

Partial specification of routing configurations

Alexander J. T. Gurney*, Limin Jia†, Anduo Wang* and Boon Thau Loo*

*University of Pennsylvania, Philadelphia, USA

†Carnegie-Mellon University, Pittsburgh, USA

Abstract—

The formal analysis of routing protocol configurations for safety properties is well established. Methods exist to identify potential protocol oscillations by analysis of the network topology and route preference information. However, if not all of this information is available, then the existing theory does not apply. We present an analysis of *partial specification* of protocol instances and apply it to eBGP and iBGP examples, so that potential oscillations can be detected from the incomplete data. This technique is applicable to the incremental design of network configurations, where some parts of the design have been specified but others are not yet known. We also anticipate that automated tools could be used to ‘fill in the blanks’ of a partial configuration in some optimal way. To that end, we show how our analysis can be used to derive constraints on an IGP weight matrix, characterizing the set of possible weights that do not lead to BGP oscillation. We propose that these integer constraints could be used as part of a link weight optimization engine, to achieve some traffic engineering goal while not violating global stability.

I. INTRODUCTION

The Internet is a network of networks, operated by entities that are often in commercial competition, but who still seek global connectivity. It is therefore essential that the routing system allow these networks to enforce policy that is at odds with what their neighbors might prefer. Unfortunately, this means that there will be some combinations of policy that are wholly incompatible, leading to protocol oscillation and other related problems. Various researchers have identified safety conditions, by which policy is restricted in order to ensure protocol convergence [1], [2], [3], [4]. But the formal treatment of these conditions assumes a model where all route preferences are known, in order for the verification to take place. In reality, many details about connectivity and policy will be unavailable. There is therefore a need for a theory which is able to speak about the correctness of configurations which are only partially known.

There are many aspects of a routing configuration which could be unspecified, including the protocol in use, the network topology, and the local configurations of each router. This paper focuses on the third of these—we assume that we are dealing with some network whose layout is known, and that the choice of routing protocol has been made (it is the Border Gateway Protocol [5]), but that protocol has not yet been fully configured. In particular, the *route preferences* may not be fully specified. We are interested in this scenario because it describes the state that a network operator may see when adjusting (configuring) parts of the router policy

configuration, while other parts of the policy are set and the network topology known.

Given the partial specification, we still want to be able to discuss the convergence of the overall system, based on the information that we do know. In this paper, we develop a formalism for modeling partial specifications based on the well-studied *stable paths problem* (SPP) model of interdomain routing [6], [2]. Whereas the conventional SPP definition gives a degree of preference to all routes, our *partial stable paths problem* definition allows some routes to have unknown preference. We are able to show that for any such instance, there are two possible outcomes. One is that the known preferences are already problematic, which means that no matter how the remaining route preferences are resolved, there is a potential oscillation. The other is that there exists at least one way to assign preferences to the remaining routes so that the resulting configuration will be guaranteed to be stable. In other words, if there is no safe way to resolve the unknown preferences, then that fact can be detected purely by examining the known preferences.

One application for analyzing partial problems is the iterative design of network configurations, proceeding from an incomplete view of how the network should behave, and then specifying more and more details, until finally a complete configuration is reached. The partial SPP definition allows us to reason about a set of complete configurations—all those whose preferences are consistent with those already given. Given a partial SPP, we can then determine the entire set of *completions* of that SPP for which convergence is guaranteed. By considering all the ways in which preferences could be assigned to the unranked paths, we can determine a precise characterization of the conditions when extensions to the partial specification are ‘safe’. The network operator can use this characterization to guide the incremental configuration.

This paper makes the following contributions.

- We propose a formal model (the partial stable paths problem) for partial specifications, and show how to discover potential oscillations in a sound manner.
- We identify a precise safety condition under which a partial specification has a safe extension.
- We explain how to use our formalism to aid policy configuration.

Roadmap The rest of this paper is organized as follows. In Section II, we present the formal definition for partial SPP, and explain how to detect possible oscillations based on the partial specification. In Section III, we use IGP weight assignment in iBGP as an example to motivate using our model to aid in-

cremental policy configuration. We formally define the safety conditions of a partial SPP in Section IV, and continue with our running example, deriving integer constraints on the IGP weight values. Finally we discuss related work (Section V) and future work (Section VI).

II. A FORMAL MODEL OF PARTIAL SPECIFICATIONS

In this section, we formally model partially specified policies in terms of a partial variant of the familiar stable paths problem. We also give a safety characterization, in terms of an auxiliary structure called the *paths digraph*.

A. The partial stable paths problem

Definition 1. A *partial stable paths problem* (PSPP) S is defined as a tuple $(G, d, \bigcup_{i \in N} P_i, \bigcup_{i \in N} \{(R_i, \leq_i)\})$, where

- $G = (N, E)$ is a graph, and N is the set of nodes and E is the set of directed arcs on N .
- d is a fixed destination node in N .
- For each node i , P_i is the set of *permitted paths* from i to d .
- For each node i , R_i is a subset of i 's permitted paths P_i , and \leq_i is a total order on R_i .

A PSPP is *well-formed* if P_d consists only of the empty path, from d to itself, and $R_d = P_d$. We only consider well-formed PSPPs. We say a PSPP is *full* if each node's total order is on the entire set of permitted paths (for all i , $R_i = P_i$). A full PSPP is an ordinary stable paths problem.

We can define a partial order among partial SPPs, and a notion of maximal element. For two PSPPs S and T , where

$$\begin{aligned} S &= (G, d, \{P_i\}_i, \{(R_i, \leq_i)\}_i) \\ T &= (G', d', \{P'_i\}_i, \{(R'_i, \leq'_i)\}_i) \end{aligned}$$

we say T extends S (written $S \leq T$) if $G = G'$, $d = d'$, and:

- For all i in N , $P_i = P'_i$.
- For all i in N , $R'_i \supseteq R_i$; and if $p \leq_i q$ for some p and q in R_i , then $p \leq'_i q$ as well.

T is called a *completion* of S , if $S \leq T$ and T is a total order. Such elements are maximal according to the partial order (there is no element strictly exceeding T).

We write $\uparrow(S)$ to denote the set of all extensions of S , formally: $\uparrow(S) = \{T \mid S \leq T\}$. These are exactly the PSPPs which can be obtained from S by giving preferences to the paths which, in S , did not have specific preferences.

B. Oscillation detection of partial specifications

Given a PSPP, one important question that we would like to answer is if the configuration is safe, as with the SPP theory. For a *partial* configuration to be safe, we mean that it is possible to complete it to a full configuration that has a unique solution (and therefore does not oscillate). If this is not possible then the partial configuration is already unsafe.

Previous work on routing stability in terms of stable paths problems has been founded on a *dispute digraph* structure [6]. Each SPP instance defines a dispute digraph instance. If the SPP instance represents a routing configuration which does not



Fig. 1. The two kinds of arc in a paths digraph.

converge, then the accompanying dispute digraph must contain a cycle; so acyclicity of the digraph is a sufficient condition for convergence.

Unfortunately, the dispute digraph definition is not ideal for our purpose of investigating partial specification, especially the safe completions of partial specification. Given PSPP instances S and T , with $S \leq T$, it is not obvious how their dispute digraphs should be related, and so it is difficult to make inferences about the stability of T from the data given in S .

We use a related definition, the *paths digraph*, based on earlier work by Sobrinho [3]. The paths digraph structure *does* have a clear relationship to the partial order on PSPP instances, and so we are able to use it in examining partial specification.

Definition 2. Given a PSPP S , the *paths digraph* is defined as follows. The nodes of the digraph are the permitted paths in S (elements in P_i). There are two kinds of directed arc:

- 1) An arc from p to q is a *transmission arc*, if p is the immediate prefix of q (i.e. $q = (ij)p$ for some nodes i and j).
- 2) An arc from p to q is a *preference arc*, if p and q are paths from the same source, and p is preferred to q .

We illustrate these arcs in Figure 1.

We say that S is *acyclic* if its paths digraph is acyclic. Note that the paths digraph is defined for an SPP, which is a full PSPP. The following theorem shows that acyclicity of the paths digraph is equivalent to acyclicity of the dispute digraph. See [7] for the proof of Theorems 1 and 2.

Theorem 1. *The dispute digraph of a PSPP S contains a cycle if and only if its paths digraph contains a cycle.*

Theorem 1 allows us to use the paths digraph to analyze safety. When we assign preferences to previously unranked paths, the corresponding operation on the paths digraph is simply to add new preference arcs. This means that we can detect cycles in the paths digraph before completing the configuration. The following theorem formalizes this idea, and characterises the instances for which we are able to ‘fill in the blanks’, and provide a complete, safe routing solution for the given partial configuration.

Theorem 2. *Let S be a PSPP. Then S has an acyclic completion if and only if S is acyclic.*

To give an example that has no safe completion, consider the PSPP in Figure 2. Its paths digraph is shown in Figure 3: it already contains a cycle, which does not go away when new preference arcs are added, and so this PSPP has no acyclic completion. No matter which preferences are chosen at node C, an oscillatory configuration called a *dispute wheel* [2] is




Fig. 2. A PSPP that does not have a completion.




Fig. 3. Paths digraph for the PSPP of Figure 2.

created. If $CAAd$ is preferred to $CEEd$ (written ' $CAAd \prec CEEd$ ') then nodes A, B and C are in a dispute wheel: node A has $ABd \prec Ad$, B has $BCEd \prec Bd$, and C has $CAAd \prec CEEd$. But if we instead choose $CEEd \prec CAAd$, there is a dispute wheel among C, E and F, since $CEEd \prec CAAd$, $EFd \prec Ed$ and $FCAd \prec Fd$. These options correspond, in the paths digraph, to adding a new preference arc between $CAAd$ and $CEEd$. This can be done in either direction, but there is still a cycle.

Since the absence of cycles from the dispute digraph of an SPP is sufficient for convergence [6], we obtain the following theorem.

Theorem 3. *If a PSPP is acyclic then it has some completion which converges.*

Proof: From Theorem 2, a PSPP which has an acyclic paths digraph must have some completion whose paths digraph is also acyclic; by Theorem 1, its dispute digraph is acyclic, and so convergence is guaranteed. ■

III. GUIDING INCREMENTAL CONFIGURATION

The ability to assess the stability of partially specified instances means that stability problems can be detected before the entire network configuration has been given. Our ultimate goal is to aid the process of finding a safe configuration that is a completion of some given partial configuration. The completion process could be incremental, where new preferences are incorporated one at a time, and in such a way as to avoid creating cycles in the paths digraph. Instead of

constructing a paths digraph and searching for cycles at each step, we can extract a condition to characterize all the ways that a given PSPP instance could be completed (so that all preferences are specified) without inducing a cycle—that is, given S , the subset of $\uparrow(S)$ containing only acyclic instances.

The advantages of obtaining such a condition include not only that it can be automatically checked, but also that it has the potential to be used by a partially or fully automated process to guide the configuration of the rest of the system.

In this section, we give an example of applying the partial SPP model to extract a safety condition; this is generalized in Section IV.

A. Stability of iBGP

Even within an autonomous system (AS), internal BGP (iBGP) presents stability concerns. In particular, interaction with the multi-exit discriminator (MED) attribute is known to be potentially problematic [8], [9], [10]. MED allows a neighboring AS to exercise control over egress point selection, overriding the normal decision procedure based on IGP distance ('hot-potato' routing). For particular choices of MED values and IGP weights, oscillations may occur.

In Figure 4, some combinations of IGP weight values will lead to oscillation and others will not. For example, if the weights are $w_{RS} = 1$, $w_{AR} = 5$, $w_{BR} = 4$, and $w_{CS} = 12$ then an oscillation will occur [8]. But this would not be the case if, say, w_{RS} were raised to 20.

B. Modeling IGP weight assignment as a PSPP

Given the problematic interactions between MED and IGP weights, it is desirable to have a process that derives IGP weights that are guaranteed not to cause such problems.

This situation is a good fit for our partial specification analysis. We will treat the IGP-induced preferences as unknown, with all other preferences (including those determined on the basis of MED) as known. Then, we can consider all of the ways in which the unknown preferences can be resolved, without creating a cycle in the paths digraph. Finally, each of these route preferences corresponds to an inequality on IGP weight values: so in the end, we will have a condition composed of linear inequalities (for which see Section IV-B).

In this example, we assume that the IGP weights in the network shown in Figure 4 are unknown. The paths digraph for this partially specified instance is shown in Figure 5. We have drawn additional unoriented preference arcs, between pairs of paths whose preferences are determined by IGP weight. In a completion, there will be *some* preference expressed, but the directions are not yet known.

The number of combinations to check may be reduced, since not all preferences will be capable of being induced by the same IGP weight settings. In this example, we can use the monotonic rule of integer arithmetic ($a < b$ if and only if $c + a < c + b$) to observe that the preference between $RAXd$ and $RBYd$ must be the same as between $SRAXd$ and $SRBYd$.

We obtain these three independent choices:

- $RAXd$ versus $RBYd$,

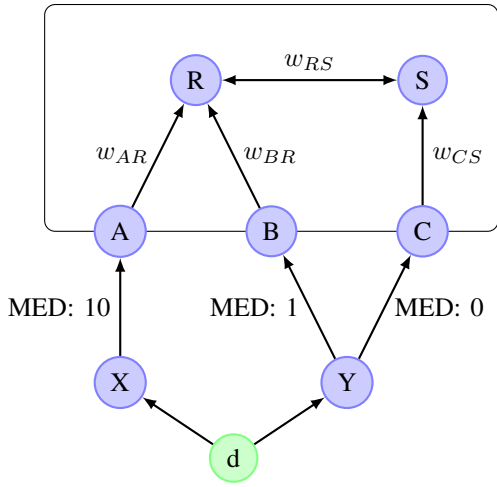


Fig. 4. Fill in the blanks: what choices of w_{AR} , w_{BR} , w_{RS} and w_{CS} ensure stability?

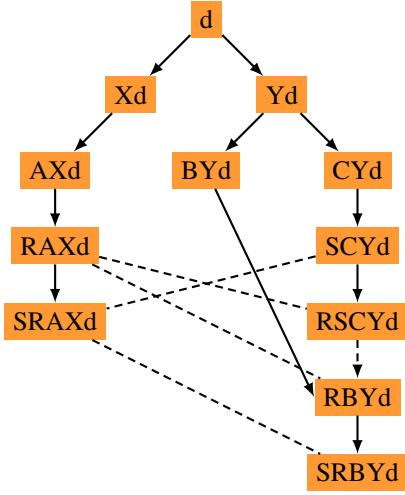


Fig. 5. Paths digraph for the example of Figure 4. Dashed lines without arrowheads indicate preference arcs whose direction is yet to be determined.

- RAXd versus RSCYd,
- SRAXd versus SCYd.

There are therefore no more than eight ways of resolving the preferences in this example. By inspection, there are three of these cases where a cycle occurs.

- SRAXd \prec SCYd, and RSCYd \prec RAXd
- RBYd \prec RAXd, and SRAXd \prec SCYd
- RBYd \prec RAXd, and RAXd \prec RSCYd.

To ensure safety, an IGP weight assignment of the network in Figure 4 needs to avoid making any of these preference combinations hold. In other words, such assignment will guarantee the path digraph to be acyclic, and by Theorem 3, the resulting routing policy converges.

IV. CAPTURING ALL SAFE COMPLETIONS

The procedure above can be generalized to derive a *cycle formula* for any PSPP. This will capture the combinations of

preference that induce cycles in the paths digraph. Its negation therefore captures the safe choices of preference.

In the case of the MED-IGP example, the unknown preferences relate to different choices of IGP weight values, which are integers. We can then translate the cycle formula into a formula on linear integer inequalities. This will describe the set of IGP weight matrices which can be considered safe.

A. The cycle formula

Definition 3. A *partially directed acyclic graph* (PDAG) is given by $G = (N, E_d, E_u)$, consisting of a set of nodes N , a set of directed arcs E_d , and a set of undirected arcs E_u , such that the graph (N, E_d) is acyclic [11].

The paths digraph of an acyclic PSPP defines a PDAG, where the undirected arcs are between any pair of nodes (from the same source to the same destination) whose relative preference is unspecified. These will be preference arcs in the completion, pointing one way or another, but for now we only know that there is *some* preference to be defined between those paths. The check for which completions are acyclic amounts to considering all possible orientations of these arcs, and finding which combinations induce cycles in the digraph.

We can trace paths and cycles in a PDAG, following both directed and undirected arcs.

Definition 4. For a path π in a PDAG, define its *enabling formula* to be the conjunction

$$E(\pi) = \bigwedge \{(p \prec q) \mid (p, q) \text{ is an unoriented arc in } \pi\}.$$

The enabling formula is intended to capture the preference decisions that need to be made, in completing a PSPP, for the given path to exist in the paths digraph. Note that $E(\pi)$ is always of finite length, since there are only finitely many unoriented arcs that π can traverse.

The formula includes the symbol ' \prec ' for strict preference. Logically, we will manipulate formulae under the assumption that \prec is the strict part of a preorder relation; in particular, it is irreflexive and transitive. For example, for all p and q , the conjunction $(p \prec q) \wedge (q \prec p)$ is equivalent to falsity.

We can now consider which combinations of preference lead to a cycle in the paths digraph.

Definition 5. For a PDAG G , define the *cycle formula* to be the disjunction

$$C(G) = \bigvee \{E(\pi) \mid \pi \text{ is a cycle in } G\}.$$

Again, this $C(G)$ is always of finite length. It tells us which combinations of preference induce cycles in the paths digraph. Its negation $\neg C(G)$ does the opposite, characterizing the preference decisions that are guaranteed to be safe, formally:

Theorem 4. Given a PSPP, S , let L be the negation of the cycle formula on its paths digraph. If L is satisfiable, then there is at least one T , with $S \leq T$, such that T satisfies L , and T converges.

Proof: If L is satisfiable, then there is at least one way of assigning preferences to the unranked paths in S , so that no cycle is created in the paths digraph. Let T be one such resulting PSPP. Then $S \leq T$, and T satisfies L , and T has no cycle in its paths digraph, so T must converge. ■

Given a PSPP S , a cycle formula can be automatically constructed using a brute-force algorithm that enumerates all possible combinations of assignments of unknown preferences. The performance of this procedure is not critical for analysis, since the formula is pre-computed. We leave optimization of such a process for future work.

B. Transfer to the integer domain

The derivation above gives us a condition in terms of path preferences. For the networking application, we need to translate these preferences back into the language of BGP. We will obtain a version of the cycle formula that makes reference to the attributes of BGP routes, as opposed to abstract relations like ' $p \prec q$ '. In doing this, we can find that some combinations of preferences that were *apparently* feasible according to the cycle formula, are in fact unobtainable. This can happen if there is no way of configuring the BGP routing policy so that those preferences are the result. Such manipulations could be automated with the aid of an SMT solver capable of handling linear arithmetic, such as Yices [12].

In Figure 5, the unknown path preferences are induced by the IGP weights. Consequently, every preference clause ' $p \prec q$ ' can be translated into an inequality between the weights of paths p and q . (If other attribute values were unknown, this complete rewriting would not necessarily be possible.) A straightforward translation is for the preference 'RAXd \prec RSCYd'. To make this happen, the weight w_{AR} must be strictly less than the sum of weights $w_{RS} + w_{CS}$. So we can translate this preference into the integer inequality $w_{AR} - w_{RS} - w_{CS} < 0$. Each of the other path preferences can be handled in the same way, so the final formula contains only conjunctions and disjunctions of integer inequalities.

Based on the attributes which follow IGP distance in the BGP decision chain, we can determine whether each inequality should be strict. This yields a more precise characterization of the solution space. For example, when route reflectors are in use, the next attribute after IGP distance is the route reflector originator identifier [13]. If we want path p to be better than path q , and the originator for p has a lower identifier than that for q , then we know that the sum of IGP weights for p must be less than *or equal to* the sum for q . If the weights were equal, then the next tiebreaker would give us the desired preference in any case. But if the originator identifiers were ordered in the opposite sense, then the required IGP distance inequality would be a strict one.

An example of integer-based simplification is as follows. It appears that one possibility for assigning preferences is to have RSCYd better than RAXd, and at the same time SRAXd better than SCYd. But this cannot happen, since the inequalities $w_{RS} + w_{CS} < w_{AR}$ and $w_{RS} + w_{AR} < w_{CS}$ cannot be simultaneously satisfied for natural number weight values.

We finally obtain

$$(w_{RS} + w_{CS} < w_{AR}) \\ \vee ((w_{AR} < w_{RS} + w_{CS}) \wedge (w_{AR} < w_{BR}))$$

as the negated formula. This identifies just two ways of avoiding a cycle in the paths digraph.

The mitigation advice in RFC 3345 is to make the link weight w_{RS} large [8], so that each reflector will select a best path from its local cluster, rather than needing to go via another reflector. This formula tells us how large that weight has to be compared to the other link weights, and also what other ways there are to avoid oscillation. For example, if w_{RS} were as small as 1, we would still converge if we set w_{CS} to 1 and w_{AR} to 3. If this could not be done due to w_{AR} also having to be small, then we could ensure convergence by increasing w_{BR} . Each option has different implications for egress point selection and internal traffic flow, which can be predicted from the routing data.

The consideration of every safe completion has therefore not only revealed additional possibilities for how to make a partial configuration safe: it has also established precise numeric bounds on the integer parameters involved.

V. RELATED WORK

Routing algebra The main alternative formalism to the SPP model is *routing algebra* [14], [4], [15]. There are various different structures used in this theory, but they share the general idea that route preferences can be considered independently from the network topology. This is, in a sense, a form of partial specification: once we have decided on the route *signatures* and on their relative preferences, and verified certain correctness properties, protocol convergence is guaranteed regardless of the layout of the network. Could our notion of partial specification of preference also be treated in an 'algebraic' way? Some steps in this direction have already been taken, since we have convergence results which apply even when some pairs of routes are *incomparable* [15], and a notion of *refinement* whereby those incomparabilities are resolved in one way or another, becoming strict preferences [16].

Indeed, this approach is very close to the 'partial SPPs' we treat in this paper; there are, however, some obstacles to using the algebraic definitions for the partial specification problem. The principal one is that in the algebraic theory, certain properties must hold of path preferences as a matter of definition: for example, the order must be transitive. Therefore, we are forbidden to even speak about certain situations where cyclic preferences occur—they simply cannot be formulated in terms of routing algebra. But we want to be able to say that certain rankings induce cyclic preferences, whereas others do not, and to capture exactly which preferences are involved in either case. We cannot do this, or at least not as straightforwardly, in the algebraic world.

For algebras satisfying the *strict inflationary* property, the set of all paths in the labelled graph can be put in preference order, where a path is always preferred over any extension of that path; and convergence is guaranteed in this case [15].

This viewpoint is compatible with the paths digraph structure, which has transmission arcs (for path extension) and preference arcs, but is required to be acyclic overall.

MED oscillation There is substantial literature on the ‘MED problem’ in BGP, extending from practical advice [8], to theoretical approaches [9], to proposed protocol changes [10]. Our partial specification analysis is aimed at understanding the shape of the parameter space for avoiding MED-induced oscillation, and assumes the traditional BGP mechanism. Knowing this, more precise advice can be given about which combinations of attribute values are desirable. So far, our approach has not been extended to multipath models, for which the oscillation problem may be alleviated [10].

VI. FUTURE WORK

We would like to provide an automated tool that could derive the cycle formula from a given partial BGP configuration (or partial SPP instance). Derivation of this formula is likely to be computationally difficult, though hopefully not intractable in practice. As presented, we are dealing with a data structure whose components are *all simple paths* in a graph, and considering *all possible* ways of orienting the undirected arcs in a partially directed graph. This can be mitigated in the case of the IGP weight example, by making use of facts about integer arithmetic.

The next part of the story is to put the formula to good use. We anticipate that our analysis will be a good fit for traffic engineering methods: we provide safety constraints in terms of the integer parameter values, and then some automated process will tune those weight values to achieve some optimal outcome, subject to the constraints. This would mean integrating our generated safety constraint into an existing traffic engineering tool. We could then assess the quality of configurations that were output, as well as the method’s performance, for various different scenarios.

Traffic engineering methods can be divided into two categories: there are heuristic techniques, making use of local search, genetic algorithms, and the like; and there are exact techniques, drawn from linear programming and related work in mathematical optimization [17]. The output of our analysis could be used in each of these settings. Firstly, the predicate that we derive can be used to screen candidate configurations in a local search process, such as tabu search [18]. In this way, the search can be restricted to only considering weight assignments that ensure stability. Secondly, the fact that the components of our predicate are linear inequalities suggests that there should be a more intelligent way to combine it with linear programming methods, though the presence of disjunctions is a complication.

Existing techniques for assigning link weights do so from the perspective of optimization, ignoring the fact that these weights are also used in BGP egress point selection. There are some approaches that attempt to capture BGP hot potato logic, by simulating the BGP decision process [19], [20], [21]. This does not directly address the stability issue: it is more concerned with the effect of BGP on the traffic matrix.

Additionally, the use of a simulator is somewhat unwieldy, compared to a pure numeric method.

There may be other occasions where the partial specification theory is well matched with an automated configuration process. The theory could also be extended to handle scenarios where other kinds of partial information are involved, such as when the permitted paths are unknown, or when some detail of the protocol rules must be inferred.

ACKNOWLEDGEMENT

We are grateful to Jennifer Rexford for discussions on the topic and potential applications of this paper. This research was supported by NSF grants CCF-0820208, CNS-0845552, CNS-1040672, CNS-1115706, and CNS-1117052; and AFOSR grant no. FA9550-08-1-0352.

REFERENCES

- [1] L. Gao and J. Rexford, “Stable internet routing without global coordination,” *IEEE/ACM Transactions on Networking*, pp. 681–692, 2001.
- [2] T. G. Griffin, F. B. Shepherd, and G. Wilfong, “The stable paths problem and interdomain routing,” *IEEE/ACM Transactions on Networking*, vol. 10, no. 2, pp. 232–243, 2002.
- [3] J. Sobrinho, “Network routing with path vector protocols: Theory and applications,” in *Proc. ACM SIGCOMM*, pp. 49–60, 2003.
- [4] T. G. Griffin and J. Sobrinho, “Metarouting,” in *Proc. ACM SIGCOMM*, pp. 1–12, 2005.
- [5] Y. Rekhter, T. Li, and S. Hares, “RFC 4271: A Border Gateway Protocol 4 (BGP-4),” 2006.
- [6] T. G. Griffin, F. B. Shepherd, and G. Wilfong, “Policy disputes in path-vector protocols,” in *Proc. IEEE ICNP*, 1999.
- [7] A. J. T. Gurney, L. Jia, A. Wang, and B. T. Loo, “Partial specification of routing configurations,” Department of Computer and Information Science, University of Pennsylvania, Tech. Rep., 2011. [Online]. Available: <http://netdb.cis.upenn.edu/papers/partial-spec-tr.pdf>
- [8] D. McPherson, V. Gill, D. Walton, and A. Retana, “RFC 3345: Border Gateway Protocol (BGP) persistent route oscillation condition,” 2002.
- [9] T. G. Griffin and G. Wilfong, “An analysis of the MED oscillation problem in BGP,” in *Proc. IEEE ICNP*, 2002.
- [10] A. Flavel and M. Roughan, “Stable and flexible iBGP,” in *Proc. ACM SIGCOMM*, 2009.
- [11] D. Dor and M. Tarsi, “A simple algorithm to construct a consistent extension of a partially oriented graph,” Computer Science Department, University of California, Los Angeles, Tech. Rep. R-185, 1992.
- [12] B. Duterte and L. de Moura, “A fast linear-arithmetic solver for DPLL(T),” in *Proc. CAV*, pp. 81–94, 2006.
- [13] T. Bates, E. Chen, and R. Chandra, “RFC 4456: BGP route reflection: an alternative to full mesh internal BGP (iBGP),” 2006.
- [14] J. Sobrinho, “An algebraic theory of dynamic network routing,” *IEEE/ACM Transactions on Networking*, vol. 13, no. 5, pp. 1160–1173, 2005.
- [15] A. J. T. Gurney, “Construction and verification of routing algebras,” Ph.D. dissertation, University of Cambridge, 2009.
- [16] A. J. T. Gurney and T. G. Griffin, “Neighbor-specific BGP: an algebraic exploration,” in *Proc. IEEE ICNP*, 2010.
- [17] N. Wang, K. H. Ho, G. Pavlou, and M. Howarth, “An overview of routing optimization for Internet traffic engineering,” *IEEE Communications Surveys*, vol. 10, no. 1, pp. 36–56, 2008.
- [18] B. Fortz and M. Thorup, “Increasing Internet capacity using local search,” *Combinatorial Optimization and Applications*, vol. 29, no. 1, pp. 13–48, 2004.
- [19] B. Quoitin and S. Uhlig, “Modeling the routing of an autonomous system with C-BGP,” *IEEE Network*, vol. 19, no. 6, pp. 12–19, 2005.
- [20] S. Cerav-Erbas, O. Delcourt, B. Fortz, and B. Quoitin, “The interaction of IGP weight optimization with BGP,” in *Proc. International Conference on Surveillance and Protection*, 2006.
- [21] S. Balon and G. Leduc, “BGP-aware IGP link weight optimization in the presence of route reflectors,” in *Proc. IEEE INFOCOM*, pp. 316–324, 2009.