Metastability Of Large Networks With Mobile Servers

F. Baccelli [1], A. Rybko [2], Senya Shlosman [2,3,4], A. Vladimirov [2]

[1]UT Austin, Department of Mathematics, USA
[2]Institute for Information Transmission Problems, RAS, Moscow
[3]Aix Marseille Université, Université de Toulon
[4]Skolkovo Institute of Science and Technology, Moscow
mailto:@smith.comvladim@iitp.ru

Abstract

We study symmetric queueing networks with moving servers and FIFO service discipline. The mean-field limit dynamics demonstrates unexpected behavior which we attribute to the metastability phenomenon. Large enough finite symmetric networks on regular graphs such as cycles are proved to be transient for arbitrary small inflow rates. However, the limiting non-linear Markov process possesses at least two stationary solutions. The proof of transience is based on martingale technique.[1]

Keywords: ad hoc network, transience, metastability, mean field

I Introduction

In this paper we consider networks with moving servers. The setting is the following: the network is living on a finite or countable graph $G = (V, E)$, at every node $v \in V$ of which one server s is located at any time. For every server, there are two incoming flows of customers: the exogenous customers, who come from the outside, and the transit customers, who come from some other servers. Every customer c coming into the network (through some initial server $s(c)$) is assigned a destination $D(c) \in V$ according to some randomized rule. If a customer c is served by a server located at $v \in V$, then it jumps to a server at the node $v' \in V$, such that $\text{dist}(v', D(c)) = \text{dist}(v, D(c)) - 1$, thereby coming closer to its destination. If there are several such v', one is chosen uniformly. There the customer c waits in the FIFO queue until his service starts. If a customer c completes his service by the server located at v, and it so happens that $\text{dist}(v, D(c))$ is 1 or 0, the customer is declared to have reached its destination and leaves the network.

The important feature of our model is that the servers of our network are themselves moving over the graph G. Namely, we suppose that any two servers s, s' located at adjacent nodes of G exchange their positions as the alarm clock associated to the edge rings. The time intervals between the rings of each alarm clock are i.i.d. exponential with rate β. When this happens, each of the two servers takes all the customer, waiting in its buffer or being served, to the new location. In particular, it can happen that after such a swap, the distance between the location of the customer c and its destination $D(c)$ increases (at most by one). We assume that the service times of all customers at all servers are i.i.d. exponential with rate 1.

The motivation for this model comes from opportunistic multihop routing in mobile ad hoc wireless networks, see [5, 9, 7, 1, 3, 4]. Within this context, the servers represent mobile wireless devices. Each device moves randomly on the graph G which represents the phase space of device locations. The random swaps represent the random mobile motions on this phase space.

Each node $v \in G$ of the phase space generates an exogenous traffic (packetized information) with rate λ_v corresponding to the exogenous customers alluded to above. Each such packet has some destination, which is some node of G. In opportunistic routing, each wireless device adopts the following greedy routing policy: any given packet scheduled for wireless transmission is sent to the neighboring node which is the closest to the packet destination. The neighbor condition represents in a simple way the wireless constraints. It implies a multihop route in general. This routing policy is the most natural one to use in view of the lack of knowledge of future random swaps.

In this paper we restrict consideration to cyclic graphs $C_K = \mathbb{Z}^1/K\mathbb{Z}^1$ and their *mean-field* versions, see below. Our main results, however, can be easily extended to much wider classes of networks.

The interest in mean-field versions is both of mathematical and practical nature. The mathematical interest of the mean-field version of a network is well documented. There are also practical motivations for analyzing such networks: their properties are crucial for understanding the long-time behavior of finite size networks.

The results we obtain look somewhat surprising. First of all, we find that for finite graphs the network is transient once the diameter of the graph is large enough. For example, consider the network on the graph C_K with Poisson inflows with rate $\lambda > 0$ at all nodes, exponential service times with rate 1, FIFO discipline and node swap rate $\beta > 0$. Then for all $K \geq K(\lambda, \beta)$ the queues at all servers tend to infinity as time grows. In words this means that the network is unstable for any λ, however small it is – once the network is large enough.

The same picture takes place for mean-field graphs C_K^N with N finite. They consist of N "parallel" copies of C_K such that two nodes in different copies are adjacent if and only if the projections of these nodes to a single copy of C_K are adjacent. However, the limiting picture, for $N = \infty$, is different: the corresponding NLM process on C_K has stationary distributions, provided $0 < \lambda \leq \lambda_{cr}(K, \beta)$, with $\lambda_{cr}(K, \beta) < \infty$ for all $K \leq \infty$. Moreover, for all $\lambda < \lambda_{cr}$ there are at least two different stationary distributions, see Sect. 4 for more details. We demonstrate results of numerical modeling that suggest existence of three equilibria in some cases.

On the other hand, the general convergence result of [2] claims the convergence of the networks on C_K^N to the one on C_K^∞ as $N \to \infty$, which seem to contradict to the statements above. The explanation of this 'contradiction' is that the convergence in [2] holds only on finite time intervals $[0, T]$.

That is, for any T there exists a value $N = N(T)$, such that the network on C_K^N is close to the limiting network on C_K^∞ for all $t \in [0, T]$, provided $N \geq N(T)$. Putting it differently, the C_K^N network behaves like the limiting C_K^∞ network – and might even look as a stationary process – for quite a long time, depending on N, but eventually it departs from such regime and gets into the divergent one. Clearly, the picture we have is an instance of metastable behavior. We believe that more can be said about the metastable phase of our networks, including the formation of critical regions of servers with oversized queues, in the spirit of statistical mechanics, see e.g. [8], but we will not elaborate here on that topic.

II Finite networks

2.1 The C_K network

The only case of a finite network we study here is the cyclic graph $C_K = \mathbb{Z}^1/K\mathbb{Z}^1$. As was mentioned, our main results proved for this graph are easily extendable to much wider classes of networks. We use notation $C_K = (V_K, E)$, where $V_K = \{1, \dots, K\}$ and $E = \{(1,2), \dots, (K-1, K), (K, 1)\}$. For simplicity we take K to be odd.

We study a continuous-time Markov process on a countable state Q, related to the graph C_K. Namely,

$$Q = \{q^v : v \in V_K\} = (V_K^*)^{V_K},$$

where V_K^* is the set of all finite words in the alphabet V_K, including the empty word $\emptyset$.

The queue $q^v \in V_K^*$ at a server located at $v \in V_K$ consists of a finite (≥ 0) number of customers which are ordered by their arrival times (FIFO service discipline) and are marked by their destinations which are vertices of the graph C_K. Since the destination of the customer is its only relevant feature, in our notations we sometime will identify the customers with their destinations.

2.1.1 Dynamics

Let us introduce the continuous-time Markov process $\mathcal{M} \ll \mathcal{M}(t)$ with the state space Q. Let h_v be the length of the queue q^v at node v. We have $q^v = \{q_1^v, \dots, q_{h_v}^v\}$ if $h_v > 0$ and $q^v = \emptyset$ if $h_v = 0$.

The following events may happen in the process $\mathcal{M}$.

An arrival event at node v changes the queue at this node. If the newly arrived customer has for its destination the node w, then the queue changes from q^v to $q^v \oplus w$, that is, to $\{q_1^v, \dots, q_{h_v}^v, w\}$ if $h_v > 0$ or from $\emptyset$ to $\{w\}$ if $h_v = 0$.

In this paper we consider the situation where each exogenous customer acquires its destination at the moment of first arrival to the system, in a translation-invariant manner: the probability to get destination w while arriving to our network at the node v depends only on $w - v$ mod K. The case $w = v$ is not excluded. We thus have the rates $\lambda_{v,w}$, $v, w \in C_K$, and the jump from q^v to $q^v \oplus w$, corresponding to the arrival to v of the exogenous customer with final destination w happens with the rate $\lambda_{v,w}$. We introduce the rate λ of exogenous customers as

$$\lambda = \sum_w \lambda_{v,w} \tag{1}$$

(according to our definitions it does not depend on v).

Each node is equipped with an independent Poisson clock with parameter 1 (the service rate). As it rings, the service of the customer q_1^v is over, provided $h_v > 0$; nothing happens if $h_v = 0$. In the former case the queue at node v changes from q^v to

$$q_-^v = \{q_2^v, \dots, q_{h_v}^v\}$$

(we also define $\emptyset_- = \emptyset$) and immediately one of the two things happen: either the customer q_1^v leaves the network, or it jumps to one of the two neighboring queues, $q^{v \pm 1}$. The customer q_1^v leaves the network only if its current position, v, is at distance ≤ 1 from its destination, i.e. iff $q_1^v = v - 1$, v, or $v + 1$. (This is just one of many possible choices we make for simplicity.) Otherwise it jumps to one of the neighboring vertices $w = v \pm 1$, which is the closest to its destination, i.e. to the one which satisfy: $\text{dist}(w, q_1^v) = \text{dist}(v, q_1^v) - 1$ (there is a unique such $w \in V_K$ since we assume K to be odd. The case of even K requires small changes).

The last type of event is the swap of two neighboring servers. Namely, there is an independent Poisson clock at each edge $uv \in E$ of C_K, with rate $\beta > 0$. As it rings, the queues at the vertices u and v swap their positions, that is,

$$q^v(t_+) = q^u(t), \qquad q^u(t_+) = q^v(t).$$

2.1.2 Submartingales

Here we introduce some martingale technique that will be used for the proof of transience of $\mathcal{M}$ for K large enough. To begin with, we label the K servers by the index $k = 1, \dots, K$; this labelling will not change during the evolution. Together with the original continuous-time Markov process $\mathcal{M}(t)$ we will consider the embedded discrete time process $M(n)$, which is the value of $\mathcal{M}(t)$ immediately after the n-th event. The state of the process M consists of the states of all K

servers and all their locations.

The general theorem below will be applied to the quantities X_n^k, which are, roughly speaking, the lengths of the queues at the servers k, $k = 1, \ldots, K$, of the process $M(\Lambda n)$. The integer parameter $\Lambda = \Lambda(K, \lambda, \beta)$ will be chosen large enough, so that, in particular, after time Λ, the locations of the servers are well mixed on the graph C_K, and the joint distribution of their location on C_K is close to the uniform one. Moreover, we want the expectations of all the differences $X_{n+1}^k - X_n^k$ to be uniformly positive.

We start with the following theorem.

Theorem 1 *Let $\mathcal{F} = \mathcal{F}_n$, $n = 0,1,\ldots$, be a filtration and let X_n^k, $k = 1, \ldots, K$, be a finite family of non-negative integer-valued submartingales adapted to $\mathcal{F}$, such that for all $k = 1, \ldots, K$, and all $n = 0,1,\ldots$, the following assumptions hold:*

(1) For some $\rho > 0$ the inequality

$$\mathbb{E}_{\mathcal{F}_n}(X_{n+1}^k - X_n^k) \geq \rho \tag{2}$$

holds whenever $X_n^k > 0$.

(2) The increments are bounded by a constant R:

$$|X_{n+1}^k - X_n^k| \leq R \quad \text{a. s.} \tag{3}$$

Then there exists an initial state $(X_0^1, \ldots, X_0^K)$ such that, with positive probability, $X_n^k \to +\infty$ as $n \to +\infty$ for all $k = 1, \ldots, K$.

In order to prove the theorem we begin with an auxiliary lemma.

Lemma 2 *Let $\mathcal{Y}^k = \{Y_n^k : n = 0,1,\ldots, k = 1, \ldots, K$, be a finite family of submartingales adapted to the same filtration $\mathcal{F}$ and such that $Y_n^k \in 0,1]$ for all k, n. Suppose also that for any $\varepsilon > 0$ there exists a $\delta > 0$ such that*

$$\mathbb{E}(Y_{n+1}^k - Y_n^k) > \delta \quad \text{once} \quad 0 < Y_n^k < 1 - \varepsilon$$

for all k and n. Suppose that the initial vector $Y_0 \in A = [0,1]^K$ is deterministic and satisfies the condition

$$\sum_{k=1}^{K} Y_0^k > K - 1., \tag{4}$$

Then, with positive probability, $Y_n^k \to 1$ as $n \to \infty$, for all $k = 1, \ldots, K$.

Proof. Since all submartingales Y^k are bounded, there is a limit $\lim_{n \to \infty} Y_n^k$ almost surely for all k, see the Martingale Convergence Theorem in [6]. The value of this limit vector with probability 1 is either the 'maximal' vertex $(1, \ldots, 1)$ of the cube A or a point a on the 'lower boundary' B of A: $B = \{a : \min_{k=1,\ldots,K} a_k = 0\}$. Indeed, for all other vectors $v \in A$, we have

$$\mathbb{E}(Y_{n+1}^k - Y_n^k) > 0 \quad \text{if} \quad Y_n^k = v_k = v, \quad k = 1, \ldots, K.$$

Note that

$$\sum_{k=1}^{K} b_k \leq K - 1 \tag{5}$$

for any vector $b \in B$. By the submartingale property, we conclude that

$$\mathbb{E}\sum_{k=1}^{K} Y_n^k \geq \sum_{k=1}^{K} Y_0^k > K - 1 \tag{6}$$

for all $n = 1, \ldots$. Inequalities (4)-(6) rule out the option that the limit of Y_n belongs to B with probability 1.

Now, in order to derive Theorem 1 from Lemma 2, we make the following change of variables for submartingales X_n^k. For a positive parameter $\alpha < 1$ we define an 'irregular lattice' $h_i \in \mathbb{R}_+$, by

$$h_0 = 0, \quad h_{i+1} = h_i + \alpha^i, \quad i = 0,1,\ldots.$$

We get $\lim_{i \to \infty} h_i = H = (1 - \alpha)^{-1} < \infty$. Now, for each $k = 1, \ldots, K$, we define the process Y_n^k on the same filtration $\mathcal{F}$ by the relation

$$Y_n^k(\omega) = h_{X_n^k(\omega)}.$$

The processes Y_n^k take values at the 'lattice' $\{h_i\}$ for $k = 1, \dots, K$. They are still submartingales if $1 - \alpha$ is small enough. Indeed, for such α the local structure of the lattice in an R-neighborhood of a given point is modified only slightly. Since $|X_{n+1}^k - X_n^k| \leq R$ and $\mathbb{E}(X_{n+1}^k - X_n^k) \geq \rho > 0$, we conclude that the submartingale property is preserved.

Then the hypothesis of Lemma 2 holds (up to a constant factor H) and Theorem 1 is proved.

2.1.3 Transience

Let us return to the process $\mathcal{M}(t)$. Suppose that the parameters $\lambda > 0$ and $\beta > 0$ are fixed. We remind the reader that our service rate is set to 1.

Theorem 3 *For each $\lambda > 0$ and $\beta > 0$, there exists $K^* \in \mathbb{Z}_+$ such that for any $K \geq K^*$ the process $\mathcal{M}$ is transient.*

Proof. First of all, we construct a discrete time Markov chain $\mathcal{D}$ on the state space Q. To define it, we start with the embedded Markov chain $M(n)$, defined earlier, and then pass to the chain $M(\Lambda n)$, with the integer Λ to be specified later. To get the chain $\mathcal{D} \ll \{\mathcal{D}_n\}$, we modify the chain $M(\Lambda n)$ as follows: if for some n at least one of the K queues is at most Λ, we add to all such queues extra customers, to make these queues to be of length exactly Λ and then stop the process forever. Otherwise we do no changes. The obtained Markov chain is denoted by $\mathcal{D}$.

We start the process $\mathcal{D}$ at some configuration Q_0 with all queues longer than Λ.

We now prove the following statement: if Λ is large enough, the queue length process $\mathcal{D}$ at any given server is a submartingale satisfying the conditions of Theorem 1, with respect to the filtration defined by our discrete-time Markov chain $M(n)$ (the individual queue length processes are clearly adapted to this filtration). This completes the proof because of Theorem 1. We need the following lemmas.

Lemma 4 *(1) Let us consider the following function $\pi(t)$ of the process $\mathcal{M}$. At each $t \geq 0$, $\pi(t)$ is the current permutation of indices of K servers with respect to indices of K nodes. Then the evolution of $\pi(t)$ is a continuous time Markov process, independent of service and arrival processes, and, as $t \to \infty$, the distribution of $\pi(t)$ converges to the uniform one on the set S_K of all permutations.*

(2) Let us fix index i and denote by $v(i, t)$ the position of the server i at time t. Then the distribution of $v(i, t)$ converges to the uniform distribution on $\{1, \dots, K\}$ as $t \to \infty$.

Proof. Let us introduce the graph structure on the permutation group S_K. Namely, we consider all the transpositions $\tau \in S_K$ corresponding to the exchanges of pairs of neigboring servers, and we call two permutations π', π'' to be connected by an edge iff $\pi' = \pi''\tau$ for some τ.

The resulting graph on S_K is connected – because G is connected. The process of migration of servers is, obviously, a random walk on this graph, that is, a reversible process. Hence, as $t \to \infty$, the distribution of permutations converges to the uniform one uniformly on all initial states. The assertion of the lemma clearly follows.

Lemma 5 *For any initial state Q_0, the probability of a customer with position $H > 0$ in the queue to leave the network after being served, tends to $3/K$ as $H \to \infty$, uniformly in Q_0.*

Proof. As the waiting time of the customer tends to infinity with $H \to \infty$, the distribution of its server on V_K tends to the uniform one on C_K (see Lemma 4). In order for the customer c to exit the network, the last server of c has to be located at this moment at one of the three nodes: $D(c) + 1, D(c)$, or $D(c) - 1$. The lemma follows.

Now we see that for all the customers in the initial queues whose positions are at least H, the mean chance of exit approaches $3/K$ as $H \to \infty$, and the rate of this approach does not depend on the particularities of the initial state Q_0, but only on H.

The next remark is that if a customer is served and then jumps to a different server, then the index j of that server is distributed almost uniformly over the remaining $K - 1$ indices. This fact follows from Lemma 5. Again, the rate of convergence is independent of Q_0 because the servers swap positions independently of anything else. So we have established a lemma, analogous to Lemma 5:

Lemma 6 *The probability of a customer with position H on server i to jump to server j tends to $1/(K - 1)$ as $H \to \infty$ uniformly in i, j, and in the initial states $Q_0 \in Q$.*

We need a third combinatorial lemma, and we start with some definitions, and then formulate and prove it. Let $\{u, v\} \subset C_K$ be an ordered pair of elements. We define the map T from the set of all such pairs into the union $V_K \cup \{*\}$, by

$$T\{u, v\} = \begin{cases} w & \text{for } w \text{ defined by } |u - w| = 1, \ |v - w| = |u - v| - 1, \\ & \text{provided } |u - v| > 1, \\ * & \text{otherwise.} \end{cases}$$

For K odd the map T is well-defined. In case $T\{u, v\} = w$ we say that a customer transits through w (on his way from u to v).

Let $D: V_K \to V_K$ be an arbitrary map. We want to compute the quantity

$$p_K = \frac{1}{K!} \sum_{\pi \in S_K, i \in V_K} \mathbb{I}_{\{T\{\pi(i), D(i)\} = \pi(j)\}}, \tag{7}$$

where S_K is the symmetric group, π runs over all permutations from S_N, while i and j are taken from some fixed labelling of the elements of V_K. Thus p_K is the probability of transit through the node $\pi(j)$ in the ensemble defined by the uniform distribution on S_K. Of course, it does not depend on j. (Note that we consider the action of S_K on pairs $\{u, v\}$ given by $\pi\{u, v\} = \{\pi u, v\}$.)

Lemma 7

$$p_K = \frac{K - 3}{K}.$$

Proof of the Lemma. Let $1, 2, \ldots, K$ be the labelling fixed; without loss of generalarily we can take $j = 1$. Instead of performing the summation in (7) over whole group S_K, we partition S_K into $(K - 2)!$ subsets A_π, and perform the summation over each A_π separately. If the result will not depent on π, we are done. Here $\pi \in S_K$, and, needless to say, for π, π' different we have either $A_\pi = A_{\pi'}$ or $A_\pi \cap A_{\pi'} = \emptyset$.

Let us describe the elements of the partition $\{A_\pi\}$. So let π is given, and the string $i_1, i_2, i_3, \ldots, i_l, i_{l+1}, \ldots, i_K$ is the result of applying the permutation π to the string $1, 2, \ldots, K$. Then we include into A_π the permutation π, and also $K - 1$ other permutations, which correspond to the cyclic permutations, e.g. we add to A_π the strings $i_K, i_1, i_2, i_3, \ldots, i_l, i_{l+1}, \ldots,$ $i_{K-1}, i_K, i_1, i_2, i_3, \ldots, i_l, i_{l+1}, \ldots,$ and so on. We call these transformations 'cyclic moves'. Now with each of K permutations already listed we include into A_π also $K - 2$ other permutations, where the element i_1 does not move, and the rest of the elements is permuted cyclically, i.e., for example from $i_{K-1}, i_K, i_1, i_2, i_3, \ldots, i_l, i_{l+1}, \ldots,$ we get $i_K, i_2, i_1, i_3, \ldots, i_l, i_{l+1}, \ldots, i_{K-1},\ i_2, i_3, i_1, \ldots, i_l, i_{l+1}, \ldots, i_{K-1}, i_K$, and so on. We call these transformations 'restricted cyclic moves'. The main property of thus defined classes of configurations is the following: Let $a \neq b \in \{1, 2, \ldots, K\}$ be two arbitrary indices, and $l \in \{2, \ldots, K\}$ be an arbitrary index, different from 1. Then in every class A_π there exists exactly one permutation π', for which $i_1 = a$ and $i_l = b$.

Given π, take the customer $l \neq 1 (= j)$, and its destination, $D(l)$. If we already know the position i_1 of customer 1 on the circle C_K, then in the class A_π there are exactly $K - 1$ elements, each

of them corresponds to a different position of the server l on C_K. If it so happens that $i_1 = D(l)$, then for no position of the server l the transit from l through i_1 happens. The same also holds if $i_1 = \left(D(l) + \frac{K-1}{2}\right) mod K$ or $i_1 = \left(D(l) + \frac{K+1}{2}\right) mod K$. For all other $K - 3$ values of i_1 the transit from l through i_1 happens precisely for one position of l (among $K - 1$ possibilities). Totally, within A_π we have $(K - 1)(K - 3)$ transit events. Since $|A_\pi| = K(K - 1)$, the lemma follows. +

End of the proof of the theorem. Now we define the submartingales X_n^k and show that they satisfy all the properties of Theorem 1. We define X_n^k to be the length of the queue of the k-th server in the process $\mathcal{D}_n$, from which the constant Λ is subtracted. Clearly, $X_n^k \geq 0$. We now show that if K and Λ are both suitably large, then the properties (1) and (2) of Theorem 1 hold.

Relation (3) is evidently satisfied with $R = \Lambda$. Let us check (2). Let us start the process M at a configuration where all the queue lengths are of the form $X_0^k + \Lambda$ with $X_0^k > 0$, $k = 1, \dots, K$. We want to show that after time Λ, we have $\mathbb{E}(X_1^k - X_0^k) \geq \rho$, for some $\rho > 0$. Let $H = H(K)$ be the time after which the distribution of the K servers is almost uniform on C_K, see Lemma 5. Before this moment, we do not know much about our network, so we bound the lengths of the queues M_H^k roughly, by $M_H^k \geq M_0^k - H$. After the time H the probability that a customer leaving a server leaves the network is almost $1/K$, and the probabilities that it jumps to the left or the right are both close to $\frac{K-1}{2K}$.

More precisely, by Lemma 7, the rate of arrival to every server after time H is almost $\lambda + (K - 3)/K$, which is higher than the exit rate, 1, provided K is large enough (namely, $K > K^* = 3/\lambda$). Hence the expected queue lengths in the process M grow linearly in time, at least after time H, which implies the existence of $\Lambda > 0$ such that $\mathbb{E}(M_\Lambda^k) \geq M_0^k + \rho$. So, Theorem 1 applies.

III Infinite networks

3.1 NLMP on $\mathbb{Z}^1$

In this section we consider the limit of the network $(\mathbb{Z}^1)^N$ as $N \to \infty$, i.e. the NLMP on $\mathbb{Z}^1$. The limit of the network C_K^N can be studied in the same way. This NLMP is described in details in [2], and we use the notations therein. Here we are interested in its stationary distributions.

The NLMP is the evolution of the measure $\otimes \mu_v$ on the states (queues q_v) of the (jumping) servers at the nodes $v \in \mathbb{Z}^1$, given by the equations

$$\frac{d}{dt}\mu_v(q_v, t) = \mathcal{A} + \mathcal{B} + \mathcal{C} + \mathcal{D} + \mathcal{E} \tag{8}$$

with

$$\mathcal{A} = -\frac{d}{dr_{i^*(q_v)}(q_v)}\mu_v(q_v, t) \tag{9}$$

being the derivative along the direction $r(q_v)$ (in our case of the exponential service time with rate 1 we have, of course, that $\frac{d}{dr_{i^*(q_v)}(q_v)}\mu_v(q_v, t) = \mu_v(q_v, t)$)

$$\mathcal{B} = \delta\left(0, \tau(e(q_v))\right)\mu_v(q_v! e(q_v), t)[\sigma_{tr}(q_v! e(q_v), q_v) + \sigma_e(q_v! e(q_v), q_v)] \tag{10}$$

where q_v is created from $q_v! e(q_v)$ by the arrival of $e(q_v)$ from v', and $\delta\left(0, \tau(e(q_v))\right)$ takes into account the fact that if the last customer $e(q_v)$ has already received some amount of service, then he cannot arrive from the outside;

$$\mathcal{C} = -\mu_v(q_v, t)\sum_{q_v'}[\sigma_{tr}(q_v, q_v') + \sigma_e(q_v, q_v')], \tag{11}$$

which corresponds to changes in queue q_v due to customers arriving from other servers and from the outside (in the notations of (1), $\sigma_e(q_v, q^v \oplus w) = \lambda_{v,w}$);

$$\mathcal{D} = \int_{q_v': q_v'! C(q_v') = q_v} d\mu_v(q_v', t)\sigma_f\left(q_v', q_v'! C(q_v')\right) - \mu_v(q_v, t)\sigma_f\left(q_v, q_v! C(q_v)\right), \tag{12}$$

where the first term describes the situation where the queue q_v arises after a customer was served

in a queue q'_v (longer by one customer), and $q'_v! \, C(q'_v) = q_v$, while the second term describes the completion of service of a customer in q_v;

$$\mathcal{E} = \sum_{v' \text{n.n.} v} \beta_{vv'} [\mu_{v'}(q_v, t) - \mu_v(q_v, t)], \tag{13}$$

where the β-s are the rates of exchange of the servers.

For the convenience of the reader we repeat the equation $(8-813)$ once more:

$$\frac{d}{dt}\mu_v(q_v, t) = -\frac{d}{dr_{i^*(q_v)}(q_v)}\mu_v(q_v, t)$$

$$+\delta\left(0, \tau(e(q_v))\right)\mu_v\left(q_v! \, e(q_v)\right)[\sigma_{tr}(q_v! \, e(q_v), q_v) + \sigma_e(q_v! \, e(q_v), q_v)]$$

$$-\mu_v(q_v, t)\sum_{q'_v} [\sigma_{tr}(q_v, q'_v) + \sigma_e(q_v, q'_v)] + \int_{q'_v : q'_v! C(q'_v) = q_v} d\mu_v(q'_v)\sigma_f\left(q'_v, q'_v! \, C(q'_v)\right) \tag{14}$$

$$-\mu_v(q_v)\sigma_f\left(q_v, q_v! \, C(q_v)\right) + \sum_{v' \text{n.n.} v} \beta_{vv'}[\mu_{v'}(q_v) - \mu_v(q_v)].$$

We are looking for the fixed points μ of the evolution (14). Then the measures δ_μ (on measures) will be stationary measures of our NLMP. Note that the dynamical system (14) might have other stationary measures (on measures) then those corresponding to the fixed points. We will simplify our setting. Namely, we make the following changes:

1. for the graph G we take the lattice $\mathbb{Z}^1$;

2. all the customers have the same class;

3. the service time distribution η is exponential, with the mean value 1;

4. the service discipline considered is FIFO;

5. the exogenous customer c arriving to the node v has for its destination the same node v, i.e. $D(v) = v$; inflow rates at all the nodes are constant, equal to λ;

6. the two servers at v, v', which are neighbors in $\mathbb{Z}^1$ can exchange their positions with the same rate $\beta \equiv \beta_{vv'}$;

The queue q_v can in this setting be identified with the sequence of destinations $D(c_i)$ of its customers. The equation for the fixed point then becomes:

$$0 = \mu_v\left(q_v! \, e(q_v)\right)[\sigma_{tr}(q_v! \, e(q_v), q_v) + \sigma_e(q_v! \, e(q_v), q_v)]$$

$$-\mu_v(q_v)\sum_{q'_v} [\sigma_{tr}(q_v, q'_v) + \lambda] + \sum_{q'_v : q'_v! C(q'_v) = q_v} \mu_v(q'_v)$$

$$-\mu_v(q_v)\mathbb{I}_{q_v \neq \emptyset} + \sum_{v' = v \pm 1} \beta[\mu_{v'}(q_v) - \mu_v(q_v)].$$

We are interested in translation-invariant solutions. In that case the queue q_v can be identified with the sequence of (signed) distances between the node v and the destinations $D(c_i)$ of its customers, so it becomes a finite integer sequence $\mathcal{N} \equiv \{n_1, \dots, n_l; n_i \in \mathbb{Z}^1\}$, where $l \geq 0$ is the length of the queue q_v. The rate of the arrival of the transit customer, $\sigma_{tr}(q_v! \, e(q_v), q_v) \equiv \sigma_{tr}([n_1, \dots, n_{l-1}], [n_1, \dots, n_{l-1}, n_l])$ is then a function of one integer, n_l, and so we adopt the notation

$$\lambda_{n_l} \equiv \sigma_{tr}([n_1, \dots, n_{l-1}], [n_1, \dots, n_{l-1}, n_l]).$$

According to our definitions, we thus have

$$\lambda_k \equiv \lambda_k(\mu) = \begin{cases} \sum_\mathcal{N} \mu(k+1, \mathcal{N}) & \text{if } k > 0, \\ \sum_\mathcal{N} \mu(k-1, \mathcal{N}) & \text{if } k < 0, \\ 0 & \text{if } k = 0. \end{cases} \tag{15}$$

In what follows we look only for states μ which have symmetric rates λ_k :

$$\lambda_k = \lambda_{-k}. \tag{16}$$

The probability $\mu_v(q_v)$ then turns into $\mu(\mathcal{N})$; note, however, that for $v' = v \pm 1$ we need to interpret $\mu_{v'}(q_v)$ as $\mu(n_1 \mp 1, \dots, n_l \mp 1)$. The equation now becomes:

$$\mu(n_1, \dots, n_{l-1})[\lambda_{n_l} + \lambda\delta(n_l, 0)] - \mu(n_1, \dots, n_l)(\sum_k \lambda_k + \lambda)$$

$$+ \sum_k \mu(k, n_1, \dots, n_l) - \mu(n_1, \dots, n_l) \mathbb{I}_{l \neq 0} \tag{17}$$
$$+ \beta[\mu(n_1 + 1, \dots, n_l + 1) + \mu(n_1 - 1, \dots, n_l - 1) - 2\mu(n_1, \dots, n_l)] = 0.$$

As we see later, the equations $(15) - (1517)$ can have several solutions, one solution or no solution, depending on the value of the parameter λ. If μ is a solution of the equations $(15) - (1517)$ for some λ, then we denote by

$$\nu(\mu) = \sum_k \lambda_k(\mu)$$

the rate of the transit customers to every node in the state μ, and by $\eta(\mu)$ the rate of the total flow to every node in the state μ:

$$\eta(\mu) = \nu(\mu) + \lambda.$$

Theorem 8 *For every positive $\eta < 1$ there exist a unique value $\lambda(\eta)$ of the exogenous flow rate λ and the state μ_η on the set of the queues $\{\mathcal{N}\}$, satisfying the equations $(15) - (1517)$ with $\lambda = \lambda(\eta)$, such that*

$$\eta(\mu_\eta) = \eta.$$

Proof. Consider the process which is described just by the relation (17), with arbitrary parameters λ_k, $k = 0, \pm 1, \dots$ and λ. This is an ordinary queuing system with a single server and with infinitely many types of customers. The customer of type k arrives with rate λ_k (and with rate $\lambda_0 + \lambda$ for $k = 0$). Consider the random variable ξ_η, which is the total time a customer spends in such a server in the stationary state. It has exponential distribution, which depends only on $\eta = \sum_k \lambda_k + \lambda$ (and which does not depend on the type of the customer), namely, $\mathbb{E}(\xi_\eta) = (1 - \eta)^{-1}$.

Suppose a customer of type k arrives to such a server. When it leaves the server, its type is changed to $k + \tau_\eta$, where τ_η is a random (integer valued) variable. That change happens due to the β-terms in (17). By symmetry, $\mathbb{E}(\tau_\eta) = 0$. The distribution of τ_η is the following. Consider a random walker $W(t)$, living on $\mathbb{Z}^1$, which starts at 0 – i.e. $W(0) = 0$, and which makes ± 1 jumps with rates β. Then $\tau_\eta = W(\xi_\eta)$.

We now are going to present a choice of the rates λ_k and λ in such a way that the equations (15) is satisfied as well. Our choice of the rates λ_k, λ is related to the stationary distribution of a certain ergodic Markov process on $\mathbb{Z}^1$, which we describe now. Define the matrix of transition probabilities $P_1 = \{\pi_{st}\}$ by $\pi_{st} = \Pr(\tau_\eta = s - t)$. Of course, this Markov chain on $\mathbb{Z}^1$ is not positive recurrent since its mean drift is zero. Let P_2 be the second Markov chain, with transition probabilities

$$\rho_{st} = \begin{cases} 1 & \text{for } t > 0, \ s = t + 1, \\ 1 & \text{for } t = 0, \ s = 1, 0, -1, \\ 1 & \text{for } t < 0, \ s = t - 1, \\ 0 & \text{in other cases.} \end{cases}$$

I.e. P_2 is non-random map of $\mathbb{Z}_1$ into itself. Consider the composition Markov chain, with transition matrix Q being the product, $Q = P_1 P_2$. This chain is, in contrast, positive recurrent (it has a positive drift towards the origin), and it has a stationary state $q = \{q_k, k \in \mathbb{Z}^1\}$. We take

$$\lambda_k = \eta q_k, \ k \neq 0; \ \lambda = \eta q_0. \tag{18}$$

The relations (15) are satisfied since the process Q describes the evolution of the type of the customer in the stationary state of the process $(15) - (1517)$.

We now state some properties of the function $\lambda(\eta)$ as the parameter η varies in $(0,1)$.

Proposition 9 *There is a $\lambda_+ > 0$ such that, for any positive $\lambda < \lambda_+$, there are at least two different velues $\eta = \eta_-(\lambda)$ and $\eta = \eta_+(\lambda)$ satisfying the relation $\lambda(\eta) = \lambda$ and such that $\eta_-(\lambda) \to 0$ and $\eta_+(\lambda) \to 1$ as $\lambda \to 0$.*

Proof. Clearly, $\lambda(\eta) \to 0$ as $\eta \to 0$. We want to argue that $\lambda(\eta) \to 0$ also when $\eta \to 1$. Indeed, in this regime every customer spends more and more time waiting in the queue, so for every k the

probability $\Pr(\xi_\eta \leq k) \to 0$ as $\eta \to 1$. Therefore the distribution of the random variable τ_η becomes more and more spread out: for every k, $\Pr(|\tau|_\eta \leq k) \to 0$ as $\eta \to 1$. Therefore the same property holds for the stationary distribution q, and the claim follows from relations (18) and Proposition 9. In particular, it means that the following equation on η:

$$\lambda(\eta) = a > 0$$

has at least two solutions for small a: the corresponding η is either small or close to 1. Indeed, this follows from the continuity of $\lambda(\eta)$.
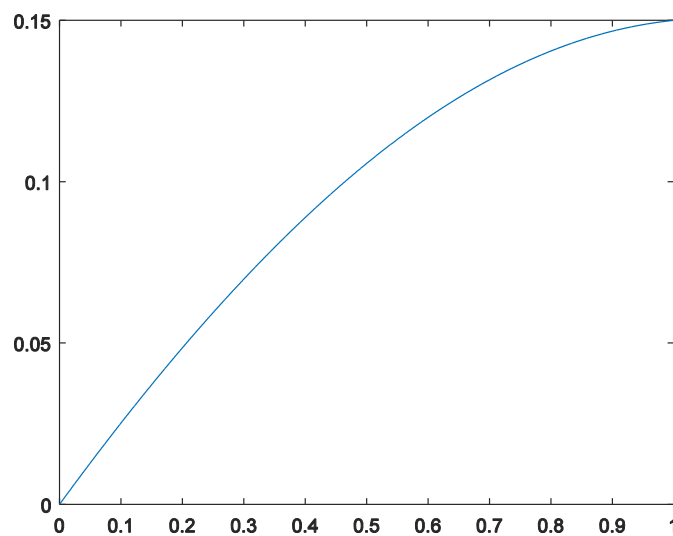
IV Some examples

Let us look at the function $\lambda(\eta)$ for some finite cyclic graphs C_K with different parameters, that is, with different randomized rules for the destination assignment. As we see, there are cases with one, two, and three equilibrium solutions.
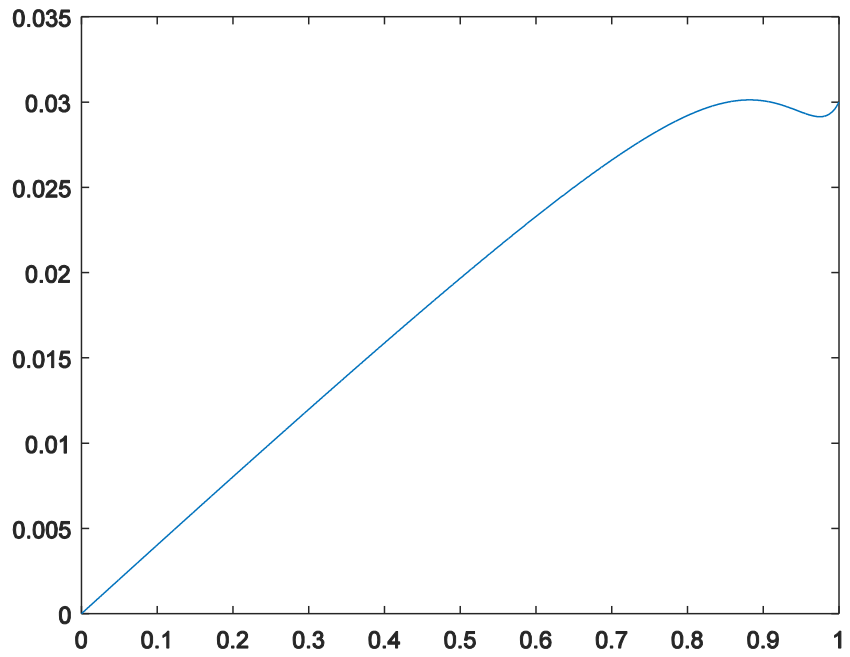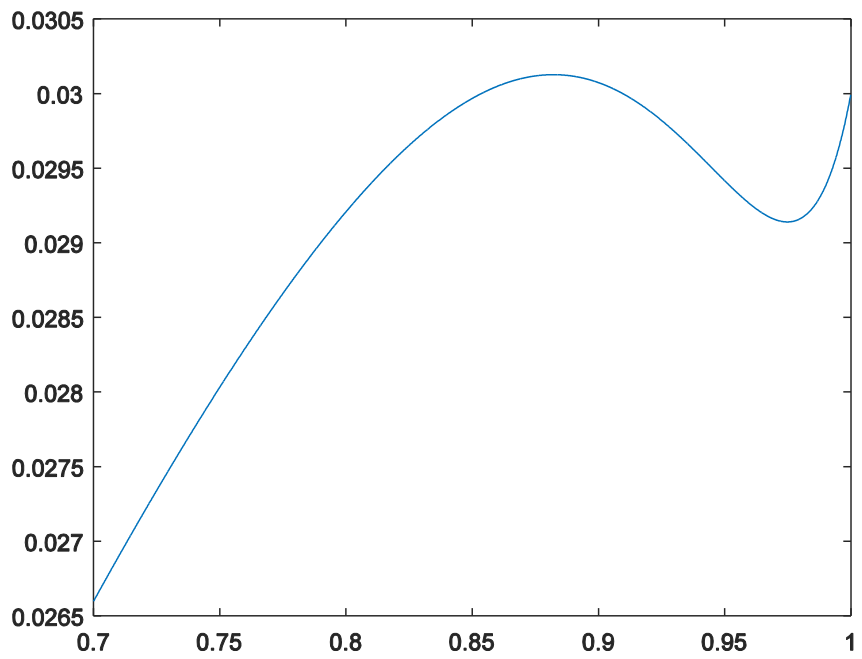
Typical case: two solutions
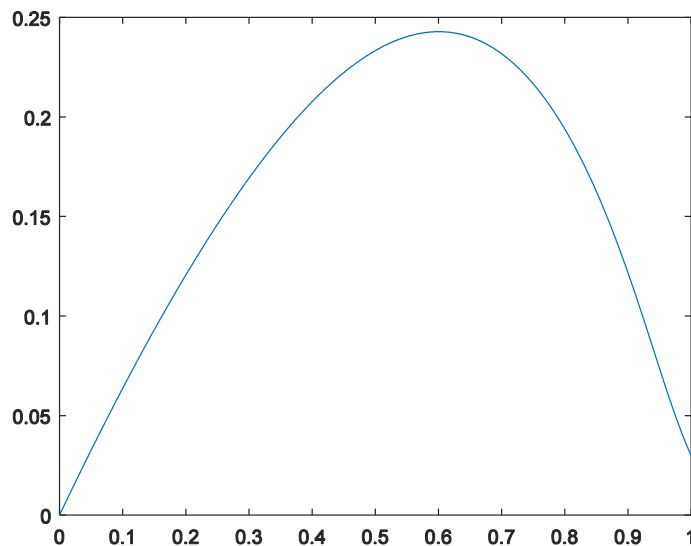
Single solution

Three solutions

Three solutions: closer look

References

[1] Emmanuel Baccelli and Charles E. Perkins. Multi-hop Ad Hoc Wireless Communication. Internet-Draft draft-baccelli-manet-multihop-communication-04, Internet Engineering Task Force, September 2014. Work in Progress.

[2] F. Baccelli, A. Rybko, and S. Shlosman. Queuing Networks with Varying Topology – A Mean-Field Approach. *ArXiv e-prints*, November 2013.

[3] François Baccelli and Bartlomiej Blaszczyszyn. Stochastic geometry and wireless networks, volume 1: Theory. *Foundations and Trends in Networking*, 3(3-4):249–449, 2009.

[4] François Baccelli and Bartlomiej Blaszczyszyn. Stochastic geometry and wireless networks, volume 2: Applications. *Foundations and Trends in Networking*, 4(1-2):1–312, 2009.

[5] Josh Broch, David A. Maltz, David B. Johnson, Yih chun Hu, and Jorjeta Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. pages 85–97, 1998.

[6] R. Durrett. *Probability: Theory and Examples*. Cambridge series on statistical and probabilistic mathematics. Cambridge University Press, 2010.

[7] C.S.R. Murthy and B.S. Manoj. *Ad Hoc wireless networks: architectures and protocols*. Prentice Hall communications engineering and emerging technologies series. Prentice Hall PTR, 2004.

[8] Roberto H Schonmann and Senya B Shlosman. Wulff droplets and the metastable relaxation of kinetic ising models. *Communications in mathematical physics*, 194(2):389–462, 1998.

[9] C.K. Toh. *Ad Hoc Mobile Wireless Networks: Protocols and Systems*. Prentice Hall PTR, 2002.