Libera Università Internazionale
degli Studi Sociali Guido Carli
PREMIO TESI D'ECCELLENZA

Routing Games, Information and Learning

Francesco Giordano

2021-2022

LUISS

Libera Università Internazionale
degli Studi Sociali Guido Carli

Working Papaer n. 8/2021-2022
Publication date: January 2024
Routing Games, Information and Learning

ISBN 979-12-5596-094-2

Luiss Academy is an imprint of
Luiss University Press – Pola Srl
Viale Pola 12, 00198 Roma
Tel. 06 85225485
E-mail lup@luiss.it
www.luissuniversitypress.it

Routing Games, Information and Learning

By Francesco Giordano

1. Introduction and Contribution

In the thesis we explore various aspects of models related to *Selfish Routing*. These models, known as *Routing Games*, provide a way to describe and simulate real-world routing scenarios, like traffic on roads and data transmission in telecommunication networks.

Initially, we delve into the traditional setup, which involves a static game without uncertainty. In this scenario, traffic flows from a group of starting points to a set of destinations. We discuss key findings related to equilibria, optimality, and efficiency in this context.

Subsequently, we investigate recent contributions that introduce elements of uncertainty. This uncertainty can manifest in different ways, such as variations in traffic demands or changes in the costs experienced by the participants. Real-life traffic situations often involve various sources of randomness, which have been the subject of prior theoretical studies. When analyzing how uncertainty influences routing behavior, we pay particular attention to incomplete information that can be revealed through the concept of Bayesian rationality.

We examine routing instances where an unknown factor affects the costs associated with edges, and the objective is to minimize the expected cost for each route. It becomes possible to learn the true underlying factor when these incomplete information games are repeatedly played, each time with a random level of demand. Notably, two types of social learning emerge in non-atomic routing games: *strong learning*, where the true network state is identified, and *weak learning*, where players learn how to make decisions as if they knew the true state.

The original theoretical contribution of this thesis is presented in the second part. A recent study by Macault et al. (2022) clarified the conditions under which *social learning* occurs in single-commodity non-atomic routing games with unlimited edge capacities. The authors demonstrated that strong learning is almost surely guaranteed if the network follows a specific topological condition called *series-parallel*. Building on this foundation, we precisely define when a network state is *identifiable* and expand the network conditions necessary to achieve almost certain learning, including in scenarios with multiple commodities and capacity limitations. This thesis extends the existing results in two distinct directions: (i) it considers models with multiple commodities, (ii) it accommodates edges with traffic capacity limits.

With strictly increasing, continuous, and unbounded costs, we find that the learning conditions differ between instances with capacity limitations and those with unlimited edge capacities. Specifically, in instances with capacity limitations, a sufficient condition for achieving learning is that the network satisfies a condition where, at full congestion, the load on each edge reaches an upper bound. This condition is considered a milder form of *capacity conservation*, and it doesn't impose any restrictions on the network's topology. In contrast, in instances with unlimited edge capacities, a sufficient condition for achieving learning is that the sub-network available to each commodity follows a *series-parallel* structure.

2. DETERMINISTIC ROUTING GAMES

A traffic scenario can be represented as a non-cooperative game played on a network. In a selfish routing game, various traffic flows, each referred to as a *commodity*, need to be directed from their sources to their respective destinations within a directed graph. The assignment of traffic to specific paths within the network is known as *flow*.

In the context of non-atomic games, the set of players is defined based on infinitely divisible traffic masses, with each player having a negligible impact on the costs associated with network edges.

A *non-atomic* instance is defined as $G=(N,d,\gamma,c)$, that is, by a network, a vector of traffic demands, a vector of edge capacities, and a vector edge cost functions. When edge capacities are unlimited, an instance is described only by the network, costs, and demands.

The game is endowed with set of available paths for the players. The amount of traffic assigned to each path defines a *flow*. Specifically, a flow is a vector that specifies, for each path, a non-negative quantity representing the traffic routed through that path. A flow vector is considered feasible when it successfully routes all traffic demand through the network while adhering to edge capacity limits.

The *load* on an edge is defined as the amount of traffic that passes through that edge. A flow is in equilibrium when it exclusively employs paths with the minimum costs, preventing any advantageous redistribution of traffic. In non-atomic games, an equilibrium flow is also referred to as a *Nash Flow* or a *Wardrop Equilibrium*. In cases with unlimited edge capacity, the costs of all utilized paths are equal, and they are lower than the costs incurred by any potential traffic distribution on unused paths.

A Wardrop equilibrium in a non-atomic routing game always exists. This is due to the fact that routing games are potential games with continuous sets of players. This can be demonstrated through the equivalence of Nash flows and local minima of a *potential function* that entirely characterizes the game's costs. Furthermore, in non-atomic scenarios, equilibrium costs are essentially unique, meaning that multiple equilibrium flows may exist, but all equilibrium flows lead to the same costs.

In contrast, an *optimal flow* is a flow that minimizes the total cost experienced by the players. Optimal flows also always exist, as they are solutions to a convex optimization problem within a closed and bounded set.

In instances with capacity limits, the Wardrop principle of equilibrium no longer holds true. In other words, equilibrium costs of paths are not guaranteed to be equal, and equilibrium loads are not always unique. The equilibrium concept for instances with capacity constraints analogous to the Wardrop Equilibrium is known as *capacitated user equilibrium*.

3. ROUTING GAMES UNDER UNCERTAINTY

This section delves into the complexities of *Social Learning* in a *Dynamic Non-atomic Routing Game* where the network setup can have either infinite or finite edge capacities. We explore the main ideas and results in a less technical manner.

To begin, we consider the presence of various sources of uncertainty that can influence routing decisions. This leads us to explore ways to model this uncertainty within deterministic routing games. In our work, we particularly focus on a model involving incomplete information, where traffic demands and costs are subject to randomness. Importantly, this randomness in costs depends solely on the *underlying network state*, which is initially unknown to the players. In this context, we examine routing instances where the true network conditions are represented through an unknown network state. Players hold a common belief about this unknown state and make routing decisions based on their beliefs.

We consider a finite set of possible network states, with an unknown random state that is realized before the game is played. Players share a common belief about the unknown state.

In the case of an instance with an unknown state, a flow vector is considered to be in equilibrium if it minimizes the expected costs of paths, given the public belief. This means that in an instance with an unknown state, a flow is at equilibrium when no commodity has a path with a strictly smaller expected cost than any other path being used.

We define a *Dynamic Non-atomic Routing Game* as a routing scenario repeated over time, where each stage game features random traffic inflow. Time is discrete, and networks are characterized by a total capacity, that can be infinite.

Traffic demands are modeled as a sequence of independent identically distributed non-negative random vectors, with each having common marginal distributions bounded by network capacities. In each period, a traffic demand is realized and publicly observed. Players choose an equilibrium flow, and information about equilibrium load and costs is shared with all players.

Through repeated play of the game, players collectively leverage the randomness in traffic demands to uncover the randomness in the network state. The public belief is updated in each stage game according to Bayes' rule, and when a state is identified through equilibrium costs, the public belief changes.

We introduce the concept of the posterior public belief, which depends on the random demands up to a generic period t. As a probability distribution, it is always bounded, making the sequence of posterior beliefs a bounded martingale. By the martingale convergence theorem, this sequence converges.

In this context, we consider two ideas of social learning: *Strong Learning* and *Weak Learning*. Strong Learning implies that the sequence of posteriors converges almost surely to a distribution assigning positive probability solely to the true state. On the other hand, Weak Learning suggests that the true state may not necessarily be discovered, but in equilibrium and asymptotically, the total traffic is routed as if the true state were known.

In the section on *Social Learning in non-atomic routing games*, we further explore the problem of social learning within a Dynamic Non-atomic Routing Game. We investigate this problem in scenarios with both infinite and finite edge capacities.

Before presenting our key results, we delve into the structural properties of capacity-constrained networks. We note that in capacitated networks, edge loads are upper-bounded, and we characterize this bound, which may not always match the edge's capacity. We redefine *feasible loads* in terms of these upper bounds and use this foundation to determine when an unknown network state is *identifiable*. This preliminary analysis is crucial for understanding when social learning is possible in networks with capacity constraints.

Ultimately, we demonstrate that social learning occurs in instances with finite capacities when the underlying network adheres to a condition on its edge capacities known as *weak capacity conservation*. Additionally, we show that social learning occurs in instances with infinite capacities if the underlying network exhibits a *series-parallel* structure in each sub-network available to its source-destination pairs.

In scenarios with finite capacities, the maximum load achievable on an edge may be less than the edge's capacity. However, this maximum load is determined by the set of edge capacities, and we characterize this *load upper bound*. This upper bound is significant in defining when an unknown network state is *identifiable* and can be learned. We establish that a load is feasible if and only if it simultaneously adheres to the upper bound for all its edges.

We then determine the *feasibility* of a flow by considering two conditions: one ensuring that edge capacities are respected, and the other ensuring that all traffic is successfully routed. We define the set of feasible flows as the set in which both conditions are met.

With a clear understanding of when a flow vector is feasible, we characterize the conditions for identifying an unknown network state. An unknown state is *identifiable* if, for all pairs of possible state realizations, there exists an edge such that, for at least one load value, there are two costs with different values in different states.

In networks with capacity constraints, equilibrium loads may not necessarily reach the edge capacity for all edges. This can hinder strong learning. Specifically, if the load value required to identify states cannot be achieved, strong learning is not possible. Our thesis provides examples where learning does not occur in this context.

4. CAPACITY CONSERVATION

In the following discussion, we introduce the concept of *Capacity Conservation*, focusing on the total capacity of a network. This total capacity represents the sum of the capacities of the edges in the smallest *cut* of the network. Specifically, a cut in a directed multigraph with a single source and destination is a set of edges that makes it impossible to route from the source to the destination without crossing this set of edges. The capacity of a cut is determined by adding up the capacities of the individual edges within it. The *smallest cut* refers to the cut with the smallest capacity, and the capacity of this smallest cut represents the network's overall capacity. It indicates the maximum amount of traffic that can be routed through the network.

In the context of a multi-commodity network, which is the combination of sub-networks available to each commodity, we can still refer to a cut as a set of edges that makes it impossible to go from any source to any destination without crossing this set. In this case as well, the capacity of the network corresponds to the capacity of the smallest cut.

To proceed, we introduce a network condition, known as *capacity conservation*, that enables social learning in a Dynamic Non-atomic Routing Game (DNRG) played on a capacitated network. Moving forward, for each internal node, we define the sets of entering edges into the node and exiting edges from the node itself. The *capacity conservation* property is satisfied if, for all internal nodes that are neither the source nor the destination, the sum of the capacities of the edges entering that node equals the sum of the capacities of the edges exiting from that node. In simpler terms, for any internal node that is not the source or destination, the sum of capacities of the edges entering that node is equal to the sum of capacities of the edges exiting from that node.

A network that fully respects capacity conservation ensures that when the traffic demand matches the network capacity, all edge loads reach their load upper bound. In other words, if the traffic demand equals the network capacity, all edge loads reach their maximum capacity. This capacity conservation property is highly desirable. Consequently, we introduce a less strict form of it, known as *weak capacity conservation*. In our definition of weak capacity conservation, for a given routing instance and for each edge, the network is said to respect *weak capacity conservation* if, when the traffic demand matches the network capacity, all edge loads reach their load upper bound.

In our dynamic model with random traffic demand, an exact match between the realization and the network capacity occurs with probability zero. However, due to the strictly increasing nature of edge costs, equilibrium loads are continuous with respect to the demand. Therefore, we can deduce that when the support of traffic demand is from zero to the network capacity, the support of equilibrium loads lies in a interval between zero and the (equilibrium) load upper bound, that is a deterministic non-negative function of the demand realization.

This argument is similarly applicable to multi-commodity instances. However, when different commodities share common edges, we must carefully consider the allocation of capacity on these shared edges. This predetermined allocation on common edges allows us to determine the capacity of the sub-network available to each commodity.

Under this construction, capacity conservation implies that the total network capacity equals the sum of the capacities of the sub-networks available to each commodity. In a multi-commodity network, we can extend the concept of weak capacity conservation to assert that when the traffic demand for each commodity matches the capacity of the sub-network available to that commodity, all edge loads reach their load upper bound. In summary, when the support of the multi-commodity traffic demand falls within the range from zero to the capacity of that subnetwork, for each commodity the support of equilibrium loads lies in an interval within zero and the equilibrium load upper bound for each edge.

5. SP NETWORKS

Next, we turn our attention to networks with potentially infinite edge capacities. We begin by introducing the concept of single-commodity *series-parallel* networks. These networks possess a specific topology that aligns with favorable properties, reducing the likelihood of learning failure.

The critical point to consider is that if a network isn't series-parallel, it contains sub-networks in which some edges remain unused in equilibrium, even when the demand is high.

To capture similar favorable properties in a multi-commodity setting, we examine scenarios where each set of paths available to any given commodity forms itself a series-parallel network. We refer to the resulting network as *U-SP*, signifying that each commodity routes its traffic on a series-parallel network.

6. LEARNING IN CAPACITY CONSTRAINED NETWORKS

To briefly recap, players initially start with a prior distribution on the network state, which is common knowledge. In each stage, a random demand is realized, and the equilibrium flow is played. All newly acquired information about the experienced costs is immediately shared with all players, who update their collective belief following Bayes' rule. Strong learning takes place when the posterior public belief almost surely converges to a degenerate distribution that assigns positive probability solely to the true state.

Now, let's outline the main results of social learning in capacity-constrained networks. Certain fundamental cost assumptions are as follows. First, two assumptions on costs: (i) edge costs strictly increase with load for all possible states, (ii) edge costs are continuous with load for all possible states. Also, the following limit condition holds: as the load on an edge approaches its capacity, the cost becomes infinitely large for all network states.

Here are the theorems that describe learning in capacity-constrained networks:

Theorem 1: Learning in Single-Commodity Capacitated Networks
In a single-commodity Dynamic Non-atomic Routing Game (DNRG) with an identifiable unknown network state, and where the network respects weak capacity con-

servation and satisfies the assumptions of above, strong learning occurs if the support of the random demand falls within the range between zero and the network capacity.

Proof Idea for Theorem 1
Under these conditions, for any potential edge load between zero and the load upper bound, there exists a demand realization that generates it. Due to the continuity of costs with respect to loads and loads with respect to demands, it's enough to reach a value within a small neighborhood of the required demand that ensures a change in belief. Because of the support condition on the demand, each of these neighborhoods has a positive probability.

Theorem 2: Learning in Multi-Commodity Capacitated Networks
In a multi-commodity DNRG with an identifiable unknown network state, where the network respects weak capacity conservation and satisfies the above assumptions, strong learning occurs if the joint distribution of the demand vector has a support such that for each commodity the support of demand of the given commodity falls within the range from zero to the capacity of the subnetwork available to that commodity.

In summary, strong learning is achieved in both single-commodity and multi-commodity networks with identifiable unknown network states that respect weak capacity conservation and satisfy the stated cost assumptions, provided that the demand falls within specific support ranges.

7. LEARNING IN NETWORKS WITH INFINITE CAPACITIES

The preliminary investigation carried out by Macault et al. (2022) concentrated on scenarios where networks have unlimited capacities, specifically in the single origin-destination case. The authors derived the following conclusions:

In a single-commodity Dynamic Non-atomic Routing Game (DNRG) with an identifiable unknown network state, assuming the network has a series-parallel (SP) topology, edges possess infinite capacities, and the above assumptions are met, the occurrence of strong learning is established when the random demand's support extends to zero to infinity.

A similar outcome is observed in instances with multiple commodities.

Theorem 1: Learning in Multi-Commodity Infinite Capacity Networks
In a multi-commodity DNRG with an identifiable unknown network state, assuming the network has a U-SP topology, edges have infinite capacities, and the assumptions of above are satisfied, strong learning is achieved when the support of the demand vector covers the entire positive real space.

The learning process remains consistent with the single origin case, where equilibrium edge loads exhibit continuity concerning demands and remain unbounded.

8. CONCLUSION

In this thesis, we delve into recent developments in the traffic assignment problem. In a selfish routing game, traffic must be directed from various sources to specific destinations within a directed graph. The distribution of traffic on the network's paths is referred to as *flow*.

The first part of the thesis studies routing games under complete information and introduces concepts like equilibrium flows, optimal flows, and routing efficiency. Subsequently, we explore models that come into play when dealing with uncertainty. These models involve routing instances in which the network's state, which affects edge costs, is unknown. In these situations, traffic is routed to minimize the expected cost on each path. Learning the true underlying state becomes feasible when these games involving incomplete information are repeatedly played, with a random level of demand in each stage. In the realm of non-atomic routing games, two types of social learning emerge: *strong learning*, where the true network state is identified, and *weak learning*, where players adapt their strategies as if they knew the true state.

The second part of the thesis is dedicated to exploring the conditions under which social learning can be achieved. We define when a network state is *identifiable* and outline the network conditions necessary for almost certain learning. With a focus on strictly increasing, continuous, and unbounded costs, we discover that the learning conditions differ between instances with capacity constraints and those with unlimited edge capacities. In capacitated instances, a sufficient condition for achieving learning is ensuring that, under full congestion, the load on each edge reaches its upper bound. This condition is a milder form of *capacity conservation*, where the total capacity entering and exiting each node remains equal. On the other hand, in instances with infinite edge capacities, learning can be achieved if the sub-network available to each commodity adheres to a *series-parallel* structure.

9. BIBLIOGRAPHY

Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. "Informational Braess' Paradox: The Effect of Information on Traffic Congestion." *Operations Research*, 66(4), 2018, pp. 893-917.

Itai Arieli and Manuel Mueller-Frank. "A General Analysis of Sequential Social Learning." *Mathematics of Operations Research*, 46(4), 2021, pp. 1235-1249.

Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated Games with Incomplete Information*. MIT Press, 1995.

Abhijit V Banerjee. "A Simple Model of Herd Behavior." *The Quarterly Journal of Economics*, 107(3), 1992, pp. 797-817.

Martin Beckmann, Charles B McGuire, and Christopher B Winsten. *Studies in the Economics of Transportation*. Cowles Commission for Research in Economics by Yale University Press, 1956.

Petra Berenbrink and Oliver Schulte. “Evolutionary Equilibrium in Bayesian Routing Games: Specialization and Niche Formation.” *Theoretical Computer Science*, 411(7-9), 2010, pp. 1054-1074.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades.” *Journal of Political Economy*, 100(5), 1992, pp. 992-1026.

Xujin Chen, Zhuo Diao, and Xiaodong Hu. “Network Characterizations for Excluding Braess’s Paradox.” *Theory of Computing Systems*, 59(4), 2016, pp. 747-780.

Roberto Cominetti. “Equilibrium Routing Under Uncertainty.” *Mathematical Programming*, 151(1), 2015, pp. 117–151.

Roberto Cominetti, Marco Scarsini, Marc Schröder, and Nicolas E Stier-Moses. “Price of Anarchy in Stochastic Atomic Congestion Games with Affine Costs.” In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 2019, pp. 579-580.

Roberto Cominetti, Marco Scarsini, Marc Schröder, and Nicolas Stier-Moses. “Convergence of Large Atomic Congestion Games.” arXiv preprint arXiv:2001.02797, 2020.

Roberto Cominetti, Valerio Dose, and Marco Scarsini. “The Price of Anarchy in Routing Games as a Function of the Demand.” *Mathematical Programming*, 2021, pp. 1-28.

José R Correa and Nicolás E Stier-Moses. “Wardrop Equilibria.” In *Encyclopedia of Operations Research and Management Science*. Wiley, 2011.

José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. “Selfish Routing in Capacitated Networks.” *Mathematics of Operations Research*, 29(4), 2004, pp. 961-976.

José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. “A Geometric Approach to the Price of Anarchy in nonatomic Congestion Games.” *Games and Economic Behavior*, 64(2), 2008, pp. 457-469.

Richard J Duffin. “Topology of Series-Parallel Networks.” *Journal of Mathematical Analysis and Applications*, 10(2), 1965, pp. 303-318.

Michael J. Evans and Jeffrey S. Rosenthal. “Probability and Statistics: The Science of Uncertainty.” *The American Statistician*, 59, 2005, pp. 276-276.

Caroline Fisk. “More Paradoxes in the Equilibrium Assignment Problem.” *Transportation Research Part B: Methodological*, 13(4), 1979, pp. 305-309.

Lester Randolph Ford and Delbert Ray Fulkerson. *Flows in Networks*. Princeton University Press, 1962.

Françoise Forges. “Games with Incomplete Information: From Repetition to Cheap Talk and Persuasion.” *Annals of Economics and Statistics*, 2020, pp. 3-30.

Drew Fudenberg, Fudenberg Drew, David K Levine, and David K Levine. *The Theory of Learning in Games*, Volume 2. MIT Press, 1998.

Martin Gairing, Burkhard Monien, and Karsten Tiemann. “Selfish Routing with Incomplete Information.” *Theory of Computing Systems*, 42(1), 2008, pp. 91-130.

Jacob K Goeree, Thomas R Palfrey, and Brian W Rogers. “Social Learning with Private and Common Values.” *Economic Theory*, 28(2), 2006, pp. 245-264.

Allan Gut. *Probability: A Graduate Course*, Volume 200-5. Springer, 2005.