

State Dependent Control of Closed Queueing Networks

Extended Abstract*

Siddhartha Banerjee
Cornell University
Ithaca, NY
sbanerjee@cornell.edu

Yash Kanoria
Columbia University
New York, NY
ykanoria@columbia.edu

Pengyu Qian
Columbia University
New York, NY
pq2124@columbia.edu

ABSTRACT

We study the design of state dependent control for a closed queueing network model, inspired by shared transportation systems such as ridesharing. In particular, we focus on the design of *assignment policies*, wherein the platform can choose which supply unit to dispatch to meet an incoming customer request. The supply unit subsequently becomes available at the destination after dropping the customer. We consider the proportion of dropped demand in steady state as the performance measure.

We propose a family of simple and explicit state dependent policies called Scaled MaxWeight (SMW) policies and prove that under the complete resource pooling (CRP) condition (analogous to a strict version of Hall's condition for bipartite matchings), any SMW policy induces an exponential decay of demand-dropping probability as the number of supply units scales to infinity. Furthermore, we show that there is an SMW policy that achieves the optimal exponent among all assignment policies, and analytically specify this policy in terms of the matrix of customer-request arrival rates. The optimal SMW policy protects structurally under-supplied locations.

CCS CONCEPTS

• **Mathematics of computing** → Queueing theory; Markov processes; Stochastic control and optimization; • **Theory of computation** → Stochastic control and optimization; • **Applied computing** → Transportation;

KEYWORDS

maximum weight policy; closed queueing network; control; Lyapunov function; shared transportation systems; large deviations

ACM Reference Format:

Siddhartha Banerjee, Yash Kanoria, and Pengyu Qian. 2018. State Dependent Control of Closed Queueing Networks: Extended Abstract. In *SIGMETRICS '18 Abstracts: ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems Abstracts, June 18–22, 2018, Irvine, CA, USA*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3219617.3219619>

*This document is an extended abstract for [1].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS '18 Abstracts, June 18–22, 2018, Irvine, CA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5846-0/18/06.

<https://doi.org/10.1145/3219617.3219619>

1 INTRODUCTION

Recently there is an increasing interest in the control of shared transportation systems such as Uber and Lyft. These platforms are dynamic two-sided markets where customer requests (demand) arrive at different physical locations stochastically over time, and supply units circulate in the system as a result of driving demand to their destinations. The platform's goal is to maximize throughput (proportion of demand fulfilled), revenue or other objectives by employing various types of controls. The main inefficiency comes from the geographic mismatch of supply and demand: when a customer arrives, she has to be matched immediately with a nearby supply unit, otherwise the customer will abandon the request due to impatience.

We study the *assignment policy* as an exclusive control lever to mitigate the loss of demand, where the platform can choose which supply unit to assign when a customer request comes in. Previous work has studied assignment decisions made in a *state independent* manner through randomized policy based on fluid limit solution. However, this approach requires exact knowledge of arrival rates (which is infeasible in practice), and fails to react to the stochastic variation in the system. Although this control guarantees asymptotic optimality in the Law of Large Numbers sense, it converges only slowly to the fluid limit. To counter these issues, we study the *state dependent* assignment control of shared transportation systems.

2 PROBLEM FORMULATION

We model the system as a closed queueing network with n servers representing physical locations, and K 'jobs' that stand for supply units. This is a common model for shared transportation systems. For the K -th system, its state at any time t is given by $\mathbf{X}^K[t]$, a vector that tracks the number of supplies at each location at time t .

- **Compatibility Graph.** For pedagogical reasons, we move to a setup where we distinguish formally between demand locations V_D where customers arrive and supply locations V_S where supply units wait (and where customers are dropped off). We encode the compatibility graph as a bipartite graph $G(V_S \cup V_D, E)$. An edge $(i, j') \in E$ represents a compatible pair of supply and demand nodes, i.e., supplies stationed at i can serve demand arriving at j' . We denote the neighborhood of a supply node $i \in V_S$ (resp. demand node $j' \in V_D$) in G as $\partial(i)$ (resp. $\partial(j')$); thus, for a supply node i , its compatible demand are given by $\partial(i) = \{j' \in V_D | (i, j') \in E\}$, and similarly for each demand node. Moreover, for any set of supply nodes $A \subseteq V_S$, we also use $\partial(A)$ to denote its demand neighborhood (and vice versa).




Figure 1: The bipartite compatibility graph for the assignment problem: On the left are supply nodes i , and on the right are demand nodes i' . Customers arrive at i' with distributions chosen according to $\phi_{i'}$; the total rate of a customer arrival at i' is thus $1^T \phi_{i'}$. Similarly, assuming no demand is dropped, the total rate a supply unit arrives at i is $1^T \phi_{(i)}$. The edges entering a node i' encode compatible (i.e., nearby) nodes that can supply node i' .

- **Demand Arrivals.** Demands with origin i' and destination j arrives at location i' according to an independent Poisson process with rate $\phi_{i'j}$. We denote the i -th column of $\phi \triangleq [\phi_{ij}]$ (i.e., the arrival rates from different origins of customers to destination i) as $\phi_{(i)}$, and the transpose of the i' -th row of ϕ (i.e., the arrival rates of customers to node i' with different destination nodes) as $\phi_{i'}$. Thus, the customer arrival rate at node i' is $1^T \phi_{i'}$, and, assuming all customers are matched, the rate of supply units arriving at node i is $1^T \phi_{(i)}$. We assume that matrix $[\phi_{ij}]$ is irreducible.
- **Dispatch Policies.** Each time a customer arrives, the platform makes an assignment decision from a compatible supply location based on the current spatial distribution of available supplies. The problem we want to study is how to design assignment policies which minimize the probability of dropped demand. For fixed K , this problem can be formulated as an average cost Markovian decision process on finite state space, and is thus known to admit an optimal stationary policy (i.e., ones where assignment decisions at time t only depend on the system state $\mathbf{X}^K[t]$).
- **System Evolution.** Equivalently, we study the discrete jump chain of the continuous-time process. At the beginning of time-slot t , the state of the system is $\mathbf{X}^K[t-1]$. Suppose the platform uses an assignment policy U , and in time-slot t , a customer arrives at origin node $o[t]$ with destination $d[t]$ (chosen from arrival matrix ϕ with probability proportional to ϕ_{ij}). If the platform assigns a supply unit from $S[t]$ (potentially $\emptyset$) under policy U , the assigned supply unit will pick up the demand and relocate to $d[t]$ instantly. Formally, we have

$$\mathbf{X}^K[t] = \begin{cases} \mathbf{X}^K[t-1] - \mathbf{e}_{S[t]} + \mathbf{e}_{d[t]}, & \text{if } S[t] \in V_S. \\ \mathbf{X}^K[t-1], & \text{if } S[t] = \emptyset. \end{cases}$$

- **Objective.** The platform's goal is to maintain adequate supply in all neighborhoods and hence meet as much as demand as possible, therefore we adopt the (global) proportion of dropped demand $P^{K,U}$ as a measure of efficiency. To obtain tight characterizations, we further consider the asymptotic regime where the number of supply units in the system K goes to infinity and focus on the decay rate of $P^{K,U}$:

$$\gamma(U) = -\liminf_{K \rightarrow \infty} \frac{1}{K} \log P^{K,U}.$$

Main Assumption: Complete Resource Pooling. A main assumption in our model is the complete resource pooling (CRP) condition. For the model considered in this paper, the CRP condition is closely related to the condition in Hall's marriage theorem in bipartite matching theory.

Assumption 1. We assume that for any $J \subsetneq V_D$ where $J \neq \emptyset$,

$$\sum_{i \in \partial(J)} 1^T \phi_{(i)} > \sum_{j' \in J} 1^T \phi_{j'}.$$

The intuition behind this assumption is clear: it assumes the system is "balanceable" in that for each subset $J \subsetneq V_D$ of demand locations we have enough supply at neighboring locations to meet the demand. We show that the CRP condition is necessary for any assignment policy to have demand dropping probability that converges to zero.

3 SCALED MAXWEIGHT POLICIES

A popular assignment policy is MaxWeight policy (hereafter referred to as vanilla MaxWeight). In our setting, MaxWeight would correspond to assigning from the compatible location with most supplies (with appropriate tie-breaking rules). The popularity of MaxWeight scheduling stems from the fact that it is known to be optimal for different metrics in various problem setting. However, in our setting, we show that vanilla MaxWeight does not achieve the optimal exponent in general.

We slightly generalize vanilla MaxWeight by attaching a positive scale-down parameter w_i to each queue $i \in V_S$, and assign from the compatible queue with largest *weighted* queue length X_i/w_i . We call this family of policies *Scaled MaxWeight (SMW) policies*, and denote SMW with parameter $\mathbf{w}$ as SMW($\mathbf{w}$).

The formal definition of SMW is as follows.

Definition 1 (Scaled MaxWeight). Given system state $\mathbf{X}[t-1]$ at the start of t -th period and for demand arriving at $o[t]$, SMW($\mathbf{w}$) assign from

$$\operatorname{argmax}_{i \in \partial(o[t])} \frac{X_i[t-1]}{w_i}$$

if $\max_{i \in \partial(o[t])} \frac{X_i[t-1]}{w_i} > 0$; otherwise the demand is dropped. If there are ties when determining the argmax, assign from the location with highest index.

4 MAIN RESULTS AND CONTRIBUTIONS

As a function of system primitives, we derive a large deviation rate-optimal assignment policy that minimizes dropped demand (maximizes throughput). Our optimal policy is strikingly simple and its parameters depend in a natural way on demand arrival rates. Our contribution is threefold:

- (1) **Achievability:** We show that for any $\mathbf{w} > 0$, SMW($\mathbf{w}$) achieves exponential decay of the demand dropping probability with exponent ¹

$$\gamma(\mathbf{w}) = \min_{J \in \mathcal{J}} \left(\mathbf{1}_{\partial(J)}^T \mathbf{w} \right) \log \left(\frac{\sum_{i' \notin J} \sum_{j \in \partial(J)} \phi_{i'j}}{\sum_{i' \in J} \sum_{j \notin \partial(J)} \phi_{i'j}} \right) > 0.$$

Here $\mathcal{J}$ includes the subsets such that the denominator inside logarithm is strictly positive². Hence all SMW policies with $\mathbf{w} > 0$ guarantee exponential decay of demand-dropping probability in K under the CRP condition³. SMW policies are simple, explicit and appear promising for practical applications.

- (2) **Converse bounds:** We provide lower bound of demand dropping probability for any assignment policy using random-walk related inequalities. We first show that for any state independent assignment policy π , it cannot achieve exponential decay rate:

$$p^{K, \pi} = \Omega \left(\frac{1}{K^2} \right),$$

which demonstrates the value of state dependent control - even a naive state dependent assignment policy with no knowledge of demand arrival rates beats the best state independent assignment policy asymptotically. Then we justify the CRP assumption by showing that it is a necessary condition for exponential convergence. Further, we obtain the surprising finding that the optimal SMW policy is, in fact, exponent-optimal among all state dependent policies: under any policy U , it must be that

$$\gamma(U) \leq \gamma^*,$$

where

$$\gamma^* = \sup_{\mathbf{w} > 0} \gamma(\mathbf{w}).$$

- (3) **Qualitative insights:** We characterize the system behavior under SMW policies as K goes to infinity, which is technically challenging since the problem remains n -dimensional even in the limit. We solve the multi-dimensional variation problem that characterizes the exponent base on a family of novel Lyapunov functions (a different one for each SMW policy). We establish the critical subset property of the problem: given a system state (in the limit), there exists a (state dependent) subset J of demand locations that are most likely to be depleted of supply in compatible locations, hence leading to demand dropping. The optimal SMW policy avoids using supply from locations compatible with J for demand arising outside of J .

¹We emphasize that for SMW policies, the limit exists, which further does not depend on the starting state $X^{K, U}[0]$.

²When $\mathcal{J} = \emptyset$ there exists a policy that never drops demand in steady state, see [1] for detail.

³Note that the argument of the logarithm has a larger numerator than denominator for every $J \subsetneq V_D$ since CRP holds, implying that $\gamma(\mathbf{w})$ is the minimum of several positive numbers, and hence is positive.

5 DISCUSSION

Before closing, we point out a few more features of our result:

Novel Lyapunov analysis for a closed queueing network. A key technical challenge we face in our closed queueing network setting is that it is a priori unclear what the “best” state is for the system to be in. This is in contrast to open queueing network settings in which the best state is typically the one in which all queues are empty, and the Lyapunov functions considered typically achieve their minimum at this state. We get around this issue via an innovative approach where we define a tailored Lyapunov function that achieves its minimum at the resting point of the SMW policy we are analyzing, and use this Lyapunov function to characterize its exponent $\gamma(\mathbf{w})$. Moreover, given the optimal choice of $\mathbf{w}$, our tailored Lyapunov function corresponding to this choice of helps us establish our converse result.

Choosing $\mathbf{w}$ where ϕ is imperfectly known. Another key issue is that of knowledge of the true arrival rates ϕ . We remark that if these rates are entirely unknown, but CRP holds, the platform can pick an SMW policy such as vanilla max weight, and be sure to achieve exponential decay of the demand dropping probability (albeit with a suboptimal exponent). This is already an improvement over the state independent control policies studied thus far – in particular, any state independent policy will drop an $\Omega(1)$ fraction of demand if there is any model misspecification whatsoever. If ϕ is not precisely known but is known to lie within some set, that setting may lend itself to a natural robust optimization problem of maximizing the demand dropping exponent, worse case over possible ϕ . We leave this as an open question.

Future directions: Computational challenges. A limitation of our main results is that it does not prescribe how to compute $\mathbf{w}$. It simply specifies a concave maximization problem that must be solved, one where the objective is the minimum over an exponential number of linear functions. We suspect that structural properties of “realistic” G and ϕ can be exploited to make this problem tractable; for the moment, we leave it as an open question.

ACKNOWLEDGMENTS

YK and PQ acknowledge support from the National Science Foundation under grant no. CMMI 1653477, and SB acknowledges support from Army Research Office grant W911NF-17-1-0094.

REFERENCES

- [1] Siddhartha Banerjee, Yash Kanoria, and Pengyu Qian. 2018. The Value of State Dependent Control in Ridesharing Systems. *arXiv preprint arXiv:1803.04959* (2018).