Fast Convergence to Wardrop Equilibria by Adaptive Sampling Methods

Gouleakis Themistoklis

June 2, 2011

Problem definition

The problem we are going to deal with has the following properties:

- The game is a selfish routing game divided into rounds.
- There is an infinite number of agents each responsible for an infinitesimal amount of traffic.
- In each round, each agent samples an alter- native routing path and compares the latency on this path with its current latency.
- In the next round all the agents have the opportunity to choose a different path (simultaneusly).

Problem:The latency of some agent may increase!
Even worse: the game may get stuck in oscillations (and never reach an equilibrium).
Solution:
Let the agents sample alternative routes at random and migrate with a probability depending on the observed latency difference.

Wardrop's traffic model

- We consider a model for selfish routing where an infinite population of agents carries an infinitesimal amount of load each
- Let E denote a set of resources (edges).
- Continuous, non-decreasing latency functions $e : [0, 1] \rightarrow R^+$.
- A set of commodities with flow demands or rates $r_i, i \in [k]$ such that $\sum_{i=1}^{k} r_i = 1$.
- For every commodity $i \in [k]$ let $P_i \subseteq 2^E$ denote a set of strategies (paths) available for commodity i.
- Let $P = \cup_{i\in[k]} P_i$ and let $L = max_{p\in P} |p|$.

An instance is symmetric if k = 1 and asymmetric otherwise. An instance is single-resource if for all $p \in P$, $|p| = 1$.

Definition: Wardrop equilibrium

A feasible flow vector $(f_p)_{p \in P}$ is at a Wardrop equilibrium for the instance Γ if for every commodity $i \in [k]$ and every $p, p' \in P_i$ with $f_p > 0$ it holds that $l_p(f) \leq l_{p'}(f)$.

Potential function:

$\Phi(f) = \sum_{e \in E} \int_0^{f_e} l(x) dx$

- The set of allocations in equilibrium coincides with the set of allocations minimizing the potential function.
- Our goal is the design of distributed rerouting policies that approximate the Wardrop equilibrium.

Shifted potential

- Observe, however, for certain instances of the routing game, Φ^* might be zero. In this case, we suggest to shift the potential by some positive additive term.
- So, we get an α-shifted potential.
- $\Phi^* + \alpha$ is strictly positive.
- This is equivalent to adding a virtual amount of to the latency observed on every path.

Definition: Relative slope

A differentiable latency function l has relative slope d at x if $l'(x) \leq d \cdot \frac{l(x)}{x}$. A latency function has relative slope d if it has relative slope d over the entire range $[0, 1]$ and a class of latency functions $\mathcal{L}$ has relative slope d if every $l \in \mathcal{L}$ has relative slope d.

Related to the derivative of $xl(x)$.
Examples: polynomials and exponentials.

Rerouting policy

- In every round, an agent is activated with constant probability $\lambda = 1/32$.
- Then he performs the following two steps:
  1. Sampling: With probability $(1-\beta)$ perform step 1(a) and with probability β perform step 1(b).
     (a) Proportional sampling: Sample path $Q \in P_i$ with probability $\frac{f_Q}{r_i}$.
     (b) Uniform sampling: Sample path $Q \in P_i$ with probability $\frac{1}{|P_i|}$.
  2. Migration: If $l_Q < l_P$, migrate to path Q with probabillity $\frac{l_p - l_Q}{d(l_p+\alpha)}$
- The parameter β must be chosen subject to the constraint

$$\beta \leq \frac{min_{p \in P} l_p(0) + \alpha}{L * max_{e \in E} max_{x \in [0,\beta]} l'_e(x)} \qquad (1)$$

Definition: Exploration - replication policy

For an instance Γ let $d \geq 1$ be an upper bound on the relative slope of the latency functions and let β be chosen as in Equation (1). For every commodity $i \in [k]$ and every path $P, Q \in \mathcal{P}_i$ with $l_Q \leq l_P$, the (α, β)-exploration-replication policy migrates a fraction of

$$\mu_{PQ} = \lambda \cdot \frac{1}{d} \left((1-\beta) \cdot \frac{f_Q}{r_i} + \beta \cdot \frac{1}{|\mathcal{P}_i|} \right) \frac{l_P - l_Q}{l_P + \alpha}$$

with $\lambda = \frac{1}{32}$ agents from path P to path Q.

Fact

Let Γ be an instance of the congestion game and let $\Gamma^{+\alpha}$ be an instance that we obtain from Γ by inserting a new resource e_P for every $P \in \mathcal{P}$ with constant latency function $l_{e_P}(x) = \alpha$. Let Φ and $\Phi^{+\alpha}$ denote the respective potential functions.

1. The (α, β)-exploration-replication policy behaves on *Gamma* precisely as the $(0, \beta)$-exploration-replication policy does on $\Gamma{+}\alpha$.
2. If $\Phi^{+\alpha}(f) \leq (1+\epsilon)(\Phi^{+\alpha})$, then $\Phi(f) \leq (1+\epsilon)\Phi + \epsilon\alpha$.

Definition

For two flow vectors f and f' of consecutive rounds, the virtual potential gain is the potential gain that would occur if the latencies were fixed at the beginning of the round, i. e.

$$V(f, f') = \sum_{e \in E} l_e(f)(f'_e = f_e)$$

By our policy, this value is always negative.

Lemma

Consider an instance Γ and the (α, β)-exploration-replication policy changing the flow vector from f to f' in one step. Then we have $\Delta\Phi = \Phi(f') - \Phi(f) \geq \frac{1}{2} \sum_{P,Q \in \mathcal{P}} \mu_{PQ}(l_Q - l_P) = \frac{V(f,f')}{2}$.

Definition: $\delta - \epsilon$ equilibrium

For a flow vector f let $\mathcal{P}^{+}(\delta) = \{P \in \mathcal{P} | l_P(f) \geq (1+\delta)\bar{l}(f)\}$ denote the set of δ-expensive strategies and let $\mathcal{P}(\delta) = \{P \in \mathcal{P} | l_P(f) \leq (1-\delta)\bar{l}(f)\}$ denote the set of δ-cheap strategies. The population f is in a $\delta - \epsilon$-equilibrium iff at most ϵ agents utilize δ-expensive and δ-cheap strategies. We write $\mathcal{P}^{+}$ and $\mathcal{P}^{-}$ if δ is clear from the context.

Theorem

Consider a symmetric congestion game Γ *and an initial flow vector* f_{init}. *For the* (α, β)*-exploration-replication policy, the number of rounds in which the population vector is not* $\delta - \epsilon$*-equilibrium w.r.t* $\Gamma^{+\alpha}$ *(as defined in Fact 3) is bounded from above by:*

$$\mathcal{O}\left(\frac{d}{\epsilon\delta^2}\log\left(\frac{\Phi(f_{init}) + \alpha}{\Phi* + \alpha}\right)\right)$$

Lemma

Consider a symmetric routing game and a flow at $\delta - \epsilon$-equilibrium . If the (α, β)-exploration-replication policy changes the average latency ℓ in one round by $\Delta > 10\lambda \cdot (2\epsilon + 2\delta + \beta)\bar{\ell}$, it reduces the potential Φ by at least $\Delta/(10(\delta + 1))$.

Definition (δ-Equilibrium)

A population vector f is at a δ-equilibrium if for every commodity $i \in [k]$ and for every $P \in \mathcal{P}_i$ it holds that $\ell_P(f) \geq \bar{\ell}_i - \delta\bar{\ell}$ and, in addition, if $f_P > 0$, $\ell_P(f) \leq \bar{\ell}_i + \delta\bar{\ell}$.

Single-resource

Theorem

Consider a symmetric single-resource instance Γ and an initial flow vector f_{init}. If $\beta \leq \epsilon/\delta$, the (α, β)-exploration-replication policy generates a configuration with potential $\Phi \leq (1+\epsilon)\Phi^* + \epsilon\alpha$ in at most

$$\mathcal{O}\left(\frac{d^{12}}{\epsilon^7}\log^4\left(\frac{|E|}{\beta}\right)\log\left(\frac{\Phi(f_{init})+\alpha}{\Phi* + \alpha}\right)\right)$$

rounds

Single-resource

Theorem

Consider a symmetric instance Γ and an initial flow vector f_{init}. If $\beta \leq \epsilon^2/(L^3\delta^2)$ then the (α, β)-exploration-replication policy generates a configuration with potential $\Phi \leq (1+\epsilon)\Phi^* + \epsilon\alpha$ in at most

$$poly\left(d, \frac{1}{\epsilon}, L\right) \frac{d^{12}}{\epsilon^7} \log^4\left(\frac{|E|}{\beta}\right) \log\left(\frac{\Phi(f_{init}) + \alpha}{\Phi * + \alpha}\right)$$

rounds

Definition (δ-ϵ-Equilibrium)

For a flow vector f, for every commodity $i \in [k]$, let $\mathcal{P}_i^+(\delta) = \{P \in \mathcal{P}_i \,|\ell_P(f) \geq \overline{\ell}_i(f) + \delta\bar{\ell}\}$ denote the set of δ-expensive strategies and let $\mathcal{P}_i^-(\delta) = \{P \in \mathcal{P}_i \,|\ell_P(f) \leq \overline{\ell}_i(f) - \delta\bar{\ell}\}$ denote the set of δ-cheap strategies. The population f is called an δ-ϵ-equilibrium iff at most ϵ agents utilize δ-expensive and δ-cheap strategies.

Theorem

Consider an asymmetric congestion game Γ *and an initial flow vector* f_{init}*. For the* (α, β)*-exploration-replication policy, the number of rounds in which the population vector is not at a* δ*-*ϵ*-equilibrium w.r.t* $\Gamma^{+\alpha}$ *(as defined in Fact 3) is bounded from above by:*

$$\mathcal{O}\left(\frac{d}{\epsilon^2\delta^2}\log\left(\frac{\Phi(f_{init})+\alpha}{\Phi* + \alpha}\right)\right)$$

In particular, this bound holds for $a = \beta = 0$ *(and hence* $\Gamma^{+\alpha} = \Gamma$*).*

Relative slope is necessary

Theorem

For every d, there exists a class L of latency functions with relative slope d together with an initial flow vector f , such that any Markovian rerouting policy monotone for L requires $\Omega(d/\sqrt{e})$ rounds in order to obtain a $(1+\epsilon)$ approximation to the optimum potential.

Sampling with static probabillities is slow

Theorem

For every m, there exist a set of resources E with $|E| = m$ and strategy set P with $|P| = 2m/4$ such that for every rerouting policy with static sampling probabilities for P there exist a set of latency functions $(l_e)_{e \in E}$ and an initial population such that the rerouting policy needs at least $\Omega(|P| log(1/\varepsilon))$ rounds to reach a $(1+\varepsilon)$-approximation of the optimal potential for the symmetric instance $\Gamma = (E, P, (l_e))$.