Available online at www.sciencedirect.com
SCIENCE DIRECT
BioSystems 85 (2006) 37–45
BioSystems
www.elsevier.com/locate/biosystems

Formal agent-based modelling of intracellular chemical interactions

Mark Pogson [a], Rod Smallwood [a], Eva Qwarnstrom [b], Mike Holcombe [a,*]

[a] *Computational Systems Biology Group, Department of Computer Science, University of Sheffield, UK*
[b] *Cell Biology Group, Division of Genomic Medicine, University of Sheffield, UK*

Received 6 February 2006; accepted 7 February 2006

Abstract

Individual-based or agent-based models have proved useful in a variety of different biological contexts. This paper presents an agent-based model using a formal computational modelling approach to model a crucial biological system—the intracellular NF-κB signalling pathway. The pathway is vital to immune response regulation, and is fundamental to basic survival in a range of species. Alterations in pathway regulation underlie many diseases, including atherosclerosis and arthritis. Our modelling of individual molecules, receptors and genes provides a more comprehensive outline of regulatory network mechanisms than previously possible with equation-based approaches. The model has been validated with data obtained from single cell experimental analysis.

Keywords: Agent-based modelling; Chemical interaction; Nuclear factor kappa B; Signalling pathway

1. Introduction

Many aspects of life involve the interaction of multiple components and subunits and the corresponding emergence of both form and function. This is true whether we are dealing with molecules within an individual cell, cells within tissue, organs within an organism or organisms within a community or ecology. Understanding how the components interact, develop and die is a crucial part of understanding the system and its lifecycle. Agent-based (sometimes called individual-based) approaches – whereby the components are represented as autonomous software artefacts that exist within a software environment – provide a mechanism for understanding their behaviour through simulation of the actual behaviour of the equivalent biological system.

Such approaches have been taken by Paton et al. and a number of other authors, particularly in the area of population modelling (Gregory et al., 2002).

2. Previous modelling approaches for the NF-κB pathway

Intracellular signalling pathways are vital to the control and regulation of cell behaviour. Complex interactions of genes, proteins and other molecules within the cell must be addressed in order to gain a better understanding of how these pathways operate (Alberts et al., 2002; Nelson, 2004; Lodish et al., 2000). Rather than simply catalogue and characterise the physical components of the cell, it is necessary to bring together this information in mathematical models to understand the functional activity of signalling pathways (Cho and Wolkenhauer, 2003).

Activation of the NF-κB pathway is controlled by inhibitors of NF-κB (IκB) proteins, which sequester the majority of NF-κB in the cytoplasm as complexes by masking their nuclear localisation signals. During

* Corresponding author. Tel.: +44 114 222 1802; fax: +44 114 222 1810.
E-mail address: m.holcombe@dcs.shef.ac.uk (M. Holcombe).

0303-2647/$ – see front matter © 2006 Published by Elsevier Ireland Ltd.
doi:10.1016/j.biosystems.2006.02.004

activation, IκB is phosphorylated by IκB kinases (IKK), causing its destruction. The newly freed NF-κB is consequently transported into the nucleus, inducing inflammatory genes, including those encoding IκB, thus regulating the pathway through negative feedback (Carlotti et al., 1999, 2001; Yang et al., 2003).

Pathway activation is tightly controlled at multiple levels. Detailed information on the parameters regulating specific steps and their impact on activation is of fundamental importance for understanding the pathway as a whole. Recently, modelling of regulation at the level of the inhibitor has been performed using differential equations (Hoffmann et al., 2002; Lipniacki et al., 2004; Nelson et al., 2004). Previous attempts to quantitatively model intracellular signalling pathways have often relied on reaction kinetics, employing ordinary differential equations (ODEs) to describe each chemical concentration with time. This implicitly assumes homogeneity, treating the cell as a well-mixed bag of chemicals Burrage et al. (2003). Due to internal structure and low numbers and non-uniform distributions of certain key molecules in the cell, this is certainly not true. While the assumption may be reasonable in some circumstances, a more detailed modelling method is necessary to achieve a better understanding of mechanical and structural effects on intracellular pathways.

Due to the complexity of signalling pathways, large numbers of linked ODEs are often necessary for a reaction kinetics model. This is not only undesirable due to the clumsiness of the description, but solutions can also be difficult, especially with varying parameters. Such models can be narrowly limited in the range in which they function properly, which is at odds with the robustness of nature. The many interdependent differential equations can be very sensitive to their initial conditions and constants, with small changes in these sometimes causing huge behavioural changes in the system, though this is obviously also dependent on the equation solver used. Solutions may be liable to describe unrealistic or impossible behaviour, such as negative concentrations, unless the initial conditions and constants are correct to a high degree of precision. Regardless of whether the necessary data are even biologically or chemically accurate, such a critical dependence is clearly unfeasible in a real cell. While differential equation models may produce useful results under certain conditions, they provide a rather incomplete view of what is actually happening in the cell (Bhalla and Iyengar, 1999).

Time delays in certain cellular processes such as transcription can have very significant effects on pathway behaviour. However, these are difficult to include in a differential equation model, and are rarely given consideration. Spatial effects on pathways are also often crucial, but again difficult to incorporate into such a model, though partial differential equations (PDEs) could theoretically be used. Despite the importance of differential equation models, their scope is limited, and they are unable to fully represent the complex interactions inside the cell (Papin et al., 2005).

Other modelling techniques have also been applied to signalling pathways, though generally with less emphasis on dynamical properties. However, most methods treat biochemical networks simply as communication systems, focussing to a greater or lesser degree on the movement of information as opposed to the physics of interactions. While this is appropriate in many circumstances, a more realistic description of the cell (rather than an abstract pathway) is necessary for certain pathways, such as the NF-κB pathway, where cell structure and mechanical stimulation are particularly important. Although several models overcome some of the problems associated with reaction kinetics models, none consider intracellular pathways as anything more than just the chemicals involved. For a more complete appreciation of many pathways, this must be rectified.

3. An agent-based approach

Agent-based modelling treats each individual component of a system as a single entity (or agent) obeying its own pre-defined rules and reacting to its environment and neighbouring agents accordingly. Various computational models exist to represent agents, though the communicating X-machine is perhaps the most powerful and intuitive (Eleftherakis et al., 2001; Balanescu et al., 1999). Agent-based modelling has recently been applied to a variety of biological systems, including insect communities and epithelial tissue (Jackson et al., 2004a, 2004b; Walker et al., 2004; Holcombe et al., 2003).

Agents can be used to represent any identifiable component of a system. For a biochemical pathway, this means that anything from a molecule to a signalling receptor to an entire chain of interactions can be modelled as an agent, thus providing a modular and extensible modelling framework that allows abstraction of detail as necessary.

An agent-based chemical model does not have the same restrictions as a differential equation model; any number (low or high, within computational limitations) and distribution of molecules can be modelled, spatial concerns can easily be accounted for, as can time delays, and individual interactions of agents do not produce the same ‘volatility’ as several interdependent differential equations. Furthermore, it is not possible for there to be

a negative number of agents, so there is no possibility for negative concentrations. An agent-based model also provides a clearer picture of what is actually occurring in the cell.

Although an agent-based model allows greater scope than a differential equation model, it also requires more details to be defined. An obvious concern is that the movement of agents must be explicitly stated. This is not a problem, however, as any reasonably random movement within an agent's confines is sufficient for the model to operate properly; depending on the resolution of the model, different movement rules will have a different effect on results, but agents must at least move around enough to regularly collide. Furthermore, as the model is developed, the movement of agents can be used to define cell structure. The movement of agents raises the concern of energy conservation, though since pathways depend on energy stored from the metabolism (which is beyond the scope of any pathway model), this is not a major problem.

While many agent types in a signalling pathway model may include abstractions of detail, it is vital that the agent-based model is able to deal with individual interactions of molecule agents with the same accuracy as reaction kinetics. As discussed, reaction kinetics treats reacting chemicals as being uniformly well mixed, and implicitly assumes that the number of molecules involved is large. In contrast, an agent-based approach can treat each chemical molecule as a single identifiable entity, as shown in Fig. 1. Thus any distribution and number of molecules is permissible.

The agent-based model must of course agree with the corresponding reaction kinetics model in the circumstance where reaction kinetics can reasonably be applied (i.e. with large numbers of molecules of well-mixed chemicals). Since information about reacting chemicals is invariably given for such a situation, and because little information exists about individual molecular interactions, it is important that the necessary data for the agent-based model can be inferred from

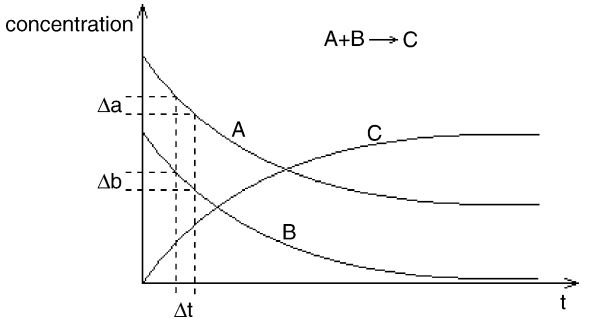

Fig. 2. Generalised graph of concentration against time t for two chemicals A and B interacting to form C. With a time step of Δt, the change in A concentration is Δa, and the change in B concentration Δb. Since one A molecule interacts with one B molecule to form one C molecule, it follows that $\Delta a = \Delta b$. Axes are linear.

reaction kinetics. This has been considered earlier by Andrews and Bray (2004), but their consideration of the below is only brief and greater emphasis is placed on more detailed interactions, which in the case of large interaction networks is prohibitively computationally intensive.

Take two reacting chemicals A and B combining to form C, with rate constant k (neglecting the reverse reaction for the moment):

$$\mathrm{A} + \mathrm{B} \xrightarrow{k} \mathrm{C} \tag{1}$$

If the concentrations of A, B and C at time t are $a(t)$, $b(t)$ and $c(t)$, respectively, then from reaction kinetics:

$$\frac{\mathrm{d}a(t)}{\mathrm{d}t} = -ka(t)b(t) \tag{2}$$

Now if the magnitude of the change in concentration of A during a time step of Δt is Δa, as shown in Fig. 2, then:

$$a(t + \Delta t) = a(t) - \Delta a \tag{3}$$

The iterative process of (3) acts as the basis for agent-based interactions, with Δt a constant time step.

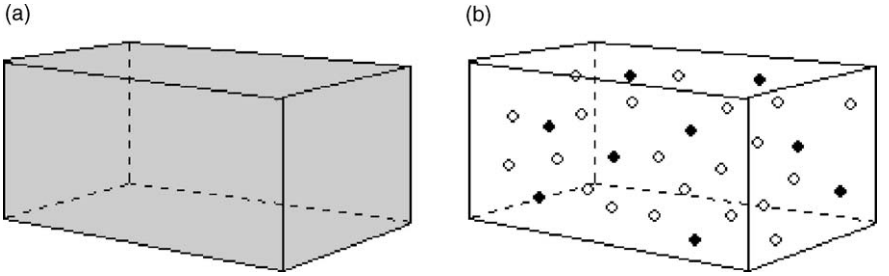

Fig. 1. Chemical reactions. Reaction kinetics differential equations treat reacting chemicals as well mixed and uniform (a), whereas the agent-based approach models each individual molecule (b).

Assuming Δt is sufficiently small, then:

$$\Delta a \approx -\Delta t \frac{\mathrm{d}a(t)}{\mathrm{d}(t)} \tag{4}$$

Substituting (2) into (4) gives:

$$\Delta a = ka(t)b(t)\Delta t \tag{5}$$

Since one A molecule interacts with one B molecule to form one C molecule, it follows that $\Delta a = \Delta b$, where Δb is the magnitude of change in concentration of B in time step Δt. Therefore, the proportion of B molecules that interact with A at each time step is:

$$\frac{\Delta b}{b(t)} = \frac{\Delta a}{b(t)} = \frac{ka(t)b(t)\Delta t}{b(t)} = ka(t)\Delta t \tag{6}$$

If V is the volume of the container, and A and B are randomly distributed, then in order for the proportion of B molecules that interact with A to equal that obtained in (6), it must hold that the total volume inside which A interacts with B (denoted V_i, as shown in Fig. 3) is:

$$V_\mathrm{i} = ka(t)\Delta t V \tag{7}$$

The number of A molecules n_a at any time is given by

$$n_\mathrm{a}(t) = 10^3 a(t) L V \tag{8}$$

where L is Avagadro's constant, and the factor of 10^3 is required if the concentration $a(t)$ is measured in molars (as is generally the case) and V measured in cubic metres.

So each A molecule can interact with a B molecule if the B molecule is within an 'interaction volume' v surrounding A of size:

$$v = \frac{V_\mathrm{i}}{n_\mathrm{a}(t)} = \frac{ka(t)\Delta t V}{10^3 a(t) L V} = \frac{k\Delta t}{10^3 L} \tag{9}$$

Assuming each individual interaction volume is spherical, then the maximum possible separation, or 'interaction radius' r, of two interacting A and B molecules is:

$$r = \sqrt[3]{\frac{3k\Delta t}{4\pi 10^3 L}} \tag{10}$$

The rate constant can now therefore be used to deduce information on local interactions. The distance r will be larger than the combined molecular radii of A and B due to the length of the time step, since molecules move around in this time. The time step used must of course be sufficiently small for this approach to hold (as assumed in (4)), though it must also be large enough for the average step length of each agent to exceed the binding radius.

Checking the result of (10): if L is length and T is time, then the dimensions (denoted by square brackets) of r are:

$$[r] = \sqrt[3]{[k][\Delta t]} = \sqrt[3]{(L^3 T^{-1})T} = L \tag{11}$$

as required.

It is worth noting that the same result is obtained regardless of whether A or B is used to calculate the interaction radius, as should be expected. However, unless A and B are both of the same concentration, it seems reasonable to question this statement: if the interaction radius is the same no matter how it was calculated, then the total interaction volume V_i will be different depending on whether A or B are 'assigned' this radius, when it might be expected that it should be the same. On the other hand, if the total interaction volume were to remain the same, then this would require the interaction radius change, which certainly should not be the case.

It is actually correct that the total interaction volume changes (recall that (7) does indeed depend on the choice of molecule). Say there are more A molecules than B. If A molecules are each assigned the interaction radius, then V_i will be larger than if B molecules were assigned the radius, but the number of molecules they seek to interact with is smaller. Likewise, if B molecules are assigned the interaction radius, then V_i is smaller, but the number of molecules they are seeking is larger. The number of

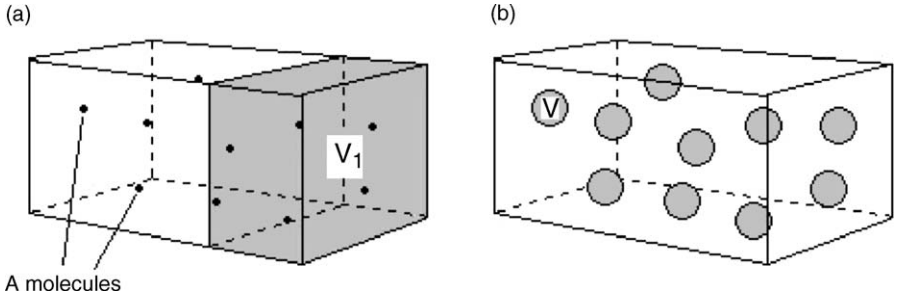

Fig. 3. Interaction volumes. In (a), the shaded volume shows the size of the total 'interaction volume' V_i; (b) shows this volume 'redistributed' over all the A molecules, giving each its own interaction volume v. B molecules are not shown for simplicity.

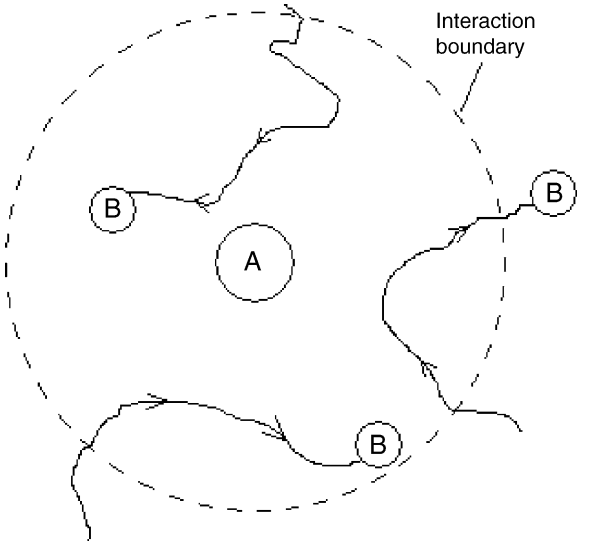

Fig. 4. Possible movement of B molecules during a time step Δt in the centre-of-mass frame of an A molecule. Since the movement and orientation of molecules is not known during the time step, a random element is required in the interaction decision.

A and B molecules in each other's interaction volume is consequently the same regardless. This is apparent in the derivation of r, though it may not be immediately obvious.

This local interaction method can be used not only for large numbers of randomly distributed molecules, but also for any molecular distributions since the interaction radius obtained should hold whatever the circumstance. Such local interactions are key to the model, and allow its expansion to deal with any interactions—not just straightforward large-scale chemical reactions.

Only the possibility of interaction has been considered so far, and not the actual process of interaction. For instance, if more than one B molecule lies inside an A molecule's interaction volume, it must have some method to 'choose' the most suitable B molecule to bind. This could be done at random, and with large numbers of molecules this would be of little consequence, but since the purpose of the method is to deal with any situation, it seems sensible to have a more physically realistic 'decision' process.

A spherical interaction volume is of course an approximation to the actual region inside which another molecule must be to interact, since the orientation of molecules is vital for the interaction. Clearly knowledge of orientation of molecules is beyond the scope of the model, especially since interactions occur during a time step where the orientations will have changed and the movement of the molecules is not exactly known, as represented in Fig. 4. So although a spherical interaction volume is used, the attraction of molecules will not be the same in all directions, hence a randomness in the decision of interaction is required to account for the mutual orientation of molecules, as well as how close they may have been during the time step.

An inverse square law is used in conjunction with the random element of the decision to account for the non-covalent attractive forces involved. Hence, if the distance to each molecule is d, and R is an appropriately scaled random number ('appropriately' meaning it is neither dominant nor insignificant in comparison to the inverse square of the distance), then the molecule with the largest value of R/d^2 is chosen to interact.

It is possible that even if a molecule is within the required distance to interact, its orientation is wrong, or the attraction is insufficient to bind. This could be modelled by including a 'threshold' over which R/d^2 must be for molecules to interact. Whilst this would affect the reaction rate since slightly fewer molecules would interact, it is an issue that should be considered. In addition to this concern, it is clear in Fig. 4 that it is possible for a molecule to have been able to interact during the time step, but due to the time at which calculations are performed it would not be considered. However, this is only a minor issue that can be neglected as being highly unlikely, and choosing an appropriately small time step can minimize the problem. The two issues should to a degree cancel out since they have an opposite effect, and if greater detail were required a more exact technique would be necessary.

So far, only association of molecules has been considered, though dissociation is of course also possible. Dissociation can easily be accounted for by making bound products separate randomly at a rate specified by the dissociation rate constant (using a uniform random distribution initially, though this could be modified as appropriate). Although this is straightforward, it may be an unnecessary consideration since the rate of dissociation is often negligible, and its inclusion may simply unnecessarily overcomplicate computation. However, some features of signalling pathways may depend on dissociation, such as the shuttling of NF-κB and IκB between the cytoplasm and nucleus in unstimulated cells, so the model must be able to deal with it if necessary.

First-order reactions can be modelled by using a similar random change as would be used for two-particle dissociation. Higher-order reactions are modelled by the relevant number of two-particle local interactions (where the two 'particles' become products of earlier interactions), as are enzyme interactions.

Just as terms are added to reaction kinetics differential equations to deal with several chemical reactions,

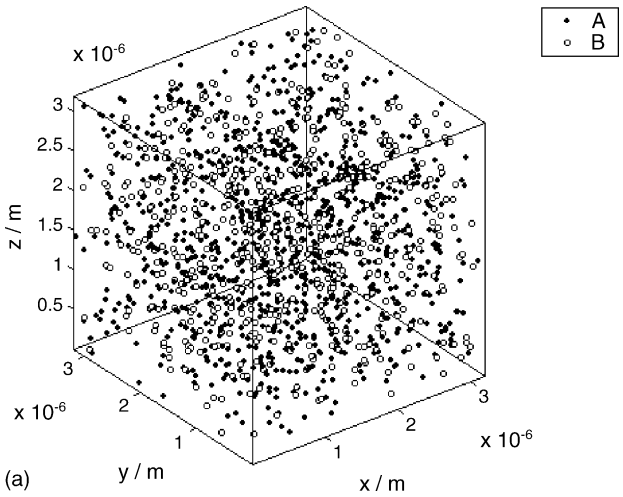

Fig. 5. Results from the chemical model. (a) Molecule agents are placed within a container, which move and interact according to their rules. (b) Left column: agent-based model results. Right column: corresponding reaction kinetics differential equation results. Top row: A + B → C. Middle row: A + B ↔ C. Bottom row: A + B → D, A + C → E.

each molecule agent can seek interaction with several relevant types of molecule, with the appropriately sized interaction radius for each type.

The method used to calculate interaction is not perfect, but it allows pathways to be modelled in greater detail than with differential equations. The discontinuity of attraction at the 'interaction radius' is not quite a true reflection of reality, but it is sufficiently accurate for the purposes of most models, and it allows rate constant data to be used to infer information for local interactions.

Results of the agent-based chemical model are in good agreement with reaction kinetics differential equations, as shown in Fig. 5. Using the method to model the NF-κB pathway also provides promising results, as shown in Fig. 6.

4. An agent-based model of the NF-κB pathway

In this model, molecular agents diffuse through the cell, binding and dissociating from other molecules, receptors and cell structures in accord with signals they send and receive from surrounding agents. Every agent is represented by a complete computational model – the communicating stream X-machine – that provides an intuitive and rigorous basis to model the functional behaviour of systems in a flexible and extensible manner. An important feature is the memory of each agent's X-machine, which contains its physical location, meaning that the number of states required to model the system is manageably small. It is essential that the agents are both biologically plausible as entities and that their behaviour is based on experimental measurements. In the model, as in reality, molecular interactions are local events that depend only on the position and current state of the molecules involved, where the state of a molecule is whether or not it is already bound. The physics of a molecule is modelled according to specific agent-based characteristics, including which types of interaction are possible. If two molecules may interact according to the rules, they must satisfy criteria on their state and proximity, derived from standard rate constants as above. If interaction occurs, the state of each agent changes to a 'bound' state, which can be reversed through random thermal separation.

TIR mediated activation is considered. Each NF-κB, IκB and IκB-kinase (IKK) molecule is an individual agent, as are the importing and exporting nuclear and cell surface receptors (Fig. 6a). Soluble extracellular agonists are not modelled as agents but treated as a whole chemical entity whose fluctuating local concentrations at the cell surface must rise above a certain level to initiate signal transduction in nearby receptor agents. Similarly, in relation to the cytoplasmic TIR domain, the local concentrations of certain molecules in the cytoplasm must be above a defined level in the vicinity of an active receptor to complete the process. Following this, a temporary agent with an internal time delay is created to account for the cascade that triggers the IκB kinases (IKKs). At the final steps of activation, an analogous temporary agent method is used to account for the translation of IκB.

The agents are contained within a spherical cell consisting of a cytoplasm and concentric spherical nucleus (Fig. 6b). Space is continuous and time is discrete in the model. The rules that govern agent movements

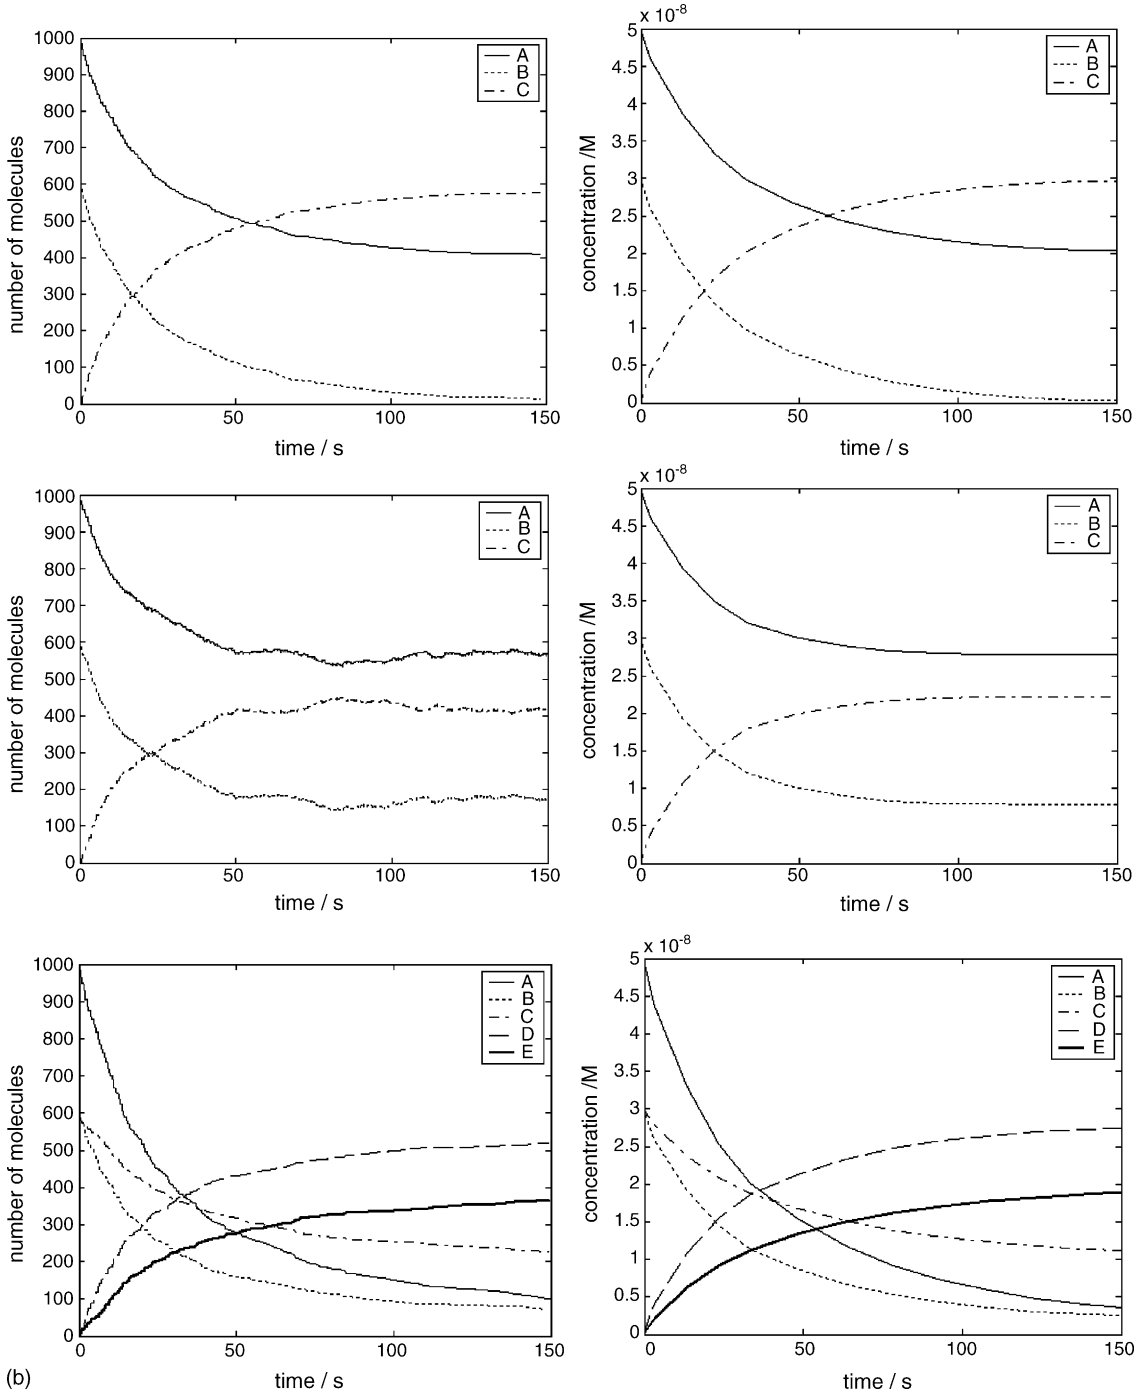

Fig. 5. (*Continued*).

incorporate cell structure by defining spatial boundaries. This allows investigation of the impact of cell shape and biomechanical effects.

The model is in good agreement with biological data obtained by real time single cell analysis (Fig. 6c and d). Continuous monitoring of signal transduction events in live cells was performed using GFP-tagged regulatory intermediates and confocal microscopy (Carlotti et al., 1999, 2001; Yang et al., 2003). Simultaneous observations of the NF-κB subunit relA and IκBα were carried out using cyan (ECFP) and yellow (EYFP) variants of GFP, respectively, as previously (Yang et al., 2003).

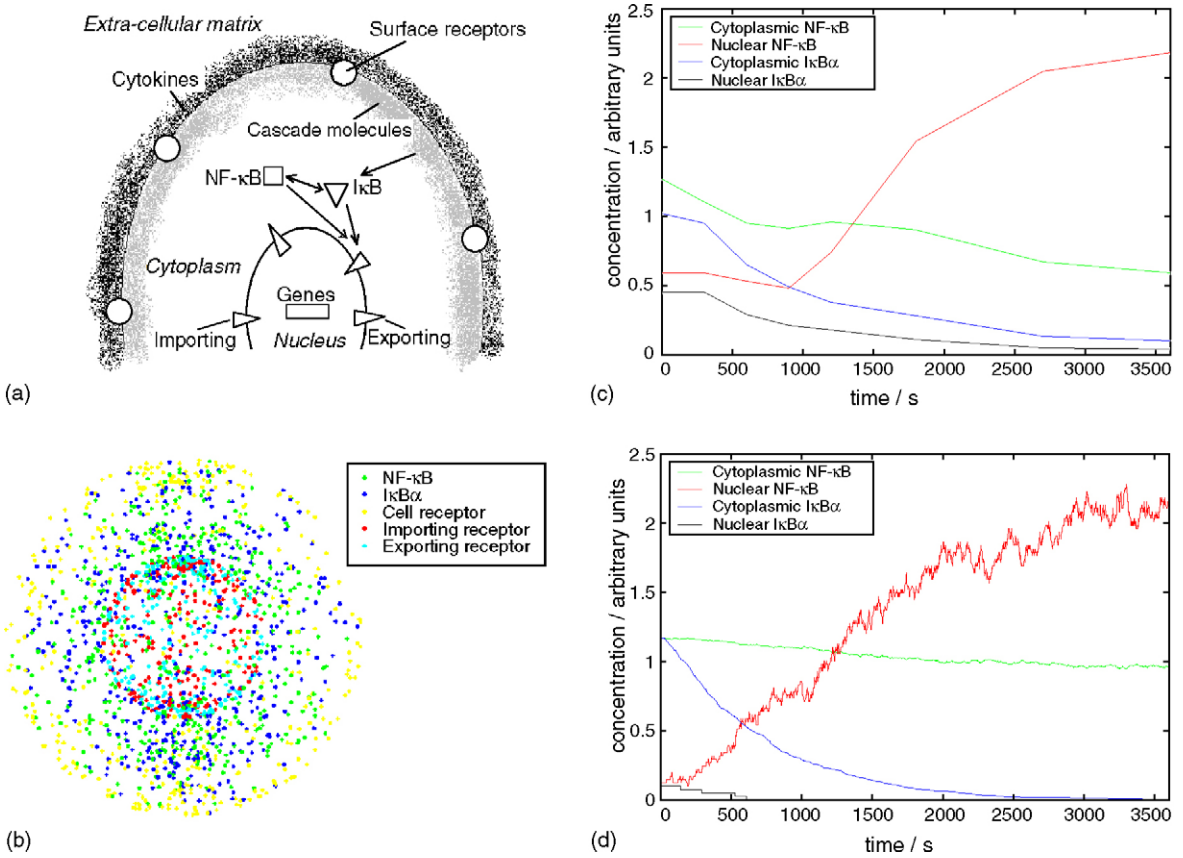

Fig. 6. Formulation and validation of the agent-based model. (a) Simplified diagram of the principal pathway agents in the model. Each agent can exist in a number of states. (b) Three-dimensional visualisation of the positions of agents in the model at a moment in time. The cell is scaled down to reduce computation, containing in the order of 1000 agents. Concentrations of molecules are based on biological data. The genes that NF-κB can activate are placed randomly along a line at the centre of the nucleus. (c) Quantitation of single cell data transfected with ECFPrelA and IκBαEYFP. Prior to stimulation, both components are located in the cytoplasm (top row). Following pathway activation, NF-κB translocates to the nucleus whilst IκBα and NF-κB–IκBα complex levels fall. (d) Model results following TIR activation over the same time period. Results are for a single cell, and demonstrate fundamental similarities with experiment.

Complex-formation was assessed by measuring fluorescence resonance energy transfer (FRET), with ECFPrelA the donor and IκBαEYFP the acceptor.

The spatial detail of the model permits explicit modelling of structural components of the cell such as the cytoskeleton, which is partly composed of actin filaments. This is the subject of a further paper.

5. Conclusions

The formal agent-based modelling paradigm illustrated here provides a useful technique for understanding many aspects of biological systems. It is a suitable model for cellular regulatory events such as the NF-κB pathway, providing a clear and intuitive mechanism to determine and explore the key features of the system. The model provides predictions about regulation of the pathway, which are unachievable by other modelling methods, and displays an appropriate robustness, allowing future investigation of the variability of cells in nature.

Agents provide a powerful framework for more detailed modelling of intracellular signalling pathways, allowing a greater understanding of their operation. The scope of the method allows newly available single-cell confocal microscopy data to be utilised to the full. The detailed mathematical structure of agent models may also allow the use of automated reasoning techniques such as model checking to properly understand the system in all possible circumstances within which it operates (Eleftherakis et al., 2001). This is an important capability that has not yet been exploited in biological modelling. The description of the individual agents can be formalised using a specially designed language based on XML and compiled directly into a general simula-

tion engine which will be developed later for parallel computer systems and GRID facilities.

Acknowledgements

The authors wish to acknowledge the support from the EPSRC/BBSRC, the Wellcome Trust and the MRC and the contribution of former colleagues.

References

Alberts, B., Johnson, A., Lewis, J., Raff, M., Roberts, K., Walter, P., 2002. Molecular Biology of the Cell, 4th ed. Garland Science.

Andrews, S.S., Bray, D., 2004. Stochastic simulation of chemical reactions with spatial resolution and single molecule detail. Phys. Biol. 1, 137–151.

Balanescu, T., Holcombe, M., Cowling, A.J., Gheorgescu, H., Georghe, M., Vertan, C., 1999. Communicating stream X-machines systems are no more than X-machines. J. Universal Comput. Sci. 5, 494–507.

Bhalla, U.S., Iyengar, R., 1999. Emergent properties of networks of biological signaling pathways. Science 283, 381–387.

Burrage, K., Burrage, P., Jeffrey, S., Pickett, T., Sidje, R., Tian, T., 2003. A grid implementation of chemical kinetic simulation methods in genetic regulation. Proceedings of the APAC03 Conference on Advanced Computing, Grid Applications and eResearch.

Carlotti, F., Chapman, R., Dower, S.K., Qwanstrom, E.E., 1999. Activation of nuclear factor κB in single living cells. J. Biol. Chem. 274, 37941–37949.

Carlotti, F., Dower, S.K., Qwanstrom, E.E., 2001. Dynamic shuttling of nuclear factor κB between the nucleus and cytoplasm as a consequence of inhibitor dissociation. J. Biol. Chem. 275, 41028–41034.

Cho, K.-H., Wolkenhauer, O., 2003. Analysis and modelling of signal transduction pathways in systems biology. Biochem. Soc. Trans. 31, 1503–1509.

Eleftherakis, G., Kefalas, P., Sotiriadou, A., 2001. XmCTL: extending temporal logic to facilitate formal verification of X-machine models. Anal. Universit. Bucuresti, Mat. Inform. 50, 79–95.

Gregory, R., Paton, R.C., Saunders, J.R., Wu, Q.H., 2002. A model of bacterial adaptability based on multiple scales of interaction: COSMIC. In: Paton, R., Belouri, H., Holcombe, M., Parish, J.H., Tateson, R. (Eds.), Computation in Cells and Tissues. Springer.

Hoffmann, A., Levchenko, A., Scott, M.L., Baltimore, D., 2002. The IκB–NF-κB signalling module: temporal control and selective gene activation. Science 298, 1241–1245.

Holcombe, M., Holcombe, L., Gheorghe, M., Talbot, N., 2003. A hybrid machine model of rice blast fungus, Magnaporthe Grisea. Biosystems 68, 223–228.

Jackson, D., Holcombe, M., Ratnieks, F., 2004a. Coupled computational simulation and empirical research into the foraging system of Pharaoh's ant (Monomorium pharaonis). Biosystems 76, 101–112.

Jackson, D., Holcombe, M., Ratnieks, F., 2004b. Trail geometry gives polarity to ant foraging networks. Nature 432, 907–909.

Lipniacki, T., Paszek, P., Brasier, A.R., Luxon, B., Kimmel, M., 2004. Mathematical model of NF-κB regulatory module. J. Theor. Biol. 228, 195–295.

Lodish, H., Berk, A., Zipursky, S.L., Matsudaira, P., Baltimore, D., Darnell, J., 2000. Molecular Cell Biology, 4th ed. W.H. Freeman and Company.

Nelson, D.E., Ihekwaba, A.E.C., Elliott, M., Johnson, J.R., Gibney, C.A., Foreman, B.E., Nelson, G., See, V., Horton, C.A., Spiller, D.G., Edwards, S.W., McDowell, H.P., Unitt, J.F., Sullivan, E., Grimley, R., Benson, N., Broomhead, D., Kell, D.B., White, M.R.H., 2004. Nuclear-cytoplasmic oscillations of a transcription factor are required for sustained gene expression in response to a stimulus. Science 306, 704–708.

Nelson, P., 2004. Biological Physics. W.H. Freeman and Company.

Papin, J.A., Hunter, T., Palsson, B.O., Subramaniam, S., 2005. Reconstruction of cellular signalling networks and analysis of their properties. Mol. Cell Biol. 6, 99–111.

Walker, D.C., Southgate, J., Hill, G., Holcombe, M., Hose, D.R., Wood, S.M., MacNeil, S., Smallwood, R.H., 2004. The epitheliome: modelling the social behaviour of cells. Biosystems 76, 89–100.

Yang, L., Ross, K., Qwanstrom, E.E., 2003. RelA control of IκBα phosphorylation: a positive feedback loop for high affinity NF-κB complexes. Biol. Chem. 278, 30881–30888.