A Fractional Model of the Border Gateway Protocol (BGP)

P.E. Haxell*
University of Waterloo
Waterloo, ON, N2L 3G1
Email: pehaxell@math.uwaterloo.ca

G.T. Wilfong
Bell Laboratories
Murray Hill, NJ
Email: gtw@research.bell-labs.com

Abstract

The Border Gateway Protocol (BGP) is the interdomain routing protocol used to exchange routing information between Autonomous Systems (ASes) in the internet today. While intradomain routing protocols such as RIP are basically distributed algorithms for solving shortest path problems, the graph theoretic problem that BGP is trying to solve is called the stable paths problem (SPP). Unfortunately, unlike shortest path problems, it has been shown that instances of SPP can fail to have a solution and so BGP can fail to converge.

We define a fractional version of SPP and show that all such instances of fractional SPP have solutions. We also show that while these solutions exist they are not necessarily half-integral.

1 Introduction

The internet consists of tens of thousands of subdomains known as Autonomous Systems (ASes) where each AS is a network of routers controlled by some administrative agent. The managers of an AS have the conflicting desires to have their AS connected to the rest of the internet (i.e., to have routes to destinations (IP addresses) in other ASes and to have other ASes know how to route traffic to the destinations that it owns) but not to allow too much traffic of other ASes to transit over their network. In order to control these issues, neighboring ASes establish contracts called *service level agreements* (SLAs) between each other. These contracts can roughly be thought of as promises to transit certain traffic for each other. Thus network operators need some way to encode in their routers, the routing policies that would let them meet the requirements of their SLAs. To do this, they use a protocol known as the Border Gateway Protocol (BGP) [RL95].

BGP can be thought of as working in the following manner. Consider some destination d where d is an IP address or more generally, a block of IP addresses. The router where d originates will announce via BGP to some of its neighbors that it "owns" d. Which neighbors it tells is a function of the economic deals it has with its neighbors. In turn, these neighbors might tell some of their neighbors that a route to d is available through them. In general, a router R might hear via BGP from a subset of its neighbors, about a number of different routes to a particular destination d. BGP then uses the encoded economic policies of the AS in which it belongs, to choose its most "preferred" route from amongst those it currently knows about. Then it selectively tells some neighbors about this most preferred route via BGP.

Unfortunately, the individual economic goals of an AS might result in route preferences that conflict with the preferences of other ASes in such a way that BGP never converges [VGE00]. In order to study this phenomenon, BGP has been modeled as a

*Research partially supported by NSERC

formal graph theoretic problem called the *Stable Paths Problem (SPP)* [GSW02]. We will describe SPP in Section 2.

It has been shown that some instances of SPP fail to have a solution. This implies that BGP can fail to converge (since in these cases there is no stable solution to which it can converge).

Motivated by some similarities between SPP and the stable matching problem [GS62] and the fact that while stable matchings do not always exist, fractional stable matchings do always exist [Tan91], we define a fractional version of SPP we call *fractional SPP*. Intuitively, SPP and fractional SPP can be thought of as follows. In both, each node has a preference ordering of (some of the) paths from itself to the specified destination. In SPP, each node v "chooses" a path P_v, and this allows other nodes to choose paths that contain P_v as a subpath. One can think of each v as putting a weight of 1 on P_v and 0 on all other paths from v to the destination. A solution to an SPP instance then is a set of choices in which each node selects a path that obeys the above subpath constraint and no node can change its mind (while the others remain fixed) and get a more preferred path. On the other hand, in fractional SPP a node can fractionally assign weights to paths from itself to the destination so that it assigns a total weight of no more than 1. When a node u puts a weight w on a path P, this constrains the weights the other nodes can put on paths containing P so that for each other node these weights total no more w. The intuition is that normally we view a node's BGP announcement of a route P to d as saying that if another node hears of P then it is permitted to send all of its traffic to d so that it passes over P. In our fractional model of BGP, we then think of allowing a node to announce routes each with a fractional weight w, $0 \leq w \leq 1$ where the total of the weights on routes offered by a node is at most 1. The weight is interpreted as meaning that the node is offering to allow any node hearing of this route to send (at most) a w fraction of its traffic to d along this route. The notion of stability in our fractional model is similar to SPP. Intuitively, a solution is one in which no node can shift some weight to more preferred paths given that the other nodes keep their weights fixed. We will show that, unlike (integral) SPP, every instance of fractional SPP has a solution.

One can view SPP as a pure strategy game whose solutions correspond to Nash equilibria. However it should be noted that it is not the case that fractional SPP is just a mixed strategy game and so we cannot conclude that stable solutions (i.e., Nash equilibria) exist simply by appealing to Nash's Theorem [Nas50]. It may also be possible to formulate SPP in terms of a cooperative game in order to study its core, but it is not necessary to do so for our purposes and we therefore choose an elementary treatment.

2 Formal Definition of SPP

The dynamic operation of BGP as outlined in Section 1, can be modeled as an equivalent static graph problem called the stable paths problem (SPP) defined below. The problems are equivalent in the sense that for any network and for any configuration of BGP (i.e., encoding of policies) on the routers in the network, BGP has a stable solution that it might converge to if and only if the corresponding instance of SPP has a solution.

We now describe SPP as it was defined in Griffin et al. [GSW02]. Let G be a graph with a distinguished node d, called the *destination*. Suppose each node $v \neq d$ in G has a *preference list* of simple paths from v to d. This is an ordered list of some (not necessarily all) such paths. We use the notation $\pi(v)$ to denote the set of paths in the preference list of v, and we write $P < P'$ to mean that P and P' are both in some $\pi(v)$ and v prefers P' to P. We describe the preferences numerically as follows: for $P \in \pi(v)$, we define $U(P)$, the utility of path P, to be c if there are $c - 1$ paths $P' \in \pi(v)$ where $P' < P$. For a path S we also define

$\pi(v, S)$ to be the set of paths in $\pi(v)$ that end with the path S. (Note that if S is just the "empty" path d then $\pi(v, S) = \pi(v)$.) A *solution* to an SPP instance is a (not necessarily spanning) arborescence T in G with sink node d, with the following stability property. Let Q be any path, and let v be its starting node. Then one of the following holds:

- $P \subseteq T$ for some $P \in \pi(v)$ with $P \geq Q$,
- there exists a proper final segment S of Q such that $S \not\subseteq T$.

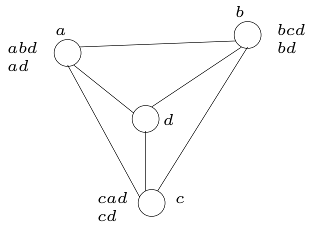

Figure 1: SPP instance with no solution.

It is known that not every instance of SPP has a solution. For example, the instance called BAD GADGET is an instance of SPP that has no solution [GW99] and is described as follows (see Figure 1): G is a copy of K_4 with vertices a, b, c, and d. Each of a, b, and c has two paths in its preference list: a has $P_1 = ad$ and $P_2 = abd$, and a prefers P_2 to P_1. The preference lists for b and c are analogous: each prefers to go through its clockwise neighbor than to go straight to d. The preference lists are shown in the figure ordered from most preferred to least preferred in order from top to bottom. It is easily seen that no solution exists for such an instance of SPP.

3 Fractional SPP

In this section we define a fractional generalization of SPP that we call *fractional SPP*. The parameters defining an instance of fractional SPP are the same as those for an instance of SPP. That is, we have a graph with a designated destination node d where each non-destination node has a preference list, i.e., an ordered list of some of the paths from itself to the destination node d. The only difference will be the definition of a solution which we now describe.

For fractional SPP we define a solution to be an assignment of a non-negative weight $w(P)$ to each path P in $\pi(v)$ for every v so that the weights satisfy the three properties listed below.

"Unity" condition: For each node v, $\sum_{P \in \pi(v)} w(P) \leq 1$.

"Tree" condition: For each node v, and each path S, we have $\sum_{P \in \pi(v,S)} w(P) \leq w(S)$.

"Stability" condition: Let Q be a path, and let v be its starting node. Then one of the following holds:

- $\sum_{P \in \pi(v)} w(P) = 1$, and each $P \in \pi(v)$ with $w(P) > 0$ is such that $P \geq Q$.
- there exists a proper final segment S of Q such that $\sum_{P \in \pi(v,S)} w(P) = w(S)$, and moreover each $P \in \pi(v, S)$ with $w(P) > 0$ is such that $P \geq Q$.

Our aim is to show that every instance of fractional SPP has a solution. Because of a technicality, we will do this in two stages. First we show that for any positive constant ϵ, every instance of fractional SPP has an ϵ-solution, which is defined as follows:

For each node v and each path $P \in \pi(v)$, we assign a non-negative weight $w(P)$ such that the following conditions hold.

"Unity" condition: $\sum_{P \in \pi(v)} w(P) \leq 1$ for each v.

"ϵ-Tree" condition: For each node v, and each path S, we have $\sum_{P \in \pi(v,S)} w(P) \leq w(S) + \epsilon$.

"ϵ-Stability" condition: Let Q be a path, and let v be its starting node. Then one of the following holds:

- $\sum_{P \in \pi(v)} w(P) = 1$, and each $P \in \pi(v)$ with $w(P) > 0$ is such that $P \geq Q$.
- there exists a proper final segment S of Q such that $\sum_{P \in \pi(v,S)} w(P) = w(S) + \epsilon$, and moreover

each $P \in \pi(v, S)$ with $w(P) > 0$ is such that $P \geq Q$.

In the next section we show that for every $\epsilon > 0$, every instance of fractional SPP has an ϵ-solution. Then in the following section we apply a standard compactness-type argument to conclude that every instance has an exact solution.

4 Approximate Solvability of Fractional SPP

The main tool in our proof is an important result due to Scarf [Sca67]. The idea of applying Scarf's Lemma to a stability type problem was used by Aharoni and Fleiner in [AF03].

THEOREM 4.1. *(Scarf's Lemma) Let $n < m$ be positive integers, let $b \in \mathbf{R}^n_+$, and let B and C be $n \times m$ matrices with the following properties:*

- *the first n columns of B form an identity matrix, and the set $\{x \in \mathbf{R}^m_+ : Bx = b\}$ is bounded,*
- *each entry c_{ik} for $k > n$ satisfies $c_{ii} < c_{ik} < c_{ij}$ for each $j \neq i$, $j \leq n$.*

Then there exists $x \in \mathbf{R}^m_+$ such that $Bx = b$ and the set of columns S of C that correspond to the support supp$(x) = \{k : x_k \neq 0\}$ of x form a dominating set. This means that for every column j, there exists a row i such that $c_{ik} \geq c_{ij}$ for every $k \in$supp(x).

We will apply Scarf's Lemma to matrices B and C defined from an instance of fractional SPP as follows.

DEFINITION 4.1. *Let a graph G with preference lists be a given instance of fractional SPP. Let $N = \{(v, P) : P \subset Q$ for some $Q \in \pi(v)\}$ (note the proper containment), and set $n = |N|$. We let $m = n + t$ where t is the number of paths in U. The $n \times t$ matrix B' is indexed by $N \cup U$, with entries as follows: the $((v, P), Q)$-entry is -1 if $P = Q$, is 1 if $Q \in \pi(v, P)$, and is 0 otherwise.*

The matrix B is formed by attaching an $n \times n$ identity matrix to B' on the left.

Let M be a number larger than the size of any preference list. The matrix C' is defined as follows: if $Q \in \pi(v, P)$ then the $((v, P), Q)$-entry is the utility c of $Q \in \pi(v, P)$, and if $Q \notin \pi(v, P)$ then the entry is M.

The matrix C is formed by attaching a $n \times n$ matrix to C' on the left, in which each diagonal entry is smaller than each entry of C', and each off-diagonal entry is larger than each entry of C'.

Let us emphasise that the empty path is used in the index set N, even though it is not an element of the union $U = \cup_v \pi(v)$. Note then in particular that when P is the empty path, there is no -1 entry in the row (v, P). Observe that the matrices B and C as defined here are of the form required in the assumptions of Scarf's Lemma.

LEMMA 4.1. *Let G with preference lists be an instance of fractional SPP, and let B, C, N, n and m be as defined in (4.1). Let $\epsilon \geq 0$ be given, and let $b(\epsilon) \in \mathbf{R}^n$ be the vector with coordinates indexed by N defined as follows: for the empty path P, each (v, P)-coordinate of $b(\epsilon)$ is 1, and for all other paths the (v, P)-coordinate is ϵ. Then each coordinate of each element x of the set $\{x \in \mathbf{R}^m_+ : Bx = b(\epsilon)\}$ lies in the interval $[0, 1 + \epsilon]$.*

Proof. Suppose the vector $(g(v_1, P_1), \ldots, g(v_s, P_t), w(P_1) \ldots, w(P_t))$ is a non-negative solution to $Bx = b(\epsilon)$. By definition of B, we have first of all (looking at rows (v, P) for the empty path $P = d$) that for each v, $g(v, d) + \sum_{Q \in \pi(v)} w(Q) = 1$. This tells us that $g(v, d) \leq 1$ and $w(Q) \leq 1$ for each v and $Q \in U$. Now for each P of length at least 1 we have $g(v, P) - w(P) + \sum_{Q \in \pi(v,P)} w(Q) = \epsilon$, telling us that each $g(v, P) \leq w(P) + \epsilon \leq 1 + \epsilon$. Thus each coordinate of the solution lies in the interval $[0, 1+\epsilon]$. ■

Our last preliminary lemma is a technical result that essentially tells us that the solution provided by Scarf's Lemma gives a stable solution to SPP.

LEMMA 4.2. *Let S be a dominating set of columns in C. Suppose S is also the support of some non-negative solution $x^*(\epsilon)$ to $Bx = b(\epsilon)$ for some $\epsilon > 0$.*

Let $x^*(\alpha) = (g_\alpha(v_1, P_1), \ldots, g_\alpha(v_s, P_t), w_\alpha(P_1) \ldots, w_\alpha(P_t))$ be a non-negative solution to $Bx = b(\alpha)$ for some $\alpha \geq 0$, whose support is contained in S. Then the weight function w_α satisfies the α-stability condition.

Proof. Let Q be a path. Let $x^*(\epsilon) = (g_\epsilon(v_1, P_1), \ldots, g_\epsilon(v_s, P_t), w_\epsilon(P_1) \ldots, w_\epsilon(P_t))$. Suppose that the column in C' indexed by Q is dominated in C in the row (v, P). Then $g_\epsilon(v, P) = 0$, as the $((v, P), (v, P))$-entry of C is smaller than all entries in C', and hence also $g_\alpha(v, P) = 0$. Therefore if $P = d$ we get $\sum_{P' \in \pi(v,P)} w_\epsilon(P') = 1$ and $\sum_{P' \in \pi(v,P)} w_\alpha(P') = 1$, and if $P \neq d$ then $\sum_{P' \in \pi(v,P)} w_\epsilon(P') = w_\epsilon(P) + \epsilon$ and $\sum_{P' \in \pi(v,P)} w_\alpha(P') = w_\alpha(P) + \alpha$. Now we claim that $Q \in \pi(v, P)$. Suppose not: then the $((v, P), Q)$ entry of C is M, and so by definition of C', none of the paths P' with $w_\epsilon(P') \neq 0$ are in $\pi(v, P)$. But in this case $\sum_{P' \in \pi(v,P)} w_\epsilon(P') = 0$, contradicting the fact that this value is 1 or $w_\epsilon(P) + \epsilon > 0$. (We remark that this is where $\epsilon > 0$ was needed.) Therefore $Q \in \pi(v, P)$. Finally, note that by definition of C', all $P' \in \pi(v, P)$ for which $w_\epsilon(P') \neq 0$ are preferred by v to Q (or are equal to Q). Hence also all $P' \in \pi(v, P)$ for which $w_\alpha(P') \neq 0$ are preferred by v to Q or are equal to Q. Thus w_α satisfies the α-stability condition. ■

Now we are ready to use Scarf's Lemma to show that every instance of fractional SPP has a ϵ-solution for any $\epsilon > 0$. As mentioned previously, the matrix C will capture the notion of stability, while the matrix B will guarantee the unity and tree conditions.

THEOREM 4.2. *Let $\epsilon > 0$, and let a graph G with preference lists be an instance I of fractional SPP. Let matrices B and C be as defined in (4.1). Then there exists a non-negative solution x^* to $Bx = b(\epsilon)$, whose support S is dominating in C. This gives an ϵ-solution of I.*

Proof. By Lemma 4.1, the set $\{x \in \mathbf{R}^m_+ : Bx = b(\epsilon)\}$ is bounded. We may therefore apply Scarf's Lemma to obtain a solution $x^* = (g(v_1, P_1), \ldots, g(v_s, P_t), w(P_1) \ldots, w(P_t))$ whose support is dominating in C. We claim that the weight function w is an ϵ-solution to I.

The unity condition follows because for each v, we have $g(v, d) + \sum_{Q \in \pi(v)} w(Q) = 1$. The ϵ-tree condition holds because for each P of length at least 1 we have $g(v, P) - w(P) + \sum_{Q \in \pi(v,P)} w(Q) = \epsilon$. To verify the ϵ-stability condition we apply Lemma 4.2 to S with $\alpha = \epsilon$ and $x^*(\epsilon) = x^*(\alpha) = x^*$. Therefore w is an ϵ-solution to I as required. ■

5 Exact Solution

In this brief section we show that Theorem 4.2 in fact implies that every instance of SPP has a solution (i.e. a 0-solution). To find it, we will just consider an infinite sequence of ϵ-solutions where ϵ tends to 0, and show that some subsequence of these solutions converges to an exact solution.

THEOREM 5.1. *Every instance I of fractional SPP has a solution.*

Proof. Let a graph G, together with preference lists of paths for each node, be the given instance I. Let the matrices B and C be as in (4.1). For the sequence $1 > 2^{-1} > 2^{-2} > \ldots$ of positive constants converging to 0, consider the sequence of vectors $b(2^{-1}), b(2^{-2}), \ldots$ as defined in Lemma 4.1.

For each $i \geq 1$, by Theorem 4.2 there is a non-negative solution $x^*(2^{-i})$ to $Bx = b(2^{-i})$, whose support is dominating in C. Let S be a subset of columns of B that occurs as the support of $x^*(2^{-i})$ for infinitely many i, and let $\epsilon_1 > \epsilon_2 > \ldots$ be the infinite subsequence of $2^{-1} > 2^{-2} > \ldots$ for which S is the support of the solution. Since $\epsilon_i < 1$ for each i, by Lemma 4.1, there exists a subsequence $\alpha_1 > \alpha_2 > \ldots$ of $\epsilon_1 > \epsilon_2 > \ldots$ such that the solutions $x^*(\alpha_1), x^*(\alpha_2), \ldots$ converge to a vector x^*, in which every coordinate lies in $[0, 2]$. The support of x^* is contained in S, and by continuity x^* is a solution to $Bx = b(0)$.

We claim that the weight function w associated with x^* is a solution to I. The unity and tree conditions follow as before from the fact that x^* is a solution to $Bx = b(0)$. To verify the stability condition we apply Lemma 4.2 with 0 in place of α and ϵ_1 in place of ϵ. Therefore w is a solution to I as required. ■

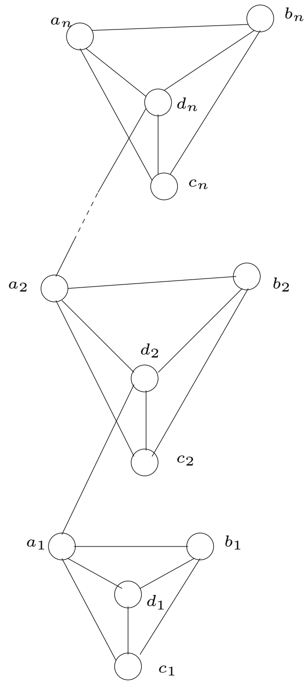

Figure 2: fractional BGP is not half-integral

6 Half-integral Solutions

Consider the fractional SPP instance I whose graph and preference lists are as shown in Figure 1. If each node assigns a weight of 1/2 to each of the two paths in its preference list, it is straightforward to verify that this is a fractional solution to I and in fact, it is the only fractional solution.

The above example might lead one to believe that perhaps there is a half-integral solution for all instances of fractional SPP just as there is for all instances of the fractional stable matching problem [Tan91]. However this is not the case as the following example illustrates.

Consider the fractional SPP instance shown in Figure 2. The preference lists for a_1, b_1 and c_1 are analogous to those for a, b and c respectively in Figure 1, i.e., they all prefer the path through their clockwise neighbor over the direct path to d_1. The node d_1 is the destination node and so we think of it as having a single path d_1 in its preference list. For $2 \leq i \leq n$ define the path P_i to be $a_{i-1}d_{i-1}a_{i-2}d_{i-2} \ldots a_1d_1$. Then the preference list of paths for d_i is d_iP_i. The preference list of paths for a_i is $a_ib_id_iP_i$ followed by $a_id_iP_i$. For b_i the preference list is $b_ic_id_iP_i$ followed by $b_id_iP_i$ and finally the preference list for c_i is $c_ia_id_iP_i$ followed by $c_id_iP_i$. Since any solution for this instance must have a_1, b_1 and c_1 assign a weight of 1/2 to each of their paths, we know in particular that the weight on $P_2 = a_1d_1$ is 1/2. Then the weight that d_2 can assign to its single path d_2P_2 is at most 1/2 since that path ends with P_2. Thus the total weight that any of a_2, b_2, c_2 can assign to their paths is at most 1/2 since they all end with d_2P_2. It is then easy to check that in fact in a solution, d_2 assigns a weight of 1/2 to its path and each of a_2, b_2, and c_2 assigns a weight of 1/4 to each of their paths. In general, a simple inductive argument shows that the only solution for this instance is when for any i, $1 \leq i \leq n$, node d_i assigns a weight of $1/2^{i-1}$ to its one path and each of a_i, b_i and c_i assigns a weight of $1/2^i$ to each of its two paths.

7 Discussion

A solution (if it exists) to an instance of SPP is a not necessarily spanning arborescence T that includes the destination d. The arborescence T determines a *routing* in the network of routers running BGP, i.e., it determines how packets to d will be routed. In particular, a packet destined for d is routed from its originating node v to d along the path from v to d in T.

In fractional SPP, a solution (which we have shown always exists) corresponds to a structure that is more general than a tree. The question arises as to whether this more complicated structure can be interpreted as some kind of "fractional routing". That is, can it be interpreted as a description of how packets are to be routed in some corresponding network of routers? One way to describe such a *fractional routing* is as follows. If a node v assigns a weight of $w(P)$ to a path P, then it will route a $w(P)$ fraction of the packets it originates destined for d so that they travel along P. This could easily be accomplished with a connection oriented style of routing such as MPLS.

Griffin et al. [GSW02] observed that much as other routing protocols such as RIP and OSPF can be viewed as distributed algorithms for solving the shortest paths problem, BGP can be viewed as a distributed algorithm for solving the stable paths problem. Then as we have seen, some instances of the stable paths problem have no solutions and so it can be the case that BGP will fail to converge. Now that we have defined a fractional stable paths problem that is always guaranteed to have a solution, it would be interesting to develop a distributed algorithm, i.e., a protocol, that would solve the fractional stable paths problem. The development of such a protocol remains an open problem.

References

[AF03] R. Aharoni and T. Fleiner. On a lemma of Scarf. *Journal of Combinatorial Theory, Series B*, 87(1):72–80, 2003.

[GS62] D. Gale and L. Shapley. College admissions and stability of marriage. *American Math Monthly*, 69(1):9–15, 1962.

[GSW02] T. Griffin, B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, 10(2):232–243, 2002.

[GW99] T. Griffin and G. Wilfong. An analysis of BGP convergence properties. In *Proceedings of SIGCOMM*, pages 277–288, 1999.

[Nas50] J. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.

[RL95] Y. Rehkter and T. Li. A Border Gateway Protocol (BGP version 4). RFC 1771, 1995.

[Sca67] H. Scarf. The core of an n person game. *Econometrica*, 35(1):50–69, 1967.

[Tan91] J. Tan. A necessary and sufficient condition for the existence of a complete stable matching. *J. Algorithms*, 12(1):154–178, 1991.

[VGE00] K. Varadhan, R. Govindan, and D. Estrin. Persistent route oscillations in inter-domain routing. *Computer Networks*, 32:1–16, 2000.