2013

A Mechanism Design Approach to Bandwidth Allocation in Tactical Data Networks

Ankur Mour

Purdue University

PURDUE UNIVERSITY
GRADUATE SCHOOL
Thesis/Dissertation Acceptance

This is to certify that the thesis/dissertation prepared

By Ankur Mour

Entitled
A Mechanism Design Approach to Bandwidth Allocation in Tactical Data Networks

For the degree of   Master of Science in Aeronautics and Astronautics

Is approved by the final examining committee:

Daniel A. DeLaurentis
　　　　　　　Chair

Jitesh H. Panchal

Seokcheon Lee

Karen Marais

To the best of my knowledge and as understood by the student in the *Research Integrity and Copyright Disclaimer (Graduate School Form 20)*, this thesis/dissertation adheres to the provisions of Purdue University's "Policy on Integrity in Research" and the use of copyrighted material.

Approved by Major Professor(s): Daniel A. DeLaurentis

Approved by: Weinong Wayne Chen   11/26/2013
　　　　　　　　Head of the Graduate Program   Date

A MECHANISM DESIGN APPROACH TO BANDWIDTH ALLOCATION IN TACTICAL DATA NETWORKS

A Thesis

Submitted to the Faculty

of

Purdue University

by

Ankur Mour

In Partial Fulfillment of the

Requirements for the Degree

of

Master of Science in Aeronautics and Astronautics

December 2013

Purdue University

West Lafayette, Indiana

Dedicated to my parents.

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Mour, Ankur M.S.A.A., Purdue University, December 2013. A Mechanism Design Approach to Bandwidth Allocation in Tactical Data Networks. Major Professor: Daniel A. DeLaurentis.

The defense sector is undergoing a phase of rapid technological advancement, in the pursuit of its goal of information superiority. This goal depends on a large network of complex interconnected systems – sensors, weapons, soldiers – linked through a maze of heterogeneous networks. The sheer scale and size of these networks prompt behaviors that go beyond conglomerations of systems or 'system-of-systems'. The lack of a central locus and disjointed, competing interests among large clusters of systems makes this characteristic of an Ultra Large Scale (ULS) system. These traits of ULS systems challenge and undermine the fundamental assumptions of today's software and system engineering approaches. In the absence of a centralized controller it is likely that system users may behave opportunistically to meet their local mission requirements, rather than the objectives of the system as a whole. In these settings, methods and tools based on economics and game theory (like Mechanism Design) are likely to play an important role in achieving globally optimal behavior, when the participants behave selfishly. Against this background, this thesis explores the potential of using computational mechanisms to govern the behavior of ultra-large-scale systems and achieve an optimal allocation of constrained computational resources

Our research focusses on improving the quality and accuracy of the common operating picture through the efficient allocation of bandwidth in tactical data networks among self-interested actors, who may resort to strategic behavior dictated by self-interest. This research problem presents the kind of challenges we anticipate when we have to deal with ULS systems and, by addressing this problem, we hope to develop a methodology which will be applicable for ULS system of the future. We build upon the previous works which investigate the application of

auction-based mechanism design to dynamic, performance-critical and resource-constrained systems of interest to the defense community.

In this thesis, we consider a scenario where a number of military platforms have been tasked with the goal of detecting and tracking targets. The sensors onboard a military platform have a partial and inaccurate view of the operating picture and need to make use of data transmitted from neighboring sensors in order to improve the accuracy of their own measurements. The communication takes place over tactical data networks with scarce bandwidth. The problem is compounded by the possibility that the local goals of military platforms might not be aligned with the global system goal. Such a scenario might occur in multi-flag, multi-platform military exercises, where the military commanders of each platform are more concerned with the well-being of their own platform over others. Therefore there is a need to design a mechanism that efficiently allocates the flow of data within the network to ensure that the resulting global performance maximizes the information gain of the entire system, despite the self-interested actions of the individual actors.

We propose a two-stage mechanism based on modified strictly-proper scoring rules, with unknown costs, whereby multiple sensor platforms can provide estimates of limited precisions and the center does not have to rely on knowledge of the actual outcome when calculating payments. In particular, our work emphasizes the importance of applying robust optimization techniques to deal with the uncertainty in the operating environment. We apply our robust optimization – based scoring rules algorithm to an agent-based model framework of the combat tactical data network, and analyze the results obtained.

Through the work we hope to demonstrate how mechanism design, perched at the intersection of game theory and microeconomics, is aptly suited to address one set of challenges of the ULS system paradigm – challenges not amenable to traditional system engineering approaches.

CHAPTER 1.   INTRODUCTION

1.1   <u>Motivation</u>

The current advances in technology have made real time information about the state of the world increasingly available. In the defense sector, this has translated into a race for information domination – to develop superior techniques to collect, fuse, analyze and exploit information to meet mission requirements. This goal depends on complex interconnected web of systems – thousands of platforms, sensors, decision makers, weapons and soldiers connected through a maze of heterogeneous networks. These systems will challenge our imaginations and push the boundaries of the very concept of today's systems and systems of systems. They will be ultra-large scale systems.

> "Our soldiers depend on software and will depend more on software in the future.
>
> The Army's success depends on software and the software industry.
>
> We need better tools to meet future challenges, and neither industry nor
>
> government is working on how to do things light-years faster and cheaper.
>
> How can future systems, which are likely to be a billion lines of code,
>
> be built reliably if we can't even get today's systems right?"
>
> — Asst. Sec Army Claude Bolton(2005)

Independent parallel research was conducted in the UK and USA to elicit the characteristics and the challenges involved in developing, maintaining and managing such large-scale software-intensive complex systems. In the UK, this effort was led by a consortium of British academics and industrial practitioners with a focus on the science and engineering of these so called "Large-Scale Complex Information Technology Systems (LSCITS)." Around the same time, a team led by Linda Northrop at the Software Engineering Institute (SEI) at Carnegie Mellon University

published a report on the challenges associated in engineering these complex socio-technical ecosystems or "Ultra Large Scale Systems" (Northrop et al., 2006).

The Software Engineering Institute (SEI) was funded to conduct a 12 month long investigation of Ultra Large Scale (ULS) systems software by the office of Assistant Secretary of the U. S. Army (Acquisition, Logistics, & Technology). SEI was tasked to come up with a proposed agenda to fund, coordinate, and conduct research for ULS systems. The ultimate goal was to the creation of a collaborative research network that would carry out the work towards solving the ULS system problem for the U. S. Department of Defense.

The primary characteristic of ULS systems is ultra-large size on any imaginable dimension:

- Lines of code
- Amount of data stored, accessed, manipulated, and refined
- Number of connections and interdependencies
- Number of hardware elements
- Number of computational elements
- Number of system purposes and user perception of these purposes
- Number of routine processes, interactions, and "emergent behaviors"
- Number of (overlapping) policy domains and enforceable mechanisms
- Number of people involved in some way

These characteristics are not without precedent and can be increasingly seen in today's systems of systems – the crucial difference is that these traits will dominate in ULS systems. The sheer scale of ULS systems will change everything. Two dimensions of scale can demonstrate this point:

- The ULS systems will be developed by and serve a growing number of human users. The human participants in the system at any time may have disjointed and competing interests. Without an incentive to provide truthful information, users could be prone to distort and misreport their private information, if it is in their interest to do so, even if this deception comes at the expense of the whole system.

- ULS systems will be necessarily decentralized in a number of ways – decentralized data, development, evolution and operational control. In these settings, it is impractical to assume that an omnipresent decision maker can be constructed that knows enough about each of the users, parts and tasks, to impose an effective solution. For the sake of analogy, one can imagine the economic distortions like supply, price, forecasting etc. that are induced by central economies. As the economics diversify, there distortions can be expected to become more severe.

1.2   Perspectives on ULS Systems

One way to understand and gauge the sheer difference in scale between traditional systems and ULS systems is to think in terms of infrastructure – cities as opposed to individual buildings. The task of building the large systems of today is akin to constructing a large monolithic building. In stark contrast, ULS systems will operate at levels of complexity comparable to cities. At first glance it might seem that a city is nothing but a collection of a large number of buildings – this is not true. As Howard Rheingold said "Cities are places of massive information flows, networks, and conduits, and myriad transitory information exchanges" (Rheingold, 1993). Cities are not conceived or built by individual organizations, nor is it specified in advance by listing requirements; but rather a city emerges and evolves with time through the lightly regulated and coordinated actions of many individuals acting locally over time.

The factors contributing to a city's success include extensive infrastructures apart from individual buildings, along with myriad mechanisms that regulate local actions in a bid to maintain coherence in the absence of central control. These mechanisms include but are not limited to, communication services, government organizations, transportation systems and emergency services, distribution of consumer goods and utilities. Cities thrive on necessities – economic and cultural. The essence of a city emerges from global mechanisms and protocols designed to stimulate growth and evolution.

In a similar vein, there is a need to shift perspective when we characterize ULS systems. Apart from the analogy of buildings and cities, another way to understand this desired change in

outlook is in terms of ecosystems—socio-technical ecosystems to be particular. Similar to a biological ecosystem, a ULS system will comprise of a capricious community of organizations, computing devices, and people, all interdependent and competing, with complex networked dependencies and intrinsic adaptive behavior in an evolving environment.

Socio-technical ecosystems include people, organizations, and technologies at all levels with significant and often competing interdependencies:

- Competition for resources.

- Organizations and stakeholders responsible for setting policies (operational, acquisition and production policies) in a bid to encourage effective use of these scarce resources to achieve system objectives.

- Services may vary based on how different automated subsystems and leaders choose to distribute available resources to missions, subject to various levels of importance. This distribution of resources needs to be optimized, by imposing appropriate incentives and rules.

- A need to develop local and global indicators of Quality of Service (QoS) to determine if the incentives are working as intended. These QoS indicators may trigger necessary changes in policies and in element and system behavior.

1.3   Systems-of-Systems vs. Ultra Large Scale Systems

Some of the characteristics of ULS systems will be in common with today's systems of systems (SoSs). Maier (1998) developed a list of characteristics that distinguish large monolithic systems from systems of systems:

- **Operational independence of elements:** Constituent systems are useful in their own right and generally operate independent of other systems.

- **Managerial independence of elements:** Component systems are acquired and managed independently; they maintain their existence independent of the SoS.

- **Evolutionary development:** An SoS is never completely, finally formed but constantly changes and comes into existence gradually.

- **Emergent behavior:** Properties appear in an SoS that are not apparent (or predictable) from the constituent systems.

- **Geographic distribution:** Components are distributed geographically such that their interactions are limited to information exchange rather than mass or energy exchange.

DeLaurentis (2005) identified three traits with important implications for modeling SoS, above and beyond those provided by Maier.

- **Networks:** Networks define the connectivity between independent systems in the SoS through rules of interaction.

- **Heterogeneity**: Constituent systems are of significantly different nature, with different elementary dynamics that operate on different time scales.

- **Trans-domain:** Effective study of SoS requires unifying knowledge across fields of study: engineering, economics, policy, operations etc.

Maier expounds upon different classes of systems of systems based on their operational and managerial independence. Maier cites the Web and national economies as examples of virtual SoSs. What he defines as virtual SoS comes closest to our understanding of ULS systems:

"Virtual systems lack a central management authority. Indeed, they lack a centrally agreed upon purpose for the system-of-systems. Large scale behavior emerges, and may be desirable, but the super-system must rely upon relatively invisible mechanisms to maintain it."

—Maier (1998)

The characteristics of SoS as outlined by Maier and DeLaurentis provide a good starting point to classify systems but are not so useful when it comes to understanding the underlying technical problems that characterize such systems. The sheer scale of ULS system is what endows them with most of their traits, and undermines today's assumptions. These traits are as follows:

Table 1-1 Characteristics of ULS systems and assumptions undermined

Trait	Description	Today's assumptions undermined
Decentralization	Decentralized in a variety of ways—decentralized data, development, evolution, and operational control.	- All conflicts must be resolved centrally and uniformly.
Inherently conflicting, unknowable, and diverse requirements	Developed and used by a wide variety of stakeholders with different, conflicting, complex, and changing needs.	- Requirements can be known in advance and change slowly. - Tradeoff decisions will be stable.
Continuous evolution and deployment	System will evolve not in phases, but continuously - new capabilities will be integrated and unused capabilities dropped, while system is operating.	- System improvements are introduced at discrete intervals.
Heterogeneous, inconsistent, and changing elements	System doesn't contain uniform homogenous parts – there are inconsistent misfits added continuously as system is extended and repaired.	- Effect of a change can be predicted sufficiently well. - Configuration information is accurate and controllable. - Components and users are fairly homogeneous.
Erosion of the people/system boundary	People will not just be users of a ULS system; they will be elements of the system, affecting its overall emergent behavior.	- Collective behavior of people is not of interest. - People are just users - Social interactions are not relevant.
Normal failures	Software and hardware failures will be norm rather than exception.	- Failures will occur infrequently. - Defects can be removed.
New paradigms for acquisition and policy	The acquisition of a ULS system will be simultaneous with operation and require new methods for control.	- A prime contractor is responsible for system development, operation, and evolution.

These highlighted characteristics are not all independent and most appear in even today's systems and systems of systems; but in ULS systems, they dominate. Each of these characteristics derives from the consequences of scale changing everything and undermines the assumptions that underlie today's software developers and acquirers. Each plays a hand in the complexity of designing ULS systems, validating their behavior, and evolving their capabilities.

1.4   Addressing ULS Systems Challenges

As mentioned earlier, the characteristics of ULS systems challenge and undermine the fundamental assumptions of today's software engineering approaches. Current approaches to defining, developing, deploying, operating, acquiring, and evolving software-intensive systems are based on perspectives that are fundamentally ill-equipped to handle the characteristics arising from ultra-large scale. For the last few decades, engineering has been the all-pervasive metaphor for creating software systems. But in ULS systems, our concern is not just limited to software; indeed we are now dealing with a veritable ecosystem of people, organizations, governance, social interaction, hardware, and software. A purely engineering requirement driven perspective will no longer suffice as the dominant metaphor. The scale, decentralization, distribution, and heterogeneity of ULS systems will present challenges in effective design, evolution, management, control, monitoring and assessment of ULS systems. In order to address these challenges, we need to shift our perspective on how we characterize problems arising from scale. We need to develop empirical methods to bring about coherence in the context of complexity and scale. An exciting place to start is at the intersection of traditional software engineering and other disciplines which deal with human factors, such as, anthropology, biology, and microeconomics.

"The older is not always a reliable model for the newer, the smaller for the larger,
or the simpler for the more complex…Making something greater than any existing thing
necessarily involves going beyond experience."

—Petroski (2005)

Pushing the Limits: New Adventures in Engineering

We discussed how in most cases ULS systems will lack a central locus of institutional control and the system participants at any time will have disjointed and competing interests. In such settings it is highly likely that the system users may behave opportunistically to meet their local mission requirements, rather than the goals and objectives of the system as a whole. Hence there is a need to provide a basis for satisfying system-wide quality goals and simultaneously also satisfy individual goals and expectations of the various stakeholders. In such cases, methods and tools based on economics and game theory will play an important role in achieving globally optimal behavior, even when the participants behave selfishly.

1.4.1   Mechanism Design

The field of *mechanism design* lies at the intersection of economics and game theory and is concerned with designing mechanisms, protocols, and institutions that are mathematically proven to satisfy certain system-wide objectives under the assumption that individuals interacting through such institutions act in a self-interested manner and may hold private information that is relevant to a required decision. Mechanism design doesn't take cooperation among agents as granted. Instead it induces cooperation as an *emergent* property of agents engaged in selfish, competitive economic behavior. This mirrors the metaphor of the invisible hand principle as conceived by Adam Smith which asserts that an individual's self-interest is ultimately economically beneficial to society as a whole (Smith, 1863). *Computational mechanism design* puts mechanism design into a computational setting and includes both the use of computers to design mechanisms and the use of mechanisms to control computing.

Mechanism design has its roots in microeconomics, where it is known as *institution design* and in game theory, where it is sometimes known as *implementation theory*. The research literature provides plenty of examples of mechanisms being put to practical use to achieve large-scale social objectives (Federico & Rahman, 2003; Hinz, 2003;Gerkey & Mataric, 2002).

Auction – based computational mechanism design has recently begun to see its application in a host of areas. One well-documented use of computational mechanism falls under the umbrella of e-commerce, where it is used for allocating computational resources. Google uses an explicitly designed auction mechanism for allocating advertising space on Web pages, returned

from keyword searches, in order to generate the main chunk of their revenue (Edelman & Ostrovsky, 2007). Supply chain optimization is another subfield in e-commerce which uses computational mechanism design (R. R. Chen, Roundy, Zhang, & Janakiraman, 2005; Sandholm, 1996).

Mechanism design finds application in problems involving the allocation of scarce resources like network bandwidth, storage capacity and power among both human and computational entities that are inclined to resort to strategic behavior dictated by guile and self-interest. McMillan provides a discussion on the consequences of designing a defective mechanism, as in the case of the New Zealand radio spectrum auction, and the importance of getting the details of mechanism design right, like in the case of the U.S. public radio spectrum auction (McMillan, 1994). As systems scale up in dimensions, interaction protocols resistant to strategic manipulation are needed that are capable of efficiently aggregating information from different parts of a system to facilitate global decision making.

Our investigation focuses on the potential of using computational mechanisms to govern the behavior of ultra-large-scale systems and achieve an optimal allocation of computational resources. In our scenario computational systems can be viewed as pseudo-economies, with computational entities competing with each other for the use of scarce computational resources to satisfy their local mission goals. We investigate the application of mechanism design to dynamic, performance-critical and resource-constrained systems of interest to the defense community.

1.5 Bandwidth Allocation in Tactical Data Networks

ULS systems will have to support warfighters at all echelons who are engaged in information-extensive activities and who must share constrained critical resources. Military group operations need all the air, ground and sea platforms participating in an operation, to work like a cohesive force. There is an obligation for military systems to be interoperable with other systems - military or civilian. We define interoperability as a system's ability to provide and accept services from other systems and use these exchanged services to operate efficiently and effectively. The

platforms in a group must share and exchange tactical data from their onboard sensors in order to establish and maintain a common operating picture of the tactical situation as shown in Figure 1-1.
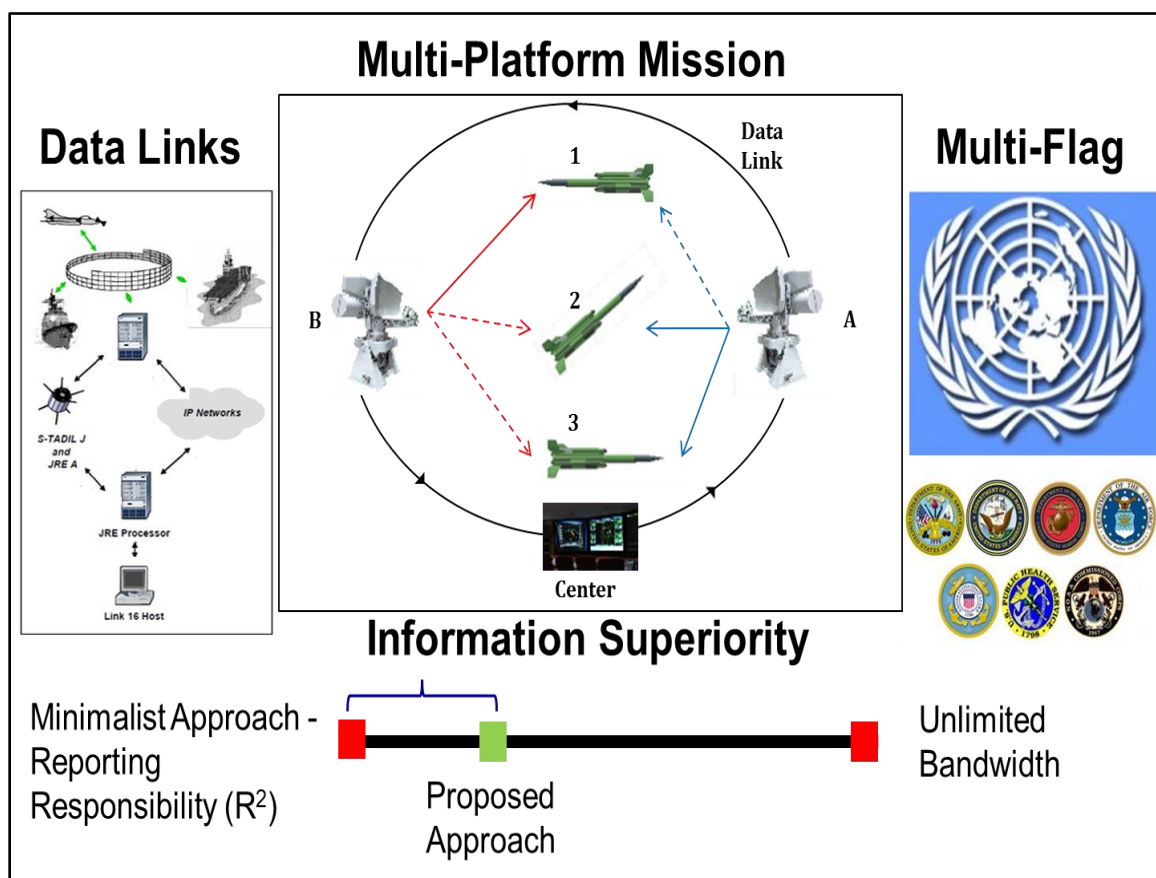
Figure 1-1 Bandwidth Allocation in Tactical Data Networks

The exchange of tactical data among the military platforms is facilitated over a standardized radio network, known as a TActical Data Information Link (TADIL). TADILSs are used to transmit measurements pertaining to both the platform themselves and the targets. As per DISA (Defense Information Systems Agency) guidance the term TADIL has been officially replaced by the generic term Tactical Data Link (TDL). However, we shall continue to refer to them as TADILs given the legacy holdover. These tactical data information links are characterized by their standard message formats and transmission formats.

The sensors onboard military platforms have an incomplete and inaccurate view of the operating picture. This is partly because each sensor can only detect and track targets within a limited region of observation immediately surrounding itself. More importantly, the navigational systems onboard the military platforms provide the sensors with an imprecise estimate of its own location. The sensors estimate the position of the targets within its region of observation by making noisy measurements of the bearing and range of the target from itself. Thus the sensors need to use data transmitted from neighboring sensors, in order to improve the accuracy of their measurements. The fusion of the sensor's own noisy information with the observations from a number of other sensors reduces the overall uncertainty in the measurements. However, the bandwidth in these tactical data networks is a scarce resource and the mission outcome can be significantly affected by decisions made in real time about which data to share. Ad-hoc Bandwidth allocation can have serious repercussions and can even jeopardize a mission.

> "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 to allow the highest-priority messages to get through. This can delay, or cancel other messages or data transmissions, which are placed into in a lower priority."
>
> —Wilson (2004)
> Network Centric Warfare: Background and Oversight Issues for Congress

The track data exchanged among the military platforms over the TADIL encapsulates the processed radar data which represents artifacts such as airplanes, helicopters, missiles, ships, boats, submarines along with various other kinds of land, sea and air based vehicles. The shared tactical data is then used by each platform to create a Common Operational Picture (COP), the accuracy of which depends on a number of factors:

- **Gridlocking:** Gridlock is a technique whereby remote tracks received from a designated reference unit are compared to local observed data to determine any

data registration corrections. This entails eliminating or minimizing sensor alignment errors, navigational position errors, sensor biases among others. If gridlocking is not implemented to provide correlation between local and remote tracks, there is the possibility of remote tracks being painted multiple times and overlapping each other.

- **Track Correlation:** Track correlation is the process of minimizing or eliminating the display of multiple tracks that represent the same artifact. Track correlation is a fundamental problem in distributed multi-target multi-sensor tracking system. It involves the selection of the most probable association between target tracks from a very large set of possibilities.

- **Reporting Responsibility (R^2) Rules:** The data link R^2 rules permit only the unit with the best quality data (position, heading, velocity, etc.) to report a surveillance track on the link. This strategy prevents multiple track reports on the link for a single object, thus minimizing the data latency. The platform selected to provide the report for a track is said to have R^2 for that track.

The Reporting Responsibility (R^2) rule is a minimal precedence based mechanism. It precludes any possibility of collaboration in building a common operating picture by disallowing the redundant reporting of a single object. Although R^2 rules minimize data latency it also reduces the diversity of the source data. Redundant reporting of objects can play a critical role in resolving ambiguities. In the context of network bandwidth, the R^2 approach can be regarded as an extreme minimalist approach.

In our research we consider the minimalist approach characterized by TADILs, such as LINK-16, as our point of departure. We start with 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 network cycle time.

In military settings, there are scenarios with multi-nation, multi-platform coalition exercises, where the military commanders of each platform may be concerned with the well-being of their own platform over others. Self-interested behavior can be a concern, even when operating

under a single flag, as platforms may overstate their need for scarce resources, like bandwidth. This need not be out of some malevolence; the platforms may be trying to do what they think is best, but may unknowingly use resources that would be more beneficial for other platforms. For example, soldiers have been known to overstate the precedence of messages with the intention of speeding up their delivery, but in doing so, they unintentionally saturate the network, reducing its effectiveness not only for them but also for others. When we allow the possibility of platforms to exhibit deceptive behavior to further their self-interests, the problem of bandwidth allocation becomes particularly complex and traditional system approaches of resource allocation are ill-equipped to address this problem. This is one of the exact challenges that we anticipate when we are dealing with Ultra Large Scale Systems, which will lack a central locus of operational or institutional control. We use the case-study of a coalition military exercise with constrained bandwidth resources and self-interested actors, to address the challenge of designing a mechanism for optimal resource allocation.

In our case study, the platform-sensors are individually owned by different stakeholders who may have conflicting goals and resort to strategic behavior marked by a combination of guile and self-interest, if it furthers their local goals. The sensors may end up operating in competitive rather than cooperative environments, and as such, may attempt to optimize their own gain from the network at a cost to the overall performance of the entire network. This is particularly true in networks where the bandwidth available for transmission of data among the sensors is limited. Individual platforms benefit from receiving data from other platforms but have no incentive for sharing it. 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.

Thus, against this background we seek to design a mechanism which can efficiently allocate a finite bandwidth, beyond what is used in a conventional R^2 approach, to enhance the common operating picture. We model the platform sensors as rational agents seeking to optimize their own utility and liable to deceptive behavior to further their objectives. Our goal is to design a mechanism that efficiently allocates the flow of data within the network to ensure that the resulting global performance maximizes the information gain of the entire system, despite the selfish actions of the individual sensors.

1.6   Research Problem

The research problem is recapped and encapsulated as follows:

We consider a scenario where a number of military platforms have been tasked with the goal of detecting and tracking targets. The sensors onboard a military platform have a partial and inaccurate operating picture and need to make use of data transmitted from neighboring sensors in order to improve the accuracy of their own measurements. The communication takes place over tactical data networks like Link 16, where the bandwidth is a scarce resource. The problem is compounded by the possibility that the military platforms may have conflicting goals and exhibit deceptive behavior to satisfy their local mission objectives, at the expense of the global system goal. Therefore there is a need to design a mechanism to optimally allocate the bandwidth in order to improve the quality of the common operating picture, despite the selfish actions of the individual platforms.

The heart of the mechanism, which we need to design, resembles a portfolio optimization. The portfolio problem assumes that a portfolio needs to be constructed consisting of a set of stocks. Each of the stocks has a return and a risk value associated with it and the objective is to determine what fraction of wealth must be invested in each stock to maximize the portfolio value. In reference to our problem scenario, the stocks represent the observations made by the sensor agents. The return value of the stock can be regarded as the information content of each observation, while the total wealth available represents the bandwidth to be allocated on the tactical data link. Thus the objective is to determine which track information to select for transmission given the fixed additional bandwidth available to maximize the total information gain. The challenges of interdependent valuations, selfish behavior, constrained resources and dynamic uncertain environments dictate that our mechanism needs to go beyond a simple portfolio optimization as highlighted in Figure 1-2.
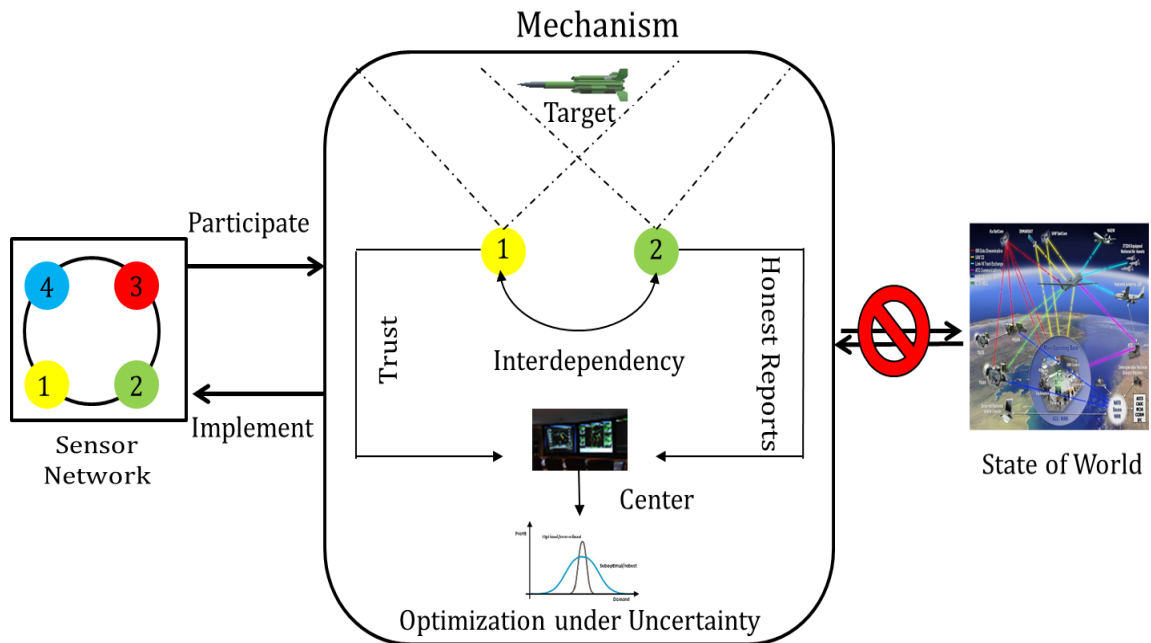
Figure 1-2 Research Framework

- Participate – Since the sensor platforms are individually owned by different stakeholders, the mechanism needs to ensure that the platforms voluntarily participate in lieu of some incentive of participation.

- Honest Reports – The sensor platforms may resort to strategic behavior to optimize their own gain from the network, at a cost to the overall performance of the entire network. In our problem domain, individual platforms benefit from receiving data from other platforms but have no incentive for sharing it. 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. The mechanism has to incentivize the platforms to truthfully reveal their track information. Without true input values, we can't construct an optimal portfolio solution.

- Interdependency - In tactical sensor networks, the observations made by the sensors are polluted by uncertainty and noise. Indeed, other sensor agents may possess information that would, if known to a particular sensor platform, affect the value it

assigns to the target. The mechanism must account for this information interdependency in the reported observations.

- State of World – The mechanism should work e ven when the center has no access to the true state of the world. The dynamic and uncertain nature of the operating environment means that the track data evolves between the time the information is reported and the time when the observation can be observed. Thus the center needs to evaluate the received reports without any knowledge of the true outcome.

- Optimization under uncertainty - The mechanism need to take into the possibility that given the dynamic operating environment, there might be some uncertainty in the reported data. The optimal portfolio solution may be completely meaningless if it is not robust to data uncertainty.

- Implementation - The mechanism has to ensure that once a target has been allocated to a platform, it invests all its resource to track the particular target. The optimal portfolio solution will fail if the participating platforms disregard their responsibility.

Thus, our research objective is to develop a mechanism design – based approach which satisfies the elicited requirements and can efficiently allocate a finite bandwidth, beyond what is used in the R^2 approach, to enhance the quality of the operating picture

1.6.1   Research Contributions

The research work described here builds upon the previous work (discussed in Chapter 2) but adds significantly to the complexity and fidelity of the problem settings. In particular, our work emphasizes the importance of applying robust optimization techniques when designing computational mechanisms. Against this background, our work makes the following contributions:

1. We demonstrate how mechanism design, perched at the intersection of game theory and microeconomics, is aptly suited to address one set of the challenges of the ULS

system paradigm – challenges not amenable to traditional system engineering approaches.

2. We identify the flaws and the shortcoming of auction-based mechanism design to address the problem of bandwidth allocation under the settings and show the viability of approaching the problem with an alternative approach within the Mechanism Design research domain, in the form of strictly proper scoring rules.

3. We develop a standalone application framework by utilizing Purdue's Agent Based Modeling tool DAF (Discrete Agent Framework) that provides a unique insight into the application of computational mechanism design in decision making.

1.7   Thesis Structure

The remainder of this thesis progresses as follows:

- In Chapter 2 we highlight and discuss the relevant literature to address trust in Multi-Agent Systems (MASs). We examine different prevalent approaches to trust and provide an introduction to the game theoretic approach to Mechanism Design. We then discuss the notion of trust within the mechanism design literature and focus on auction based models for truth elicitation.

- In Chapter 3, we provide a description of the Multi-Agent System framework we employ to address the research problem. We introduce the agent-based modeling tool Discrete Agent Framework (DAF) and describe the agent-based model created in DAF to emulate the combat tactical data network and construct the common operating picture.

- In Chapter 4, we introduce the concept of scoring rules and how scoring rules can be used as a methodology to address the shortcomings of auction-based mechanism design. We describe and analyze our scoring rules – based two-stage mechanism with unknown costs whereby multiple sensor platforms can provide estimates of limited precisions and

the center does not have to rely on knowledge of the actual outcome when calculating payments.

- In Chapter 5, we present the concept of robust optimization and provide the Bertsimas-Sim formulation of the robust portfolio optimization problem that we adopt for our research problem.

- In Chapter 6, we investigate the application of our modified strictly proper scoring rules based mechanism to the agent-based model of the tactical data network. The effects of the lack of access to true outcome, deceptive behavior on the part of the agents and the protection level of the robust optimization are evaluated and studied.

- In Chapter 7, we summarize the key findings from this work and discuss the potential avenues of future research to extend the scope of the current work.

CHAPTER 2.   LITERATURE REVIEW

Trust is a fundamental concern in ultra-large-scale systems. Trust is the willingness of a party to be vulnerable to the actions of another party based on the expectation that the other will perform a particular action important to the trustor, irrespective of the ability to monitor or control the other party (Mayer, Davis, & Schoorman, 1995). In ULS systems it is necessary to ensure that trust exists at the heart of all interactions between the system participants. Trust in the mechanism motivates agents to honestly reveal their private information to other agents in the system, which is instrumental to achieving the global objective. In this chapter, we discuss how trust can be addressed in Multi-agent systems. We examine two different prevalent approaches to trust - at the individual level and at the system level. We provide an introduction to the game theoretic approach to Mechanism Design and present the important possibility and impossibility results in literature. We then discuss the notion of trust within the mechanism design literature and focus on auction based models for truth elicitation.

2.1   Trust

The proliferation of narrow intra-disciplinary definitions of trust and the multiple interpretations of trust in everyday life, have led to literature confusion regarding the meaning of trust (McKnight & Chervany, 2002). Bigley & Pearce (1998) chronicled the different uses of the word trust showing both how various definitions are similar and how they diverge. McKnight & Chervany (2002) have defined inter-personal disposition to trust constructs from psychology and economics as well as institution-based trust constructs from sociology.

- *Trust-related behavior* indicates that a person voluntarily depends on another person with a feeling of relative security, even though negative consequences are possible (Lewis & Weigert, 1985). In the context of Ultra Large Scale systems, we use the

information-sharing construct of trust-related behavior. *Information sharing* represents trust-related behavior because it makes one vulnerable to the actions of the trustee with respect to the information (Mishra, 1996).

- However, the object of trust may involve situations and structures instead of people. *Institution – based trust* means one believes, with feelings of relative security, that favorable conditions are in place that are conducive to situational success in a risky endeavor or aspect of one's life (Shapiro, 1987). We adopt the *structural assurance* sub-construct of institution – based trust. *Structural Assurance* means one securely believes that protective structure - guarantees, contracts, regulations, promises, legal recourse, processes, or procedures - are in place that are conducive to situational success.

ULS systems will, in many cases, lack a central locus of operational or institutional control and the participants in the system, at any time, will have disjointed and competing interests. In such settings, it is highly plausible that the system users may behave opportunistically to meet their own local mission requirements, irrespective of the goals and objectives of other participants or the system as a whole. Thus, when devising a mechanism or protocol for such systems, one needs to ensure that the actions undertaken by the agents along with the allocation of resources, result in both system and individual level goals being satisfied. The construction of such a mechanism becomes difficult as participating agents are free to choose what information to convey to other agents, which agent to interact with, and when to interact with the environment and other agents. Maximization of an agent's individual utility is dependent on receiving perfect information of the environment, including information on other agents. However this desired state of perfect information is nearly impossible to achieve within practical contexts, due to limiting computational capabilities, storage capacities, and network bandwidth. In such uncertain environments, it becomes necessary for agents to have to be able to trust each other in their interactions within MAS.

Trust models can be broadly divided into two different levels:

- Individual level trust models – Agents are endowed with reasoning ability to form opinions on the honesty and reliability of their counterparts.

- System level trust models – Designing protocols and mechanisms whereby agents are forced to be truthful by the rules of engagement.

The two approaches are not necessarily disparate and may, in fact, be viewed as complementary. Individual level trust models cannot possibly reconcile with the inherent uncertainty in the MAS and must rely on system-level trust models to reduce this uncertainty. On the other hand, in system-level trust models, there is an intrinsic tradeoff between trust and efficiency and thus the decision making process on the part of the agent can be guided using individual level trust models.

2.2   Individual – level Trust Models

Individual level trust models aim to guide the decision making process of an agent in deciding on how, when, and whom to interact with. In these models, an agent has two alternatives when attempting to choose which agent to trust. First, it has the option to directly interact with the other agents in order to draw their inferences after several encounters. Direct interaction leads us to consider methods by which agents can learn or evolve better strategies to deal with honest and dishonest agents, such that payoffs are maximized in the long run. The second alternative is to interact with other entities indirectly by referring to third-party information in order to make a decision. This invariably requires agents to develop methods to reliably acquire and reason about the information gathered from other agents. In addition, the elements of both these approaches can be combined in order to introduce more theoretical foundations based on a probability theory framework. We can also take a higher level view of trust in the form of socio-cognitive models which involves taking the knowledge of motivations of other agents for granted and proposes ways to reason about these motivations.

2.2.1   Learning Models

In learning models, we consider trust as an emergent property that comes from direct interactions among agents. We assume that agents interact with one another multiple times rather than through one - shot interaction. It is further assumed that agents have an incentive to defect (Dasgupta, 2000). Trust is achieved on the basis of consistency in performing according to agreed - upon exchanges. Although defection could result in higher payoffs for the defecting party and loss of utility to the other party, it reduces the possibility of future interactions since the losing agent would typically attempt to avoid risking future utility losses.

An often cited example that features defection is the Axelrod's tournament revolving around the Prisoners' Dilemma (Axelrod, 1980). The Prisoner's Dilemma is a game involving two prisoners, in two different rooms, having to separately choose between two moves, either "cooperate" or "defect". We can summarize the game situation through Table 2-1.

Table 2-1 Prisoners' Dilemma

	Prisoner B cooperates	**Prisoner B defects**
Prisoner A cooperates	Each serves 5 years	Prisoner A: 10 years Prisoner B: goes free
Prisoner A defects	Prisoner B: 10 years Prisoner A: goes free	Each serves 5 years

The tournament is of *N* step prisoners' dilemma in which participants have to choose their mutual strategy again and again, and retain memory of their previous encounters. Within controlled settings, the most successful deterministic strategy is 'tit-for-tat' which simply cooperates on the first iteration of the game and after that, the player does what his or her opponent did on the previous move. If two agents adopt tit-for-tat they end up with the highest payoffs compared to all other strategies. However, when faced with other selfish strategies, tit-for-tat does not get the maximum payoff, as it loses on the first encounter.

Wu & Sun (2001) showed that if agents are allowed to adapt to their environment, they could use reinforcement learning to elicit trustworthiness. Agents learn from their current actions, and choose their next state to maximize the feedback or utility that they receive from their actions. In the case of multi-agent bidding, they clearly demonstrate that the evolution of strategies facilitate the isolation of malicious agents from good agents who are keen to cooperate. However, while strictly applying to the bidding context, their model does not take into account the fact that there might be some short-term utility loss in cooperating with third parties.

In this context, Sen & Sajja (2002) addressed some of these issues and showed how reciprocity can emerge when the agents learn to predict that they will receive future benefits if they cooperate. They demonstrate that collaborating liars have more to gain after a few iterations in environments with large numbers of philanthropic agents. However this necessitates the presence of large numbers of agent interaction, to build the requisite trust. Mui, Mohtashemi, & Halberstadt (2002) took this notion forward and designed a probabilistic trust model which could identify the minimum number of encounters needed to elicit trust. In the cases where this probabilistic threshold is unreachable, we need to resort to other techniques to establish trust.

Up to this point, all the above learning and evolutionary models of multi-agent strategic interactions assume that agents have complete and perfect information about each and every entity in the MAS. This is clearly a utopian scenario and it is highly unlikely that these strict assumptions will translate to real-life scenarios where agents have a partial and incomplete view of the common operating picture. In real-life scenarios, agents do not have the option to adapt and refine their strategy through multiple encounters or evaluating all the possible actions available. Hence we need to look at other individual-level trust models which are more amenable to real-world scenarios.

2.2.2   Reputation Models

In scenarios where there is a lack of prior information about an agent's incentives and there are limited interactions among agents, reputation models are another way to achieve trust. Reputation can be defined as the opinion or view of someone about something. We distinguish between trust and reputation in the sense that trust is derived from direct interactions while

reputation is acquired from the environment or other agents and ultimately leads to trust. In MAS, reputation is particularly useful when there are a large number of interacting agents. Reputation models are prevalent in online markets and services, stock-trading and auctions which enable agents to refer to third party information to guide their decision making.

The research on reputation models can be broadly divided into two different areas: retrieval of ratings and aggregation of ratings from other agents. Both these approaches borrow the concept of a social network from sociology which draws an analogy between the arrangement of agents in networks and that of humans in societies. An agent can act as a reputation source by validating the result of an interaction between agents, and transmitting the relevant information (like performance ratings) through this network of social relationships, thus giving rise to the concept of reputation.

In this vein, Yu and Singh tackle the problem of retrieving ratings from a social network through the use of referrals (Yu & Singh, 2000), propose a method of representing a social network (Singh, Yu, & Venkatraman, 2001) and then present ways to aggregate the information through the network (Yu & Singh, 2003). In a MAS referral network, agents can act as referral – providers or referral – followers. Aggregation of ratings from other agents is another field of interest within the reputation trust domain. For example, on websites like eBay, ratings of +1 or −1 value along with textual information get aggregated to generate overall rating. Such simplistic accumulation may give rise to problems when some agents either do not return ratings or when they do, manipulate the system by misreporting their information. Yu and Singh's model demonstrates the power of referrals through the Dempster Shater's theory of evidence (Shafer, 1976) in aggregating information obtained from referrals while coping with the lack of information (Yu & Singh, 2002). Schillo, Funk, & Rovatsos (2000) developed a model that demonstrated how agents can cope with lying witnesses in their environment and use witness information to reason effectively against lying agents.

However, all the reputation models we discussed here tend to simplify direct interactions and often fail to contextualize such interactions relative to the witnesses and other interacting agents. All these approaches require knowledge about the result of the interaction after it is

finished, which might not be possible within the highly uncertain environment that the ULS systems operate in. In particular, requiring precise information about the state of the world after an interaction is complete is often not feasible as agents may not have access to the information in order to evaluate the referrals.

2.2.3   Socio Cognitive Models

One of the common shortcomings of the two methods for modeling individual trust that we discussed in the previous sections, is that they rely on the outcomes of interactions to assess trustworthiness of an agent. However, there are other subjective considerations that go into this analysis, which if considered, would provide a more holistic analysis of the traits of the agents in the network (Dasgupta, 2000; Gambetta, 1990). For example, an agent can use the supplies and skill-set available to their counterparts to subjectively assess if it can indeed use these to carry out an agreed task.

Castelfranchi & Falcone (1998, 2000) and Falcone & Castelfranchi (2001) underline the advantages of employing a cognitive view of trust rather than just a narrow quantitative view. They operate within the context of task delegation where an agent X seeks to delegate a task to another agent Y. Agent X needs to gauge the trustworthiness of agent Y by taking into account the different beliefs it has regarding the motivations of agent Y – Competence belief, Willingness belief, Persistence belief and Motivation belief. But their approach is inspired by human beings, who do not always act rationally as opposed to the agents in MAS which are assumed to be rational entities. Braynov & Sandholm (2002) address this issue by employing a rational approach to model an opponent's trust within the context of non-enforceable contracts. They clearly show that in those instances where one agent can precisely estimate the trust of its counterpart, it leads to maximum utility gain and trade between the agents while an incorrect estimation of the same trust, leads to inefficient resource allocation, and hence loss of utility for the two agents.

The socio-cognitive approach is a comparatively newer method to modeling individual level trust, which takes a high-level view of the field. However, this approach is lacking in the rational grounding which learning and reputation models exhibit. If we were to consider these three

modes together rather than in isolation, they could contribute to provide an integrated holistic evaluation of trust at the individual level. However, given the limited computational capacities of the participating agents it might not be feasible gather information from all possible sources in the environment or to take into account all factors contributing to the trust in their counterparts.

2.2.4 Probabilistic Trust Models

One of the common characteristics of the individual trust models - learning and reputation based as well as the socio-cognitive - that we have discussed so far, is that none of them have a sound foundation in statistics. Probabilistic Trust models address this issue by continuing this trend of combining elements from both learning and reputation models, but based on the theoretical foundations of a probability theory framework.

Ismail & Josang (2002) proposed using beta probability density functions to combine feedback from the users on an online system that provides reputation ratings, by having members rate the performance of the other members in the community. Yuan & Sung (2004) proposed a model that can take into account both public reputations as well as the accounts of private interaction. Jones, Janicke, & Cau (2009) proposed a methodology of combining multiple models of individual trust by suggesting the use of 'trust communities'. These communities would facilitate sets of agents combining their personal observations and generating higher quality collective information.

However, even these probabilistic – based trust models inherently assume that agents will always invest all the resources at their disposal in generating and providing feedback regarding their experiences, which violates our requirement of selfish behavior. Secondly, they assume that the operating conditions of the agents would always be conducive for agents to report their observations, which rational and self-interested agents may decide against, to save resources such as bandwidth or power. Thirdly, protocols that use multiple trust models would not be able to continuously monitor the evolving parameters of the dynamic environment that characterizes ULS systems. Finally, even these individual trust models fail to take into account elements of outcome uncertainty, by its reliance on precise knowledge of the common operating picture

once an interaction is over. In conclusion, although the probabilistic trust models can be viewed as a step in the right direction, as they are more robust than the other individual trust models, they are still a long way from addressing our key research requirements.

Given these limitations and shortcomings, it is more prudent to generate protocols and mechanisms that, instead, force the agents to be trustworthy in all their interactions. In this way, these rules of engagement can compensate for the limited scope of individual-level trust models.

2.3   System – level Trust Models

In the context of Ultra-Large Scale systems, it is clear that we need a more systematic approach to induce trust that focuses on the design of protocols and mechanisms to guide the interactions among agents. Such mechanisms can be diverse and include auctions, voting, contract nets, bargaining and market mechanisms, to name some. We need to dictate certain rules of engagement that prevent deception and collusion between participants, and enable an agent to trust one another. These rules endeavor to provide desirable global properties such as investing sufficient resources to complete their allocated task and truthfully reporting their observations. For our purposes, trust and truthful reporting are closely related as we assume that a sensor successfully executes its task when it invests necessary resources in generating an observation and truthfully reports them to other agents.

The rules of engagement and the constraints on interactions can be imposed in different ways. In our research we concern ourselves with two distinct approaches. First, by devising truth eliciting interaction protocols; protocols which guarantee that agents are always better off by being truthful and stand to gain no additional utility by deception. Second by developing reputation mechanisms to foster trust; mechanism that spread an agent's reputation of being truthful or a liar throughout the system. We will take a look at both these approaches in the section below.

2.3.1 Truth-elicitation Protocols

Truth-eliciting interaction protocols have garnered a lot of attention, and a number of different protocols and mechanisms have been devised over the years. These protocols seek to prevent agents from colluding or manipulating the system by imposing constraints dictating the interaction among agents and the information shared during such interactions. An agent stands to gain no additional utility by colluding or lying, if it adheres to the afore-mentioned protocols. The most widely used mechanism in literature is auctions, which introduce trust in the Multi-agent systems from different perspectives. The most common single-sided auctions include the English auction, Dutch auction, First price and Second price auctions. In this section we shall confine our discussion to the first three types of auction and discuss Second Price auctions, after we have introduced the concept of Mechanism Design.

The English auction is perhaps the most widely known kind of auction and has been historically used in famous auction houses like Christie's and Sotheby's (McCabe, Rassenti, & Smith, 1990). The auctioneer starts off the auction by announcing an opening bid, and each bidder is free to raise their bid until no bidder is willing to raise the bid any further, thus ending the auction. The item gets sold to the highest bidder at a price equal to his or her bid. The English auction is technically an open ascending price auction as opposed to a Dutch auction which is an open descending price auction. The Dutch auction is the antithesis of English auctions, whereby the auctioneer begins with a high asking price, and subsequently lowers it until some participant is willing to accept the last price announced by the auctioneer. The item gets sold to the highest bidder at a price equal to the last announced price. The first-priced auction is a sealed bid auction, and hence different from the two open auctions described before, as it involves agents submitting their bids with no knowledge of others' bids. In this type of auction all bidders simultaneously submit sealed bids so that no bidder knows the bid of any other participant. The highest bidder wins the auction and pays the price they submitted.

The Dutch and English auctions induce truthful elicitation on the auctioneer's part since the bids are made publicly and the winner and the winning price are all public information. However, none of the three auctions ensure that bidders reveal their true valuation for the objects on auction. The dominant strategy in Dutch and First-price-sealed-bid auctions is to either reveal or

bid a lower valuation, and in case of the English Auction, the dominant strategy is to bid only a small amount more than the current highest bid till one's true valuation is reached. Also the single-sided auctions are susceptive to bidder collusion, since agents can collaborate to cheat the mechanism by sharing information about their bids. For example in English auctions, the bidders and auctioneer may collude to artificially inflate the ask price thereby forcing the bids to go really high. In a similar vein, bidders may collude to withhold their bids in a Dutch auction until the ask price hits rock bottom.

Cryptographic techniques have been proposed as a security mechanism to prevent bidder collusion in auctions, though it can also result in increased computational costs (Brandt & others, 2002; Brandt, 2001). The auctions also disregard the possibility of multiple encounters among agents. As we discussed in Section 2.2.1 trustworthy behavior can be induced if agents understand that they stand to lose utility in future interactions or prevented from engaging in future interactions. However, open multi-agent systems allow agents to move about and interact with other agents in the system. This allows malicious agents to move from group to group and exploit trustworthy agents as they move around. Hence there is a need to make agents share their ratings of their opponents with the other agents in the system, to prevent malicious agents from exploiting the openness of MASs.

2.3.2   Reputation Mechanisms

In Section 2.2.2 we discussed how reputation models can make agents share their ratings of their counterparts, which can then be gathered and aggregated to be shared with other agents in the system. However the reputation models at the individual trust level fail to account for the selfishness of the participating agents and that the agents will share information only if it is in their best interests. Reputation Mechanisms seek to remedy this shortcoming by modeling agents' reputation at the system level. Reputation mechanisms induce truthful ratings from agents, store and aggregate the ratings, and then publicize these ratings, all at the system level.

Zacharia & Maes (2000) outline the essential traits for good system-level reputation mechanisms and present two reputation systems, SPORAS and HISTOS, for online communities like forums, mailing lists and chat-rooms. In their proposed mechanism SPORAS, the reputation

of the new agents entering the system is set at a minimum value and as these agents receive ratings, their reputation value subsequently increases. HISTOS, the enhancement to SPORAS, also takes into account group dynamics and the reputation of each agent depends on how that agent rates other agents. More importantly, both these reputation mechanisms are robust to collusion. However, neither of these systems penalizes agents deliberately subverting the system by giving false good reports in an attempt to build good reputation since positive ratings are valued more.

Jurca & Faltings (2003a, 2003b) address this shortcoming and induce truthfulness by making side payments to agents when they share feedback, thereby making it rational for them to share reputation information. Agents can sell or even purchase reports to and from the information agents in the system. They seek to make their model robust to lying witnesses by having information agents pay one agent for reports only if they match the next report given by another agent. However, what they gain in robustness is lost through bidder-collusion as all the entities can collude to lie in their reports. Dellarocas (2002) introduced a more realistic feedback mechanism called *Good Will Hunting,* which is specifically tailored for trading environments. The mechanism induces sellers to reveal the true quality of their goods through the threat of biased future reporting of quality of goods to be sold. Buyers are given rebates on future transactions to induce truthful reporting. Their framework oversimplifies the trading model and does not account for buyers colluding to sabotage a seller's reputation.

Faced with these shortcomings, we shift our focus to a more fundamental approach that guarantees incentive compatibility (truthful reporting) and individual rationality (voluntary participation) through certain payment schemes. This is achievable through the application of techniques from the field of mechanism design, which we will expound upon in the next section.

2.4   Mechanism Design

Mechanism Design, a sub-field of microeconomics and game theory, concerns itself with the design of mechanisms, protocols, and institutions that are mathematically proven to satisfy certain system-wide objectives. It assumes that individuals interacting through such institutions

act in a self-interested manner and may hold private information that is relevant to a required decision. Mechanism Design has been used to design protocols for different areas:

- Peer-to-peer systems (M. Chen, Yang, & Liu, 2004)
- Sensor fusion (Dang, Dash, Rogers, & Jennings, 2006)
- Network routing (Holzman & others, 2003)
- Electricity markets (Federico & Rahman, 2003; Hinz, 2003)
- Allocating network capacity (Anderson, Kelly, & Steinberg, 2006; Anshelevich et al., 2004)
- Allocating processor cycles for scientific computing on worldwide grid (M. Chen et al., 2004)
- Allocating tasks for autonomous robots (Gerkey & Mataric, 2002)

This section presents traditional Mechanism Design, which aims to satisfy certain economic criteria (such as efficiently allocating resources, maximizing revenues or having a fair system) given the setting of selfish agents in interactive decision making.

2.4.1 Basic Definitions

We start by introducing the basic concepts of Game Theory, which is intimately linked with mechanism design. Game Theory concerns itself with the study of the strategic interactions in a system of self-interested agents.

Consider a set of N individual agents $I = \{i_1, i_2, \ldots, i_N\}$. Each agent is characterized by its type $\theta_i \in \Theta_i$ from a set of possible types Θ_i which determines the preferences of an agent over different outcomes of a game. Although θ_i represents private information available only to i, we assume, as it is standard, that it is drawn from a commonly known joint distribution. An agent's preferences over outcomes $o \in O$, for a set O of outcomes, can then be expressed in terms of a utility function that is parameterized on the type. These utility functions represent the Von Neumann – Morgenstern utility functions which we discuss in detail in Section 4.1.1. Let $u_i(o, \theta_i)$ denote the utility of agent i for outcome $o \in O$, given type u_i. Agent i prefers outcome o_1 over o_2 when $u_i(o_1, \theta_i) > u_i(o_2, \theta_i)$.

The fundamental concept of agent choice in game theory is expressed as a strategy, which is defined as follows.

Strategy

A strategy is a complete contingent plan, or decision rule, that defines the action an agent will select in every distinguishable state of the world. $s_i(\theta_i)$ denotes the strategy of agent i given type θ_i, where $\sum_i$ is the set of all possible strategies available to an agent. Formally, a strategy for agent i is a mapping from Θ_i to M_i ($s_i : \Theta_i \rightarrow M_i$), where M_i, is the message set for agent i, and includes any messages that the agent communicates.

Strategy Profile

A strategy profile is a vector of the strategies available to an agent i based on its type. $s_i(\theta_i) \in s$ denotes the strategy of agent i given type θ_i, where $\sum_i$ is the set of all possible strategies available to an agent. $s \in M = M_1 \times M_2 \times \ldots \times M_N$, where M is the set of joint messages.

Utility

The utility $u_i (s_1, s_2, \ldots, s_I, \theta_i)$ of agent i given its type θ_i and the strategy profile $s = (s_1, s_2, \ldots, s_I)$ selected by each agent, determines its base preferences over different outcomes in the world. In our work we focus only on quasi-linear utility as it ensures that agents can transfer utility through monetary side payments.

An agent's utility function is known to other agents in the system and represents a preference relation over different pairs of outcomes given the type θ_i. In game theory the fundamental model of agent rationality is based on maximization of expected utility. Given the agent's preferences θ_i over outcomes, and its beliefs about other agents' strategies, it will always select a strategy that maximizes its expected utility.

2.4.2   Solution Concepts

Computing the equilibrium outcome of any game with two or more rational agents depends on the structure and the assumptions about the preferences of the agents and the information available to each agent. Our goal is to design a mechanism so that certain desired system-wide properties (e.g. truthfulness, efficiency, and fairness) emerge as an equilibrium outcome from the interaction among the sensor platforms. Game Theory provides an array of different solution concepts. Each solution concept differs in the assumptions about agent rationality and the knowledge each agent has regarding the preferences of all the other agents in the system. We highlight these concepts below:

Nash Equilibrium

The Nash equilibrium states that in equilibrium, every agent will select a strategy to maximize its own utility given the strategy of every other agent. A strategy profile $s = (s_1, s_2, \ldots, s_I)$ is in Nash equilibrium if every agent maximizes its expected utility, for every i,

$$u_i \left(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i \right) \geq u_i \left(s_i'(\theta_i), s_{-i}(\theta_{-i}), \theta_i \right) \tag{2.1}$$

$\forall \, s_i' \neq s_i.$

Game theorists have been using the Nash equilibrium as one of the fundamental solution concepts. However, the Nash solution concept is very restrictive because of the assumptions imposed on the information available to each agent as well as each agent's beliefs about other agents. Each agent must have the same common perfect information about the preferences of every other agent in the model and all agents must select the same Nash equilibrium. Nash equilibrium doesn't hold in games with multiple equilibriums.

Dominant Strategy Equilibrium

The Dominant Strategy Equilibrium is a stronger solution concept which states that in equilibrium, every agent will have the same utility-maximizing strategy irrespective of the

strategies of all other agents. A strategy s_i is a dominant strategy if it weakly maximizes the agent's expected utility for all possible strategies of other agents,

$$u_i \ (\ s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \ \geq \ u_i \ (\ s_i'(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \tag{2.2}$$

$\forall \ s_i' \ \neq \ s_i, s_{-i} \ \in \ \sum_{-i}.$

Unlike the Nash solution concept which makes strong assumptions about the information available to agents about each other, dominant-strategy equilibrium makes no such assumptions. This makes dominant strategy a robust solution concept. In fact it doesn't even require agents to believe that others will choose its own optimal strategy in a rational manner. Hence, it is understandable why dominant strategy implementations of social choice functions are preferable over Nash implementations, in the context of mechanism design.

Bayesian-Nash equilibrium

The Bayesian-Nash Equilibrium states that in equilibrium, every agent selects an expected utility – maximizing strategy in equilibrium with the other agents' expected-utility maximizing strategies. A strategy profile $s = (\ s_1(.), s_2(.), …, s_I(.))$ is in Bayesian- Nash equilibrium if for every agent i and all preferences $\theta_i \ \in \ \Theta_i$

$$u_i \ (\ s_i(\theta_i), s_{-i}(.), \theta_i) \ \geq \ u_i \ (\ s_i'(\theta_i), s_{-i}(.), \theta_i) \tag{2.3}$$

$\forall \ s_i'(.) \neq \ s_{i()}$ and u_i denotes the expected utility over distribution of agent types $F(\theta)$

Bayesian-Nash is a stronger solution concept than the Nash equilibrium but weaker as compared to dominant strategy equilibrium, even though it makes reasonable assumptions about agent information.

Looking ahead to mechanism design, we can declare the preferable ordering of the different implementation concepts $Dominant \ > \ Bayesian - Nash \ > \ Nash.$ An ideal mechanism is

one which provides agents with a dominant strategy and implements a solution to the multi-agent distributed optimization problem.

2.4.3   Social Choice Function

The social choice function allows the selection of the optimal outcome given agent types and thus satisfies the system-wide goal of the mechanism. Social choice function $f : \Theta_1 \times \Theta_2 \times \ldots \times \Theta_I \rightarrow O$ chooses an outcome $f(\theta) \in O$, given types $\theta = (\theta_1, \theta_2, \ldots, \theta_I)$. The mechanism design problem is to implement the solution to the SCF.

The social choice function selects an alternative from a set of alternatives, given everyone's preferences. The goal of the mechanism design is to implement a social choice function that satisfies certain desired properties. For instance, in our research problem a good choice of the social choice function can be a function that maximizes the total utility of all the sensor platforms and the information gain of the common operating picture, when allocating the bandwidth.

In our discussion we shall assume that agents have quasi-linear utility functions and are risk-neutral. A quasi-linear utility function for agent i with type θ_i can be defined as

$$u_i(o, \theta_i) = v_i(x, \theta_i) - p_i \qquad (2.4)$$

where outcome o defines a choice $x \in K$ from a discrete choice set and a payment p_i by the agent, with valuation function $v_i(x)$.

With quasi-linear agent preferences we can separate the outcome of a social choice function into a choice $x(\theta) \in K$ and a payment $p(\theta)$ made by each agent i:

$$f(\theta) = (x(\theta), p_1(\theta), p_2(\theta), \ldots, p_I(\theta)) \qquad (2.5)$$

for preferences $\theta = (\theta_1, \theta_2, \ldots, \theta_I)$

2.4.3.1 Properties of SCF

- **Pareto Optimality:** A social choice function is said to be Pareto optimal or efficient if it implements outcomes for which no alternative outcome is strongly preferred by at least one agent, and weakly preferred by all other agents. SCF $f(\theta)$ is Pareto Optimal if for every $ó \neq f(\theta)$, and all preferences $\theta = (\theta_1, \theta_2, \dots, \theta_I)$

$$u_i(ó, \theta_i) > u_i(o, \theta_i) \qquad (2.6)$$
$$\Rightarrow \ for \ all \ j \in I \ , u_j(ó, \theta_i) < \ u_j(o, \theta_i)$$

- **Allocative Efficiency:** Allocative efficiency maximizes the total value over all agents. Social choice function $f(\theta) = \big(x(\theta), p(\theta)\big)$ is allocatively efficient if for all preferences $\theta = (\theta_1, \theta_2, \dots, \theta_I)$

$$\sum_{i=1}^{I} v_i(x(\theta), \theta_i) \geq \sum_{i=1}^{I} v_i(\acute{x}, \theta_i), \qquad (2.7)$$

for all $x' \in K$

- **Budget Balanced:** Budget Balance implies there are no net transfers in or out of the system. Social choice function $f(\theta) = \big(x(\theta), p(\theta)\big)$ is budget balanced if for all preferences $\theta = (\theta_1, \theta_2, \dots, \theta_I)$

$$\sum_{i=1}^{I} p_i = 0 \qquad (2.8)$$

- **Weak Budget Balanced:** Weak Budget Balance implies there can be a net payment made from agents to the mechanism, but no net payment from the mechanism to the agents.

$$\sum_{i=1}^{I} p_i > 0$$

<div align="right">(2.9)</div>

A mechanism $M = (\sum_1, \sum_2, \dots, \sum_I, g(.))$ defines the set of strategies $\sum_i$ available to each agent, and an outcome rule $g : \sum_1 \times \sum_2 \times \sum_I \rightarrow O$, such that $g(s)$ is the outcome implemented by the mechanism for strategy profiles $s = (s_1, s_2, \dots, s_I)$

Given mechanism M with outcome function $g(.)$, we say that a mechanism implements social choice function $f(\theta)$ if the outcome computed with equilibrium agent strategies is a solution to the social choice function for all possible agent preferences.

A mechanism has property P if it implements a social choice function with property P.

- **Ex-post Pareto optimal:** Mechanism M is Ex-post Pareto optimal if it implements a Pareto optimal social choice function $f(\theta)$ over specific agent types.

- **Ex-ante Pareto optimal:** Mechanism M is Ex-ante Pareto optimal if there is no outcome that at least one agent strictly prefers and all other agents weakly prefer in expectation.

- **Efficient mechanism:** Mechanism M is efficient if it implements an allocatively-efficient social choice function $f(\theta)$.

- **Ex-post Budget Balanced:** Mechanism M is Ex-post Budget Balanced if the equilibrium net transfers to the mechanism are non-negative for all agent preferences, every time.

- **Ex-ante Budget Balanced:** Mechanism M is Ex-ante Budget Balanced if the equilibrium net transfers to the mechanism are balanced in expectation for a distribution over agent preferences.

- **Individual Rationality (IR):** Mechanism M is said to be individually rational if the agents participate in the mechanism voluntarily in lieu of some incentive of participation.

Depending on which mechanism stage the agent chooses to participate in, there are three individually rational strategy profiles.

- **Ex-post Individual Rationality:** In this strategy profile $s(\theta)$ agents choose whether to stay or leave after announcing their types and learning an outcome from the set of feasible outcomes. If $\bar{u}(\theta_i)$ is the utility for opting out, then in an ex-post IR mechanism:

$$u_i \left(f(s_i(\theta_i), s_{-i}(.)), \theta_i \right) \geq \bar{u}(\theta_i) \tag{2.10}$$

$\forall \, \theta_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i}$ where $\Theta_{-i} = \Theta_1 \times \dots \times \Theta_{i-1} \times \Theta_{i+1} \times \dots \times \Theta_n$

- **Interim Individual Rationality:** In this strategy profile $s(\theta)$ agents choose whether to stay or leave after announcing their types but before the outcome is calculated

$$U(\theta_i | f) = E[u_i \left(f(s_i(\theta_i), s_{-i}(.)), \theta_i \right)] \geq \bar{u}(\theta_i) \tag{2.11}$$

- **Ex-ante Individual Rationality:** In this strategy profile $s(\theta)$ an agent must make its decision before knowing its type, therefore it must know the types' prior distribution

$$U(\theta_i) = E[u_i \left(f(s_i(\theta_i), s_{-i}(.)), \theta_i \right)] \geq E(\bar{u}(\theta_i)) \tag{2.12}$$

An ex-post IR strategy profile satisfies both the interim and ex-ante conditions. Interim IR is a stronger profile as opposed to ex-ante, as the agents in interim IR agree to

participate after learning their types, while the agents in the ex-ante IR don't have access to this information when they decide to participate in the mechanism.

One last important mechanism property, defined for direct-revelation mechanisms, is incentive-compatibility.

2.4.4   Direct Revelation Mechanism

A direct-revelation mechanism $M = (\Theta_1, \Theta_2, \dots, \Theta_I, g(.))$ restricts the strategy set $\sum_i = \Theta_i$ for all i, and has outcome rule $g : \Theta_1 \times \dots \times \Theta_I \rightarrow O$ which selects an outcome $g(\hat{\theta})$ based on reported preferences $\hat{\theta} = (\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_I)$. In other words, in a direct-revelation mechanism, the strategy of agent i is to report its type $\hat{\theta}_1 = s_i(\theta_i)$, based on its actual preferences θ_i

2.4.4.1   Incentive Compatible Mechanism

A mechanism is incentive compatible if it truthfully implements a social choice function. For direct revelation mechanisms, the social choice function $f(.)$ is truthfully implementable, if the mechanism has equilibrium $s(\theta) \in M$ in which $s(\theta_i) = \theta_i$ for all $\theta_i \in \Theta_i$. We can distinguish between three different equilibrium strategy profiles:

- **Dominant Strategy Incentive Compatibility:** A direct revelation mechanism M is dominant strategy incentive compatible (or strategy proof) if truth revelation is a dominant strategy equilibrium.

- **Ex-post Nash Incentive Compatibility:** A direct revelation mechanism M is Ex-post Nash incentive compatible if truth revelation is ex-post Nash equilibrium.

- **Bayesian Nash Incentive Compatibility:** A direct revelation mechanism M is Bayesian Nash incentive compatible if truth revelation is an Bayesian Nash equilibrium of the game induced by the mechanism

The *revelation principle* is the cornerstone of mechanism design and this principle is defined for direct mechanisms in conjunction with incentive compatibility. According to the revelation principle, any mechanism can be transformed into an equivalent incentive-compatible direct-revelation mechanism, under weak conditions, such that it implements the same social-choice function. This leads to the central possibility and impossibility results of mechanism design.

2.4.4.2   Dominant Strategy Revelation Principle

Suppose there exists a mechanism (direct or otherwise) M that implements the social-choice function $f(.)$ in dominant strategies. Then $f(.)$ is truthfully implementable in dominant strategy (strategy-proof mechanism).

The revelation principle implies that we just need to restrict attention to truth-revealing direct-revelation mechanisms.   The significance of dominant-strategy revelation principle can be understood when we are tasked to identify which social choice functions are implementable in dominant strategies. According to the revelation principle we need only to identify those functions $f(.)$ for which truth-revelation is a dominant strategy for agents in a direct-revelation mechanism with outcome rule $g(.) = f(.)$.

We provide a discussion on the key impossibility and possibility results within the Mechanism Design literature in Appendix A. The possibility result most pertinent to our research is the Vickrey-Clarke-Groves (VCG) Mechanism which is individually rational, incentive compatible and allocatively efficient (but not budget-balanced). In fact, the VCG mechanism is the only family of mechanisms that implement an efficient and individually-rational SCF where truth-telling is a dominant strategy, amongst direct-revelation mechanisms. The VCG mechanism achieves its strategy-proofness through its payment scheme whereby an agent's utility is aligned with that agent's marginal contribution to the mechanism.

We will now turn our attention to reviewing different auction-based mechanisms. We have already discussed and rejected three single-sided auctions earlier – English, Dutch and First Price. In the next discussion we briefly discuss the Second-Price sealed bid auction introduced earlier.

We then turn to the more generic and popular family of mechanisms- the VCG mechanism and highlight the shortcomings of this mechanism, with respect to our research. We also review the various modified-VCG auction-based mechanisms proposed in literature.

<div align="center">2.5   <u>Trust in Mechanism Design</u></div>

The *sealed-bid second-price auction* is also known as Vickrey auction. Vickrey auction is a special case of the family of VCG mechanisms, for allocating one item. It is identical to the sealed first-price auction in that all bidders submit sealed bids individually but the winner pays the second-highest bid rather than their highest winning bid. We can express this for the case, with bids b_1 and b_2 indicating the first- and second- highest bids respectively. In Vickrey auction the item is sold to the item with the highest bid (b_1), for a price computed as $b_1 - (b_1 - b_2) = b_2$ i.e. the second-highest bid.

The Vickrey auctions helps in understanding the intuition behind the strategy-proofness of the VCG mechanisms. In the Vickrey auction revealing one's true valuation is a dominant strategy as one's bid only dictates the range of prices one is willing to accept, not the actual price it will have to pay. The price that an agent pays remains independent of its bid price. Thus in case an agent knows the second-highest bid, it can still bid it's true value because it just has to pay enough to out-bid the other agent. The mechanism is weak budget-balance because the second-highest bid price is always non-negative, and the mechanism is individually-rational as the second-highest bid price is not more than the highest bid price, which in equilibrium, is equal to the winning bidder's value.

Yet in the context of this research, the VCG mechanism has several shortcomings:

1. **The VCG auction induces truth elicitation on the part of bidders, but not the auctioneer :-**

The VCG auction cannot prevent the auctioneer from lying, as the auctioneer could ask for a higher price than the second price. The bidders have no access to the sealed and private bids of the others and the winner has no option but to pay the requested price. M. Hsu & Soo (2002)

address this drawback by randomly delegating one of the bidders to be the auctioneer and use a public blackboard to publish the agents' bids. This blackboard serves the purpose of allowing the first and second highest bidders to verify the results, but introduces complexity to the mechanism and compromises the robustness of the network. Yokoo & Suzuki (2004) have proposed an encryption method that guarantees secure bidding without endangering the privacy of bids. However it introduces significant latency in the communication as it involves multiple overlays of calculations and is not ideal for real-time networks.

2. **VCG mechanism assumes that the bidder's valuations are not interdependent on others' valuations :-**

In tactical sensor networks, the observations made by the sensor agents are polluted by uncertainty and noise. As a result, individual sensors have a limited and partial view of the common operating picture. Indeed, other sensor agents may possess information that would, if known to a particular bidder, affect the value he assigns to the object. The resulting information structure is one of *interdependent values* (Krishna, 2002). A naïve extension of the VCG mechanism is known not to work in the case of interdependent valuations.

In the last decade, auctions based on interdependent valuations have garnered significant attention among researchers. Some of the most relevant work include (Dasgupta & Maskin, 2000; Jehiel & Moldovanu, 2001; Krishna, 2002).

Krishna (2002) looked at direct mechanisms for a number of scenarios: single and multiple items, two and more buyers as well as single dimensional and multi-dimensional signals. For the base scenario of multi-bidders single-item with single-dimensional signals, he showed the existence of efficient allocations. In this direct mechanism, the bidding agents would submit their interdependent valuation functions along with their private signals to a central auctioneer. If these valuations satisfied certain conditions, the auctioneer would be able to decide the efficient allocation of the item and the payment scheme was devised to incentivize the agents to reveal their signals truthfully. In the case of multi-items and multi-dimensional signals, efficient allocations could not be achieved. Krishna concluded that the problems associated with

interdependent valuations do not arise from the multiplicity of the items per se but rather from the multiplicity of the dimension of the signals received by the bidders.

Dasgupta & Maskin (2000) approached the problem of interdependent valuation in the case of single item and single-dimensional signals, from a different perspective. In their mechanism the agents, instead of submitting their valuation functions and observed signals to the auctioneer, would instead make contingent bids – bidder A would submit a range of bids that specified its bid when bidder B bid a particular value and so forth. This makes the bidding process complex for the case of single item, which becomes even more complex when multiple items need to be allocated. This ends up complicating the mechanism because instead of just revealing their valuation function and signals as in Krishna's mechanism, the agents now have to submit bids based on what other agents might bid.

Krishna postulated that it is the multiplicity of the dimension of the signals which leads to inefficient allocations in direct mechanisms. Jehiel & Moldovanu (2001) took this notion forward and showed that it is not only VCG mechanisms that are affected by this particular shortcoming. Instead, in an interdependent valuation setting, no one-stage mechanism could achieve both efficiency and incentive compatibility for procurement of estimates from multiple sources.

Mezzetti (2004) addressed this challenge to a certain extent and showed that an efficient allocation with multidimensional types is possible, if (a) values are privately realized by the agents once an allocation is made and (b) two-stage mechanisms can be adopted in which the payments are made contingent on realized values reported in a second stage. Mezzetti designed a two-stage mechanism: in the first stage the agents would submit their reports to the auctioneer, who would in turn determine the allocation of the items among the bidding agents. In the second stage the agents observe their payments and receive the final transfers from the auctioneer. However the equilibrium relies on agents correctly reporting their realized values – this poses an issue for our research scenario, as the agents will have a tendency to under-report the quality of their data, so that the bandwidth is allocated to the other agents in the scenario for the transmission of their data.

Klein et al. (2008) addressed this pertinent issue of under-reporting valuation. They assume since the auctioneer (or the trusted center in their case) is located within the network and is capable of observing the quality of the data broadcasted, the trusted center could infer for each agent, its value from the realized outcome. Although their mechanism deals with agents' under-reporting the quality of their information, it has no means of handling agents degrading their measurements and reporting data of poor quality. What this translates into is that the agents can get away with investing few or none of its resources in generating the information, which it can subsequently report truthfully thereby subverting the allocative efficiency of the mechanism. Another drawback is that the payments made to the agents are based on the realized outcome, rather than the expected one. Thus their proposed mechanism cannot handle selfish behavior on the part of the agents nor the uncertainty that characterizes the ULS system environment.

3. VCG mechanism can't ensure budget-balance

VCG mechanisms are not budget-balanced and cannot operate without intervention. It needs funds to be regularly deposited with the auctioneer in order to maintain wealth for the agents. Although budget balance is not one of our research objectives, it is a desirable property for the mechanism to operate independently without relying on external sources or intervention.

Cavallo (2006) proposed a redistribution mechanism to achieve budget balance in a modified VCG mechanism. This redistribution mechanism is applied once the center computes the allocation of single or multiple items among two or more bidders, in order to achieve budget balance asymptotically. For the case of ten or more agents it can achieve redistribution of more than 70% of the VCG surplus. Petcu, Faltings, & Parkes (2006) proposed two distributed mechanisms to achieve different degrees of budget balance: redistributing the VCG tax to attain weak budget balance and achieving exact budget balance at the expense of optimality. In fact, sacrificing optimality or individual rationality for budget-balance is a natural consequence of the impossibility results (Appendix A). However, within the framework proposed by Petcu et al. (2006) the agents are assumed to have precise and perfect valuations, which do not hold up in operating environment characterized by uncertainty.

In conclusion, we have reviewed several auction-models that relate to our research problem, and pointed out the inherent flaws that characterize auction-based mechanisms. Thus we need to shift our focus from the realms of auction based models, to another promising alternative approach within the Mechanism Design research domain, in the form of strictly proper scoring rules.

CHAPTER 3.   MULTI-AGENT SYSTEMS FRAMEWORK

In this chapter, we provide an overview of the Multi-Agent System framework used to generate the surrogate model for our research problem. We introduce the agent-based modeling tool Discrete Agent Framework (DAF), developed at Purdue, and describe the agent-based model framework created in DAF to emulate the combat tactical data network and construct the common operating picture.

3.1   Multi-Agent System Framework

We need to generate a surrogate model for the real world which exhibits the necessary fidelity and complexity to study the application of mechanism design in a practical context. To this end, we use the Multi-Agent System framework to model the tactical data network dynamics, and incentivized interactions between participating agents and the resultant quality of information gained due to fusion of sensor data. The model goes a long way in helping us translate the research problem and objectives into real-world settings and eliciting the requirements of the desired mechanism.

Over the past few decades, the field of Multi-Agent System (MAS) has garnered quite a lot of attention in science and engineering. It has become particularly important in different aspects of computer science, like robotics, distributed systems and artificial life and intelligence. The field of MAS is uniquely suited to account for operational independence of agents, along with phenomena related to free will, competition, consensus, communication, belief, knowledge and deception. It is an ideal framework to model the decentralization, uncertainty, evolution, and heterogeneity that characterize ultra large scale systems.

3.1.1 Agent Characteristics

Before we can proceed further with our discussion on MAS, it is necessary to define what we mean by the terms 'agent' and 'multi-agent system'. Wooldridge (2009) gave a definition that is oft-cited: "an *agent is a computer system that is situated in some environment, and that is capable of autonomous action in the environment in order to meet its design objective".* Russell, Norvig, Canny, Malik, & Edwards (1995) define an agent as "*anything that can perceive its environment through sensors and act upon that environment through actuators*". Ferber (1999) floated the notion of a minimal common definition which is widely accepted as the most detailed and definitive description of an agent:

An agent is a physical or virtual entity

    a)   *which is capable of acting in an environment,*

    b)   *which can communicate directly with other agents,*

    c)   *which is driven by a set of tendencies (in form of individual objectives or of a satisfaction / survival function which it tries to optimize),*

    d)   *which possesses resources of its own,*

    e)   *which is capable of perceiving its environment (but to a limited extent),*

    f)   *which has only a partial representation of this environment (and perhaps none at all),*

    g)   *which possesses skills and can offer services,*

    h)   *which may be able to reproduce itself,*

    i)   *whose behavior tends towards satisfying its objectives, taking account of the resources and skills available to it and depending on its perception, its representations and the communications it receives.*

The definition as provided by Ferber (1999) in (a) - (i) represent a set of notional behaviors, namely, autonomy, communication, intentionality, reactivity, flexibility, learning and self-actuation. Each of the terms in this definition is significant. Agents exhibit a host of different properties and characteristics, which enable us to classify agents using different schemes. We discuss some of these distinguishing properties next.

Autonomy

Perhaps, the most important and defining characteristic of an agent is the capability to act autonomously. By autonomous, we mean the agents are directed by a set of inherent tendencies and not by external commands originating from another agent or a user. In fact, agents are free to accede to, or, reject requests from other agents. Agents are not passive entities, merely responding to its environment and other agents; but are actively initiating actions to achieve certain individual goals or optimize certain satisfaction or survival functions.

Adaptive/Reactive

An adaptive agent has the ability to learn from past experiences and modify or adapt its behavior based on its accumulated experiences. By extension of adaptability, agents require some memory in order to store and learn from experiences. An agent may have predefined rules or some abstract complex mechanism that allows it to adapt to its environment. The concept of adaption is applicable in both the individual and group cases for agents.

Reactive agents, on the other hand, are directed by their individual goals and satisfaction or survival functions and are incapable of adapting to its environment. Reactive agents have no representation of their environment or of other agents, and no sentience to adapt to changing environment. Nevertheless, reactive agents are interesting in that they can constitute groups or colonies, which are capable of adapting to their environment.

Goal-Oriented/Utility-based

An agent may be goal oriented which indicates that its actions are directed towards achieving a given or computed goal. Goal oriented agents don't have any utility or objective functions to optimize but have motivation mechanisms pushing them towards accomplishing a goal.
Utility – based agents have predefined satisfaction or survival functions which it attempts to optimize. A goal specifies a crude distinction between success and failure while utility functions tend to be more generic and provide a measure of success for a given state. Utility – based

agents have to make decisions comparing choice between conflicting goals, and choice between likelihood of success and importance of goal.

Physical / Virtual

Physical entities are agents which act in the real world. It can take the form of physical objects like aircrafts or robots. Virtual entities, such as software modules, have no physical presence in the real world.

Social

An agent is social if it is capable of dynamically interacting with other agents in the system. Agents are endowed with the ability to distinguish the features of the agents with which they communicate. The designer may explicitly model the communication protocols, which define the rules of interaction with other agents as well as the environment.

Agents may exhibit a host of different characteristics apart from the ones encompassed in Ferber's minimal definition. Agents may be *mobile*, capable of transporting itself from one machine to other. Agents explicitly designed to collect, classify and filter information are called *info-gathering* agents. Agents may have *character*, i.e. a believable personality and emotional state. Agents are *temporally continuous* and run as a continuous process. Some agents which require assistant users are termed as *interface agents*.

Various schemes exist for agent classification: Agents may be classified according to the tasks they perform, or the range and sensitivity of their senses, or their control structure, or how much internal sate they possess, or the environment, in which they operate, or the range and effectiveness of their actions. And, there are many, many more such classification schemes we can choose from.

In modeling languages, like Unified Modeling Language (UML), agent classification is important as it provides a basis for representing the common underlying features and/or capabilities for

each and every agent. Odell, Nodine, & Levy (2005) describe UML constructs for classifying agents based on the capabilities that they have from their physical implementation (Agent Physical Classifiers) and from their current activities (Agent Role Classifiers).

- **Agent Physical Classifier:** It defines the sets of core, or primitive, features that all agents possess. Every agent must be classified according to some fixed Physical Classifier.

- **Agent Role Classifier:** It classifies agents according to the roles they are capable of playing at any given time. Agents can change roles over time (dynamic classification) and be associated with multiple roles at the same point in time (multiple classifications).

Franklin & Graesser (1997) have provided a hierarchical classification of agents based on the properties of an agent. We highlight some of the important agent characteristics, relevant to our research problem, with the aid of an interactive flowchart in Figure 3-1.
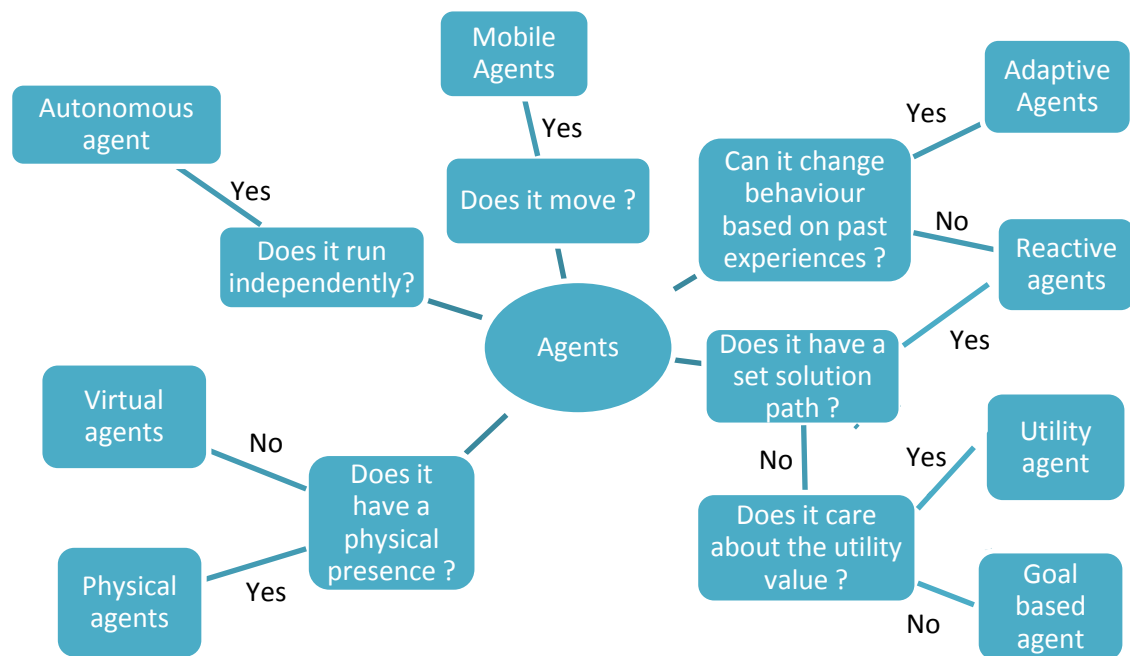
Figure 3-1 Classification of Agents

3.1.2 Multi-Agent Systems

Multi-agents systems are even harder to define than agents. 'Multi-agent system' is commonly used as an umbrella term for different types of systems, because of its application in different fields. Once again we fall back to the minimalist definition as provided by Ferber (1999).

The term 'multi-agent system' (MAS) is applied to a system comprising the following elements:

a) *An environment, E, that is, a space which generally has a volume.*

b) *A set of objectives, O. These objects are situated, that is to say, it is possible at a given moment to associate any object with a position in E. These objects are passive, that is, they can be perceived, created, destroyed and modified by the agents.*

c) *An assembly of agents, A, which are specific objects (A $\subseteq$ O), representing the active entities of the system.*

d) *An assembly of relations, R, which link objects (and thus agents) to each other.*

e) *An assembly of operations, Op, making it possible for the agents of A to perceive, produce, consume, transform and manipulate objects from O.*

f) *Operators with the task of representing the application of these operations and the reaction of the world to this attempt at modification, which we shall call the laws of the universe.*

Given the different definitions and fields of applications, there exist a plethora of different terms for the same methodology. Whereas most authors refer to these models as multi agent models, others talk about agent-based models. Similarly, the terms *multi-agent simulation* and *Agent-Based Modeling* (ABM) are used interchangeably.

Vlassis (2007) lists six fundamental aspects of MAS that distinguishes it from single agent systems. We present these fundamental traits with a short description of each, in Table 3-1 . We capture how our research problem of bandwidth allocation in tactical defense scenarios and by extension, Ultra Large Scale (ULS) systems, exhibits these characteristics as well.

Table 3-1 Fundamental Characteristics of MAS

Characteristic	Description	Research problem
Agent Design	Includes heterogeneous agents whose characteristics and behaviors vary in their extent or sophistication.	Consists of heterogeneous agents like targets, sensor platforms etc. which affect all functional aspects from perception to decision making.
Environment	Unlike the static environment of single-agent systems, MAS environment appears dynamic from the agent's point of view due of the presence of multiple agents.	The military platforms in our research problem operate in an uncertain and dynamically evolving environment, which is a primary characteristic of ULS systems.
Perception	Agents have only a partial representation of their environment and have no overall perception of what is happening.	The sensors onboard the military platforms have an incomplete and inaccurate view of the operating picture and have to fuse their perceptions to improve the accuracy of information.
Control	In MAS the control is decentralized for reasons of robustness. Coordinating the actions of the agents is a challenge, which is the subject of game theory.	The sensor platforms make their own decisions independently, and the mechanism coordinates the actions of the multiple agents.
Knowledge	Common-knowledge which indicates how much one knows about the current world, and what every agent knows that every other agent know about and so on.	The solution concept of the game theoretic formulation depends on the common knowledge. The revelation principle in MD lets us restrict our attention to strategy-free mechanisms.
Communication	Agents can both receive and transmit messages, which allows for coordination and negotiation.	Rules of interaction are specified in the form of protocols and message formats, like Link 16.

Weiss (1999) provides a comprehensive and detailed overview of the architecture of multi-agent systems with the potential ranges of their attributes. We encapsulate the highlights in Table 3-2:

Table 3-2 MAS architecture attributes

	Attribute	Range and Types
Agents	Number	Two or more
	Uniformity	Homogenous/Heterogeneous
	Goals	Contradictory/Complementary
	Architecture	Reactive/Deliberative
	Abilities	Simple/ Complex
Interactions	Frequency	Low/High
	Persistence	Short/Long term
	Level	Signal passing/knowledge intensive
	Pattern	Decentralized / Hierarchical
	Variability	Fixed/Dynamic
	Purpose	Competitive/ Cooperative
Environment	Predictability	Foreseeable/Unforeseeable
	Accessibility	Limited/Unlimited
	Dynamics	Fixed/Variable
	Diversity	Poor/Rich
	Resource availability	Restricted/Sufficient

3.1.3   Short History of MAS

Compiling a concise, definitive history of MAS is a challenging task given that the roots of this methodology are immersed in a host of different disciplines. For a more detailed analysis of the same, the interested reader should peruse Bousquet & Le Page, 2004; Ferber, 1999; Wooldridge, 2009.

Originally, the roots of MAS stem from the field of Artificial Intelligence (AI) where it first appeared in mid-nineteenth century but didn't garner much attention until the mid-1980s. The

AI community was more focused on single agent: its autonomy and its environment. Bousquet & Le Page (2004) credit the field of Distributed Artificial Intelligence (DAI) as the root of MAS. Researchers like Huhns & Stephens (1999) focused more on the organization of multiple agent interactions. The later researchers took inspiration from biology and created hybrid architectures based on reactive and reasoning behavior. Langton (1989) worked on the Artificial Life theory, a field based on physics and the general context of the sciences of complexity. Research of MAS moved independently and simultaneously until about the early 1990s.With the advent of the internet and electronic commerce, the interest in MAS skyrocketed in the mid-1990s and at the same time researchers started to reformulate certain questions in the social and natural sciences, based on MAS. All in all, the history of multi-agent systems was influenced by their multidisciplinary nature and has influences from computer science, natural sciences, cognitive psychology, sociology, linguistics and other social sciences.

3.1.4   Applications of MAS

Ferber (1999) differentiates the applications of Multi-Agent Systems into five main categories as depicted in Figure 3-2.
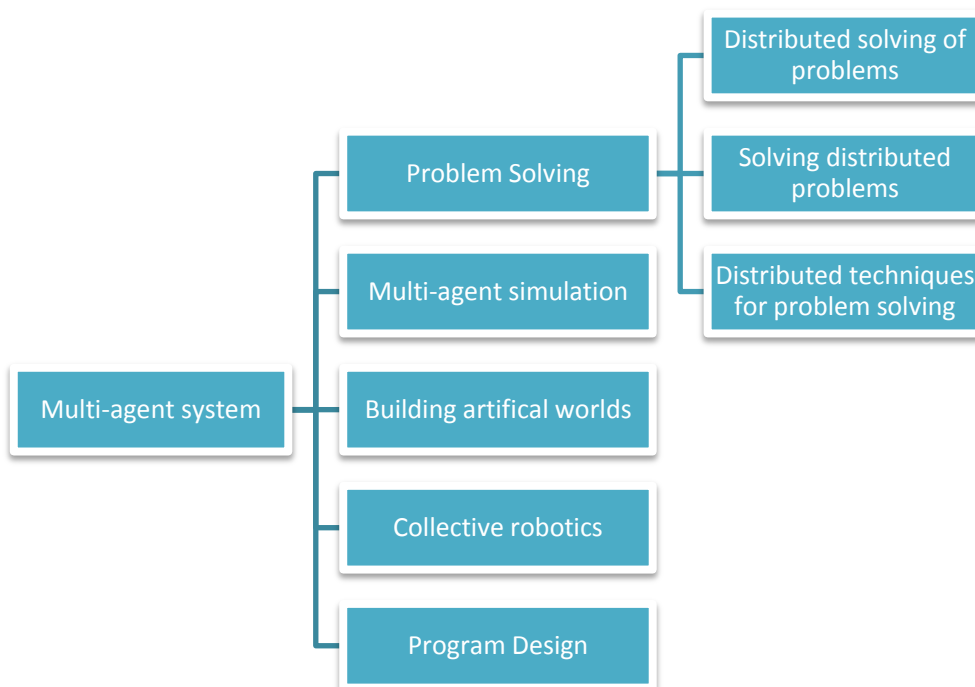
Figure 3-2 Applications of MAS

3.2   Agent – Based Modeling

Agent-Based Modeling (ABM) or Multi-agent simulation is a powerful simulation modeling technique where a system is modeled as a collection of autonomous agents. ABM is a relatively new approach to modeling the dynamics of complex systems and complex adaptive systems. The ABM mindset advocates modeling a system from the perspective of individual constituent entities. By modeling systems from ground-up, the range of diversity of attributes and behaviors among the heterogeneous agents give rise to organization and behavior of the system. Even a simple ABM with few agents and interaction rules, can give rise to complex behavioral patterns. Unanticipated and non-programmed patterns, behaviors, and structures emerge from agent interactions and provide valuable information of the system the ABM sought to emulate. ABM provides a methodology to model social systems with autonomous adaptive rational agents, which are not amenable to rigorous mathematical modeling.

ABMs can be characterized by four constructs (Jiang & Gimblett, 2002):

- **Agents:** Set of all simulated heterogeneous behavioral entities.
- **Objects:** Set of all represented passive entities that do not react to stimuli.
- **Environment:** Topological space where agents and objects are located and signals propagated.
- **Communications:** Set of all possible communications between entities.

3.2.1   Benefits of ABM

Bonabeau (2002) captures the benefits of ABM over other modeling techniques in three concise statements: ABM captures emergent phenomena; ABM provides natural description of the system; and ABM is flexible.

1. **Emergence:-**

Perhaps, the biggest advantage of ABM is its unique ability to capture non-intuitive emergent phenomena resulting from the interaction among agents. The guiding principle behind emergent phenomena is the same as that behind system-of-systems: "The whole is more than a

sum of its constituent parts." Emergence can be counterintuitive: a traffic jam moving in a different direction opposite to the cars that cause it. Conway's Game of Life (Conway, 1970) provides an example of emergent behavior. The properties of emergence are decoupled from the properties of the agent or its interactions and this makes them difficult to predict. ABM is typically used whenever there is a potential for emergent phenomena: individual behavior is non-linear and is characterized by thresholds or non-linear coupling and exhibits non-markovian behavior or hysteresis. ABM is particularly suited to capture systems where aggregation will simply not work. Aggregate differential equations assume global homogenous behavior and tend to even out fluctuations. However under certain conditions, these fluctuations can be amplified and show significant deviation from the predicted aggregate behavior.

2. Natural Description:-

ABM is the most natural technique for modeling systems composed of behavioral actors. The ABM is not an abstract holistic overview description of a system – ABM simulates a system from the perspective of the actions of the individual constituent entities. ABM should be used when systems are more naturally described through activities rather than processes, or when describing the complex individual behavior of agents renders the different equations as intractable.

3. Flexibility:-

With ABM, one doesn't need to know the exact agent complexity and description ahead of time. ABM exhibits flexibility in multiple dimensions. One can change the number of agents in the simulation, and the levels of description or complexity. ABM also provides a natural framework to work with the agent behavior, rules of interaction, agent rationality and adaptivity and others. The main intuition behind using ABM is not to optimize but rather seek the adaptive nature of agent rules and behavior.

3.2.2   Use Cases of ABM

Axtell (2000) argues that there exist three distinct uses of ABM techniques. We illustrate each of these use cases with the help of examples provided in the original paper:

1. **Models that can be formulated and completely solved**

This usage is closest conceptually to traditional simulation in operations research. This use arises when a social process can be explicitly formulated in the form of soluble equations. If these equations can be solved numerically, then the ABM acts as a Monte-Carlo analysis and if they are soluble analytically, the ABM acts as a tool for presenting the mathematical results.

For example, consider the classical OR simulation of a bank teller line. This is a queuing model and no general analytical solution is known for arbitrary distributions of arrivals and service times. Therefore, the queuing process is commonly simulated via the Monte Carlo method and distributions of waiting times and server utilization result. However, this is completely equivalent to actually instantiating a population of agents, giving them heterogeneous arrival times according to some distribution, and then running the agent-based and studying the queue lengths that emerge, over time building up the entire waiting time distribution function.

2. **Models that are partially soluble**

The most important use of agent models is for this class of problem where one can describe the system process using mathematical equations, but they are not completely soluble. In this case the ABM serves as a way to gain insight into the functioning of the model. It can shed light on solution structure and properties of the model. What's more, once a model has been created it not only provides information about the stability or the equilibria of the solution but rather entire trajectories of the solution. Furthermore, the computational model can be used to test dependence of the results on assumptions and values of parameters and even provide counter-examples.

In models of traffic, important output statistics are the distributions of jams by size and lifetime. Agent-based models have been created to study the dynamic aspects of traffic e.g., Nagel &

Rasmussen (1994). *These models are capable of reproducing real world data with high fidelity* (Casti, 1997). *In particular, on crowded roads it is known that local flow-rate data can be highly non-stationary. Differential equation (fluid mechanical) models of traffic have a difficult time of capturing this feature of the data, as well as faithfully representing other transient and dynamical properties of real world traffic. However, these phenomena do emerge in large-scale, massively parallel (i.e., agent) computational models of traffic. Additionally, the jamming distributions that arise in these models display a kind of universality also seen in statistical physics. That is, the macro-statistics of the systems are relatively insensitive to the agent specifications, i.e., many reasonable models of driving behavior produce the same distribution of traffic jams!*

3. Models that are intractable

When a social process model is either apparently or provably intractable, then trying to express the same through a mathematical framework is an exercise in futility. In such instances, ABM is the only available technique for systematic analysis and a viable substitute for formal mathematical analysis.

For example, it is well-known that there does not exist closed form solutions to certain relatively simple differential equations in terms of elementary functions. When a problem is intractable in this way it has nothing to do with its complexity. Rather, it is an artifact of the limited explorations undertaken to date in the infinite library of functions. In such circumstances one makes recourse to numerical solution. But there are also instances in which numerical solution would appear to be essentially intractable, not in the sense of being impossible but merely not useful. This occurs when governing equations are highly nonlinear. When such circumstances arise in computational physics, particle models can sometimes be advantageous. The same is true of agent-based computational models in the social sciences: if it is hard to make any real progress solely by analytical manipulation, agent models may prove useful.

In each of these three cases, ABM can assist with:

- Visual output and symbolic check of the mathematical solutions

- Introduce bounded rationality, and study sensitivity of solutions

- Explain system dynamics and observed phenomena

3.2.3   Applications of ABM

Agent-based modeling has been used in quite a variety of application ranging from physical sciences to social sciences, from biological sciences to management sciences. Bonabeau (2002) delineates the application areas of ABM into four distinct categories using real-world examples:

1. Organizations: Organizational design and operational risk.
2. Markets: Strategic simulation, stock market and software agents.
3. Diffusion: Innovation diffusion and adoption dynamics.
4. Flows: Traffic, customer, and evacuation flow management.

Our discussions indicate that ABM is uniquely suited to model ultra-large-scale, decentralized systems as it provides a natural framework for describing systems comprising of autonomous rational and self-interested agents interacting with each other in a dynamic and uncertain environment.

3.3   ABM as a Tool for Modeling ULS systems

Agent-based modeling is an important tool for the engineering of large-scale decentralized complex systems. ABM is increasingly the tool of choice for modeling systems as well as system-of-systems, as it provides systems engineers the means to investigate alternative architectures and gain an understanding of the impact of the behaviors of individual systems on emergent behaviors. We believe that these properties make ABM an ideal tool to analyze ULS systems as well.

3.3.1   Brief Literature Review

The recent increase in availability of software packages for agent-based simulation and the increased understanding that agent-based modeling is well suited for modeling complex decentralized systems has resulted in many applications of agent-based modeling. Kilicay-Ergin & Dagli (2008) describe the use of AnyLogic agent-based simulation software to model the

behavior of alternative system architectures for financial markets. J. Hsu, Price, Clymer, Garcia-Jr, & Gonzalez (2009) describe using OpEMCSS software to simulate the behavior of a complex adaptive system that they call the World Model. Giachetti, Marcelli, Cifuentes, & Rojas (2013) describe a simulation that uses Java-based CybelePro software to model the performance of a human-robot team as an agent-based system.

3.3.2   Comparison of ABM to Other Approaches

Various methods and tools for simulation of complex processes exist; however, they primarily fall into the main categories of equation-based system dynamics, discrete event simulation, and agent-based modeling. The following section provides comparative discussion on the differences between agent-based models and the other methods.

System dynamics is defined as the "the study of information-feedback characteristics of industrial activity to show how organizational structure, amplification (in policies), and time delays (in decisions and actions) interact to influence the success of the enterprise" (Jay, 1958). Typically, system dynamics represents an aggregate level of performance as continuous differential equations. Analysts employ these models support strategic-level decision-making and to develop an overarching view of long-term trends in the dynamics of an enterprise. System dynamics has been widely used across a range of applications that range from socio-economic to engineering systems, and aims to reduce complex behaviors to their most aggregate forms assuming that adequate, structured representations of the behaviors exist.

Borshchev & Filippov (2004) and Schieritz & Milling (2003) provide a comprehensive comparison of simulations that use systems dynamics versus simulations that use agent-based modeling. System dynamics focuses on a top-down, aggregate modeling that typically uses continuous-form representations (feedback loops) of system processes; agent-based models are based on discrete agent-specific logic rules that take a bottom-up approach to simulation. Agent-based modeling provides a means of connecting micro-level behaviors to the macro level of a system whereas systems dynamics link system structures to system behavior. The main difference is in the ability of agent-based models to capture emergent behaviors. The agent-based setting

allows flexible interactions between individual agents, which results in non-intuitive dynamic modes being generated. System dynamics reduces the possibility of exploring emergent phenomena due to the natural filtering of these modes that occurs through enforcement of aggregate equations over populations of individuals within the system, and the establishment of a rigid flow structure.

Discrete Event Simulation is a method used to model real world processes as a series of interconnected discrete events that are functional processes. These processes are typically at the low- to mid-level state of abstraction in the hierarchy of interconnected systems and do not consider performance characteristics of the individual elements that execute these processes themselves. The focus of a process-centric simulation here is naturally well suited to applications where processes are the critical aspect of analysis such as in healthcare (e.g., patient flow), manufacturing (e.g., production floor processes layout), and logistics (e.g., distribution processes at a hub).

As with the system dynamics approach, the discrete event simulation approach is a top-down approach that models aggregate behaviors of processes. While discrete methods use a powerful and intuitive representation of processes in a system, they are mainly intended to model and represent finite interactions where the underlying structure of the process is already known. They share the focus with systems dynamics of modeling top-down characteristics of a system and assume pre-defined structures and aggregations of macro behaviors. In contrast, agent-based models are able to more generally represent individual entities that drive the discrete events and allow for possible emergent behaviors that are not otherwise apparent from the aggregated discrete dynamics of a system.

3.3.3 DAFApproach to ABM

Purdue University developed the Discrete Agent Framework (DAF) for agent-based modeling in 2010 to enable easy application in multiple domains. Developed in object-oriented *MATLAB*, this engine provides the foundation to build agent-based simulation models to explore various architecture configurations for large scale complex interconnected systems and evaluate their

performance. DAF also enables coordinated development, verification, and validation of the system architecture through selective failure simulation.

DAF allows the modeling effort to focus on the systems architecture itself, reducing the computational modeling overhead. The first major application of DAF was through a sponsored research initiative of the Missile Defense Agency of the US Department of Defense to examine and model a Ballistic Missile Defense System (BMDS) by simulating it as a collection of functions (executed by agents).

DAF views system architecture as a collection of agents that are connected by communication links. In practice, each agent in DAF is an in-code application of a formal model developed from research and of communication links that emulate real or proposed communication standards. This approach allows a DAF user to follow Maier's communication-centric architecting approach (Maier, 1998) by using the different possibilities of linking these agents as a means to distinguish one architecture from another. DAF can be used to generate and evaluate a wide variety of architectures by defining the functional capabilities and behaviors of agents and the communication links between agents. A representative implementation of DAF involves generating architecture alternatives, then simulating them to identify the configuration that provided the best balance of efficiency and reliability.

There are certainly many agent-based modeling packages available and each provides many combinations of capability, ease-of-use, and availability. For example, NetLogo is a Java-based package that provides multi-platform complex system simulation, but also comes with a large database of sample models and implementations (Wilensky, 1999). SWARM is another open source Objective C/Java-driven simulation system for modeling complex systems through discrete event simulation. Initial development work on the SWARM system indicates that an object-oriented development environment is ideal for building agent-based simulations (Minar, Burkhart, Langton, & Askenazi, 1996). The primary differentiator between DAF and other packages is that DAF is MATLAB-based, and as a result, any DAF application can utilize the many mathematical, statistical, and visualization tools built into MATLAB or the many supported and third-party toolboxes associated with MATLAB. Additionally, the widespread use of MATLAB in

research and industry reduces the time required to learn how to use DAF and the time to apply it to a particular project.

What makes DAF an ideal tool for emulating tactical data links for our research problem is that the standard communication protocols such as TCP (Transmission Control Protocol), SDP (Session Description Protocol) and UDP (User Datagram Protocol) have already been modeled in DAF. Anticipating the need of modeling communication protocols for a reliable and accurate simulation, the developers of DAF created Communication Agent or "comms" agent. The comms agent serves as a middle-man between two agents and forwards messages from sender to destination based on the latency corresponding to a communication protocol. Figure 3-3 illustrates the placement of the communication agent in an architecture, where the dotted line indicates the communication link without a comms agent.

While the comms agent facilitates the implementation of communication latency in message passing, it also provides the ability to model adverse effects of cyber-attacks and impact of encryption.
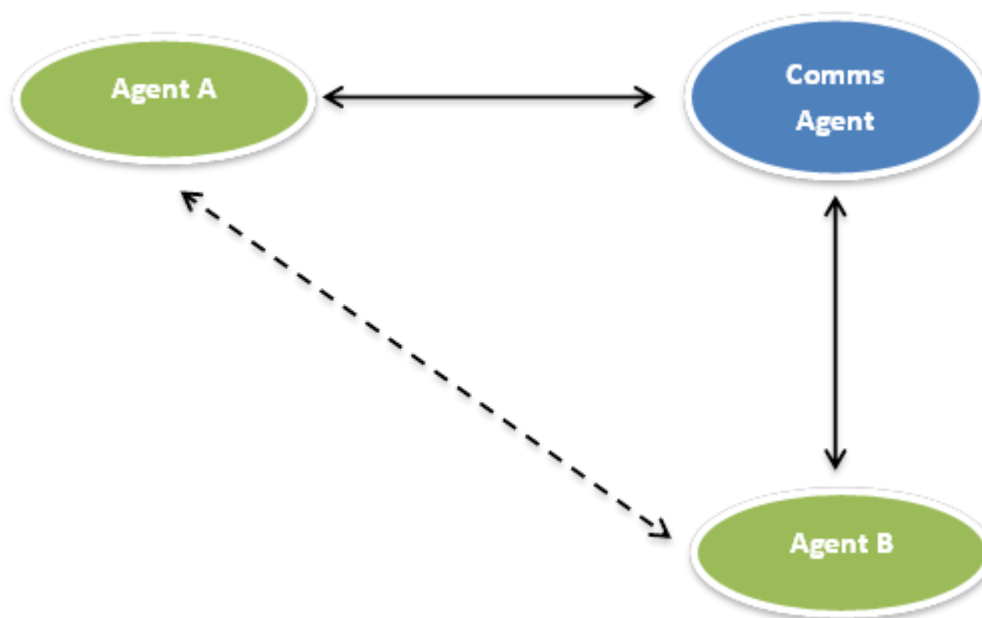
Figure 3-3 Communication architecture in DAF

The current implementation of the communications agent supports three different communications protocols: TCP, UDP, and SDP. These three standards model different levels of message delivery guarantees and latency. A comparison of the three modeled protocols is provided in Table 3-3:

Table 3-3 Default Protocols in DAF

Protocols	Description	Pros	Cons
TCP	Re-transmits message until delivered	Guaranteed message delivery (no packet loss)	High latency
UDP	Does not re-transmit message	Low latency	Unreliable message delivery (high packet loss)
SDP	Re-transmits message 'n' times (1<n<∞)	Lower packet loss than UDP and lower latency than TCP	Higher packet loss than TCP and higher latency than UDP

The communications agent also provides the ability to model the physical link type that exists in the real world, such as satellite communications and fiber optic cable. This is advantageous since link type determines bandwidth, distance-based link latency, and probability of link loss. Table 3-4 below provides a brief comparison of the three supported link types.

Table 3-4 Default Links in DAF

Link type	Link latency (ms)	Bandwidth
Fiber link	0.5 - 12.5	1 Gbps
Wireless line-of-sight	0.1 - 2.2	1400 Kbps
SATCOM	250	1400 Kbps

DAF is uniquely suited for modeling the elementary functionality of Link 16 communication protocol between agents, as it applies to our problem scenario. We can plug in the appropriate values of bandwidth, latency, and probability of link loss to model the characteristics of Link 16.

3.4   Research Framework

We develop an agent-based simulation model framework within DAF to capture the performance of a group of military platforms tasked with the goal of detecting and tracking targets. Figure 3-4 captures the paper model of the ABM featuring three distinct categories of agents: targets, sensor platforms and the trusted center (auctioneer).

The sensors onboard the military platforms can detect and classify targets within its region of observation. The platforms transmit the measurements from their sensors to the trusted center, which allocates the targets among the platforms for tracking. Each platform then broadcasts its observations over the tactical data link to all other platforms in the simulation.

We will highlight the key concepts of platform behavior and interaction, by examining the research application framework, one step at a time.
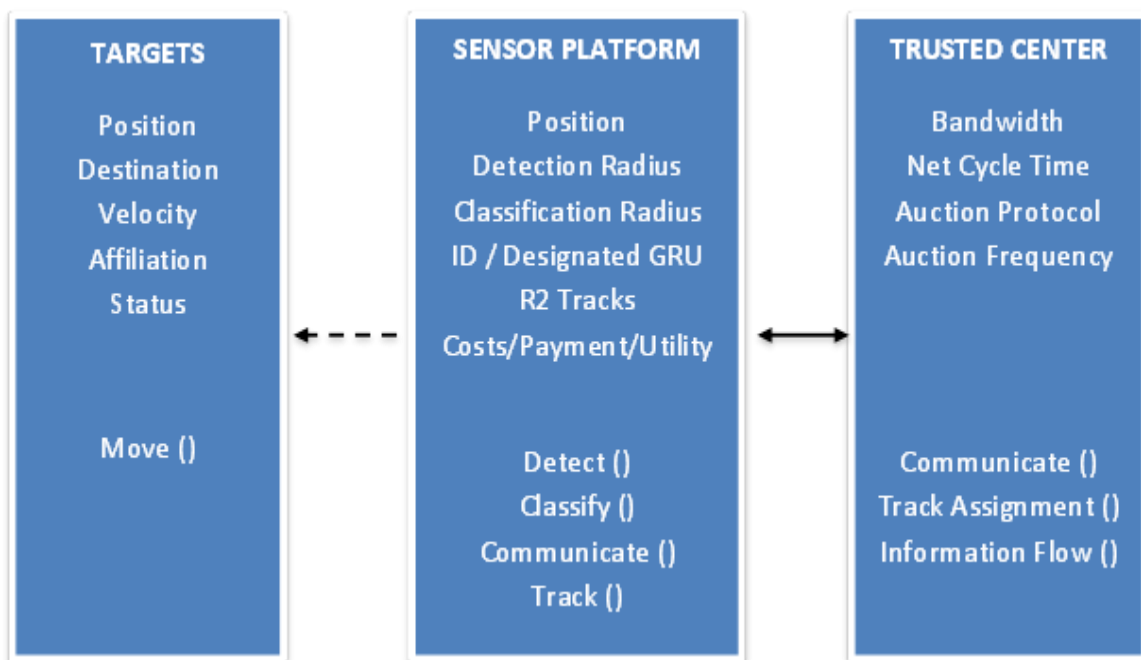
Figure 3-4 Paper Model of the agent-based simulation

3.4.1   Agents

1. Targets:

The targets are a heterogeneous group of autonomous agents that continuously traverse the simulation map.   At the start of simulation the targets are randomly distributed and then move around the map, according to their inherent dynamics. These autonomous agents have different speeds, destinations, sizes as well as affiliations. Not all target tracks are enemy agents: the simulation features an eclectic mix of friendly, neutral and hostile agents. The sensors onboard the military platforms track all the targets, but assigns higher priority to tracking the hostile targets. Keeping in mind the evolving nature of the ULSS environment, current targets may leave the simulation and new targets may enter at any time.

2. Sensor Platforms:

The military platforms are equipped with onboard sensors with a fixed spherical observation and classification sensor envelopes. The platforms can detect any target within the observation radius, with a predefined probability of detection. However at this distance, it is not possible for the platforms to classify the targets. When the targets are within the smaller classification radius, the platforms can categorize the targets as friendly, neutral or hostile entities.

Each sensor platform transmits its own position as well as the track data observed by the onboard sensors over the tactical data link. However, the platforms are doing this from their own frame of reference and not accounting for navigation errors, radar orientation, and so on. Thus there are errors associated with each platform's own position information as well as the track data. We can visualize this as an elliptical region around a target's true position, which is computed by taking into account the covariance associated with the radar, navigation and bearing errors.
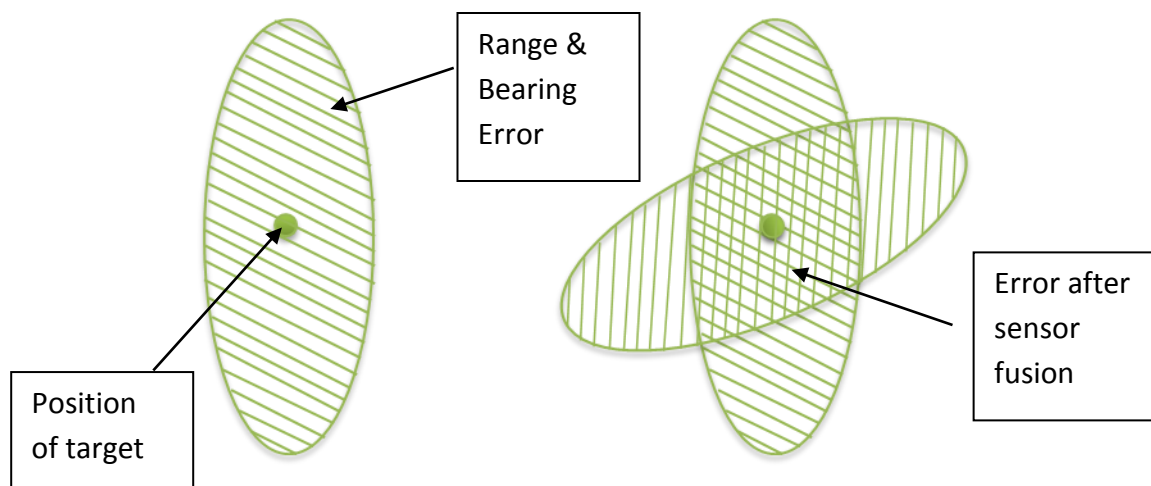
Figure 3-5 Reduction in Uncertainty due to data fusion

Figure 3-5 illustrates the importance of sensor fusion. The left side of the figure indicates the position of the target as estimated by one sensor platform. The elliptical region denotes the range and bearing errors associated with the sensor measurements. When this estimate is fused with the position estimate provided by another sensor platform, the region of uncertainty is reduced (shaded region on the right) and it is possible to predict the target position with higher accuracy. This demonstrates the shortcoming of using Reporting Responsibility (R^2) mechanism, which doesn't allow for redundant reporting of a single object.

One of the sensor platforms in the simulation is assigned the role of the "Grid Reference Unit." Generally speaking, in Navy tactical networks the Aegis Class ships play the role of the GRU, as it possesses the highest quality track data. The GRU's coordinate system is adopted as the truth and the GRU almost universally is assigned the reporting responsibility (R^2) for any track within its classification radius. Since we are operating in an environment with endemic uncertainty, the role of GRU may be played by different platform agents, at different points in the simulation.

3. **Auctioneer:**

The role of the auctioneer or the Network Control Station is assigned to a virtual agent. The auctioneer receives the tactical data over the network from each sensor platform. In the

minimalist R^2 approach, the auctioneer indicates to each sensor platform when it has a transmit opportunity, and the communication proceeds in a round-robin fashion. The time it takes to complete one such round of communication is calculated as the Network Cycle Time (NCT). After each platform transmits its measurements, the auctioneer distributes the reporting responsibility for each target within the observation region, among the platforms in the simulation. The platform selected to provide the report for a track is said to have R^2 for that track. At the next transmit opportunity it is the task of the platform to broadcast its R^2 data.

3.4.2   Communication Network

Link 16 has been the designated DoD primary tactical data link since October 1994 and is assumed as the baseline for our application framework. The official NATO *Standardization Agreement (STANAG) 5516 Edition 3,* defines the specifications for Link 16, and as such is the governing document with respect to Link 16 network management, messages and procedures (STANAG, 1999). A detailed discussion on Link 16 has been included in Appendix B.

Zhao-xiong, Xing, Xue-min, & Jing-lun (2010) have modeled the Link 16 architecture using NS-2, Zhao, Chen, & Lu (2008) simulate the tactical data link in QualNet, Yu, Kuang, Wang, & Liu (2005) use OPNET, while Cruz (2004) use NETWARS to study the performance of Link 16. However these models are computationally expensive and are not suited in our application framework, as our emphasis is on the mechanism design aspect of the problem, and not, on the data links. Hence we use the in-built communication agent within DAF in order to emulate the TADIL. In his thesis, Stinson (2003) captures the performance of Internet Protocol over Link 16 and provides comparison of Link 16 parameters with legacy TADILs. We use his work as the reference for creating a Link 16 communication agent in DAF, with corresponding values for bandwidth, latency and probability of link loss.

3.4.3   Information Valuation Metric

In order to facilitate the selection of sensors with the highest quality target observation, we adopt the information valuation metric as used by Rogers, Dash, Jennings, Reece, & Roberts (2006). In more detail, each sensor has an imprecise estimate of its own global coordinate

position expressed as a joint Gaussian estimate with mean (x, y) and covariance P_0. The center makes imprecise measurements of the range and bearing to multiple targets within its region of interest. For a single target at (x_t, y_t), the sensor makes a noisy measurement of range with mean r and variance σ_r^2, and bearing with mean θ and variance σ_θ^2. We can calculate the total uncertainty in the global coordinate position of the target, in terms of the covariance matrix P given by:

$$P^{-1} = \begin{pmatrix} P_0^{-1} & 0 \\ 0 & 0 \end{pmatrix} + dHR^{-1}dH^T$$

where,

$$R = \begin{pmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_\theta^2 \end{pmatrix} \tag{3.1}$$

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

dH is the Jacobian of the observation model H.

The information content can then be computed as the trace of the inverse of their covariance matrix:

$$I = Tr(dHR^{-1}dH^T) \tag{3.2}$$

The information valuation metric is additive when two independent observations are fused together. Thus given the information content of another sensor's observation of the same target, the total information content of the fused observations can be computed by simply adding the two information valuation measures. The more precise the measurement, the smaller the covariance ellipse, and consequently the greater the information content of the covariance matrix. The representation can now be extended to any given number and distribution of targets, allowing the sensors to value the information content of their observations.

3.4.4   Bandwidth Allocation Using R^2 Rules

At the start of the simulation, the auctioneer indicates to the first sensor platform that it now has the opportunity to transmit its own position and the track data of all the targets in its region of observation. The communication proceeds in a round-robin fashion and each platform gets the opportunity to transmit its own measurements. The time taken to complete one such round of transmission is denoted as Network Cycle Time (NCT). Once all the platforms have finished transmitting, the auctioneer analyzes the received messages. The auctioneer distributes the reporting responsibility for each target among the platforms in the simulation, by taking into account the following parameters:

- The designated GRU platform
- The information content for each target position
- The classification of the targets as hostile or otherwise

Once the auctioneer finishes allocating the targets among the platform, it broadcasts the final allocation to all the sensor platforms. At the next transmit opportunity it is the responsibility of the platforms to transmit the R^2 data of its assigned targets to all the other platforms in the simulation. The Reporting Responsibility rule assumes that once a target has been assigned to a platform, the platform will invest all the resources at its disposal, to track that particular target. However, this assumption may not hold, if we allow for platform agents to act in a selfish and self-interested manner. The R^2 assignment is repeated after a predefined number of cycles. Since the track information doesn't change very often, we only assign targets to the platforms after 10-15 cycles. When any platform loses a track for which it has R^2, an auction is initiated in the next transmission cycle, to assign that track to a different platform.

3.4.5   Auctioning Additional Bandwidth

By its very nature, the Reporting Responsibility (R^2) rule is an extreme minimalist mechanism. By using R^2 rule, the total network cycle time is reduced to its minimum. However it precludes any possibility of collaboration in building a common operating picture by disallowing the redundant reporting of a single object. When a track is not transmitted, we lose the opportunity to fuse its

data. Hence we need a mechanism which allows for the recovery of the highest gain in information for a given quantum of additional bandwidth.

The total additional bandwidth to be allocated can be computed at the start of each auction cycle in terms of Network Cycle Time. For instance, let the NCT for the transmission of R^2 data be 2.3 seconds, and we decide to allocate some additional bandwidth among the sensor platforms to increase the NCT to 2.5 seconds. These additional $2.5 - 2.3 = 0.2$ seconds can be used to transmit additional track data and fuse it with the existing R^2 data. The latency introduced by the increase in Network Cycle Time will be offset by the gain in information content. The mechanism will allocate additional bandwidth corresponding to the increase in NCT to maximize the gain in information.

We need to find which track information to select for transmission given the fixed additional bandwidth available to maximize the total information gain. However, our mechanism needs to go beyond a simple portfolio problem and satisfy the requirements as captured in Table 3-5.

Table 3-5 Mechanism Requirements

Properties	Mechanism Requirements
Individual Rationality	The mechanism should ensure that the platforms voluntarily participate in the mechanism.
Incentive Compatibility	The mechanism has to incentivize the sensor platforms to truthfully reveal their track information.
Interdependency	The mechanism must account for the information interdependency in the reported observations.
Uncertainty	The mechanism needs to take into account that there might be some uncertainty in the reported data.
Implementation	The mechanism has to ensure the platform invests all its resource to track the assigned targets.
Lack of access to outcome	The mechanism should work even when the center has no access to the true state of the world.

Hence as mechanism designers, we need to design a protocol that will ensure that each sensor platform voluntarily participates in the mechanism, truthfully reports their interdependent data, which the auctioneer can implement in an uncertain environment without any access to the true outcome, and the sensor platforms will then honestly execute their responsibilities. We shall introduce our mechanism, which satisfies the six requirements elicited above, over the next two chapters.

CHAPTER 4.   SCORING RULES

In this chapter we discuss the Von Neumann–Morgenstern utility theorem and motivate the use of scoring rules in decision making under uncertainty. We then introduce the concept of strictly proper scoring rules and highlight the various instances of implementation of scoring rules within the Mechanism Design Literature. We discuss and analyze four different instantiations of modified strictly proper scoring rules, viz. Quadratic, Spherical, Logarithmic and Parametric family. We finally introduce our proposed mechanism and discuss its economic properties.

4.1   Decision Making Under Uncertainty

In their seminal work "Theory of Games and Economic Behavior" Neumann & Morgenstern (1944) provided the basis for modern-day game theory. The Von Neumann - Morgenstern utility theory provides the foundation for using utilities to represent preferences.

4.1.1   Von Neumann - Morgenstern Utilities

We focus on the preferences of a player over a set of outcomes denoted by X. For any two outcomes $x_1, x_2 \in X$, we can express the player's preferences over the two outcomes in the form of binary relationships:

- $x_1 \succcurlyeq x_2$: Outcome x_1 is weakly preferred to x_2
- $x_1 \succ x_2$: Outcome x_1 is strictly preferred to x_2
- $x_1 \sim x_2$: Outcome x_1 is equally preferred as x_2

For the set of outcomes $X = \{x_1, x_2, \dots, x_N\}$ we can associate a probability distribution (lottery) $[p_1 : x_1 \; ; \; p_2 : x_2; \; \dots; \; p_N : x_N]$ such that

$$p_j \geq 0 \; for \; j = 1, 2, \ldots, N \; and \; \sum_{j=1}^{N} p_j = 1 \tag{4.1}$$

We are now in a position to highlight the six axioms of Utility Theory as enunciated by Von Neumann and Morgenstern.

1. **Completeness**

   The completeness property induces an ordering on X using the preference relations:

   $$\forall x_1, x_2 \in X, x_1 \succ x_2 \; or \; x_2 \succ x_1 \; or \; x_1 \sim x_2 \tag{4.2}$$

2. **Transitivity**

   The transitivity property assumes that preference is consistent across any three outcomes:

   $$\forall x_1, x_2, x_3 \in X, x_1 \succcurlyeq x_2 \; and \; x_2 \succcurlyeq x_3 \Rightarrow x_1 \succcurlyeq x_3 \tag{4.3}$$

3. **Substitutability**

   The Substitutability property dictates the conditions for substitutability of outcomes over which the player has equal preference. Thus if $x_1 \sim x_2$, the player is indifferent to lotteries $[p : x_1; p_3 : x_3; \ldots; p_N : x_N]$ and $[p : x_2; p_3 : x_3; \ldots; p_N : x_N]$ as long as

   $$p + \sum_{j=3}^{N} p_j = 1 \tag{4.4}$$

4. **Continuity**

   The property of continuity states that the upper and lower contour sets of a preference relation over lotteries is closed.

   $$\forall x_1, x_2, x_3 \in X, x_1 \succ x_2 \; and \; x_2 \succ x_3$$
   $$\Rightarrow \exists \, p \in [0,1] \ni x_2 \sim [p : x_1; 1 - p : x_2] \tag{4.5}$$

5. **Decomposability**

The decomposability property dictates the condition under which the player is indifferent between two or more lotteries. Let σ be a lottery over X and $P_\sigma(x_i)$ denote the probability that x_i is selected by σ. Then according to the decomposability property the player is indifferent between two lotteries σ_1 and σ_2 if,

$$P_{\sigma_1}(x_i) = P_{\sigma_2}(x_i) \; \forall \; x_i \; \in X \Rightarrow \; \sigma_1 \sim \sigma_2 \tag{4.6}$$

6. **Monotonicity**

The property of monotonicity means that a lottery which assigns a higher probability to the player's preferred outcome is preferred to one which assigns a lower probability to its preferred outcome, as long as the other outcomes remain unchanged.

$$\forall \; x_1, x_2 \; \in \; X, x_1 \; > \; x_2 \; and \; 1 \; \geq \; p \; > \; q \; \geq \; 0$$
$$\Rightarrow [p : x_1; \; 1 - p : x_2] \; > \; [q : x_1; \; 1 - q : x_2] \tag{4.7}$$

Von Neumann - Morgenstern Theorem

Given a set of outcomes X and a preference relation $\geqslant$ on X that satisfies completeness, transitivity, substitutability, decomposability, monotonicity and continuity, there exists a utility function $u : X \rightarrow R$ with the following properties:

$$u(x_1) \geq \; u(x_2) \; iff \; x_1 \; > x_2$$
$$u([p_1 : x_1 \; ; \; p_2 : x_2; \; ...; \; p_N : x_N]) = \sum_{j=1}^{N} p_j u(x_j) \tag{4.8}$$

4.1.2   Expected Utility Theory

The expected utility values i.e. the weighted sums obtained by adding the utility values of outcomes multiplied by their respective probabilities, is the key to decision-making. A rational decision maker chooses between risky or uncertain lotteries by comparing their expected utility values. The expected utility theory approach to decision making under uncertainty rests on the assumption that all uncertainty can be expressed in terms of numerical probabilities (De Groot,

1970; Neumann & Morgenstern, 1944; L. Savage, 1972). According to Von Neumann and Morgenstern, these probabilities were "objective" and derived from relative "frequencies in the long run."

"Probability has often been visualized as a subjective concept more or less in the nature of estimation. Since we propose to use it in constructing an individual, numerical estimation of utility, the above view of probability would not serve our purpose. The simplest procedure is, therefore, to insist upon the alternative, perfectly well founded interpretation of probability as frequency in the long run."

- Neumann & Morgenstern (1944)

However, in most real-world contexts it is not possible to objectively assign probabilities to uncertain events due to the lack of data. A formal decision analysis can only be undertaken if the decision maker constructs a subjective probability distribution reflecting his beliefs about the relative likelihoods of the concerned event (De Finetti, 1937; Ramsey, 1931).

"According to the subjective, or personal, interpretation of probability, the probability that a person assigns to a possible outcome or some process represents his own judgment of the likelihood that the outcome will be obtained. This judgment will be based on that person's beliefs and information about the process. Another person, who may have different beliefs or different information, may assign a different probability to the same outcome."

- De Groot (1970)

One of the fundamental concerns with the use of subjective probabilities in decision-making scenarios is the inability to determine if the probability quoted by a person actually corresponds to the person's belief or judgment. *Scoring rules* provide a technique to encourage assessors to declare their subjective probabilities in accordance with their judgments, by rewarding or penalizing the assessor. Scoring Rules involve awarding the assessors with a score based on the assessor's stated probabilities and the actual event that transpires. Thus scoring rules play a twin role in *probability assessment* and *probability evaluation*. In *probability assessment,* scoring rules encourage the assessor to be 'honest', i.e., to make his statements correspond to his

judgments. In *probability evaluation*, scoring rules are used to evaluate the goodness of the probabilities.

4.2 Strictly Proper Scoring Rules

Scoring Rules have been proposed as a methodology to address the shortcomings of auction-based mechanism design in probability assessment and evaluation. Scoring rules were introduced independently by Brier (1950), De Finetti (1962) and Good (1952) for the purpose of expected value maximization. Recently, many researchers have used scoring rules as a viable mechanism to address the challenges of interdependent valuations in systems with self-interested participating actors.

Scoring rules are used to assess the accuracy of probabilistic forecasts, by awarding a score based on the forecast and the event that materializes. Much of the methodology of scoring rules was developed by atmospheric scientists for use for evaluating the accuracy of forecasts (Peterson, Snapper, & Murphy, 1972). Meteorologists frequently need to make observations about a current event or forecasts about events that will happen in the future. These forecasts or observations are probabilistic estimates, which need to be evaluated based on the actual event that transpires. Scoring Rules provide a method to evaluate the probabilistic forecasts, assign a numerical score to the forecasters, and rank competing forecast techniques. This set of tasks is often referred to as *forecast verification.*

Although forecast verification is not directly relevant to our research problem, it provides us with a framework wherein the agents are encouraged to make careful assessments and to be honest (Garthwaite, Kadane, & O'hagan, 2005). The agents are awarded based on the accuracy of their reported observations. The closer the observation value to the actual value, the higher the score assigned to the agent. A rational agent which seeks to maximize its utility is incentivized to invest its resources in making accurate high-quality assessments and reporting them truthfully.

In this section, we discuss the background on scoring rules and the mathematical formulations involved. We talk about how scoring rules have been used in the literature as reputation mechanisms and the four most popular continuous scoring rules.

4.2.1   Background on Strictly Proper Scoring Rules

Gneiting & Raftery (2004) provided a formal definition of scoring rules, which highlight the difference between *proper* and *strictly proper* scoring rules.

Let us consider a case where the forecaster quotes the predictive distribution P and the event x materializes. The forecaster's reward is denoted by the function $S(P, x)$ which can take values in the extended real line $R = (-\infty; \infty)$. Suppose, then, that the forecaster's best judgment is the distributional forecast Q and the expected value of $S(P, \cdot)$ $E_{x \sim Q}\big(S(P, x)\big)$ is written as $S(P, Q)$. As apparent, the forecaster has no incentive to predict any $P \neq Q$, and is encouraged to quote his true belief $P = Q$. If $S(Q, Q) \geq S(P, Q)$ with equality, if and only if $P = Q$, the scoring rule is said to be $strictly\ proper$. If $S(Q, Q) \geq S(P, Q)$ for all P and Q, the scoring rule is said to be $proper$.

To put this definition into perspective, a *strictly proper scoring rule* is the one in which the forecaster can maximize his score by reporting exactly his or her true beliefs about the situation. In the case of a *proper scoring rule,* although the forecaster gets the maximum score by reporting his or her true beliefs, it may be possible to get the same score by reporting something else. In our work, we are interested in strictly proper scoring rules.

L. J. Savage (1971) and Schervish (1989) provided representations that characterize scoring rules for probabilistic forecasts of categorical and binary variables. Let us consider the sample space $\Omega = \{1, 2, \dots m\}$ consisting of m mutually exclusive events, and represent the probabilistic forecast with the vector $(p_1, p_2, \dots, p_m)$. We consider the convex class P

$$P = \{p = (p_1, p_2, \dots, p_m) : p_1, p_2, \dots, p_m \geq 0,$$
$$p_1 + p_2 + \dots + p_m = 1\}$$

(4.9)

We can now define the scoring rule S as a collection of m functions,

$$S(\cdot, i) : P \rightarrow R, \quad i = 1,2, \dots ,m \quad\quad (4.10)$$

Thus, when the forecaster quotes the probability vector p and the event i transpires, the forecaster is awarded a score of $S(p;\ i)$. Before proceeding forward with Savage's Theorem, we need to define a couple of concepts:

1. Regular Scoring Rule:

   A scoring rule S for categorical forecasts is *regular* if $S(\cdot, i)$ is real-valued for $i = 1,2, \dots ,m$, except possibly that $S(p, i) = -\infty$ if $p_i = 0$.

2. Sub-gradient of a Convex function:

   If $G : P \rightarrow R$ is a convex function, the vector $G'(p) = (\ G'_1(p),\ G'_2(p), \dots ,G'_m(p)\ )$ is called the sub-gradient of G at the point $p \in P$ if

   $$G(q) \geq G(p) + <G'(p), q - p> \quad\quad (4.11)$$

   for all $q \in P$, where $<\cdot,\cdot>$ denotes the standard scalar product. We assume that the components of $G'(p)$ are real-valued, except that we permit $G'_i(p) = -\infty$ if $p_i = 0$.
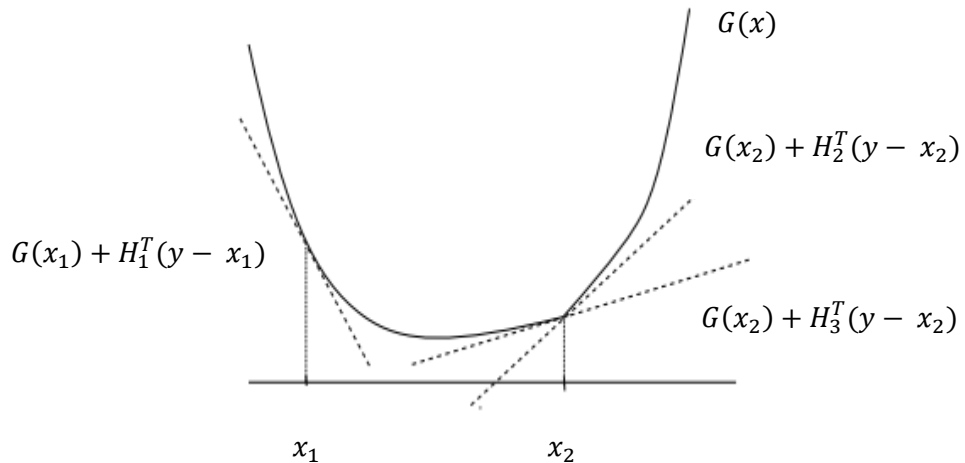
Figure 4-1 Sub-gradient of a Convex function at two distinct points

At the point x_1, the convex function $G(x)$ is differentiable and H_1 (which is the derivative of G at x_1) is the unique sub-gradient at x_1. At the point x_2, $G(x)$ is not differentiable. At x_2, the function $G(x)$ has many sub-gradients: two sub-gradients H_2 and H_3 are shown in Figure 4-1.

Savage's Theorem:

A regular scoring rule S for categorical forecasts is strictly proper if and only if

$$S(p, i) = G(p) - < G'(p), p > + G'_i(p) \tag{4.12}$$

for $i = 1, 2, \ldots, m$ where $G : P_m \to R$ is a (strictly) convex function and $G'(p)$ is a sub-gradient of G at the point p, for all $p \in P_m$.

Rephrasing this, a regular scoring rule S is strictly proper if and only if the expected score function $G(p) = S(p, p)$ is strictly convex on P_m and the vector with components $S(p, i)$ for $i = 1, 2, \ldots, m$ is a sub-gradient of G at the point p, for all $p \in P_m$.

The classic case of *yes or no* forecast is more revealing. We restrict the sample space to $\Omega = \{0,1\}$ and the probability forecast is $p \in [0,1]$ for *yes*. The scoring rule can then be reduced to a pair of functions

$$S(\cdot, 1) : [0,1] \to R \text{ and } S(\cdot, 0) : [0,1] \to R \tag{4.13}$$

Thus $S(p, 1)$ represents the score if the forecaster assigns the probability p to an event which actually transpires, and $S(p, 0)$ is the score if the forecaster quotes probability p and the event does not materialize. According to the theorem, every strictly proper scoring rule can be represented in the form:

$$\begin{aligned} S(p, 1) &= G(p) + (1 - p)G'(p) \\ S(p, 0) &= G(p) - pG'(p) \end{aligned} \tag{4.14}$$

where $G : [0,1] \rightarrow R$ is a strictly convex function and $G'(p)$ is a sub-gradient of G at the point $p \in [0,1]$. The sub-gradient $G'(p)$ is real-valued, except that we permit $G'(0) = -\infty$ and $G'(1) = -\infty$ For the case when G is differentiable, then $G'(p)$ is unique and equals the derivative of G at p.

We encapsulate the three most popular scoring rules (Quadratic, Spherical and Logarithmic) for both the categorical case and binary case in Table 4-1.

Table 4-1 Popular Scoring Rules

Scoring Rule	Categorical		Binary		
	$G(p)$	$S(p,i)$	$G(p)$	$S(p,1)$	$S(p,0)$
Quadratic	$\sum_{j=1}^{m} p_j^2 - 1$	$-\sum_{j=1}^{m} p_j^2 - 1 + 2p_i$	$2p^2 - 2p$	$-2p^2 + 4p - 2$	$-2p^2$
Spherical	$\left(\sum_{j=1}^{m} p_j^\alpha\right)^{\frac{1}{\alpha}}$	$\dfrac{p_i^{\alpha-1}}{\left(\sum_{j=1}^{m} p_j^\alpha\right)^{\frac{\alpha-1}{\alpha}}}$	$(p - p^2)^{1/2}$	$\dfrac{-(1-p)^{\frac{1}{2}}}{2(p^{1/2})}$	$\dfrac{-2(p^{1/2})}{(1-p)^{\frac{1}{2}}}$
Logarithmic	$\sum_{j=1}^{m} p_j \ln p_j$	$\ln p_i$	$p \ln p + (1-p)\ln(1-p)$	$\ln p$	$\ln(1-p)$

For our research problem, the use of binary or categorical strictly proper scoring rules is not applicable as the distributed information is represented by continuous distributions. Hence, we derive the continuous counterpart for the binary quadratic, spherical and logarithmic scoring rules in Table 4-2. We replace the discrete probability value p with the probability density function $r(x)$ of the continuous random variable x.

Table 4-2 Scoring Rules for continuous variables

	Quadratic	Spherical	Logarithmic
Scoring Rule $S(r(x))$	$2r(x) - \displaystyle\int_{-\infty}^{\infty} r(x)^2 dx$	$\dfrac{r(x)}{(\int_{-\infty}^{\infty} r(x)^2 dx)^{1/2}}$	$\ln r(x)$

The Quadratic or the Brier rule, which was rigorously described by Selten (1998), is a special case of a parametric scoring rule. The parametric scoring rule, which is included within the power rule family, can be denoted by,

$$S = k(r)^{k-1} - (k-1) \sum_{j=1}^{m} r^k - 1 \qquad (4.15)$$

where k is a real number such that $k > 1$. For $k = 2$, the parametric scoring rule becomes the quadratic rule.

We shall restrict our discussion to these four scoring rules – quadratic, spherical, logarithmic and parametric – as we can analytically derive and express their expected values, in closed forms in the latter half of this chapter. In Chapter 6 we shall discuss the results of applying the scoring rules to the research application framework.

4.2.2   Applications of Strictly Proper Scoring Rules

In the last four decades, strictly proper scoring rules have found application in quite a few fields:

- Accounting (Wright, 1988)
- Bayesian statistics (L. J. Savage, 1971)
- Business (Holstein, 1972)
- Climate Prediction (Gneiting & Raftery, 2005; Palmer, 2002)
- Computer (Miller, Pratt, Zeckhauser, & Johnson, 2007)
- Education (Echternacht, 1972)
- Finance (Shiller, Kon-Ya, & Tsutsui, 1996)
- Macroeconomic forecasting (Garratt, Lee, Pesaran, & Shin, 2003)
- Medicine (Spiegelhalter, 1986)
- Politics (Tetlock, 2005)
- Psychology (McClelland & Bolger, 1994)
- Stochastic finance (Duffie & Pan, 1997)
- Others (Church et al., 2006; Hanson, 2002; Prelec, 2004)

Miller et al. (2007) have proposed the use of strictly proper scoring rules in order to get agents to truthfully report a probabilistic estimate and then commit costly resources into generating the observations. In their scoring rules – based mechanism, agents provided a rating for a product or service, and their score is calculated based on how close its rating is to the expected ratings provided by other agents in the system. In this mechanism, each agent will seek to honestly report their observations to maximize their expected utility, provided the other agents are also truthful. This makes it a Nash equilibrium solution. Jurca & Faltings (2005) have also used continuous scoring rules in their reputation mechanism, in which they divide the agents into pairs of two and have them rate each other. However in both these mechanisms, truthful reporting by the agents is not a unique equilibrium strategy and other Nash equilibria might exist which do not have the desired properties. Although computing all the Nash equilibria in any mechanism is notoriously difficult, we can still say that there exist multiple Nash equilibria in which the agents can collude to provide dishonest feedback. In their later work, Jurca & Faltings (2007) introduced a small group of agents which will always be truthful and thus prevent collusion among agents. All the agents' reports are rated against one of these trusted agents' reports and the payments are designed such that an agent maximizes its utility by reporting honestly. This eliminates the undesired Nash equilibria and truthful reporting becomes a unique equilibrium. However their solution still leaves much to desire, as it compromises the network robustness.

Proper scoring rules have also been used to address the principal-agent problem (Grossman & Hart, 1983; Rogerson, 1985) which features a contractor interested in purchasing information and a contractee who is selected to supply that information. Zohar & Rosenschein (2008) proposed two different mechanisms to address this problem, but they base it on the unrealistic assumption that there is a common probability distribution among agents, and that this distribution expresses 'close' notions of the governing probability distributions. Another common drawback of all the schemes we have discussed so far is that they all assume that the cost of generating information is public information. But in our problem, the cost is a function of the accuracy of the information, and is a private value known only to each individual agent. Hence none of the discussed approaches are applicable to our work.

Despite the various shortcomings, the proper scoring rules still address some of the issues we encountered with auction-based mechanism.

- Truth-elicitation is induced among agents, even without access to information about the realized outcome. This is achieved by utilizing the reports submitted by other agents, to compute the payment to a particular agent.

- Agents are no longer restricted to report only discrete probability distributions, like the probability of success or failure and can now even submit continuous probability distributions. Hence the lack of consistency and interoperability of observations can be addressed by modeling the reported observations as continuous distributions.

4.2.3 Application of Strictly Proper Scoring Rules in Mechanism Design

We demonstrate the features of strictly proper scoring rules through a practical application similar to our desired implementation scenario. We follow the methodology as laid out by Miller et al. (2007) wherein we relax the assumption that the costs involved in generating an observation is a private value known only to each individual agent. Instead, we assume that these costs are known to the auctioneer.

As we discussed earlier, one of the highlights of the strictly proper scoring rules, is that we can model the agent's reported observations as continuous distributions. This removes the ambiguity in the reported observations by applying standard distributions like the Gaussian distribution. Hence we model the agent's noisy private measurement, x, as Gaussian random variable,

$$x \sim N\left(x_0, {1}/{\theta}\right) \tag{4.16}$$

where, x_0 is the true state of the observable and θ is the precision of the observation.

One of the drawbacks of the auction-based MD was that it did not account for agents not investing all their available resources in generating the observations. Miller et al. combat this

issue through the introduction of scaling parameters. They show that the affine transformation of the scoring rules, does not affect the inherent properties of the scoring rules, like, incentive compatibility. Since the costs are known to the auctioneer, the scaling parameters can be selected to incentivize an agent to generate and then report the observations.

If we denote the scoring rule by the function $S(x_0; x, \theta)$ and the expected score as $\bar{S}(\theta)$ then we can formulate the expected payment as

$$\bar{P}(\theta) = \alpha \bar{S}(\theta) + \beta \tag{4.17}$$

where α and β are the scaling parameters.

Thus the expected utility of the agent can be calculated as

$$\bar{U}(\theta) = \alpha \bar{S}(\theta) + \beta - c(\theta) \tag{4.18}$$

where $c(\theta)$ is the cost of generating an observation with precision θ.

We can now select the parameter α to maximize the agent's expected utility when it generates and reports truthfully its observation with precision θ_0.

$$\frac{d\bar{U}}{d\theta}|\theta_0 = 0 \ \Rightarrow \alpha = \frac{c'(\theta)}{\bar{S}'(\theta)} \tag{4.19}$$

The next desired property is individual rationality, i.e. ensuring that the agent will always derive a non-negative utility by participating in the mechanism. Thus we select the remaining parameter β in a way which ensures that the agent will be willing to incur the cost of producing a forecast, since the expected utility is always positive. Since the auctioneer knows the costs involved in generating an observation, we can make the agents indifferent between generating an observation or not, by equating

$$\bar{U}(\theta_0) = 0 \ \Rightarrow \beta = c(\theta_0) - \frac{c'(\theta_0)}{\bar{S}'(\theta_0)} \bar{S}(\theta_0) \tag{4.20}$$

We can now compare the four different scoring rules – quadratic, spherical, logarithmic and parametric - by replacing the general probability density functions, with Gaussian distributions. We also calculate the expected values along with the parameter expressions for the strictly proper scoring rules.

<p style="text-align:center">Table 4-3 Scoring Rules for Gaussian distributions</p>

	Quadratic	Spherical	Logarithmic	Parametric
$S(x_0; x, \theta)$	$2N - \left(\dfrac{\theta}{4\pi}\right)^{1/2}$	$\left(\dfrac{4\pi}{\theta}\right)^{1/4} N$	$\log N$	$kN^{(k-1)} - \dfrac{k-1}{\sqrt{k}}\left(\dfrac{2\pi}{\theta}\right)^{(1-k)/2}$
$\bar{S}(\theta)$	$\left(\dfrac{\theta}{4\pi}\right)^{1/2}$	$\left(\dfrac{\theta}{4\pi}\right)^{1/4}$	$\dfrac{1}{2}\log\left(\dfrac{\theta}{2\pi}\right) - \dfrac{1}{2}$	$\dfrac{1}{\sqrt{k}}\left(\dfrac{2\pi}{\theta}\right)^{(1-k)/2}$
$\bar{S}'(\theta)$	$\dfrac{1}{4\sqrt{\theta\pi}}$	$\dfrac{1}{4}\left(\dfrac{1}{4\pi\theta^3}\right)^{1/4}$	$\dfrac{1}{2\theta}$	$\dfrac{k-1}{2\theta\sqrt{k}}\left(\dfrac{2\pi}{\theta}\right)^{(1-k)/2}$
α	$4c'(\theta_0)\sqrt{\theta\pi}$	$4c'(\theta_0)(4\pi\theta^3)^{\frac{1}{4}}$	$2c'(\theta_0)\theta_0$	$\dfrac{2c'(\theta_0)\theta_0\sqrt{k}}{k-1}\left(\dfrac{\theta_0}{2\pi}\right)^{\frac{1-k}{2}}$
β	$c(\theta_0) - 2\theta_0 c'(\theta_0)$	$c(\theta_0) - 4\theta_0 c'(\theta_0)$	$\dfrac{c(\theta_0) - 2\theta_0 c'(\theta_0)}{\left(\dfrac{1}{2}\log\left(\dfrac{\theta_0}{2\pi}\right) - \dfrac{1}{2}\right)}$	$c(\theta_0) - \dfrac{2\theta_0 c'(\theta_0)}{k-1}$

where N represents the Gaussian distribution $N\left(x_0; x, \frac{1}{\theta}\right)$

An important property of strictly proper scoring rules is the concavity of the expected scoring rule function. This concavity will form the basis of the proofs we present in the next section. However, this property does not hold for any or all continuous strictly proper scoring rules, as we can deduce for the parametric scoring family. For values of $k > 3$ the second derivative of $\bar{S}(\theta)$ is positive and hence the expected score function becomes convex.

$$\bar{S}''(\theta) = \frac{(k-1)(k-3)}{4\sqrt{k}\theta^2}\left(\frac{2\pi}{\theta}\right)^{(1-k)/2} \tag{4.21}$$

A convex expected scoring rule does not incentivize an agent to produce the observation at the given precision θ_0. Hence the parameter k is restricted to the space $(1,3)$.

4.3   Proposed Mechanism

Having established the background of scaled strictly proper scoring rules, we can now use the methodology as outlined by Miller et al.. We first enumerate the deficiencies of the scoring rules discussed so far, and the need for a different mechanism to address these deficiencies.

1.   Unknown Costs

In the previous section, we stated that the common drawback of all the continuous strictly proper scoring rules we discussed, is that they all assume that the cost of generating information is public information. We ourselves assumed that these costs are known to the auctioneer in order to derive the analytical expressions for quadratic, spherical, logarithmic and parametric rules.

However, in our research problem, the costs of generation of an observation, represents the private information known only to each individual agent. These costs represent the amount of resources and time the sensors invest in generating their information as well as the technological capacity of the sensors. Hence we need to develop a mechanism which can deal with the lack of knowledge of the costs involved.

2.   Multiple sources of information

The reporting responsibility (R^2) rules permit only one agent with the best quality observation to report a surveillance track on the data link. However, it is plausible that there might be instances when no one agent can generate track data of the required precision value. In our research framework we calculate the precision values using the information valuation metric discussed in Section 3.4.3. This might occur due to the lack of resources available to an agent to generate the observation for a given track. In such settings, it is prudent to fuse information from multiple sources in order to generate real-world data of sufficient accuracy. Thus there is a need for the

mechanism to be able to operate in an environment where the agents have limitations in the values of the information content they can provide.

3. Uncertainty in environment

The R^2 rules assume that once a target has been assigned to a platform, the platform will invest all the resources at its disposal, to track that particular target. The strictly proper scoring rules methodology addresses this issue by awarding the agents based on how close their observation value is to the actual value that is observed once the event transpires. Hence, the closer the observation value to the actual value, the higher the score assigned to the agent. But ULS systems operate in dynamic and uncertain environment which constantly evolves between the time the information is reported, and the time when the observation can be observed. Hence, we need a mechanism that accounts for such uncertainties where the center is unable to evaluate the received reports.

With these requirements in mind we adopt a modified version of the two-stage mechanism proposed by Papakonstantinou, Rogers, Gerding, & Jennings (2011) based on continuous strictly proper scoring rules. In the first stage, the trusted center (auctioneer) elicits the unknown costs of the agents and preselects a subset of agents that can provide the information at the lowest costs. In the next stage the preselected agents are induced to reveal their observation, using a payment scheme based on the fused reported estimates rather than the true outcome. Through appropriately scaled and modified strictly proper scoring rules, the mechanism induces the agents to deploy the resources at their disposal into generating the data at the required information content level and then reporting the data honestly.

4.3.1   Setting up the Mechanism

We set up the background of our model against which we describe the two-stage mechanism. First, we discuss the problem of eliciting information from multiple sources and then evaluating that information, without knowledge of the outcome. We show how the modified strictly proper scoring rules, handle the lack of knowledge of the outcome while preserving the property of incentive compatibility. We shall also describe the assumptions of the mechanism that are critical to its application to our model.

4.3.1.1   MAS Model

We consider a scenario with N rational agents that are capable of producing a noisy and inaccurate observation x_i with a precision θ_i. There is a limit on the maximum precision that an agent can provide, represented as θ_i^c. Thus the precision of the agent's observation can lie anywhere between 0 and θ_i^c. The trusted center collects and fuses these observations, to obtain a required precision denoted by θ_0. There is no additional utility derived if the precision exceeds θ_0. The agent's private observations are modeled as Gaussian variable $x_i \sim N\left(x_0, \frac{1}{\theta_i}\right)$ where x_0 represents the true state of the world. This value is unknown to the agents and the trusted center.

We use the methodology proposed by DeGroot & Schervish (2002) to fuse the information from two or more sources. If there are k unbiased and conditional independent estimates $\{x_1, x_2, \dots, x_k\}$ with precision values $\{\theta_1, \theta_2, \dots, \theta_k\}$ they can be fused into a single estimate $\left(\bar{x}, \frac{1}{\bar{\theta}}\right)$ where $\bar{x}$ is the mean and $\bar{\theta}$ is the precision.

$$\bar{x} = \frac{\sum_{i=1}^{k} x_i \theta_i}{\sum_{i=1}^{k} \theta_i} \quad \text{and} \quad \bar{\theta} = \sum_{i=1}^{k} \theta_i \tag{4.22}$$

The fusion of the sensor information ensures that the fused precision is higher than any of the individual sensor's precision $\bar{\theta} \geq \theta_i$. However care must be taken to ensure that the agents report their observations truthfully, as fusion of misreported estimates with truthful estimates will, in fact, degrade the quality of the information of the state of the world. Hence, the agents must be given incentives to honestly report their observations, even when the center has no access to the true state of the world to evaluate the reported observations. As we specified earlier, the traditional strictly proper scoring rules are unequipped to cope with the center's lack of knowledge, and hence we introduce the modified scoring rules proposed by Papakonstantinou, Rogers, Gerding, & Jennings (2011).

4.3.1.2   Modified Scoring Rules

In the modified strictly proper scoring rules, the trusted center fuses the observations from all the other agents and excludes the agent whose reported observation is being evaluated. Thus using the same k unbiased and conditional independent estimates we introduced earlier, we fuse k_{-i} probabilistic estimates $\{x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_k\}$ with precision values $\{\theta_1, \theta_2, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_k\}$ into a single estimate $\left(\bar{x}_{-i}, 1/\bar{\theta}_{-i}\right)$ where $\bar{x}_{-i}$ is the mean and $\bar{\theta}_{-i}$ is the precision.

$$\bar{x}_{-i} = \frac{\sum_{j=1}^{k_{-i}} x_i \theta_i}{\sum_{j=1}^{k} \theta_i} \quad \text{and} \quad \bar{\theta}_{-i} = \sum_{j=1}^{k_{-i}} \theta_i \tag{4.23}$$

An agent seeking to maximize its utility will need to consider, its belief about the observations of all the other agents in the system, when reporting its own observation. The expected scoring rule for an agent i which is denoted by $\bar{S}(\bar{x}_{-i}; x_i, \theta_i)$ is not maximized at $\hat{\theta}_i = \theta_i$ but rather at $\hat{\theta}_i = \theta_i + \bar{\theta}_{-i}$. We can denote this mathematically as

$$\bar{S}\left(\bar{x}_{-i}; \hat{x}_i, \hat{\theta}_i\right) =$$
$$\int_{-\infty}^{\infty} N\left(\bar{x}_{-i}; x_i, 1/\theta_i + 1/\bar{\theta}_{-i}\right) S\left(\bar{x}_{-i} \mid N\left(\bar{x}_i; \hat{x}_i, 1/\hat{\theta}_i\right)\right) d\bar{x}_{-i} \tag{4.24}$$

where $N\left(\bar{x}_{-i}; x_i, 1/\theta_i + 1/\bar{\theta}_{-i}\right)$ represents the distribution of true estimate by agent i and $N\left(\bar{x}_i; \hat{x}_i, 1/\hat{\theta}_i\right)$ represents the distribution of reported estimate of agent i. The expected score of agent i is maximized at $\hat{\theta}_i = \theta_i + \bar{\theta}_{-i}$, since at this value of the reported precision, the two Gaussian distributions are identical.

We have stated that an agent i's utility is maximized when it reports its precision $\hat{\theta}_i = \theta_i + \bar{\theta}_{-i}$. However, the precision of the observations generated by the other agents is a private value, and is not accessible to agent i. But the trusted center has access to both θ_i and $\bar{\theta}_{-i}$ and hence the center can modify the scoring rule such that the agent i only reports θ_i but its payment is

calculated using θ_i and $\bar{\theta}_{-i}$ .Thus we introduce the concept of modified scoring rule $S\left(\bar{x}_{-i}; \widehat{x}_\iota, \widehat{\theta}_\iota + \bar{\theta}_{-i}\right)$ whose expected value can be expressed as

$$\bar{S}\left(\bar{x}_{-i}; \widehat{x}_\iota, \widehat{\theta}_\iota + \bar{\theta}_{-i}\right) = \int_{-\infty}^{\infty} N\left(\bar{x}_{-i}; x_i, \frac{1}{\theta_i} + \frac{1}{\bar{\theta}_{-i}}\right)$$
$$S\left(\bar{x}_{-i} \left| N\left(\bar{x}_i; \widehat{x}_\iota, \frac{1}{\widehat{\theta}_\iota} + \frac{1}{+\bar{\theta}_{-i}}\right)\right.\right) d\bar{x}_{-i}$$

(4.25)

We can easily show how the modified scoring rules induce truthful elicitation as a Nash equilibrium solution. We recall the definition of a strictly proper scoring rule, which states that for a scoring rule function $S(Q, R)$ its expected value is maximized when $Q = R$. In our expected value formulation, Q and R represent the reported and true Gaussian distribution of the probability estimates. Thus,

$$N\left(\bar{x}_i; x_i, \frac{1}{\theta_i} + \frac{1}{\bar{\theta}_{-i}}\right) \equiv N\left(\bar{x}_i; \widehat{x}_\iota, \frac{1}{\widehat{\theta}_\iota} + \frac{1}{+\bar{\theta}_{-i}}\right)$$
$$\Rightarrow \widehat{x}_\iota = x_i \quad and \quad \widehat{\theta}_\iota = \theta_i$$

(4.26)

Thus, based on our modified strictly proper scoring rule, an agent can maximize its expected score and by extension, its expected payment by truthfully its observations, assuming that the other agents in the system also honestly report their observations. This makes truthful revelation a Nash equilibrium and the optimal strategy for all the agents in the system.

4.3.1.3   Assumptions

The cost of generating an observation represents the private information known only to each individual agent, and is denoted as $c(\theta)$. The cost is a function of the precision of the observations. We make two critical assumptions regarding the cost function:

1. The cost function is convex, i.e. $c''(\theta) \geq 0$. This is a realistic assumption in all instances, where increasing precisions lead to diminishing returns. We assume the cost functions

to be linear which guarantees convexity. This mechanism will be equally valid for non-linear cost functions as well, provided we ensure that the cost functions are convex.

2. Although each agent can have different cost functions, we assume that the different cost functions, and the derivatives thereof, do not cross. What this indicates, is that the ordering of the cost functions and their derivates, is preserved over all precisions.

These assumptions are critical to proving the individual rationality as well as incentive compatibility of the mechanisms. Papakonstantinou, Rogers, Gerding, & Jennings (2008) have shown that no mechanism based on strictly proper scoring rules can achieve these two economic properties without these assumptions being satisfied for the private cost functions.

4.3.2   Proposed Mechanism

Jehiel & Moldovanu (2001) showed that in an interdependent valuation setting, no standard one-stage mechanism could achieve both efficiency and incentive compatibility for the procurement of estimates from multiple sources. Mezzetti (2004) addressed this challenge to a certain extent and showed that an efficient allocation with multidimensional types is possible, if (a) values are privately realized by the agents once an allocation is made and (b) two-stage mechanisms can be adopted in which the payments are made contingent on realized values reported in a second stage.

In accordance with Mezzetti (2004), we design a two-stage mechanism based on modified scaled strictly proper scoring rules. In the first stage, the center preselects M of the N available agents based on the reported cost functions and then identifies a subset of the M preselected agents to generate the observations. In the second stage, the center announces the payment scheme based on scoring rules, which incentivize the agents to truthfully generate and report their observations to the center.

In the first step (Step 1.1) of the first stage of the mechanism, the center requests all agents in the system to report their cost functions and to reveal its private maximum information content. In practice, the center only needs the cost function and the derivatives of the cost function at

the reported maximum information content value. However for the sake of convenience, the agents are asked to reveal their entire cost functions. In the next step (Step 1.2) the center preselects M agents from the N available agents through one single reverse $(M + 1)^{th}$ auction (i.e. an auction where the highest M bidders win and pay a uniform price determined by the $(M + 1)^{th}$ price). Ideally the center should divide all the available agents into groups of n agents, and then initiate $(m + 1)^{th}$ multiple price auctions ($1 \leq m < n$) to pre-select M agents. However, we make use of the result from Papakonstantinou, Rogers, Gerding, & Jennings (2010) which states that for linear cost function settings, the center can minimize its expected payment by setting $n = N$ and $m = M$.

In the second stage, each of the agents selected in the first stage are asked to produce their observations at their reported maximum information content levels. The center announces the modified strictly proper scoring rule with parameters α_j and β_j. For an agent j, these parameters are formulated based on the fused reported information content of every agents apart from j and the cost associated with the single reverse $(M + 1)^{th}$ auction, from the first stage. In Step 2.2, each of the selected agents produce and report their observations to the center, which in turn, calculates the payments based on the modified strictly proper scoring rule $S_j\left(\bar{\bar{x}}_{-j};\ \hat{x}_j, \hat{\theta}_j + \bar{\theta}_{-j}\right)$.

Having outlined the mechanism, in detail, we now proceed to state the properties of the mechanism which induce the agents to reveal their cost functions and their maximum information content values honestly and subsequently produce and truthfully report their observations to the center.

4.3.3   Economic Properties of the Mechanism

We discuss the economic properties of the modified strictly proper scoring rules based mechanism design algorithm. The interested reader is directed to Papakonstantinou et al. (2008, 2010, 2011) for the proofs of the same.

1. **First Stage**

   1.1. The trusted center asks $N \geq 2$ agents to report their cost functions $\hat{c}_i(\theta)$ and their maximum information content $\hat{\theta}_i^c$, for all agents $i \in \{1,2,\dots,N\}$

   1.2. The center selects M $(1 \leq M < N)$ agents with the lowest costs, associates these agents with the $(M+1)^{th}$ cost and discards the rest of the agents.

2. **Second Stage**

   2.1. The center asks agent j, selected in Step 1.2, to generate the observations and presents it with a modified strictly proper scoring rule with parameters α_j and β_j

   $$\alpha_j = \frac{c'_{M+1}(\hat{\theta}_j^c)}{\bar{S}'(\hat{\theta}_j^c, \bar{\theta}_{-j})}$$

   $$\beta_j = c_{M+1}(\hat{\theta}_j^c) - \frac{c'_{M+1}(\hat{\theta}_j^c)}{\bar{S}'(\hat{\theta}_j^c, \bar{\theta}_{-j})} \bar{S}(\hat{\theta}_j^c, \bar{\theta}_{-j})$$

   (4.27)

   where c_{M+1} is the $(M+1)^{th}$ cost identified in stage 1.2 and $\bar{\theta}_{-j}$ is the fused information content values of all agents that are asked to generate observations, except agent j.

   2.2. Each of these agents will produce an estimate x_j with information content θ_j and report $(\hat{x}_j, \hat{\theta}_j)$ to the center which issues the payment

   $$P_j(\bar{x}_{-j}; \hat{x}_j, \hat{\theta}_j + \bar{\theta}_{-j}) = \alpha_j S_j(\bar{x}_{-j}; \hat{x}_j, \hat{\theta}_j + \bar{\theta}_{-j}) + \beta_j$$

   (4.28)

1. **The mechanism is incentive compatible with respect to the agents' reported costs.**

Truthful reporting of the cost function of an agent in the first stage is a weakly dominant strategy. This is valid in both cases when an agent's misreporting effects or does not effects whether it is preselected or not. If the agent was pre-selected by misreporting but would not have been if it was truthful, we can show that the expected utility is strictly negative. On the other hand, if the agent was not pre-selected by misreporting but would have been if it was truthful, the agent's utility becomes zero and hence it has no incentive to misreport its cost function.

2. **The mechanism is interim individually rational.**

The mechanism is interim individually rational. What this means is that the expected utility is non-negative, but there might be instances when there is a wide-discrepancy between the agent's actual reported precision and the maximum estimated precision, which results in negative payment. Technically we could make the mechanism ex-post individually rational by setting a very high value for β but this would violate the incentive compatibility property.

3. **Truthful reporting of the maximum precisions and estimates is a Nash equilibrium, with the reported maximum precision being the actual precision of the estimate.**

We proved in Section 4.2.1 how for the modified strictly proper scoring rules, truthful revelation is a Nash equilibrium and the optimal strategy for all the agents in the system. A preselected agent's utility is maximized when it reports the actual precision of its produced estimate as its maximum precision, given that all other agents do the same. Thus an agent will truthfully report its maximum precision, and then produce an estimate of precision equal to its reported precision, and report the estimated precision truthfully.

4. **The probability of achieving the required precision by the center increases as M increases.**

The number of pre-selected agents M is a variable and can be adjusted by the designer. There exists a trade-off between the expected payment made by the center, and the probability of achieving the required precision. We can show that for any distributions of the reported precisions, the probability of achieving the required precision increases, as the number of pre-selected agents increase.

5. **There can be no incentive compatible mechanism regarding the agents' cost functions revealed when the cost functions overlap.**

This property follows from the convexity of the cost functions, which implies that if the true costs of an agent making the prediction is higher than the costs used for scaling the scoring function, then the agent's utility will always be negative. Thus when the cost functions overlap, an agent will be able to do better by misreporting and losing, rather than by truthfully reporting and winning. Hence in order to preserve the incentive compatibility, we need the cost functions to not cross each other at any point.

6. **In a setting with linear cost functions, for a given probability of achieving θ_0, the center minimizes its expected total payment when $n = N$ and $m = M$**

This result is valid for the cases when we assume that the cost functions of the agents are linear functions of the precision. Since we associate the $(m + 1)^{th}$ cost with the group of m agents selected, we can show that the expected value of the $(m + 1)^{th}$ cost is lowest in expectations when $n = N$ and $m = M$. This, by extension, ensures that the total expected payment made by the center is minimized. However, this result only holds in the case of linear cost functions. If we have non-linear cost function, we will need to iteratively preselect the agents. This is done by first dividing the N agents into groups of n, and then conducting multiple reverse $(m + 1)^{th}$ price auctions to preselect the M agents.

7. **The sum of the expected agents' payments is independent of the knowledge of the actual outcome**

This property compares the mechanism with the unknown outcome to the one where the actual outcome is known, and concludes that in both the cases the total expected payment made by the center is the same. Hence, the lack of knowledge of the actual outcome for the mechanisms has no impact on the expected payments.

4.4   Next Steps

In this chapter, we proposed a mechanism based on the modified scaled strictly proper scoring rules which overcomes the shortcomings of the auction-based mechanism design and addresses five out of the six requirements we highlighted in Chapter 3. Specifically the mechanism is individually rational and incentive compatible and works in uncertain and dynamic environment, where the costs of generating observations are private values unknown to the center and the true outcome cannot be observed by either the agents or the center. The mechanism incentivizes the agents to truthfully reveal their interdependent track information and the center then fuses different observations to increase the accuracy of the common operating picture. The strictly proper scoring rules ensure that an agent's payment is dependent on the accuracy of its reported observations, and hence guarantees that the agent will invest all its resources in generating the observations.

The proposed mechanism is instrumental in selecting the agents to provide observations for a single target, above and beyond, those reported by R^2. However, in our multi-agent system there are numerous targets but only a limited bandwidth to transmit the additional information on the target positions. Thus we need a methodology to decide which targets to select in order to ensure that we obtain the highest gain in information for a given quantum of additional bandwidth, given the inherent uncertainty in the reported data. In the next chapter we develop and propose a robust optimization technique, which can optimize the track selection, in uncertain and evolving environments.

CHAPTER 5.   ROBUST OPTIMIZATION APPROACH

In this chapter, we provide a robust optimization formulation to decide which sensor-target pair should be selected for transmission, given the inherent uncertainty in the reported data. We discuss the various robust techniques in literature and present the highlights of the Bertsimas - Sim linear optimization framework which we adopt for our problem. Finally, we outline the final steps of our modified scaled strictly proper scoring rules mechanism, which is formulated as a robust zero-one portfolio optimization problem.

5.1   Introduction

The modified scaled strictly proper scoring rules mechanism selects a set of agents to provide observations for a single target, above and beyond, those reported by R^2. Since there are numerous targets in the system, we end up with different sets of sensor agents for each target. However we can only allocate a limited bandwidth for transmitting information over the tactical data network. Hence, we need a methodology to decide which sensor-target pair should be selected for transmission to ensure that we can obtain the highest gain in information for a given quantum of additional bandwidth. The problem is compounded by the inherent uncertainty in the information content of the observations. Deterministic optimization techniques that rely on nominal data, no longer work in these settings. Robust techniques provide an attractive choice in addressing the feasibility and optimality of the optimization solution, given the data uncertainty.

Traditionally problems in mathematically programming are solved under the assumption that the input data is precisely known and invariant. However, in reality, the input data, in most cases, may be different from those assumed, which makes the original optimal solution, sub-optimal or even infeasible. In such scenarios, it is prudent to design solution approaches that are

"robust" and immune to uncertainty in the input data. Robust optimization is a technique that is specifically designed to handle data uncertainty and produce near-optimal, feasible and robust solutions, under several data realizations.

Robust optimization differs significantly from methods like stochastic programming, which requires the assumption of knowledge of the input data distributions. In stochastic programming approach, the uncertain parameters are replaced with the expected values of the parameters, and the new nominal problem is then solved. Robust optimization is also different from the dynamic programming approach, which models the uncertain parameters as random variables. The dynamic programming approach suffers from the curse of dimensionality and assumes the distributions of the uncertain parameters are available. Bertsimas & Thiele (2006) showed that the robust optimization techniques lead to high-quality solutions and often outperform their dynamic programming-based counterparts.

5.2   A Brief History

The origin of stochastic linear programming, or linear programming under uncertainty, can be traced back to two seminal papers written independently by Dantzig (1955) and Beale (1955). Dantzig's work was inspired by an earlier paper he coauthored in which he proposes that linear programming methods should be extended to include the case of uncertain demands for the problem of optimal allocation of a carrier fleet to airline routes to meet an anticipated demand distribution (Ferguson & Dantzig, 1954).

> "In retrospect, it is interesting to note that the original problem that started my research is still outstanding -- namely the problem of planning or scheduling dynamically over time, particularly planning dynamically under uncertainty. If such a problem could be successfully solved it could eventually through better planning contribute to the well-being and stability of the world."
>
> -George Dantzig (1991)

A robust approach to addressing linear optimization problems with uncertainty in the input data was proposed in the early 1970s. The robust optimization approach accepts near-optimal

solutions for nominal values of the input data to ensure that the solution obtained continues to be feasible and near-optimal, when the data changes.

The first inroads were made by Soyster (1973) who proposed a linear optimization model that optimizes the model for the worst-case realizations of the data elements. In other words, each data element was assumed to take its most extreme value to ensure that the solution was always feasible, albeit at the cost of optimality.

Ben-Tal & Nemirovski (2000) argued that the Soyster's model produces solutions that are too conservative and sacrifices too much optimality just to ensure the solution robustness. They address the issue of over-conservatism by proposing the use of ellipsoidal uncertainties, which solves the robust counterparts of the linear problems in the form of conic quadratic problems. However this approach has practical drawbacks as it leads to non-linear, although convex, models, which are computationally expensive to solve. The biggest drawback to the Ben-Tal/Nemirovski approach is that it does not extend well to integer formulations. When the decision variables are enforced to be integers, the problem becomes a nonlinear optimization problem, which is inherently non-convex and extremely difficult to optimize efficiently. This makes the formulation especially ill-suited for discrete optimization problems, like our portfolio problem.

Bertsimas & Sim (2004) proposed a new approach for robust linear optimization which retains a linear framework, and at the same time, provides deterministic and probabilistic guarantees against constraint violations. Their methodology is based on the premise, that only a small subset of data elements takes their worst-case values at the same time. They argue that robustness comes at a certain associated cost, and they provide a parameter Γ to control the degree of robustness of the solution. The parameter Γ guarantees a feasible solution for instances in which fewer than Γ parameters take their worst-case values. The approach even provides a probabilistic guarantee, that if more than Γ parameters change, the robust solution will still be feasible to a high degree of probability. Since their robust formulation retains the linear nature of the problem, it can be extended to discrete optimization problems, like

knapsack and portfolio problems. For these reasons, we shall adopt the Bertsimas - Sim approach, for addressing our research problem.

5.3 Linear Programming Problems

Let us consider a standard nominal linear optimization problem:

$$\text{maximize:} \quad c'x$$

$$\text{subject to:} \quad Ax \leq b \tag{5.1}$$

$$l \leq x \leq u$$

Here $c = [c_1, c_2, \dots, c_N]$ denote the objective function coefficients; the matrix A and the vector b represent the data in the constraints imposed on the decision variables $x = [x_1, x_2, \dots, x_N]^T$. The vector x is a feasible solution if it satisifies the constraints imposed by A and b. We assume that the data uncertainty only effects the matrix A, since we can always change the objective function to maximize z, then incorporate the new constraint $z - c'x \leq 0$ into the existing constraint $Ax \leq b$.

Each entry of the constraint matrix A, a_{ij} is modeled as a symmetric variable $\tilde{a}_{ij}$. We do not assume any probability distribution for the random variable $\tilde{a}_{ij}$ but restrict it to only take values in the interval $[\,a_{ij} - \hat{a}_{ij}, a_{ij} + \hat{a}_{ij}\,]$. Specifically a_{ij} denotes the nominal mean value of the distribution and $\hat{a}_{ij}$ is the half-interval of $\tilde{a}_{ij}$. Associated with the uncertain data $\tilde{a}_{ij}$, we define another random variable $\tilde{\eta}_{ij} = (\tilde{a}_{ij} - a_{ij})/\hat{a}_{ij}$ which obeys an unknown but uniform distribution, taking values in the range $[-1,1]$. We let J_i represent the set of coefficients in row i subject to parameter uncertainty, i.e., $\tilde{a}_{ij}$, j $\in J_i$ takes value from the symmetric distribution.

5.3.1 Soyster's Approach

Thus, under this uncertainty model, we can present Soyster's robust formulation as follows:

$$\text{maximize:} \quad c'x$$

$$\text{subject to:} \quad \sum_j a_{ij} x_j + \sum_{j \in J_i} \hat{a}_{ij} y_j \le b_i \qquad \forall i$$

$$-y_j \le x_j \le y_j \qquad \forall j$$

$$l \le x \le u$$

$$y \ge 0$$

$$(5.2)$$

The robust optimal solution x^* of this formulation will remain feasible for every possible realizations of $\tilde{a}_{ij}$. It can be seen that the Soyster's approach is the most conservative in practice, in the sense that the use of the worst-case values results in far-from optimal solutions, for many realizations of $\tilde{a}_{ij}$.

5.3.2 Robust Approach of Ben-Tal and Nemirovski

Ben-Tal & Nemirovski (2000) propose the following formulation

$$\text{maximize:} \quad c'x$$

$$\text{subject to:} \quad \sum_j a_{ij} x_j + \sum_{j \in J_i} \hat{a}_{ij} y_j + \Omega_i \sqrt{\sum_{j \in J_i} \hat{a}_{ij}^2 z_{ij}^2} \le$$

$$b_i \qquad \forall i$$

$$y_{ij} \le x_j - z_{ij} \le y_{ij} \qquad \forall\, i,j \in J_i$$

$$l \le x \le u$$

$$y \ge 0$$

$$(5.3)$$

This model is less conservative than Soyster's model. The authors have proved that in this formulation, the probability of the i^{th} constraint being violated is at most $\exp(-\Omega_i^2/2)$. However the non-linearity of the model makes it a less than desirable approach for solving discrete optimization problems.

5.3.3 Robust Approach of Bertsimas and Sim

Bertsimas and Sim (2002) introduce a parameter Γ_i, not necessarily integer, that can take values in the interval $[0, |J_i|]$. Intuitively speaking, it is unlikely that all $|J_i|$ coefficients will assume its worst-case value. The formulation provides protection against the case when up to $\lfloor \Gamma_i \rfloor$ of these

coefficients are allowed to assume their extreme values and one coefficient a_{it} is allowed to change by $(\Gamma_i - \lfloor \Gamma_i \rfloor)\hat{a}_{it}$. The corresponding robust non-linear formulation can be written as

$$
\begin{aligned}
\text{maximize:} \quad & c'x \\
\text{subject to:} \quad & \max_{\{S_i \cup \{t_i\} | S_i \subseteq J_i, |S_i| = \lfloor \Gamma_i \rfloor, \, t_i \in J_i \setminus S_i\}} \left\{ \sum_{j \in S_i} \hat{a}_{ij} \, y_j + \right. \\
& \left. (\Gamma_i - \lfloor \Gamma_i \rfloor)\hat{a}_{it_i} y_t \right\} + \sum_j a_{ij} x_j \; \leq b_i \qquad \forall i \\
& -y_j \leq x_j \leq y_j \qquad\qquad\qquad\qquad\quad \forall j \\
& l \leq x \leq u \\
& y \geq 0
\end{aligned}
\tag{5.4}
$$

In the above formulation S_i represents the set of uncertain parameters that take their extreme values such that $|S_i| = \lfloor \Gamma_i \rfloor$, for the i^{th} constraint. The second term in the i^{th} constraint is a protection function, which uses the parameter Γ_i to offer various levels of protection. $\Gamma_i = 0$ represents the deterministic case while $\Gamma_i = |J_i|$ reduces the formulation to Sosyster's method.

Bertsimas & Sim (2004) prove that the non-linear formulation can be reformulated as a linear optimization model.

$$
\begin{aligned}
\text{maximize:} \quad & c'x \\
\text{subject to:} \quad & \sum_j a_{ij} x_j + z_i \Gamma_i + \sum_{j \in J_i} p_{ij} \leq b_i \qquad && \forall \, i \\
& z_i + p_{ij} \geq \hat{a}_{ij} y_j \qquad && \forall \, i,j \in J_i \\
& -y_j \leq x_j \leq y_j \qquad && \forall \, j \\
& l_j \leq x_j \leq u_j \qquad && \forall \, j \\
& p_{ij} \geq 0 \qquad && \forall \, i,j \in J_i \\
& y_j \geq 0 \qquad && \forall j \\
& z_i \geq 0 \qquad && \forall i
\end{aligned}
\tag{5.5}
$$

Thus, the Bertsiamas and Sim is as flexible as the one proposed by Ben-Tal & Nemirovski (1998, 1999, 2000), El Ghaoui & Lebret (1997), El Ghaoui, Oustry, & Lebret (1998) among others. Since

the robust optimization form preserves the linearity of the problem, it is computationally more tractable than non-linear problem formulations. This is especially true for discrete optimization problems. Their formulation provides probabilistic guarantees of the feasibility of the constraints when more than Γ coefficients take their worst value. Bertsiamas and Sim also provide the corresponding mathematical formulation when the data are correlated, i.e., there are finite number of sources of data uncertainty that affects all the data. They report numerical results for portfolio optimization problem, knapsack problem, supply chain management problem and network flow problem.

Marla (2007) has provided a discussion of the advantages and disadvantages of the Bertsimas - Sim problem.

Advantages:

1. It is applicable to linear and integer programs.

2. Linear integer programs retain their linearity, but contain more variables, thereby minimally degrading the tractability of the problem.

3. The formulation does not make any assumptions regarding the probability distributions of the uncertain data; it captures uncertainty through symmetric bounds of variation alone.

4. The 'level of robustness' can be adjusted by varying the parameter Γ, thereby providing measures of the changes in the planned objective function with changes in the protection level.

5. The correlated - data model can capture simple correlations between uncertain data in one constraint equation, but not across constraint equations.

Limitations:

1. The problem does not provide any guidelines for choosing Γ, which means that the problem has to be resolved multiple times for different values of Γ_i, for each i. For large scale problems, this may pose computational challenges.

2. It assumes symmetric bounded distribution about the nominal values for the uncertain parameters.

3. The model has no provision to account for known probability distributions.

4. The probability bounds for constraint violation can only be derived for each constraint and not for the system as a whole.

5.4   Robust Optimization Framework for the Proposed Mechanism

The Bertsimas Sim framework is well-suited to address our research problem, despite the cited drawbacks. Since we are working under a single bandwidth constraint, we can change the protection level by varying Γ to determine the change in the cost of robustness as a function of the protection level. Also the probabilistic bound derived for the bandwidth constraint, would be applicable to the whole system.

Thus we introduce the subsequent steps of the second stage to our mechanism, based on the Bertsimas Sim formulation of a robust portfolio problem.

---

2. Second Stage (continued)

2.3   For each target $\in K$, where K is the set of all targets in the system, the trusted center selects a set of agents L_k to report their observations $(\hat{x}_{jk}\,, \hat{\theta}_{jk})$ to the center, $j \in L_k$. For each observation $\hat{x}_{jk}$, the center calculates the nominal information content θ_{jk} and the bound $\hat{\theta}_{jk}$.

2.4   Based on the total NCT, the maximum limit of the NCT, and the time needed to transmit each observation, the center calculates the total number of target-agent pairs N_{auc} to be selected for transmission. The center solves the portfolio optimization problem

$$\text{maximize:} \quad \sum_{k \in K ,\, j \in L_k} \theta_{jk} z_{jk}$$

$$\text{subject to:} \quad \sum_{k \in K ,\, j \in L_k} z_{jk} = N_{auc} \qquad\qquad (5.6)$$

$$z_{jk} \in \{0,1\}$$

2.5   The agents selected through the solution of the optimization problem, are asked to transmit the observations on their allocated target to each agent in the simulation, until the next auction process.

5.5   Robust Portfolio Problem

The portfolio problem assumes that a portfolio needs to be constructed consisting of a set of stocks. Each of the stocks has a return and a risk value associated with it and the objective of the problem is to determine the fraction of wealth that must be invested in each stock, to maximize the portfolio value.

In reference to our research problem, the stocks represent the observations made by the sensor agents. The return value of the stock can be interpreted as the information content of each observation; the risk value indicates the uncertainty in the information content while the total wealth available for investment signifies the maximum bandwidth available for transmission on the tactical data link.

We reformulate the traditional portfolio problem as zero-one portfolio problem where the decision to be made is whether a particular stock should be included in the portfolio or not. The portfolio problem that we need to optimize is the following discrete optimization problem:

$$\text{maximize:} \quad \Sigma_{k \in K, j \in L_k} \theta_{jk} z_{jk}$$

$$\text{subject to:} \quad \Sigma_{k \in K, j \in L_k} z_{jk} = N_{auc} \quad \quad (5.7)$$

$$z_{jk} \in \{0,1\}$$

where,

K	Set of targets in the system
L_k	Set of agents selected through the proper-scoring rules algorithm for target k
θ_{jk}	Information content of the observation made by agent j of target k
N_{auc}	The total number of agent-target pairs that can selected for transmission
z_{jk}	Binary decision variable, which takes value 1 if agent j is selected to transmit information of target k and 0 otherwise

Regarding the uncertainty model for data, we assume that the information content is uncertain. In other words, we model the information content θ_{jk} as a random variable $\tilde{\theta}_{jk}$ that has a symmetric distribution in the interval $[\theta_{jk} - \hat{\theta}_{jk}, \ \theta_{jk} + \hat{\theta}_{jk}]$. θ_{jk} is the expected information gain, while $\hat{\theta}_{jk}$ is a measure of the uncertainty of the information content.

Using the Bertsimas Sim framework for reformulating the portfolio problem:

$$\text{maximize:} \quad y$$

$$\text{subject to:} \quad y \leq \sum_{k \in K, j \in L_k} \theta_{jk} z_{jk} - \beta(z_{jk}, \Gamma)$$

$$\sum_{k \in K, j \in L_k} z_{jk} = N_{auc}$$

$$z_{jk} \in \{0, 1\} \tag{5.8}$$

$$\beta(z_{jk}, \Gamma) = \max_{\{S \cup \{t\} | S \subseteq J, |S| = \lfloor \Gamma \rfloor, t \in J \setminus S\}}$$

$$\left\{ \sum_{jk \in S} \hat{\theta}_{jk} z_{jk} + (\Gamma - \lfloor \Gamma \rfloor) \hat{\theta}_{jk_t} z_{jk_t} \right\}$$

In this setting, Γ is the level of protection that can be adjusted for the robust portfolio optimization. We can solve this linear discrete optimization problem through state-of-the-art LP solvers. For our research work, we use Gurobi Optimizer 5.5 to solve the problem at hand to optimality.

5.6   Conclusion

In this chapter, we have finished the formulation of the framework that we introduced in Chapter 4. The proposed mechanism, based on modified scaled strictly proper scoring rules, two-stage mechanisms and robust optimization, addresses all the six requirements we outlined at the end of Chapter 1 and Chapter 3. Specifically, we have designed a protocol that ensures that each agent participates and truthfully reports their interdependent track information, based on which the auctioneer can allocate targets to the agents by taking into account the uncertainty in the information content, and the selected agents invest their resources to generate and truthfully report the observations on their allocated targets. We are now in a position to apply our mechanism to the agent-based model described in Chapter 3, and discuss the numerical results and evaluate the performance of the mechanism.

CHAPTER 6.   RESULTS

In this chapter, we investigate the application of our modified strictly proper scoring rules based mechanism to the agent-based model of the tactical data network. We first discuss the Graphical-User Interface (GUI) of the ABM created in DAF (Discrete Agent Framework), and then study the behavior of the mechanism under different settings. The effects of the lack of access to true outcomes, deceptive behavior on the part of the sensor platforms and the protection level of the robust optimization is evaluated and studied.

6.1   Agent-based Model GUI

We develop an agent-based simulation model within DAF to capture the performance of a group of military platforms, tasked with the goal of detecting and tracking targets. At the core of the ABM is a track data generator that provides the raw track data. Each platform adjusts the raw data to reflect the position, heading, and affiliation and the error bounds as measured by the onboard sensors.

Figure 6-1 highlights the main features of the DAF-based ABM Graphical-User Interface (GUI). The application framework provides significant dynamism and plenty of features that can be customized to generate unique scenarios. We first discuss these customizable features, located in the bottom left corner of Figure 6-1.

1. **Number of Targets**

    The total number of initial targets to populate the scenario can be specified here. The selected targets have different origin and destination points, dynamics and affiliations – Hostile, Friendly or Neutral.
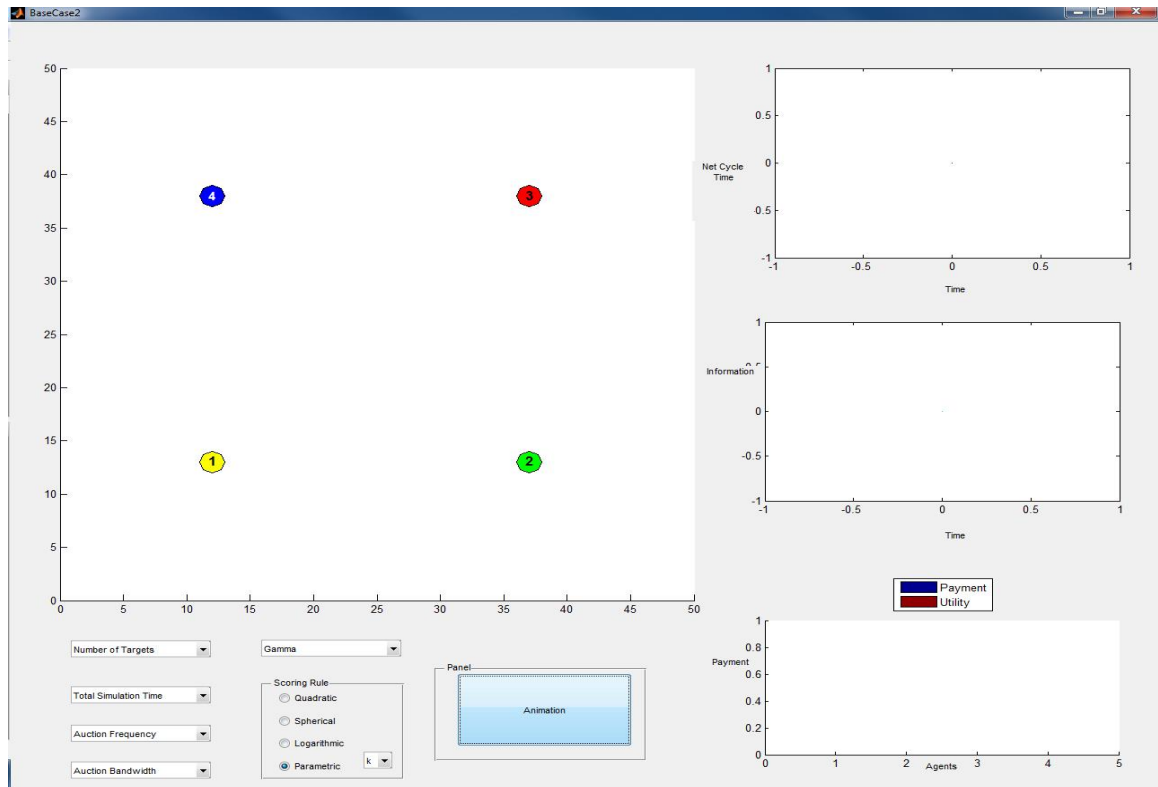
Figure 6-1 Application GUI

2. Total Simulation Time

This parameter allows the total time duration of the simulation to be specified.

3. Auction Frequency

Auction frequency can be varied from 2 to 20 cycles, where N means an auction is initiated every N cycles. Since the mechanism is implemented over two distinct stages spread over two transmission cycles, a minimum auction frequency of 2 needs to be selected for a successful auction to be conducted.

4. Auction Bandwidth

This option allows for the selection of the maximum Net Cycle Time (NCT) or Tracks that can be auctioned. The baseline NCT required to transmit the R^2 data depends on the number of tracks in the simulation as well as the configuration of the tracks

generated. The auction bandwidth represents the difference between maximum NCT and baseline NCT, which is auctioned for the transmission of additional track information to improve the quality of the common operating picture.

5. **Gamma**

Gamma indicates the protection level of the robust optimization framework that is used to select the subset of track information to be transmitted, over and above the R^2 data. A higher value of gamma provides lower returns on uncertainty adjusted information value with lower probability of constraint violation while a lower value of gamma provides higher returns with more risk.

6. **Scoring Rule**

The simulation framework is designed to work with four different modified strictly proper scoring rules discussed in Chapter 5. We shall evaluate the expected and actual payments, variance of payments and utility values for the different scoring rules – Quadratic, Spherical, Logarithmic and Parametric. For the parametric scoring rule, an additional parameter has to be selected and the value of the parameter can lie in the interval (1,3).

Once all the input values have been specified the simulation is executed by pressing the 'Animation' button. The top-left window provides a real-time track display of the simulation scenario with four symmetrically positioned platforms and a distributed set of targets. The three windows on the right allow the study of the behavior of the mechanism for different scenarios. We simulate a demo scenario in the application GUI with the following inputs:

- Number of Targets = 16 units
- Total Simulation Time = 50 time steps
- Auction Frequency = 6 cycles
- Auction Bandwidth = 10 units

- Gamma   =   0
- Scoring Rule   =   Quadratic

We provide a snapshot of the simulation in Figure 6-2, in order to discuss the highlights of the captured behavior. We start from the main runtime window, and then examine the different plots in the right from top to bottom.
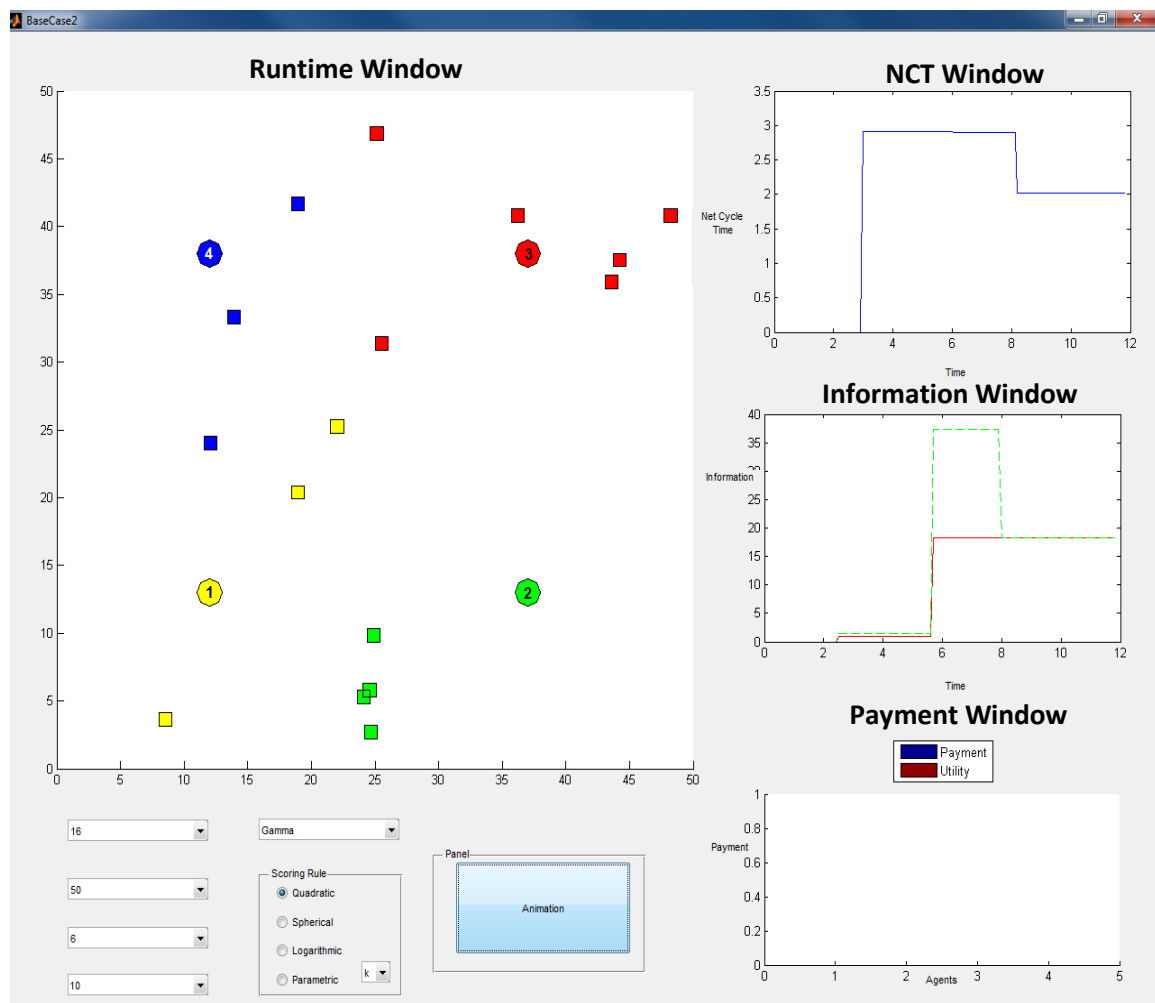

Figure 6-2 Snapshot of a simulation in progress

1. **Runtime Window**

   The main run-time simulation window displays the four platforms located symmetrically in the 50 – by - 50 miles region. The platforms are numbered 1, 2, 3 and 4 and the role of the Grid Reference Unit (GRU) is played by Platform 3. The GRU is typically the platform possessing the highest quality track data and the GRU almost universally is assigned the reporting responsibility (R^2) for any track within its classification radius. The tracks are spread across the simulation map and are designated by small squares. The color of these squares indicates the sensor that has been assigned R^2 for that particular track. For instance, the yellow tracks are being tracked by Platform 1, the green ones by Platform 2, the red targets by Platform 3 and Platform 4 is tracking the blue-colored tracks. Since an auction has not yet been conducted, there is only a single platform tracking any target.

2. **NCT Window**

   The top-right window shows the net cycle time against the current simulation time. In the first cycle, each platform sends the track information for all the targets within its observation region and this round of transmission corresponds to a higher NCT of approximately 3 seconds. Once the auctioneer assigns R^2 of the tracks to the sensor platforms, the platforms only broadcast the track data for their assigned targets and the NCT drops to 2 seconds in the subsequent cycles.

3. **Information Window**

   The middle window displays the information content of the tracks being transmitted. It shows three distinct information values: the maximum information achievable (dashed green line), the current information being transmitted (red line) and the baseline R^2 information (blue line). The information value is computed as a function of the covariance of track data as outlined in Section 3.4.3. The dashed green line includes redundant reporting on the targets that are not selected for transmission while the blue

and red lines overlap, since only the R^2 information is being transmitted. The difference between the green and the blue lines indicates the loss in information that occurs for R^2 mechanism which doesn't allow for redundant reporting of a single object. When a target is not transmitted, the opportunity to fuse its data is also lost. Hence the goal of the mechanism is to recover the maximum information for a given quantum of additional transmission time.

4. Payment Window

The payment and utility for each round of auction is displayed in the window at the bottom-right corner. The payments made to each platform are shown using the blue-colored bars while the red bars indicate the utility values gained by each platform. These values are determined by the scoring rules selected and can take both positive and negative values depending on the reported maximum and actual information content values of the track data. Since the auction mechanism has not been conducted, no payments have been made to the platforms yet.

The platforms continue to track and transmit information for the assigned targets. Once six transmission cycles have elapsed, an auction is initiated based on the input parameters. The auction mechanism follows the two stages as outlined in Chapters 4 and 5, and a snapshot of the post-auction simulation is captured in Figure 6-3.

The main run-time simulation window now shows multiple targets being tracked by one or more platforms. Three of these targets are being tracked by three sensors while four targets are being tracked by two sensors. Since we specified the auction bandwidth parameter input corresponding to ten units, the mechanism selects ten additional transmission tracks. This selection is based on a host of factors – the information value, the track affiliation, protection level gamma and the uncertainty in the information values.
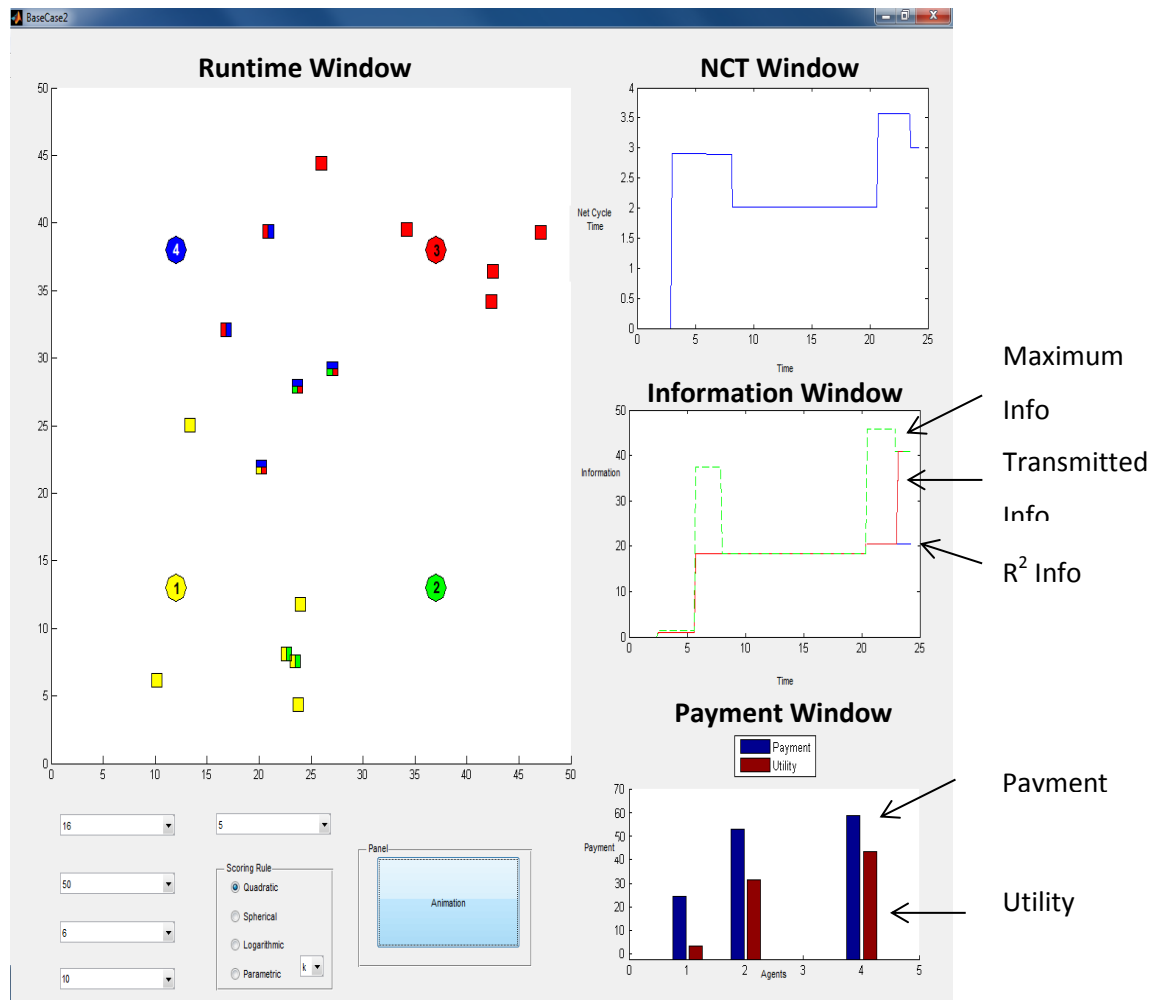
Figure 6-3 Outputs of the simulation

The Net Cycle Time window shows a peak corresponding to the time when the auction mechanism was executed. In the first stage of the mechanism each platform transmits the maximum information content values as well as its cost functions for all the targets in its observation region to the center. This results in the first peak corresponding to a NCT of 3.5 seconds. Once the remaining stages of the mechanism have been executed, a small subset of ten tracks is selected for transmission and the NCT comes down to 3 seconds.

The Information window shows the divergence of the three different metrics used to measure information content of the tracks being transmitted. In the first stage of the auction, all sensor platforms transmit the maximum information values for all targets within their detection radius

and the dashed green line corresponding to the maximum available information shoots up. However the transmitted information (red line) and the R^2 information (blue line) continue at the baseline level. Once the second stage of the mechanism has been executed, the transmitted information spikes to a higher level corresponding to the selected subset of targets. The difference between the red and blue lines indicates the track information recovered as a result of the auction mechanism.

The payment window shows the reimbursement handed out to the participating platforms in the auction mechanism. The utility is measured as difference of the payment and the costs incurred for generating and reporting the track data. An interesting facet of the simulation is the effect of the mechanism on the platform that has been assigned the role of GRU. As apparent from Figure 6-3 the Sensor Platform 3 (the designated GRU) has payment and utility values of zero. This asymmetry can be attributed to the fact that the GRU universally is assigned the reporting responsibility (R^2) for any track within its classification radius. Hence it cannot contribute any additional track data towards the auction mechanism which makes it a perpetual receiver of non - R^2 track information from the other sensor platforms. This represents a seemingly anomalous yet correct real world behavior representation (Klein et al., 2008).

6.2   Sensitivity Analysis

The simulation framework exhibits significant dynamism by offering various parameters which can be adjusted to vary the characteristics of the simulation scenario. The trade-space obtained by assuming different initial conditions can be inexhaustibly large and hence for the sake of expediency, we fix the values of certain parameters. In all the following simulations, the number of initial targets in the simulation is fixed at sixteen, the auction bandwidth is assigned the maximum value of forty-eight, the auction frequency is maintained at six cycles and the simulation is carried out for fifty time steps. We assume the cost as a linear function of information content, represented as $c_i(\theta) = c_i * \theta$, where the value of c_i is selected from a uniform distribution $c_i \in U(1, 1.5)$. This allows us to focus our discussion on the variables of interest – scoring rules, protection levels and the maximum number of preselected sensor platforms M.

6.2.1   Maximum Number of Preselected Sensor Platforms (M)

In the first stage of the proposed mechanism, the trusted center requests all sensor platforms in the simulation to report their cost functions and to reveal their private maximum information content. The center then preselects M sensor platforms from the N available sensor platforms with the lowest cost functions through one single reverse $(M + 1)^{th}$ auction. In our simulation scenario, since there are four sensor platforms and the R^2 tracks have already been pre-assigned, there are only $N = 3$ sensor platforms available for selection for transmitting $non-R^2$ track data. Thus M can take three distinct integer values $M \in [1, 3]$. The number of pre-selected sensor platforms M dictates the maximum number of sensor platforms that can track any one target. For example, for $M = 2$, a maximum of 3 platforms can be assigned to one target, one for R^2-track data and two for non-R^2 track information.

We first observe the impact of M on the total information flow in the simulation and the net cycle time. The results are shown in Figure 6-4.

Figure 6-4 represents the variation of the transmitted information and the net cycle time as the maximum number of pre-selected platforms (M) is varied from 1 to 3. The baseline case of R^2 is represented using the blue diamonds; the maximum information case is represented using the green triangles, while the red squares represent the case corresponding to the subset of selected track data. For $M = 1$, only one platform can be selected to transmit the non - R^2 track data for each target. The average net cycle time corresponding to this scenario is 3.012 seconds and located approximately midway between the baseline value (2.127 seconds) and the maximum value (3.782 seconds). In the next scenario M is fixed as 2 and thus two platforms can be selected per target for $non - R^2$ track data transmission. In this instance, the transmitted information value is much closer to the maximum achievable situational awareness, though it also corresponds to a higher average net cycle time of 3.472 seconds. When M is fixed at 3, it is possible to achieve the highest possible gain in the overall Common Operating Picture (COP) quality. This case corresponds to the highest mean NCT of 3.782 seconds.
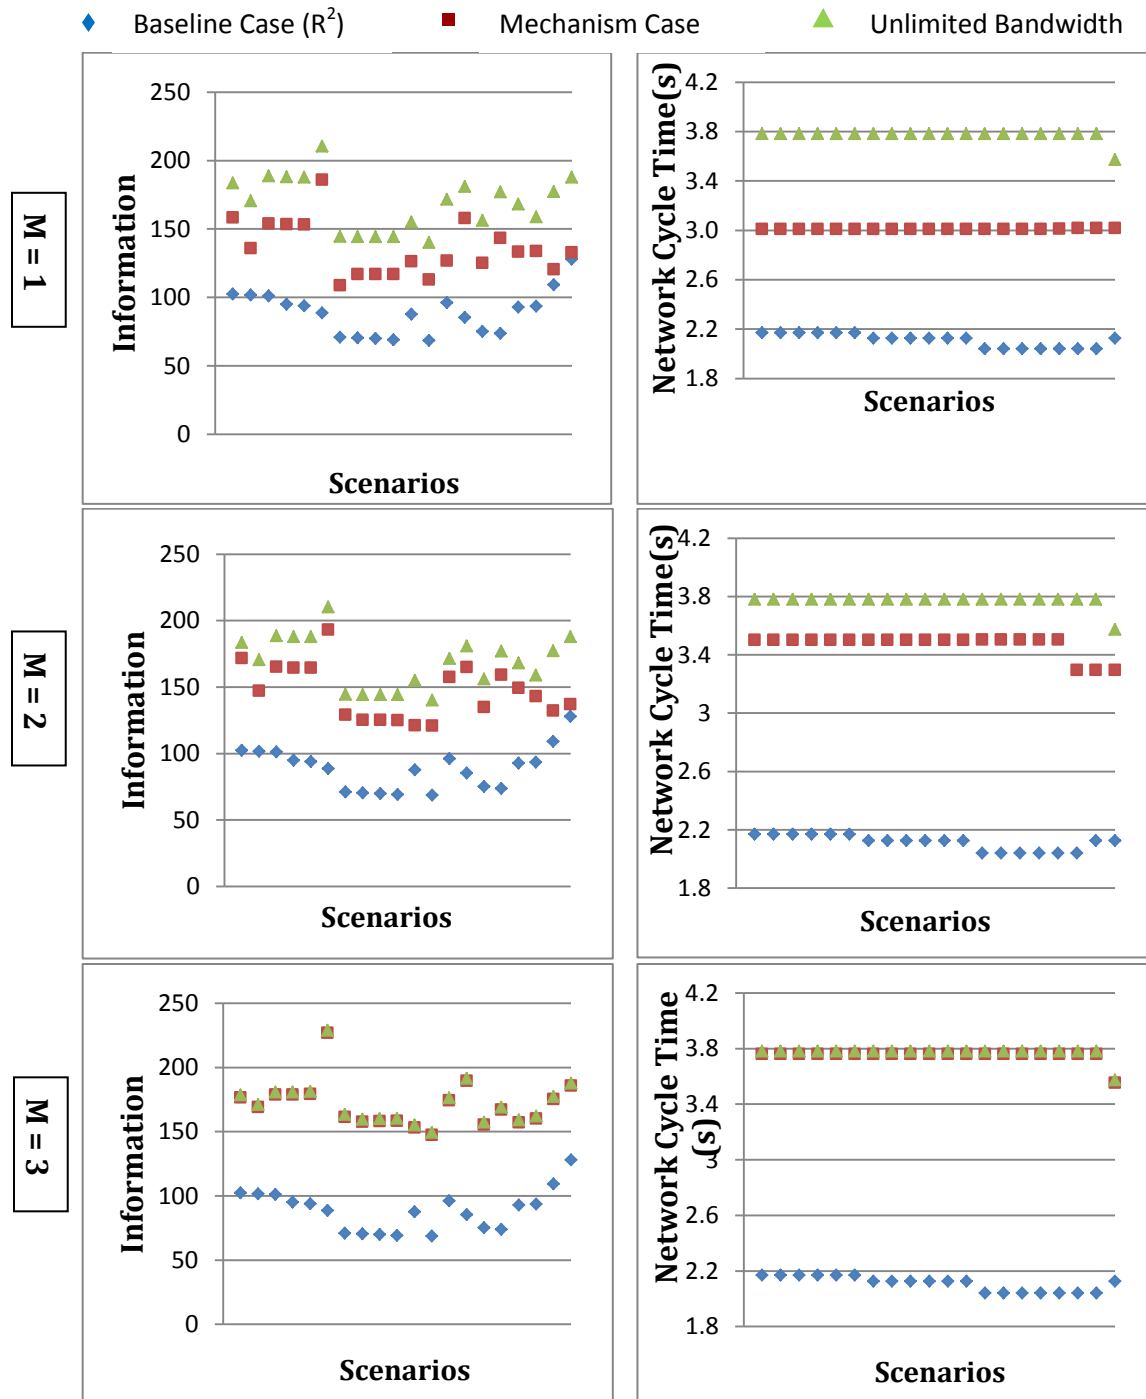
Figure 6-4 Information and Network Cycle Time for different values of M

The selection of the value of M is dictated by a tradeoff between information content and information latency. Increasing the value of M allows additional track data to be transmitted over the tactical data network, though it also results in increased latency between successive track updates. A few guiding factors that should be taken into consideration when choosing M are listed below:

1. The gain in the information for a given quantum of NCT
2. The total expected payment made by the center

We capture the interaction of these factors with the number of pre-selected sensor platforms in Figure 6-5.

Figure 6-5(a) shows how the additional gain in the information flow per second of increase in the Net Cycle Time changes with the maximum number of pre-selected sensor platforms. $M = 1$ results in the highest recovery of non - R^2 track information for a given quantum of extra NCT. This value goes down for the case $M = 2$, indicating that the increase in the transmitted information is marginal as compared to the extra latency induced in the network. For the case when all the sensor platforms are allowed to track a single target, the gain in the information content per quantum of bandwidth improves slightly, but does not reach the value corresponding to $M = 1$. However as the error bounds indicate, there is considerable uncertainty associated with these values as the gain in information per unit second is dependent on the distribution of the targets over the simulation map. Thus caution must be exercised when using this as a metric for selecting the value of M.

Figure 6-5 (b) plots the expected total payment made by the center against the maximum number of pre-selected sensor platforms. We consider the payments made under different scoring rules – Quadratic, Spherical, Logarithmic and the limiting cases of the Parametric ($k \rightarrow 1$ and $k \rightarrow 3$). For the sake of clarity we omit the standard errors and plot only the mean values of the expected payment. The variance of the expected payment will be discussed in greater depth in the next section.
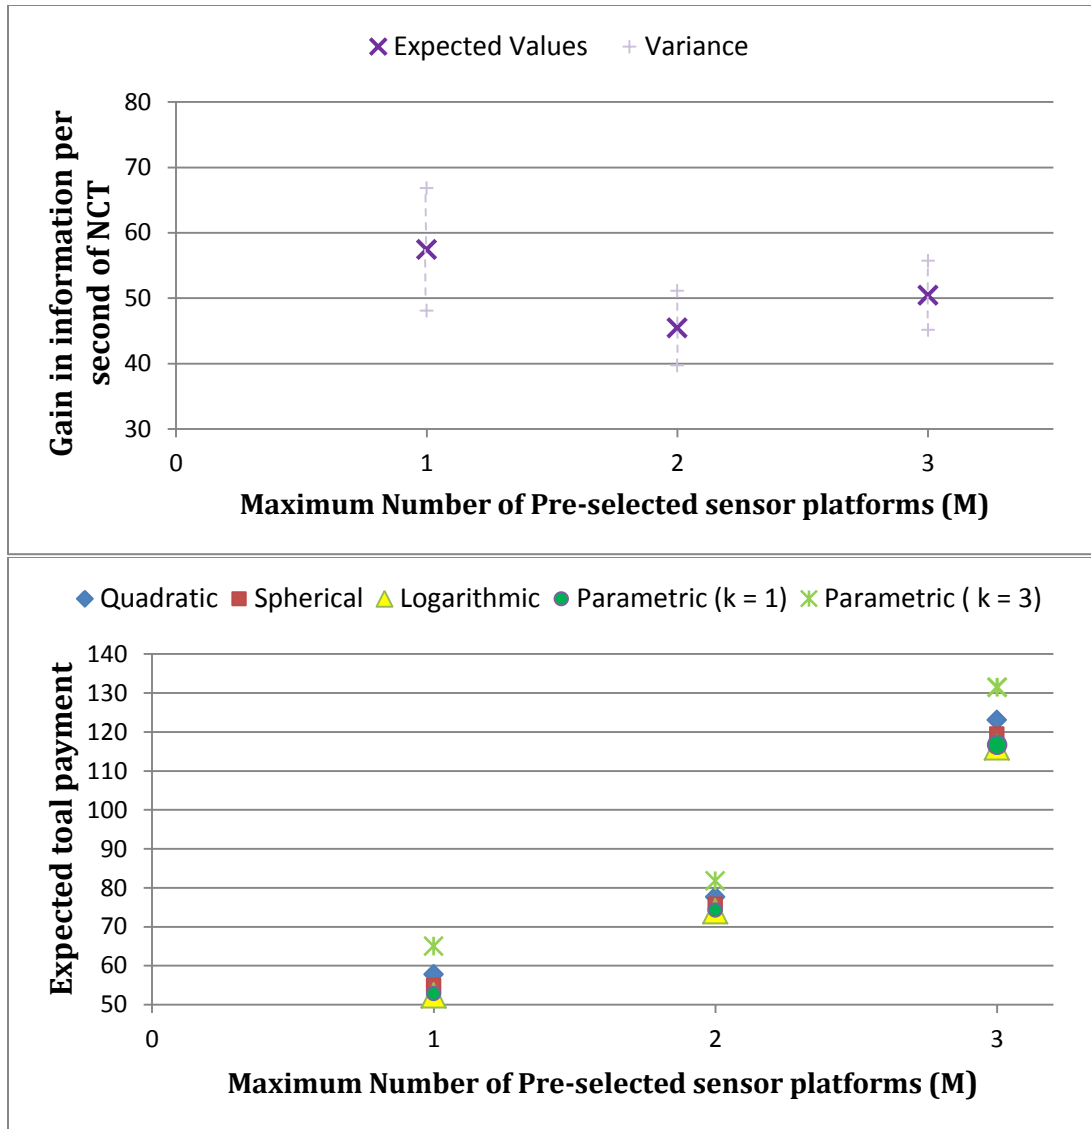
Figure 6-5 (a) Gain in Information per second of NCT vs. M (b) Expected Total Payment vs. M

From the graph in Figure 6-5(b) we can observe that by setting $M = 1$ the center minimizes its expected payment as it directly chooses the sensor platform with the lowest cost function. This case represents a lower payment bound for the mechanism for each instance of the implemented scoring rule, because the center selects the parameters α and β to ensure that the expected utility of the platforms is only slightly greater than 0. Logarithmic scoring rule and the limiting case of the Parametric scoring rule ($k \rightarrow 1$) results in the lowest expected payment for every value of M, while the other limiting case of the Parametric scoring rule ($k \rightarrow 3$) results in the highest expected payment. If the cost distributions and maximum information content

values are known a priori, the value of M can be selected through multiple simulations. But it is unreasonable to expect that these distributions will be public knowledge in a real-world scenario. It is, therefore, prudent to select the highest value of M which ensures that the mechanism captures the highest possible gain in the overall common operating picture quality, even if it results in the center incurring the highest expected payment.

6.2.2   Scoring Rules

In this section, we turn our attention to the different scoring rules. Specifically, we empirically evaluate the parametric scoring rule and compare it to the quadratic, spherical and logarithmic scoring rule, for the parameter space of $k = (1,3)$. We discuss how the parametric scoring rule converges to one of the other scoring rules for different values of the parameter k, along with the advantages and disadvantages of the various scoring rules.

In order to facilitate the discussion on the comparison of the four scoring rules, we generate the plots of the total expected payment made by the center, the variance of the payment and the minimum payment for the parameter space of $k = (1,3)$. We present these results in Figure 6-6.

Figure 6-6 illustrates that the payment scheme based on the logarithmic scoring rule results in the center making the lowest expected payments to the sensor platforms. The logarithmic scoring rule also results in the lowest variance in the payments made by the center. This would seem to indicate that the logarithmic scoring rule enjoys a significant advantage over the quadratic and spherical scoring rules.

Another distinctive trait that can be observed in Figure 6-6 is that the expected payment and the payment variance that results from the logarithmic, spherical and quadratic scoring rules is the same as those based on the parametric scoring rule, for values of the parameter $k \rightarrow 1, k = 1.5$ and $k = 2$ respectively. Thus for $k \rightarrow 1$, the expected payment and the variance of the payment based on the parametric scoring rule are asymptotically equal to the expected payment and the payment variance of the logarithmic scoring rule. Furthermore the same results apply for $k = 1.5$ and the spherical scoring rule, where the expected payment and the
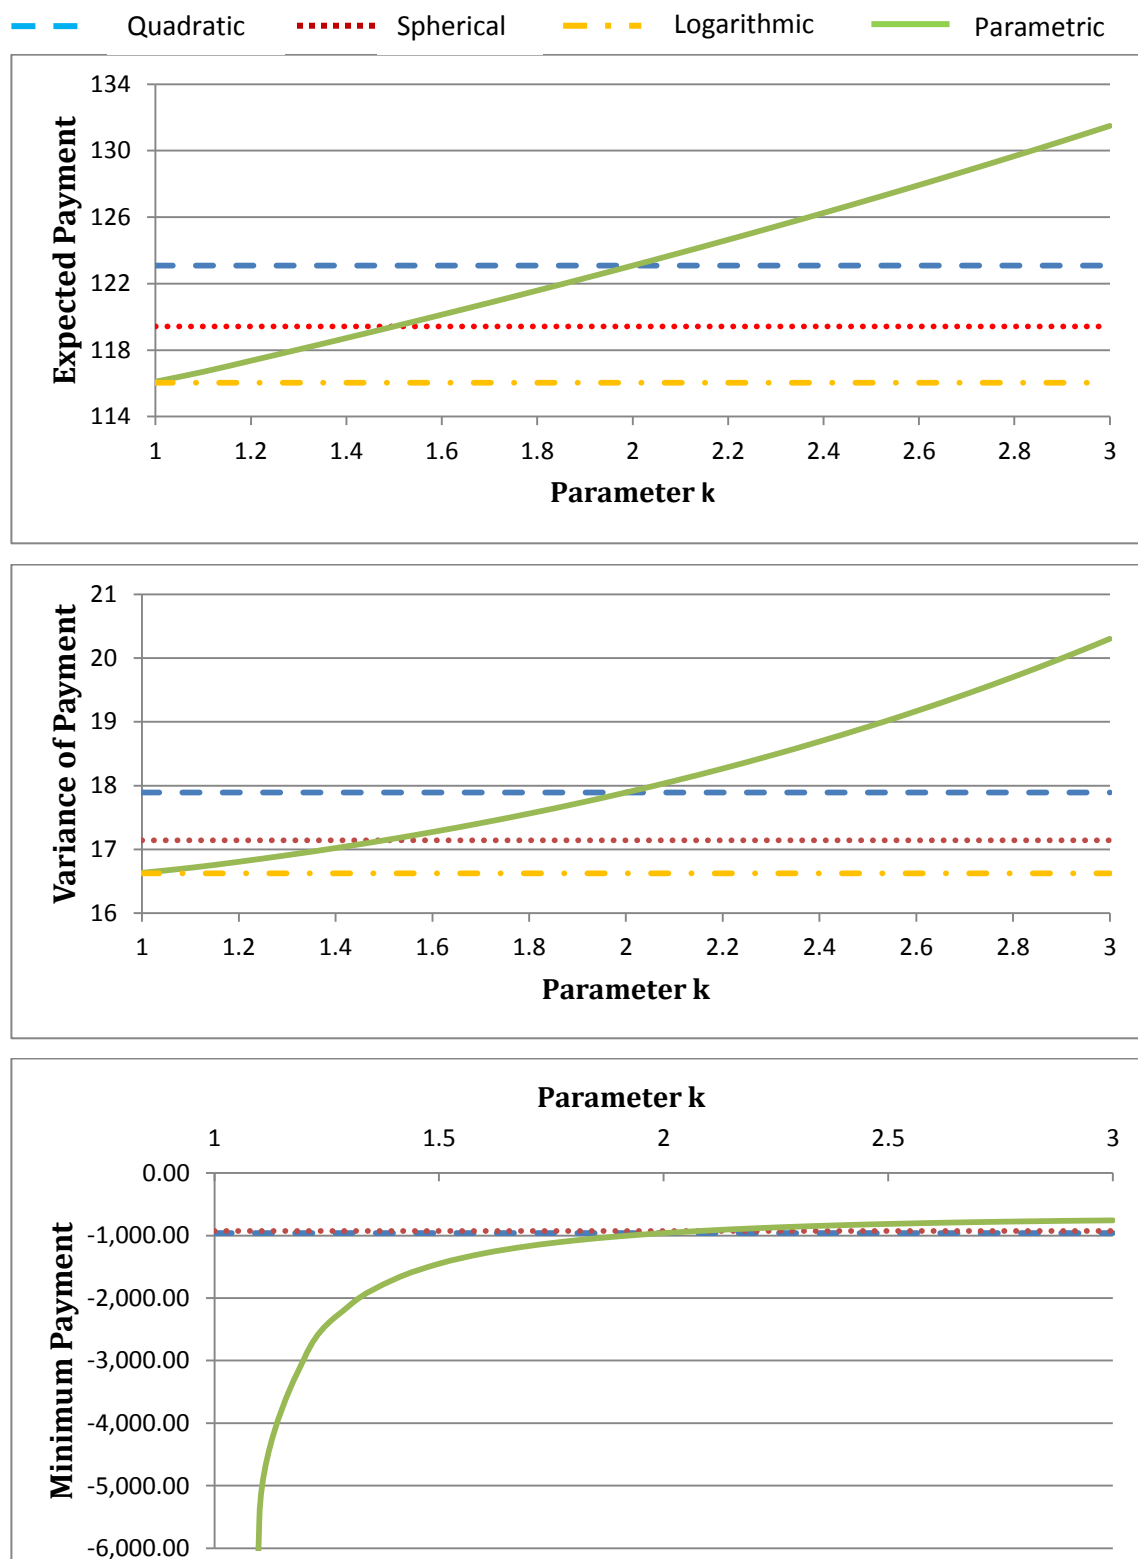
Figure 6-6 (a) Expected Payment (b) Variance of Payment (c) Minimum Payment vs. Different Scoring Rules

variance of the payment made by the center is the same. When $k = 2$, the values of the expected payment and the payment variance for the parametric scoring rule converges to the corresponding values obtained from the quadratic scoring rule. This result serves as a validation for the analytical derivation where the parametric scoring rule takes the same expression for the expected payments as the logarithmic, spherical and quadratic scoring rules, for values of the parameter $k \to 1$, $k = 1.5$ and $k = 2$ respectively.

Hence the parametric scoring rule seems to share the advantages exhibited by the logarithmic scoring rule. However the parametric scoring rule enjoys a significant edge when we consider the lower bounds on the payments. To analyze this result further, we consider the analytical expressions for the parametric and logarithmic scoring rule we obtained in Chapter 4.

$$
\begin{aligned}
S_{logarithmic} &= \log N\left(\bar{x}_i; x_i, \frac{1}{\theta_i} + \frac{1}{\bar{\theta}_{-i}}\right) \\
P_{logarithmic} &= 2c'^{(\theta_0)\theta_0} * S_{logarithmic} + c(\theta_0) - 2\theta_0 c'(\theta_0) \\
&\qquad * \left(\frac{1}{2}\log\left(\frac{\theta_0}{2\pi}\right) - \frac{1}{2}\right) \\
S_{parametric} &= kN\left(\bar{x}_i; x_i, \frac{1}{\theta_i} + \frac{1}{\bar{\theta}_{-i}}\right)^{(k-1)} \\
&\qquad - \frac{k-1}{\sqrt{k}}\left(\frac{2\pi}{\theta}\right)^{(1-k)/2} \\
P_{parametric} &= \frac{2c'(\theta_0)\theta_0\sqrt{k}}{k-1}\left(\frac{\theta_0}{2\pi}\right)^{(1-k)/2} * S_{parametric} + c(\theta_0) \\
&\qquad - \frac{2\theta_0 c'(\theta_0)}{k-1}
\end{aligned}
\tag{6.1}
$$

In the modified scaled strictly proper scoring rules, the trusted center fuses the observations from all the other sensor platforms and excludes the sensor platform whose reported observation is being evaluated. Thus if the reported observations of the evaluated sensor platform is far from those reported by the other sensor platforms, then the probability distribution function of the platform goes to 0 (i.e. $N\left(\bar{x}_i; x_i, \theta_i^{-1} + \bar{\theta}_{-i}^{-1}\right) = 0$). In the logarithmic case, the scoring rule approaches negative infinity and thus the logarithmic scoring

rule does not have a finite lower bound of the payment. In the case of parametric scoring rule as $k \rightarrow 1$, the scoring rule goes to 0 as the probability distribution function goes to 0. The payment based on the parametric scoring rule also does not have a finite lower bound, since the coefficient of the scoring rule function approaches negative infinity when $k \rightarrow 1$. Therefore both the logarithmic scoring rule and the limiting case of the parametric scoring rule family results in large negative payments when the sensor platforms produce imprecise observations. Nevertheless, if the parameter is chosen judiciously, the effect of the platform's imprecise estimate can be minimized for the parametric scoring rule family.

Figure 6-6 plots the lower bounds of the payment of the parametric scoring rule for the parameter space $k = (1, 3)$ against the spherical and quadratic scoring rule. The logarithmic scoring rule and the limiting case of the parameter scoring rule family ($k \rightarrow 1$) does not have a finite lower bound and hence are omitted from the figure. It can be observed that the minimum payments increase for the family of parametric scoring rules as the value of the parameter increases. For $k = 2$, the lower bound is similar for the quadratic and the parametric scoring rule, like we saw in the case of expected payments and the variance of payment. But the same is not true for $k = 1.5$ and the spherical rule, as the lower bound on the payment for the spherical scoring rule is higher than that of the quadratic rule.

The selection of an appropriate value of k for the parametric scoring rule is, hence, dictated by a tradeoff between expected payment, variance of the payment and the lower bound on the payment. From the simulations, a value of $k \in [1.1, 1.5]$ appears to be a judicious compromise between the different factors. This set of parameter values produces low expected payments and variance of payments which are close to the ones obtained from the logarithmic scoring rule, and at the same time, imposes a finite lower bound on the minimum payments. The exact choice of the parameter will depend on the overall objectives of the mechanism and the mechanism designer.

6.2.3   Protection Level

This section illustrates the robustness of the performance of the simulation for different values of the protection level. In all the preceding results the value of the conservation parameter,

Gamma, was assumed to be zero and the nominal values of the information content was used for all computations. In this scenario, we model the information as a random variable with a symmetric distribution where the uncertainty is calculated as the difference between the estimated and the actual information content. We use the same inputs for the application framework as in the previous sections, except for the Auction Bandwidth which is decreased to 18 units. The motivation behind this is to observe how the composition of the portfolio of the selected sensor-target pairs changes with the protection level. This case study is indifferent to the choice of the scoring rule, as our interest lies solely with the expected information flow and the probability of achieving this expected information flow. We expect to see similar results for different values of M (maximum number of preselected sensor platforms) and the various scoring rules, as the portfolio optimization problem only takes in the distribution of the information valuation measures as input.

We first solve the robust portfolio optimization problem using the Bertsimas-Sim formulation as outlined in Section 5.5 for different values of protection level (Γ).

Figure 6-7(a) shows the expected information flow and the uncertainty-adjusted information flow in the mechanism for different values of Gamma. The uncertainty-adjusted information is the value of the objective function that the Bertsimas-Sim framework seeks to optimize. It represents the difference between the expected information flow and the uncertainty function when at most Γ variables are allowed to take their worst values. The graph illustrates the phase transitions that occur as the value of Γ increase.

- For $\Gamma \leq 6$, both the expected information flow and the uncertainty adjusted information flow are insensitive to the protection level. This can be attributed to the difference between the total number of target-sensor pairs available for selection ($= 24$) and the maximum auction bandwidth ($= 18$). In this phase, the solution indicates that the worst information values will be taken by the target-sensor pairs which are not a part of the solution portfolio.

- For $6 < \Gamma \le 24$, the uncertainty adjusted information flow decreases as the protection level increases. In this phase, the number of sensor platforms that report observed information content values different from their estimated information content values is allowed to increase with Γ. Once the value of gamma becomes equal to the total number of target-sensor pairs available for selection $(= 24)$, the uncertainty adjusted information reaches its minimum value and doesn't decrease any further.

- For $6 < \Gamma \le 24$, the expected information flow shows sudden phase transitions for values of $\Gamma = 10, 15, 16$ and 20. These transition points for the expected information flow coincides with the protection levels where the composition of the portfolio changes.

- For $\Gamma \ge 24$, the portfolio is composed of target-sensor pairs with the largest uncertainty-adjusted information values. This represents the ultra-conservative solution given by Soyster's method and for this phase both the expected and uncertainty-adjusted returns are insensitive to increase in Γ.

As we discussed in Chapter 5, one of the main attractions of the Bertsimas-Sim framework for discrete robust optimization problem is that it provides probabilistic bounds of constraint violation. The formulation provides a theoretical bound on the fraction of portfolios with information flow values which will fall below the robust solution of the uncertainty adjusted information value. We plot this probability of underperforming as a function of the protection level Γ in

Figure 6-7(b). For low protection levels, the probability of the portfolio solution falling below the robust solution is quite high. As the protection levels increase the probability of underperforming decreases by several orders of magnitude.
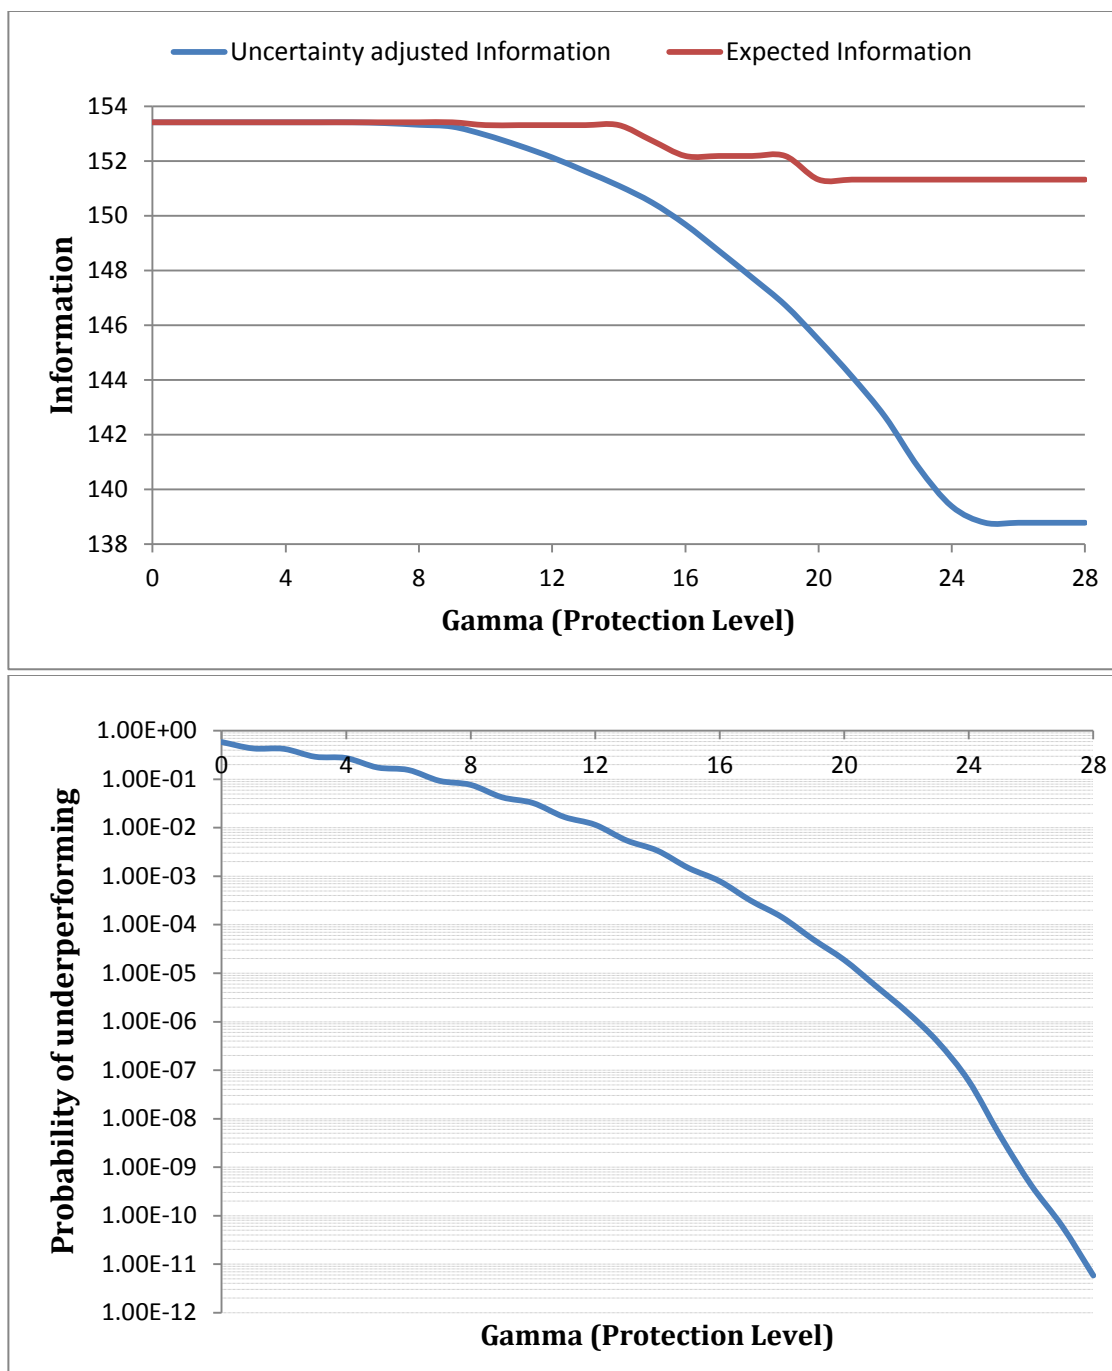
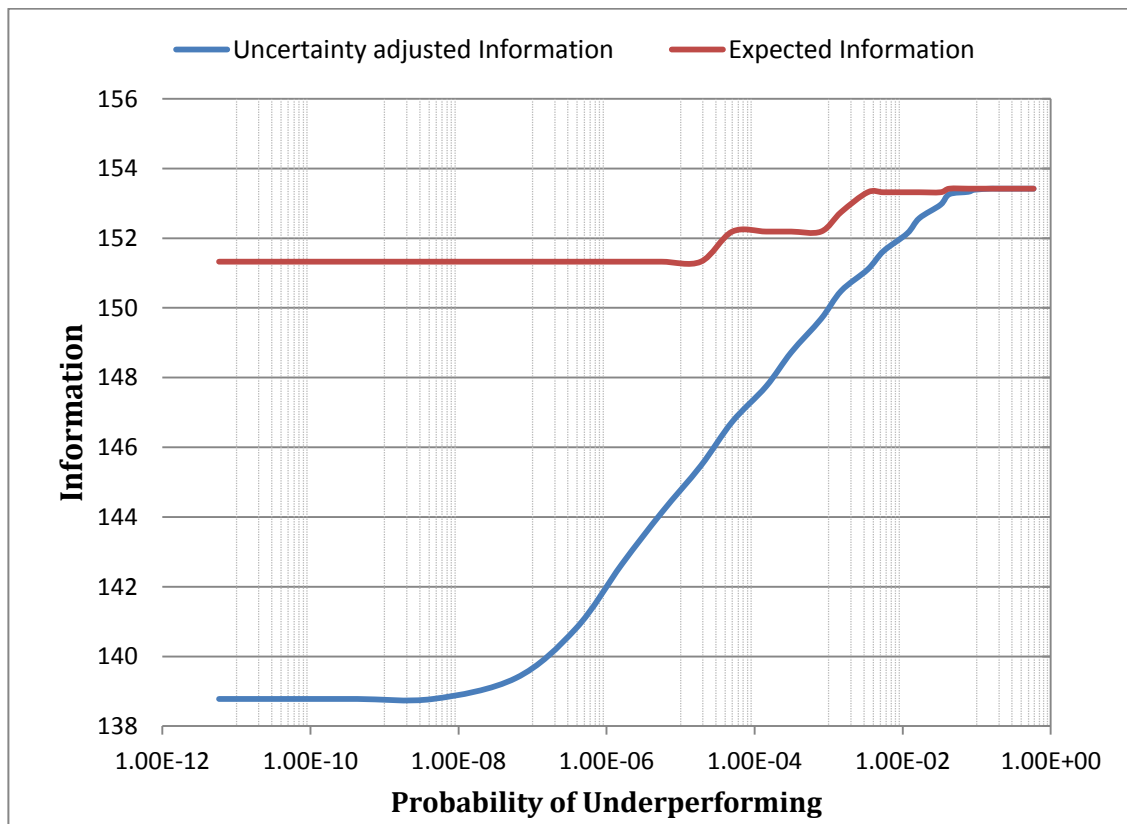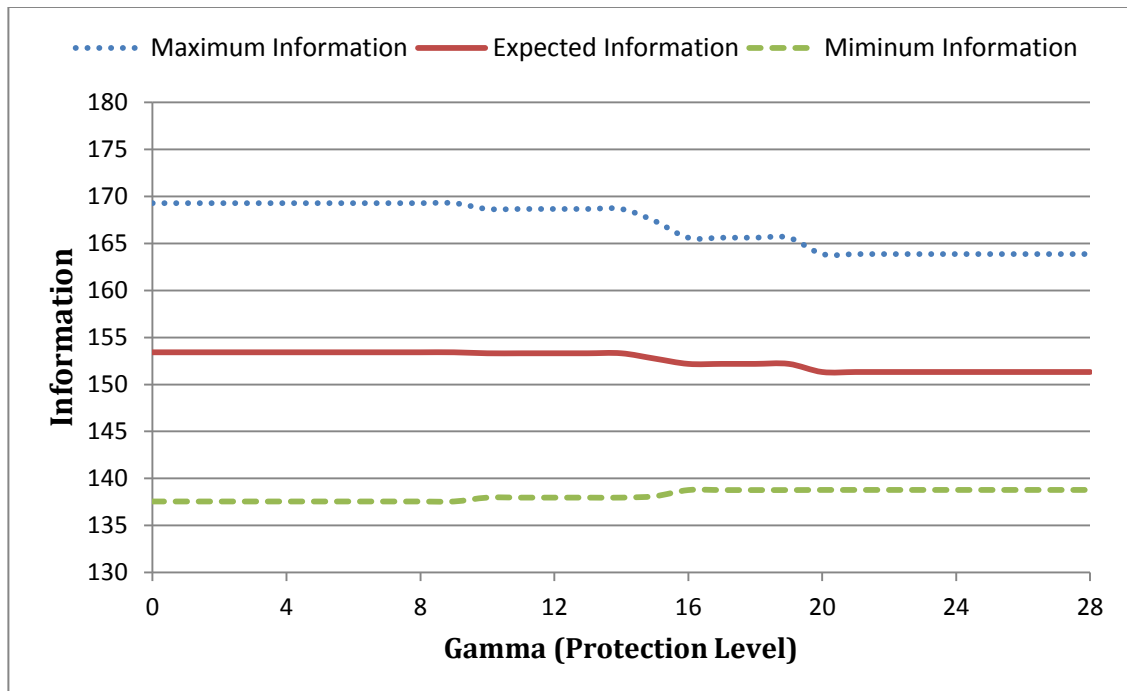Figure 6-7 (a) Information, (b) Probability of Underperforming vs. Protection level

Figure 6-8 (a) Minimum, Expected and Maximum information Flow vs. Protection Level (b) Information vs. Probability of Underperforming

In Figure 6-8 we present the aggregated result of the simulations indicating the trade-off between the portfolio return (information flow) and the risk (probability of underperforming). Figure 6-8(a) plots the values of the maximum, expected and minimum information flow against the protection level. As the value of Γ increases, the maximum and expected information flow values decreases, whilst the minimum values increase. Figure 6-8(b) shows the same results captured in Figure 6-8(a) of the expected information flow and the uncertainty-adjusted information, but against the associated risk, instead of Γ. Hence the figure accurately captures the tradeoff between risk and return.

The robust optimization formulation provides the center with a methodology to decide which sensor-target pair to select for transmission in a robust fashion that provides a deterministic guarantee of obtaining the highest gain in information for a given quantum of additional bandwidth. Moreover it endows the auction center with the flexibility to adjust the level of conservatism of the robust solutions in terms of probabilistic limits of underperformance of the selected sensor platforms for track data transmission.

6.3   Empirical Evaluation of Lack of Access to Outcome

We now turn our attention to the impact of the lack of access to outcome when determining the payments made to the platforms by the center. As we discussed in Section 4.2 for real-world scenarios operating in dynamic and uncertain environments, the state of the world changes between the time the information is reported and the time when the observation can be observed. Nevertheless in our simulation framework we can create a scenario where the center has access to information on all the tracks in the simulation. The center can use this information to evaluate the observations reported by the sensor platforms. This scenario allows us to examine the penalty that the center has to incur in the pre-selected sensor platform's payments as a consequence of its inaccessibility to the true outcome.

In
Figure 6-9 we calculate the total expected payment and the variance in the total payment made by the center for each of the four scoring rules. For the case of the parametric scoring rule, we restrict our attention to the parameter space $k = [1.1\,,1.5]$ we selected earlier in the chapter.

We only consider the case of $M = 1$, as we obtain the same trend in the results for a higher number of pre-selected platforms.

Figure 6-9(a) illustrates the two cases, the one where the center has access to the true outcomes and the other where the center uses the fused estimates to calculate the payments, using different markers. For both these cases, the total expected payment that the platforms expect to derive is the same. The lack of knowledge of the actual outcome of the mechanism has no impact on the expected payments. This result is valid for each of the four scoring rules – quadratic, spherical, logarithmic and parametric. This is in accordance with the economic property of the modified strictly proper scoring rules based mechanism design algorithm that was outlined in Section 4.2.3. The center does not incur any penalty on the expected payments for its lack of access to the true outcome. Papakonstantinou et al. (2011) have explained this seemingly anomalous behavior as a consequence of the modified strictly proper scoring rules mechanism itself. The property of incentive compatibility of the mechanism mandates that a pre-selected sensor platform's expected payment is equivalent to the cost used for scaling the scoring rules.

The lack of knowledge of the actual outcome does, however, have an impact on the variance of the total payment as demonstrated in
Figure 6-9(b). The variance of the payments is lower in the case where the center has access to the true outcomes as compared to the case where the center uses the fused estimates to calculate the payments. This is observed for all the different classes of scoring rules. The higher variance in the total payments can be attributed to the increased variance of the platforms' reported estimates. Thus the uncertainty introduced in the mechanism due to the lack of access to the real-world outcome introduces externalities in the payments made by the center. In spite of this drawback, the modified strictly proper scoring rules mechanism provides a robust framework which can adapt to the dynamic and evolving environment that characterizes tactical operations.
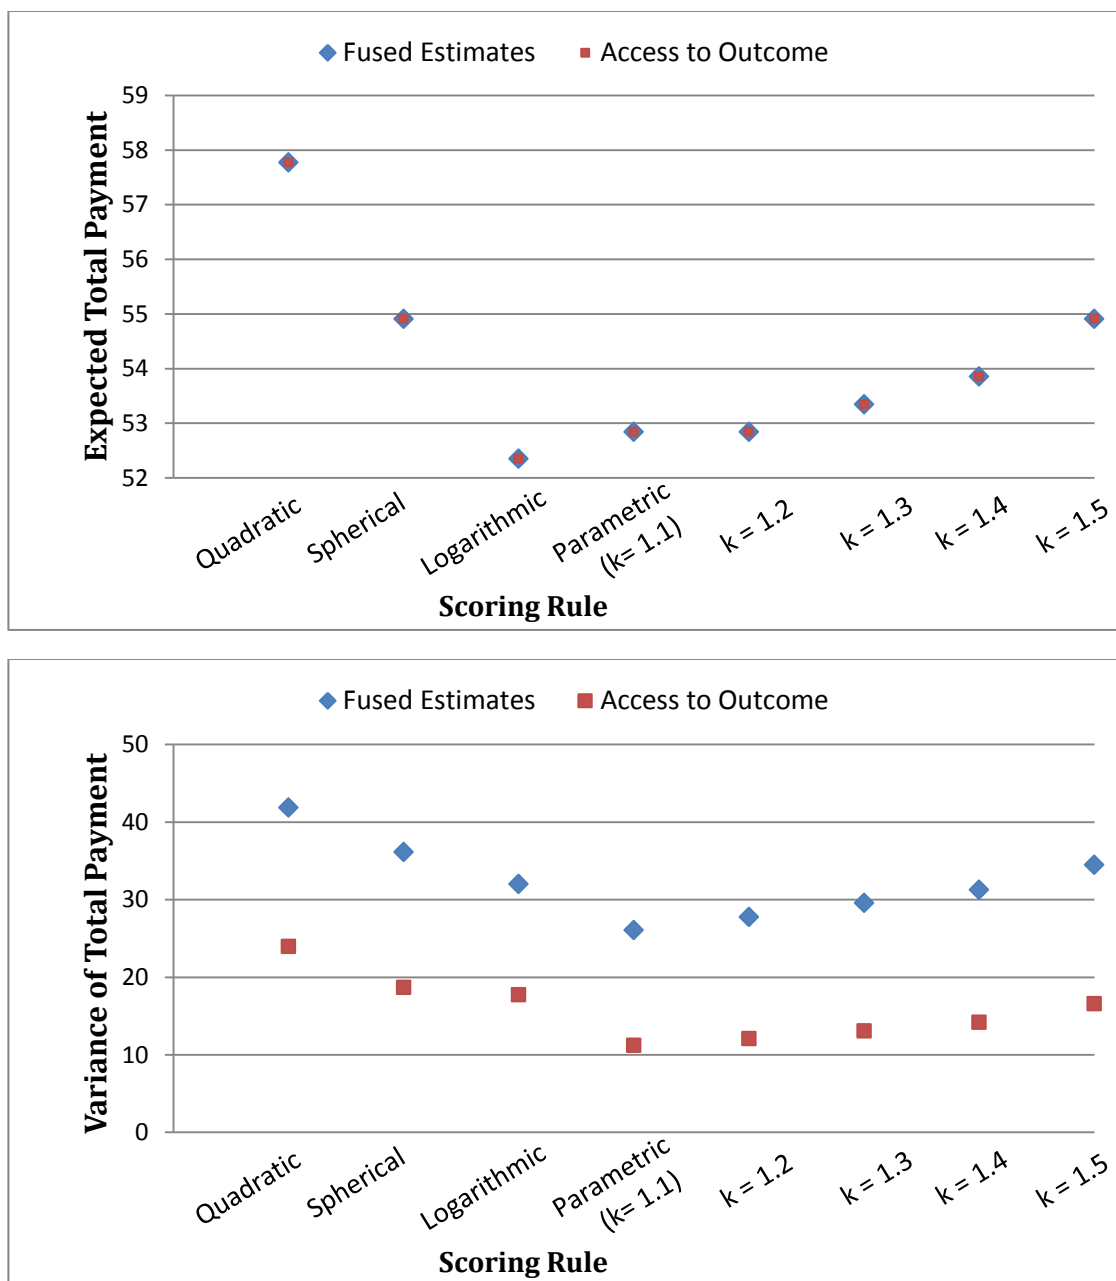
Figure 6-9 (a) Expected Total Payment and (b) variance of Total payment vs. Different Scoring Rules

6.4   Empirical Validation of Incentive Compatibility

In this section, we investigate the behavior of the mechanism when one of the sensor platforms resorts to deceptive behavior. Specifically we allow one of the sensor platforms to over-report its maximum information content values and the quality of its observed track data. This scenario allows us to extend our study of the robustness of the mechanism beyond information uncertainty, to cases when the underlying assumption of sensor platform rationality is violated.

In Section 4.2.3.1, a formal argument was provided for the incentive compatibility property of the modified strictly proper scoring rules in this mechanism. Thus an agent seeking to maximize its utility would truthfully report its maximum precision, and then produce an estimate of precision equal to its reported precision, and report the estimated precision truthfully ,assuming, of course that the sensor platform is rational. Notwithstanding the formal argument, it is instructive to study the stability of the mechanism when the underlying assumption of sensor platform rationality is violated, as it provides empirical validation of incentive compatibility.

In this set of scenarios, we assume that Sensor Platform 4 inflates the information valuation metric values for all the reported targets by a factor of 5. This would lead to the auctioneer awarding more tracks to Sensor platform 4, both R^2 and non - R^2. We compare the payments received by Sensor Platform 4 as well as the other platforms, with the corresponding truthful case when all sensor platforms honestly report their track data. The results of this comparison are shown in

Figure 6-10.

Figure 6-10 shows the payments received by the sensor platforms for both the truthful reporting and over-reporting scenarios, under different scoring rules. The over-reporting of the track data quality by Platform 4 ensures that it gets reporting responsibility for an increased number of targets. This would imply that the payment received by Platform 4 would decrease. However the results show that for all the four instances of the scoring rule, Platform 4 ends up receiving a negative payment, which is equivalent to making a payment. The over-reporting has no impact on Sensor Platform 3's payment as it acts as the Grid Reference Unit in both the cases. The GRU is assigned the reporting responsibility (R^2) for all tracks within its classification radius, and

neither receives nor makes any payments to the center. The other two sensor platforms in the mechanism also experience a decrease in the payment due to the deceptive behavior exhibited by Platform 4. This is an intuitive outcome as the lack of the access to the true outcome for the center results in each platform's payment being dependent on the observations of all the other platforms in the application framework, including the lying ones. The utility loss due to the dishonest reporting by the Platform 4 is presented in Table 6-1.

Table 6-1 Decrease in Utility Values for different scoring rules

	Decrease in Utility Values			
	Sensor Platform1	Sensor Platform 2	Sensor Platform 3	Sensor Platform 4
Quadratic	42.28	20.75	0	86.51
Spherical	41.40	19.80	0	83.89
Logarithmic	27.65	10.95	0	69.26
Parametric (k = 1.1)	29.34	12.65	0	70.64

An interesting point to note in Table 6.1 is that the drop in utility experienced by Platform 4 is much higher than what the other platforms experience. This is different from "spiteful bidding' which is a well - documented susceptibility of the Vickrey – Clarke - Groves mechanism (Brandt, Sandholm, & Shoham, 2005). Spiteful bidding occurs when an agent lies about the quality of its track data and accepts a decrease in its utility, if the utility of some other honest agent decreases more severely than its own. Our proposed scoring rule mechanism addresses this shortcoming of spiteful bidding by imposing higher penalties on the sensor platform which resorts to deceitful behavior.

The sensor platform 4 is worse-off in the over-reporting case than it would have been had it honestly reported its observations. This result serves as an empirical validation for the theoretical argument given for incentive compatibility in Chapter 4.
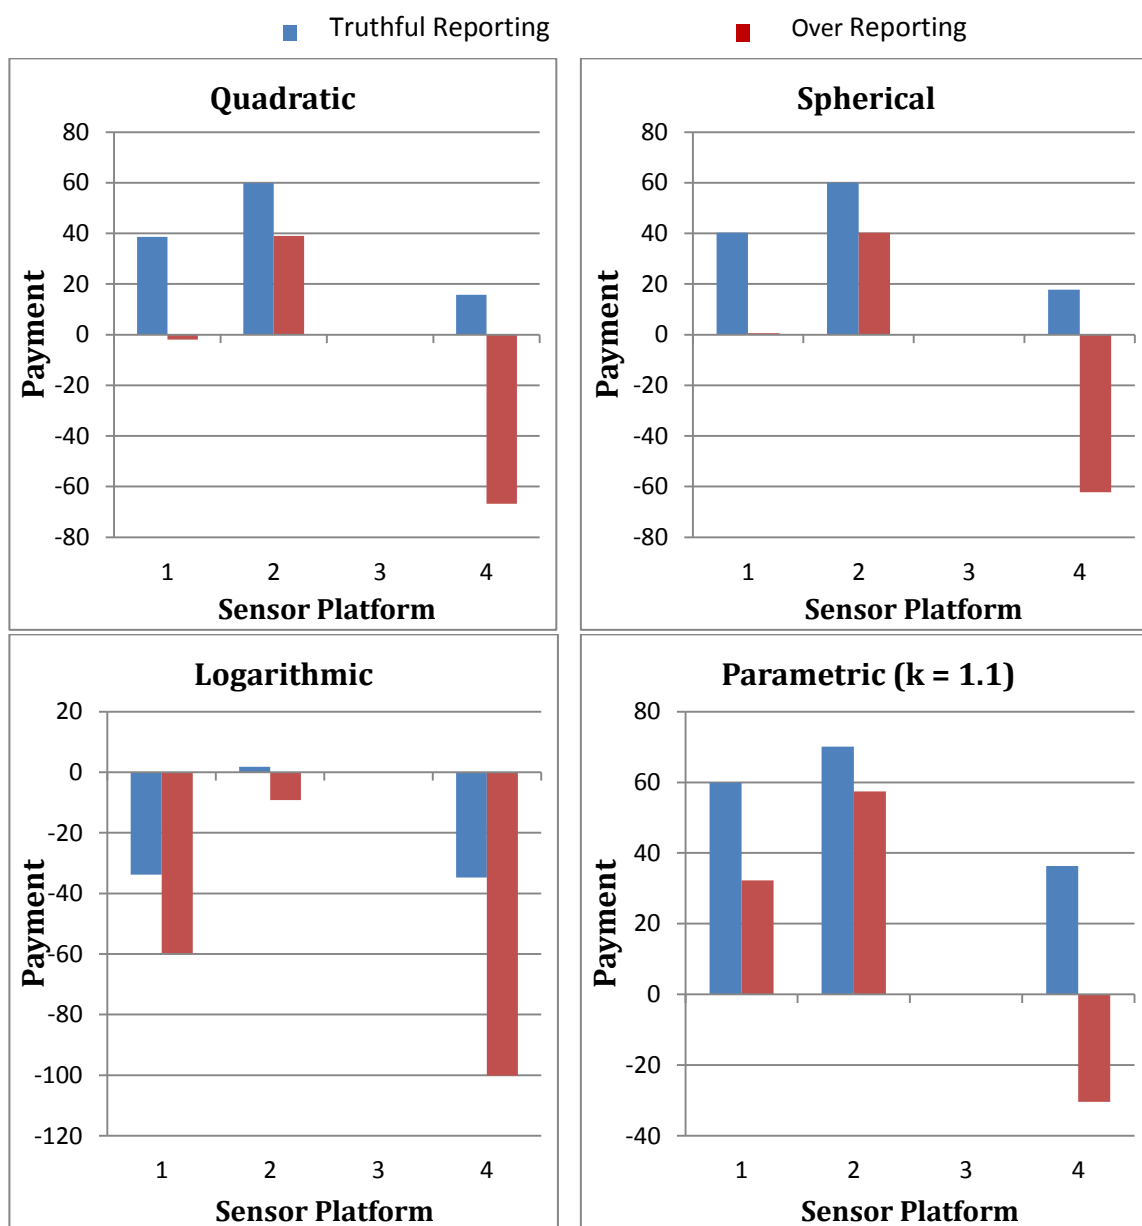
Figure 6-10 Payment to all sensor platforms for the case when Platform Sensor 4 over-reports

6.5 <u>Conclusions</u>

In conclusion, we have successfully applied our proposed robust-optimization based strictly proper-scoring rules algorithm to the research application framework. The standalone application framework boasts of significant dynamism and customizable features that allows for the generation of unique scenarios. Depending on the overall objectives of the mechanism, the designer can choose the number of pre–selected sensor platforms, the scoring rules and the protection level so as to best recover the most valuable information for a given quantum of extra bandwidth. We also show that the center does not incur any penalty on the expected payments for its lack of access to the true outcome and that it is in the best interest of the sensor platform to honestly produce an estimate of its maximum information content and report the estimated information content truthfully.

CHAPTER 7.   CONCLUSIONS AND FUTURE WORK

We summarize the contributions of the research towards developing a mechanism design framework for bandwidth allocation in tactical data networks in this chapter. We capture the highlights of the motivation behind the problem, the use of scoring rules within the mechanism design domain to address the problem, creating an agent-based application framework and the results of applying the robust optimization scoring rules mechanism to the research application framework. We highlight some of the other areas of application for our proposed mechanism, beyond the tactical data exchange environment. Finally, we discuss some possible avenues of future research to extend the scope of the current work.

7.1   Summary

The goal of this thesis was to explore the potential of using computational mechanisms to govern the behavior of ultra-large-scale systems and achieve an optimal allocation of constrained computational resources. We asserted that the advances in technology have fueled the race for information dominance in the defense sector, which relies on complex interconnected web of systems to meet its objectives. The sheer scale and size of these systems prompt behaviors that go beyond conglomerations of systems or 'system-of-systems', and provide the first glimpse of the Ultra Large Scale (ULS) systems of the future.

We highlighted how the dominant characteristics of these ULS systems challenge and undermine the fundamental assumptions of today's software and system engineering approaches. In most cases ULS systems will lack a central locus of institutional control and system users may behave opportunistically to meet their local mission requirements, rather than the goals of the system as a whole. We thus motivated the need to provide a basis for satisfying system-wide quality goals and simultaneously also satisfy the individual goals and expectations of the various stakeholders. In such cases, methods and tools based on.

economics and game theory will play an important role in achieving globally optimal behavior, even when the participants behave selfishly.

We introduced one such method Mechanism Design, which lies at the intersection of microeconomics and game theory. The field of mechanism design concerns itself with the design of protocols and institutions that are mathematically proven to satisfy certain system-wide objectives. It assumes that the individuals interacting through such institutions are capable of acting in a self-interested manner and may hold private information relevant to satisfying the system objective. Mechanism design has been widely used in problems involving the allocation of scarce resources among both human and computational entities which are inclined to resort to strategic behavior.

In our work, we focused on dynamic, performance-critical and resource-constrained systems of interest to the defense community. Specifically we considered a scenario where a number of military platforms have been tasked with the goal of detecting and tracking targets. Military platforms need to share and exchange tactical data from their onboard sensors, in order to establish and maintain a common operating picture of the tactical situation. This exchange of tactical data is facilitated over standardized radio networks, known as tactical data information links (TADILs), which have limited bandwidth. Our objective is to improve the quality and accuracy of the common operating picture through the efficient allocation of a finite amount of bandwidth in the tactical data networks. The problem is compounded by the possibility that the local goals of military platforms might not be aligned with the global system goal. Such a scenario might occur in multi-flag, multi-platform military exercises, where the military commanders of each platform are more concerned with the well-being of their own platform over others. Therefore there is a need to design a mechanism that efficiently allocates the flow of data within the network to ensure that the resulting global performance maximizes the information gain of the entire system, despite the self-interested actions of the individual actors. This research problem presents the kind of challenges we anticipate when we have to deal with ULS systems and by addressing this problem, we sought to develop a methodology which will be applicable for ULS systems of the future.

Information sharing is dictated by the Reporting Responsibility (R^2) rule which is used to select the military platform with the best quality data (position, velocity, etc.) to report a surveillance track on the TADIL. By its very nature, the Reporting Responsibility (R^2) rule is an extreme minimalist mechanism which precludes any possibility of collaboration by disallowing the redundant reporting of a single object. We emphasized the need for a mechanism which would allow for the recovery of the highest gain in information for a given quantum of additional bandwidth. The heart of this mechanism would resemble a portfolio optimization problem where the objective would be to select sensor – target pairs that maximize the information gain, given the bandwidth constraint. However our mechanism would need to go beyond a simple portfolio problem and handle the requirements of Incentive Compatibility, Individual Rationality, Interdependency, Information Uncertainty, Implementation and Lack of knowledge of the true state of the world.

In order to address the challenges of unknown costs, selfish behavior, constrained resources and dynamic environments, we decided to formulate the problem as a multi-agent systems (MAS) problem. We stressed on the need to generate a surrogate model for the real-world operation which exhibits the necessary fidelity and complexity to study the application of mechanism design in a practical context. To this end, we leveraged the Discrete – Agent Framework (DAF) developed at Purdue University in order to design a Multi-Agent System application framework. This framework allows us to capture the performance of a group of military platforms tasked with the goal of detecting and tracking targets. We emulated the elementary functionality of Link 16 communication protocol that is used as the tactical data network in the simulation.

When agents operate in evolving uncertain environments, it becomes imperative for trust to exist at the heart of all interactions between the agents. We discussed how trust can be addressed in multi-agent systems using two different prevalent approaches - at the individual level and at the system level. After considering learning and reputation trust models, socio-cognitive and probabilistic trust models, and various truth-elicitation protocols we concluded that they fail to adequately address all our research requirements. We then shifted our attention to the more fundamental approach of mechanism design which guarantees incentive compatibility (truthful reporting) and individual rationality (voluntary participation) through

certain payment schemes. We highlighted the drawbacks of the most popular auction-based mechanisms - the Vickrey-Clarke-Groves (VCG) mechanism. Hence we changed our focus from the realms of auction based models, to another promising alternative approach within the Mechanism Design research domain, in the form of strictly proper scoring rules.

Scoring Rules have been traditionally used to assess the accuracy of probabilistic forecasts, by awarding a score based on the forecast and the actual outcome. We restricted our attention to strictly proper scoring rules in which agents can maximize their score and payment by truthfully reporting their private observations or probabilistic estimates. Strictly proper scoring rules were scaled through the introduction of appropriate scaling parameters to ensure that agents are motivated to invest all their available resources in generating the observations.

We adopted a modified version of the two-stage mechanism based on continuous strictly proper scoring rules, proposed by Papakonstantinou et al. (2011). In the first stage, the trusted center (auctioneer) elicits the unknown costs of the agents and preselects a subset of agents that could provide the information at the lowest costs. In the next stage the preselected agents are induced to reveal their observations, using a payment scheme based on the fused reported estimates rather than the true outcome. The mechanism is individually rational and incentive compatible and works in uncertain and dynamic environments, where the costs of generating observations are private values unknown to the center and the true outcome cannot be observed by either the agents or the center. Through appropriately scaled and modified strictly proper scoring rules, the mechanism ensures that an agent's payment is dependent on the accuracy of its reported observations, and hence guarantees that the agent will invest all its resources in generating the observations. We tackled the problem of uncertainty in the information content of the observations by formulating it as a robust portfolio optimization problem. We adopted the Bertsimas-Sim linear framework to solve the portfolio problem, which provides a protection level parameter Γ to control the degree of robustness of the solution.

We thus designed a mechanism that ensures that each sensor platform truthfully reports their track information, based on which the auctioneer can allocate targets to the sensor agents by taking into account the uncertainty in the information content, and the selected sensor agents

invest their resources to generate and truthfully report the observations on the allocated targets. We applied our modified scaled strictly proper scoring rules based mechanism to the ABM framework to study the behavior of the mechanism under different settings.

We computed the transmitted information and the corresponding Network Cycle Time (NCT), as the maximum number of pre-selected agents (M) is varied from 1 to 3. We showed that the choice of the parameter M is dictated by a tradeoff between information content and information latency as well as a tradeoff between the total expected payment and the marginal information gain per second of NCT. We asserted that it is prudent to select the highest value of M which ensures that the mechanism captures the highest possible gain in the overall common operating picture quality, even if it results in the center incurring the highest expected payment.

We empirically evaluated the parametric soring rule and compared it to the quadratic, spherical and parametric scoring rule, for the parameter space of $k = (1,3)$. We discussed how the parametric scoring rule converges to one of the other scoring rules for different values of the parameter k. The logarithmic scoring rule resulted in the lowest expected payment and payment variance, but it did not have a finite lower bound on the payment. The parametric scoring rule shared the same advantages and disadvantages of logarithmic scoring rule, but the effect of the platform's imprecise estimate could be minimized for the parametric scoring rule family by choosing the parameter k carefully. We argued that a parameter value of $k = [1.1, 1.5]$ appears a judicious compromise as it retains the low expected payments and variance of payments of the logarithmic scoring rule, and at the same time, imposes a finite lower bound on the minimum payments.

We also analyzed the effect of the protection level (Γ) on the robust optimization framework. A higher value of Γ provides a lower uncertainty adjusted information value with correspondingly lower risks while a lower value of Γ provides higher returns with more risk. The robust optimization formulation provides the center with a methodology to decide which sensor-target pair to select for transmission and also provides a deterministic guarantee of obtaining the highest gain in information for a given quantum of additional bandwidth.

The impact of the lack of access to the outcome when determining the payments was also determined. We proved empirically that the center does not incur any penalty on the expected payments for its lack of access to the true outcome. The lack of knowledge of the actual outcome does, however, have an impact on the variance of the total payment. The variance of the payments is lower in the case where the center has access to the true outcomes as compared to the case where the center uses the fused estimates to calculate the payments.

We extended our study of the robustness of the mechanism beyond information uncertainty, to cases when the underlying assumption of agent rationality is violated. Specifically we allowed one of the sensor platforms to over-report its maximum information content values and the quality of its observed track data. We showed that when a sensor platform over-reports its information content it experiences a decrease in utility and that a sensor platform is always better off by honestly reporting its observations. This result served as an empirical validation of the incentive compatibility property of the mechanism.

7.2   Applications

Although we have studied the application of computational mechanism design for bandwidth allocation in tactical data links, the empirical aspects of the research work suggest that the applicability of the proposed mechanism goes beyond tactical data links. In fact, the intended goal of this work was to study the potential of using computational mechanism design to develop a methodology to handle the anticipated challenges of ULS system of the future. The proposed framework is envisioned as application agnostic and with minor tweaks, should be amenable to any setting which involves exchange of information or services between buyers and sellers. Some of these settings include environmental monitoring sensor networks, citizen sensor networks, electronic product contract manufacturing and e-commerce applications.

Sensor Networks have been touted as one of the most promising technologies for the next few decades (Chong & Kumar, 2003; Lesser, Ortiz, & Tambe, 2003; Srivastava, Culler, & Estrin, 2004). The emergence of small inexpensive sensors based on Microelectromechanical system (MEMS) has resulted in a boom in the development of small experimental sensor networks. These networks have found ubiquitous use for a wide range of real world applications which include

monitoring remote or hostile environments (Martinez, Hart, & Ong, 2004), predicting floods in river estuaries (De Roure, 2005), radiation monitoring (Nemzek, Dreicer, Torney, & Warnock, 2004), etc. One example of such applications is the GlacsWeb project which involves the use of hundreds of tiny battery powered sensors embedded into the ice of the Briksdalsbreen glacier in Norway (Martinez et al., 2004). These sensors are capable of measuring pressure, temperature and orientation of the glacier ice. A base-station polls the sensors at fixed intervals and each sensor forwards their measurements using low power radio-based transmission. The data gets aggregated further at reference stations and sent over standard internet protocols for analysis. The sensors have a partial and inaccurate view of their operating environment and face constrained power and bandwidth capabilities. In such hostile environments, neither the base nor the reference stations have access to the true outcome or the ground truth. There is a possibility of selfish behavior emerging among these sensors where the local goals of preserving their own battery power may conflict with the global objective of information aggregation. Again, a trade-off between taking measurements and communicating these measurements is bound to surface.

Another increasingly common trend is the emergence of the so-called "smart cities" which leverage the information and communication network infrastructure to take an edge in urban competitiveness. Wireless sensor networks, especially citizen sensor networks, are one of the few technologies that are fuelling smart cities. Citizen sensor networks allow citizens to register and connect their mobile devices, in real-time, to feed data into open online information databases (Goodchild, 2007; Sheth, 2009). The information can range from observations of actual events like measuring noise or environmental parameters to probabilistic estimates of weather or traffic forecasts. Examples include Xively, NoiseTube, Wikisensing which facilitate online collaboration between users (Maisonneuve, Stevens, Niessen, & Steels, 2009; Silva, Ghanem, & Guo, 2012; Yamanoue, Oda, & Shimozono, 2013). The developers or information-buyers can connect to these databases and build apps based on the data. As high speed mobile internet and smartphones proliferate, the citizens sensor networks are expected to find increasingly commercial applications. Thus we can expect the information providers, whose contributions are currently voluntary and altruistic, to take a more self-interested turn. In fact this trend can already be observed in traffic monitoring services like Inrix and TrafficCast which

aggregate the information from citizen sensor networks and sell it at a fee (Sadri & Gossain, 2010; Schrank, Lomax, & Turner, 2010). Citizen sensor networks foreshadow some of the key characteristics of ULS systems and there is a high probability of participants exhibiting self-interested behavior as citizen sensor networks become commercialized.

Another field where Mechanism Design has found significant application is that of e-commerce. Given that our scoring rules mechanism handles the exchange of information between buyers and sellers, it is easily extendable to e-commerce applications. Online forums, markets, and communities typically employ ratings–based or rankings–based mechanisms for evaluating items or services. The large scale nature of such systems, the inherent competitiveness and the financial opportunities involved makes the emergence of selfish behavior among users a realistic possibility. Another field of interest is the electronic product contract manufacturing which has become a US $100 billion business worldwide (Deshpande, Schwarz, Atallah, Blanton, & Frikken, 2011). Mechanism Design is used to facilitate the selection of contract manufacturers by Original Equipment Manufacturer (OEM), and to determine which party is in charge of procurement of the individual components that make up the OEM's product. Given the self-interested nature of the participants, the lack of access to the true costs and the uncertainty involved in development schedules and budgets, we could employ the scoring-rules mechanism within the Secure Price Masking (SPM) algorithm to address these highlighted issues.

7.3   Future Work

In this section we highlight a few areas of future research which represent a small sample of the various avenues that can be explored within the context of scoring rules in mechanism design. Some of these lines of research are critical to ensure wider applicability of the mechanism to address similar problems in some of the application areas discussed in the previous section.

1. **Auctioneer Trust**

Apart from open-auctions like the Dutch and the English auctions, the various auctions discussed in Chapter 2 (e.g. the second-price auction and the VCG auction) are susceptible to manipulation on the part of the auctioneer. In the first stage of the proposed mechanism, the center conducts a reverse $(M + 1)$ price auction to choose the set of pre-selected agents and

then uses the $(M + 1)^{th}$ cost in order to determine the scaling parameters for the different scoring rules. Since the cost functions represent the private information known only to each of the sensor platforms, the auctioneer can exploit its knowledge of the cost functions to take advantage of the mechanism. Hypothetically speaking, the center may use a lower cost function that the $(M + 1)^{th}$ cost when designing the scoring rules and calculating the payments. This is similar to the VCG auction we discussed earlier where the auctioneer could ask for a higher price than the second-highest bid from the auction winner. In Chapter 2, we discussed two approaches that have been proposed in literature to address this particular shortcoming. The first was the public blackboard mechanism proposed by M. Hsu & Soo (2002) which we deemed as unsuitable for our research problem, as it assumed that information sellers can also act as information buyers. Yokoo & Suzuki (2004) proposed the use of encrypted protocols for the transmission of information between the agents and the center, which introduces significant latency in the communication as it involves multiple overlays of calculations and renders it unsuitable for real-time applications. Brandt (2001) and Brandt & others (2002) devised a collusion proof auction mechanism which ensures the privacy and correctness of any $(M + 1)^{th}$ price auction by distributing the calculation of the selling price between the individual buyers using some cryptographic techniques. Although the protocol is computationally expensive for a large number of agents, it could prove a good starting point for future research. Thus extending the notion of trust within the mechanism, from the agents to the center, could count as one of the possible areas of future work.

2. Virtual Currency

The scoring-rules based mechanism rewards each platform for transmitting information on the tracks in the scenario. The payment is in terms of virtual currency and there is a need for understanding the full implications of the use of virtual currencies as a means to incentivize the agents to be truthful in the observations.

"We are interested in social engineering for machines. We want to understand the kinds of negotiation protocols, and punitive and incentive mechanisms that would cause individual designers to build machines that act in particular ways. Since we assume that the agents' designers are basically interested in their own goals, we want to find interaction techniques

that are 'stable', that make it worthwhile for the agent designer not to have

machine deviate from the target behavior."

- Rosenschein & Zlotkin (1995)

In our research application framework of military group operations, there are two primary sources of incentives:

- **Utility:** The military platforms are modeled as rational agents seeking to maximize their own utilities. The payments made to the military platforms in the mechanism must have some real-world implications. This can take the form of after-action rewards like promotion or paid vacation, to motivate the platforms to invest all the resources at their disposal in generating the observations.

- **Survivability:** The survival of an individual military platform is inextricably linked to the survival of the entire group. This inherent logic also serves as a source of incentive for each platform to invest its resources in tracking the targets assigned to it and then reporting these observations truthfully.

We need to understand the role of virtual currency in the context of social institutions and to elicit user preferences in order to ensure that the payment scheme embedded in the mechanism is practical and proper.

3. Dynamic Environments

There is a need to evaluate the applicability of the mechanism to more dynamic application frameworks. For our research application, we treated the problem as a sequence of static events changing from cycle-to-cycle. Use case scenarios might arise where more than one network cycle needs to be taken into account for planning across multiple periods. We did not account for situations where readings from past cycles are good approximations for the current cycle which might lead to multiple simultaneous network cycles, and adding extra complexity to the problem. Further research is necessary to assess the applicability of the mechanism in such environments.

4. Bidder Collusion

We argued that ULS systems operate in dynamic and uncertain environment that is constantly evolving. The state of the world changes between the time the information is reported and the time when the observation can be evaluated. In the absence of access to the true outcome, the auction center relies on the fused estimate obtained from the agents themselves in order to evaluate their reported observations. This makes the mechanism susceptible to collusion among the agents. Jurca & Faltings (2007) showed that bidder collusion is an expected outcome of any mechanism which calculates the payments made to an agent based on other agents' reported observations. They propose introducing a small group of agents which will always be truthful and thus prevent collusion among agents, eliminating the undesired Nash equilibria. We showed the infeasibility of their solution within the context of our research problem as it introduces additional complexity and contradicts the envisioned use of the mechanism in networks where the center has no external means of evaluating the agent's observations. Designing an incentive –compatible mechanism which is robust to collusion among agents is a promising area for future research.

LIST OF REFERENCES

LIST OF REFERENCES

Anderson, E., Kelly, F., & Steinberg, R. (2006). A contract and balancing mechanism for sharing capacity in a communication network. *Management science*, *52*, 39–53.

Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., & Roughgarden, T. (2004). The price of stability for network design with selfish agents. *IEEE Symposium on Foundations of Computer Science* (pp. 295–304).

Arrow, K. J. (1977). *The property rights doctrine and demand revelation under incomplete information*. Institute for Mathematical Studies in the Social Sciences, Stanford University.

Axelrod, R. (1980). More effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, *24*, 379–403.

Axtell, R. (2000). Why agents?: on the varied motivations for agent computing in the social sciences.

Beale, E. (1955). On minimizing a convex function subject to linear inequalities. *Journal of the Royal Statistical Society. Series B (Methodological)*, 173–184.

Ben-Tal, A., & Nemirovski, A. (1998). Robust convex optimization. *Mathematics of Operations Research*, *23*, 769–805.

Ben-Tal, A., & Nemirovski, A. (1999). Robust solutions of uncertain linear programs. *Operations research letters*, *25*, 1–13.

Ben-Tal, A., & Nemirovski, A. (2000). Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming*, *88*, 411–424.

Bertsimas, D., & Sim, M. (2004). The price of robustness. *Operations research*, *52*, 35–53.

Bertsimas, D., & Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations Research*, *54*, 150–168.

Bigley, G. A., & Pearce, J. L. (1998). Straining for shared meaning in organization science: Problems of trust and distrust. *Academy of Management Review*, *23*, 405–421.

Blanchette Jr, S. (2005). *US Army Acquisition-The Program Executive Officer Perspective*. DTIC Document.

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, *99*, 7280–7287.

Borshchev, A., & Filippov, A. (2004). From system dynamics and discrete event to practical agent based modeling: reasons, techniques, tools. *Proceedings of the 22nd international conference of the system dynamics society*.

Bousquet, F., & Le Page, C. (2004). Multi-agent simulations and ecosystem management: a review. *Ecological modelling*, *176*, 313–332.

Brandt, F. (2001). Cryptographic protocols for secure second-price auctions. *Cooperative Information Agents V* (pp. 154–165). Springer.

Brandt, F., & others. (2002). A verifiable, bidder-resolved auction protocol. *Proceedings of the 5th International Workshop on Deception, Fraud and Trust in Agent Societies (Special Track on Privacy and Protection with Multi-Agent Systems)* (pp. 18–25). Citeseer.

Brandt, F., Sandholm, T., & Shoham, Y. (2005). Spiteful Bidding in Sealed-Bid Auctions. *Computing and Markets*.

Braynov, S., & Sandholm, T. (2002). Contracting with uncertain level of trust. *Computational Intelligence*, *18*, 501–514.

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly weather review*, *78*, 1–3.

Castelfranchi, C., & Falcone, R. (1998). Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. *Multi Agent Systems, 1998. Proceedings. International Conference on* (pp. 72–79). IEEE.

Castelfranchi, C., & Falcone, R. (2000). Trust is much more than subjective probability: Mental components and sources of trust. *System Sciences, 2000. Proceedings of the 33rd Annual Hawaii International Conference on* (p. 10–pp). IEEE.

Casti, J. (1997). Would-be worlds: how simulation is changing the world of science. New York: Wiley.

Cavallo, R. (2006). Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems* (pp. 882–889). ACM.

Chen, M., Yang, G., & Liu, X. (2004). Gridmarket: A practical, efficient market balancing resource for Grid and P2P computing. *Grid and Cooperative Computing* (pp. 612–619). Springer.

Chen, R. R., Roundy, R. O., Zhang, R. Q., & Janakiraman, G. (2005). Efficient auction mechanisms for supply chain procurement. *Management Science*, *51*, 467–482.

Chong, C.-Y., & Kumar, S. P. (2003). Sensor networks: evolution, opportunities, and challenges. *Proceedings of the IEEE*, *91*, 1247–1256.

Church, N. I., Dallal, H. J., Masson, J., Mowat, N. A. G., Johnston, D. A., Radin, E., Turner, M., et al. (2006). Validity of the Rockall scoring system after endoscopic therapy for bleeding peptic ulcer: a prospective cohort study. *Gastrointestinal endoscopy*, *63*, 606–612.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public choice*, *11*, 17–33.

Conway, J. (1970). The game of life. *Scientific American*, *223*, 4.

Cruz, C. I. (2004). *Netwars Based Study of a Joint Stars Link-16 Network*. DTIC Document.

D Aspremont, C., & Gérard–Varet, L.-A. (1982). Bayesian incentive compatible beliefs. *Journal of Mathematical Economics*, *10*, 83–103.

Dang, V. D., Dash, R. K., Rogers, A., & Jennings, N. R. (2006). Overlapping coalition formation for efficient data fusion in multi-sensor networks. *AAAI* (Vol. 6, pp. 635–640).

Dantzig, G. B. (1955). Linear programming under uncertainty. *Management science*, *1*, 197–206.

Dantzig, G. B. (1991). Linear programming. *History of Mathematical Programming*, 19–31.

Dasgupta, P. (2000). Trust as a commodity. *Trust: Making and breaking cooperative relations*, *4*, 49–72.

Dasgupta, P., & Maskin, E. (2000). Efficient auctions. *The Quarterly Journal of Economics*, *115*, 341–388.

De Finetti, B. (1937). La prevision: ses lois logiques, ses sources subjectives. *Annales de l'institut Henri Poincaré* (Vol. 7, pp. 1–68). Presses universitaires de France.

De Finetti, B. (1962). Does it make sense to speak of good probability appraisers. *The Scientist Speculates: An Anthology of Partly-Baked Ideas*. New York: Wiley.

De Groot, M. H. (1970). Optimal statistical decisions.

De Roure, D. (2005). Floodnet: A new flood warning system. *Royal Academy of Engineering Quarterly*, *23*, 48–51.

DeGroot, M. H., & Schervish, M. J. (2002). Student solutions manual for probability and statistics.

DeLaurentis, D. (2005). Understanding transportation as a system-of-systems design problem. *43rd AIAA Aerospace Sciences Meeting and Exhibit* (Vol. 1).

Dellarocas, C. (2002). Towards incentive-compatible reputation management. *Workshop on Deception, Fraud and Trust in Agent Societies, Bologna, Italy* (pp. 26–40).

Deshpande, V., Schwarz, L. B., Atallah, M. J., Blanton, M., & Frikken, K. B. (2011). Outsourcing Manufacturing: Secure Price-Masking Mechanisms for Purchasing Component Parts. *Production and Operations Management*, *20*, 165–180.

Duffie, D., & Pan, J. (1997). An overview of value at risk. *The Journal of derivatives*, *4*, 7–49.

Echternacht, G. J. (1972). The use of confidence testing in objective tests. *Review of Educational Research*, *42*, 217–236.

Edelman, B., & Ostrovsky, M. (2007). Strategic bidder behavior in sponsored search auctions. *Decision support systems*, *43*, 192–198.

El Ghaoui, L., & Lebret, H. (1997). Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, *18*, 1035–1064.

El Ghaoui, L., Oustry, F., & Lebret, H. (1998). Robust solutions to uncertain semidefinite programs. *SIAM Journal on Optimization*, *9*, 33–52.

Falcone, R., & Castelfranchi, C. (2001). Social trust: A cognitive approach. *Trust and deception in virtual societies* (pp. 55–90). Springer.

Federico, G., & Rahman, D. (2003). Bidding in an electricity pay-as-bid auction. *Journal of Regulatory Economics*, *24*, 175–211.

Ferber, J. (1999). *Multi-agent systems: an introduction to distributed artificial intelligence* (Vol. 1). Addison-Wesley Reading.

Ferguson, A. R., & Dantzig, G. B. (1954). *The Problem of Routing Aircraft, a Mathematical Solution*. DTIC Document.

Franklin, S., & Graesser, A. (1997). Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents. *Intelligent agents III agent theories, architectures, and languages* (pp. 21–35). Springer.

Gambetta, D. (1990). Trust: Making and breaking cooperative relations.

Garratt, A., Lee, K., Pesaran, M. H., & Shin, Y. (2003). Forecast Uncertainties in Macroeconomic Modeling. *Journal of the American Statistical Association*, *98*.

Garthwaite, P. H., Kadane, J. B., & O'hagan, A. (2005). Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association*, *100*, 680–701.

Gerkey, B. P., & Mataric, M. J. (2002). Sold!: Auction methods for multirobot coordination. *Robotics and Automation, IEEE Transactions on*, *18*, 758–768.

Giachetti, R. E., Marcelli, V., Cifuentes, J., & Rojas, J. A. (2013). An agent-based simulation model of human-robot team performance in military environments. *Systems Engineering*, *16*, 15–28.

Gneiting, T., & Raftery, A. E. (2004). *Strictly proper scoring rules, prediction*. and estimation. Technical Report 463, Department of Statistics, University of Washington.

Gneiting, T., & Raftery, A. E. (2005). Weather forecasting with ensemble methods. *Science*, *310*, 248–249.

Good, I. J. (1952). Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 107–114.

Goodchild, M. F. (2007). Citizens as sensors: the world of volunteered geography. *GeoJournal*, *69*, 211–221.

Grossman, S. J., & Hart, O. D. (1983). An analysis of the principal-agent problem. *Econometrica: Journal of the Econometric Society*, 7–45.

Groves, T. (1973). Incentives in teams. *Econometrica: Journal of the Econometric Society*, 617–631.

Hanson, R. (2002). Logarithmic market scoring rules for modular combinatorial information aggregation. *George Mason University*.

Hinz, J. (2003). Optimal bid strategies for electricity auctions. *Mathematical Methods of Operations Research*, *57*, 157–171.

Holstein, C.-A. S. Staël von. (1972). Probabilistic forecasting: An experiment related to the stock market. *Organizational Behavior and Human Performance*, *8*, 139–158.

Holzman, R., & others. (2003). Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences*, *46*, 193–205.

Hsu, J., Price, M., Clymer, J. R., Garcia-Jr, J., & Gonzalez, E. (2009). Agent-Based Modeling the Emergent Behavior of a System-of-Systems (pp. 1–1).

Hsu, M., & Soo, V. (2002). A secure multi-agent Vickrey auction scheme. *Workshop on Deception, Fraud and Trust in Agent Societies, AAMAS 2002, Bologna, Italy* (pp. 86–91).

Huhns, M. N., & Stephens, L. M. (1999). 2 Multiagent Systems and Societies of Agents.

Hurwicz, L. (1973). The design of mechanisms for resource allocation. *The American Economic Review*, *63*, 1–30.

Ismail, R., & Josang, A. (2002). The beta reputation system.

Jay, F. (1958). Industrial dynamics: a major breakthrough for decision makers. *Harvard Business Review*, *36*, 37–65.

Jehiel, P., & Moldovanu, B. (2001). Efficient design with interdependent valuations. *Econometrica*, *69*, 1237–1259.

Jiang, B., & Gimblett, H. R. (2002). An agent-based approach to environmental and urban systems within geographic information systems. *Integrating Geographic Information Systems and Agent-based Modeling Technniques for Simulating Social and Ecological Processes, Santa Fe Studies in the Sciences of Complexity*, 171–190.

Jones, K., Janicke, H., & Cau, A. (2009). A property based framework for trust and reputation in mobile computing. *Advanced Information Networking and Applications Workshops, 2009. WAINA'09. International Conference on* (pp. 1031–1036). IEEE.

Jurca, R., & Faltings, B. (2003a). Towards incentive-compatible reputation management. *Trust, Reputation, and Security: Theories and Practice* (pp. 138–147). Springer.

Jurca, R., & Faltings, B. (2003b). An incentive compatible reputation mechanism. *E-Commerce, 2003. CEC 2003. IEEE International Conference on* (pp. 285–292). IEEE.

Jurca, R., & Faltings, B. (2005). Reputation-based service level agreements for web services. *Service-Oriented Computing-ICSOC 2005* (pp. 396–409). Springer.

Jurca, R., & Faltings, B. (2007). Collusion-resistant, incentive-compatible feedback payments. *Proceedings of the 8th ACM conference on Electronic commerce* (pp. 200–209). ACM.

Kilicay-Ergin, N., & Dagli, C. (2008). Executable modeling for system of systems architecting: An artificial life framework. *Systems Conference, 2008 2nd Annual IEEE* (pp. 1–5). IEEE.

Klein, M., Moreno, G. A., Parkes, D. C., Plakosh, D., Wallnau, K., & Seuken, S. (2008). Handling interdependent values in an auction mechanism for bandwidth allocation in tactical data networks. *Proceedings of ACM SIGCOMM 2008 Workshop on Economics of Networked Systems (NetEcon 2008)*.

Krishna, V. (2002). Auction theory. 2002. Academic press.

Langton, C. G. (1989). Artificial life: the proceedings of an Interdisciplinary Workshop on the

Synthesis and Simulation of Living Systems, held September, 1987. *Los Alamos, New*

Mexico. Addison-Wesley Publishing Company.

Lesser, V., Ortiz, C. L., & Tambe, M. (2003). *Distributed sensor networks: A multiagent*

perspective (Vol. 9). Springer.

Lewis, J. D., & Weigert, A. (1985). Trust as a social reality. *Social forces*, *63*, 967–985.

Maier, M. W. (1998). Architecting principles for systems-of-systems. *Systems Engineering*, *1*,

267–284.

Maisonneuve, N., Stevens, M., Niessen, M. E., & Steels, L. (2009). NoiseTube: Measuring and

mapping noise pollution with mobile phones. *Information Technologies in*

Environmental Engineering (pp. 215–228). Springer.

Marla, L. (2007). *Robust optimization for network-based resource allocation problems under*

uncertainty. Massachusetts Institute of Technology.

Martinez, K., Hart, J. K., & Ong, R. (2004). Environmental sensor networks. *Computer*, *37*, 50–56.

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust.

Academy of management review, *20*, 709–734.

McCabe, K. A., Rassenti, S. J., & Smith, V. L. (1990). Auction institutional design: Theory and

behavior of simultaneous multiple-unit generalizations of the Dutch and English

auctions. *The American Economic Review*, *80*, 1276–1283.

McClelland, A. G., & Bolger, F. (1994). The calibration of subjective probability: Theories and

models 1980–94.

McKnight, D. H., & Chervany, N. L. (2002). What trust means in e-commerce customer relationships: an interdisciplinary conceptual typology. *International journal of electronic commerce*, *6*, 35–60.

McMillan, J. (1994). Selling spectrum rights. *The Journal of Economic Perspectives*, *8*, 145–162.

Mezzetti, C. (2004). Mechanism design with interdependent valuations: Efficiency. *Econometrica*, *72*, 1617–1626.

Miller, N. H., Pratt, J. W., Zeckhauser, R. J., & Johnson, S. (2007). Mechanism design with multidimensional, continuous types and interdependent valuations. *Journal of Economic Theory*, *136*, 476–496.

Minar, N., Burkhart, R., Langton, C., & Askenazi, M. (1996). The swarm simulation system: A toolkit for building multi-agent simulations. Santa Fe Institute Santa Fe.

Mishra, A. K. (1996). Organizational responses to crisis: The centrality of trust.

Mui, L., Mohtashemi, M., & Halberstadt, A. (2002). Notions of reputation in multi-agents systems: a review. *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1* (pp. 280–287). ACM.

Myerson, R. B., & Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of economic theory*, *29*, 265–281.

Nagel, K., & Rasmussen, S. (1994). Traffic At the Edge of Chaos.

Nemzek, R. J., Dreicer, J. S., Torney, D. C., & Warnock, T. T. (2004). Distributed sensor networks for detection of mobile radioactive sources. *Nuclear Science, IEEE Transactions on*, *51*, 1693–1700.

Neumann, J. von, & Morgenstern, O. (1944). Theory of Gamesand Economic Behavior. Princeton, NJ: Princeton University Press.

Northrop, L., Feiler, P., Gabriel, R. P., Goodenough, J., Linger, R., Longstaff, T., Kazman, R., et al. (2006). {Ultra-Large-Scale Systems}-The Software Challenge of the Future.

Odell, J., Nodine, M., & Levy, R. (2005). A metamodel for agents, roles, and groups. *Agent-Oriented Software Engineering V* (pp. 78–92). Springer.

Palmer, T. (2002). The economic value of ensemble forecasts as a tool for risk assessment: From days to decades. *Quarterly Journal of the Royal Meteorological Society*, *128*, 747–774.

Papakonstantinou, A., Rogers, A., Gerding, E. H., & Jennings, N. R. (2008). A Truthful Two-Stage Mechanism for Eliciting Probabilistic Estimates with Unknown Costs. *ECAI* (pp. 448–452).

Papakonstantinou, A., Rogers, A., Gerding, E. H., & Jennings, N. R. (2010). Mechanism design for eliciting probabilistic estimates from multiple suppliers with unknown costs and limited precision. *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets* (pp. 102–116). Springer.

Papakonstantinou, A., Rogers, A., Gerding, E. H., & Jennings, N. R. (2011). Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artificial Intelligence*, *175*, 648–672.

Petcu, A., Faltings, B., & Parkes, D. C. (2006). MDPOP: Faithful distributed implementation of efficient social choice problems. *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems* (pp. 1397–1404). ACM.

Peterson, C. R., Snapper, K. J., & Murphy, A. H. (1972). Credible Interval Temperature Forecasts. *Bulletin of the American Meteorological Society*, *53*, 966–1073.

Petroski, H. (2005). *Pushing the limits: new adventures in engineering*. Random House Digital, Inc.

Prelec, D. (2004). A Bayesian truth serum for subjective data. *Science*, *306*, 462–466.

Ramsey, F. P. (1931). Truth and probability (1926). *The foundations of mathematics and other logical essays*, 156–198.

Rheingold, H. (1993). *The virtual community: Homesteading on the electronic frontier*. Basic Books.

Rogers, A., Dash, R. K., Jennings, N., Reece, S., & Roberts, S. (2006). Computational mechanism design for information fusion within sensor networks. *Information Fusion, 2006 9th International Conference on* (pp. 1–7). IEEE.

Rogerson, W. P. (1985). Repeated moral hazard. *Econometrica: Journal of the Econometric Society*, 69–76.

Rosenschein, J. S., & Zlotkin, G. (1995). Rules of Encounter: designing conventions for automated negotiation among computers. 1994. The MIT Press, Cambridge, MA, USA.

Russell, S. J., Norvig, P., Canny, J. F., Malik, J. M., & Edwards, D. D. (1995). *Artificial intelligence: a modern approach* (Vol. 74). Prentice hall Englewood Cliffs.

Sadri, P., & Gossain, V. (2010). Systems and methods for determining a relevance rank for a point of interest. Google Patents.

Sandholm, T. W. (1996). Limitations of the Vickrey auction in computational multiagent systems. *Proceedings of the Second International Conference on Multiagent Systems (ICMAS-96)* (pp. 299–306).

Satterthwaite, M. A. (1975). Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, *10*, 187–217.

Savage, L. (1972). J.(1954) The Foundations of Statistics. New York: John Wiley and Sons.

Savage, L. J. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, *66*, 783–801.

Schervish, M. J. (1989). A general method for comparing probability assessors. *The Annals of Statistics*, *17*, 1856–1879.

Schieritz, N., & Milling, P. M. (2003). Modeling the forest or modeling the trees. *Proceedings of the 21st International Conference of the System Dynamics Society* (pp. 20–24).

Schillo, M., Funk, P., & Rovatsos, M. (2000). Using trust for detecting deceitful agents in artificial societies. *Applied Artificial Intelligence*, *14*, 825–848.

Schrank, D., Lomax, T., & Turner, S. (2010). TTIs 2010 Urban Mobility Report Powered by INRIX Traffic Data. *Texas Transportation Institute, The Texas A&M University System*.

Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, *1*, 43–62.

Sen, S., & Sajja, N. (2002). Robustness of reputation-based trust: Boolean case. *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1* (pp. 288–293). ACM.

Shafer, G. (1976). *A mathematical theory of evidence* (Vol. 1). Princeton university press Princeton.

Shapiro, S. P. (1987). The social control of impersonal trust. *American journal of Sociology*, 623–658.

Sheth, A. (2009). Citizen sensing, social signals, and enriching human experience. *Internet Computing, IEEE*, *13*, 87–92.

Shiller, R. J., Kon-Ya, F., & Tsutsui, Y. (1996). Why did the Nikkei crash? Expanding the scope of expectations data collection. *The Review of Economics and Statistics*, 156–164.

Silva, D., Ghanem, M., & Guo, Y. (2012). WikiSensing: An Online Collaborative Approach for Sensor Data Management. *Sensors*, *12*, 13295–13332.

Singh, M. P., Yu, B., & Venkatraman, M. (2001). Community-based service location. *Communications of the ACM*, *44*, 49–54.

Smith, A. (1863). *An Inquiry into the Nature and Causes of the Wealth of Nations*. A. and C. Black.

Soyster, A. L. (1973). Technical Note Convex Programming with Set-Inclusive Constraints and Applications to Inexact Linear Programming. *Operations research*, *21*, 1154–1157.

Spiegelhalter, D. J. (1986). Probabilistic prediction in patient management and clinical trials. *Statistics in medicine*, *5*, 421–433.

Srivastava, M., Culler, D., & Estrin, D. (2004). Guest editors' introduction: overview of sensor networks. *Computer*, *37*, 41–49.

STANAG, N. (1999). 5516 Edition 3, Tactical Data Exchange-Link 16. *The MITRE Corporation Documentation*, *1*, 88–1O6.

Stinson, C. W. (2003). *Internet Protocol (IP) Over Link-16*. DTIC Document.

Tetlock, P. E. (2005). *Expert political judgment: How good is it? How can we know?* Princeton University Press.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, *16*, 8–37.

Vlassis, N. (2007). A concise introduction to multiagent systems and distributed artificial intelligence. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, *1*, 1–71.

Weiss, G. (1999). *Multiagent systems: a modern approach to distributed artificial intelligence*. The MIT press.

Wilensky, U. (1999). {NetLogo}.

Wilson, C. (2004). Network Centric Warfare: Background and Oversight Issues for Congress. DTIC
Document.

Wooldridge, M. (2009). *An introduction to multiagent systems*. Wiley. com.

Wright, W. F. (1988). Empirical comparison of subjective probability elicitation methods*.
Contemporary Accounting Research, *5*, 47–57.

Wu, D., & Sun, Y. (2001). The emergence of trust in multi-agent bidding: a computational
approach. *System Sciences, 2001. Proceedings of the 34th Annual Hawaii International
Conference on* (p. 8–pp). IEEE.

Yamanoue, T., Oda, K., & Shimozono, K. (2013). Demonstration of a Loosely Coupled M2M
System using Arduino, Android and Wiki Software.

Yokoo, M., & Suzuki, K. (2004). Secure generalized vickrey auction without third-party servers.
Financial Cryptography (pp. 132–146). Springer.

Yu, B., & Singh, M. P. (2000). A social mechanism of reputation management in electronic
communities. *Cooperative Information Agents IV-The Future of Information Agents in
Cyberspace* (pp. 154–165). Springer.

Yu, B., & Singh, M. P. (2002). An evidential model of distributed reputation management.
*Proceedings of the first international joint conference on Autonomous agents and
multiagent systems: part 1* (pp. 294–301). ACM.

Yu, B., & Singh, M. P. (2003). Searching social networks. *Proceedings of the second international
joint conference on Autonomous agents and multiagent systems* (pp. 65–72). ACM.

YU, X., KUANG, J., WANG, H., & LIU, Q. (2005). Modeling and Simulation of JTIDS Network Using
OPNET. *Computer Engineering*, *31*, 108–110.

Yuan, S., & Sung, H. (2004). A learning-enabled integrative trust model for e-markets. *Applied Artificial Intelligence*, *18*, 69–95.

Zacharia, G., & Maes, P. (2000). Trust management through reputation mechanisms. *Applied Artificial Intelligence*, *14*, 881–907.

ZHAO, C., CHEN, L., & LU, J. (2008). Simulation of Link 11 System Based on QualNet [J]. *Telecommunication Engineering*, *2*, 019.

Zhao-xiong, H., Xing, L., Xue-min, L., & Jing-lun, Z. (2010). Link-16 model architecture for multiple nets simulation in NS-2. *Industrial Engineering and Engineering Management (IEEM), 2010 IEEE International Conference on* (pp. 1645–1649). IEEE.

Zohar, A., & Rosenschein, J. S. (2008). Mechanisms for information elicitation. *Artificial Intelligence*, *172*, 1917–1939.

APPENDICES

Appendix A   Possibility and Impossibility Results in Mechanism Design

Appendix A briefly covers some of the key *impossibility results* and *possibility results* in Mechanism Design. Impossibility results prove the impossibility of implementing specific Social Choice Functions (SCFs) under particular solution concepts. Possibility results discuss two mechanisms which achieve different sets of the SCF properties.

The basic approach to show impossibility is through a conflict across the assumptions and conditions – we first assume direct-revelation and incentive-compatibility, then we express the desiderata of an outcome rule as a bunch of mathematical conditions and finally we show a conflict across the conditions.

Impossibility results are dictated by conditions on the assumptions about the environment, the agent preferences as well as the equilibrium solution concept. Hence it is necessary to introduce and define some new concepts.

- **Dictatorial SCF:** A social-choice function is said to be dictatorial if one or more agents always receive its most-preferred or one of its most-preferred alternatives.

- **General Preferences:** Preferences θ_i are general when they establish a complete and transitive preference ordering on outcomes. An ordering is complete if for all $o_1, o_2 \in O$, we have $o_1 > o_2$ or $o_2 > o_1$ (or both). An ordering is transitive if for all $o_1, o_2, o_3 \in O$, if $o_1 > o_2$ and $o_2 > o_3$ then $o_1 > o_3$.

- **Coalition-Proof Mechanisms:** A mechanism M can be caled coalition-proof if truth revelation is a dominant strategy for any agent coalition, where the coalition is allowed to make side-payments and even redistribute items once the mechanism is over.

- **General Environment:** A general environment is the environment in which there is a discrete set of possible outcomes O and agents have general preferences.

- **Simple Exchange Environment:** A simple exchange environment is the environment in which there the sellers selling single units of the same good.

We are now in a position to describe the three main impossibility results.

Gibbard-Satterthwaite Impossibility Theorem

The Gibbard-Satterthwaite Impossibility Theorem (Satterthwaite, 1975) states that -

In a setting consisting of

- *agents with general utility preferences*
- *a finite set of outcomes O, with more than 3 possible outcomes (i.e. $|O| > 3$)*

a social-choice function is dominant-strategy implementable if and only if it is dictatorial.

The above theorem is limiting in many respects because when coupled with the revelation principle, it implies no mechanism can be implemented based on dominant strategy given the general conditions. Great care must be exercised when interpreting impossibility results such as Gibbard-Satterthwaite, as they do not necessarily continue to hold in restricted environments. Thus one way to circumvent this impossibility result is to impose restrictions on agents, environment and solution concepts:

- Simple *exchange* environment
- Additional constraints on agent preferences (e.g. quasi-linear utility)
- Weaker implementation concepts, e.g. Bayesian-Nash implementation

Hurwicz Impossibility Theorem

The Hurwicz Impossibility Theorem (Hurwicz, 1973) states that -

There does not exist any incentive-compatible mechanism that implements a SCF that is efficient and budget-balanced in dominant strategy equilibrium, even with quasi-linear preferences.

This result is equally constraining, as it implies that for dominant strategy solutions we can't attain both allocative efficiency and budget- balance in a simple exchange economy. Thus we need to sacrifice some desired property, and this is usually implemented in the following manner:

- Sacrifice strong budget-balance to achieve strategy-proofness, efficiency and weak budget- balance via the Vickrey-Clarke-Groves mechanisms.

- Use a weaker implementation concept, Bayesian-Nash equilibrium, so as to achieve budget balance, efficiency and incentive compatibility via d'Aspremont-Gerard-Varet-Arrow (dAGVA) mechanism.

However, one can't achieve all the desiderata – incentive compatibility, budget-balance, efficiency and individual rationality – in the dAGVA mechanism due to the Myerson-Satterthwaite Impossibility Theorem.

Myerson-Satterthwaite Impossibility Theorem

The Myerson-Satterthwaite Impossibility Theorem (Myerson & Satterthwaite, 1983) states that -

There does not exist any mechanism that implements a SCF that is efficient, budget-balanced and individually rational in Bayesian-Nash strategy equilibrium, even with quasi-linear preferences.

What the theorem translates into is we can only achieve at most two of desired properties – Efficiency, Individual Rationality, and Budget Balance - in a market with quasi-linear agent preferences, even with Bayesian-Nash implementation.

In our quest to design auction-based mechanism designs, the impossibility results severely restrict the available options - we can either opt for the VCG mechanism which allows us achieve efficiency, incentive-compatibility and individual rationality under dominant strategies or the alternative is to select a weaker solution concept of Bayesian-Nash equilibrium so as to achieve budget-balance at the expense of individual rationality via the dAGVA mechanism.

<div align="center">**Possibility Results**</div>

We will now turn our attention to the possibility results and present two mechanisms –Vickrey-Clarke-Groves (VCG) and dAGVA that achieve different sets of the SCF. The mechanisms are similar in that both achieve incentive-compatibility and efficiency. However, whilst the VCG mechanism is individually-rational but not budget-balanced, the dAGVA mechanism is budget-balanced but not individually-rational.

Vickrey-Clarke-Groves Mechanism

In their seminal work Clarke (1971), Groves (1973) and Vickrey (1961) presented a family of direct mechanisms, called the Vickrey-Clarke-Groves mechanisms, for problems in which agents' preferences are quasi-linear. The VCG mechanism implements an efficient and individually-rational SCF where truth-telling is a dominant strategy.

The VCG mechanisms are the only allocatively-efficient and strategy-proof mechanisms for agents with quasi-linear preferences and general valuation functions, amongst all direct-revelation mechanisms.

The outcome function of the VCG mechanism is specified by an allocation rule M and a payment function r. The VCG auction typically proceeds as outlined below:

1. The auctioneer lists the set of items M for sale.
2. Each agent i will reports its valuation function $v_i(K, \theta_i)$ for all the allocations $K \in \mathrm{K}$ in the set of all possible sets of the items in M.

3. Each agent i will reports its type $\hat{\theta}_i$ as well.

4. The auctioneer will subsequently find the efficient allocation by solving the following equation:

$$\hat{K}^* = \arg \max_{K \in \mathrm{K}} \sum_{i \in I} v_i(K, \hat{\theta}_i)$$
(A.1)

5. Finally the auctioneer will compute the transfer r_i from each agent as:

$$r_i = \left[\max_{K \in \mathrm{K}} \sum_{j \in -i} v_j(K, \hat{\theta}_j) \right] - \left[\sum_{j \in -i} v_j(\hat{K}^*, \hat{\theta}_j) \right]$$
(A.2)

The VCG mechanism achieves its strategy-proofness through its payment scheme. This scheme squarely aligns the utility attained by an agent with that agent's marginal contribution to the mechanism. We can calculate the utility an agent derives with the efficient allocation and payment scheme:

$$u_i(K, \hat{\theta}_i) = v_i(\hat{K}^*, \theta_i) - r_i(\hat{K}^*, \hat{\theta}_i)$$

$$= v_i(\hat{K}^*, \theta_i) - \left[\max_{K \in \mathrm{K}} \sum_{j \in -i} v_j(K, \hat{\theta}_j) \right]$$

$$- \left[\sum_{j \in -i} v_j(\hat{K}^*, \hat{\theta}_j) \right]$$
(A.3)

Thus the VCG scheme, in essence, internalizes the externality, by aligning an agent's private goal of maximizing its local incentive with the global system goal of efficient allocation. Hence it is it the best interest of an agent given its own true preferences and the reports of other agents, for the mechanism to select the best system-wide solution.

But the VCG mechanism is not budget balanced. When implemented in an auction setting, the VCG mechanism results in a surplus for the auctioneer. Budget balance is an important criteria, and the dAGVA mechanism achieves budget-balance but at the expense of individual rationality.

dAGVA Mechanism

The dAGVA also known as the "expected Groves" mechanism, was proposed by Arrow (1977) and d Aspremont & Gérard–Varet (1982). The dAGVA mechanism circumvents the Hurwicz impossibility and achieves efficiency and budget-balance but in Bayesian-Nash equilibrium and not dominant-strategy equilibrium. The dAGVA mechanism is not individual-rational, as is apparent from the Myerson - Satterthwaite impossibility theorem.

The dAGVA mechanism is ex-ante individual-rational, Bayesian-Nash incentive-compatible, efficient and (strong) budget-balanced with quasi-linear agent preferences.

The allocation rule in dAGVA is the same as in the VCG mechanism but differs crucially in its payment scheme. In more detail, the dAGVA mechanism proceeds as follows:

1. The auctioneer lists the set of items M for sale.
2. Each agent i will reports its valuation function $v_i(K, \theta_i)$ for all the allocation $K \in K$.
3. Each agent *i* will reports its type $\hat{\theta}_i$ as well.
4. The auctioneer will subsequently find the efficient allocation by solving the following equation:

$$\hat{K}^* = \arg \max_{K \in K} \sum_{i \in I} v_i(K, \hat{\theta}_i) \tag{A.4}$$

5. It also computes each transfer r_i from each agent as:

$$r_i = x_i(\hat{\theta}_{-i}) - E_{\theta_{-i}} \left[\max_{K \in K} \sum_{j \in -i} v_j(K(\hat{\theta}_i, \theta_{-i}), \hat{\theta}_j) \right] \tag{A.5}$$

The *expected* utility derived by an agent derives can be expressed as:

$$u_i(K, \hat{\theta}_i) = E_{\theta_{-i}}\left[\max_{K \in \mathrm{K}} v_i(K(\hat{\theta}_i, \theta_{-i}), \hat{\theta}_i\right] - r_i(\hat{K}, \hat{\theta}_i)$$

$$= E_{\theta_{-i}}\left[\max_{K \in \mathrm{K}} \sum_{i \in I} v_i(K(\hat{\theta}_i, \theta_{-i}), \hat{\theta}_i)\right] - x_i(\hat{\theta}_{-i})$$

(A.6)

The first term, in the above equation, is the expected total value for agents $j \neq i$ when they tell the truth and agent i announces type $\hat{\theta}_i$ while $x_i(\hat{\theta}_{-i})$ is an arbitrary function on the types of agents. The dAGVA mechanism makes it possible to choose the $x_i(\hat{\theta}_{-i})$ function in such a way so as to satisfy budget-balance, i.e. $\sum_{i \in I} r_i = 0$ which implies:

$$\sum_{i \in I}\left(x_i(\hat{\theta}_{-i}) - E_{\theta_{-i}}\left[\max_{K \in \mathrm{K}} \sum_{j \in -i} v_j(K(\hat{\theta}_i, \theta_{-i}), \hat{\theta}_j)\right]\right) = 0$$

(A.7)

One possible choice of $x_i(\hat{\theta}_{-i})$ could be the average of the negative part of the transfer of all the other agents.

$$x_i(\hat{\theta}_{-i})$$

$$= \frac{1}{|I| - 1}\sum_{j \in -i}\left(E_{\theta_{-j}}\left[\max_{K \in \mathrm{K}} \sum_{j \in -i} v_j(K(\hat{\theta}_i, \theta_{-i}), \hat{\theta}_j)\right]\right)$$

(A.8)

The properties of incentive compatibility, efficiency and budget- balance make the dAGVA mechanism seem like an attractive option. However, the dAGVA mechanism has quite a few drawbacks:

- It fails to satisfy the individual rationality constraint
- Bayesian-Nash is a weaker solution concept than dominant-strategy
- It puts too much emphasis on agent information-revelation

The results of our discussion of possibility and impossibility results have been surmised in Table A-1:

Table A-1 Main possibility and impossibility results in Mechanism Design

Result Name	Preferences	Solution Concept	Environment	Effic-iency	Budget Balance	Individual Rationality
Gibbard Satterthwaite	General	Dominant	General	No	No	No
Hurwicz	Quasi-linear	Dominant	Simple Exchange	No	No	--
Myerson Satterthwaite	Quasi-linear	Bayesian-Nash	Simple Exchange	No	No	No
VCG	Quasi-linear	Dominant	Simple Exchange	Yes	No	Yes
dAGVA	Quasi-linear	Bayesian-Nash	Simple Exchange	Yes	Yes	No

Appendix B   Link 16

Tactical data links involve transmissions of bit-oriented digital information which are exchanged via data links known as Tactical Digital Information Links (TADIL). TADIL J, also known as Link 16, has been the designated DoD primary tactical data link since October 1994 and is assumed as the baseline for the current work. This section highlights the key concepts involved with TADILs and specifically, TADIL J.

TADIL J is a communication, navigation, and identification system that supports information exchange between tactical command, control, communications, computers, and intelligence (C4I) systems. The radio transmission and reception component of TADIL J is the Joint Tactical Information Distribution System (JTIDS) or its successor, the Multifunctional Information Distribution System (MIDS). This high-capacity, ultra high frequency (UHF), line of sight (LOS), frequency hopping data communications terminals provide secure, jam-resistant voice and digital data exchange.

JTIDS/MIDS terminals operate on the principal of time division multiple access (TDMA), wherein time slots are allocated among all TADIL J network participants for the transmission and reception of data. TDMA eliminates the requirement for a net control station (NCS) by providing nodeless communication network architecture. There are 1536 time slots in a 12 second frame, after which the pattern of time slots repeats. A platform transmits specific message types, such as position, voice, surveillance, electronic warfare, etc., on the time slots assigned to that platform and message type; these assignments have network-wide significance as the intended receivers must have those same time slots available to receive from that transmitter. If a platform is receiving signals from multiple transmitters during the same time slot, it will correctly decode the transmission from the strongest transmitter and discard the rest.

The above description fits many different types of TDMA systems but Link 16 adds unique features on top of this basic architecture which impact the survivability problem. Transmissions are fast frequency-hopped over a range of different frequencies at a high rate, with a different frequency hopping pattern in each time slot. The frequency hopping pattern is controlled by three variables, a crypto key, a net number, and the time of day (as represented by specific time

slot identifiers). The net number makes Link 16 more flexible than a basic TDMA system as it allows different platforms to transmit on the same time slot, but different net number, without interfering; this is referred to as multi-netting. For example, platform A can transmit voice on a given slot using net 1 while platform B transmits surveillance on the same time slot but using net 2. The two transmissions will not interfere because, with different net numbers, they will be transmitted using different frequency hopping patterns. However, on a given time slot, a platform can only send or receive from a single net; multi-netting is used primarily to allow non-overlapping communities of interest to use the same time slots, freeing other time slots for communication needs shared by all platforms.
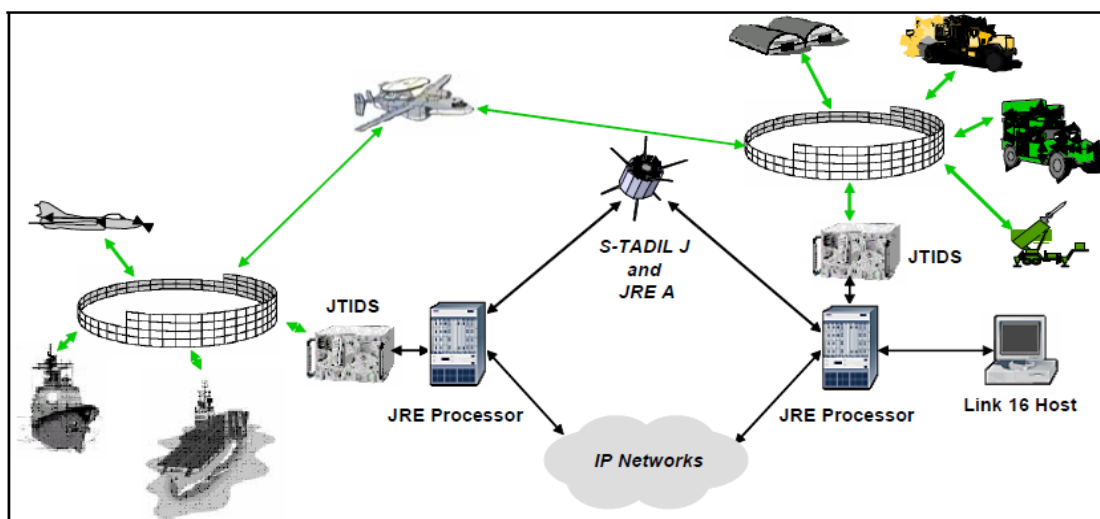

Figure B-1 Link 16 Architecture

A Link 16 time slot can be assigned to transmit, receive, or relay. A relay assignment means that the time slot will retransmit information received (or transmitted) on a previous time slot, designated at the time the network is designed. Relay functionality allows Link 16 networks to extend over large geographical areas, even though the range of frequencies used by Link 16 limit its direct transmissions to Line-Of-Sight (LOS) receivers, a distance that varies with aircraft altitude. Platforms with relay assignments allow recipients to overcome LOS constraints and receive data from distant platforms, or from platforms such as low-flying aircraft, that may be masked by terrain.