SBM2–5-2011 notes

course page

October 17, 2011

1 Chemosensor modeling in space - SBM3-2011

We want to model the spatial distribution of chemosensor clusters; specifically their observed periodic distribution. We consider a grid model with periodic boundary on the y axis, and the x axis representing E. Coli's pole to pole axis (the cylinder one). We insert new receptors at a rate which scales up with the number of free sites on the grid (one can imagine a regulation in eColi doing this), and we move them around to simulate diffusion. The grid grows in the x direction.

The model is taken from *Wang et al. Self-organized periodicity of protein clusters in growing bacteria, 2008.*

§Outline -

We will consider in turn:
- the model state space (ie how you describe the state of the system at a given time)
- event types (aka transitions between states)
- event rates (average number of times they are triggered per time unit)
- some small computation to ensure that the receptor density is kept roughly constant
- see an example of a simple diffusion equation and solution and how one can use this to understand why clusters appear mostly at middle points
- discuss in general the principle of Metropolis simulation which is used in this model (as in many others) and is based on the introduction of an energy functional
- discuss the *price of lost freedom* (here α is a generic energy-based technique to penalize glueing, and the ensueing loss of freedom)
- discuss how one can make a rule-based version of this model and why

§state space -

Here we have a grid and any place can be either busy or free. This gives:

$$S(t) = L_x(t) \times L_y \to \{0, 1\}$$

Interestingly S changes over time (growth). A state is an *allocation* $\sigma(i, j) \in \{0, 1\}$ for (i, j) points on the grid, $0 \leq i < L_x(t)$, $0 \leq j < L_y$.

We will say a receptor is *free* when it has no neighbouring receptor on the grid (so a 1 with 0s around in our encoding).

It is useful also to consider the projection to $\mathbb{N}^2$ defined as $\sigma \mapsto (N, n)$ where:
- $N = L_x L_y$ is the total number of sites,
- n is the total number of receptors (ie of 1s in σ),

Equally, we coud project on L_x, n as it contains the same information (given L_y).

From which we can derive $\rho(t) = n(t)/N(t)$ for the *receptor density*, which in the model will be kept around 0.2. Talking about which, we have to be careful. As we inject matter and grow space at the same time, we might grow too quickly and get a very dilute universe (not very interesting for a model that serves as a tool to study cluster formation). If on the other hand the density is too high, we will surely have just one large cluster (not interesting either). So do we (or does eColi) want an expansive/contracting universe? 20% density seems a reasonable target. But how do we fix the density? We will do a small calculation next paragraph to get a computational handle on that.

§growth and average density -

We have exponential growth in direction x, $\partial_t L_x(t) = \gamma L_x(t)$.

By the growth law, the average cell length (in some length unit) is:

$$L_x(t) = L_x(0)e^{\gamma t} \tag{1}$$

and so the average total number of places (or sites) is:

$$N(t) = N(0)e^{\gamma t} \tag{2}$$

On the other hand, still on average:

$$\partial_t n(t) = k(N(t) - n(t)) \tag{3}$$

with solution (C a constant):

$$n(t) = Ce^{-kt} + \frac{k}{k+\gamma}N(t) \tag{4}$$

Hence the steady state value of the density is $\dfrac{k}{k+\gamma} = (1 + \gamma/k)^{-1} = 0.2$.

Verify this makes sense to you qualitatively. Why does it only depend on γ and k. It is a piece of the dynamics that happens entirely in the projection mentioned earlier. Why is it important to control the density? Is it observable biologically?

§Elementary events in the model - rates

We have events of the following type:
- *growth* at rate γ, a grid column is duplicated
- *insertion* of new receptors, at rate k, one receptor is added to a free site on the grid
- *diffusion* of single receptors, at rate δ, one free receptor diffuses to a nearby free site on the grid
- *binding/unbinding* with energy E based on 1) a coupling term J, ie $-J < 0$ per bond (or

neighbour-neighbour instance) and 2) an entropy cost for loss of freedom αJ, ie $\alpha J > 0$ per non-free receptor

In the model diffusion and binding are seen as one combined event. The diffusion/binding event type is implemented by a Metropolis algorithm, where the probability of a transition depends on the ΔE incurred by the transition. We do not need δ for Metropolis, as we will see. In a more refined model, one could (but would not have to) distinguish these events.

Caevat: it is important to understand that rates γ, k, δ (derived from D) are microscopic rates, ie per event of the given type; not to be confused with the rates at which the event type itself fires.

§diffusion - 1D case

We can do another small calculation to understand why new clusters appear at middle points between older ones.

Suppose a 1-dimensional space defined $x \in \mathbb{R}$. We have that

$$n(x,t) \propto t^{-1/2} e^{-x^2/4Dt} \tag{5}$$

is a solution to the diffusion equation with an initial condition $n(x,0) = \delta(x)$ (Dirac function) expressing the presence of a definite amount of stuff at the start (as opposed to a continuous injection of new matter; Dirichlet condition vs Neumann condition):

$$\partial_t n(x,t) = D\partial_{x^2} n(x,t) \tag{6}$$

as both sides of the equation are equal to $(e^{-x^2/(4Dt)}(x^2 - 2Dt))/(4Dt^{5/2})$ (as one can verify from alpha).

The missing constant is given by the constraint that $\int_{-\infty}^{+\infty} t^{-1/2} e^{-x^2/4Dt}\, dx = n_0$, which says that nothing disappears or appears, things just diffuse.

If we set $\lambda = \sqrt{2Dt}$ for the *diffusion length*, so that $n(x,t) \propto t^{-1/2} e^{-x^2/2\lambda^2}$, we recognise a Gaussian with standard deviation λ, which we can interpret as the typical distance to origin for a particle travelling for t time units (randomness slows it by a square root so to speak).

Eg if $D = 0.018\mu m^2 s^{-1}$, then after $2.5s$, a particle has typically moved $\sqrt{0.18}\mu m \simeq 0.4\mu m$ away from its original position, so about 1/5 of EColi's length.

§diffusion - 1D case with sinks

It is easy to see that the steady state solution to diffusion $\partial_{xx} n(x) = 0$ on an interval $[0, \ell]$, with sinks at both ends, meaning meaning $n(0,t) = n(\ell,t) = 0$ (clusters on each side are considered as sinks), is:

$$n(x) = k/2D \cdot ((\ell/2)^2 - (x - \ell/2)^2) \tag{7}$$

so peaks at $\ell/2$ which explains (or goes some way to explain) why clusters appears at middle points.

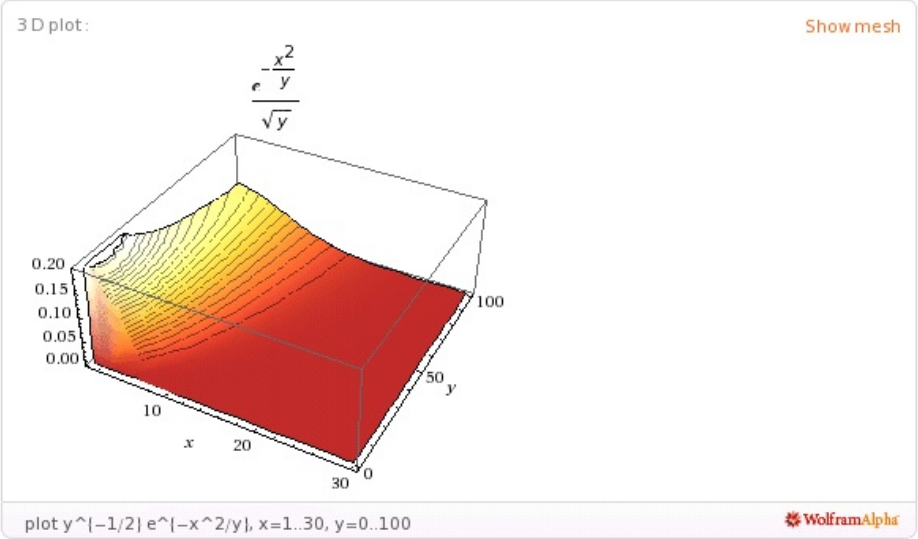

Figure 1: 1d diffusion, y is time.

§energy deltas - insertions and diffusions

The key to the simulation (Metropolis or rule-based) is to be able to compute the variation of energy consecutive to an event. According to our definition, the total energy in state x is:

$$E(x) = \begin{pmatrix} -J \\ \alpha J \end{pmatrix} \cdot \begin{pmatrix} \beta(x) \\ n - \phi(x) \end{pmatrix} \tag{8}$$

where n is the number of nodes in x, $\phi(x) \leq n$ is the number of free nodes in x, $\beta(x)$ is the number of *bonds* in x.

Recall that by convention, the lower the energy the more likely the state. We see that the two energy terms above are pulling in opposite directions (as always). The first wants many bonds to be formed (≤ 4), the second wants many free receptors to be created (≤ 4). The parameter α arbitrates between the two possibilities. When $\alpha = 0$, the energy function is a two-dimensional Ising model, for which the critical interaction strength is known to be $J_c \approx 1.763$ (Onsager, 1943). Above that one can have multiple islands/groups of 1s and 0s (Weiss domains). If α becomes large, losing freedom becomes more expensive, and clusters will be smaller.

We can compute the energy landscape, ie the ΔEs. In general, along a transition from x to y (by insertion or diffusion) we have:

$$\Delta E = \begin{pmatrix} -J \\ \alpha J \end{pmatrix} \cdot \begin{pmatrix} \beta(y) - \beta(x) \\ n(y) - n(x) + \phi(x) - \phi(y) \end{pmatrix}$$

which gives for an insertion:

$$\Delta E_i = \begin{pmatrix} -J \\ \alpha J \end{pmatrix} \cdot \begin{pmatrix} \beta_i \\ \phi_i \end{pmatrix}$$

where $0 \leq \beta_i \leq 4$ is the number of neighbours of the inserted receptor, $\phi_i = 1 + \phi(x) - \phi(y)$ is the number of receptor 'trapped' by the inserted one (himself included).

For a diffusion:

$$\Delta E_d = \begin{pmatrix} -J \\ \alpha J \end{pmatrix} \cdot \begin{pmatrix} \beta_d \\ \phi_d \end{pmatrix}$$

where $-3 \leq \beta_d \leq 3$ is the number of bonds created by the moving receptor minus the number broken; $\phi_d = \phi(x) - \phi(y)$.

We will now make a complete computation of the above in the 1d case with a circular grid.

§balances - insertion case

To minimize the size of neighbourhoods that one must incorporate in the rules below, it is convenient to store the state of freedom/business as a state of the agent itself (we could even store its *current degree*), here an anonymous agent () with two sites one for f/b the state of freedom of the agent, one for $0/1$ the presence of a receptor.

This gives the following balances for insertions (1d case, no boundary):

$$\begin{array}{ll} (f,1)(0)(f,1) \to (b,1)(b,1)(b,1) & -2J + 3\alpha J \\ (f,1)(0)(b,1) \to (b,1)(b,1)(b,1) & -2J + 2\alpha J \\ (b,1)(0)(f,1) \to (b,1)(b,1)(b,1) & -2J + 2\alpha J \\ (b,1)(0)(b,1) \to (b,1)(b,1)(b,1) & -2J + \alpha J \end{array}$$

$$\begin{array}{ll} (f,1)(0)(0) \to (b,1)(b,1)(0) & -J + 2\alpha J \\ (0)(0)(f,1) \to (0)(b,1)(b,1) & -J + 2\alpha J \\ (b,1)(0)(0) \to (b,1)(b,1)(0) & -J + \alpha J \\ (0)(0)(b,1) \to (0)(b,1)(b,1) & -J + \alpha J \end{array}$$

$$(0)(0)(0) \to (0)(f,1)(0) \quad 0$$

This gives the following balances for diffusions (1d case, no boundary, diffusion to the left):

$$\begin{array}{ll} (f,1)(0)(b,1)(b,1) \to (b,1)(b,1)(0)(?,1) & \alpha J \text{ if } ? = b,\ 0 \text{ else} \\ (b,1)(0)(b,1)(b,1) \to (b,1)(b,1)(0)(?,1) & 0 \text{ if } ? = b,\ -\alpha J \text{ else} \end{array}$$

$$\begin{array}{ll} (0)(0)(b,1)(b,1) \to (0)(f,1)(0)(?,1) & J - \alpha J \text{ if } ? = b,\ J - 2\alpha J \text{ else} \\ (f,1)(0)(f,1)(0) \to (b,1)(b,1)(0)(0) & -J + 2\alpha J \\ (b,1)(0)(f,1)(0) \to (b,1)(b,1)(0)(0) & -J + \alpha J \end{array}$$

$$(0)(0)(f,1)(0) \to (0)(f,1)(0)(0) \quad 0$$

Likewise one has diffusion to the right.

We see that one needs the degree to completely eliminate the need to compute the ΔE when applying a move. One might call this a static compilation of the Metropolis moves. More later.

With our numerical values, $\alpha = 0.5$, $-J + 2\alpha J = 0$ so the formation of a dimer is energy neutral (see diffusion step 4 above).

§Numerically - mind the units

We have:
- $J, \alpha = 4, 0.5$ so well above the Onsager threshold

- $L_y = 50$ that is fixed (does not have to be!)
- $L_x(0) = 20$
- $\gamma, k = 8 \times 10^{-6}, 2 \times 10^{-6} s^{-1}$

For an explicit modeling of diffusion, $D = 0.018 \mu m^2 s^{-1}$ which is related to the discrete diffusion coefficient δ by $h^2\delta = D$, $h = \ell_y/L_y$, ℓ_y the actual dimension along y (we could take x as well) approximately $\ell_y = 3\mu m$.

Numbers taken from *Moran et al. SnapShot: key numbers in biology. Cell (2010) vol. 141:7 p1262.*

2 Quick reminder on CTMCs

A *random exponential time* of parameter $\lambda > 0$ is an $[0, +\infty)$-valued random variable T such that $p(T > t) = \exp(-\lambda t)$. Thus, the density of T is $\lambda \exp(-\lambda t)$, for $t \geq 0$; and T's mean is $\int_0^{+\infty} \lambda \exp(-\lambda t) t \, dt = \lambda^{-1}$.

Suppose given a set X which is at most countably infinite, and a rate function $q(x, y) \in \mathbb{R}^+$, for x, y in X, and $x \neq y$.

The *transition graph* or the support of q, written $|q|$, is the binary relation, or the directed graph, on X which contains (x, y) iff $q(x, y) > 0$.

We suppose $|q|$ has finite out-degree (this also called being image-finite).

We can define a continuous-time Markov chain over X in the following way. When the chain is at x in X, for each of the finitely many ys such that $q(x, y) > 0$, draw a random exponential time $\tau(x, y)$ with parameter $q(x, y)$; advance time by $\tau = \min \tau(x, y)$, and jump to the (almost surely) unique y such that $\tau(x, y) = \tau$.

The idea is that all possible next states compete, and the higher the rate of $q(x, y)$, the more likely it is that y will be the next state. It is easy to calculate that the probability to jump to y is actually $q(x, y)/\sum_z q(x, z)$; and that for small ts, the probability to jump to x within t is equivalent to $q(x, y)t$, hence one can think of $q(x, y)$ as the rate at which one jumps from x to y.

Note that for the above definition to make sense it is important to suppose as we have done that $|q|$ is image-finite. We will also suppose thereafter that $|q|$ is symmetric (not to be confused with the much stronger assumption that q is a symmetric function, ie $q(x, y) = q(y, x)$), and define $\rho(x, y) = q(y, x)/q(x, y)$ when either (equivalently both) of $q(x, y)$ and $q(y, x)$ are > 0.

2.1 Equilibrium

Now, on to the definition of an equilibrium that will be our central concern here.

Consider a function p defined on X and with values in $\mathbb{R}^+$. One says p is an *equilibrium* for q if p is not everywhere zero, and:
- [*detailed balance*] for all $(x, y) \in |q|$, $p(x)q(x, y) = q(y, x)p(y)$
- [*convergence*] $Z = \sum_X p(x) < +\infty$

If such a p exists, we can obtain a probability on X by normalizing p as p/Z. Naturally, if X is finite the second condition always holds.

The detailed balance condition implies that p, construed as a probabilistic state of the system,

is a fixed point of the action of the chain q, and as $|q|$ is symmetric, regardless of the initial state, the chain will converge to p.

3 The Metropolis-Hastings dtMC structure

We assume:
- an energy function $E : X \to \mathbb{R}$ on a countable state space X
- an *a priori* symmetric irreflexive ($\alpha(x,x) = 0$) Markov kernel α on X

With this data, we can define a discrete-time Markov chain (dtMC) as a process that mostly but not always follows lines of decreasing energy (gradient-driven). The idea is that α will select a priori a candidate transition which we then might accept with a certain probability. If energy diminishes along that transition (modulo an α-related correction, see below) we take it certainly; if it increases we take it with a decreasing probability (as we are reluctant to climb up the energy landscape).

Note that the α kernel defines the transition graph.

Define for $\alpha(x,y) > 0$ (equivalently $\alpha(y,x) > 0$):

$$\begin{array}{lcll} \beta(x,y) & = & \min(1, e^{E(x)}\alpha(x,y)^{-1} \cdot e^{-E(y)}\alpha(y,x)) & \text{acceptance probability} \\ p(x,y) & = & \alpha(x,y)\beta(x,y) & \text{total motion probability} \end{array}$$

In case of a rejection, we stay at x, which completes the definition of the transition function:

$$p(x,x) = 1 - \textstyle\sum_{y \neq x} \alpha(x,y)\beta(x,y) \quad \text{null event probability}$$

So - supposing wlog $\beta(x,y) \leq 1$ - we do get discrete detailed balance[1] wrt E:

$$\begin{array}{rcl} p(y,x)/p(x,y) & = & \alpha(y,x)/\alpha(x,y) \cdot \beta(y,x)/\beta(x,y) \\ & = & \alpha(y,x)/\alpha(x,y) \cdot e^{-E(x)}\alpha(x,y) \cdot e^{E(y)}\alpha(y,x)^{-1} \\ & = & e^{E(y)} \cdot e^{-E(x)} \end{array}$$

Note that this probability equipment preserves the underlying transition graph defined by α - as $\beta(x,y)$ is never zero. In particular, it is still symmetric (hence irreducible; it also preserves aperiodicity). Only the transition probabilities are altered (and loops introduced). Unless x is a local max for E, $\sum_{y \neq x} p(x,y) < 1$, and the remaining mass is converted into $p(x,x)$ the probability of staying put at x.

If α is uniform and the underlying transition graph is regular (with finite and non-zero out-degree, as in for instance our 1d model above where every node has degree 2), then $\beta(x,y)$ accepts certainly if $\Delta E := E(y) - E(x) \leq 0$ - or else with probability $\exp(-\Delta E) \leq 1$. This case is called Metropolis.

One can extend the above to $\Delta E = +\infty$ - then $\beta(x,y) = 0$ for infinite ys, ie rejection is certain. The underlying transition graph gets restricted (eg one can refuse to jump off a domain in the plane).

One can also extend this to non-symmetric α, in which case one needs to reject certainly if $\alpha(y,x) = 0$ (meaning β corrects the lack of symmetry by cancelling assymetric jumps).

[1] meaning for all x, y $p(x)P(x,y) = p(y)P(y,x)$, which implies $pP = p$ the invariance of p

One can also deal with non irreflexive α, ie $\alpha(x,x) > 0$ for some x, but when is this useful?

One can use α to localize the moves, optimize the price of computing ΔE, and favour a priori ys with a lower energy if possible. One can think of α as a heuristics -if say the goal is to sample from the equilibrium probability.

§a continuous-time version -

We start with the same ingredients, X, α, E except that α is now an arbitrary rate function on $X \times X \setminus \Delta_X$. We define as in the discrete-time case:

$$\begin{aligned} q(x,y) &= \alpha(x,y)\beta(x,y) \\ \beta(x,y) &= \min(1, \alpha(y,x)/\alpha(x,y)\exp(-(E(y)-E(x))) \end{aligned}$$

hence, supposing wlog $\alpha(y,x)/\alpha(x,y)\exp(-(E(y)-E(x)) \leq 1$ (equivalently $\alpha(x,y)/\alpha(y,x)\exp(-(E(x)-E(y)) \geq 1$)

$$\begin{aligned} q(y,x)/q(x,y) &= \alpha(y,x)\alpha(x,y)^{-1}\alpha(x,y)/\alpha(y,x)\exp(-(E(x)-E(y)) \\ &= \exp(E(y)-E(x)) \end{aligned}$$

so we have detailed balance.

As α in this case is any function, it is always possible to choose $\alpha = 1$, and everything simplifies to $q(x,y) = \min(1, \exp(-(E(y)-E(x)))$.

§a common mistake - go ostinato

Suppose one repeats the $\alpha\beta$ protocol without incrementing time. Eg write until$(\omega_i, 0 \leq w_i \leq 1)$ for the obstinate 2-step postselection (as a for loop as it almost surely terminates), where:
- 1) one draws uniformly an i (in a finite non-empty set, say $\{1,\ldots,m\}$) and
- 2) accepts ω_i with probability w_i, or else goto 1 and advance time.

This picks up ω_i with probability $p(\omega_i) = w_i/\sum_j w_j$, as this probability $p(\omega_i)$ satisfies the recursive relation (supposing $|\Omega| = m < \infty$):

$$p(\omega_i) = w_i/m + 1/m(\textstyle\sum_j(1-w_j))p(\omega_i)$$

Getting to pick ω_i will happen after some mean time which one can also compute - this will increase as w_i decreases.

Now, supposing $E(x) \geq E(y)$ ie $w(x,y) = 1$:

$$\begin{aligned} p(y,x)/p(x,y) &= w(y,x)/\textstyle\sum_{y'} w(y,y') \cdot \sum_{y'} w(x,y')/w(x,y) \\ &= \exp(E(y)-E(x)) \cdot \textstyle\sum_{y'} w(x,y')/\sum_{y'} w(y,y') \end{aligned}$$

so that - if $\sum_{y'} w(x,y') = \sum_{y'} w(y,y')$ for any two neighbours x, y - E is an equilibrium for $p(x,y)$ (on the initial component). But there is no reason why this should be true in general with this assignment! Eg if x is the center of a star and has energy zero, while peripheral nodes have energy 1, then this is saying that $d(x) = 1/e$ which is absurd - perhaps in the limit where many neighbours have lower energy, therefore contributing a 1, the equation is approximately satisfied?

§an example - to understand the difference between the two versions

Suppose X is finite and the transition graph underlying α is a star-shaped graph with a hub x_0, and n branches; suppose $E = 0$ everywhere.

Let us consider first the discrete case with $\alpha(x)$ the uniform probability. Then $p(x_0, y) = 1/n \cdot 1 = 1/n$, $p(y, x_0) = 1 \cdot 1/n = 1/n < 1$, $p(y, y) = 1 - 1/n$. One stays put a lot in a peripheric node. The limit is the uniform probability by construction (as $E = 0$ means uniform), modulo the problem of periodicity. We see also that it is not always possible to pick $\alpha(x, y) = \alpha(y, x)$ when the transition graph does not have constant outdegree. If we were to go obstinato (see above), then the limit probability would be $1/2$ at x_0 and $1/2n$ elsewhere; not converging to the intended uniform probability.

Consider now the ctMC case, with $\alpha = 1$. Then $q(x_0, y) = 1 = q(y, x_0)$, again this converges by construction to the uniform probability. But now there are no rejections, it is just that the process stays naturally longer on a peripheric point where the frequency is 1, than at the hub where it is n.

4 return to model

§Description of the Metropolis hybrid simulation -

First, as we have noticed already we can give a self-consistent description of the projected MC on L_x, n:

$$L_x, n \to^{\gamma L_x} L_x + 1, n$$
$$L_x, n \to^{k(L_x L_y - n)} L_x, n + 1$$

Rates are indicated on top of the transition. This is a pure ctMC which we can run and will describe self-consistently the evolution of the population and grid size.

Now to run our simulation, we must be able to mix a discrete-time MC as given by the Metropolis part of the algorithm, and the continuous-time dynamics of growth and insertions (which have just written in the projected space).

One solution, in this particular case, is to exploit the above ctMC. This will give the (random) times at which the grid should grow (and then we pick up uniformly the place of insertion of a new column), and receptors should be inserted (again we pick the position of insertion uniformly, although they will incur different ΔEs as we have seen earlier).

Here is another, and completely generic method to do this.

§hybrid dt/ctMC structures -

Suppose we are given at each state σ a partition $\mathcal{E}(\sigma) = \mathcal{E}_c(\sigma) + \mathcal{E}_d(\sigma)$ of all events which are possible at σ (the event horizon).

Given a double dt/ctMC structure on the same state space, and τ a fixed time increment (per discrete step), we can define a hybrid dt/ctMC protocol, one step of which is as follows.

Starting in state σ:
- draw all continous random times, $t(e)$, $e \in \mathcal{E}_c(\sigma)$
- pick the smallest $t(e_0) := \inf_{e \in \mathcal{E}_c} t(e)$

- if $t(e_0) < \tau$ execute e_0 (a CTMC step) with result in a new state σ' and a time increment δt
- else execute a discrete step on $\mathcal{E}_d(\sigma)$ which results in a new state σ' and a time increment τ

§simulation -

We can apply our general scheme. Each state of the system σ has a partition as above: the continuous events are growths and insertions, the discrete ones are diffusions/bindings; we pick $\tau = 1$ which means the reference time unit is a MC step. Or we can opt for the ctMC version of Metropolis and just add the transitions for growths and insertions.

If we return to the simulation output, between time $[7.10^5, 8.10^5]$, we see that each column insertion event had about one opportunity to apply (as $\gamma \sim 10^{-5}$), and as we have about L_x such events (this is an under approximation), this means that the size should approximately (more than) double in this time interval. Indeed this is what we see in the result (a 'typical' run of the simulation). This shows that rates k, γ are measured in a time unit which is that of an MC step.

To know what that means in terms of clock time units, we could calibrate by knowing the doubling time for an eColi.

Figure 2: One Metropolis simulation of the model of spatial clustering. MC steps are in multiple of 10^4.

§What if we do not want to do Metropolis? - Rule-based version

We have seen Kappa last session; we also have an extension Spatial Kappa (due to Donal Stewart, used in our last 2 award-winning IGEM projects). As the state space $L_x(t) \times L_y \to \{0, 1\}$ grows with time, we would have to adapt this, extend it further. Another problem is to accomodate *exclusivity*, that is to say the fact that any location in the grid can be populated by at most one object (whereas there is no such restriction in Spatial Kappa at the moment). So let's keep with Kappa for now.

In principle it is easy. We just turn each case of the ΔE analysis into its own rule and have to ensure that the rate ratio $\rho(x, y)$ is the correct value, that is $\exp(\Delta E)$; convergence to the correct invariant measure is guaranteed as long as one stays in the reversible subset of rules; dynamics is different though (more realistic).

Exercise: Do it. Reproduce the growthless model in Kappa. Hint: 1) describe the grid itself as a Kappa object; 2) chose rates so that the energy deltas are consistent, which means rule have to be able to tell the ΔE (which means they have to incorporate enough to compute the freeness of any node involved, see the calculation made earlier).

How would you model growth?

Here is a question for us to ponder later. How is the Metropolis dynamics close to a rule-based one? Can we pick up rule rates such that they coincide (in the absence of binary rules perhaps?).

§non-ambiguous ernergy balance - using multisets of sites

Modeling the grid in Kappa using multisets of sites, where left, right, up, down are indistinguished would increase the symmetry hence decrease the number of rules to write; for diffusion for instance the above would be enough. There is a tool compiling this to normal Kappa.

Consider the following energy shapes (using multisets of sites to lighten the notation):
- bond: $(o_1, a^1), (o_1, a^1)$ with price $-J$
- free agents, that is to say all "stars" with at least one occupied neighbouring site:

$$\begin{array}{l}(o_1, a^1, a^2, a^3, a^4), (o_1, a^1), (o_1, a^2), (o_1, a^3), (o_1, a^4)/4! + \\ (o_1, a^1, a^2, a^3, a^4), (o_1, a^1), (o_1, a^2), (o_1, a^3), (o_0, a^4)/3! + \\ (o_1, a^1, a^2, a^3, a^4), (o_1, a^1), (o_1, a^2), (o_0, a^3), (o_0, a^4)/2!^2 + \\ (o_1, a^1, a^2, a^3, a^4), (o_1, a^1), (o_0, a^2), (o_0, a^3), (o_0, a^4)/3!\end{array}$$

each with price αJ. Note the symmetry discounts (as we count shape embeddings).

We could also count free stars $(o_1, a^1, a^2, a^3, a^4), (o_0, a^1), (o_0, a^2), (o_0, a^3), (o_0, a^4)/4!$ and subtract this to $(o_1) = n$.

If we know the value of these observables, then we know the energy of the system.

Hence, the formal constraint that rules must fullfil to have a *non-ambiguous energy balance* is that every variation of the energy shape counts incurred by applying the rule should be computable from the rule itself. The rule has to provide enough context to compute its ΔE.

Else, we have to discover its rate per each event at run-time, and it will depend on the event.