Shortest Paths with Bundles and Non-additive Weights Is Hard

Paul W. Goldberg* and Antony McCabe

Department of Computer Science, University of Liverpool, U.K.
{P.W.Goldberg,A.McCabe}@liverpool.ac.uk

Abstract. In a standard path auction, all of the edges in a graph are sold as separate entities, each edge having a single cost. We consider a generalisation in which a graph is partitioned and each subset of edges has a unique owner. We show that if the owner is allowed to apply a non-additive pricing structure then the winner determination problem becomes NP-hard (in contrast with the quadratic time algorithm for the standard additive pricing model). We show that this holds even if the owners have subsets of only 2 edges. For subadditive pricing (e.g. *volume discounts*), there is a trivial approximation ratio of the size of the largest subset. Where the size of the subsets is unbounded then we show that approximation to within a $\Omega(\log n)$ factor is hard. For the superadditive case we show that approximation with a factor of n^ϵ for any $\epsilon > 0$ is hard even when the subsets are of size at most 2.

1 Introduction

One of the most commonly studied types of *set-system* procurement auction is the path auction (e.g. [2,21,9,16]). In a typical path auction each seller, or *agent*, is represented by an edge in a graph and the *feasible sets* (the sets of sellers that are suitable to the buyer) are exactly those that contain a path between two specified vertices of the graph. This models a number of reasonable settings, such as routing over the internet or in a transport network.

In the standard setting, each of the edges is considered as if it were a separate entity and takes part in the auction as such. However, it seems reasonable to assume that multiple edges may actually be controlled by a single entity. In the examples of Internet routing or transport networks, it seems particularly likely that this assumption will hold. Therefore, we propose a model that incorporates the idea that some entity may control the pricing of a set of individual edges, and we allow them to price *bundles* of edges in a non-additive way. It is quite common, in economic and related literature, to study the concept of 'volume discounts' or 'bundling services' (see, e.g., [20,19]) — when the price charged for a collection of goods or services is lower than the sum of the individual prices. In many situations, producing large quantities of some commodity is more efficient

* Supported by EPSRC Grant EP/G069239/1 "Efficient Decentralised Approaches in Algorithmic Game Theory".

P.G. Spirakis and M. Serna (Eds.): CIAC 2013, LNCS 7878, pp. 264–275, 2013.

than producing smaller quantities and hence has lower unit cost. More directly, two reasonable motivations that may apply here are, firstly, that a path auction may actually be representative of something other than a network layout. It could be a model, for instance, to describe the components needed to produce some complicated commodity, such as an electronic device, where edges represent some component, and a path through the graph represents a collection of components that are sufficient to build the device. Alternatively, it is reasonable to assume, in general, that there may be some overhead to making each single purchase (such as the overhead of preparing bids, performing background and credit checks, or drawing up contracts). In this setting, the overhead may be substantial for selling a single item but could be much reduced for subsequent items. With this motivation it seems reasonable to allow a seller to alter their prices depending on the total number of edges they sell.

If we take volume discounts (or bundling) to an extreme level, we have a situation where a single price may be charged in order to purchase access to any or all of the edges owned by some entity. This can be particularly attractive in cases where there is surplus capacity on the network, such as at off-peak times, and it is the aim of the seller to create demand for their service that may otherwise not exist. A type of path auction where agents can own multiple edges in a graph was studied by Du et al. [7]. They show that allowing agents to manipulate which edges they declare ownership of may have disastrous effects on the payments that are made (more precisely, an exponential *frugality ratio* [16,18]). Hence we do not consider this, and simply assume that information on the ownership of edges is public knowledge.

Perhaps the most commonly seen auction mechanism is the Vickrey-Clarke-Groves (VCG) mechanism [22,5,13]. While this can be applied to any set-system auction, it requires that an optimal feasible set be chosen as 'winners'. However, this may not always be possible in polynomial time and this intractability is generally seen as extremely undesirable. There has also been attention in describing polynomial-time mechanisms, based on approximation algorithms, for a variety of settings [1,8,3] in which finding the optimal set (and hence running VCG) cannot be computed in polynomial time. Shortest path problems can be solved efficiently (in quadratic time [6]) and it is the aim of this paper to consider the complexity of finding an optimal solution when we apply bundling (or 'volume discounts'). Unfortunately, our results are largely negative. We find that even the most basic form of discount that we consider — a simple "buy at least x edges and get a discount of d" scheme — results in the winner-determination problem being NP-hard. This precludes running the VCG mechanism in polynomial time, which makes its use undesirable. We note that where the discounts are small relative to the costs then just ignoring discounts would be close to optimal. For the case of large discounts we show limits on the approximation ratio that may be achieved (subject to certain complexity theoretic assumptions).

Instead of having only negative 'discounts' we also consider positive values, which we call *volume supplements*. While superadditive pricing may not seem as readily justifiable as the subadditive case, it is worth studying for completeness.

We are able to show that the winner determination problem for superadditive pricing is also NP-hard, and it is also hard to find an approximate solution which is within a factor of n^ϵ of the optimal for any $\epsilon > 0$ where the supplement values are large in relation to the weights.

2 Preliminaries

Our model involves finding the shortest path between two specified vertices of a graph, taking into account the volume discounts that may be offered by the edge owners. In order to represent 'ownership' we partition the edges of a graph into disjoint sets, which we call *bundles*. Each bundle has a *discount vector*, which specifies a *discount value* depending on the number of edges in that bundle which are in the chosen path.

To compute a *discounted-weight* for a given path we add the weight of all of the edges in the path to the (negative) discount values for each of the bundles. Hence the aim of the problem is to discover the path between the specified vertices with the lowest discounted-weight, and we will see this is hard via a reduction from MINIMUM SET COVER, a well-known NP-hard problem [11].

Name SHORTEST PATH WITH DISCOUNTS
Instance A graph $G = (V, E)$, positive weight function w on edges, two distinct vertices s, t, a collection, $Z = \{Z_1, \ldots, Z_m\}$ of subsets (or *bundles*) of E ($Z_i \subseteq E$) such that $\forall e \in E$, e occurs in exactly one Z_i and a set of discount vectors $\mathcal{D} = \{\boldsymbol{d}^1, \ldots, \boldsymbol{d}^m\}$ (one for each Z_i) such that $d^i_j \leq 0$ for $j \in \{0, \ldots, |Z_i|\}$ (observe that, for ease of notation, this vector includes a discount value for zero edges).
Output The subset $P \subseteq E$ with minimum discounted-weight, given by $\sum_{e \in P} w(e) + \sum_{Z_i \in Z} d^i_{|Z_i \cap P|}$, that contains a path from s to t.

Name EXACT COVER BY 3-SETS
Instance A finite set X containing exactly $3n$ elements; a collection, C, of subsets of X each of which contains exactly 3 elements.
Output Does C contain an exact cover for X, i.e. a sub-collection of 3-element sets $D \subseteq C$ such that each element of X occurs in exactly one subset in D?

Name MINIMUM SET COVER
Instance A finite set X and a collection, C, of subsets of X.
Output What is the size of the minimum cover of X using C?

3 Complexity

We firstly show a simple reduction, from MINIMUM SET COVER, which shows hardness and also preserves approximation. Informally, we consider a multigraph on a line and show how the sets that 'cover' elements of a ground set (in a set cover) can be simulated with bundles of edges that cover gaps in the line which represents the ground set. The discounts are arranged such that the cost of every non-empty bundle of edges is 1, regardless of the specific number of edges used.

Lemma 1. MINIMUM SET COVER *is polynomial-time reducible to* SHORTEST PATH WITH DISCOUNTS.

Proof. Taking an instance I of MINIMUM SET COVER, let $m = |X|$ and let $n = |C|$. Build an instance $I' = (G, Z)$ of SHORTEST PATH WITH DISCOUNTS as follows. (We present G as a multigraph, although we will see that this assumption may be removed later on.) Let $n = |C|$ and let $m = |X|$. Create $m + 1$ vertices, labeled $V_1, \ldots, V_{m+1}$, and let $s = V_1$ and $t = V_{m+1}$. For each $C_i \in C$, for all $j \in C_i$, add an edge $e_{i,j} = (V_j, V_{j+1})$ and add all of the $e_{i,j}$ edges to bundle Z_i. Let $w(e) = 1$ for all $e \in E$ be the weight function, and let the discount vectors be given by $\boldsymbol{d^i} = (0, 0, -1, -2, \ldots, 1 - |Z_i|)$. Hence for a path P containing a non-zero number of edges from some bundle Z_i then $w(Z_i \cap P) + d^i_{|Z_I \cap P|} = 1$, so there is a contribution of exactly 1 to the discounted-weight for each included Z_i. Therefore, for any set of bundles $S \subseteq Z$ containing a path between s and t, the discounted-weight may be given by $|S|$.

An example of this reduction is shown in Fig. 1 and Fig. 2. Fig. 1 describes a MINIMUM SET COVER instance, for which the minimum cover is of size 2 (given by the set $D = \{C_1, C_4\}$). In the corresponding instance of SHORTEST PATH WITH DISCOUNTS in Fig. 2, the lowest discounted-weight path is of weight 2, consisting of the path $P = \{e_{1,1}, e_{1,2}, e_{1,5}, e_{4,3}, e_{4,4}, e_{4,6}\}$ (shown as dashed lines), and hence the bundles $S = \{Z_1, Z_4\}$.

$X = \{1, 2, 3, 4, 5, 6\}$				
$C_1 = \{1, 2, 5\}$	$C_2 = \{1, 3, 4\}$	$C_3 = \{2, 3, 6\}$	$C_4 = \{3, 4, 6\}$	$C_5 = \{4, 5, 6\}$

Fig. 1. Example of MINIMUM SET COVER

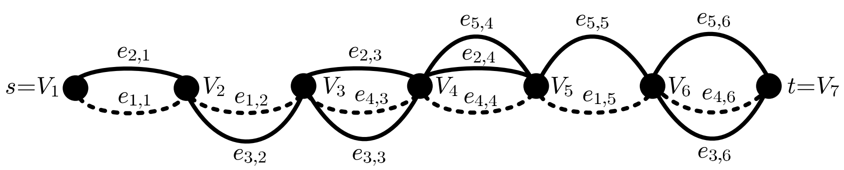

$Z_1 = \{e_{1,1}, e_{1,2}, e_{1,5}\}$, $Z_2 = \{e_{2,1}, e_{2,3}, e_{2,4}\}$, $Z_3 = \{e_{3,2}, e_{3,3}, e_{3,6}\}$,
$Z_4 = \{e_{4,3}, e_{4,4}, e_{4,6}\}$, $Z_5 = \{e_{5,4}, e_{5,5}, e_{5,6}\}$

Fig. 2. SHORTEST PATH WITH DISCOUNTS construction from Fig. 1

Returning to the proof, a set S contains a path between s and t if and only if it contains an edge (V_j, V_{j+1}) for all $j \in \{1, \ldots, m\}$. Hence, where S contains such a path, let $D \subseteq C$ be the set containing C_i if and only if $Z_i \in S$. Hence D contains a cover of X and $|D| \leq |S|$ — as an edge from (V_j, V_{j+1}) in some $Z_i \in S$ implies an element j in some $C_i \in D$. Similarly, where D contains a cover of X, then the corresponding set S contains a path from s to t and $|S| \leq |D|$. Following this translation, we have $|S| = |D|$ and the size of the minimum cover of X can be computed by finding a path P with lowest discounted weight (and hence the number of bundles in S which contain the edges in P). □

The reduction in Lemma 1 gives an instance of SHORTEST PATH WITH DISCOUNTS that is typically a multigraph. This can be amended to construct a simple graph, by replacing each edge with two edges in series (adding one new vertex shared by each pair of edges).

Lemma 1 shows that SHORTEST PATH WITH DISCOUNTS is NP-hard, in general, provided that the cardinality of the bundles is at least 3 (as a generalisation of EXACT COVER BY 3-SETS). As bundles of size 1 correspond to a polynomial-time computable shortest-path [6], we now consider the remaining case where the bundles are of size at most 2 — and show that this also results in an NP-hard problem, even when the discount values are arbitrarily small.

In Fig. 3 we see an example of the proposed reduction — Fig. 1 is encoded as an instance of SHORTEST PATH WITH DISCOUNTS; where the shortest discounted path is given by $P = \{e'_{1,1}, e'_{1,2}, e'_{1,5}, e_{2,0}, e'_{2,0}, e_{2,x}, e_{3,0}, e'_{3,0}, e_{3,x}, e'_{4,3}, e'_{4,4}, e'_{4,6}, e_{5,0}, e'_{5,0}, e_{5,x}, e_{1,1}, e_{1,2}, e_{4,3}, e_{4,4}, e_{1,5}, e_{4,6}\}$, which is marked as a dashed line. Observe that Fig. 1 is a 'yes' instance of EXACT COVER BY 3-SETS and a path of discounted-weight $21 + 9d$ exists only because just two 'upper' paths between V_1 and V_6 (containing $e'_{1,...}$ and $e'_{4,...}$) are sufficient to give a discount on all edges in a path from V_6 to t. A 'no' instance would require three or more upper paths, which must give a discounted-weight greater than $21 + 9d$.

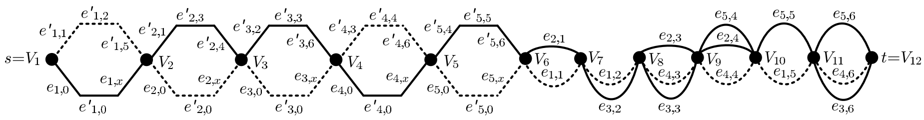

$$
\begin{aligned}
Z = \{&\{e_{1,1}, e'_{1,1}\}, \{e_{1,2}, e'_{1,2}\}, \{e_{1,5}, e'_{1,5}\}, \{e_{1,0}, e'_{1,0}\}, \{e_{1,x}\}, \{e_{2,1}, e'_{2,1}\}, \{e_{2,3}, e'_{2,3}\}, \\
&\{e_{2,4}, e'_{2,4}\}, \{e_{2,0}, e'_{2,0}\}, \{e_{2,x}\}, \{e_{3,2}, e'_{3,2}\}, \{e_{3,3}, e'_{3,3}\}, \{e_{3,6}, e'_{3,6}\}, \{e_{3,0}, e'_{3,0}\}, \\
&\{e_{3,x}\}, \{e_{4,3}, e'_{4,3}\}, \{e_{4,4}, e'_{4,4}\}, \{e_{4,6}, e'_{4,6}\}, \{e_{4,0}, e'_{4,0}\}, \{e_{4,x}\}, \{e_{5,4}, e'_{5,4}\}, \\
&\{e_{5,5}, e'_{5,5}\}, \{e_{5,6}, e'_{5,6}\}, \{e_{5,0}, e'_{5,0}\}, \{e_{5,x}\}\}
\end{aligned}
$$

Fig. 3. SHORTEST PATH WITH DISCOUNTS translated from Fig. 1 as described in Theorem 1. The dashed line marks an optimal solution.

Theorem 1. SHORTEST PATH WITH DISCOUNTS *is NP-hard even when each bundle has size at most* 2 *and there is a single small discount value* $d < 0$.

Proof. Let $I = (X, C)$ be an instance of EXACT COVER BY 3-SETS, and build an instance $I' = (G, Z)$ of SHORTEST PATH WITH DISCOUNTS as follows. Let $n = |C|$ and let $m = |X|$. Create $n + m + 1$ new vertices, labeled $V_1, \ldots, V_{n+m+1}$, and let $s = V_1$ and $t = V_{n+m+1}$. Let $w(e) = 1$ for every edge $e \in G$.

For each of the subcollections $C_i \in C$, let the elements in C_i be given by $\{x, y, z\}$. Insert three edges in series $e'_{i,x}$,$e'_{i,y}$,$e'_{i,z}$ making an 'upper' path from V_i to V_{i+1} in G (adding unlabelled vertices, as needed). Next, insert three additional edges $e_{i,x} = (V_{n+x}, V_{n+x+1})$,$e_{i,y} = (V_{n+y}, V_{n+y+1})$,$e_{i,z} = (V_{n+z}, V_{n+z+1})$. Now create three bundles $Z_{i,x} = \{e'_{i,x}, e_{i,x}\}$, $Z_{i,y} = \{e'_{i,y}, e_{i,y}\}$, and $Z_{i,z} = \{e'_{i,z}, e_{i,z}\}$ and add these to Z.

For each $i \in \{1, \dots, n\}$ create a second 'lower' path from V_i to V_{i+1} consisting of three new edges, in series, and add two new bundles to Z, one which contains two of these three new edges and another bundle containing the remaining edge. Let d be some small value, $-1 \leq d < 0$ for a discount parameter and let the discount vectors be fixed, $\boldsymbol{d}^j = (0, 0, d)$ for all $j \in Z$. (The purpose of the 'lower' paths are to create an alternative path from each V_i to V_{i+1} which has a discounted-weight of $3 + d$ — as two of the edges are in the same bundle, exactly one will be discounted.)

Taking $P \subseteq E$ to be the lowest discounted-weight path between s and t, let $S \subseteq Z$ be the minimal set of bundles that contain all of the edges in P. We can assume that all of the edges in P between V_{n+1} and V_{n+m+1} are 'discounted' edges (i.e. they are in some bundle with another edge that is also in P). This is without loss of generality, as for any path P containing an undiscounted edge e, there is another path P', having the same discounted-weight, that includes a discounted edge e' in place of e. To verify this, we know that edge e is in a bundle with another edge e'. As e is not discounted, we know that e' is not in the path. Hence, there is some 'lower' path chosen between V_1 and V_{n+1} that avoids e' with discounted-weight $3+d$. Create a path P' by including the 'upper' path that includes e', which increases the discounted-weight by $-d$ on this path, but also decreases the discounted-weight of edge e by $-d$, hence paths P and P' have the same discounted-weight.

Observe that any minimal path from s to t contains exactly $3n + m$ edges. Furthermore, we have seen that there is a minimum path P where the m edges from V_{n+1} to V_{n+m+1} are all discounted. Hence, we are interested in how many of the upper paths must be selected between V_1 and V_{n+1}, as these add to the weight of the solution when compared to the lower paths. Assume a collection D that contains a minimum cover of X; now consider a set of bundles S such that for every $C_i \in D$ all of $Z_{i,x}, Z_{i,y}, Z_{i,z}$ are present in S. Observe that S contains a path from V_{n+1} and V_{n+m+1}, or else D does not cover every element in X. Furthermore, S contains an upper path between two vertices in $V_i, \dots, V_{i+1}$ only if D contains the corresponding set C_i. Let P be the path obtained from all the bundles in S, as well as 'lower' paths between any vertices in $\{V_1, \dots, V_{n+1}\}$ that were not connected by S. We can then determine that the total discounted-weight of path P may be given by $3n+m+nd+md-|D|d$ (discounts are applied to m edges from V_{n+1} to V_{n+m+1}, and $n - |D|$ edges in the lower paths from V_1 to V_{n+1}). We can see that P is, in fact, a minimum discounted-weight path, as any lower weight path must have more lower paths chosen, yet still discounts all of the edges between V_{n+1} and V_{n+m+1}, which would imply a cover of X exists which is smaller than D, giving a contradiction.

We now claim that the path from s to t with lowest discounted-weight will have a weight of $3n + m + nd + md - (m/3)d$ if and only if C contains an exact cover for X, i.e. there is a sub-collection D of 3-element sets $D = (D_1, ..., D_n)$

such that each element of X occurs in exactly one subset in D. Observe that if C contains an exact cover, there is a set D such that $|D| = m/3$ and every element in X is contained in some C_i in D. Now consider the path P that is created, as described, by augmenting the set of bundles S (by adding lower paths, where needed). We have seen that the discounted-weight of P may be given by $3n+m+nd+md-|D|d$; hence where an exact cover exists then P has weight of $3n+m+nd+md-|m/3|d$. If no exact cover exists then the discounted-weight of P is at least $3n+m+nd+md-(|m/3|+1)d$.

Hence we have the decision problem "Is there a path in I' with discounted-weight of $3n+m+nd+md-(m/3)d$?" which has a 'yes' answer if and only if instance I contains an exact cover by 3-sets, which is NP-hard to compute. As computing a minimum discounted-weight path must answer this question (we have seen that no lower weight path can exist), then we see that computing the minimum discounted-weight path is also NP-hard. □

Thus we have seen NP-hardness in the case where one edge in a bundle is discounted, and one is not. If we assume a slightly more general scheme, in which a discount is only applied after x edges are purchased, it is a simple observation to extend this hardness result. In the construction given in Theorem 1, simply simulate every edge $e \in G$ by $x-1$ edges in series, which are contained in the same bundle. Hence, in this new graph, selecting any two of these new simulated edges will trigger the discount value.

It is also worth noting, again, that although this reduction was presented for simplicity as a multigraph, it would also work as a simple graph. All of the edges between V_{n+1} to V_{n+m+1} could be replaced by two edges, in series, having a new undiscounted edge added to the original edge. This would simply add a constant weight of m to every possible minimal path, and hence the complexity remains unchanged.

4 Inapproximability Results

If the discount values are small, relative to the weights, then simply ignoring the discount values and computing a shortest path would give a good polynomial-time approximation of the optimal solution. (Assuming a maximum weight of 1 and a discount value of d, then $1/(1+d)$ is an upper bound on the approximation ratio.) However, if we assume larger discount values, we can see inapproximability results, from Lemma 1. If we assume an unweighted case ($\forall e \in E, w(e) = 1$) and allow large discount values, then we have a scenario in which the sets of edges can be thought of as bundled together, where a single price applies to all non-empty subsets of a bundle.

We now discuss the implications of Lemma 1 in terms of the inapproximability of finding a minimum discounted-weight path, although this is primarily a short review of known set cover approximations.

4.1 Bounded Cardinality

Let k be an upper bound on the cardinality of the bundles; i.e. no single owner has more than k edges, and call this k-CARDINALITY SHORTEST PATH WITH DISCOUNTS. MINIMUM k-SET COVER is the variation in which the cardinality of all sets in C are bounded from above by a constant k; and hence it is a corollary of Lemma 1 that MINIMUM k-SET COVER may be computed exactly with k-CARDINALITY SHORTEST PATH WITH DISCOUNTS.

It is known that the MINIMUM k-SET COVER can be approximated to within a constant factor, but that no better than a constant factor is possible unless P=NP. (see, e.g., [10,14]). Fairly obviously, we can get a k-approximation with the 'bundling' discount scheme, by ignoring the discounts and simply computing the shortest path regardless. It is worth comparing this with the approximation ratios that are known for MINIMUM k-SET COVER. The best approximation ratio for MINIMUM k-SET COVER remains something of an open problem — the best known result for the weighted case is currently $H_k - \frac{k-1}{8k^9}$, where $H_k = \sum_{1,\ldots k} \frac{1}{k}$ is the $k-$th harmonic number [14], and hence finding some better approximation ratio for k-CARDINALITY SHORTEST PATH WITH DISCOUNTS would likely be a difficult, but significant, result.

Theorem 2. *No polynomial-time algorithm with an approximation ratio of less than* $(1-\epsilon)\ln n$ *for any* $\epsilon > 0$ *exists for* SHORTEST PATH WITH DISCOUNTS *unless* $NP \subset D_{TIME}(n^{\log\log n})$.

Proof. In the proof of Lemma 1, we show that all possible solutions to a MINIMUM SET COVER problem exist with identical weight in an instance of SHORTEST PATH WITH DISCOUNTS and vice-versa, so it is approximation-preserving. Hence any approximation ratio for SHORTEST PATH WITH DISCOUNTS would give the same ratio for MINIMUM SET COVER. It was shown by Feige in [10], for MINIMUM SET COVER, that no approximation ratio of $(1-\epsilon)\ln n$ exists for any $\epsilon > 0$ under the assumption that $NP \subset D_{TIME}(n^{\log\log n})$, hence this also holds for SHORTEST PATH WITH DISCOUNTS. □

5 Superadditive Pricing

We now consider the alternative setting, where instead of a (negative) 'discount' value, we introduce a (positive) 'supplement'. We call this problem SHORTEST PATH WITH SUPPLEMENTS, which differs from SHORTEST PATH WITH DISCOUNTS in that each of the 'discount' vectors has only positive values $d^i_{|T|} \geq 0$ (we call these 'supplement' values).

We will see, firstly that even for small supplement values (any $d > 0$) that finding the optimal solution is NP-hard. If we have larger supplement values then we will also see strong inapproximability results. Hardness is shown with a reduction from EXACT COVER BY 3-SETS, as follows. The reduction is demonstrated with the example from Fig. 1, which is shown in Fig. 4. Observe that the shortest path has supplemented weight of 6 — which may be described by either $(1,2,5),(3,4,6)$ or $(3,4,6),(1,2,5)$.

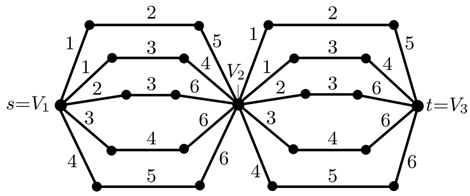

Fig. 4. Fig. 1 As a SHORTEST PATH WITH SUPPLEMENTS problem. $Z = \{1, 2, 3, 4, 5, 6\}$; edges are labeled with the 'owner'.

Theorem 3. SHORTEST PATH WITH SUPPLEMENTS *is NP-hard, even with bundles of size at most* 2*.*

Proof. In the first stage, we will ignore the restriction on the size of the bundles. Take the input X, C to an instance I of EXACT COVER BY 3-SETS. Let $m = |X|$, and build an instance I' of SHORTEST PATH WITH SUPPLEMENTS as follows. Let $Z' = X$ be the set of owners, and let $V' = \{V'_1, \ldots, V'_{m/3+1}\}$. For each triple $C_i \in C$ let the elements in C_i be given by $\{x, y, z\}$, and insert $m/3$ paths of length 3, one between each V'_j and V'_{j+1} for all $j \in \{1, \ldots, m/3\}$, and for each of these paths, let the owners of the edges be given by one each of $\{x, y, z\}$ (the ordering is unimportant). Let $m' = |E'|$, let there be a small supplement parameter, $1/m^2 > d > 0$, and let the supplement vectors be fixed, $\boldsymbol{d}^i = (0, 0, d, \ldots, d)$.

Observe that the shortest path from s to t contains exactly m edges. Also observe that, where a shortest path exists of weight m, this has exactly m edges and that no edges share the same owner. Where an exact cover by 3-sets exists, then there are exactly $m/3$ subsets of C that contain all of the elements in X. Hence there are $m/3$ paths (each of length 3) from V_1 to V_{m+1} which can be selected, such that no owner appears for more than one edge. Hence a shortest path of weight m exists in instance I' if and only if the instance I is a 'yes' instance of EXACT COVER BY 3-SETS, and SHORTEST PATH WITH SUPPLEMENTS is NP-hard.

Now, we can see that this still holds where the bundles are of at most size 2. For the graph $G' = V', E'$, build an input I'' with graph $G'' = V'', E''$. Replace every edge e' in E' by a path of length $m' - 1$, and label each of these edges $e''_{(e,j)}$ for all $j \in E' \setminus \{e'\}$. For every bundle Z_i in the input, consider every pair of edges $(u, v) \in Z_i$ and add a new bundle to Z'' containing the two edges labeled $e''_{(u,v)}$ and $e''_{(v,u)}$. Hence for any path $P' \subseteq E'$ of size t there is a corresponding path $P'' \subseteq E''$ (with the replaced edges) of size $(m' - 1)t$. Also observe that P'' contains two edges in the same bundle, $e''_{(u,v)}$ and $e''_{(v,u)}$ if and only if P' contained two edges (u, v) in the same bundle. Hence for every shortest path P' in I' with supplemented weight given by $|P| + xd$, the shortest path in I'' has supplemented weight $|P|(m' - 1) + xd$; hence the solution to instance I'' gives a solution to instance I', and thus instance I, and SHORTEST PATH WITH SUPPLEMENTS is NP-hard even with edge bundles of size at most 2. $\square$

As in the subadditive case, it is easy to observe that where the supplement values are small then ignoring the discounts gives a good approximation ratio (assuming a maximum weight of 1 and a maximum supplement value of d, then d is an upper bound on the approximation ratio, as the optimal solution has a supplemented weight of at least the 'found' solution, and the supplemented weight of this found solution is bounded from above by a factor of d). However, if we allow supplement values to be as large as m, then we will see strong inapproximability results, from SHORTEST PATH WITH FORBIDDEN PAIRS.

Name SHORTEST PATH WITH FORBIDDEN PAIRS
Instance A weighted graph $G = (V, E, w)$ with two distinct vertices s, t and a collection, $F = \{F_1, \ldots, F_m\}$ of pairs of vertices.
Output The minimum weight path $P \subset E$ from s to t such that for every edge $e \in P$ adjacent to vertex x there is no other edge $e' \in P$ adjacent to vertex y where there is some $F_i \in F = \{x, y\}$.

It is known that no polynomial time approximation algorithm for SHORTEST PATH WITH FORBIDDEN PAIRS exists with an approximation ratio of n^ϵ for any $\epsilon > 0$, unless P=NP [15].

Theorem 4. *No polynomial-time algorithm with an approximation ratio of less than n^ϵ for any $\epsilon > 0$ exists for* SHORTEST PATH WITH SUPPLEMENTS *unless P=NP.*

Proof. Taking the input to a SHORTEST PATH WITH FORBIDDEN PAIRS, build an instance of SHORTEST PATH WITH SUPPLEMENTS with $G' = G$ (let $n = |V|$ and $m = |E|$). For each pair of forbidden vertices $\{x, y\}$ assume an edge $e \in E$ that is adjacent to one of these vertices (let it be x) and add a bundle Z_i containing this edge e and all edges that are adjacent to y. Assume an unweighted graph, i.e. $w(e) = 1$ for all $e \in E$, and that discount values are large; $d \geq n$; hence any SHORTEST PATH WITH SUPPLEMENTS is exactly a shortest path that contains no two edges in the same bundle (unless no such path exists, in which case the total weight is greater than n). Observe that any path through vertex x contains an edge adjacent to x, and hence if it also contains an edge adjacent to y then there is a bundle containing both edges and the supplemented weight is at least $n + 1$ (from the construction of bundles). Similarly, any path with supplemented weight of less than n contains no two edges in the same bundle, and hence does not visit the forbidden pair $\{x, y\}$.

Another instance I'' can be created with bundles of size at most 2 in the same way as in the proof of Theorem 3, by replacing each edge with m edges in series. Let the weight of each edge be $w''(e'') = 1/m$. Hence every SHORTEST PATH WITH SUPPLEMENTS (with weight of less than m) gives a path of the same size in I which avoids forbidden pairs. Let n'' be the number of vertices in I'', and observe that $n'' \leq n^3$. Any approximation algorithm for I'' that gives an approximation ratio of $\alpha = n''^\epsilon$ (with $\epsilon > 0$) for I'' would give an approximation ratio of $\alpha = n^{\epsilon/3}$ for SHORTEST PATH WITH FORBIDDEN PAIRS, which would imply P=NP. □

6 Conclusion and Open Problems

We gave a model for generalising shortest-path problems that appears to have relevance to economic theory. As 'volume discounts' appear in a wide range of literature, it seems an obvious goal to try and combine them with the well-studied path auction. However, the results we have seen are quite negative — even the very simplest discounting scheme, of "buy at least x edges and get a d discount" results in the winner determination problem becoming NP-hard, and hence the VCG mechanism may not be a practical solution.

Our only alternative mechanism, of ignoring the discounts, is obviously quite unsatisfactory — we have identified that allowing volume discounts is a desirable feature. We have also shown that better than $\Omega(\log(n))$ approximation ratios may not be possible, where the discount parameters may be large. This is a reasonable assumption in many settings, such as the supply of services when there is relatively little overhead, for example the use of transport or network infrastructures that have surplus capacity.

Considering the problem as k-CARDINALITY SHORTEST PATH WITH DISCOUNTS, we do not know how to improve on the approximation ratio of k, or if a ratio of k might be close to optimal. However, there is scope for the existence of an approximation algorithm with an improved approximation ratio. An approximation algorithm which does better than the naive approach of ignoring the discounts would likely be a fundamental contribution to a better auction mechanism. There could be economic benefits to being able to run a tractable procurement auction which takes advantage of these volume discounts, and this is an area that might benefit from future study.

The reductions of Theorem 1 and Theorem 3 show that finding an exact solution is NP-hard, even for series-parallel graphs. We may also be interested in discovering if there are certain other properties of the problem that are required for hardness results — for example, the distance between edges that are bundled together. However, it is easy to observe in Lemma 1 that the solutions to a single-pathed multigraph give exactly the solutions to a set-cover problem of the same size, and hence there exists a H_k-approximation for this class of graph. If the general case does not give rise to a better than k approximation, there may be interesting results in finding classes of graph that have a better approximation.

It is also worth noting that there is a difference in the approximation of the subadditive and superadditive cases. While the discount parameters may be reasonably limited by the weights of the edges, the supplement parameter does not have this limitation — and so we find that approximation may be harder in the superadditive case when the supplement parameter is large.

Differences in approximation can be particularly important in terms of auction mechanisms; the frugality ratio of an auction mechanism is often used to measure its performance in terms of payment. It was shown in [12] that the frugality ratio (of a truthful approximation mechanism) may be upper bounded a factor of the approximation ratio; hence a good approximation ratio is important in making an attractive auction mechanism. Mechanisms with a much better frugality ratio than VCG have been shown to exist for path auctions (e.g. [16,4,17]). It is a natural question to ask if similar approaches could improve frugality for this auction model. By combining those approaches with a good approximation algorithm, it may be possible to create a particularly desirable auction mechanism — that is both tractable and has good frugality in this setting.

References

1. Archer, A., Papadimitriou, C., Talwar, K., Tardos, E.: An approximate truthful mechanism for combinatorial auctions with single parameter agents. In: Procs. of 14th SODA, pp. 205–214 (2003)
2. Archer, A., Tardos, E.: Frugal path mechanisms. In: Procs. of 13th SODA, pp. 991–999 (2002)
3. Briest, P., Krysta, P., Vöcking, B.: Approximation techniques for utilitarian mechanism design. SIAM J. Comput. 40(6), 1587–1622 (2011)
4. Chen, N., Elkind, E., Gravin, N., Petrov, F.: Frugal mechanism design via spectral techniques. In: Procs. of 51st FOCS, pp. 755–764 (2010)
5. Clarke, E.H.: Multipart pricing of public goods. Public Choice 11(1) (1971)
6. Dijkstra, E.W.: A note on two problems in connexion with graphs. Numerische Mathematik 1, 269–271 (1959)
7. Du, Y., Sami, R., Shi, Y.: Path auctions with multiple edge ownership. Theoretical Computer Science 411(1), 293–300 (2010)
8. Elkind, E., Goldberg, L.A., Goldberg, P.W.: Frugality ratios and improved truthful mechanisms for vertex cover. In: Procs. of the 8th ACM-EC, pp. 336–345 (2007)
9. Elkind, E., Sahai, A., Steiglitz, K.: Frugality in path auctions. In: Procs. of 15th SODA, pp. 701–709 (2004)
10. Feige, U.: A threshold of ln n for approximating set cover. Journal of the ACM 45, 314–318 (1998)
11. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness (1979)
12. Goldberg, P.W., McCabe, A.: Commodity auctions and frugality ratios. In: Serna, M. (ed.) SAGT 2012. LNCS, vol. 7615, pp. 180–191. Springer, Heidelberg (2012)
13. Groves, T.: Incentives in teams. Econometrica 41(4), 617–631 (1973)
14. Hassin, R., Levin, A.: A better-than-greedy approximation algorithm for the minimum set cover problem. SIAM Journal on Computing 35(1), 189–200 (2005)
15. Kann, V.: Polynomially bounded minimization problems that are hard to approximate. Nord. J. Comput. 1(3), 317–331 (1994)
16. Karlin, A.R., Kempe, D., Tamir, T.: Beyond VCG: Frugality of truthful mechanisms. In: Procs. of 46th FOCS, pp. 615–626 (2005)
17. Kempe, D., Salek, M., Moore, C.: Frugal and truthful auctions for vertex covers, flows and cuts. In: Procs. of 51st FOCS, pp. 745–754 (2010)
18. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press (2007)
19. Phillips, R.: Pricing and Revenue Optimization. Stanford University Press (1995)
20. Stremersch, S., Tellis, G.J.: Strategic Bundling of Products and Prices: A New Synthesis for Marketing. Journal of Marketing 66(1), 55–72 (2000, 2002)
21. Talwar, K.: The price of truth: Frugality in truthful mechanisms. In: Alt, H., Habib, M. (eds.) STACS 2003. LNCS, vol. 2607, pp. 608–619. Springer, Heidelberg (2003)
22. Vickrey, W.: Counterspeculation, Auctions, and Competitive Sealed Tenders. The Journal of Finance 16(1), 8–37 (1961)