HAL
open science

The Power of Two Choices on Graphs: the Pair-Approximation is Accurate

Nicolas Gast

To cite this version:

Nicolas Gast. The Power of Two Choices on Graphs: the Pair-Approximation is Accurate. Workshop on MAthematical performance Modeling and Analysis, Jun 2015, Portland, United States. hal-01199271

HAL Id: hal-01199271

https://inria.hal.science/hal-01199271v1

Submitted on 15 Sep 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Power of Two Choices on Graphs: the Pair-Approximation is Accurate*

Nicolas Gast
Inria
Univ. Grenoble Alpes
CNRS, LIG, F-38000 Grenoble, France
nicolas.gast@inria.fr

The power of two-choice is a well-known paradigm to improve load balancing where each incoming task is allocated to the least loaded of two servers picked at random among a collection of n servers [6, 4]. We study the power of two-choice in a setting where the two servers are not picked independently at random but are connected by an edge in an underlying graph. Our problem is motivated by systems in which choices are geometrically constrained (see the model of bike-sharing systems introduced in [1, Section 4]).

We study a dynamic setting in which jobs leave the system after being served by a server to which is was allocated. Our focus is when each server has few neighbors (typically 2 to 4) for which an mean-field approximation is not accurate. The static counterpart of our model is studied in [2] in which it is shown by counting the number of arrivals on an edge that the power of two-choice does not hold when the degree is small. This technique cannot be used for studying the dynamic setting as the departures induce long-range dependence. The process is N-dimensional and has no product-form stationary distribution. An exact analytic solution seems out of reach. We use pair-approximation, a technique wide-spread in biology [5]. We build the equations and show that they describe accurately the steady-state of the system. Our results show that, even in a graph of degree 2, choosing between two neighboring improve dramatically the performance compared to a random allocation.

1. GEOMETRIC TWO-CHOICE MODEL

Our system is composed of n identical servers that are connected by an undirected graph (V, E), where the set of vertexes is the set of servers $V = \{1 \dots n\}$. Each server serves jobs at rate μ and uses a first-come first-serve discipline. Jobs arrive in the system at rate $n\lambda$. For each incoming job, one server, say s_1, is picked uniformly at random among the n servers. Then, another server s_2 is picked uniformly at random among the neighbors of s_1. The job is then allocated to the server s_1 or s_2 that has the least number of jobs (ties are broken at random). This allocation scheme is similar to the one of [2]. We denote the load by $\rho = \lambda/\mu$ and assume that $\rho < 1$. We now describe a few examples that we will explore numerically in Section 3.

*This work is partially supported by the EU project QUANTICOL, 600708.

Copyright is held by author/owner(s).

Example 0: The complete graph (classical two-choice model). The classical two-choice model [4, 6] corresponds to a complete graph: any server is a neighbor of any other server. Jobs are allocated to the least loaded of two servers chosen independently among the n servers. This model is the most studied as there are almost-closed-form results for this model (see Section 2.1).

Example 1: Two choices on a ring. In this scenario, two stations s_1 and s_2 are neighbors if $s_1 = s_2 \pm 1$ (modulo n). An equivalent representation of this model is to consider that for each $s \in \{1 \dots n\}$, jobs arrive at rate λ and are allocated to the shortest queue among the one of s or $s+1$ (modulo n). This is illustrated on Figure 1.

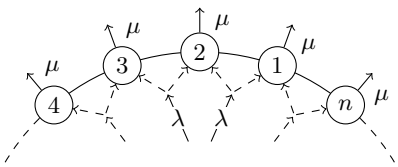

Figure 1: Ring model. Jobs arrive at rate $N\lambda$ in the system. For each arrival, $i \in \{1 \dots n\}$ is picked at random and the job is allocated to the server i or $i+1$ (mod n) that has the least number of jobs.

Example 2: Two choices on a 2D torus. In this case, we assume that $\sqrt{n}$ is an integer. The n servers are placed on a $2D$ grid. The position of a server s is represented by its coordinates (x_s, y_s). It has four neighbors: $(x_s \pm 1, y_s)$ and $(x_s, y_s \pm 1)$ (modulo $\sqrt{n}$). This case is shown on Figure 2a.

Example 3: Regular random graph. In Section 3, we also simulate random graphs with fixed degree. For each server, k neighbors are picked at random (for simplicity of generation, we allow self loop, like node 5 of Figure 2b). The interaction graph remains constant during the simulation.

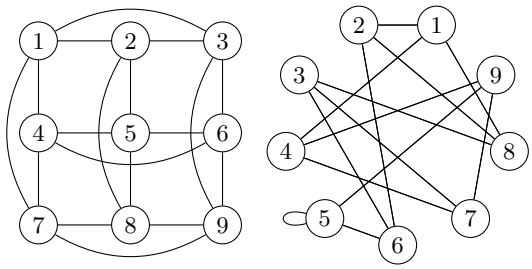

(a) 2D torus (b) Fixed degree $k = 3$

Figure 2: 2D torus and random graph model

2. THE PAIR APPROXIMATION EQUATIONS

In this section, we first recall some results and basic methodology for the classical two-choice model (on a complete graph). Then we develop our pair-approximation equations.

2.1 The mean-field approximation

When the interaction graph is complete, all servers are exchangeable. For $i \in \{0, 1, 2, \dots\}$, let $X_i(t)$ be the proportion of servers that have i jobs at time t. There is a departure from a server with i jobs at rate $\mu X_i(t)$. When there is an arrival, the two chosen servers have i and j jobs with probability $X_i(t)X_j(t)$. If $i = j$ or $i < j$, the job is allocated on a server with i jobs. If $j < i$, the job is allocated on a server with j jobs. As a result, there is an arrival in a server with i jobs at rate $\lambda(X_i(t))^2 + 2\lambda \sum_{j=i+1}^{\infty} X_i(t)X_j(t) = 2\lambda X_i(t)P_i(t)$, where $P_i(t) = X_i(t)/2 + \sum_{j=i+1}^{\infty} X_j(t)$ is the probability that an arrival on a server with i jobs is allocated to this server.

This shows that $X(t) = (X_0(t), X_1(t), \dots)$ is a density dependent population process [3]. A classical approximation of $X(t)$ is to consider the following system of ordinary differential equations, called the mean-field equation:

$$\frac{dx_i}{dt} = \mu(x_{i+1} - x_i \mathbf{1}_{\{i>0\}}) + 2\lambda(\mathbf{1}_{\{i>0\}} p_{i-1} x_{i-1} - p_i x_i), \quad (1)$$

where $p_i = x_i/2 + \sum_{j=i+1}^{\infty} x_j$. By using results from [3], it is shown in [4] that, as n goes to infinity, the stationary distribution of the system concentrates on the unique fixed point of (1) which is such that the number of servers having i or more jobs equals $\rho^{2^i - 2}$. This fixed point is a very good approximation of the steady state of the original system, even for $n = 100$. This method can be generalized to $d \geq 2$ choices in which case $\rho^{2^i - 1}$ becomes $\rho^{(d^i - 1)/(d-1)}$. Note that when jobs are allocated a server at random ($d = 1$), the proportion of servers with i or more jobs is ρ^i. The power of two-choice refers to the fact that two choices improves the situation by an exponential factor compared to one but three or more only improves marginally compared to two.

2.2 The pair-approximation

We now consider a general interaction graph in which all nodes have the same degree k. Let $Y_{i,j}(t)$ the proportion of connected pairs of servers that have (i, j) jobs and $X_i(t) = \sum_j Y_{ij}(t)$ the proportion of servers that have i jobs. When the graph is complete, $Y_{ij}(t) = X_i(t)X_j(t)$, which implies that $X(t)$ is a density dependent process. When the graph is not complete, this does not hold. A randomly chosen neighbor of a randomly chosen server having i jobs has j jobs with probability $Y_{ij}(t)/X_i(t)$. Hence, an arrival on a server that has i jobs is allocated to this server with probability $Q_i(t) = (Y_{ii}(t)/2 + \sum_{j=i+1}^{\infty} Y_{ij}(t))/X_i(t)$.

We now look at the evolution of $Y_{ij}(t)$. Let (i, j) be the state of a pair of servers connected by an edge. This state becomes $(i-1, j)$ when there is a departure on i, which occurs at rate μ if $i \geq 1$. It becomes $(i+1, j)$ when there is an arrival on i. This can be caused by two types of events: (a) arrival on the edge – if each node has k neighbors, an edge (i, j) is chosen at rate $2\lambda/k$ and the packet is allocated to the first server with probability $a(i, j) = 1$ if $i < j$, $a(i, i) = 1/2$ and $a(i, j) = 0$ if $i > j$ – or (b) arrival on another neighbor of the first server – each other neighbor of i that has state ℓ induces an arrival on i at rate $2\lambda a(i, \ell)/k$. Let $Z_{\ell,i,j}(t)$ be the proportion of connected triplets of stations having state (ℓ, i, j). The arrivals on the first server of a pair (i, j) from one of the $k - 1$ other neighbors occur at rate $2\lambda(k - 1)R_{ij}(t)/k$, where $R_{ij}(t) = (Z_{i,i,j}(t)/2 + \sum_{\ell=i+1}^{\infty} Z_{\ell,i,j}(t))/Y_{ij}(t)$.

This shows that, as $X(t)$, the process $Y(t)$ is not a density dependent process because the rates of its transitions involve quantities that depend on triplets. In what follows, we consider a density dependent population process that is an approximation of the original process and has the same transitions but with different rates: in all the rates that involve a quantity $R_{ij}(t)$, we replace this quantity[1] by $Q_i(t)$. This approximation is called the pair-approximation and leads to the following differential equation for y_{ij}:

$$\begin{aligned} \frac{dy_{ij}}{dt} &= \mu\left(y_{i+1,j} - y_{i,j}\mathbf{1}_{\{i>0\}} + y_{i,j+1} - y_{i,j}\mathbf{1}_{\{j>0\}}\right) \quad (2) \\ &+ \frac{2\lambda}{k}\left(y_{i-1,j}a(i-1,j) + y_{i,j-1}a(j-1,i) - y_{ij}\right) \\ &+ \frac{2\lambda(k-1)}{k}\left(q_{i-1}y_{i-1,j}\mathbf{1}_{\{i>0\}} + q_{j-1}y_{i,j-1}\mathbf{1}_{\{j>0\}} - (q_i+q_j)y_{i,j}\right) \end{aligned}$$

The first line of this equation corresponds to the rate of changes of the proportion of pairs (i, j) induced by the departures; the second line is for the arrival on the pair (i, j) and the last line on the arrival on the neighbors or i and j. The first two lines are exact while the last involves the approximation $z_{\ell,i,j} \approx y_{i,\ell}y_{i,j}/x_i$.

3. NUMERICAL EVALUATION

In this section, we compare numerically the steady states of the three examples of Section 1 with the fixed-point of the pair-approximation equation (2). We did the comparison for values from $\rho = 0.5$ to $\rho = 0.99$ and only a subset of the results are reported here. All tested values show that the pair-approximation provides an excellent approximation of the stationary distribution.

The computation of the steady-state distribution of examples 1, 2 and 3 of Section 1 is obtained by running a discrete-event simulator. In all cases, we simulate a system with $n = 1000$ servers for a total $T = 10^{11}$ events. Comparisons with smaller values of T indicate that $T = 10^{11}$ is enough to reach the steady-state. The fixed point of the pair-approximation equations is computed by integrating numerically the system of differential equations (2).

In Figure 3, we report the steady-state probability x_i that a given server has i jobs as a function of i. Each plot compares five curves: two are obtained by simulation – (ring/random graph with fixed degree $k = 2$) for the first two plots and (2D torus/random graph with $k = 4$) for the last one –, and two are the fluid approximations of the model (mean field and pair approximation with $k = 2$ or $k = 4$). The last curve corresponds to a model without choices (each server is an independent M/M/1 queue) and is here for comparison. These results show that the pair-approximation predicts very accurately the general shape of the distribution of the simulated model, which are far from both the one-choice and the mean-field approximation[2]. The tail of the distribution, however, does not seem to be correct, even if it is much closer for the pair-approximation than for the

[1] Alternatively, the same equation can be obtained by replacing $Z_{\ell,i,j}$ by $Y_{i,\ell}(t)Y_{i,j}(t)/X_i(t)$.

[2] Recall that the mean-field approximation is very accurate for the complete graph but here as the graph is sparse.

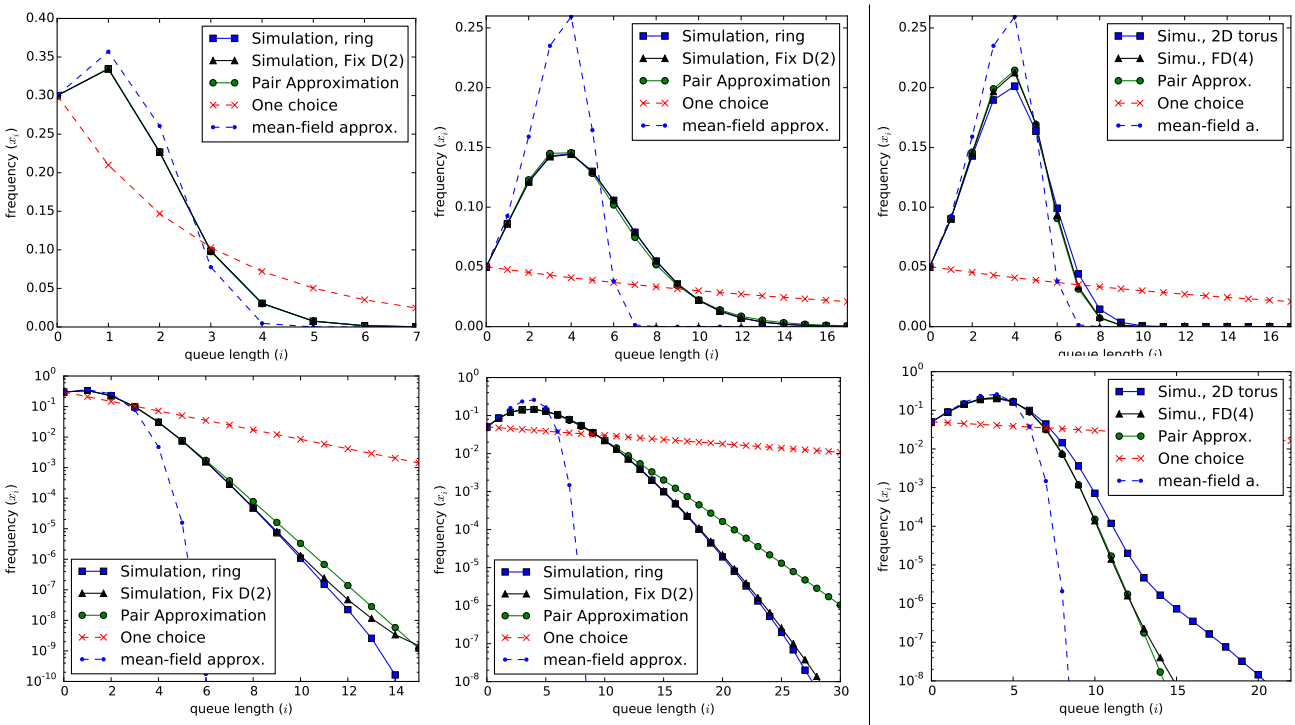

$\lambda = 0.7$, Ring or Fixed deg. $(k = 2)$ $\lambda = 0.95$, Ring or Fixed deg. $(k = 2)$ $\lambda = 0.95$, 2D torus or Fix. deg. $(k = 4)$

Figure 3: Steady-state probability for a server to have i jobs as a function of i. We compare values obtained by simulations and two fluid approximations (mean-field and pair-approximation). Second line is in log-scale.

mean-field model.

In Figure 4, we also report the average queue length as a function of ρ for the various models. This shows that, when the simulated interaction graph has a fixed degree k, the pair-approximation with the same value k is a good approximation of the average queue length. Also, when $k = 4$, the average queue length is already very close to the mean-field approximation. We also simulated a Erdos-Renyi interaction model in which two nodes are connected with probability k/n, for $k \in \{2, 4\}$. The results (not reported here) show that in this case the pair-approximation with constant k is not a good approximation. The reason is that in an Erdos-Renyi graph, all nodes do not have the same degree. There is a non-negligible proportion of nodes with degree 0 that drives the performance close to the one-choice case.

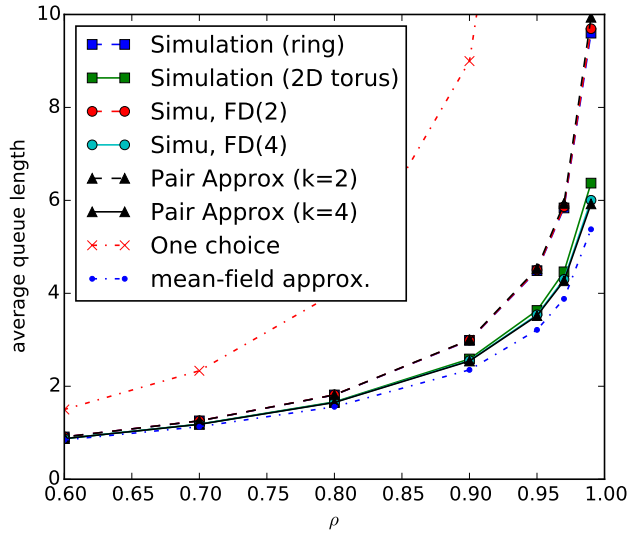

Figure 4: Average queue length as a function of ρ. For a given k, the pair approximation is a good approximation of the graph with constant degree k.

4. REFERENCES

[1] C. Fricker and N. Gast. Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity. *EURO Journal on Transportation and Logistics*, pages 1–31, 2014.

[2] K. Kenthapadi and R. Panigrahy. Balanced allocation on graphs. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 434–443. SIAM, 2006.

[3] T. G. Kurtz. Strong approximation theorems for density dependent markov chains. *Stochastic Processes and Their Applications*, 6(3):223–240, 1978.

[4] M. Mitzenmacher. *The power of two choices in randomized load balancing*. PhD thesis, University of California at Berkeley, Department of Computer Science, 1996.

[5] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502–505, 2006.

[6] N. D. Vvedenskaya, R. L. Dobrushin, and F. I. Karpelevich. Queueing system with selection of the shortest of two queues: An asymptotic approach. *Problemy Peredachi Informatsii*, 32(1):20–34, 1996.