Topics in Game Theory
Spring 2009
Lecture 5: Modeling BGP as A Game

Lecturer: Aaron Jaggard
Scribe: Aleksandar Nikolov

March 9, 2009

1 Overview

In the last lecture we formally defined stability in the Border Gateway Protocol (BGP) as the Stable Paths Problem (SPP). We showed that computing a stable paths assignment is NP-Complete, and proved a condition (dispute wheel freeness) under which stable paths do exist and are unique. In this lecture, we will formulate routing as a game, and formulate applicable solution concepts. We will explore under what conditions the routing game is well-behaved.

2 One-round formulation

In the real world routing is usually an interaction between independent and not necessarily cooperative entities (e.g. ISP's), motivated to maximize their individual utility. For this reason, routing is a good candidate to be modeled and analyzed as a game.

We start with a setting similar to that of the stable paths problem. Let $G = (V, E)$ be an undirected rooted graph. The set of vertices is $V = \{d, 1, \ldots, n\}$, where d is the destination vertex. Let $\mathcal{P}$ be the set of paths from i to d. For every $i \in 1, \ldots, n$ there exists a *valuation function* $\nu_i : \mathcal{P}_i \to \mathbb{Z}$, s.t. $P_1, P_2 \in \mathcal{P}_i, P_1 \neq P_2 \Rightarrow \nu_i(P_1) \neq \nu_i(P_2)$. Let also $N(i)$ denote the set of neighbor nodes of i.

We can now define a one-round routing game [LSZ08]:

- For a player i, the set of strategies $S_i = N(i)$, i.e. i can choose to pick one of its neighbors to route through.

- The utility of player i is

$$u_i = \begin{cases} v_i(P), & \text{if } P \text{ is induced by the routing choices of } 1, \ldots, n \\ 0, & \text{otherwise} \end{cases}.$$

The Nash equilibria in this game are precisely the stable solutions in the equivalent SPP formulation.

Recall that finding a stable solution in SPP is an NP-Complete problem. With the one-round game formulation, we want to know the communication complexity of reaching an equilibrium.

Theorem 1 ([LSZ08]) *If all nodes strictly prefer paths of length at most x to paths of length greater than x, then determining whether there exists a pure Nash equilibrium in the one-round game requires $\Omega(2^x)$ bits of communication between the source nodes in the worst case.*

3 Multiround games

3.1 Setting

The one-round game allows little interaction between players, and for this reason is very similar to the centralized SPP formulation. The following multiround extensions raises more questions pertaining to game theory and Nash equilibria.

The setting for the multiround game is the same as for the one-round version: n players, a destination d, an undirected graph G on $d, 1, \ldots, n$, and valuation functions v_i. Utility is also defined as before. However, the game consists of infinitely many rounds. Specifically, in each round a scheduler selects ≥ 1 nodes to play which:

- learn updates from neighbors;
- pick a neighbor to forward to; the neighbor should be the most preferred one according to the valuation function v_i, i.e. it should induce the most preferred path given the other players' current choices;
- announce routes to d.

Note that the players are not restricted to know the path they announce or to announce the path they actually selected to use. I.e. they can invent paths as well as be untruthful about their choices.

Given the valuation function v_i, a player i's strategy's should define import and export as a function of v_i. We have a notion of 'fair' strategy, which we call *best response dynamics* (BRD). It is defined by three conditions:

1. learn most recent update from all neighbors;
2. pick best path P according to v_i (this information is private);

3. announce P to **all** neighbors.

Note that any strategy should be fully specified. If we want to define a strategy which deviates from BRD, we should still specify unambiguously what paths are announced and to what subset of the node's neighbors.

3.2 Ex post Nash equilibrium

This game deviates from the games we've considered so far. We want to keep v_i private and still have a notion of equilibrium. Nash equilibrium is defined only for perfect information games. Therefore, we need a different solution concept.

We begin with a definition from mechanism design:

Definition 1 *A distributed mechanism specification* $d_m(g, \Sigma, s^m)$ *consists of the following components:*

- g *is the output function which maps strategies to outputs;*
- $\Sigma = \Sigma_1 \times \ldots \times \Sigma_n$ *is the space of feasible strategies (*Σ_i *is the set of feasible strategies for player* i*);*
- $s^m \in \Sigma$ *is the suggested strategy.*

Our goal is to design a mechanism in which all players have an incentive to play the suggested strategy s^m. With routing, the suggested strategy is the best response dynamics as defined above. Intuitively, we'd like all players to follow a transparent and honest policy in choosing and announcing routes.

In order to keep each player's valuation function private, we define for each i a private type θ_i. A strategy $s_i \in \Sigma_i$ maps the private type to behavior. Given a vector of private types $\theta = (\theta_1, \ldots, \theta_n)$, and $s \in \Sigma$, the outcome of the game is $g(s(\theta))$. We write the utility of i as $u_i(g(s(\theta)), \theta_i)$. In order to implement the private valuation function, we assume that utility depends not only on the outcome of the game but also on θ_i.

We are now ready to introduce a new solution concept, applicable to the network community and to our problem in particular. On a higher level, we want players to run our protocol (best response dynamics), regardless of their private types. Specific behavior will depend on θ_i, but in a way specified by the protocol.

Definition 2 ([SP04]) *A strategy vector* s^* *is an ex post Nash equilibrium if*

$$\forall i, \forall \theta_i, \forall \theta_{-i}, \forall s_i' \neq s_i^* : u_i(g(s^*(\theta)) \geq u_i(g(s_i'(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)$$

Intuitively, we require that no matter what the private types are, no player can improve by changing its strategy. This solution concept lives between dominant strategy and Nash equilibrium. In a Nash equilibrium each player is assumed to have full knowledge of the fixed strategies of the other players. In a dominant strategy, on the other hand, a player is oblivious of the others' strategies. In ex post Nash, a player has partial information about the other players: the private type which determines the valuation function is not known but the mapping from private type to behavior is public.

A related concept in the protocol world is the notion of incentive compatible strategies.

Definition 3 *We say* $s \in \bigcap_i \Sigma_i$ *is incentive-compatible in ex post Nash if* s^* *(with* $\forall i : s_i^* = s$*) is an ex post Nash equilibrium.*

Incentive compatible strategies capture the notion of a single universal strategy that induces an equilibrium when all players follow it.

There also are stronger notions of equilibrium than ex post Nash. For example, we usually want to guarantee resistance against collusion: no group of players should be able to increase the utility of a single member without decreasing the utility of another. Another interesting question is whether in a system in which honesty is not enforced, ν_i should depend on the path i gets, or on the path it thinks it gets. This question is irrelevant if we only allow one node to deviate from the best response dynamics, but becomes interesting once we allow the possibility of multiple players deviating at the same time.

3.3 Counter example for BRD

Unlike one-round games, in the multiround case with ex post Nash, dispute wheel freeness does not guarantee an equilibrium. Consider the following example.

Example

The network in Figure 1 satisfies the Gao-Rexford conditions (here the directions on the edges indicate customer-provider relationships). Hence, we can guarantee a unique stable paths solution: m forwards to 1, 1 forwards to 2, and 2 forwards to d (the destination node). However, m (the "manipulator" node) can deviate from the best response dynamics to get a higher utility than the with the stable paths solution above:

1. m announces md (which does not exist) to 2;
2. 2 chooses to forward to m;
3. 2 announces $2md$ to 1;
4. 1 announces $1d$ to m.

By this example, the best response dynamics are not incentive compatible in ex post Nash.

3.4 Positive results

In the general case, best response dynamics are not guaranteed to induce an equilibrium. For this reason, we need to explore what restrictions we need

Figure 1: A network consistent with the Gao-Rexford conditions that is not in equilibrium

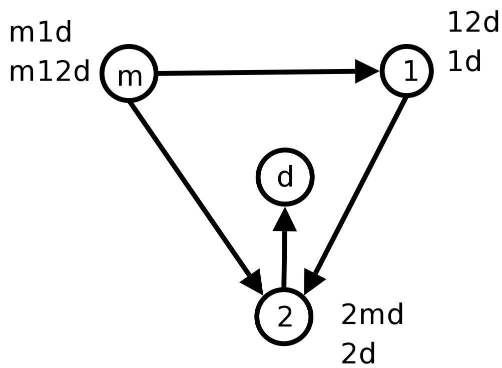

to put on the valuation functions of players in order to make BRD incentive compatible in ex post Nash.

Definition 4 *Route verification holds if a node can verify that any route advertised to it is the extension of a route advertised to the neighbor from which it was learned and ignore the route if it's not.*

A protocol which satisfies route verification is S-BGP [KLS00].

Theorem 2 ([LSZ08]) *Route verification and dispute wheel freeness imply that best response dynamics are incentive compatible in ex post Nash. Also the protocol requires no monetary transfers and is collusion-proof.*

By "no monetary transfers" we mean that we do not need to pay players in order to follow the prescribed strategy. By "collusion-proof" we mean that no group of players can increase one player's utility without decreasing another's.

Definition 5 *A network instance is policy consistent if*

$$\forall u, \forall P, Q \in \mathcal{P}^u : \nu_u(P) \leq \nu_u(Q) \textit{ and } wP, wQ \in \mathcal{P}^w \Rightarrow \nu_w(wP) \leq \nu_w(wQ).$$

Theorem 3 ([FRS06]) *Dispute wheel freeness and policy consistency imply that best response dynamics converge to a social welfare-maximizing path assignment and are incentive-compatible in ex post Nash with no monetary transfers.*

Social welfare in the theorem above is defined as $\sum_j \nu_j(P_j)$.

4 Decoupling signaling and forwarding

We saw that path verification gives us some pretty strong guarantees for the stability of BGP. Nevertheless, we want to know if a player can do something else to deviate from the protocol in order to benefit, even though path verification is used. One approach is to announce a known path which is inferior to the chosen one. Then players can try to guess what paths the neighbors like, but use the ones they themselves like.

4.1 Set equilibrium

We will introduce a different, more relaxed notion of equilibrium, due to Lavi and Nisan [LN05]. We call this solution concept set equilibrium.

Definition 6 *A strategy is BGP-compliant if it satisfies conditions 1 and 2 from best response dynamics, and specifies some sort of filtering mechanism (as long as all path announcements match the selected paths).*

A strategy vector s is in equilibrium if $s \in \Sigma' = \Sigma'_1 \times \ldots \times \Sigma'_n$, i.e. if players move within their sets of BGP-compliant strategies but not outside.

4.2 Model and some positive results

We'll redefine the valuation function so that it accounts for signaling (who to announce paths to) as well as forwarding preferences:

$$v_i = \sigma_i + \phi_i$$

Here ϕ_i is the forwarding preference and is defined the same way we have defined the valuation function so far. The new component, σ_i, depends on the routing tree rooted at node i (assuming the protocol converges) and reflects signaling preferences. This model assumes that players are not only interested in the routing path they are using, but also in attracting traffic from certain other players.

We will first show a static example from [JRW08]:

- $\forall v : \phi_v : \mathcal{P}^v \to \mathbb{Z}$;
- $\forall v, \forall w \in N(v) : \sigma_{v,w} : \mathcal{P}^v \to \mathbb{Z}$ (gives a preference value for each path announced from v to w).

We introduce the following notation:

1. The signaling assignment from v to w is $\sigma(v, w)$ $(v \in V, w \in N(v))$.
2. $\mathcal{K}_\sigma(v) = \{v\sigma(w, v) : w \in N(v), v\sigma(w, v) \in \mathcal{P}^v\}$.

We want $\sigma(v, w)$ to minimize $\sigma_{v,w}$ subject to the paths in $\mathcal{K}_\sigma(v)$.

We can consider the forwarding digraph induced by σ. Specifically, this is the graph we get by looking at $\mathcal{K}_\sigma(v)$ and for each node choosing the path that maximizes ϕ_v. Unfortunately, even with Gao-Rexford conditions we can have a stable signaling solution, but a forwarding loop.

Next, we are interested in what happens to the routing game when the utility functions include signaling. It is easy to construct an example in which if a manipulator m has a benefit from more nodes going through it, decoupled from path benefit, conditions are violated.

Theorem 4 ([GHJ+08]) *Assume the following conditions are satisfied:*

- *There is no dispute wheel, and everyone but* v *uses BGP-compliant strategies;*
- v *wants to maximize the number of nodes routing through it in the final solution;*
- *all nodes but* v *use next-hop preferences and all-or-nothing export;*
- *either:*
  - *we have policy consistency and consistent export,*
  - *or path verification is implemented.*

Then v *has a BGP-compliant strategy with all-or-nothing export. Specifically, export all is such a strategy.*

A few terms need clarification:

- **next-hop preferences**: preferences only depend on the neighbor itself;
- **all-or-nothing export**: export to nodes does not depend on the chosen path.

5 Assignments

The following assignment is due Monday, March 23rd.

Each student should prepare:

- 2-3 page outline of the project write-up/presentation;
- slides - at best most slides for the presentation should be completed; an outline for the remaining slides is acceptable.

The outline should be helpful to find out how much material there is to present. Think about how many and which proofs you should include. Concentrate on the essence of the problems, and on the main interesting points which can lead to future work.

References

[BOGMR85] Michael Ben-Or, Oded Goldreich, Silvio Micali, and Ronald L. Rivest. A fair protocol for signing contracts (extended abstract). In *Proceedings of the 12th Colloquium on Automata, Languages and Programming*, pages 43–52, London, UK, 1985. Springer-Verlag.

[Cha81] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–90, 1981.

[FRS06] Joan Feigenbaum, Vijay Ramachandran, and Michael Schapira. Incentive-compatible interdomain routing. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 130–139, New York, NY, USA, 2006. ACM.

[GHJ+08] Sharon Goldberg, Shai Halevi, Aaron D. Jaggard, Vijay Ramachandran, and Rebecca N. Wright. Rationality and traffic attraction: incentives for honest path announcements in bgp. *SIGCOMM Comput. Commun. Rev.*, 38(4):267–278, 2008.

[JRW08] Aaron D. Jaggard, Vijay Ramachandran, and Rebecca N. Wright. Towards a realistic model of incentives in interdomain routing: Decoupling forwarding from signaling. Technical report, DIMACS, 2008.

[KLS00] Stephen Kent, Charles Lynn, and Karen Seo. Secure border gateway protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, 2000.

[LN05] Ron Lavi and Noam Nisan. Online ascending auctions for gradually expiring items. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1146–1155, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[LSZ08] Hagay Levin, Michael Schapira, and Aviv Zohar. Interdomain routing and games. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 57–66, New York, NY, USA, 2008. ACM.

[SP04] Jeffrey Shneidman and David C. Parkes. Specification faithfulness in networks with rational nodes. In *PODC '04: Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 88–97, New York, NY, USA, 2004. ACM.