

Constrained Bandwidth Allocation in Multi-Sensor Information Fusion: A Mechanism Design Approach

R.K. Dash, A. Rogers and N.R. Jennings

Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK
{rkd02r,acr,nrj}@ecs.soton.ac.uk

S. Reece and S. Roberts

Engineering Science
University of Oxford
Oxford, OX1 3PJ, UK
{reece,sjrob}@robots.ox.ac.uk

Abstract – *Sensor networks are increasingly seen as a solution for a large number of environmental, security and military monitoring tasks. Typically, in these networks, noisy data from a number of local sensors is fused to reduce the uncertainty in the global picture. A central issue in this information fusion is the decision of what data should be shared between sensors, in order to maximise the global gain in information, when the bandwidth of the communication network is limited. In this paper, we study the problem from a selfish agent perspective. We show how the uncertainty in the measurement of an event can be cast as a utility function derived from the Kalman filter. We then use the tools of mechanism design to engineer an incentive-compatible mechanism that allows rational selfish agents to individually maximise their own utility, whilst ensuring that the overall utility of the system is also maximised. We apply the mechanism to multi-sensor target detection and consider the complexity of finding an efficient solution with broadcast communication protocols.*

1 Introduction

Multi-sensor networks (MSN) are being deployed in a wide variety of application areas ranging from military sensing to environmental monitoring and traffic control. Such networks consist of a number of sensors connected via a communication network. Each sensor is able to sense their local environment, but can also make use of data transmitted from neighbouring sensors in order to improve the accuracy of their own measurements (e.g. by combining their own noisy observation, with the observations from a number of other sensors, in order to reduce the final uncertainty).

Typically, however, the communication network that the sensors use to transmit and receive data has a limited capacity. As a result, not all of the available information can be distributed to all of the sensors. Thus, a key question posed by information fusion within multi-sensor networks, is how to make most effective use of this limited communication bandwidth. To date, work addressing this issue, has mainly

concentrated on using cooperation among distributed sensors (see section 2 for details). At its core, this approach involves determining the exchanges of observed data between the sensors that results in the maximum gain in information across the whole sensor network.

However, this approach overlooks the fact that in some applications each sensor may be *individually-owned* by different stakeholders. Such scenarios occur in applications like traffic control where each sensor is owned by a particular vehicle, in picosatellite projects where multiple companies own very small satellites monitoring a certain area, and in multi-nation multi-platform military exercises, where the military commanders of each platform are concerned with the well-being of their own platform over others. In such cases, the sensors may be operating in competitive rather than cooperative environments, and, as such, may attempt to optimise their own gain from the network, at a cost to the overall performance of the entire network. This is especially true in networks where the bandwidth available for transmission of data between the sensors is limited. In such cases, sensors may not be willing to transmit any data at all, since doing so yields no immediate benefit and reduces the amount of bandwidth that may be used to receive data from other sensors.

Thus, against this background, we seek to design a system that allows selfish sensors to exchange data within a bandwidth constrained network. We wish to ensure that the resulting global performance maximises the information gain of the entire system despite the selfish actions of individual sensors. To this end, we turn to the field of mechanism design which concerns itself with the analysis and design of systems in which the interactions between strategic, autonomous and rational agents leads to predictable global outcomes [1]. We thus model the sensors as agents who are seeking to optimise their own utility and our aim is then to find a *mechanism* (i.e an allocation rule and payment scheme that is contingent on the reports from each agent) that efficiently allocates the flow of data within the network.

We approach the problem by first describing a typical scenario, based on target detection, where information fusion needs to be carried out in a multi-sensor network. We then formulate a valuation function that characterises the value that an agent places on a particular piece of data originating from other agents. We then seek a mechanism that uses this valuation function to determine the optimal allocation of bandwidth within the network. Classically, within multi-agent systems, Vickrey-Clarke-Groves (VCG) mechanisms have been used to perform efficient resource allocation when agents have private independent valuations. However, the valuation function that we derive results in interdependent valuations since knowledge about a particular measurements precision affects the value that another agent places on it, but this knowledge is only privy to the agent having carried out the measurement. Thus, these mechanisms cannot be applied in a straightforward manner, and in tackling this problem, we advance the state of the art in the following ways:

1. We formulate a novel valuation function based around the information form of the Kalman filter [2] since this is the simplest and most elegant way of fusing different measurements of the same observation. This function equates the valuation to the expected gain in information when data from a number of sources is fused.
2. We extend the Generalised VCG mechanism developed in [3] to deal with the interdependent valuations of data found in our scenario and captured by our valuation function. We prove that this mechanism is *incentive-compatible* (i.e. it incentivises agents to truthfully reveal their signals) and thus achieves efficiency. Furthermore, we prove that the mechanism is *individually rational* (i.e. agents in the system derive a positive utility) and thus selfish agents will rather join the mechanism than opt out of it.
3. We show that the complexity of calculating the efficient allocation is dependent on the network topology and communication protocol. Specifically, in the case of a broadcast communication protocol we show that the calculation can be solved exactly in polynomial time.

The rest of the paper is organised as follows. Section 2 describes related research that has been carried out in this field. In section 3 we describe a motivating problem scenario that captures the essential features of information fusion in multi-sensor networks and we develop a valuation function for agents in this scenario. In section 4 we use the tools of mechanism design to engineer a system that takes into consideration selfish agents having interdependent valuations. We should note here that this mechanism is not specific to our problem scenario but can deal with *any* problem scenario where the agents have interdependent valuations. In section 5 we prove that the mechanism developed

is efficient, incentive-compatible and individually rational. We also discuss the complexity of calculating the efficient allocation under an example communication protocol. We conclude and discuss future work in section 6.

2 Related Work

Our work in this paper draws on two strands of research: information fusion in multi-sensor networks and mechanism design under interdependent valuations.

Traditionally, work addressing information fusion in multi-sensor networks has focused on using cooperative agents to achieve good system-wide performance [4, 5, 6]. In these cases, the agents are assumed to truthfully communicate whatever information they have gathered to other agents in the system. Competitive agents have been considered in the context of multi-sensor networks but this work has typically focused on methods to control some global resource, such as minimising the power consumption within the network [7, 8]. To date, therefore, the problem of maximising the overall information gain in a network of selfish sensors, where there are bandwidth constraints, has received little attention.

In order to address this problem a measure of information must be defined. These measures are often represented as a ‘distance’ and a number are commonly used in information fusion [9]. For instance, the Kullback-Liebler distance measures the impact of communicated data on both the precision and the mean of the prior probability density function. The Mahalanobis distance measures the impact that communicated data may have relative to the uncertainty in the prior. However both are functions of the data and the data has to be communicated to the agent before it can calculate the impact it will have on its own information. In contrast, for our mechanism to work, the agents must be able to value the data prior to its actual transmission. We thus adopt another common measure of uncertainty, namely the precision (or inverse covariance) of the data and derive a valuation function based on this measure [2].

This valuation function necessarily results in interdependent valuations and the work on mechanism design presented here draws on research in interdependent value auctions. In this area, Krishna [3] has considered efficient allocations for single items with single-dimensional signals. An alternative approach has been proposed by Dasgupta and Maskin who have also considered the case of two items, again with single-dimensional signals [10]. Our scenario is equivalent to an auction with multiple items with single-dimensional signals. Thus we adopt the approach of Krishna (due to the less stringent information requirements imposed on the agents and the more straightforward bidding protocol) and we extend it to cover this new case.




Figure 1: A sample scenario of a multi-sensor network.

3 Target Detection Scenario

In this section, we first describe our exemplar target detection scenario. We identify the important features of this scenario that we then formalise as a multi-agent system problem. We then develop the valuation function based around the information form of the Kalman filter.

3.1 Problem Description

We consider a scenario where a number of sensors are tasked with detecting targets. The sensors each have a partial and inaccurate view of the world and need to communicate with each other in order to increase this accuracy. The ‘view of the world’ in this case is a view of the target passing in the region that the sensors are monitoring. The communication network that the sensors use is constrained by a limited bandwidth. Therefore there is a need to globally decide on how to optimally allocate this bandwidth in order to best satisfy the sensors’ overall goal of forming an accurate view of the world.

In more detail, each sensor has two regions which they consider. There is a region of observation (ROO) which is the region where they can observe targets and a region of interest (ROI) which they wish to monitor. Figure 1 depicts a typical instance of a scenario where the ROI of sensor 1 is shown and there is a target within this ROI. We can observe that agent 1 can already know about this event in its ROI since this overlaps (as it usually does) with its ROO. However, due to noise inherent in the measurement process, agent 1 will have some uncertainty in its observation (e.g. the position, type or speed of the target may be described by a probability distribution rather than an absolute value). Agent 1 can however decrease this uncertainty by gaining data about the target from other agents, namely agents 3 and 5 (which also have the target in their ROO). However, if agent 1 can only receive data from one of these two agents due to bandwidth limitations, it will then have to decide as to which agent to gain the data from. This decision making process is further complicated if the other agents also have to make similar decisions. Thus different flows of data (i.e. descriptions of which sensors will transmit data and along

which path this data will flow) will yield different results in terms of the total reduction of the uncertainty (or the equivalent increase of information). Given this, the high level representation of our problem is then *to allocate the flow of data within the bandwidth constraints imposed by the communication network so as to optimise the overall gain in information each sensor has about its ROI*.

We tackle this problem by first modelling it as a multi-agent system. Each sensor is then viewed as an agent i , within a set of agents \mathcal{I} , which has data inf_i and a function $x_i(inf_i)$ that characterises how accurate this data is. The data has a size and thus a bandwidth requirement of bw_i . We initially consider the simplest communication protocol that exhibits a bandwidth constraint, and thus, we assume a broadcast protocol whereby any sensor can simultaneously transmit to all other sensors. The total bandwidth available for the transmission of the data is such that only a subset of the sensors can actually transmit their data.

In order to characterise this problem we first need to make a few assumptions about this scenario:

- The time taken in calculating the allocation of data and in communicating between agents is small compared to the time taken for another target to appear. This allows us time frames where the mechanism can be implemented.
- The agents have perfect and common knowledge about the sensor-network topology and their neighbours. This removes the problem of neighbour discovery in communication systems.

These assumptions thus permit us to concentrate solely on the issue of allocating the flow of data under the bandwidth constraints. We now need a way for each agent to value the data received from different agents based around the measure of data accuracy, x_i .

3.2 Valuation Function

We now develop our valuation function based on the information form of the Kalman filter. Now, in the standard Kalman filter, observations are of the form $z(t) = H(t)y(t) + n(t)$ where $y(t)$ is the state of the system at time t , $H(t)$ is the linear observation model and $n(t)$ is a zero mean random variable drawn from a normal distribution with variance R . The covariance update component, $P^{-1}(t | t)$, of the information form of the Kalman filter for N observations is:

$$P^{-1}(t | t) = P^{-1}(t | t-1) + \sum_{j=1}^N H^T(j)R^{-1}(j)H(j) \quad (1)$$

The summation in the above expression represents the decrease in covariance and thus the gain in information at time t when all the N observations are fused. In the case of our problem the value of receiving data from another agent can

thus be represented by the gain in information this observation engenders.

In order to achieve an efficient allocation, this gain in information must be calculated from the measure of the data accuracy prior to actually fusing it. Thus, we can represent the measure of accuracy of a data, x_j , as its covariance which is calculated from the covariance of its observation, $R(j)$:

$$x_j = H^T(j)R^{-1}(j)H(j) \quad (2)$$

Thus the gain in information of agent i , when all relevant data is transmitted to it and is fused, can be expressed as a sum of this measure of accuracy provided by each of the other agents:

$$v_i(\mathbf{x}) = x_i + \sum_{j \in -i} x_j \quad (3)$$

where $-i = \mathcal{I} \setminus i$.

Equations 2 and 3 thus cast our valuation function in the Kalman filter form. However, we need to modify this so as to incorporate the characteristics of our scenario. This is because all observations may not fall in an agent's ROI and furthermore an agent may not be able to receive all the data as a result of the bandwidth constraints of the communication network. Defining α_{ij} as the probability that the data observed by agent j is relevant to agent i and a vector \mathbf{f} as describing the flow of data in the network, then the expected valuation is :

$$\bar{v}_i(\mathbf{x}, f) = x_i + \sum_{j \in -i} f_{ij} \alpha_{ij} x_j \quad (4)$$

By slight abuse of notation, we shall hereafter refer to the expected valuation $\bar{v}_i(\cdot)$ as $v(\cdot)$.

From the valuation function, we can observe that the valuation of an agent i depends on x_j , which are signals measured by other agents. This puts us firmly in the realm of interdependent valuations. As a result, there are two conditions which are necessary in order to achieve an efficient allocation when considering selfish agents [11, 10]. Firstly:

$$\frac{\partial v_i(f, \mathbf{x})}{\partial x_j} = 0 \quad \forall i, j \in \mathcal{I} \quad (5)$$

and secondly:

$$\frac{\partial v_i(f, \mathbf{x})}{\partial x_i} > \frac{\partial v_j(f, \mathbf{x})}{\partial x_i} \quad \forall i, j \in \mathcal{I}, i \neq j \quad (6)$$

The first condition is automatically satisfied in our case since new data cannot decrease information. In the case of the second condition, it implies that we need to restrict an agent's ROI to its ROO. Otherwise, there may be an event outside its ROO which falls in its ROI such that data from another agent has a greater effect on its utility than its own data. This condition is necessary because otherwise selfish agents may profitably lie about their observed data and

derive positive utility from it. Furthermore, the overlap between the agent's ROO's must be such that this condition is satisfied (i.e. $\sum_{j \in -i} \alpha_{ij} < 1$). This means that the ROO of any agent can not be entirely overlapped by the ROO of other agents (i.e. no agent is redundant).

4 The Mechanism

As we have seen in section 3.2, the agents' utilities are interdependent, but, nevertheless, they can still be used to develop an efficient auction since they satisfy the conditions of equations 5 and 6. Furthermore, the agents only provide a single-dimensional measure, x_i , for the information they have and thus the impossibility result in [12] does not apply. Given this, we now detail a mechanism that is incentive-compatible, efficient and individually-rational for the case of multiple goods with single dimensional signals. Our mechanism proceeds as follows:

1. Each agent i transmits to the centre its valuation function $v_i(f, \mathbf{x})$ for all the possible allocations of the information flow $f \in \mathcal{F}$ where \mathcal{F} is the set of all feasible flows.
2. Each agent i also transmits its observed signal \hat{x}_i .¹
3. The centre then computes the optimal allocation f_0^* which is calculated as:

$$f_0^* = \arg \max_{f \in \mathcal{F}} \left(\sum_{i \in \mathcal{I}} v_i(f, \hat{\mathbf{x}}) \right) \quad (7)$$

4. The centre also calculates the payment r_i made by each agent i . To do this, the centre first finds the m next best allocations as the signal x_i is decreased until the presence of i makes no difference to the allocations. That is, find allocations $f_1^* \dots f_m^*$ and the signal values z_i^l such that:

$$\begin{aligned} z_i^l &= \inf \left\{ y_i : \sum_{i \in \mathcal{I}} v_i(f_l^*, y_i, \mathbf{x}_{-i}) \right. \\ &= \left. \sum_{i \in \mathcal{I}} v_i(f_{l+1}^*, y_i, \mathbf{x}_{-i}) \right\} \end{aligned} \quad (8)$$

(where each allocation f_l^* is different) until

$$\begin{aligned} z_i^m &= \inf \left\{ y_i : \sum_{i \in \mathcal{I}} v_i(f_{m-1}^*, y_i, \mathbf{x}_{-i}) \right. \\ &= \left. \sum_{i \in \mathcal{I}} v_i(f_m^*, y_i, \mathbf{x}_{-i}) \right\} \end{aligned} \quad (9)$$

¹Of course, \hat{x}_i may not be equal to x_i since the agent may strategise and lie about this value. However, we prove in section 5 that it is the best strategy for the agent (i.e. it maximises the utility of the agent) to set $\hat{x}_i = x_i$ and report this value truthfully.

where the allocation f_m^* is the optimal allocation when i does not exist i.e.

$$f_m^* = \arg \max_{f \in \mathcal{F}} \sum_{j \in \mathcal{I} \setminus i} v_j(f, \mathbf{x})$$

Then the transfer to buyer i is

$$r_i = \sum_{l=0}^{m-1} \left[\sum_{j \in \mathcal{I} \setminus i} v_j(f_l^*, z_i^l, \mathbf{x}_{-i}) - \sum_{j \in \mathcal{I} \setminus i} v_j(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right] \quad (10)$$

The above scheme rests upon making an agent derive a utility equal to the marginal contribution that its presence makes to the whole system of agents (which is the same intuition as used in the VCG [13]). Thus the additional part of this mechanism is to take into account the effect that an agent's signal x_i has on the overall utility of the system.

This mechanism is general in that it can be also applied to the case of independent valuations. In our scenario, such valuations arise when the regions of observation of the sensors do not overlap, and the agents are simply collecting, rather than combining, observations. In this case, the mechanism reduces to the well-known multiple-good VCG mechanism if we take the case of private values i.e when $v_i(\mathbf{x}, \cdot) = v_i(x_i)$. Then the optimal allocation is:

$$f_o^* = \arg \max_{f \in \mathcal{F}} \left(\sum_{i \in \mathcal{I}} v_i(f, x_i) \right) \quad (11)$$

To calculate the payment scheme, we first note that since x_i only affects i , then repeatedly decreasing x_i until the stopping condition on equation 9 does not change the valuation of the other agents $-i$ on the different allocations. This then implies that in the payment, as computed by equation 10, all the terms cancel each other except for the first and last leading to a payment of:

$$r_i = \sum_{j \in \mathcal{I} \setminus i} v_j(f_0^*, x_j) - \sum_{j \in \mathcal{I} \setminus i} v_j(f_m^*, x_j) \quad (12)$$

This is exactly the payment scheme used in the classical VCG mechanism.

5 Properties of the Mechanism

In this section, we discuss both the economic and computational properties of our mechanism and their implication for our multi-sensor network scenario.

5.1 Economic Properties

The mechanism that we present maintains the same desirable economic properties of the VCG mechanism, but as discussed earlier, does so in the face of interdependent, rather than independent, valuations. Specifically:

1. The mechanism is *incentive-compatible* in ex-post Nash Equilibrium. It is a best response strategy for the agents to reveal their types (x_i) truthfully even after they have complete information about the types of the other agents (\mathbf{x}). This is a strong equilibrium in that agents do not need to form prior beliefs about how other agents will behave [13] and they do not need to indulge in any strategic behaviour.
2. The mechanism is *individually rational*. A selfish agent will join the mechanism rather than opt out of it, since the utility derived by the agents within the mechanism is never less than that of agents outside the mechanism.
3. The mechanism is *efficient*. Since the agents truthfully reveal their types to the centre, the optimal allocation of communication bandwidth can be calculated and thus the information gain of the network as a whole is maximised despite the selfish actions of the individual agents.

For completeness, proofs of these results are included in the appendix, whilst a more general and detailed discussion of these results are presented in [14].

5.2 Computational Properties

In order to achieve the economic properties discussed in the previous section, the centre is required to compute the optimal allocation of bandwidth in equations 7, 8 and 9. This involves solving a combinatorial optimisation problem, and thus, the complexity of this task determines the computational load on the centre. Here we analyse the complexity when considering a broadcast communication network (the simplest network configuration which exhibits a bandwidth constraint).

In a broadcast communication protocol, any individual agent may simultaneously transmit its data to all other agents. To prevent several agents transmitting simultaneously, and thus interfering with each other, each time frame is typically divided into time slots of differing length and individuals are allocated particular slots depending on their bandwidth requirements. When the total bandwidth required by the agents exceeds that available, only a subset of agents are able to transmit their data. The role of the optimisation is to select the subset that maximises the total gain in information of the sensors.

In order to examine the complexity of this optimisation task, we can cast it as an instance of the knapsack problem, $\{(s_1, v_1), (s_2, v_2), (s_3, v_3), \dots, (s_n, v_n), K\}$, where K is the knapsack capacity and s_j and v_j are the size and value of item j . The goal is to choose a subset of items of maximum value with the total size at most K . In our case, s_j represents the bandwidth requirement of each agent, bw_j , and K represents the total bandwidth available, bw_{total} . Since

we are assuming that the total bandwidth available is insufficient to allow all the agents to transmit their data, then $\sum_j bw_j > bw_{\text{total}}$.

In order to map our problem into a knapsack problem, we must express the value which each agent contributes to the total by transmitting its data. Since this is a broadcast communication protocol, then we know that if $f_{ij} = 1$ for any i , then $f_{ij} = 1 \forall i \in \mathcal{I}$. Hence, if an agent j is transmitting data, then each of the other agents $\mathcal{I} \setminus j$ will derive a utility of $\alpha_{ij}x_j$. Thus, the value of each item, v_j , is represented by the total utility that it contributes to all other agents:

$$v_j = \sum_{\mathcal{I} \setminus j} \alpha_{ij}x_j. \quad (13)$$

Having thus cast the problem as a knapsack problem, we can use standard results to solve it. Although the knapsack problem is known to be *NP*-Complete, it is possible to calculate an exact solution in pseudo-polynomial time using dynamic programming [15]. This algorithm scales linearly with the number of agents and thus the optimum can be calculated efficiently, even in very large sensor networks.

6 Conclusions and Future Work

In this paper, we have addressed the problem of information fusion within a bandwidth-limited multi-sensor network. Since in many real applications the sensors are individually owned by different stakeholders, we have modelled them as selfish rational agents who are seeking to maximise their own gain through exchanging and receiving data from other sensors within the network. We presented a motivating problem scenario based of multi-sensor target detection and developed utility functions for these sensors based on the information form of the Kalman filter.

In order to ensure good global performance of the sensor network under our assumption of selfish sensors, we used a mechanism design approach to incentivise selfish sensors to act in a manner which results in the maximum information gain across the entire network. Since the utility function of sensors within a multi-sensor network are necessarily interdependent, we needed to extend the standard VCG mechanism to deal with this whilst conserving its desirable properties, namely efficiency, incentive-compatibility and individual rationality. The resulting mechanism was proved to have these properties and is applicable to any scenario where multiple goods (i.e. in this case observations of targets) are to be allocated. We showed that the mechanism deals with both interdependent and independent valuation scenarios, and reduces to the well-known VCG mechanism in the latter case. Finally, we analysed the computational complexity of finding the optimal solution required by the mechanism. In the case of a broadcast communication protocol, we showed that it reduces to the knapsack problem and thus an exact solution can be found in polynomial time.

As future work, we are investigating how we can apply the techniques developed here in richer information fusion

scenarios. One such extension is understanding how introducing multi-dimensionality to the valuation function leads to a sub-optimal allocation [12]. Another is understanding how we can extend the valuation function derived here to handle instances in which the agents communicate previously fused information as well as direct observations.

Our ultimate aim in this work is to extend the mechanism into a distributed one in which there is no central allocator of bandwidth. These distributed mechanisms are attractive since they have a number of advantages over centralised one in terms of robustness to failure and scalability. However, such a mechanism will need to consider how selfish agents might manipulate the communication of messages and the computation of the allocation for their own gain. As such, according to the revelation principle, the existence of the mechanism that we described here, is a necessary condition for the existence an equivalent distributed mechanism that has the same properties [16].

Acknowledgements

Rajdeep K. Dash is funded by a BAE Systems studentship and an ORS scholarship. This research was undertaken as part of the ARGUS II DARP (Defence and Aerospace Research Partnership). This is a collaborative project involving BAE SYSTEMS, QinetiQ, Rolls-Royce, Oxford University and Southampton University, and is funded by the industrial partners together with the EPSRC, MoD and Dti. The authors would like to thank Andy Wright and lead Rezek for helpful discussions.

References

- [1] A. MasColell, M. Whinston, and J.R. Green, *Microeconomic Theory*, Oxford University Press, 1995.
- [2] J. Manyika and H.F. Durrant-Whyte, *Data Fusion and Sensor Management: A Decentralized Information-Theoretic Approach*, Ellis Horwood, 1997.
- [3] V. Krishna, *Auction Theory*, Academic Press, 2002.
- [4] V. Lesser, C. Ortiz, and M. Tambe, Eds., *Distributed Sensor Networks: a multiagent perspective*, Kluwer Publishing, 2003.
- [5] S. H. Clearwater, Ed., *Market-Based Control - A paradigm for distributed resource allocation*, World Scientific, 1996.
- [6] M. Rosencrantz, G. Gordon, and S. Thrun, "Decentralized sensor fusion with distributed particle filters," in *Proceedings of the 19th Annual Conference on Uncertainty in Artificial Intelligence (UAI-03)*, San Francisco, CA, 2003, pp. 493–500, Morgan Kaufmann Publishers.

- [7] B. P. Gerkey and M. J. Mataric, "Sold!: Auction methods for multi-robot coordination," *IEEE Transactions on Robotics and Automation, Special Issue on Multi-robot Systems*, vol. 18, no. 5, pp. 758 – 768, October 2002.
- [8] A. Rogers, E. David, and N. R. Jennings, "Self-organized routing for wireless micro-sensor networks," *IEEE Transactions on Systems, Man, and Cybernetics (Part A)*, vol. 35, no. 3, pp. 349–359, 2005.
- [9] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration for tracking applications," *IEEE Signal Processing Mag.*, pp. 61–72, March 2002.
- [10] P. Dasgupta and E. Maskin, "Efficient auctions," *Quarterly Journal of Economics*, vol. 115, pp. 341–388, 2000.
- [11] R. Mirrlees, "An exploration in the theory of optimum income taxation," *Review of Economic Studies*, vol. 38, pp. 175–208, 1971.
- [12] P. Jehiel and B. Moldovanu, "Efficient design with interdependent valuations," *Econometrica*, vol. 69, no. 5, pp. 1237–59, 2001.
- [13] R. K. Dash, D. C. Parkes, and N. R. Jennings, "Computational mechanism design: A call to arms," *IEEE Intelligent Systems*, vol. 18, no. 6, pp. 40–47, 2003.
- [14] R. K. Dash, A. Rogers, and N. R. Jennings, "A mechanism for multiple goods and interdependent valuations," in *Proceedings of the 6th International Workshop on Agent-Mediated E-Commerce*, New York, USA, 2004, pp. 197–210.
- [15] M. R. Garey and D. S. Johnson, *Computers and Intractability - A Guide to the Theory of NP-Completeness.*, Freeman, San Francisco, 1979.
- [16] R.B. Myerson, "Incentive compatibility and the bargaining problem," *Econometrica*, vol. 47, pp. 61–73, 1979.

Appendix

Proof of Economic Properties

In this section we present proofs of the economic properties of the mechanism, namely that it is *incentive-compatible*, *efficient* and *individually rational*.

Proposition 1. *The mechanism is incentive-compatible in ex-post Nash Equilibrium.*

A mechanism is *incentive-compatible* in ex-post Nash Equilibrium if it is a best response strategy for the players to reveal their types truthfully even after they have complete information about the signal values \mathbf{x} .

Proof. Let $v_{-i}(\cdot) = \sum_{j \in -i} (v_j(\cdot))$ and $v_{\mathcal{I}}(\cdot) = \sum_{i \in \mathcal{I}} (v_i(\cdot))$. Suppose now that all players except i report their signals truthfully (i.e. $\hat{\mathbf{x}}_{-i} = \mathbf{x}_{-i}$). Let the optimal allocation when i reports truthfully be f_0^* . We can then analyse the utility $u_i(\cdot)$ that agent i derives by reporting a certain \hat{x}_i . There are two cases that should be analysed namely when $\hat{x}_i < x_i$ and $\hat{x}_i > x_i$. The utility of an agent on reporting $\hat{x}_i = x_i$ is :

$$u_i(f_0^*, \mathbf{x}) = v_i(f_0^*, \mathbf{x}) + \sum_{l=0}^{m-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \quad (14)$$

Now suppose an agent reports $\hat{x}_i \neq x_i$ but this does not change the optimal allocation f_0^* implemented. Then, $u_i(f_0^*, \mathbf{x}) = u_i(f_0^*, \hat{x}_i, \mathbf{x}_{-i})$. This is because, then the agent will derive the same value $v_i(f_0^*, \mathbf{x})$ if the allocation does not change and the payment will be the same as the signals $z_i^0 \dots z_i^m$ are computed by the centre. Now consider the case that an agent reports $\hat{x}_i < x_i$ such that this changes the allocation. Then some other optimal allocation, which is necessarily one of the allocations f_1^*, \dots, f_m is implemented. Denote the resulting allocation when $\hat{x}_i < x_i$ as f_n^* , i.e. $z_i^{n-1} < \hat{x}_i \leq z_i^n$.

The utility that the agent gets from this new allocation is then:

$$u_i(f_n^*, \mathbf{x}) = v_i(f_n^*, \mathbf{x}) + \sum_{l=n}^{m-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \quad (15)$$

The difference, $D_n = u_i(f_0^*, \mathbf{x}) - u_i(f_n^*, \mathbf{x})$ between truthful reporting and over-reporting (as given by equations 14 and 15 respectively) is:

$$\begin{aligned} D_n &= v_i(f_0^*, \mathbf{x}) - v_i(f_n^*, \mathbf{x}) \\ &+ \sum_{l=0}^{n-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, z_i^0, \mathbf{x}_{-i}) - v_{-i}(f_n^*, z_i^n, \mathbf{x}_{-i}) \\ &- v_i(f_n^*, \mathbf{x}) + \sum_{l=1}^n \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_l^*, z_i^{l+1}, \mathbf{x}_{-i}) \right) \end{aligned}$$

Since $\frac{\partial v_{-i}(f_i^*, \mathbf{x})}{\partial x_i} \geq 0$, we thus have:

$$\begin{aligned} D_n &\geq v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, z_i^0, \mathbf{x}_{-i}) \\ &- v_{-i}(f_n^*, z_i^n, \mathbf{x}_{-i}) - v_i(f_n^*, \mathbf{x}) \end{aligned}$$

But, given that $\frac{\partial v_{-i}(\cdot)}{\partial x_i} < \frac{\partial v_i(\cdot)}{\partial x_i}$, we then have that:

$$D_n > v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, \mathbf{x}) - v_{-i}(f_n^*, z_i^n, \mathbf{x}_{-i}) - v_i(f_n^*, z_i^0, \mathbf{x}_{-i})$$

Using the condition that $\frac{\partial v_{-i}(f_l^*, \mathbf{x})}{\partial x_i} \geq 0$ again leads to:

$$D_n > v_{\mathcal{I}}(f_0^*, \mathbf{x}) - v_{\mathcal{I}}(f_n^*, z_i^0, \mathbf{x}_{-i}) > v_{\mathcal{I}}(f_0^*, z_i^0, \mathbf{x}_{-i}) - v_{\mathcal{I}}(f_n^*, z_i^0, \mathbf{x}_{-i})$$

By the optimality of f_0^* (i.e. f_0^* is always calculated such that it maximises the sum of valuations as per equation 7), and we thus have $D_n > 0$.

On the other hand, if an agent reports $\hat{x}_i > x_i$ and this induces an allocation f_{-n}^* , then the utility it derives is :

$$u_i(f_{-n}^*, \mathbf{x}) = v_i(f_{-n}^*, \mathbf{x}) + \sum_{l=-n}^{m-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \quad (16)$$

The difference, $D_{-n} = u_i(f_0^*, \mathbf{x}) - u_i(f_{-n}^*, \mathbf{x})$ between truthful reporting and under reporting (as given by equations 14 and 16 respectively) is

$$\begin{aligned} D_{-n} &= v_i(f_0^*, \mathbf{x}) - v_i(f_{-n}^*, \mathbf{x}) \\ &\quad - \sum_{l=-n}^{-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, z_i^{-1}, \mathbf{x}_{-i}) - v_i(f_{-n}^*, \mathbf{x}) \\ &\quad - v_{-i}(f_{-n}^*, z_i^{-n}, \mathbf{x}_{-i}) \\ &\quad + \sum_{l=-n+1}^{-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^{l+1}, \mathbf{x}_{-i}) \right) \end{aligned}$$

Since $\frac{\partial v_{-i}(f_l^*, \mathbf{x})}{\partial x_i} \geq 0$, we thus have:

$$D_{-n} \geq v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, z_i^{-1}, \mathbf{x}_{-i}) - v_i(f_{-n}^*, \mathbf{x}) - v_{-i}(f_{-n}^*, z_i^{-n}, \mathbf{x}_{-i})$$

Then:

$$D_{-n} \geq v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, \mathbf{x}) - v_i(f_{-n}^*, \mathbf{x}) - v_{-i}(f_{-n}^*, z_i^{-n}, \mathbf{x}_{-i})$$

But $\frac{\partial v_{-i}(\cdot)}{\partial x_i} < \frac{\partial v_i(\cdot)}{\partial x_i}$ and thus:

$$\begin{aligned} D_{-n} &> v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, \mathbf{x}) \\ &\quad - v_i(f_{-n}^*, z_i^{-n}, \mathbf{x}_{-i}) - v_{-i}(f_{-n}^*, \mathbf{x}) \\ &> v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, z_i^{-n}, \mathbf{x}_{-i}) \\ &\quad - v_i(f_{-n}^*, \mathbf{x}) - v_{-i}(f_{-n}^*, \mathbf{x}) \\ &> v_{\mathcal{I}}(f_0^*, \mathbf{x}) + v_{-i}(f_{-n}^*, \mathbf{x}) \end{aligned}$$

Again, by optimality of f_0^* , we have $v_{\mathcal{I}}(f_0^*, \mathbf{x}) - v_{\mathcal{I}}(f_{-n}^*, \mathbf{x}) \geq 0$ and hence $D_{-n} > 0$. We thus see that i derives highest utility when reporting $\hat{x}_i = x_i$. \square

Proposition 2. *The mechanism is efficient.*

This implies that the centre finds the outcome such that

$$f^* = \arg \max_f \sum_{j \in \mathcal{I}} u_j(f, \mathbf{x}) \quad (17)$$

Proof. The above is a result of the incentive-compatibility of the mechanism. Since the goal of the centre is to achieve efficiency, then given truthful reports, the centre will achieve efficiency. \square

Proposition 3. *The mechanism is individually rational.*

A mechanism is *individually rational* if there is an incentive for agents to join it rather than opting out of it. We begin by assuming that the utility an agent derives from not joining the mechanism is 0. Then, we need to prove that the utility an agent derives in the mechanism is always ≥ 0 .

Proof. Given that the agents are incentivised to report truthfully, an agent i derives utility:

$$\begin{aligned} u_i(f_0^*, \mathbf{x}) &= v_i(f_0^*, \mathbf{x}) + \sum_{l=0}^{m-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_0^*, \mathbf{x}) + \sum_{l=0}^{m-1} \left(v_{\mathcal{I}}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{\mathcal{I}}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) - \sum_{l=0}^{m-1} \left(v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_i(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \end{aligned} \quad (18)$$

Since $v_{\mathcal{I}}(f_l^*, z_i^l, \mathbf{x}_{-i}) = v_{\mathcal{I}}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i})$ (from equation 8):

$$\begin{aligned} u_i(f_0^*, \mathbf{x}) &= v_i(f_0^*, \mathbf{x}) - \sum_{l=0}^{m-1} \left(v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_i(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_0^*, \mathbf{x}) - v_i(f_0^*, z_i^0, \mathbf{x}_{-i}) + v_i(f_m^*, z_i^m, \mathbf{x}_{-i}) \\ &\quad + \sum_{l=1}^{m-1} \left(v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_i(f_l^*, z_i^{l+1}, \mathbf{x}_{-i}) \right) \end{aligned} \quad (19)$$

From equation 9, $v_i(f_m^*, z_i^m, \mathbf{x}_{-i}) = 0$. Furthermore, since $\frac{\partial u_i(K, \mathbf{x})}{\partial x_i} \geq 0$, then $\sum_{l=1}^{m-1} \left(v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_i(f_l^*, z_i^{l+1}, \mathbf{x}_{-i}) \right) > 0$ and $v_i(f_0^*, \mathbf{x}) - v_i(f_0^*, z_i^0, \mathbf{x}_{-i}) > 0$. Hence $u_i(f_0^*, \mathbf{x}) > 0$. \square