# Subjective-Cost Policy Routing*

Joan Feigenbaum[1,**], David R. Karger[2,***], Vahab S. Mirrokni[3,***], and Rahul Sami[4,†]

[1] Yale University Computer Science Department,
51 Prospect St., New Haven, CT 06520, USA
feigenbaum@cs.yale.edu

[2] MIT Computer Science and Artificial Intelligence Laboratory,
32 Vassar St, Cambridge, MA 02139, USA
karger@csail.mit.edu

[3] MIT Computer Science and Artificial Intelligence Laboratory,
32 Vassar St, Cambridge, MA 02139, USA
mirrokni@csail.mit.edu

[4] University of Michigan School of Information,
1075 Beal Ave., Ann Arbor, MI 48109, USA
rsami@si.umich.edu

**Abstract.** We study a model of interdomain routing in which autonomous systems' (ASes') routing policies are based on *subjective* cost assessments of alternative routes. The routes are constrained by the requirement that all routes to a given destination must be confluent. We show that it is NP-hard to determine whether there is a set of stable routes. We also show that it is NP-hard to find a set of confluent routes that minimizes the total subjective cost; it is hard even to approximate minimum cost closely. These hardness results hold even for very restricted classes of subjective costs.

We then consider a model in which the subjective costs are based on the relative importance ASes place on a small number of objective cost measures. We show that a small number of confluent routing trees is sufficient for each AS to have a route that nearly minimizes its subjective cost. We show that this scheme is trivially strategyproof and that it can be computed easily with a distributed algorithm that does not require major changes to the Border Gateway Protocol. Furthermore, we prove a lower bound on the number of trees required to contain a $(1 + \epsilon)$-approximately optimal route for each node and show that our scheme is nearly optimal in this respect.

## 1 Introduction

The Internet is divided into many *Autonomous Systems* (ASes). Loosely speaking, each AS is a subnetwork that is administered by a single organization. The task of routing between different ASes in the Internet is called *interdomain routing*.

* This work was supported by the DoD University Research Initiative (URI) administered by the Office of Naval Research under Grant N00014-01-1-0795.

** Supported in part by ONR grant N00014-01-1-0795 and NSF grant ANI-0207399.

*** Supported in part by NSF grant ANI-0225660.

† Supported by NSF grant ITR-0219018. Most of this work was done while the author was at MIT.

X. Deng and Y. Ye (Eds.): WINE 2005, LNCS 3828, pp. 174–183, 2005.

Currently, the only widely used protocol for interdomain routing is the Border Gateway Protocol (BGP). BGP allows an AS to "advertise" routes it currently uses to neighboring ASes. An AS $i$ with many neighbors may thus receive advertisements of many different routes to a given destination $j$. It must then select one of these available routes as the route it will use to send its traffic; subsequently, $i$ can advertise this chosen route (prefixed by $i$ itself) to all *its* neighbors. Proceeding in this manner, every AS in the Internet can eventually discover at least one route to destination $j$.

Thus, one of the key decisions an AS has to make is that of *route selection*: Given all the currently available routes to destination $j$, which one is traffic sent on? At first glance, it may seem as though ASes would always prefer the shortest route; in practice, however, AS preferences are greatly influenced by other factors, including perceived reliability and existing commercial relationships. For this reason, BGP allows ASes complete freedom to pick a route according to their own *routing policies*. The resulting routing scheme is called *policy-based routing*, or *policy routing* for short. However, BGP does place one important constraint on routing: It stipulates that an AS can only advertise a route that the advertising AS itself currently uses. This is because of the way traffic is routed in the Internet: Routers examine the destination of incoming packets and simply forward the packet to the next hop on the current route to that destination. At a given time, each AS typically has exactly one active route to the destination. Thus, the set of all ASes' routes to a given destination AS $j$ must be *confluent*, *i.e.*, they must form a tree rooted at $j$.

The policy-routing aspect of interdomain routing has recently received a lot of attention from researchers. Varadhan *et al.* [1] observed that general policy routing could lead to route oscillations. Griffin, Shepherd, and Wilfong [2, 3] studied the following abstract model of general policy routing: Each AS $i$'s policy is represented by a preference ordering over all possible routes to a given destination $j$. At any given time, $i$ inspects the routes all of its neighbors are advertising to $j$ and picks the one that is ranked highest. AS $i$ then advertises this route (prefixed by $i$ itself) to all its neighbors. Griffin *et al.* proved that, in such a scenario, BGP may not converge to a set of *stable paths*; the routes might keep oscillating as ASes continuously change their selection in response to their neighbors' changes. They further showed that, given a network and a set of route preferences, it is NP-complete to determine whether a set of stable paths exists. In recent work, Feamster *et al.* [4] showed that instability can arise even for restricted routing policies.

Feigenbaum *et al.* [5] extended the model of [2] by including cardinal preferences instead of preference orderings. Specifically, they assume that AS $i$ conceptually assigns each potential route a monetary value and then ranks routes according to their value. The advantage of working with cardinal preferences is that a set of paths can be stabilized by making payments to some of the ASes: Although the ASes' *a priori* preferences may have led to oscillation (in the absence of payments), ASes preferences can be changed if they receive more money for using a less valuable route. This is the basis for the mechanism-design approach to routing, which seeks to structure incentives so as to achieve a stable, globally optimal set of routes; see [5] for further details. In the context of policy routing, the most natural global goal is to select a set of confluent routes that maximizes the *total welfare* (the sum of all ASes' values for their selected routes). However, Feigenbaum *et al.* showed that, for general valuation functions, it is NP-hard to find a welfare-maximizing set of routes; it is even NP-hard to approximate the maximum welfare to within a factor of $n^{\frac{1}{4}-\epsilon}$, where $n$ is the number of nodes. Thus, in this model too, general routing policies lead to computationally intractable problems.

The natural approach to get around the intractability results is to restrict either the network or the routing policy. Restricting the network alone does not appear to be a very promising direction, because the hardness results hold even for fairly simple networks that cannot be excluded without excluding many "Internet-like" networks. This has led researchers to turn to restricted classes of preferences that can express a wide class of routing policies that ASes use in practice. Feigenbaum *et al.* [5] study *next-hop* preferences – preferences in which an AS $i$'s value for a path depends only on the next AS on the path – and show that, in this case, a welfare-maximizing set of routes can be found in polynomial time. Next-hop preferences can capture the effects of $i$'s having different commercial relationships with neighboring ASes. Similarly, in the ordinal-preference model, Gao and Rexford [6] show that, with the current hierarchical Internet structure, BGP is certain to converge to a set of stable paths as long as every AS prefers a customer route (*i.e.*, a route in which the next hop is one of its customers) over a peer or provider route; this can also be viewed as a next-hop restriction on preferences.

However, there are many useful policies that cannot be expressed in terms of next-hop preferences alone. In this paper, we study other classes of routing policies that capture realistic AS preferences. For example, an AS $i$ might wish to avoid any route that goes through AS $k$, either because it perceives $k$ to be unreliable or because $k$ is a malicious competitor who would like to drop all of $i$'s traffic. This leads to the *forbidden-set* class of routing policies: For each AS $i$, there is a set of ASes $S_i$ such that $i$ prefers any route that avoids $S_i$ over any route that uses a node in $S_i$. We can then ask the following questions: (1) If each node uses a forbidden-set routing policy, will BGP converge to a set of stable paths?, and (2) Can we find a welfare-maximizing routing tree, *i.e.*, a set of confluent routes that maximizes the number of nodes $i$ whose routes do *not* intersect the sets $S_i$? If the latter optimization problem were tractable, then this class of routing policies would be a candidate for a mechanism-design solution as in [7].

Forbidden-set policies (and many others) can be framed in terms of *subjective costs*: Each AS $i$ assigns a cost $c_i(k)$ to every other AS $k$. Then, the "cost" perceived by AS $k$ for a route $P$ is $\sum_{k\in P} c_i(k)$; AS $i$ prefers routes with lower subjective cost. Subjective-cost routing is a natural generalization of lowest-cost routing (in which there is a single objective measure of cost that all ASes agree upon). It is well known that lowest-cost routes can be computed easily, and hence we hope that some more general class of subjective-cost routing policies will also be tractable.

However, we find that even very restricted subsets of subjective-cost policies lead to intractable optimization problems: We show that, if all ASes rank paths based on subjective-cost assignments, it is still possible to have an instance in which there is no stable-path solution. Further, given a network and subjective costs, it is NP-complete to determine whether there is a set of stable paths. Moreover, the NP-completeness reduction only requires subjective costs in the range $\{0, 1, 2\}$ for each node. In the cardinal utility model, the outlook is not much brighter: We show that, even if all subjective costs are either $0$ or $1$, it is NP-hard to find a set of routes that maximizes the overall welfare; indeed, it is NP-hard even to approximate maximum welfare to within *any* factor. The forbidden-set routing policies can be formulated in terms of $0$-$1$ subjective costs, and hence optimizing for this class is also difficult. We then turn to subjective costs with bounded ratios. We show that, if the subjective costs are restricted to lie in the range $[1, 2]$, the problem of finding a confluent tree with minimum total subjective cost is APX-hard; thus finding a solution that is within a $(1 + \epsilon)$ factor of optimal is intractable. In this case, however, an unweighted shortest-path tree provides a trivial 2-approximation to the optimization problem.

In light of all these hardness results, we consider a more restricted scenario in which the differing subjective cost assignments arise from differences in the relative importance placed on two *objective* metrics, such as latency and reliability. Thus, we suppose that every path $P$ has two objective costs $l_1(P)$ and $l_2(P)$. We assume that AS $i$ evaluates the cost of path $P$ as the convex combination $\lambda_i l_1(P) + (1 - \lambda_i)l_2(P)$, where $\lambda_i \in [0, 1]$ reflects the importance $i$ places on the first metric. Here, too, it is NP-hard to find a routing tree that closely approximates the maximum welfare. However, if we slightly relax the constraint that each AS stores only a single route to the destination, we show that it is possible to find a nearly optimal route, as follows. Given any $\epsilon > 0$, we can find a set of $O(\log n)$ trees[1] rooted at $j$ with the following property: If each AS $i$ chooses the route it likes best among the $O(\log n)$ alternatives, the overall welfare is within a $(1 + \epsilon)$ factor of optimal. This solution can be implemented by replacing each destination with a set of $O(\log n)$ logical destinations and then finding a lowest-cost routing tree to each of these logical destinations. The results generalize to the convex combinations of $d > 2$ objective metrics; $O(4^d \log^{d-1} n)$ trees are required in this case.This scheme is trivially strategyproof, and, further, it can be implemented with a "BGP-based" algorithm, *i.e.*, an algorithm with similar data structures and communication patterns to BGP (*cf.* [7, 5]).

The rest of this paper is structured as follows: In section 2, we introduce the subjective-cost model of routing preferences. In section 3, we study the stable-paths problem for path rankings based on subjective costs. In sections 4 and 5, we study the problem of finding a routing tree that minimizes the total subjective cost. Due to space restrictions, the proofs have been omitted from this extended abstract; they will appear in the final version of the paper.

## 2 Subjective-Cost Model for Policy Routing

In this section, we present the subjective-cost model of AS preferences. The model involves each AS $i$'s assigning a cost $c_i(k)$ to every other AS $k$. These costs are *subjective*, because there is no requirement that $c_i(.)$ and $c_k(.)$ be consistent. We assume that each subjective cost $c_i(k)$ is non-negative. The total cost of an AS $i$ for a route $P_{ij}$ to destination $j$ is

$$c_i(P_{ij}) = \sum_{k \in P_{ij}} c_i(k) .$$

[1] The dependence on $\epsilon$ is detailed in Section 5.

Here, the notation $k \in P_{ij}$ is used to indicate that $k$ is a *transit node* on the the path $P_{ij}$; $i$ and $j$ are thus excluded from the summation. AS $i$ wants to use a route $P_{ij}$ that minimizes the cost $c_i(P_{ij})$.

The subjective-cost model can be used to express a wide range of preferences, but it does place some restrictions on AS preferences. For instance, an AS $i$ cannot prefer a path $P$ over a path $P'$ whose nodes are a strict subset of $P$. The class of preferences that can be expressed as subjective costs includes:

- Lowest-cost routing
  If $c_i(k)$ is the actual cost of transiting AS $k$, minimizing the path cost is exactly lowest-cost routing.
- Routing with a forbidden set
  Let $c_i(.)$ take the following form: If $k \in S_i$, $c_i(k) = 1$, else $c_i(k) = 0$. Then any route that avoids ASes in $S_i$ is preferred by $i$ over any route that involves an AS in $S_i$.

Subjective costs can form the basis for either ordinal preferences or cardinal utilities. In section 3, we study the stable-paths problem for path rankings based on subjective costs. In sections 4 and 5, we study the problem of finding a routing tree that minimizes the total subjective cost.

## 3 Stable Paths with Subjective Costs

The *Stable Paths Problem (SPP)*, introduced by Griffin *et al.* [3], is defined as follows. We are given a graph with a specified destination node $j$. Each other node $i$ represents an AS; there is an edge between two nodes if and only if they exchange routing information with each other. Thus, a path from $i$ to $j$ in the graph corresponds to a potential route from AS $i$ to the destination. Each AS $i$ ranks all potential routes to destination $j$. A *route assignment* is a specification of a path $P_{ij}$ for each AS $i$ such that the union of all the routes forms a tree rooted at $j$ (*i.e.*, the confluence property is satisfied). A route assignment is called *stable* if, for every AS $i$, the following property holds: For every neighbor $a$ of $i$, AS $i$ does not strictly prefer the path $aP_{aj}$ over the path $P_{ij}$; in other words, $i$ would not want to change its current route to any of the other routes currently being advertised by its neighbors. The stable-paths problem is *solvable* if there is a stable route assignment.

Griffin *et al.* [2, 3] have shown that there are instances of SPP that are unsolvable, and, further, that it is NP-complete to determine whether a given SPP is solvable. Their constructions used preferences that cannot be directly expressed as subjective-cost preferences. This leads us to hope that, for subjective-cost preferences, the stable-paths problem might be tractable. Unfortunately, this is not the case. In this section, we prove that these hardness results extend to subjective-cost preferences.

Assume that the rankings assigned by ASes are based on an underlying subjective-cost assignment. Then, the stable paths problem can be viewed in terms of a strategic game, as follows: The players of this game are the ASs. Given a graph $G(V, E)$ with a specific destination $j$ and a subjective-cost function $c : V(G) \times V(G) \to \mathcal{R}$, the *next-hop game* is defined as follows. ASes correspond to the vertices of graph $G$. The strategy

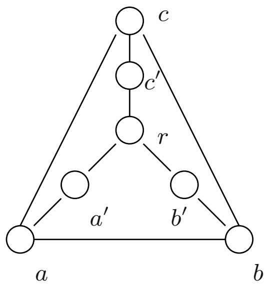


**Fig. 1.** A bad triangle

space for AS $i$ is the set $N(i)$ of neighboring nodes in the graph; thus, AS $i$'s picking the route advertised by a neighboring AS $a$ corresponds to $i$'s playing strategy $a$. Given a vector of strategies (one for each player), the cost incurred by player $i$ is the subjective cost of its route to the destination; if there is no route from $i$ to the destination, $i$'s cost is $\infty$. A vector of strategies is a pure-strategy Nash equilibrium if, given the strategies of all the other ASes, no AS could decrease its subjective cost by changing its strategy. A pure-strategy Nash-equilibrium strategy profile must result in every AS's having some route to $j$, and, hence, it must correspond to a valid route assignment. Thus, proving that an SPP is solvable is equivalent to proving that the corresponding next-hop game has a pure-strategy Nash equilibrium.

**Definition 1.** *The* bad triangle *is defined as follows. It is a graph $G$ with vertex set $\{a, b, c, a', b', c', r\}$ and edge set $\{aa', a'r, cc', c'r, bb', b'r, ab, bc, ca\}$. Set $c_a(c) = c_a(c') = 0$; $c_b(a) = c_b(a') = 0$; and $c_c(b) = c_c(b') = 0$. All other subjective costs are set to* 1*. A bad triangle is shown in Figure 1. (This construction is based on the* bad gadget *defined in [2].)*

In the bad triangle, AS $a$ prefers the path $(a, b, b', r)$ to the path $(a, a', r)$, AS $b$ prefers the path $(b, c, c', r)$ to the path $(b, b', r)$, and AS $c$ prefers the path $(c, a, a', r)$ to the path $(c, c', r)$. It follows from the arguments in [2] that this network is not solvable.

We now show that, as in the case of unrestricted routing policies, it is NP-complete to determine if an SPP based on subjective-cost preferences is solvable.

**Theorem 1.** *Given an instance of the next-hop game, it is NP-complete to decide whether it has a pure Nash equilibrium or not.*

The proof is based on the corresponding NP-completeness proof in [2].

## 4 The Minimum Subjective-Cost Tree (MSCT) Problem

In this section, we assume that the subjective cost $c_i(k)$ is an actual monetary amount that is measured in the same unit across all ASes. A natural overall goal is then to minimize the sum of subjective costs, *i.e.*, to pick a set of routes $\{P_{ij}\}$ that minimizes $\sum_j \sum_i c_i(P_{ij})$. However, there is a constraint that all the routes to a single destination $j$ must form a tree, because the packets are actually sent by forwarding. This constraint applies independently to each destination, and so we can consider the simpler problem of routing to a single destination $j$.

Thus, we can frame the subjective-cost minimization problem as:

**Subjective-cost minimization:** We are given a graph $G$, a set of cost functions $\{c_i(.)\}$, and a specific destination $j$. We want to find a set of routes $\{P_{ij}\}$ and payments $p_i$ to each AS $i$ such that:

1. The routes $\{P_{ij}\}$ form a tree rooted at $j$.
2. Among all such trees, the selected tree minimizes the sum $\sum_i \sum_{k \in P_{ij}} c_i(k)$.

We first prove that, for arbitrary cost functions, the MSCT problem is NP-hard to approximate within any multiplicative factor. Let $c_{\max} = \max_{v,u \in V(G)} c_v(u)$ and $c_{\min} = \min_{v,u \in V(G)} c_v(u)$. Then, we have the following result:

**Theorem 2.** *It is NP-hard to approximate the MSCT problem within a factor better than $\frac{c_{\max}}{c_{\min} n^2}$, where $n$ is the number of vertices. In particular, it is NP-hard to approximate MSCT within any factor if $c_{\min} = 0$ and $c_{\max} > 0$.*

Note that the above theorem does not show hardness for the special cases in which $\frac{c_{\max}}{c_{\min}}$ is not large. This may be a reasonable restriction; however, we now show that this also yields an intractable optimization problem. In particular, we study the special case in which all subjective costs are either 1 or 2. We call this problem the $(1,2)$-MSCT problem. In the following, we give a hardness result for the $(1,2)$-MSCT problem.

**Theorem 3.** *The $(1,2)$-MSCT problem is APX-Hard.*

Theorem 3 shows that, for sufficiently small $\epsilon$, it is hard to find a $(1+\epsilon)$-approximation for the $(1,2)$-MSCT problem. However, we note that finding a 2-approximation is easy: Simply ignore the costs, and construct an unweighted shortest-path tree with destination $j$. This is optimal to within a factor of 2, because the number of nodes on the shortest path from $i$ to $j$ is a lower bound on the subjective cost $c_i(P_{ij})$ for any path $P_{ij}$ from $i$ to $j$.

## 5 An Alternative Model: Subjective Choice of Metrics

In this section, we consider a more restricted preference model. We assume that there are multiple objective metrics on routes (*e.g.*, cost and latency), and ASes' preferences differ only in the relative importance they accord to different metrics. This is a non-trivial restriction only when the number of objective metrics is small; here, we first consider the case in which there are only two objective metrics on a route. The results are generalized to $d > 2$ objective metrics in Section 5.1.

Formally, suppose that any transit AS $k$ has two associated objective "length" values $l_1(k)$ and $l_2(k)$. Both the length values can be extended to additive path metrics, *i.e.*, we can define $l_1(P_{ij}) = \sum_{k \in P_{ij}} l_1(k)$ and $l_2(P_{ij}) = \sum_{k \in P_{ij}} l_2(k)$. Note that we use the term "metric" for the ease of presentation and that we do not impose the triangle equality on the length functions $l_1$ and $l_2$.

Each AS $i$ has a private parameter $\lambda_i$, $0 \leq \lambda_i \leq 1$. AS $i$'s subjective cost for the route $P_{ij}$ is given by $c_i(P_{ij}) = \lambda_i l_1(P_{ij}) + (1 - \lambda_i) l_2(P_{ij})$, *i.e.*, AS $i$'s preferences are modeled as a convex combination of the two path metrics.

It is easy to show that the APX-hardness proof for the $(1, 2)$-MSCT problem (Theorem 3) can be adapted to the two-metric routing problem as well:

**Theorem 4.** *In the subjective-metric model, it is APX-hard to find a tree $T$ that minimizes total subjective cost.*

We now investigate whether relaxing the confluent-tree routing constraint would lead to stronger results. If we allowed the routes to be completely arbitrary, then clearly we could have optimal routing: Each AS could simply use the route it liked the best. However, supporting these routes would either require source routing (*i.e.*, the packet header contains a full path) or a massive increase in storage at each router to record the forwarding link for each source and destination. Instead, we ask whether we can get positive results with only a small growth in routers' space requirements.

Our approach is to use a *small number* $r$ of confluent routing trees $T_1, T_2, \cdots T_r$ to each destination $j$. Then, each AS $i$ evaluates its subjective cost to $j$ in each of the routing trees and picks a tree $T_{t_i}$ that minimizes this subjective cost. AS $i$ then marks each packet it sends with the header $< j, t_i >$. Each AS *en route* stores its route to $j$ along each tree $T_j$; thus, it can inspect the header of each incoming packet and forward along the appropriate route.

We can prove the following result:

**Theorem 5.** *Suppose that, for transit AS $k$ , $l_1(k)$ and $l_2(k)$ are integers bounded by a polynomial, i.e., $l_1(k), l_2(k) < n^c$ for some constant c. Then, for any given $\epsilon > 0$, there is a set of routing trees $T_1, T_2, \cdots T_r$ with $r = O(\frac{1}{\epsilon}[\log n + \log(\frac{1}{\epsilon})])$ such that:*

*For each AS i, there is a tree $T_{t_i}$ such that $c_i(T_{t_i}) \leq (1 + \epsilon) c_i(P^*_{ij})$, where $P^*_{ij}$ is the minimum-subjective-cost route from $i$ to $j$.*

*Further, this set of trees can be constructed in polynomial time.*

We now sketch the tree construction used in the proof.

Let $\alpha = (1 + \epsilon)$. Each tree $T_t$ in our collection is the shortest-path tree for a specific convex combination of the two metrics. We name the trees after the metrics they optimize:

$T_\infty$**:** $l_1(\cdot)$, with ties broken by minimum $l_2(\cdot)$.
$T_{-\infty}$**:** $l_2(\cdot)$, with ties broken by minimum $l_1(\cdot)$.
$T_t$**:** $l_t(\cdot) = \frac{\alpha^t}{1+\alpha^t} l_1(\cdot) + \frac{1}{1+\alpha^t} l_2(\cdot)$ for $t \in \{-k, -(k-1) \ldots, -1, 0, 1, \cdots k\}$, where $k = \lceil \log_\alpha (2\epsilon^{-1} n^{c+1}) \rceil$.

There are several points worth noting about this scheme: (1) It achieves a result that is slightly stronger than our initial goal – it approximately maximizes each individual node's welfare, not just the sum of all nodes' welfare. (2) The computation of the trees is oblivious to the nodes' preference information. Thus, if we assume that the objective costs are common knowledge (or verifiable), this scheme is trivially a strategyproof mechanism. (3) Each tree computation involves computing lowest-cost routes for a specific objective metric. Thus, it is easily computed within the framework of BGP itself. (In the terminology of Feigenbaum *et al.* [7, 5], there is a natural BGP-based distributed algorithm for this scheme.)

We now prove a corresponding lower bound that shows that Theorem 5 is nearly optimal.

**Theorem 6.** *Let* $\epsilon > 0$ *be given. There is a family of instances of the subjective-metric routing problem, with all weights in* $[0, n^c]$ *for some constant c, such that the following property holds:*

> *Any set of routing trees that contains a* $(1+\epsilon)$*-approximately optimal path* $P_{ij}$ *for each* $i$ *must have* $\Omega(\log n/\epsilon)$ *trees.*

*(Here,* $n$ *is the number of nodes of the network.)*

### 5.1 Generalization to More Than 2 Metrics

In this section, we show that Theorems 5 and 6 generalize to the case in which there are $d > 2$ objective metrics, and an AS's subjective cost is a convex combination of these metrics.

**Theorem 7.** *Suppose that, for transit AS* $k$ *, all lengths* $l_1(k), l_2(k) \ldots, l_d(k)$ *are integers bounded by a polynomial, i.e. ,* $l_j(k) < n^c$ *for some constant c. Then, for any given* $\epsilon > 0$*, there is a set of routing trees* $T_1, T_2, \cdots T_r$ *with* $r = O(4^d [\frac{(c+1)\log n + \log(\frac{2}{\epsilon})}{\log(1+\epsilon)}]^{d-1})$ *such that:*

> *For each AS* $i$*, there is a tree* $T_{t_i}$ *such that* $c_i(T_{t_i}) \leq (1+\epsilon) c_i(P^*_{ij})$*, where* $P^*_{ij}$ *is the minimum-subjective-cost route from* $i$ *to* $j$*.*

*Further, this set of trees can be constructed in polynomial time for any constant* $d$*.*

**Theorem 8.** *Let* $\epsilon > 0$ *be given and* $d > 2$ *be given. There is a family of instances of the subjective-metric routing problem, with all weights in* $[0, n^c]$ *for some constant c, such that the following property holds:*

> *Any set of routing trees that contains a* $(1+\epsilon)$*-approximately optimal path for each* $i$ *must have* $\Omega((\frac{\log n}{d \log d + d \log(1+\epsilon)})^{d-1})$ *trees.*

*(Here,* $n$ *is the number of nodes of the network.)*

## 6 Conclusion

In this paper, we have studied classes of ordinal and cardinal preferences based on subjective costs. The subjective-cost preference model is intuitively appealing, and it is very expressive. However, our results show that, even if the costs are restricted to a very small range, unstructured subjectivity leads to intractable problems in both models: NP-completeness of the stable paths problem for ordinal preferences and APX-hardness of the minimum subjective-cost tree problem for cardinal preferences.

The root cause of these hardness results appears to be the high dimension of the space of AS preferences. Thus, it is necessary to work with models that provide a more consistent global structure. In Section 5, we consider the case in which there are two objective cost metrics, and ASes differ in the relative importance they place on the first metric. For example, ASes may agree on the latency and packet-loss rate of each node in the network but have subjective opinions about the relative importance of latency and loss rate. Thus, in this model, the space of all AS *types* is one-dimensional. We showed that it is possible to select a small number ($O(\frac{1}{\epsilon}[\log n + \log(\frac{1}{\epsilon})])$ for a $(1+\epsilon)$-approximation) of representative types such that every ASes' preferences are closely approximated by one of the representatives; then, by picking a *set* of routing trees, each of which is optimized for a specific representative type, we can guarantee each AS a route that $(1+\epsilon)$-approximately minimizes its subjective cost. Further, this scheme is easy to implement, even in the distributed-computing context: Each destination can be replaced by a small number of logical destinations, and a lowest-cost routing algorithm (*e.g.*, the Bellman-Ford algorithm) can be used for each logical destination.

It is also possible that other models that restrict the subjectivity of the costs in some way may yield positive results. For example, the nodes' subjective costs for a given transit node $k$ are random variables drawn from a specific distribution. Finding such models that are both realistic and tractable is an interesting avenue for future research.

## Acknowledgements

We thank Shay Kutten for suggesting the forbidden-set routing problem, and Hari Balakrishnan for pointers to the literature. We are grateful to Andrew Twigg and Vijay Ramachandran for pointing out errors in an earlier version of this paper.

## References

1. Varadhan, K., Govindan, R., Estrin, D.: Persistent route oscillations in inter-domain routing. Technical Report 96-631, USC/ISI (1996)
2. Griffin, T., Wilfong, G.: An analysis of BGP convergence properties. In: Proceedings of ACM SIGCOMM '99. (1999) 277–288
3. Griffin, T., Shepherd, B., Wilfong, G.: The stable paths problem and interdomain routing. IEEE/ACM Transactions on Networking **10** (2002) 232–243
4. Feamster, N., Johari, R., Balakrishnan, H.: The implications of autonomy for stable policy routing. In: Proceedings of the ACM SIGCOMM Conference (SIGCOMM '05). (2005) 25–36
5. Feigenbaum, J., Sami, R., Shenker, S.: Mechanism design for policy routing. In: Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing (PODC '04). (2004) 11–20
6. Gao, L., Rexford, J.: Stable internet routing without global coordination. IEEE/ACM Transactions on Networking **9** (2001) 681–692
7. Feigenbaum, J., Papadimitriou, C., Sami, R., Shenker, S.: A BGP-based mechanism for lowest-cost routing. Distributed Computing **18** (2005) 61–72