



Handling Interdependent Values in an Auction Mechanism for Bandwidth Allocation in Tactical Data Networks

Mark Klein
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213
mk@sei.cmu.edu

Gabriel A. Moreno
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213
gmoreno@sei.cmu.edu

David C. Parkes
Harvard School of Engineering
and Applied Sciences
Cambridge, MA 02138
parkes@eecs.harvard.edu

Daniel Plakosh
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213
dplakosh@sei.cmu.edu

Sven Seuken
Harvard School of Engineering
and Applied Sciences
Cambridge, MA 02138
seuken@eecs.harvard.edu

Kurt Wallnau
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213
kcw@sei.cmu.edu

ABSTRACT

We consider a tactical data network with limited bandwidth, in which each agent is tracking objects and may have value for receiving data from other agents. The agents are self-interested and would prefer to receive data than share data. Each agent has private information about the quality of its data and can misreport this quality and degrade or otherwise decline to share its data. The problem is one of *interdependent value* mechanism design because the value to one agent for the broadcast of data on an object depends on the quality of the data, which is privately known to the sender. A recent two-stage mechanism due to Mezzetti (2004) can be modified to our setting. Our mechanism achieves efficient bandwidth allocation and provides incentive compatibility by conditioning payments on the realized value for data shared between agents.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics;
C.2.0 [Computer-Communication Networks]: General

General Terms

Design, Economics, Theory

1. INTRODUCTION

Almost all military group operations rely on the platforms (air, sea, and ground) involved in a mission to act as a cohesive force. Thus, the platforms must establish and maintain a common understanding of the tactical situation. Tactical data is often exchanged between the platforms using a standardized radio network, commonly called a tactical data in-

formation link (TADIL) [5]. However, the *bandwidth* in these tactical data networks is a very scarce resource and the mission outcome can be significantly affected by decisions made in real-time about what data to share. Thus, *ad hoc bandwidth allocation* can seriously jeopardize the mission, e.g., when position updates on the enemy are delayed.¹

Major kinds of tactical data exchanged between platforms include their own positions and movements and the trajectory (or “track”) data for objects of interest. Track data is processed radar data which typically represents real objects such as airplanes, helicopters, ships, and land vehicles. The quality of the common operational picture critically depends on the effective sharing of this data. In standard practice, each track is assigned to the *single* platform with the best quality data for that track (the so-called “reporting responsibility” (R2) rules). Over a single network cycle, each responsible platform sends data on its assigned tracks.

These rules are designed to minimize latency by making the network cycle time (the time between two consecutive data exchanges on the same object) as short as possible, since data on each track is broadcast only once. However, this is an extreme minimalist approach because there is no attempt to improve quality by sharing additional data on the same track. We take as our starting point the premise that additional communication per network cycle can significantly improve the quality of the combined data, and by enough to warrant the additional latency that comes from a longer cycle time. We assume that some underlying latency constraint determines the maximum network cycle time which translates into a maximum amount of bandwidth available for data exchange. Our goal is to allocate all available bandwidth per network cycle such that the value to all platforms from the shared data is maximized.

In our domain, the problem of bandwidth allocation is particularly complex because we also allow for the *self-interest*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NetEcon’08, August 22, 2008, Seattle, Washington, USA.
Copyright 2008 ACM 978-1-60558-179-8/08/08 ...\$5.00.

¹A Research Service report for the U.S. Congress notes that “When the supply of bandwidth becomes inadequate during combat, military operations officers have sometimes been forced to subjectively prioritize the transmission of messages. They do this by literally pulling the plug temporarily on some radio or computer switching equipment in order to free up enough bandwidth [...]”[15]

of individual platforms. Even in military settings some self-interest is a concern because individual platforms have local goals and benefit from receiving data but have no value for sending data. Thus, we can expect a tendency for platforms to under-represent the quality of their data so that the bandwidth is allocated to the transmission of data by others.

Deceptive behavior is a particular concern when coalition forces are involved in a mission. For example, even within the Navy, there are many sub-groups such as fleets (Atlantic, Pacific and Asiatic) or commands (Surface Force, Submarine Force, Naval Air Force). Undoubtedly, given the size and scope of the Navy, it is not possible for every sub-group to serve exactly the same agenda. Thus we must account for the possibility of deceptive behavior, even under one flag.

1.1 A Simple Example

To understand the need for incentive design in this domain it is helpful to consider the following example. We assume that the quality of data is private information but that there is a commonly known value function that maps this quality into value, measured in units of an appropriate currency.

EXAMPLE 1. *Consider a scenario with two ships A and B and one object of interest. Ship A knows that its information about the object is of a quality that makes it worth value 5 to any other ship and ship B knows its data is of a quality that makes it worth value 2 to any other ship. Suppose that only one piece of data can be broadcast per cycle. When truthful, ship A will be asked to broadcast its data because that maximizes the total value improvement. However, A would prefer to under-report its quality (e.g. to 1) so that B will be asked to broadcast its data instead.*

Because of this apparent conflict of interests, we consider the design of mechanisms where participants are rewarded for their particular role in sharing the data through payments, such that all participants' incentives are aligned. From the perspective of mechanism design, the problem considered here is challenging because of the presence of *interdependent values*. The value of one participant for an allocation of bandwidth depends on private information held by other participants, namely on the quality of their observed data. Moreover, our framework is one of "multi-dimensional" private information because each participant can potentially have information on multiple objects. The following example illustrates how a standard Groves mechanism [6] generally fails in domains with interdependent values:

EXAMPLE 2. *Let's again consider the scenario from Example 1 with ships A and B and one object of interest. With a Groves mechanism, if ships truthfully report the quality of their data, A will be asked to broadcast data to B and A will be paid 5, i.e., the value to B based on the reports about quality. Notice, however, that A can drive this payment arbitrarily high by over-claiming the quality of its data.*

1.2 Related Work

Designing useful mechanisms in the presence of value interdependence is a familiar challenge in mechanism design theory. Dasgupta and Maskin [2] consider the design of efficient, interdependent value mechanisms for the single-dimensional setting. Ito and Parkes [7] extend their methods to a single-minded combinatorial auction environment. Constantin et al. [1] have studied the problem of dynamic mechanisms in a setting with interdependent values. See

Krishna [10] for an accessible introduction to the problem of interdependent valuations.

Dash et al. [3, 13] have previously studied the problem of interdependent valuations in the context of data fusion in sensor networks but with the restriction that each participant tracks at most a single object. We dispense with this restriction in our work. Moreover, for a mechanism to implement the efficient outcome in an interdependent value setting, it is generally necessary that the agents' value functions satisfy the *single crossing condition* [2, 10]. It is unclear to us that this condition will hold in their setting.² These authors also assume that the mechanism has the power to enforce an outcome to ensure that participants actually follow through with the required communication actions. In our mechanism, the incentives are aligned so that participants will *choose* to obey the proposed outcome. In this sense we achieve the *faithfulness* property of *distributed* mechanism design [14, 4]. The main additional assumption that we make over those in Dash et al. [3] is that the mechanism center can observe the quality of the data transmitted.

While it is possible to design efficient mechanisms for interdependent value settings with one-dimensional signals when the single crossing condition is satisfied, there exist strong impossibility results by Jehiel and Moldovanu [8] for the multi-dimensional case that would also apply to our domain.³ However, this only applies when one is constrained to a "standard mechanism design paradigm" of determining payments simultaneously with determining the outcome. In our work, we modify a mechanism due to Mezzetti [11] which circumvents this impossibility result and also no longer requires the single crossing condition to be satisfied.

The result is an efficient (i.e., value-maximizing), incentive compatible mechanism for allocating bandwidth. Our mechanism provides a general solution to a practically motivated model of data fusion in sensor networks, where we can dispense with some of the restrictive assumptions that limit earlier work. The key innovation in Mezzetti [11] is to define payments based on the actual value from the *realized outcome* rather than on agents' reports of private information. This allows Mezzetti to achieve efficiency despite interdependent valuations. Our mechanism provides more robust properties than that of Mezzetti, as we discuss below.

²The single-crossing condition (SCC) can be understood as follows. Consider the value to a ship for an allocation as its *increase* in value over-and-above its value for the null allocation in which no data is shared. Given this, for a particular allocation of bandwidth the SCC requires that the marginal increase in value to ship A for a marginal increase in the quality of its own data should be greater than the marginal increase in value to some other ship B. But this condition fails: ship A has *zero* marginal value for an increase in the quality of its own data for this allocation because it has already internalized the value from its own data and only has value for receiving new data from other ships.

³With just one object in the environment, each agent prefers that another agent has high quality data about the object than low quality data. No such total ordering exists with multi-dimensional signals: it is unclear whether high quality data on object A but low quality data on B is preferred to low quality data on object A but high quality data on B. The participants in our problem have multi-dimensional private signals because they each track multiple objects. However, it is not possible to decompose the problem across objects because the resource constraint links the allocation decision on all objects.

Complementing our theoretical results, we have also developed an application framework that exhibits sufficient scale and dynamic complexity to study the feasibility of computational mechanism design in this setting. The framework emulates a tactical data network (based loosely on the Navy LINK-11) and includes much of what is required to construct a common operational picture from radar sensor data.⁴

2. PRELIMINARIES

We consider participants $i \in N = \{1, \dots, n\}$ (henceforth called agents) and objects of interest $j \in G = \{1, \dots, m\}$. The *full type* $\psi_i = (s_i, q_i, d_i) \in \Psi_i$ represents all private information of an agent, where $s_i = (s_{i1}, \dots, s_{im}) \in \{0, 1\}^m$, with $s_{ij} = 1$ when object j is in the *region of observation* of agent i , $q_i = (q_{i1}, \dots, q_{im}) \in \mathbb{R}_{\geq 0}^m$, where q_{ij} is the *quality* of i 's observed data about object $j \in G$, and $d_i = (d_{i1}, \dots, d_{im}) \in \mathcal{D}^m$, where d_{ij} is the data agent i actually observed on object j and $\mathcal{D}$ is the set of possible data items that can be observed for one object. We define the *reduced type* (or simply *type*) $\theta_i = (s_i, q_i) \in \Theta_i$ to capture agent i 's private information s_i and q_i that is of small size and can be reported efficiently in contrast to the data d_i . We let $\theta_{ij} = (s_{ij}, q_{ij})$ and require $s_{ij} = 0 \Rightarrow q_{ij} = 0$. We let $\Theta = \times_{i \in N} \Theta_i$ denote agents' joint type space, and $\theta \in \Theta$ denotes a particular joint type.

Agents have intrinsic value for receiving data on objects in their region of observation. We allow this value to depend on the quality of the received data and the quality of the data they already have. Consider agent i with type $\theta_i = (s_i, q_i)$, and some *other* agent $k \neq i$ holding data on object $j \in G$ with quality q_{kj} . We let $w(q_{kj}, \theta_{ij})$ denote the value to agent i for receiving a single data item of quality q_{kj} on object j . We assume that $s_{ij} = 0 \Rightarrow w(\cdot, \theta_{ij}) = 0$.

The assumption that $s_{ij} = 0 \Rightarrow w(\cdot, \theta_{ij}) = 0$ is not essential, but we use it because we assume the point of sharing data is to improve the quality of information about a track for a ship that already had responsibility for that track. Additionally, we make the following monotonicity assumption:

Assumption 1. For all agents i, k , for all objects j : $w(q'_{kj}, \theta_{ij}) \geq w(q_{kj}, \theta_{ij})$ for all $q'_{kj} \geq q_{kj}$.

The goal of the mechanism that we propose will be to determine which data should be shared by which agents, to make a bandwidth allocation decision. We use $z \in \{0, 1\}^{nm}$ to denote this allocation decision, where $z_{ij} \in \{0, 1\}$ denotes whether agent i shall broadcast its data on object j .⁵ An allocation decision z , together with the agents' data d induces an *outcome* $x(z, d) \in \mathcal{D}^{nm}$, where $x_{ij}(z, d) \in \mathcal{D}$ is the data broadcast by agent i about object j , and where ε denotes no data:

$$x_{ij}(z, d) = \begin{cases} d_{ij} & , \text{ if } z_{ij} = 1 \\ \varepsilon & , \text{ otherwise} \end{cases} \quad (1)$$

Of course, $x(z, d)$ is only the true outcome if the agents are in fact holding data d and obey the allocation decision z . Later we will allow agents to deviate, first in reporting their types, and then in their use of bandwidth and in the data that they choose to broadcast.

⁴The implementation has been packaged for use by external research collaborators and is available to researchers at the Naval Postgraduate School (for more details see [9]).

⁵Note that we actually need a *schedule* of communication and not just an allocation. This detail, however, is not important for the mechanism design, so we stick with z .

We will propose a mechanism that 1) must be able to determine agents' *expected value* given allocation z and agents' types θ , and 2) must be able to determine agents' *realized value* given a particular outcome x . For this purpose we define a *known value function, common to all agents* that can take either allocations or outcomes as its argument.⁶

First, we define agent i 's *value function on allocations* as $v_i : \{0, 1\}^{nm} \times \Theta \rightarrow \mathbb{R}_{\geq 0}$, where $v_i(z, \theta)$ denotes agent i 's expected value given allocation z and joint type θ . For this version of the paper we make the following *additivity assumption* on all value functions:

Assumption 2. The value function is *additive* across all items and only depends on the quality of the data, i.e.:

$$v_i(z, \theta) = \sum_{k \in N, k \neq i} \sum_{j \in G} z_{kj} \cdot w(q_{kj}, \theta_{ij}) \quad (2)$$

Note that a simple value function that fulfills the additivity assumption can be constructed via the Fisher information metric. Because the Fisher information is additive, the additivity of the value function follows immediately; see Rogers et al. [13] for a discussion. It is a simple matter to devise a more complex valuation model without affecting the proposed mechanism. The main effect would be to make the optimization problem that defines the social choice function more complex.

When the agents receive data d , they can evaluate it and determine its quality. We define the quality function $Q : \mathcal{D}^{nm} \rightarrow \mathbb{R}_{\geq 0}^{nm}$, where $Q(d)$ denotes the quality matrix q corresponding to data d and $Q_{kj}(d)$ denotes quality q_{kj} of data item d_{kj} . We normalize $Q(\cdot)$ such that $Q(\varepsilon) = 0$. Now we can define an agent's *realized value* for outcome $x(z, d)$:

$$v_i(x(z, d), \theta) = \sum_{k \in N, k \neq i} \sum_{j \in G} w(Q_{kj}(x(z, d)), \theta_{ij}) \quad (3)$$

Given a *bandwidth constraint* $B \in \mathbb{Z}_{>0}$, and agents' joint type $\theta = (s, q)$, we can define the social planner's goal as determining the optimal (i.e., *efficient*) bandwidth allocation $z^*(\theta)$ which, together with the agents' data d , induces outcome $x(z^*(\theta), d) \in \mathcal{D}^{nm}$, such that the total value to all agents is maximized:

$$z^*(\theta) \in \arg \max_z \sum_{i \in N} \sum_{k \in N, k \neq i} \sum_{j \in G} z_{kj} \cdot w(q_{kj}, \theta_{ij}) \quad (4)$$

$$\text{s.t.} \quad \sum_{k \in N} \sum_{j \in G} z_{kj} \leq B. \quad (5)$$

The formulation assumes that broadcasting any data item takes one unit of bandwidth. For each component $z_{kj} \in \{0, 1\}$, the total objective value contribution simplifies to $V_{kj}(\theta) = \sum_{i \neq k} w(q_{kj}, \theta_{ij})$, which is the total value accrued by agents other than k when k broadcasts its data for object j . Given this, the objective function is equivalent to:

$$z^*(\theta) \in \arg \max_z \sum_{k \in N} \sum_{j \in G} z_{kj} V_{kj}(\theta). \quad (6)$$

Thus, the social planner's problem can be solved in time $O(mn \log(mn))$ by sorting all allocation components z_{kj} and then picking the B ones with the highest value $V_{kj}(\theta)$.

⁶It is not essential to our mechanism that the value function is known, but we adopt this convention here to keep the presentation as simple as possible. Note that the value function is still parameterized via an agent's private type θ_i .

In designing a mechanism, we assume that within the network, there is a *trusted center* (or platform). This center is responsible for collecting agents' claims about their types. We make the following assumption:

Assumption 3. The trusted center is able to observe the data broadcast by the agents and interpret its quality.

Furthermore, the mechanism requires a form of currency.

Assumption 4. The trusted center can make and collect payments in some currency to and from agents and agents' value gain due to received data can be measured in units of this currency. We assume that the agents have quasilinear utility, i.e., given payment transfer $t_i \in \mathbb{R}$ to agent i and realized outcome $x(z, d)$, agent i 's utility is:

$$u_i(x(z, d), \theta) = v_i(x(z, d), \theta) + t_i. \quad (7)$$

It is important that $v_i(x(z, d), \theta)$ is a correct representation of value in the units of this currency. Without this *intrinsic* value an agent would only seek to improve its payment and not care about the data received. We can relax this in an alternate model by allowing agents to explicitly report their complete valuation functions, even including statements that they have zero value for data.

In the military context, while each ship is at sea, we suppose the trusted center can perform *accounting* for the total payment due to each ship (each ship can in turn publish a policy on how accrued payment is to be distributed amongst its personnel.) Upon completion of a mission, we then need for this accrued payment to be converted into something of value. This value can be realized in the form of promotions and other rewards, for example getting a longer leave or higher priority in getting a desirable posting in a future mission.⁷ This incentive structure could be established as part of planning a mission or a sequence of sub-missions.

3. TWO-STAGE MECHANISM DESIGN

We propose a new mechanism for bandwidth allocation that operates in two stages, motivated by Mezzetti [11]. In stage 1, agents report their reduced type to the center. Because agents are self-interested, we distinguish between an agent's *report* about its reduced type $\hat{\theta}_i = (\hat{s}_i, \hat{q}_i)$ and its actual reduced type θ_i . Based on all agents' reports $\hat{\theta}$, the center computes the optimal bandwidth allocation $z^*(\hat{\theta})$. Because the center cannot enforce an outcome, it only suggests the allocation $z^*(\hat{\theta})$ to the agents. In stage 2, the agents broadcast part of their data d_i to all other agents and then the center makes payments contingent on the broadcast data. However, an agent can choose to broadcast other data than it actually holds. We will denote the *broadcast data* agent i pretends to hold as $\bar{d}_i$. Furthermore, an agent could choose to use the bandwidth differently than suggested by $z^*(\hat{\theta})$. Thus, each agent adopts a *bandwidth use* $\bar{z}_i$.

In stage 1, we allow agent's type reports $\hat{\theta}_i$ to include both overclaims and underclaims about the quality of the

data they are holding. However, for stage 2 we assume that agents cannot broadcast data of higher quality than they actually have. More formally:

Assumption 5. An agent cannot increase the quality of its data, i.e., $Q(\bar{d}_{ij}) \leq Q(d_{ij})$ for all agents i and objects j .

In our domain, this assumption seems well-motivated because an agent could always add noise to its data, thus degrade the quality of its data, without anyone ever noticing it. However, when an agent tried to fabricate data that appears to be of higher quality, there is a non-negligible probability that this would eventually be detected by the trusted center or other agents and then trigger an investigation with the possibility of severe punishments.

3.1 A new Bandwidth Allocation Mechanism

We assume that the trusted center hosts the mechanism and does not deviate from the rules of the mechanism.

Mechanism 1 (Bandwidth Allocation Mechanism)

- 1a. Each agent makes a report $\hat{\theta}_i \in \Theta_i$ about its type.
- 1b. The center computes the optimal bandwidth allocation $z^*(\hat{\theta}) \in \{0, 1\}^{nm}$ and announces z^* to all agents.
- 2a. Each agent adopts broadcast data $\bar{d}_i$ and bandwidth use $\bar{z}_i$ which induces outcome $x(\bar{z}, \bar{d})$.
- 2b. The center observes $x(\bar{z}, \bar{d}) \in \mathcal{D}^{nm}$, and determines payments for each agent i :

$$t_i(x(\bar{z}, \bar{d}), \hat{\theta}_{-i}) = \quad (8)$$

$$\sum_{k \neq i} v_k(x(\bar{z}, \bar{d}), \hat{\theta}) - \sum_{k \neq i} v_k(z^*(\hat{\theta}_{-i}), \hat{\theta}) \quad (9)$$

The information required in the second term in line (9) is provided in the first stage of the mechanism via type reports reports $\hat{\theta}_{-i} = (\hat{\theta}_1, \dots, \hat{\theta}_{i-1}, \hat{\theta}_{i+1}, \dots, \hat{\theta}_n)$. The additional information required in the first term is available because the center is able to observe $x(\bar{z}, \bar{d})$ by Assumption 3.⁸

In establishing the incentive properties of the mechanism we have two main items to establish in equilibrium:

- In step 1a, agents choose to make truthful claims $\hat{\theta}_i = \theta_i$.
- In step 2a, agents choose to follow the recommended use of bandwidth $z^*(\hat{\theta})$ and also choose to broadcast ungraded data on the proposed objects.

When adopting both of these desired behaviors we say an agent is *faithful* [14]. The strategy in this mechanism allows for deviations $\hat{\theta}_i$ (in step 1a) and $\bar{z}_i$ and $\bar{d}_i$ (in step 2a).

DEFINITION 1. The *faithful strategy* is an *ex post Nash equilibrium* of the mechanism whenever

$$u_i(x(z^*(\theta), d), \theta) \geq u_i(x(\bar{z}, \bar{d}), \theta), \quad (10)$$

for all i , all θ_i , all θ_{-i} , all $\hat{\theta}_i$, all d_i , all d_{-i} , all $\bar{d}_i$, all $\bar{z}_i$, where $\bar{z} = (\bar{z}_i, z_{-i}^*(\hat{\theta}_i, \theta_{-i}))$ and $\bar{d} = (\bar{d}_i, d_{-i})$.

This requires that every agent maximizes its utility by being faithful as long as the other agents are also faithful and whatever the type of the other agents (note that this is a stronger equilibrium notion than Bayes-Nash equilibrium).

⁸If the mechanism cannot directly observe the quality of data actually realized (in step 2b), then this step can be replaced by a consensus vote taken from the other agents. We defer an analysis of this to a longer version of the paper.

⁷Navy personnel are rewarded based on how well they perform. For example, after action investigations and punishments (negative rewards) are common for unsuccessful missions and it is reasonable to assume that positive rewards are also possible. Of course, the appropriate sharing of data has a direct impact on the performance of a mission. Thus, rewarding the personnel based on the payments accrued in the mechanism is consistent with the Navy's general policy.

THEOREM 1. *Mechanism 1 is ex post Nash faithful and efficient (allocates bandwidth to maximize agents' values).*

PROOF. Assume agents $\neq i$ are faithful and consider the utility to agent i given that it can deviate in $(\hat{\theta}_i, \bar{z}_i, \bar{d}_i)$:

$$v_i(x(\bar{z}_i, z_{-i}^*(\hat{\theta}_i, \theta_{-i}), \bar{d}_i, d_{-i}), \theta) \quad (11)$$

$$+ \sum_{k \neq i} v_k(x(\bar{z}_i, z_{-i}^*(\hat{\theta}_i, \theta_{-i}), \bar{d}_i, d_{-i}), \theta), \quad (12)$$

where the first term is i 's value realized from the data received from other agents and the second term is the first term in the payment equation (9). We can ignore the second term in the payment equation (9) because it is completely independent of agent i 's strategy. Given that i can't upgrade its data (Assumption 5) and given that others' value increases with the quality of data (Assumption 1), we first observe that we can stipulate $\bar{d}_i = d_i$, i.e. given the use of bandwidth $\bar{z}_i$, agent i should broadcast non-degraded data. This does not affect Eq. (11), and increases the value of Eq. (12) whatever the choice of $\bar{z}_i$ and $\hat{\theta}_i$. In considering the joint effect of $\bar{z}_i$ and $\hat{\theta}_i$, agent i now seeks to maximize

$$v_i(x(\bar{z}_i, z_{-i}^*(\hat{\theta}_i, \theta_{-i}), d), \theta) + \sum_{k \neq i} v_k(x(\bar{z}_i, z_{-i}^*(\hat{\theta}_i, \theta_{-i}), d), \theta) \quad (13)$$

Note that the combined effect of $\bar{z}_i$ and $\hat{\theta}_i$ is limited to its influence on the bandwidth allocation in this expression. Now, since the other agents are faithful, when reporting true type $\hat{\theta}_i = \theta_i$ in step 1b of the mechanism, the center chooses allocation $z = z^*(\theta_i, \theta_{-i})$ across all feasible z to maximize

$$v_i(z, \theta) + \sum_{k \neq i} v_k(z, \theta). \quad (14)$$

Now, with truthful type reports θ , we can derive that an allocation z that maximizes equation (14) at the same time maximizes equation (13) because:

$$v_i(z, \theta) = \sum_{k \in N, k \neq i} \sum_{j \in G} z_{kj} \cdot w(q_{kj}, \theta_{ij}) \quad (15)$$

$$= \sum_{k \in N, k \neq i} \sum_{j \in G} z_{kj} \cdot w(Q_{kj}(x(z, d)), \theta_{ij}) \quad (16)$$

$$= \sum_{k \in N, k \neq i} \sum_{j \in G} w(Q_{kj}(x(z, d)), \theta_{ij}) \quad (17)$$

$$= v_i(x(z, d), \theta), \quad (18)$$

Equation (16) follows because with truthful reports and $\bar{d}_i = d_i$ we have that $Q_{kj}(x(z, d)) = q_{kj}$ where $z_{kj} = 1$. Equation (17) follows because $x_{kj}(z, d) = \varepsilon$ where $z_{kj} = 0$ and $Q(\varepsilon, \cdot) = 0$, thus we can drop the multiplication with z_{kj} . Therefore, the best strategy for agent i is to be faithful, i.e., to report $\hat{\theta}_i = \theta_i$, subsequently follow $\bar{z}_i = z_i^*(\theta)$ and report non-degraded data d_i . This completes the proof, when we also observe that the outcome of the mechanism is efficient with faithful agents. ■

This equilibrium is also robust to mistakes by other agents. As long as agent i reports its own type truthfully in step 1a, its best-response in step 2a is to be faithful whatever the strategies of other agents, including deviations by other agents in both step 1a and step 2a. The only way in which faithfulness is not a subgame perfect equilibrium here (i.e.

the best strategy even off the equilibrium path), is when agent i itself deviates from reporting $\hat{\theta}_i = \theta_i$ in step 1a, in which case agent i might then be able to provide the other agents with more value in step 2a by transmitting alternate data (e.g., data for which it had understated the quality).

EXAMPLE 3. Consider the example from earlier and introduce an additional ship C with data of quality that gives value 1 to the other ships. Suppose the ships report true claims about the quality of their data. The center proposes that ship A uses the bandwidth. In equilibrium, ship A broadcasts its data without degradation and the center observes that the data provides value 5 to ships B and C . Without ship A , based on reports from B and C , ship B would share its data and the total value would be 2. Ship A receives payment $10-2=8$. For ship B , with its presence ship C is observed to obtain value 5. Without ship B , based on reports, ship C would obtain value 5. Ship B receives no payment. The same analysis holds for ship C . Ship A can no longer improve its utility by overstating the quality of its data because its payment would remain 8 since it depends on the observed quality of the data actually broadcast. Similarly, ship B is currently receiving utility of 5 (from the data of ship A), and with a high claim of quality (e.g. 7) it would be asked to share its data but would only receive a payment of $4-5=-1$ because it cannot improve the quality of its actual data; its utility would fall to -1 .

It is somewhat intricate to see that faithfulness is not a dominant strategy. Note that in the above example, when ship B overstates its quality (e.g. 7) then ship A 's utility is $2 + (2-7)=-3$. This is because it receives data of quality 2, ship C 's value is 2, and the value to the ships other than A without A present only based on reported quality (which is overstated) would be 7 (with ship B sending data to ship C). But by reporting its own quality as 20, A would be asked to share its data and receive a higher utility of $10-7=3$.

COROLLARY 1. *In the ex post Nash equilibrium of Mechanism 1, each agent's payment is the total amount by which its presence in the system improves the value to other agents.*

PROOF. In equilibrium, the payment is simply

$$\sum_{k \neq i} v_k(x(z^*(\theta), d), \theta) - \sum_{k \neq i} v_k(z^*(\theta_{-i}), \theta), \quad (19)$$

which establishes the result given that $z^*(\theta_{-i})$ allocates the available bandwidth to maximize the total value to agents $\neq i$ when i is not present. ■

It is obvious that an additional agent in the system can only increase but never decrease the total value accrued because in the worst case the center could always not allocate to the agent. It follows that the mechanism is *ex post individually rational*, such that no agent's utility is ever negative.

To understand the budget properties, notice that in EXAMPLE 3, ship A receives a payment of 8 in equilibrium and thus the mechanism runs at a deficit. However, a budget deficit does not seem very problematic in our context of a military deployment, because it is expected that agents will generate credits through their participation in the mechanism. However, this problem would need to be addressed in commercial applications. Our last example illustrates that sometimes, agents can also be required to *make* payments.

EXAMPLE 4. Consider three ships A, B and C , and three objects 1, 2, 3, where only one piece of data can be broadcast. Ship A has data quality that provides value 5, 0, 0 for objects 1, 2 and 3 respectively to the other ships, similarly ship B has data quality that provides value 7, 8 and 10 and ship C data quality that provides value 8, 7 and 5. Ship A does not have objects 2 and 3 in its region of observation. In equilibrium, ship C shares data on object 1 with ships A and B . The payment to ship A is $8 - 10 = -2$, 8 since ship B gets value 8 from the data of ship C and -10 since ship C would get data of value 10 from ship B without the presence of A . Thus, ship A imposes an externality on ships B and C and must make a payment. The other payments in this example are $8 - 8 = 0$ by ship B and $16 - 7 = 9$ by ship C .

As this example illustrates, in our domain, we would need to take steps to achieve fairness across agents because different ships may have different inherent data collection capabilities. For example, depending on a ship's radar equipment, some baseline payment (positive or negative) could be set *a priori* before a mission.

3.2 Relation to the Mezzetti Mechanism

Mezzetti [11] recognized that the impossibility results related to mechanism design with interdependent values can be avoided if (a) values are privately realized by the agents once an allocation is made, and (b) two-stage mechanisms can be adopted in which payments are made contingent on realized values that are reported in a second stage.

The analog in our setting is that recipients of data can determine its quality and thus the realized value. But before the allocation of bandwidth is determined, the mechanism can only rely on reports on quality of data from prospective senders. In the simple context of EXAMPLE 2, the mechanism proposed by Mezzetti would receive claims about quality from A and B and then ask A to broadcast data, because it is of higher quality and thus of more value to B than B 's data is to A . Then B would be asked to make a report about the value of the data received and A would receive this as a payment. Here we see the fragility of the mechanism: B is indifferent at the second stage between reporting the true value realized from the data or any kind of misreport. Moreover, collusive deviations are easy, with B over-stating value in return for a cut of the payments received by A .

While conceptually similar to that of Mezzetti, our mechanism avoids these problems. The equilibrium properties do not rely on agents correctly reporting their realized values. We can avoid this because it is reasonable to assume that the trusted center is situated within the network and can observe the quality of broadcast data, and thus infer for each agent its value from the realized outcome. Unlike Mezzetti, we are also able to collect VCG-like "charge-back" payments (the second term in payment Eq. (9)) for a similar reason: the agents' reports reveal enough information to the center to determine the total value excluding any particular agent.

4. CONCLUSION

We have introduced an incentive compatible mechanism for efficient bandwidth allocation in a tactical data network. Inherent to every data sharing domain is the problem with interdependent valuations. The mechanism we propose is based on a two-stage mechanism introduced by Mezzetti [11] and circumvents the well-known impossibility results for interdependent value settings.

Future work should consider in particular the dynamics of the problem. The current solution is stationary, with the same communication tasks performed every network cycle. In practice, forward-looking schedules are necessary (for efficiency and incentive compatibility purposes) which requires a dynamic mechanism design approach [12]. Furthermore, we consider making the tradeoff between latency and data quality endogenous, with agents reporting preferences about this tradeoff. We also intend to allow for non-linear valuations, and to couple this with private valuation functions. This research also raises interesting questions about the use of payment schemes in operational contexts such as those presented by military domains. The best way to provide "after-action" rewards commensurate with the aggregate payments accrued during the execution of repeated mechanism instantiations remains to be determined.

5. REFERENCES

- [1] F. Constantin and D. C. Parkes. On revenue-optimal dynamic auctions for bidders with interdependent values. In *Proc. of the Workshop on Agent Mediated Electronic Commerce IX (AMEC'IX)*, Honolulu, Hawaii, May 2007.
- [2] P. Dasgupta and E. S. Maskin. Efficient auctions. *Quarterly Journal of Economics*, 115:341–388, 2000.
- [3] R. Dash, A. Rogers, S. Reece, S. Roberts, and N. R. Jennings. Constrained bandwidth allocation in multi-sensor information fusion: a mechanism design approach. In *Proc. of the 8th Int. Conf. on Information Fusion*, July 2005.
- [4] J. Feigenbaum, M. Schapira, and S. Shenker. Distributed Algorithmic Mechanism Design. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani (eds.), Cambridge University Press, 2007.
- [5] J. Golliday, C. Leslie. Data link communications in tactical air command and control systems. *IEEE Journal on Selected Areas in Communication*, 3(5):779–791, 1985.
- [6] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [7] T. Ito and D. C. Parkes. Instantiating the contingent bids model of truthful interdependent value auctions. In *Proc. of the 6th Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2006.
- [8] P. Jehiel and B. Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69:1237–1259, 2001.
- [9] M. Klein, D. Plakosh, and K. Wallnau. Using the Vickrey-Clarke-Groves auction mechanism for enhanced bandwidth allocation in tactical data networks. Technical report CMU/SEI-2008-TR-004, CMU Software Engineering Institute, 2008.
- [10] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [11] C. Mezzetti. Mechanism design with interdependent valuations: Efficiency. *Econometrica*, pages 1617–1626, 2004.
- [12] D. C. Parkes. Online Mechanisms. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani (eds.), Cambridge University Press, 2007.
- [13] A. Rogers, R. K. Dash, N. R. Jennings, S. Reece, and S. Roberts. Computational mechanism design for information fusion within sensor networks. In *Proc. of the 9th Int. Conf. on Information Fusion*, July 2006.
- [14] J. Shneidman, D. C. Parkes, and L. Massoulié. Faithfulness in Internet algorithms. In *Proc. of the SIGCOMM Workshop on Practice and Theory of Incentives and Game Theory in Networked Systems (PINS)*, September 2004.
- [15] C. Wilson. Network centric warfare: Background and oversight issues for congress. CRS Report RL32411, The Library of Congress, June 2004.