Mechanisms for a Spatially Distributed Market[1]

Moshe Babaioff[2], *Noam Nisan and Elan Pavlov*[3]
{ mosheb, noam, elan@cs.huji.ac.il }

School of Computer Science and Engineering
The Hebrew University of Jerusalem, Jerusalem 91904, Israel

Abstract

We consider the problem of a spatially distributed market with strategic agents. In this problem a single good is traded in a set of independent markets, where shipment between markets is possible but incurs a cost. The problem has previously been studied in the non-strategic case, in which it can be analyzed and solved as a min-cost-flow problem. We consider the case where buyers and sellers are strategic. Our first result gives a double characterization of the VCG prices, first as distances in a certain residue graph and second as the minimal (for buyers) and maximal (for sellers) equilibrium prices. This provides a computationally efficient, individually rational and incentive compatible welfare maximizing mechanism. This mechanism is, necessarily, not budget balanced and we provide also a budget-balanced mechanism (which is also computationally efficient, incentive compatible, and individually rational) that achieves high welfare. Some of our results extend to the cases where buyers and sellers have arbitrary convex demand and supply functions and to the case where transportation is controlled by strategic agents as well.

[1]Partial version of this paper appeared in the Fifth ACM Conference on Electronic Commerce, 2004.

[2]This research was supported by grants from the Israeli Ministry of Science, the Israeli Academy of Sciences and the USA-Israel Bi-national Science Foundation. The first author (Babaioff) was also supported by Yeshaya Horowitz Association. The third author (Pavlov) was also supported the Evergrow project of the EU and part of the work was done while on a visit to the Exystence Institute

[3]The authors wish to thank Rachel Chen and Max Shen for fruitful discussions.

1 Introduction

With the recent emergence of the Internet as a central platform for commerce, we have seen a very large number of online markets and auctions on the Web. While clearly financial markets have been computerized for quite some time now, we have recently seen many more types of computerized markets for goods, services, obligations, etc. While, in principle, computerized markets or auctions are no different than the classical non-computerized variants, it is well known that significant differences emerge due to the opportunities created by new economy of scale, complexity, speed, software agents and other factors. Additionally, many significant issues that were traditionally handled by human conventions and intuitions must be made formal so that software can handle them. Indeed there is a very large body of academic work attempting to deal with these new issues of the world of computerized markets and auctions (see [8] for a survey).

In the real world, interaction between markets is everywhere. A market for some good is very much affected by many related markets: other markets for the same good, markets for goods that are complements or substitutes to it, markets for its production factors, etc. These interactions between markets are all governed and magically handled by the famous ”invisible hand” (and perhaps some government regulation) resulting, presumably, in a socially optimal global outcome. In the world of electronic commerce, much of this ”invisible hand” must be designed, analyzed, and programmed.

One can think of three majors classes of interactions between markets:

- Between different markets for the same good [17, 5] (or at least, nearly equivalent goods). E.g. between the different markets around the world for buying and selling oil.
- Between markets for different goods that are complements or substitutes of each other. E.g. between the oil market and the gas market or between the oil market and the car market. To some extent, combinatorial auctions from the ”simultaneous ascending auctions” family [10] [14] may be viewed as being conducted over a set of markets, each for a single good.
- Between different markets along the supply chain. E.g. between the oil market and a plastics market. This was discussed in [3, 4, 20].

The focus of this paper is the first, and probably the simplest type of interaction: the interaction between different markets for the same good. This scenario has been studied in the economics literature and is termed a ”spatially distributed market” – viewing the different markets as residing in a different geographic locations, but being logically parts of a single distributed global market [11]. Clearly, in a computerized world, the physical location is not the issue, but rather the ”virtual location”. What makes ”virtual locations” distinct from each other is a cost associated with moving the goods between these locations. Potential examples include the oil market mentioned above; different stock markets (either in different countries or between various ECN’s (Electronic Communication Networks)[4]; digital content like movies with a significant communication overhead; or independently administered markets, where transfers between them incur some administrative transaction cost.

Our model consists of a network of markets, where goods may be transported between the markets for a price. In most of the paper we assume that the transportation costs are known and that the markets follow protocols (i.e. are “obedient” and not strategic). Buyers and sellers have a valuation for a unit of the traded good, where buyers are willing to buy a unit for at most their valuation, and sellers to sell a unit for at least their valuation. The difference between our work and most of the previous work is that we assume *strategic agents* in the game-theoretic sense of mechanism design [12], rather than

[4]see for example Island site at http://www.island.com and Instinet site at http://www.instinet.com

the price-taking agents that were considered in most economic literature e.g. [2]. To the best of our knowledge, only a few recent papers considered strategic agents in a spatially distributed market. The work by Roundy et al. [17] considered one-sided auctions which do not exhibit most of the issues of the two-sided case studied here. The work by Chu and Shen [5] considered two-sided auctions and provides some economic results similar to ours for a similar model, but without considering the computational aspects of such mechanisms.

We start by analyzing the non-strategic case. The allocation in a spatially distributed market consists of specifying which buyers and sellers trade in each market, as well as the vector of inter-market transfers of goods needed to clear all markets. We first show how finding the globally efficient allocation can be reduced to a min-cost-flow problem, thus obtaining an efficient algorithm for the problem. (The work of Roundy et al. [17] did not consider the algorithmic aspects, but this reduction seems to be known in similar settings). Trade in each market is executed at the market price, and the vector of market prices is of central interest. In equilibrium, the prices in the markets are consistent with the shipment costs and the agents are satisfied with the allocation given the prices. We apply well known results from the theory of min-cost-flow problem to obtain the first and second welfare theorems for this setting:

Theorem 1 *(The Two Welfare Theorems for SDM)*

- *(First Welfare Theorem for SDM) If an allocation and a price vector are in a spatial price equilibrium then the allocation is efficient.*
- *(Second Welfare Theorem for SDM) For any efficient allocation there exists a price vector such that is in a spatial price equilibrium with it.*

We then move on to the strategic setting. In this setting we start by characterizing the VCG prices. We obtain a double characterization of these prices: as distances in the residual graph obtained in the reduction above, and as the maximum (for sellers) and minimum (for buyers) equilibrium prices in the markets:

Theorem 2 *(SDM VCG prices characterizations)*

1. *Denote by d_i the distance from the sink node to market M_i in the residual graph, and by D_i the distance from market M_i to the sink node in the residual graph. Then in the VCG payment scheme, every seller in market M_i receives d_i, and every buyer in market M_i pays $-D_i$.*
2. *The set of price vectors that are in equilibrium with the efficient allocation form a complete lattice. The vector of VCG prices that the buyers pay is the minimal element of the lattice. The vector of VCG prices paid to the sellers is the maximal element of the lattice.*

This VCG pricing scheme yields an incentive compatible and individually rational pricing scheme. The characterization above allows efficient computation of these prices, and thus provides a (a) computationally efficient, (b) incentive compatible, (c) individually rational, and (d) socially efficient mechanism for the spatially distributed markets problem. Unfortunately, this mechanism is not budget balanced, and indeed adding the budget balance property is impossible without relaxing one of the conditions (b–d) [15].

Our next mechanism is budget-balanced as well as (a) computationally efficient, (b) incentive compatible, (c) individually rational, but slightly relaxes condition (d) of social efficiency. The idea is to slightly reduce the trade from the optimally efficient allocation. This idea was first used by McAfee [13] for double auctions and was later used by Babaioff and Nisan [3] and Babaioff and Walsh [4] for supply chain formation. Chu and Shen [5] have used a similar idea for double auction with pair related costs (a model that is equivalent to a market for each agent), and provided similar economic results, but without considering the computational aspects of such mechanisms. Our main observation is that the global trade needs only be reduced by a very small amount to achieve our goals: a single unit reduction in each component of the network of markets that we call a *Commercial Relationship Component* (a set of markets that have a direct or indirect trade). We thus present a "trade reduction mechanism" with the following properties:

Theorem 3 *The "trade reduction mechanism" is (a) computationally efficient, (b) incentive compatible, (c) individually rational, and (d) budget balanced. For any values of the agents valuations for which the efficient allocation is non-empty, the efficiency is bounded* [5] *by*

$$\frac{\mathbf{V}(A^{TR}(v))}{\mathbf{V}(A^{*}(v))} \geq \min_{\gamma\in\Gamma} \frac{|\gamma|-1}{|\gamma|}$$

where $\mathbf{V}(A^{TR}(v))$ *is the value of the TRM allocation,* $\mathbf{V}(A^{*}(v))$ *is the value of the efficient allocation,* Γ *is the set of CRCs and* $|\gamma|$ *is the size of trade (number of sellers/buyers) in the CRC* $\gamma \in \Gamma$.

Furthermore, if the number of markets is fixed and if the valuations/costs of all buyers/sellers are chosen from the same distribution, then when the number of agents go to infinity, the efficiency converges to perfect efficiency (1) in expectation.

Structure of the Paper: In section 2 we formally present our model and our notations. Section 3 deals with the non-strategic case, and provides the basic reduction to the min-cost flow problem. Section 4 provides the analysis of VCG prices, while section 5 describes the trade-reduction mechanism. In section 6 we shortly describe ideas for future extensions. Appendix A summarizes results we need from the theory of min-cost-flows. Appendix B presents extensions of the model and characterizes the VCG payments for these extensions. All proofs have been omitted from the body of the paper and appear in Appendix C.

2 Model

In this section we present our model for a ***Spatially Distributed Market (SDM)***. In section 2.1 we consider the case of non strategic agents, where all information is public. In section 2.2 we extend the model to the case of strategic agents with quasi-linear utility functions, which have private independent values and act in order to maximize their personal utility. This model will be our main interest in this paper.

2.1 Spatially Distributed Market Model with Non-Strategic Agents

We now describe the model for a spatially distributed market with non-strategic agents. In this model, a global market for a single good is constructed from a set of k ***markets*** $M_1,\ldots,M_k$ each in a different location. These markets are the nodes of a simple directed graph representing the possible commercial relationships between the markets. If the good can be shipped from M_i to M_j then there is a directed edge (arc) (M_i,M_j) in the graph. For each edge (M_i,M_j) there is an integral ***cost*** of $c_{i,j} \geq 0$, which is the cost of shipment of one unit of the good from market M_i to market M_j along the edge (M_i,M_j) (we allow the cost of edge (M_i,M_j) to differ from the cost of (M_j,M_i), meaning $c_{i,j} \neq c_{j,i}$). We assume that this cost

[5] Actually it is possible in some cases to achieve a slightly more efficient allocation but there exist cases where this is the best that can be achieved

is exogenous and publicly known. The number of nodes in this graph is k and the number of edges is denoted by m.

We assume that the ***capacity*** of any edge between two markets is infinite, so any amount of good can be shipped from any market to any other market (we relax this assumption when we consider transportation that is controlled by strategic carriers, see section B.2 in the Appendix). We denote the set of all agents as $\mathbf{N}$. The agents are divided to markets, and to ***sellers*** and ***buyers*** in each market. The set of buyers in market M_i is $\mathbf{B}_i$, and set of sellers in that market is $\mathbf{S}_i$. We mark the set of all buyers by $\mathbf{B}$ ($\mathbf{B} = \cup_i \mathbf{B}_i$), and the set of all sellers by $\mathbf{S}$ ($\mathbf{S} = \cup_i \mathbf{S}_i$). The set of agents is therefore $\mathbf{N} = \mathbf{B} \cup \mathbf{S}$. Each buyer (seller) is a single parameter agent, which means that she wants to buy (sell) a single unit of the good in one particular market, and has one parameter that represents the value (cost) that she gets from trading. A buyer (seller) ***trades*** if she buys (sells) a unit of the good in her market. Excess demand (supply) in a given market will be matched to a surplus of supply (demand) in other markets, by ***shipment (inter market trade)*** of goods from one market to the other. We denote by x the ***shipment vector***, where $x_{i,j} \geq 0$ is the number of units shipped from market M_i to market M_j along the edge (M_i, M_j). There is a ***trade*** between market M_i and market M_j if there is a shipment of goods from M_i to M_j ($x_{i,j} > 0$).

The valuation of buyer $b \in \mathbf{B}$ from buying a unit of a good is $v_b \geq 0$. The cost for seller $s \in \mathbf{S}$ for selling her unit of the good is $c_s \geq 0$ (so her valuations for trading is $v_s = -c_s \leq 0$). Any agent has a valuation of zero if she does not trade. In the first part of this paper we assume that the valuations/costs are publicly known, we change this assumption when we consider strategic agents. By a slight abuse of notation we sometimes identify the cost of a seller or the valuation of a buyer with her identity (so we say $v_b \in \mathbf{B}$ for a buyer with valuation v_b). It will be clear from the context if we are referring to the agent or to the cost/valuation. The vector of all agents valuations will be marked as v.

Given a set of agents and their valuations v, an ***allocation*** $\mathbf{A} = (T, x)$ is the set of trading agents T and a shipment vector x, for which the feasibility condition described below holds. Allocation $\mathbf{A}$ is ***feasible*** iff each market is ***materially balanced***, that is for each market M_i

$$|\mathbf{S}_i \cap T| + \sum_{j:(M_j, M_i) \in E} x_{j,i} = |\mathbf{B}_i \cap T| + \sum_{j:(M_i, M_j) \in E} x_{i,j}$$

The RHS is the number of units arrived to market M_i, and the LHS is the number of units leaving market M_i ($|\mathbf{B}_i \cap T \cap M_i|$ is the number of trading buyers in market M_i, and $|\mathbf{S}_i \cap T|$ is the number of trading sellers in market M_i).

The ***value*** $\mathbf{V}(\mathbf{A})$ ***of an allocation*** $\mathbf{A}$ is the sum the agent values in $\mathbf{A}$, minus the shipment cost:

$$\mathbf{V}(\mathbf{A}) \equiv \sum_{w \in T} v_w - \sum_{(M_i, M_j) \in E} x_{i,j} c_{i,j} = \sum_{v_b \in T \cap \mathbf{B}} v_b - \sum_{c_s \in T \cap \mathbf{S}} c_s - \sum_{(M_i, M_j) \in E} x_{i,j} c_{i,j}$$

The value of an allocation $\mathbf{A}$ excluding the value of agent i is $\mathbf{V}_{-w}(\mathbf{A}) \equiv \mathbf{V}(\mathbf{A}) - v_w$. An allocation $A^*(v)$ is ***efficient*** if it is feasible and has a maximal value with respect to v, over all feasible allocations. The ***efficiency*** of allocation $\mathbf{A}$ is $\frac{\mathbf{V}(\mathbf{A})}{\mathbf{V}(A^*(v))}$. The ***Spatial Market Value Maximization Problem (SMVMP)*** is the problem of finding an efficient allocation for a given spatially distributed market problem.

For ease of exposition and in order to simplify the analysis we will assume throughout the paper that no two allocations have the same value (so there is a unique efficient allocation). We can break ties between allocations by lexicographical order on the identities of the agents. We refer the reader to e.g., [4] for the technical details.

Any trade mechanism should set the allocation as well as the agents payments for the goods they buy or sell. We assume a quasi-linear utility model, which means that a trading buyer $v_b \in \mathbf{B}$ obtains a utility of $v_b - p_b$ by buying a unit and paying p_b. A trading seller $c_s \in \mathbf{S}$ obtains a utility of $p_s - c_s$ by selling her unit and receiving p_s. When a non-trading buyer (seller) pays (receives) 0, she has a utility of 0. The agents are rational (self-interested), utility maximizer entities, so each agent will always trade in a way that maximize her utility.

We are interested in the case where prices are set per market and are independent of particular agents. Let p_i be the price of a unit in market M_i (this is the price that a trading buyer in the market pays for a unit, and a trading seller in the market receives for her unit). Let $\vec{p}$ be the vector of k market prices (one per market). We now define when an allocation and a price vector are in equilibrium (note that we consider the case of static equilibrium, and we do not handle any dynamic behavior over time.). In section 3 we present a strong connection between SMVMP and this equilibrium concept.

Definition 4 *An allocation $\mathbf{A} = (T, x)$ and a vector of market prices $\vec{p}$ are in **Spatial Price Equilibrium (SPE)** iff the allocation and price vector satisfy the following equilibrium conditions.*

1. *For any edge $(M_i, M_j) \in E$:*
   - $p_i + c_{i,j} \geq p_j$.
   - *If $x_{i,j} > 0$ then $p_i + c_{i,j} = p_j$.*
2. *For any market M_i*
   - *For any buyer $v_b \in \mathbf{B}_i$, if $v_b > p_i$ then $v_b \in T$, and if $v_b < p_i$ then $v_b \notin T$.*
   - *For any seller $c_s \in \mathbf{S}_i$, if $c_s < p_i$ then $c_s \in T$, and if $c_s > p_i$ then $c_s \notin T$.*

The first condition on the market edges states that the price in market M_j is never higher then the price of importing the good from market M_i. The second condition states that if there is a trade between market M_i and market M_j then it costs the same to buy a unit in market M_j as to buy a unit in market M_i and ship it to M_j. Note that since $x_{i,j} \geq 0$ we can infer that for any edge $(M_i, M_j) \in E$, if $p_i + c_{i,j} > p_j$ then $x_{i,j} = 0$. The condition on the agents matches the behavior of self-interested agents, where each agent trades if she gains from trading, and does not trade if she loses.

We can divide the graph to components, each containing a set of markets for which trade is possible between them (these are the connected components of the corresponding undirected graph). Any such component can be dealt with separately (since no trade is possible between the components). So we assume without loss of generality that there is only one component and that $k \leq m \leq k^2$, where m is the number of edges.

2.2 Spatially Distributed Market Model with Strategic Agents

We now extend the SDM model to the case that the value v_w of each agent w is ***private***, and is independent of the value of the other agents. The agents are rational (self-interested), utility maximizers entities, so if it is possible the agents will manipulate the allocation and prices in their favor (the markets are assumed to be obedient and act according to the defined protocol, only the agents act strategically).

This work belongs to the field of Mechanism Design (MD) and in particular Algorithmic Mechanism Design (AMD) (for background on MD refer to [12], and for AMD refer to [16]). We present the basic definitions we need from MD.

A ***mechanism*** for SDM is constructed from an allocation rule and a payment rule. An ***allocation rule*** is a function that given a set of agent values, outputs an allocation. A payment rule is a function that given a set of agent values, outputs the payment from each agent. The mechanism is executed on the set of the agents ***reported values*** which might be different from the agents true values. We are interested in mechanisms that encourage the agents to report their true values.

A mechanism is ***incentive compatible (IC)*** in dominant strategies, if for any agent and for any values of the other agents, bidding truthfully maximizes the agent utility over all her possible bids. For such a mechanism, truth telling is a dominant strategy equilibrium. Such a mechanism is also ***individually rational (IR)*** if no agent has negative utility by participating in the mechanism. A mechanism is ***efficient*** if for any values of the agents, the efficient allocation is a dominant strategy equilibrium. A mechanism is ex-post (weakly[6]) ***budget-balanced*** if for any values of the agents, the sum of all payments from the trading buyers is not less then the sum of all the payments to the trading sellers and the shipment costs. This means that for any allocation $\mathbf{A} = (T,x)$ picked by the mechanism,

$$\sum_{v_b \in \mathbf{B} \cap T} p_b - \sum_{c_s \in \mathbf{S} \cap T} p_s - \sum_{(M_i,M_j) \in E} x_{i,j} c_{i,j} \geq 0$$

where p_b is the payment from buyer v_b, p_s is the payment to seller c_s and $x_{i,j}$ is the shipment on edge $(M_i,M_j) \in E$ with cost $c_{i,j}$.

2.3 Background - Minimum Cost Flow

To derive our results we use a reduction to the well known ***Minimum Cost Flow Problem (MCFP)***. In this section we present the definition of the problem, Appendix A presents some more background on the MCFP from the book by Ahuja et al. [1]. The appendix contains the results for the MCFP needed to derive our results.

Let $G = (V,E)$ be a directed network with cost $c_{i,j}$ and capacity $u_{i,j} \geq 0$ associated with every arc $(i,j) \in E$. We associate with the node $i \in V$ a number b_i which indicates its supply or demand. The ***Minimum Cost Flow Problem (MCFP)*** can be stated as follows:

$$Minimize \sum_{(i,j) \in E} c_{i,j} x_{i,j}$$

subject to

$$\sum_{j:(i,j) \in E} x_{i,j} - \sum_{j:(j,i) \in E} x_{j,i} = b_i \ for\ all\ i \in V$$

$$0 \leq x_{i,j} \leq u_{i,j}$$

Where $x_{i,j}$ is the flow on arc $(i,j) \in E$. Let C denote the largest magnitude of any arc cost, and let U denote the largest magnitude of any supply/demand or finite capacity. We assume that all data (cost, supply/demand, capacity) are integral and that $\sum_{i \in V} b_i = 0$.

The ***residual network*** $G(x)$ corresponding to the flow x is defined as follows. We replace each arc $(i,j) \in E$ by two arcs, (i,j) and (j,i). The arc (i,j) has cost $c_{i,j}$ and ***residual capacity*** $u_{i,j} - x_{i,j}$, and the arc (j,i) has a cost $c_{j,i} = -c_{i,j}$ and residual capacity $x_{i,j}$. The residual network consists only of arcs with positive residual capacity.

3 Spatial Price Equilibrium with non-strategic agents

In this section we present a reduction of SMVMP for non-strategic agents, to the minimal cost flow problem. We also present the relationship between SMVMP and Spatial Price Equilibrium (SPE). Figure 1(a) presents an example of SDM that will be used throughout the paper.

[6]Weakly budget-balanced means that there might be an exact budget balance (balance of 0) or a surplus

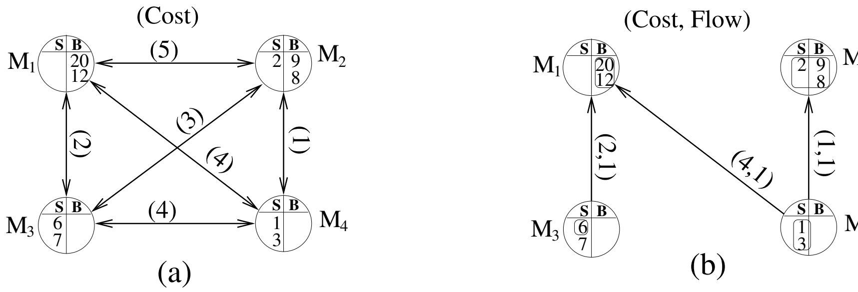

Figure 1: An example of SDM with 4 markets. Each market is a node and contains the sellers and buyers under the respective S and B columns. (a) The SDM problem, the cost of each edge is marked near the edge. (b) The efficient allocation for that problem, trading agents are circled.

We now turn to present a reduction from SMVMP (the problem of finding the efficient allocation) to the minimum cost flow problem. The reduction between the two problems is well known.

The efficient allocation is calculated by solving a MCFP on the following graph. We copy the markets graph (the market nodes and the edges between them). We add another sink node indexed by 0 and named node $\mathbf{Z}$ to the graph. For each buyer v_b with value $v_b \geq 0$ in market M_j, we add an arc with capacity one and cost of $-v_b$ from the market to the sink node. For each seller c_s with cost $c_s \geq 0$ in market M_i, we add an arc with capacity one and cost of c_s from the sink node to market M_i. Note that the graph we have built is not a simple graph. There are multiple parallel arcs from each market to the sink node corresponding to buyers, and from the sink to each market, corresponding to the sellers. All nodes have supply/demand of zero ($b_i = 0$ for any node i), so any feasible flow is a circulation, and there is always a solution to this MCFP.

The graph of the MCFP and the optimal flow for the example of Figure 1(a) are presented in Figure 2(a) and Figure 2(b) respectively. The infinite capacity of the edges between the markets is marked by "inf".

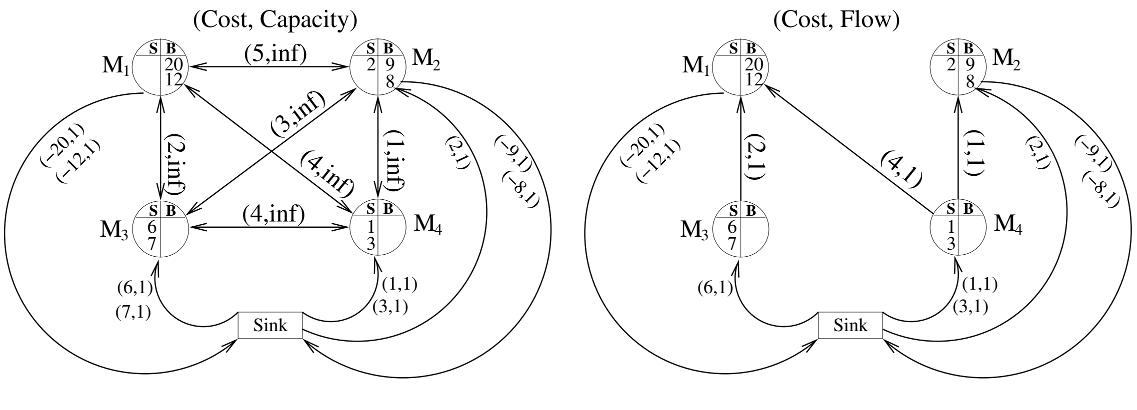

(a) The MCFP graph

(b) The Minimum Cost Flow

Figure 2: (a) The MCFP graph and (b) The optimal flow, for the example presented in Figure 1(a). Only edges with non zero flow are presented. Multiple edges from/to the sink to/from the same market are shown as a single edge with multiple costs.

We find the minimal cost flow x^* for this graph and recover the allocation $\mathbf{A} = (T, x)$ from this flow in the following way. We note that since all data is integral, a integral efficient flow can be found. In particular, the flow on each agent's arc is either zero or one. We define the set of trading agents T to be the set of agents for which there is a flow of 1 on their corresponding edge (A buyer receives a unit of the good iff there is a flow of one on her edge, and a seller sells her unit of the good iff there is a flow of one on her edge). The shipment vector x equals to the minimal cost flow vector x^* on the inter market edges. A flow of $x^*_{i,j}$ from market M_i to market M_j means that $x_{i,j} = x^*_{i,j}$ units of the good are being shipped from market M_i to market M_j.

We apply the above algorithm on the minimum cost flow of Figure 2(b) to recover the maximal value allocation presented in Figure 1(b). The trading agents are circled and the shipment of goods are marked on the edges. The only non trading agent is 7 in market M_3, since there is no flow on her edge.

Using this reduction we can easily derive the first and second welfare theorems for SDM. The first theorem is derived from the complementary slackness optimality conditions for the MCFP. To prove the second theorem, we show that the distances from the sink node in the residual graph are equilibrium prices.

Theorem 5 *(First Welfare Theorem for SDM) If an allocation* $\mathbf{A}$ *and price vector* $\vec{p}$ *are in a spatial price equilibrium then the allocation* $\mathbf{A}$ *is efficient.*

Proof. See section C.1 in the Appendix. □

Theorem 6 *(Second Welfare Theorem for SDM) If an allocation* $\mathbf{A}$ *is efficient then there exists a price vector* $\vec{p}$ *such that* $\mathbf{A}$ *and* $\vec{p}$ *are in a spatial price equilibrium.*

Proof. The theorem is a direct result of Lemma 25 presented in appendix C.1. □

4 An Efficient Mechanism for Strategic Agents

In this section we present our results for SDM with **strategic** agents (the cost/value of each agent is private information). The well known VCG mechanism ([6, 9, 19]) is a mechanism that is IR, IC and efficient (but is not budget balanced) for SDM. We first characterize the VCG payments for SDM as distances in some residual graph, this enables us to present a computationally efficient algorithm for the VCG mechanism. We also characterize the VCG payments for SDM as the extreme elements in the lattice of prices that are in equilibrium with the efficient allocation. We characterize the VCG payments for two model extensions in Appendix B.

The VCG allocation rule picks the allocation that maximizes the efficiency for the reported values. We present the general VCG payment scheme and apply it to the SDM. This scheme causes the agents to bid truthfully. Since the mechanism is IC, we assume that the agents reported bids are the same as the agents true values (so the bids are also denoted by v). The VCG payment of agent w with respect to the bids v (which are the true values in equilibrium) is defined as

$$VCG_w(v) \equiv \mathbf{V}(A^*(v_{-w})) - \mathbf{V}_{-w}(A^*(v)) \tag{1}$$

Where v_{-w} is the vector of bids of all agents but w.

Intuitively, $VCG_w(v)$ is the "harm" done by agent w to the other agents by bidding v_w, so this payment scheme internalizes the externalities. Observe that $VCG_w(v) \leq v_w$ and that $VCG_w(v) = 0$ if w is not in the efficient allocation.

The following fact summarizes the well known properties of any VCG mechanism in particular for SDM.

Fact 7 *The VCG mechanism for SDM is incentive compatible, individually rational and efficient. It also has a budget deficit.*

We present two characterizations of VCG payments for SDM, the first as distances in the residual graph, and the second as the extreme elements in the lattice of prices that are in equilibrium with the efficient allocation. For the first characterization as distances in the residual graph we need some notation.

Let $d(n1,n2)$ be the distance (shortest path length) from node $n1$ to node $n2$ in the residual graph $G(x^*)$, where the edge length is its cost. For any two nodes $n1$ and $n2$, $d(n1,n2)$ is well defined since there are no negative cycles in the residual graph $G(x^*)$ (by Theorem 20). For the special case that one of the nodes is the sink node we present the following notation. Let $d_i = d(0,i)$ be the distance from the sink to market M_i in the residual graph, and let $D_i = d(i,0)$ be the distance from market M_i to the sink node. We denote by $\vec{d}$ the vector of distances from the sink to the markets, and by $\vec{D}$ the vector of distances from the markets to the sink node.

The residual graph for the optimal flow shown at Figure 2(b) is presented in Figure 3. Note that all the agents' edges except the edge of cost 7 from the sink to market M_3 have been replaced by reversed residual edges with opposite sign, since all of the corresponding edges have a flow of maximal capacity (1). The distance from the sink to each market d_i, and the negative of the distance from each node to the sink $-D_i$, are presented near each of the market nodes.

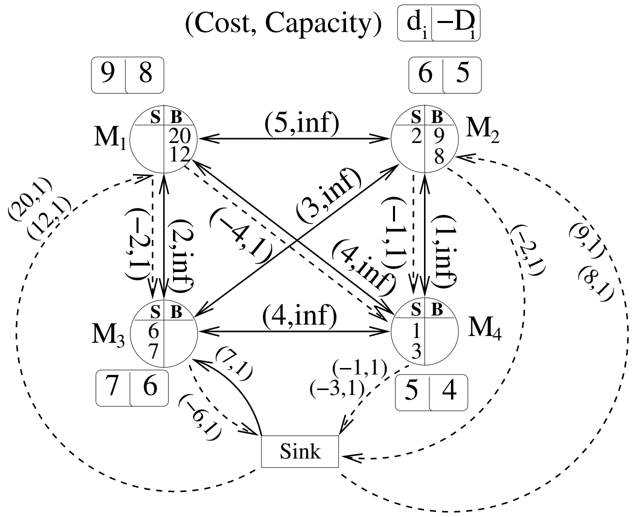

Figure 3: The residual graph for the optimal flow shown at Figure 2(b). For each market the distance from the sink (d_i) and the minus of the distance to the sink ($-D_i$) are shown next to the market.

Theorem 8 *The VCG payments for trading agents are the following.*

- *For any market M_i, any trading seller in market M_i receives d_i.*
- *For any market M_j, any trading buyer in market M_j pays $-D_j$.*

Proof. See section C.2.1 in the Appendix. □

Note that in any market, all trading buyers pay the same price, and all trading sellers receive the same price. But, a trading buyer typically pays a different (smaller) price than a trading seller in the same market receives.

Our characterization of the VCG payments as distances provides a computationally efficient algorithm for the VCG mechanism. The running time of the algorithm is dominated by the efficient allocation calculation, which can be solved as a convex minimum cost flow. This can be done in polynomial time, assuming that the values/costs in each market are sorted [7] For n agents in k markets, with m edges between the markets, and maximal value/cost of agents and edges C we derive the following result.

Proposition 9 *Assuming that the values/costs in each market are sorted, there exists an algorithm that calculates the VCG mechanism for SDM that runs in time $O((m+k\ log\ k)\ m\ log\ (C+n))$.*

Proof. See section C.2.3 in the Appendix. □

Next we show that the VCG payments are the extreme elements in the lattice of prices that are in spatial price equilibrium with the efficient allocation. The question of whether the VCG prices are the extreme equilibrium prices was considered by Gul and Stacchetti [10] for combinatorial auctions. Our results are similar to the results they present for combinatorial auctions with single improvement properties.

Definition 10 *Let $\vec{p},\vec{q}$ be two price vectors of length k.*
*Their **join** $\vec{r}=\vec{p}\vee\vec{q}$ is defined as $r_i=\min\{p_i,q_i\}$ for each $1\le i\le k$.*
*Their **meet** $\vec{s}=\vec{p}\wedge\vec{q}$ is defined as $s_i=\max\{p_i,q_i\}$ for each $1\le i\le k$.*

Definition 11 *A set of price vectors P is a **lattice** if $\vec{p},\vec{q}\in P$ then $\vec{p}\vee\vec{q},\vec{p}\wedge\vec{q}\in P$. A lattice is **complete** if for any $W\subset P$ it holds that $\bigvee(W),\bigwedge(W)\in P$, where $\bigvee(W)_i=\inf\{p_i|p\in W\}$ and $\bigwedge(W)_i=\sup\{p_i|p\in W\}$.*

We are now ready to state the theorem.

Theorem 12 *The set of price vectors that are in SPE with the efficient allocation form a complete lattice P.*

- *The vector of VCG prices paid to the sellers, $\vec{d}$, is the maximal element of the lattice, that is $\bigwedge(P)=\vec{d}$.*
- *The vector of VCG prices that the buyers pay, $-\vec{D}$, is the minimal element of the lattice, that is $\bigvee(P)=-\vec{D}$.*

Proof. See section C.2.2 in the Appendix. □

5 A Budget Balanced Mechanism

Section 4 presented a VCG mechanism for the spatially distributed market. The problem is that like all VCG mechanisms for bilateral trade, the payments of the buyers and the sellers result in a budget deficit as shown by Myerson and Satterthwaite [15]. For SDM this means that the total payment from the buyers is less then the total payment to the sellers plus the cost of transportation. In this section we suggest a budget balanced, individual rational and truth telling mechanism with high efficiency which we call the *Trade Reduction Mechanism (TRM)*. Our ultimate goal is to maximize total value subject to the requirement that the budget is balanced. The TRM mechanism achieves high total value subject to BB.

Our basic tool is to leave some efficiency on the table and get budget balance by truth telling payments. This idea was first used by McAfee [13] for double auctions and was later used by Babaioff and Nisan [3] and Babaioff and Walsh [4] for supply chain formation. In a trade reduction allocation we discard efficient trade(s) and pick a sub optimal allocation in order to achieve budget balance. In our case, since the market

[7] We assume that the values are sorted since we are interested in the communication costs, and sorting can be done locally in each market

is distributed it is not clear what trades(s) to reduce, how to decide on payments and what trades will take place after the reduction. The Trade Reduction Mechanism we propose offers a solution for these questions.

To define the TRM mechanism we need the following definitions.

Definition 13 *Markets M_i, M_j are in a* commercial relationship *(CR) if there is trade between M_i and M_j or between M_j and M_i in the efficient allocation.*

Definition 14 *The Commercial Relationship Component (CRC) of a market M_i is the transitive closure of the commercial relationship property.*

In essence the CRC of market M_i contains all of the markets with which M_i has an direct or indirect commercial relationship.

Our main result in this section is the somewhat surprising fact that a *single* trade reduction in each CRC suffices to achieve an individually rational and incentive compatible mechanism that is budget balanced (BB). In order to formally define what we mean by a "single trade reduction", we first define some sub graph of the residual graph. We use it to find the allocation and the payments of the trade reduction mechanism. This is a directed graph with length (cost) on each edge (we do not need the capacity on the edges).

Definition 15 *The* reduced residual graph (RRG) *is a graph consisting of all nodes of the residual graph and the following subset of the edges, each with its cost in the residual graph serving as the edge length.*

- *For each edge (M_i, M_j) such that there is flow on the edge, we add the edge with its cost and its reversed residual edge with the negated cost.*
- *For each market M_i we add the residual edges corresponding to the trading buyer with the minimal valuation (if such buyer exists) and the trading seller with the maximal cost (if such seller exists).*

Note that all edges between CRCs as well as edges corresponding to non trading agents are not in the RRG[8]. Also, the only edges that are retained in the RRG are the edges belonging to the lowest value trading agents for each class (buyers or sellers) in each market.

The reduced residual graph of the residual graph of Figure 3 is shown in Figure 4(a).

We now formally define the allocation and the payments of the TRM.

The TRM allocation

Given the values of the agents v, the allocation of the trade reduction mechanism $A^{TR}(v)$ is as follows:

- Calculate the efficient allocation using the MCFP graph as shown in section 3, and find the residual graph $G(x^*)$ and the RRG.
- For each CRC calculate the *minimal positive cycle* in the RRG and remove it from the allocation.

The minimal positive cycle in Figure 4(a) is drawn with thicker lines. The resulting trade reduction allocation after removing this cycle is shown in Figure 4(b). Note that the trade was rearranged along the cycle, so market M_4 now ships two units to market M_1 instead of only one as in the efficient trade.

The TRM payments

Like the VCG payments, the payments in the TRM are also calculated as distances in a graph, but not in the residual graph but rather in the reduced residual graph.

[8]So the CRCs can be calculated as the connected components (see [7]) of the undirected graph corresponding to the RRG without the sink and the agents edges.

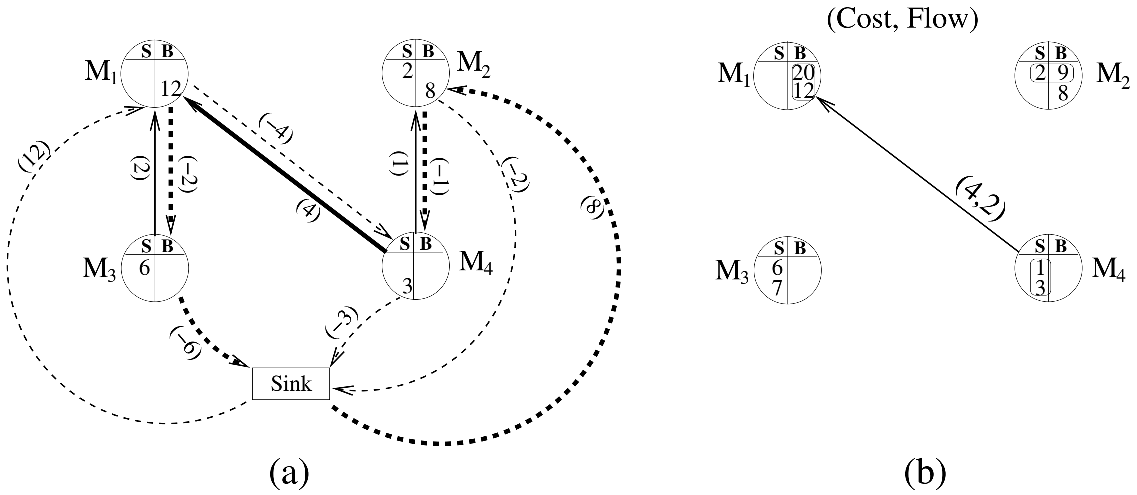

Figure 4: (a) The reduced residual graph for the residual graph presented in Figure 3. The minimal positive cycle is drawn with thicker lines. (b) The trade reduction allocation created by removing the minimal positive cycle.

We mark the distance between nodes $n1$ and $n2$ in the RRG as $\tilde{d}(n1,n2)$. For any two nodes $n1$ and $n2$, $\tilde{d}(n1,n2)$ is well defined since the RRG is a sub graph of the residual graph which has non negative cycles. Let $\tilde{d_i} = \tilde{d}(0,i)$ be the distance from the sink to market M_i in the RRG, and let $\tilde{D_i} = \tilde{d}(i,0)$ be the distance from market M_i to the sink node in the RRG. Note that these distances are not necessarily the same as the distances in the residual graph since some edges (like the non trading agents edges) are removed.

Definition 16 *The TRM payments are defined as follows.*

- *For any market M_i, any trading seller in market M_i receives $-\tilde{D}_i$.*
- *For any market M_j, any trading buyer in market M_j pays $\tilde{d}_j$.*

For a commercial relationship component γ, let the ***trade size*** $|\gamma|$ be the number of buyers (sellers) trading in γ (note that both are the same since there is no trade between different CRCs). We denote by Γ the set of all the CRCs.

Theorem 17 *The TRM mechanism is individually rational, incentive compatible, budget balanced and the efficiency lost of the mechanism in each CRC is at most one over the trade size in the CRC. Formally, if the efficient allocation is non-empty*

$$\frac{\mathbf{V}(A^{TR}(v))}{\mathbf{V}(A^{*}(v))} \geq \min_{\gamma\in\Gamma} \frac{|\gamma|-1}{|\gamma|}$$

where $A^{TR}(v)$ is the TRM allocation for agents with values v.

Assuming that the values/costs in each market are sorted, there exists an algorithm that calculates the mechanism output with running time of $O((m + k\ log\ k)\ m\ log\ (C+n))$.

Proof. See section C.3 in the appendix. □

6 Conclusion and Further Research

In this paper we have considered the problem of spatially distributed market. We have presented the two welfare theorems for SDM. We have characterized the VCG payments both as distances in the residual graph and as the extreme elements of a lattice of prices. We have also presented the TRM mechanism that is IR, IC, BB and has high efficiency. We have presented computationally efficient algorithms for both the VCG and TRM mechanisms. For a fixed number of markets, these algorithms are dependent on the logarithmic of the number of agents.

This work focused on mechanisms where the agents are buyers and sellers in the different markets. In future research more attention should be given to the model with strategic carriers that we have briefly discussed in section B.2. An interesting open question is to find an IR, IC, BB and highly efficient mechanism for this model. Another challenge is to further extend this model to the case that the carriers as well as the buyers and sellers have multi unit demand/supply.

References

[1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[2] Simon Anderson and Maxim Engers. Spatial competition with price-taking firms. *Economica*, 61:125–36, May 1994.

[3] Moshe Babaioff and Noam Nisan. Concurrent auctions across the supply chain. In *Third ACM Conference on Electronic Commerce*, pages 1–10, 2001. Extended version to appear in JAIR, 2004.

[4] Moshe Babaioff and William E. Walsh. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. In *Fourth ACM Conference on Electronic Commerce*, pages 64–75, 2003. Extended version to appear in DSS, 2004.

[5] Leon Y. Chu and Zuo-Jun Max Shen. Dominant strategy double auction with pair-related costs. University of Florida, Working Paper, 2003.

[6] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, Fall 1971.

[7] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. MIT Press/McGraw-Hill, 1990.

[8] Sven de Vries and Rakesh Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, 15(3):284–309, 2003.

[9] Theodore Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.

[10] Frauk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.

[11] Peter H. Lindert and Jeffrey G. Williamson. Does globalization make the world more unequal. NBER 8228 Working Paper, 2001.

[12] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[13] R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56:434–450, 1992.

[14] Paul Milgrom. Putting auction theory to work: The simultaneous ascending auction. *The Journal of Political Economy*, 108(2):245–272, April 2000.

[15] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

[16] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1/2):166–196, April/May 2001.

[17] Rabin Roundy, Rachel Chen, Ganesh Janakriraman, and Rachel Q. Zhang. Efficient auction mechanisms for supply chain procurement. Technical Report 1287, School of Operations Research and Industrial Engineering, Cornell University, 2001.

[18] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.

[19] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[20] William E. Walsh, Michael P. Wellman, and Fredrik Ygge. Combinatorial auctions for supply chain formation. In *Second ACM Conference on Electronic Commerce*, pages 260–269, 2000.

A Background - Minimum Cost Flow

In this section we present the minimum cost flow problem and its solution, as well as some technical characteristics of the conditions for optimality of the flow that will be used to derive our results. Although the results are standard we present them for completeness. Our notation follows the book by Ahuja et al. [1].

A.1 The Minimum Cost Flow Problem and its Solution

In this section we present the minimum cost flow problem and its solution.

A.1.1 The Minimum Cost Flow Problem

Let $G=(V,E)$ be a directed network with cost $c_{i,j}$ and capacity $u_{i,j}\geq 0$ associated with every arc $(i,j)\in E$. We associate with the node $i\in V$ a number b_i which indicates its supply or demand. The ***Minimum Cost Flow Problem (MCFP)*** can be stated as follows:

$$Minimize \sum_{(i,j)\in E} c_{i,j}x_{i,j}$$

subject to

$$\sum_{j:(i,j)\in E} x_{i,j} - \sum_{j:(j,i)\in E} x_{j,i} = b_i \; for\; all\; i\in V$$

$$0\leq x_{i,j}\leq u_{i,j}$$

Where $x_{i,j}$ is the flow on arc $(i,j)\in E$. Let C denote the largest magnitude of any arc cost, and let U denote the largest magnitude of any supply/demand or finite capacity. We assume that all data (cost, supply/demand, capacity) are integral and that $\sum_{i\in V} b_i = 0$. Under these assumptions there is an polynomial algorithm that either states that such a flow does not exist or finds an **integral** minimal cost flow (Theorem 9.10 at [1]).

A ***Convex Minimal Cost Flow Problem (CMCFP)*** is a generalization of MCFP where on each edge (i,j), instead of a constant cost $c_{i,j}$ per unit, there is a convex function, $C_{i,j}(x_{i,j})$. In this work we focus on the case that these functions are piecewise linear with integral break points. We are looking for an integral solution to the problem. We refer the reader to [1] chapter 14 for more details about CMCFP.

A.1.2 Algorithms for MCFP and CMCFP

There is a vast literature on algorithms for solving the MCFP and the CMCFP (see the books [1, 18] for background and references to more literature), in this section we bring the result that we later use. Polynomial time algorithms for both problems are known, and the capacity scaling algorithm by Minoux (see section 14.5 in [1]) solves the CMCFP with the same running time it solves the MCFP.

Theorem 18 *The capacity scaling algorithm outputs an integral solution to an integral CMCFP of graph $G=(V,E)$ in running time $O((|E| log\, U)S(|V|,|E|,C))$, where $S(|V|,|E|,C)$ is the time needed to solve a Shortest Path Problem with non-negative arc lengths with $|V|$ nodes, $|E|$ edges and maximal edge cost of C.*

The Fibonacci heap implementation for Dijkstra's algorithm for the Shortest Path Problem has a strongly polynomial running time of $O(|E|+|V|\ log|V|)$.

A.2 The Residual Graph and Optimality Theorems

We will later need some technical characteristics of flow optimality, for which we need a few more definitions. The ***residual network*** $G(x)$ corresponding to the flow x is defined as follows. We replace each arc $(i,j)\in E$ by two arcs, (i,j) and (j,i). The arc (i,j) has cost $c_{i,j}$ and ***residual capacity*** $u_{i,j}-x_{i,j}$, and the arc (j,i) has a cost $c_{j,i}=-c_{i,j}$ and residual capacity $x_{i,j}$. The residual network consists only of arcs with positive residual capacity.

We associate a real number π_i, unrestricted in sign, with each node $i\in V$, and refer to this as the node ***potential*** (these potentials are parameters of the dual problem, we refer the interested reader to [1] for details). We define the ***reduced cost*** of an arc $(i,j)\in E$ as $c^{\pi}_{i,j}=c_{i,j}-\pi_i+\pi_j$.

The following observation, which is a direct result of a telescopic summation, will later be helpful.

Observation 19 *(Property 9.2 at [1]) For any directed path P from node r to node l, $\sum_{(i,j)\in P} c^{\pi}_{i,j} = \sum_{(i,j)\in P} c_{i,j} - \pi_r + \pi_l$*

The following theorems characterize necessary and sufficient conditions for optimality of the flow. We cite the theorems without proofs, we refer the reader to [1] for proofs. The following two theorems characterize optimality on the residual graph.

Theorem 20 *(Theorem 9.1 at [1] - Negative Cycle Optimality Conditions) A feasible solution x^* is an optimal solution of the minimum cost flow problem if and only if is satisfies the negative cycle optimality conditions: namely, the residual network $G(x^*)$ contains no negative cost (directed) cycle.*

Theorem 21 *(Theorem 9.3 at [1] - Reduced Cost Optimality Conditions) A feasible solution x^* is an optimal solution of the minimum cost flow problem if and only if some set of node potentials π satisfy the following reduced cost optimality conditions: $c^{\pi}_{i,j}\geq 0$ for all arc (i,j) in $G(x^*)$.*

The next theorem characterizes optimality on the original graph.

Theorem 22 *(Theorem 9.4 at [1] - Complementary Slackness Optimality Conditions) A feasible solution x^* is an optimal solution of the minimum cost flow problem if and only if for some set of node potentials π, the reduced costs and the flow values satisfy the following complementary slackness optimality conditions for ever arc $(i,j) \in E$:*

- *If $c^{\pi}_{i,j} > 0$ then $x^*_{i,j} = 0$.*
- *If $c^{\pi}_{i,j} < 0$ then $x^*_{i,j} = u_{i,j}$.*
- *$0 < x^*_{i,j} < u_{i,j}$ then $c^{\pi}_{i,j} = 0$.*

Note that the last condition is redundant and can be inferred from the previous two and the fact that $x_{i,j} \geq 0$.

The last theorem we cite enables decomposition of flow to cycles.

Theorem 23 *(Theorem 3.7 at [1] - Augmenting Cycle Theorem) Let x and x^o be any two feasible solutions of a network flow problem. Then x equals x^o plus the flow of at most m directed cycles in $G(x^o)$. Furthermore, the cost of x equals the cost of x^o plus the cost of flow on the augmenting cycles.*

B Extended Models and Their VCG Payments Characterization

Below we present two extensions to our model. We characterize the VCG payments for these extended models, and present algorithms to calculate the VCG mechanism.

B.1 Agents with multi unit demand and supply

A natural generalization of the model we have presented above, is a model where buyers and sellers can bid for multiple units in their markets. In this section we characterize the VCG payments in the case that each buyer has a decreasing valuation function, and each seller has an increasing cost function. So for example, a buyer can bid to buy the first unit for \$10, the second for \$7, and the third for \$3. This constraint is a natural one and can be thought of as the law of diminishing returns.

The allocation for this more general model is calculated by reduction to the original formulation. We find the minimum cost flow in a graph where the bid for each of the units of each of the agents is handled as a different agent in our original formulation. We can do that since each agent has a decreasing value (increasing cost) function.

More care should be taken when we calculate prices, since we do not want an agent to manipulate her payment by changing her losing bids. Each agent should pay the "harm" done by her bid to the other agents (real agents). If a seller in market M_i sells q units of the good, then the price she should pay is the cost of shipping additional q units of the good from the sink node to her market, without taking into account her losing bids. Similarly a buyer which buys q units of the good pays the cost of shipping q units from her market to the sink node, without taking into account her losing bids.

We note that these payments can be calculated faster in the following way. Assume that Q is the maximal number of units sold by any single seller in market M_i. For any $q \leq Q$ we calculate the minimal cost of shipping additional q units from the sink node to market M_i, not using agents in market M_i. This is done by greedily picking q lowest cost paths from the sink node to market M_i. The price that a seller the sells q units is the price of the q lowest paths from the sink node to market M_i, picked from both the paths that we calculated from M_i to sink and the direct edges of other agents from M_i to sink (trading sellers and non trading buyers). Similar calculations can be done for the buyers.

B.2 Carriers bidding for shipment privilege

Another natural generalization of the model we have presented above, is a model where the cost of shipment between markets is not exogenous. In this model there are also carriers bidding for the privilege of performing a shipment of a unit from one market to the other. Each carrier is actually selling a shipment of one unit of good between the two locations, and she reports her cost of that shipment to the auction before it runs.

The efficient allocation is calculated by solving the same MCFP as before, after replacing the edge from market M_i to market M_j by multiple edges of capacity one and the costs reported by the shipment agents. Note that now the edges between the markets are capacitated. Also note that this is also a CMCFP, so the algorithm presented in section C.2.3 still works.

The prices for the sellers and buyers in each market are calculated as before. The payment that each shipment agent from market M_i to market M_j pays is the minimal distance $d(i,j)$ in the residual graph. This is the alternative shipment route to that agent. Note that this payment is the same for all shipment agent from market M_i to M_j, so we only need to calculate this once per edge. All these prices can be calculated by solving one All Pairs Shortest Path problem. This problem can be solved by many different algorithms, for example the Floyd-Warshall algorithm runs in time $\Theta(|V|^3)$ (see [7]) which is less then the allocation calculation running time. This means that the mechanism for this extended model has the same running time as the mechanism for the basic model, this running time was presented in Proposition 32.

Naturally, the results from section B.1 and this section can be combined to an efficient mechanism for the case of agents with decreasing value (convex value) function for multiple units, and shipping agents having an increasing cost function for shipment privilege on the edges.

C Proofs

C.1 The Two Welfare Theorems

Theorem 24 *(First Welfare Theorem for SDM) If an allocation* $\mathbf{A}$ *and price vector* $\vec{p}$ *are in a spatial price equilibrium then the allocation is efficient.*

Proof. The proof of the theorem follows from the complementary slackness optimality conditions of Theorem 22, where we define the potentials as $\pi_i = -p_i$ for $1 \leq i \leq k$ and $\pi_0 = 0$ (the sink has a potential of 0 for normalization purposes).

We denote the shipment vector of the allocation by x. We should verify that the complementary slackness optimality conditions holds for any edge in the graph.

We first consider the edges between two markets. From SPE, for any arc $(M_i, M_j) \in E$:

- $p_i + c_{i,j} \geq p_j$ therefore $c^{\pi}_{i,j} = c_{i,j} - \pi_i + \pi_j \geq 0$.
- if $x_{i,j} > 0$ then $p_i + c_{i,j} = p_j$ so $c^{\pi}_{i,j} = c_{i,j} - \pi_i + \pi_j = 0$.

We conclude that the complementary slackness optimality conditions holds for any edge between two markets.

We now consider the buyers edges (from the buyer market to the sink). From SPE, for any market M_j and buyer b (with value b) in that market, if $b > p_j$ then b buys a unit, and if $b < p_j$ she does not buy a unit. This means that for the buyer edge e_b from her market to the sink with cost $-b$ we have that:

- If $c^{\pi}_{e_b} = -b - (-p_j) + 0 < 0$ then the flow on e_b is 1 and equals the capacity.

- If $c^{\pi}_{e_b} = -b - (-p_j) + 0 > 0$ then the flow on e_b is 0.

So the complementary slackness optimality conditions holds for any buyer edge. Similar argument shows that it also holds for any seller edge.

So the complementary slackness optimality conditions holds for any edge and by Theorem 22 we conclude that the flow is optimal and therefore the corresponding allocation is efficient. □

The Second Welfare Theorem for SDM is a result of the following lemma.

Lemma 25 *If an allocation* **A** *is efficient then it is in a SPE with the price vector $\vec{p} = \vec{d}$ ($p_i = d_i$, where d_i is the distance from the sink to the market node in the residual graph of the flow corresponding to the allocation* **A***).*

Proof. We look at the (optimal) flow x^* which corresponds to the allocation **A**. The price in market M_i is defined to be $p_i = d_i$, the distance from the sink to the market node in the residual graph (this distance is well defined since there are no negative cost cycles by Theorem 20). For the sink this distance is 0. From the shortest path optimality conditions, for any edge (M_i, M_j) in $G(x^*)$ we have $d_j \le d_i + c_{i,j}$, therefore $p_j \le p_i + c_{i,j}$. If $x_{i,j} > 0$ then both the edge and the residual reversed edge with negated cost are in the residual graph, so $p_j \le p_i + c_{i,j}$ and $p_i \le p_j - c_{i,j}$ therefore $p_j = p_i + c_{i,j}$. So the SPE conditions for the markets hold.

We now turn to show that the SPE conditions for the agents also hold. We should show that for any market M_i and agent w with value v_w in that market, if $v_w > p_i$ then a unit is assigned to w, and if $v_w < p_i$ then no unit is assigned to w.

In the residual graph, for any agent w in market M_i with value v_w, there is an edge with cost v_w from the sink to M_i if a unit is assigned to her (remains in the seller hands or sold to the buyer), and an edge with with cost $-v_w$ from M_i to the sink if no unit is assigned to her.

The distance from the sink to M_i is d_i, therefore if there is an edge from the sink to M_i with cost v_w then $d_i \le v_w$ by the definition of distance. We conclude that if $p_i = d_i > v_w$ then no unit is assigned to w.

We now show that if $v_w > p_i = d_i$ then a unit is assigned to w. Assume that no unit is assigned to w, then there is an edge with cost $-v_w$ from M_i to the sink. The cycle along the shortest path from the sink to M_i and back to the sink on the edge with cost $-v_w$ has a cost of $d_i + (-v_w) < 0$ which is contradiction to the fact that the residual graph has no negative cycles (Theorem 20).

Since all conditions for SPE holds, we conclude that the allocation with the defined prices are in SPE. □

C.2 VCG Payments and Algorithms

In this section we present the proofs of our double characterizations for SDM VCG payments, we also present a computationally efficient algorithm for the VCG mechanism.

Note that Theorem 26 implies that all trading sellers in the same market receive the same payment, and all trading buyers in the same market pay the same (but, a buyer typically pays a different price than a seller in the same market receives). So, for any trading seller in M_i, we denote the value she pays as $VCG^s_i(v)$, and for any trading buyer in M_j, we denote the value she pays as $VCG^b_j(v)$.

C.2.1 VCG Payments Characterization as Distances

Theorem 26 *The VCG payments for trading agents are the following.*

- *For any market M_i, any trading seller in market M_i receives d_i, that is $VCG^s_i(v) = -d_i$.*

- *For any market M_j, any trading buyer in market M_j pays $VCG_j^b(v) = -D_j$.*

Proof. The proof of the theorem is a result of supply/demand sensitivity analysis for the MCFP. Removing a trading seller/buyer causes the market of the agent to have a deficit/surplus of one unit. In order to retain feasibility, a unit must be shipped from/to the sink node to/from this market (by cost minimization the rest of the flow remains unchanged).

We first show that for any trading seller in market M_i, $VCG_i^s(v) = -d_i$. By the definition of the VCG payment, the VCG payment of the agent is the change in the other agents total value, caused by the bid of that agent. Removing a trading seller in market M_i creates a deficit of one unit of good in market M_i. Covering this deficit requires pushing an additional flow of one unit from the sink node to market M_i. In order to minimize the cost, this flow must be pushed along the shortest path from the sink to market M_i in the residual graph. Therefore a seller in market M_i receives d_i.

Similar argument shows that for any trading buyer in market M_j, $VCG_j^b(v) = -D_j$. The extra unit in market M_j that is left if we remove the trading buyer must flow to the sink node along a shortest path in order to create a minimum cost flow, therefore the payment by the buyer is as stated. □

C.2.2 VCG Payments Characterization as Lattice Elements

Theorem 27 *The set of price vectors that are in SPE with the efficient allocation form a complete lattice.*

- *The vector of VCG prices that the buyers pay, $\overrightarrow{VCG^b(v)}$, is the minimal element of the lattice.*
- *The vector of VCG prices paid to the sellers, $-\overrightarrow{VCG^s(v)}$ is the maximal element of the lattice.*

Proof. By Lemma 28 the set of price vectors that are in SPE with the efficient allocation form a complete lattice. By Lemma 25, $-\overrightarrow{VCG^s(v)} = \vec{d}$ is an element of the lattice, similar argument to the one presented in that lemma shows that $\overrightarrow{VCG^b(v)} = -\vec{D}$ is also an element of the lattice. By Lemma 29 $\overrightarrow{VCG^b(v)}$ is the minimal element of the lattice and $-\overrightarrow{VCG^s(v)}$ is the maximal element of the lattice. □

Lemma 28 *The set of price vectors that are in SPE with the efficient allocation form a complete lattice.*

Proof. Assume that $\vec{p}$ and $\vec{q}$ are two vectors that are in SPE with the efficient allocation. We show that $\vec{p} \wedge \vec{q}$ and $\vec{p} \vee \vec{q}$ are also in SPE with the efficient allocation. By SPE of $\vec{p}$ and $\vec{q}$ we know that for any two markets M_i, M_j:

1. $p_i + c_{i,j} \geq p_j$ and $q_i + c_{i,j} \geq q_j$.
2. If $x_{i,j} > 0$ then $p_i + c_{i,j} = p_j$ and $q_i + c_{i,j} = q_j$.

From (1) we know that $\max\{p_i, q_i\} + c_{i,j} \geq \max\{p_j, q_j\}$. From (2) we know that if $x_{i,j} > 0$ then $p_i > q_i$ iff $p_j > q_j$. Therefore if $x_{i,j} > 0$ then $\max\{p_i, q_i\} + c_{i,j} = \max\{p_j, q_j\}$. We conclude that both $\vec{p} \wedge \vec{q}$ and the efficient allocation are in SPE. A similar proof works for $\vec{p} \vee \vec{q}$.

The claim that the lattice is complete is a direct result from the fact that if $p_i^w + c_{i,j} \geq p_j^w$ for any $w \in W$ then $\inf_{w \in W} p_i^w + c_{i,j} \geq \inf_{w \in W} p_j^w$, and if $p_i^w + c_{i,j} = p_j^w$ for any $w \in W$ then $\inf_{w \in W} p_i^w + c_{i,j} = \inf_{w \in W} p_j^w$. □

Next we show that for any price vector that is in SPE with the efficient allocation, each trading seller receives the maximal equilibrium price in her market and each trading buyer pays the minimal equilibrium price in her market.

Lemma 29 *Let $\vec{p}$ be a price vector in SPE with the efficient allocation.*

- *For any market M_i, $p_i \leq -VCG_i^s(v)$.*
- *For any market M_j, $p_j \geq VCG_j^b(v)$.*

Proof. Since $\vec{p}$ is in SPE with the efficient allocation, then the node potentials defined as $\pi = -\vec{p}$ (see Theorem 6) satisfy the reduced cost optimality conditions presented in Theorem 21. This means that $c^{\pi}_{i,j} \geq 0$ for all arc (i,j) in $G(x^*)$, and therefore for any path P, $\sum_{(i,j)\in P} c^{\pi}_{i,j} \geq 0$. By Observation 19, for any directed path P from node r to node l,

$$0 \leq \sum_{(i,j)\in P} c^{\pi}_{i,j} = \sum_{(i,j)\in P} c_{i,j} - \pi_r + \pi_l$$

We assume that losing agents pay 0, so the potential of the sink node is 0, that is $\pi_0 = 0$. Applying the above observation on the sink market ($r = 0$) we conclude that for any market M_l and any path P from the sink to market M_l, $\sum_{(i,j)\in P} c_{i,j} \geq -\pi_l$. In particular, $-\pi_l$ is bounded from above by the shortest path from the sink to market M_l (with the costs as edge lengths), so $d_l \geq -\pi_l$. We conclude that for every $\vec{p}$ in equilibrium and for any market l, $-VCG_l^s(v) = d_l \geq -\pi_l = p_l$. This means that the VCG price received by the seller is the maximal possible equilibrium price in her market (the potential, which she receives, is the negative of the price in her market, which is what she pays).

Similarly, if we take $l = 0$, then for any market M_r and any path P from market M_r to the sink node 0, $\sum_{(i,j)\in P} c_{i,j} \geq \pi_r$. In particular, π_r is bounded from above by the shortest path from market M_r to the sink node, so $D_r \geq \pi_r$. We conclude that for every $\vec{p}$ in equilibrium and for every market r, $VCG_r^b(v) = -D_r \leq -\pi_r = p_r$. This means that the VCG price that a buyer pays is the minimal possible equilibrium price in her market. □

C.2.3 Efficient Algorithm for VCG Mechanism

In section 3 we have seen that the efficient allocation for SDM can be calculated by solving a minimal cost flow problem. The MCFP we have built can be presented as a CMCFP, since we can replace the multiple edges between the sink and each of the market nodes with a single edge with a piecewise linear convex cost function. The value of the function for t units of flow is the t lowest cost (highest value) bid in that market. The reason that the function is convex is that the minimum cost flow always picks lower cost agents before higher cost agents, as otherwise the cost can be reduced. We are assuming that the values/costs in each market are sorted, and do not add the time needed for sorting when we state the algorithm running time.

The graph of the CMCFP we build is a simple graph, it has $k+1$ nodes and $m+2k$ edges, which by our assumption that $k \leq m$ is $O(m)$ edges. The maximal edge cost is $C = \max\{\max_{w\in N} |v_w|, \max_{(M_i,M_j)\in E} c_{i,j}\}$. The capacity of the edges between the markets is infinite, and the capacity of the edges from/to the sink is essentially bounded by n (the number of agents). The largest magnitude of any node supply/demand is also bounded by n (it is not zero since supply/demand in the nodes is created when the graph is transformed to a graph with non-negative edges. The magnitude of any supply/demand is at most n), so $U = \max\{C, n\}$ which means that $U = O(C+n)$.

From the results for CMCFP (Theorem 18) we derive the following corollary.

Corollary 30 *The efficient allocation for a SDM can be calculated by solving a single convex minimal cost flow problem on a graph with $k+1$ nodes, $m+2k$ edges, maximal cost $C = \max\{\max_{w\in N} |v_w|, \max_{(M_i,M_j)\in E} c_{i,j}\}$ and largest magnitude of any node supply/demand or finite capacity of $U = \max\{C, n\}$.*

There exists an algorithm that outputs the efficient allocation that runs in time $O((m + k \log k)\, m \log (C + n))$.

This algorithm is polynomial in k (since $m \leq k^2$), $log\ C$ and $log\ n$. Note that the running time grows polynomially in $log\ n$ and not in n. This is important since it is reasonable to assume that the number of agents is significantly greater then the number of markets.

To calculate the prices we need to find the shortest path from each market to the sink and from the sink to each of the markets in the residual graph. Note that can remove all edges of agents that are not the lowest value trading agent and the highest value no-trading agent in each market. So we need to solve two single source shortest path problems with arbitrary edge lengths. The first calculates all the sellers payments by running the algorithm from the sink node on the residual graph. The second calculates all the buyers payments by running the algorithm from the sink node in the reverse residual graph (after reversing the direction of all arcs in the residual graph). Note that the shortest path from the sink to each market in the reversed graph has the same length as the shortest path from the market to the sink node in the residual graph.

The FIFO implementation of the label correcting algorithm for the single source shortest path problems has a running time of $O(|V||E|)$ (see [1] chapter 5).

We summarize these in the following corollary.

Corollary 31 *Given the efficient allocation, the VCG payments can be calculated by solving two single source shortest path problems with graph of $k+1$ nodes and $m+2k$ edges.*

Given the efficient allocation the VCG payments can be calculated in time $O(m\ k)$.

Note that the efficient allocation calculation dominates the payments calculation. So from Corollary 31 and Corollary 30 we conclude the following proposition.

Proposition 32 *There exists an algorithm that calculates the VCG mechanism for SDM that runs in time $O((m+k\ log\ k)\ m\ log\ (C+n))$.*

C.3 The TRM properties

In this section we prove Theorem 17. In section C.3.1 we prove that the TRM auction is IR, IC and BB, and in section C.3.2 we prove that it is highly efficient. In section C.3.3 we present and analyze an algorithm that calculates the TRM mechanism.

C.3.1 TRM is IR, IC and BB

In this section we prove that the TRM auction is IR, IC and BB. Note that any incentive compatible mechanism that is normalized (losing agents pay 0) is also individually rational, since for any agent, bidding truthfully ensures non negative utility. So for our normalized mechanism, to prove IR it is sufficient to prove that for each agent, truthful bidding maximizes her utility (IC). After proving that the TRM is IC, we also prove that the TRM is BB.

We use a well known characterization of IC mechanisms in the case that all agents are single parameter agents (the valuation of each agent is fully described by a single parameter, which in our case is her valuation for trading) We present the characterization in the context of SDM. We start with a definition. A allocation rule is ***bid monotonic*** if for any agent w, if w trades when she bids x, then she also trades if she bids $y > x$ (she never becomes a non trading agent by improving her bid).

Proposition 33 *If an allocation rule is bid monotonic then for each agent w there exists a **critical value** CV_w, such that if w bids more than CV_w she trades and if w bids less than CV_w she does not trade (where the bids of all other agents are fixed).*

Now we are ready to state the characterization of IC mechanisms for single parameter agents. The characterization gives necessary and sufficient conditions for IC.

Theorem 34 *A normalized mechanism for single parameter agents is IC if and only if its allocation rule is bid monotonic and the payment of each agent is her critical value for trading.*

Proof. Case if: Assume that the mechanism is normalized, and the allocation rule is bid monotonic. Trading agents pay their critical value for trading, so any agent w with value v_w pays CV_w if she trades (wins).

To prove IC we prove that w cannot improve her utility by misreporting her value. Consider the case that agent w wins the auction by bidding her true value v_w. Her utility is non-negative since by the definition of the critical value, $v_w \geq CV_w$, hence her utility is $v_w - CV_w \geq 0$. If w bids untruthfully and loses, then she gets zero utility, which is not better than her utility with a truthful bid. If w bids untruthfully and wins the auction, then since she still pays her critical value CV_w her utility remains the same.

Now consider the case in which w loses the auction by bidding truthfully. Her utility is zero and $v_w \leq CV_w$ by the definition of the critical value. If w bids untruthfully and loses, her utility remains zero. If w bids untruthfully and wins, her utility is $v_w - CV_w \leq 0$.

In both cases, we have shown that w agent cannot improve her utility by bidding untruthfully, thus proving that the mechanism is IC.

Case only if: Since we do not use this in the paper, we do not present the proof. □

The TRM mechanism is normalized. So in order to prove that it is IC we should show that it is bid monotonic and that the payments are by critical values.

Proposition 35 *The TRM allocation rule is bid monotonic.*

Proof. We should show that a trading agent never becomes a loser by improving her bid. The efficient allocation remains the same if agent w improves her bid. Therefore the CRC of w remains the same. If w is not is the minimal positive cycle in her CRC in the first place, then it is not in this cycle if she improves her bid, so she still trades as we wanted to prove. □

So to prove that the mechanism is IC, it is now sufficient to prove that the payments are by the critical values.

Lemma 36 *The critical values for the agents in the TRM mechanism are the following.*

- *For any market M_i, the critical value for any trading seller in market M_i is $\tilde{D}_i$.*
- *For any market M_j, the critical value for any trading buyer in market M_j is $\tilde{d}_j$.*

Proof. The proof of this lemma is a direct result of Lemma 41 and Lemma 42. □

To prove Lemma 41 and Lemma 42 we need a few observations.

Lemma 37 *For any two markets M_i, M_j in the same CRC, there exists a shortest path between M_i and M_j in the RRG that does not pass through the sink and is also a shortest path between M_i and M_j in the residual graph. Therefore, for any two markets M_i and M_j in the same CRC, $d(i,j) = \tilde{d}(i,j)$.*

Proof. Let $P_{i,j}$ be a shortest path between market M_i and M_j in the residual graph, by definition its length is $d(i,j)$. Let $P'_{i,j}$ be a shortest path between market M_i and M_j in the RRG, by definition its length is $\tilde{d}(i,j)$. Let $\tilde{P}_{i,j}$ be a shortest path between market M_i and M_j in the RRG not passing through the sink node. We mark the length of this path as $|\tilde{P}_{i,j}|$.

Since the capacity of the edges between the markets is infinite, then for each edge along the path $\tilde{P}_{i,j}$ in the RRG, there is a reversed edge of negated cost in the RRG. Therefore there is a path $\tilde{P}_{j,i}$ of cost $-|\tilde{P}_{i,j}|$ from M_j to M_i in the RRG. If $|\tilde{P}_{i,j}| > \tilde{d}(i,j)$ then the cycle along $P'_{i,j}$ and $\tilde{P}_{j,i}$ has a cost $\tilde{d}(i,j) - |\tilde{P}_{i,j}| < 0$ which is a contradiction to Theorem 20 that states that there are no negative cycles in the residual graph (and therefore in the RRG). We conclude that there is a shortest path ($\tilde{P}_{i,j}$) from M_i to M_j in the RRG that does not pass through the sink, and $\tilde{d}(i,j) = |\tilde{P}_{i,j}|$.

Suppose that $d(i,j) \neq \tilde{d}(i,j)$. Since the RRG is a sub graph of the residual graph, it is impossible that $d(i,j) > \tilde{d}(i,j)$. Now assume that $d(i,j) < \tilde{d}(i,j)$. Any path in the RRG also exists in the residual graph, therefore if $d(i,j) < \tilde{d}(i,j)$ then there is a cycle of cost $d(i,j) - \tilde{d}(i,j) < 0$ (along $P_{i,j}$ and $\tilde{P}_{j,i}$) in the residual graph. Again, this contradicts the fact that there are no negative cycles in the residual graph. □

Note that under the assumption that no two allocations have the same value, there is a unique shortest path (both in the residual graph and in the RRG) between any two markets in the same CRC. So from the lemma we conclude that the shortest path in the RRG between any two markets does not pass through the sink node.

A simple consequence of the previous lemma is the following:

Corollary 38 *For any two markets M_i, M_j in the same CRC, $\tilde{d}(j,i) = -\tilde{d}(i,j)$.*

Proof. Assume in contradiction that $\tilde{d}(j,i) > -\tilde{d}(i,j)$. By Lemma 37 there exist a shortest path P between M_i and M_j of length $\tilde{d}(i,j)$ that does not pass through the sink. Since the capacity of the edges between the markets is infinite, there is a reverse path P' from M_j to M_i of length $-\tilde{d}(i,j)$. So we have found a path P' from M_j to M_i with length shorter then the distance between the two markets, which is a contradiction. Similar argument shows that it is never the case that $\tilde{d}(j,i) < -\tilde{d}(i,j)$ □

In the following, we abuse the notation and use the same marking for a trading agent, the market she belongs to, and to the cost of the residual edge (with capacity 1) corresponding to that agent (the meaning will be clear from the context). For example, a trading buyer C_b corresponds to a residual edge from the sink to market C_b (the market of that buyer), and has a cost of C_b. Note that the cost C_b equals to the buyer's value for a unit of the good. A seller C_s corresponds to a residual edge from market C_s (the market of that seller) to the sink, and has a cost of C_s. Note that the cost C_s equals to the negative of the seller's cost for her unit. The distance in the RRG between the markets of buyer C_b and seller C_s will be marked as $\tilde{d}(C_b, C_s)$.

We now look at the structure of the minimal positive cycle. We start with the following lemma:

Lemma 39 *For any CRC γ, the minimal positive cycle in γ visits the sink node exactly once.*

For any CRC γ, the minimal positive cycle in γ contains one (residual edge corresponding to the) buyer C_b and one (residual edge corresponding to the) seller C_s, both of them trade in the efficient allocation. The cost of the minimal positive cycle in γ is $C_b + \tilde{d}(C_b, C_s) + C_s$.

Proof. We say that a cycle is a ***true cycle*** in the residual graph, if the cycle is not constructed from a path and its reversed (negated cost) path.

By our assumption that there are no two allocation with the same value, there are no zero cost true cycles in the residual graph, so the RRG without the sink has no zero cost true cycles. There are no negative cycles in the RRG (By Theorem 20). There are also no positive cycles in the RRG without the sink node, since if such cycle exists, the reversed negated cycle exists in the RRG, and it has a negative cost which is again a contradiction. This means that the RRG without the sink node has no true cycles (of any cost). Therefore, any cycle in the RRG which does not visit the sink node, is not a true cycle and has zero cost by Corollary 38. We conclude that for any CRC, the minimal positive cycle visits the sink node at least once.

Any cycle that visits the sink node more than once can be split to sub cycles, each starting and ending at the sink node. From the above we conclude that if the minimal positive cycle in a CRC visits the sink more than once, it can be split to sub cycles, each starting and ending at the sink node, and each cycle has a strictly positive cost Each of those cycles has a lower cost than the original cycle. This contradicts our assumption that the cycle is the minimal positive cycle in its CRC. We conclude that the minimal positive cycle of any CRC visits the sink exactly once.

Recall that all the edges in the RRG leaving the sink node corresponds to trading buyers, and all the edges in the RRG arriving to the sink node corresponds to trading sellers. So, the minimal positive cycle of any CRC includes a single residual edge corresponding to a trading buyer and a single residual edge corresponding to a trading seller.

The cost of the minimal positive cycle in a CRC γ that includes buyer C_b and seller C_s is $C_b + \tilde{d}(C_b, C_s) + C_s$, since any sub path of a minimal cycle must be of length equal to the distance between its source and destination nodes, otherwise it can be shorten. □

We now characterize the distances in the RRG. We show that for any CRC γ that includes buyer C_b and seller C_s in its minimal positive cycle, the shortest path to any market in that CRC ha to pass through the edge C_b, and the shortest path from any market in that CRC pass through the edge C_s.

Lemma 40 *Let γ be a CRC with a minimal positive cycle that includes buyer C_b and seller C_s.*

For any market B in γ,

$$\tilde{d}(Z,B) = C_b + \tilde{d}(C_b,C_s) + \tilde{d}(C_s,B) = C_b + \tilde{d}(C_b,B)$$

For any market S in γ,

$$\tilde{d}(S,Z) = \tilde{d}(S,C_b) + \tilde{d}(C_b,C_s) + C_s = \tilde{d}(S,C_s) + C_s$$

Proof. For any market B in γ, $\tilde{d}(Z,B) = C_b + \tilde{d}(C_b,C_s) + \tilde{d}(C_s,B)$ since

- From the definition of distance, $\tilde{d}(Z,B) \leq C_b + \tilde{d}(C_b,C_s) + \tilde{d}(C_s,B)$, since the path from the sink (Z) to market B through the edge with residual cost C_b, and through market C_s, is a possible path.
- From the minimality of the cycle including edges C_b and C_s, $\tilde{d}(Z,B) + \tilde{d}(B,C_s) + C_s \geq C_b + \tilde{d}(C_b,C_s) + C_s$. By reducing $\tilde{d}(B,C_s) + C_s$ from both sides and noting that $-\tilde{d}(B,C_s) = \tilde{d}(C_s,B)$ (by Corollary 38), we get $\tilde{d}(Z,B) \geq C_b + \tilde{d}(C_b,C_s) + \tilde{d}(C_s,B)$.

Finally note that $\tilde{d}(C_b,C_s) + \tilde{d}(C_s,B) = \tilde{d}(C_b,B)$ since any sub path of a shortest path is also a shortest path.

The proof that for any market S in γ, $\tilde{d}(S,Z) = \tilde{d}(S,C_b) + \tilde{d}(C_b,C_s) + C_s = \tilde{d}(S,C_s) + C_s$, is similar to the above proof. □

We are now ready to prove that the distances stated in Lemma 36 are indeed the critical values for trading in the TRM.

Lemma 41 *For any buyer B that trades in the TRM mechanism, $\tilde{d}(Z,B)$ is her critical value for trading.*

Proof. We should show that if buyer B bids a value of B (fixing the other agents bids) then,

- if $B > \tilde{d}(Z,B)$ then B trades.

- if $B < \tilde{d}(Z,B)$ then B does not trade.

Since B trades in the efficient allocation, then there is an edge of cost L_B with $L_B \leq B$ from the sink to the market of B (the edge L_B corresponds to the lowest value trading buyer in the market of B). By the definition of distance, $B \geq L_B \geq \tilde{d}(Z,B)$. We conclude that if $B < \tilde{d}(Z,B)$ then B does not trade.

On the other hand, let γ be the CRC of B with minimal positive cycle including C_b and C_s. If $B > \tilde{d}(Z,B)$, then $B > C_b + \tilde{d}(C_b,C_s) + \tilde{d}(C_s,B)$ (by Lemma 40). By adding $\tilde{d}(B,C_s) + C_s$ to both sides and noting that $\tilde{d}(B,C_s) = -\tilde{d}(C_s,B)$ (by Corollary 38), we get

$$B + \tilde{d}(B,C_s) + C_s > C_b + \tilde{d}(C_b,C_s) + C_s \tag{2}$$

The minimal positive cycle including B must include a seller in the same CRC γ (by Lemma 39). For any seller S' in γ (with residual edge cost S') that trades in the efficient allocation, since the cycle including C_s and C_b is the minimal positive cycle,

$$C_b + \tilde{d}(C_b,C_s) + \tilde{d}(C_s,B) + \tilde{d}(B,S') + S' \geq C_b + \tilde{d}(C_b,C_s) + C_s$$

Therefore by adding $-C_b - \tilde{d}(C_b,C_s) + B + \tilde{d}(B,C_s)$ and noting that $\tilde{d}(B,C_s) = -\tilde{d}(C_s,B)$ (by Corollary 38), we get

$$B + \tilde{d}(B,S') + S' \geq B + \tilde{d}(B,C_s) + C_s \tag{3}$$

By combining equations 2 and 3, for any seller S' in γ trading in the efficient allocation

$$B + \tilde{d}(B,S') + S' > C_b + \tilde{d}(C_b,C_s) + C_s$$

We conclude that if $B > \tilde{d}(Z,B)$, then for any seller S' in γ trading in the efficient allocation, the cycle including B and S' is not the minimal positive cycle. Therefore B also trades in the TRM allocation. □

Lemma 42 *For any seller S that trades in the TRM mechanism, $\tilde{d}(S,Z)$ is her critical value for trading.*

Proof. We omit the proof since it is similar to the proof of Lemma 41. □

We have proved that the TRM is IC, thus also IR. We now prove that the TRM mechanism is budget balanced (i.e, has a non negative budget).

Lemma 43 *The TRM mechanism is budget balanced.*

Proof. To prove budget balance it suffices to prove that every trade yields a non-negative net balance. We look at one trade between seller in market M_i and buyer in market M_j and show that it has a non-negative net balance. The trading seller in market M_i pays $\tilde{D}_i$, and the trading buyer in market M_j pays $\tilde{d}_j$. The total balance from the trade including the shipment cost of $\tilde{d}(i,j)$ is $\tilde{D}_i + \tilde{d}(i,j) + \tilde{d}_j = \tilde{d}(0,i) + \tilde{d}(i,j) + \tilde{d}(j,0)$. But this is the weight of a cycle in the reduced residual graph (going from the sink to market M_i, then to M_j and back to the sink)! Now recall that there are no negative weight cycles in the reduced residual graph (since there are no negative weight cycles in the residual graph by Theorem 20), therefore the total balance is non-negative. □

C.3.2 TRM efficiency bound

In this section we prove the lower bound on the efficiency of the TRM. The proof is based on a similar proof presented by Babaioff and Walsh in [4], where they bound the efficiency of a trade reduction mechanism for a supply chain model.

We look at a commercial relationship component $\gamma \in \Gamma$ as a restriction of the residual graph to the sink and the markets of the CRC. Let $\mathbf{V}(\mathbf{A})|_\gamma$ be the value of the allocation $\mathbf{A}$ in the sub graph of $\gamma \in \Gamma$, that is $\mathbf{V}(\mathbf{A})|_\gamma \equiv \sum_{w\in \mathbf{A}\cap\gamma} v_w - \sum_{(M_i,M_j)\in E\cap\gamma} x_{i,j}c_{i,j}$.

Lemma 44 *Let v be any vector of agents values with non-empty efficient allocation $A^*(v)$ and trade reduction allocation $A^{TR}(v)$, then:*

$$\frac{\mathbf{V}(A^{TR}(v))}{\mathbf{V}(A^*(v))} \geq \min_{\gamma\in\Gamma} \frac{|\gamma|-1}{|\gamma|}$$

Proof. We partition the allocation by the CRCs. We have $\mathbf{V}(A^*(v)) = \sum_{\gamma\in\Gamma} \mathbf{V}(A^*(v))|_\gamma$ and that $\mathbf{V}(A^{TR}(v)) = \sum_{\gamma\in\Gamma} \mathbf{V}(A^{TR}(v))|_\gamma$. For each CRC $\gamma \in \Gamma$ we also have that $\mathbf{V}(A^*(v))|_\gamma \geq \mathbf{V}(A^{TR}(v))|_\gamma \geq 0$. By applying Lemma 45 and Lemma 46 we get:

$$\frac{\mathbf{V}(A^{TR}(v))}{\mathbf{V}(A^*(v))} = \frac{\sum_{\gamma\in\Gamma}\mathbf{V}(A^{TR}(v))|_\gamma}{\sum_{\gamma\in\Gamma}\mathbf{V}(A^*(v))|_\gamma} \geq \min_{\gamma\in\Gamma} \frac{\mathbf{V}(A^{TR}(v))|_\gamma}{\mathbf{V}(A^*(v))|_\gamma} \geq \min_{\gamma\in\Gamma} \frac{|\gamma|-1}{|\gamma|}$$

which is what we wanted to prove. $\square$

Lemma 45 *For any set of indexes m and pairs R_m and O_m such that $0 \leq R_m \leq O_m$:*

$$\frac{\sum_m R_m}{\sum_m O_m} \geq \min_m \left(\frac{R_m}{O_m}\right)$$

Proof. Let k be the index of elements that minimize the ratio $\frac{R_m}{O_m}$. For every m, $\frac{R_m}{O_m} \geq \frac{R_k}{O_k}$, therefore for every m, $O_k * R_m \geq R_k * O_m$. Summing over m we get $O_k * (\sum_m R_m) \geq R_k * (\sum_m O_m)$. Hence, $\frac{\sum_m R_m}{\sum_m O_m} \geq \frac{R_k}{O_k} = \min_m \frac{R_m}{O_m}$. $\square$

Lemma 46 *For any CRC $\gamma \in \Gamma$,*

$$\frac{\mathbf{V}(A^{TR}(v))|_\gamma}{\mathbf{V}(A^*(v))|_\gamma} \geq \frac{|\gamma|-1}{|\gamma|}$$

Proof. By Theorem 23 the trade reduction allocation can be partitioned to $|\gamma|-1$ cycles, $X_1,\ldots,X_{|\gamma|-1}$ such that $\mathbf{V}(A^{TR}(v))|_\gamma = \sum_{i=1}^{|\gamma|-1} V(X_i)$. By the definition of the trade reduction allocation, the efficient allocation has one additional cycle $X_{|\gamma|}$ such that for all $i \in \{1,\ldots,|\gamma|-1\}$, $X_{|\gamma|} \leq X_i$ and $\mathbf{V}(A^*(v))|_\gamma = \sum_{i=1}^{|\gamma|} V(X_i)$. By applying Lemma 47 we get:

$$\frac{\mathbf{V}(A^{TR}(v))|_\gamma}{\mathbf{V}(A^*(v))|_\gamma} \geq \frac{|\gamma|-1}{|\gamma|}$$

$\square$

Lemma 47 *Let $n \in \mathbf{Z}^+$, $m \in \{1,\ldots,n\}$, and $X_i \in \mathbf{R}^+$ for all $i \in \{1,\ldots,n\}$. If $X_i \geq X_m$ for all $i < m$ and $X_i \leq X_m$ for all $i > m$, then*

$$\frac{\sum_{i=1}^m X_i}{\sum_{i=1}^n X_i} \geq \frac{m}{n}$$

Proof. The proof is by induction on n for any fixed m. For any n such that $n \geq m$ we prove the claim by induction on n.

If $n = m$ the claim is true since we have 1 on both sides of the inequality. Now assume that we have proven the claim for some n_0 such that $n_0 \geq m$, to prove the claim for $n_0 + 1$. By the induction hypothesis,

$$\frac{\sum_{i=1}^{m} X_i}{\sum_{i=1}^{n_0} X_i} \geq \frac{m}{n_0},$$

hence $n_0 \sum_{i=1}^{m} X_i \geq m \sum_{i=1}^{n_0} X_i$.

Since $X_i \geq X_m \geq X_{n_0+1}$ for all $i \leq m$, we have $\sum_{i=1}^{m} X_i \geq m X_{n_0+1}$. Using the induction hypothesis we get by summation

$$n_0 \sum_{i=1}^{m} X_i + \sum_{i=1}^{m} X_i \geq m \sum_{i=1}^{n_0} X_i + m X_{n_0+1}$$

therefore

$$\frac{\sum_{i=1}^{m} X_i}{\sum_{i=1}^{n_0+1} X_i} = \frac{\sum_{i=1}^{m} X_i}{\sum_{i=1}^{n_0} X_i + X_{n_0+1}} \geq \frac{m}{n_0 + 1}$$

which is what we wanted to prove. □

C.3.3 Algorithm for TRM

In this section we prove that the TRM mechanism can be calculated in time polynomial in $k, log\ n$ and $log\ C$, assuming that the values/costs in each market are sorted.

Lemma 48 *Assuming that the values/costs in each market are sorted, there exists an algorithm to calculate the TRM output in time $O((m + k\ log\ k)\ m\ log\ (C + n))$.*

Proof. In order to calculate the TRM allocation, we first find the efficient allocation. This can be done in time $O((m + k\ log\ k)\ m\ log\ (C + n))$ by Theorem 32. Now we need to remove the minimal positive cycle in each CRC. This can be done by solving two shortest path problems on the RRG. This is true since by Lemma 39, the minimal positive cycle has a cost of $C_b + \tilde{d}(C_b, C_s) + C_s$. Note that $\tilde{d}(Z, C_s) = C_b + \tilde{d}(C_b, C_s)$ and $\tilde{d}(C_s, Z) = C_s$ since the cycle is minimal. So by finding the distance from/to the sink node in the RRG, and finding the node with the minimal sum of the two, we can find the minimal cycle in each CRC. These two shortest path problems can be solved in time $O((k + 1)(m + 2k)) = O(m\ k)$ (see Corollary 31).

Finally note that we have also calculated the payments of the TRM, since by Definition 16 the payments are exactly the distances we have calculated. □