = \{\mathcal{C}_1, \mathcal{C}_2\}$ and $\mathcal{O} = \{\mathcal{O}_1\}$. In the case of N three-state channels, the induced partitioning on the expanded state-space is a mapping of the 3^N states $s = s^1 s^2 \dots s^N \in \mathcal{S}^{(N)}$ to one of the 2^N partitions $p = p^1 p^2 \dots p^N \in \mathcal{P}^{(N)}$, where $\mathcal{S}^{(N)}$ and $\mathcal{P}^{(N)}$ are the Cartesian products $\times_{n=1}^N \mathcal{S}$ and $\times_{n=1}^N \mathcal{P}$,

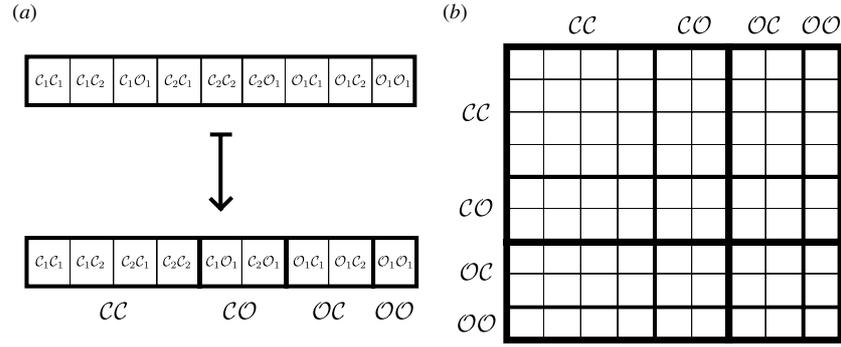


Figure 3. (a) Permutation of states and partition structure for $N = 2$ three-state channels under the closed/open partitioning strategy $\mathcal{P} = \{\mathcal{C}, \mathcal{O}\}$, where $\mathcal{C} = \{c_1, c_2\}$ and $\mathcal{O} = \{o_1\}$. The induced partitioning on $\mathcal{S}^{(2)}$ is thus $\mathcal{P}^{(2)} = \{\mathcal{CC}, \mathcal{CO}, \mathcal{OC}, \mathcal{OO}\}$, with the states ordered lexicographically in both $\mathcal{S}^{(2)}$ and each partition. (b) Block structure of the expanded generator matrix $Q^{(2)}$ when permuted in this manner. The thickness of the lines denotes the hierarchical structure of the partitioning.

respectively. In figure 3(b), lexicographical ordering of partitions—i.e., $\mathcal{P}^{(2)} = \{\mathcal{CC}, \mathcal{CO}, \mathcal{OC}, \mathcal{OO}\}$ —and states within each partition leads to a permuted generator matrix composed of 16 blocks, with each block denoted by $Q^{(2)}[p, q]$. Here and in the general case ($N > 2$), the block $Q^{(N)}[p, q]$ contains transitions from partition p to partition q with each block having a Kronecker-structured representation similar to equation (7).

More specifically, the diagonal blocks of the permuted generator matrix $Q^{(N)}[p, p]$ that correspond to transitions within each partition are given by

$$Q^{(N)}[p, p] = \bigoplus_{n=1}^N X_{\infty}[p^n, p^n] + \sum_{i,j=1}^N \bigotimes_{n=1}^N X_{ij}^n[p^n, p^n]. \quad (10)$$

In this expression, each matrix $X_{\infty}[p^n, p^n]$ contains the rows and columns of X_{∞} corresponding to the states in p^n (and similarly for each $X_{ij}^n[p^n, p^n]$), where X_{∞} and X_{ij}^n are defined in equation (8). For each off-diagonal block of the permuted generator matrix $Q^{(N)}[p, q]$ ($p \neq q$) that corresponds to a transition between partitions, we write $p^n = q^n$ for $n \neq k$ and $p^k \neq q^k$ where $k \in \{1, 2, \dots, N\}$ is the index of the channel changing state. Using this notation, the nonzero off-diagonal blocks of the permuted generator matrix can be written as

$$Q^{(N)}[p, q] = \left(\bigotimes_{n=1}^{k-1} I_{|p^n|} \right) \otimes X_{\infty}[p^k, q^k] \otimes \left(\bigotimes_{n=k+1}^N I_{|p^n|} \right) + \sum_{i,j=1}^N \bigotimes_{n=1}^N X_{ij}^n[p^n, q^n], \quad (11)$$

where $X_{\infty}[p^k, q^k]$ ($X_{ij}^n[p^n, q^n]$) contains the rows and columns of X_{∞} (X_{ij}^n) corresponding to the states in p^k and q^k (p^n and q^n), $|p^n|$ denotes the size of partition p^n and $I_{|p^n|}$ denotes a $|p^n| \times |p^n|$ identity matrix. Equations (10) and (11) are an example of a *hierarchical Kronecker representation* (see [44] for review).

3.2. Approximate numerical methods

The hierarchical Kronecker representation of equation (10) can be used to realize approximate solution techniques

that often drastically reduce the computational effort while introducing only small approximation errors. A promising concept introduced in [45] is to represent components of the iteration vector by Kronecker products of vectors of much smaller dimension. Consistent with the hierarchical Kronecker representation discussed above, the iteration vector $\pi^{(N)}$ is partitioned into $|\mathcal{P}^{(N)}|$ components (see figure 3(a)) and each component is either represented in an exact manner (detailed representation) or in an approximate manner (compositional representation). The compositional representation of the portion of π (dropping the superscript (N) for clarity) corresponding to partition p , denoted by $\pi[p]$, is given by

$$\pi[p] = \alpha_p \bigotimes_{n=1}^N \pi_p^n \quad (12)$$

where π_p^n is a vector of dimension $1 \times |p^n|$ with elements summing to unity and α_p is a non-negative constant that scales the probability mass of the $\pi[p]$ in π such that $\sum_{p \in \mathcal{P}^{(N)}} \alpha_p = 1$.

As mentioned above, a compositional representation such as equation (12) is flexible in the sense that the user may choose the partitions to be represented exactly versus approximately. Alternatively, the choice of exact versus approximate representation may be adjusted adaptively during the iterative solution process, e.g., using a detailed representation for those partitions that accumulate most of the probability mass [45]. Perhaps most importantly, the partitioning strategy \mathcal{P} can significantly influence computational efficiency and the quality of the approximation. For example, a partitioning strategy that includes every state of the single-channel model as a partition—e.g., $\mathcal{P} = \{c_1\}\{c_2\}\{o_1\}$ for the three-state model of figure 2(a)—leads to an induced partitioning $\mathcal{P}^{(N)}$ that maps every state of the Ca^{2+} release site model to a distinct partition (i.e. $\pi[p] = \alpha_p$ with $\pi_p^n = 1$ in equation (12)). This strategy results in no approximation error, but unnecessarily increases both storage and run time compared to an exact iterative method. In the other extreme, a partitioning strategy that includes every state of the single-channel model in one partition—e.g., $\mathcal{P} = \{c_1 c_2 o_1\}$ —can only result in low approximation error when the stationary distribution of the Ca^{2+} release site model is well approximated

Table 2. Benchmark calculations for 10 three-state channels computed using Linux PCs with dual-core 3.8 GHz EM64T Xeon processors and 8 GB RAM solving equation (7). Description of solvers: JOR, Jacobi over-relaxation method; JOR_AD, the method of Jacobi with aggregation/disaggregation; ARNOLDI, the method of Arnoldi; BICGSTAB, the biconjugate gradient stabilized method; PRE_ARNOLDI, the method of Arnoldi with Neumann pre-conditioning; BSOR_BICGSTAB, the biconjugate gradient stability method with block-successive over-relaxation pre-conditioning; ML_JOR_F_DYN, multi-level method with JOR smoother, F cycle and dynamic ordering.

Solver	Max Res	Sum Res	CPU (s)	Wall (s)	Iters
JOR	9.49×10^{-13}	5.16×10^{-12}	279	279	1840
JOR_AD	9.44×10^{-13}	5.13×10^{-12}	415	415	1550
ARNOLDI	2.42×10^{-13}	4.04×10^{-11}	214	215	1440
BICGSTAB	8.66×10^{-13}	4.89×10^{-11}	146	148	602
PRE_ARNOLDI	8.62×10^{-15}	1.82×10^{-12}	26	27	160
BSOR_BICGSTAB	8.22×10^{-15}	5.29×10^{-13}	19	19	52
ML_JOR_F_DYN	5.87×10^{-13}	1.68×10^{-10}	15	15	46

by the Kronecker factorization $\pi = \otimes_{n=1}^N \pi^n$. This is unlikely because it implies that the N channels at the Ca^{2+} release site are gating independently, i.e., not interacting via the buffered diffusion of Ca^{2+} . Below we identify and discuss optimal partitioning strategies for both the three- and six-state models of figure 2. In general, we find that more refined partitioning leads to better approximation, but at the expense of storage requirements and computational efficiency (see section 4.4).

3.3. Abstract Petri net notation (APNN) toolbox

A number of software tools are available that implement methods for Kronecker representations. We selected the Abstract Petri net notation toolbox [46] and its numerical solution package Nsolve because of its rich variety of numerical techniques for the steady-state analysis of Markov chains. Nsolve provides more than 70 different iterative numerical methods and allows the user to define SAN descriptors of the form of equations (7) and (8) through ASCII file input [46–49].

4. Results

4.1. Benchmarked exact methods

In order to investigate the numerical techniques that work best in combination with the Kronecker representation of our Ca^{2+} release site models (equations (7) and (8)), we wrote a script for the software environment MATLAB that takes a specific Ca^{2+} release site model—defined by K_+ , K_- , $e_{\mathcal{O}}$, c_{∞} and C as defined in section 2—and produces the input files needed to interface with Nsolve. Using release sites composed of 10 three-state channels (figure 2(a)), we performed a preliminary study to determine which of the 70-plus numerical methods implemented in Nsolve were compatible with equations (7) and (8).

Table 2 lists seven solvers that converged in less than 20 min CPU time with a maximum residual $\|\pi^{(N)} Q^{(N)}\|_{\infty}$ less than 10^{-12} for a release site composed of 10 three-state channels (see supplementary material stacks.iop.org/PhysBio/5/036003 for additional results). For each method we report the maximum residual when convergence is achieved, the sum of the residuals $\|\pi^{(N)} Q^{(N)}\|_1$, the CPU and wall clock times (in seconds) and

the total number of iterations performed. Because randomly selected channel positions (cf figure 1) have an impact on the interaction matrix C , the generator matrix $Q^{(N)}$ and the performance of solvers, the release site ultrastructure was identical for each calculation and thus the rows of table 2 can be directly compared. We find that the traditional Jacobi over-relaxation method (JOR) works well for this problem with $3^{10} = 59\,049$ states, but the addition of aggregation/disaggregation (AD) steps is not particularly helpful. The separable preconditioner (PRE) of Buchholz [48] and the block SOR preconditioner (BSOR) are very effective and help to reduce solution times to less than 50 s for several projection methods including ARNOLDI and the biconjugate gradient stability method (BICGSTAB). A multi-level (ML) solver with a JOR smoother, DYN ordering and F cycle gives the best results [40, 41].

4.2. Scalability of exact methods

In the previous section we benchmarked the efficiency of several different algorithms that can be used to solve for the stationary distribution of Ca^{2+} release site models. To determine how these results depend on problem size, we chose representatives of four classes of solvers (JOR, PRE_ARNOLDI, BSOR_BICGSTAB and ML_JOR_F_DYN) that worked well for release sites composed of 10 three-state channels (see table 2). Using these four solvers, figure 4 shows the wall clock time required for convergence ($\|\pi^{(N)} Q^{(N)}\|_{\infty} < 10^{-12}$) as a function of the number of channels (N) for both the three- and six-state models (circles and squares, respectively). Because the N channels in each Ca^{2+} release site simulation have randomly chosen spatial positions that may influence the time to convergence, figure 4 shows both the mean and standard deviation (error bars) of the wall clock time for five different release site configurations. Note that for each value of N in figure 4, the radius of each Ca^{2+} release site was chosen so that stochastic Ca^{2+} excitability was observed.

Figure 4 shows that the time until convergence is shorter when the Ca^{2+} release site is composed of three-state as opposed to six-state channels regardless of the numerical method used (compare circles to squares). Consistent with table 2 we find that for large values of N the ML_JOR_F_DYN (black) method requires the least time, followed by BSOR_BICGSTAB (dark gray), PRE_ARNOLDI (light gray) and

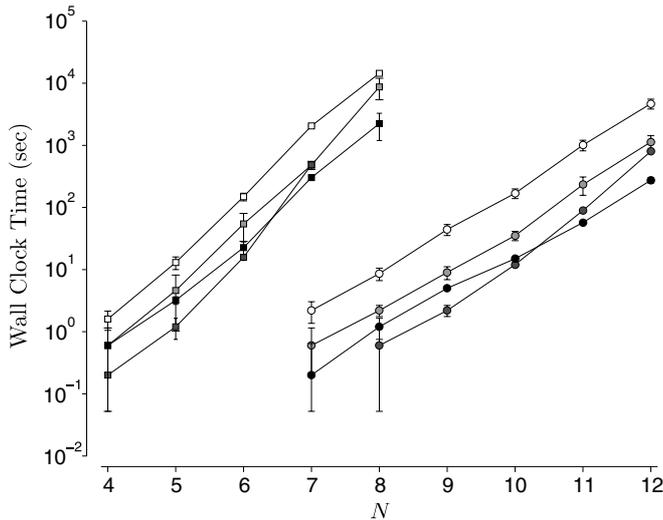


Figure 4. Circles and error bars show the mean \pm SD of wall clock time for five release site configurations of the three-state model (figure 2(a)) using: JOR (white), PRE_ARNOLDI (light gray), BSOR_BICGSTAB (dark gray) and ML_JOR_F_DYN (black). Squares and error bars give results for the six-state model (figure 2(b)). Single-channel parameters as in figure 2. Calculations performed using 2.66 GHz dual-core Intel Xeon processors and 2 GB RAM. Reproduced with permission from [50].

finally JOR (white). Though there are important differences in the speed of the four solvers, the wall clock time until convergence is proportional to the number of states M in a single-channel model, that is, the slope of each line in figure 4 for $M = 6$ is nearly double that of the corresponding lines when $M = 3$.

We also found substantial differences in the amount of memory needed to run those solvers (not shown). While simple methods like JOR allocate space mainly for a few iteration vectors, Krylov subspace methods like ARNOLDI use more vectors (20 in the default Nsolve configuration) and this can be prohibitive for large models. For projection methods such as BICGSTAB that operate on a fixed and small set of vectors, we observe that the space for auxiliary data structures and vectors is on the order of seven to ten iteration vectors for these models. In general, we find that the iterative numerical methods that incorporate pre-conditioning (e.g., PRE_ARNOLDI and BSOR_BICGSTAB) are quite fast compared to more traditional relaxation techniques such as JOR. However, the power of pre-conditioning is only evident when problem size is less than some threshold that depends upon memory limitations. On the other hand, multi-level (ML) methods are constructed to take advantage of the Kronecker representation and to have very modest memory requirements. This is consistent with our experiments that indicate ML methods have the greatest potential to scale well with problem size (black symbols in figure 4), whether that be an increase in the number of channels (N) or the number of states per channel (M).

4.3. Exact methods versus Monte Carlo simulation

Although there may be problem size limitations, we expected that the stationary distribution of our Ca^{2+} release site models

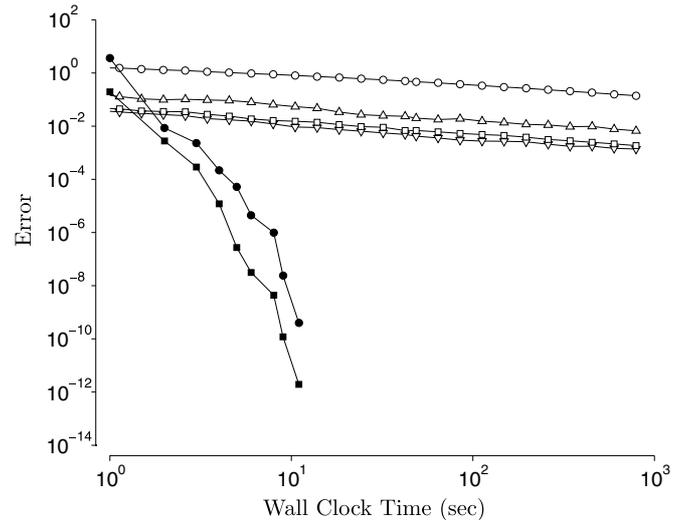


Figure 5. Convergence of response measures for a release site composed of 10 three-state channels using ML_JOR_F_DYN and Monte Carlo (filled and open symbols, respectively). Circles and squares give 1- and ∞ -norms of the residual errors, upper pointing triangles give the relative error in the puff/spark *Score* for Monte Carlo (mean of 50 simulations shown) compared with the *Score* given by ML_JOR_F_DYN upon convergence. Similarly, the lower pointing triangles give the relative error in the probability that all N channels are closed. Parameters as in figure 1. Reproduced with permission from [50].

could be found more quickly using iterative methods than Monte Carlo simulation. This is confirmed in the convergence results of figure 5 using a release site composed of 10 three-state channels, the multi-level solver ML_JOR_F_DYN (filled symbols) and Monte Carlo simulation beginning with all N channels in state C_1 that averaged 1260 transitions per second (open symbols).

The open squares and circles of figure 5 show the maximum and sum of the residuals ($\|\pi^{(N)} Q^{(N)}\|_\infty$ and $\|\pi^{(N)} Q^{(N)}\|_1$, respectively) averaged over 50 simulations. As expected, the residuals associated with the Monte Carlo simulations converge much slower than those obtained with ML_JOR_F_DYN. Interestingly, figure 5 shows that even coarse response measures can be more quickly obtained using numerical iterative methods than Monte Carlo simulation. In the Monte Carlo simulations, the relative errors of the puff/spark *Score* (upwards pointing triangles) and the probability that all N channels were closed (downwards pointing triangles) converge at essentially the same rate as the maximum residual error $\|\pi^{(N)} Q^{(N)}\|_\infty$ (open squares).

4.4. Benchmarked approximate methods

In sections 4.1–4.3 we identified several exact solvers that perform well when using the SAN descriptor (equations (7) and (8)) to solve for the stationary distribution of a Ca^{2+} release site model. However, we have found that these techniques are not applicable when the number of states in the release site model (M^N) becomes large and so it is necessary to consider approximate methods and various partitioning strategies (recall section 3.2). The APNN toolbox contains

Table 3. Optimal partitioning strategies listed with their relative errors and memory requirements for both $N = 3$ three- and six-state channels with states $\mathcal{S} = \{C_1, C_2, O_1\}$ and $\mathcal{S} = \{C_1, C_2, R_1, R_2, R_3, O_1\}$, respectively (see figure 2). The last column also lists relative memory requirements for larger N used in section 4.6. For the three-state model, the optimal \mathcal{P} is the best of 1, 3, 1 possibilities when $|\mathcal{P}| = 1, 2, 3$. For the six-state model, the optimal \mathcal{P} is the best of 1, 31, 90, 65, 15, 1 possibilities when $|\mathcal{P}| = 1, 2, 3, 4, 5, 6$.

$ \mathcal{P} $	Optimal \mathcal{P} for the three-state model	ϵ_π	ϵ_{score}	$\nu(3)$	$\nu(12)$
1	$\{C_1 C_2 O_1\}$	1.66	9.38×10^{-1}	3.33×10^{-1}	6.77×10^{-5}
2	$\{C_1 C_2\}\{O_1\}$	5.40×10^{-3}	3.50×10^{-3}	1.33	1.39×10^{-1}
3	$\{C_1\}\{C_2\}\{O_1\}$	6.00×10^{-8}	4.38×10^{-9}	3.00	1.20×10^1
$ \mathcal{P} $	Optimal \mathcal{P} for the six-state model	ϵ_π	ϵ_{score}	$\nu(3)$	$\nu(8)$
1	$\{C_1 C_2 R_1 R_2 R_3 O_1\}$	4.84×10^{-1}	5.40×10^{-1}	8.30×10^{-2}	2.86×10^{-5}
2	$\{C_1 C_2 R_1 R_2 R_3\}\{O_1\}$	3.50×10^{-2}	1.36×10^{-2}	3.33×10^{-1}	3.70×10^{-3}
3	$\{C_1 C_2\}\{R_1 R_2 R_3\}\{O_1\}$	7.10×10^{-3}	5.50×10^{-3}	7.50×10^{-1}	6.25×10^{-2}
4	$\{C_1\}\{C_2\}\{R_1 R_2 R_3\}\{O_1\}$	1.90×10^{-4}	7.60×10^{-6}	1.33	4.68×10^{-1}
5	$\{C_1\}\{C_2\}\{R_1 R_2\}\{R_3\}\{O_1\}$	5.07×10^{-7}	4.12×10^{-8}	2.08	2.23
6	$\{C_1\}\{C_2\}\{R_1\}\{R_2\}\{R_3\}\{O_1\}$	4.21×10^{-8}	4.16×10^{-8}	3.00	8.00

an implementation of only one approximate method, namely APP_POWER, with its iterative engine based on the power method (POWER). In this section we perform a preliminary study using APP_POWER on $N = 3$ channels to evaluate the 5 possible partitioning strategies for the three-state single-channel model of figure 2(a) and the 203 possible partitioning strategies for the six-state model of figure 2(b).

For the three-state model the minimum number of partitions is $|\mathcal{P}| = 1$ and the maximum number of partitions is $|\mathcal{P}| = 3$, and in both cases there is only one possible partitioning strategy (see the first column of table 3). Three partitioning strategies are possible when $|\mathcal{P}| = 2$ including $\mathcal{P} = \{C_1 C_2\}\{O_1\}$, $\{C_1\}\{C_2 O_1\}$ and $\{C_1 O_1\}\{C_2\}$, the first of which we found to be optimal in the sense of having minimum relative error for both the full stationary distribution (ϵ_π) and the puff/spark Score (ϵ_{score}). Table 3 also shows the amount of memory required by the approximate method to store $\pi^{(N)}$ when each partitioning strategy is employed. For clarity this is presented as a relative quantity,

$$\nu(N) = \frac{Z_A}{Z_E} = N \left(\frac{|\mathcal{P}|}{M} \right)^{N-1},$$

where M is the number of states in the single-channel model, N is the number of channels at the release site and $Z_A = NM|\mathcal{P}|^{N-1}$ and $Z_E = M^N$ are the memory requirements for the approximate and exact methods, respectively. As suggested in section 3.2, the maximum number of partitions ($|\mathcal{P}| = 3$) would not be used in practice because it would yield results equivalent to an exact method but require more storage ($\nu(N) > 1$). The minimum number of partitions ($|\mathcal{P}| = 1$) would not be used because of excessive error (ϵ_π and ϵ_{score} large). However, for a large number of channels (e.g., $N = 12$) and the optimal $|\mathcal{P}| = 2$ partitioning strategy $\mathcal{P} = \{C_1 C_2\}\{O_1\}$, the relative memory requirement of the approximate method and approximation errors are acceptable ($\nu(12) = 0.14$, $\epsilon_\pi = 0.0054$ and $\epsilon_{\text{score}} = 0.0035$). Interestingly, the optimal $|\mathcal{P}| = 2$ partitioning strategy associates the two closed states of the single-channel model and isolates the open state. This makes intuitive sense given the central role of the closed and open aggregated classes of states in the coupling of channels in the Ca^{2+} release site model (recall section 2).

Table 3 shows the results of a similar study of partitioning strategies for $N = 3$ six-state single-channel models. In this case there are 203 possible partitioning strategies and the best of the 31 possible $|\mathcal{P}| = 2$ strategies is $\mathcal{P} = \{C_1 C_2 R_1 R_2 R_3\}\{O_1\}$. Again, we find the optimal $|\mathcal{P}| = 2$ strategy partitions closed and open states of the single-channel model and, interestingly, the optimal $|\mathcal{P}| = 3$ strategy separately partitions closed and refractory states, $\mathcal{P} = \{C_1 C_2\}\{R_1 R_2 R_3\}\{O_1\}$. While states C_1, C_2, R_1, R_2 and R_3 are similar in that channels in these states do not increase the $[\text{Ca}^{2+}]$ experienced by neighboring channels, including all the refractory states in one partition may work well because the sojourn time in states R_1, R_2 and R_3 is 3–300 times longer than the sojourn time in states C_1 and C_2 (see appendix C).

4.5. Error in approximate methods

In table 3 we identified the optimal partitioning strategies for $N = 3$ coupled three- and six-state channels. To determine if these results generalize for larger problems, we tested the accuracy of several different partitioning strategies for Ca^{2+} release sites with $N = 12$ three-state channels or $N = 8$ six-state channels. In both cases we consider the optimal $|\mathcal{P}| = 2$ partitioning strategy that separates closed and open states. This will be denoted below by \mathcal{C}/\mathcal{O} where it is understood that this refers to $\mathcal{P} = \{C_1 C_2\}\{O_1\}$ for the three-state model and $\mathcal{P} = \{C_1 C_2 R_1 R_2 R_3\}\{O_1\}$ for the six-state model. In the case of the six-state model we also evaluate the optimal $|\mathcal{P}| = 3$ partitioning strategy $\mathcal{P} = \{C_1 C_2\}\{R_1 R_2 R_3\}\{O_1\}$ (denoted by $\mathcal{C}/\mathcal{R}/\mathcal{O}$). We focus on these strategies because they perform well for $N = 3$ (see table 3). But as discussed above, these strategies are consistent with salient properties of these single-channel models (e.g., two conductance levels and, in the case of the six-state model, fast Ca^{2+} activation and slow Ca^{2+} inactivation).

Using $N = 12$ three-state channels, figure 6(a) shows the probability distribution of the number of open channels (rightmost panel) calculated exactly using the high-performing multi-level method ML_JOR_F_DYN (black bars) described in sections 3.1 and 4.1. Figure 6(a) also shows the probability distribution of $N_{\mathcal{O}}$ calculated approximately using the method APP_POWER with the \mathcal{C}/\mathcal{O} partitioning strategy (white bars)

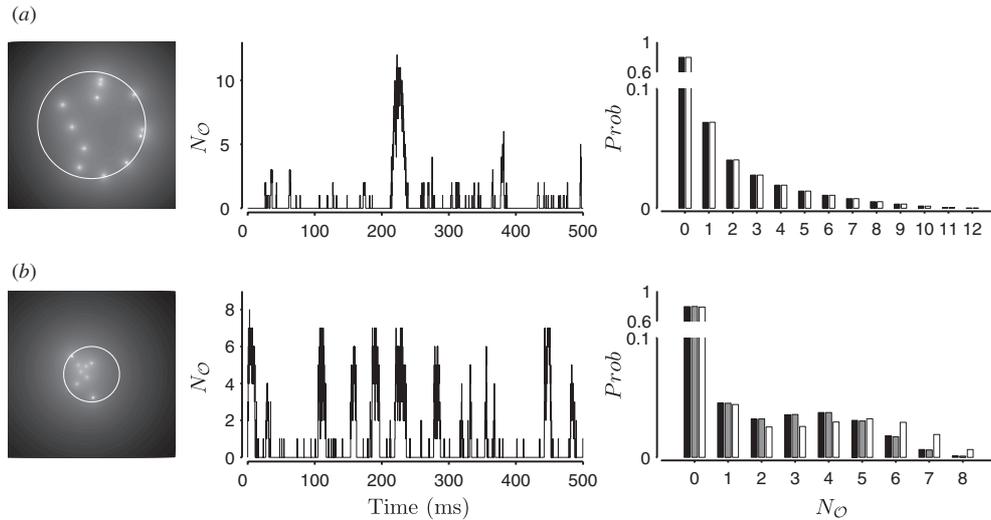


Figure 6. (a) Statistics for a release site composed of 12 three-state channels. Left: local $[\text{Ca}^{2+}]$ near $3 \times 3 \mu\text{m}^2$ ER membrane modeled as in figure 1. Middle: localized Ca^{2+} elevations reminiscent of Ca^{2+} puffs/sparks. Right: probability distribution of the number of open channels calculated exactly using ML_JOR_F_DYN (black bars) and approximately using APP_POWER with \mathcal{C}/\mathcal{O} partitioning (white bars). (b) Statistics as in (a) for 8 six-state channels with black bars denoting ML_JOR_F_DYN and white and grey bars denoting APP_POWER with \mathcal{C}/\mathcal{O} and $\mathcal{C}/\mathcal{R}/\mathcal{O}$ partitioning, respectively.

as described in sections 3.2 and 4.4. A similar study using $N = 8$ six-state channels is shown in figure 6(b) where black, white and grey bars show results obtained by the exact ML_JOR_F_DYN method, APP_POWER with the \mathcal{C}/\mathcal{O} partitioning strategy and APP_POWER with the $\mathcal{C}/\mathcal{R}/\mathcal{O}$ partitioning strategy, respectively. For the three-state model with no refractory states, the \mathcal{C}/\mathcal{O} partitioning strategy well approximates the exact results (compare black and white bars in figure 6(a)); however, this relatively coarse level of partitioning does not perform as well for the six-state model with three refractory states (compare black and white bars in figure 6(b)). For the six-state model we find that the finer partitioning strategy $\mathcal{C}/\mathcal{R}/\mathcal{O}$ results in a better approximation of the steady-state distribution of $N_{\mathcal{O}}$ (compare black and grey bars).

4.6. Scalability of approximate methods versus Monte Carlo simulation

In the previous section, we showed that for $N = 8$ six-state channels the accuracy of the APP_POWER method improves with more refined partitioning when the response measure of interest is the probability distribution of the number of open channels. While the $\mathcal{C}/\mathcal{R}/\mathcal{O}$ partitioning strategy performed well in its approximation of this particular response measure, its practical value depends on the time to convergence in comparison to exact calculation. Figure 7(a) shows the wall clock time required for the convergence of $\pi^{(N)}$ as a function of the number of channels (N) for the six-state single-channel model. Convergence is achieved when

$$\max_{p \in \mathcal{P}^{(N)}} \max_n \|\pi_{p,l}^n - \pi_{p,l-1}^n\|_{\infty} \leq 10^{-8},$$

where $\pi_{p,l}^n$ is the l th iterate of the n th factor of partition $\pi[p]$ in equation (12). As expected, figure 7(a) shows that results obtained using the approximate method APP_POWER and \mathcal{C}/\mathcal{O} partitioning converge faster than those assuming $\mathcal{C}/\mathcal{R}/\mathcal{O}$

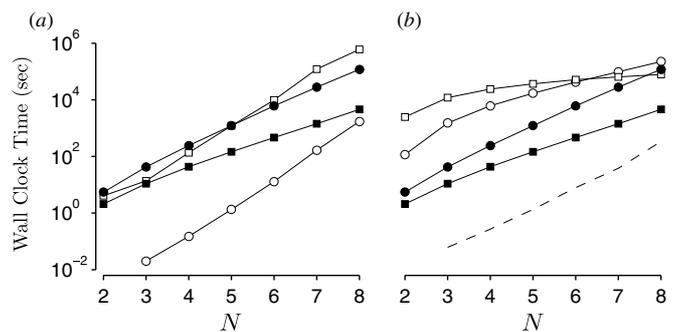


Figure 7. (a) Filled and open symbols show the wall clock time for the six-state model using approximate and exact methods, respectively. Approximate results are shown for two levels of partitioning (\mathcal{C}/\mathcal{O} , squares and $\mathcal{C}/\mathcal{R}/\mathcal{O}$, circles) with the APP_POWER method. Exact solutions are calculated using the POWER method (squares) and the ML_JOR_F_DYN method (circles). (b) Results as in (a) with open symbols corresponding to the Monte Carlo estimates of two coarse response measures: the distribution of the number of open channels (circles) and the distribution of probability across the M states of an arbitrarily selected individual channel (squares). The dashed line shows the projected performance of an approximate multi-level solver that uses ML_JOR_F_DYN rather than POWER as its iterative engine.

partitioning for all N considered (compare filled squares to filled circles); however, we expect the $|\mathcal{P}| = 3$ partitioning to give a better approximation than the $|\mathcal{P}| = 2$ partitioning (recall figure 6(b)). Figure 7(a) also shows the convergence time of the exact ML_JOR_F_DYN method (open circles) and the exact POWER method (open squares). Although POWER was excluded from table 2 due to slow convergence, it is appropriate to compare the convergence times of APP_POWER and POWER because the methods have similar iterative engines. While the extra overhead in using APP_POWER with $\mathcal{C}/\mathcal{R}/\mathcal{O}$ partitioning slows its convergence for small problem size relative to the

exact POWER method, the approximate method becomes faster than its exact counterpart when the release site has $N \geq 5$ channels. In general, we find that approximate methods have potential to scale well with problem size; the slopes of these curves appear to be related to the partitioning refinement (i.e., $|\mathcal{P}| = 2, 3$), whereas the exact method curves have slopes proportional to the size of the single-channel model ($M = 6$). Figure 7(b) compares the wall clock times for the approximate APP_POWER method (closed symbols as in figure 7(a)) to Monte Carlo simulation (open symbols) where measures of interest are the probability distribution of the number of open channels (circles) and the distribution of probability across the $M = 6$ states of an arbitrarily selected individual channel (squares). For equitable comparison, the Monte Carlo calculations were terminated when the confidence intervals for each bin in a particular distribution fell below the average residual bin error of APP_POWER (determined by comparison to the exact ML_JOR_F_DYN result). Figure 7(b) shows that when the release site is composed of $N < 8$ channels the APP_POWER method with $\mathcal{C}/\mathcal{R}/\mathcal{O}$ partitioning (filled circles) converges more quickly than Monte Carlo calculations. However, it appears that for the response measures considered here the time required for Monte Carlo estimates scales better for larger problem sizes (compare open square to filled circle at $N = 8$). When a finer measure is desired (e.g., the entire stationary distribution $\pi^{(N)}$), we would expect APP_POWER to outperform Monte Carlo simulation even when N is large (see the conclusions).

5. Conclusion and outlook

We have presented a Kronecker-structured representation for Ca^{2+} release sites composed of Ca^{2+} -regulated Ca^{2+} channels under the assumption that these channels interact instantaneously via the buffered diffusion of intracellular Ca^{2+} (section 2). Because informative response measures such as the puff/spark *Score* can be determined if the steady-state probability of each release site configuration is known, we have identified iterative numerical solution techniques that perform well in this biophysical context.

While the benchmark stationary distribution calculations presented here all utilize the Kronecker structure of the Ca^{2+} release site SAN descriptor, we find significant performance differences among iterative solution methods (table 2). When it is possible to obtain an exact solution, multi-level methods provide excellent convergence with modest additional memory requirements for the Kronecker representation. When the available main memory permits, BSOR-preconditioned projection methods such as TFQMR and BICGSTAB are also effective, as is the method of Arnoldi combined with a simple preconditioner. In case of tight memory constraints, the Jacobi and Gauss–Seidel iterations are also possible (but slower). When these numerical iterative methods apply, they outperform our implementation of Monte Carlo simulation for estimates of response measures such as the puff/spark *Score* and the probability distribution of the number of open channels (figure 5).

Using the approximate method APP_POWER, we determined the optimal partitioning strategy for a given number of partitions ($|\mathcal{P}| = 1, 2, 3, \dots$). As shown in table 3, the optimal partitioning strategy for $|\mathcal{P}| = 2$ or 3 is often one of the two intuitive groupings of states: \mathcal{C}/\mathcal{O} or $\mathcal{C}/\mathcal{R}/\mathcal{O}$. Using these partitioning strategies, the approximate method APP_POWER shows better scalability than the exact methods (figure 7(a)). In particular, APP_POWER with $\mathcal{C}/\mathcal{R}/\mathcal{O}$ partitioning has sufficient accuracy and competitive runtime compared to its exact counterpart POWER.

While Monte Carlo simulation shows better scalability for the convergence of specific coarse response measures, APP_POWER is preferred for problems of modest size ($N \leq 7$). Figure 7(b) also suggests that an approximate method will outperform Monte Carlo when finer response measures such as the full stationary distribution are desired. For a release site composed of NM -state channels, the Monte Carlo estimates of the stationary distribution require storage for M^N states, while the approximate method only stores $NM|\mathcal{P}|^{N-1}$ entries in its approximate representation of the stationary distribution, where $|\mathcal{P}|$ is the number of partitions. Given this storage savings for large problems and the scalability of APP_POWER relative to its exact counterpart (POWER), our results suggest that a promising avenue of future research is the implementation of an approximate method with a multi-level solver as its iterative engine. Because the convergence time of ML_JOR_F_DYN is several orders of magnitude faster than APP_POWER (figure 7(a)), an approximate multi-level solver could potentially outperform the Monte Carlo estimates of course response measures for the problem sizes considered here as well as larger N (see dashed line in figure 7(b)). Also, since single-channel models of IP₃R s and RyRs can be significantly more complicated than the three- and six-state models that are the focus of this paper, development of an approximate multi-level solver would represent an important contribution to the numerical analysis of the stochastic gating of instantaneously coupled Ca^{2+} -regulated Ca^{2+} channels. Finally, since some puff and spark statistics—such as puff/spark duration and inter-event interval distributions—cannot be determined from $\pi^{(N)}$, it is also important to determine if transient analysis can be accelerated using the Kronecker structure of the Ca^{2+} release site SAN descriptor (equations (7) and (8)).

The focus of this paper is the development of computationally efficient methods for stationary distribution calculations when channel positions are explicitly modeled. As shown in figure 1 and prior work [32], simulated Ca^{2+} puffs and sparks can be sensitive to the details of release site ultrastructure, and the exact and approximate numerical solution methods benchmarked here are appropriate in this situation. For Ca^{2+} release site models that are insensitive to channel position, one may perform mean-field calculations that do not account for the details of release site ultrastructure, provided the method maintains a distinction between each channel's substantial influence on its own stochastic gating and the collective contribution of elevated $[\text{Ca}^{2+}]$ from neighboring channels [13, 32]. In this manner, the state-space for a Ca^{2+} release site composed of NM -state channels can be reduced

from M^N to $(M + N - 1)!/N!(M - 1)!$ (cf table 1). The identification of structured representations for release sites composed of mean-field coupled channels is an important topic for further research.

As mentioned above, single-channel models of IP₃R s and RyRs can be significantly more complicated than the three- and six-state models that are the focus of this paper. For example, a recent IP₃R model includes 14 states, 6 of which are open [51], and the well-known De Young–Keizer IP₃R model includes 4 eight-state subunits for a total of 330 distinguishable states [52]. Unfortunately, we are currently unable to benchmark exact numerical solution methods for these more complex and realistic models (cf figure 4), because the state-space explosion for such Ca²⁺ release sites is overwhelming for large M (see table 1). Although beyond the scope of this paper, it might be possible to develop a SAN descriptor similar to equations (7) and (8) in which the elementary matrices K_- and K_+ correspond to a single channel subunit (thereby reducing M and increasing N). However, it is unclear whether iterative approximate numerical solution methods that utilize such a Kronecker representation would be more efficient than methods using the SAN descriptor presented here. In any case, development of an approximate multi-level solver (as discussed above) appears to be an important preliminary step in the development of numerical solution methods that can outperform Monte Carlo estimates for release sites composed of complex single-channel models.

Acknowledgments

The authors thank Buchholz and Dayar for sharing their implementation of Nsolve. This material is based upon work supported by the National Science Foundation under grant nos 0133132 and 0443843. Smith gratefully acknowledges a research leave during academic year 2007–2008 supported by the College of William and Mary and a long-term visitor position at the Mathematical Biosciences Institute at the Ohio State University. Some results of sections 4.1–4.3 have previously appeared in the form of conference proceedings [50].

Appendix A. Buffered diffusion of Ca²⁺ and the coupling matrix

To specify the values of the N -by- N coupling matrix C , we assume channels are localized on a planar ER membrane ($z = 0$). If we write $\mathbf{r}_i = x_i\hat{x} + y_i\hat{y}$ as the position of the pore of channel i , then assuming one high-concentration Ca²⁺ buffer, the local [Ca²⁺] at position $\mathbf{r} = x\hat{x} + y\hat{y} + z\hat{z}$ given by the ‘steady-state excess buffer approximation’ is [53, 36]

$$c(\mathbf{r}) = c_\infty + \sum_i \frac{\sigma_i}{2\pi D|\mathbf{r}_i - \mathbf{r}|} e^{-|\mathbf{r}_i - \mathbf{r}|/\lambda}, \quad (\text{A.1})$$

where σ_i is the source amplitude of channel i and λ is the buffer length constant. If we assume identical source amplitudes,

$$\sigma_i(t) = \begin{cases} 0 & \text{channel } i \text{ closed} \\ \sigma_O & \text{channel } i \text{ open,} \end{cases}$$

and write \mathbf{a}_j as the position of the Ca²⁺ regulatory site for channel j , the increase in Ca²⁺ experienced by channel j when channel i is open is given by

$$c_{ij} = \frac{\sigma_O}{2\pi D|\mathbf{r}_i - \mathbf{a}_j|} e^{-|\mathbf{r}_i - \mathbf{a}_j|/\lambda}.$$

Assuming the regulatory sites are located a small distance r_d above the channel pores, we write $\mathbf{a}_j = x_j\hat{x} + y_j\hat{y} + r_d\hat{z}$ and $r_{ij} = |\mathbf{r}_i - \mathbf{a}_j|$, so $r_{jj} = |\mathbf{r}_j - \mathbf{a}_j| = r_d$. Thus, the off-diagonal elements of the coupling matrix $C = (c_{ij})$ are

$$c_{ij} = \frac{\sigma_O}{2\pi D r_{ij}} e^{-r_{ij}/\lambda} \quad (i \neq j), \quad (\text{A.2})$$

and the diagonal elements of C are identical and given by

$$c_{ii} = c_d = \frac{\sigma_O}{2\pi D r_d} e^{-r_d/\lambda}. \quad (\text{A.3})$$

Note that $r_{ij} = r_{ji}$ implies that the interaction matrix is symmetric ($c_{ij} = c_{ji}$).

Appendix B. Instantaneous coupling and superposition of Ca²⁺-mediated interactions

Throughout this paper we assume that the formation and collapse of Ca²⁺ microdomains is fast compared to channel gating. This assumption allows specification of the Ca²⁺ concentration experienced by the Ca²⁺-regulatory site of each channel as an instantaneous function of the state of the release site (see appendix A). For simplicity, we use the ‘excess buffer approximation’ to determine these local Ca²⁺ concentrations, but other representations of Ca²⁺ buffering could be employed (for review, see [36]).

The SAN descriptor for N coupled Ca²⁺-regulated Ca²⁺ channels given by equations (7) and (8) assumes that Ca²⁺-mediated interactions between channels can be superposed, that is, the local [Ca²⁺] experienced by channel j can be written as $c_\infty + \sum_i \gamma_i c_{ij}$, where c_∞ is the background [Ca²⁺], c_{ij} is the increase in [Ca²⁺] experienced by channel j when channel i is open (e.g., equations (A.2) and (A.3)) and $\gamma_i = 0$ or 1 when channel i is closed or open, respectively. This superposition of Ca²⁺-mediated interactions occurs when the partial differential equations representing the buffered diffusion of intracellular Ca²⁺ are linear, as is the case for the excess buffer approximation [35, 54, 36].

While the time-scale of domain formation is extremely rapid (microseconds), the slower time-scale for domain collapse (tens of milliseconds) can lead to slow Ca²⁺ feedback on channel gating that is not represented in Ca²⁺ release site models that assume instantaneous coupling [22, 55]. Even when the formation and collapse of the Ca²⁺ microdomain is not fast compared to channel gating, instantaneous coupling may be assumed for convenience, so long as it is understood that this formalism may distort the relationship between single-channel kinetics and the stochastic dynamics of Ca²⁺ release sites. When release site ultrastructure is important and instantaneous coupling is not assumed, the exact and approximate numerical solution methods discussed in this paper do not apply, and stationary distributions must be estimated in a conventional fashion using Monte Carlo

simulation (see, e.g. [14]). In the case of mean-field coupled channels, the probability density of domain $[Ca^{2+}]$ jointly distributed with the state of the release site can be calculated directly as the steady state of a system of advection–reaction equations (a master equation) [22]. Using bivariate densities this approach can be extended to account for the effect of luminal depletion on the stochastic dynamics of Ca^{2+} release sites [55]. However, it is unclear if the Ca^{2+} release site SAN descriptor (equations (7) and (8)) and the exact and approximate numerical solution methods investigated in this paper could be used to accelerate calculations for which instantaneous coupling is not assumed.

Appendix C. The topology of the six-state model

The topology of the six-state model in figure 2(b) is derived from a two-subunit Ca^{2+} channel model that includes both fast Ca^{2+} activation and slower Ca^{2+} inactivation. Individual subunits are assumed to have topology

$$\mathcal{C} \rightleftharpoons \mathcal{O} \rightleftharpoons \mathcal{R} \quad (C.1)$$

where \cdot denotes a Ca^{2+} binding step and \mathcal{C} , \mathcal{O} and \mathcal{R} represent closed, permissive and refractory (i.e., long-lived closed) states, respectively. Expanding to obtain the topology of a two-subunit channel yields

$$\begin{array}{c} \mathcal{C}\mathcal{C} \rightleftharpoons \mathcal{C}\mathcal{O} \rightleftharpoons \mathcal{C}\mathcal{R} \\ \cdot \downarrow \uparrow \quad \uparrow \downarrow \cdot \\ \mathcal{O}\mathcal{O} \rightleftharpoons \mathcal{O}\mathcal{R} \rightleftharpoons \mathcal{R}\mathcal{R} \end{array} \quad (C.2)$$

where Ca^{2+} binding steps are denoted as above. Because we assume identical (but not independent) subunits we may lump equivalent states—e.g., $\mathcal{C}\mathcal{O}$ and $\mathcal{O}\mathcal{C}$ —leading to the six distinguishable states shown in equation (C.2). The six-state model shown in figure 2(b) is produced by identifying each of the six states of equation (C.2) as closed, open or refractory,

$$\begin{array}{c} \mathcal{C}_1 \rightleftharpoons \mathcal{C}_2 \rightleftharpoons \mathcal{R}_1 \\ \cdot \downarrow \uparrow \quad \uparrow \downarrow \cdot \\ \mathcal{O}_1 \rightleftharpoons \mathcal{R}_2 \rightleftharpoons \mathcal{R}_3 \end{array} \quad (C.3)$$

where we have assumed that both subunits must be in the permissive state for the channel to be open and, consequently, $\mathcal{O}\mathcal{O} = \mathcal{O}_1$ is the only open state. While the remaining states are not open, we designate states $\mathcal{C}\mathcal{C} = \mathcal{C}_1$ and $\mathcal{C}\mathcal{O} = \mathcal{O}\mathcal{C} = \mathcal{C}_2$ as closed and $\mathcal{C}\mathcal{R} = \mathcal{R}\mathcal{C} = \mathcal{R}_1$, $\mathcal{O}\mathcal{R} = \mathcal{R}\mathcal{O} = \mathcal{R}_2$ and $\mathcal{R}\mathcal{R} = \mathcal{R}_3$ as refractory because the sojourn time in states \mathcal{R}_1 , \mathcal{R}_2 and \mathcal{R}_3 is over three times longer than the sojourn time in states \mathcal{C}_1 and \mathcal{C}_2 at the background $[Ca^{2+}]$ of $c_\infty = 0.05 \mu\text{M}$. This separation of times scales is even larger at the higher $[Ca^{2+}]$ experienced in a Ca^{2+} release site when neighboring channels are open.

References

- [1] Colquhoun D and Hawkes A 1995 A Q-matrix cookbook: how to write only one program to calculate the single-channel and macroscopic predictions for any kinetic mechanism *Single-Channel Recording* ed B Sakmann and E Neher (New York: Plenum) pp 589–633
- [2] Smith G 2002 Modeling the stochastic gating of ion channels *Computational Cell Biology* ed C Fall, E Marland, J Wagner and J Tyson (Berlin: Springer) pp 291–325
- [3] Ball F, Milne R, Tame I and Yeo G 1997 Superposition of interacting aggregated continuous time Markov chains *Adv. Appl. Probab.* **29** 56–91
- [4] Ball F and Yeo G 1999 Superposition of spatially interacting aggregated continuous time Markov chains *Methodol. Comput. Appl. Probab.* **2** 93–116
- [5] Duke T A and Bray D 1999 Heightened sensitivity of a lattice of membrane receptors *Proc. Natl Acad. Sci. USA* **96** 10104–8
- [6] Duke T A, Noverre N L and Bray D 2001 Conformational spread in a ring of proteins: a stochastic approach to allosteric *J. Mol. Biol.* **308** 541–53
- [7] Bray D and Duke T 2004 Conformational spread: the propagation of allosteric states in large multiprotein complexes *Annu. Rev. Biophys. Biomol. Struct.* **33** 53–73
- [8] Cheng H, Lederer W and Cannell M 1993 Ca^{2+} sparks: elementary events underlying excitation–contraction coupling in heart muscle *Science* **262** 740–4
- [9] Cheng H, Lederer M, Lederer W and Cannell M 1996 Ca^{2+} sparks and Ca^{2+} i waves in cardiac myocytes *Am. J. Physiol.* **270** C148–59
- [10] Yao Y, Choi J and Parker I 1995 Quantal puffs of intracellular Ca^{2+} evoked by inositol trisphosphate in *Xenopus* oocytes *J. Physiol.* **482** 533–53
- [11] Parker I, Choi J and Yao Y 1996 Elementary events of IP_3 -induced Ca^{2+} liberation in *Xenopus* oocytes: hot spots, puffs and blips *Cell Calcium* **20** 105–21
- [12] Berridge M 1997 Elementary and global aspects of Ca^{2+} signalling *J. Phys. (Lond.)* **499** 291–306
- [13] Nguyen V, Mathias R and Smith G 2005 A stochastic automata network descriptor for Markov chain models of instantaneously-coupled intracellular Ca^{2+} channels *Bull. Math. Biol.* **67** 393–432
- [14] Swillens S, Dupont G, Combettes L and Champeil P 1999 From Ca^{2+} blips to Ca^{2+} puffs: theoretical analysis of the requirements for interchannel communication *Proc. Natl Acad. Sci. USA* **96** 13750–5
- [15] Rios E and Stern M 1997 Ca^{2+} in close quarters: microdomain feedback in excitation–contraction coupling and other cell biological phenomena *Annu. Rev. Biophys. Biomol. Struct.* **26** 47–82
- [16] Swillens S, Champeil P, Combettes L and Dupont G 1998 Stochastic simulation of a single inositol 1,4,5-trisphosphate-sensitive Ca^{2+} channel reveals repetitive openings during ‘blip-like’ Ca^{2+} transients *Cell Calcium* **23** 291–302
- [17] Stern M et al 1999 Local control models of cardiac excitation–contraction coupling. A possible role for allosteric interactions between ryanodine receptors *J. Gen. Physiol.* **113** 469–89
- [18] Shuai J and Jung P 2002 Stochastic properties of Ca^{2+} release of inositol 1,4,5-trisphosphate receptor clusters *Biophys. J.* **83** 87–97
- [19] Rengifo J, Rosales R, Gonzalez A, Cheng H, Stern M and Rios E 2002 Intracellular Ca^{2+} release as irreversible Markov process *Biophys. J.* **83** 2511–21
- [20] Shuai J and Jung P 2003 Optimal ion channel clustering for intracellular calcium signaling *Proc. Natl Acad. Sci. USA* **100** 506–10
- [21] Hinch R 2004 A mathematical analysis of the generation and termination of calcium sparks *Biophys. J.* **86** 1293–307
- [22] Mazzag B, Tignanelli C and Smith G 2005 The effect of residual Ca^{2+} on the stochastic gating of Ca^{2+} -regulated Ca^{2+} channels *J. Theor. Biol.* **235** 121–50

- [23] DeRemigio H and Smith G 2005 The dynamics of stochastic attrition viewed as an absorption time on a terminating Markov chain *Cell Calcium* **38** 73–86
- [24] Groff J and Smith G 2008 Ryanodine receptor allosteric coupling and the dynamics of calcium sparks *Biophys. J.* **95** 135–54
- [25] Fill M and Copello J 2002 Ryanodine receptor calcium release channels *Physiol. Rev.* **82** 893–922
- [26] Hagar R, Burgstahler A, Nathanson M and Ehrlich B 1998 Type III IP₃ receptor channel stays open in the presence of increased Ca²⁺ *Nature* **396** 81–4
- [27] Mak D and Foskett J 1997 Single-channel kinetics, inactivation, and spatial distribution of inositol trisphosphate (IP₃) receptors in *Xenopus* oocyte nucleus *J. Gen. Physiol.* **109** 571–87
- [28] Moraru I, Kaftan E, Ehrlich B and Watras J 1999 Regulation of type 1 inositol 1,4,5-trisphosphate-gated Ca²⁺ channels by IP₃ and Ca²⁺: simulation of single channel kinetics based on ligand binding and electrophysiological analysis *J. Gen. Physiol.* **113** 837–49
- [29] Ramos-Franco J, Fill M and Mignery G 1998 Isoform-specific function of single inositol 1,4,5-trisphosphate receptor channels *Biophys. J.* **75** 834–9
- [30] Bezprozvanny I, Watras J and Ehrlich B 1991 Bell-shaped Ca²⁺-response curves of Ins(1,4,5)P₃- and Ca²⁺-gated channels from endoplasmic reticulum of cerebellum *Nature* **351** 751–4
- [31] Groff J and Smith G 2008 Calcium-dependent inactivation and the dynamics of calcium puffs and sparks *J. Theor. Biol.* at press doi:10.1016/j.jtbi.2008.03.026
- [32] DeRemigio H, Groff J and Smith G 2008 Calcium release site ultrastructure and the dynamics of puffs and sparks *Math. Med. Biol.* at press. doi:10.1093/imammb/dqn004
- [33] Smith G 2002 An extended DeYoung–Keizer-like IP₃ receptor model that accounts for domain Ca²⁺-mediated inactivation *Recent Research Developments in Biophysical Chemistry* vol 2 ed C Condat and A Baruzzi (India: Research Signpost) pp 37–55
- [34] Bezprozvanny I 1994 Theoretical analysis of Ca²⁺ wave propagation based on inositol (1,4,5)-trisphosphate receptor functional properties *Cell Calcium* **16** 151–66
- [35] Naraghi M and Neher E 1997 Linearized buffered Ca²⁺ diffusion in microdomains and its implications for calculation of [Ca²⁺] at the mouth of a Ca²⁺ channel *J. Neurosci.* **17** 6961–73
- [36] Smith G, Dai L, Muira R and Sherman A 2001 Asymptotic analysis of equations for the buffered diffusion of intracellular Ca²⁺ *SIAM J. Appl. Math.* **61** 1816–38
- [37] Stewart W 1994 *Introduction to the Numerical Solution of Markov Chains* (Princeton: Princeton University Press)
- [38] Fernandes P, Plateau B and Stewart W J 1998 Efficient descriptor-vector multiplications in stochastic automata networks *J. ACM* **45** 381–414
- [39] Gillespie D T 1976 A general method for numerically simulating the stochastic time evolution of coupled chemical reactions *J. Comput. Phys.* **22** 403–34
- [40] Buchholz P and Dayar T 2004 Comparison of multilevel methods for Kronecker-based Markovian representations *Computing* **73** 349–71
- [41] Buchholz P and Dayar T 2007 On the convergence of a class of multilevel methods for large, sparse Markov chains *SIAM Matrix Anal. Appl.* **29** 1025–49
- [42] Buchholz P 2000 Multilevel solutions for structured Markov chains *SIAM J. Matrix Anal. Appl.* **22** 342–57
- [43] Horton G and Leutenegger S T 1994 A multi-level solution algorithm for steady-state Markov chains *SIGMETRICS Perform. Eval. Rev.* **22** 191–200
- [44] Buchholz P and Kemper P 2004 Kronecker based matrix representations for large Markov models *Validation of Stochastic Systems (Lecture Notes in Computer Science vol 2925)* ed C Baier, B R Haverkort, H Hermanns, J P Katoen and M Siegle (Berlin: Springer) pp 256–95
- [45] Buchholz P 2004 Adaptive decomposition and approximation for the analysis of stochastic Petri nets *Perform. Eval.* **56** 23–52
- [46] Buchholz P and Kemper P 1999 A toolbox for the analysis of discrete event dynamic systems in N Halbwachs and D Peled (eds) *Computer Aided Verification (CAV '99)* (Berlin: Springer) pp 483–6
- [47] Buchholz P and Dayar T 2005 Block SOR preconditioned projection methods for Kronecker structured Markovian representations *SIAM J. Sci. Comput.* **26** 1289–313
- [48] Buchholz P 1999 Projection methods for the analysis of stochastic automata networks *Numerical Solution of Markov Chains* ed B Plateau, W Stewart and M Silva (Zaragoza: Prensas Universitarias de Zaragoza) pp 149–68
- [49] Buchholz P 1997 An aggregation–disaggregation algorithm for stochastic automata networks *Prob. Eng. Inform. Sci.* **11** 229–53
- [50] DeRemigio H, Kemper P, LaMar M and Smith G 2008 Markov chain models of coupled intracellular calcium channels: Kronecker structured representations and benchmark stationary distribution calculations *Pacific Symp. Biocomput.* **13** 354–65
- [51] Fraiman D and Dawson S 2004 A model of IP₃ receptor with a luminal Ca²⁺ binding site: stochastic simulations and analysis *Cell Calcium* **35** 403–13
- [52] De Young G and Keizer J 1992 A single-pool inositol 1,4,5-trisphosphate-receptor-based model for agonist-stimulated oscillations in Ca²⁺ concentration *Proc. Natl Acad. Sci. USA* **89** 9895–9
- [53] Neher E 1986 Concentration profiles of intracellular Ca²⁺ in the presence of diffusible chelator *Exp. Brain Res.* **14** 80–96
- [54] Neher E 1998 Usefulness and limitations of linear approximations to the understanding of Ca²⁺ signals *Cell Calcium* **24** 345
- [55] Huertas M and Smith G 2007 The dynamics of luminal depletion and the stochastic gating of Ca²⁺-activated Ca²⁺ channels and release sites *J. Theor. Biol.* **246** 332–54