

Microbial Gaia: A new model for the evolution of environmental regulation

Hywel T.P. Williams & Timothy M. Lenton

Earth System Modelling Group, School of Environmental Sciences,

University of East Anglia, Norwich NR4 7TJ UK

{h.williams,t.lenton}@uea.ac.uk

September 19, 2007

Whether free-living, aggregated in biofilms, or packaged into higher organisms, microbes are the most prevalent form of life on Earth. They make up more than half the biomass on Earth and occupy every conceivable niche in the biosphere. Their metabolic diversity and adaptability means that microbes can live where other life is impossible, in extreme environments such as deep-sea thermal vents, beneath the Arctic and Antarctic ice-caps, and even inside other organisms. Microbial life is heavily implicated in all the biogeochemical cycles, dwarfing the contribution of multicellular organisms — which in any case have microbes at the root of their metabolism. Mitochondria and chloroplasts, the main energy-producing organelles in animal and plant cells, both have evolutionary origins as once-free bacteria annexed into the greater whole. (Of course, the architect of this theory of endosymbiosis is none other than the co-founder of Gaia theory, Lynn Margulis.) Inside the body (but outside of the cell) the microbial gut flora of many higher species is an essential part of the metabolic processes of the host organism; cows could not digest cellulose with the incredible microbial diversity of the rumen [8], while humans rely on bacteria to digest various carbohydrates and lipids [9]. In short, Gaia's metabolism is microbial.

In recent years technological advances have led to a much greater understanding of the microbial world. We now understand (a little) how microbes metabolise, grow, and reproduce, and comprehend (a bit) how microbial genetics works, complete with lateral gene transfer

that allows changes in cell DNA during the lifetime of the organism, a mechanism that seems quite alien compared to genetics in higher organisms. But our understanding of how microbial ecology interacts with the environment is one of the least understood, and most important, areas of current research. By virtue of their metabolic diversity, microbes are able to perform complex chemical transformations when multiple species collaborate. Artificially assembled microbial communities have been used to perform multi-step reactions for bioremediation of environmental pollutants (e.g., breaking down 3-chloroaniline [10]) and leaching of metal ores (e.g., copper [8]). The collaborative nature of these manufactured communities mimics (in spirit if not complexity) the tight integration of naturally occurring microbial communities such as biofilms, in which many species form three-dimensional physical structures on surfaces that allow regulation of the microenvironment. Biofilms (also known as microbial mats) have been around for a very long time. Fossilised microbial mats called stromatolites have been found that are 3.5 billion years old [8]. This long history of environmental alteration suggests that understanding the dynamics of microbial ecosystems will be key to understanding Gaia. The question of how multiple microbial species can cooperate in biofilms to collectively improve their local environment, in apparent contravention of selfish Darwinian principles, has obvious parallels in the Gaia debate.

The fundamental question of how system-wide regulation can emerge from evolving lower-level components, be it at a global scale as with Gaia, or a smaller one as with ponds or biofilms, remains to be answered. It seems plausible that similar principles are at work at all scales, so if you can understand the microcosm, you at least have a chance of understanding the macrocosm. To this end, we have been developing a new mathematical model of microbial evolution [13, 16]. The Flask model simulates a well-mixed flask of liquid in which is suspended a population of ‘microbes’. Each microbe has genetic loci specifying traits for nutrient consumption/release and its interaction with the abiotic environment. Every microbe has a preferred abiotic environmental state in which its metabolic rate is maximised, while microbe growth alters the abiotic environment via cost-free by-products of metabolism. Microbes grow by converting consumed nutrients into biomass, and reproduce by splitting when a certain biomass threshold is reached. During each reproduction event, the parental genotype is inherited by the offspring individual with a low probability of mutation at each locus. Microbes interact with their shared environment, removing and adding materials dur-

ing metabolism at a rate determined by the state of the environment. The ecological success of a microbe species is determined by its fit with the environment, in terms of nutrient supply and abiotic conditions. Successful species will grow and diversify, while unsuccessful species will starve and die off. Thus we have a highly abstracted, but versatile, mathematical model of microbial growth and evolution in a changing environment.

The Flask model is designed to address some of the key questions about the evolution of regulation. The original Daisyworld model showed that global regulation does not require top-down control or teleology [12]. Later variations of the model have since shown that in some circumstances natural selection is compatible with regulation, while in others it is not; broadly speaking regulation occurs when daisy albedo can adapt, but not with unconstrained adaptation of optimal growth temperature [17]. However, one key assumption of the Daisyworld model means that it represents a special case in evolutionary terms. In Daisyworld, properties that confer local fitness also contribute to global regulation. If this assumption is relaxed, regulation does not emerge [13, 17]. In the real world, it is conceivable that traits can arise which offer individual-level selective advantage but have a destabilising effect on the global environment. Most commonly, traits which alter the environment may be selectively neutral for individuals, because most biogeochemical cycles involve atmospheric and oceanic transports which have mixing rates sufficient to remove the local effects of environmental alteration.

In many cases, environment-altering traits are simply by-products of other traits that are selected [11]. For example, mammals breathe out CO_2 when they respire, but this a simple consequence of respiration; having a respiring metabolism may be selected for, but expelling CO_2 as waste is not. The question for Gaia thus becomes, ‘how can regulation arise from environment-altering traits of organisms that are selectively neutral or that are by-products of other selected traits?’. The by-product view of regulation avoids the problem of cheats that might destroy regulation, because by-products are not altruistically contributed to some collective regulatory enterprise. But other theoretical problems arise with the concept of by-product regulation, such as why we should expect to see any regulation if there is no adaptive mechanism to select for ‘good’ regulatory traits. It is possible that regulating negative feedbacks can emerge from the local interactions of organisms with their environment [4, 7], but equally possible that destabilising positive feedbacks could appear [2, 3].

In the Flask model, environment-altering traits may contribute to or oppose regulation and are selectively neutral at the individual level. In a single well-mixed flask, we observe the reliable emergence of nutrient recycling loops [1, 16]. Microbes are constrained not to be able to consume a nutrient that they also release as waste, but it is common for one species to excrete a nutrient that another consumes. Although recycling loops may involve many different species, the simplest loop would involve just two species, where species 1 consumes A and releases B, and species 2 consumes B and releases A, as with (e.g.) oxygenic photosynthesisers and aerobic respirers in the real world. In the Flask model, the nutrient consumption of a microbe is scaled down if the nutrients it requires are limited, so it makes ‘good evolutionary sense’ to depend on as few nutrients as possible in order to reduce the risk of limitation. Given that populations grow exponentially in times of plenty, nutrient limitation is the normal state of the system, so selection pressure on consumption patterns is a powerful adaptive force. We commonly observe the evolution of single-nutrient consumers, which form ecosystems where nutrients are recycled many times before being lost to the system. The theoretical carrying capacity of the system would be reached when all supplied nutrients are consumed and efficiently recycled, so that the only loss of material is the unrecoverable ‘heat loss’ due to metabolic inefficiency. In practice the system typically adapts to operate close to this limit.

With single well-mixed flask ecosystems, we also observe a number of characteristic dynamics in the interaction between the population and the abiotic environment [16]. When abiotic preferences are allowed to evolve, the population loosely tracks the slowly changing abiotic environment, as individuals that are unsuited to current conditions are displaced by better-adapted competitors. When preferences are fixed, the population can drive the environment into inhospitable conditions to which it is unable to adapt, causing a fall in population sizes. In most cases this is a gradual process which is unlikely to drive the community completely extinct, because negative feedbacks on growth [4] reduce the size of ‘harmful’ populations as the environment worsens. However, there are circumstances where ‘rebel’ species can cause a rapid environmental change resulting in total extinction. Populations are typically converged around a shared preferred environmental state, either because this is enforced (with fixed preferences) or because of genetic convergence of the relevant traits (with evolvable preferences). This makes the population vulnerable to sudden environmental

change. This can occur if a mutant appears that can exploit stocks of previously unutilised nutrients, allowing it to grow rapidly. Since abiotic environmental effects are by-products of growth, the rapidly expanding species will cause a large and sudden shift in the state of the abiotic environment away from the state to which the community is adapted. Any species which cannot adapt quickly enough – including the rebel species causing the problem – will go extinct. While diversity of environmental preferences might allow some species to survive, fixed or converged preferences mean that it is possible for the entire community to go extinct simultaneously in this way. (As an aside, this pattern of endogeneously caused extinction is rather alarming when compared to a certain real-world species who are causing rapid environmental change by exploiting a previously unavailable energy source...)

To address the key outstanding question of how (or if) we can reconcile global regulation with local adaptation, we have recently developed a spatial version of the Flask model, in which a number of flasks are connected in a ring topology [15]. Each flask is locally well-mixed as before, but in the new model we implement a slow rate of diffusive mixing between neighbouring flasks. This creates a simple scheme that gives properties of spatial structure and environmental heterogeneity, which we feel may be crucial to the appearance of environmental regulation. The liquid transferred between flasks carries with it materials and microbes. Material transfer (nutrients and abiotic factors) reduces any between-flask environment gradients. Microbe transfer between local populations carries genetic information that changes species composition and thus alters the nature of each community's collective interaction with its local environment. For simplicity, in the current version of the spatial model we impose universal fixed environmental preferences on all microbes.

One obvious change that occurs when we move from the single well-mixed flask to the multi-flask spatial system is a marked increase in robustness against extinction [15]. Endogeneously created extinctions are relatively frequent in the single-flask system, but are virtually eliminated in a similar spatial system. The main reason for this is that spatial structure creates environmental heterogeneity. If a rebel species appears that causes a local extinction, it is unlikely to spread quickly enough to cause a global extinction since the slow rate of between-flask mixing acts as a brake. By the time a destructive species reaches locations around the ring that are far from its origin, it is likely that the original denuded location will have been recolonised from a surviving neighbour population. This time lag

afforded by spatial structure also increases the time available for other species to adapt to the environmental changes caused by a rebel species. These refugia-recolonisation dynamics form an additional mechanism for robustness that operates alongside the negative feedbacks on growth that already reduce the likelihood of endogeneously caused extinctions.

Robustness from spatial structure is not a new concept and has been observed in a variety of systems. Of more interest to Gaia researchers is the observation that once life becomes established in the spatial model, the environment is consistently regulated to habitable conditions, even in the face of severe external perturbations. To understand how this regulation occurs, we must first of all note that the increased metabolic throughput resulting from nutrient recycling means that the biota in the model are the main control on the abiotic environment, which at carrying capacity dominates over the external flux of abiotic factors. The second key observation is the existence of feedbacks on growth [4]. As previously noted, these feedbacks on growth often prevent a population from driving itself extinct — since the worsening environment reduces the size of the population causing the environment-degrading effect — but they also play a role in achieving system-wide regulation by creating a form of between-flask competition based on environmental factors.

When the abiotic environment is favourable, populations expand exponentially to reach carrying capacity and selection pressure on nutrient consumption patterns is the main adaptive force. Genes controlling environment-altering traits are not selected and are subject to genetic drift, while the abiotic environment also drifts in response to the stochastic forcing from the biota. However, when the environment drifts far enough from the optimum conditions for growth, metabolic rate is reduced, microbes begin to starve, and nutrients build up in the environment. The abiotic environment becomes the key factor limiting growth and traits affecting the environment are now subject to selection pressure. Each flask is locally well-mixed, so no microbe can gain selective advantage over its local competitors via its environmental alteration. However, the collective metabolic action of a local population on its environment can improve or degrade growing conditions, affecting the size of the population. Since the abiotic environment is limiting, any local population that improves its environment will grow rapidly and expand to colonise neighbouring locations, while degrading populations starve and die out. Since the net environmental effect of a population has a genetic basis (in the genotypes of its constituent individuals) this between-flask competition creates a form

of higher-level selection that acts on the collective environment-altering properties of local communities, leading to community species compositions that have overall improving environmental effects. Over time, the prevalence of environment-improving communities steers the global environment towards more hospitable conditions.

Invoking higher-level selection to explain Gaian global regulation may seem a bit like jumping from the frying pan into the fire, but the form of higher-level selection we infer is a restricted form and supported by careful analysis. In previous work we have conducted experiments using artificial selection to shape the properties of single-flask communities [14], but the situation here is quite different. Artificial selection imposes both the unit of selection and the method of transmission, both of which are emergent from system dynamics in the current scenario. No long term higher-level adaptation is implied; local populations are constantly shifting due to mutations and movement of microbes, so there is no steady population genotype in which selection might fix beneficial traits. There are no community ‘generations’ and the spread of ‘good’ communities is by differential proliferation rather than explicit reproduction. Transport by diffusive mixing means community reassembly is uncertain and thus the chance of reliable transmission is weak. These factors combine to make community-level inheritance a very noisy process. However, there is enough persistence in the local communities for selection based on collective environment-altering traits to provide a steering force that moves the system away from the harsh boundaries of the habitable region and towards the ‘green fields’ of the nutrient-limited regime, where the populations expand to carrying capacity. Once the system moves back into the nutrient-limited regime, traits affecting the environment become selectively neutral once again, until the system eventually drifts too far from the growth optimum and the cycle repeats.

An interesting question is what will happen when preferences for the optimal abiotic environment are allowed to mutate – will the population regulate or adapt? We have run the spatial model with evolvable preferences, and find that the ‘error’ between the abiotic environment and the mean population preference is consistently reduced. Long periods of stability occur in which environmental state remains at a fixed level, but it is hard in the evolvable-preference scenario to separate the contribution to stability of environmental regulation from that of organism adaptation. Our current formulation of the spatial model with fixed preferences can be argued to reflect the real world, where adaptation is ultimately con-

strained by the laws of physics and chemistry [5]. For example, any photosynthetic reaction has optimal conditions of temperature, light intensity, CO₂ supply, etc., in which its rate is maximised. While there is some flexibility for metabolism to adapt to different environmental conditions, there remain hard thermodynamic constraints on any metabolic reaction (and on the space of possible metabolisms) that make a single optimal environment a reasonable working assumption.

The Flask model suggests a new mechanism by which environmental regulation can be reconciled with natural selection. The novel concept of multiple selection regimes operating at different levels and leading to the emergence of regulation will require further investigation backed by empirical evidence, but we feel that the model we propose is theoretically robust and widely applicable. In future work we will explore the possibility of regulation with adapting environmental preferences, hoping to understand our preliminary observation of ‘regulatory epochs’ of environmental stability. We are already extending the model to use a more accurate and complete chemical model than the current abstract formulation, with the initial aim of developing a simple adaptive box model of Archaean biogeochemistry. Another project is to develop adaptive modules based on the Flask scheme that can be used with existing Earth system models (e.g., the GENIE model [6]), where adaptation has traditionally been hard to incorporate. This challenging project promises significant pay-offs if successful, by linking evolutionary theory with well-modelled biogeochemical processes. A longer term goal is to compare the Flask model to real-life flasks containing real live microbes. This avenue is one of the most exciting areas for future research and is an arena where some of our Gaian hypotheses may be empirically tested.

References

- [1] K. Downing and P. Zvirinsky. The simulated evolution of biochemical guilds: Reconciling Gaia theory and natural selection. *Artificial Life*, 5(4):291–318, 1999.
- [2] J. W. Kirchner. The Gaia hypothesis: Fact, theory, and wishful thinking. *Climatic Change*, 52:391–408, 2002.
- [3] J. W. Kirchner. The Gaia hypothesis: Conjectures and refutations. *Climatic Change*, 58:21–45, 2003.

- [4] T. M. Lenton. Gaia and natural selection. *Nature*, 394:439–447, 1998.
- [5] T. M. Lenton and J. E. Lovelock. Daisyworld is Darwinian: constraints on adaptation are important for planetary self-regulation. *Journal of Theoretical Biology*, 206:109–114, 2000.
- [6] T. M. Lenton, R. Marsh, A. R. Price, D. J. Lunt, Y. Aksenov, J. D. Annan, T. Cooper-Chadwick, S. J. Cox, N. R. Edwards, S. Goswami, J. C. Hargreaves, P. P. Harris, Z. Jiao, V. N. Livina, A. J. Payne, I. C. Rutt, J. G. Shepherd, P. J. Valdes, G. Williams, M. S. Williamson, , and A. Yool. Effects of atmospheric dynamics and ocean resolution on bi-stability of the thermohaline circulation examined using the Grid ENabled Integrated Earth system modelling (GENIE) framework. *Climate Dynamics*, 2007.
- [7] T. M. Lenton and D. M. Wilkinson. Developing the Gaia theory: A response to the criticisms of Kirchner and Volk. *Climatic Change*, 58:1–12, 2003.
- [8] M. T. Madigan and J. M. Martinko. *Brock Biology of Microorganisms (11th edition)*. Pearson Prentice Hall, New Jersey, 2006.
- [9] C.L. Sears. A dynamic partnership: Celebrating our gut flora. *Anaerobe*, 11(5):247–251, 2005.
- [10] W. Swenson, J. Arendt, and D. S. Wilson. Artificial selection of microbial ecosystems for 3-chloroaniline biodegradation. *Environmental Microbiology*, 2(5):564–571, 2000.
- [11] T. Volk. *Gaia’s Body: Toward a Physiology of Earth*. Springer-Verlag, New York, 1998.
- [12] A. J. Watson and J. E. Lovelock. Biological homeostasis of the global environment: the parable of Daisyworld. *Tellus*, 35B:284–289, 1983.
- [13] H. T. P. Williams. *Homeostatic Adaptive Networks*. PhD thesis, University of Leeds, Leeds, UK, 2006.
- [14] H. T. P. Williams and T. M. Lenton. Artificial selection of simulated microbial ecosystems. *PNAS*, 104:8918–8923, 2007.
- [15] H. T. P. Williams and T. M. Lenton. Evolution of environmental regulation in a simulated network of microbial ecosystems. *In preparation*, 2007.

- [16] H. T. P. Williams and T. M. Lenton. The Flask model: emergence of nutrient-recycling microbial ecosystems and their disruption by environment-altering ‘rebel’ organisms. *Oikos*, 116:1087–1105, 2007.
- [17] A. J. Wood, G. J. Ackland, J. Dyke, H. T. P. Williams, and T. M. Lenton. Daisyworld: a review. *Reviews of Geophysics*, page in press, 2007.