Available online at www.sciencedirect.com
ScienceDirect
Theoretical Population Biology 70 (2006) 387–400
Theoretical Population Biology
www.elsevier.com/locate/tpb

Spatial effects favour the evolution of niche construction

Matt Silver, Ezequiel Di Paolo*

Evolutionary and Adaptive Systems Group, Centre for Computational Neuroscience and Robotics, University of Sussex, Brighton BN1 9QH, UK

Received 26 April 2005
Available online 26 August 2006

Abstract

We present an individual-based, spatial implementation of an existing two-locus population genetic model of niche construction. Our analysis reveals that, across a broad range of conditions, niche-construction traits can drive themselves to fixation by simultaneously generating selection that favours 'recipient' trait alleles and linkage disequilibrium between niche-construction and recipient trait alleles. The effect of spatiality is key, since it is the local, resource-mediated interaction between recipient and niche-constructing loci which gives rise to gene linkage. Spatial clustering effects point to a possible mechanism by which an initially rare recipient trait whose selection depends on niche construction could establish in an otherwise hostile environment. The same mechanism could also lead to the spread of an established niche-constructing colony. Similar phenomena are observed in the spatial modelling of two species 'engineering webs'. Here, the activities of two niche-constructing species can combine to drive a particular recipient trait to fixation, or in certain circumstances, maintain the presence of polymorphisms through the preservation of otherwise deleterious alleles. This may have some relevance to ecosystem stability and the maintenance of genetic variation, where the frequencies of key resources are affected by the niche-constructing activities of more than one species. Our model suggests that the stability of multi-species webs in natural populations may increase as the complexity of species–environment interactions increases.

Keywords: Niche construction; Spatial effects; Individual-based model; Ecosystem stability; Engineering web; Resource-mediated interaction.

1. Introduction

Environmental change has long been acknowledged as a major factor affecting organic evolution (Van Valen, 1973). More recently, attention has focused on the evolutionary consequences of the role organisms play in influencing their own environment in the short and long term—a phenomenon known as *niche construction* (Lewontin, 1983; Odling-Smee et al., 2003). It has been suggested that such organism-environment feedback can affect evolutionary dynamics in a variety of ways, from unexpected order, through to instability, chaos and sudden extinction (Robertson, 1991).

Organisms alter their environment in a myriad ways. The changes so wrought can dramatically affect the fitness of individuals in current and future generations. For example, the web maintenance and defence behaviours of many spiders depend on their inheriting the genes for web construction, whereas earthworms are uniquely adapted to the changes in soil structure and chemistry brought about by the cumulative effect of the niche-constructing activities of previous generations. Other examples of niche construction include the construction of beaver dams and termite mounds, the storing of food, regulation of soil temperature, acidity and humidity (Odling-Smee, 1996), and the enhanced flammability exhibited by some plant species (Mutch, 1970; Schwilk, 2003).

Laland et al. (1996, 1999) have used a two-locus, population genetic model to analyse the joint evolution of a niche-constructing trait, and a so-called recipient trait whose selection depends in part on the environmental change wrought by the population's own niche construction. They found that niche construction could be a 'potent evolutionary agent', generating forms of evolutionary 'momentum' and 'inertia' and leading to the fixation of otherwise deleterious alleles, or supporting

*Corresponding author.
E-mail addresses: mjs22@sussex.ac.uk (M. Silver), ezequiel@sussex.ac.uk (E. Di Paolo).

0040-5809/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.tpb.2006.08.003

stable polymorphisms where none would otherwise be expected.

Such shifts in a population's genetic makeup can have a significant effect on its subsequent evolutionary trajectory, since the abundance of initially rare phenotypes, alleles, or species may sometimes determine the long-term response of an evolutionary system to change (Holt, 1995).

Of particular interest is the way in which their model showed how niche construction is able to generate linkage disequilibrium (LD) between niche constructing and recipient trait loci. In this respect, niche construction can be compared with other, similar mechanisms which modify the selective environment that a gene experiences, for example, epistasis and maternal and indirect genetic effects (Wolf et al., 2000; Brodie, 2005). A key difference is that niche construction links gene loci through their interaction with biotic and/or abiotic environmental components.

Niche construction is also likely to be important in an ecological context, since a species' niche-constructing activities may well have an impact on other species with which it shares an environment (Odling-Smee et al., 2003). The term 'ecosystem engineering' has been used to describe the activities of organisms that modify other species' environments, thereby directly or indirectly controlling the availability of resources (Jones et al., 1994, 1997). These interactions fall outside the more familiar realm of, for example, predator–prey interactions and food webs, since they are defined as being non-trophic. Examples include the creation of living space in soil cavities caused by root growth and the effects of chemical exudates from lichens on rock erosion (Shachak et al., 1987). Ecosystem engineering may also illuminate current debates on the relation between ecological dynamics and biodiversity (Hector et al., 2000; Huston et al., 2000; Wardle et al., 2000; see also reviews in Tilman and Kareiva, 1997; Loreau et al., 2002), by highlighting the potentially stabilising effects of non-biotic factors. Jones et al. predict that for large ecosystems over long timescales, engineering 'webs' are likely to increase species richness and abundance, with implications for community and ecosystem stability. The connection between niche construction and ecosystem engineering suggests a possible mechanism for a link between ecosystem complexity and stability at both the evolutionary and ecological level, by for example providing adaptive interpretations of resource acquisition and utilisation (Holt, 1995; Laland et al., 1999; Reusch et al., 2005).

In this paper, we explore the effect of introducing spatiality, finite populations and stochasticity in an individual-based version of Laland et al.'s genetic model of niche construction. It is anticipated that the effect of introducing such factors may affect evolutionary outcomes when compared with non-spatial, 'mixed-medium' models, as has been found elsewhere (Durrett & Levin, 1994a, b; Krakauer & Pagel, 1995; Nakamaru et al., 1997; van Baalen & Rand, 1998; Di Paolo, 2000). To maintain parsimony, all the remaining assumptions that went into the original model remain purposely unchanged. Finally, our model is extended to a two-species engineering web with the addition of another species. We study the simplest cases of non-trophic ecological interaction corresponding to the three possible combinations of positive and negative niche construction on a shared resource such as the regulation of soil pH (Jones et al., 1997). In order to isolate the effects of niche engineering on gene frequencies, we explicitly do not include factors such as competitive spatial exclusion in this initial model. Our intention here is to extend Laland et al.'s model in a number of ecologically interesting ways while acknowledging the exploratory character of this task.

2. Methods

We implement a spatial, individual-based version of the genetic model of niche construction analysed by Laland et al. (1996, 1999). The population consists of a finite number of diploid individuals, each with two, diallelic loci: a *resource-dependent* locus, **A**, with alleles A and a; and a *niche-constructing* locus, **E**, with alleles E and e.

Individual genotypes are arranged in an $n \times n$ square lattice with wrap-around (toroidal) boundaries. Each lattice point (x_i, y_j) is occupied by a single individual $\{\mathbf{A}_{ij}, \mathbf{E}_{ij}\}$ and has an associated, local environmental resource frequency $\{R_{ij}\}$. Each individual has eight nearest neighbours.

Alleles at an individual's **A** locus make a contribution to fitness that is in part a function of the local resource frequency. Resource frequencies are subject to change as a result of (i) the niche-constructing activities of individuals in the population, and (ii) independent processes of depletion and renewal. An individual's capacity for niche construction depends on the frequency of E alleles at its **E** locus. The spatial distribution of the resource, $\boldsymbol{R}$ at time t is given by the following recursion:

$$\boldsymbol{R}_t = \lambda_1 \boldsymbol{R}_{t-1}[1 - \gamma \boldsymbol{p}(E)] + \lambda_2 \boldsymbol{p}(E) + \lambda_3. \quad (1)$$

A scalar version of this equation is applied at each lattice point. All variables and coefficients are dimensionless. $\boldsymbol{p}(E)$ is the spatial frequency distribution of E alleles in the population at time t–1. At any particular lattice point, this takes one of the values $\{0, 1/2, 1\}$ corresponding to the three possible allelic combinations at an individual's **E** locus, respectively $\{ee, Ee, EE\}$. λ_1 is a coefficient of independent depletion; γ and λ_2 are coefficients of negative and positive niche construction, corresponding to niche-constructing activity that, respectively, decreases or increases $\boldsymbol{R}$; λ_3 is a coefficient of independent renewal. Local resource frequency is thus a function of independent processes of depletion and renewal, and of the cumulative effect of local niche-construction activity over preceding generations.

We assume that $0 < \lambda_1, \lambda_2, \lambda_3, \gamma < 1$ and $\lambda_1 + \lambda_2 + \lambda_3 \leqslant 1$, so that the local resource frequency, R, can take any real value between 0, corresponding to complete absence of the

resource, and 1, corresponding to resource 'saturation'. Thus $\boldsymbol{R}$ represents an environmental resource whose frequency is constrained within physical or density-dependent limits, for example pH in water, accumulation of detritus or soil humidity.

Genotype fitness is determined entirely by the alleles at an individual's **A** locus (the situation where niche-constructing activity is the subject of direct selection is not considered here). Fitness can depend both on resource frequency and on selection from an external source. The fitness of the three possible allelic combinations at **A** are as follows:

$$\begin{aligned} f_{AA} &= \alpha_2 + \varepsilon R, \\ f_{Aa} &= 1 + \varepsilon\sqrt{R(1-R)}, \\ f_{aa} &= \beta_2 + \varepsilon(1-R), \end{aligned} \tag{2}$$

where R is the local resource frequency (defined below). The first terms in each of these fitness relations correspond to *fixed fitness components*, representing the effect of external selection operating at **A**. The second terms refer to the *frequency-dependent components* of selection. Each A allele makes a contribution to genotype fitness proportional to $\sqrt{R}$, and each a allele a contribution proportional to $\sqrt{(1-R)}$. The coefficient of proportionality ε determines the strength (relative to external selection), and direction of resource frequency-dependent selection, with positive ε indicating increased environmental resource levels will favour the A allele. This relationship between fitness and resource frequency follows that of Laland et al. (1996, 1999), to facilitate comparison between our individual-based model and the original. This particular form of fitness function minimally captures the functionality of the system we wish to study, namely the co-evolution of niche construction and recipient traits in a population that is not itself biased towards producing niche construction. An alternative fitness model in which heterozygotes were made fully neutral with respect to the resource, with fitness intermediate between the two homozygotes, was also tested under the case of no external selection (i.e. $\alpha_2 = \beta_2 = 1$). Spatial and non-spatial simulations showed no significant deviation from the results presented here (data not shown). Laland et al. (1996) also tested different powers of resource dependence—with the frequency-dependent components (RHS of (2)) raised to the power f, with $f = 0.5$, 1 and 2. They found their results to be qualitatively unaffected by the choice of f, this time in the case of external selection at **A** ($\alpha_2 \neq \beta_2$). Taken together, these further investigations suggest that, within the constraints of the niche-constructing system we are trying to model, the actual form of the fitness dependence may not be fundamental here. We should also mention that a simpler haploid model might also be expected to reveal effects similar to those observed here. Once again, our aim is not to explore the significance of every possible variable in Laland's model. Instead we seek to isolate a few evolutionary and ecologically interesting changes, while keeping the rest of the model intact, in order to compare with the original.

Offspring are formed by combining the genes from two 'parent' genotypes. Independent segregation is assumed throughout, i.e. offspring receive one (random) allele from each parent at each locus. Individuals are able to mate with any of their eight nearest neighbours, plus themselves (neighbourhood size of $8+1$), so that the possibility of self-fertilisation is included.

Offspring (viability) fitness is determined with reference to the resource level, R, at one (randomly selected) parent's location. We consider this to be the simplest mechanism that maintains the link between an offspring's viability and the effect on resource level of the previous generation's niche construction. One alternative would have been to include a model of resource diffusion, so that a local or regional correlation between niche-constructing activity (or lack of it) would also be conserved. Since different forms of resource transport represent rather specific natural scenarios, we have decided to implement the simplest situation in this first instance. Clearly this represents a particularly strong form of 'localised ecological inheritance' in that offspring fitness is determined entirely by parental resources. There are many forms of niche construction which fit this model, one example being that of out crossing plants where seed dispersal distance is less than pollen dispersal distance so that an individual's environment is close to its maternal one (Galloway, 2005). However, 'looser' forms of ecological inheritance, corresponding to different bio-geographical scenarios are clearly also possible (see Discussion).

Probability of offspring survival is proportional to fitness, with a finite probability that even fitter offspring do not survive. Dispersal and mating neighbourhoods are assumed to be identical, so that surviving offspring are placed in one of the $8+1$ cells in the first parent's neighbourhood, replacing the original occupant. In this way fitter offspring may spread out to colonise neighbouring cells, while an effective carrying capacity equal to the total population size is maintained. As a limitation of choosing a grid model with full constant occupancy we are immediately discarding the possibility of studying density dependence effects (see Discussion).

Parsimony with Laland et al.'s population genetic model requires that the distribution of resources is updated once every 'generation'. In the spatial model, a generation is defined as n^2 random couplings, so that there will be significant overlap between one generation and the next. Results are compared with a non-spatial control in which spatial associations between genotypes and resource are effectively disregarded, with individuals and resource locations picked at random from the general population.

2.1. Single-species model

Two cases are studied in detail, corresponding to *no external selection* and *external selection at* **A** (Laland et al.,

1996, 1999). In the no external selection case, genotype fitness is determined entirely by the frequency R of resource encountered in the environment, i.e. $\alpha_2 = \beta_2 = 1$. Parameters affecting resource dynamics are as described in Laland et al. (1999, p. 10244; graph (a)), corresponding to positive niche construction with independent renewal and depletion of resources. (Note that positive niche construction with $\lambda_2 > 0$ and $\gamma = 0$ is defined by Eq. (1) as phenotypic activity that *increases* resource frequency.) $\varepsilon = 0.3$ so that positive niche-construction favours the A allele. In the external selection at **A** case, parameters correspond to positive niche construction favouring the a allele, with strong external selection favouring A (Laland et al., 1999, p. 10244; graph (e)). Here niche construction *opposes* an external source of selection.

2.2. Two-species 'engineering web' model

In a second analysis, the single-species model is extended to include an additional species whose fitness also depends on the frequency of the same environmental resource. $\boldsymbol{R}$ now potentially depends on the niche-constructing activity of *two* species, S_1 and S_2, as well as independent processes of renewal and depletion. Each grid cell is occupied by a single genotype from each species. Each species' fitness depends only on local resource frequency, corresponding to the previous 'no external selection' case.

Three types of engineering web are considered, corresponding to different combinations of species capable of either positive or negative niche construction. In 'web 1', only S_1 is able to affect resource frequency through positive niche construction, depending on frequency of E alleles at $\mathbf{E_1}$. Here the resource recursion is given by

$$\boldsymbol{R}_t = \lambda_1 \boldsymbol{R}_{t-1} + \lambda_2 \boldsymbol{p}(E_1) + \lambda_3. \tag{3}$$

In 'web 2', both species are able to affect resource frequency through positive niche construction, determined by the local frequency (i.e., the frequency at each grid position) of E alleles at $\mathbf{E}_1$ and $\mathbf{E}_2$. The resource recursion is given by

$$\boldsymbol{R}_t = \lambda_1 \boldsymbol{R}_{t-1} + \lambda_2 \left[\frac{\boldsymbol{p}(E_1) + \boldsymbol{p}(E_2)}{2}\right] + \lambda_3. \tag{4}$$

In 'web 3', each species' niche construction works in opposition to the other's. In this case, resource dynamics are given by

$$\boldsymbol{R}_t = \lambda_1 \boldsymbol{R}_{t-1}[1 - \gamma \boldsymbol{p}(E_2)] + \lambda_2 \boldsymbol{p}(E_1) + \lambda_3. \tag{5}$$

Experimental outcomes are measured over a range of initial mean E allele frequencies ($p(E)$). In each case $p(E_1) = p(E_2)$ at $t = 0$. Parameters are chosen so that there is a single, unstable polymorphic equilibrium at **A** for both species when $p(E_1) = p(E_2) = 0.5$ at $t = 0$. All other initial E allele frequencies lead to fixation to one or the other **A** allele for both species.

The time at which either **A** allele in either species first fixes is measured for each engineering web across a range of initial E allele frequencies. This provides one measure of web stability, since those webs in which both species remain polymorphic at **A** retain the capacity to adapt to continuing fluctuations in 'external' resource frequency. A web in which either species has fixed to either **A** genotype loses this ability. A similar point with respect to polymorphisms in a single species is made in Laland et al. (1996, pp. 307–308). (See also Holt, 1995). Note that the requirement that both species maintain polymorphisms at A is in fact a rather stringent one, since in the present 'no external selection' case, there are no stable polymorphic equilibria.

3. Results

3.1. Single species model

For an infinite, non-spatial population with no external selection at **A**, Eq. (1) has an unstable equilibrium for an average, population-wide E allele concentration, $p(E)$ of 0.25, corresponding to an average of one E allele per two diploid individuals. Populations with E allele frequencies above or below this will fix to either the A or a allele, respectively.

Data from the individual-based simulation of spatial populations ranging in size from 400 (20×20 grid) to 6400 (80×80) are plotted in Fig. 1, together with data for a non-spatial 'control' population of 6400 individuals. Predictions of the infinite, non-spatial population genetic model are also included for comparison.

Results for the non-spatial control are close to those predicted by the infinite, population-genetic model. The majority of test populations fix either to A or a, for initial $p(E)$, respectively, above or below 0.25. With $p(E) = 0.25$, approximately half the test populations fix to A, with the remainder fixing to a. The rounded, 's-shape' characteristic of this curve is the result of genetic drift due to finite-population sampling error. This tips a minority of test populations to fixation of the alternate **A** allele close to $p(E) = 0.25$. As expected, the effect is less pronounced for larger population sizes (data not shown).

For *spatial* populations, fixation of the A ('recipient trait') allele is observed at significantly lower initial E allele frequencies. This effect is accentuated in larger populations. Observation of evolutionary dynamics in the run up to equilibrium reveals a marked difference from the non-spatial model, with strong linkage between A and E alleles, and local 'clustering' of niche-constructing (AA/EE and AA/Ee) genotypes, followed by fixation or loss at *both* loci (e and a alleles are similarly linked).

The evolution of LD for typical spatial and non-spatial populations of 1600 individuals, with initial $pA = 0.5$ and initial $pE = 0.15$, is illustrated in Fig. 2. $p_A$loc and $p_E$loc are the frequencies of individuals with at least one A allele, and the frequency of individuals with at least one E allele, respectively. LD is measured using a diploid version of Lewontin's disequilibrium coefficient (see the Appendix).

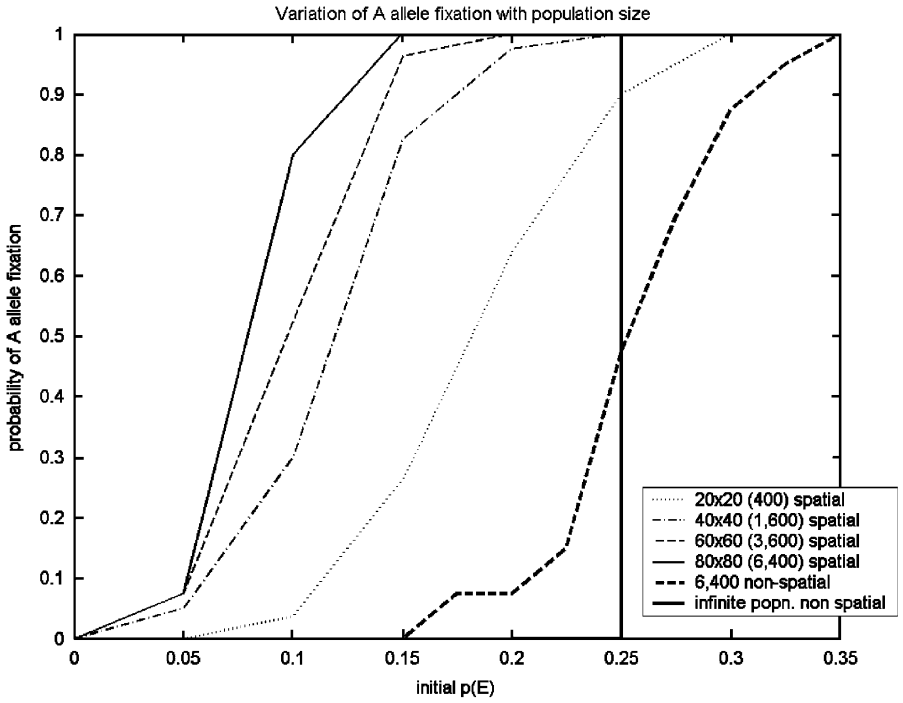

Fig. 1. Variation of probability of A allele fixation with (spatial) population sizes of 400, 1600, 3600 and 6400. No external selection at **A**. $\lambda_1 = 0.7$; $\lambda_2 = 0.2$; $\lambda_3 = 0.1$; $\gamma = 0$; $\alpha_2 = \beta_2 = 1$. Non-spatial individual-based and population genetic data are included for comparison. Probabilities are based on 80 trial runs (3 smallest population sizes) and 40 trial runs (largest population size of 6400). In each case, allele A fixes to the right of the line, and allele a to the left.

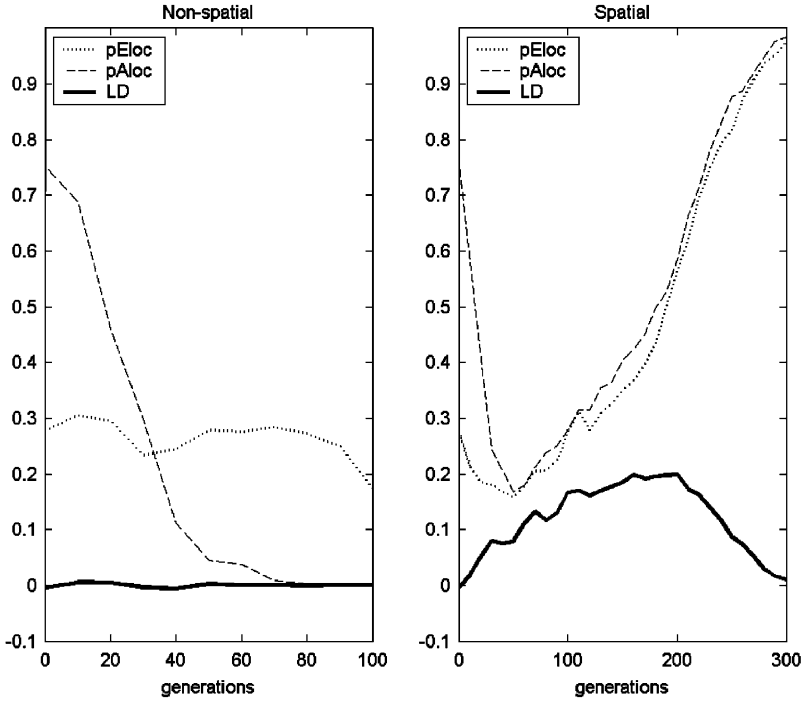

Fig. 2. Evolution of linkage disequilibrium (LD) in a population of 1600 individuals (40×40 spatial grid) with initial $pE = 0.15$ and initial $pA = 0.5$ (see Appendix for a derivation of the relationship between initial vaules of pA and pE, and $p_A$loc and $p_E$loc).

In the non-spatial case, there is no significant linkage between loci. pA (and $p_A$loc) quickly falls to zero and LD ≈ 0 throughout. pE (and $p_E$loc) drifts downward due to finite population sampling error. In the spatial case, LD is again initially zero, corresponding to a random distribution of A and E alleles. However, LD rapidly increases with the evolution of a 'patchy' distribution of linked A and E alleles. pA (and $p_A$loc) shows a rapid initial decline as most

small niche-constructing patches disappear, until a single cluster (or amalgamation of smaller clusters) reaches sufficient size to first stabilise, and then rapidly expand, driving both E and A to fixation. Once A and E allele frequencies pass 0.5, the expected frequency of individuals with at least one A and one E allele begins to rise, and $LD \to 0$. LD is almost at a maximum—with the vast majority of E alleles in the population linked to at least one A allele (and vice versa)—from the point at which a single, stable niche-constructing cluster has formed (around generation 100—also see Fig. 4). This reflects the fact that niche-constructing alleles are now subject to strong A locus-dependent selection since they affect genotype fitness through their cumulative impact on *local* resource frequency. For example, at low E allele concentrations, E alleles rapidly vanish, outside of niche-constructing clusters, despite there being no direct selection pressure at **E**. This is because low levels of niche construction mean that local resource levels are low, and a alleles dominate. Any individual carrying a niche-constructing E allele in such a region is selected against, since the resulting increase in local resource levels means that offspring fitness is reduced. (A similar argument applies to linkage within a cluster.)

Without spatiality, there can be no association between the loci. This time there is no correlation between possession of an E allele and reduced fitness, since the association between the niche-constructing effect of each E allele and the fitness of the individual carrying that allele is broken. This is because offspring fitness is calculated against the resource level at a random location—as it must be for the model to be non-spatial. There is therefore no linkage, and E allele frequencies just drift.

In the case where there is an additional source of external selection acting on **A** which opposes the effect of niche construction, rising initial E allele frequency is this time associated with decreasing A allele frequency ($p(A)$) at equilibrium, since positive niche construction now favours a alleles. As predicted by the population genetic model, fixation to A occurs only for initial $p(E)$ very close to zero in non-spatial simulations. For all other initial $p(E)$, heterozygotes advantage at **A** leads to polymorphic equilibria (Fig. 3, left hand graph). Here drift due to sampling error leads to wide fluctuations in $p(A)$, measured at 750 generations. When the effect of drift is excluded (by plotting $p(A)$ against $p(E)$ *at 750 generations*), the non-spatial data are in close agreement with the predictions of a non-spatial population genetic analysis (data not shown).

The introduction of spatial interactions and a localised resource (Fig. 3, right hand graph) has a similar effect as in the previous 'no external selection' case, with clustering effects and resource-mediated coupling between **A** and **E** loci. E allele frequency tends to increase from its initial value, for all initial $p(E) > 0.1$, and there is a significant (and increasing) likelihood that E becomes fixed for E with initial $p(E) > 0.4$, with a corresponding rise in the frequency of niche-constructing aa/EE genotypes at 750 generations. The curve of mean $p(A)$ against initial $p(E)$ is essentially flat for $p(E)$ above 0.4, with test populations converging towards the high $p(E)$–low $p(A)$ region of phase space typical of non-spatial populations with very high initial

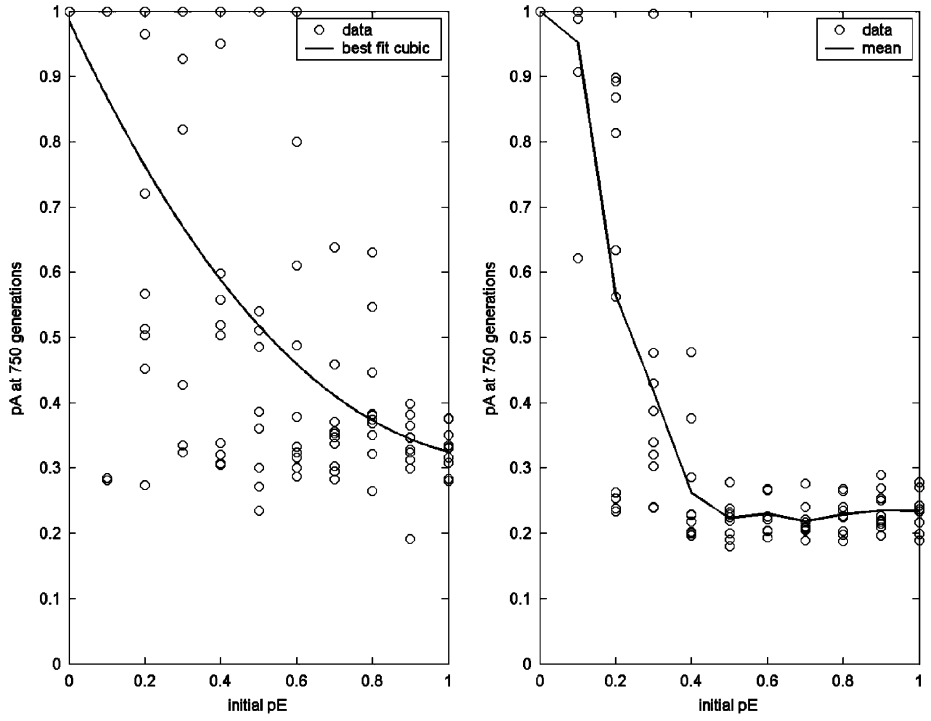

Fig. 3. Individual-based simulation of non-spatial (left) and spatial (right) populations where niche construction opposes external selection at **A**. $\alpha_2 = 1.1$; $\beta_2 = 0.9$; $\lambda_1 = 0.67$; $\lambda_2 = 0.083$; $\lambda_3 = 0.25$; $\gamma = 0$. Population = 3600 (60×60 spatial grid). Results are for 10 trial populations at each initial $p(E)$ value.

$p(E)$. As initial $p(E) \to 0$, clustering is rare and drift to $p(E) = 0$ dominates.

In short, niche-constructing (aa/EE) genotypes occur with higher frequency and across a wider range of initial conditions, when compared with non-spatial simulations. Further investigation suggests that for large populations over long timescales, fixation to E is likely to occur for the majority of populations, over an even wider range of initial conditions (data not shown).

3.1.1. Cluster formation

Strong linkage disequilibrium between niche-construction and recipient trait alleles does not by itself explain the increased frequency of niche-constructing behaviour compared with non-spatial simulations. In fact, the fixation of E alleles in otherwise unfavourable environments is seen to depend on the formation of spatial *clusters* of niche constructors (Fig. 4).

Two factors appear crucial to stable cluster formation: (i) the chance 'seeding' of a nascent niche-constructing cluster, i.e. the random presence of similar niche-constructing individuals at a close distance whose effect on local resource frequency initially favours their own propagation, and (ii) the passing of a minimum cluster size threshold, below which such clusters tend to break apart (Fig. 5). Close to the threshold, clusters may survive for prolonged periods without either expanding or disintegrating; above the threshold, clusters expand rapidly, and populations fix to the niche-constructing genotype at both loci.

The spatial distribution of resource, A and E alleles all appear to play a part in the chance seeding of a cluster. In the 'no external selection at **A**' case, raised frequencies of either A or E alleles increase the fitness of genotypes containing E or A alleles, respectively, leading to raised resource levels. This hints at a possible explanation for the relationship between probability of A fixation and population size (Fig. 1), since larger populations will, by chance, contain more regions where 'higher' concentrations of E and A alleles coincide, increasing the probability of cluster formation.

A possible rationale for the existence of a cluster size threshold concerns the behaviour of individuals close to the cluster boundary. In the 'no external selection at **A**' case, E alleles associated with a particular cluster tend to be more 'spread out' than the associated A alleles. This is because time lags in the accumulation (or depletion) of resources through niche construction mean that (i) e alleles are able to invade and spread into areas with high concentrations of A alleles and raised resource levels without significantly reducing A fitness in the short term; and similarly (ii) E alleles are able to invade and spread into areas with high concentrations of a alleles and low resource levels for a number of generations before resource levels rise to the extent that a fitness is reduced. Every cluster is therefore surrounded by a 'fringe' of relatively high numbers of heterozygote $AaEe$ individuals formed from matings between individuals either side of the boundary (see Fig. 4).

If average concentrations of E alleles in this boundary layer rise above the levels necessary for niche construction to favour the AA zygote (i.e. localised $p(E) > 0.25$), the cluster will expand. If average $p(E)$ falls below this level, the cluster will contract and disintegrate.

E alleles frequencies at the boundary depend on three factors:

1. the rate of migration, through reproduction, of E alleles from within the cluster,
2. the rate of migration, through reproduction, of e alleles from outside the cluster and,
3. the length of time that E alleles can persist at the boundary while resource levels are low (< 0.5). This in turn will depend on the precise form of the frequency-dependent fitness relations and coefficients governing resource dynamics (Eq. (1)). (For example, if the rate of resource accumulation through niche construction (λ_2) is too high, then selection pressure against E alleles (due to LD) in the boundary layer will also be high. The chance of E alleles combining with recipient trait A alleles before extinction is therefore reduced).

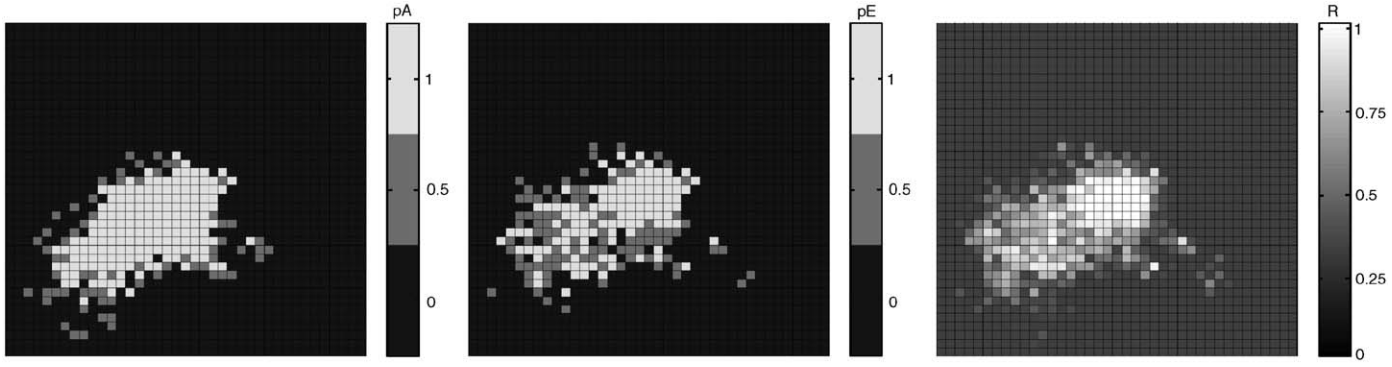

Fig. 4. Spatial simulation: evolution of a niche-constructing 'cluster' in a 30×30 population, with initial $p(E) = 0.15$, after 100 generations. Snapshot of distribution of (left) A alleles, (centre) E alleles and (right) resource. Strong linkage disequilibrium is apparent.

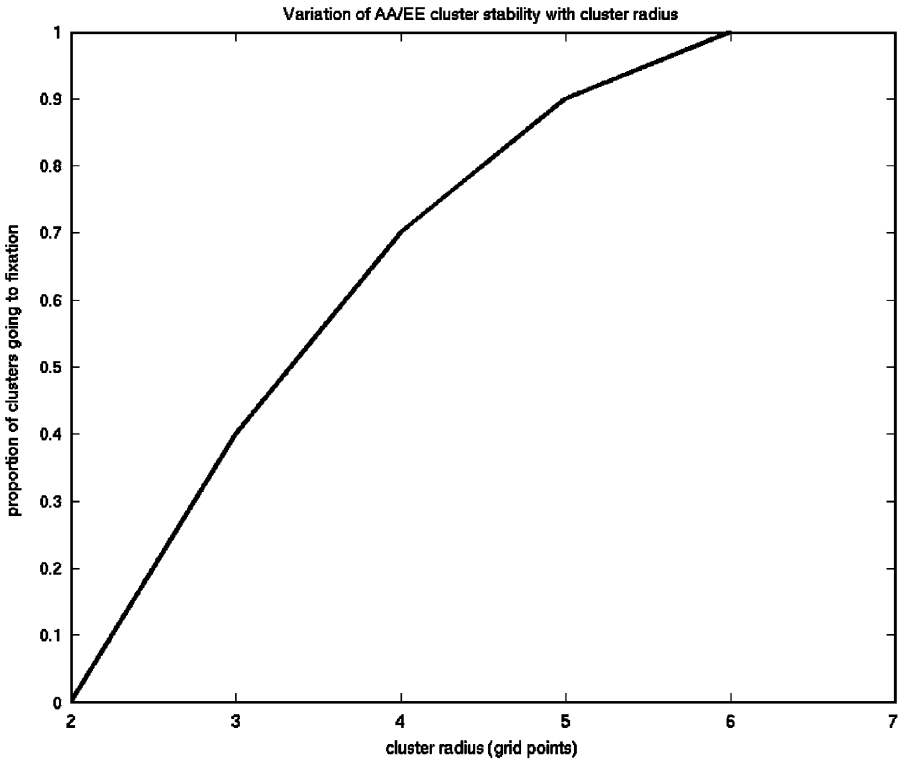

Fig. 5. No external selection at **A**: variation of stability of a niche-constructing AA/EE cluster with cluster size.

For a given set of parameters, the quantity described in 3 is finite and fixed. The migration rates described in 1 and 2 are also fixed, and are related to the average rate of movement between adjacent cells. However, for any given cluster shape existing in a majority non-niche-constructing population, the magnitude of 2 will always be greater than the magnitude of 1, purely for geometrical reasons. Consider, for example, a circular cluster of niche-constructing individuals. The 'outward-facing' surface of the boundary layer (C_a) will always be larger than the 'inner-facing' surface (C_A), adjacent to the cluster core (Fig. 6).

The ratio of these two quantities is given by

$$\frac{C_a}{C_A} = \frac{r_A + d_h}{r_A}. \tag{6}$$

Since the thickness of the boundary layer (d_h) is observed to vary very little, $C_a/C_A \to 1$ as cluster size increases ($r_A > d_h$). Thus as r_A increases, there will come a point at which the net immigration of A alleles into the boundary layer will be sufficient to combine with 'persisting' E alleles (due to the process described in 3 above) and tip resource frequency above the level required to sustain niche-constructing genotypes.

Thus, it is the ability of niche-constructing E alleles to persist for short periods of time in unfavourable (low R) environments that enables the subsequent colonisation of such environments by genotypes dependent on niche construction. Persistence times are necessarily short, since E alleles quickly drive resource levels too high for their own survival. The neighbouring presence of recipient trait A alleles is crucial to the subsequent establishing of a niche-constructing colony.

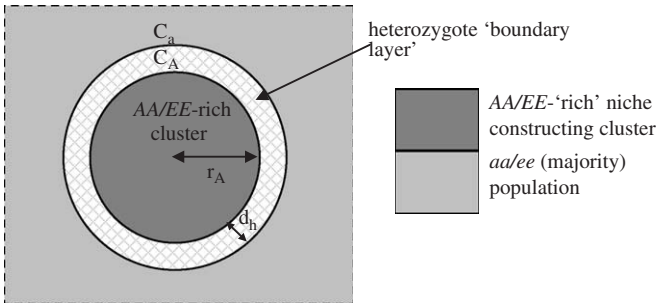

Fig. 6. Qualitative distribution of alleles within and close to the boundary of a niche-constructing cluster.

Similar reasoning can be applied to the case where external selection acts on **A**, though clustering at **A** is now more 'diffuse', since equilibria at this locus are polymorphic.

The precise figures in 1–3 are difficult to quantify. We simply suggest that, although cluster boundaries are never regular, statistically speaking, nascent clusters will resemble that shown in Fig. 6 (the niche-constructing cluster shown in Fig. 4 is a typical example). The line of reasoning sketched above then indicates that such formations will give rise to a cluster size stability threshold, beyond which cluster expansion is likely to be observed.

3.2. Two-species engineering web

For an infinite, non-spatial population with the parameters indicated in Section 2.2, Eqs. (3–5) indicate an unstable polymorphic equilibrium for all webs at initial $p(E_n) = 0.5$ ($n = 1, 2$), corresponding to a uniform resource frequency of 0.5. At this point, all three possible genotypes for each species have equal fitness (Eq. (2)), and neither species fixes for either allele at **A**. Either side of this point both species rapidly fix as either A alleles (initial $p(E_n) > 0.5$) or a alleles (initial $p(E_n) < 0.5$) dominate.

Results from the individual-based simulation of finite-population, non-spatial engineering webs is presented in the left hand graph of Fig. 7. Engineering webs 1 and 2 (where one or both species are capable of positive niche construction) demonstrate behaviour similar to that expected for an infinite population, though finite-population effects mean that times to first fixation are longer (for initial $p(E_n) \neq 0.5$), and populations with initial $p(E_n) = 0.5$ now fix to one or other allele since stochastic sampling error tips $p(E)$ (and resource frequency) away from 0.5.

With engineering web 3 (where each species can engage in niche construction that opposes the other), mean time to fixation at $p(E_n) = 0.5$ is now even more rapid, since, the roughly half of populations that first fix to the *aa* diplotype converge more rapidly as a result of negative niche construction. For initial E allele frequencies above and below this point, equilibrium resource levels are, respectively, higher or lower than 0.5, with fixation to A or a. Lack of symmetry is explained by the fact that the relationship between resource level and $p(E)$ is no longer linear (Eq. (5)).

Results from individual-based *spatial* simulations are presented in the middle graph of Fig. 7. Behaviour of webs 1 and 2 is similar to the non-spatial case, except that clustering effects mean that times to first fixation are slightly higher in those marginal environments where niche-constructing alleles would otherwise be lost. This effect is particularly pronounced close to initial $p(E) = 0.5$. Here niche-constructing clusters frequently give rise to the (comparatively slower) fixation of A_nA_n/E_nE_n (positive niche-constructing) genotypes in the roughly half of cases where resource levels drop below 0.5 due to sampling error.

The behaviour of web 3 however shows significant quantitative and qualitative differences compared with non-spatial simulations. Times to first fixation are now considerably higher over a wide range of initial E allele frequencies from approximately 0.2–0.4 (when compared with webs 1 and 2). This is interesting since opposing forces of positive and negative niche construction are no longer balanced away from $p(E) = 0.5$, suggesting that the observed 'stabilising' effect is to some extent independent of the choice of parameters in Eq. (5).

The persistence of polymorphic populations in both species is due to the formation of quasi-stable, spatial structures with distinctive topologies (see e.g. Fig. 8), which are able to maintain a dynamic equilibrium between colonies of positive and negative niche constructors. It is only when initial E allele frequencies are very high or very low that one species rapidly fixes to its niche-constructing genotype, while the other loses niche-constructing alleles before being able to form 'defensive' clusters. This is because resource frequencies are too high or too low for clusters of negative or positive niche constructors to establish.

This 'stabilising' effect is enhanced in larger populations, with both an increased *range* of initial E allele frequencies over which stability is enhanced, and increased *times* to first fixation (right hand graph of Fig. 7). Non-spatial populations show no such variation (data not shown).

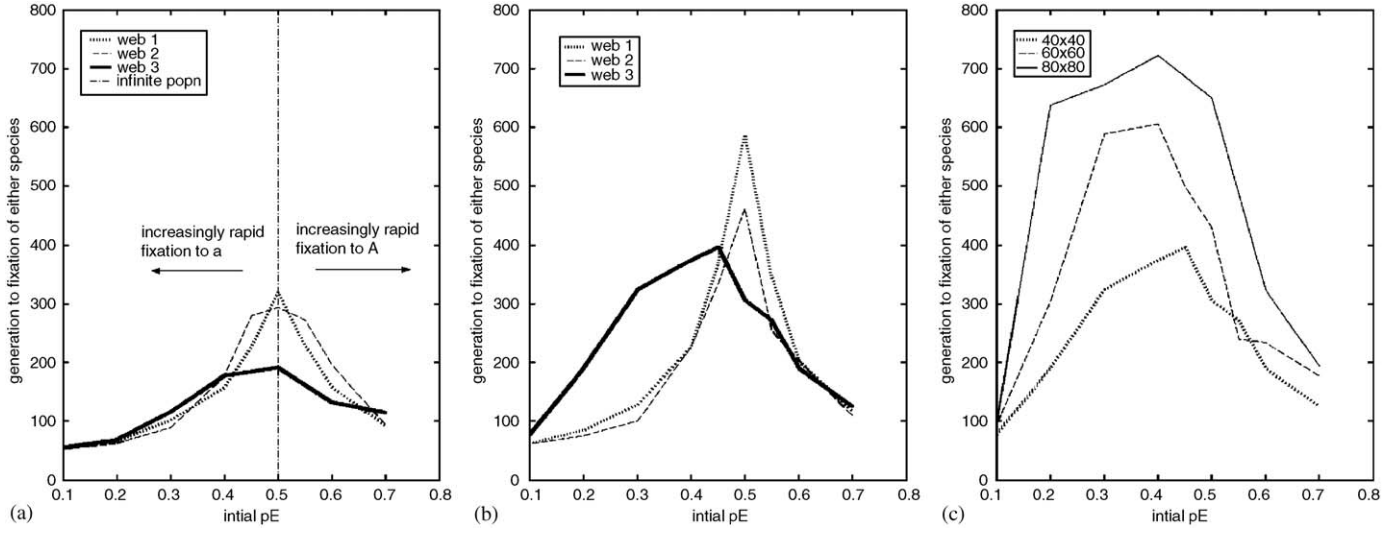

Fig. 7. Comparison of three niche-constructing engineering webs for (a) non-spatial and (b) spatial (40×40) populations of 2×1600 individuals. Engineering webs 1 and 2: $\lambda_1 = 0.7$; $\lambda_2 = 0.1$; $\lambda_3 = 0.1$. Engineering web 3: $\lambda_1 = 0.7$; $\lambda_2 = 0.2$; $\lambda_3 = 0.1$; $\gamma = 0.286$. (c) Engineering web 3—effect of varying population size. In each case, initial $p(E_1) =$ initial $p(E_2)$. Results are average across 10 trial populations.

3.2.1. Web 3—understanding web dynamics

When S_1 and S_2 are able to, respectively, increase and decrease resource frequency, a 'push–pull' effect is evident. At intermediate frequencies of niche-constructing alleles, an initially random distribution of alleles at $\mathbf{A_nE_n}$ quickly coalesce into separate niche-constructing clusters that occupy distinct regions in space. The expansion of either species' niche-constructing genotype is then limited by the niche-constructing activity of the other via their opposing effects on local resource frequencies, resulting in a relatively stable dynamic equilibrium between the two species. The situation is symmetric—each cluster is dominated by one niche-constructor species and effectively "defended" against invasion from the other niche-constructor species, owing to LD and local control of resource levels. The result is less a case of competitive exclusion than of 'coordination', where each species reduces for the other the advantage of occupying spatial regions inhabited by niche constructors of the first. These two factors can only take place if the populations are spatially distributed.

With relatively high concentrations of E alleles, a banded structure is the most stable geometry (Fig. 8). Protrusions of one niche-constructing genotype into the area occupied by the other are seen to rapidly disappear. Such behaviour ties in with the cluster expansion hypothesis given in the preceding section. When E allele concentrations are low in one (or both) species, roughly circular clusters of niche constructors are observed to form.

As populations increase in size, the probability that the necessary conditions for cluster formation will obtain, increases for both species. This is the case even where the prevailing resource frequency is unfavourable to one of the niche-constructing genotypes—e.g. at relatively high or low E allele concentrations. The number of populations where only one species of niche constructor dominates is then reduced, and the stabilising 'push–pull' effect becomes more common over a wider range of initial $p(E)$. Once again the fact that web 3 is able to sustain polymorphisms across a *range* of $p(E)$ suggests that the observed stabilising effect may apply across a wider range of parameter space than has been considered here.

At the margins, it is the chance seeding of stabilising niche-constructing clusters of sufficient size that determines whether or not polymorphic populations are able to establish. It is therefore interesting to study the case where the chance of establishing such clusters is increased—for example where niche construction enhances the resource frequency-dependent fitness of an already established recipient trait. Preliminary investigations suggest that web stability is then further enhanced, with polymorphisms being sustained in both populations, over a wider range of initial conditions (data not shown).

Type-3 webs also show an interesting response to the introduction of external environmental perturbations in the form of periodic fluctuations in the coefficient of independent renewal, λ_3. The rate of change of resource frequency is reduced on both upward and downward phases of the

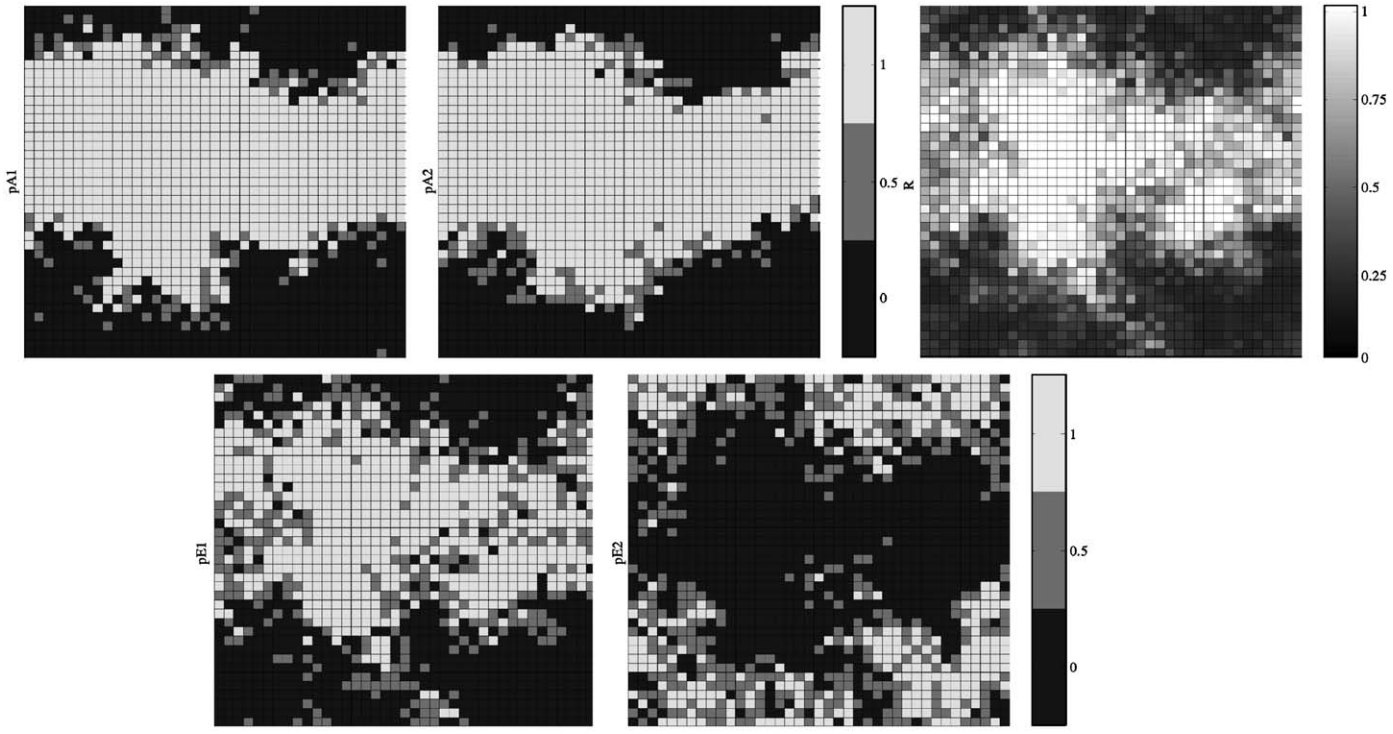

Fig. 8. Engineering web 3: two-species, quasi-stable, band-like structure. 40×40 spatial population with initial $p(E_1) =$ initial $p(E_2) = 0.4$ after 180 generations. This structure is common when both types of niche-constructing alleles are initially abundant. Regions of S_1 AA/EE positive niche constructors, and S_2 aa/EE negative niche constructors are clearly visible, as is the corresponding resource distribution.

resource cycle, when compared with the 'external' fluctuations. This is because the niche-constructing activity of the declining species in each case tends to restrict the growth of the increasing species, resulting in damped oscillations in resource frequency. In this way a selective response at the species level gives rise to a form of 'homeostatic control' at the ecosystem level. Webs of type-1 and type-2 show no such effect (data not shown).

4. Discussion

Two cases were studied in the single species model: one in which the evolution of a recipient trait depends only on the frequency of an environmental resource; and the other in which this evolution depends additionally on an external source of selection acting in opposition to resource frequency-dependent selection. In both cases, spatial effects lead to (i) a significant increase in the frequency of niche-constructing behaviour, with niche-constructing E alleles frequently going to fixation; and (ii) an increase in the range of initial E allele frequencies over which such behaviour occurs. As a consequence, the recipient trait associated with raised levels of niche construction is also observed to fix (no external selection) or to increase in frequency (external selection), over a wider range of initial $p(E)$.

Two major factors drive this change. First is the presence of a localised resource distribution which gives rise to strong linkage-disequilibrium between niche-construction and recipient trait alleles. LD occurs because individual (**A**-locus dependent) fitness is affected by an individual's (**E**-locus dependent) niche-constructing behaviour, due to the latter's effect on local resource levels. Linkage equilibrium cannot be maintained because spatial structure effectively prevents mixing at a global level. This positive genotype-environment feedback loop is key to the driving of recipient and niche-constructing traits to fixation. In this respect, niche construction is to be contrasted with other sources of epistatic interaction—here gene linkage is dependent on the population's impact on external, environmental resource frequency.

Second is the establishment of small spatial clusters of niche-constructing alleles. The subsequent expansion of these clusters is due to the ability of niche-constructing alleles to persist, for short periods of time, in initially unfavourable environments. This then facilitates the colonisation of genotypes where the recipient trait associated with niche-constructing alleles would otherwise be lost.

This sequence of events points to a possible mechanism by which an initially rare recipient trait whose existence depends on niche construction could establish in an otherwise hostile environment. The same mechanism could also apply to the spread of an established niche-constructing colony. This may have parallels with empirical and theoretical work on bacterial 'proactive invaders'—organisms that appear to manufacture conditions to facilitate their own invasion into a novel environment (currently being studied by Francois Taddei and colleagues at the Necker Medical School, Paris—K.N. Laland, personal communication). Crucial to such a process in natural scenarios will be the relative timescales of local diffusion of niche-constructing effects on resources and the survival rate of invaders if niche construction were not present. This effect is likely to be more pronounced in cases where the ecological resource has a naturally large scale of diffusion (not the situation modelled here), and yet remains contained spatially so that niche constructors still reap the benefits of their activity (e.g., in ponds and lakes).

Although the results presented here necessarily refer to a very limited region of the set of all possible parameters, the observed phenomena may in fact have more general applicability. Consider for example the 'external selection at **A**' case. This could represent the evolution of a niche-constructing trait which acts to dampen the effect of statistical variations in a resource due to external environmental perturbations (Laland et al., 1996; Lewontin, 1983). Since such adaptations are in a sense fundamental to the maintenance of life, the kinds of evolutionary dynamics observed here in the environmentally mediated evolution of recipient and niche-constructing traits are likely to be observed across a wide range of prevailing environmental conditions. This idea could be further explored by modelling the co-evolution of $p(E)$ and λ_2 (or γ) with the other parameters controlling independent depletion and renewal.

The introduction of spatial structure also has a significant effect in two-species models. In the 'no external selection' case, where two species niche construct in a way that alters resource frequency in the same direction, there is a 'bootstrapping' effect, with each species able to drive both itself and the other to fixation of niche-constructing behaviour and the recipient trait that depends on it.

The situation where each species instead engages in niche-constructing activity that opposes the other leads to an increased chance that polymorphisms will be maintained in each. This situation can arise whenever a species exploits a resource produced by another species (Jones et al., 1994, 1997; Odling-Smee et al., 2003). Our model suggests that clustering effects can sustain polymorphisms over a wide range of initial E allele frequencies—an outcome not predicted by a non-spatial, mixed-medium approximation. This is particularly interesting, since it suggests that opposing forces of positive and negative niche construction need not be finely balanced for stabilising effects to occur. The stability of polymorphisms is also observed to increase with increasing population size, and in the situation where interactions occur between two *established* niche-constructing species. Since equilibrium resource frequency depends on a combination of the frequency of niche-constructing alleles, and the rates of individual-based and independent resource accumulation and depletion (Eq. (5)), the indication is that stabilising effects are likely to be observed across a wider range of the

parameter space than has been considered here. This warrants further investigation as such effects could contribute to the presence of widespread genetic variation in natural populations.

The two-species evolutionary dynamics described above may have some relevance to resource dynamics and species diversity in multi-species ecosystems. Models of engineering webs 1 and 2 show how the activities of niche-constructing species can combine to drive a particular recipient trait to fixation. The effects of such positive feedback can produce instability in the long term (Robertson, 1991). In contrast, the model of engineering web 3 shows how interactions between species and their environment can also promote the kind of equilibria which may render an ecosystem more orderly through the preservation of otherwise deleterious alleles (Laland et al., 1999). In this case, positive feedback loops generated by niche construction may generate increased variation (and response to selection) within an ecosystem by maintaining polymorphisms, with the potential to affect long-term evolutionary and ecological dynamics (Crespi, 2004). Furthermore, the two-species system gives rise to a form of 'environmental homeostasis' in which the effects of perturbations in the 'external' environment are reduced. This is reminiscent of other models demonstrating global regulative effects (Lansing et al., 1998; Lenton, 1998).

Taken together, these findings suggest that where the frequencies of key resources are affected by the niche-constructing activities of more than one species, stable polymorphisms may in certain circumstances be more likely to be maintained. Under the present two-species model, stabilising effects are only observed in the presence of both positive and negative niche constructors. It would be interesting to study the effect of introducing more complexity into the web, for example by adding more species with a range of positive and negative niche-constructing effects. It is possible, contrary to what might be expected, that such added complexity might still lead to enhanced stability, with not two but many species 'pulling' resource levels in different directions over a range of environmental conditions. The suggestion is speculative, but a similar mechanism, known as 'rein control' (Clynes, 1969) has been implicated in other cases of biological stability (Saunders et al., 1998). In an ecological context, Harvey (2004) has suggested that environmentally mediated interactions between large numbers of separate systems, roughly the equivalent of extending our model to the case of multiple species, will tend to exert a degree of environmental control at the whole-system level. Provided that the numbers of 'pushers' and 'pullers' are roughly balanced, this stabilising effect is found to be largely independent of precise parameter values. The possibility that a similar effect might be observed in a multi-species version of the current model is worthy of further investigation.

The magnitude and complexity of real species–resource interactions is not well understood. The kinds of stabilising effects described here are most likely to occur where a resource plays a significant part in the lifecycles of many different species, or where multiple species exert some degree of (local) control over the same resource. Examples might include water, heat and light, and perhaps those 'substrates amenable to biogeomorphic action', such as the digging, scraping, burrowing, boring and chemical erosion of soils and sediments (Jones et al., 1997). On a global scale, such resources would certainly include the major biogenic gases, the regulation of which is central to the stability of the biosphere (Lenton, 1998). Any assessment of the significance of such effects in real ecosystems would require the gathering of empirical evidence on multi-species–resource interactions. Such evidence would go beyond data about species abundance and trophic relations, to include some estimation of non-biotic resources and non-trophic interactions, making it difficult to obtain.

The sensitivity of our results to variations in the strength of the association between genetic and ecological inheritance might also be explored. The current model uses a particularly strong form of ecological inheritance, with offspring fitness determined entirely by parental resource levels. In this respect our model bears some resemblance to studies of the influence of environmentally mediated maternal effects on gene frequencies (e.g. Wolf, 2000; Mousseau & Fox, 1998), although in the current model offspring resource levels are additionally affected by multi-agent, multi-generational and external factors. It would be interesting to explore the effect of weakening the association between parental and offspring resource by, for example, varying the range and spatial decay of parental niche-constructing effects, and the magnitude of their effect on offspring fitness or by introducing different modes of resource transport. We would expect any weakening of the association between genetic and ecological inheritance to produce a similar reduction in spatial clustering effects and gene linkage. The robustness of the clustering mechanism to variation of the mating and dispersal kernels might also be explored further.

Parsimony with the original model meant that the possibility of extinctions (local or global) was excluded. More sophisticated models including population dynamics and different individual states may uncover novel phenomena requiring a qualification of the results presented here. A strong form of density dependence is also implicit in our model, since cell occupancy is restricted to a single individual (of each species in the two species model). The effect of multiple cell occupancy merits further investigation. Finally, the relative merits of discrete, lattice models, as opposed to patchy or continuous models for the introduction of spatiality is well covered in the literature Durrett & Levin, 1994b; Van Baalen & Rand, 1998). Alternative methods of

implementing spatiality may well produce effects not apparent in the current model.

While our model inevitably contains a number of simplifying assumptions, we suggest that it serves to illustrate some of the complex, non-linear, ecological and evolutionary interactions that may occur in spatial populations of niche constructors. Moreover, the model's flexibility affords further comparative study by exploring the extensions suggested above.

Acknowledgments

This work has benefited from very constructive feedback by two anonymous reviewers. We would like to acknowledge Kevin N. Laland, John Odling-Smee and Marcus W. Feldman for their insightful comments.

Appendix

An estimate of LD is given by a diploid version of Lewontin's disequilibrium coefficient (Lewontin and Kojima, 1960):

$$\mathrm{LD} = \mathrm{p}_{AE}\mathrm{loc} - \mathrm{p}_{A}\mathrm{loc} \times \mathrm{p}_{E}\mathrm{loc}, \tag{A.1}$$

where p_{AE}loc is the frequency of individuals with at least one E allele at **E**, and at least one A allele at **A**; and p_{A}loc and p_{E}loc are the frequencies of individuals with at least one A allele, and the frequency of individuals with at least one E allele, respectively. For a large population with no linkage between **A** and **E**, i.e. with a random distribution of A and E alleles; $\mathrm{p}_{A}\mathrm{Eloc} = \mathrm{p}_{A}\mathrm{loc} \times \mathrm{p}_{E}\mathrm{loc}$, and $\mathrm{LD} = 0$.

The relationship between initial p_A, p_E, p_{A}loc and p_{E}loc for an initially random distribution of alleles can be derived from simple probability theory:

$$\begin{aligned}\mathrm{p}_{A}\mathrm{loc} &= p_{AA} + p_{Aa} + p_{aA} \\ &= [p_A \times p_A] + 2 \times [p_A \times (1 - p_A)], \text{ and}\end{aligned}$$

$$\begin{aligned}\mathrm{p}_{E}\mathrm{loc} &= p_{EE} + p_{Ee} + p_{eE} \\ &= [p_E \times p_E] + 2 \times [p_E \times (1 - p_E)].\end{aligned}$$

Substituting $p_A = 0.5$, and $p_E = 0.15$ gives values for initial p_{A}loc and initial p_{E}loc of 0.75 and 0.28, respectively—very close to the measured initial values illustrated in Fig. 2.

It's important to note that although LD varies between 0 and 0.2 in the spatial case, this actually represents a situation in which almost all A and E alleles in the population are linked. This follows from Eq. (A.1). For example, at generation 100, the measured values for $\mathrm{p}_{A}\mathrm{loc} = 0.279$, $\mathrm{p}_{E}\mathrm{loc} = 0.278$ and the number of linked loci, $\mathrm{p}_{AE}\mathrm{loc} = 0.244$; i.e. a large majority of (the minority of) individuals which have at least one E allele also have at least one A allele (see also Fig. 4). For a population in linkage equilibrium, the expected frequency of linked loci under random distribution is $\mathrm{p}_{A}\mathrm{loc} \times \mathrm{p}_{A}\mathrm{loc} = 0.0776$. Thus the figure for LD (from (x)) is 0.244—0.0776 = 0.166, represents a population in which more than 3 times as many individuals have linked A and E alleles, compared with a population in linkage equilibrium.

References

Brodie, E.D., 2005. Caution: niche construction ahead. Evolution 59 (1), 249–251.

Clynes, M., 1969. Cybernetic implications of rein control in perceptual and conceptual organization. Ann. NY Acad. Sci. 156, 629–670.

Crespi, B.J., 2004. Vicious circles: positive feedback in major evolutionary and ecological transitions. Trends Ecol. Evol. 19 (12), 627–633.

Di Paolo, E.A., 2000. Ecological symmetry breaking can favour the evolution of altruism in an action-response game. J. Theor. Biol. 203, 135–152.

Durrett, R., Levin, S., 1994a. The importance of being discrete (and Spatial). Theor. Popul. Biol. 46, 363–394.

Durrett, R., Levin, S., 1994b. Stochastic spatial models: a user's guide to ecological applications. Philos. Trans. R. Soc. Lond. B 343, 329–350.

Galloway, L.F., 2005. Maternal effects provide phenotypic adaptation to local environmental conditions. New Phytol. 166, 93.

Harvey, I., 2004. Homeostasis and rein control: from daisyworld to active perception. In: Proceedings of the Ninth International Conference on Simulation and Synthesis of Living Systems, ALIFE 9, Boston, September 12–15th, 2004. MIT Press.

Hector, A., et al., 2000. No consistent effect of plant diversity on productivity-response. Science 289, 1255a.

Holt, R.D., 1995. Linking species and ecosystems: where's Darwin?'. In: Jones, C.G., Lawton, J.H. (Eds.), Linking Species and Ecosystems. Chapman & Hall, London.

Huston, M.A., Aarssen, L.W., Austin, M.P., et al., 2000. No consistent effect of plant diversity on production. Science 289, 1255.

Jones, C.G., Lawton, J.H., Shachak, M., 1994. The population dynamics of ecosystem engineers. Oikos 76, 273–283.

Jones, C.G., Lawton, J.H., Shachak, M., 1997. Positive and negative effects of organisms as physical ecosystem engineers. Ecology 78 (7), 1946–1957.

Krakauer, D.C., Pagel, M., 1995. Spatial structure and the evolution of honest cost-free signalling. Proc. R. Soc. Lond. B 260, 365–372.

Laland, K.N., Odling-Smee, F.J., Feldman, M.W., 1996. The evolutionary consequences of niche construction: a theoretical investigation using two-locus theory. J. Theor. Biol. 9, 293–316.

Laland, K.N., Odling-Smee, F.J., Feldman, M.W., 1999. Evolutionary consequences of niche construction and their implications for ecology. Proc. Natl. Acad. Sci. USA 96, 10242–10247.

Lansing, J.S., Kremer, J.N., Smuts, B.B., 1998. System-dependent selection, ecological feedback and the emergence of functional structure in ecosystems. J. Theor. Biol. 192, 377–391.

Lenton, T.M., 1998. Gaia and natural selection. Nature 394, 439–447.

Lewontin, R.C., 1983. Gene, organism and environment. In: Bendall, D.S. (Ed.), Evolution from Molecules to Men. Cambridge University Press, Cambridge, UK.

Lewontin, R.C., Kojima, K., 1960. The evolutionary dynamics of complex polymorphisms. Evolution 14, 450–472.

Loreau, M., Naeem, S., Inchausti, P. (Eds.), 2002. Biodiversity and ecosystem functioning. Oxford University Press, Oxford.

Mousseau, T.A., Fox, C.W., 1998. The adaptive significance of maternal effects. Trends Ecol. Evol. 13, 403–407.

Mutch, R.W., 1970. Wildland fires and ecosystems: a hypothesis. Ecology 51, 1046–1051.

Nakamaru, M., Matsuda, H., Iwasa, Y., 1997. The evolution of cooperation in a lattice-structured population. J. Theor. Biol. 184, 65–81.

Odling-Smee, F.J., 1996. Niche construction. Am. Nat. 147, 641–648.

Odling-Smee, F.J., Laland, K.N., Feldman, M.W., 2003. Niche construction. The neglected process in evolution. Monographs in Population Biology, vol. 37. Princeton University Press, Princeton.

Robertson, D.S., 1991. Feedback theory and Darwinian evolution. J. Theor. Biol. 152, 469–484.

Reusch, T.B.H., Ehlers, A., Hämmerli, A., Worm, B., 2005. Ecosystem recovery after climatic extremes enhanced by genotypic diversity. Proc. Natl. Acad. Sci. USA 102, 2826–2831.

Shachak, M., Jones, C.G., Granot, Y., 1987. Herbivory in rocks and the weathering of a desert. Science 236, 1098–1099.

Saunders, P.T., Koeslag, J.H., Wessels, J.A., 1998. Integral rein control in physiology. J. Theor. Biol. 194, 163–173.

Schwilk, D., 2003. Flammability is a niche construction trait: canopy architecture affects fire intensity. Am. Nat. 162, 725–733.

Tilman, D., Kareiva, P. (Eds.), 1997. The Role of Space in Population Dynamics and Interspecific Interactions. Princeton University Press, Princeton.

Van Baalen, M., Rand, D.A., 1998. The unit of selection in viscous populations and the evolution of altruism. J. Theor. Biol. 193, 631–648.

Van Valen, L., 1973. A new evolutionary law. Evol. Theor. 1, 1–30.

Wardle, D.A., et al., 2000. Biodiversity and ecosystem function: an issue in ecology. Bull. Ecol. Soc. Am. 81, 235–239.

Wolf, J.B., 2000. Gene interactions from maternal effects. Evolution 54 (6), 1882–1898.

Wolf, J.B., Brodie, E.D., Wade, M.J. (Eds.), 2000. Epistasis and the Evolutionary Process. Oxford University Press, Oxford.