TOPICS IN GAME THEORY
SPRING 2009

Lecture 4: Interdomain Routing and the Stable Paths Problem
Lecturer: Aaron Jaggard
Scribe: Paul Ringstad
February 9, 2009

Overview

In the current iteration of the Internet, the Border Gateway Protocol (BGP) is the sole protocol used for routing between domains. BGP allows each autonomous system to define its own preference for routing policy rather than strictly adhering to a distance-based policy. Analysis of the robustness of policy choices in BGP has implications towards the overall efficiency and functionality of the Internet. Thus, it makes sense to embark in a study of the interdomain routing with regard to the BGP protocol.

BGP Model

The first step in analyzing BGP is defining a suitable model. For this purpose, we have an undirected graph, $G(V, E)$ in which we want to find some route to destination d. Nodes can be considered autonomous systems or subnetworks (i.e. ISPs). d can be thought of as a destination address. The originator of a packet sent on the network announces a destination and each node in the network has preferences for routing policies through which to send the packet en route to d. The main question we seek to answer via this model is, when can we expect good behavior in the system to arise? And, in regards to game theory, how can we define a mechanism that maintains good behavior?

The process by which routing policies are set in BGP can be thought of as the following:

(1) d advertises itself
(2) $\forall v \in V, v \neq d$
- iteratively receive updates about paths to d
- receive status updates
- choose best path, update the forwarding table
- announce best path to neighbors

There are an infinite number of possible policies for a given graph.

Simple Example. An example preference scheme might be that i prefers to route $i(i+1)d$ to get to the destination. See Figure 1 for a graphic of this preference scheme.

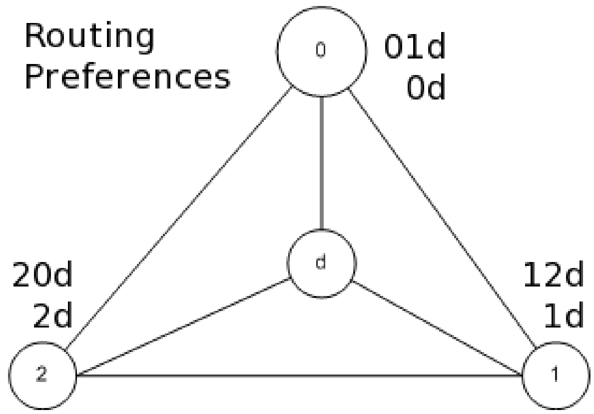

FIGURE 1. d is the destination

FORMALIZATION OF THE STABLE PATHS PROBLEM (SPP)

Undirected graph G=(V,E)

- For each $v \in V$, P^v is set of permitted paths from v to destination vertex d
- For each $v \in V$, there is a ranking function defined over P^v. If $\lambda^v(P_1) < \lambda^v(P_2)$ then P_2 is a more preferred permitted path than P_1.
- Empty path, $\epsilon \in P^v$ is permitted, and is ranked lowest,$\lambda^v(\epsilon) = 0$, $\lambda(P) > 0$.
- $P_1, P_2 \in P^v \implies \lambda^v(P_1) \neq \lambda^v(P_2)$

We want to assign permitted paths to vertices, such that $u \in V$, $\pi(u) = P^u$, where π is the function mapping vertices to paths. Ideally the preferred routes for the nodes will be such that they are stable,

$$\pi(v) = v\pi(u)$$

and

$$\lambda^v(\pi(v)) = max_u \lambda^v(v\pi(u))$$

where $u \in N(v)$, u is a neighbor of v.

The general question we seek to answer for the graph is 'Is there a stable path assignment?'. A stable path solution is checkable in polynomial time, but the finding a solution is NP-Hard. However, there are approximation strategies for Nash equilibrium which are applicable. It should be noted that there are differences exist between this problem statement and that of minimum spanning tree (see Shepherd, Griffin, Wilfong for more on that).

REDUCING SPP TO 3-SAT

In order to show that finding a solution to the SPP is NP-complete, we first note that checking a solution is possible in polynomial time. The remainder of the proof relies on the reduction of the problem to 3-SAT, a known NP-complete problem.

By reducing the problem to that of a 3-SAT, path assignment is equated with truth assignment. A truth assignment indicates that there is a stable assignment of policies.

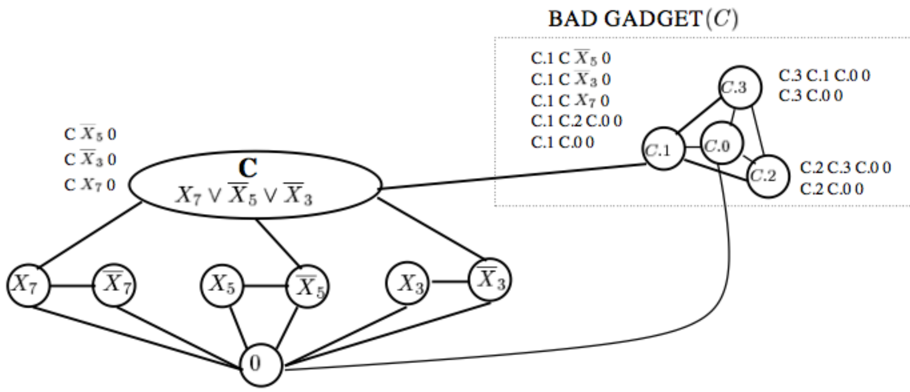

FIGURE 2. 3-SAT

The intuition behind the form of the reduction is as follows. Clauses, as indicated by the C, are composed of a disjunction of literals. In Figure 2, the clause shown is the disjunction $\bar{X}_7 \vee \bar{X}_5 \vee \bar{X}_3$. The literals in the clause determine the whether there is a link between the "variable assignment gadget" at that literal and the clause. The variable assignment gadget is the pair of X_i and $\bar{X}_i$ that are connected with the destination. Figure 3 shows the variable assignment gadget with X_i set to true in 3b and X_i set to false in 3c. A BAD GADGET is such that it has no stable path assignment unless at least one of the literals in the connected clause is set to true. Therefore, if there is a stable assignment it must be the case that $\pi(C.1) = (C.1 l_j 0)$ for some literal in the clause.

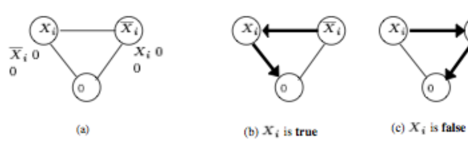

FIGURE 3. Variable assignment gadget

Example of Stable Assignment. See Figure 4 and Figure 5 for examples of satisfied 3-SAT formulations.

DISPUTE WHEELS

In determining a heuristic approach to constructing stable path assignments in a greedy manner, the presence of dispute wheels in the network indicates that there does not exist a unique solution.

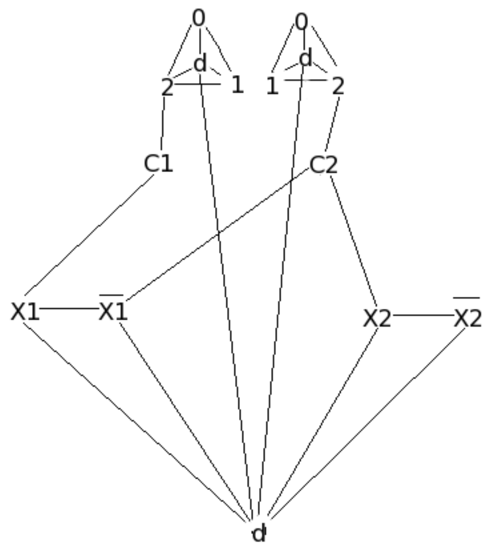

FIGURE 4. 3-SAT with clause $C1 = X_1$ and $C2 = \hat{X}_1 \vee X_2$.

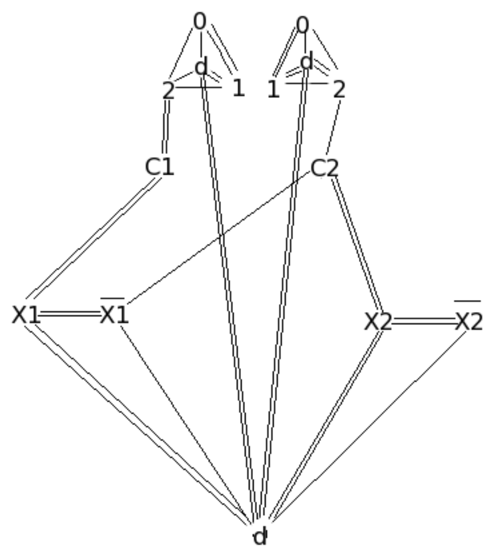

FIGURE 5. Assignment that is stable for graph in Figure 4. Doubled edges indicate policy choice.

0.1. **Definition.** A dispute wheel in an instance of SPP is a collection on nodes, $0, .., n-1$, of "spoke paths" $Q_{i=0}^{n-1}$, and of "rim paths", $R_{i}{}_{i=0}^{n-1}$. Figure 6 shows a visual a dispute wheel, where $i \in |V|$, and i modulo n.

- R_i is a path from from i to $i+1$
- $Q_i \in P^i$, Q_i is a permitted path for i
- $R_i Q_{i+1} \in P^i$
- $\lambda^i(Q_i) \leq \lambda^i(R_i Q_{i+1})$, the path $R_i Q_{i+1}$ is preferred to Q_i

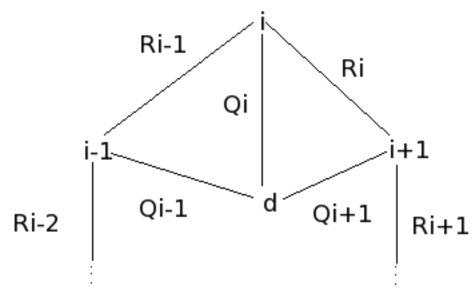

FIGURE 6. Dispute Wheel

Theorem 1. *If an instance of SPP has no dispute wheel then it has a unique stable path assignment.*

Proof. Assume that G has no dispute wheel and has two distinct solutions, $\pi_1 = (P_1, ...P_{n-1})$ and $\pi_2 = (Q_1, ..., Q_{n-1}$. We can think of there being two trees rooted at the destination, T_1 and T_2. Let H be the intersecting edges of the two trees, such that $H = (V, E(T_1) \bigcap E(T_2))$. Let T be the component of H containing the destination. Therefore, any edge entering $V(T)$ is either in T_1 or T_2 but not both. Moreover, $V - V(T)$ is non-empty because T_1 and T_2 are distinct trees, and at least one of the trees have an edge entering $V(T)$. Let u, v be any pair of vertices in T_1 such that v is in T and u is not. Then u must be in T_2, and π_2 must be non-empty for u since the empty path can not be preferred to $(u, v)Q_v$. In other words, since u is not in T, then u must have some other preferred route to the destination, and therefore must be in T_2. We can imagine that there exist other such vertices u_0, v_0, for which $u_0 \notin V(T)$ and $v_0 \in V(T)$, for which u_0 has a path to the destination in T_2. This path must be of the form $R_0(u_1, v_1)Q_1$ where

- $u_1 \notin V(T)$, $v_1 \in V(T)$ and Q_1 is the unique path in T from v_1 to the destination
- R_0 is a path from u_0 to u_1 in T_2, existing in the set $V - V(T)$
- R_0 has at least one edge

To show that this implies a dispute wheel we need to show that for each i,

$$\lambda^{u_i}((u_i, v_i)Q_i) \leq \lambda^{u_i}((u_i, v_i)Q_i)$$

That is, each i prefers routing along the "rim paths" rather than directly to the destination. If this inequality did not hold, then

$$\lambda^{u_i}((u_i, v_i)Q_i > \lambda^{u_i}((u_i, v_i)Q_i)$$

which implies that T_2 is unstable. □

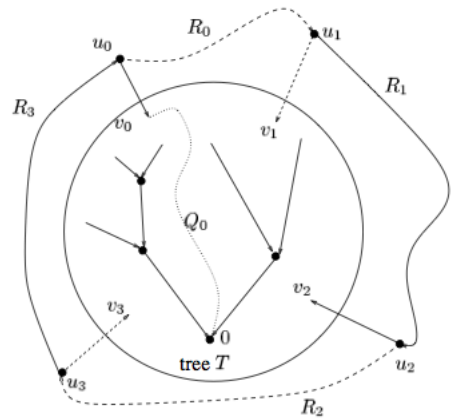

FIGURE 7. Graphic for proof of unique solution if no dispute wheel exists

A Toy Example. See Figures 8 and 9 for an example of an embedded dispute wheel.

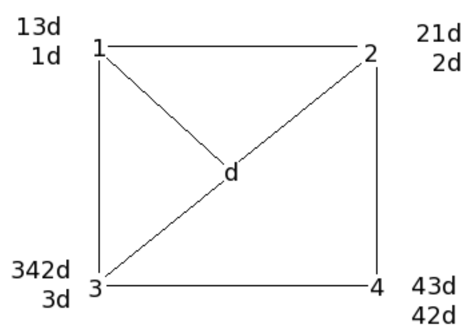

FIGURE 8. Example of an embedded dispute wheel

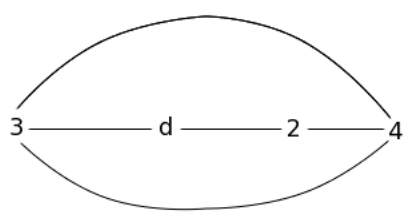

FIGURE 9. Dispute wheel embedded in Figure 8

Theorem 2. *An SPP instance is dispute wheel free iff it is equivalent to an SPP instance in which $\lambda^v(vP) > \lambda^u(P)$. The equivalence is such that $I \equiv \hat{I}$, $G \equiv \hat{G}$ if*

- $\forall v \in P^v = \hat{P}^v$
- $\lambda^v(P) = \hat{\lambda}^v(P)$
- $\lambda^v(P) > \lambda^v(Q)$ *iff* $\hat{\lambda}^v(P) > \hat{\lambda}^v(Q)$

Theorem 3. *An SPP instance is solvable if it contains no dispute wheel*

(see Shepherd, Griffin, and Wilfong paper for proof of both the above theorems)

Concluding Remarks

Alluding to the problem statement and the purpose of this research, we want to be able to say something about the robustness and stability of the BGP protocol in the context of the Internet. Here we show that if we assume some reasonable conditions on the connectivity within the graph we can guarantee that no dispute wheels exist and therefore a stable assignment exists. These conditions are exactly the Gao-Rexford conditions as seen below.

- customer → provider (cost is positive)
- peer → peer (cost is 0)
- There are no customer → provider cycles. This is intuitive in that providers will not purchase bandwidth from customer.

Additionally, there are forbidden configurations, which can be thought of as a filter on the notifications between nodes.

- As a peer do not announce paths from peer to peer or peer to provider, that is, if you're a peer do not handle other peers' traffic
- Do not announce paths from provider to provider
- Do not announce paths from providers to peers

Finally, the conditions force a preference to routes learned from customers, which can be thought of as a filter on what actions can be taken given announcements.

If all of these conditions hold, there is no dispute wheel (as shown by collaboration with Shepherd, Griffin and Wilfong). This formulation can be extended to an arbitrary number of classes with the same non-dispute wheel quality remaining.

References

- Stable internet routing without global coordination. (Gao,Rexford)
- The Stable Paths Problem and Interdomain Routing. (Shepherd, Griffin, Wilfong)

Projects. If you haven't chosen a project you need to soon and email your choices to the instructors. Presentations will start in 3 weeks on March 30th. Presentations will be 1 hour per presenter, so, a group of two people should present for two hours.

Additional papers on interdomain routing.

- Differential privacy via mechanism design (McSherry, Talwar)
- On the value of value of private information (Papadimitriou, Kleinberg, Vazrani)
- On hardness of oscillation on BGP (Fabrikant, Papadimitriou)