

Vickrey Prices and Shortest Paths: What is an edge worth?

John Hershberger

Mentor Graphics Corp.
8005 SW Boeckman Road
Wilsonville, OR 97070, USA
john_hershberger@mentor.com

Subhash Suri*

Computer Science Department
University of California
Santa Barbara, CA 93106, USA
suri@cs.ucsb.edu

Abstract

We solve a shortest path problem that is motivated by recent interest in pricing networks or other computational resources. Informally, how much is an edge in a network worth to a user who wants to send data between two nodes along a shortest path? If the network is a decentralized entity, such as the Internet, in which multiple self-interested agents own different parts of the network, then auction-based pricing seems appropriate. A celebrated result from auction theory shows that the use of Vickrey pricing motivates the owners of the network resources to bid truthfully. In Vickrey's scheme, each agent is compensated in proportion to the marginal utility he brings to the auction. In the context of shortest path routing, an edge's utility is the value by which it lowers the length of the shortest path—the difference between the shortest path lengths with and without the edge. Our problem is to compute these marginal values for all the edges of the network efficiently. The naïve method requires solving the single-source shortest path problem up to n times, for an n -node network. We show that the Vickrey prices for all the edges can be computed in the same asymptotic time complexity as one single-source shortest path problem. This solves an open problem posed by Nisan and Ronen [12].

1. Introduction

Shortest paths are fundamental in many areas of computer science, operations research, and engineering. Their applications include network and electrical routing, transportation, robot motion planning, critical path computation in scheduling, etc. In addition, shortest paths also provide

a unifying framework for many optimization problems such as knapsack, sequence alignment in molecular biology, inscribed polygon construction, and length-limited Huffman-coding, etc. (Eppstein [4] is a good reference for shortest paths and their applications.) Most complex applications of the shortest path problem, however, require more than just the calculation of a single shortest path. In some applications, the desired path might be subject to additional constraints that are hard to quantify. In others it might be useful to examine not just *the* shortest but a larger set of “short paths.” In some applications, it is desirable to see how the shortest path is influenced by various system parameters, through a “sensitivity analysis.” Our problem belongs to this last category. We wish to determine, for each edge e in a graph, what effect e 's deletion has on the shortest path between two given nodes.

Our problem is motivated by recent interest in pricing networks and computing resources, which in turn is prompted by the prominent role the Internet has come to play in our lives. One of the distinguishing characteristics of the Internet is that it involves interaction among multiple (often very many) self-interested participants. These participants (organizations, people, computers, software), called “agents” in the AI terminology, cannot be trusted to follow the rules of a protocol, especially if deviating from the protocol is beneficial to the agent. Thus, unlike traditional distributed computing protocols, a protocol for these new settings must be designed explicitly to account for willing manipulation by the users. We briefly describe an example that helps illustrate this phenomenon.

One of the most famous distributed protocols on the Internet is the Transmission Control Protocol (TCP), implemented on each host on the Internet. An important feature of this protocol is its *congestion control* mechanism. The protocol uses packet loss as an indication of network congestion, and is designed to reduce the sending host's transmission rate (using a fairly aggressive exponential back-

*Subhash Suri's research on this paper was partially supported by National Science Foundation grants CCR-9901958 and ANI-9813723.

off). It then gradually increases the transmission rate until another sign of congestion is detected. This cycle of increasing and then decreasing the transmission rate allows the protocol to discover and utilize whatever bandwidth is available between two communicating hosts, while at the same time sharing the overall resource among many such pairs.

TCP is *self-regulating*, meaning that it assumes that individual hosts will respond to congestion exactly as the designers of the protocol intended. But a self-interested host (agent) has motivation *not* to decrease its sending rate in the hope that *others* will reduce their rate, eliminating the congestion, while he can continue to enjoy the higher rate. In one extreme case, no one follows the TCP congestion rules, and the system crashes; in another extreme case, a few misbehaving users enjoy an unfair share of the network resources, while the rule-abiding majority of users suffer. Because a protocol like TCP is easily manipulated, many researchers have proposed game-theoretic and price-based mechanisms to share bandwidth and other network resources [5, 13, 19].

In this context, a natural economic question is this: how much is an edge in a network worth to a user who wants to send data between two nodes along a shortest path? If the network is a decentralized entity, such as the Internet, in which multiple self-interested agents own different parts of the network, then an auction is often the best mechanism to determine the utility of various network elements. A celebrated result from auction theory shows that the use of Vickrey pricing motivates the agents to bid truthfully. In Vickrey’s scheme, each agent is compensated in proportion to the marginal utility he brings to the auction. The insight of Vickrey is that although agents have an incentive to lie about their costs in the hope of receiving larger compensation from the network, making an edge’s payment depend only on the declarations of *other* agents eliminates this manipulative element.¹

Suppose we are interested in discovering the shortest path from node x to node y in a network G , whose links are owned by self-interested agents. We assume that agents bid on *individual* links—that is, either each agent owns at most one link, or if an agent owns multiple links, he bids on each independently. (We do not consider the setting where an agent can make strategic bids on *subsets* of links. In

¹The Vickrey mechanism is a generalization of the well-known sealed bid *second price auction*, in which an object is sold to the highest bidder, but the winner pays a price equal to the runner-up’s bid. This auction protocol is known to be truthful, in that a rational agent’s best bidding strategy is to bid his true valuation. Thus, in a distributed network where network links belongs to rational, self-interested agents, Vickrey pricing elicits truthful responses from agents, leading to economic efficiency in the network—the shortest paths use agents with lowest costs (even though actual costs are private, not public, information). This truthfulness, however, does come at a cost: the mechanism may need to subsidize the agents. For a proof of the Vickrey mechanism’s truthfulness, see [11].

such *bundle auctions*, even determining the winning bids is NP-complete, but under a restricted setting Bikhchandani et al. [2] solve the Vickrey payment problem using linear programming.) The network employs the Vickrey pricing mechanism to elicit agents’ *true* preferences (costs). The payment p^e made to an edge e is determined as follows:

$$p^e = \begin{cases} d(x, y; G \setminus e) - d(x, y; G|_{e=0}) & \text{if } e \text{ is on the shortest path} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

That is, if edge e does not belong to the shortest path in G , then its agent receives zero payment. Otherwise, the payment to e is the difference between the cost of the shortest path without e , and the cost of the shortest path assuming e is free.

This formulation is from Nisan and Ronen [12], who posed the following question: what is the computational complexity of determining all the Vickrey payments? The straightforward method requires computing the x - y shortest paths up to n times: once in G , and once in $G \setminus e$, for each e that belongs to the shortest path in G . The payment function expressed in equation (1) requires two shortest path computations, but the second term can easily be deduced from the shortest path distance $d(x, y; G)$, for each e that belongs to the shortest path. In a graph with n nodes and m edges, each shortest path can be computed in $O(n \log n + m)$ time using Dijkstra’s algorithm, if all the edges have non-negative costs, or in $O(nm)$ time using the Bellman-Ford algorithm, if the network has negative cost edges but no negative cycles [3]. Since there are at most $n - 1$ edges on the shortest path from x to y , the naïve method’s total cost for computing the payments to all the agents is $O(n^2 \log n + nm)$ for non-negative cost networks, and $O(n^2m)$ for networks with negative edge costs.

Our main result is an algorithm to compute the Vickrey payments to all the agents in essentially the same time bound as one single-source shortest path computation. Our algorithm builds two shortest path trees, one based on x and the other based on y , and computes the payment term $d(x, y; G \setminus e)$ for each edge e by combining parts of these trees. The total time complexity is $O(m + n \log n)$ if the edge costs are non-negative, and $O(nm)$ otherwise.

2. Related work

Researchers in multi-agent systems, or distributed AI, have studied cooperation and competition among “software agents.” Many of these papers use ideas from mechanism design to analyze strategies and responses of these agents in negotiations and resource allocations [14, 15, 20]. Others use market-based ideas to solve distributed computational problems [21, 23]. Researchers in networking have

proposed game-theoretic techniques to deal with congestion control in the Internet [5, 6, 13, 19].

Our paper is motivated by the algorithmic mechanism design paper of Nisan and Ronen [12]. They investigate computational complexity and algorithmic issues in mechanism design, and raise some intriguing problems. Specifically, they asked the question that forms the basis of our work: Can the payment functions of the Vickrey mechanism be computed faster than n invocations of the optimization problem? The payment function computation for the network routing problem is equivalent to the following: given a directed graph G , and two specified nodes x and y , determine, for each edge e in the graph, the effect on $d(x, y)$ of deleting e .

Our problem is related to the topic of “sensitivity analysis” in operations research. In sensitivity analysis, the goal is to determine the robustness of a solution: how much the system parameters can be perturbed before the solution changes. For instance, the sensitivity analysis of the minimum spanning tree requires computing for each edge e the amount $\delta(e)$ by which the cost of e must change before the minimum spanning tree changes; $\delta(e)$ is positive if e is part of the MST, and negative otherwise. Tarjan [22] presents an $O(m\alpha(m, n))$ time algorithm for calculating $\delta(e)$ for all edges of a graph with n nodes and m edges, where α is a functional inverse of Ackermann’s function. Tarjan also presents a similar result for performing the sensitivity analysis of a shortest path tree. In our problem, however, we are not interested in computing the cost threshold of an edge, but rather in deleting an edge, and then finding the new shortest path. To the best of our knowledge, all the known methods for this type of sensitivity analysis of shortest paths require $\Omega(m)$ work per shortest path edge [1]. Our problem also has some similarity to the k -shortest paths problem studied by Eppstein [4], but requires different techniques.

Finally, Bikhchandani et al. [2] and Schummer and Vohra [18] have considered general auction settings where the Vickrey payments correspond to *dual variables* in a linear program. Their results depend on a combinatorial condition, which they call the “agents are substitutes” condition. It turns out that the minimum spanning tree problem and the assignment problem satisfy the “agents are substitutes” condition, and therefore the Vickrey payments for those problems can be determined efficiently. However, the “agents are substitutes” condition does not hold for the shortest path problem, so the methods of [2] and [18] do not apply to our setting.

3. Shortest path preliminaries

We assume that our network is modeled by a graph $G = (V, E)$, with $|V| = n$ and $|E| = m$. Each edge $e \in E$ has an associated cost $c(e)$. We consider both directed and

undirected graphs, and present our algorithm for undirected graphs first, since some of the details are simpler. We assume that a pair of vertices u and v has at most one edge connecting them, but this is not a necessary restriction for our algorithms: the algorithms work just as well when there are multiple parallel edges joining pairs of vertices.

A path in G is a sequence of edges, such that consecutive edges share a common vertex, and each vertex is incident to at most two path edges. The total cost of a path in G is the sum of the costs of the edges on the path. The *shortest path* between two vertices a and b , denoted by $path(a, b)$, is the path joining a to b , assuming one exists, that has minimum cost. The distance between a and b , denoted $d(a, b)$, is the length of $path(a, b)$, or infinity if no path exists.²

We denote shortest paths and distances when an edge e has been removed from the graph G by the notations $path(a, b; G \setminus e)$ and $d(a, b; G \setminus e)$. The distance in the full graph $d(a, b)$ is shorthand for $d(a, b; G)$; likewise $path(a, b)$ is shorthand for $path(a, b; G)$.

There are two distinguished vertices x and y in the graph, called the *source* and the *target* vertices. The shortest path joining them is $path(x, y) = (v_1, v_2, \dots, v_k)$, where $v_1 = x$ and $v_k = y$. Recall that we want to compute, for each $i \in \{1, \dots, k-1\}$, the length of the shortest path from x to y that does not use the edge $e_i = (v_i, v_{i+1})$, which we call the *x - y distance omitting e_i* . This is precisely the term $d(x, y; G \setminus e_i)$ in the payment function.

We can analyze the shortest paths from x to y in terms of the set of edges crossing a *cut*. In later sections we will choose the cut according to the structure of the graph, but for now let us simply consider any partition of the vertex set V into two sets V_x and V_y such that $x \in V_x$ and $y \in V_y$. The set of edges crossing the cut is denoted by $E_{cut} = E(V_x, V_y)$. Each edge $(u, v) \in E_{cut}$ has $u \in V_x$ and $v \in V_y$. Any path from x to y must include at least one edge from E_{cut} . Therefore, we can express our problem as follows: For each $e_i = (v_i, v_{i+1})$ and some cut (V_x, V_y) possibly dependent on e_i , compute

$$d(x, y; G \setminus e_i) = \min_{\substack{(u,v) \in E_{cut} \\ (u,v) \neq e_i}} \left(\begin{array}{l} d(x, u; G \setminus e_i) + \\ c(u, v) + \\ d(v, y; G \setminus e_i) \end{array} \right) \quad (2)$$

In Section 4 we apply this expression to compute the x - y distance omitting each edge of $path(x, y)$ in the special case in which G is undirected and all the vertices of V lie on $path(x, y)$. In Section 5 we solve the problem for general undirected graphs, and in Section 6 we solve the problem for directed graphs.

²For convenience we assume that all path lengths are distinct, so the shortest path between any two vertices is unique. This condition is easy to enforce by a symbolic perturbation of the edge costs.

4. Payment computation in path graphs

We illustrate our ideas by solving the special case in which $path(x, y)$ includes all the vertices of V . The additional structure in this case makes it particularly easy to compute the distances in $G \setminus e$ required by equation (2). We define the cut (V_x, V_y) based on the natural partition of $path(x, y) = (v_1, v_2, \dots, v_n)$ induced by the removal of the edge $e_i = (v_i, v_{i+1})$. We choose $V_x = \{v_1, \dots, v_i\}$ and $V_y = \{v_{i+1}, \dots, v_n\}$. To simplify the notation, we define $E_i = E(V_x, V_y)$ for e_i . See Figure 1.




Figure 1. The edges of E_i cross the dashed vertical line.

For a given edge $e_i = (v_i, v_{i+1})$ to be removed, consider a cut edge $(u, v) \in E_i$. Because $path(x, u)$ is contained in V_x , $d(x, u) = d(x, u; G \setminus e_i)$. Likewise, because $path(v, y)$ is contained in V_y , $d(v, y) = d(v, y; G \setminus e_i)$. These distances are simply the lengths of the subpaths of $path(x, y)$ connecting each point to x and y . Thus equation (2) reduces to

$$d(x, y; G \setminus e_i) = \min_{\substack{(u,v) \in E_i \\ (u,v) \neq e_i}} d(x, u) + c(u, v) + d(v, y). \quad (3)$$

To compute the shortest x - y path omitting each edge of $path(x, y)$ efficiently, we evaluate equation (3) for each $e_i = (v_i, v_{i+1})$ in sequence from $i = 1$ to $n - 1$. For a given i , we minimize the quantity $d(x, u) + c(u, v) + d(v, y)$ over all $(u, v) \in E_i \setminus e_i$. But the difference between E_i and E_{i+1} is easy to compute: it consists of exactly those edges with one endpoint at v_{i+1} . To produce E_{i+1} from E_i , we add to E_i all edges whose left endpoint is v_{i+1} , and we remove from E_i all edges whose right endpoint is v_{i+1} .

To formalize this, let $left(e)$ and $right(e)$ be the indices of the endpoints of e in $path(x, y)$, with $left(e) < right(e)$. In later sections, when the endpoints of e do not necessarily lie on $path(x, y)$, we will redefine $left(e)$ and $right(e)$ to be indices such that $e \in E_i$ for all $left(e) \leq i < right(e)$. We perform the following algorithm:

Path Algorithm

1. Let L and R be k -element arrays whose elements are sets of edges, initially empty.
Let Q be a priority queue of (weight, edge) pairs, indexed by weight, initially empty.
 2. For each $e \in E \setminus path(x, y)$
 - If $left(e) < right(e)$, put e into $L[left(e)]$ and $R[right(e)]$.
 3. For $i = 1$ to $k - 1$
 - (a) For each $e = (u, v) \in L[i]$
 - Insert (w, e) into Q , with weight $w = d(x, u; G \setminus e_i) + c(u, v) + d(v, y; G \setminus e_i)$.
 - (b) Remove from Q all (w, e) pairs with $e \in R[i]$.
 - (c) Report the minimum weight in Q as the x - y distance omitting e_i .
-

At step i , Q contains the edges of $E_i \setminus e_i$, and so the x - y distance omitting e_i is correctly computed. Only the priority queue operations take non-constant time. A naïve priority queue implementation gives a running time of $O(m \log m)$.

We can improve this time complexity by using a Fibonacci heap for Q with at most $k - 1$ nodes, numbered from 2 to k [3]. The j 'th node stores the minimum-weight edge in the current cut that belongs to $R[j]$. We initialize the Fibonacci heap to have $k - 1$ nodes, each with weight ∞ . In step 3a, for each $e \in L[i]$, we compare its weight w with the weight of the heap node with index $j = right(e)$; if e has the lesser weight, we perform a DecreaseKey operation on node j and reset its edge to be e . (We do not need to do anything if e has the greater weight, because the edge stored at node j will be in the cut just as long as e .) In step 3b, we delete node i from Q —the minimization in step 3a means that this node is the representative of all the edges in $R[i]$. Thus the heap contains only a subset of the edges that would be stored in a naïve implementation of Q , but it is guaranteed to contain the minimum-weight element of the cut set.

The Fibonacci heap implementation performs $O(m)$ DecreaseKey operations, but only $O(n)$ inserts, deletes, and FindMin operations. Only the delete operation takes $O(\log n)$ time in the Fibonacci heap; the other three operations take $O(1)$ amortized time apiece. Thus, the total time spent on priority queue operations using the Fibonacci heap implementation is $O(n \log n + m)$.

5. Networks with undirected edges

In a general undirected graph, not all vertices lie on $path(x, y)$. This means that the structure of the shortest paths from x to other vertices is unrelated to the structure of the shortest paths from those other vertices to y . The *shortest path tree* with source x is the union of all the shortest paths from x to other vertices in V . Since we assume uniqueness of shortest paths, this union of paths is indeed a tree. Each vertex v has a unique parent u in the tree; $path(x, v)$ is obtained by concatenating $path(x, u)$ with the edge (u, v) . Let us denote the shortest path tree with source x by X .

We can also define a shortest path tree with sink y , which we denote by Y . This is the union of all shortest paths from vertices in V to the destination y . (Since we are assuming G is undirected, Y is identical in structure to the shortest path tree with source y . For directed graphs, this is not true.) The shortest path trees X and Y can be computed in $O(n \log n + m)$ time using Dijkstra's algorithm and Fibonacci heaps [3], or in $O(m \log n)$ time using simple data structures.

In the case of a path graph, as discussed in Section 4, $X = Y = path(x, y)$, and the removal of an edge e splits X and Y into identical components. However, for general undirected graphs, this is not true. For example, in Figure 2, the vertices in the upper branch of the graph lie in opposite components of $X \setminus e$ and $Y \setminus e$.




Figure 2. X and Y have different structures.

To compute the x - y distance omitting each edge $e_i = (v_i, v_{i+1})$ in $path(x, y) = (v_1, \dots, v_k)$, we define the cut (V_x, V_y) based on the shortest path tree X . Because $path(x, y)$ is contained in X , the removal of e_i splits X into two components; we choose the component containing x to be V_x , and the complement to be V_y . To be more specific in the determination of V_x and V_y , we assign vertices of V to *blocks* based on their position in the shortest path tree X . If we delete all the edges of $path(x, y)$ from X , the

vertices connected to v_i in the remaining forest form block B_i . If $u \in B_i$, we define $block(u) = i$. Thus for a given edge $e_i = (v_i, v_{i+1}) \in path(x, y)$, $V_x = \cup_{j=1}^i B_j$, and $V_y = \cup_{j=i+1}^k B_j$. See Figure 3.




Figure 3. Removing $path(x, y)$ from X defines blocks B_i .

Consider computing the x - y distance omitting e_i according to equation (2). For all $u \in V_x$, $path(x, u)$ is contained in V_x by definition, so we have $d(x, u) = d(x, u; G \setminus e_i)$. Because the partition of X induced by deleting e_i is not the same as the corresponding partition of Y (Figure 2), it is not obvious that $path(v, y)$ is contained in V_y for all $v \in V_y$. Nevertheless, it turns out that $path(v, y)$ does not use e_i , which is what we need.

Lemma 1 *Let v be a vertex in $V_y = \cup_{j=i+1}^k B_j$ for some $e_i = (v_i, v_{i+1})$. Then $d(v, y) = d(v, y; G \setminus e_i)$.*




Figure 4. Unlike the shaded path, $path(v, y)$ cannot include e_i , because $path(v_{i+1}, v)$ is contained in V_y .

Proof: The proof is by contradiction. Suppose that $path(v, y)$ uses the edge e_i . It must traverse e_i in the forward direction, from v_i to v_{i+1} , because the shortest path from v_{i+1} to y is fully contained in V_y and does not traverse e_i . Then $path(v, y)$ is the concatenation of $path(v, v_{i+1})$ with $path(v_{i+1}, y)$, and the first subpath contains vertex v_i (which is a vertex of V_x) in its interior (see Figure 4). On the other hand, because $v \in V_y$, the shortest path tree X shows that $path(v_{i+1}, v)$ is completely contained in V_y . Since G is undirected, $path(v, v_{i+1})$ is just the reversal of

$path(v_{i+1}, v)$. But one contains v_i and one does not, a contradiction. Therefore $path(v, y)$ does not contain e_i , and the lemma is established. ■

We have (almost) reduced the general undirected graph case to the case in which all vertices lie along $path(x, y)$. For an edge $e = (u, v) \notin path(x, y)$, we define $left(e) = block(u)$ and $right(e) = block(v)$, assuming $block(u) \leq block(v)$. This ensures that $e \in E_i$ if and only if $left(e) \leq i < right(e)$, and we can apply the algorithm of Section 4 directly. For each edge $(u, v) \in E_i$, the distance $d(x, u)$ is computed using the shortest path tree X , and the distance $d(v, y)$ is computed using Y .

6. Directed networks

When G is directed, things become more complicated. For example, the shortest path tree with source s is not the same as the shortest path tree with sink s . To get the shortest path tree Y with sink y , we must reverse the orientation of every edge in E and compute the shortest path tree with source y in this modified tree.

If G contains only edges with non-negative costs, we can compute the shortest path tree using Dijkstra's algorithm in $O(n \log n + m)$ time, as in the undirected case. However, directed graphs may contain negative-cost edges; so long as there are no negative-cost cycles, it still makes sense to compute shortest paths. If G contains negative-cost edges, Dijkstra's algorithm is not applicable, and we must use a less efficient $O(nm)$ algorithm to compute shortest path trees [3].




Figure 5. Lemma 1 is false for directed graphs.

Even after the shortest path trees X and Y are computed, computing the x - y distance omitting each edge on $path(x, y)$ is still more complicated than in the undirected case. The chief difficulty is that Lemma 1 does not hold for directed graphs. Figure 5 shows an example in which vertex v belongs to the component of $X \setminus e$ that contains y , but the shortest path $path(v, y)$ contains the edge e . Fortunately, it turns out that we can finesse our way around the failure of Lemma 1. As the following lemma shows, we do not need to minimize over all the edges in $E(V_x, V_y)$ in equation (2), and the edges that we do need don't violate Lemma 1.

Lemma 2 Let V_x and V_y be the components of X induced by removing an edge $e \in path(x, y)$. Then $path(x, y; G \setminus e)$ includes exactly one edge of $E(V_x, V_y)$.




Figure 6. u is the last vertex of P in V_x . Shaded $path(x, u)$ is contained in V_x .

Proof: Consider any path P connecting x to y in $G \setminus e$, and let u be the last vertex of P in V_x . (The path P may pass from V_x to V_y several times, but we choose the last such transition. See Figure 6.) The shortest path from x to u in G is contained in V_x , since V_x is a subtree of X containing both x and u . Therefore $path(x, u) = path(x, u; G \setminus e)$, and we can shorten P by replacing the portion of P up to u by $path(x, u)$. The only edge of $E(V_x, V_y)$ in this shorter path is the one immediately following u . It follows that the shortest path from x to y in $G \setminus e$ must contain exactly one edge of $E(V_x, V_y)$. ■

A simple corollary of this lemma is the fact that if (u, v) is the single edge of $path(x, y; G \setminus e)$ in $E(V_x, V_y)$, then $path(v, y; G \setminus e) = path(v, y)$. It follows that in the minimization of equation (2), we do not need to consider any edge (u, v) of $E(V_x, V_y)$ such that $path(v, y)$ contains any vertex of V_x . Consequently, the minimization set of equation (2) can be reduced as follows:

$$d(x, y; G \setminus e) = \min_{\substack{(u,v) \in E(V_x, V_y) \\ (u,v) \neq e \\ path(v,y) \cap V_x = \emptyset}} \left(\begin{array}{l} d(x, u; G \setminus e) + \\ c(u, v) + \\ d(v, y; G \setminus e) \end{array} \right) \quad (4)$$

To filter out edges that violate the condition on $path(v, y)$, we label each vertex $v \in V$ according to the lowest-indexed block of X that $path(v, y)$ passes through. Define $minblock(v)$ to be the smallest i such that $path(v, y)$ contains a vertex of block B_i . That is, $minblock(v) = \min_{w \in path(v,y)} block(w)$. See Figure 7. We can compute $minblock(v)$ for all vertices v in $O(n)$ time by a pre-order traversal of Y starting from y . For any edge (u, v) , $minblock(u)$ is just $\min(block(u), minblock(v))$, and the preorder traversal visits v before u .

For a directed edge $e = (u, v) \notin path(x, y)$, we define $left(e) = block(u)$ and $right(e) = minblock(v)$. With




Figure 7. $block(v) = 5$, but $minblock(v) = 3$.

these definitions of $left(e)$ and $right(e)$, we ensure that e belongs to the minimization set of equation (4) for e_i if and only if $left(e) \leq i < right(e)$, and we can apply the algorithm of Section 4. The distances $d(x, u)$ and $d(v, y)$ are available from the shortest path trees X and Y . We have established our main result:

Theorem 3 *Given a directed network G with m edges and a pair of vertices (x, y) , we can compute $d(x, y; G \setminus e)$ for each edge $e \in path(x, y)$ in total time $O(n \log n + m)$ plus the time to compute a shortest path tree in G .*

This theorem allows us to compute the Vickrey payments for all edges of a shortest path in a network, as given in equation (1), in the same asymptotic time as is needed to compute the shortest path itself.

7. Concluding remarks

With the emergence of the Internet as a global platform for communication, computation, and commerce, there is an increased need to design efficient protocols that motivate self-interested agents to cooperate. Example applications include resource allocation in computational grids [24], market-based protocols for scheduling or task allocation [21, 23], and congestion control in the Internet [5, 6, 10, 19]. One of the most celebrated results in the field of mechanism design is the Vickrey (or Vickrey-Clarke-Groves) protocol, which uses a payment scheme to motivate selfish agents to bid truthfully.

In this paper, we focused on the algorithmic aspect of computing the Vickrey payments in the context of shortest path routing in an internet, where multiple self-interested agents own portions of the network. Naïvely, computing payment functions for n agents requires n shortest path computations. Our main result shows that this computational overhead can be significantly reduced—the payments are computable in the same asymptotic time as a single shortest path tree. Our algorithm is quite simple, and uses only some elementary properties of shortest paths.

We believe our algorithm will have applications to other graph problems as well. For example, we have recently used these ideas to compute the k simple (loopless) shortest paths in the same asymptotic time as k single-source shortest path

tree computations (paper in preparation). The running time of our algorithm is an improvement by a factor of $\Omega(n)$ over the previous best results, which date back to Lawler’s and Yen’s algorithms of the early seventies [8, 25, 26].

Many interesting and challenging problems remain in the still nascent field of algorithmic mechanism design. For instance, many applications in distributed computing may require designing new mechanisms [12, 21, 24]. There are also many important problems in which computing Vickrey payments requires solving NP-complete problems [9, 16, 17]. In those cases, it would be interesting to use the techniques of approximation or randomization to design new polynomial-time mechanisms. The work of Bikhchandani et al. [2] and Schummer and Vohra [18] also suggests a promising direction for further exploration.

References

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ, 1993.
- [2] S. Bikhchandani, S. de Vries, R. Vohra, and J. Schummer. Linear Programming and Vickrey Auctions. *Proceedings of the IMA workshop on e-auctions and markets*, 2001.
- [3] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 1990.
- [4] D. Eppstein. Finding the k shortest paths. *SIAM J. Computing*, 28:652–673, 1998.
- [5] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35:1969–1985, 1999.
- [6] R. Karp, E. Koutsoupias, C. Papadimitriou, and S. Shenker. Optimization problems in congestion control. In *Proc. 41st Annu. IEEE Sympos. Found. Comput. Sci.*, 2000.
- [7] V. King. A simpler minimum spanning tree verification algorithm. *Algorithmica*, 18(2):263–270, 1997.
- [8] E. L. Lawler. A procedure for computing the K best solutions to discrete optimization problems and its application to the shortest path problem. *Management Science*, 18, pp. 401–405, 1972.
- [9] D. Lehmann, L. O’Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *Proc. ACM Conference on Electronic Commerce*, 2000.

- [10] J. K. MacKie-Mason and H. R. Varian. Pricing congestible network resources. *IEEE Journal of Selected Areas in Communications*, 1995.
- [11] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [12] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. 31st Annu. ACM Sympos. Theory Comput.*, 1999.
- [13] A. Odlyzko. A modest proposal for preventing Internet congestion. <http://www.research.att.com/~amo/doc/modest.proposal.ps>, 1997.
- [14] J. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiations Among Computers*. MIT Press, 1994.
- [15] T. Sandholm. Distributed rational decision making. In *Introduction to Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence*. MIT Press, 1999.
- [16] T. Sandholm and S. Suri. Improved algorithm for optimal winner determination in combinatorial auctions and generalizations. In *AAAI 17th National Conference on Artificial Intelligence*, 2000.
- [17] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. CABOB: A Fast Optimal Algorithm for Combinatorial Auctions. In *IJCAI 17th International Joint Conference on Artificial Intelligence 2001*.
- [18] J. Schummer and R. Vohra. Auctions for Procuring Options. Northwestern University Technical Report, 2001.
- [19] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *Telecommunications Policy*, pages 183–201, 1996.
- [20] Y. Shoham and K. Tanaka. A dynamic theory of incentives in multi-agent systems. In *Intl. Joint Conf. on Artificial Intelligence*, 1997.
- [21] W. E. Walsh and M. P. Wellman. A market protocol for decentralized task allocation. In *Proc. 3rd International Conference on Multi-Agent Systems*, 1998.
- [22] R. E. Tarjan. Sensitivity analysis of minimum spanning trees and shortest path trees. *IPL*, 14 (1), pp. 30–33, 1982.
- [23] W. E. Walsh, M. P. Wellman, P. R. Wurman, and J. K. MacKie-Mason. Auction protocols for decentralized scheduling. In *Proc. 18th International Conference on Distributed Computing Systems*, 1998.
- [24] R. Wolski, J. S. Plank, J. Brevik, and T. Bryan. G-commerce: Market formulations controlling resource allocation on the computational grid. In *IPDPS*, 2001.
- [25] J. Y. Yen. Finding the K shortest loopless paths in a network. *Management Science*, 17, pp. 712–716, 1971.
- [26] J. Y. Yen. Another algorithm for finding the K shortest loopless network paths. *Proc. of 41st Mtg. Operations Research Society of America*, 20, 1972.