# APPROVAL SHEET

| | |
|---|---|
| Title of Dissertation: | Opportunistic Bartering of Digital Goods and Services in Pervasive Environments |
| Name of Candidate: | Olga Vladi Ratsimor<br>Doctor of Philosophy, 2007 |
| Dissertation and Abstract Approved: | ____________________________<br>Dr. Tim Finin<br>Professor<br>Department of Computer Science and Electrical Engineering |
| Date Approved: | ____________________________ |

Major: Computer Science.

Professional publications:

- Sethuram Balaji Kodeswaran, Olga Ratsimor, Filip Perich, and Anupam Joshi, "Utilizing Semantic Tags for Policy Based Networking", IEEE Global TeleComm. Conf. (GLOBCOM-07).

- Olga Ratsimor, Dipanjan Chakraborty, Anugeetha Kunjithapatham, Deepali Khushraj, Anupam Joshi, Timothy Finin, Yelena Yesha, "Service Discovery in Agent-based Pervasive Computing Environments", Extended Version, Journal on Mobile Networking and Applications (MONET), Special issue on Mobile and Pervasive Commerce, 2003.

- Olga Ratsimor, Anupam Joshi, Timothy Finin, Yelena Yesha, “Intelligent Ad Hoc Marketing within Hotspot Networks”, Technical Report, Nov 2003.
- Olga Ratsimor, Anupam Joshi, Timothy Finin, Yelena Yesha, “eNcentive: A Framework for Intelligent Marketing in Mobile Peer-To-Peer Environments”, The 5th International Conference on Electronic Commerce (ICEC), Oct 2003.
- Sethuram Balaji Kodeswaran, Olga Ratsimor, Anupam Joshi, Timothy Finin, Yelena Yesha, “Using Peer-to-Peer Data Routing for Infrastructure-based Wireless Networks”, IEEE International Conference on Pervasive Computing and Communications (PerCom), March 2003.
- Olga Ratsimor, Sethuram Balaji Kodeswaran, Anupam Joshi, Timothy Finin, Yelena Yesha, “Combining Infrastructure and Ad hoc Collaboration For Data Management in Mobile Wireless Networks”, Paper, ACM Workshop on “Ad-hoc Communications and Collaboration in Ubiquitous Computing Environments”, Nov 2002.
- Sethuram Balaji Kodeswaran, Olga Ratsimor, Anupam Joshi, Timothy Finin, Yelena Yesha, “Numi: A Framework for Collaborative Data Management in a Network of InfoStations”, UMBC SRC, Nov 2002.
- Olga Ratsimor, Dipanjan Chakraborty, Sovrin Tolia, Deepali Kushraj, Anugeetha Kunjithapatham, Gaurav Gupta, Anupam Joshi, Timothy Finin, “Allia: Alliance-based Service Discovery for Ad-Hoc Environments”, ACM Mobile Commerce Workshop, September, 2002.
- Timothy Finin, Anupam Joshi, Lalana Kagal, Olga Ratsimor, Sasikanth Avancha, Vlad Korolev, Harry Chen, Filip Perich, R. Scott Cost, “Intelligent Agents for Mobile and Embedded Devices”, Special issue of the International Journal on Cooperative Information Systems, 2002.
- Timothy Finin, Anupam Joshi, Lalana Kagal, Olga Ratsimor, Vlad Korolev, and Harry Chen, “Information Agents for Mobile and Embedded Devices”, Paper, Fifth International Workshop Cooperative Information Agents, Modena, Italy, September 6-8, 2001.
- Olga Ratsimor, Vladimir Korolev, Anupam Joshi, and Timothy Finin, “Agents2Go: An Infrastructure for Location-Dependent Service Discovery in the Mobile Electronic Commerce Environment”, Paper, ACM Mobile Commerce Workshop, July, 2001.

Professional positions held:

- Research Scientist, (present).

  Intelligent Automation, Inc., Rockville, MD

- Summer Research Intern, (May 2006 - August 2006).

  GE's Global Research Center, PDS Pervasive Decisioning Systems Lab, Niskayuna, New York

- Graduate Research Assistant, (January 2000 - May 2007).

  Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore County

- Lecturer, (September 2005 - December 2005).

  Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore County

- Lecturer, (June 2001 - July 2001)

  Information Science and Systems, Morgan State University, Baltimore, Maryland

- Graduate Teaching Assistant, (June 1999 - December 1999), (June 2000 - August 2000).

  Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore

# ABSTRACT

Title of Dissertation: Opportunistic Bartering of Digital Goods and Services in Pervasive Environments

Olga Vladi Ratsimor, 2007

Dissertation directed by: Dr. Tim Finin
Professor
Department of Computer Science and Electrical Engineering

The vision of mobile personal devices querying peers in their environment for information such as local restaurant recommendations or directions to the closest gas station, or traffic and weather updates has long been a goal of the pervasive research community. However, considering the diversity and the personal nature of devices participating in pervasive environments it is not feasible to assume that these interactions and collaborations will take place with out economically-driven motivating incentives.

This dissertation presents a novel bartering communication model that provides an underlying framework for incentives for collaborations in mobile pervasive environments by supporting opportunistic serendipitous peer-to-peer bartering for digital goods such as ring tones, MP3's and podcasts.

To demonstrate viability and advantages of this innovative bartering approach, we compare and contrast the performances of two conventional, frequently employed, peer-to-peer interaction approaches namely Altruists and FreeRiders against two collaborative strategies that employ the Double Coincidence of Wants paradigm from the domain of barter exchanges. In particular, we present our communication framework that represents these collaborative strategies through a set of interaction policies that reflect these strategies. Furthermore, we present a set of results from our in-depth simulation studies that compare these strategies. We examine the operation of the nodes employing our framework and executing these four distinct strategies and specifically, we compare the performances of the nodes executing these strategies in homogeneous and heterogeneous networks of mobile devices. We also examine the effects of adding InfoStations to these net-

works. For each of the strategies, we observe levels of gains and losses that nodes experience as result of collaborative digital good exchanges. We also evaluate communication overhead that nodes incur while looking for possible collaborative exchange. Furthermore, this dissertation offers an in-depth study of the swarm-like inter-strategy dynamics in heterogeneous networks populated with diverse nodes displaying varying levels of collaborative interaction attitudes. Further, the bartering framework is extended by incorporating value -sensitive bartering models that incorporate digital goods and content valuations into the bartering exchange process. In addition, the bartering model is extended by integration of socially influenced collaborative interaction that exploit role based social relationships between mobile peers that populate dynamic mobile environments.

Taken as a whole, the novel research work presented in this dissertation offers the first comprehensive effort that employs and models opportunistic bartering-based collaborative methodology in the context of serendipitous encounters in dynamic mobile peer-to-peer pervasive environments where mobile entities negotiate and exchange digital goods and content.

# Opportunistic Bartering of Digital Goods and Services in Pervasive Environments

by
Olga Vladi Ratsimor

Dissertation submitted to the Faculty of the Graduate School
of the University of Maryland in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2007

*This dissertation is lovingly dedicated to my husband, my Dad and my Mom.*

# Acknowledgement

First and foremost, I would like to thank my advisors, Dr. Timothy Finin, Dr. Anupam Joshi and Dr. Yelena Yesha for their guidance and support throughout my graduate career and during the completion of this thesis. Their words of advice have been invaluable. I am also extremely grateful to the members of my committee, Dr. Filip Perich, Dr. Milton Halem and Mrs. Marisa Viveros. I would like to acknowledge the students in the Department of Computer Science, and in particular, the students of the eBiquity Research Group.

Finally and most importantly, I am extremely grateful to my parents, my husband, family members and friends, who have always wished the best for me and showered me with their encouragements and moral support. To my husband, Balaji, without your optimism, reassurance and support on so many fronts over the past few years this project would not have been possible. Thank you for being there when I needed you.

# TABLE OF CONTENTS

# LIST OF TABLES

# LIST OF FIGURES

# Chapter I

# INTRODUCTION

*The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it."* - Mark Weiser[112]

## I.A Overview

With the explosion of wireless technologies [108], it is becoming clear that mobile computing is slowly but surely becoming a dominant new culture [89]. People have quickly become more and more reliant on the flexibility and autonomy that mobile devices can provide. Mobile users no longer can imagine their lives without their cell phones, PDAs, MP3 players, digital cameras, etc. This mobile culture is in a fluidic state of nonstop evolution [38, 100]. New hardware, new software and new services are being developed and broadly used. Old applications and old services are being transformed to keep up with this mobile evolution. In addition, a new class of digital goods and services is starting to evolve [103]. Distribution and sales of cell phone ring tones, MP3 music files, podcasts, mobile games, and electronic coupons is developing into a separate industry [35]. Mobile users equipped with personal devices are actively accruing theses goods and services. Many users are motivated by a goal of personalization of their off-the-shelf equipment [41]. Others are looking for entertainment in a compact, pocket-sized form that can be ubiquitously present with them throughout the day. As this mobile revolution unravels, it is becoming clear that the next big step to truly embed pervasive computing in to everyday lives of regular individuals is seamless, automated, personalized and context aware collaborations between these personal devices. The vision of mobile personal devices querying peers in the environment for information such as local restaurant recommendations or directions to the closest gas station, or traffic and weather updates has long been a goal of the pervasive research community [77, 75]. However along the way of envisioning these interactions and collaborations, we have lost track of reasons that would prompt personal devices that have limited resources to cooperate with their peers in the environment. Personal devices have technical limitations such as limited computing power, limited battery life etc. and most important of all, they are designed to serve the needs of a single individual.

### I.A.1 Problem Statement

Considering the diversity and the personal nature of devices participating in pervasive environments, is the current model of altruistic ad hoc collaborations still the most effective model to motivate collaborations? Is it valid to assume that devices in pervasive environments will collaborate regardless of their technical limitations? Is there a communication and collaboration model that can be employed by personal mobile devices to enhance the productivity of opportunistic peer-to-peer collaboration?

Consider the following scenario:

*Sam and Ellie are undergraduate students at UMBC. Both, Sam and Ellie are taking "The Introduction to Computer Science" class. Though they are registered for different sections, both of their classes are covering the same material. Sam missed one of the classes and he is interested in getting the notes for that lecture. Ellie has the lecture notes that Sam is looking for but she is reluctant to give them away. After all, she has put a lot of effort in composing those notes. Sam proposes an idea that would motivate Ellie to give him a copy of the notes. He suggests that in exchange for the much needed notes, he will share with her a useful technique that his lecturer taught during the last lecture. Ellie is intrigued. She has been having hard time solving one of the homework problems and she feels that this technique could help her with solving that problem. Sam and Ellie exchange lecture notes for the technique description. This collaboration proved to be beneficial to both of the students.*

This social collaboration scenario is not far from the desired collaborations that have been envisioned in mobile pervasive environments. In our work, we employ this bartering model to stimulate opportunistic collaborations in pervasive environments.

### I.A.2 Thesis Statement

***Opportunistic collaborations in mobile pervasive environments driven by self motivated economic goals yield improved cooperation efficiency and productivity for such interactions.***

## I.B Contributions

### I.B.1 Conceptual Methodology for Collaborative Trading in Pervasive Environments

The primary objective of this thesis is the development of a conceptual bartering methodology that reflects the context of cooperative exchanges and trade-based interactions that occur as a result of the serendipitous encounters between peers in dynamic mobile pervasive environments.

### I.B.2 Framework for Opportunistic Collaborative Interaction in MANETs

To solve the problem described in the previous section and to apply the developed methodology, this thesis introduces an architectural framework that provides support for opportunistic peer-to-peer collaborative interactions in mobile pervasive environments. The framework allows nodes to express the preferred collaboration attitude through a set of collaborative policies that best reflect users cooperation philosophy. In addition, our framework facilitates advertisement and discovery of digital goods and content, policy based negotiations and transaction management. This framework allow nodes to collaborate in both homogeneous environments where all of the nodes have similar collaboration policies and in heterogeneous environments which are populated with nodes of varying collaboration levels.

### I.B.3 Characterization Study of Collaborative Interactions in Homogeneous and Heterogeneous MANETs

An important contribution of our work is an in-depth study of opportunistic collaborative interactions in pervasive environments. In particular, we examine effectiveness of collaborative approaches/strategies in both homogeneous and heterogeneous environments. Homogeneous environments are populated with nodes that share similar collaboration philosophy. Thus we examine the effectiveness of the interaction strategies as a baseline for measuring collaborative behavior. Furthermore, in heterogeneous networks where nodes exhibit varying levels of cooperation, we evaluate the effectiveness of the opportunistic interactions and inter-strategy dynamics that occur in these swarm-like collaborative exchanges.

### I.B.4 Development of Value-Sensitive Bartering Communication Models and Investment-Based Trading Approach for MANETs

An extension of our bartering collaboration methodology is the development of a value-sensitive bartering model that employs environmental conditions to derive valuation for digital content in the environment. The valuations are quantified into distinct categories that symbolize the worthiness and the significance of the digital content. This thesis also provides an in-depth study of a set of specific implementations of the value-sensitive bartering model including supply and demand sensitive valuation models and personalized valuation models. Furthermore, as an extension to the value-sensitive collaboration model, we have developed an investment-based bartering strategy that considers valuation of digital content and relaxes the strict constrains of conventional bartering methods that relies on the principal of *coincidence of double wants*.

### I.B.5 Socially Influenced Collaborations in MANETs

Consequently, to expand this space of possible deals during bartering, our framework exploits role-based relationships such as social relationships or trust-based relationships among owners of the mobile peer devices. In particular, our framework reflects significance of the social relationships by exhibiting different levels of cooperation during the bartering process.

## I.C Research Impact

With proliferation of mobile technology and the rapid expansion in the industry of mobile digital content, the interaction patterns and approaches are soon to become a crucial aspect of mobile pervasive environments. Promoting collaborative exchanges of services and information among small personal device in such environments is an important component of successful peer-to-peer interactions. Considering currently utilized conventional methods of collaborations and the levels of collaborative interactions that take place in mobile pervasive environments, it is clear that there is a clear deficit of productive peer-to-peer exchanges in this class of environments. This dissertations explores an alternative method of cooperation that enhances the collaboration process and stimulates peer-to-peer interactions and exchange of digital content, goods and services. This dissertation presents the first comprehensive research work that employes and models opportunistic bartering-based collaborative methodology in the context of serendipitous encounters in dynamic mobile peer-to-peer pervasive environments. This dissertation also shows that incorporation of barter-driven cooperation policies into the collaborative interaction process drastically improves the productivity and effectiveness for the individual mobile nodes and the efficiency of the environment as whole. In addition, the bartering approach provides mobile nodes with resilience to the harsh conditions that are possible in heterogeneous environments which are populated with a wide range of nodes with varying collaborative approaches. This dissertation also shows that mobile nodes do not need to employ very sophisticated methods of bartering to achieve a high level of effectiveness. In particular, the valuation of digital content is not a crucial factor in establishing effective bartering interactions. Though the nodes could attempt to employ sophisticated valuation mechanisms such as personalized good valuations, these extensions are not critical to the overall effectiveness of the battering process. On the other hand, social role based collaborative exchanges can promote inter-circle productivity thus furthering the effectiveness of the inter group collaborations. Overall, this dissertation presents an approach that promotes efficient and productive interactions in dynamic mobile environments and provides nodes with an interaction method that protects them from inefficient philanthropic interactions and predatory free riding behavior that is prevalent in mobile heterogeneous peer-to-peer environments.

## I.D Dissertation Outline

This dissertation advances the field of collaborative peer to peer interactions in mobile pervasive environments. The following is a chapter by chapter overview of this dissertation.

**Chapter I** is an introductory chapter that briefly describes the issues that effect collaborative exchanges in pervasive environments. The chapter outlines the research contributions of this dissertation and expected impact of this research.

**Chapter II** presents the key motivating factors that drive the research work presented in this dissertation. We also describe a set of collaborative interactions that are representative of the presented approach.

**Chapter III** describes related work that was done in the areas of mobile pervasive computing, mobile commerce, context awareness and other related domains.

**Chapter IV** provides detailed description of our framework that facilitates the collaborative interactions. This chapter also presents the collaboration approaches and the policies that are employed by the nodes during the exchange process.

**Chapter V** describes the results of the in-depth study of the presented framework. In particular, this chapter establishes a baseline for performance of the four collaboration strategies described in Chapter IV. This chapter also considers the effect of adding InfoStations to the mobile collaborative environments and looks at the effect of this environmental change.

**Chapter VI** presents the results of the inter-strategy interactions that occur in heterogeneous environments. In particular, this chapter looks at the interactions in “evenly mixed” networks where each of the strategies has equal representation in the network population. In addition, we present the results of our time-based simulation study that models the heterogeneous network in a swarm-line interactions and clearly shows the advantages of bartering collaboration approach.

**Chapter VII** presents the results of the inter-strategy interactions in heterogeneous networks where the population of the network has a strong domination by one of the considered strategies. This chapter further highlights the strength and versatility of the presented bartering collaboration approach.

**Chapter VIII** describes a set of extensions to the bartering collaboration model that take into consideration aspects of good valuations. This set of extensions include valuation models that rely on demand-sensitive valuation of goods, demand and supply sensitive valuation of and finally considers a model where valuations are personalized to match user preferences and interests. In addition, this chapter describes an investment based extension that enables mobile node executing the bartering framework to acquire goods for future transactions where these goods are used as bartering tokens.

**Chapter IX** describes an extension to our bartering model that relies on the social role based interactions to improve collaborative exchanges in the environment.

**Chapter X** is a concluding chapter that summarizes the key results of this dissertation and presents the overview of the accomplishments of our research work.

## Chapter II

# MOTIVATION

In the last few decades, we have witnessed dramatic changes in computing technology and people's perception of this technology. Many recent technological advances have managed to deeply embed themselves into every day lives of regular individuals. For instance, small personal devices are no longer perceived to be luxury articles but are thought of as an every day necessity. Visions of these personal devices coming together and forming pervasive computing environments have been a focus of many prior and ongoing research projects [62, 14, 110, 60, 97]. Majority of the work has focused on establishing communication between these devices[33], discovering services on these devices [54, 87, 77], embedding personalization and context awareness into the services and their management [64, 61, 77]. However, amidst this important work, there have been limited effort in investigating approaches that would incentivize much discussed opportunistic peer-to-peer collaborations. In this dissertation, we present an alternative approach aimed at stimulating peer-to-peer collaboration [62, 102, 96]. We present a novel bartering communication model that promotes opportunistic exchange of digital goods, services and content in the context of serendipitous encounters in dynamic mobile pervasive environments.

To better illustrate the problem, consider the current status of the mobile pervasive environments that currently surround us. A typical office environments, a typical shopping mall environments or a typical city street environments are heavily populated with mobile personal devices such as PDAs, smart phones, MP3 players, laptops. Users of these mobile devices are often interested in obtaining digital content that either addresses their immediate contextual needs or delivers entertainment or helps them customize and improve operations of their personal device. However, many of the devices do not attempt to interact or assist other mobile peers that are in their communication proximity. The key reasons behind this impediment is that there are no motivating factors that promote user collaborations. In fact, these environments, mobile users are more likely to either attempt to contact their cell phone service provider or browse through the vast amount of information available on the web rather than to attempt to search for a needed digital content amongst the near by peers. Interestingly relying on the geographically collocated mobile peers can prove to be more effective. The rational behind this channels of communications is that it is common occurrences, that people, who are in proximity of one another, are likely to share some personal or contextual interests[11].

For example, when two drivers, traveling in the opposite directions, encounter one another at an intersection both are likely to share a common interest in the traffic conditions that they are about to encounter. Another example is a student searching for a new song from *Red Hot Chili Peppers* while moving through a university campus. Chances of him meeting another student that has a match for his search are very high. So, this communication channel is particularly effective for acquiring context relevant content that is hard to acquire through the traditional communication channels. Our research assumes that mobile devices and their owners exhibit "rational behavior" and are reluctant to give away digital goods and other resources that they are in control of, unless there is a stimulating, rewarding incentive that prompts these collaborations. Thus, returning to previous examples, the driver will not be willing to give away the traffic data however; the driver would be willing to trade the traffic data for some other context-relevant good such as information concerning seating availability in a popular local restaurant. Similarly, two students are more likely to exchange MP3s rather than give them away. Thus, the bartering communication model is well suited for opportunistic serendipitous peer-to-peer interactions and exchanges. In the following Section, we will further motivate bartering model by grounding it to a particular set of example.

## II.A Sample Scenarios

Consider a dynamic mobile computing environment which is currently populated with four mobile devices: *D1*, *D2*, *D3* and *D4*. Each of the devices has a set of services that it is interested in acquiring and a set services that it is willing to offer to others. For example, device *D1* is interested in acquiring *Song-A* and could offer *RingTones-B*, *C* and *D*. One option is that all of the devices satisfy each others needs and blindly and generously offer up their available resources to each other. Another option is that devices could possibly sell needed services to one another. And third option is that devices attempt to exchange a service for a service and a good for a good in order to address their needs and wants. Issues associated with the first two options are discussed in a later Sections. Let's look in detail at the third option.

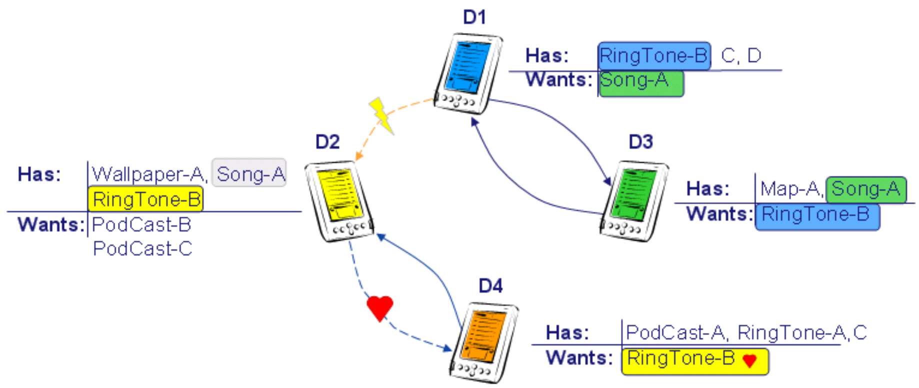


Figure II.1: Basic Bartering Scenarios

## II.A.1 Basic Collaborative Bartering Scenario

Suppose, device *D1* is interested in acquiring *Song-A* (Figure II.1). *D1* discovers that *D2* and *D3* are the devices that could potentially offer the song it is interested in. It starts by contacting *D2* and proposing an exchange of *Song-A* for a *RingTone-B*. Device *D2* responds with a rejection. *D2*'s reason for rejection is that it is not interested in ring tones. After analyzing the reason for rejection *D1* terminates negotiations with *D2*. *D1* attempts to approach *D3* with same proposal. *D3* is interested in one of the ring tones that *D1* has to offer. *D3* responds with a positive reply. *D1* and *D3* conduct a transaction and exchange *Song-A* and *RingTone-B*.

## II.A.2 Relationship-Based Collaboration Scenario

Suppose, device *D2* and device *D4* are owned by two sisters (that like each other). Both devices are aware of this social relationship. Now suppose, device *D4* is interested in acquiring *RingTone-B*. *D4* determines that *D1* and *D2* have this ring tone. *D4*, by default, will prefer to barter with the sibling device since there is a long history of interactions and familiarity between these two nodes. *D4* composes proposal requesting *RingTone-B* and in exchange offering either *PodCast-A* or a *RingTone-A* or *C*. *D2* is not interested in the podcast or the ring tones since *D2* is only looking for *Wallpaper-B* which *D4* does not have. However, *D2*, instead of sending a rejection to *D4*, considers the social connection between the owners of the devices and offers *D4* much desired *RingTone-B* with out anything in return.

### II.A.3 Investment Scenario

Devices could also acquire services that they are not planning to personally use. These services could be treated as an investment and could be bartered off at a later time.

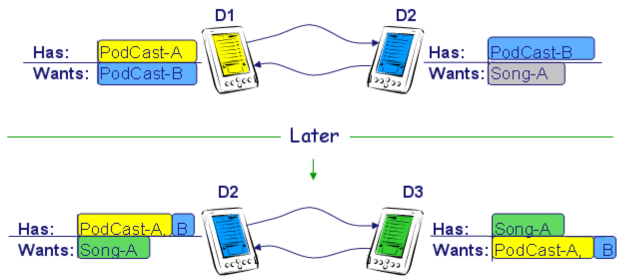


Figure II.2: Acquiring Services for Later Resale

Suppose, device *D1* is interested in acquiring *PodCast-B* (Figure II.2). It discovers that the only device that is capable of offering this podcast is *D2*. Unfortunately, *D2* is not personally interested in *PodCast-A* that *D1* is offering in return. However, despite personal disinterest in the *PodCast-A*, *D2* identifies this ring tone to be in great demand by other devices. So, *D2* looks at *PodCast-A* as an investment that could be later cashed out for another service that *D2* needs. Thus, *D2* accepts the proposal from *D1* and exchanges *PodCast-B* for *PodCast-A*. So now, *D2* has *PodCast-A* and *B* but it is still looking for *Song-A*. At a later time, *D2* meets *D3*. *D2* identifies that *D3* has *Song-A* that *D2* was looking for some time. *D2* composes a proposal that suggests an exchange of *Song-A* for *PodCast-A* and *PodCast-B*. *D3* analyzes the proposal and decides though individually each of the podcasts are not as valuable as the song however their combined value is adequate. Thus after consideration *D3* decides to go through with the exchange. Clearly, *D2*'s the investment into *PodCast-A* paid off. *D3* did not agree to exchange *Song-A* for just one ring tone.

To realize the above described scenarios, we need to address a set of challenges that will affect peer-to-peer collaborations in pervasive environments. In this dissertation, we argue that a bartering communication model is a viable approach to promote collaboration and exchange of digital goods and services.

In the subsequent parts of this chapter, we present the challenges of collaborations and interactions in mobile pervasive environments and also describe alternative approaches that can be used to support exchanges of digital goods and content in mobile environments.

## II.B Challenges of Collaborations in Pervasive Environments

### II.B.1 Free Riders in Pervasive Peer-to-Peer Environments

In Chapter III, we discuss prior and ongoing research in cooperation and collaboration of peers in conventional peer-to-peer settings. Similar to conventional peer-to-peer systems, pervasive environments best function when devices agree to collaborate and coordinate with each other. One of the common visions of pervasive environment interactions is devices asking one another to share their digital goods, services and content to enhance the social experience of their owners. Request for traffic updates, restaurant recommendations, search for cheap gas station are all frequently discussed applications of pervasive environments [42, 94, 28, 104]. Parallels can be drawn between collaborations involved in these applications and collaborations in peer-to-peer music file sharing systems. However, unlike conventional peer-to-peer systems, pervasive environments are populated with small personal devices that are tightly constrained in their resources and computing power. Under these circumstances, inefficient and uneconomical collaborations are far more damaging to the well being of the individual devices and to the computing environment as a whole. Analogous to the conventional peer-to-peer systems, collaborations in pervasive environment are not immune to the problem of free riders [83, 65, 12]. In fact, the limited computing power and the personal nature of the devices that are involved in the collaborations further stimulates these devices to disregard ongoing partnerships and cooperative teamwork efforts in the environment [92, 89, 99, 34]. An example of free rider behavior in pervasive environments is a device that, in an attempt to conserve it's computing resources and power, is ignoring and dropping service requests and other communications from other devices while sending out service and content discovery requests for its own benefit. If a large enough percentage of devices sabotage collaborations, then this incapacitates the computing environment, leaving other cooperative devices potentially isolated and unable to complete their tasks. Free riders are very taxing on collaborative computing environments [39, 36, 26].

To deal with this issue, environments need to have strong incentives that would attract devices with limited resources and limited computing power into honest personally meditated collaborations. The challenge lies in design and development of a communication management component that is capable of collaboration opportunities and exploiting them to derive the benefit. In order to achieve that, management components on these devices need to be able to compute benefits that they would receive from collaborations, so that devices are not unnecessarily overtaxed and over burdened. Incentives need to be embedded into the communication model to ensure that collaborations are done out of self beneficial economically motivated goals.

### II.B.2 Information Noise Makers

Pervasive environments and personal devices are not exempt from spamming and other heavy communication services [114, 16]. Such services could be regarded by many nodes in the environment as useless noise makers. Intense environmental noise has a strong potential of hindering collaborative efforts in the environment by overwhelming the devices with extra information. In such "noisy" environments, devices must be capable of filtering out communications and interactions that are perceived to be of no interest or importance and still be able to participate in relatively useful information exchange and discover desired services. The challenge lies in development of a communication management component that is capable of identifying importance of the communications in a personalized manner [86, 85]. In essence, a communication management component needs to be able to determine the value of incoming data and offered services and make a decision of how critical is this interaction to the state of the device [78, 47]. Acquisition and retention of digital goods, content and services should be done on the basis of perceived value and desire and not on the basis of availability.

### II.B.3 Unpredictable Quality of Services and Goods

When devices collaborate and communicate, they acquire and collect information, goods and services. Depending on the environment and context, it is possible that received data and services are not what were originally expected [53]. Upon acquisition, a device might identify the good or a service to be useless and there is no point it keeping and maintaining this good or service. Similarly, it is possible that when a good or a service was acquired, it was believed to be a very useful, important, and a valuable item. However, over a period of time, it became clear that the potential usefulness or need for this good or service was overestimated [44]. Maintenance and upkeep of this good or a service is consuming limited valuable resources. Yet another possibility is that an obtained good or service was acquired for a short term use. Once the intended time expires, the device should be capable of disposing of the service in an optimal way. Another possibility is that the service is valuable only in a particular rarely occurring context, which is not expected to occur in the near future. Once again, the device should be able to identify such goods and services and possibly dispose of them in an optimal fashion. On the other hand, devices should be able to identify and protect goods and services that are perceived to be very valuable but are rarely used. In addition to that, once a good or a service is deemed insignificant, there are several logical approaches that could be taken to free the device form this good or service. One simple approach is to remove the service and cut the losses of resources that are being expended on maintenance of this good or a service. Another alternate approach is to try to sell the service or to trade it away in exchange for another good. In essence, this approach could be titled as “one mans trash is another mans treasure”.

### II.B.4 Anonymous Collaborations

Conventional views of mobile devices in pervasive environments assume that all of the communications and collaborations are anonymous. Social relationships between the owners of the devices are not propagated into of the collaboration and communication models that are employed. There is a number of reputation and trust based mechanisms that are used to identify cooperative and uncooperative devices in the environment [75, 76, 74]. These mechanisms are powerful and can be combined with real human social relationships. Such relationships are significant since they can be used to explain and justify the levels of cooperation that personal devices exhibit during peer-to-peer collaborations. There are a set of tools such as LinkedIn.com [1], orkut.com [2], mySpaces.com [3] and FaceBook.com [4] that represent such social relationships. Removal of anonymity from collaborations brings additional flexibility to these strategies in pervasive environments. Collaboration strategies between “friendly” devices can be more relaxed since there is the factor of long term relationship between the devices. The challenge lies in developing mechanisms that allow personal devices to follow and execute personalized collaboration strategies that reflect existing social relationships between their owners. The communication management component needs to be able to reason over the relationship status and incorporate the characteristics of the relationship into the communications and negotiations.

## II.C Currency and Micropayment Based Alternative Approach

A possible alternative approach that could address some of the challenges satiated above is a computational economy approach that is based on use of currency and micropayments. In this approach, every digital good, service or information is assigned a dollar value or an abstract system currently value. Other currency, including internal currency or a credits system can also be used to implement this approach. Devices that are interested in obtaining digital content, a good or a service that reflects particular local and temporal context can purchase these items from the mobile peers that surrounded them. This approach would allow nodes to acquire most relevant content. This approach would also ensure that the collaborative nodes in the environment are protected from the predatory interactions initiated by the by free riders. Overall, this economic approach provides the flexibility of interactions. Devices would acquire only the data that they would truly need or be willing to pay for. In essence, this pricing and purchasing approach would create a miniature economy within computing environments.

Unfortunately, every economic approach that involves currency requires a set of mechanisms that address issues of payment processing and currency management. Micropayment mechanisms would need to be used to finalize transactions. However, micropayments have a number of well known shortcomings such as "mental transaction cost" and "user anxiety" [101, 105, 71, 73]. These shortcomings would certainly impact the over all performance of the computational microeconomic approach. Furthermore, the traditional microeconomic approaches rely on the existence of the well connected framework that overlooks these aspects of micropayment related management. In addition the conventional micropayment approach does not take into consideration issues of context sensitive, personalized interactions and social relationships of device owners. The drawbacks of the micropayment mechanisms would also apply to environments that use abstract currency.

The environments that use abstract currency would need to use mechanisms that handle issues of inflation and deflation of that currency. Clearly this will be a complex and taxing process for the mobile nodes that have limited battery and limited computational resources. Though, this economic approach addresses the problem of free riders and noise makers in the computing environment, it does not consider the issues of dynamic changes of perceived value of data.

Finally, there would be a need to develop a pricing mechanism that would be used to describe every item in the environment. This mechanism would need to provide a set of critical functions to the nodes in order to be an effective tool that can be used for dynamic peer-to-peer interactions in the fast pace mobile environment. In addition, this mechanism would have to be light weight to allow nodes to install it and execute its functions on the limited mobile platforms. This mechanism would need to be able to take a set of parameters into account. In particular, this pricing mechanism would need to capture the context sensitive characteristics of the digital content such as the geographical aspects of the data and temporal restrictions and limitations such as data expirations and data staleness. In addition, this mechanism would need to be aware of the personal user settings and preferences, such as upcoming context for this user and personal interests and partialities. Also, the pricing mechanism would have to be able to capture the environmental characteristics such as current and upcoming trends, preferences and interests of the general network population. To achieve this, this mechanism would need to employ complex learning techniques and be able to quickly adapt to the fast changing mobile environments and constantly evolving user context. Clearly, this complex solution would impose a major burden onto the mobile devices that have limited battery and limited computational resources. This heavy computational footprint would incapacitate typical mobile device and thus will not be able to provide robust and effective pricing information for the constantly evolving digital content and fast changing mobile environments.

The currency and the micropayments based approach is capable of delivering a flexible comprehensive toolset for the conventional environments that have strong presents of well established centralized authorities. However, the footprint and robustness of these interaction methods are not well suited for the context of opportunistic, serendipitous peer-to-peer interactions in the pervasive mobile environments that lack these centralized, well established, reliable authorities.

## II.D Bartering Approach

In contrast to the currency and the micropayments based approach, the bartering based approach is well suited for opportunistic collaborative interactions in mobile pervasive environments. Our bartering framework provides mobile nodes with the simple yet powerful policy driven mechanisms that allow mobile nodes to select and employ collaborative interaction method that best reflects their attitude towards collaboration process. In essence, the collaboration attitude are encapsulated in strategies that are represented through a set of well defined negotiation policies. This approach allows nodes to customize their interactions to better reflect their personal preferences and partialities. In addition, our bartering framework is further enhances by the use of the value sensitive collaboration models that acknowledge the perceived value of goods and services. Our bartering framework provides nodes with tools that allow nodes to classify the digital content according to the levels of importance and incorporated this categorization into the bartering process. Our framework also explores the social relationships between the owners of the mobile nodes and incorporates these relationships into the collaborative interaction process.

There are a number of issues that affect bartering interactions. One of the more prominent and characteristic issues associated with the bartering process is the issue of the *double coincidence of wants*. In fact, this issue is one of the major drawbacks of bartering interactions in any type of environment [49]. Basically, in order to have a successful trade, the participants need to desire each other's goods and services. This constraint, in its strictest form, dramatically limits the pool of possible solutions. However, to promote bartering in pervasive environments, our framework also considers mechanisms that broaden the set of possible solutions to facilitate a trade. In particular, we consider a bartering strategy that involves nodes in the exchanges where the bartering partner needs to show some level of reciprocity. This strategy does not require "even" trading. This relaxed approach improves the set of available trading options of goods and services during collaborative process. We also employ social relationship dependent strategy. In essence, the nodes executing this strategy relax the strict constraints of equality of the trade during the exchange with friendly devices. In addition, our framework facilitates socially motivated interactions, Incorporation of this relationship based bartering approach further widens the set of possible solutions since devices have an option of being more tolerant and not driving hard bargains during an exchange with a friendly device.

## II.E Summary

In this chapter, we describe the challenges of collaborations in pervasive environments. We motivate our approach by describing a set of basic collaboration scenarios. We describe a set of challenges that affect traditional pervasive environments. We also discuss alternative solutions to deal with the described challenges. We compare our bartering approach with conventional approaches. And finally, we conclude with a discussion of challenges that need to be addressed to ensure optimal approach to manage collaborations in pervasive environments.

Chapter III

# BACKGROUND AND RELATED WORK

## III.A Introduction

In this chapter, we survey related work, both to point out contributions of previous researchers and to place our proposed contributions in the proper context. We organize this survey around the main themes of our research:

- Pervasive Environment,
  - Context-aware Computing,
  - Context Mediation,
  - Ambient Services,
  - Mobile P2P Computing
- Incentive Mechanisms in P2P Systems
- Mobile Commerce
- Electronic Goods and Services
- Bartering

## III.B Pervasive Environments

Pervasive environments are populated with mobile personal devices such as cell phones PDAs and laptops. These devices host personal applications such as calendars, contact managers and business type applications. But in addition to traditional personal applications, a new set of mobile personalized entertainment services are emerging and rapidly moving onto personal mobile devices. Explosive growth of computational capabilities, dramatic increase in demand for communication and information services and development of integration technologies has resulted in the development of new and effective pervasive computing and communications paradigms. Technologies such as Bluetooth[5] and 802.11[6] enable and empower development to an entire class of applications that allow users to interact, collaborate and share information in a seamless and spontaneous way.

### III.B.1 Mobile P2P Frameworks

In the following sections, we describe frameworks and architectures that were built to facilitate peer to peer interaction in pervasive environments. In addition to the frameworks described below, as a part of our preliminary work, we have developed the *Numi* project. The *Numi* framework employs collaborative agents to facilitate p2p data routing to enrich service environments for mobile devices that operate in infrastructure-based wireless networks. Detailed description of this project can be found in Chapter 4.

#### Infostations

The Infostations project experimented with the notion of data hoarding in mobile environments. They argued that even though cellular voice and data networks facilitate “any time, anywhere” connectivity, they are expensive and offer low bandwidth. Infostation networks[88] have often been suggested as a viable alternative to meet the needs of mobile applications. An infostation network consists of a set of towers offering short-range high bandwidth radio coverage. They offer high-speed discontinuous coverage, which is inherently low cost. Network access is available to users that are passing in close proximity to an Infostation. In this sense, the infostation is similar to a base station coupled with an information server such that the base station provides the network connectivity while the information server handles the data requests. A mobile device thus experiences areas of connectivity (when close to an infostation) and areas of disconnection (when there is no infostation nearby). Specialized data link protocols have been suggested for allowing devices to communicate with such Infostations [33].

**MoGATU**

MoGATU is a lightweight peer-to-peer data management architecture for pervasive environments. The MoGATU framework facilitates serendipitous querying and data management in mobile ad-hoc environments [78] [77]. The MoGATU project regards devices as semi autonomous entities which are guided in their interactions by personal user profiles and the context that devices operate in. MoGATU uses a contract-based transaction model. Information about users is described in the personal profile that are represented in a rich semantic language. Each of the objects is described in terms of ”beliefs”, ”desires”, and ”intentions”. MoGATU introduces data-based routing algorithms and semantic-based data caching and replication algorithms. These algorithms enable mobile devices to utilize their data-intensive vicinities. MoGATU devices also use automated interactions in attempt to obtain data relevant to the user’s ”intentions” and ”desires” which are encoded in the user profiles. MoGATUs transaction models is based on contract net principles for peer-to-peer interaction.

**Proem**

Proem is a peer to peer middleware framework that allows deployment of mobile and ad hoc applications [57]. The key goal of Proem is the support of wearable communities [32]. Wearable communities emerge when sufficient number of people that use wearble computing devices come together and use their devices to communicate and interact with one another [55]. The framework was developed to augment face to face interactions that happen between people that come in contact with one another. Proem relies on personal area networks to create a “digital sphere”. When individuals come in close proximity to each other, their “digital spheres” overlap and this enables devices to communicate. During this stage, devices are free to exchange information and access each others services. This communication can potentially enhance social interactions between individuals and further promote face to face contact. Once the devices move apart, the connection is severed and the spheres detach and end communications. Match making, friend discovery, and file sharing are the sample applications that can be built on top of Proem. The Proem middleware consists of three main components: an application runtime environment, a set of middleware services, and a protocol stack [56]. The application runtime environment in essence is a peerlet engine running Proem compliant applications called peerlets. Peerlets employ an event based model and can be added to the peerlet engine and removed from the engine at runtime. Proem employs a set of protocols to facilitate interactions: transport, presence, data sharing, and community building protocols are used to facilitate interactions between devices. Proem also employs a set of managers to provide basic services for presence management, profile management, data space management and community management. It also provides mechanisms to log encounters and propagate events.

### III.B.2 Context Awareness

Context awareness is a powerful concept that empowers many applications that wound not otherwise be able to operate in pervasive environments. Context aware applications exploit environmental factors such as: time, user's location, current and future events, other users in the environment etc. Describing context, evaluating it, identifying the most relevant information and then reasoning over it has been a focus of many research projects. The aspect of filtering through the contextual information is the most relevant aspect for our proposed approach.

One of the subfields of context aware computing is the use of ambient intelligence [97]. It exploits location relevant contexts. Ambient services are linked to the surrounding physical environment. They have geographical boundaries of relevance and utility. Ambient services are highly applicable in the domain of location based mobile commerce [64]. An example of a mCommerce ambient service is a mobile advertisement service that distributes ads to potential customers that are a short distance away form the store. This type of context dependent advertisement is very effective since it is highly relevant to consumer's proximity to the store.

Contextual mediation[22] can also be used to filter out contextual noise that overwhelms a mobile device that operates in pervasive environments. Contextual mediation [21]is a form of application aware adaptation that is used to manage contextual information requests. During a contextual mediation process, an application selects the most relevant and the most appropriate subset of available contextual data and delivers it to the user. This contextual data is described using semantic representation and a set of attributes. The utility of the data varies between "no interest" to "most interesting". Preferences that reflect users interests are used to narrow the set of contextually relevant data.

#### myCampus Framework

*myCampus* framework [93] is one of the more prominent and well developed projects that incorporates and robustly exploits context wariness in pervasive environments. *myCampus* framework uses *eWallet* management component that is running on uses personal devices. Users personal resources are modeled as semantic web services. Personal and contextual resources and services are represented in using ontologies. Task specific agents can use *eWallet* to reason over the context and application specific reasoning. *myCampus* framework gives special attention to the aspects of privacy for context sensitive applications. *myCampus* project allows automated management of various context dependent applications such as scheduling of a meeting or privacy sensitive inquiry about locations of another user.

## III.C Incentive Mechanisms in P2P Systems

A number of peer to peer file sharing systems, have been developed in recent years. The more prominent examples of music file sharing systems that have proven to be immensely popular are *Napster*, *Gnutella* and *KaZaA* [7, 8, 9]. Participants of these networks can offer songs from their musical collections for others to download. Participants can also search and upload songs from other peers on the network.

However, majority of the peer to peer systems in their original implementations suffer from the free rider problems. Many of the participants did not contribute resources to and only consumed them [12, 83, 65]. A number of approaches have been developed to encourage contributions to collaborations in peer to peer file sharing systems. The approaches involved currency based solutions, and payment based incentives and the use of credit systems.

Below we describe several research projects that address issues of free rider problems in peer-to-peer systems.

#### MojoNation

*MojoNation*[48] was one of the earlier designs that addressed the issues of free riding, freeloading and denial of service attacks in peer to peer systems. It also offered mechanism for balancing resource supply and demand. *MojoNation* followed a free economy approach of using internal currency called Mojo. The *MojoNation* barter system combined a "*digital token*" micropayment system with *peer-to-peer microcredit*. Initially each joining user was given an amount of currency. The user could spend the currency on purchasing the files from other peers. Users earned *Mojo currency* by sharing their disk space, processing power, bandwidth etc. All of the transactions were cleared by the centralized authority of a central bank.

**Karma**

*Karma* is a framework for P2P resource sharing[111]. The *Karma* project addressed issues of free riding in p2p applications through the use of a decentralized currency. Decentralization is one of the key features of *Karma*. This project employs a peer-to-peer scheme for tracking karma currency transfers. This approach protects the systems against malicious attempts to corrupt and alter the currency balance in the system. *Karma* also utilizes a secure exchange mechanism that ensures that participants can not counterfeit karma currency. This project also addresses issue of inflation and deflation to regulate the supply of *karma currency* on the network. These mechanisms provide significant improvement over the unregulated approach of *MojoNation* system. Similar to *MojoNation*, *Karma* employs a reward mechanism to incentivize peers in the system to contribute their resources.

**N-Way Exchange Based Approach**

[13] explored an alternative bartering approach to resolve the free rider problems. Users directly traded resources between themselves. Transfer priority is given to the users that conduct exchanges rather than simple consumption. Simple consumption is permitted only if there are no peers that are willing to participate in the exchange. The authors also explored n-way exchange mechanisms. In the n-way exchanges users formed rings of N peers. Each peer is served by its predecessor and service the successor in the ring.]

**Bittorrent**

A file sharing system, called BitTorrent[27], addressed free riding problem thorugh the use of barteirng model. Bittornet's aproach relied on didviding files into fiexed size compontnes. The componets are stored at multiple peers.

**Our Proposed Approach**

Unlike our proposed approach, above described systems did not consider value of the entities that are being exchanged. These approaches also did not take into consideration social relationships between the users of the system.

## III.D mCommerce

Commerce has long been a critical component that stimulates collaborations[89]. Economically motivated collaborations use incentives of pay off to promote trade and other interactions and collaborations. Our research primarily focuses on collaborations in pervasive environments. Thus, the most relevant form of commerce is mobile commerce. There is a grate diversity of mobile applications that fall the under umbrella of mobile commerce [109]. Mobile financial applications, mobile advertising, mobile inventory management, product location[106], proactive service management[110], wireless business re-engineering, mobile actions, mobile entertainment, vehicular mobile commerce[107], are all considered to belong to the mobile commerce domain [92, 109].

We have done significant amount of preliminary work in the development of mobile commerce middleware that operate in pervasive environments. In particular, we have developed the *eNcentive framework* which focuses on the issues of delivering mobile advertisements to the users in pervasive environments[85]. We have also developed the *Agents2Go* framework that supports location-dependent service discovery in mobile electronic commerce environments[86]. We also have developed the *Allia* framework that uses alliance-based service discovery to discover mCommerece services in ad-hoc environments[84, 87]. Detailed description of theses projects can be found in Chapter4.

A framework for mCommerce depends on four general levels of operations: wireless network infrastructure, mobile middleware, wireless user infrastructure and finally, mCommerce applications. Below, we describe a set of mobile commerce middleware projects and implementations are most relevant to our proposed research.

#### EasiShop - Context Aware Shopping

The *EasiShop* framework[50, 51] explores the concept of cross merchant product comparison shopping through the use of personal mobile agents and virtual market places. The framework assists consumers with their shopping experience. It envisions consumers that are quipped with personal devices that are running *EasiShop Shopping Agent*. The *Shopping Agent* is a mobile agent that is responsible for managing the shopping list and finding the best suited products that match consumers preferences. As, consumers walk by the retailers locations, thier *Shopping Agents* enter the *EasiShop Catchment Zones* where they come in contact with the retailers *Retailer Agents* that is hosted by the wireless device in the shops. If a *Shopping Agent* identifies a product that a consumer is interested in, then the interaction is moved to the *Market Place*.

*Market Place* acts as a forum where consumers and retailers can interact with one another independent of their physical movements. *Market Place* categorizes products into relevant sets which are represented as *stalls*. At a particular *stall* a consumers *Shopping Agent* has a chance to cross-compare the products and their prices. This enables the *Shopping Agent* to make informed decisions and derive recommendation that could be passed to the consumer.

**iClouds - Mobile Advertising System**

*iClouds*[45] is a framework that was developed to support dissemination of advertisements to mobile devices. This framework envisions formation of *information clouds* when several mobile devices come in proximity of each other. Devices are enabled to exchange electronic advertisements that are supplied by various merchants. Anonymous bonus points are used to reward active users that propagate advertisements. Devices keep track of users interests in data structures referred to as *iLists*. *iHave list* contains ads that user has and *iWish list* contains the list of types of ads that user is interested in. During encounters, devices exchange their *iLists*. If a match is found, the ad is moved to the device that desires the ad. The ads propagate from a device to a device creating a chain. The ad keeps information about every device that participated in its propagation. If the ad materializes into a purchase, then every device that assisted with propagation is rewarded with bonus points. The bonus points are managed by the *Mediator*. The *Mediator* in essence is a central database that coordinates merchants and customers. To protect customers privacy dynamic network data and encryption of application layer information is provided during the building of the chain.

**Our Proposed Approach**

We propose to exploit economically motivated incentives to promote collaboration in pervasive environments. Our approach relies on bartering as the exchange mechanism. Economical motivated mobile personal devices are driven to exchange goods and services in an attempt to increase their personal net worth and provide needed goods and services to their users. Personalized valuation of goods and services is at the center of our approach. We also propose to employ personalized bartering strategies that acknowledge and exploit personal relationships between users of the devices. Detailed description of our approach can be found in Chapter 5.

## III.E Digital Goods and Services

MP3s, podcasts, ring tones, screen savers, mobile games, wallpaper, video clips and electronic coupons are invading everyday lives of average consumers. Digital goods, services and information are becoming common mainstream entities. The use of digital goods is becoming part of daily routine for many people. Currently, for a small fee, average consumers can easily acquire any one of the above listed digital goods, store it on a personal device and have a change to enjoy it at any moment they whish. Some goods, such as ring tones, wallpaper and screen savers are used by many individuals as a statement of their individuality and personality. Other digital goods and service such as games, MP3s, podcasts and video clips bring entertainment and leisure. Development of this new class of goods and services has resulted in the development and emergence of a new digital economy. Buying ring tones and MP3s is becoming a matter of social norm. Users can log in into one of hundreds of websites, listen to hundreds of previews of numerous types and various genre of ring tones and just for a couple of dollars they can buy one and install it their cell phone in a matter of seconds.

However, behind all the fun of immensely popular digital goods and services, lies the serious and complex nature of digital goods and services. Formally, a digital good can be defined as: *"a payoff-relevant bitstring that affects the utility of or a payoff to some individual in the economy"*[82]. Digital goods are discrete since they are distributed and acquired in integer amounts. Unprotected digital goods can be easily duplicated and distributed. A copy of a digital good is usually another independent digital good. Digital goods are non-rival goods which means that consumption of this good by one individual or a device does not diminish the amount that is available to others. A digital good is an experience good. In order to get an understanding of quality and content of the goods, one needs to gain access to that good. Digital goods can be fragile and non fragile [82]. A fragile good is a good that will significantly denticulate in value if even small portion of the good is lost. A non fragile good is a good that will not significantly suffer in economic value even if a small portion of that good is lost. All these aspects of the nature of the digital goods have an effect on the distribution and acquisition mechanisms and strategies employed.

This dynamic nature of the new digital economy does not come with out controversies. At the early acceptance stages, sharing of MP3 files was a common practice. Music industry claims that is lost millions of dollars due to the violation of the of the copy rights of the songs that were freely distributed by users of p2p systems such as Napster and Gnutella. This controversy brought into focus issue of high costs involved in production of the original first copy of a digital good.

### III.E.1 Original Cost of Digital Goods

Information, digital goods and services are inherently expensive to produce. On the other hand duplication and reproduction of already existing information, good or service quite inexpensive [98]. For example, if a device wants to produce a traffic report for a particular stretch of road, it needs to collect a substantial amount of sensory data from that stretch of road, analyze it and may be even compare it with previous records. Finally, it needs to assemble the data and convert it into a report format. This device needs to spend a considerable amount of computing and battery power while it's sensors are collecting and processing the data and compiling it into a concise report. On the other hand, if the device wanted to share this information with another device, all it needs to do is to clone this information and transmit it to the other device. Clearly, the cost of producing an additional copy is insignificant compared to the "sunk cost" of the production of the first original copy. In fact, there is no limit to the production of additional copies. If a device can produce one copy, it can produce hundreds of copies. In addition, devices that receive the copies of information can themselves make duplicates and start to distribute this information. This distribution will generate fast growing competition amongst sellers. The end result of this replication and competition will push information price to zero. This clearly contributes to the diminished importance of original cost during valuation of services and data. In pervasive environments, production of the first good is notably more expensive since, the devices that frequently operate in these environments have very limited resources and computational power. To address the information price issue, devices could employ Digital Rights Management (DRM) mechanisms. These mechanisms restrict duplication and potentially slow down the effect of the rapid and uncontrolled information distribution. Unless DRM mechanisms are employed the price of services, information and digital goods in the environment will eventually be pushed to zero.

### III.E.2 DRM - Digital Rights Management

Recent controversial developments in p2p file sharing applications exposed complexities involved in managing and distributing electronic goods and services. In addition to that, these controversies brought to the forefront, the importance of Digital Rights Management (DRM) mechanisms. DRM mechanisms are designed to protect the rights of the creators of original goods and service. Typically, DRM mechanisms incorporate encryption, conditional access, copy control mechanisms, and media identification and tracing mechanisms such as water marks [91].

## III.F Bartering

Bartering is one of the forms of trading [62]. It is also frequently referred to as a “pure exchange”. Bartering is a method of trading services and information directly for one another (without the use of money or other similar unit of account or medium of exchange). In a barter exchange, one good is traded directly for another. Other forms of trading are: bargaining, bidding, auctions, clearing and contracts [72]. The process of bartering shares many common principals of other trading techniques. In bartering, sellers of a good or service are worse-off when there is a large number of other sellers of a similar good or service. Buyers of a good or service will prefer to have as many sellers as possible for the good or service they are buying. Buyers of a good or service prefer to have as few other buyers as possible. Sellers of a good or service would like to have as many buyers for their good or service as possible.

Bartering is an ancient form of trading. It is widely believed that it superseded use of money and currency. Bartering is still commonly used in various social and economical interactions particularly when the infrastructure that facilitates currency based exchanges fails. Recently, during the first few days in the aftermath of hurricane Katrina that hit the city of New Orleans, when communications were down and banks were not operational, local residents and business owners conducted bartering transactions to acquire needed resources.

Bartering has proven to be a reliable tool of exchange when the conducting transactions become unreasonably expensive. The absentce of exchange infrastructure or high costs of conducting transactions can be a motivation factor to barter. This is frequently the case in pervasive environments. Access to bank accounts from a personal device such as a cell phone for a micropayment is not a very complex process. However, the ratio of the transaction cost to the average amount involved in payments is very high for many average consumers. In fact, micropayment mechanisms have long been criticized for high transaction cost and high consumer anxiety [71, 105]. Bartering can alleviate some of the transaction costs in pervasive environments. A user of a mobile personal device seeking particular MP3 would rather give one of his or her MP3s as a payment and not deal with conducting a monetary transaction that would involve bank transfers or a credit card transaction.

Clearly, barter exchange has a set of issues that can hinder a successful exchange. One such issue is *double coincidence of wants*, another is potential of asymmetric knowledge of the quality goods that are being exchanged.

**Double Coincidence of Wants**

One of the critical issues that impacts success of bartering is *double coincidence of wants*. The concept of *double coincidence of wants* is the key behind the traditional economic definitions of a bartering process. The phrase "*double coincidence of wants*" was used in Jevons (1875)[49]. *"[T]he first difficulty in barter is to find two persons whose disposable possessions mutually suit each other's wants. There may be many people wanting, and many possessing those things wanted; but to allow of an act of barter there must be a double coincidence, which will rarely happen."* The translation of this concept into the bartering in pervasive environment is the following. Suppose *Device-1* that has a *Service-A* and wants *Service-B* meets another *Device-2* that has *Service-B* and wants *Service-A*. This convenient coincidence allows *Device-1* and *2* to exchange services and satisfy each others needs and wants. General belief in conventional economics is that this type of coincidence is very uncommon and thus making it difficult to trade.

**Our Proposed Approach**

We propose relaxation of *double coincidence of wants* through use of methods and strategies that can be employed in pervasive environments. To increase the space of possible deals, we propose to exploit social relationships between the owners of the personal devices that operate in pervasive environments. These methods and strategies are discussed in Chapter 5.

**Asymmetric Knowledge**

Sellers of the good or service have full knowledge of the quality of that good or service. On the other hand, a buyer has only partial knowledge that has been provided by the seller. This unequal and asymmetrical knowledge of the quality of goods and service that are being exchanged adds complexity and anxiety which can influence the exchange.

**Our Proposed Approach**

We propose to exploit social relationships of the owners of devices as a guarantee of the quality of service. Devices that belong to users that know each other are unlikely to intentionally mislead and misrepresent the quality of the good that they are offering. Clearly if a relationship is not established, then there is a risk of abuse of asymmetrical knowledge. However, in pervasive environments, an argument can be made that devices that come in contact in the pervasive environment most likely belong to individuals that have some social link to one anther. It is also reasonable to assume that these devices will come in contact again. This close knit, personal nature of pervasive environments will discourage abuse of asymmetry of knowledge of quality of good or service.

## III.G Summary of Related Work

In this chapter, we describe related work to provide the foundation that is necessary to describe our research approach and our expected contributions. We discuss the unique nature of pervasive environments and devices that operate in them. We also describe p2p systems and their approach to addressing the free riding problems during resource sharing. We describe mCommerce approaches to stimulate collaborations in mobile pervasive environments. We also describe and discuss aspects of digital goods and services. In this chapter, we also introduce and motivate the use of bartering exchanges. We also examine the issues that effect the success of bartering. We lay the foundation necessary to motive the use of bartering to stimulate collaborations in pervasive environments.

Chapter IV

# THE MBARTER FRAMEWORK

## IV.A A Typical Opportunistic Bartering Scenario

Consider the following scenario. Bob is an adamant mobile user who recently purchased a trendy new phone and is interested in acquiring a set of ring tones, screen savers, a few images and other digital goods that reflect his strong support for the local football team. His new phone came preloaded with a set of default ring tones and screen savers and Bob is willing and ready to trade them away in an attempt to get new digital goods that better reflect this personal interests. As Bob moves through the geographical region, he meets other mobile users that he can potentially trade with. During his search, Bob comes in radio range with Jane. As their devices discover each others presence, Bob's device initiates a bartering protocol in an attempt to discover a possible trade. Bob's device starts out by sending Jane's device a list of desired goods and a list of goods that can be used as payment. Jane is also interested in bartering so, Jane's device examines the two lists and compares them with Jane's lists of desires and potential "give away" goods. It identifies that it has an image that will interest Bob and it also identifies that Bob has a hip new polyphonic ringtone that interests Jane. Since there is a match, Jane's device consults Jane's collaboration policies. Jane's collaboration policy can be categorized as *DCoW* policy i.e. Jane is reluctant to give goods away with out receiving a similar good in return. Since collaboration with Bob is an exchange of goods that complies with Jane's collaboration policy, Jane's device proposes this exchange to Bob's device. Upon receiving this proposal Bob's device examines the proposal by comparing it against Bob's collaboration policies. Bob's calibration policy is similar to Jane's, so Bob's device approves the exchanges. This approval is followed by the transfer of goods which is reconfirmed by a set of acknowledgment messages. In this manner, Bob receives the picture that he is interested in and Jane gets a new ring tone. This type of exchange is beneficial to both mobile users since they are able to acquire needed goods in an efficient manner.

Clearly, the outcome of the bartering interactions very much depends on the collaboration policies that are employed by the mobile users. Consequently, there is a wide range of policies that the devices can employ to help users negotiate and barter in mobile peer-to-peer environments. The combinations of these policies create collaboration strategies that the nodes follow as general guidelines during such bartering encounters. In fact, one can think of these strategies as points along "the collaboration continuum" [96, 95]. Our research considers four basic strategies from this continuum: a *FreeRider*, an *Altruist*(also referred to as philanthropist [96]), a *Weak Double Coincidence of Wants* (*WDCoW*) and a *Double Coincidence of Wants* (*DCoW*). The *FreeRider* strategy and the *Altruist* strategy are at the opposite ends of the continuum. These two strategies have been extensively used in a number of collaboration applications including the traditional wired peer-to-peer systems [12, 80, 70]. The *WDCoW* and the *DCoW* strategies lie along the continuum, where the *WDCoW* is on the more cooperative side of the continuum and *DCoW* is on the less cooperative side of the continuum. Both of these strategies are the derivatives of fundamental concepts from the barter exchange domain.

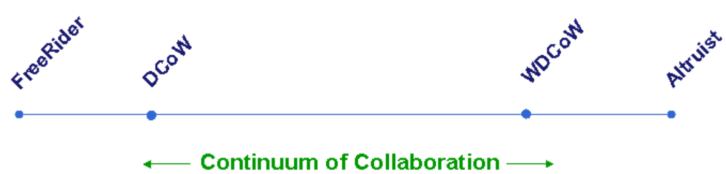


Figure IV.1: Strategies in the Collaboration Continuum.

To better understand the acceptable outcomes for each of the strategies, consider the following set of short examples.

An uncooperative *FreeRider* node would only be willing to participate in interactions where:

- This *FreeRider* node gets 5 goods and pay 0 goods.

A less cooperative *DCoW* would be willing to participate in interactions where:

- this *DCoW* node gets 3 goods and gives away 3 goods,
- this *DCoW* node gets 5 goods and gives away 0 goods (i.e. *DCoW* can act as a benefiting *FreeRider*),
- this *DCoW* node gets 5 goods and gives away 3 good (i.e. *DCoW* can act as a benefiting *WDCoW* / *Altruist*).

A more cooperative *WDCoW* node would be willing to participate in interactions where:

- this *WDCoW* node gets 5 goods and gives away 1 good,
- this *WDCoW* node gets 1 good and gives away 5 goods,
- this *WDCoW* node gets 5 goods and gives away 0 goods (i.e. *WDCoW* can act as a benefiting *FreeRider* when dealing with another *Altruist*).

On the other hand, the most cooperative *Altruist* node would be willing to participate in interactions where:

- this *Altruist* node gets 3 goods and gives away 5 goods,
- this *Altruist* node gets 0 goods and gives away 5 goods (i.e. *Altruist* helps *FreeRider*),
- this *Altruist* node gets 5 goods and gives away 0 goods (i.e. *Altruist* can act as a benefiting *FreeRider*).

The formal descriptions of the policies for these four strategies are defined in Section IV.C.

## IV.B The Framework Architecture

To study the proposed bartering collaboration model, we have developed a framework that allows discovery of common interests, facilitates negotiations and supports exchange and transfer of digital goods.

### IV.B.1 Bartering Structures

Our framework uses two basic structures to facilitate bartering interactions: an *iHave List* and an *iWant List*. An ***iHave List*** is set of goods and services that a device is willing to disclose. Digital goods and services on this list can be used as payment in bartering transactions. An ***iWant List*** is set of goods and services that a device desires and is actively searching for. When a good or service from this list is acquired, the record is moved to the node's *iHave List*. These lists are central to the bartering communications. The nodes advertise/exchange these lists (or subsets of these lists) to find common points of interests. Upon a "match", a node invokes its bartering strategy by generating an appropriate Proposal. In the current design and implementation, there are no privacy mechanisms that protect the *iWant List* and the *iHave List* from being snooped (by the nodes that have no "matches"). However this design limitation can be addressed with existing well established mechanisms. For example, to provide a level of discreetness, *iWant* and *iHave List* can be indexed using Bloom filters [19, 116]. Peer nodes can query these indices in attempt to establish a possible "match".

We envision that a real world version of the proposed *bartering communication model* would rely on user's personal profile and user specified policies on what goods to add to the *iWant List*. Clearly, this is an interesting research area and there is a lot previous and ongoing work in this research rich domain [58, 78, 23].

### IV.B.2 Bartering Communication Protocol

In an attempt to find a bartering partner, a node generates an *Inquiry* message that gets broadcast to the nodes one hop neighbors. This node is referred to as an *Inquiring node* (or an ***Inquirer***). A *Proposing node* (or a ***Proposer***) is a node that generates the *Proposal* message in response to an incoming *Inquiry* message.

Nodes don't concurrently negotiate with multiple nodes. A state model within our bartering protocol ensures that a node participates in only one session at a time. A node can be in one of the following states: ***Free***, ***Proposed***, ***Accepted***, ***ACKed Accept*** and ***Sent Goods***. A node does not enter in to a new session unless it is in a *Free State*. A ***State Manager*** is responsible for maintaining the states of the node. The *State Manager* also sets and keeps track of timers for each *non-Free* state. The timers are used to prevent locking of the node in any unfinished sessions that could occur in dynamic ad hoc environments. If a timer is triggered and the node is still in the same state for the same transaction, the *State Manager* concludes that this session is not progressing (either due to the fact that nodes moved out of range or some other errors such as network congestion). As a result of this error, the *State Manager* abandons this session and resets the node's state to *Free*.

Our Bartering Protocol uses eight message types to communicate between nodes: *Inquiry*, *Proposal*, *Accept*, *Reject*, *Busy*, *ACK Accept*, *Delivery* and *ACK Delivery*.

An ***Inquiry*** message is broadcast by the *Inquirer* (when it is in a *Free* state). This message is broadcast to the node's one hop neighbors approximately every 15 seconds. It is also possible to extend our simulation study and allow nodes to set a larger *TTL*. An *Inquiry* message is composed of a time stamp, *Session ID*, *iWant List* and *iHave List*. A ***Session ID*** is a unique number that represents the session (it is derived from the *node's ID*).

Upon receiving an *Inquiry* message, a node's *State Manager* checks if the node is in a *Free* state. If not, then the *Inquiry* is discarded. Otherwise, the node's ***Proposal Composition Manager*** examines the *Inquiry* using its ***Proposal Composition Policy***. If there is no suitable "match", the *Inquiry* is dropped. Otherwise, the

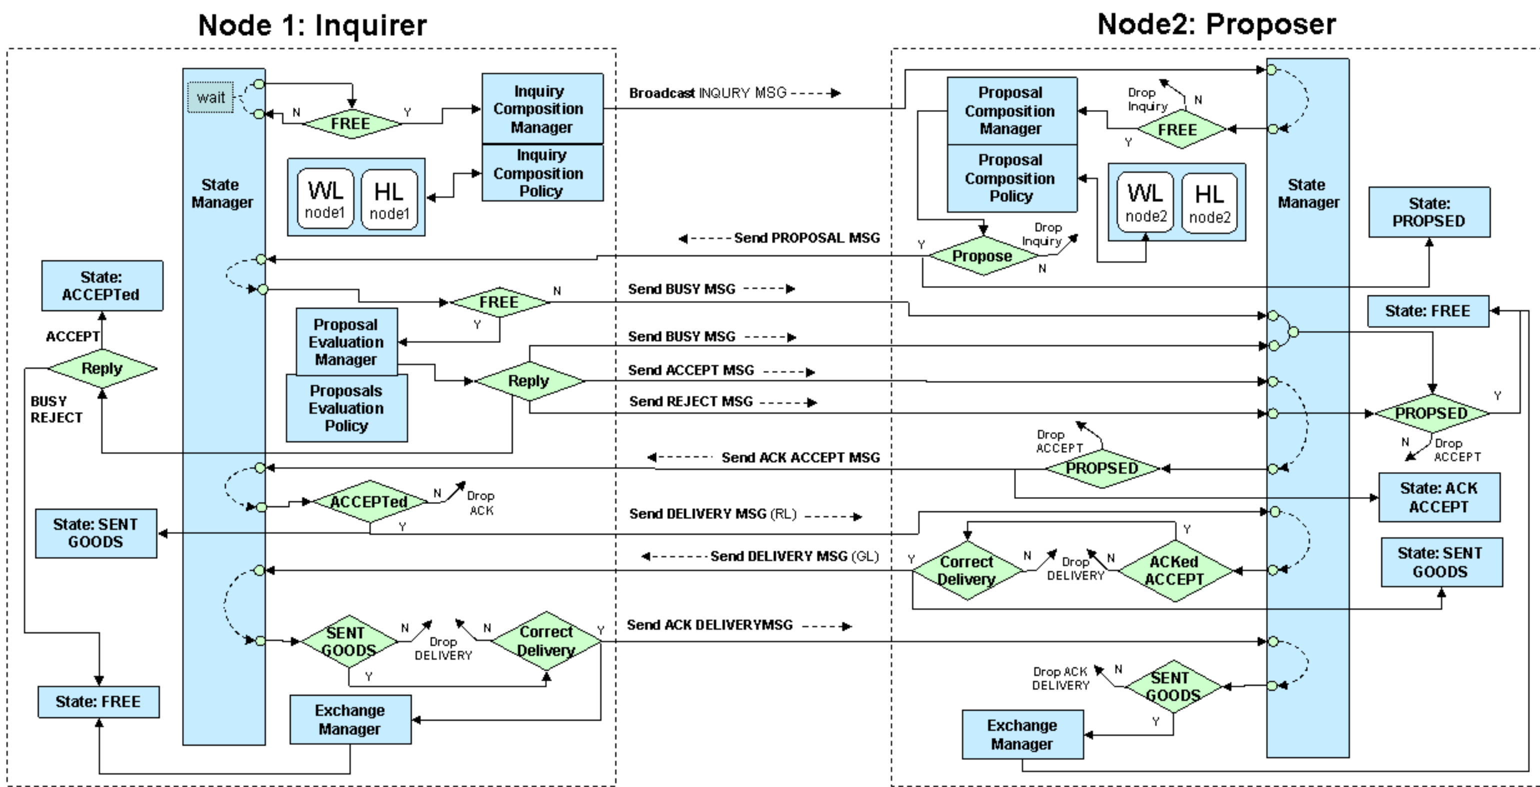


Figure IV.2: Bartering Communication Protocol Flow

*Manager* generates an appropriate ***Proposal***. This *Proposal* is compared with all the rejected *proposals* that this node may already have sent to this *Inquiring* node. If it matches with any of the prior *Rejections*, then the *Proposal* is discarded and node does not respond to the *Inquiry*. This comparison ensures that the node does not repeatedly propose the same exchange. If it is a new *Proposal*, then the *Proposing* node sends it to the Inquirer. All *Proposal* messages consist of a *Give List* and *Receive List*, *Proposer's* node ID, *Inquirer's* node ID, *Session ID* from the *Inquiry* message and the timestamp of the *Inquiry* message. A ***Give List*** is a subset of the intersection of the *iHave List* of the *Proposer* (the *Proposing* node) and the *iWant List* of the *Inquirer* (the *Inquiring* node) $GL \in (HL_R \cap WL_I)$. Basically, a *Give List* contains a set of goods that the *Proposing* node is willing to offer to the node that is looking for the goods. A ***Receive List*** is a subset of the intersection of the *iHave List* of the *Inquirer* and the *iWant List* of the *Proposer* $RL \in (HL_I \cap WL_R)$. Basically, the *Receive List* contains a set of goods that the *Proposer* wants to receive as a payment for the goods in the *Give List* that it is offering to the *Inquirer*. A ***Proposal Composition Policy*** is applied to these *Lists* to insure compliance with the node's strategy. From the point of view of the *Proposer*, a *Receive List* and a *Give List*, together, constitute a potential "match" between the needs and haves of the two nodes.

Each *Inquiring* node has a ***Proposal Evaluation Manager*** which is responsible for examining the incoming *Proposals*. This *Manager* uses the ***Proposal Evaluation Policy*** to determine if the incoming *Proposal* is acceptable or not. The *Proposal Evaluation Policy* examines the *Receive List* and the *Give List* and checks whether the proposed *match* complies with the *Inquirer's* strategy. If the *Proposal* is not acceptable a *Reject* message is generated and sent to the *Proposer*. All acceptable *Proposals* are collected by the *Proposal Evaluation Manager* for a period of time. The *Manager* examines this collection of *Proposals* and using *Proposal Evaluation Policy*, identifies the "best" *Proposal*. This "best" *Proposal* is accepted and the winning *Proposer* is notified by an *Accept* message. The rest of the *Proposals* are discarded and the loosing *Proposing* nodes are sent *Busy* messages. These *Busy* messages, unlike *Reject* messages, do not discourage the *Proposers* from sending similar *Proposals* at a later time.

A ***Reject*** message, an ***Accept*** message, and a ***Busy*** message is sent by the *Inquirer* in response to a *Proposal* from the *Proposer*. All three message types contain *Proposer's* node *ID*, *Inquirer's* node *ID*, *Session ID* and the timestamp of the *Inquiry* message.

Upon receiving a *Reject*, a *Busy* or an *Accept* message, the *State Manager* at a *Proposer* checks if it's state is set to *Proposed* (for the correct session). If not, then the message is discarded. Otherwise, if *Reject* or *Busy* message was received, the *State Manager* resets its state to *Free*. For a *Reject* message, the *Proposer* (for future *Proposal* compositions and comparisons) records the specifics of the rejected *Proposal*. Note that, a rejected *Proposal* will not be resent to the same *Inquirer*. This reduces traffic on the network and prevents nodes from unnecessarily re-proposing already failed attempts of exchange. In case of the *Busy* message, the node assumes that the *Inquirer* is in the middle of another exchange session. Upon receiving *Accept* message, the *Proposer* issues an ***ACK Accept*** message to the *Inquirer* and sets it state to *ACKed Accept*. Once the *Inquirer* gets this *ACK Accept*, its *State Manager* checks if the node is in the state of *Accepted* (for this session). If not, then the *ACK Accept* message is dropped. Otherwise, the node composes and sends a ***Delivery*** message, to the *Proposer* and sets its state to *Sent Goods*. A *Delivery* message is composed of the *Receive Lists* that was agreed on by the two nodes, *Proposer's* node *ID*, *Inquirer's* node *ID*, *Session ID* and the timestamp of the *Inquiry* message.

Once the *Proposer* receives a *Delivery*, its *State Manager* checks if the node is in the state of *ACKed Accept*. If not, the *Delivery* is dropped. Otherwise, the *Proposer* checks if the *Delivery's Receive List* complies with the reached agreement. If it does not comply, then the node drops the *Delivery* and resets its state to *Free*. Otherwise, the *Proposer* composes and sends a *Delivery* message to the *Inquirer* and resets its state to *Sent Goods*. In contrast to the first *Delivery*, this *Delivery* contains the *Give Lists* (and not the *Receive List*), it also contains *Proposer's* node *ID*, *Inquirer's* node *ID*, *Session ID* and the timestamp of the *Inquiry* message.

Once the *Inquirer* receives a *Delivery* message, its *State Manager* checks if the node is in the state of *Sent Goods*. If not, then the *Delivery* is dropped. Otherwise, the node checks if the *Delivery's Give List* complies with the previously sent agreement. If not, then the *Delivery* is dropped and the *Inquirer's* state is reset to *Free*. Otherwise, the contents of the *Delivery's Give List* are transferred into the node's *iHave List* and are also removed from the *iWant List*. If the node is operating in a *DRM* (*Digital Rights Management*) setting, then the goods that were sent to the *Proposer* are removed from the *iHave List*. Once the transfer is complete, the node composes an ***ACK Delivery*** message and sends it to the *Proposer*. This message contains the *Proposer's* node *ID*, the *Inquirer's* node *ID*, the *Session ID* and the timestamp of the *Inquiry* message.

Upon receiving an *ACK Delivery*, the *Proposer* checks if the node is in the state of *Sent Goods*. If not, then the *ACK Delivery* is dropped. Otherwise, the node transfers the content of the *Delivery's Receive Lists* into its *iHave List* and removes the newly acquired goods from it's *iWant List*. If the node is operating in a *DRM* setting, then the goods that were sent to the *Inquirer* are removed from the *iHave List*. Finally, the node sets its state to *Free*.

## IV.C Formal Definitions of Collaboration Strategies

We have implemented four basic strategies: *FreeRider*, *Altruist*, Weak Double Coincidence of Wants (*WD-CoW*) and Double Coincidence of Wants (*DCoW*). A node's strategy is defined by the policies used to compose *Proposals* (*Proposal Composition Policy*) and to respond to the received *Proposal* (*Proposal Evaluation Policy*). During collaborations, none of the nodes are aware of the strategies of their neighbors.

### IV.C.1 A *FreeRider* Strategy

The ***Free Rider*** strategy is a well covered subject in P2P systems such as Gnutella [8, 12] and Kazaa[9]. Our interpretation of this behavior does not deviate from the conventional view of *FreeRiders*. The motto of a *FreeRider* node is *"get goods with out paying for them"*. A node that executes this strategy composes an *Inquiry* such that it conceals the node's true *iHave List*. The *Inquiry* contains an empty *iHave List* and a copy of the *iWant List* of the node. As a result, the only *Proposals* that a *FreeRider* node will ever receive in response to its *Inquiry* will only benefit this free riding node: ($\mid GL \mid > 0 \;\; \&\& \;\; RL = \emptyset$ ). Thus, following the *Proposal Evaluation Policy* for a *FreeRider Strategy*, all of the incoming *Proposals* will be accepted by the *FreeRider*. Similarly, upon receiving an *Inquiry* from another node, a *FreeRider* will follow its *Proposal Composition Policy* and will examine only the *iHave List* of this incoming *Inquiry*. If it finds an intersection between the *Inquiry's iHave List* and the *FreeRider's iWant List* ($\mid RL \mid > 0$) then the node composes a *Proposal* where it's ($RL = HL_R \cap WL_I \;\; \&\& \;\; GL = \emptyset$), otherwise it ignores this *Inquiry*. The *Proposal Composition Policy* and the *Proposal Evaluation Policy* ensure that a *FreeRider* node only *accepts* goods and never *gives away* its goods to other nodes in the environment. Thus, a *FreeRider* is a purely selfish node.

**FreeRider's Proposal Composition Policy:**

$\big(|RL| > 0 \quad \&\& \quad |GL| == 0\big)$

**FreeRider' Proposal Evaluation Policy:**

Accepts: $\big(\, |RL| > 0 \quad \&\& \quad |GL| == 0\big)$

Rejects: $\big(\, |GL| > 0\big)$

### IV.C.2 An *Altruist* Strategy

In contrast to the *FreeRider* strategy, an ***Altruist*** is a purely selfless node. An *Altruist* will help/barter with all other nodes in the environment (including *FreeRiders*). The motto of an *Altruist* node is *"if I can help you I will, if you can help me please do so"*. One can also think of this strategy as a "good neighbor" strategy. An *Altruist's Inquiry* will contain its complete *iHave List* and *iWant List*. If this *Inquiry* generates a *Proposal*, an *Altruist* will *Accept* this *Proposal* regardless of whether this *Altruist* node benefits from the *Proposal* or not. In essence, the *Proposal Evaluation Policy* of an *Altruist* is to accept *Proposals* that are: $\big((|GL| \geq 0 \quad \&\& \quad |RL| > 0) \quad \| \quad (|GL| > 0 \quad \&\& \quad |RL| \geq 0)\big)$. This ensures that an *Altruist* never rejects an incoming *Proposal*. Similarly, upon receiving an *Inquiry* from another node, an *Altruist* will generate a *Proposal* if it can compose a *Proposal* such that: $\big((|GL| \geq 0 \quad \&\& \quad |RL| > 0) \quad \| \quad (|GL| > 0 \quad \&\& \quad |RL| \geq 0)\big)$. This *Proposal Composition Policy* insures that a node responds to all *Inquires* where one of the nodes (regardless of which one) can benefit from this exchange. These policies ensure that an *Altruist* collaborates with all the nodes in the environment regardless of their strategy. Thus, an *Altruist* is a purely selfless node.

**Altruist's Proposal Composition Policy:**

$\big((|GL| \geq 0 \quad \&\& \quad |RL| > 0) \quad \| \quad (|GL| > 0 \quad \&\& \quad |RL| \geq 0)\big)$

**Altruist's Proposal Evaluation Policy:**

Accepts: $\big((|GL| \geq 0 \quad \&\& \quad |RL| > 0) \quad \| \quad (|GL| > 0 \quad \&\& \quad |RL| \geq 0)\big)$
Rejects: *NA*

### IV.C.3 A *Weak Double Coincidence of Wants* Strategy

A node executing the ***WDCoW*** strategy relaxes the strict "double coincidence of wants". The node's motto is *"to barter only when both of nodes can benefit from the exchange (*)"*. A *WDCoW's Inquiry* contains its complete *iHave List* and *iWant List*. If this *Inquiry* generates a *Proposal* from another node, the node will consult it's *Proposal Evaluation Policy* and will check if $\big(|GL| \geq 0 \quad \&\& \quad |RL| > 0\big)$. This node will not respond to *Inquiries* where $\big(RL = \emptyset\big)$. Thus, a *WDCoW* will ignore all *Inquiries* from *FreeRiders*. A *WDCoW* will also ignore some of the *Inquiries* from *Altruists*. This will occur only if an *Altruists* is interested in a set of goods that the *WDCoW* has $\big(|RL| > 0\big)$ in its *iHave List*; however, the *Altruist* is only node that benefits from this exchange ($GL = \emptyset$). However, a *WDCoW* will take advantage of *Altruists'* offers where $\big(|GL| > 0 \quad \&\& \quad RL = \emptyset\big)$. In essence, a *WDCoW* will agree to an exchange where an *Altruist* is willing to give goods to this *DCoW* and the *Altruist* wants nothing back from this *WDCoW*. (*) Thus, collaborations between *WDCoWs* and *Altruists* are sometimes *asymmetrical*. When it comes to generating *Proposals* in response to incoming *Inquires*, a *WDCoW's* policy of *Proposal Composition* is to generate *Proposals* such that $\big(|HL_I \cap WL_R| > 0 \quad \&\& \quad |HL_R \cap WL_I| > 0\big)$. Thus, a *WDCoW* will never propose an exchange where only one of the nodes benefits from the collaboration (regardless of the other node's strategies). Thus, a *WDCoW* will never propose to a *FreeRider*. In addition, any *Proposal* generated by a *DCoW* will always be accepted by another *WDCoW* node or by an *Altruist* node.

**WDCoW's Proposal Composition Policy:**

$\big(|RL| > 0 \quad \&\& \quad |GL| > 0\big)$

**WDCoW' Proposal Evaluation Policy:**

Accepts: $\big(|RL| > 0\big)$

Rejects: $\big(|RL| == 0\big)$

### IV.C.4 A *Double Coincidence of Wants* Strategy

A *DCoW* strategy is one where a node looks for "even" exchanges. This node's motto is *"to barter only when both of nodes can benefit (*) from the exchange and the benefits are EQUAL"*. In essence, a node with this strategy counts the number goods that it would receive from this collaboration and the number of goods it needs to give up during this collaboration. A *DCoW's Inquiry* contains its complete *iHave List* and *iWant List*. If this *Inquiry* generates a *Proposal* from another node, the *DCoW* node consults its *Proposal Evaluation Policy* and checks if $\big(|RL| \geq |GL| \geq 0 \quad \&\& \quad |RL| > 0\big)$.

**DCoW's Proposal Composition Policy:**

$\big(|RL| == |GL| > 0\big)$

**DCoW' Proposal Evaluation Policy:**

Accept: $\big(|RL| \geq |GL| \geq 0 \quad \&\& \quad |RL| > 0\big)$

Reject: $\big(|GL| > |RL| \geq 0\big)$

### IV.C.5 Summary of Bartering Policies

The Table IV.1 summarizes the formal description of the policies for each of the four strategies.

#### Symmetries and Asymmetries of Collaborations

In each of the strategies, the *Proposal Composition Policy* and *Proposal Evaluation Policy* are symmetrical. When nodes of the same strategy collaborate, they never issue *Proposals* that would be rejected by the *Proposal Evaluation Policy* of the *Inquiring* node. For example, a *WDCoW* node never rejects a *Proposal* from another *WDCoW* node. Note that a *FreeRider* never generates *Proposals* addressed to another *FreeRider*.

## IV.D Stylized Models

To further describe interactions between the above described strategies, we present stylized model that offer probability based description of the *Gains* that result for the interaction between two nodes executing a particular strategy.

- $\boldsymbol{V}$ represents the total *Gain* of the system/environment
- $\boldsymbol{i}$ and $\boldsymbol{j}$ represent goods that are being exchanged during the transaction
- $\boldsymbol{M_{BA}}$ represents the probability that node *A* meets node *B* such that node *B* has something that node *A* wants
- $\boldsymbol{P_{A_w i}}$ represents the probability that node *A* wants good $\boldsymbol{i}$
- $\boldsymbol{P_{A_h i}}$ represents the probability that node *A* has good $\boldsymbol{i}$
- $\boldsymbol{X_{Ai}}$ represents the perceived *Value* of good $\boldsymbol{i}$ by node *A*
- $\boldsymbol{Y_{Ai}}$ represents the *Cost* incurred by acquiring and maintaining good $\boldsymbol{i}$ by node *A*
- $\boldsymbol{Z}$ represents the overhead interaction *Cost* of conducting the transaction.

Table IV.1: Places for Collaboration Strategies

| | Proposal Composition Policy | ProposalEvaluation Policy | | Inquiry Reply from Nodes |
|---|---|---|---|---|
| | | Accept | Reject | |
| **FreeRider** | $(\lvert RL\rvert > 0 \ \ \&\& \ \ \lvert GL\rvert == 0)$ | $(\lvert RL\rvert > 0 \ \ \&\& \ \ \lvert GL\rvert == 0)$ | $(\lvert GL\rvert > 0)$ | *Altruist* |
| **Altruist** | $((\lvert GL\rvert \geq 0 \ \ \&\& \ \ \lvert RL\rvert > 0) \ \ \lVert \ \ (\lvert GL\rvert > 0 \ \ \&\& \ \ \lvert RL\rvert \geq 0))$ | $((\lvert GL\rvert \geq 0 \ \ \&\& \ \ \lvert RL\rvert > 0) \ \ \lVert \ \ (\lvert GL\rvert > 0 \ \ \&\& \ \ \lvert RL\rvert \geq 0))$ | $\emptyset$ | *Altruist, WDCoW, DCoW* |
| **WDCoW** | $(\lvert RL\rvert > 0 \ \ \&\& \ \ \lvert GL\rvert > 0)$ | $(\lvert RL\rvert > 0)$ | $(\lvert RL\rvert == 0)$ | *Altruist, WDCoW, DCoW* |
| **DCoW** | $(\lvert RL\rvert == \lvert GL\rvert > 0)$ | $(\lvert RL\rvert \geq \lvert GL\rvert \geq 0 \ \ \&\& \ \ \lvert RL\rvert > 0)$ | $(\lvert GL\rvert > \lvert RL\rvert \geq 0)$ | *Altruist, WDCoW, DCoW* |

### IV.D.1 FreeRider meets FreeRider

Since *FreeRider* do not interact with one another, the *Gain* of the system is an empty set.

$V = \emptyset$

### IV.D.2 FreeRider meets Altruist

As described in the earlier sections, when a *FreeRider* meets an *Altruist*, the *FreeRider* experiences positive *Gain* at the expense of the *Altruist*. In order for a transaction to occur when two such nodes meet, the *Altruist* must have something that the *FreeRider* wants. Using the terms described above, the system *Gain* from such an interaction can be modeled as:

$V = M_{AB}\big\{\sum_{i=0}^{n} P_{A_w i}P_{B_h i}(X_{Ai} - Z_{AB}) - \sum_{i=0}^{n} P_{A_w i}P_{B_h i}(Y_{Bi} + Z_{AB})\big\}$

The probability that a *FreeRider* node *A* will meet an *Altruist B* such that *B* has something that *A* wants is captured in $M_{AB}$. The first summation represents the *Gain* of the *FreeRider* node for all goods that *A* wants that *B* has and will offer up (since it is an *Altruist*). The second summation represents the *Loss* of the *Altruist* node (due to the goods given away).

### IV.D.3 FreeRider meets WDCoW or DCoW

When *FreeRiders* and *WDCoWs* meet these nodes do not interact thus, the system experiences no such transactions. Similarly, when *FreeRiders* and *DCoWs* meet they do not directly interact with each other. Thus, the universe of such transactions is an empty set.

$V = \emptyset$

### IV.D.4 Altruist meets Altruist

When two *Altruists* encounter one another the transaction occurs when one of the nodes is interested in the content possessed by the other node. The exchange policies of these nodes allow nodes to collaborate when $\big((|GL| \geq 0 \quad \&\& \quad |RL| > 0) \quad || \quad (|GL| > 0 \quad \&\& \quad |RL| \geq 0)\big)$.

These all of the *Altruist*-to-*Altruists* interactions can be reduced into a *FreeRiders*-to-*Altruist* type interactions when viewed from each node's perspective. Essentially, the exchange only depends on a single set of coincidence of "want" and "have". Using the terms described above, the system *Gain* from such an interaction can be modeled as:

$$V = M_{AB}\big\{\sum_{i=0}^{n} P_{A_wi}P_{B_hi}(X_{Ai} - Z_{AB}) - \sum_{i=0}^{n} P_{A_wi}P_{B_hi}(Y_{Bi} + Z_{BA})\big\}$$
$$+M_{BA}\big\{\sum_{j=0}^{n} P_{B_wj}P_{A_hj}(X_{Bj} - Z_{BA}) - \sum_{j=0}^{n} P_{B_wj}P_{A_hj}(Y_{Aj} + Z_{AB})\big\}$$

The first term of the equation represents the *Gain* of the *Altruist A* node for all goods that *A* wants that *Altruist B* has and will offer up. The second term of the equation represents the *Gain* of the *Altruist B* for all goods that *A* has to offer to *Altruist B*.

### IV.D.5 Altruist meets WDCoW

Consider a meeting between an *Altruist* node and a *WDCoW* node. Let the *Altruist* be a node *A* and let the *WDCoW* be a node *B*. The interactions between these nodes are guarded by the exchanges policies of these nodes. Using the terms described above, the system *Gain* from the interactions between *Altruists* and *WDCoWs* can be modeled as:

$$V = M_{BA}\big\{- \sum_{i=0}^{n} P_{A_hi}P_{B_wi}(Y_{Ai} + Z_{AB}) + \sum_{i=0}^{n} P_{B_wi}P_{A_hi}(X_{Bi} - Z_{BA})\big\} +$$
$$+\{M_{AB}M_{BA}\{\sum_{i \neq j}^{n} P_{A_wi}P_{B_hi}P_{B_wj}P_{A_hj}(X_{Ai} - Y_{Bj} - Z_{AB})$$
$$+ \sum_{i \neq j}^{n} P_{A_wi}P_{B_hi}P_{B_wj}P_{A_hj}(X_{Bj} - Y_{Ai} - Z_{BA})\}\}$$
$$* \{M_{AB}\{\sum_{i=0}^{n} P_{A_wi}P_{B_hi}(X_{Ai} - Z_{AB}) - \sum_{i=0}^{n} P_{A_wi}P_{B_hi}(Y_{Bi} + Z_{BA})\}\}$$
$$+\{M_{AB}M_{BA}\{\sum_{i \neq j}^{n} P_{A_wi}P_{B_hi}P_{B_wj}P_{A_hj}(X_{Ai} - Y_{Bj} - Z_{AB})$$
$$+ \sum_{i \neq j}^{n} P_{A_wi}P_{B_hi}P_{B_wj}P_{A_hj}(X_{Bj} - Y_{Ai} - Z_{BA})\}$$
$$*M_{BA}\{- \sum_{j=0}^{n} P_{B_wj}P_{A_hj}(Y_{Aj} + Z_{AB}) + \sum_{j=0}^{n} P_{B_wj}P_{A_hj}(X_{Bj} - Z_{BA})\}\}$$

One possible transaction that can occur between these two nodes is where the *Altruist* proposes to "help out" the *WDCoW* node by giving a set of goods to this *WDCoW* with out any the *WDCoW* reciprocating this action. This philanthropic transaction can be described as a *free riding* interaction. Essentially, from the prospective of the *Altruist* node, the node *B* acts as a *free riding* node. Using the terms described above, the system *Gain* from this interaction is represented by the first term of the equation presented above.

Another possible interaction can be characterized as an exchange of a set of $K$ goods that node *A* wants and node *B* has for another set of $N$ goods that node *B* wants and node *A* has (such that $N \cup K = \emptyset$ and $N > K > 0$). This transaction can be reduced into an "even" exchange transaction complemented by the *free riding* transaction. Essentially, the "even" exchange is represented by the transfer of $|K|$ number of goods from the $K$ set to the node *A* and the transfer of $|K|$ number of goods from the $N$ set to the node *B*. This transaction is further complimented by the *free riding* interaction which involves transfer of $|N - K|$ number of goods from the $N$ set to the node *B*. This transaction can be characterized as an interaction where the node *WDCoW* benefits from the transaction more than the *Altruist*. This *Gains* from this interaction are modeled by the second term of the equation presented above.

Finally, the third possible interaction is the interaction where the *Altruist* node receives a greter benefit from the transaction than the *WDCoW* node. Similar to the previously described interaction, this transaction can be reduced to an "even" trade interaction complemented by the *free riding* interaction. This interaction is modeled as a third term of the equation presented above.

### IV.D.6 Altruist meets DCoW

Consider a meeting between an *Altruist* and a *DCoW*. Let the *Altruist* be a node *A* and let the *DCoW* be a node *B*. Using the terms described above, the system *Gain* from interactions between *Altruists* and *DCoWs* can be modeled as shown below.

$$V = M_{BA}\{-\sum_{i=0}^{n} P_{A_h i} P_{B_w i}(Y_{Ai} + Z_{AB}) + \sum_{i=0}^{n} P_{B_w i} P_{A_h i}(X_{Bi} - Z_{BA})\}+$$
$$+\{M_{AB}M_{BA}\{\sum_{i\neq j}^{n} P_{A_w i} P_{B_h i} P_{B_w j} P_{A_h j}(X_{Ai} - Y_{Bj} - Z_{AB})$$
$$+\sum_{i\neq j}^{n} P_{A_w i} P_{B_h i} P_{B_w j} P_{A_h j}(X_{Bj} - Y_{Ai} - Z_{BA})\}$$
$$*M_{BA}\{-\sum_{j=0}^{n} P_{B_w j} P_{A_h j}(Y_{Aj} + Z_{AB}) + \sum_{j=0}^{n} P_{B_w j} P_{A_h j}(X_{Bj} - Z_{BA})\}\}$$

As described in the previous sections the *DCoW* nodes are intolerant of philanthropic behavior when they are on the loosing end of the transactions. Thus, the only *free riding* type interaction can occur if the *Altruist* is the *Initiator* of such transaction. The first term of this equation presented below models this *free riding* type interaction.

Another possible transaction that can be initiated by the *Altruist* is the transaction that delivers a greater benefit to the *DCoW* node *B* rather than to the *Altruistic* node *A*. As mentioned earlier, the *Proposal Evaluation Policy* of the *DCoW* accepts the transactions where the *DCoW* node befits more than the other collaborating partner. This transaction can be modeled as an "even" exchange transaction complemented by the *free riding* type interaction. This particular type of transaction between the *Altruists* and *DCoW* is modeled by the second term of the equation presented above.

Finally, the interaction could also be initiated by the *DCoW* node. In this case, the *Proposal Composition Policy* dictates the use of the "even" exchange method of collaboration. This particular type of interaction is also modeled by the second term of the equation presented above.

### IV.D.7 WDCoW meets WDCoW

Consider a meeting between a *WDCoW* and another *WDCoW*. Let the first *WDCoW* be a node *A* and let the second *WDCoW* node be a node *B*. Using the terms described above, the system *Gain* from interactions between these two *WDCoWs* can be modeled by equation presented below.

As described in the previous sections, the *WDCoW* nodes are intolerant of being on the loosing end of the pure philanthropic transaction. However, they participate in interactions where there are unequal levels of reciprocity. Thus, the *free riding* type interaction will never occur between these nodes.

Basically, there are two types of transactions that can occur between the two *WDCoW*. In the first category of transactions, the node *A* is the node that benefits the most from the exchange. While in the seconds category of the transactions, the node *B* benefits the most from the exchange. The first term of this equation models the first transaction type. While, the second term models the second transaction type. Both of these transactions can be subdivided into an "even" exchange interaction complemented by the *free riding* type interaction.

$V = \{M_{AB}M_{BA}\{\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h j}(X_{Ai} - Y_{Bj} - Z_{AB})$

$+\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h j}(X_{Bj} - Y_{Ai} - Z_{BA})\}\}$

$*\{M_{AB}\{\sum_{i=0}^{n} P_{A_w i}P_{B_h i}(X_{Ai} - Z_{AB}) - \sum_{i=0}^{n} P_{A_w i}P_{B_h i}(Y_{Bi} + Z_{BA})\}\}$

$+\{M_{AB}M_{BA}\{\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h j}(X_{Ai} - Y_{Bj} - Z_{AB})$

$+\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h j}(X_{Bj} - Y_{Ai} - Z_{BA})\}$

$*M_{BA}\{-\sum_{j=0}^{n} P_{B_w j}P_{A_h j}(Y_{Aj} + Z_{AB}) + \sum_{j=0}^{n} P_{B_w j}P_{A_h j}(X_{Bj} - Z_{BA})\}\}$

### IV.D.8 WDCoW meets DCoW

Now, lets consider a meeting between a *WDCoW* and a *DCoW*. Let the *WDCoW* be a node *A* and let the *DCoW* node be a node *B*. The exchange policies of both of these nodes are un-acceptable to the purely philanthropic exchanges where these nodes are at the loosing end of the transaction. However, as previously stated, the *DCoW* nodes are tolerant of the interactions that do not deliver the equal levels of reciprocity.

$V = \{M_{AB}M_{BA}\{\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h j}(X_{Ai} - Y_{Bj} - Z_{AB})$

$+\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h j}(X_{Bj} - Y_{Ai} - Z_{BA})\}\}$

$*\{M_{AB}\{\sum_{i=0}^{n} P_{A_w i}P_{B_h i}(X_{Ai} - Z_{AB}) - \sum_{i=0}^{n} P_{A_w i}P_{B_h i}(Y_{Bi} + Z_{BA})\}\}$

$+M_{AB}M_{BA}\{\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h j}(X_{Ai} - Y_{Bj} - Z_{AB})$

$+\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h j}(X_{Bj} - Y_{Ai} - Z_{BA})\}$

Taking these policies into consideration, there are two possible interactions that can take place when the node *A* and the node *B* meet. The first type of interaction is initiated by the *WDCoW* node. This transaction is composed such that the node *A* receives greater benefit for the transaction than the node *B*. This interaction is modeled as an “even” exchange transaction complimented by the *free riding* type interaction benefiting node *B*. The first term of the equitation presented above models the system *Gains* from this type of interaction. Note that, if the *WDCoW* node generates a *Proposal* that delivers this *WDCoW* a greater benefit than the benefit received by the *DCoW* node then this *Proposal* is rejected and does not contribute to the *Gains* of the system.

Finally, the type of interaction that is initiated by the *DCoW* node reflects the *DCoW's* collaborative attitude. Essentially, these interactions are initiated in accordance with the *DCoW's Proposal Composition Policy* and thus can be described as an “even” exchange interaction. The system *Gains* form this interaction are modeled by the second term of the above presented equation.

### IV.D.9 DCoW meets DCoW

Finally, lets consider a meeting between a*DCoW* and another *DCoW*. Let the first *DCoW* be a node *A* and let the second *DCoW* node be a node *B*. As described in the previous sections, the *DCoW* nodes are intolerant of philanthropic behavior. In addition, these nodes also demand even levels of reciprocity. Thus, considering the *Proposal Composition Policy* of the participating nodes the only exchange that can occur between the two *DCoWs* are “even” exchange interactions. Therefore, in order for a transaction to occur when two such nodes meet, the node *A* must “have” something that the node *B* “wants” and the *B* must “have” something that the node *A* “wants”. Using the terms described above, the system *Gain* from such an interaction can be modeled as:

$$V = M_{AB}M_{BA}\big\{\sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h J}(X_{Ai} - Y_{Bj} - Z_{AB})$$
$$+ \sum_{i\neq j}^{n} P_{A_w i}P_{B_h i}P_{B_w j}P_{A_h J}(X_{Bj} - Y_{Ai} - Z_{AB})\big\}$$

## IV.E The Framework Evaluation Metric

To validate the viability of our approach and to evaluate the performance of our communication model, we have developed a simulation for a large scale deployment of nodes executing one of the collaboration strategies. We have used GloMoSim [117] as our modeling tool (detailed configuration description of GloMoSim setup is in Section IV.E.1).

To establish a baseline for the performance of each of the strategies, initially, we consider four homogeneous networks (running one of the four strategies). This allows us to compare the effectiveness of each of the strategies and determine specific characteristics that are descriptive of the performance of each of the considered collaboration strategies.

In addition to basic homogeneous networks, we also examine the influence that *InfoStations* can inflict on each of the four homogeneous networks. An *InfoStation* is a wireless "information kiosk" that is capable of providing data services and high speed connectivity to mobile wireless devices that come into its communication range [33, 88, 115, 37]. This addition of InfoStations to our MANETs is motivated by our desire to simulate the recently emerging trend of wireless access points(using 802.11 technology[6]) and providing near by wireless users with high-speed, relatively cheap wireless access. In the context of our work, *InfoStations* inject additional digital goods into the environment and act as ultra cooperative data rich wireless "information kiosks"[54].

Consequently, we examine heterogeneous networks that are composed of a mixture of nodes running different strategies. We start by examining an *evenly mixed network* where, the network population is evenly divided into four quarters, each running one of the four strategies (i.e. there is no majority or minority). We continue our study by exploring networks that have a distinct majority (40% of network population) running one of the strategies and three evenly sized minorities (20% of network population) each running one of the other three strategies. For example we look at a mix where population is split into 40% of *Altruists*, 20% of *FreeRiders*, 20% of *WDCoWs* and 20% of *DCoWs*. Essentially, we examine thirteen distinct networks: four homogeneous networks, four homogeneous networks with five *InfoStations*, evenly mixed network and finally the four networks with a distinct majority and minorities.

In addition to the above described network population configurations we also examine effects of employing the value-based evaluation paradigm of digital goods and content. Furthermore, we study the effects of the role based social relationships on the bartering process. The details of these extensions are in Chapter VIII and Chapter IX.

As we examine the above described networks, we look at the following factors and behaviors to identify the key aspects that are representative of each of the considered strategies and network population compositions:

- **Average Gains** and **Average Losses** - the average *Gain* represents the average number of goods acquired by a node as a result of collaborations that this node participated in. Similarly, the average *Loss* represents the average number of goods a node had to give away during its collaborative exchanges and interactions.
- **Average Transaction Count** - represents the average number of transactions that a node participated in.
- **Average Size of Transactions** - represents the average number of goods that are involved in an exchange or interaction.
- **Average Number of Peer-to-Peer Transactions** - represents the average number of transactions that a node participated in where the bartering partner is another mobile node. This comes in to play when there are InfoStations present in the environment.
- **Average Size of Peer-to-Peer Transactions** - represents the average number of goods that are involved in a p2p exchange. This comes in to play when *InfoStations* are added to the environment.
- **Outcome of Collaborative Interactions**- not all interactions end in a successful exchange. Some *Proposals* could get rejected, or get a *Busy* reply and transactions could be terminated due to node movements or due to the congestion of the network.
- **Unfulfilled Wishes** - To offer additional measurement of productivity of each of the strategies we count number of wishes that are still on the *iWish List*. We also examine the nodes progress over a period of time.

The effectiveness of a strategy is directly proportional to the *Gains* that a node experiences during its collaborations. The effectiveness is also inversely proportional to the *Losses* that a node incurs due to these collaborations. In addition, communication related overhead has an impact on the effectiveness of a strategy. For example, a strategy is more efficient if a node participates in a few large volume (or high *Gain*) transactions than many small ones to minimize protocol overhead. Similarly, the interactions that end in a *Proposal Rejection* are very taxing on the *Proposing* node since it unproductively worked on creating a “match and sending the *Proposal*.

### IV.E.1 Framework Evaluation Configurations

To evaluate our communication model and study the scalability of our framework which employs the proposed digital good exchange protocol, we have built a prototype simulation using GloMoSim 2.02 simulation tool [117]. The following table provides GloMoSim configuration parameters that were used in our simulation studies.

| **GloMoSim Parameters** | |
|---|---|
| **Number of Nodes** | 30-70 (+10) |
| **Number of InfoStations** | 0 - 5 |
| **Total Simulation Time** | 100 min |
| **Terrain Size** | 1000m x 1000m |
| **Node Speed** | 9-13 mps (20-30 mph), Pause 10 sec |
| **Mobility Model** | Random Waypoint |
| **Range for Node and InfoStations** | 30 m |
| **Average Node Degree** | 0.12 - 0.23 |

#### Collaborative Environment Parameters

In addition to the GloMoSim configuration parameters, our simulation study has the following application level configurations:

- There are 450 unique goods in the universe.
- All nodes start out with 50 goods in *iHave List* (randomly selected from 450).
- All nodes also start out with 50 goods in *iWant List*, randomly selected according to the Pareto distribution [31, 24, 15] where $f(x) = 450 * (x^{-0.8309})$. FigureIV.3 shows the $f(x)$.
- All *InfoStations* have 450 goods in their *iHave List* and 0 in their *iWant List*.

- There is no good duplication or replication by the nodes. However, *InfoStations* have unlimited number of copies of every good. Thus as two peers interact, a node gains goods that the other node loses. Therefore, in a network with no *InfoStations* the total number of goods in the environment does not change.

- All nodes and *InfoStations* have a cap of giving only 25 goods per transaction.

- Every node (if *Free*) broadcast its *iHave List* and its *iWant List* every 14sec + 500 mil-sec.

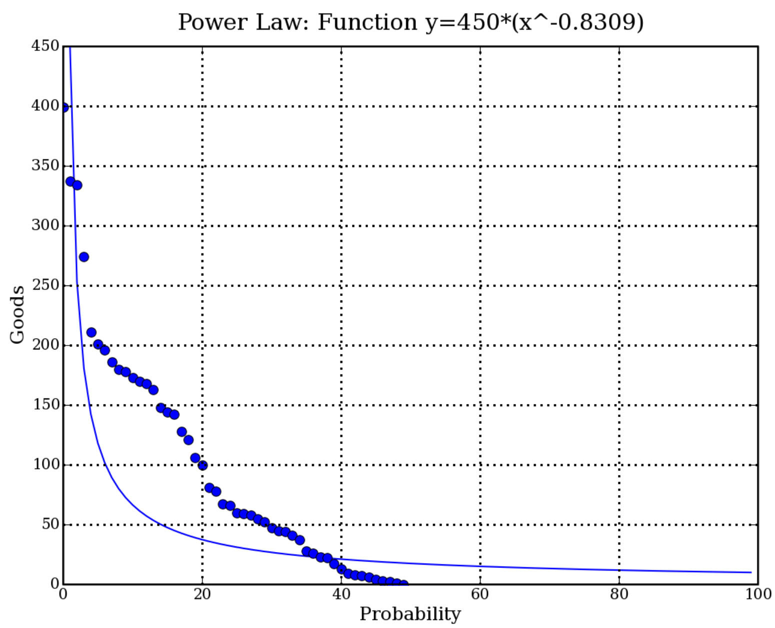


Figure IV.3: Pareto distribution

## IV.F Summary and Discussion

This chapter provides detailed description of our framework that facilitates the collaborative interactions. This chapter also presents the collaboration approaches and the policies that are employed by the nodes during the exchange process. In particular this chapter describes formal policies for collaborations for the two conventional interaction strategies of FreeRiders and Altruists and for the two bartering interaction of Double Coincidence of Wants (DCoW) and weak DCoW. In the consequent chapters, we will preset a characterization study that evaluates each of the strategies parented in this chapter.

**Chapter V**

# COLLABORATIVE EXCHANGES IN HOMOGENEOUS NETWORKS

In this Chapter, we will look at collaborations in homogeneous networks populated with mobile nodes executing one of the four exchange strategies described in Section IV.C and compare performances of each of the four strategies according to the metrics defined in Section IV.E. We also look at the impact of adding *InfoStations* to these homogeneous networks.

## V.A Comparing Gains and Losses in Homogeneous Networks

As mentioned in Chapter IV, each of the strategies displays different levels of cooperation when nodes interacts. As expected, these levels of cooperation permeate into the levels of *Gains* and *Losses* that nodes experience as a result of these interactions. For example, the *Altruists* are the most cooperative nodes in the "collaboration continuum" while the *FreeRiders* are the least cooperative. Figure V.1 clearly shows that the *Altruists* enjoy the highest *Gains* while the *FreeRiders* have absolutely no *Gains* at all (for purposes of graph clarity, the *Gains* of the *FreeRiders* are not in the graph). Similarly, the *Gains* of the *WDCoWs* and the *DCoWs* are reflective of the stands of these strategies in the "collaboration continuum". The lines in Figure V.1 in addition to representing *Gains*, also represent the *Losses* of the collaborative strategies. However, in homogeneous networks with no *InfoStations*, the average *Loss* equals to the average *Gain* since, the nodes receive goods from their peers. Since, in these networks, one node's *Gain* is another node's *Loss*, the average *Gain* is equivalent to the average *Loss*.

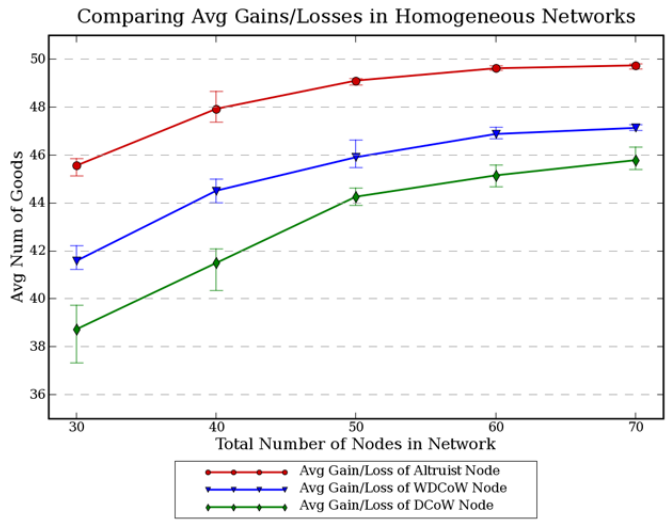


Figure V.1: Comparing the Average Gains/Losses of Nodes Running Different Strategies in Homogeneous Networks

Now, lets examine the effects of adding *InfoStations* to the homogeneous networks. FigureV.2 shows the *Gains* for the nodes in the networks with no *InfoStations* (represented by the dashed lines) and the *Gains* for the nodes in the networks with five *InfoStations* (represented by the solid lines). Regardless of the presence of the *InfoStations*, the levels of *Gain* that these nodes experience are closely related to the levels of collaboration exhibited by the nodes. For example, a *FreeRider* node, running the least collaborative strategy, gains significantly less (even with five *InfoStations* in the network) compared to the other nodes with more collaborative strategies (with no *InfoStations* in the network). On the other hand, the *Altruist* nodes, which are running the most collaborative strategy, receive the highest *Gains*. Similarly, the *WDCoW* strategy is also a cooperative strategy; however, it is not as cooperative as the *Altruist* strategy. The *WDCoW's Gains* are slightly lower but are still very similar to the *Gains* of the *Altruists* (particularly, when the *WDCoWs* operate in the network with five *InfoStations*). Finally, the *Gains* of the less collaborative *DCOW* nodes fall short of the *Gains* of the *Altruists* and the *WDCoWs* but, clearly surpass the *Gains* of the *FreeRiders* (even in the network with five *InfoStation*). Also, for the three collaborative strategies, the nodes in the larger networks see higher *Gains* since there are more opportunities for these nodes to interact with their peers. As *InfoStations* are added to the homogeneous network, the *Gains* of all four strategies increase. As shown in Figure V.2, the *FreeRiders* benefit the most from the presence of the *InfoStations*. For the other three collaborative strategies, the benefits are very similar.

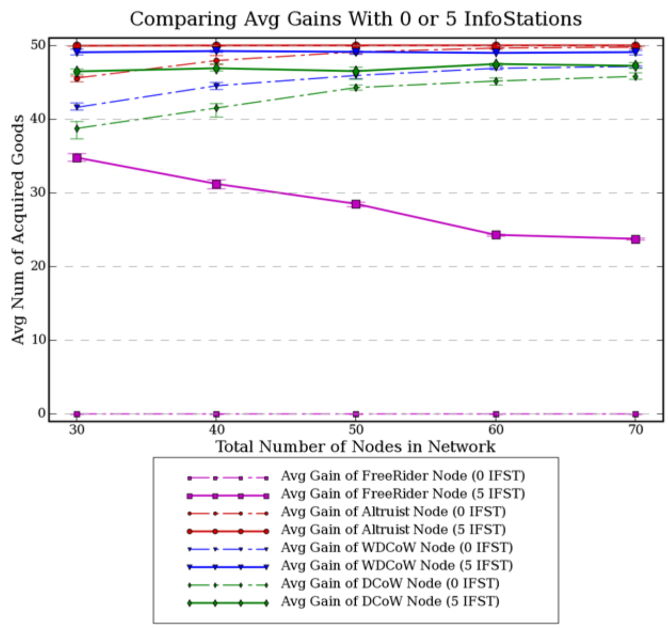


Figure V.2: Comparing the Average Gains of Nodes Running Different Strategies with 0 or 5 InfoStations

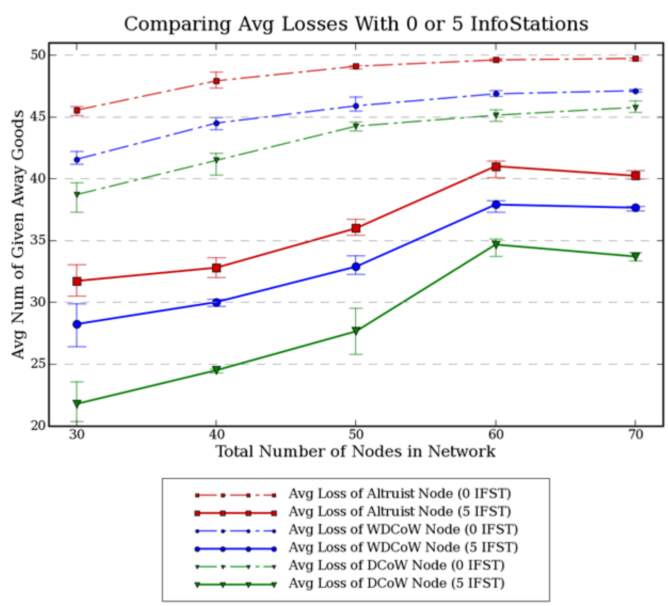


Figure V.3: Comparing the Average Losses of Nodes Running Different Strategies with 0 or 5 InfoStations

In addition to the *Gains*, mobile nodes also experience *Losses* which are shown in Figure V.3. In homogeneous networks with no *InfoStations*, the average *Losses* equal to the average *Gains* since, the nodes receive goods from their peers. Once the *InfoStations* are added to the homogeneous networks, the average *Losses* decline due to the fact that the nodes are able to get digital goods from the *InfoStations* in addition to their peers. So, let's look at the *Losses* of the nodes in the networks with five *InfoStations*. Similar to the *Gains*, the *Losses* are related to the levels of cooperation. The *FreeRiders* experience no *Losses* as they do not give anything away (for graph clarity, they are not part of the graph). On the other hand, the most cooperative *Altruists* experience the highest *Losses*. However, the *InfoStations* provide some relief to the *Altruists*. The *Losses* of the *WDCoW* and the *DCoW* nodes are not as high as the *Losses* of the *Altruists* but also reflect their levels of cooperation. All of these findings are very much in line with original expectations.

## V.B Communication Overhead in Homogeneous Networks

In addition to the *Gains* and the *Losses*, the effectiveness of every exchange strategy can be influenced by the communication-related overheads. Each successful transaction is preceded by an exchange of control messages. If the transaction has a small return, the overhead from these communications can have a significant impact on the efficiency of the exchange strategy. Figure V.4 shows the average number of successful transactions while, Figure V.5 shows the average size of these transactions. Let's first consider the homogeneous networks with no *InfoStations*. The most interesting finding that comes out of these two graphs is that the nodes executing the *WDCoW* strategy have the lowest communication overhead in homogeneous networks. The *WDCoW* nodes participate in the smallest number of transactions and have the largest average transaction size. This makes the nodes that are running the DCoW strategy the most efficient collaborators. On the other hand, the *Altruists* have the largest communication overhead since; they have the largest average transaction size and a large number of transactions. Thus, when comparing overheads of the *Altruists* and the *WDCoWs*, the *WDCoW* strategy is more efficient than the *Altruistic* strategy. Now let's look at the communication overhead of the *DCoW*. Interestingly, the *DCoWs* participate in more transactions than the *WDCoWs*; however, the average size of the *DCoWs'* tractions is small compared to the average size of the *WDCoWs'* transactions. These overheads make the *DCoW* strategy a less efficient collaborative strategy than the *WDCoW* strategy. Now, let's compare the overheads of the *DCoWs* and the *Altruists*. The average transaction count of the *DCoWs* is similar to the average transaction count of the *Altruists*. However, the *DCoWs* show better average transaction size. Thus, the *DCoWs* have a lower transaction overhead than the *Altruists*.

Now, let's look at the homogeneous networks with five *InfoStations*. All of the cooperative strategies experience a similar jump in the average transaction size. This relatively evenly distributed improvement is attributed to the fact that the *InfoStations* generously give away goods to any requesting node. However, the presence of the *InfoStations* does not evenly effect the average number of transactions. The *WDCoWs* and the *DCoWs* see a moderate decrease in the average transaction size while the *Altruists* see a significant decrease. Thus, when adding *InfoStations* to the homogeneous networks, the *Altruists'* communication costs significantly decrease compared to the two bartering strategies. The key reasons behind the differences in these communication overheads lie in the *Proposal Composition Policies* and the *Proposal Evaluation Policies* of each of the strategies. The *Altruists'* exchange policies make nodes very responsive and overly cooperative; while, the *WDCoWs'* policies make nodes more focused on the *Gain*-related cost effectiveness of the ongoing exchange. On the other hand, the *DCoWs'* policies make nodes consider the fairness of the ongoing exchange which naturally limits the size of the transactions.

In essence, when it comes to deriving communication related efficiency, the *WDCoW* strategy is the most economical strategy across the board. It fares very well when the nodes operate in the homogeneous networks without *InfoStations*. Upon adding of the *InfoStations*, the *WDCoW's* efficiency increases (since the number of transactions further decreases and the average transaction size gets a significant jump). On the other hand, nodes running the *Altruist* strategy prove to be the least efficient communicators particularly, in the network with no *InfoStations*. The *Altruists* participate in a larger number of small-sized transactions thus, increasing their communications overhead. However, when adding *InfoStations*, the *Altruists* see the most dramatic improvement in communication costs. Finally, the nodes executing the *DCoW* strategy fare well without the help from the *InfoStations*. With the presence of the *InfoStations*,the *DCoWs* and the *Altruists* have similar communication-related overheads; however, neither one of theses two strategies is as efficient as the *WDCoW* strategy.

### V.B.1 Communication Overheads in Peer-to-Peer Transactions

To better understand the initial conclusions about the communication related costs, we take a closer look at the communication overheads generated specifically by the peer-to-peer exchanges. Examining this subcategory of communication overhead allows us to better understand the impact of the *InfoStations* in homogeneous networks. Figure V.6 and Figure V.7 show the overhead differences that the nodes experience when the *Info-*

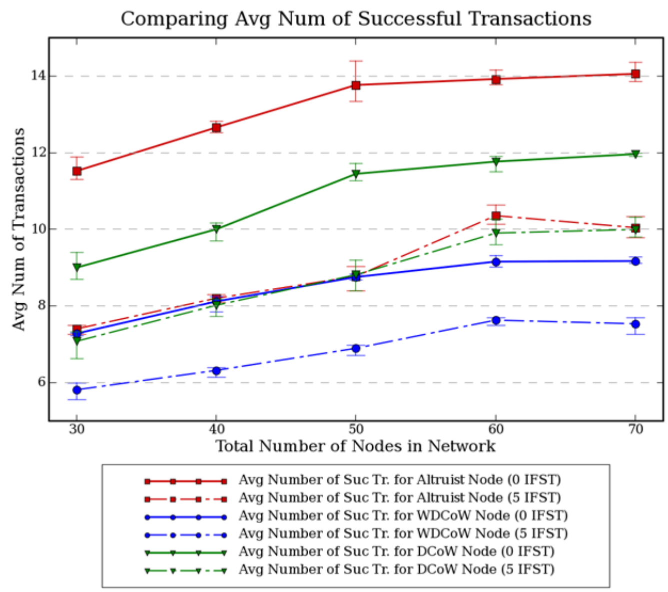


Figure V.4: Comparing the Average Number of Successful Transactions of Nodes Running Different Strategies with 0 or 5 InfoStations.

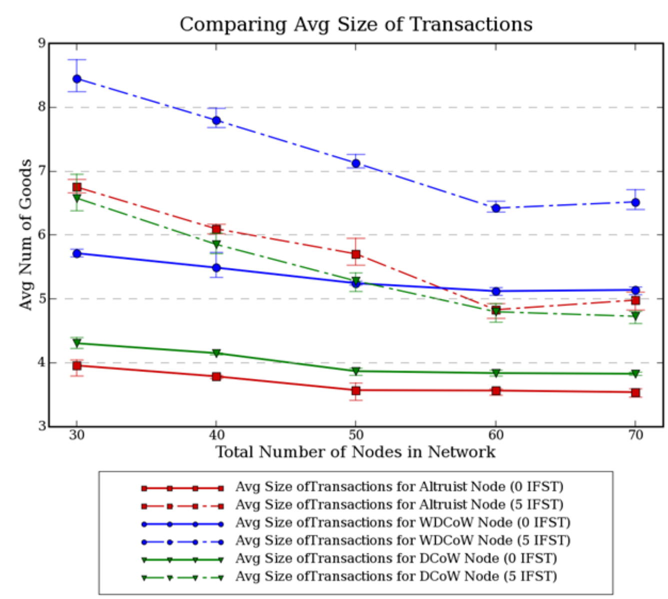


Figure V.5: Comparing the Average Size of Successful Transaction of Nodes Running Different Strategies with 0 or 5 InfoStations.

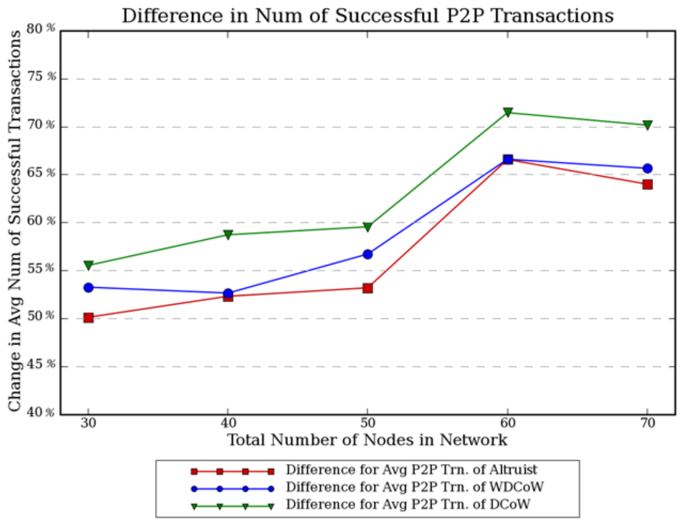


Figure V.6: Differences in the Average Number of Successful Peer-to-Peer Transactions Between 0 InfoStations and 5 InfoStations.

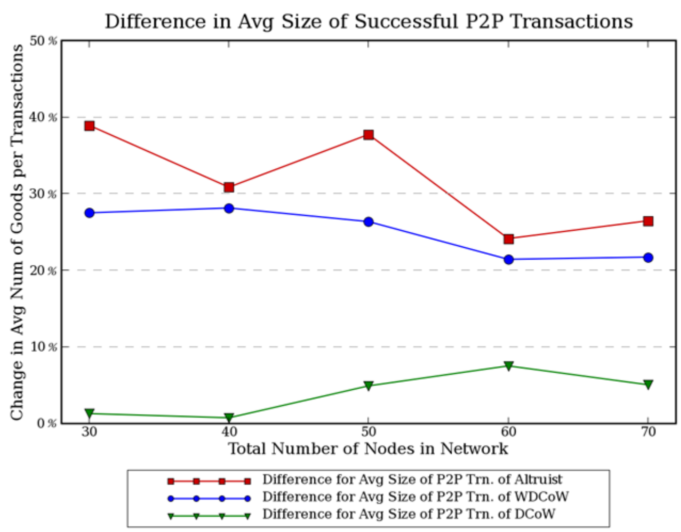


Figure V.7: Differences in the Average Size of Successful Peer-to-Peer Transactions Between 0 InfoStations and 5 InfoStations.

*Stations* are added to the homogeneous networks. Note that, the *Altruists* experience the largest improvement in their peer-to-peer communication costs. They see the largest relative decline in the peer-to-peer transaction count and also experience the largest relative increase in the average peer-to-peer transaction size. Basically, the presence of the *InfoStations* has the most influence on the peer-to-peer exchanges between the nodes executing the *Altruist* strategy. In contrast, the *DCoWs* experience the lowest relative declines in the peer-to-peer communication overheads. The relative increase of their average *DCoW* peer-to-peer transaction size shows the lowest improvement. In addition, the relative rate of the *DCoW's* peer-to-peer transaction count remains relatively high compared to the other collaborative strategies. Thus, the nodes executing the *DCoW* strategy do not take advantage of the *InfoStations* as the *Altruists* do. The *WDCoWs'* communication overheads undergo a greater improvement than the overheads of the *DCoWs*. The *WDCoWs'* see a relatively big jump in their average peer-to-peer transaction size and also see a relatively promising decline in a number of peer-to-peer exchanges. All of the above described changes in the overhead costs are the result of the collaboration policies of each of the considered strategies. The *Altruists'* policies are the most open and supportive of collaborative exchanges thus, they are more likely to be affected by the changes in the environment while the *DCoWs'* policies are more restrictive thus, this make the nodes executing this strategy less likely to be affected by the environmental changes such as the presence of the *InfoStations* in the network.

### V.B.2 Transaction Failures

So far, we have considered communication overheads that occur as a result of successful transactions. However, in mobile wireless environments, transactions can also fail. These failures can occur due to the node's motion (nodes move out of range) or due to the noise in the wireless medium. These failures add to the communication-related inefficiencies which is not present in the more traditional, less dynamic collaborative environments. Figure V.8 shows the average number of failed transactions for each of the three cooperative strategies. First, let's look at the networks with no *InfoStations* (represented by the solid lines). The *Altruists* suffer the highest rates of transaction failures while the *WDCoWs* experience the lowest failure rate. These differences can be attributed to the number of transactions that each of the strategies participates in. Once the *InfoStations* are added, the failure rates continue to reflect communication levels of each of the strategies. These results are in line with our expectations and further compliment the findings of already described communication overheads.

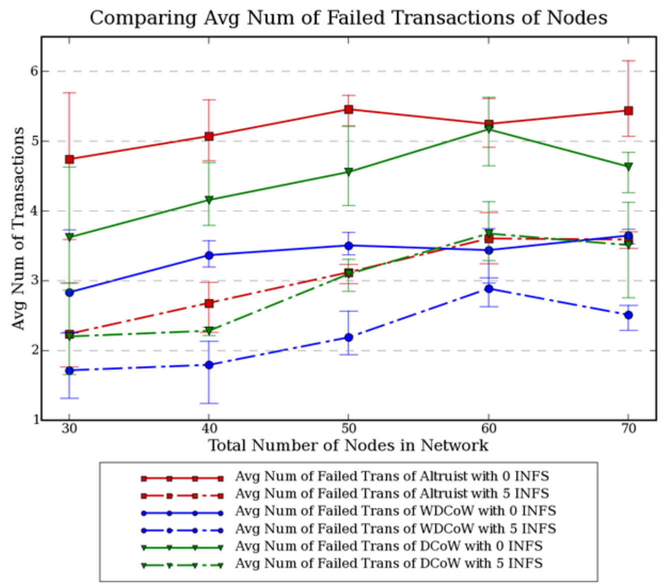


Figure V.8: Comparing the Average Number of Failed Transactions of Nodes Running Different Strategies with 0 or 5 InfoStations

## V.C Summary and Discussion

To conclude, in homogeneous networks the *WDCoW* strategy is the most efficient strategy among the four discussed strategies. It delivers relatively high *Gains* and incurs relatively tolerable *Losses*. It also has the lowest communication overhead and lowest failure rate. The *DCoW* strategy fairs well by delivering modest *Gains* while suffering the lowest *Losses* amongst the three collaborative strategies. It also incurs tolerable communication overhead. On the other hand, the *Altruists* deliver the highest *Gains*, they do so by incurring the largest *Losses* and also incur a very high communication overhead.

## Chapter VI

# COLLABORATIVE EXCHANGES IN EVENLY MIXED HETEROGENEOUS NETWORKS

In Chapter V, we have examined the *Gains*, the *Losses* and the communication overheads in homogeneous networks. In this Chapter, we analyze the performance of each of the strategies presented in Chapter IV in *evenly mixed* heterogeneous networks. An *evenly mixed* network, is a network where each of the four strategies described in Section IV.C have an equal representation of 25% of the network population (i.e. there are no majorities or minorities). Examining this particular population configuration allows us to identify how each of the four network subgroups effect one anther. We start by examining the *evenly mixed* networks of varying sizes after 100 minutes of simulation and later move to analysis of time-based evaluation of *evenly mixed* networks where we look at the progress of 70 nodes during 3 hours of interactions (in increments of 20 minutes).

## VI.A Evenly Mixed Network

Similar to the analysis of performance of each of the strategies in the homogeneous networks (described in Chapter V), we examine the *Gains*, the *Losses* and the communication overheads of each of the subgroups in our *evenly mixed* networks. In particular, we look at the inter-strategy collaborative interactions and identify the benefits of the proposed bartering approach while contrasting bartering with conventional collaboration approaches.

### VI.A.1 Gains and Losses in Evenly Mixed Networks

FigureVI.1 provides a unique view of the benefits of the *DCoW* and the *WDCoW* bartering strategies over the conventional *Altruistic* and *FreeRiding* strategies. FigureVI.1 shows the *Gains* (represented by the solid lines) and the *Losses* (represented by the dashed lines) of an *evenly mixed* heterogeneous network. First, let's look at the *Gains*. The highest *Gains* are achieved by the two bartering strategies. In fact, the *Gains* of the *DCoWs* and the *WDCoWs* are very similar. This is due to the fact that these two strategies have very similar exchange policies. In contrast, the *Gains* of *FreeRiders* are significantly lower than the *Gains* of the other collaborative strategies. As mentioned in Section IV.C, *FreeRiders* are able to acquire goods only from *Altruists*. Since the other collaborative strategies are not as philanthropic as the *Altruists* are, the *FreeRiders* have limited number of opportunities to acquire the needed goods. Finally, the *Gains* of the *Altruist* are not as outstanding as they were in homogeneous networks. In fact, *Altruists' Gains* are the lowest amongst the three cooperative strategies. This disappointing performance is related to unbeneficial interactions with other nodes in the network in particular with *FreeRiders*.

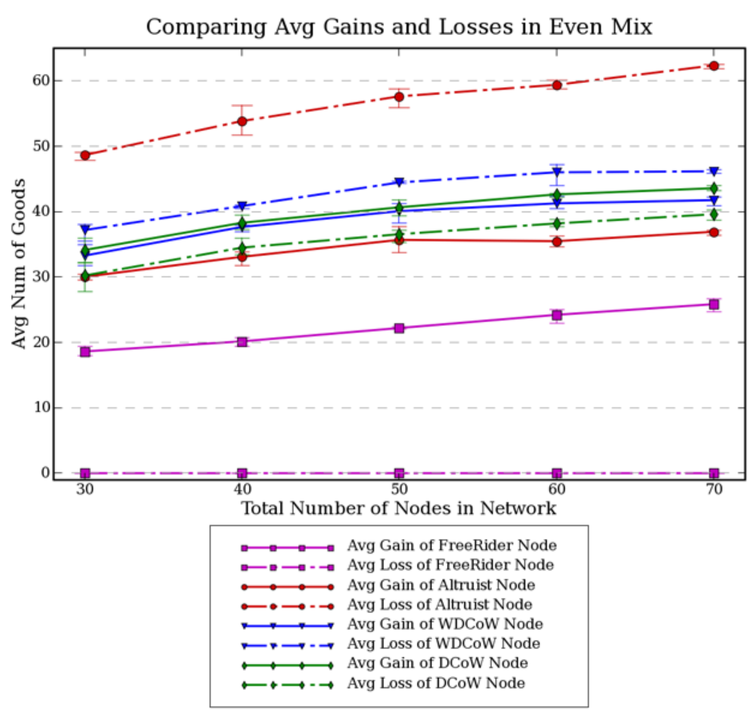


Figure VI.1: Comparing the Average Gains and Losses of Nodes Running Different Strategies in Evenly Mixed Networks

When examining the *Losses*, the *Altruists* and the *FreeRiders* stand out since they represent the two extremes. Firstly, an average *Altruist* pays a hefty price for executing a highly cooperative strategy. The *Altruists* experience the largest *Losses* due to the fact that they indiscriminately participate in every possible exchange regardless of the benefits and losses of these interactions. These *Losses* weaken *Altruists* and contribute to their lower *Gains*. In essence, the other three more cautious strategies exploit *Altruists'* liberal exchange policies and weaken Altruists ability to collaborate. Furthermore, *Altruists' Losses* are significantly higher than their *Gains*. This discrepancy makes this philanthropic strategy even less appealing. In contrast to the *Altruists*, the *FreeRiders* get away without any *Losses*. However, since *FreeRiders' Gains* are also very low, this strategy is not particularly appealing as well. Now, let's look at the *Losses* of the nodes running the *WDCoW* and the *DCoW* strategies. Both of the bartering strategies experience lower *Losses* than the conventional *Altruist* strategy since the bartering nodes are more cautious during every exchange. When comparing the *WDCoWs* and the *DCoWs*, the more liberal *WDCoWs* have higher *Losses* than the conservative *DCoWs*. Similarly, when contrasting their *Losses* with their *Gains*, the *WDCoWs' Losses* are slightly higher than their *Gains*, while the *DCoWs' Losses* are lower than their *Gains*. This is inline with our expectations, since the *WDCoWs* potentially can participate in exchanges where they give away more goods than they receive. On the other hand, the *DCoWs'* policies insure that these nodes are never on the losing end of the deal. From the graph we can determine that, both the *WDCoW* and the *DCoW* can be considered to be effective strategies that provide reasonable *Gains* while incurring tolerable *Losses*. In essence, in regards to *Gains* and *Losses*, in evenly mixed heterogonous networks, the two bartering strategies exhibit better performance than the non-bartering strategies.

### VI.A.2 Communication Overheads in Evenly Mixed Networks

As mentioned earlier, the communication overhead is an important component of collaborations in mobile peer-to-peer environments. To better understand the interaction between the nodes in the heterogeneous network, we take a closer look at the transactions in the *evenly mixed* network. FigureVI.2 and FigureVI.3 show how nodes in each of the population subgroups interact with their peer that are executing different strategies.

As described in Section IV.B.2, the outcome of the collaborative exchanges is very dependent on the collaboration policies of the two nodes involved in the interaction. The content of the *Proposal* is very dependent on the *Proposal Composition Policy* of the proposing node. Similarly, the response to the *Proposal* is dependent on the *Proposal Evaluation Policy* of the *Inquiring* node.

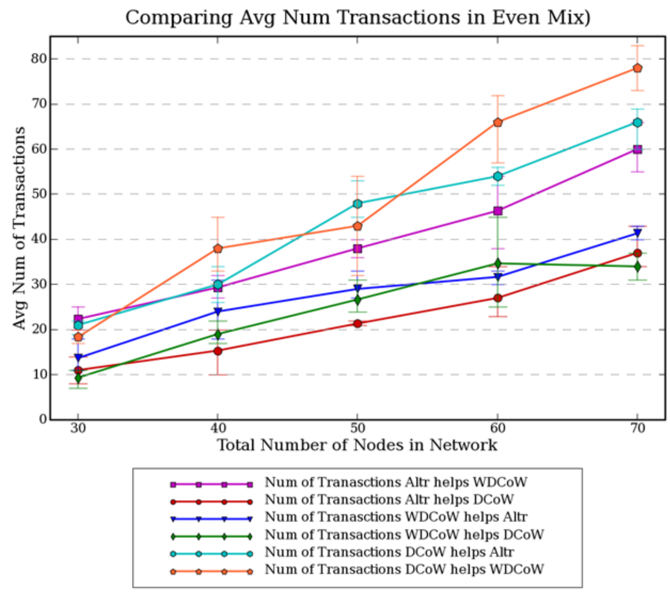


Figure VI.2: Comparing Average Number of Successful Transactions of Nodes Running Different Strategies in Evenly Mixed Networks

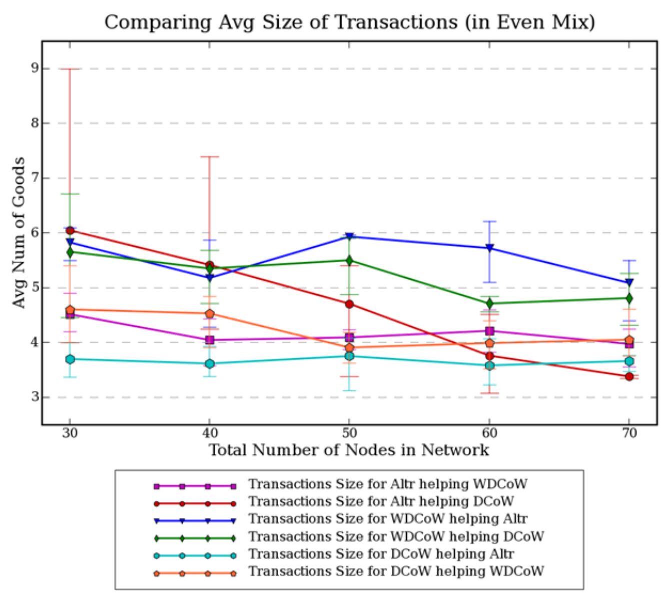


Figure VI.3: Comparing Average Size of Successful Transactions of Nodes Running Different Strategies in Evenly Mixed Networks

We start by examining interactions between the *DCoWs* and the *WDCoWs* depicted in FigureVI.2 and FigureVI.3. If a *WDCoW* is responding to a *DCoWs' Inquiry*, then the *Proposal* is created by the *WDCoW*. As discussed previously in Section V.B, *WDCoWs* tend to generate *Proposals* that are relatively large in size. However, the *DCoWs' Proposal Evaluation Policy* is very prone to *Rejections*. Considering these two aspects, the transactions between these two nodes, initiated by the *WDCoW*, are not common but when they do occur, they have high transaction size. On the other hand, transactions where a *DCoW* responds to a *WDCoW Inquiry* are relatively common, since *DCoW's Proposal Composition Policy* generates *Proposals* that are always accepted by the *DCoWs* (and also by the *Altruists*). However, referring to our previous discussions in Section V.B, the average size of these exchanges is not particularly large since the *DCoWs* insist on "even" exchange.

Now, let's consider transactions between the *Altruists* and the *WDCoWs*. The *Altruists* help the *WDCoWs* more often than the *WDCoWs* help the *Altruists*. This mismatch is once again related to composition and evaluation policies of the nodes. In addition to mismatch, there is a transaction size difference between these transactions. The transactions where the *WDCoWs* make *Proposals* have a larger average transaction size compared to transactions where the *Altruists* make *Proposals*. This behavior is consistent with the strategies' behaviors exhibited in homogeneous networks.

Finally, let's examine transactions between the *Altruists* and the *DCoWs*. Since *Proposal Evaluation Policy* of the nodes executing the *DCoW* strategy is very strict, there are very few transactions where *Altruists* initiate the transaction. As we will discuss later, there are a lot of *Rejections* that take place during these interactions. Also, the average size of these transactions is relatively low. On the other hand, the number of transactions where the *DCoWs* initiate transactions with the *Altruists* is high. This is due to the fact that any *Proposal* that could be generated by the *DCoWs* is acceptable to the *Altruists*. As expected, the average size of these transactions is also low since, the *DCoWs* always insist on an "even" exchange.

In addition to examining the communication cost of nodes collaborating with other strategies, we also consider costs of collaborations within their own subgroup. FigureVI.4 and FigureVI.5 show how nodes behave when they collaborate with their own "kind". As expected, the *Altruists* continue to have large number of small-volume transactions while, the *WDCoWs* continue to participate in few large volume transactions. The *DCoWs* also fare well, since their transaction costs are lower than the *Altruists* but greater than the *WDCoWs*. These results are consistent with results from interactions in the homogeneous network.

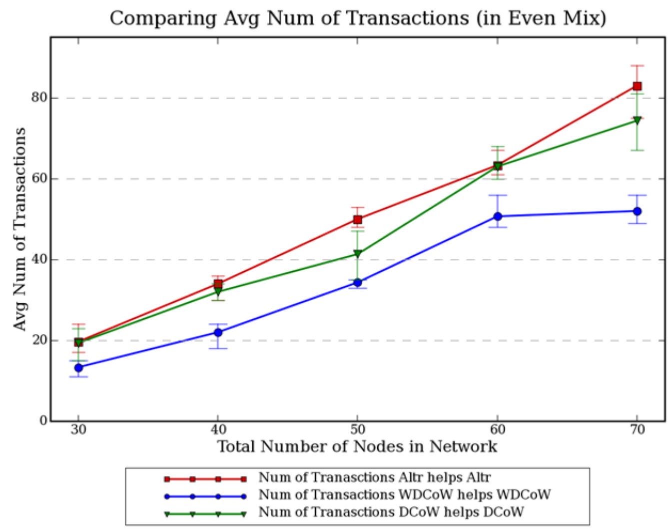


Figure VI.4: Comparing Average Number of Successful Transactions of Nodes Running Different Strategies in Evenly Mixed Networks

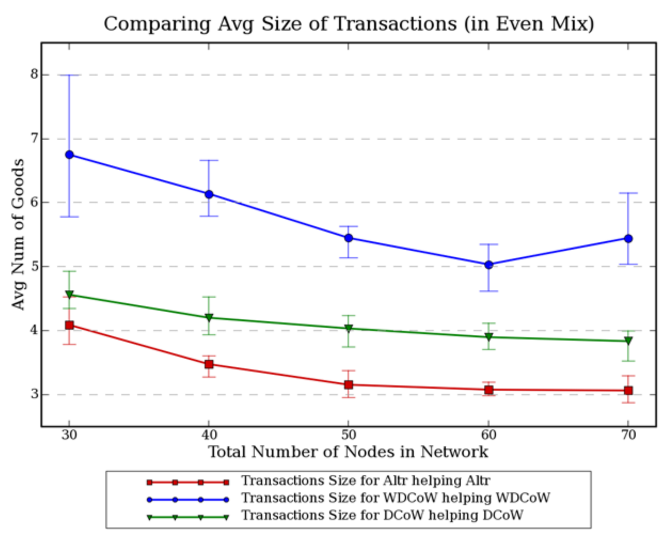


Figure VI.5: Comparing Average Size of Successful Transactions of Nodes Running Different Strategies in Evenly Mixed Networks

### VI.A.3 Rejections and Transactions Failures in Evenly Mixed Networks

Unlike interactions in homogeneous networks, in heterogeneous networks interactions can end with *Rejections*. Out of the four strategies there are only three that receive *Rejections* to their *Proposals*: the *FreeRider*, the *Altruist* and the *WDCoW*. Also, the two bartering strategies are the only two strategies that generate the *Rejection* messages. Let's look at the details of sending and receiving *Rejection* messages. The *FreeRiders* get *Rejections* only from the *WDCoWs* and the *DCoWs*. They never get *Rejections* from the *Altruists*. The *Altruists* can receive *Rejections* from the *WDCoWs*. This occurs when an *Altruist* makes a "freerider-type" *Proposal*. The *Altruists* can also receive *Rejections* from the *DCoWs*. This occurs when an *Altruist* requests to receive more goods than it is planning to give. Similarly, the *DCoWs* reject the *WDCoWs'* "non-even" *Proposals* that would benefit the *WDCoWs* more than the *DCoWs*. The formal definition of the policies driving the rejection generating process can be found in Table IV.1.

Though there are only two strategies that reply with *Rejection*, after 100 min of simulation, the levels of *Rejections* in an *evenly mixed* network is significantly higher than the levels of successful transactions. FigureVI.6 shows the average number of *Rejections* that an average node, experiences in an *evenly mixed* network. Since generating a *Proposal* is a relatively heavy computational process, for a node to receive a *Rejection* in response to its *Proposal* is a very taxing computational overhead. If a node receives a *Rejection* then the work that was done to create this *Proposal* becomes an unnecessarily wasted effort.

Figure VI.6 shows that, when comparing the rates of *Rejections* in an *evenly mixed* heterogeneous network, the *DCoWs* on average send the most number of *Rejections*. This is not surprising, since there is a significant portion of the population that generates *Proposals* that do not comply with *DCoW Proposal Evaluation Policy*. Similarly, the *FreeRiders* receive the highest number of *Rejections*. This is due to the fact that the *Proposal Composition Policy* of the *FreeRider* generates *Proposals* that are likely to be rejected by a half of the population of the *evenly mixed* networks. The *Altruists* receive more *Rejections* than the *WDCoWs* due to the fact that they can generate "freerider-type" *Proposals* or *Proposals* that benefit the *Altruist* more than the other node. A *WDCoW* never generates "freerider-type" *Proposals*; however, it generates "non-even" *Proposal* and thus it can receive *Rejections* from the *DCoWs*. Comparing the rejection rates of the *Altruists* and the *WDCoWs*, the *WDCoWs* have the smaller overhead associated with receiving *Rejections*. Also, the rejection rate is much higher in larger size networks than in the smaller networks. This increase is related to the fact that as the density of the network increases, larger set of free riding nodes have more opportunities to attempt to interact with bartering nodes.

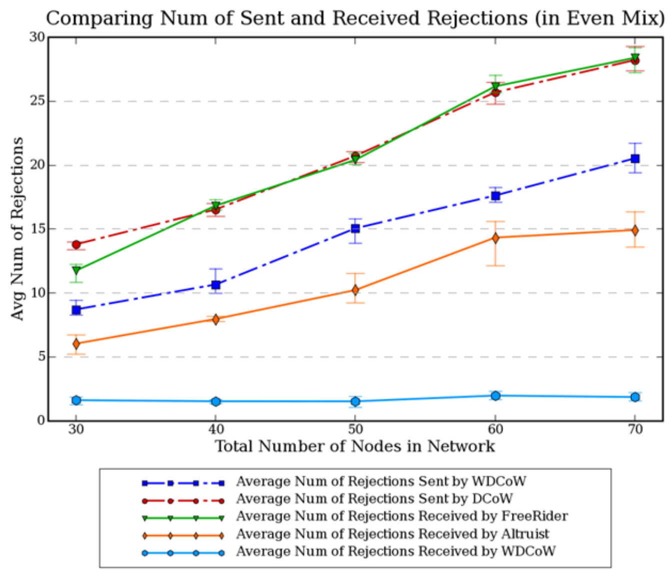


Figure VI.6: Comparing Number of Sent and Received Rejections of Nodes Running Different Strategies in Evenly Mixed Networks

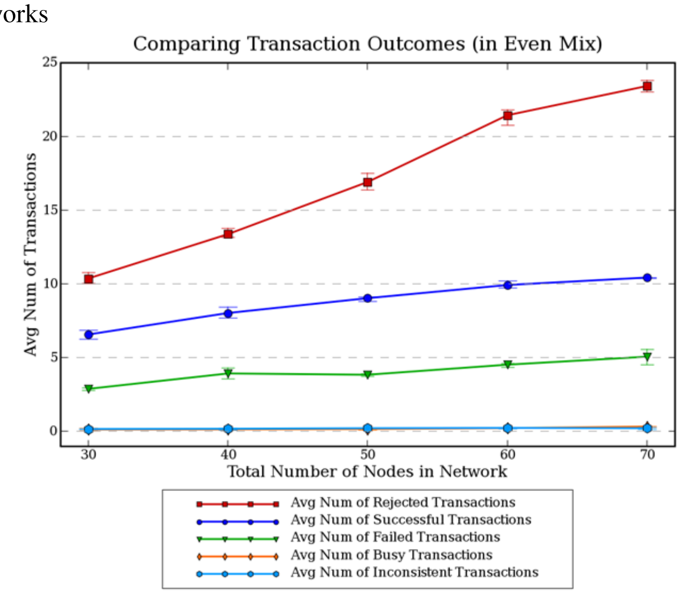


Figure VI.7: Comparing Transaction Outcomes for Nodes Running Different Strategies in Evenly Mixed Networks

Similar to homogeneous networks, in *evenly mixed* heterogeneous networks, a node can have failed transactions. These failures are attributed to nodes moving out of range and to unreliability of wireless communication medium. FigureVI.7 shows that though the failed transactions (represented by the green line) are common, they are less frequent than the *Rejections*. Finally, as previously discussed in Section IV.B.2, a transaction can also have a *Busy* outcome and an *Inconsistent* outcome. Both of these outcomes are relatively rare and have small impact on the communication costs. Thus the *Rejections* are the primary source of interaction inefficacies and failed transaction are the secondary source. Also unlike rejection rates, the failed transaction rates are not as affected by the size of the network. This implies that the primary cause behind the failed transactions is the mobility of the nodes and not the noise and congestion of the wireless medium.

Brining together all of the above described analysis, we conclude that the bartering strategies are more efficient during their operations in evenly mixed network. Both, the *DCoWs* and the *WDCoWs* get good gains while experiencing tolerable *Losses*. Bartering nodes also incur lower communication overheads. Therefore, the bartering approach is more efficient than conventional free riding or altruistic strategies.

## VI.B Time Based Evaluation of Strategies in Evenly Mixed Network

In addition to tracking node interactions in heterogeneous networks of varying sizes, we also track node performance over time. This evaluation allows us to understand *swarm*-like dynamics [113, 20, 52] of the cooperative interactions that occur in these heterogeneous networks. We examine the progress that the nodes in each of the subgroups of an *evenly mixed* 70-node network go through in a span of 3 hours of simulation time in increments of 20 minutes.

### VI.B.1 Counting the Unfulfilled Wishes

The primary goal for each of the nodes in the network is to fulfill as many wishes as possible on the nodes *iWant List*. Measuring the number of unfulfilled wishes gives us the insight into productivity and effectiveness of each of the subgroups of the population. Measurement of unfulfilled wishes also provides us with opportunity to evaluate the dynamics of the subgroup interactions and observe the inter-strategy collaboration dynamics. Figure VI.8 shows the number of nodes in each of the subgroups of the *evenly mixed* network that have unfulfilled wishes in their *iWant Lists*. In particular, this Figure shows how close the nodes get to obtaining all of the goods on their *iWant Lists*. We have subdivided nodes into threshold-like categories according to the size of their *iWant List* at the particular time (marked on the $x$ axis on the graph). The nodes with empty *iWant List* are in the first subcategory (represented by red bars), followed by the second subcategory which contains nodes with 1 to 5 wishes in their *iWant List* (represented by orange bars), and so on. The last broad subcategory contains a set of nodes with more than 25 unfulfilled wishes on their *iWant List* (i.e. nodes have less than half of their wishes fulfilled).

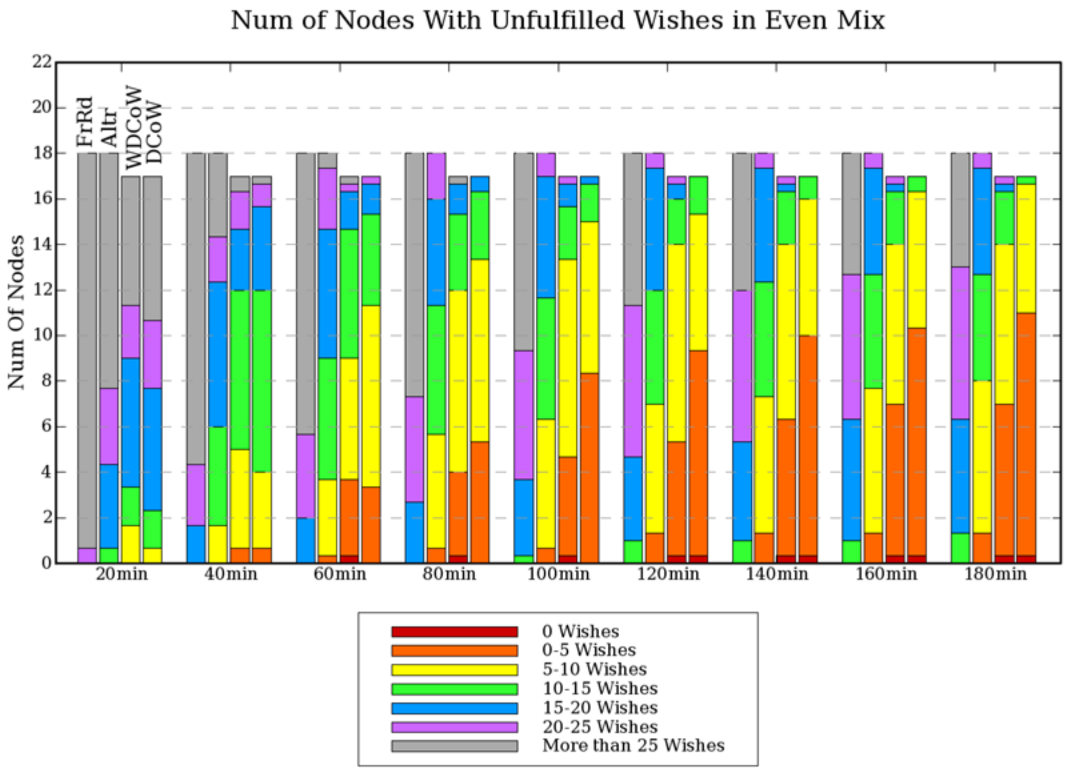


Figure VI.8: Number of Nodes With Unfulfilled Wishes in Evenly Mixed Networks

As expected, the *FreeRiders* which are represented by the left most column in each of the column sets are the least productive nodes. Note that, after the first 20 min of simulation, a significant majority of the *FreeRiders* belong to the broad subcategory of nodes with more than 25 wishes in their *iWant List*. As the time progresses, the *FreeRiders* make progress in fulfilling their wishes. After 180 minutes of simulation time, the majority of the *FreeRiders* have less than 20 unfulfilled wishes on their *iWant List*. However, despite all of this progress, these nodes still deserve to be categorized as the most unproductive part of the network population. Let's look at the progress of the *Altruists*. Interestingly, the fulfillment rate of the *Altruists* is disappointing. These nodes make very limited progress during the simulation. This is due to the fact that, the *Altruists* are overwhelmed by the *FreeRiders*. In fact, progress made by *FreeRiders* is solely attributed to the generous philanthropic collaboration policies of the *Altruists*. In essence, the progress of the *Altruists* stalls during the second part of the time-based simulation. Both of these conventional exchange strategies are clearly underperforming compared to the proposed bartering approach.

Let's investigate the effectiveness of the bartering strategies. When examining the *DCoWs'* performance it is clear that these nodes are executing the most effective collaboration strategy. The *DCoWs* make a steady progress and in 180 min majority of the nodes have less than five wishes on their *iWant List*. These accomplishments can be attributed to the cautious trading policies of the *DCoWs*. In essence, these nodes insist on an "*even*" or a self-benefiting exchange which protects them from inefficient interactions which are common amongst the nodes that are executing one of the other three strategies. Similarly, the *WDCoW* nodes, which execute the other bartering strategy, also make good progress. However, this progress is not as distinguished as the progress of *DCoWs*. This promising performance of *WDCoWs* is also related to the exchange policies of the nodes. The *WDCoWs'* policies provide some level of protection from uneven exchanges that the *Altruists* and the *FreeRiders* propose. The "long term" effectiveness of the bartering strategies can be explained by the nodes' ability to protect their interests from the dynamics of the interactions between the nodes that are executing conventional altruistic and free riding strategies.

### VI.B.2 Counting Goods in iHave Lists

To better understand the "long term" effectiveness of the bartering strategies, we further examine the details of our time-based experiment by analyzing the growth and declines of the nodes' *iHave Lists*. The size of an *iHave List* impacts node's ability to find collaboration partners. Thus, tracking the changes of the average list size gives us the insight into the nodes' collaboration readiness. FigureVI.9 tracks the progress of the average size of the *iHave Lists* for each of the subgroups of the mobile network population. Every node in the network starts out with 50 goods in their *iHave List* (represented by the dashed black line in the graph). The first 20 minutes of the simulation show the initial inclinations of each of the subcategories in the network population. The *Altruists* experience a major decline in the number of goods in their *iHave Lists*. Their philanthropic policies predispose these nodes to giving away their digital content. The declines continue as the simulation progresses. At the very end of the simulation, the *Altruists* have the lowest number of goods in their *iHave List*. These declines incapacitate *Altruists* since they have less and less to offer to the other trading nodes in the network. In contrast, the *FreeRiders* are getting "wealthier" with time. Since they never give away any of their goods, they collectively act as a "black hole" for digital goods circulating in the network. In essence, the *FreeRiders* only accumulate goods and never release them back into the network. Since *Altruists* are the only subgroup of the population that interacts with the *FreeRiders*, the *Altruists* suffer the greatest declines in the average size of their *iHave Lists*. Both, the *FreeRiders* and the *Altruists* undergo major changes in the average size of their *iHave Lists* in the span of the simulation. In contrast, the average size of the *iHave Lists* of the bartering strategies remain relatively stable through out the time of the simulation. In essence, both, the *WDCoWs* and the *DCoW* are able to maintain the rates of their *iHave Lists*. The *DCoWs* get a minor increase that comes from interactions with the *Altruists* and the *WDCoWs*. On the other hand, the *WDCoWs* gets a minor decrease in the average size of their *iHave List*. This decrease is attributed to interactions with the *Altruists* and the *DCoWs*. Over the span of the simulation, the average size of the *DCoWs' iHave List* grows (very slowly) while for the *WDCoWs* average size drops (very slowly). This stability allows nodes to maintain their collaboration appeal to the other nodes in the network.

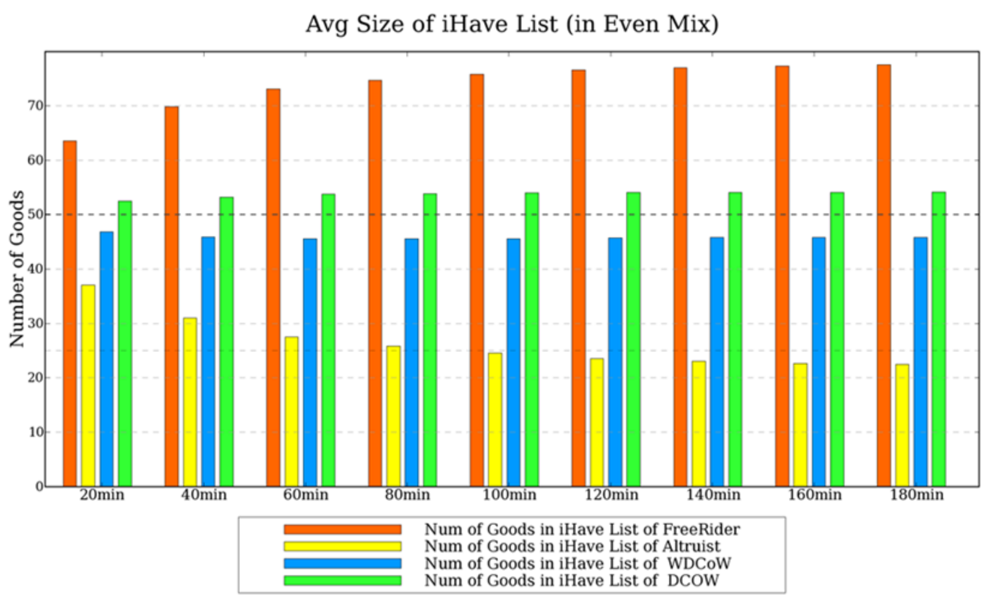


Figure VI.9: Comparing the Average Size of iHave Lists of Nodes Running Different Strategies in Evenly Mixed Networks

In essence, the *swarm*-like dynamics of this micro-economy show that the *FreeRiders* eventually consume significant percentage of the goods that the *Altruists* possesses. This in term transforms the *Altruists* into *FreeRider-like* nodes. As a significant portion of the population starts to exhibit strong *free riding* tendencies, the bartering nodes find fewer bartering opportunities. However, these bartering strategies also protect and isolate the nodes from depleting their *iHave List* and allow them to productively interact with the network population and thus fulfilling a significant portion of the wishes on their *iWant List*.

### VI.B.3 Counting Number of Transactions

To further understand the intricacies of the outcome of our time-based experiment, lets look at the number of transactions that take place in the *evenly mixed* network. Figure VI.10 shows both the heterogeneous transaction counts (the bottom part of the graph) and the homogeneous transaction counts (the top of the graph). Note that, a significant percentage of transactions occur during the first part of the simulation. During this period, the *Altruists* conduct the largest number of their transactions. A second column (from the left) in each of the column groups in the lower part of the graph represents the average count of transaction initiated by the *Altruists*. The majority of these transactions are attributed to the *Altruists* interactions with the *FreeRiders*. The first column of each of the column groups in the graph represent the average number of transactions initiated by the nodes executing the *FreeRider* strategy. All of these transactions are interactions with the *Altruists*. These two transaction counts show the severity of the impact that the *FreeRider* have on the *Altruists*. Clearly, the majority of transactions conducted by the *Altruists* are dedicated to servicing the *FreeRiders*. In fact, throughout this time-based simulation, the *Altruists* interact more with the *FreeRiders* than with any other nodes running collaborative strategies. This volume of interactions is attributed to the fact that *FreeRiders* are unable to engage nodes from any other subset of the network population. Thus, the *FreeRiders* are not as preoccupied and have more opportunities to interact with the *Altruists* which are the only set of population that responds to their *Inquires* and *Proposals*. Note, as the time progresses the *Altruists* continue to show high rates of transactions that service the *FreeRiders*. In fact, towards the end of the simulation, the interaction activity of the *Altruists'* is completely designated to transacting with the *FreeRiders*. This interaction pattern eventually converts *Altruists* into the *FreeRiders* thus splitting network population into the *free riding* nodes and the bartering nodes.

Lets look at transaction counts of the bartering strategies. The *DCoWs* tend to participate in more transactions than the *WDCoWs*. During the first half of the simulation, the *DCoWs* transactions count is split into three relatively even parts: transaction with the *Altruists*, transaction with the *WDCoWs* and transaction with peers of its own subgroup. In essence, these nodes do not differentiate between the subgroups. These non-preferential interaction patterns are related to the primary concern of the DCoWs which is the preservation of the evenness of the transaction. In the second half of this time-based experiment, the *DCoWs* reduce interactions with *Altruists* since the *Altruists* transform into *FreeRider* type nodes. In fact there are very few transactions conducted by the *DCoWs*. This can also be explained by the fact that, these nodes have reached a very high wish fulfillment rate and have very few goods in their *iWant List*.

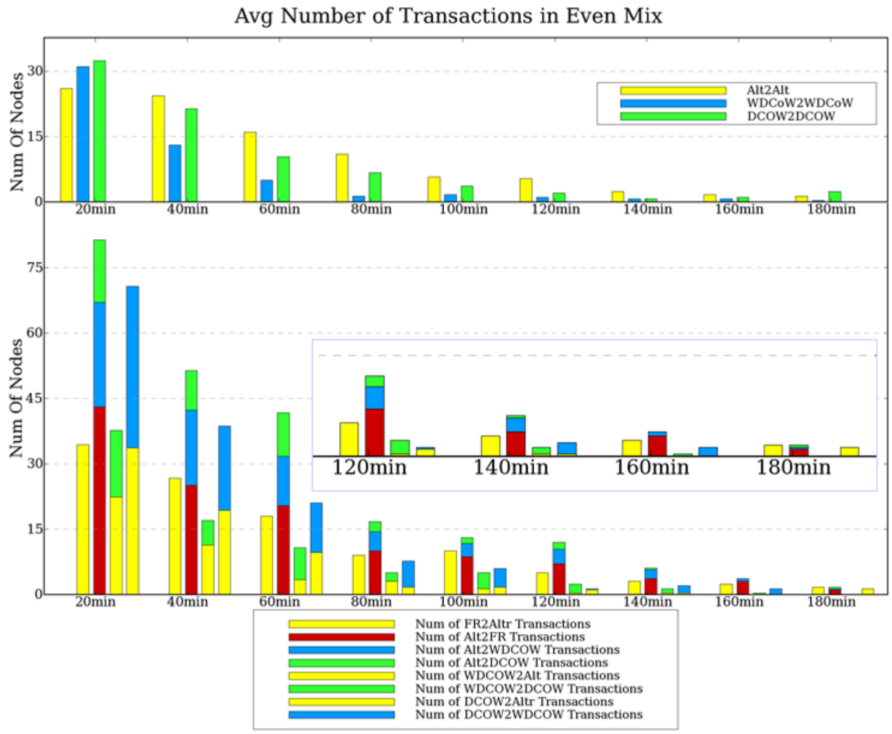


Figure VI.10: Comparing the Average Number of Transactions for Nodes Running Different Strategies in Evenly Mixed Networks

In contrast to *DCoWs*, the *WDCoWs* conduct very few transactions. Through out the simulation, their average transaction count is frequently lower than the average transaction count of the *FreeRiders*. As mentioned in Section V.B and Section VI.A.2, the *WDCoWs* tend to have relatively high average transactions size. In the time-based experiment, this materializes into the fact that the *DCoWs* are able to acquire needed goods in few transactions. Looking at the Figure VI.8 from Section VI.B.1, it is clear that the *DCoWs* wish fulfillment rates are high which basically shows that this strategy also exhibits “long term” efficiency.

## VI.B.4 Counting Number of Rejections

In addition to the successful transactions, nodes also send and receive *Rejections*. FigureVI.11 shows the rates of the *rejected Proposals* in our time-based experiment. The graph shows both aspects of *Rejections*: number of sent *Rejections* and number of received *Rejections*. As discussed in Section IV.C, the *Rejections* are sent only by the *DCoWs* and the *WDCoWs*. *Rejections* can be received by the *FreeRiders*, the *Altruists* and the *WDCoWs*. The left two bars of each of the groups in Figure VI.11 represent the average number of sent *Rejections*. Similarly, the right three bars represent the average numbers of received *Rejections*.

This graph complements the graph in Figure VI.10 since, it shows how bartering strategies are able to protect themselves from unbeneficial transactions by generating *Rejections* to all of the *Proposals* from the *FreeRiders* and to some of the *Proposals* from the *Altruists*. In particular, the *DCoWs* generate the largest number of *Rejections*. They generate most *Rejections* during the early stage of the simulation and through out the simulation, continue to lead in the *Rejection*-sending rates. Not surprisingly, the *FreeRiders* are the main recipients of the *Rejections*. In fact, they lead in rate of receiving *Rejections* throughout the simulation. The *Altruists* are not as bothered by the incoming *Rejections* at the start of the simulation. However, towards the end of the simulation, the *Altruists' Rejection*-receiving rates are almost as high as the *FreeRiders Rejection*-receiving rates. This, again, confirms the previous findings that in this swarm like micro-economy, the *Altruist* eventually become the *FreeRiders* and the bartering strategies are forced to reject collaborations that are initiated by the conventional interaction strategies.

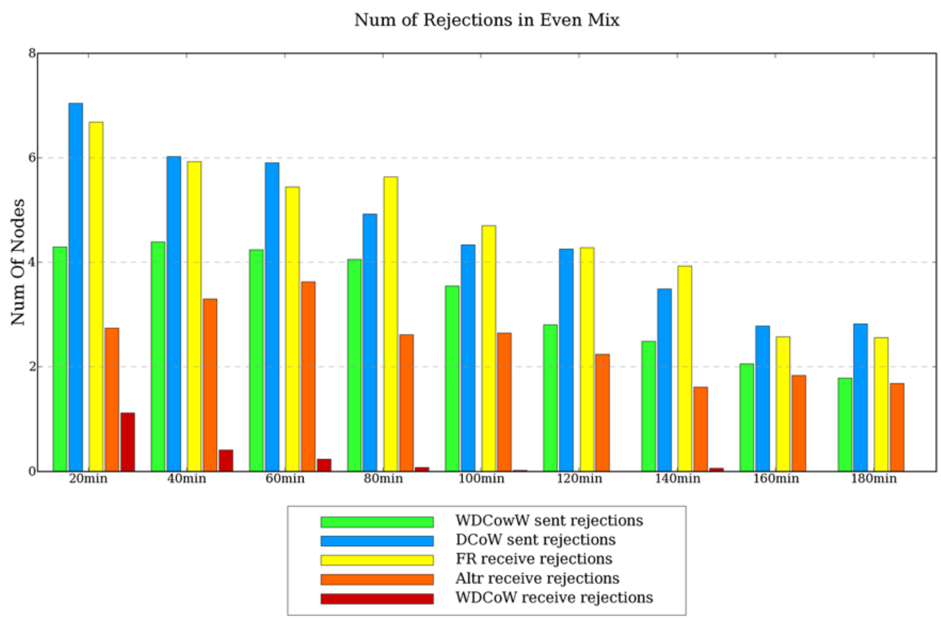


Figure VI.11: Comparing the Average Number of Rejections for Nodes Running Different Strategies in Evenly Mixed Networks

## VI.C Summary and Discussion

In this Chapter, we have examined heterogeneous networks that are evenly subdivided into four subgroups. Mobile nodes in each of the subgroups are executing one of the collaboration strategies described in Chapter IV. This network population configuration is designed to bring out the key aspects of the inter-strategy collaboration patterns that further highlight the key differences between conventional collaboration approaches and the proposed bartering collaboration approach. This Chapter clearly highlights the *Gains* and *Losses* related benefits of the bartering approach in such networks. This Chapter also underlines the importance of the bartering approach in the mobile environments where the opportunistic collaborations are not only burdened by the communication overheads that are directly related to the aspects of the dynamic wireless environment but also are challenged by the unpredictability of encounters with other mobile peers that do not necessarily share the same collaboration strategy. This Chapter also examines the *swarm*-like network dynamics that take place in the heterogeneous networks. In particular, the opportunistic interactions in the mixed environments lead to a split of the network into two distinct groups of *free rider type* nodes and bartering nodes. In essence, the nodes that are executing the conventional altruistic strategies inevitably get overpowered by the other nodes in the network population and start to act as *free riders*.

## Chapter VII

# COLLABORATIVE EXCHANGES IN STRATEGY-DOMINATED HETEROGENEOUS NETWORKS

In Chapter VI, we have examined the *Gains*, the *Losses* and the communication overheads in *evenly mixed* heterogeneous networks. In this Chapter, we analyze the performance of the collaboration strategies (presented in Chapter IV) in heterogeneous networks that are *"dominated"* by one of the strategies. In particular, we consider four heterogeneous networks: a *FreeRider-dominated* network, an *Altruist-dominated* network, a *WDCoW-dominated* network and finally, a *DCoW-dominated* network. Each of these networks has a *dominating* strategy presence such that, the majority strategy represents 40% of the network population while, the other strategies are 20% each. Studying strategy dominated networks gives us a better understanding of inter-strategy communication patterns and influences that strategies exert on one another [79, 29, 69]. In addition, studying these dominated networks gives us an opportunity to examine the more realistic network population compositions that are more likely to occur in the real world situations.

## VII.A Gains and Losses in Strategy-Dominated Heterogeneous Networks

Now that we have looked at performance of each of the strategies in the *evenly mixed* heterogeneous network, let's consider networks *dominated* by one of the strategies. This Section in detail will look at the *Gains* and

*Losses* that each of the four strategies experiences in the dominated networks.

### VII.A.1 Gains and Losses in FreeRider-Dominated Heterogeneous Networks

FigureVII.1 shows *Gains* and *Losses* for the network *dominated* by the *FreeRiders*. As previously discussed in Sections IV.C and VI.A.1, the *Altruists* form the only subset of the population that are willing to interact with the *FreeRiders*. In the *FreeRider-dominated* network, the ratio of *FreeRiding* part of the population to the *Altruistic* part of the population is two to one. This ratio creates an extremely unfavorable environment for both, the *FreeRiders* and the *Altruists*. Essentially, the issue is that there is a small set of suppliers and a large set of consumers. Thus, as expected, both of the subgroups of the population experience very low *Gains*. In fact, out of the three collaborative minorities, the *Altruists* get the lowest *Gains*. In addition to the low *Gains*, the *Altruists* also experience extremely high *Losses*. These *Losses* are directly related to the population ratios. In regards to the *Gains* of the *FreeRiders*, the *FreeRiding* nodes continue to see disappointing levels of *Gains* and, as mentioned before, they do not incur any *Losses*. Basically, the large majority of *FreeRiders* incapacitates both, itself and the *Altruistic* minority that supports the *FreeRiders*.

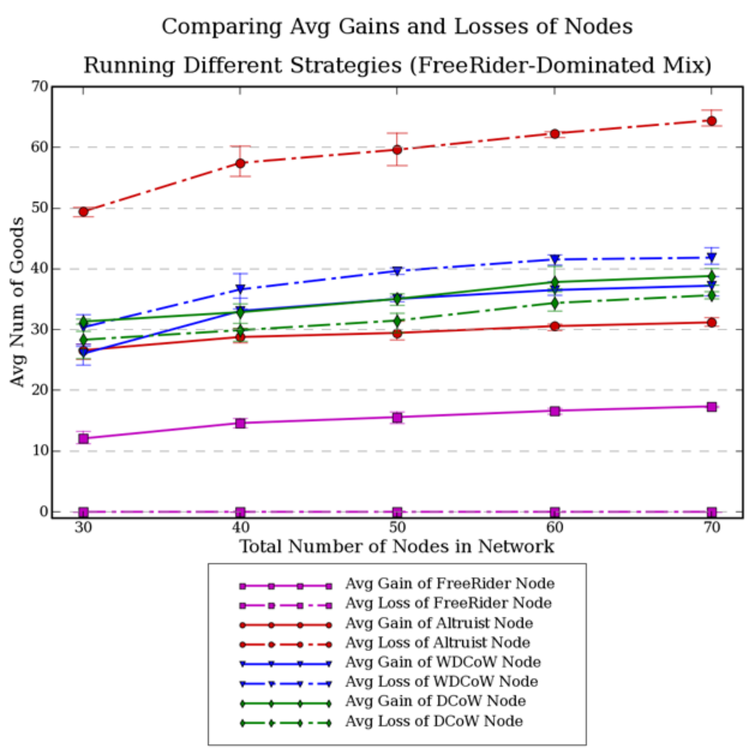


Figure VII.1: Comparing the Average Gains and Losses of Nodes Running Different Strategies in FreeRider-Dominated Heterogeneous Networks

Interestingly, in the *FreeRider-dominated* network, the bartering *DCoW* minority is able to protect itself from the *Free Riding* majority. The *DCoWs* survive with the highest *Gains* in this network configuration and, as before, have tolerable *Losses*. Similarly, the *WDCoWs* fair well since, they show promising *Gains* and also suffer tolerable *Losses*. Thus, in the mobile networks dominated by the *Free Riding* majority, the bartering approach proves to be relatively resilient. To conclude, due to the large presence of uncooperative nodes in the population, the *FreeRider-dominated* networks deliver limited levels of productive collaborations.

### VII.A.2 Gains and Losses in Altruist-Dominated Heterogeneous Networks

FigureVII.2 shows *Gains* and *Losses* for the network dominated by the *Altruists*. Since *Altruists* are selfless and ultra cooperative nodes, all subgroups of the population in the network experience high levels of *Gain*. In fact, all the three collaborative strategies exhibit similar levels of *Gains*. The *DCoW* minority shows the highest *Gains* and the lowest *Losses*. Similarly, the performance of the*WDCoW* minority is consistent with it's previous performances. The *Altruistic* majority has similar or slightly lower *Gains* than the bartering nodes; however, these *Gains* are offset by the highest *Losses* in the network.

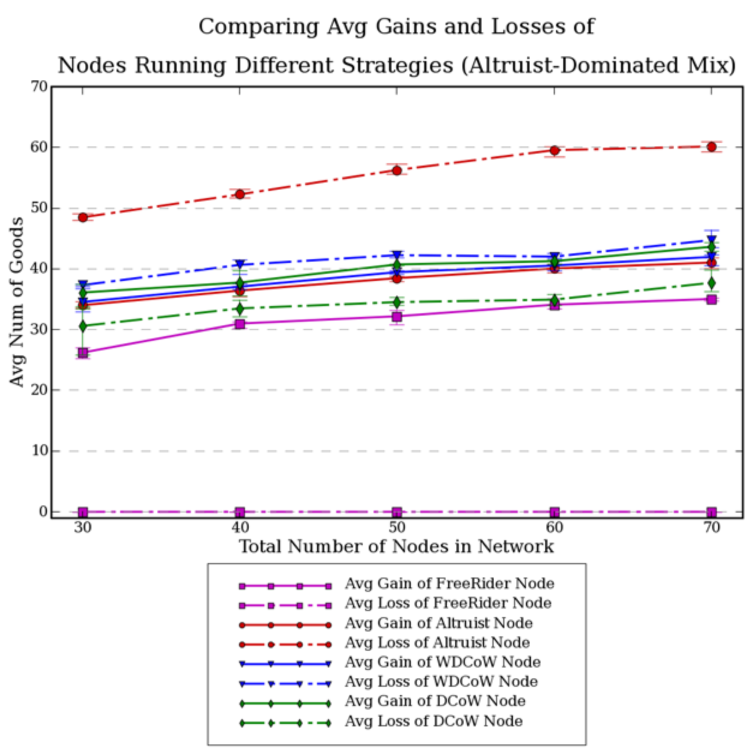


Figure VII.2: Comparing the Average Gains and Losses of Nodes Running Different Strategies in Altruist-Dominated Heterogeneous Networks

Interestingly, the *FreeRider* minority distinguishes itself by exhibiting uncharacteristically high *Gains*. This improvement can be attributed to the fact that, this particular network configuration is composed of a limited number of *FreeRiders* and a relatively large population of highly accommodating *Altruists*. This ratio is very beneficial to the *FreeRiding* minority since, these nodes get plenty of opportunities to collect goods from their only set of collaborative partners. In addition, the limited number of *FreeRiders* means that, on average, these nodes encounter less competition from one another. Furthermore, the bartering minorities present limited opposition in regards to interactions with *Altruists*. In fact, nodes form the *WDCoW* minority occasionally enrich the *Altruists* majority by participating in exchanges that benefit the *Altruist* nodes more than the *WDCoW* nodes.

### VII.A.3 Gains and Losses in WDCoW-Dominated and DCoW-Dominated Heterogeneous Networks

FigureVII.3 shows *Gains* and *Losses* for the network dominated by the *WDCoW* while, FigureVII.4 shows *Gains* and *Losses* for the *DCoW-dominated* network. In both of these networks, there is a strong resemblance in the levels of *Gains* for both bartering strategies. Since, the *DCoWs* are less forgiving, the *Gain* levels for all of the cooperative strategies are slightly lower than in the *WDCoW-dominated* network. This is consistent with previously described findings and also with our expectations.

Also, levels of *Gain* for the *Altruists* are relatively low and unlike the previous two networks the *Gains* in these networks are similar to the *Gains* in the *evenly mixed* network. In the *DCoW-dominated* network, all of the collaborative strategies see a slight decrease in the levels of *Losses*. This relative decline is attributed to the bartering policies of the *DCoWs* which focus on the evenness of the exchange thus frequently rejecting *Proposals* that are composed by the *Altruists*. Essentially, when the majority of nodes in the network employ the bartering communications approach for their interactions, the levels of collaborations in this network are higher than in the networks dominated by nodes executing conventional collaborative interaction methods.

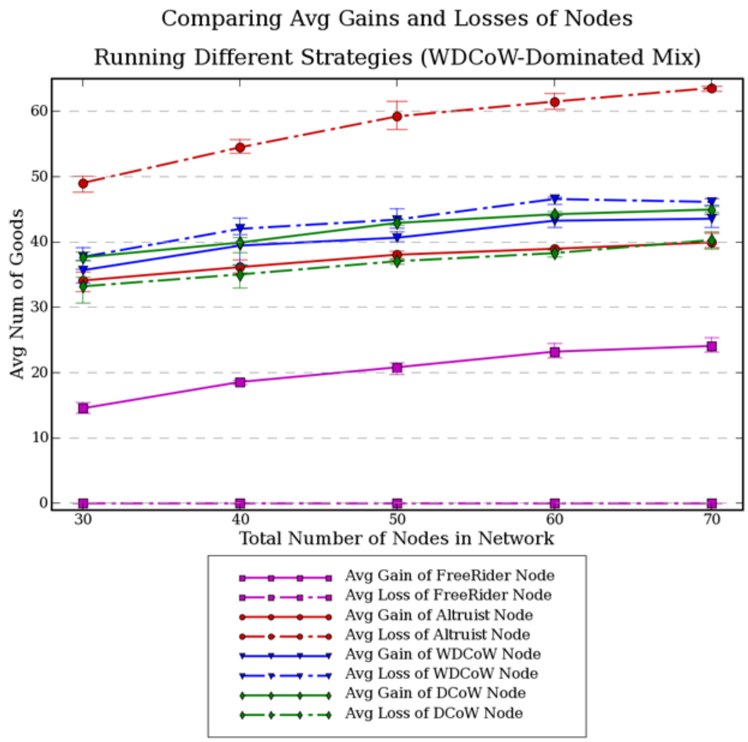


Figure VII.3: Comparing the Average Gains and Losses of Nodes Running Different Strategies in the WDCoW-dominated Heterogeneous Networks

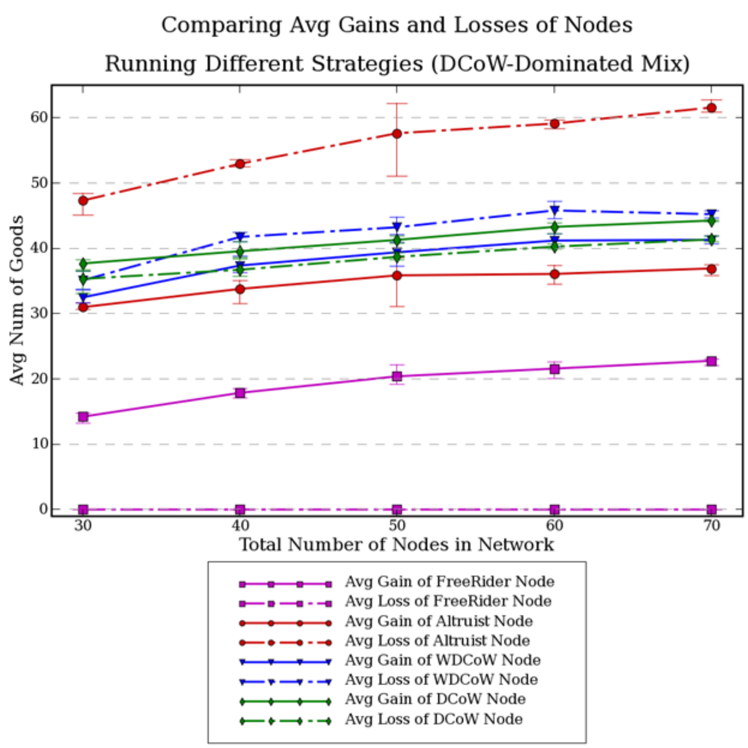


Figure VII.4: Comparing the Average Gains and Losses of Nodes Running Different Strategies in the DCoW-dominated Heterogeneous Networks

## VII.B Comparing Strategy Performance in Heterogeneous Networks

Each of the strategies described in Chapter IV.C has different inter-strategy collaborative advantages and disadvantages. In Chapter VI and Section VII.A, we have described how network population can influence the strategy performances. In this Section, we look at the performance of each of the strategies and attempt to identify the network population compositions where the nodes that are executing this strategy experience the best and the worst performance. In essence, this Section will attempt to answer the following question: *"For a node executing a particular strategy what should the network population composition be so that this node can get the highest* Gains *and incur the lowest Losses?"*

### VII.B.1 Performance of FreeRiders in Heterogeneous Networks

Lets identify the network population composition that deliver the best and the worst performance to the nodes executing the *FreeRiders* strategy. FigureVII.5 shows the *Gains* that the *FreeRiders* experience in heterogeneous networks. As mentioned earlier, the *FreeRiders* can acquire digital goods only from the *Altruistic* subset of the population. So, as expected, *FreeRiders* experience the highest levels of *Gains* in the *Altruist-dominated* network where, the ratio of *Altruist* nodes to *Free Riding* nodes is the highest. The *FreeRiders* also experience relatively high *Gains* in the *evenly mixed* network. This again relates to the ratio of the *Altruists* and the *FreeRiders* in the network. However, in the network dominated by the bartering nodes, the levels of *Gain* of the *FreeRiders* are significantly lower than in the network with *Altruistic* majority or in *evenly mixed* networks. The collaborative exchange policies of the bartering majorities in these networks are very intolerant of *free riding* behavior. Thus, the *FreeRiders* have limited number of opportunities to interact with the network population. Similarly, in the network dominated by *FreeRiders*, the *Gains* of *FreeRiders* are very low, particularly in the larger network. These low rates are attributed to the fact that a large number of *FreeRiders* (40% of the network population) is competing against each other for the opportunity to collaborate with a small number of *Altruists* (20% of the network population). So to summarize, the *FreeRiders* exhibit the best performance in the *Altruists-dominated* networks. The worst performance is exhibited in the *FreeRider-dominated* networks.

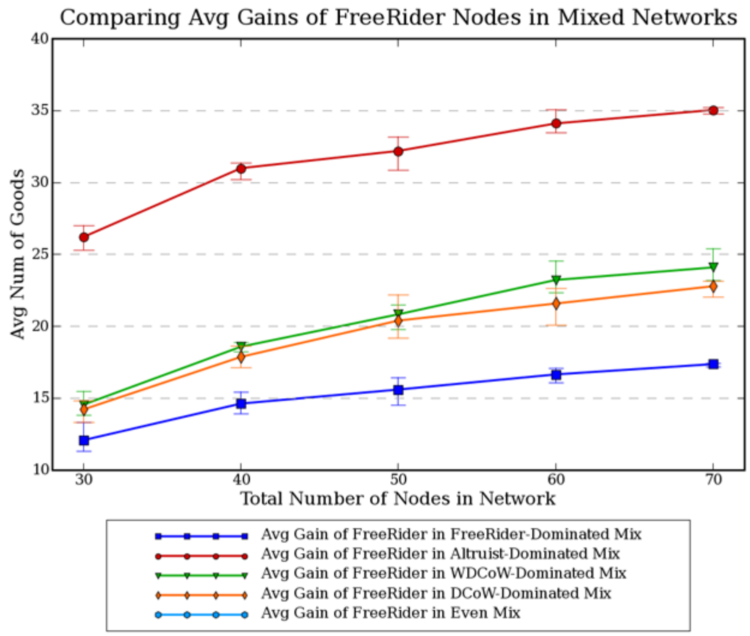


Figure VII.5: Comparing the Average Gains of FreeRiders in Heterogeneous Networks

### VII.B.2 Performance of Altruists in Heterogeneous Networks

Lets identify the heterogeneous networks where *Altruists* experience the best and the worst levels of productivity. FigureVII.6 shows the *Gains*(represented by the solid lines) and *Losses* (represented by the dashed lines) that *Altruists* experience in five heterogeneous networks. The collaboration policies of the Altruists make them the most philanthropic nodes in the environment. Thus, not surprisingly, the *Altruists* receive the lowest *Gains* and the highest *Losses* in a *FreeRider-dominated* network where the *Free Riding* majority takes full advantage of the liberal collaboration policies of the *Altruists*. The *Altruists'* productivity is at its peak when the *Altruistic* nodes are surrounded by their own kind. In the *Altruist-dominated* network, the average *Gains* are at their highest and the average *Losses* are at their lowest. Clearly, the philanthropists do well when they are surrounded by other philanthropists. In the *DCoW-dominated* networks, where there is the greatest concern with "fairness of exchange", the *Altruists* receive unexceptional *Gains*. However, the average levels of *Loss* are relatively low. Thus, the exchange policies employed by the *DCoW* majority have a clear positive effect on the productivity levels of collaborative *Altruists*. Interestingly, the *DCoW* strategy, which is the other bartering strategy, has a less prominent impact on the *Losses* of the *Altruists*. However, the *DCoW-dominated* network population delivers the second best level of *Gains* amongst the five heterogeneous networks. So to summarize, the *Altruists* exhibit the best performance in the *Altruists-dominated* networks where they are surrounded by their philanthropic peers. Interestingly, the *Altruists* also have very high levels of productive interactions in the *WDCoW-dominated* networks. The worst performance is exhibited in the *FreeRider-dominated* networks.

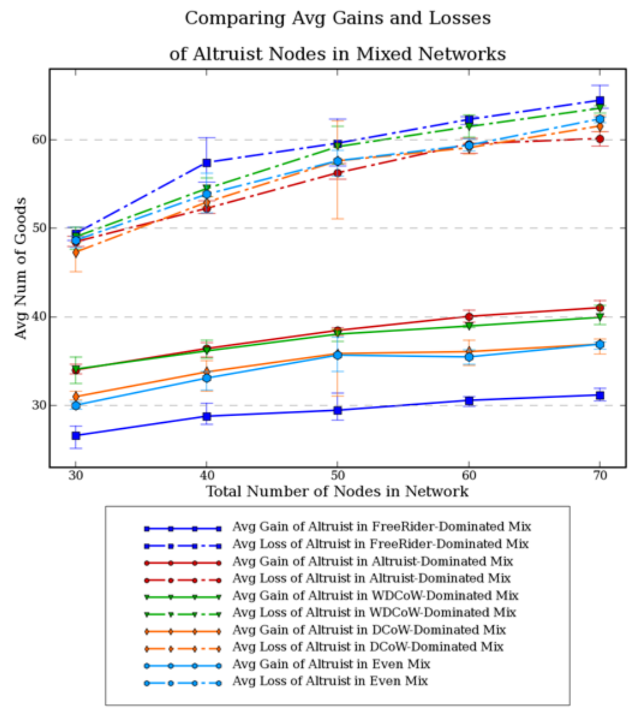


Figure VII.6: Comparing the Average Gains and Losses of Altruists in Heterogeneous Networks

### VII.B.3 Performance of WDCoWs in Heterogeneous Networks

Now lets identify the network populations that are best and worst suited for the nodes executing the *WDCoW* strategy. FigureVII.7 and VII.7 show the *Gains* and the *Losses* of the *WDCoWs* operating in five heterogeneous networks.

Similar to *Altruists*, the *FreeRider-dominated* network provides the *WDCoWs* with the worst conditions for of opportunistic bartering and collaborative peer to peer exchanges. The average *Gains* and the average *Losses* of the *WDCoWs* are the lowest in the *FreeRider-dominated* network. The *WDCoW* nodes are very unproductive in this network configuration.

In regards to the *Altruist-dominated* network, the *WDCoWs* exhibit unremarkable *Gains* but they also show low *Losses*. Coupling these *Losses* and these *Gains*, it is clear that the *WDCoWs* take advantage of the philanthropic nature of *Altruists*. Interestingly, the *WDCoWs* exhibit the highest *Gains* in the network dominated by their own kind. However, the *WDCoWs-dominated* network also delivers very high *Losses*

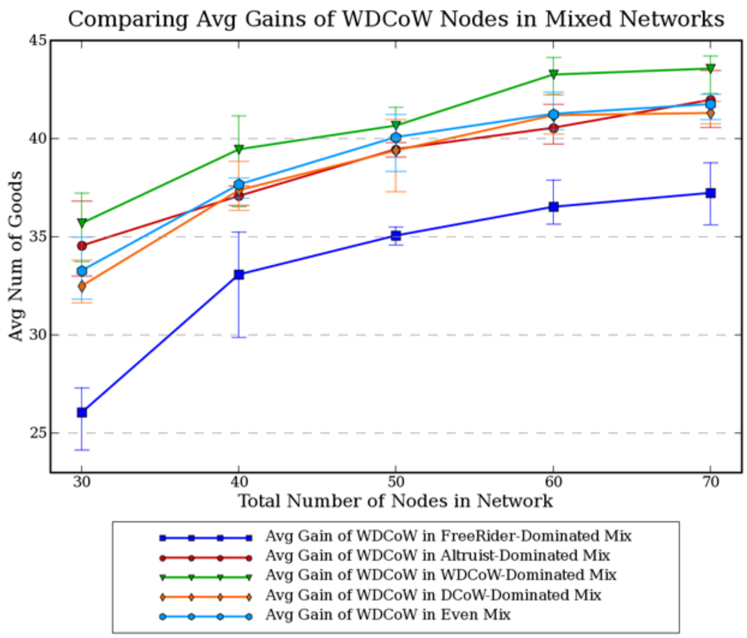


Figure VII.7: Comparing the Average Gains of WDCoW in Heterogeneous Networks

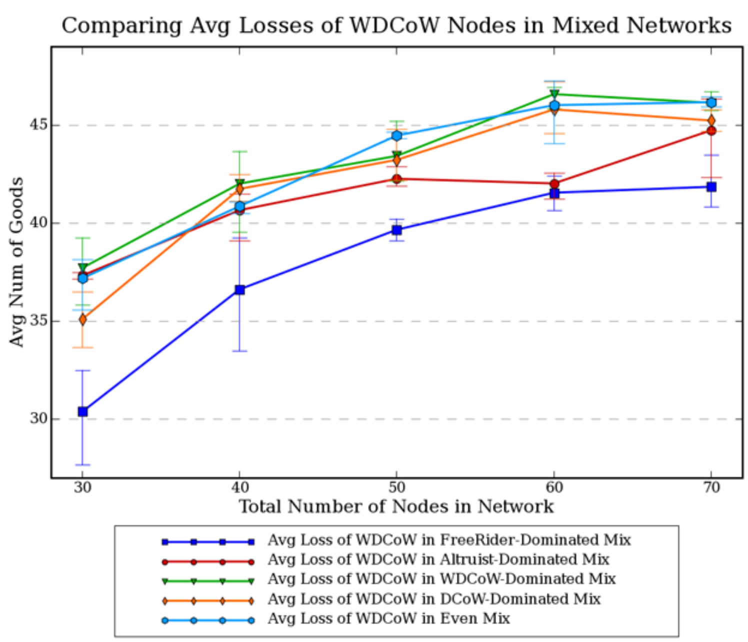


Figure VII.8: Comparing the Average Losses of WDCoW in Heterogeneous Networks which is related to the fact that majority of this network population insists on bartering interactions and does not initiate philanthropic transactions. The *Loss* levels in other cooperative networks are relatively high and similar to each other. This similarity and consistency is related to the versatility of the *WDCoW* strategy. The exchange policies of this strategy enable the *WDCoW* nodes to deliver consistent performance in a wide range of network populations. Similar to other strategies, the *WDCoWs* exhibits the worst performance in the *FreeRider-dominated* networks.

### VII.B.4 Performance of DCoWs in Heterogeneous Networks

FigureVII.9 and VII.9 compares the *Gains* and the *Losses* of the *DCoW* nodes. Let's examine which of the network configurations is the best suited network for the *DCoW* nodes. Similar to the *Altruists* and the *WDCoWs*, the *DCoW* nodes get the lowest *Gains* and the lowest *Losses* when they operate in the *FreeRider-dominated* networks. So the *FreeRider-dominated* network is once again the worst network configuration for collaborative interactions. The *Gains* in the other networks are higher and relatively similar. However, at a closer look, the highest *Gains* are archived in the *WDCoW-dominated* networks. This slight edge can be contributed to the fact that, the *WDCoWs* can generate *Proposals* that benefit the *Inquiring* node more than the *Proposing* node. Thus, the *DCoWs* occasionally get involved in transactions with *WDCoWs* where they *gain* more than they give away. This aspect of the exchange process elevates the *Gains* of the *DCoWs* in the *WDCoW-dominated* networks over the *Gains* of the *DCoWs* in the *DCoW-dominated* networks. The *Losses* in this network configuration are also relatively high, partially because there are a lot more opportunities to barter. Interestingly, the *DCoW* nodes suffer the highest *Losses* when they are surrounded by their own kind. These levels of *Losses* are related to the rigid bartering policies that are executed by the *DCoWs*. When these nodes initiate transactions, the size of the *Receive List* is equal to the size of the *Give List* and as a result, the *DCoW* nodes in the *DCoW-dominated* networks tend to loose as much as they receive. The *DCoWs* fair well in the *Altruist-dominated* networks. In this network configuration, the *Gains* are relatively low but so are the *Losses*. Overall, the *DCoW* nodes exhibit consistent performance in the networks that are largely populated with other cooperative nodes. The stability and consistency of the *DCoW* nodes and the *WDCoW* nodes shows that the bartering approach is a robust collaborative methodology which is well suited for opportunistic collaborations and exchanges in mobile peer-to-peer heterogeneous environments.

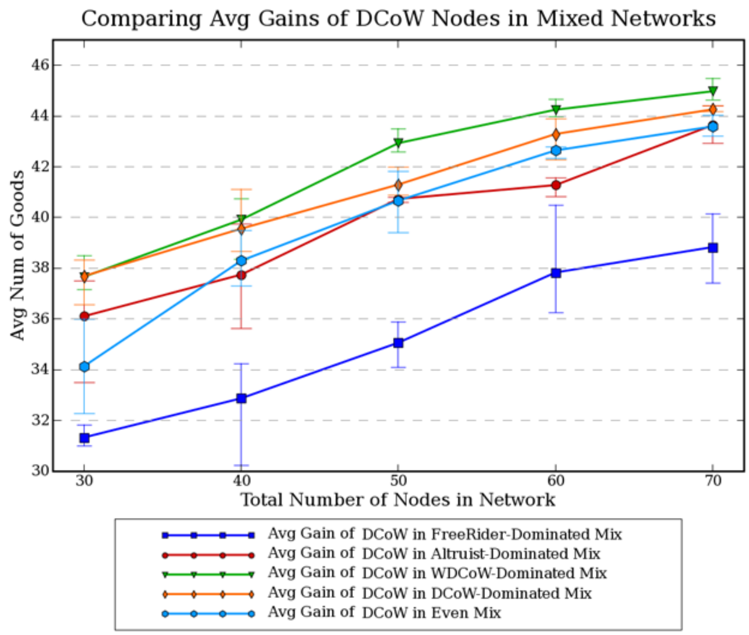


Figure VII.9: Comparing the Average Gains of the WDCoW in Heterogeneous Networks

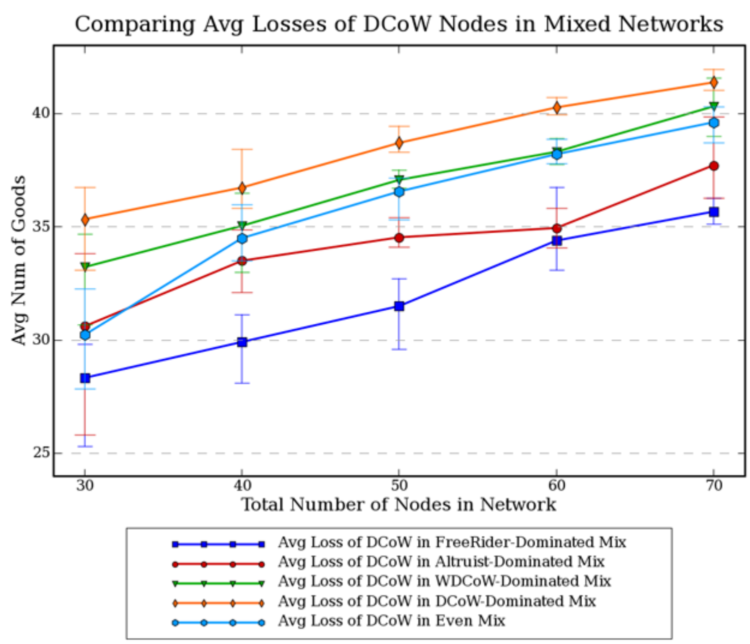


Figure VII.10: Comparing the Average Losses of the DCoW in Heterogeneous Networks

## VII.C Comparing Communication Overhead in Strategy-Dominated Heterogeneous Networks

As discussed earlier, communication overheads can have significant impact on the node's efficiency. However, in addition to considering the efficiency of a single node, we also need to examine the efficiency of the network as a whole. In this section, we compare the average communication overheads for each of the four *dominated* heterogeneous networks. FigureVII.11 shows the average number of successful transactions in these networks, while, FigureVII.12 shows the average size of these transactions.

As expected, the *FreeRider-dominated* networks have the lowest transaction count while, the *Altruist-dominated* networks have the highest transaction count. Also, the average transaction size in both of these networks is the smallest compared to other heterogeneous network mixes. Clearly, the prominent characteristics for these conventional strategies have an adverse affect on the overall network efficiently. In contrast, the *evenly mixed* networks, the *WDCoW-dominated* and the *DCoW-dominated* networks have very similar transaction counts. However, the *WDCoW-dominated* networks display the highest average transactions size. This distinction, sets the *WDCoW-dominated* networks aside by delivering the lower overall communication overheads. This encouraging performance is related to the *WDCoW* exchange policies which are well balanced since they are not too rigid nor too liberal. In essence, these policies are not at either extreme of the collaboration continuum. Thus, in the *WDCoW-dominated* networks, the average network operating costs are relatively low and the productivity of the nodes, in terms of average levels of *Gains*, is relatively high. Similar to the *WDCoW-dominated* network, the *DCoW-dominated* networks also exhibit low operating costs since these networks have a low transaction count and a high average transaction size. By looking at these heterogeneous networks, we can conclude that communication overheads in the networks dominated by the nodes that are utilizing the proposed bartering approach are low compared to the networks dominated by the conventional strategies.

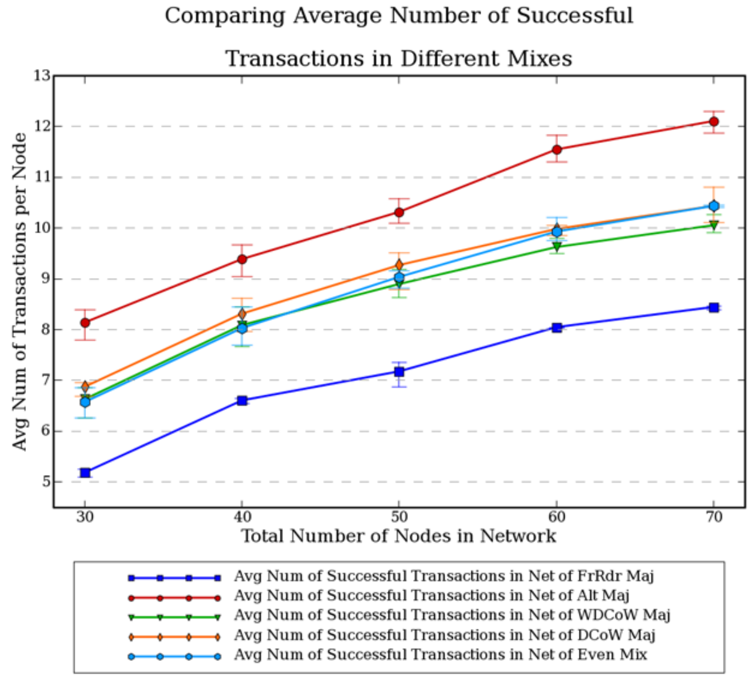


Figure VII.11: Comparing the Average Gains and Losses of Nodes Running Different Strategies in Evenly Mixed Heterogeneous Networks

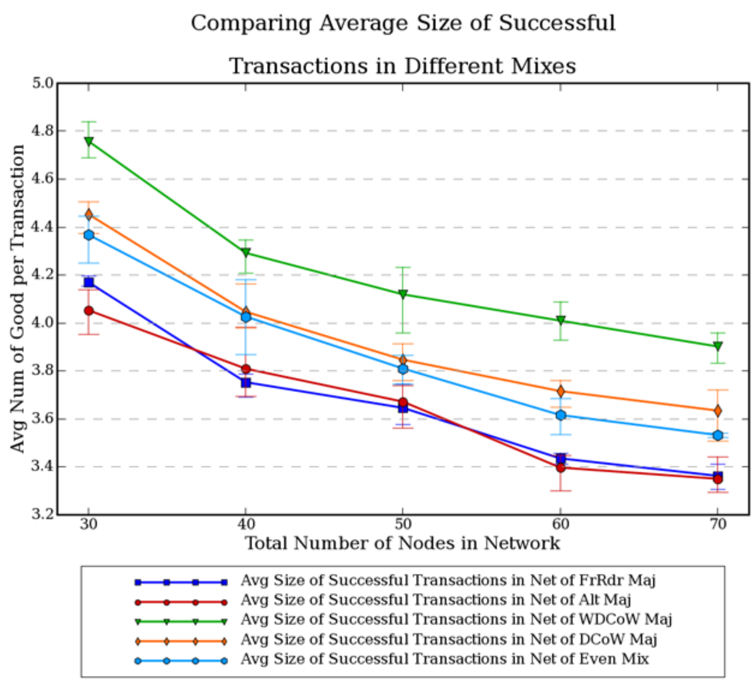


Figure VII.12: Comparing the Average Gains and Losses of Nodes Running Different Strategies in Evenly Mixed Heterogeneous Networks

### VII.C.1 Comparing Rejections in Strategy-Dominated Heterogeneous Networks

As mentioned earlier, *Rejections* are one of the major contributors to the communication overhead related costs. FigureVII.13 shows levels of rejected *Proposals* in heterogeneous strategy-dominated networks. When comparing these levels of *Rejections* in these networks, it is not surprising that the *FreeRider-dominated* networks have the highest *Rejection* rates while, the *Altruist-dominated* networks have the lowest *Rejection* rates. Also, as expected, the *DCoW-dominated* networks experience relatively high *Rejection* rates. These high rates are attributed to the *DCoW's Proposal Evaluation Policy* which rejects all unbeneficial exchanges. Interestingly, both, *evenly mixed* networks and the *WDCoW-dominated* networks have very similar *Rejection* rates. This similarity can be attributed to the fact that *WDCoWs'* collaboration policies do not obstruct collaborative interactions.

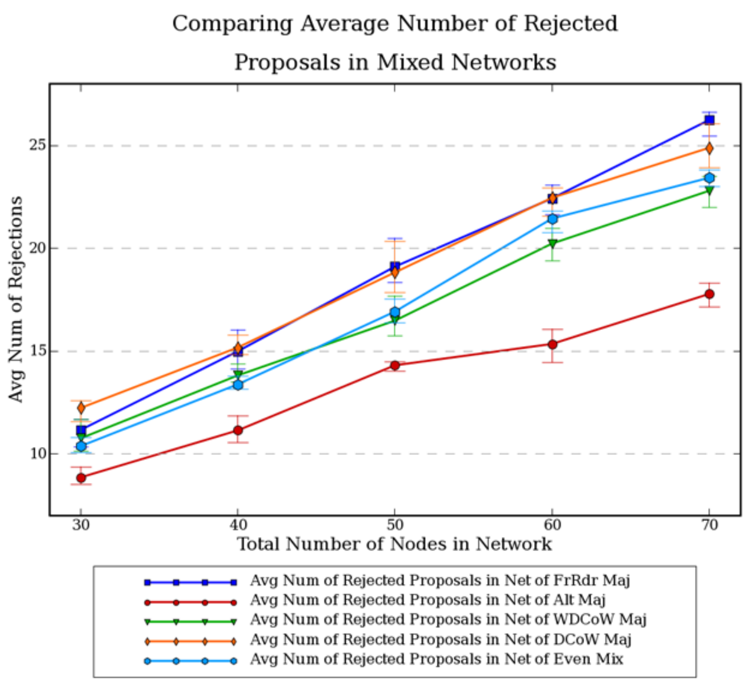


Figure VII.13: Comparing the Average Gains and Losses of Nodes Running Different Strategies in Evenly Mixed Heterogeneous Networks

## VII.D Summary and Discussion

To conclude, the *WDCoW* and the *DCoW* bartering strategies once again prove to be effective collaborative interaction methods in heterogeneous networks. These strategies show good levels of resilience to the diverse network conditions since the nodes that are executing these strategies deliver relatively high levels of *Gains* and experience relatively modest and tolerable levels of *Losses*. These strategies also protect nodes from excessive operational costs. Nodes that are executing the *WDCoW* and the *DCoW* strategies also exhibit lower communication overheads by either having few large-volume transactions or by participating in more even and balanced exchanges that protect these nodes from recklessly giving away their goods to other nodes in the environment. In contrast, the very cooperative *Altruists* and the extremely selfish *FreeRiders* are not as effective in these heterogeneous environments. On average, the philanthropic nodes that are executing the *Altruistic* strategy incur dramatic *Losses* since other nodes tend to take advantage of their liberal collaboration policies. While, the *FreeRiders*, unwilling to share their goods, are sidelined by *Rejections* from the less liberal nodes and thus have very limited *Gains*. Also, the *Altruists* incur major communication related overhead since they mainly participate in great number of low-volume transactions. Likewise, due to the very high *Rejection* rate, the *FreeRiders* get very high communication overheads as well. Thus, bartering exchange policies provide versatile collaboration mechanisms that are well suited for a wide range of network configurations.

# Chapter VIII

# VALUE-BASED BARTERING AND INVESTMENTS APPROACH

In the previous Chapters, we have described the key concepts of the proposed bartering collaboration model and contrasted this model with conventional widely used *Free Riding* and *Altruistic* collaboration approaches [12, 96, 95]. In this Chapter, we examine two extensions to the bartering model. We take a close look a the valuation sensitive bartering and at the investment enhanced trading.

## VIII.A Value Based Bartering

The bartering collaboration model presented and analyzed in Chapters IV, V and VI makes an assumption that all of the digital goods and digital content in the mobile environment are of the same equal value. In essence, when nodes setup a collaborative exchange they do not consider the importance and the significance of the goods being given away or being purchased. We can will refer to this initial valuation model as the ***Equal Valuation*** model (***EV***). Clearly, this model does not capture the dynamic personal nature of the mobile peer-to-peer environments. These environments are populated with diverse set of users with varying interests and partiality. Users in these environments assign different levels of worthiness and significance to the digital content that they desire, possesses. and consume. Valuing all digital goods with an equal value eliminates an important dimension of the peer-to-peer collaborations. However adding value-sensitive aspect to the bartering collaborations adds a layer of complexity that can hinder nodes' collaboration productivity by limiting collaboration opportunities.

### VIII.A.1 Valuation Models for Digital Goods

To challenge the use of our initial *Equal Valuation* model, we have developed a set of more descriptive valuation models that allow us to measure the effects of valuation sensitive bartering for digital goods and content during the opportunistic peer-to-peer bartering exchanges. In particular, we have developed three valuation models and incorporated them into our collaborative exchange framework described in Chapter IV. These three models are: the *Demand Sensitive Valuation* model, the *Demand and Supply Sensitive Valuation* model and finally the *Personalized Valuation* model. Each of these models is incorporated into the proposed bartering framework and can be used to enrich the bartering collaboration process. Each model uses a particular method to develop a set of values that represent valuation of each good and associate these value with every digital good in the environment. Each of these three models uses its particular approach to measure the importance of every good in the environment. All of the values, in every one of these models are incorporated into the bartering process. With these valuation models, during the battering process, the nodes do not consider the number of goods that are being exchanged but instead examine the value of the ongoing transaction. This allows nodes to add extra layer of awareness that captures yet another dimension of collaborations and exchanges in mobile pervasive environments. From the framework implementation perspective, the *Proposing node* ensures that the value of the *Receive List* in the *Proposal* is equal to the value of the *Give List* of this *Proposal*.

Before we give the detailed description of the three proposed valuation models, we need to present the details of how our framework interprets the concept of "value" of a digital good. Identifying and capturing the valuation of a digital good in a qunatifiable discrete number that represents this valuation is a changing process. Debates about what is a "value" and what is the best way to capture this concept of "value" has been extensively looked at in a number of scientific domains including Social Sciences, Economics and Ethics [90, 98, 81, 46]. Though there are a number of ways to interpret valuation concept, for the purposes of our research, we have limited our interpretation of "value" by treating it as a quantifiable unit of measurement. Specifically, for the purpose of our proposed framework, we sort all of the digital goods into ten threshold-like categories. Goods belonging to the "category-1" are considered to be the least valuable goods in the environment. Goods belonging to the "category-10" are considered to be the most valuable goods in the environment. This approach delivers a well defined scale that is used to unambiguously quantify the valuation for every good in the network. This categorization approach is different from the conventional approach of assigning a currency based "price tag". With categorization, not only we can capture the relative differences between the good values but also limit the inflation and deflation of values in the network. In essence, every one of our three valuation models normalizes the values of the goods by categorizing them in to one of these categories. Our current implementation does not handle issues of value inflation and deflation. However, normalizing the values and reassigning goods to new categories can be used to address issues of value inflation and deflation. Also, for purposes of comparison and analysis, we have identified "category-5" to the be the only category used the networks employing the *Equal Valuation* model.

Now that we have defined our method of quantifying the value, we move on to describing the valuation modes that our bartering framework can employ during the peer-to-peer exchanges.

### Demand Sensitive Valuation

The ***Demand Sensitive Valuation*** model (***DSV***) is a model that computes values for each of the goods in the environment based on the level demand for every good. This model is designed with a basic assumption that, the valuation of each good should be directly related to the demand that nodes have for this good. We assume that, the values can be externally computed at the beginning of the collaboration cycle and are shared with all of the nodes at the beginning of this collaboration cycle. In essence, this model represents a "categorized global valuation list" that is built by evaluating the demands of every node in the environment. For example, if the environment contains a great number of node that are looking to acquire a particular good then that good is categorized into a "high value" category. On the other hand, if a good is in a very low demand then it is categorized into a "low value" category.

Note that in our current simulation evaluation, the demand is explicitly stated in the *iWant List* of every node. One can envision a system where, mobile nodes periodically synchronize their *iWant Lists* with a central server (or a set of distributed servers) which could identity the current demand trends and assign appropriate values that reflect this demand. This synchronization assumption is not unreasonable, since currently many users of mobile technology already frequently synchronize their mobile devices with some server.

As we have described in Section IV.E.1, our current simulation evaluation of the bartering collaboration framework employs *Power Law distribution*[31, 24, 15] which is used to generate list of wishes for every *iWant List* in the network. This uneven distribution of wishes also creates unevenly distributed demand for the digital content in the network. Thus, when demand sensitive values are generated, they also reflect these uneven levels of desire for each of the goods in the environment.

Below is a *pseudo code* used to derive the value related categorization of good in our simulation evaluation of *DSV* model. Note that the $DemandCount$ is a list-type data structure that is populated with a count of the number of times each good is listed in the *iWant Lists* of the nodes in the network. Basically, this list contains the demand levels for every good. Similarly, the $maxDemandCount$ and the $minDemandCount$ represent the count for the least demanded good and the count for the most demanded good (respectively) in the environment. Using these two parameters, the valuations are normalized into the previously described categories.

**Algorithm 1** *Computing Global Valuation List Based on Demand*

**Input:** Array of $DemandCount$ for digital goods, $maxDemandCount$, $minDemandCount$
**Output:** Array of Digital Good Valuations

$coef = \frac{10-1}{maxDemandCount - minDemandCount}$

**for** every $Good$ in $ListOfGoods$ **do**

$Value = coef * (DemandCount[GoodID] - minDem)$
$Value = \mathsf{round}(Value + 1)$
$GlobalListOfValues[GoodID] = Value$

**end for**

**Demand and Supply Sensitive Valuation**

The ***Demand and Supply Sensitive Valuation*** model (***DSSV***) builds valuation categories in the similar manner as the *DSV* model. The key assumption used in this particular model is that the valuations of the digital goods and content are not only dependent on how many nodes are looking for these goods but also dependent on the levels of available supply of these goods in the environment. For example, the goods that are in high demand and low supply are categorized into a “high value” category. On the other hand, if a good is widely available and very few nodes are looking to acquire this good then this good is categorized into a “low value” category. Thus, this model captures additional aspects of valuation that further improves the categorization mechanism used to identify valuations of the goods in the network.

Note that, in our current simulation evaluation, the demand can be determined by looking at the nodes' *iWant Lists* and the supply can be determined by looking at the nodes' *iHave Lists*. Similar to *DSV* model, at the beginning of the collaboration cycle, nodes are bootstrapped with a set of *DSSV*-compliant categorized valuation list that reflect the significance of every good in the network.

Below is a *pseudo code* used to derive value related categorization of good in our simulation evaluation of *DSSV* model. Note that, the $DemandCount$ is a list-type data structure that is populated with a count of the number of times each good is listed in the *iWant Lists* and the $SupplyCount$ is a list-type data structure that is populated with a count of the number of times each good is listed in the *iHave Lists*. Essentially, these lists contain the demand and supply levels in the network.

**Algorithm 2** *Computing Global Valuation List Based on Supply and Demand*

**Input:** Array of $DemandCount$, Array of $SupplyCount$, $maxDemandCount$, $minDemandCount$
**Output:** Array of Digital Good Valuations

$coef = \frac{10-1}{maxDemandCount - minDemandCount}$

**for** every $Good$ in $ListOfGoods$ **do**

  **if** $SupplyCount[GoodID] != 0$ **then**
    $ratio = \frac{DemandCount[GoodID]}{SupplyCount[GoodID]}$
  **else**
    $ratio = maxDemandCount$
  **end if**

  $Value = coef * (ratio - minDemandCount)$
  $Value = \mathsf{round}(Value + 1)$
  $GlobalListOfValues[GoodID] = Value$

**end for**

#### Personalized Demand and Supply Sensitive Valuation

The ***Personalized Demand and Supply Sensitive Valuation*** model (***PDSSV***) is a model that takes it one step farther than the *DSSV* model. In addition to being sensitive to the demand and supply levels, this model assumes that not all users have the same value for the same good. In essence, the assumption is that users have different preferences and tastes for the digital content that exists in the network. The *PDSSV* assumes that the valuations are similar but not necessarily totally identical to the globally available categorized *DSSV* list. These deviations and preferences in the valuation method better reflect the heterogeneous nature of the mobile environments that are inhabited by the users with diverse interests. Our framework assumes that nodes are already bootstrapped with well expressed profiles that reflect personal preferences and interests of the mobile users. There are a number of ways to acquire the profiles, they can be learned and derived from user's day to day activities[18, 17], or they can be built from the widely popular online websites such as orkut.com[2], mySpaces.com [3] and FaceBook.com[4].

Similar to the *DSV* and the *DSSV* models, the *PDSSV* model also assumes that there is a synchronization of the demand and supply information. This model relies on the distribution of the global valuation categorization list. However, once each node receives this categorization list, it reexamines this list and modifies the values to better suit its' personal preferences. Essentially, nodes consider the global valuation categorization list to be a guideline. Each node takes these guideline into account and modifies valuations to better reflect its preferences.

From the prospective of our simulation evaluation study, the networks that are using the *PDSSV* model, at the beginning of the collaboration cycle assemble a global valuation list according to the *DSSV* model. During the bootstrapping, each node receives the global categorization list. Each node has an option of personalizing these valuations in this list to better reflect the user's personal preferences. The node also has an option of leaving some or all of the values unchanged. Personalizing the valuations is essentially done through re-categorizing the valuations of the goods. The re-categorizing is limited to upgrading to a category directly above the original category or by downgrading to the category directly below the original category. Currently, in our simulation, a node at random generates one of the following three integers $[-1, 0, 1]$ which symbolize the action that the node takes in regards to re-categorizing. Note that, goods in the valuation "category-1" cannot be further downgraded and the goods in "category-10" can not be further upgraded. This personalization approach allows nodes to have similar but not identical valuations of the goods in the environment. Additional valuation methods and models can be added to the framework . However, any of the valuation models would need to overcome the issues of additional complexity that valuation sensitive approach brings in to the bartering process.

In Chapter V, we have already described behavior of the homogeneous *DCoW* networks. In the following Sections, we will compare the performances of the nodes executing the *DCoW* strategy while employing the valuation sensitive bartering models which we described above.

### VIII.A.2 Gains in Valuation Sensitive Bartering

To better understand the effects of introducing the valuation concept into the bartering process, we compare the *Value Gains* that the nodes experience in the homogeneous bartering network executing the *DCoW* strategy. The ***Value Gain*** of a node represents how much overall value, on average, did the node acquire during the simulation. We examine the *Value Gains* in the four valuation sensitive modes: the *EV* model, the *DSV* model, the *DSSV* model and finally in the *PDSSV* model. FigureVIII.1 shows *Value Gains* which are normalized for the purposes of comparison.

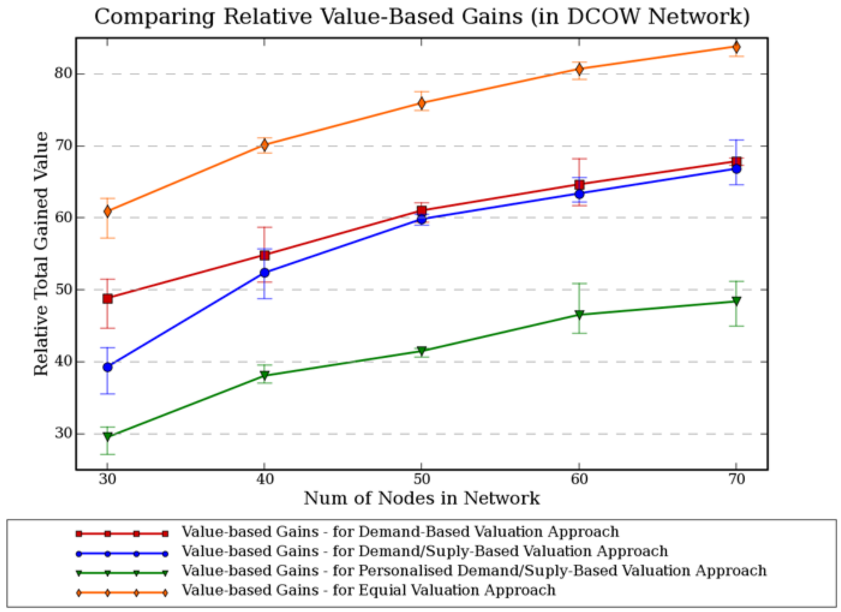


Figure VIII.1: Gains in Valuation Sensitive Approaches

Note that, the most productive set of nodes are the nodes that are employing the *EV* model. The reason behind such high *Value Gains* for this model are related to the fact that this model does not add valuation-related noise to the bartering process. The evenness of the good count is the primary concern of the nodes executing the *EV* model. The other three valuations models add valuation categorization thus forcing nodes to employ value-sensitive "filtration" which impose additional restrictions on the terms of the exchange. When comparing the three value-sensitive models, the *PDSSV* model stands out as the least productive model. This model uses personal valuation approach which introduces additional noise into the value categorization process. This noise resonates into the bartering transactions and further complicates the matching process. This more complex matching process affects the overall productivity of the nodes and eventually materializes into lower *Value Gains* for the *PDSSV* networks. The *DSV* model and the *DSSV* model both show relatively similar levels of the *Value Gains* with a sight edge to the *DSV* model. The reasons for this similarity are related to the communication overheads which are discussed in the next Section.

### VIII.A.3 Impact of Transaction Overheads on Valuation Sensitive Bartering

The importance of communication overhead was previously discussed in Sections V.B and VI.A.2. In the case of valuation sensitive bartering, there is an additional parameter of the *average transaction valuation* that can be used to describe effectiveness collaboration. In this Section, we also discuss the average transaction valuation which represents the average total value of the average transaction. The average transaction valuation complements the average size of the transaction parameter. The greater the average transaction valuation the more effective is the exchange.

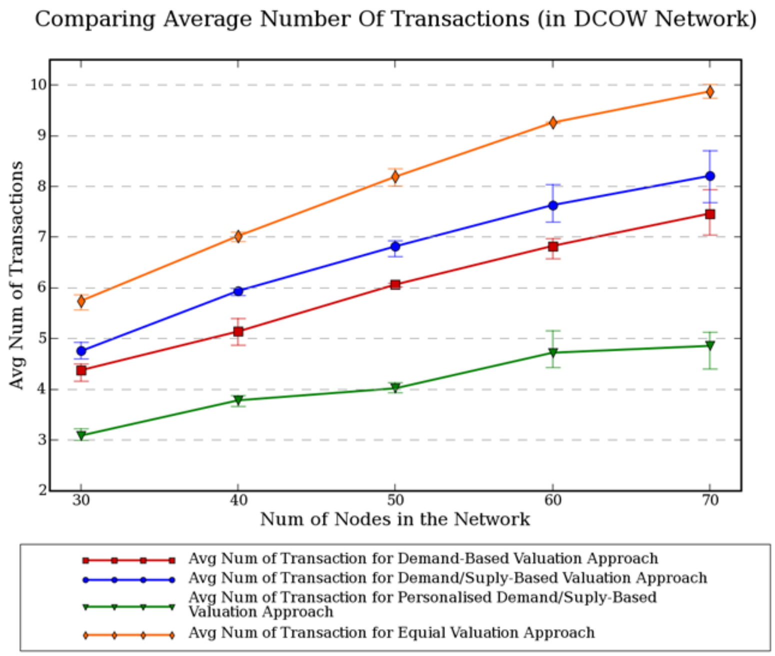


Figure VIII.2: Transaction Count in Value-Based Approach

FigureVIII.2 shows the average number of transactions, while FigureVIII.3 and FigureVIII.4 show the average size of bartering transactions and the average valuations of the bartering transactions. Note that the average transaction valuation are normalized for purposes of comparison.

Lets start our examination of the communication overhead by looking at the performance of the simplistic *EV* model. In the *EV* networks, the nodes conduct the largest number of transactions with the smallest average transaction size. The average transaction valuation is also very low. As mentioned earlier, all of these factors contribute to the bartering inefficiency. Thus, the *EV* model is the least efficient collaboration model. Combining this analyses with the analyses of *Value Gains*, it is clear that, the main reason behind the high *Value Gains* exhibited by the *EV* nodes needs to be attributed to the very high quantity of the transactions and not to the nodes' efficient bartering behavior.

On the other hand, the networks that are employing the *PDSSV* approach exhibit relatively efficient bartering behavior since their transaction count is low and the average transaction size and the average transaction valuations are both relatively high. However, this efficient behavior is not exhibited in a sufficient quantity

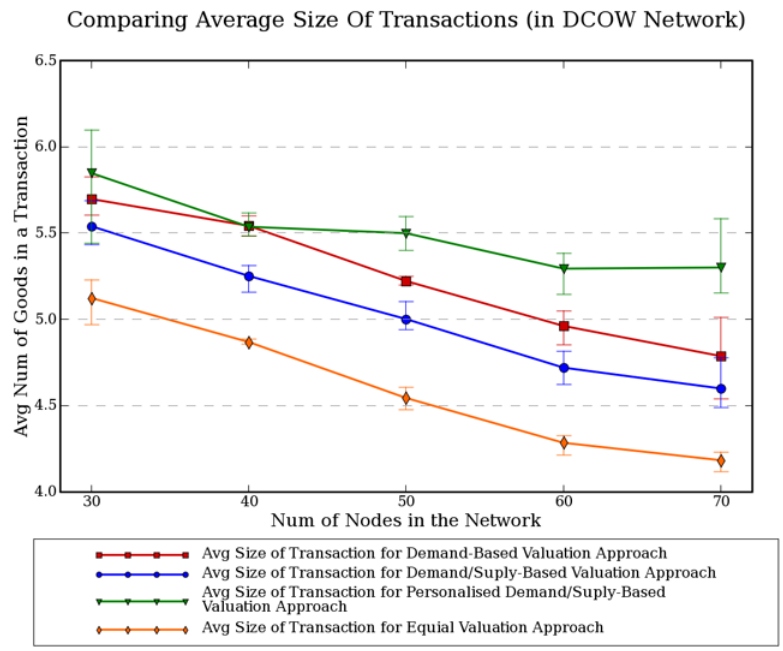


Figure VIII.3: Transaction Size for Value-Based Approach

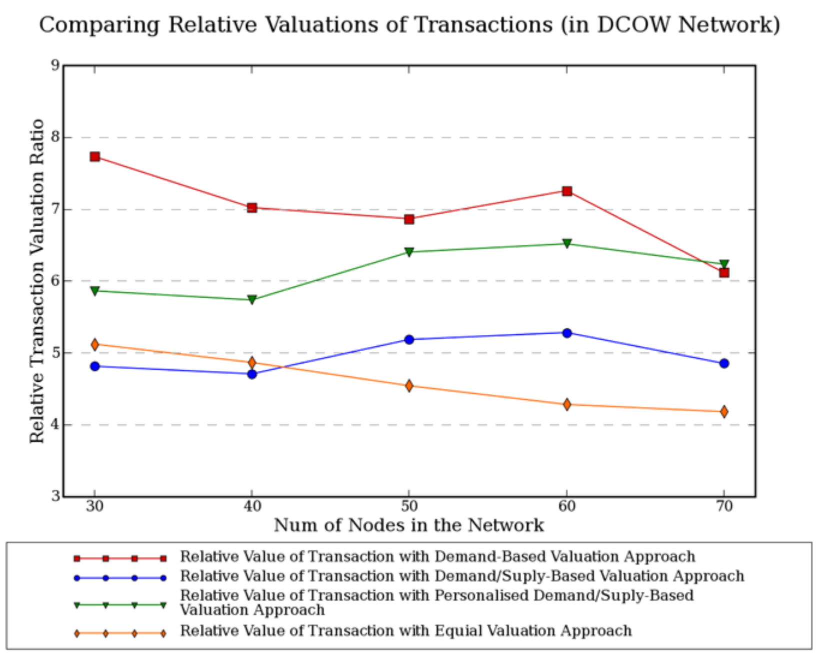


Figure VIII.4: Transaction Valuation for Value-Based Approach to contribute to the overall network productivity. As mentioned earlier, the noise generated by the personal value re-categorization limits the pool of possible exchange deals which in turn, hinders the overall network productivity.

Now, lets compare the performance of the *DSV* and the *DSSV* networks. Though the average transaction size and the average transaction valuations of the *DSV* networks are higher than the size and valuation of the transactions of the *DSSV* networks, the transaction count shows the opposite trend. The *DSSV* networks have greater average transaction count than the *DSV* networks. Lack of a clear-cut difference contributes to the similarity in the levels of the productivity that are exhibited by these networks. Thus, based on just these parameters, there is no clear-cut benefit of using the *DSSV* model over the *DSV* model. They deliver a relatively even performance.

## VIII.B Investment-Based Approach

In attempt to improve the average *Value Gains* and increase nodes collaboration related effectiveness, we have developed an extension to the value-sensitive bartering approach described in the previous section. This extension is built on the concept of investment [41, 103, 102] which involves purchasing goods with intention to resell these goods. The application of this concept to our framework involves nodes purchasing digital content that is not part of the nodes' initial *iWant List* in hope of later re-trading this content for goods that are in the nodes' original *iWant List*. Essentially, the nodes take a risk and acquire not-needed goods for purposes of later trading them away for the desired content. The key contribution of our work is to employ this investment based approach to improve opportunistic bartering collaborations in mobile peer-to-peer environments.

As described in Section IV.C, nodes executing the *DCoW* strategy insist on even trades where the good count of the *Receive List* is equal to the good count of the *Give List*. In case of the value-sensitive bartering, the basic trading philosophy is modified such that the bartering nodes compare the total value of the *Receive List* with the total value of the *Give List* and not the cardinality of these lists. At the start of the transactions, the *Proposing node* establishes an initial "match" and then using its *Proposal Composition Policy*, it trims the *Lists* to be of an equal value. Clearly, trimming the lists reduces the transaction size and thus effects bartering communication efficiency. The occasions where the *Give List* needs to undergo trimming are frequent. Using the proposed investment based trading extension, the nodes take an opportunity and instead of trimming the

*Give List*, attempt to expand the *Receive List*. This expansion is done to equate the value of the *Receive List* with the value of the *Give List*. The goods that are added to the *Receive List* are selectively picked for the *iHave List* of the bartering partner. The selection strategy is simple, the *Proposing node* identifies the cardinality and the total value of the goods that were originally planed to be trimmed away. The *Proposer* looks through the *iHave List* of the bartering partner and attempts to find a set of goods such that this set has a smaller cardinality and a similar value as the trim set. In essence, the node tries to exchanges it's "unwanted" goods for a smaller set of more valuable "unwanted" goods. The reasoning is that, if the node reduces the cardinality of its *iHave List* and also increases the total value of the goods in this *iHave List*, it will have a better chance of giving "unwanted" goods way and acquiring "needed" goods. Basically, the philosophy is that in the environment where the value is derived from the network-based demand it is better to have few valuable goods than many less valuable goods. The hope is that other nodes in the network will find these goods desirable and will be willing to pay for them with the goods that this node is originally interested in acquiring. Clearly, there is a risk of acquiring a good and not being able to trade it away. In the next Section, we will look at the performance of the nodes employing the investment enhanced trading.

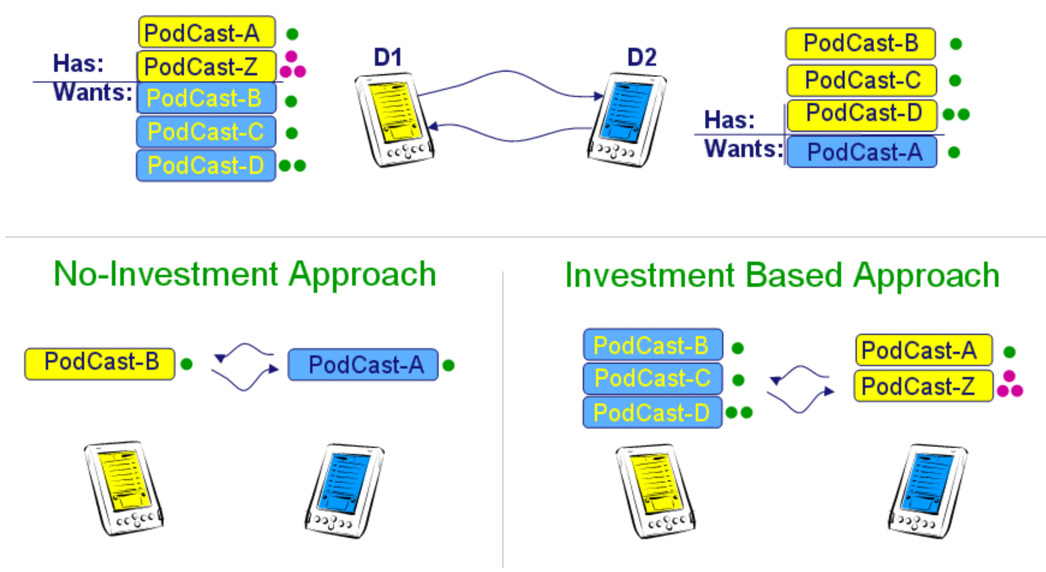


Figure VIII.5: Scenario for Investment-Based Approach

To simplify our analysis description we will refer to the value-sensitive networks that employ investment based extension as the *iDSV*, the *iDSSV* and the *iPDSSV* networks.

### VIII.B.1 Valuation Gains in Investment Based Approach

Lets start our analysis by comparing the average *Value Gains* of the networks employing the investment enhancements of the value-sensitive bartering. In Section VIII.A.1, we have described a set of valuation models. FigureVIII.6 shows the average *Value Gains* for each of the three value-sensitive models with and without investment based extension. Firstly, the main trend is that all three models saw a positive improvement in their *Value Gains* when employing the proposed investment based extension. Both, the *iDSV* and the *iDSSV* networks experience similar improvements, while the *PSSSV*, which was lagging behind, goes through the most dramatic improvment that elevates its *Value Gains* to the levels of the other two value-sensitive networks. The reasons for this major improvement in the *iPSSSV* networks will be discussed later. However, an interesting conclusion that can be derived from this graph is that even with this investment based extension the value-sensitive trading is still delivering the lower *Value Gains* than the initially used *Equal Value* model. Clearly, the value-sensitive trading adds a significant communication complexity that hinders the productivity of collaborative interactions. Though, the investment based extension helps the value-sensitive trading to reduce the effects of the filtering, these improvements are not substantial enough to negate the drawbacks of "filtering".

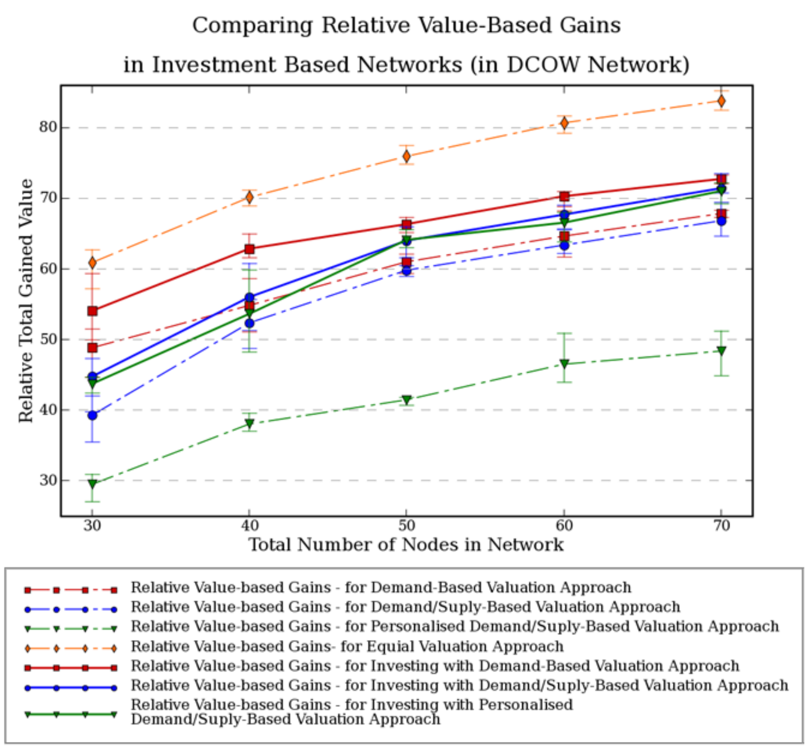


Figure VIII.6: Gains in Investment-Based Approach

### VIII.B.2 Communication Overheads of Investment-Based Approach

To better understand the dynamics of the investment based trading in mobile peer-to-peer environments, lets take a closer look at the transaction count and the average size and value of the average transactions occurring in these networks.

FigureVIII.7 shows that investment based trading has a mixed impacts on the value-sensitive networks. In particular the *iPDSSV* networks see a major increase in the transaction count while the *iDSV* and the *iDSSV* networks experience a minor drop in the average number of transactions. FigureVIII.8 shows that, when the networks employ investment based extension, the average transaction size in the *iDSV* and the *iDSSV* networks experience a relatively even increase. However, FigureVIII.9 shows that the *iPDSSV* networks undergo a sizable increase in average valuation of their transactions, while both the *iDSV* and the *iDSSV* networks experience moderate and relatively similar increase for the average transaction value.

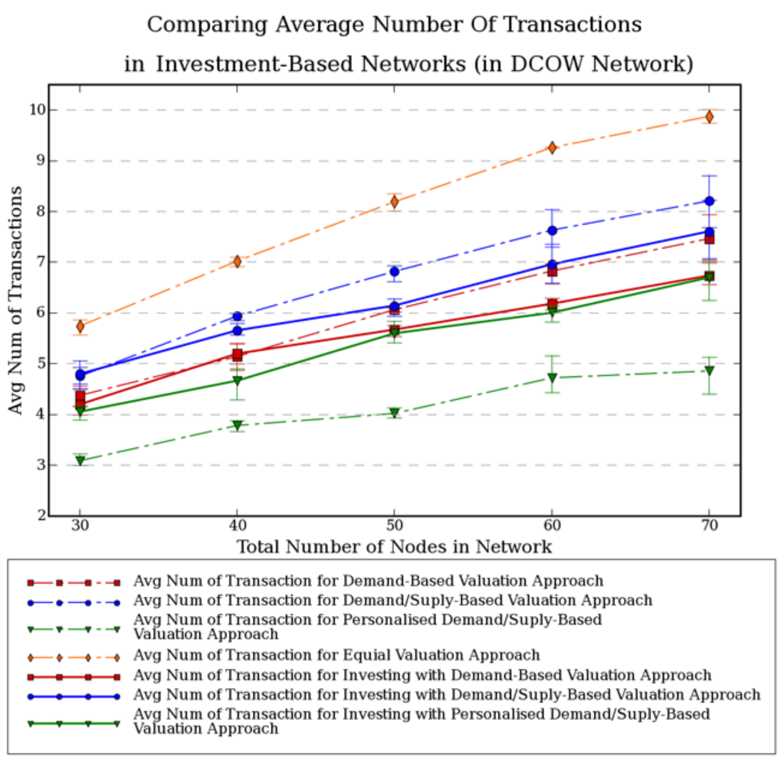


Figure VIII.7: Transaction Count in Investment-Based Approach

From these three graphs, it is clear that the dramatic increase in *Value Gains* of the *iPDSSV* networks are attributed to the increase in the overall collaboration volume in these networks. Essentially, investment based approach provides *iPDSSV* nodes with the extended flexibility such that these nodes are able to participate in a greater number of transactions that, on average, are of a higher value than the transactions in the originally

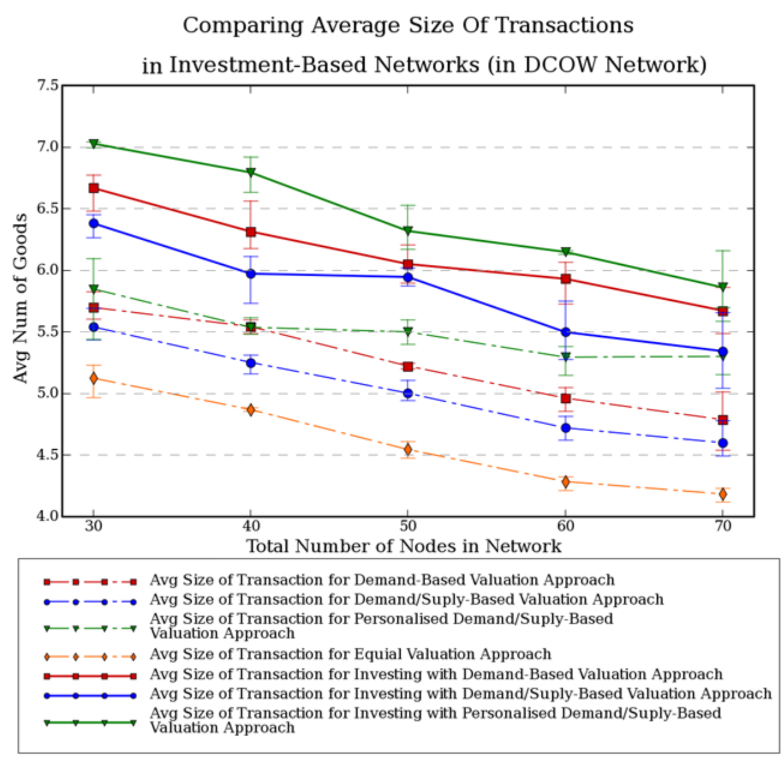


Figure VIII.8: Average Transaction Size in Investment-Based Approach

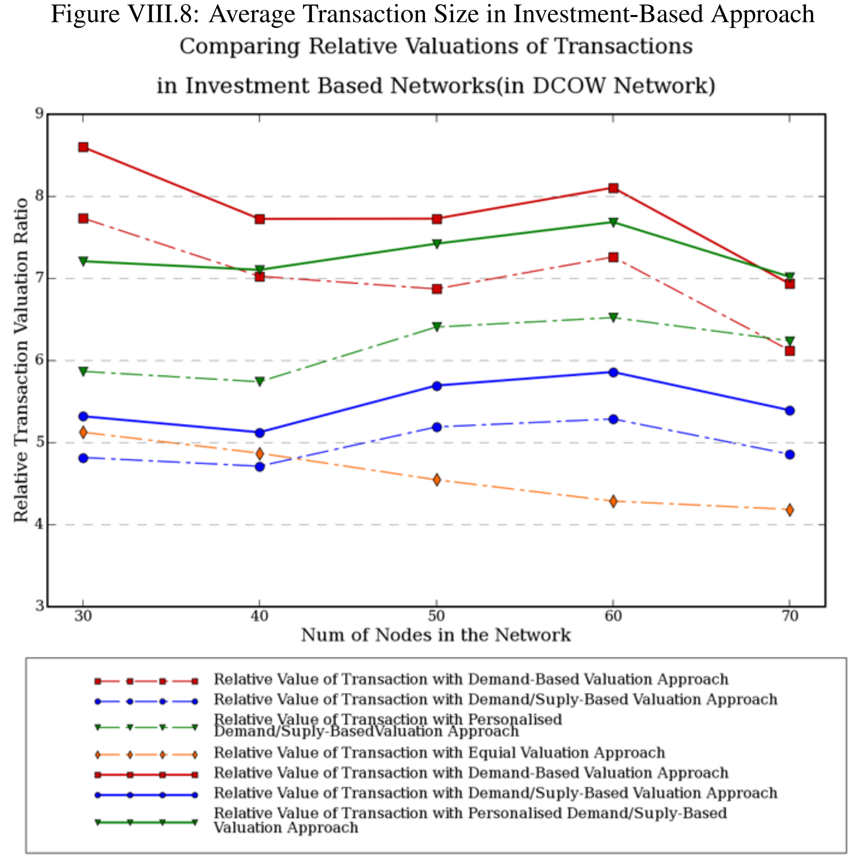


Figure VIII.9: Average Transaction Value in Investment-Based Approach described *PDSSV* networks. In contrast, the *iDSV* networks and the *iDSSV* networks experience the increase in their *Value Gains* due to the general increase in overall efficiency of the collaborative transactions.

### VIII.B.3 Impact of Failed Investments

Similar to concept of investing in the real world, the proposed investment approach can also result in "failed investments". It is unreasonable to expect that every investment will be a successful investment. Thus, there is distinct possibility that a mobile node acquires a set of goods as an investment and does not encounter any other nodes that are interested in any of these "investment" good. This accumulation of unwanted goods clearly can add up and impact the overhead and operational costs of the bartering nodes. On the other hand, since the unwanted goods are replaced by another set of unwanted goods the impact of this overhead is not significant. FigureVIII.10 show the normalized valuation of the failed investments. Clearly, the *iDSSV* has the best record of turning around investment since it has the lowest failure rate.

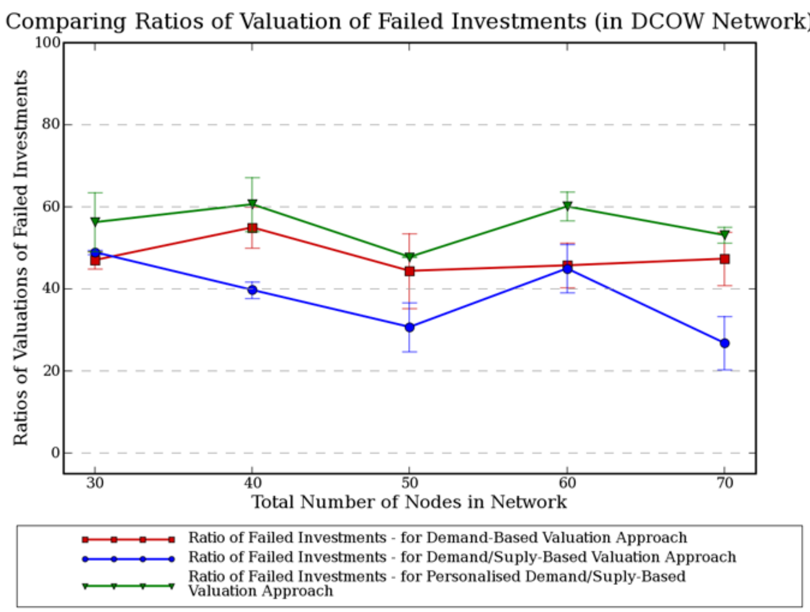


Figure VIII.10: Impact of Failed Investments - Value Count

On the other hand the *iDSV* is not as effective as the other value-sensitive networks. The main reason behind this difference is the fact that valuation categorization of the *iDSSV* relies on both the network-based demand and supply as opposed to the *iDSV* derives its categorization only from the levels of demand. The *iPDSSV* networks display the highest failure rate compared to the other value-sensitive bartering networks. The primary reason for this poor performance is that, the nodes in the *iPDSSV* networks employ the "fuzzy" categorization of content valuations which adds more uncertainty to the investment decision making. However, despite the high the failure rate, the *iPDSSV* model experiences a very strong improvement in the *iPDSSV's Value Gain* performance. Clearly, for the *iPDSSV* networks the benefits of the outweigh the costs.

To conclude, the investment based trading delivers major improvements to the collaborative interactions in the *iPDSSV* networks. This extension provides flexibility and long term collaboration benefits to the otherwise excessively complex valuation methodology. Furthermore, the investment based trading also improves the performance of the *iDSV* and the *iDSSV* networks.

## VIII.C Summary and Discussion

In this Chapter we have described two extensions to the proposed bartering communication model. We have analyzed the value based bartering and the investment based trading approaches. In particular, we have described three value-sensitive digital good categorization methods. The *Demand Sensitive Valuation* model, employed Demand based valuations. The *Demand and Supply Sensitive Valuation* model relied on both demand and supply to derive the valuations. Finally, the *Personalized Valuation* model extended the *DSSV* model by customizing the values to better reflect personal preferences of the mobile users.

## Chapter IX

# SOCIALLY INFLUENCED BARTERING

In previous Chapter VIII, we have described a set of extensions to the bartering collaboration model that take into consideration aspects of good valuations. In this Chapter, we describe another extension to our bartering model that relies on the social role based interactions to improve collaborative exchanges in the environment.

Social interaction motivated by collaboration are an important factor in mobile environments [89, 30, 15]. The long lasting nature of social relationships [11, 25] can be leveraged to provide short term flexibility [81] by allowing nodes to adjust their strict bartering attitude to a more relaxed and kindhearted attitude when they meet a mobile peer who is part to their social group. To expand the limited space of possible deals and exchanges that are considered during bartering process, our framework exploits social relationships between owners of the mobile peer devices. In particular, our framework reflects significance of the social relationships by exhibiting different levels of cooperation during the bartering process. This approach allows nodes to act as philanthropic *Altruists* during the exchanges with members of the node's social circle while still maintaining a strict bartering attitude towards the nodes that do not belong to this social circle.

As previously mentioned in Sections II.B.4 and VIII.A.1, there are a number of methods that can be used to bootstrap the social network information on to the mobile devices [66, 67, 68]. Current well developed web based social networks such as LinkedIn.com [1], orkut.com [2], mySpaces.com [3] and FaceBook.com [4] can be used to extract the needed social information. Our framework does not provide tools for this extraction and assume that the mobile nodes are already bootstrapped with the needed data.

## IX.A Bartering Environments without Social Networks

Let's first consider homogeneous bartering environments where the nodes do not acknowledge social relationships during the bartering collaborations. In fact, we can look at this environment as a baseline (for the future socially aware interactions) since this environment experiences only simple non-role based interactions. In essence, every node in this environment treats all of its mobile peers as "strangers". In these environments, the nodes are focused on the immediate goal of an "even" exchange and treat all interactions as anonymous transactions. To measure the productivity of such network, we examine the trends in levels of the *unfulfilled wishes* in the network population. As we have previously described in section VI.B, the time based evaluation of the levels of unfulfilled wishes provides a beneficial perspective into the swarm-like dynamics of the peer-to-peer interactions. Figure IX.1 show the results of the time-based study of this homogeneous network configuration. As the time progresses, the bartering mobile nodes eventually reach high levels of fulfillment.

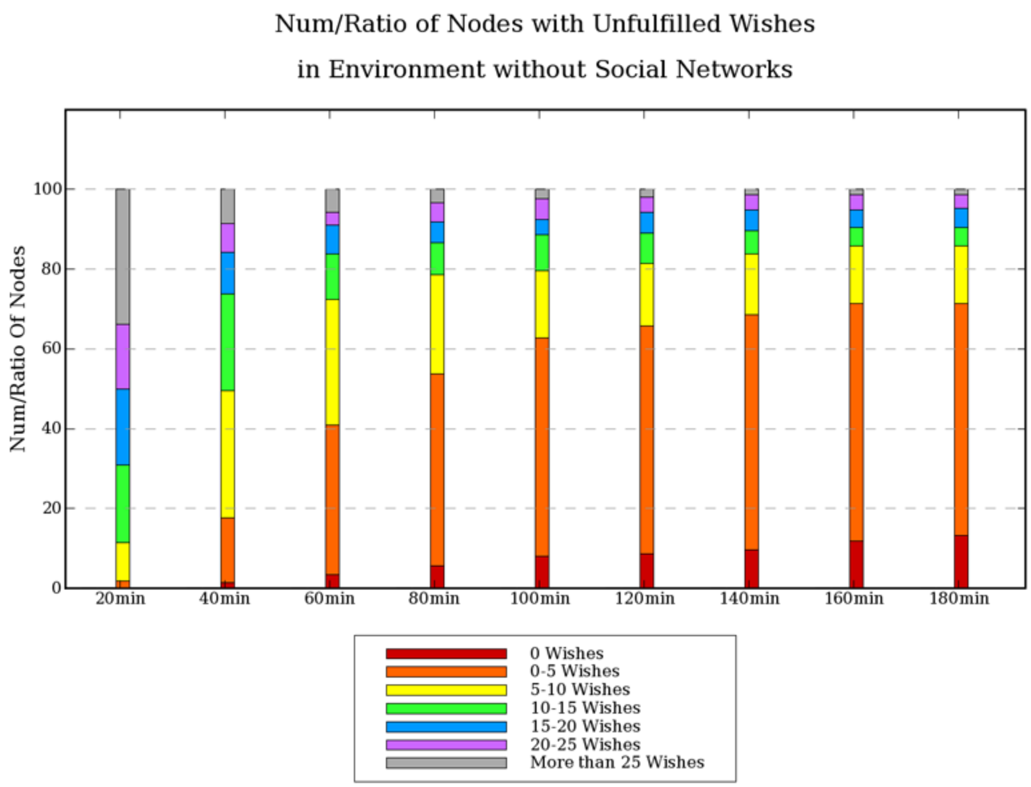


Figure IX.1: Number/Ratio of Nodes with Unfulfilled Wishes in Environment without Social Networks

## IX.B Performance of Bartering Environments Employing Social Networks

Section IX.A provides a baseline for the environment that employs anonymous collaborations. In this section, we contrast this baseline performance with socially sensitive collaborative interactions. In particular, we look at the two network configurations.

The first environment configurations is populated with two distinct social groups. The first social group is a larger group comprised of 28 nodes which is 40% of network population. The second group is a significantly smaller group and it is composed of just 7 nodes which is 10% of network population. In addition to these two groups, the network contains a set of nodes that do not identify with any other node in the environment. These unassociated nodes represent the largest part of the network environment of 35 nodes, which is 50% of network population. This population configuration allows us to look at the dynamics between the three distinct groups that have various level of support from their peers.

The second environment configuration is composed of the four distinct social groups of nodes and a set of unassociated nodes. The first groups is the largest group in the network. It is comprised of 21 nodes, which is 30% of network population. The other three groups are of an equal size of 14 nodes each, which is equivalent to 20% of the network population. Finally, the unassociated nodes are represented by 7 nodes, which is 10% of the network population. This network configuration is reflective of environments that have one relatively strong majority and a set of similar sized minorities. The unassociated nodes reflect the presence of transient "strangers" in the network.

### IX.B.1 Performance of Environments with Population Subdivided into Two Social Groups

Lets first examine the performance of the network with two distinct social groups and a large set of unassociated nodes. Figure IX.2 shows the number of nodes and their levels of success, we rely on these levels of unfulfilled wishes as a way to gauge the productivity of the groups. In addition, Figure IX.3 shows the normalized productivity ratios for each of the groups. From both of these Figures, it is clear that the larger social group shows a greater level of productivity than the other subgroups of the population. Also, the smaller network shows comparatively low level of productivity. As expected, the unassociated nodes are not as effective as the nodes in the large network group. However, interestingly, the unassociated nodes that represent the largest part of the network population show better performance than the small social group. This is performance can be explained by the environment partitioning that occurs through the time of the simulation. This partitioning is evident when evaluating transaction counts.

Figures IX.4 and IX.5 provide an insight into the interactions between the subgroups by showing the transaction rates of each of the subgroups. Essentially, the larger social group reaches a high fulfillment level relatively quickly. And as this process takes place, the nodes in this group have less motivation to interact with the another subparts of the population. Thus, these behavioral dynamics isolate the smaller social subgroup and the unassisted nodes. Thus, these nodes are sectioned off into a separate small sub-environment. Furthermore, the limited size of this remaining environment that is not part of the larger social group impedes on the interactions dynamics of the residual part of the network population. Essentially, the residual sub-environment contains only a few bartering partners that are interested in the collaborative exchanges.

Clearly, the social group that is the strong majority of the network is the most successful sub-part of the population. The altruistic interactions between the nodes belonging to the same large social circle allow this subgroup to achieve the highest levels of satisfaction since, the digital goods circulating within this subgroup are freely shared. This larger group, in essence, forms its own micro-environment and isolates the rest of the population since the nodes of this subgroup become less interested in collaborations.

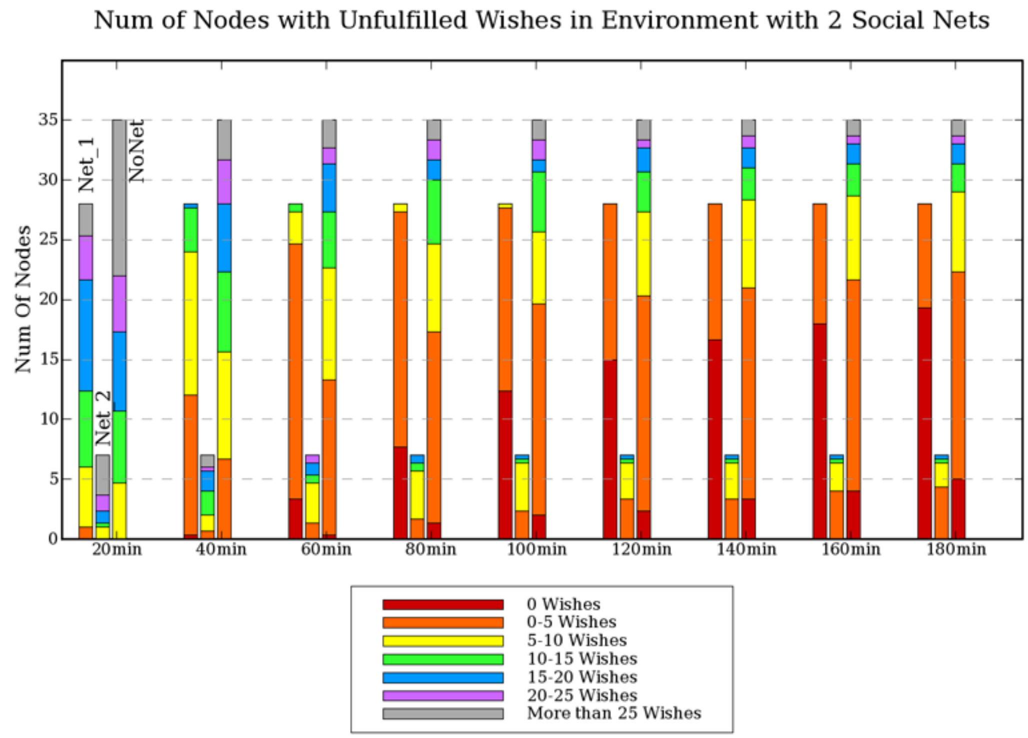


Figure IX.2: Number of Nodes with Unfulfilled Wishes in Environment with 2 Social Nets

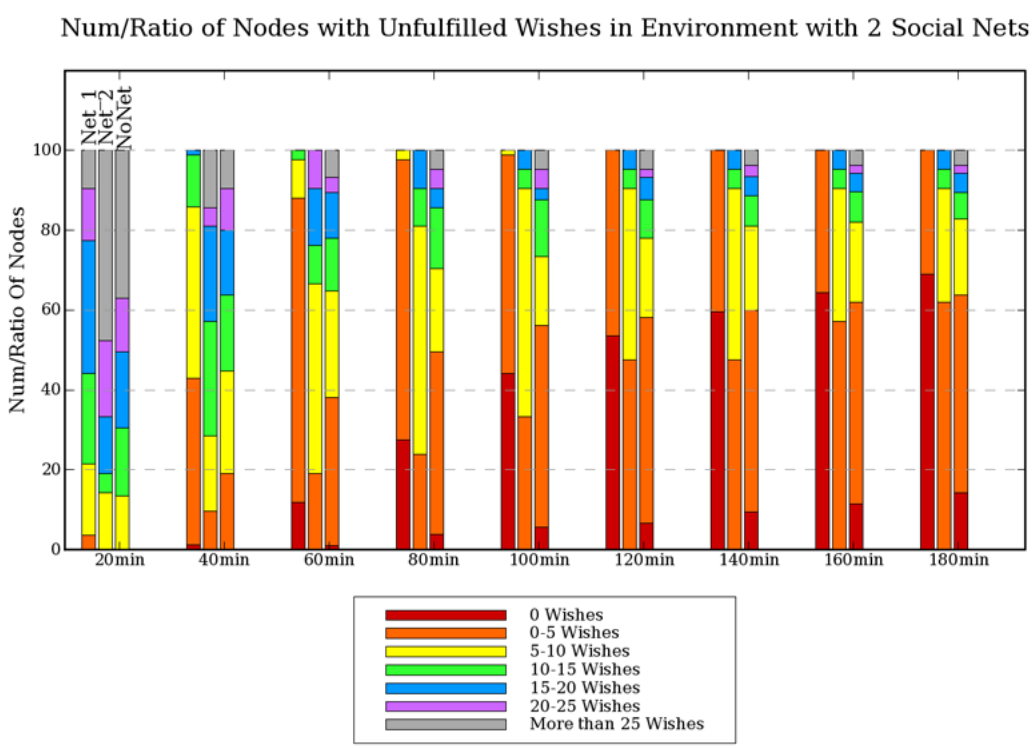


Figure IX.3: Number/Ratio of Nodes with Unfulfilled Wishes in Environment with 2 Social Nets

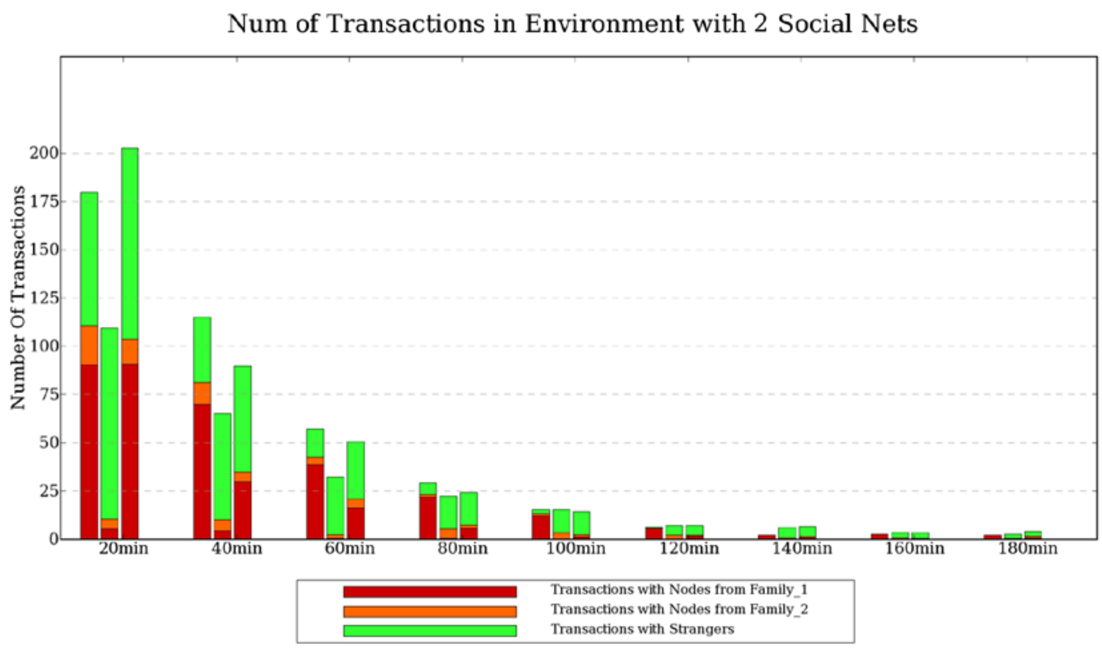


Figure IX.4: Number of Transactions in Environment with 2 Social Nets

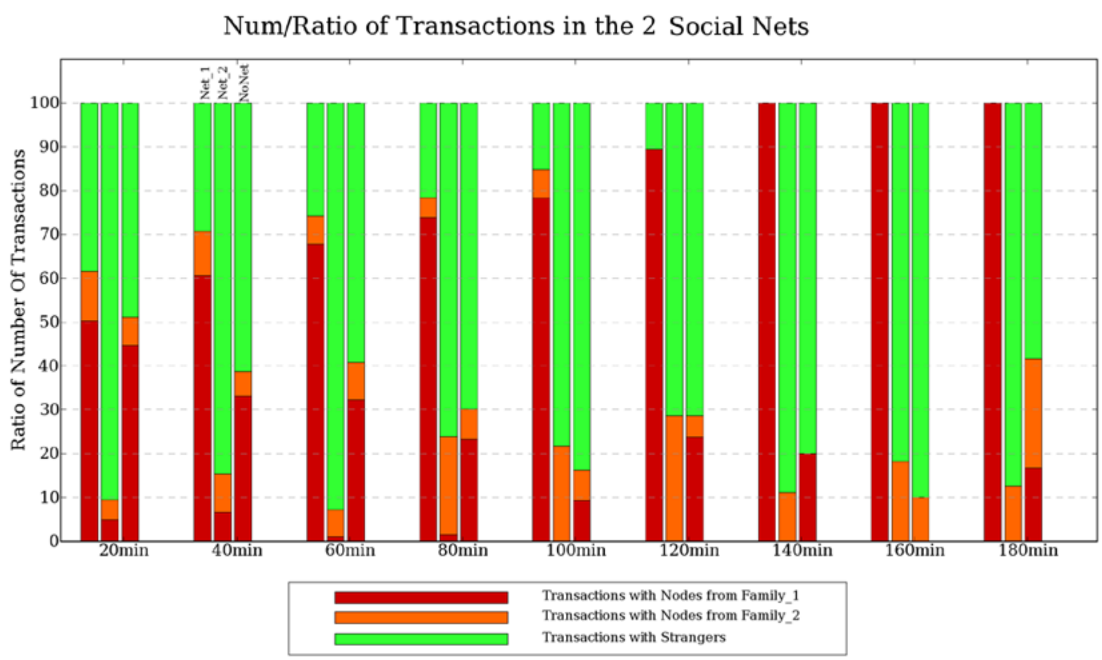


Figure IX.5: Number/Ratio of Transactions in Environment with 2 Social Nets

### IX.B.2 Performance of Environments with Population Subdivided into Four Social Groups

Now, let's examine the performance of the second network configuration where the population is comprised of four distinct social groups and a small set of unassociated nodes. Figure IX.6 and IX.7 show the productivity levels of nodes in this network configuration. While, Figure IX.8 and IX.9 show the transaction rates between these groups. Note that Figure IX.7 and IX.9 provide the normalized levels of data presented in Figure IX.6 and IX.8.

As in the first network configuration, the larger social group of this network is the most productive subpart of the population. This group relatively quickly distinguishes itself by achieving a high level of fulfillment. Looking at the transaction rates from Figure IX.8 and IX.9, we can see that this group conducts a large portion of it's transactions with it's social circle. Also, these graphs show that the this larger group is, by far, the busiest transacting group that interacts with the rest of the environment, particularly at the early parts of the simulation. As the simulation progresses, we observe the overall drop in interaction rates. This can be attributed to the high levels of fulfillment of the larger subgroup. Essentially, the larger subgroup looses interest in the collaborative process and thus isolates itself from the rest of the network population.

Now, let's examine the performance of the three smaller social groups which are each comprised of only 14 nodes. These groups resemble small social cliques of peers [68]. Initially, these nodes are not as productive as the nodes from large social group. However, these nodes become more productive at the later time of the simulation. The fulfillment rate of each of these three groups are relatively lower compared to the larger group. When comparing these smaller groups between themselves, one of the groups (represented by the third column in the graphs) has a slight edge compared to the other two. This edge can be explained by the slightly higher transaction count in the early part of the time based simulation.

Finally, let's consider the performance of unassociated nodes. These nodes are a very small minority and are meant to reflect the presence of the transient "strangers" that are not associated with any of the node in the environment. These nodes are clearly the least productive nodes in this network configuration. They display the lowest levels of fulfillment throughout the time based simulation. As the large social group and the three small group progress through their *iWant* Lists, they become less and less interested in interactions with external nodes. The unassociated nodes are not able to gain enough momentum during the time based simulation to join the fast pace of interactions between the nodes in the social groups. Figure IX.8 and IX.9 show low transaction rates of the unassociated nodes through out the simulation. In very time slice, the

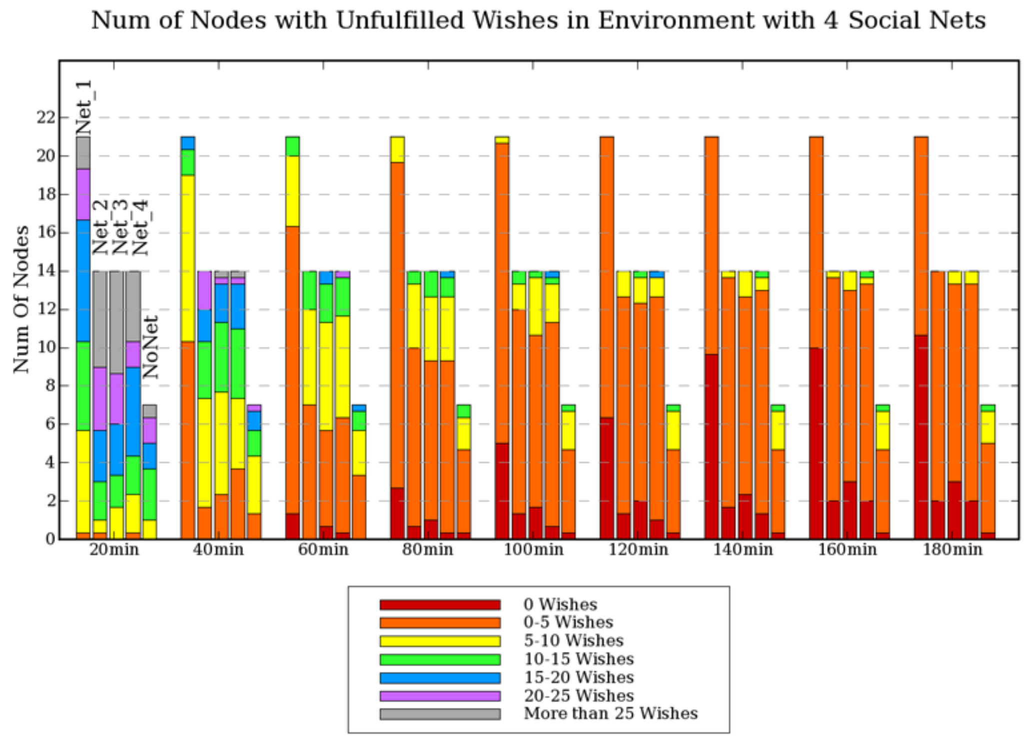


Figure IX.6: Number of Nodes with Unfulfilled Wishes in Environment with 4 Social Nets

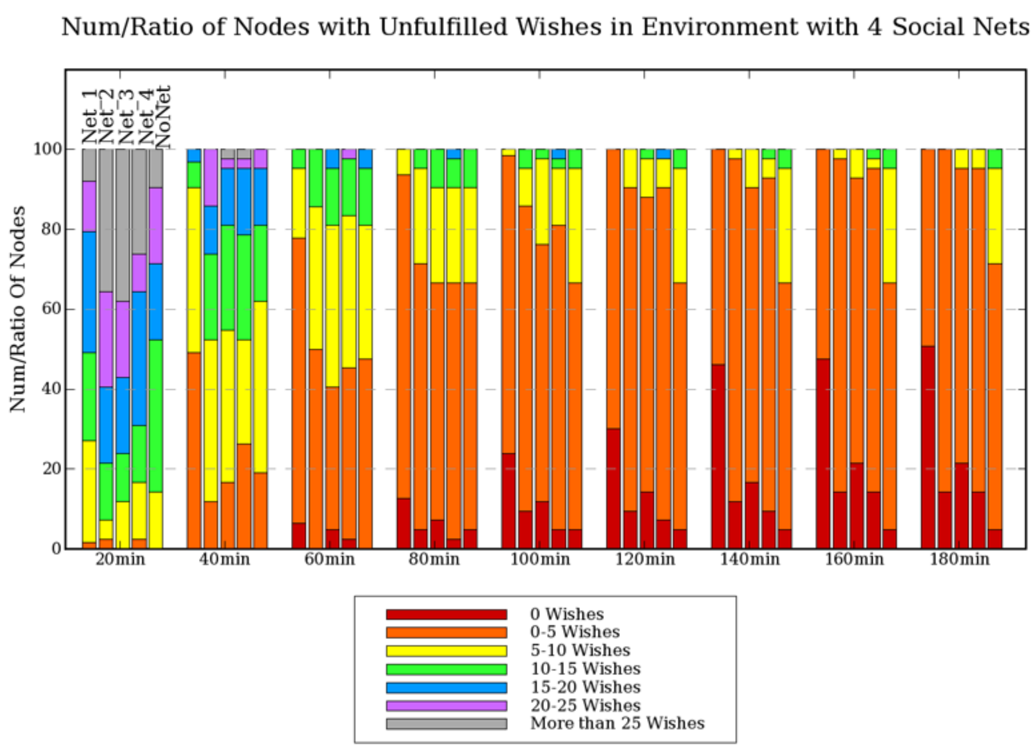


Figure IX.7: Number/Ratio of Nodes with Unfulfilled Wishes in Environment with 4 Social Nets

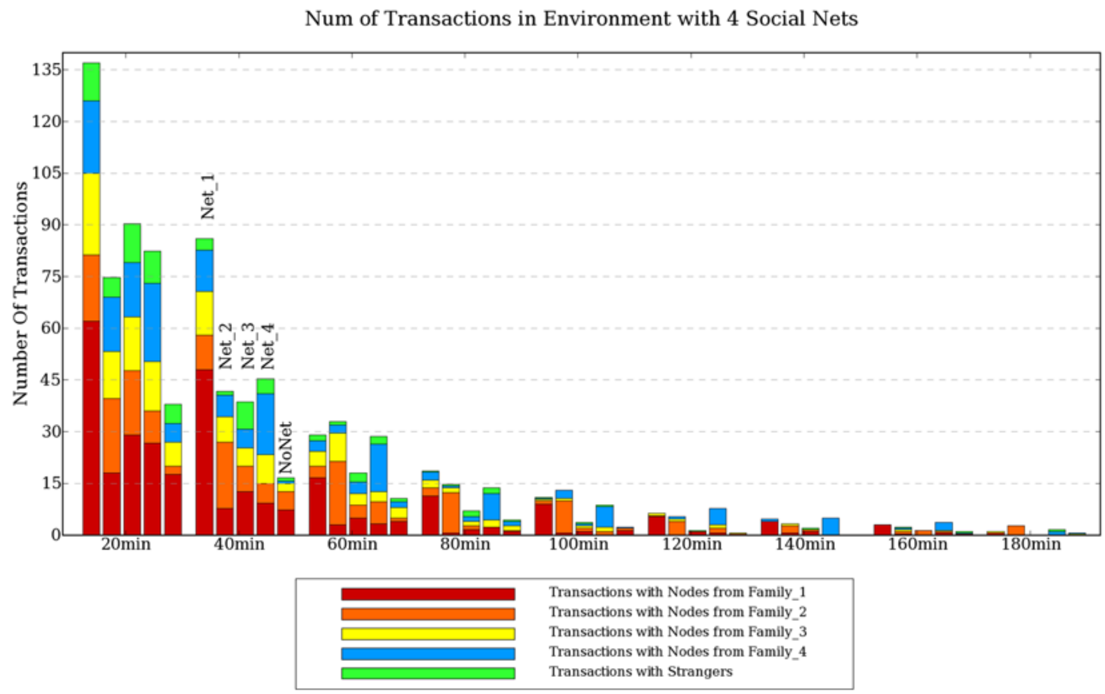


Figure IX.8: Number of Transactions in Environment with 4 Social Nets

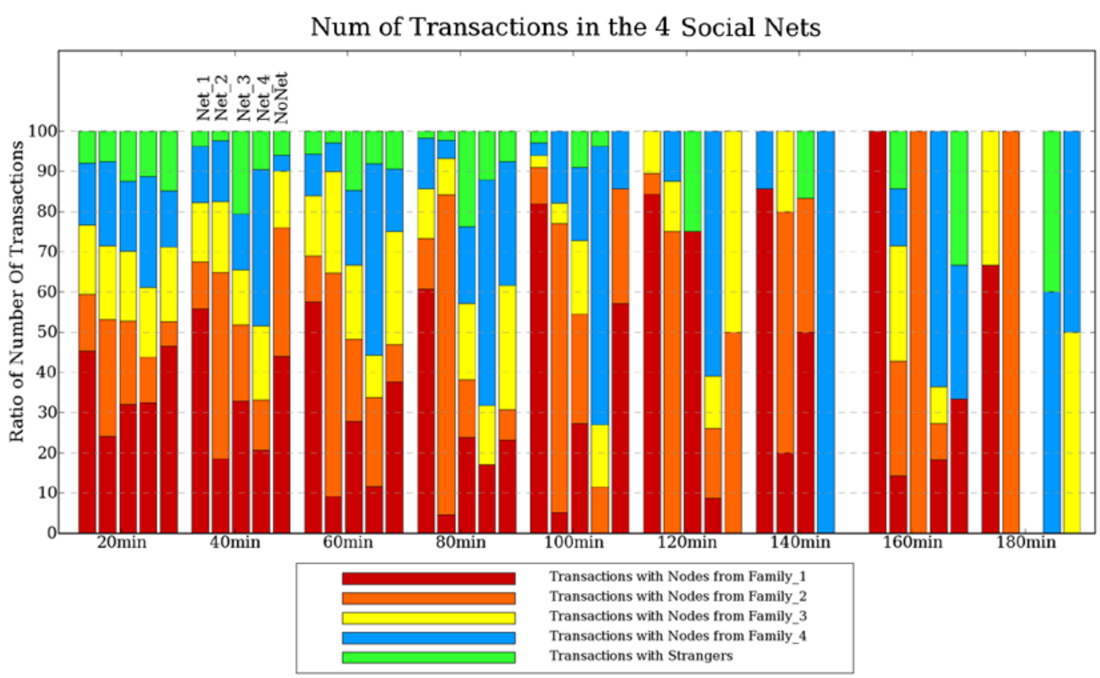


Figure IX.9: Number/Ratio of Transactions in Environment with 4 Social Nets unassociated nodes are experiencing the lowest transaction rates. When comparing performance of these unassociated node with the performance of the baseline network presented in Section IX.A, it is clear that these nodes are unproductive. In essence, these nodes are sidelined during the collaboration process.

To summaries, the second network configuration showed the expected levels of productivity for each of the subparts of the population. Clearly, the larger the social group is the fastest in achieving high fulfillment rates. The unassociated nodes clearly suffer in this population configuration, since they are unable to participate in the fast pace collaborative exchanges between the socially connected nodes.

## IX.C Summary and Discussion

In this Chapter, we have described an extension to our bartering communication model that relies on the social role based interactions. We have analyzed this extension by evaluating the swarm-like dynamics of the inter-group interactions. The size of the social group has a major impact on the productivity of the group and the network as a whole. The large social groups are displaying high levels of wish fulfillment while the unassisted nodes are more likely to be sidelined particularly when they are in the network where most of the nodes are associated with one of the social groups in the environment.

## Chapter X

# CONCLUSION

This chapter summarize the research work and the research contributions presented in this dissertation. This chapter also presents possible future research directions in the relevant research domains.

## X.A Research Summary

Promoting collaborative interaction in mobile pervasive environments is an important aspect that lies at the forefront of the research work presented in this dissertation. Our work focuses of development of collaborative bartering methodology that promotes and encourages cooperative interactions and exchanges of digital goods, services and content. This dissertation presents the first comprehensive research work that employs and models an opportunistic bartering-based collaborative methodology in the context of dynamic mobile peer-to-peer pervasive environments which frequently lack a centralized coordination authority that is traditionally present in the conventional well, connected more and stable computing environments. This bartering-based collaborative methodology is well suited for such environments since it does not require external coordination and extensive transaction-related management that could be associated with negotiation and exchange of digital content during serendipitous encounters that are innate to mobile pervasive environments.

In addition, this dissertation also presents a framework that provides mechanisms for negotiation and bartering exchange that are necessary for opportunistic peer-to-peer interactions in dynamic mobile environments. The presented framework provides mobile nodes with an interaction protocol and a set of cooperation strategies that can be employed by these nodes during the interaction process. These strategies are represented through a set of policies that reflect the collaborative interaction attitude intended to be taken by the nodes. In particular, this dissertation focuses on comparison of the two conventional strategies of “free riding” and “altruistic” behaviors frequently observed in conventional peer-to-peer systems [12, 83, 65] against the novel bartering-based approach derived form the presented opportunistic collaborative methodology. Specifically, the bartering-based approach is represented by the two trading strategies that reflect the key characteristics of the bartering exchange paradigm [49]. The *Weak Double Coincidence of Wants* (*WDCoW*) strategy represents a “good neighbor” approach were nodes are looking for a presence of a minimum levels of reciprocity during the collaborative exchanges. While, the *Double Coincidence of Wants* (*DCoW*) strategy represents a stricter approach that insists on “even” levels of reciprocation during the trading exchanges.

Furthermore, this dissertation offers an in-depth study of the effectiveness, the communication and co-operation overhead costs of the above described collaborative approaches. In particular, a set of simulation results that model peer-to-peer interactions in homogeneous environments are described in Chapter V. These results serve as a baseline for the study of the inter-strategy interactions that occur in the heterogeneous networks populated by nodes with varying levels of cooperation-relevant attitudes and approaches. Moreover, a time-based simulation study offers detailed performance finding and analysis of swarm-like dynamics that occur in heterogeneous populations of nodes with varying cooperative attitudes. Furthermore, this dissertation also presents results of collaborative inter-strategy interactions in heterogeneous networks where the population of the network has a strong domination by one of the considered strategies. These results further highlight the strength and versatility of the presented bartering collaboration methodology when it is applied in the context of serendipitous encounters of small personal mobile devices.

In addition, this dissertation offers a set of extensions of the original concept of bartering based interactions by exploring the concept of valuation sensitive bartering. Essentially, the framework is extended with a set of valuation models that derive value of the digital goods and content and incorporate these valuations into the bartering process. In particular, the framework employs the *Demand Sensitive Valuation* model, the *Demand and Supply Sensitive Valuation* model and finally the *Personalized Valuation* model. An in-depth study of the effects of these valuation models on the collaborative effectiveness of the bartering networks is presented in Chapter VIII. Moreover, the bartering framework further extends the concept of valuation sensitive interactions by employing an investment-based trading approach which facilitates nodes with mechanisms for acquiring valuable digital content with an intent to later trade it away for originally desired set of goods. Furthermore, this dissertation explores a concept of social role based interactions in the context of dynamic mobile peer-to-peer pervasive environments by adding an extension that correlates collaboration levels to the levels of social relationships between the collaborating nodes.

Taken as a whole, this dissertation delivers a comprehensive view of a novel bartering-based collaborative method which is designed for the context of serendipitous opportunistic collaborative interactions in dynamic mobile peer-to-peer pervasive environments.

## X.B Future Research Directions

The following categories represent potential directions for future research work.

- Impact of digital rights management on collaborative exchanges and interactions.
- Incorporation of node movement into cooperation attitude selection.
- Effects of uncertainty of quality of digital goods and content on collaborative exchanges.

### X.B.1 Impact of Digital Rights Management on Collaborative Interactions

Collaborative interactions can be effected by the constraints related to the transfer of ownership rights of the digital content. There is a wide range of Digital Rights Management mechanisms and approaches that can guard the process of transferring of the digital content [43, 63, 91]. One such approach involves complete denial of transfer. This approach represent one of the extremes in the wide range of available DRM mechanisms [59, 40]. Another strict DRM approach requires total surrender of ownership rights during the transfer of the content. This strict approach is commonly applied to commercially generated copyrighted digital content such as MP3 files and other entertainment and gaming digital content. In contrast, the approach that provides nodes with total freedom of replication allows these nodes to freely exchange and replicate digital goods and content. This approach is commonly used in cases of user generated content or advertisement related content that is designed to reach a larger audience of consumers. In addition to these two extreme approaches, there are a number of methods that attempt to strike balance between complete freedom of replication and totally ridged restrictions. One such example is the “squirting” approach which is employed by the Zune MP3 player [10]. The “squirting” approach limits the replication of content to three copies thus allowing nodes to have some flexibility when they attempt to replicate and distribute the content.

This wide range of DRM approaches can have an impact on the collaborative interactions between mobile devices. Ridged approaches that restrict transfer of rights can dramatically hinder the interactive process in any collaborative environment. On the other hand, the infinite duplication and uncontrolled replication of content can result in degradation in quality of digital content available in the network. These issues are not thoroughly addressed in this dissertation since we make an explicit assumption that goods are transferable and are node duplicated during the transaction. Researching impact of DRM approach can potentially be an interesting extension to the work presented in this dissertation.

### X.B.2 Incorporating Expected Future Node's Movements into The Bartering Process

The research work presented in this dissertation does not consider proactive context awareness of the collaborative environments. Adding an extension that considers the node's upcoming route would be an interesting expansion of this research work. For example, a mobile node can identify the fact that it is moving towards a more populated area and thus adjust it's collaborative attitudes and policies accordingly. The node could also exploit its awareness of the desires, interests and predictable predispositions of the upcoming network population. Thus the node could adjust its current trading patters and attitudes by taking this upcoming context into consideration. This proactive context awareness can have an impact on the negotiating strategy selection and thus have an impact on the overall collaboration productivity.

### X.B.3 Uncertainty of The Quality of Digital Goods and Content

The research work presented in this dissertation does not take into account aspects of the quality of digital content. This work makes an assumption that the content is well described and properly formatted. Though this assumption is not unreasonable, removal of this assumption presents an interesting extension. Development of a set of mechanisms that guard nodes from exchanging incomplete or incorrectly marked digital goods and content would have a positive impact on the collaborative process and would allow nodes to make more informed decisions when establishing a collaborative exchange.

## X.C Conclusion

This research work presented in this dissertation delivers a novel opportunistic barter-based collaboration method for the context of serendipitous encounters in dynamic mobile pervasive environments. This dissertation presents a framework employing bartering methods and also delivers an in-depth study of effects of the collaborative attitudes in dynamic mobile environments.

# BIBLIOGRAPHY


[1] LinkedIn, World Wide Web, http://www.LinkedIn.com. II.B.4, IX

[2] Orkut, World Wide Web, http://www.orkut.com. II.B.4, VIII.A.1, IX

[3] mySpace, World Wide Web, http://www.myspace.com. II.B.4, VIII.A.1, IX

[4] FaceBook, World Wide Web, http://www.facebook.com. II.B.4, VIII.A.1, IX

[5] Bluetooth White Paper, WWW, http://www.bluetooth.com/developer/whitepaper/. III.B

[6] IEEE 802.11TM WIRELESS LOCAL AREA NETWORKS, WWW, http://grouper.ieee.org/groups/802/11/. III.B, IV.E

[7] Napster, World Wide Web, http://www.napster.com. III.C

[8] Gnutella, World Wide Web, http://www.gnutella.com. III.C, IV.C.1

[9] KaZaA, World Wide Web, http://www.kazaa.com. III.C, IV.C.1

[10] Zune, World Wide Web, http://www.zune.net. X.B.1

[11] *The familiar stranger: anxiety, comfort, and play in public places*, New York, NY, USA, 2004. ACM Press. II, IX

[12] E. Adar and B. A. Huberman. Free riding on gnutella. *First Monday*, Sept. 2000. II.B.1, III.C, IV.A, IV.C.1, VIII, X.A

[13] K. G. Anagnostakis and M. B. Greenwald. Exchange-based incentive mechanisms for peer-to-peer file sharing. In *ICDCS '04: Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS'04)*, pages 524–533, Washington, DC, USA, 2004. IEEE Computer Society. III.C

[14] S. Avancha, P. D'Souza, F. Perich, A. Joshi, and Y. Yesha. P2P M-Commerce in Pervasive Environments. *ACM SIGecom Exchanges*, 3(4):1–9, January 2003. II

[15] A.-L. Barabasi. *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*. Plume Books, April 2003. IV.E.1, VIII.A.1, IX

[16] S. J. Barnes. Wireless digital advertising: Nature and implications. In *International Journal of Advertising*, pages Vol. 21, 399–420, 2002. II.B.2

[17] D. Billsus and M. Pazzani. Revising user profiles: The search for interesting web sites. In *Proceedings of the 3nd International Workshop on Multistrategy Learning*, pages 239–250, Harper's Ferry, W.Va., May 1996. VIII.A.1

[18] D. Billsus and M. J. Pazzani. User modeling for adaptive news access. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research*, 10(2-3):147–180, 2000. DOI 10.1023/A:1026501525781. VIII.A.1

[19] B. H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970. IV.B.1

[20] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Santa Fe Institute Studies in the Sciences of Complexity. Oxford University Press, 1999. VI.B

[21] D. Chalmers. *Contextual Mediation to Support Ubiquitous Computing*. PhD thesis, Imperial College London, 2002. III.B.2

[22] D. Chalmers, N. Dulay, and M. Sloman. A framework for contextual mediation in mobile and ubiquitous computing applied to the context-aware adaptation of maps. *Personal Ubiquitous Comput.*, 8(1):1–18, 2004. III.B.2

[23] L. Chen and K. Sycara. WebMate: A personal agent for browsing and searching. In K. P. Sycara and M. Wooldridge, editors, *Proceedings of the 2nd International Conference on Autonomous Agents (Agents'98)*, pages 132–139, New York, 9–13, 1998. ACM Press. IV.B.1

[24] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. The origin of power laws in internet topologies revisited. In *Proceedings of IEEE INFOCOM*, June 2002. IV.E.1, VIII.A.1

[25] C.-M. Chiu, M.-H. Hsu, and E. T. G. Wang. Understanding knowledge sharing in virtual communities: an integration of social capital and social cognitive theories. *Decis. Support Syst.*, 42(3):1872–1888, 2006. IX

[26] S. H. Clearwater, editor. *Market-based Control: A Paradigm for Distributed Resource Allocation*. World Scientific Publishing Co. Pte. Ltd., London, 1996. II.B.1

[27] B. Cohen. Incentives build robustness in BitTorrent. In *Workshop on Economics of Peer-to-Peer Systems, Berkeley, CA, USA*, June 2003. III.C

[28] A. Dey, G. D. Abowd, and D. Salber. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction*, 16:97–166, 2001. II.B.1

[29] P. Domel, A. Lingnau, and O. Drobnik. Mobile agent interaction in heterogeneous environments. In *Mobile Agents*, pages 136–148, 1997. VII

[30] T. Erickson and W. A. Kellogg. Social translucence: an approach to designing systems that support social processes. *ACM Trans. Comput.-Hum. Interact.*, 7(1):59–83, March 2000. IX

[31] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proceedings of ACM SIGCOMM*, September 1999. IV.E.1, VIII.A.1

[32] S. Fickas, G. Kortuem, J. Schneider, Z. Segall, and J. Suruda. When cyborgs meet: Building communities of cooperating wearable agents, October 1999. III.B.1

[33] G. Wu and C. W. Chu and K. Wine, J. Evans and R. Frenkiel. Winmac: A novel transmission protocol for infostations., 1999. II, III.B.1, IV.E

[34] D. Garlan, D. P. Siewiorek, A. Smailagic, and P. Steenkiste. Project aura: Toward distraction-free pervasive computing. *IEEE PERVASIVE COMPUTING*, 1(2):22–31, 2002. II.B.1

[35] A. Gershman. Ubiquitous commerce - always on, always aware, always pro-active. In *SAINT*, pages 37–38, 2002. I.A

[36] N. S. Glance and T. Hogg. Computational social dilemmas. Technical report, Xerox PARC, 1995. II.B.1

[37] D. J. Goodman, J. Borras, N. B. Mandayam, and R. Yates. INFOSTATIONS: a new system model for data and messaging services. In *IEEE VTC*, pages 969–973, Phoenix, Arizona, May 1997. IV.E

[38] R. Grinter and M. Eldridge. Wan2tlk?: everyday text messaging. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 441–448, New York, NY, USA, 2003. ACM Press. I.A

[39] T. Groves and J. Ledyard. Optimal allocation of public goods: A solution to the 'free rider' problem. *Econometrica*, 45:783–809, 1977. II.B.1

[40] C. A. Gunter, S. Weeks, and A. K. Wright. Models and languages for digital rights. In *HICSS*, 2001. X.B.1

[41] R. E. Hall. *Digital Dealing: How eMarkets Are Transforming the Economy*. W. W. Norton, New York, 2002. I.A, VIII.B

[42] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. *Wireless Networks*, 8(2/3):187–197, 2002. II.B.1

[43] F. Hartung and F. Ramme. Digital rights management and watermarking of multimedia content for m-commerce applications. *Communications Magazine, IEEE*, 38(11):78–84, 2000. X.B.1

[44] K. A. Hawick and H. A. James. Middleware for context sensitive mobile applications. In *ACSW Frontiers '03: Proceedings of the Australasian information security workshop conference on ACSW frontiers 2003*, pages 133–141, Darlinghurst, Australia, Australia, 2003. Australian Computer Society, Inc. II.B.3

[45] A. Heinemann and T. Straub. An anonymous bonus point system for mobile commerce based on word-of-mouth recommendation. In *Proceedings of the 19th Annual ACM Symposium on Applied Computing (ACM SAC 2004) Nicosia, Cyprus*, 2004. III.D

[46] J. Hirshleifer. The Private and Social Value of Information and the Reward to Inventive Activity. *American Economic Review*, 61(4):561–574, 1971. VIII.A.1

[47] T. Imielinski and B. R. Badrinath. Mobile wireless computing: Challenges in data management. *Communications of the ACM*, 37(10):18–28, 1994. II.B.2

[48] A. Z. Industries. MojoNation. http://www.mojonation.com/. III.C

[49] W. S. Jevons. *Money and the Mechanism of Exchange*. D. Appleton and Company, 1875. II.D, III.F, X.A

[50] K. Keegan and G. O'Hare. Easishop: Context sensitive shopping for the mobile user through mobile agent technology, 2002. III.D

[51] S. Keegan and G. M. P. O'Hare. Easishop: Enabling ucommerce through intelligent mobile agent technologies. In *Proceedings of 5th International Workshop, Mobile Agents for Telecommunication Applications, MATA 2003, Marakech, Morocco, October 8-10, 2003, Proceedings*, pages 200–209, 2003. III.D

[52] J. Kennedy and R. Mendes. Population structure and particle swarm performance. In D. B. Fogel, M. A. El-Sharkawi, X. Yao, G. Greenwood, H. Iba, P. Marrow, and M. Shackleton, editors, *Proceedings of the 2002 Congress on Evolutionary Computation CEC2002*, pages 1671–1676. IEEE Press, 2002. VI.B

[53] J. Kjeldskov. "just-in-place" information for mobile device interfaces. In *Mobile HCI '02: Proceedings of the 4th International Symposium on Mobile Human-Computer Interaction*, pages 271–275, London, UK, 2002. Springer-Verlag. II.B.3

[54] S. B. Kodeswaran, O. Ratsimor, A. Joshi, T. Finin, and Y. Yesha. Using peer-to-peer data routing for infrastructure-based wireless networks. In *IEEE International Conference on Pervasive Computing and Communications (PerCom 2003)*, Fort Worth, TX, USA, March 2003. II, IV.E

[55] G. Kortuem. *A methodology and software platform for building wearable communities*. PhD thesis, University of Oregon, 2002. Adviser-Zary Segall. III.B.1

[56] G. Kortuem. Proem: a middleware platform for mobile peer-to-peer computing. *SIGMOBILE Mob. Comput. Commun. Rev.*, 6(4):62–64, 2002. III.B.1

[57] G. Kortuem, J. Schneider, D. Preuitt, T. G. Thompson, S. Fickas, and Z. Segall. When peer-to-peer comes face-to-face: Collaborative peer-to-peer computing in mobile ad hoc networks. In *Proceedings 2001 International Conference on Peer-to-Peer Computing*, Aug, 2001. III.B.1

[58] G. Kortuem, Z. Segall, and T. C. Thompson. Close encounters: Supporting mobile collaboration through interchange of user profiles. *Lecture Notes in Computer Science*, 1707:171–182, 1999. IV.B.1

[59] S. Kwok, S. Cheung, K. Wong, K. Tsang, S. Lui, and K. Tam. Integration of digital rights management into the Internet Open Trading Protocol. *Decision Support Systems*, 34(4):413–425, 2003. X.B.1

[60] O. B. Kwon and N. Sadeh. Applying case-based reasoning and multi-agent intelligent system to context-aware comparative shopping. *Decis. Support Syst.*, 37(2):199–213, 2004. II

[61] H. Lei, D. Chakraborty, H. Chang, M. J. Dikun, T. Heath, J. S. Li, N. Nayak, and Y. Patnaik. Contextual collaboration: Platform and applications. In *SCC '04: Proceedings of the 2004 IEEE International Conference on Services Computing*, pages 197–206, Washington, DC, USA, 2004. IEEE Computer Society. II

[62] T.-P. Liang and J.-S. Huang. A framework for applying intelligent agents to support electronic trading. *Decis. Support Syst.*, 28(4):305–317, 2000. II, III.F

[63] Q. Liu, R. Safavi-Naini, and N. P. Sheppard. Digital rights management for content distribution. In C. Johnson, P. Montague, and C. Steketee, editors, *First Australasian Information Security Workshop (AISW2003)*, volume 21 of *CRPIT*, pages 49–58, Adelaide, Australia, 2003. ACS. X.B.1

[64] S. W. Loke and A. Zaslavsky. Integrated ambient services as enhancement to physical marketplaces. In *HICSS '04: Proceedings of the Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04) - Track 9*, page 90284.3, Washington, DC, USA, 2004. IEEE Computer Society. II, III.B.2

[65] S. K. Mario. Incentives for combatting freeriding on p2p networks., 2003. II.B.1, III.C, X.A

[66] P. Mika. Bootstrapping the FOAF-web: An Experiment in Social Network Mining. In *1st Workshop on Friend of a Friend, Social Networking and the Semantic Web*, Galway, Ireland, September 2004. IX

[67] P. Mika. Social Networks and the Semantic Web: An Experiment in Online Social Network Analysis. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*, Beijing, China, September 2004. IX

[68] P. Mika. Flink: Semantic Web Technology for the Extraction and Analysis of Social Networks. *Journal of Web Semantics*, 3(2), 2005. IX, IX.B.2

[69] R. Murphy, C. Lisetti, L. Irish, R. Tardif, and A. Gage. Emotion-based control of cooperating heterogeneous mobile robots, 2001. VII

[70] N. Ntarmos and P. Triantafillou. Aesop: Altruism-endowed self organizing peers. In *3rd Int. Workshop on Databases, Information Systems and Peer-to-Peer Computing*, 2004. IV.A

[71] A. Odlyzko. The case against micropayments. In *Financial Cryptography 2003*, Minneapolis, MN, Feb. 2003. Springer Verlag. Lecture Notes in Computer Science. II.C, III.F

[72] S. Ontanon and E. Plaza. Cooperative case bartering for case-based reasoning agents. In *CCIA '02: Proceedings of the 5th Catalonian Conference on AI*, pages 294–308, London, UK, 2002. Springer-Verlag. III.F

[73] I. Papaefstahiou and C. Manifavas. Evaluation of micropayment transaction costs. Journal of Electronic Commerce Research 5 (2). II.C

[74] J. Parker, A. Patwardhan, F. Perich, A. Joshi, and T. Finin. Trust in Pervasive Computing. *Mobile Middleware*, pages 473–496, February 2006. II.B.4

[75] A. Patwardhan. *Trustworthy data management in pervasive environments*. PhD thesis, University of Maryland, Baltimore County (UMBC), 2007. Adviser-Anupam Joshi. I.A, II.B.4

[76] A. Patwardhan, A. Joshi, T. Finin, and Y. Yesha. A Data Intensive Reputation Management Scheme for Vehicular Ad Hoc Networks. In *Proceedings of the Second International Workshop on Vehicle-to-Vehicle Communications*. IEEE, July 2006. II.B.4

[77] F. Perich. *On peer-to-peer data management in pervasive computing environments*. PhD thesis, University of Maryland, Baltimore County (UMBC), 2004. Adviser-Anupam Joshi. I.A, II, III.B.1

[78] F. Perich, S. Avancha, D. Chakraborty, A. Joshi, and Y. Yesha. Profile driven data management for pervasive environments. In *3th International Conference on Database and Expert Systems Applications (DEXA 2002)*, Aix en Provence, France, September 2002. II.B.2, III.B.1, IV.B.1

[79] E. Pitoura and B. Bhargava. A framework for providing consistent and recoverable agent-based access to heterogeneous mobile databases. *ACM SIGMOD Record*, 24(3):44–49, Sep 1995. VII

[80] J. Pouwelse, P. Garbacki, J. W. A. Bakker, J. Yang, A. Iosup, D. Epema, M.Reinders, M. van Steen, and H. Sips. Tribler: A social-based peer-to-peer system. In *5th Int'l Workshop on Peer-to-Peer Systems (IPTPS)*, Feb 2006. IV.A

[81] D. G. Pruitt. *Negotiation Behavior*. Academic Press, 1981. VIII.A.1, IX

[82] D. Quah. Digital goods and the new economy. CEP Discussion Papers dp0563, Centre for Economic Performance, LSE, March 2003. III.E

[83] L. Ramaswamy and L. Liu. Free riding: A new challenge to peer-to-peer file sharing systems, 2003. II.B.1, III.C, X.A

[84] O. Ratsimor, D. Chakraborty, S. Tolia, D. Kushraj, A. Kunjithapatham, G. Gupta, A. Joshi, and T. Finin. Allia: Alliance-based service discovery for ad-hoc environments, 2002. III.D

[85] O. Ratsimor, T. Finin, A. Joshi, and Y. Yesha. encentive: a framework for intelligent marketing in mobile peer-to-peer environments. In *ICEC '03: Proceedings of the 5th international conference on Electronic commerce*, pages 87–94, New York, NY, USA, 2003. ACM Press. II.B.2, III.D

[86] O. Ratsimor, V. Korolev, A. Joshi, and T. Finin. Agents2go: An infrastructure for location-dependent service discovery in the mobile electronic commerce environment. In *ACM Mobile Commerce Workshop held in conjunction with MobiCom 2001.*, Rome, Italy., July 2001. II.B.2, III.D

[87] O. V. Ratsimor, D. Chakraborty, A. Joshi, T. Finin, and Y. Yesha. Service discovery in agent-based in pervasive computing environments. *Journal on Mobile Networking and Applications*, November 2003. Special issue on Mobile and Pervasive Commerce. II, III.D

[88] R.H. Frenkiel, B.R. Badrinath, J. Borras, and R. Yates,. The infostations challenge: Balancing cost and ubiquity in delivering wireless data. *IEEE Personal Communications*, 7(2):pp.66–71, April 2000. III.B.1, IV.E

[89] H. Rheingold. *Smart Mobs: The Next Social Revolution*. Perseus Publishing, 2002. I.A, II.B.1, III.D, IX

[90] M. Rokeach. *The Nature of Human Values*. The Free Press, New York, 1973. VIII.A.1

[91] W. Rosenblatt, S. Mooney, and W. Trippe. *Digital Rights Management: Business and Technology*. John Wiley & Sons, Inc., New York, NY, USA, 2001. III.E.2, X.B.1

[92] N. Sadeh. *Mobile Commerce: New Technologies, Services and Business Models*. Wiley, 2002. II.B.1, III.D

[93] N. Sadeh, T.-C. Chan, L. Van, O. Kwon, and K. Takizawa. Creating an open agent environment for context-aware m-commerce. In Burg, Dale, Finin, Nakashima, Padgham, Sierra, and Willmott, editors, *Agentcities: Challenges in Open Agent Environments*, pages 152–158, LNAI, Springer Verlag, 2003. III.B.2

[94] B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *IEEE Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, US, 1994. II.B.1

[95] M. Sekaran and S. Sen. To help or not to help. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, pages 736–741, Pittsburgh, PA, 1994. IV.A, VIII

[96] S. Sen. Reciprocity: A foundational principle for promoting cooperative behavior among self-interested agents. In V. Lesser, editor, *Proceedings of the First International Conference on Multiagent Systems*, pages 322–329. MIT Press, 1996. II, IV.A, VIII

[97] N. Shadbolt. Ambient intelligence. *IEEE Intelligent Systems*, 18(4):2–3, 2003. II, III.B.2

[98] C. Shapiro and H. R. Varian. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press, 1999. III.E.1, VIII.A.1

[99] J. P. Sousa and D. Garlan. Aura: an Architectural Framework for User Mobility in Ubiquitous Computing Environments. In *Proceeding of the 3rd Working IEEE/IFIP Conference on Software Architecture*, Montreal, August 25-31 2002. II.B.1

[100] L. Srivastava. Mobile phones and the evolution of social behaviour. *Behaviour and Information Technology*, 24:111–129(19), March-April 2005. I.A

[101] N. Szabo. Micropayments and mental transaction costs. II.C

[102] D. Tapscott. *The digital economy: promise and peril in the age of networked intelligence*. McGraw-Hill, New York, NY, USA, 1996. II, VIII.B

[103] D. Tapscott and R. McQueen. *Digital Economy: Competing in the New Networked Economy*. McGraw-Hill Education, Boston, 1995. I.A, VIII.B

[104] D. Touzet, J.-M. Menaud, M. Banâtre, P. Couderc, and F. Weis. SIDE Surfer: a Spontaneous Information Discovery and Exchange System. In *Proceedings of the Second International Workshop on Ubiquitous Computing and Communications (WUCC'2001)*, Barcelona, Spain, September 2001. II.B.1

[105] N. van Someren. The practical problems of implementing micromint. In *FC '01: Proceedings of the 5th International Conference on Financial Cryptography*, pages 41–50, London, UK, 2002. Springer-Verlag. II.C, III.F

[106] U. Varshney. Location management for mobile commerce applications in wireless Internet environment. *ACM Transactions on Internet Technology*, 3(3):236–255, Aug. 2003. III.D

[107] U. Varshney. Vehicular mobile commerce. *Computer*, 37(12):116–118, 2004. III.D

[108] U. Varshney and R. Vetter. Emerging mobile and wireless networks. *Commun. ACM*, 43(6):73–81, 2000. I.A

[109] U. Varshney and R. Vetter. Mobile commerce: Framework, applications and networking support. *Mob. Netw. Appl.*, 7(3):185–198, 2002. III.D

[110] U. Varshney, R. J. Vetter, and R. Kalakota. Mobile Commerce: A New Frontier. *Computer*, pages 32–38, October 2000. II, III.D

[111] V. Vishnumurthy, S. Chandrakumar, and E. G. Sirer. KARMA: A secure economic framework for peer-to-peer resource sharing. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems*, Berkeley, California, USA, June 2003. III.C

[112] M. Weiser. The computer for the 21st century. *Human-computer interaction: toward the year 2000*, pages 933–940, 1995. I

[113] T. White and B. Pagurek. Emergent behaviour and mobile agents. In *Proceedings of the workshop on Mobile Agents in the Context of Competition and Cooperation at Autonomous Agents '99*, 1999. VI.B

[114] T. Yamakami. Impact from mobile spam mail on mobile internet services . In *Parallel and Distributed Processing and Applications, International Symposium, ISPA 2003, Aizu, Japan, July 2-4, 2003, Proceedings*, pages 179–184, 2003. II.B.2

[115] T. Ye, H.-A. Jacobsen, and R. H. Katz. In proceedings of the second international conference on mobile data management (mdm 2001). In *Mobile Computing and Networking*, pages 109–120, 1998. IV.E

[116] E. Yoneki and J. Bacon. An Adaptive Approach to Content-Based Subscription in Mobile Ad Hoc Networks. In *Proc. 2nd IEEE International Conference on Pervasive Computing and Communications, Workshop on Mobile Peer-to-Peer Computing*, pages 92–97, Orlando, USA, March 2004. IV.B.1

[117] X. Zeng, R. Bagrodia, and M. Gerla. GloMoSim: a Library for Parallel Simulation for Large-scale Wireless Networks. In *proc. 12th Workshop on Parallel and Distributed Simulations (PADS '98)*, 1998. IV.E, IV.E.1