

Spatial self-organization in a cyclic resource–species model

Stephen M. Krone*, Yongtao Guan

Department of Mathematics, University of Idaho, Moscow, ID 83844-1103, USA

Received 15 March 2005; received in revised form 24 October 2005; accepted 2 November 2005

Available online 19 December 2005

Abstract

Biological communities are remarkable in their ability to form cooperative ensembles that lead to coexistence through various types of niche partitioning, usually intimately tied to spatial structure. This is especially true in microbial settings where differential expression and regulation of genes allows members of a given species to alter their lifestyle so as to fill a functional role within the community. The resulting species interactions can involve feedback, as in the case of some bacterial consortia that participate in the cooperative degradation of a given resource in a succession of steps and in such a way that certain “later” species provide catalytic support for the primary degrader. We seek to capture the essential features of such spatially extended biological systems by introducing a lattice-based stochastic spatial model (interacting particle system) with cyclic local dynamics. Here, a given site progresses through a sequence of resource and species states in a prescribed order. Furthermore, this succession of states (at a site) is assumed to form a cyclic pattern due to a natural feedback mechanism. We explore conditions under which all the species are able to coexist and consider the extent to which this coexistence requires the development of spatio-temporal patterns, including spiral waves. This self-organization, if it occurs, results when synchronization of the dynamics at the microscopic level leads to macroscopic patterns. These patterns result in consumer-driven resource fluctuations that generate a form of spatio-temporal niche partitioning. As with most models of this complexity, we employ a mixture of mathematical analysis and simulations to develop an understanding of the resulting dynamics.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Stochastic spatial model; Interacting particle system; CA; Resources; Microbial diversity; Cyclic; Spiral waves

1. Introduction

Mathematical models and observations from macroscopic ecology tell us that spatial structure can have an enormous impact on the stability and diversity of interacting species. Bacteria typically live in structured environments such as biofilms and soils and, as in the case of biofilms, this structure is often largely the work of the bacteria themselves. In addition to this spatial self-organization, microbial communities often play an active role in helping to structure their geochemical environments.

Through their phenomenal diversity and adaptability, populations in microbial communities find ways to cooperate to efficiently consume nutrients. Pelz et al. (1999) and Bradshaw et al. (1994) give examples of such communities that, as biological units, are robust and able

to extract the maximal energetic benefit from available nutrients. This cooperativity can work in concert with the afore-mentioned spatial organization.

We consider, as a representative of the class of spatially explicit cyclic resource–species dynamics to be treated here, a microbial scenario in which a number of different bacterial species coexist by cooperatively degrading a single nutrient in a succession of steps and providing catalytic support for the primary degrader. This “catalytic support,” which can be anything that facilitates the growth of the primary degrader, introduces feedback into the system and provides a mechanism for spatial self-organization that can give rise to coexistence of the species. Ours is a stochastic spatial model that is individual based with discrete spatial structure and continuous time. Such models fall under the heading of interacting particle systems or (asynchronously updated) probabilistic cellular automata.

Several factors make it a challenge to observe and understand the complex dynamics in microbial communities, and

*Corresponding author. Tel.: +1 208 885 6317; fax: +1 208 885 5843.

E-mail address: krone@uidaho.edu (S.M. Krone).

this is where mathematical models can lead to significant advances. Pelz et al. (1999) point out that, while metabolic pathways in individual microorganisms have been well studied since 1950, very little is known about multi-species metabolic networks in bacterial communities, especially those found in nature. In addition, the spatio-temporal dynamics of microbial communities are notoriously difficult to observe in any detail. This is partly due to the fact that most sampling protocols are invasive and destructive. In the process of trying to extract crude estimates of the diversity of a sample, the microscopic spatial structure is typically broken down (by homogenizing the sample) and one obtains instead only a spatially averaged view, often at a single time point. Theoretical advances can play an important role in determining the types of species interactions that lead to stable coexistence in a dynamic spatial setting, and they can reveal signatures of such interactions that make them more easily recognizable. For example, microscopic mechanisms sometimes lead to macroscopic patterns. While the latter do not uniquely determine the former, they often suggest hypotheses that can be tested in the lab.

The complex spatial nonlinear dynamics described by our model lead to interesting properties. We are particularly concerned with conditions leading to persistence of the species and the emergence of macroscopic spatio-temporal patterns from the microscopic dynamics. Our simulations imply, for example, that for certain parameter settings the species cannot coexist in the absence of spatial patterning. In these instances, coexistence is enhanced by the self-organization of species (and their resources). The fact that the resources can become patterned in such a process suggests potential mechanisms for detecting spatial structure in microbial communities.

2. The model

To motivate the model, let us imagine a collection of bacterial species that take part in the successive degradation of a single nutrient. Assume that these bacterial players live on a surface, say at a rock–water interface. The compound to be degraded adsorbs onto the surface from the bulk fluid. One species (the primary degrader) is able to break down this initial compound, taking in nutrients in the process. The product left over is then further broken down by a second species, leaving a second product, and so on. This sequence of steps is known as cross-feeding (Rosenzweig et al., 1994; Doebeli, 2002). The last species in the sequence we are modeling, after taking up its share of the nutrient, then provides some form of catalytic support that enhances the growth of the first species. This last step in the dynamics causes feedback and thus results in a cyclic progression of states at each individual site. Such feedback mechanisms can arise in many ways. The last species, for example, could remove a toxin produced by any one of the previous species that would otherwise kill or inhibit the growth of the first species. Pelz et al. (1999) discuss such a

situation; they study a microbial consortium with four primary species that are involved in carbon sharing and with one of the species scavenging toxic metabolites that would kill the primary degrader if they were allowed to accumulate. While their system does not constitute an exact match to our model, it gives a nice example of the type of feedback we are modeling. As an alternative mechanism producing feedback, the last species might simply create conditions under which the primary nutrient can adsorb from the bulk fluid.

The above scenario is meant to provide some intuition for the model which follows. There are many plausible systems that will provide the same dynamics. Our goal is to understand how such spatially explicit cyclic systems behave and to give some indication of the stability properties that make them likely to arise in the first place.

The stochastic spatial model introduced here is constructed on a two-dimensional grid (the square lattice \mathbb{Z}^2 or a finite sublattice), with periodic boundary conditions imposed when the space is finite. Each site of the lattice can be in any one of the following states:

$$R_1, S_1, R_2, S_2, \dots, R_L, S_L,$$

where L denotes the number of species, S_i is the i th species, and R_i is the i th resource (or stage in the decomposition of the initial compound). If a site is occupied by an individual from species S_i , this individual will reproduce at rate β_i and deposit its (single) offspring onto a randomly chosen neighboring site; only if the offspring lands on a site in state R_i will it be viable. In such an event, the neighboring site will harbor the i th species on the i th resource; this, however, will be denoted by S_i instead of the bulkier $R_i + S_i$ (and we will simply say that the site is occupied by the i th species, greatly simplifying our explanations later). Thus, species S_i consumes or degrades its “required resource” R_i to produce its “product resource” R_{i+1} , the subscript being understood cyclically (i.e. $R_{L+1} = R_1$). To keep the model from having too many states, we assume that a species will disappear from a site once it has converted its required resource to its product resource. (This could be due to the removal of the species as a result of starvation or being out-competed by the next species.) Thus, if a site is in state S_i , it will change to state R_{i+1} at constant rate δ_i .

In summary, a given site proceeds through a cyclic progression of states

$$R_1 \rightarrow S_1 \rightarrow R_2 \rightarrow S_2 \rightarrow \dots \rightarrow R_L \rightarrow S_L \rightarrow R_1 \rightarrow \dots$$

The transitions at such a site x come in two flavors:

$$R_i \rightarrow S_i \quad \text{at rate } \beta_i f_i(x) \quad (1)$$

and

$$S_i \rightarrow R_{i+1} \quad \text{at rate } \delta_i, \quad (2)$$

where $f_i(x)$ denotes the fraction of species S_i in some neighborhood about the site x . We will refer to β_i and δ_i as the growth and death rates for species S_i .




Fig. 1. The cyclic local dynamics at site x for the model with three species and three resources.

The above transitions and the cyclic nature of the progression of states at a given site are illustrated in Fig. 1. The sizes of the neighborhoods are important and will be discussed below. We say the transition in (1) is “by contact” since it depends on the states at nearby sites, requiring the presence of at least one S_i near site x . The transition in (2) is “spontaneous” in that it happens automatically, independently of the states at the surrounding sites. We remark that since ours is a stochastic model, the word “rate” always refers to exponential waiting times; i.e. some event happens at rate λ if the amount of time it takes is an exponential random variable with rate λ (mean $1/\lambda$). Fig. 1 illustrates the transition rates at a given site. This is an example of what we call cyclic local dynamics, our model being additionally characterized by the alternating of contact and spontaneous transitions. Note that the model reduces to the basic contact process (cf. Durrett, 1988) in the case of one species and one resource ($L = 1$).

To complete the description of the contact transition rates above, we must specify the neighborhoods used to calculate the fraction of neighboring sites in a given state. In our simulations, we treated each of the following three neighborhood types for a site x :

- 4 nearest neighbors: the four adjacent sites north, south, east, and west of the site x ;
- 8 nearest neighbors: the sites in the 3×3 box centered at x (except for x) made up of the four sites above plus the four diagonally adjacent sites;
- 24 neighbors: the sites in the 5×5 box centered at x (except for x).

Otherwise put, the 4 nearest neighbors are the sites within L^1 distance one of x (i.e. the von Neumann neighborhood of x), the 8 nearest neighbors are the sites within L^∞ distance one of x (i.e. the Moore neighborhood of x), and 24 neighbors correspond to the sites within L^∞ distance two of x .

Now that the local dynamics are specified, the goal is to understand how these translate into global behavior by keeping track of the entire configuration of states on the spatial grid. Such stochastic spatial models are sometimes called *interacting particle systems* (IPS) or (asynchronously updated) stochastic *cellular automata* (CA). While the latter term is more familiar to biologists, some confusion can occur if one fails to distinguish between the different

types of CA model. Hereafter, we will use the term IPS model.

3. Particle system behavior

In this section, we first characterize the different patterns and behaviors of our model. Then we investigate the corresponding parameter regions that yield these different behaviors through extensive simulations of the particle system.

The model exhibits configurations that can be classified by four major patterns: frozen, synchronous waves, asynchronous waves, and mixing. These patterns are shown in Fig. 2 and are understood to refer to patterns in a simulation that persist for a long time. A *frozen* pattern results when all the species die out, leaving only resources; once all species are absent, the configuration can no longer change. The other three types of patterns involve survival of all species. *Synchronous waves* include spiral waves and traveling waves that are spatially large (compared to the grid size) and temporally synchronized. *Asynchronous waves* consist of transient spiral waves and traveling waves that are not large (compared to the grid size) and temporally unsynchronized. *Mixing* refers to spatially homogeneous configurations and corresponds to species survival in the absence of spatial patterning. The classification of simulation runs into the four patterns was based on the coefficient of variation (cf. Hassell and May, 1988), as described below.

We remark that the difference between synchronous and asynchronous waves is one of degree, with asynchronous waves amounting to synchronous waves that exist only on small length scales. This coincides with our understanding based on the eigenvalue criterion given in Section 4.

All simulations were run on a lattice of size 300×300 . Each run started with a random initial configuration and data were collected after 2000 units of time, where one “unit of time” corresponds to an average of 90,000 (i.e. $300 \cdot 300$) possible random events: on average, each site has the opportunity to update once during a unit of time. To fully investigate the parameter space for this model is difficult, if not impossible, because of the large number of parameters ($2L$ different parameters when there are L species). However, we can obtain important information by investigating *two special cases*. One is the *symmetric case* in which all species growth rates are equal and all death rates equal 1. The other is an *asymmetric case* in which we allow the growth rate of the first species to differ from the (majority) common growth rate of all other species. These cases are discussed in more detail below.

The afore-mentioned patterns should not be confused with short-lived transient patterns such as traveling waves of infection in a contact process or epidemic model. Once such a wave passes, the system relaxes to a steady state distribution that does not accommodate these waves. In our cyclic resource–species model, the patterns we refer to are recurring. Of course, any frozen state is absorbing (in




Fig. 2. Four different patterns for an 8-state version of the model. Top left: frozen; top right: synchronous waves; bottom left: asynchronous waves; bottom right: mixing. The colors designate species (green, red, light blue, blue) and resources (yellow, gray, purple, brown). For example, yellow represents the product resource of the green species and the required resource of the red species.

the sense that it can never change) and as soon as one species dies the whole system is doomed to freeze. Since we are using a finite lattice, eventually the process will freeze, but this will take a very long time on our 300×300 lattice—much longer than the 2000 units of time on which our simulation results are based. The patterns we are referring to appear in simulations over the course of a reasonable length of time (from a few seconds to many hours). Thus, strictly speaking, they refer to quasi-stationary distributions for the particle system on a finite lattice. If one prefers to think of the particle system on an infinite lattice, these quasi-stationary states look like pieces of the stationary distribution for the infinite system.

To summarize the patterns as a function of parameter settings, the coefficient of variation was calculated for each given run at time step 2000 by partitioning the 300×300 lattice into 400 non-overlapping squares of size 15×15 , and computing the frequencies of each state in each such square. Based on these values, we computed a sample mean, μ_i , and a sample standard deviation, σ_i , for each state i . The coefficient of variation, $CV_i = \sigma_i/\mu_i$, for each state i was then plotted as a function of the parameter values (data not shown). On such a plot, the CV values corresponding to the different species were tightly clustered (due to the narrow

width of species waves), so we assigned patterns based on CV values of the first species according to

- frozen: $CV = 0$
- mixing: $0 < CV \leq 0.5$
- async waves: $0.5 < CV \leq 1$
- sync waves: $CV > 1$.

These data are summarized in Figs. 3 and 4, where for each parameter setting we ran the simulation with different random seeds three times and the final pattern was decided by “majority rule.” The thresholds were chosen so that the mapping of CV value to pattern conformed to visual observations. Changing the thresholds by a small amount results in a slight displacement of the boundaries, but the overall effect is the same. It should be noted that, while the CV criterion provides a systematic method for assigning a pattern to a given spatial configuration, it is no substitute for direct observation of the simulation. For example, when species densities are very low and hence synchronous waves are not possible, the CV can be unduly influenced by the small mean. This could result in the improper classification of an asynchronous wave as a synchronous wave if the threshold for synchronous waves is not sufficiently high.

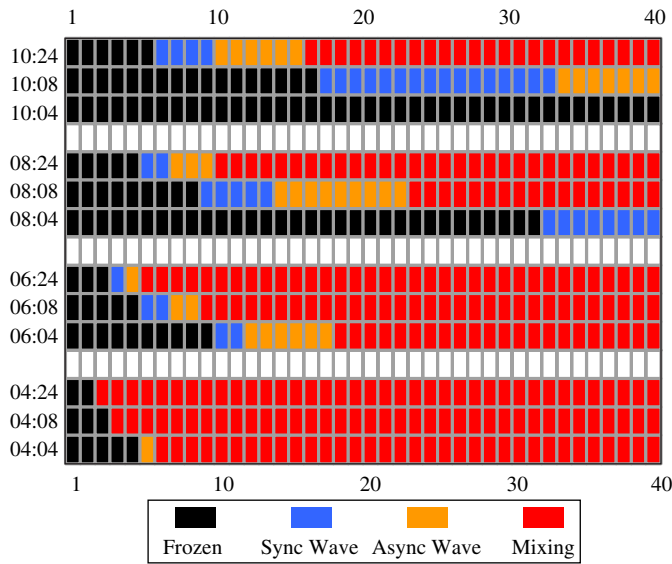


Fig. 3. (Symmetric case) Simulation results for IPS model with all species having the same growth rate β , and all death rates $\delta = 1$. The horizontal axis denotes the common growth rate β (sampled at integer values) and the vertical axis has labels of the form $N:r$, where $N = 2L$ denotes the total number of states and r denotes the neighborhood size. The outcome reported for a given values of β and $N:r$ corresponds to the majority outcome from 3 runs starting with different random configurations.

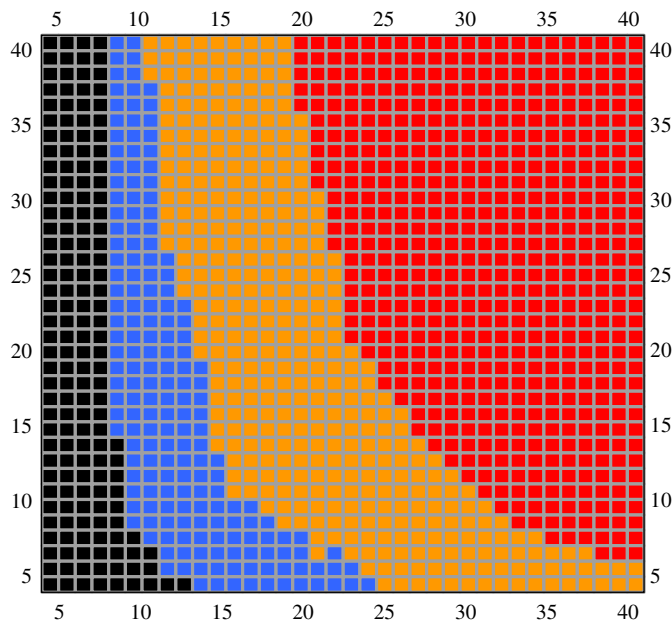


Fig. 4. (Asymmetric case) Simulation results for IPS model with all species having the same death rate $\delta = 1$, but the growth rate of species 1 is (possibly) different from the common growth rate of species $2, \dots, L$. The horizontal axis denotes the majority growth rate β and the vertical axis denotes the growth rate β_1 of the first species. The number of states is $2L = 8$ and the neighborhood size is also 8. Each small square in the figure corresponds to the majority outcome from 3 runs starting with different random configurations and the growth rates assuming integer values. The color scheme is the same as in Fig. 3.

3.1. Symmetric case

The results of the simulations in the symmetric case are summarized in Fig. 3. We treated four different values for the number of states: $N = 2L = 4, 6, 8, 10$. For each of these, we considered three neighborhood sizes: $r = 4, 8, 24$. Once these two parameters were chosen, we ran simulations for a large range of values of the common growth rate β for 2000 units of time.

Growth rate. First consider the simulation results as a function of β , for fixed values of N and r . Small values of β lead to frozen patterns. As β increases, the patterns transform first to synchronous waves (in certain cases), then asynchronous waves, and, for large β , mixing.

Neighborhood size. To see the effects of neighborhood size, compare the three bands corresponding to each fixed value of N . First of all, larger neighborhood size leads to species survival (and hence coexistence) for a wider range of β values. This is a simple consequence of the fact that a larger range provides a given species a better chance to find its required resource. Also, larger values of r create more mixing of different types; this leads to less spatial structure in the sense that the wave clusters are smaller when r is larger. We will see below that this reduction in the amount of spatial structure can happen for other reasons, as well.

Number of states. The most striking observation about the effect of the number of states is that for $N = 4$ (i.e. two species) there are no synchronous waves and hence no spirals. In fact, the same is also true when $N = 2$ since, in that case, our model reduces to the single species contact process. Synchronous wave patterns are possible for all values of $N > 4$, provided β and r are suitably chosen. Similar observations have been made regarding catalytic hypercycle models where it is known that spiral patterns arise if and only if the number of states in the cycle is greater than or equal to 5; see Dieckmann et al. (2000). Durrett and Griffeath (1993) also observed spiral waves for a suitably large number of states in two different cyclic spatial models. One is the Greenberg–Hastings model (one contact transition and all others spontaneous), and the other can be thought of as a generalized rock–scissors–paper model (all transitions are by contact); in each case, they use threshold updating for the contact transitions—a feature that encourages spiral formation. Bramson and Griffeath (1989) discovered another phase transition involving the number of states for a one-dimensional version of the generalized rock–scissors–paper model; they showed that the system freezes when the number of states exceeds four, and otherwise keeps changing. Other interesting aspects relating to spiral formation in cyclic spatial systems can be found in the references in the above papers. Of course, for our model, the number of states is always even. In the next section, we will give a simple explanation for our $N > 4$ threshold based on a linearization argument for a related system of reaction–diffusion equations (RDEs). As the number of states increases, one needs larger values of β to generate synchronous waves and to survive in general, especially when the neighborhood size is not large.

Initial configurations. For the above discussion, it was assumed that the initial configuration was random, with each site independently assigned one of the N states and probability $1/N$ per state. Thus any patterns that formed, such as spirals, were the result of self-organization. In some cases, parameters that led to species extinction, and hence a frozen pattern, could produce species persistence and spiral waves under appropriate initial configurations. This happens for β 's in the frozen region that are not too far from the synchronous wave region. For such β , once enough spatial structure is present it can be maintained. There are several ways to get such “hard-to-start spirals.” If we start with a random initial configuration, the species all die out, but before doing so the remaining resources will have become somewhat clustered; if we then “seed” the frozen configuration by randomly sprinkling in a small number of each species to the system, spiral waves will form. Such behavior has been observed in other cyclic systems (Comins et al., 1992). Often, only a very small amount of this pre-structuring is needed to get the spiral waves started. Another way to generate spirals from parameters that lead to frozen patterns under random initial configurations is to “train” the process, starting with a larger neighborhood size until synchronous waves start to form, and then decreasing the neighborhood size in the simulation. If synchronous waves have begun to form under the larger neighborhood size, then (with such a configuration as the initial state of the process) a reduction in neighborhood size causes waves to become more focused and less fuzzy. It is easy to understand how, for certain parameters, only initial configurations with some structure would lead to persistence and patterns. For example, in the model with eight states and neighborhoods of size 4, a random configuration has low probability of presenting a given species with its required resource. This results in the various species using up the random bits of resource available to them and then dying out before they can reproduce again due to the lack of required nutrient in their neighborhoods.

The synchronous waves come in two forms: (segments of) circular traveling waves and spiral waves. Along a reactive wavefront (i.e., a species invading its required resource), we typically have a high density of the required resource. This allows a traveling wave that is locally similar to an infectious wave in an epidemic model. In the wake of this reactive wavefront, the invading species leaves behind its product resource, thus laying the foundation for the next species. The shape theorem of IPS (Durrett, 1988) provides the essential ingredients for the dynamics of a traveling reactive wavefront started by a small localized cluster of a given species set in a uniform environment that has all other sites holding the required resource for that species. Such a wavefront is roughly circular in shape with the radius expanding at what is effectively a constant speed. Of course, our model is much more complicated than this. We have multiple species and resources, leaving an inhomogeneous mix of states in which a wavefront must spread, if at all. Instead of finding itself in a uniform field of the required

resource, a given species will typically encounter sites that are holding other resources or other species as well as clusters of its required resource. The result is that the species is trying to spread in a fragmented environment. In the language of IPS, a species must spread on a (space–time) percolation cluster of its required resource. The more fragmented and sparse this percolation cluster is, the less likely it becomes that species members can reproduce before dying (i.e. before the site changes from the species state to the product resource). The only way the species can persist is to spread, and the only way it can spread in such a fragmented environment is to have a growth rate that is sufficiently large compared to its death rate. For example, a contact process in a random environment (that has some sites inaccessible to the species) has a higher required growth rate than a contact process that has no such restrictions.

For our process, the cyclic nature of the local dynamics makes this effect of growth rate a double-edged sword. Large growth rates enable a species to quickly gobble up bits of its required resource almost as soon as they appear. This, however, leaves behind a fragmented landscape for the next species and leads to a breaking up the waves. This is why simulations with large β 's lead to mixing patterns and not synchronous waves.

With this insight, we can now recap the simulation results that appear in Fig. 3. If β is too small, the species members cannot spread fast enough and they die out, leaving behind a frozen state consisting only of resources. If we increase β far enough, the species will be able to persist and coexist by synchronizing their spread with other nearby members of the same species and leaving behind clustered regions of the resource required by the next species. In some cases, this structure is self-organized, even from a random initial configuration. For slightly smaller β , this synchronization may only be possible if the process starts with some structure. As we increase β more and more, the length scales on which the waves operate get smaller and smaller as wavefronts consistently encounter fragmented resources that cause them to split, thus removing the potential for spiral wave development. Asynchronous waves can be thought of as synchronous waves whose length scales are very small. As β increases still further, the species have no trouble persisting. Not only do they not require spatio-temporal synchronization to persist, but the high growth rates preclude such patterning and the result is a mixing configuration.

The dynamics of spiral waves have received much attention in the literature on RDEs and, to a lesser extent, in the literature on stochastic spatial models; cf. Grindrod (1996), Murray (1989), Pálsson and Cox (1996), Savill et al. (1997), and Dieckmann et al. (2000). We limit ourselves to a few comments on the spiral waves observed here. Each spiral has $N = 2L$ arms, appropriately ordered and forming a confluence at the eye of the spiral. A collision between two spirals occurs when two waves of the same species approach each other from different directions, consuming the same resource. Such collisions often break off pieces of the spiral, but as long

as the eye is intact the spiral can persist. The spiral waves in the IPS model are much more irregular than those that arise in RDEs and patch models; see next section.

3.2. Asymmetric case

The results of the asymmetric simulations are summarized in Fig. 4. Many of the above comments apply to this case.

We see that the patterns are strongly influenced by the majority growth rate. However, an increase in the growth rate of species 1 can destroy or weaken the spatial structure due to fragmentation of the resource for the second species, much like what was explained about large β for the symmetric case. We remark that the diagonal region in this figure corresponds to $\beta_1 = \beta$ and hence to the 8:8 symmetric case treated above. Notice that the results do not coincide exactly. This is because the stochastic nature of the simulations (these are from separate runs) allows for different behaviors on individual runs, especially near the boundary regions. This sensitivity can also be seen in the differential equation connection below in regions for which eigenvalues have positive real parts very close to zero.

4. Reaction–diffusion equations and linearization

Consider the following RDE, where u_i denotes the density of resource R_i , v_i the density of species S_i , and Δ is the Laplacian:

$$\begin{cases} \frac{\partial u_1}{\partial t} = \Delta u_1 + \delta_L v_L - \beta_1 u_1 v_1, & \frac{\partial v_1}{\partial t} = \Delta v_1 - \delta_1 v_1 + \beta_1 u_1 v_1, \\ \frac{\partial u_2}{\partial t} = \Delta u_2 + \delta_1 v_1 - \beta_2 u_2 v_2, & \frac{\partial v_2}{\partial t} = \Delta v_2 - \delta_2 v_2 + \beta_2 u_2 v_2, \\ \vdots & \vdots \\ \frac{\partial u_L}{\partial t} = \Delta u_L + \delta_{L-1} v_{L-1} - \beta_L u_L v_L, & \frac{\partial v_L}{\partial t} = \Delta v_L - \delta_L v_L + \beta_L u_L v_L. \end{cases} \quad (3)$$

This is a partial differential equation that is related to the particle system in a natural way via a “fast-stirring” limit (Durrett and Neuhauser, 1994; Krone, 2004). Roughly speaking, this says that if, in our original IPS model, between state changes at different sites we were to allow the “particles” to jiggle around a bit by exchanging the states at some nearby sites, then the above RDE (with the same rates as in the IPS) would serve as a good approximation to the large-scale spatial dynamics of the particle system. By this last part, we mean that the density at a spatial point for the solution to the RDE is approximately equal to the average density for a suitably chosen group of sites in some spatial window in the particle system. The above references can be consulted for a more thorough discussion on the relation between the IPS model and the RDE. For the purposes of the present paper, it is enough to realize that the RDE can be used to provide information about the IPS.

In Fig. 5 we see the behavior of the RDE (3) in the case of $N = 6$ states when all growth rates are the same, for two different values of the common growth rate. We solved the RDE numerically, starting with a random initial condition (in the numerical scheme) and periodic boundary. The initial condition introduces small inhomogeneities that are able to grow for certain parameters. When we set $\beta = 3.8$, spiraling and traveling waves emerge and persist. When $\beta = 5.8$, the patterns die out to a spatially homogeneous steady state. Fig. 5 portrays the density of one of the resources; the bands of resource are wider than those of the species, so they are easier to see. Because of the cyclic nature of the model and the symmetry in the parameters, it does not really matter which density is plotted.

Setting $X = (u_1, v_1, \dots, u_L, v_L)$, we can write the RDE in vector form

$$\frac{\partial X}{\partial t} = \Delta X + f(X). \quad (4)$$

The embedded ODE for this RDE, sometimes referred to as the mean-field ODE for the IPS model, is

$$\frac{dX}{dt} = f(X). \quad (5)$$

We next attempt to use the RDE to predict results for the IPS simulations. We first use a simple linearization




Fig. 5. The density of resource R_1 , with higher density corresponding to lighter gray level, in the numerical solution of the RDE (3). Here, we have $L = 3$ with all species having the same growth rate β ; see symmetric case 2 below. The middle figure shows the random initial configuration. In the left figure, $\beta = 3.8$ and spiral waves develop. In the right figure, $\beta = 5.8$ and all spatial structure dies out to yield a homogeneous distribution.

argument to show that eigenvalues of the linear system govern the behavior of the RDE. We then apply this linearization argument to study the two special cases treated in the particle system simulations: symmetric and asymmetric. The study of these two mathematically tractable cases not only gives us some insight into the IPS model, but also enables us to compare the results of the differential equation and particle system models.

4.1. Linearization

Suppose the ODE system (5) has a unique interior equilibrium $X^* = (u_1^*, v_1^*, \dots, u_L^*, v_L^*)$; i.e. $f(X^*) = 0$ and $u_i^* > 0, v_i^* > 0, i = 1, \dots, L$. We linearize the reaction term of (4) about the spatially homogeneous solution X^* of the RDE to get

$$\frac{\partial Y}{\partial t} = \Delta Y + JY, \tag{6}$$

where Y is a small perturbation and $J = J_f(X^*)$ is the Jacobian of f at the equilibrium X^* . Eq. (6) has a solution of form

$$Y(x, t) = \mathbf{c}_1 e^{i\sqrt{J}x} + \mathbf{c}_2 e^{Jt}, \tag{7}$$

where \mathbf{c}_1 and \mathbf{c}_2 are vectors that depend on the initial condition. This suggests that the behavior of the RDE (3) might be controlled, at least for small perturbations of the interior equilibrium, by the Jacobian for the embedded ODE (5). (For readers unfamiliar with taking functions of matrices, an exponential function of a matrix is obtained by writing the corresponding series expansion (Hirsch and Smale, 1974); the square root of a matrix involves its spectral decomposition and is constructed from a diagonal matrix made up of square roots of the eigenvalues and two other matrices that involve left and right eigenvectors of the original matrix (Hoffman and Kunze, 1971).) Rather than concerning ourselves with making the above equation rigorous, we simply take it as a heuristic for the method that we now propose. (But see Section 4.4.) We claim that the types of behavior in the IPS model seen in Fig. 3 can be predicted by considering the eigenvalues of the mean-field ODE (5) at the interior equilibrium. In particular, we seek to explain why $N > 4$ is required to get spiral waves. Our approach will be to show that, up to a scaling of the parameters, the IPS simulations agree with the predictions based on the RDE. One should not expect a β -for- β agreement between these two models when the particle system has finite interaction range; the fast-stirring limit RDE (which does have the same parameters as the original IPS model) only holds in a limiting sense. For example, in the one-dimensional contact process the critical growth rate is about 1.65, while the critical growth rate for the corresponding RDE is 1. What is important is that the two model types have the same qualitative behavior.

When the Jacobian matrix J has no eigenvalues with positive real part, the spatial inhomogeneities in the initial configuration do not grow. However, when the matrix J

has eigenvalues with positive real part, certain wave solutions initiated by the spatial inhomogeneities can be amplified and propagate through space. Below, we investigate several representative cases that illustrate the connection between the linearization and the IPS results.

4.2. Symmetric case studies

As in the IPS model, the symmetric case here refers to parameters $\delta_1 = \dots = \delta_L = 1$ and $\beta_1 = \dots = \beta_L = \beta$. It is easy to see that system (5) has only one interior equilibrium point X^* with components $u_1 = \dots = u_L = 1/\beta$ and $v_1 = \dots = v_L = 1/L - 1/\beta$. To have all species surviving at equilibrium, and hence all densities u_i and v_i positive, we require $\beta > L$. To see how the number of species affects the behavior of the model, we consider two different values of L .

Case 1 ($L = 2$): In this case, the total number of states in the system is 4 and

$$J_f(X^*) = \begin{pmatrix} 1 - \frac{\beta}{2} & -1 & 0 & 1 \\ -1 + \frac{\beta}{2} & 0 & 0 & 0 \\ 0 & 1 & 1 - \frac{\beta}{2} & -1 \\ 0 & 0 & -1 + \frac{\beta}{2} & 0 \end{pmatrix}. \tag{8}$$

Since $L = 2, \beta > 2$ is required for the existence of an interior equilibrium. The real parts of all the eigenvalues of matrix (8) with respect to different β are plotted in Fig. 6.

There is a zero eigenvalue for any β since the cyclic dynamics makes matrix (8) singular. There is also a negative real eigenvalue and two complex conjugate eigenvalues with negative real parts when $\beta < 18$; for $\beta \geq 18$, all eigenvalues are real and the three that are non-zero are negative. The lack of positive real part in the eigenvalues implies that system (5) is stable near the equilibrium X^* for any $\beta > 2$. The dissipative nature of the system explains why there are no synchronized waves in the corresponding particle system, as reported in Fig. 3.




Fig. 6. The real parts of all eigenvalues of matrix (8) as a function of $\beta > 2$.




Fig. 7. The real parts of all eigenvalues of matrix (9) as a function of $\beta > 3$.

Case 2 ($L = 3$): In this case, the total number of states in the system is 6 and

$$J_f(X^*) = \begin{pmatrix} 1 - \frac{\beta}{3} & -1 & 0 & 0 & 0 & 1 \\ -1 + \frac{\beta}{3} & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 - \frac{\beta}{3} & -1 & 0 & 0 \\ 0 & 0 & -1 + \frac{\beta}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 - \frac{\beta}{3} & -1 \\ 0 & 0 & 0 & 0 & -1 + \frac{\beta}{3} & 0 \end{pmatrix}. \tag{9}$$

Notice that in this case, $\beta > 3$ is required to get positive equilibrium densities. The real parts of all the eigenvalues of matrix (9), as functions of $\beta > 3$, are plotted in Fig. 7.

Compared to the 4-state case, a major difference in the 6-state case is that for a certain parameter range ($3 < \beta < 4.5$) there are eigenvalues with real part larger than 0. Therefore, within this parameter range, the system is no longer stable at the interior equilibrium and trajectories approach what is approximately a heteroclinic orbit, as shown in Fig. 8(a). When $\beta > 4.5$, the interior equilibrium is stable; Fig. 8(b) shows a spiraling in trajectory when $\beta = 4.8$. Interestingly, $\beta = 4.5$ is a bifurcation point where we observe a limit cycle, as shown in Fig. 8(c). This, again, agrees with the behavior of the IPS model (up to a scaling of the parameters) described in Fig. 3. Notice, in particular, that the parameter region corresponding to coexistence via synchronous pattern formation is adjacent to the parameter values that do not permit coexistence. The trajectories shown in Fig. 8 are representative, with the many other starting states we tried yielding similar behavior.

4.3. Asymmetric case study

In the asymmetric case, the first growth rate β_1 is different from other majority growth rates (β) that are equal. For comparison with the particle system results, we study the case $L = 4$. The interior equilibrium of system (5) is given by $u_1 = 1/\beta_1$ and $u_2 = u_3 = u_4 = 1/\beta$, $v_1 = v_2 = v_3 = v_4 = 1/4 - 1/4\beta_1 - 3/4\beta$, where positive species densities require a constraint on β and β_1 in the form of positivity of the parameter

$$a \equiv \frac{1}{4} - \frac{1}{4\beta_1} - \frac{3}{4\beta}. \tag{10}$$

In terms of this parameter, the corresponding Jacobian matrix is

$$J_f(X^*) = \begin{pmatrix} -a\beta_1 & -1 & 0 & 0 & 0 & 0 & 0 & 1 \\ a\beta_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -a\beta & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & a\beta & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -a\beta & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & a\beta & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -a\beta & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & a\beta & 0 \end{pmatrix}. \tag{11}$$

Fig. 9 illustrates the linearization results for the asymmetric model. Notice that the hyperbolic curve separating the black and white regions comes from Eq. (10). One should compare this figure to Fig. 4 to appreciate how well the linearization and eigenvalue argument explains and predicts the particle system behavior in the asymmetric case.

4.4. A coupled map lattice interpretation

Another perspective on the connection between the behavior of the IPS model and the mean-field ODE can be gleaned from a kind of coupled map lattice interpretation. For specificity, we restrict this discussion to the symmetric case discussed in Sections 3.1 and 4.2.

The global (average) behavior of the IPS model is given approximately by the mean-field ODE if the neighborhood size (i.e. the interaction range) is so large that all sites are neighbors (Durrett, 1995). For smaller interaction ranges, the ODE still provides some information about global averages, although the parameters no longer match exactly and stochastic effects play a more prominent role. Now imagine that the lattice is partitioned into a collection of spatially distributed habitat patches and, with the above approximation in mind, the dynamics within a patch proceed according to the mean-field ODE and nearby patches interact to couple the whole system. Of course, the interaction ranges are not very large and the ‘‘patches’’ overlap considerably, but such simple approximations often lead to significant insight.

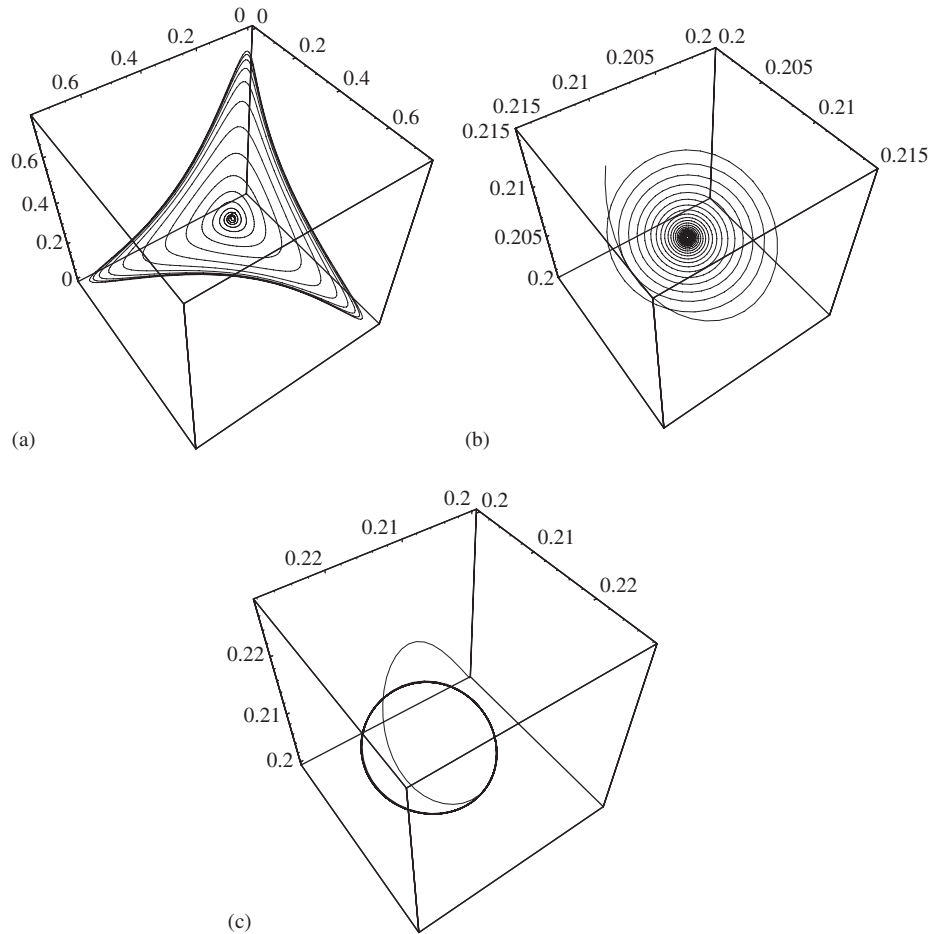


Fig. 8. Representative trajectories plotting the densities of the three resources for different values of the common growth rate: (a) $\beta = 3.8$ giving spiraling-out (nearly) heteroclinic orbit; (b) $\beta = 4.8$ giving trajectory spiraling into stable interior equilibrium; (c) $\beta = 4.5$ giving a limit cycle.

For large values of β , the mean-field ODE has a globally attracting interior equilibrium (cf. Fig. 8(b)). This means that the patches will have type densities that are fairly stable; in particular, each species will be present at positive density with high probability. So even if a patch happens to be missing a species, it has the opportunity to quickly regain that species from surrounding patches and re-establish the approximate equilibrium. This suggests that large β leads to mixed configurations, in agreement with our simulations of the IPS model.

For intermediate values of β , one obtains eigenvalues with positive real part. The interior equilibrium of the mean-field ODE is then unstable with trajectories approaching something close to a heteroclinic orbit (cf. Fig. 8(a)), with episodes of domination by a single resource (and low densities of all other resources and species) followed by rapid invasion of the preferred species (the one having the dominant resource as its required resource) and subsequent decay to domination by the next resource in the cycle. This means that the patches spend most of their time with type densities that are dominated by a single resource. If this domination is strong, in the sense that the “corners” of the trajectories in Fig. 8(a) are close to the coordinate axes, then

to keep the species from going extinct the patches must synchronize the surges of species invasion in a way that leads to waves of species spreading into regions dominated by a single resource. If this happens, then there will be contiguous regions dominated by the product resource, setting the stage for the next species in the cycle, etc.

5. Discussion

Bacterial communities structure their environments to enhance persistence and coexistence. Biofilms provide a ubiquitous example of such structuring. Another common situation of structuring in natural communities is seen in bacterial mats. This is a striking example of niche partitioning in which different microbial species organize themselves in horizontal slabs along nutrient or electron acceptor gradients; here, exogenous heterogeneity in the environment seems to be the main force driving the spatial organization rather than the purely self-organized pattern formation in our system. Looking at a cross-section of a bacterial mat, one sees with the unaided eye sharp boundaries between different species. This niche partitioning is essentially constant in time. The spatio-temporal niche




Fig. 9. The largest real part for the eight eigenvalues of matrix (11) is plotted as a function of the majority growth rate β (horizontal axis) and the growth rate, β_1 , of the first species (vertical axis). The dark region corresponds to parameter values that result in extinction of all species (i.e. no interior equilibrium and hence a frozen pattern); the gray region on the right is where none of the eigenvalues has positive real part (i.e. the largest real part is zero) and corresponds to mixing; in the remaining region, there are eigenvalues with positive real part, and the gray level is proportional to the size of the largest such real part, with lighter color corresponding to larger real part. This last region corresponds to synchronous waves, with lighter color implying larger length scales for the waves.

partitioning seen in the synchronous waves of our resource–species model can be thought of as space–time slabs of different species and resources; they are, however, significantly thinner than the static bands of a bacterial mat.

In the model presented here, the length scales of the macroscopic patterns are large (~ 1 mm) compared to the sizes of the bacterial cells (~ 1 μm). To see where these lengths come from in the IPS model, just note that lattice points correspond to the locations of individual cells, so accounting for a cell size of about one or two μm and intercellular distances also of the order of μm 's, we can think of the real distance between lattice points as being roughly 1–10 μm . With hundreds (or thousands) of lattice points per side for a typical simulation grid, we should consider our viewing window to be several mm on a side, and hence this is also the length scale for patterns that extend across a large portion of this window.

This spatial structuring of the resources may lead to protocols for detecting such cyclic dynamics in the field or in laboratory bacterial communities. When spatial patterning is strongest, the band width of resources is much larger than that for species, so methods that can distinguish between the different resources may be able to pick up the tell-tale patterns even without the species being present.

The length scale estimates mentioned above should also help to direct this. In general, methods that allow one to use “macroscopic” observations, say of large-scale spatial patterns, to acquire knowledge about the underlying microscopic mechanisms can be quite useful.

To our knowledge, microbial systems with the specific dynamics described here have not yet been observed, although the types of feedback mechanisms we assume are known to exist and cyclic dynamics are quite common in nature. Furthermore, cross-feeding appears to arise quite naturally in microbial systems, even in the absence of spatial structure. For example, Rosenzweig et al. (1994) observed the evolution of such a polymorphism in a chemostat population of *E. coli* initiated with a single clone. As the population evolved to be dominated by individuals that could more quickly exploit the single limiting nutrient (glucose), the ability to assimilate secondary metabolites (acetate and glycerol) was rendered less efficient, thus opening the way for mutant strains that specialized on these secondary resources. Doebeli (2002) gives a nice mathematical account of the evolution of cross-feeding polymorphisms. Despite all this evidence, it is not surprising that microbial systems with the cyclic spatial dynamics presented here have not been observed. This is in large part due to the fact that microbial observations on such a fine spatial scale are not often undertaken, and in laboratory communities where it is easier to gain spatial resolution through confocal microscopy and fluorescent markers, one rarely deals with communities of more than a few species. We consider this a challenge to microbial ecologists.

This is an excellent area for the blending of theoretical and empirical studies in microbial ecology. A few other examples of theoretical and empirical work combined at the microbial level are Durrett and Levin (1997), Kerr et al. (2002), and Wei and Krone (2005). With individual-based lattice models, we can effectively “see” at an individual cell level the effects of certain mechanisms. This leads to biological hypotheses that can be tested empirically. Our simulations show the types of patterns and length scales one should look for and also point out that it is enough to detect the nutrient patterning. This work also shows that spatial patterns in microbial (and other) systems can be a result of self-organization and not necessarily the result of inhomogeneities in the substrate. It is an example of microbial communities structuring both their abiotic and biotic environments.

In situations for which synchronous waves arise, another interesting feature holds. Even if the species were to be driven to extinction by external forces, say due to fluctuations in the environment (that are not modeled), the consumer-driven structuring of the resources leads to the potential for rapid re-establishment of the microbial community and its spatial structure provided a small influx of immigrants is able to re-seed the prepared environment before this structure erodes; cf. the comments about “seeding” and hard-to-start spirals in Section 3. This suggests an additional robustness to such cyclic systems.

All of these features help to illustrate why cyclic dynamics are likely to be prevalent in natural spatially structured systems and, in particular, why they could be a regular component of microbial communities. More empirical evidence is needed and the results of this paper illustrate the types of spatial patterns that one might look for. We intend to elaborate on the themes of this paper elsewhere by including evolutionary dynamics in our model.

Acknowledgements

We thank Larry Forney, Zaid Abdo, and Holger Heuer for discussions. Forney also provided extensive comments on an earlier version of the manuscript, improving both presentation and biological precision. Finally, we thank three anonymous reviewers for their helpful suggestions. SMK and YG were supported in part by NSF Grant EPS-00-80935. Some of the computational aspects of this work were enabled by NIH Grants P20 RR016454 and P20 RR16448 from the INBRE and COBRE programs, respectively, of the National Center for Research Resources.

References

- Bradshaw, D.J., Homer, K.A., Marsh, P.D., Beighton, D., 1994. Metabolic cooperation in oral microbial communities during growth on mucin. *Microbiology* 140, 3407–3412.
- Bramson, M., Griffeath, D., 1989. Flux and fixation in cyclic particle systems. *Ann. Probab.* 17, 26–45.
- Comins, H.N., Hassell, M.P., May, R.M., 1992. The spatial dynamics of host–parasitoid systems. *J. Anim. Ecol.* 61, 735–748.
- Dieckmann, U., Law, R., Metz, J.A.J. (Eds.), 2000. *The Geometry of Ecological Interactions: Simplifying Spatial Complexity*. Cambridge University Press, Cambridge.
- Doebeli, M., 2002. A model for the evolutionary dynamics of cross-feeding polymorphisms in microorganisms. *Popul. Ecol.* 44, 59–70.
- Durrett, R., 1988. *Lecture Notes on Particle Systems and Percolation*. Wadsworth, Belmont, CA.
- Durrett, R., 1995. *Ten Lectures on Particle Systems*. Lecture Notes in Mathematics, vol. 1608. Springer, Berlin.
- Durrett, R., Griffeath, D., 1993. Asymptotic behavior of excitable cellular automata. *Exp. Math.* 2, 183–208.
- Durrett, R., Levin, S.A., 1997. Allelopathy in spatially distributed populations. *J. Theor. Biol.* 185, 165–172.
- Durrett, R., Neuhauser, C., 1994. Particle systems and reaction–diffusion equations. *Ann. Probab.* 22, 289–333.
- Grindrod, P., 1996. *The Theory and Applications of Reaction–Diffusion Equations: Patterns and Waves*, second ed. Oxford University Press, Oxford.
- Hassell, M.P., May, R.M., 1988. Spatial heterogeneity and the dynamics of parasitoid–host systems. *Ann. Zool. Fenn.* 25, 55–61.
- Hirsch, M.W., Smale, S., 1974. *Differential Equations, Dynamical Systems, and Linear Algebra*. Academic Press, San Diego.
- Hoffman, K., Kunze, R., 1971. *Linear Algebra*, second ed. Prentice-Hall, Englewood Cliffs, NJ.
- Kerr, B., Riley, M.A., Feldman, M.W., Bohannan, B.J.M., 2002. Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature* 418, 171–174.
- Krone, S.M., 2004. Spatial models: stochastic and deterministic. *Math. Comput. Mod.* 40, 393–409.
- Murray, J.D., 1989. *Mathematical Biology*. Springer, Berlin.
- Pálsson, E., Cox, E.C., 1996. Origin and evolution of circular waves and spirals in *Dictyostelium discoideum* territories. *Proc. Natl Acad. Sci. USA* 95, 1151–1155.
- Pelz, O., Tesar, M., Wittich, R.-M., Moore, E.R.B., Timmis, K.N., Abraham, W.-R., 1999. Towards elucidation of microbial community metabolic pathways: unravelling the network of carbon sharing in a pollutant-degrading bacterial consortium by immunocapture and isotopic ratio mass spectrometry. *Environ. Microbiol.* 1, 167–174.
- Rosenzweig, R.F., Sharp, R.R., Treves, D.S., Adams, J., 1994. Microbial evolution in a simple unstructured environment: genetic differentiation in *Escherichia coli*. *Genetics* 137, 903–917.
- Savill, N.J., Rohani, P., Hogeweg, P., 1997. Self-reinforcing spatial patterns enslave evolution in a host–parasitoid system. *J. Theor. Biol.* 188, 11–20.
- Wei, W., Krone, S.M., 2005. Spatial invasion by a mutant pathogen. *J. Theor. Biol.* 236, 335–348.