Technical Report:	***CSM-405***

PhD in Open Constraint Satisfaction:

SSCM, A Position Statement

By

Timothy Gosling

Contents

1.1 EXECUTIVE SUMMARY 2
1.2 INTRODUCTION 3
1.3 POSITION STATEMENT 4
1.3.1 Motivation and objectives 4
1.3.2 Current state and future work 4
1.3.3 Relationship to other work 5
1.3.3.1 The Model 5
1.3.3.2 Modelled Situation and Scenarios 6
1.3.3.3 Strategies 7
1.3.3.4 Strategy, simulation and optimisation 7
1.4 REFERENCES 9

1.1 Executive Summary

The Simple Supply Chain Model (SSCM) provides a way to fully specify a supply chain by way of its participants starting knowledge and the way in which they may communicate. This work attempts to model certain supply chain scenarios, develop strategies to tackle them and optimise these strategies using evolutionary computation (EC) within a market simulator. This work is important because, with the expansion of the internet and computing power, the drive to move business online is every increasing and the benefits from automating business interaction are potentially considerable. Customers and businesses wishing to lower the costs or improve their profits and having access to many potential partners online could reap great rewards by reaching agreement with those partners through automated negotiation. Success would not only be a matter of making good deals but communicating with many people.

The model the work is based upon is closely related to the original Trading Agent Competition (TAC) game and is in some ways conceptually similar to Distributed Constraint Satisfaction (DCS), Distributed Artificial Intelligence (DAI). The model is less related to complex market architectures like MAGET or more complex planning and scheduling systems. The strategies employed communicate through a form of alternating offers protocol and don't make use of more complex agent communication languages. Negotiation is used for participants to come to agreements with one another so that the work has some relation to much of bargaining and negotiation research. The SSCM tends however to concern a situation in which there is a great deal of uncertainty and no rigid rules of conduct beyond the communications protocol. The strategy optimisation part of the work is conducted within a simple market simulation and makes use of Population Based Incremental Learning (a statistical approach to EC that marries reinforcement learning and Genetic Algorithms).

The future direction of this work is in modelling more interesting problems under SSCM, developing further strategies and developing more effective analysis tools for interpretation of results.

1.2 Introduction

This document is intended to provide a statement of position of the PhD work currently being undertaken. The document consists of two elements. The first element is to provide a discussion of the work and its relationship with various areas of interest and work being undertaken by others. The second element of this document provides a summary of the work carried out so far and the future direction. It is anticipated that this document will contribute both towards the final thesis and upcoming supervisory boards.

1.3 Position Statement

1.3.1 Motivation and objectives

At present various electronic market places, auctions and negotiation systems exist, in the near future full electronic supply chains will be possible and indeed desirable to improve efficiency [1].
This situation, however, presents a problem. While humans are good at negotiations and situation analysis they are less able to handle large volumes of information and numbers of transactions. What is need is a computer-based system or strategy for handling these situations. The strategy does not need to be the perfect negotiator, although it must be competent, but it must be able to deal with more negotiations more rapidly than a human operator could.

The first step in the process of creating this computer-based strategy is to develop a system for modelling the supply chains. Once the model has been developed it will be possible to describe specific supply chain situations or scenarios and begin the process of strategy development and analysis. A primary objective of the work is thus to create this modelling tool and to model various situations with it.

Once this objective has been met (at least in part) a further objective is to develop strategies that are capable of handling the modelled situations. These situations would require a middleman to be able to make a profit in a situation where customers are continually requesting bundles of products and may need to be negotiated with, suppliers must be negotiated with and there is a limit to both the communication capacity available and the amount of information about the market place. The strategies must be developed in light of the challenges initially stated and in light of the further objectives listed below.

A final objective in the work is to attempt to optimise these strategies through the use of evolutionary computation within a simulation system. This will necessitate both the consideration of which evolutionary computation systems are appropriate and the development of an effective simulation system within which these systems maybe used.

1.3.2 Current state and future work

At present the primary objective of developing a model for supply chains has been fulfilled with the creation of the Simple Supply Chain Model (SSCM) [2][3]. This model was partially inspired by the Trading Agent Competition (TAC) [10] and incorporates three types of supply chain participants (customers, suppliers and middlemen) each with various parameters including limited communication capacity. The model allows for the description of a wide range of scenarios including variations to the communication system being used by the participants. No constraints are placed upon the strategies each participant may use and indeed the model is intended to provide only a description of a system and the knowledge each participant starts with, nothing more.

The second related objective of modelling various problems has been partially fulfilled [2][4]. An initial model of a simply travel agent type system has been made. From this specific scenarios have been developed in which the use of this model has been further constrained and some restrictions placed on agent behaviour. No further models have been developed and thus work has focussed on this initial realisation – further modelling should be done, possibly in relation to bandwidth trading (a known area of interest for the sponsors). A formal method for the description of the scenarios arising from the initial model has not been developed although these scenarios have been described clearly.

The development of a strategy to tackle the scenarios has been undertaken and some initial experiments made within a simulation system [2][5][6]. This system was developed with the objective of allowing strategy optimisation via evolutionary computation and some results have been obtained about how the strategy adapts to different starting conditions within the scenario. Further work has been undertaken in regards to another scenario including the adaptation of the initial strategy to handle the increased complexity of the situations being faced.

One further challenge in the work has been that of data analysis. The experimental system produces a considerable amount of information of various sorts and some work has been done on the development of analysis tools. This work is discussed in a separate report [39].

1.3.3 Relationship to other work

There are a considerable number of areas of research that could be considered and are relevant to the work being undertaken. Since the concerns supply chain management research in this area would seem to be of interest, likewise the distributed computation nature of the problem would suggest work with multi agent systems, distributed problems solving, distributed optimisation/constrain satisfaction and so on would be relevant. Since we are dealing with trading systems work into bargaining, auctions and various market simulations would seem important. Finally the intended use of evolutionary computation further opens up this area as related and important. This is a huge amount of potentially relevant work. To provide a picture of how this work fits in to the big picture the following discussion is organised by project objective rather than topic area, the aim of this approach is to try and make clear the link between different parts of the project and related fields and lessen the potential confusion when a single area affects many elements of the work.

1.3.3.1 The Model

The SSCM provides a way to describe basic supply chain situation in their entirety. The SSCM does not attempt to dictate participant behaviour but it does restrict the way in which participant communicate. There isn't a restriction on the way the communications system is used; this is again left to the individual participants.

The SSCM is related to a number of different things. Firstly it is associated quite closely to the TAC game [10] by which it was partially inspired. The original TAC game provides a travel agent with a number of customers that require their orders to be satisfied within a certain time and a starting allocation of goods. Travel agents are able to satisfy customer order using their pre-supplied goods and by obtaining goods through a number of auctions. The game is purely a simulation system and no formal model exists for analysis, the simulation system was opened up to competing teams in an attempt to find effective strategies. The starting position described by the SSCM and the simulation system certainly share common elements however the SSCM doesn't specify agent behaviour with regards to customers and suppliers, something TAC did. The freedom available to the travel agents in TAC and middlemen in SSCM in somewhat similar, neither specifies the precise behaviour but both define the communications system. SSCM can effectively model the TAC situation by replacing the communications protocol with that used by TAC. Specifying the customer and supplier behaviour to be used in conjunction with the described SSCM situation and you have a good approximation of the TAC game. Further discussion of the similarities between TAC and the model can be found below.

During the development of the SSCM it has been intended that customer-middleman interactions essentially form part of a satisfaction problem while middleman-supplier interactions effectively relate to an optimisation problem. Thus work in distributed constrain satisfaction (DSC) [12] and by association constraint satisfaction (CS) [11] may also be relevant to the model. Constraint satisfaction models problems in which a number of variables require their values to be specified without violating certain constraints. CS problems (CSPs) are tackled using a central solving system based on many possible techniques. Distributed constraint satisfaction moves the traditional CSP into a distributed domain where different variables are maintained by computationally independent entities. DCS defines problems in a similar way to the SSCM, without restriction on the action of individual participants but some control over what messages are sent and how they are propagated. DCS deals specifically with constraint satisfaction in a cooperative environment where information is shared. SSCM deals with a considerably less formal environment in which cooperation is through mutual advantage and globally good solutions may not be found if it is not to the advantage of individual participants. The communication schemes within SSCM are tailored to bargaining and limited negotiation rather than relating of constraint information. While there are certainly surface similarities between SSCM and DCS there is considerable divergence in terms of intention and application.

There is a considerable body of work dedicated to various bargaining, negotiation and auction models [26][13][23][14]. These models in general deal with one to one interactions or one to many interactions and in the context of SSCM primarily relate to how participants may communicate. The models range from quite simple bargaining situations such as Rubenstein's complete information game [15] to more complex negotiation schemes where multiple issues are under consideration and maybe added or removed [16]. These model while important from a communication point of view do not capture the larger scope dealt with by the SSCM. While the Rubenstein's bargaining game (and related work [25][26]) or more relevantly Jennings auction and negotiation [17][18][19] work effectively deal with one to one interactions over a single or multiple issue they do not capture the larger context within which that interaction is placed. The SSCM is intended to capture that context. Thus while many of these model deal with the individual interactions in more detail their conceptual level is different. A further point of relevance and importance is that the SSCM is far removed from complete information games such as Rubenstein's. Participants have very little starting information at all and only the communication protocol to fall back on to make any headway. Rubesnstein's and other bargaining and negotiation models often assume far more information from price distributions to probable participant behaviour, this is something the SSCM does not do. In this sense SSCM has far more in common with other incomplete information games [25][26].

Complex market architectures also exist; one example of these is MAGMA [20] another is MAGNET [21]. MAGMA attempts to provide a complete framework for agent based electronic commerce; it deals with issues from individual agent behaviour to banking and goods handling. The MAGMA architecture while broad in scope and capable of implementing SSCM agents and scenarios impose more rigid constraints on agent behaviour than does SSCM; in effect the implementation issues within SSCM and MAGMA overlap but the underlying approach and intention is different. MAGMA could form the basis of an SSCM communication solution. MAGNET is a system focussed on the development of (possibly interlinked) contracts for the fulfilment of agent plans. It encompasses and extends the scope of the MAGMA system, providing a set of agent APIs for use within a regulated market place. Other agent market places exist such as AuctionBot [37] and Kasbah [36], these are effectively focused on how to alleviate users of having to deal with the negotiation process themselves.

Supply Chain Management deals with complex real-world supply chains and the problems faced by organisations trying to optimise them [29][35]. The SSCM is considerably less complex than the situations dealt with in traditional supply chain management although its results maybe applicable in some cases. Further the SSCM is directed towards a future with greater integration and automation of processes while SCM focuses on providing tools for decision-making and planning.

Distributed Artificial Intelligence (DAI) research has a considerable interest in negotiation between agents and the application of game theory to problem solving. The use of negotiation within DAI tends to follow two general approaches each with their own view of the world; the first is that of task assignment with high-level objectives in mind and that second is that of local negotiation between agents where a global view may not be feasible [33]. In DAI system however negotiation is used between essentially cooperative agents in a manner not too dissimilar to DCS [38]. The cooperative nature of these problems and their primary use with planning systems makes them incompatible with and distinct from the SSCM [34].

The communication protocol used with SSCM is intended to be generally adaptable and certainly not entirely fixed – in general however some variation on the alternating offers protocol [23]. has been considered most likely. This is discussed further below.

1.3.3.2 Modelled Situation and Scenarios

The modelled situation currently being investigated is a travel agent scenario similar to that in TAC [10]. There are however a number of important differences. Firstly customers are independent entities rather than game server controlled – they can (potentially) negotiate to some limited degree with the middlemen and indeed talk to multiple middlemen if initially told about more than one. TACs insistence that all customer preferences are revealed to middlemen at the beginning of the game does not hold true for SSCM – customers may communicate at any given time with the middlemen they know about and initial communications do not translate into complete information about their preferences. Suppliers in TAC are dealt with via auctions with certain rules applied about how those auctions and the related suppliers behave, SSCM suppliers deal with middlemen in a series of one to one negotiations that can include multiple goods as well as price. Success in TAC is centrally measured based upon each travel agents ability to fulfil the supplied customers requirements, SSCMs measure of success is with regard to the individual middlemen's profitability with the hope that this may prove globally optimal. As with TAC transport times for products are not considered, customers have negotiation cut-off times but delivery details and exceptions are considered outside the model.

Communication between participants is an extremely important issue when modelling a particular problem; the mechanisms and language used has a direct effect on the nature of the agents and their ability to deal with one another. The approach taken to communication thus far has been to make use of a variation on the alternating offers protocol [23]. Any participant's agent may make an offer for goods at a price, the receiving agent may either make a counter offer, accept or reject the proposal. This scheme while simple allows the research to focus on other areas and allows the agent strategies to be somewhat simpler and computationally less expensive. Other mechanisms could be used, possibly making use of more complex agent communication languages (ACLs) [24]. Argumentation as a mechanism for improving negotiation outcomes by persuading or threatening is currently an active area of research [7][8][9] and falls into this 'more complex' category. While argumentation and other systems may be a useful direction to take the work in the future this is not something that is currently intended.

1.3.3.3 Strategies

The agent architecture and strategies developed to tackle the scenarios are relatively simple, instead of a planning based approach a system somewhat akin to reactive systems is used. The negotiation mechanism and (to a lesser extent) the communication system are largely independent entities within the agent strategies and these will be discussed first.

As has been previously mentioned a considerable body of work exists within the context of bargaining. While complete information games have had perfect strategies devised ([15]) and various incomplete games have too ([25][26]), this is largely unhelpful in this context. Various reasons for this exist; part of this is due to the potential (perceived) irrationality of participants (remember the SSCM does not assume rationality), partial due to the potential multi-issue nature of negotiations but mostly due to the degree of incomplete information the agents face and the dynamics of the situation. The agent involved in one negotiation is highly likely to be involved in others, many of which will be related; the ability of the agent to know even it's own item valuation is limited and indeed it may be better to stop negotiations on short notice if conditions change. The middleman's customers, with whom it may also be negotiating, primarily determine these conditions.

The negotiation mechanism finally used was based on Matos's work ([22]) due to its high degree of flexibility and limited ability to track other player's negotiation strategies (these mechanisms have also been used in the auction domain [27]). This is critical in giving the agents a range of approaches to use in tackling different negotiators. A further reason for using this negotiation mechanism was its prior use with evolutionary computation, something intended here too.

1.3.3.4 Strategy, simulation and optimisation

The simulation system created for the work is designed to test SSCM scenario strategies and allow for their optimisation using evolutionary computation. The basic system concept is built around a feedback loop; SSCM agents exist within a market and are able to communicate with the SSCM applied restrictions, a central controller provides timing and configuration functions. The middleman agents within this market are highly parameterised and are essentially what we are interested in. The customer and supplier agents while parameterisable are so to a lesser extent and essentially make up the environment within which the middlemen reside. The middlemen maybe specifically defined or configured from the controller; controller configuration is used in conjunction with evolutionary learning. The customers may be specifically defined or defined from the central controller within set limits. The system can be run as a one off or as an ongoing market (the normal mode of operation). In continuous mode new customers may be added over time to maintain the pressure on middlemen to act effectively. The central controller is able to obtain middleman evaluations based on their earning, inventories and market prices (derived from supplier data).

Evolutionary computation is a powerful optimisation tool that has been applied to many domains including negotiation and auction research [22][27]. The evolutionary component of the simulation system is based around Population Based Incremental Learning (PBIL) [28]. PBIL was selected as it has been successfully applied to various problems ([28][30][31][32]) and allows learning to be undertaken using a small population of middlemen, large populations are impractical due to limited computational resources. PBIL combines evolutionary notions with those of reinforcement learning, the traditional population essentially being replaced by a probability distribution. Within the simulation system middlemen agents may be configured from the PBIL distribution; evaluation of middlemen performance over time may then be used to update the distribution and ineffective middlemen reconfigured. Effectiveness of one middleman is measured relative to the rest of the population. Other learning approaches may have been possible (including basic GA's) but PBIL's ability to work with small (sample) population sizes over repeated rounds was an important consideration.

1.4 References

[1] On Agent-Mediated Electronic Commerce, Minghua He, et al, 2003
[2] The Simple Supply Chain Model and Evolutionary Computation, Tim Gosling, CEC2003
[3] TechnicalReport1 - The SSCM, Tim Gosling and Edward Tsang
[4] TechnicalReport2 – SSCM Scenarios
[5] TechnicalReport3 – Scenario One Strategy
[6] TechnicalReport4 – Implementing Scenario One Strategies And The Use Of Evolutionary Computation
[7] Argumentation-Based Negotiation, Iyad Rahwan, Sarvapali D. Ramchurn, Nicholas R. Jennings, Peter McBurney, Simon Parsons, Liz Sonenberg
[8] Persuasive Negotiation for Autonomous Agents: A Rhetorical Approach, Sarvapali D. Ramchurn, Nicholas R. Jennings, Carles Sierra
[9] Reaching agreement through argumentation: a logical model and implementation, Sarit Kraus, Katia Sycara and Amir Evenchik
[10] The 2001 Trading Agent Competition, Michael P. Wellman, Amy Greenwald, Peter Stone, Peter R.Wurman, AAAI 2002
[11] Foundations of Constraint Satisfaction, Edward Tsang, ISBN 0-12-701610-4, 1993
[12] The Distributed Constraint Satisfaction Problem: Formalization and Algorithms, Makoto Yokoo, Edmund H. Durfee, Toru Ishida and Kazuhiro Kuwabara, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, VOL. 10, NO. 5, SEPTEMBER/OCTOBER 1998
[13] Distributed Rational Decision Making, Thomas Sandholm, Multiagent Systems, p201-258, MIT Press 1999
[14] Auction Theory: A Guide to the Literature, Paul Klemperer, 1999
[15] A Bargaining Model with Incomplete Information About Time Preferences, Ariel Rubinstein, Econometrica, Vol 53. No 5. 1151-1172
[16] Deliberative Automated Negotiators Using Fuzzy Similarities, C. Sierra, P. Faratin, N.R. Jennings
[17] A Heuristic Bidding Strategy for Multiple Heterogeneous Auctions, Patricia Anthony and Nicholas R. Jennings
[18] Developing a Bidding Agent for Multiple Heterogeneous Auctions, Patricia Anthony and Nicholas R. Jennings
[19] An agenda-based framework for multi-issue negotiation, Shaheen S. Fatima, Michael Wooldridge, Nicholas R. Jennings
[20] MAGMA: An Agent-Based Virtual Market for Electronic Commerce, Marksim Tsvetovatyy, Maria Gini, Bamshad Mobasher and Zbigiew Wieckowski
[21] MAGNET: A Multi-Agent Planning System for Contract Execution, Marksim Tsvetovatyy, Bamshad Mobasher and Maria Gini, AAAI 1998
[22] Determining Successful Negotiation Strategies: An Evolutionary Approach, Noyda Matos, Carles Sierra, Nick R. Jennings
[23] Strategic Negotiation in Multiagent Environments, Sarit Kraus
[24] The current landscape of Agent Communication Languages, Yannis Labrou, Tim Finin and Yun Peng, IEEE Intelligent Systems, volume 14, number 2, March/April, 1999.
[25] Bargaining with Incomplete Information, Shaheen S. Fatima, Michael Wooldridge and Nick R. Jennings, 2003
[26] Bargaining with Incomplete Information, Lawrence M. Ausubel, Peter Cramton, and Raymond J. Deneckere, Handbook of Game Theory, Vol. 3, Amsterdam: Elsevier Science B.V., chapter 50, 2002.
[27] Evolving Bidding Strategies for Multiple Auctions, Patricia Anthony and Nicholas R. Jennings, ECAI 2002
[28] Population-Based Incremental Learning: A Method for Integrating Genetic Search Based Function Optimization and Competitive Learning, Shumeet Baluja, 1994
[29] Supply Chain Modelling Using Simulation, Yoon Chang and Harris Makatsoris, International Journal of Simulation, Vol 2, No 1. 2001
[30] R. Sukthankar, S. Baluja, J. Hancock, Multiple Adaptive Agents for Tactical Driving, Applied Intelligence (1998)

[31] M. Sebag, A. Ducoulombier, Extending population-based incremental learning to continuous search spaces, In Th. Back, G. Eiben, M. Schoenauer, and H.-P. Schwefel, editors, Proceedings of the 5 th Conference on Parallel Problems Solving from Nature, pages 418--427. Springer Verlag, (1998)
[32] I. Inza, M. Merino, P. Larra~naga, J. Quiroga, B. Sierra, M. Girala, Feature Subset Selection by Population-Based Incremental Learning. A case study in the survival of cirrhotic patients treated with TIPS, Tehnical Report no. EHU-KZAA-IK-1/99, University of the Basque Country, Spain, (1999)
[33] Mechanisms for Automated Negotiation in State Oriented Domains, Gilad Zlotkin and Jeffrey S. Rosenchien, Journal of Artificial Intelligence Research 1996 p163-238
[34] Can't find the paper… Talks about non-cooperative negotiation in DAI
[35] A negotiation-based Multi-agent System for Supply Chain Management, Ye Chen, Yun Peng, Tim Finin, Yannis Labrou, Scott Cost, Bill Chu' , Jian Yao' , Rongming Sun and Bob Wilhelm
[36] Kasbah: An Agent Marketplace for Buying and Selling Goods, Anthony Chavez and Pattie Maes
[37] The Michigan Internet Auctionbot: A configurable auction server for software agents. P. Wurman, M. Wellmand and W. Walsh. In Second International Conference on Autonomous Agents, 301-308, 1998
[38] A Machine Learning Approach to Automated Negotiation and Prospects for Electronic Commerce, Jim R. Oliver
[39] Technical Report 6 - SSCM Experimental Results Analysis Systems, Tim Gosling