

UCFTAC: A Control Based Supply Chain Management Trading Agent

Ghaith Haddad, Brent Horine and Ladislau Bölöni

School of Electrical Engineering and Computer Science
University of Central Florida, Orlando, FL 32816-2450

Abstract

We describe UCFTAC, a trading agent based on control based principles, which has participated with success in 2006 Supply Chain Management Trading Agent Competition. The UCFTAC approach uses a sophisticated procurement strategy and an aggressive price control feature. It achieved competitive results in the TAC-SCM 2006 qualification round. Our design consists of a loosely coupled set of agent-managers, which exchange hints with each other and respond to direct commands from a central controller. We discuss the results obtained by the UCFTAC agent in the qualifying round of the trading agent competition.

Introduction

The Trading Agent Competition's Supply Chain Management game (TAC-SCM) (Collins *et al.* 2006) provides an effective testing ground for developers creating agents seeking to optimize the supply chain management of organizations. In this paper, we discuss the principles and implementation of the agent, named UCFTAC, developed by the authors for the 2006 TAC-SCM competition. While initially starting with very weak results, as a result of the tuning and optimization efforts, the performance of the agent steadily increased, finally allowing the agent to qualify for the finals.

The game consists of several automated agents competing as virtual companies manufacturing computers over the span of a simulated year. Agents respond to customer bids, competing for orders. They must also procure component parts from vendors at minimal cost and appropriate delivery times. Finally, they make decisions on how many of the 16 possible systems they build each day in their factory.

Design, Optimization and Tuning of the UCFTAC agent

The UCFTAC agent is composed of three main modules: the Customer Manager, the Inventory Manager, and the Production Manager. A central control agent issues commands to each of the modules as illustrated in Figure 1. The central control of the UCFTAC agent receives requests for quotes from the customers and hands them to the Customer Manager for pricing. The Inventory Manager receives hints about

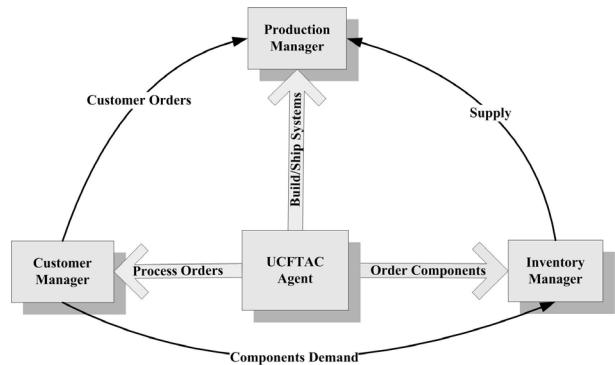


Figure 1: The flow of data and commands among the components of the UCFTAC Block agents

the component demand from the customer manager, and together with the information from the current inventory, ensures that adequate supplies are received from suppliers. The Production Manager receives information about the supply from the Inventory Manager together with customer order information from the Customer Manager, and upon command from the central control will build and ship the customer orders.

Because of the high degree of uncertainty and randomness built into the game, the authors chose a feedback control for ordering parts and setting bid prices. Many of the parameters, analogous to the gain and bandwidth of a feedback loop, are empirically determined, and manually adjusted based upon the results of the first few games of the competition. Other researchers have also employed a control approach (Kiekintveld *et al.* 2004) (Bell *et al.* 2004) generally for component ordering.

Bidding Strategy

The pricing strategy utilizes cost history to establish the estimated cost, bid history and an overall scale factor to set a target price. The approach aggressively reacts to market forces, adjusting the bid price to maximize the number of won bids at a favorable price. An end of year sale flushes the remaining inventory at the end of the game.

The success of this strategy is dependent upon a feature of

the game in which the customers always place all orders the day after receiving a bid. Because of this, the agent receives nearly immediate feedback on its bidding success. When the agent receives the customer orders, it matches them up with their corresponding bids and flags them as *won*. The remaining bids are flagged as *lost*. The agent begins the pricing process by finding the average *won* price, the average *lost* price, and the corresponding quantities for the last 10 bids issued. This data is then interpolated to find a new price, according to Equation(1). During the first days of the game, no bid history is available. Instead, default prices for components are used to generate a default price for the system with a nominal but aggressive margin designed to quickly generate orders for the factory.

$$TP = WP + \frac{(LP - WP)NW}{NW + NL} \quad (1)$$

where TP is the target price, WP is the winning price, LP is the loosing price, NW is the number of winning bids, and NL is the number of losing bids.

The estimated cost is the sum of the average bill of materials component costs for the last two orders for each component ID. The provisional bid price is the average of the target price and the estimated cost added to a margin proportional to one third of the fractional number of total bids won across all products as in Equation(2).

$$BP = \frac{EC + TP}{2} \left(1 + \frac{NW'}{3(NW' + NL')} \right) \quad (2)$$

where BP is the bid price, EC is the estimated cost, NW is the number of winnig bids across all SKUs, and NL is the number of losing bids across all SKUs. The ' indicates that the winning percentage is calculated over all of the products.

Finally, the predicted price is adjusted to remain within the bounds of a market price report produced periodically as a feature of the game.

Component Purchasing Strategy

As the game includes severe penalties for late deliveries, timely component procurement is critical. On the other hand, demanding short lead times from the vendors results in unacceptably high prices. To address this dilemma, the components are procured from multiple vendors with multiple delivery times. The preferred vendor (based upon quoted prices) is granted 70%-80% of the required quantities depending upon how many vendors exist for the component. The remaining quantities are evenly distributed among the remaining vendors. Furthermore, the items are scheduled in three shipments due in 3 days, 5 days and 6 days respectively, with the quantities front loaded. In addition, the order quantities are smoothed over three days in a form of a moving average to avoid large swings in order quantities which we found triggered higher vendor prices. Initial quantities for the beginning of the game, minimum thresholds, and expedited deliveries are applied as adjustments to the normal behavior.

Qualifying Round Results and Observations

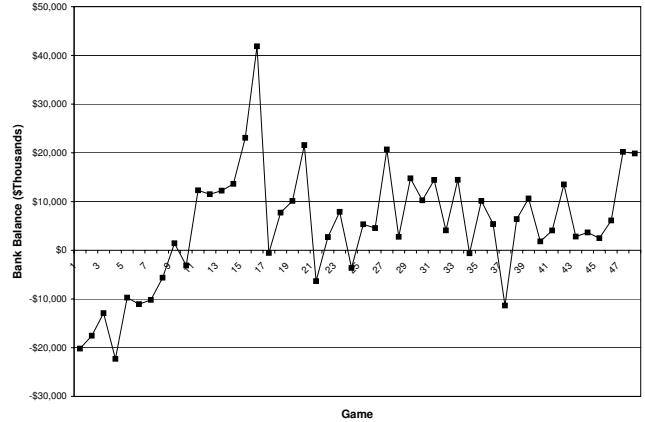


Figure 2: UCFTAC qualifying results for server TAC3

The first round of the competition is the Qualifying Round. During this round, teams typically refine their agents once they evaluate their performance in a competitive environment. UCFTAC placed 16 out of 23 in the Qualifying Round with a average bank balance of \$4.871 Million. Careful observation of the games indicates that the evolved approach performed significantly better than the initial approach. By subtracting the initial games and extrapolating the results, one can conclude that the UCFTAC team would have had an average bank balance of \$8.44 Million, placing near 11th place, had the evolved approach been used from the beginning. The round was won by TacTex with \$14.24 Million. Figure 2 illustrates the bank balance gained for each game on server 3 as an example. The figure shows that at the beginning of the qualifying round the original version of the UCFTAC showed weak performance. After the introduction of improved strategies the performance of the agent has improved. However, in this phase the agent showed an inconsistent behavior, scoring very high in some games but very low on others.

References

- Bell, S.; Benisch, M.; Bentall, M.; Greenwald, A.; and Tschantz, M. C. 2004. Multi-period online optimization in tac scm: The supplier offer acceptance problem. In *AAMAS-04 Workshop on Trading Agent Design and Analysis*.
- Collins, J.; Arunachalam, R.; Sadeh, N.; Eriksson, J.; Finne, N.; and Janson, S. 2006. The supply chain management game for the 2006 trading agent competition. http://www.sics.se/tac/tac06scmspec_v16.pdf.
- Kiekintveld, C.; Wellman, M. P.; Singh, S. P.; Estelle, J.; Vorobeychik, Y.; Soni, V.; and Rudary, M. R. 2004. Distributed feedback control for decision making on supply chains. In *ICAPS*, 384–392.