Proceedings of the 2007 Winter Simulation Conference
S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, eds.

MULTISCALE MODELS OF BACTERIAL POPULATIONS

Michael Lees

School of Mechanical, Materials and Manufacturing Engineering
University of Nottingham
Nottingham, U.K.

Brian Logan

School of Computer Science
University of Nottingham
Nottingham, U.K.

John King

School of Mathematical Sciences
University of Nottingham
Nottingham, U.K.

ABSTRACT

We present a hybrid model of the interactions within (multiple-species) populations of bacteria in a developing biofilm which integrates continuum models of population processes (e.g., diffusion of substrates and signalling molecules) with individual-based models of cellular processes (notably growth, division, displacement, and up-regulation). The cell level models combine both aggregated models of continuous processes (growth, division and displacement) for small collections of cells and individual-cell level models of quorum sensing molecule (QSM) sensing, production and up-regulation which encompass both stochastic and discrete processes. The use of both aggregated and individual models of cellular processes allows the resolution of the model to be tailored for a particular modelling problem, while at the same time remaining computationally tractable.

1 INTRODUCTION

In recent years major advances have been made in understanding the gene and signalling networks that control the behaviour of individual cells, and the need to understand the implications of these breakthroughs at the population level is increasingly widely recognised. However, while there has been a significant amount of work on continuum and qualitative (process calculi) models of cell level processes on the one hand and both continuum and individual-based models of population scale effects on the other, there has been relatively little work which attempts to span these scales. Accounting adequately for the relevant subcellular behaviour in a population of millions of distinct, diverse individuals in order to bridge the scales presents significant modelling and simulation challenges but offers the potential for significant benefits for biology and medicine.

In this paper we present a hybrid model of the interactions within (multiple-species) populations of bacteria in a developing biofilm which integrates continuum models of population processes (e.g., diffusion of substrates and signalling molecules) with individual-based models of cellular processes (notably growth, division, displacement, and up-regulation). The cell level models in particular are novel in combining aggregated models of continuous processes (growth, division and displacement) for small collections of cells, and individual-cell level models of quorum sensing molecule (QSM) sensing, production and up-regulation which encompass both stochastic and discrete processes. In contrast to previous work, e.g., Picioreanu, Kreft, and van Loosdrecht (2004), the use of both aggregated and individual models of cellular processes allows the resolution of the model to be tailored for a particular modelling problem, while at the same time remaining computationally tractable. More generally, the approach embodied in our model provides a multiscale framework for modelling interactions between cells, which spans from the cellular level to the population level.

2 BACKGROUND

We focus on multiscale models describing the interactions within (multiple-species) populations of bacteria in a developing 'biofilm'. Biofilms comprise communities of diverse individuals which may interact in both cooperative and non-cooperative fashions. They thus provide comparatively simple (and experimentally relatively well characterised) systems in which variety and selfish, spiteful, 'altruistic' and mutually beneficial behaviours all have scope to flour-

ish as a heterogeneous population develops (for a precise classification of such matters see, for example, West et al. 2006). They are sufficiently complex to exemplify many of the *generic* features of multi-cellular behaviour, without such complexity becoming overwhelming. They are in addition of enormous environmental, industrial and medical importance.

Although relatively simple in biological terms, the interactions within a biofilm span a vast range of spatial scales from sub-cellular to population, with scope for generating a huge variety of emergent behaviour. Cellular inter-relationships, even in single-species populations, are themselves highly complex, with signalling systems, such as quorum sensing, able to lead to coordinated changes in phenotype (see, for example, Winzer, Hardie, and Williams 2002). These quorum-sensing systems, whereby the bacteria communicate to monitor their population size and control their behaviour, are increasingly being understood in terms of the subcellular interactions which govern the production of the relevant signalling molecules. Moreover, mathematical models of these processes are increasingly becoming established and validated, together with those of the corresponding macroscale behaviour (transport of signalling molecules and nutrient, biofilm growth etc.).

Macroscale processes are susceptible to continuum modelling approaches, e.g., systems of (possibly stochastic) differential equations and differential-delay equations. However interactions between individuals, such as signalling between neighbours belonging to distinct bacterial strains, leads to phenomena that cannot readily be captured by traditional multiscale mathematical procedures such as homogenisation. Individual-based models are becoming increasingly widely used in this type of context (see, for example, Kreft, Booth, and Wimpenny 1998; Ginovart, Lopez, and Valls 2002; Krawczyk, Dzwinel, and Yuen 2003; Chang et al. 2003; Picioreanu, Kreft, and van Loosdrecht 2004; Paton et al. 2004). Such models use agent-based simulation techniques to investigate the interactions between (often relatively small) groups of cells and their environment. However, existing work in this area has tended to focus on the emergence of complex organisation in biofilms, with the individual cells being treated as 'black boxes' as far as possible, e.g., the use of cellular-automata approaches to describe biofilm growth.

However, to investigate many phenomena of interest, it is essential that such agent-based models incorporate in an appropriate way information about the macroscale behaviour and their results must in turn be coupled back into the rules adopted in the cell-scale modelling. Accounting adequately for the relevant subcellular behaviour in a population of millions of distinct, diverse individuals in order to bridge the scales presents significant modelling and simulation challenges (for example, to fully investigate the effect of QSM on colony development it is necessary to simulate bacterial colonies consisting of millions or tens of millions of bacteria) but offers the potential for significant benefits for biology and medicine. The issues which arise are not only of great practical significance, but are also representative of some widely-recognised multiscale modelling challenges, relating to how differences in the behaviour of individuals and stochastic effects each manifest themselves at the population scale.

In the next section we present a multiscale biofilm model which integrates continuum models of population scale processes with individual-based models of cellular level processes. It thus attempts to incorporate, albeit in a simplified way, the types of complex signalling pathways and gene networks which are currently the subject of a vast amount of experimental study within a population model. At the same time it seeks to address the multiscale and integrative challenges associated with embedding agent-based models within continuous fields associated with nutrient, quorum-sensing molecules and proliferative pressures.

3 THE MODEL

The model system is similar to that described in Picioreanu, Kreft, and van Loosdrecht (2004), and consists of a 3D 'biofilm reactor' with two compartments, bulk liquid and biofilm (see Figure 1). The bulk liquid compartment contains a (well mixed) solution of S different soluble substrates at constant concentrations. The biofilm compartment contains B different types of biomass which grow on a planar support inside a rectangular box with periodic x and y boundaries. In addition to the biomass and soluble substrates, the biofilm compartment contains a single type of extracellular polysaccharide (EPS) and Q different types of quorum sensing molecule. The biofilm and bulk liquid compartments are in contact and exchange solutes only by diffusion. Substrate and biomass which move beyond the x and y boundaries reappear at the opposite boundary. Bacteria, substrates, QSMs and other material are assumed to be washed away once they reach the z boundary (detachment layer).

In what follows, we focus on the multiscale aspects of the model, i.e., the integration of population level processes with individual-based models of cellular-level processes. A complete description of the model can be found in King, Lees, and Logan (2006).

3.1 Bacterial Particles

For efficiency of computation, individual cells are aggregated into bacterial *particles* as in Picioreanu, Kreft, and van Loosdrecht (2004). Each particle represents one or more cells of a single bacterial strain. The smallest possible particle is equivalent to a single cell; the largest possible particle is that which will fit in a voxel (see section 3.2).

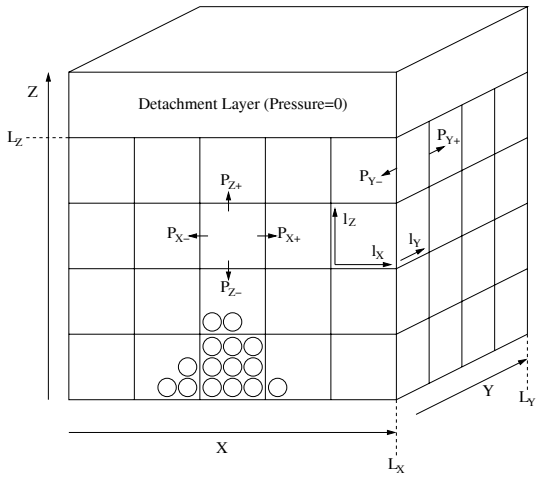

Figure 1: Computational domain.

As a particle consumes substrate its radius increases until it reaches a user-specified maximum particle radius, R, at which point it divides resulting in the creation of a new particle. The maximum size of a particle, i.e., its size at division, thus determines the resolution of the model. Particles of all biomass types (including EPS) are assumed to have the same maximum radius, but their density and hence the number of bacterial cells they represent depends on the biomass type. For strains of bacteria with larger cells the corresponding particles will contain fewer cells; conversely strains with smaller cells will have particles that contain a larger number of cells.

Particles allow the use of aggregate models of continuous processes (growth, division and displacement) for small collections of cells, and also facilitate visualisation of the relative proportions of different biomass types within a voxel. However some processes must be modelled at the level of individual cells. For example, quorum sensing molecule sensing, production and up-regulation encompasses both stochastic and discrete processes, and must be modelled at the level of individual cells. Particles therefore also contain information about the state of the individual cells they represent. Cells within a particle exist in one of two different states: up-regulated and down-regulated. Particles keep track of the number of up-regulated and down-regulated cells they currently contain and cells can change from one state to another at each timestep, in response to the level of QSMs (see section 4.5).

Up-regulated cells produce extracellular polysaccharide (EPS). The EPS produced by cells in biomass particles is aggregated into EPS particles. EPS particles are assumed to have mass M_{EPS}, where M_{EPS} is the mass of a particle of radius R and density ρ_{EPS}. The number of EPS particles at any given point is therefore simply the quantity of EPS mod M_{EPS}.

3.2 Voxels

The biofilm compartment is discretised into sub-compartments or *voxels* containing particles, substrate and signalling molecules. The size of the computational domain (L_X, L_Y, L_Z) is assumed to be an integer multiple of the size of a voxel l_X. Substrate and QSM concentrations are assumed to be uniform across each individual voxel, and the upper bound on the size of a voxel is chosen such that the substrate and QSM concentration values are 'reasonably close' to the continuous values. The size of voxels, l_X, is chosen appropriately for the system to be modelled, with smaller values (criteria being deduced from the corresponding continuum models, i.e., the size of a voxel is taken to be much smaller than the diffusion length determined from the corresponding reaction-diffusion equation) giving greater resolution at increased computational and communication cost. However the voxels are typically fairly large in relation to the size of a cell, e.g. each voxel may contain of order 10^2 particles or 10^4 cells. The bulk liquid compartment is represented by a single point for the purposes of discretisation, and this point is adjacent to all the voxels immediately below the detachment layer.

Each voxel contains zero or more particles of each biomass type (including EPS). The particles in a voxel exert a 'pressure' on the particles in the neighbouring voxels which is a function of the relative number of particles in the voxels, and these pressures are used to displace particles during the division of biomass. Each voxel has six adjacent voxels, connected at each face, which are considered in determining relative pressures, and into which particles may be displaced. Voxels have a pre-determined maximum particle capacity, N, and the pressure in the voxel is considered to be infinite when this maximum is reached. N is calculated using l_X and the maximum radius of a particle, R, assuming simple cubic packing. EPS particles behave in the same way as biomass particles for the purposes of the pressure calculation. The arguments used in developing this pressure model are again based on the continuum modelling, in this case building on multiphase formulations for growing populations such as those described in Byrne et al. (2003).

Each particle has a notional 3D position within its containing voxel which is used for visualisation purposes (see Figure 2). These notional positions are chosen such that the particles do not overlap. The pressure model and maximum particle size are chosen to ensure that there is enough free space in the voxel for this to be possible. Note that the resolution of the model with respect to substrate and QSM concentrations is determined by the size of voxels, l_X. The resolution of the model with respect to the distribution of biomass is also determined by the size of a voxel, in that the mass of each biomass type in each voxel is known. The 3D positions of particles simply make it easier to visualise the distribution of different types of biomass.

At any given point the state of the model is specified by: the amount and distribution of each type of biomass, the number of up- and down-regulated cells of each biomass type, the amount and distribution of extracellular polysaccharide, the concentration and distribution of each type of substrate and quorum sensing molecule, and the pressure distribution.

4 MODEL EVOLUTION

There are three main processes which determine the evolution of the model: the diffusion of substrate and QSM from voxel to voxel, the displacement of particles between voxels in response to proliferative pressures, and changes in the state of the particles themselves in response to the substrate and QSM concentrations in their containing voxel. The transport of substrate and QSM between voxels corresponds precisely to a simple central-difference discretisation of the relevant continuum reaction-diffusion equations. The pressure model underlying particle displacement builds on multiphase formulations for growing populations as described in Byrne et al. (2003). Particles are modelled as agents and implement a simple model of growth and division similar to that in Kreft, Booth, and Wimpenny (1998), and up-regulation (e.g., the production of extracellular polysaccharide) in the presence of QSM (Ward et al. 2003). These processes interact: particles consume substrate and produce QSM, leading to transport associated with the diffusion gradients. Consumption of substrate results in particle growth, which in turn results in increased pressures and particle displacement. Finally, the number of cells in a voxel determines QSM production and hence QSM concentration and up-regulation.

4.1 Diffusion

Diffusion is performed globally over all voxels. The diffusion algorithm iterates over each voxel using the concentration of substrate in the neighbouring voxels to determine the change in concentration at the current voxel. For each substrate, s, the change in concentration at a voxel e as a result of diffusion is:

$$\frac{dc_{s,e}}{dt} = D_s \left(\sum^{i} \Delta c_{s,i} \right) \frac{dt}{(dx)^2} \tag{1}$$

where $c_{s,e}$ is the concentration of substrate s is voxel e, $\sum^i \Delta c_{s,i}$ is the difference in concentration between the voxel e and each of its neighbours ($i = 1, \ldots, 6$) and D_s is the diffusivity of substrate s. The bulk liquid compartment is assumed to have constant concentration and the substratum has the same concentration as the voxels in the bottom layer of voxels. Thus for computational convenience the bulk liquid is taken to be well-stirred (avoiding the need to treat advective and diffusive transport there, together with any boundary layer effects at the fluid/biofilm interface), and negligible nutrient is assumed to be able to penetrate deep into the biofilm. The x and y boundaries are treated as periodic, i.e., the concentration at $(-1, y, z)$ is the same as the concentration at (L_X, y, z). The volume fraction of cells is ignored when doing diffusion calculations—the uptake of substrate by cells (see section 4.2) is much more important.

The substrate concentration at the next timestep, $c'_{s,e}$ can then be calculated as:

$$c'_{s,e} = c_{s,e} + \frac{dc_{s,e}}{dt}. \tag{2}$$

The diffusion algorithm repeatedly computes the concentration in each voxel using equation (2) until the maximum change in concentration for any voxel is less than a predefined constant:

$$\frac{dc_{s,e}}{dt} < \delta c_{max}.$$

At the voxel level, quorum sensing molecules are treated as a substrate, i.e., the QSM concentration $a_{q,e}$, is constant across a voxel and diffuses between voxels. The QSM concentration(s) in a voxel at a given timestep are therefore given by the amount of quorum sensing molecule produced by all particles in the voxel (see section 4.5) and by diffusion of signalling molecule between surrounding voxels. Note that while the evolution of the QSM concentration is determined from the time-dependent problem (cf. 1), it rapidly reaches a quasi-steady state.

4.2 Growth

The growth of each particle is a function of the substrate concentration(s) in the voxel containing the particle. The model of particle growth comprises three separate processes: uptake, metabolism (creation of new biomass) and maintenance (Kreft, Booth, and Wimpenny 1998).

At each time-step each particle consumes an amount of substrate proportional to the concentration of substrate in the voxel and the mass of the particle. The metabolism of substrate into new biomass is given by:

$$\frac{dm_j}{dt} = k_{s,j} Y_{max} \tag{3}$$

where m_j is the mass of the particle j, $k_{s,j}$ is the consumption of substrate s by particle j, and Y_{max} is the yield at the maximum growth rate. The kinetics are taken to be described by the Best equation (Koch 1997). Maintenance is modelled as consumption of biomass, and is proportional to the mass of the particle:

$$-\frac{dm_j}{dt} = m_j g Y_{max} \tag{4}$$

where g is the apparent maintenance rate at zero growth. The overall growth (i.e., change is mass) of a single particle j is therefore:

$$\frac{dm_j}{dt} = (k_{s,j} Y_{max}) - (m_j g Y_{max}). \quad (5)$$

At lower concentrations, particles grows more slowly, and at sufficiently low concentrations they start to shrink (when they are unable to consume sufficient substrate to satisfy their maintenance requirement). At present we do not consider the death of particles: instead particles continually shrink until the uptake and maintenance balance and the cell becomes dormant.

As noted in section 3.1, a single particle represents a collection of cells some of which may be up-regulated and some down-regulated. When the number of cells represented by a particle increases as a result of growth, i.e., when m_j increases by the average mass of a cell of the appropriate biomass type, a new down-regulated cell is "created", by incrementing the number of down-regulated cells in the particle.

4.3 Division

Particle division occurs when the mass of a particle exceeds the user-specified maximum particle mass:

$$m_j > \frac{4}{3}\pi R^3 \rho_b \quad (6)$$

where ρ_b is the density of biomass of type b. At this point, the original particle is split and a new daughter particle is created in the same voxel as the original particle. (The actual 3D position of the particle is not relevant, since it will be adjusted for the purposes of visualisation.) The mass of the daughter particle is randomly chosen between 0.4 and 0.6 of the mass of the particle at division, with the original particle retaining the remainder of the mass. The daughter particle is also allocated (up and down regulated) cells from the original particle in proportion to its mass. The cells transferred from parent to daughter are chosen randomly using a uniform distribution.

4.4 Displacement

As we are not currently considering real positions of particles within voxels (except for visualisation purposes), spreading does not occur within a single voxel. However, spreading and displacement may occur between voxels.

Particles are transferred between voxels if the difference in pressure between the voxels is large enough. The pressure in a voxel e, p_e is given by:

$$p_e = \frac{n_e}{N - n_e} \quad (7)$$

where n_e is the total number of particles of all biomass types (including EPS) in the voxel and N is the maximum number of particles in a voxel at close packing. The pressure in the substratum is assumed to be infinite, and so particles cannot disperse down through the bottom layer. Particle pressure in the bulk liquid is assumed to be zero, so particles can transfer freely into the bulk liquid.

Each type of biomass has a *transfer coefficient*, T_b which specifies how easy it is for biomass of that type to be displaced. The number of particles of biomass of type b to be displaced from a voxel e to the neighbouring voxel e' is then

$$\Delta n_{b,e\to e'} = T_b \times (p_e - p_{e'}) \times (n_{b,e} - n_{b,e'}). \quad (8)$$

The total number of particles of biomass type b to be displaced out of the voxel e is given by

$$\Delta n_{b,e} = \sum^{e'} \Delta n_{b,e\to e'}, \quad (9)$$

and the total number of particles of all biomass types to be displaced out of the voxel e is given by

$$\Delta n_e = \sum_{b=1}^{B} \sum^{e'} \Delta n_{b,e\to e'}. \quad (10)$$

The individual particle(s) of each biomass type to displace are chosen randomly.

T_b must be large enough that particles will displace faster than the maximum particle division rate when the voxel is full and small enough that no more than n_e particles are transferred at any timestep. If T_b is too large, too many particles will be transferred at each timestep. For example, consider a situation where a voxel contains 15 particles of biomass type b and all of its neighbouring voxels are empty. The pressure gradient is equal in all directions and so the same number of particles will be transferred in each direction. However, if $\Delta n_{b,e\to e'} > 3$ there will be insufficient particles to transfer in each direction. Conversely, if T_b is too small, voxels will become over-full. As such, it can be difficult to determine appropriate values of T_b in advance. Moreover, the model outlined above is susceptible to discretisation effects when the number of particles in a voxel is small.

To prevent over sensitivity to values of T_b, we limit the fraction of the particles in a voxel that can be transferred in any given timestep to N_Δ, and set the maximum number of particles that can be transferred to be $\Delta n_{e,max} = \min(\Delta n_e, N_\Delta n_e)$. ($N_\Delta$ is currently 25%.) The

number of particles of each biomass type to transfer, $\Delta n_{b,e}$, are then scaled in proportion $\Delta n_{b,e} = \Delta n_{b,e} \frac{\Delta n_{e,max}}{\Delta n_e}$. Full voxels are also handled specially: if $n_e > N$, the total number of particles to transfer out of the voxel, Δn_e, is increased to be at least $n_e - N$.

The appropriate number of particles of each biomass type are then selected at random for transfer at this timestep. To avoid bias in the direction in which particles are transferred with small values of n_e, the direction in which a particle is displaced is chosen probabilistically, with the probability of transferring a particle to a neighbouring voxel e' being $\Delta p_{e,e'} / \sum^{e_i} \Delta p_{e,e_i}$, where the e_i are the neighbouring voxels of e such that $\Delta p_{e,e_i} > 0$. Particles which are displaced beyond the boundary layer, i.e., L_Z above the substratum, are simply discarded.

4.5 Inter-Cell Signalling

Quorum sensing molecules (QSM) are generated by particles and provide a form of cell to cell communication known as *quorum sensing*. The molecules, which are typically different for each strain of bacteria, control a number of aspects of bacterial growth and development, including bioluminescence, population expansion by swarming, virulence, and the production of extracellular polysaccharides.

The quorum sensing mechanism involves the QSMs triggering increased expression of certain genes in the bacterium. One of the genes codes for the QSM itself, creating a positive feedback loop. The QSM therefore functions as an 'autoinducer', and bacteria will create more of the same QSM when they are surrounded by it. A cell that is in a QSM triggered state is referred to as 'up-regulated', and one that is not is referred to as 'down-regulated'. A QSM can combine with a down-regulated cell to produce an up-regulated cell and an up-regulated cell can spontaneously revert to being down-regulated (by the loss of the bound QSM). A down regulated cell produces a QSM q at a low (basal) rate $Z_{q,d}$. Once the cell becomes up-regulated it produces QSM at a much higher rate $Z_{q,u}$ ($Z_{q,u} >> 100 Z_{q,d}$). The amount of QSM produced by a particle is determined by the relative number of up-regulated and down-regulated cells within the particle.

Other signalling molecules function as inhibitors, which prevent an autoinducer combining with a cell, or prevent up-regulation when a cell combines with the autoinducer. Inhibitors therefore restrict the production of QSM by the bacteria. Inhibitors are particularly interesting from a biological point of view as they allow control of up-regulation and hence development of the bacterial colony. The relative ease with which a QSM q and the corresponding inhibitor $\bar{q}$ can combine with a cell is denoted by γ_q and $\gamma_{\bar{q}}$. In what follows, for simplicity we assume that $\gamma_q = \gamma_{\bar{q}} = \gamma$. In the presence of inhibitor, the probability of a cell changing from down to up-regulated is given by:

$$P(up) = \alpha \frac{a_{q,e}}{1 + (\gamma(a_{q,e} + a_{\bar{q},e}))} dt \tag{11}$$

where $a_{q,e}$ is the concentration of QSM q in voxel e, $a_{\bar{q},e}$ is the concentration of inhibitor $\bar{q}$ and α is the conversion rate of down-regulated cells to up-regulated cells due to QSM binding. A cell reverts from up-regulated to down-regulated with probability:

$$P(down) = \beta \frac{1 + (\gamma a_{\bar{q},e})}{1 + (\gamma(a_{q,e} + a_{\bar{q},e}))} dt \tag{12}$$

where β is the spontaneous down-regulation rate.

Up-regulated cells produce extracellular polysaccharide (EPS) at a constant rate and particles therefore produce EPS at a rate proportional to the number of up-regulated cells they currently contain. The EPS produced is aggregated into new EPS particles. Whenever the amount of EPS in the voxel increases by the maximum mass of a particle, a new EPS particle is created in the voxel. New EPS particles behave in the same way as other particle types for the purposes of pressure calculation, and like biomass particles can be displaced into surrounding voxels.

5 IMPLEMENTATION

To facilitate distribution, the implementation of the model is decomposed into a 'diffusion module' (which handles diffusion calculations) and one or more 'model regions'. Each model region processes one or more voxels, and handles the growth and division of particles within voxels, and the displacement of particles between voxels. In addition there is a visualisation module for run-time monitoring of the simulation progress and post-simulation analysis of results (see Figure 2). The model region and visualisation modules are implemented using the Mason agent toolkit (Luke et al. 2005). The diffusion module is written in Java. Interaction between modules can be by means of procedure calls, HLA (IEEE 2000) (RTI) calls, or Grid invocations. The HLA distribution uses the DMSO RTI 1.3NGv6 Java bindings, and the Grid distribution is based on HLA_GRID (Xie et al. 2005), which allows HLA-compliant simulators to be instantiated and linked using Grid services. In the interests of simplicity, we focus here on the non-distributed case in which the modules communicate via procedure calls. We assume that each model region processes a single voxel, and use the terms voxel and model region interchangeably. (In reality, model regions typically process more than one voxel, and facilitate distribution by mapping 'non-local' voxel to voxel communication into HLA or Grid invocations.) See Lees, Logan, and King (2007) for an overview of the model distribution.

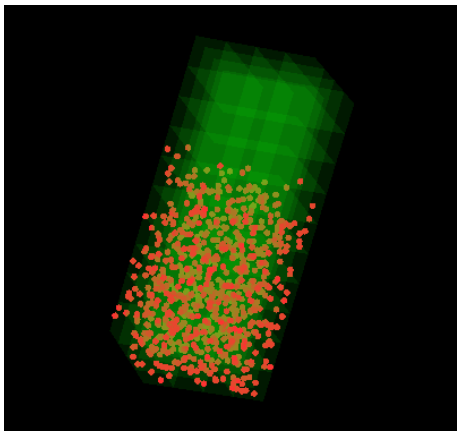

Figure 2: Simulation visualisation.

5.1 Model Timestep

In this section we give a high level description of a single iteration of one timestep of the model. This illustrates the operation of both the diffusion module and the voxels and how they interact to update the state of the model. The state of a voxel at each timestep is given by the number of particles of each biomass type, the concentration of each substrate, the consumption of each substrate at this timestep, the concentration of each quorum sensing molecule, and the production of each QSM at this timestep.

The processing of the voxels at timestep t occurs in two phases. The first involves the execution of the voxels to calculate the consumption of substrate by particles, particle growth and the number of particles following division of the biomass. Processing of phase one within each voxel itself occurs in three steps. Firstly, the growth step increases the mass of each particle within each voxel given the concentration of substrate for this timestep in the voxel. (For $t = 0$, the concentrations and number of particles are taken as parameters of the simulation.) This also gives the total consumption of all substrates by all particles in this voxel for this timestep. The second step computes the production of QSMs by each particle in the voxel. The third step is particle division: each particle which reached the maximum allowable mass during the growth step is split into two particles, increasing the number of particles in the voxel. Each voxel then sends the consumption of each type of substrate and the amount of QSM produced by its particles to the diffusion module. At the same time, each voxel sends its current particle counts to each of its neighbouring voxels.

The second phase of the timestep involves executing the diffusion module once it has received substrate and QSM information from all voxels. The diffusion module uses the substrate consumption for this timestep together with the diffusion algorithm to calculate the new substrate and QSM concentrations for the next timestep. The diffusion module then sends each voxel the substrate and QSM concentrations for the next, $t+1$, timestep.

In parallel with the execution of the diffusion module, each voxel also executes a displacement step, which uses the difference in pressure between the voxel and each of its adjacent voxels (which each voxel calculates using the number of particles in each of its neighbouring voxels) to determine displacement of particles between voxels. The (possibly empty) list of displaced particles is then sent to each of the neighbouring voxels. A snapshot of the state of a particle for migration purposes consists of its biomass type, mass, the number of up-regulated cells, the voxel from which the particle is being migrated and the direction in which it is being migrated. Once each voxel has received a list of transfer particles from all its neighbours the voxel updates its particle counts for the next timestep.

The timestep is then incremented and the cycle repeats with the voxels using the newly calculated concentrations and particle counts. While the scheme could be made significantly more efficient by exploiting the disparities in timescales between nutrient transport and growth (for example), for the purposes of this study we have for simplicity chosen to place all the effects on a similar footing. The resulting inefficiency is to some extent offset by the flexibility of the code (in particular allowing additional effects readily to be incorporated without reconsidering the simulation approach).

6 PRELIMINARY RESULTS

In this section we present some preliminary experimental results from the current version of the model. Our initial experiments have focused on the role of QSM in the development of bacterial colonies, and in particular how QSM inhibitors can be used to prevent the up-regulation of cells within a colony (Koerber et al. 2002).

To allow comparison with an existing analytic model, our experiments used a single species of bacteria, a single substrate, and two types of signalling molecules. We stress that the numbers of biomass types, substrates and QSMs were chosen purely for experimental convenience and do not reflect limitations of the underlying simulator. The parameter values for the bacterial growth and division models were taken from Kreft, Booth, and Wimpenny (1998) and the parameters for the inter-cellular signalling and up-regulation model from Koerber et al. (2002). The experiments used a $17 \times 17 \times 17$ micron voxel as in Picioreanu, Kreft, and van Loosdrecht (2004), a maximum cell radius of 0.756 microns, and a maximum particle radius of $5\times$ the radius of a cell, giving up to 1371 cells per voxel. The total size of the model was $68 \times 68 \times 1700$ microns, i.e., 4 voxels in the x and y dimensions and 100 voxels in the z dimension. (The x and y dimensions are periodic, so the model only needs to be large enough in x and y to allow reasonable spreading of

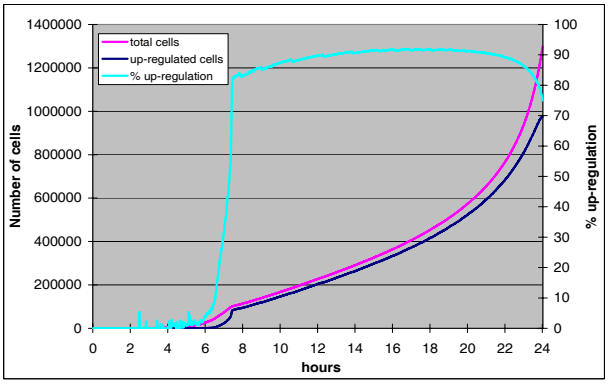

Figure 3: Up-regulation of the colony.

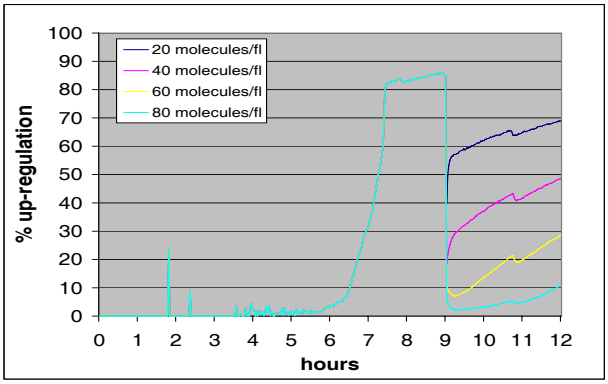

Figure 4: Effect of inhibitor on up-regulation.

particles.) The model timestep was 0.1 seconds (determined by diffusion requirements), and the simulations ran for 24 hours of simulated time, i.e., 864,000 timesteps.

The initial population consisted of 80 particles (each representing a single cell) randomly distributed on the planar support. The maximum number of particles during the simulation was 16,000 representing approximately 1.3 million bacterial cells (see below). The runs required about 40 hours elapsed time on a single processor (P4, 2.4GHz) with 2GB of physical memory.

Figure 3 shows the total number of cells, the number of up-regulated cells and the proportion of up-regulated cells over the course of the simulation for a model system without inhibitor. As can be seen there is a marked increase in up-regulation after about 7 hours, rising from less than 10% up-regulated cells at 6 hours to about 85% up-regulated cells at 8 hours. At this point, the apparent threshold concentration of QSM is reached (reflecting high bacterial density adjacent to the planar support), triggering positive feedback in the production of QSM and a corresponding sharp rise in QSM concentration (not shown). These results are qualitatively similar to those for the continuum model presented in Ward et al. (2001).

Once the critical cell density is reached, the number of up-regulated cells increases exponentially (in line with the total number of cells) until the biofilm approaches the detachment layer at the top of the model. During this period, the proportion of up-regulated cells continues to climb slowly, reaching a peak at about 18 hours, before declining to about 70% up-regulation at 24 hours. At 24 hours, all the voxels below the detachment layer contain particles, and a significant fraction of the new particles produced are immediately washed away. In the region immediately below the detachment layer, QSM concentration is low (as the molecules are continuously washed away at the detachment layer), with the result that fewer cells in this region are up-regulated. The percentage of cells which are up-regulated therefore declines as the biofilm approaches the detachment layer. Moreover, towards the end of the run, the consumption of substrate in the biofilm below the detachment layer reduces the diffusion of substrate to the particles in the lower layers. As a result, the particles in the lower layers start to shrink, as they are unable to obtain enough substrate to meet their maintenance requirements. This reduces the number of cells in each particle, and, as QSM production is determined by the number of cells rather than the number of particles, QSM production in the lower layers of the biofilm declines, resulting in a lower overall concentration of QSM.

Figure 4 shows the effect of introducing a signalling molecule that functions as an inhibitor. Different concentrations of inhibitor were added to the bulk liquid after 9 hours of simulation time, i.e., after up-regulation of the colony has occurred. As expected, the highest concentration of inhibitor (80 molecules fl^{-1}) has the greatest effect on up-regulation, reducing the percentage of up-regulated cells from about 85% to less than 5%. Lower concentrations, have a less marked effect. For example, at a concentration of 20 molecules fl^{-1}, the inhibitor only reduces up-regulation by about 40%.

The results in Figure 4 clearly indicate that inhibition of up-regulation is possible in our example population. However, in therapeutic applications, the amount of inhibitor used (corresponding to e.g., an antibiotic) is of critical importance. In future work, we plan to investigate the relationship between amount of inhibitor required and the time at which it is introduced, and the maximum thickness of the biofilm thickness which allows inhibition of up-regulation.

7 SUMMARY

We have presented a generic multiscale framework for modelling populations of cells, which spans from the cellular level to the population level. Our model integrates population level processes (e.g., diffusion of substrates and signalling molecules) with individual-based models of cellular-level processes (notably growth, division, displacement, and up-regulation). Our approach is individual-based in the sense that the model tracks the state of individual cells and ag-

gregations of cells (particles) over time, and in utilising multiagent simulation tools and techniques as the basis of the simulator implementation. In contrast to previous work, e.g., Picioreanu, Kreft, and van Loosdrecht (2004), it incorporates both aggregated and individual models of cellular processes, allowing the resolution of the model to be tailored for a particular modelling problem, while at the same time remaining computationally tractable. For example, in experiments to investigate the effect of QSM inhibitor on the up-regulation of the population, we have successfully simulated models containing 10^6 cells (10^4 particles) in near real time on a single processor.

The model aims extend the state of the art in biofilm modelling by providing support for hypothesis testing, e.g., "What would happen if we starved bacteria of this strain or in this region of nutrient?" Such predictions are still qualitative, but could be tested in the lab. While the cell models used in the current prototype are somewhat simplistic, the approach provides a generic framework into which different types of cells and more complex models of gene expression and signalling pathways (which are currently the subject of a vast amount of experimental study) can be plugged. More detailed cell level models are likely to impact simulation performance, but we believe this can be addressed through the support for distributed simulation are already incorporated in the model.

ACKNOWLEDGMENTS

This work was supported by BBSRC project number BB/D006619/1 and by EPSRC projects GR/S82862/01, EP/C549406/1 and EP/C549414/1.

REFERENCES

Byrne, H. M., J. R. King, D. L. S. McElwain, and L. Preziosi. 2003. A two-phase model of solid tumour growth. *Applied Mathematics Letters* 16 (4): 567–573.

Chang, I., E. S. Gilbert, N. Eliashberg, and J. D. Keasling. 2003. A three-dimensional, stochastic simulation of biofilm growth and transport-related factors that affect structure. *Microbiology* 149:2859–2871.

Ginovart, M., D. Lopez, and J. Valls. 2002. INDISIM, an individual-based discrete simulation model to study bacterial cultures. *Journal of Theoretical Biology* 214:305–319.

IEEE 2000. IEEE Standard for modeling and simulation (M&S) High Level Architecture (HLA) — Framework and rules. IEEE. (IEEE Standard No.: 1516-2000).

King, J., M. Lees, and B. Logan. 2006. Agent-based and continuum modelling of populations of cells. Technical report, University of Nottingham.

Koch, A. L. 1997, September. Microbial physiology and ecology of slow growth. *Microbiology and Molecular Biology Reviews* 61 (3): 305–318.

Koerber, A. J., J. R. King, J. P. Ward, P. Williams, J. M. Croft, and R. E. Sockett. 2002. A mathematical model of partial-thickness burn-wound infection by *Pseudomonas aeruginosa*: Quorum sensing and the build-up to invasion. *Bulletin of Mathematical Biology* 64:239–259.

Krawczyk, K., W. Dzwinel, and D. A. Yuen. 2003. Nonlinear development of bacterial colony modeled with cellular automata and agent objects. *International Journal of Modern Physics C* 14 (10): 1385–1404.

Kreft, J.-U., G. Booth, and J. W. T. Wimpenny. 1998. BacSim, a simulator for individual-based modelling of bacterial colony growth. *Microbiology* 144:3275–3287.

Lees, M., B. Logan, and J. King. 2007. The architecture and implementation of the BacGrid simulator. Technical report, University of Nottingham.

Luke, S., C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan. 2005. MASON: A multiagent simulation environment. *Simulation* 81 (7): 517–527.

Paton, R., R. Gregory, C. Vlachos, J. Saunders, and H. Wu. 2004. Evolvable social agents for bacterial systems modeling. *IEEE Transactions on Nanobioscience* 3 (3): 208–216.

Picioreanu, C., J.-U. Kreft, and M. C. M. van Loosdrecht. 2004. Particle-based multidimensional multispecies biofilm model. *Applied and Environmental Microbiology* 70 (5): 3024–3040.

Ward, J. P., J. R. King, A. J. Koerber, J. M. Croft, R. E. Sockett, and P. Williams. 2003. Early development and quorum sensing in bacterial biofilms. *Journal of Mathematical Biology* (47): 23–55.

Ward, J. P., J. R. King, A. J. Koerber, P. Williams, J. M. Croft, and R. E. Sockett. 2001. Mathematical modelling of quorum sensing in bacteria. *Mathematical Medicine and Biology* 18 (3): 263–292.

West, S. A., A. S. Griffin, A. Gardner, and S. P. Diggle. 2006. Social evolution theory for microorganisms. *Nature Reviews Microbiology* 4 (8): 597–607.

Winzer, K., K. R. Hardie, and P. Williams. 2002. Bacterial cell-to-cell communication: sorry, can't talk now—gone to lunch! *Current Opinion in Microbiology* 5 (2): 216–222.

Xie, Y., Y. M. Teo, W. Cai, and S. Turner. 2005. Service provisioning for HLA-based distributed simulation on the Grid. In *Proceedings of the Nineteenth ACM/IEEE/SCS Workshop on Principles of Advanced and Distributed Simulation (PADS 2005)*, 282–291. Monterey, CA, USA.

AUTHOR BIOGRAPHIES

MICHAEL LEES is a postdoctoral research fellow currently in the School of Mechanical, Materials and Manufacturing Engineering at the University of Nottingham. His research interests are in the area of adaptive optimistic synchronisation of large scale simulations of multiagent systems.

BRIAN LOGAN is a lecturer in the School of Computer Science at the University of Nottingham, UK. His research interests are in the area of agent-based systems, including the simulation of multiagent systems and applications of agent-based simulation.

JOHN KING is a professor and head of the Division of Applied Mathematics in the School of Mathematical Sciences at the University of Nottingham. His research interests are in the area of multiscale biological modelling and associated mathematical techniques, including intra- and inter-cellular signalling, such as the role of quorum sensing in biofilm development and in infections and their treatment.