

Computational Mechanism Design for Information Fusion within Sensor Networks

A. Rogers, R. K. Dash and N. R. Jennings
School of Electronics and Computer Science
University of Southampton
Southampton, UK.
{acr,rkd02r,nrj}@ecs.soton.ac.uk

S. Reece and S. Roberts
Department of Engineering Science
University of Oxford
Oxford, UK.
{reece,sjrob}@robots.ox.ac.uk

Abstract - *Conventional centralised information fusion and control architectures will be challenged by developments in sensor networks that allow sophisticated autonomous sensors, owned by different stakeholders with individual goals, to interact and share information. Given this, we advocate the use of tools and techniques from computational mechanism design (CMD), a field at the intersection of computer science, game theory and economics, to address the challenges posed by these networks. In particular, CMD allows us to engineer networks with desirable system-wide properties, in which sensors act as rational selfish agents, each attempting to fulfill their own individual goals through the exchange of observations and information. In this paper, we present our work developing such networks. Specifically, we discuss our development of a generic and principled information valuation metric for sensor networks and we report our experiences applying it within a real world information fusion sensor network scenario.*

Keywords: sensor network, information economy, multi-agent system, computational mechanism design, data fusion.

1 Introduction

Wireless micro-sensor networks, in which small cheap sensors are distributed within large and possibly hostile environments, are increasingly seen as attractive solutions to the problem of wide-area monitoring in environmental, security and military applications [1, 2, 3]. To date, most research has considered small experimental networks, in which a centralised data collection architecture has been implemented. That is, the individual sensors have little autonomy and simply send their data, observations or information to a central point at which it is fused, aggregated or acted upon.

However, this approach is likely to be challenged within future networks. As sensors become more sophisticated and are able to reconfigure themselves to make more effective use of their local resources (e.g. managing their own sampling and reporting schedule to make the most effective use of battery power), it becomes increasingly likely that attempting to opti-

mise the performance of the entire network using a centralised control regime will run into severe processing and communication bottlenecks. In addition, as sensors become ubiquitous within the environment, and thus sensors owned by different stakeholders have the possibility of exchanging observations and information (in order to resolve uncertainty in their own observations, to extend their observation range or to relay observations to a third party) it becomes increasingly difficult to identify where such a centralised point of control can be located.

Thus, against this background, in this paper we report on the challenges of applying agent-based control within sensor networks [4]. Specifically, we present our work using techniques from *computational mechanism design* [5], a field at the intersection of computer science, game theory and economics, to engineer sensor networks with predictable and desirable properties that can operate effectively, even when the sensors within the network have individual goals that they are acting autonomously and selfishly to achieve. We describe a generic and principled metric for valuing observations made by these sensors, and we show that this metric allows us to use standard results from mechanism design to engineer these systems. We present an example of these techniques in operation in a simulated aerial surveillance scenario in which sensors with individual goals make efficient use of a limited global resource (in this case, the limited bandwidth of the communication network that connects the sensors together).

The remainder of this paper is organised as follows. In section 2 we describe in more detail current sensor networks and discuss the advances that will challenge the conventional control architectures. In section 3 we present our valuation metric and describe how this metric is applied within computational mechanism design. In section 4 we present a detailed description of our aerial surveillance scenario, and finally, we discuss future work and conclude in section 5.

2 Sensor Networks

Recent years have seen an explosion in the development of small experimental sensor networks. These have been demonstrated for purposes as diverse as ra-




Figure 1: Schematic of the wireless network used in the Glacsweb project and a photograph of the actual sensors embedded within the glacier.

diation monitoring [2], detecting sniper fire in urban combat environments [3] and helping to detect and predict floods in river estuaries [6]. Since these networks have typically served as test-beds to enable the development of the sensor technology, most have used conventional data collection and distribution architectures in which sensor data is collected, fused (or aggregated) and analysed at a single central control point. An example of such a network is the GlacsWeb project (see figure 1) [1]. Here, small battery powered sensors have been embedded into the ice of the Briksdalsbreen glacier in Norway. Each sensor measures the temperature, pressure and orientation of the ice and then employs a store-and-forward data transmission protocol, whereby, using low power radio transmissions, they are polled at regular intervals for data by a base station mounted on the surface of the glacier. The data is then aggregated and forwarded onto a reference station and, from here, is sent via standard internet protocols to a dedicated server for recording and analysis (available online at <http://leo.ecs.soton.ac.uk/glacsweb/>).

However, advances in sensor technology are beginning to challenge this architecture. As sensors become cheaper and easier to deploy, the size of these networks will undoubtedly increase. In addition, as the sophistication and processing power of the sensors increases, they are likely to transition from being simple passive devices that record and transmit local measurements, to become devices that can actively adjust their configuration to make best use of local resources or observations. However, it is unlikely that these configuration decisions can be made independently. There will almost certainly be dependences between the sensors. For example, sensors may be able to independently adjust their own sampling and reporting schedules to make more efficient use of their limited battery power, but if they employ ‘multi-hop’ routing, they must also coordinate with other sensors to ensure that these sensors will be active and that there will be a viable transmission route available when they come to transmit [7, 8]. In addition, sensors may be able to fulfil their sampling requirements, whilst simultaneously avoiding making and transmitting redundant information, by opportunistically and dynamically making use of the measurements made by other nearby sensors.

In these settings, finding the optimal configuration of each sensor within the sensor network becomes a large combinatorial optimisation problem. It thus becomes likely that the capacity of the communication network over which control and data messages must be sent and the processing speed of the central control point become significant bottlenecks that limit the effectiveness of the entire network. This is particularly true in dynamic environments where many of the most interesting applications lie, and, as such, decentralised control regimes that allow the individual sensors to make autonomous local decisions are particularly attractive.

The rationale for these decentralised approaches is further strengthened by the observation that as these sensor networks become ubiquitous within the environment, it is likely that individual sensors within a particular network may benefit from interacting with sensors that are owned by different stakeholders. Indeed, these sensors may have been deployed at a different time and for different purposes. For example, a farmer using a sensor network to measure local weather conditions (e.g. rainfall, cloud cover, humidity and soil moisture levels), in order to control an irrigation system, may benefit from using measurements from neighbouring landowners in order to track changing local weather patterns. Moreover, local householders, without environmental sensors of their own, may also wish to make use of information from these sensor networks, in order to manage their own climate control systems.

Now, these different stakeholders may have very different aims or goals; one may wish to conserve the power and battery life of their own sensors, whilst another may wish to gain the maximum possible information in the shortest possible period of time. Thus, processing and communication bottlenecks aside, the presence of multiple stakeholders with different aims makes it impossible to identify where a centralised controller should be located, or on whose behalf it should be acting. As such, we not only require a decentralised control protocol, but we also require one that allows these individual stakeholders to work autonomously toward their own individual goals, whilst excluding the possibility that sensors owned by any one stakeholder can take unfair advantage of other stakeholders (e.g.

by using the communication and power resources of the other stakeholder to satisfy their own goals without providing anything in exchange or by monopolising system-wide resources to the detriment of other stakeholders).

It is this latter requirement that motivates our use of techniques from computational mechanism design [5]. These techniques are attractive as the economic metaphor on which they are predicated naturally describes systems composed of autonomous agents that are each attempting to satisfy their own individual goals under a decentralised control regime. To date, they have been investigated in applications such as resource allocation within Grid computing [9, 10] and routing allocation within communication networks [11]. However, in order to apply them within the setting of sensor networks in which observations or data (rather than tasks or resources) are exchanged between individuals, a valuation metric is required, and we discuss the valuation metric that we use in the next section.

3 Information Valuation Metric

Computational mechanism design provides a principled framework with which to design systems that exhibit desirable global properties, despite the selfish actions and goals of the constituent parts [5]. It is an extension of the economic field of mechanism design and addresses the additional challenges imposed by a computational setting (i.e. agents that are computationally limited, communication that is not cost or error free and settings that are open and dynamic). At its core, is the notion that agents hold or require valued items, and are seeking to maximise their own utility through the exchange of these items. In the real world, these items may be goods or services, and thus they will have real monetary value. When applied in the context of sensors that are exchanging information or observations within a network, we require a means for the sensors to place a value on any piece of information or observation.

Whilst there are many ad hoc means by which a value can be determined within any particular setting [12], we use a principled valuation metric that is based on the Fisher information of each observation [13]. As such, rather than being restricted to a single setting, this metric is generic in that it can be applied in any context where sensors make and exchange imprecise observations and thus must deal with uncertainty. In addition, this measure is particularly attractive as under the rules of probabilistic inference, when two independent observations are fused, the Fisher information of the fused estimate is simply the sum of the Fisher information of the two individual observations. Thus, if a sensor makes an observation with information content I_1 and also receives an observations from another sensor with information content I_2 , the total information when these two observations are fused together is simply $I_1 + I_2$. This notion is important as conventional mechanism design assumes that valuations

of goods are additive.

To illustrate this measure, we can initially consider a simple one-dimensional case in which the sensors make a single measurement of some physical characteristic of the environment. Since this observation will typically be noisy or imprecise, it is common to model the measurement as a Gaussian, whose mean reflects the estimated value of the physical characteristic being measured, and whose variance, σ^2 , describes how accurate this measurement is. In the case of this Gaussian distribution, the Fisher information is simply the inverse of this variance:

$$I = \frac{1}{\sigma^2} \quad (1)$$

Thus, as the noise in the measurement decreases, and thus the variance of the measurement decreases, the information content of the estimate increases. Hence, we can use this value, the inverse of the variance, as a principled measure of the value or information content of this observation. In the next section we present a more complex setting and describe how we use the same information metric to value observations in an aerial surveillance scenario.

4 Example Information Fusion Sensor Network Scenario

To illustrate the use and generality of the valuation metric, we now describe how it was used in a real world aerial surveillance scenario such as that posed by disaster relief. In such scenarios, multiple emergency response agencies with aerial vehicles of different capabilities, must interact in order to locate casualties. As such, each aircraft has an individual goal (e.g. to maximise their own information) that may be subject to rapid changes due to local factors (e.g. a change in the capability or priority of the aircraft). Thus, the mechanisms presented here provide a decentralised control regime that allows the aircraft to act autonomously toward their own individual goals, whilst also ensuring that the system as a whole displays desirable and predictable properties.

Thus, in this case, the sensor network is formed from sensors that are mounted onboard these aerial vehicles (see figure 2). Each sensor within the network is tasked with detecting and tracking multiple targets within a *region of observation* immediately surrounding itself and is provided with an imprecise estimate of its own location by the navigation system of the aerial vehicles to which it is mounted. Within its region of observation, the sensor is able to estimate the position of each target by making noisy or imprecise measurements of the range and bearing of the target from itself. However, in order to better resolve the uncertainty in these position estimates, the sensors must acquire target observations from neighbouring sensors and then fuse these observations with their own.

The sensors are connected together via a communication network that has a limited bandwidth, and, thus, there must be some coordination to determine which observations should be sent to which sensors in

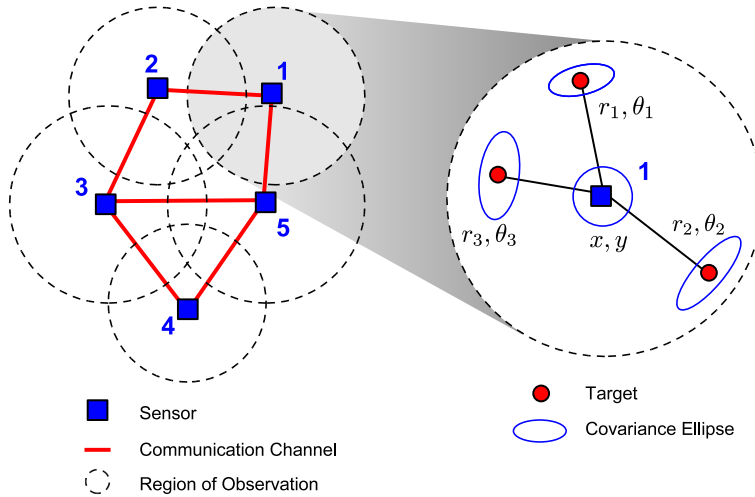


Figure 2: Schematic diagram of the network used in our information fusion sensor network application and an exploded view of sensor 1 making noisy or imprecise range and bearing measurements of several targets within its region of observation.

order to make best use of this limited global resource. Now, in conventional architectures, such coordination is easy to achieve since each sensor is willingly working toward the goals of the entire system. However, in our scenario, each sensor may be owned by a different stakeholder and is thus selfishly seeking to maximise the accuracy of their own target position estimates. In this case, there is no incentive for any sensor to transmit observations to its neighbours. Indeed, each sensor has a positive disinclination to do so, since transmitting observations will occupy valuable bandwidth which may be used to receive observations from other agents. Thus our mechanism must incentivise the sharing of observations and ensure that the global resource of communication bandwidth is used effectively, despite the selfish goals and actions of the participating sensors.

4.1 Information Valuation Metric

As discussed earlier, in order to apply the tools of mechanism design we must first apply our valuation metric. Clearly, in this case, it is more complex than the previous one dimensional example. Rather than a single environmental parameter, each sensor has an imprecise estimate of its own location, and imprecise measurements of the range and bearing to multiple targets within its own region of observation (see figure 2). The uncertainty of these measurements cannot be represented by a single value of variance, but must be represented by a covariance matrix (and as in figure 2, can be shown graphically as a corresponding set of covariance ellipses).

In more detail, we assume that a sensor has an estimate of its own global coordinate position that is expressed as a joint Gaussian estimate with mean (x, y) and covariance P_0 . If we consider a single target at (x_t, y_t) , we assume that the sensor makes a noisy range measurement with mean r and variance σ_r^2 , and a noisy bearing measurement with mean θ and variance σ_θ^2 .

The sensor represents its estimate of its own location and that of the target by a 4-vector $(\hat{x}, \hat{y}, \hat{x}_t, \hat{y}_t)$. Now, the total uncertainty in the global coordinate position of the target, given the uncertainty in the sensors position and its noisy range and bearing measurements, can be expressed as a covariance matrix P given by:

$$P^{-1} = \begin{pmatrix} P_0^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} + dHR^{-1}dH^T \quad (2)$$

where:

$$R = \begin{pmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_\theta^2 \end{pmatrix} \quad (3)$$

and dH is the Jacobian of the observation model H where:

$$H(x, y, x_t, y_t) = \begin{pmatrix} \sqrt{(x - x_t)^2 + (y - y_t)^2} \\ \arctan\left(\frac{y - y_t}{x - x_t}\right) \end{pmatrix} \quad (4)$$

As before, the information content of the observations can be calculated as a scalar value by simply calculating the trace of the inverse of their covariance matrix:

$$I = Tr(dHR^{-1}dH^T) . \quad (5)$$

Again, this measure is additive when two independent observations are fused together. Now extending this representation to any given number and distribution of targets, allows our sensors to value the information content of their observation. As before, the more precise the measurements, the smaller the covariance ellipses and thus the greater the information content of the covariance matrix. In addition, given the information content of an observation from another sensor, the sensor is able to simply sum this value with the information content of its own observations, in order to calculate what the total information content will be after fusing these two observations together.

4.2 Mechanism Design

Having derived the valuation metric, we can apply the mechanism. Our choice here depends on the details of the scenario and on the global system-wide properties that we wish to enforce. In this case, we wish to ensure that the global bandwidth resource is used efficiently (that is, ensure that given the limited bandwidth, the information gain of the entire network is maximised), thus we impose a mechanism whereby sensors are called upon to privately reveal the information content of observations to an auctioneer. This auctioneer then allocates the limited bandwidth of the communication network to those sensors whose observations will yield the highest system-wide information gain.

However, since each sensor is individually attempting to maximise its information regarding its own region of observation, with a simple mechanism there is an opportunity for a sensor to behave strategically (e.g. by understating the information content of its own observations, in an attempt to ensure that bandwidth is allocated to other sensors whose observations it can make use of or by overstating it, in order to deny bandwidth to other sensors). Such strategic behaviour is generally undesirable since it reduces the overall efficiency of the network and is computationally expensive for the individual sensors.

Thus, we have focused our research onto a sub-class of mechanisms that are said to be strategy-proof or incentive compatible [5]. That is, within the mechanism, the sensors have a dominant strategy (one which they should adopt regardless of the behaviour of other sensors) to truthfully reveal their private information regarding the value of observations to the auctioneer. To date, we have explored two settings which differ in respect of who provides valuation of the observations. In the first, the originator of the information provides the valuation, and in the second, the destination sensor provides the valuation.

4.2.1 Originator Sensor Valuation

In this setting, we have considered a sensor network that uses a broadcast communication protocol. That is, each sensor transmits its observations to all others, but the total amount that can be transmitted in any time frame is constrained. Under this protocol, the sensors provide the auctioneer with a valuation for the entire set of their target observations. The auctioneer uses this information, along with knowledge of which observations are relevant to which sensor, to calculate the optimal allocation of bandwidth that maximises the information gain of the entire network [14]. In exchange for sharing observations with other sensors, the originating sensor receives payments. These payments are calculated by the auctioneer in order that the sensors maximise their utility (expressed as the sum of their own individual information gain and payments that they receive) when they truthfully report their valuations to the auctioneer¹.

¹Such payments are most commonly encountered within Vickrey-Clarke-Groves (VCG) mechanisms where they reflect

4.2.2 Destination Sensor Valuation

In more recent work, we have extended this mechanism so that rather than transmitting their whole set of observations, sensors may transmit and exchange smaller subsets of observations. In addition, we have relaxed the assumption of a broadcast communication protocol, and assume that the sensors are connected by a communication network that allows peer-to-peer communication, with an overall total bandwidth limit.

In doing so, we have moved the requirement to value observations from the originator sensor to the receiving sensor. That is, as part of the auction mechanism, the receiving sensor receives a sample of the observations from another sensor, and is required to place a non-zero valuation on observations of targets that are of interest. Since, observations of multiple targets are correlated through the imprecise estimate of the sensors' own location, the valuations of different combinations of targets are super-additive, and thus, this results in a combinatorial auction. Again, an incentive compatible payment scheme ensures that the individual sensors maximise their own utility by truthfully reporting their valuations to the auctioneer.

This mechanism makes more efficient use of the limited communication bandwidth, but the actual operation of the allocation mechanism is more costly in terms of communication and computation. Thus, the allocation is only performed periodically when a significant change in the distribution of the targets occurs (specifically, when a target enters or leaves the region of observation of any of the sensors).

4.3 Java Simulation

In order to evaluate and demonstrate these mechanisms we have implemented them within a Java simulation where each agent runs in an asynchronous thread in real time. Figures 3 and 4 shows the system running with a 3D and an overhead view where the allocation of bandwidth and the results of the auction process can be tracked (a Java applet version of this system is available at <http://www.ecs.soton.ac.uk/~acr/demonstrator/>). At the specific instance in time shown in figure 4, the bandwidth is severely limited (the graph in the top right of the screen shows a red line indicating the bandwidth required for full exchange of observations, a green line indicating the current bandwidth limit and a blue line that indicates the bandwidth that is actually being used) and, thus, although a target falls into the region of observation of both sensors 2 and 3, there is insufficient bandwidth for these sensors to exchange observations (allocated bandwidth is indicated by the yellow lines between sensors). However, where observations are being exchanged (e.g. between sensors 1 and 3), the uncertainty in the position of the targets is resolved, and thus the covariance ellipses around these targets are greatly reduced in size and there is an increase in the

the marginal contribution that each agent makes to the overall system [5]

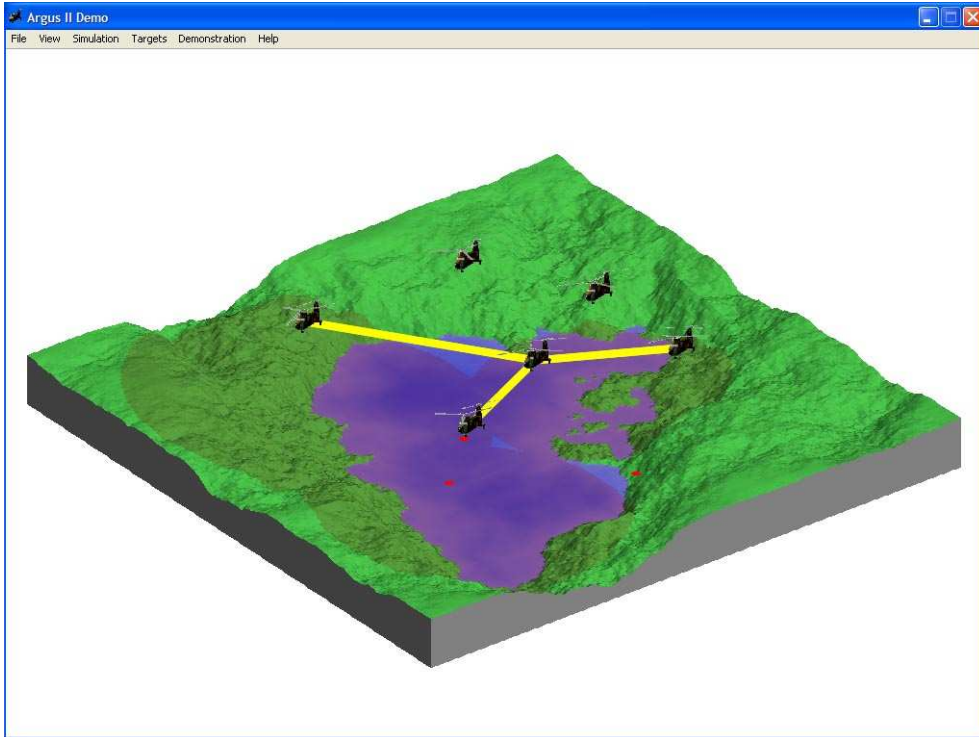


Figure 3: Example 3D view of the sensor network system showing the current communication allocation.

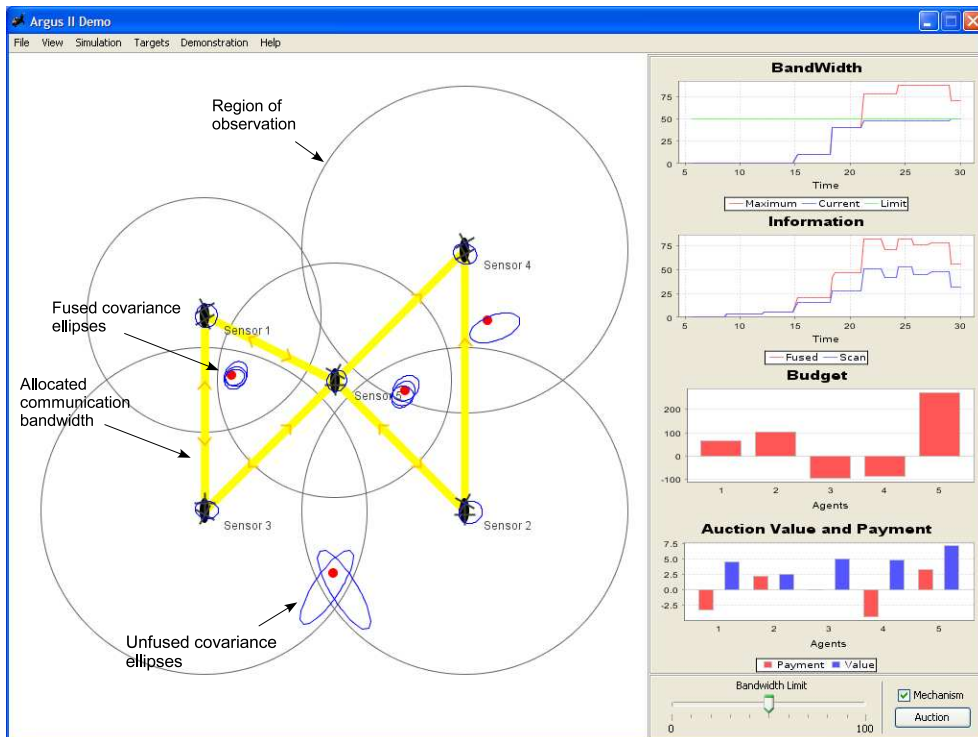


Figure 4: Example sensor network system showing the auction allocation in process and the resulting communication allocation.

total information gain of the network (shown in the second graph on the right).

The value of information that each sensor receives from other sensors and the payments that they receive in exchange for transmitting their own observations are shown in the bar-graph at the bottom right of the display (note that sensors 1 and 4 both have negative payments since they are currently receiving more information than they are transmitting; indeed, sensor 4 is transmitting no information at all). When sensors truthfully reveal the information content of their observations, they maximise their individual information gain and maintain their budget of currency (shown on the right of the display). However, a sensor that does not adopt this strategy (due to faulty, strategic or malicious behaviour), will not achieve these aims and will gradually deplete its budget. Such sensors can be recognised and removed from the network, thus incentivising the truthful reporting that is necessary to ensure that the constrained bandwidth of the sensor network is allocated to achieve the system-wide goal of maximising the information gain of the entire sensor network.

5 Conclusions

The work described in this paper has shown how we can apply agent-based control, and specifically mechanism design, within sensor networks, in which the constituent sensors are owned by different stakeholders that have their own goals and objectives. Our system has a principled currency grounded in information theory and uses computational mechanism design to engineer a sensor network with a particular desirable system-wide property (in this case, the efficient use of limited bandwidth capacity), despite the selfish goals and actions of the individual sensors. Our future work in this area is following two parallel tracks. In one track we are investigating the use of decentralised mechanisms in which the role of the auctioneer is distributed amongst the participating agents. In the second track, we are extending the range of scenarios considered by including information fusion whereby the sensors share not only observations, but state estimates that may contain correlated information. In this setting, we are no longer attempting to make the most efficient use of a global resource, but instead, we are seeking to configure the sensors (orientating them to track particular sub-sets of targets) such that the entire network achieves the maximal uncorrelated information gain. Whilst this research is ongoing, it is clear that if we are to build the type of sensor networks described in this paper, then computational mechanism design offers an invaluable and powerful set of tools and techniques.

Acknowledgments

The work reported on here undertaken as part of the ARGUS II DARP (Defence and Aerospace Research Partnership). This is a collaborative project involving

BAE SYSTEMS, QinetiQ, Rolls-Royce, the University of Oxford and the University of Southampton, and is funded by the industrial partners together with the EPSRC, MoD and Dti.

References

- [1] K. Martinez, J. Hart, and R. Ong. Environmental sensor networks. *IEEE Computer*, 37(8):50–56, 2004.
- [2] R. J. Nemzek, J. S. Dreicer, D. C. Torney, and T. T. Warnock. Distributed sensor networks for detection of mobile radioactive sources. *IEEE Transactions on Systems, Man, and Cybernetics - Part A*, 51(4):1693–1700, 2004.
- [3] G. Simon, M. Maróti, and Á. Lédeczi. Sensor network-based countersniper system. In *Proceedings of The Conference on Embedded Networked Sensor Systems*, pages 1–12. ACM Press, 2004.
- [4] N. R. Jennings. An agent-based approach for building complex software systems. *Communications of the Association for Computing Machinery*, 44(4):35–41, 2001.
- [5] R. K. Dash, D. C. Parkes, and N. R. Jennings. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.
- [6] D. De Roure. Floodnet: A new flood warning system. *Ingenia*, (23):48–51, June 2005.
- [7] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy efficient communication protocol for wireless microsensor networks. In *Proceedings of the IEEE Hawaii International Conference on System Sciences*, pages 4–7, 2000.
- [8] Y. Yu, D. Estrin, and R. Govindan. Geographical and energy-aware routing: A recursive data dissemination protocol for wireless sensor networks. Technical report, Computer Science Department, UCLA.
- [9] R. Wolski, J. S. Plank, J. Brevik, and T. Bryan. Analyzing market-based resource allocation strategies for the computational grid. *International Journal of High Performance Computing Applications*, 15(3):258–281, 2001.
- [10] S. H. Clearwater, editor. *Market-based control: A paradigm for distributed resource allocation*. World Scientific Publishing Co., Inc., NJ, USA, 1996.
- [11] 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*, 35(3):349–359, 2005.
- [12] V. Lesser, C. Ortiz, and M. Tambe, editors. *Distributed sensor networks: A multiagent perspective*. Kluwer Publishing, 2003.
- [13] J. Manyika and H. Durrant-Whyte. *Data fusion and sensor management: A decentralized information-theoretic approach*. Ellis Horwood, 1994.
- [14] R. K. Dash, A. Rogers, S. Reece, S. Roberts, and N. R. Jennings. Constrained bandwidth allocation in multi-sensor information fusion: A mechanism design approach. In *Proceedings of The Eighth International Conference on Information Fusion*, Philadelphia, USA, 2005.