Shadow Prices vs. Vickrey Prices in Multipath Routing

Parthasarathy Ramanujam, Zongpeng Li, Lisa Higham
Department of Computer Science, University of Calgary
{pkramanu, zongpeng, higham}@ucalgary.ca

***Abstract*—Shadow price and Vickrey price are two classic metrics that can be applied to measure the relative importance of links in a communication network. Each metric has been extensively investigated and enjoys important applications. We study the underlying connections between these two metrics with seemingly different definitions, under a general mathematical model of multipath multi-session multicast routing. We show that Vickrey prices provide upper-bounds for shadow prices in general, and the fine granularity version of Vickrey price, unit Vickrey price, equals exactly the maximum shadow price. We further design an efficient algorithm that computes all-link max/min shadow prices and unit Vickrey prices simultaneously, for unicast routing, reducing the complexity of a straightforward algorithm by an order of $O(|E|)$.**

Index Terms: Algorithms; Game Theory; Mathematical Programming/Optimization; Multicast; Routing; Shadow Price; Vickrey Price.

I. Introduction

How important is a link in a communication network, and how can such importance of links be algorithmically evaluated in the most efficient way? The studies in this paper center around these two basic yet fundamental questions within the context of multi-path routing of information flows. Informally, a link's "importance" is associated with its contribution towards the routing task, and depends on network parameters such as the topology of the network, capacities and costs of links, the form of applications supported by routing (*e.g.*, unicast *vs.* multicast), the location of senders/receivers and the location of the link itself. Two classic metrics for measuring link importance, shadow price [1]–[3] and Vickrey price [4]–[7], have been proposed and studied extensively, in mostly separate directions, and both enjoy a wide range of applications in fields such as network game theory, network mechanism design and congestion control [1], [3], [5], [8].

The *Vickrey price* of a link measures its "added value" to the routing task, i.e., the increase in routing cost that would occur with the removal of this link [5]–[7]. A fine granularity version, *unit Vickrey price*, is the increase in routing cost if a link's capacity is decremented by one unit. The former pertains to those where where links function or fail as a whole [5], [6], and the latter applies to scenarios where link capacities may fluctuate due to background traffic or capacity management. Conversely, one may also define the *unit Vickrey gain*, measuring cost reduction due to link capacity increment. Vickrey prices and gains help the network operator determine which links are more critical and which less so, which links are over-provisioned in capacity and which links most deserve augmentation. The application of Vickrey price goes beyond these natural interpretations. In network game theory, the Vickrey-Clarke-Groves(VCG) mechanism [4] utilizes Vickrey prices for incentive engineering. It has been shown that in various versions of optimal traffic routing [3], [8], paying a rational link its claimed cost plus its Vickrey price economically motivates a truthful report in cost.

Shadow prices, from a primal-dual perspective, are optimal dual solutions where routing is modeled in the primal program [2], [3]. By complementary slackness [2], a shadow price vector is valid if and only if there is a corresponding primal (routing flow), such that only saturated links have a non-zero shadow price and the cost of each utilized routing paths is of equal minimum cost among all paths available [3], [8]. The importance and wide range of applications of shadow prices are manifested by its multi-facet nature: from an economic perspective, shadow price is closely related to the market-clearing price of a resource that consumers are willing to pay; from an optimization perspective, shadow price represents the rate of change in the objective function upon infinitesimally small perturbations at a constraint [2]; from a geometric perspective, shadow prices can be interpreted as subgradients of the objective function along the dimension of resource provisioning changes [2]; from a game theoretic perspective, the maximum shadow price gives the opportunity cost of a resource owner who truthfully reveals her own cost; from a mechanism design perspective, shadow prices are equivalent to stabilizing taxes/tolls that enforce optimal flow routing [3], [8], [9]; from a congestion control perspective, shadow prices are valuable indications of link congestion [1], [10]. Fig. 1 and Fig. 2 illustrate the concept of Vickrey price and shadow price respectively.

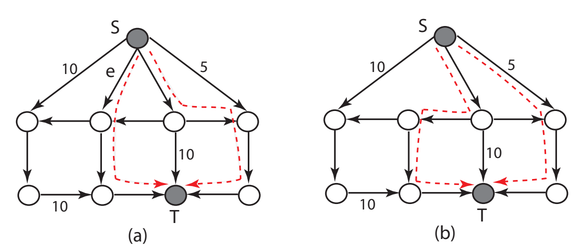

Fig. 1. (a) Optimal routing for throughput 2 from S to T in Network G (all links have capacity 1; all links have cost 1 unless labeled). Optimal routing cost is 7. (b) Optimal routing in the network with e removed, $G - e$. Optimal routing cost increases to 11. Therefore, the Vickrey price of link e is $11 - 7 = 4$.

Despite seemingly different definitions, originating from different motivations and having different applications, shadow

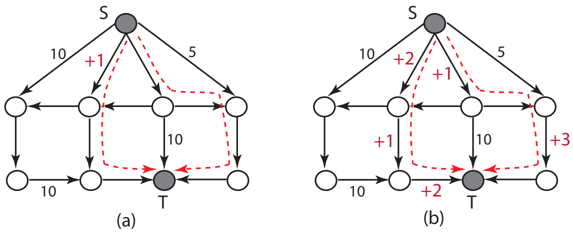

Fig. 2. Two different sets of shadow prices for the same optimal flow are shown in (a) and (b) (shadow prices listed with the prefix '+', 0 otherwise). With shadow price added to original link cost, the two paths in the flow have the same cost; costs on other paths are higher.

prices and Vickrey prices both model the relative importance of links. How closely are they related to each other? By exploiting the underlying combinatorial structure of the multipath routing problem, we provide an informative characterization of their relation. We model multipath multi-session routing using the mathematical optimization framework, and apply duality and complementary slackness to show that all shadow prices are always upper-bounded by the (unit) Vickrey price, and lower-bounded by the unit Vickrey gain. Furthermore, the maximum and minimum shadow prices of a link exactly equal its unit Vickrey price and unit Vickrey gain respectively, in a class of routing problems that have the property we refer to as *saturation-critical*. This class includes all unicast routing problems, and most multicast routing problems.

Since shadow price is not unique, it is natural to determine its range by computing the maximum and minimum shadow prices respectively. For Vickrey prices, Nisan and Ronen [11] proposed the question, *what is the computational complexity of determining all Vickrey prices*, as an important open problem in algorithmic mechanism design. We design an efficient algorithm for computing all-link unit Vickrey prices/gains (and therefore all-link max/min shadow prices), for multipath unicast routing. Our algorithm solves optimal routing once in the beginning, then proceeds to compute unit Vickrey prices for all links simultaneously, based on an all-pairs shortest path computation in the residual network of the optimal flow. It reduces the complexity of $O(|E|)$ min-cost flow computations of a straightforward algorithm to $O(1)$ min-cost flow computations, which is necessary for computing just one-link unit Vickrey price and is the best one can hope for. This advances previous work of Hershberger and Suri [6], [7] on the subject of shortest path, which is a special case of single-path unicast routing.

In the rest of the paper, we review related research in Sec. II, and present the problem model and preliminaries in Sec. III. The relations between shadow prices and Vickrey prices are established in Sec. IV, the joint algorithm is presented in Sec. V and Sec. VI concludes the paper.

II. Previous Research

The economic concept of shadow price has spawned many applications in networking research in recent years. One category is related to the "stabilizing" effect of shadow prices. By taking the primal-dual framework and applying complementary slackness or Karush-Kuhn-Tucker (KKT) conditions, one can often show that optimal flow routing becomes a stable solution (Nash/Wardrop Equilibrium) if shadow prices are applied as link taxes or tolls. Fleischer *et al.* [8] extended such a result from one session (network flows) to multiple sessions (multi-commodity flows). Bhadra *et al.* [9] and Li [3] generalized the result from unicast to multicast. Another category of applications is related to rate control. A link computes its shadow price representing the current tightness of its bandwidth supply. Each source collects shadow prices along its routing paths and uses the aggregated price as an important reference in adjusting its sending rate. Such a rate control algorithm is amenable to distributed implementations, and often has mathematically proven convergence, optimality and fairness guarantees [1]. Connections between such shadow price based rate control and congestion control algorithms implemented in the Transport Control Protocol (TCP) on the Internet have also been examined [1], [10].

The Vickrey-Clarke-Groves method [4] has been at the focal point of strategyproof mechanism design [5], [12]. Informally, the essence of a VCG mechanism is to pay each used agent its declared cost plus its added value. The added value is the added utility or reduced cost that arises due to the existence of this agent, and is referred to as the *Vickrey price* [6], [7]. An agent's dominant strategy under a VCG mechanism is to report her true utility/cost, *i.e.*, there is no economic incentive for her to lie. Therefore Vickrey prices are important in incentive engineering and strategyproof mechanism design. The major drawback of a VCG mechanism is that it runs a (potentially large) budget deficiency. Further research exists in controlling such budget deficiency and in designing strategyproof mechanisms of a non-VCG nature [5], [12].

While game theory, both in economics and in networks, has been extensively studied in the past decades, its computational aspect has attracted research attention only recently and has witnessed less progresses. In particular, we understand much better how to apply Vickrey prices than how to compute them. A straightforward algorithm for all-link Vickrey prices has a high overall complexity due to redundant computation. Hershberger and Suri [6], [7] study the problem of removing such redundancy for efficient computation of all-link Vickrey prices, for the shortest path problem. This problem is equivalent to computing the unit-Vickrey price of all links in an unit capacity network with respect to a unicast routing problem of flow rate 1. Based on graph theoretic techniques, they successfully reduced the required number of shortest path computations from $O(|V|)$ to $O(1)$. Algorithm 5.1 in this paper computes the unit Vickrey price of all links with respect to unicast routing with arbitrary throughput in networks with arbitrary link capacities. In comparison with the naive algorithm, it achieves a reduction in complexity of up to a factor of $O(|E|)$. Thus, in comparison, we study the simpler problem of computing unit Vickrey price of all the links but in the more general setting of unicast.

III. Network Model and Preliminaries

We model a computer network as a directed graph $G = \langle V_G, E_G, c_G, w_G \rangle$, where V_G is the set of nodes (computers, routers and switches), E_G is the set of edges (communication links), $c_G \in Z_+^{E_G}$ is the link capacity vector, and $w_G \in Q_+^{E_G}$ is the link cost vector. Here, Z_+ and Q_+ represent the set of positive integers and positive rational numbers respectively. We drop the subscript G when the graph being used is clear from the context. A link is represented either as e or as $\overrightarrow{uv}$.

The routing problem we consider is multi-session multicast, denoted by (G, M). M is a collection of 3-tuples, where each tuple is of the form $\langle s, R, d \rangle$. Here, $s \in V_G$ is the source node, $R \subseteq V_G$ are the receivers and $d \in Z_+$ is the required multicast throughput. A *feasible* solution to the routing problem requires a separate feasible solution for each $m \in M$ such that in addition, the combined solutions, respect the capacity constraints of the network: $\sum_{m \in M} f_m \leq c_G$, where f_m is a feasible solution for $m \in M$. The cost of a feasible solution $f = \sum_m f_m$ is denoted $|f| = \sum_e w(e)f(e)$. An *optimal* solution to the routing problem (G, M) is a feasible solution whose cost is the minimum and its cost is called the optimal routing cost, denoted $W(G, M)$.

If the routing problem (G, M) has no feasible solution, $W(G, M)$=∞ by convention. We use $G_{c(e)-\delta}$ to denote the network resulting from reducing the capacity of e by a scalar δ. If δ=$c(e)$, we simply write $G{-}e$. The *Vickrey price* of link $e$$\in$$E_G$ is therefore $W(G - e, M) - W(G, M)$. Similarly, the *unit Vickrey price* and *unit Vickrey gain* of link e are respectively $W(G_{c(e)-1}, M) - W(G, M)$ and $W(G, M) - W(G_{c(e)+1}, M)$.

We adopt the definition of *shadow price* provided within the linear programming framework [1]–[3], [8]. The LP's that model the minimum cost multi-session multicast can be found in [13] and are not presented here due to space limits. We refer to the LP that models minimum cost multi-session multicast as *the primal*. Denote the minimum cost multi-session multicast linear program for the routing problem (G, M) by $LP_p(G, M)$ and its dual by $LP_d(G, M)$. For a link e, its capacity bound will be part of the constraints in the primal and the dual variable that corresponds to this constraint, denoted as $t(e)$, is a shadow price of link e. Every optimal solution to the dual provides a shadow price to each link in the network. A vector consisting of these shadow prices is a *shadow price vector*. Following the notation of Li, we denote by $y_{m,r}(e)$ the cost share of receiver r in session m on link e.

IV. The Relation between Shadow Prices and Vickrey Prices

We first prove that any shadow price of a link provides a lower-bound for its Vickrey price. We then combine the combinatorial properties of multi-session flows with linear programming duality to prove a theorem that establishes an equality relation between the maximum shadow price and the unit Vickrey price in a class of routing problems we call 'saturation-critical', which includes most practical routing problems. We finally discuss implications and potential applications of this result.

A. Shadow Prices Provide Lower-Bounds for Vickrey Prices

Following Li [3], we say that a cost allocation together with a tax vector (y, t) *enforces* an optimal multi-session multicast flow $f^*(= \sum_m f_m^*)$, if together f^* and (y, t) satisfy: (1) *stability*, f^* is a Nash flow under (y, t), (2) *budget balance*, $(w(e) + t^*(e))f^*(e) = \sum_m \sum_{r \in R_m} y_{m,r}(e) f_m(e), \forall e$.

Li [3] showed that an optimal multicast flow f^* remains optimal even when the cost vector w is replaced by another cost vector $w' = w + t^*$, where t^* is some shadow price vector. This result can be extended to multi-session multicast in a straightforward fashion [13].

Theorem 4.1: If f^* is an optimal multi-session multicast flow and (y^*, t^*) are respectively a cost allocation vector and a shadow price vector pair to the dual of the minimum cost multi-session multicast linear program, then (y^*, t^*) enforces f^*.

We use Theorem 4.1 to show that the shadow price of a link is less than or equal to its Vickrey price.

Theorem 4.2: With respect to a given multi-session multicast routing problem (G, M), the Vickrey price of each link e is lower-bounded by $c(e)t^*(e)$, where $t^*(e)$ is any shadow price of link e.

Proof: If $t^*(e)$=0, we get the relationship trivially. Therefore, assume $t^*(e)$>0. Let f_G^* and f_{G-e}^* each be an optimal flow of rate d in G and $G{-}e$, respectively. Theorem 4.1 implies that under the cost metric $w + t^*$, the optimal flow f_G^* is still of minimum cost. In particular, we have:

$$\sum_{\overrightarrow{uv}} f_{G-e}^*(\overrightarrow{uv})(w(\overrightarrow{uv}) + t^*(\overrightarrow{uv})) \geq \sum_{\overrightarrow{uv}} f_G^*(\overrightarrow{uv})(w(\overrightarrow{uv}) + t^*(\overrightarrow{uv}))$$

which implies:

$$\begin{aligned} &\sum_{\overrightarrow{uv}} w(\overrightarrow{uv}) f_{G-e}^*(\overrightarrow{uv}) - \sum_{\overrightarrow{uv}} w(\overrightarrow{uv}) f_G^*(\overrightarrow{uv}) \\ &\quad \geq_1 \quad f_G^*(e)t^*(e) + \sum_{\overrightarrow{uv} \neq e} (f_G^*(\overrightarrow{uv}) - f_{G-e}^*(\overrightarrow{uv}))t^*(\overrightarrow{uv}) \\ &\quad =_2 \quad c(e)t^*(e) + \sum_{\overrightarrow{uv} \neq e} (f_G^*(\overrightarrow{uv}) - f_{G-e}^*(\overrightarrow{uv}))t^*(\overrightarrow{uv}) \\ &\quad \geq_3 \quad c(e)t^*(e) \end{aligned}$$

In the derivations above, $\sum_{\overrightarrow{uv}} w(\overrightarrow{uv}) f_{G-e}^*(\overrightarrow{uv}) - \sum_{\overrightarrow{uv}} w(\overrightarrow{uv}) f_G^*(\overrightarrow{uv})$ is the Vickrey price of e by definition. $\geq_1$ is by simple arithmetic manipulation and the observation that since $c_{G-e}(e) = 0$, $f_{G-e}^*(e) = 0$. $=_2$ is by complementary slackness: since we assume $t^*(e) > 0$, we know $f_G^*(e) = c_G(e)$, *i.e.*, only saturated links can have non-zero shadow prices. $\geq_3$ is also due to complementary slackness: either $t^*(\overrightarrow{uv}) = 0$, or $f_G^*(\overrightarrow{uv}) = c_G(\overrightarrow{uv}) = c_{G-e}(\overrightarrow{uv}) \geq f_{G-e}^*(\overrightarrow{uv})$, hence $\sum_{\overrightarrow{uv} \neq e} (f_G^*(\overrightarrow{uv}) - f_{G-e}^*(\overrightarrow{uv}))t^*(\overrightarrow{uv})$ is always non-negative.

□

B. The Equivalence between Max Shadow Price and Unit Vickrey Price

Similar to Theorem 4.2, it can be shown that any shadow price is upper-bounded by the unit Vickrey price. Is this bound tight? Can the maximum shadow price always reach the unit Vickrey price? We settle this question in the affirmative.

We first prove two lemmas. The first one states that if a link e is saturated in f^*, and the desired throughput d is still feasible after reducing the capacity $c(e)$ by one, then there is a new optimal flow in the reduced-capacity network that saturates e.

Lemma 4.1: If there is an optimal solution f^* to (G, M) that saturates edge e and if $(G_{c(e)-1}, M)$ is feasible, then some optimal solution to $(G_{c(e)-1}, M)$ saturates edge e.

Proof: Let f' be an optimal solution to $(G_{c(e)-1}, M)$. Let θ satisfy $c(e) - 1 = \theta f'(e) + (1 - \theta)c(e)$. Let $f'' = \theta f' + (1 - \theta)f^*$. Then, f'' saturates e and is a feasible solution to $(G_{c(e)-1}, M)$. Because $|f^*| \leq |f'|$, $|f''| \leq |f'|$. But, f' is an optimal solution to $(G_{c(e)-1}, M)$ and hence $|f'| \leq |f''|$. Therefore, $|f''| = |f'|$. We have constructed f'' which is an optimal solution to $(G_{c(e)-1}, M)$ which saturates link e. □

Denote by $G_{w(e)+\delta}$ the network resulting from increasing the cost of e by δ. The second lemma states that if the cost of an edge e is increased by a value less than its unit Vickrey price, resulting in a new graph $G_{w(e)+\delta}$, then any optimal solution saturating e in (G, M) is an optimal solution to $(G_{w(e)+\delta}, M)$.

Lemma 4.2: An optimal solution f^* to (G, M) remains optimal to $(G_{w(e)+\delta}, M)$ provided $f^*(e) = c(e)$ and $\delta \leq W(G_{c(e)-1}, M) - W(G, M)$.

Proof: $W(G, M) + c(e)\delta$ is the cost of f^* in $G_{w(e)+\delta}$. Let f' be an optimal flow in $G_{w(e)+\delta}$ and by way of contradiction assume $|f'| < W(G, M) + c(e)\delta$. Let $f'(e) = c(e) - \beta$. Let γ be the cost of f' in G. That is, $\gamma = W(G_{w(e)+\delta}, M) - (c(e) - \beta)\delta$. Observe that f' is a feasible solution to $(G_{c(e)-\beta}, M)$ and its cost remains γ. So:

$$W(G_{c(e)-\beta}, M) \leq \gamma \qquad (1)$$

Since the cost of the optimal solution is a convex function of $c(\vec{uv})$, for any $\vec{uv} \in E_G$, we have $W(G_{c(e)-\beta}, M) \geq W(G, M) + \beta(W(G_{c(e)-1}, M) - W(G, M))$ Therefore,

$$W(G_{c(e)-\beta}, M) \geq W(G, M) + \beta\delta \qquad (2)$$

Combining equations 1 and 2, $W(G, M) + c(e)\delta \leq \gamma + (c(e) - \beta)\delta = |f'|$. However, this contradicts the assumption that $|f'| < W(G, M) + c(e)\delta$. □

We are now ready to present and prove the following equality relation between shadow prices and Vickrey prices:

Theorem 4.3: Let (G, M) be any multi-session multicast problem. For any edge $e \in E_G$, if there is an optimal solution to (G, M) that saturates e, then the unit Vickrey price of link e equals the maximum shadow price at e.

Proof: Let f^* be an optimal solution to (G, M) such that $f^*(e) = c(e)$. Let, δ be the unit Vickrey price of edge e, that is, $\delta = W(G_{c(e)-1}, M) - W(G, M)$. By Lemma 4.2, f^* is also an optimal solution to (G', M), where $G' = G_{w(e)+\delta}$. For clarity, denote by f^*_G the flow f^* in G and by $f^*_{G'}$ the flow f^* in G'. Hence,

$$\begin{aligned} W(G', M) &= W(G, M) + c(e)\delta \\ &= W(G_{c(e)-1}, M) + (c(e) - 1)\delta \end{aligned}$$

Hence, any optimal solution to $(G_{c(e)-1}, M)$ that saturates edge e is also optimal to $W(G', M)$. Lemma 4.1 guarantees that $(G_{c(e)-1}, M)$ has an optimal solution that saturates e. It follows from this that the unit Vickrey price of edge e with respect to (G', M) is 0. Thus, we know from Theorem 4.2 that, in every shadow price vector to $(G_{w(e)+\delta}, M)$, the entry at e is 0.

Let $y^*_{G'}$, a cost allocation vector, and $t^*_{G'}$, a shadow price vector, be part of the same optimal solution, $\mathcal{S}'$, to $LP_d(G', M)$. By complementary slackness, $t^*_{G'}(\vec{uv})(c(\vec{uv}) - f^*_{G'}(\vec{uv})) = 0$, $\forall\ \vec{uv}$. Define t^*_G such that $t^*_G(\vec{uv}) = t^*_{G'}(\vec{uv})$ $\forall\ \vec{uv} \neq e$, $t^*_G(e) = \delta$. We claim that t^*_G is a shadow price vector with respect to (G, M). Consider $\mathcal{S}$ obtained by replacing $t^*_{G'}$ in $\mathcal{S}'$ with t^*_G. Since $\mathcal{S}'$ is in complementary slackness with $f^*_{G'}$, $\mathcal{S}$ will also be in complementary slackness with $f^*_{G'}$, and hence f^*_G, if the conditions involving t^*_G are satisfied. As $t^*_{G'}$ is part of a feasible solution to $LP_d(G', M)$, at each $\vec{uv}$, $t^*_{G'}$ satisfies $y^*_{G'}(\vec{uv}) \leq w_{G'}(\vec{uv}) + t^*_{G'}(\vec{uv})$. Further note that, $\forall\ \vec{uv}$, $w_G(\vec{uv}) + t^*_G(\vec{uv}) = w_{G'}(\vec{uv}) + t^*_{G'}(\vec{uv})$. Therefore, $\forall\ \vec{uv}$, t^*_G satisfies $y^*_G(\vec{uv}) \leq w_G(\vec{uv}) + t^*_G(\vec{uv})$. By construction, at every edge except e, t^*_G satisfies complementary slackness with f^*_G i.e. $t^*_G(\vec{uv})(c(\vec{uv}) - f^*_G(\vec{uv})) = 0$, $\forall\ \vec{uv} \neq e$. We also know that $f^*_G(e) = c(e)$. Hence, $t^*_G(e)(c(e) - f^*_G(e)) = 0$. So $\mathcal{S}$ and f^*_G satisfy complementary slackness at every link, and t^*_G is part of an optimal solution to $LP_d(G, M)$. Furthermore, $t^*_G(e) = \delta$. Thus there exists one shadow price vector for which the unit Vickrey price equals the shadow price of edge e. Since unit Vickrey price upper-bounds all shadow prices, we conclude that the unit Vickrey price and the maximum shadow price are equal at e. □

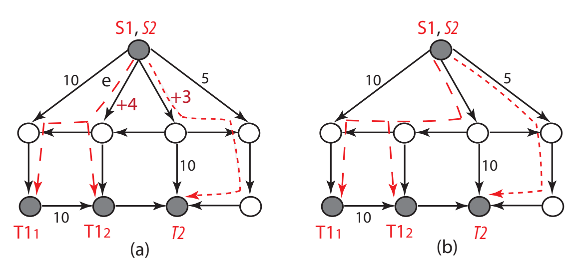

Fig. 3. (a) Optimal routing for two sessions in network G (same as in Fig. 1): tree multicast with source $S1$, receivers $T1_1$ and $T1_2$ and throughput 1; unicast with source $S2$, receiver $T2$ and throughput 1. Optimal routing cost is 8 and max shadow price of edge e is 4. (b) Min-cost routing for the same two sessions in $G - e$. Optimal routing cost in $G - e$ is 12 and hence unit Vickrey price is 4.

Fig. 3 illustrates the equality relation in an example network. We next define a *saturation-critical* routing problem as one in which each link e satisfies either (a) the unit Vickrey price and the max shadow price are trivially equal at e because they are both zero, or (b) there exists an optimal flow f^* such that $f^*(e) = c(e)$. Then following Theorem 4.3, we have:

Corollary 4.1: If (G, M) is saturation-critical, then $\forall e \in E_G$, the maximum shadow price equals the unit Vickrey price.

Due to the totally unimodular property of a node-link incidence matrix, all unicast problems always have integral optimal routing solutions [2]. For multicast, most randomly constructed networks satisfy this property. Contrived exceptions have been identified for network-coded multicast [14]; the question remains open for multicast without network coding. All routing problems for which an integral optimal solution exists satisfy the saturation-critical property. Therefore, all unicast routing problems and most multicast routing problems are saturation-critical and hence the equality between maximum shadow price and unit Vickrey price holds.

C. The Relation between Shadow Prices and Vickrey Gain

The unit Vickrey gain problem arises naturally when a network operator wishes to answer questions such as: *how much cost can be saved by augmenting a link e with an extra unit capacity*, or *which link should be chosen to furnish with extra bandwidth*? Using the same techniques above, one can show that a corresponding equality exists between the unit Vickrey gain and the min shadow price. The only difference is that this relation is true for all routing problems and not just saturation-critical routing problems. So, the range of shadow prices at a link with respect to a saturation-critical routing problem falls exactly between the unit Vickrey gain and the unit Vickrey price.

V. Algorithm Design

Algorithm 5.1 takes as input a single-session unicast problem (G, U) ($U = \langle s, r, d \rangle$) with integral inputs and outputs the unit Vickrey price and unit Vickrey gain (and therefore implicitly, by Theorem 4.2, all link shadow price ranges) for all links simultaneously. Denote the unit Vickrey price and unit Vickrey gain of link $\overrightarrow{uv}$ by UVP($\overrightarrow{uv}$) and UVG($\overrightarrow{uv}$) respectively.

Algorithm 5.1

1) Compute f^*, an integral optimal flow to (G, U).
2) Compute the residual network $G_R = G \ominus f^*$
3) Run the all pair shortest path algorithm in G_R under metric $w_{G \ominus f^*}$. Let $l_{u,v}$ denote the shortest path length from u to v.
4) For all $\overrightarrow{uv}$ such that $f^*(\overrightarrow{uv}) \neq c(\overrightarrow{uv})$, output:
   - UVP($\overrightarrow{uv}$) $=$ UVG($\overrightarrow{uv}$) $= 0$
5) For all $\overrightarrow{uv}$ such that $f^*(\overrightarrow{uv}) = c(\overrightarrow{uv})$, output:
   - UVP($\overrightarrow{uv}$) $= l_{u,v} - w(\overrightarrow{uv})$
   - UVG($\overrightarrow{uv}$) $= -l_{v,u} - w(\overrightarrow{uv})$

The complexity of our algorithm is dominated by a min-cost flow computation followed by an all-pair shortest path computation. The algorithmic complexity of the later is believed to be no worse than of the former. So, we conclude that the complexity of our algorithm is equivalent to $O(1)$ min-cost flow computations [13]. This contrasts the $(2|E|+1)$ min-cost flow computations required by a naive algorithm that computes the values by definition. It also advances previous work by Hershberger and Suri [6], [7] that provides a reduction by a factor of $O(|V|)$, for all-link Vickrey prices in the shortest path problem (a special case of single-path unicast routing).

VI. Conclusions

Our contributions in this paper are: (1) a lower-bounded relation between shadow prices and Vickrey prices, (2) an equality relation between the max/min shadow price and the unit Vickrey price/gain, and (3) a new algorithm (without proof) for computing all-link unit Vickrey prices and gains for unicast routing. Two directions are open for future research: efficient algorithm design for computing all-link Vickrey prices in unicast networks and efficient algorithm design for computing all-link (unit) Vickrey prices in multicast networks.

References

[1] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of the Operations Research Society*, vol. 49, no. 3, pp. 237–252, 1998.

[2] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Mineola, NY: Dover Publications, 1998.

[3] Z. Li, "Min-Cost Multicast of Selfish Information Flows," in *Proceedings of IEEE INFOCOM*, 2007.

[4] W. Vickrey, "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.

[5] W. Wang, X.-Y. Li, and Y. Wang, "Truthful Multicast in Selfish Wireless Networks," in *Proceedings of ACM MobiCom*, 2004.

[6] J. Hershberger and S. Suri, "Vickrey Prices and Shortest Paths: What is an Edge Worth?" in *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, 2001.

[7] ——, "Erratum to 'Vickrey Prices and Shortest Paths: What is an Edge Worth?'," in *Proc. of the 45th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2003.

[8] L. Fleischer, K. Jain, and M. Mahdian, "Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games," in *Proc. of the 45th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2004.

[9] S. Bhadra, S. Shakkottai, and P. Gupta, "Min-cost Selfish Multicast with Network Coding," *IEEE Transactions on Information Theory*, vol. 52, p. 11, 2006.

[10] S. H. Low, "A Duality Model of TCP and Queue Management Algorithms," *IEEE/ACM Transactions on Networking*, vol. 11, no. 4, pp. 525–536, 2003.

[11] N. Nisan and A. Ronen, "Algorithmic Mechanism Design," in *Proc. of the 31st ACM Symposium on the Theory of Computing (STOC)*, 1999.

[12] M.-Y. Kao, X.-Y. Li, and W. Wang, "Towards Truthful Mechanisms for Binary Demand Games: A General Framework," in *Proceedings of ACM Conference on Electronic Commerce (EC)*, 2005.

[13] P. Ramanujam, Z. Li, and L. Higham, "Shadow Prices vs. Vickrey Prices in Multipath Routing: Respective Significance, Underlying Connections and Joint Algorithms," Department of Computer Science, University of Calgary, Canada, Tech. Rep., January 2009.

[14] Z. Li, B. Li, and L. C. Lau, "On Achieving Optimal Multicast Throughput in Undirected Networks," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2467–2485, June 2006.