

Supply Chain Management League (Standard): An Overview

SCML Organizing Committee:

Y. Mohammed, A. Greenwald, K. Fujita, M. Klein, S. Morinaga, S. Nakadai

March 18, 2025

Abstract

This document provides an overview of the Automated Negotiation Agent Competition (ANAC) Supply Chain Management League Standard track (SCML-Standard). The game is intended to further research on agent negotiation. As such, the game design emphasizes negotiation and de-emphasizes operations (e.g., production, scheduling, etc.).

N.B. There are two tracks in SCML 2024. [This document pertains only to the Standard track.](#)

The SCM Standard world simulates a supply chain consisting of multiple factories that buy raw materials from, and sell final products to, one another. The factories are managed by autonomous agents. These agents are assigned a target quantity (drawn at random) to either buy or sell. They then negotiate with other agents to reach agreements, which become binding contracts that specify the terms of trade.

A simulation comprises multiple days, during which the Standard game is played. All agents have the same goal, namely to turn the maximum possible profit at the end of all days. The agent that turns the highest profit over all simulations wins. Learning is permitted from one day to the next during a single simulation; however, learning is not permitted across simulations.

Products ~~There are three product types: a raw material, an intermediate product, and a final product.~~ There are n product types: a raw material (product 0), a set of intermediate products (products 1 : $n - 2$), and a final product (product $n - 1$).

Production ~~There are two manufacturing processes, one for converting the raw material into the intermediate product, and a second for converting the intermediate product to the final product.~~ There are $n - 1$ manufacturing processes (for n products), one for converting each product (product i) to the next in the chain (product $i + 1$).

Factories Factories convert input products into output products by running their manufacturing processes on their production lines. All processes run convert exactly one unit, instantaneously, at a predefined cost.

Production Graph and Negotiations ~~Factories are organized in two layers L_0 and L_1 (see Figure 1). L_0 factories receive exogenous contracts to buy the input (raw material), and then negotiate with L_1 factories to sell them the intermediate product. L_1 factories receive exogenous contracts to sell their output (final product), and then negotiate with L_0 factories to buy the intermediate product.~~ Factories are organized in $n - 1$ layers (for n products) L_i (for $0 \leq i \leq n - 2$) (see Figure 1). L_0 factories receive exogenous contracts to buy the input (raw material). L_1 factories receive exogenous contracts to sell their output (final product). These exogenous contracts form the supply and demand in the market. L_i factories negotiate with L_{i+1} factories for selling product i . All negotiations in a single day happen simultaneously.




Figure 1: Example SCML-Standard world with 5 products (i.e. 4 manufacturing processes and 4 levels of factories). Each factory is represented by an agent, whose goal to negotiate buy and sell contracts that maximize profits.

Agents The agents in the SCM world function as **factory managers**. They negotiate to reach agreements to buy and sell intermediate products, which automatically become binding as contracts.

Negotiation Protocol Agreements are negotiated using a variant of the bilateral **alternating offers protocol**, typical of ANAC competitions. Each offer specifies a buyer, a seller, a quantity, and a unit price. The sequences of offers and counteroffers in a negotiation are private to the negotiating parties.

Negotiation Issues All negotiations concern three issues: quantity, delivery day and price.

Delivery Day an integer between zero (representing) current day and H representing the future horizon for negotiation (i.e. number of days in the future to deliver the product). H is called the negotiation horizon.

Quantity: an integer between 1 and the factory’s number of production lines multiplied by some configuration parameter $\sigma \geq 1$.

Unit Price: an integer between $\lceil \kappa \text{tp}(s) \rceil - 1$ and $\lceil \kappa \text{tp}(s) \rceil$, $\lfloor (1 - \kappa) \text{tp}(s) \rfloor$ and $\lceil (1 + \kappa) \text{tp}(s) \rceil$, where $0 < \kappa \leq 1$ is a configuration parameter (e.g., 0.1) and tp is the trading price of the intermediate product.

Utility Functions An agent’s utility function represents its profits for the day not taking into account any future contracts. As such, it is simply the total revenue it receives from any sales less its total expenses, the latter of which includes the contracted cost of the input product as well as the agent’s private production costs, storage costs, and shortfall penalties.

N.B. While each agent’s production costs, storage costs, and shortfall penalties are private information, the distributions from which these values are sampled are common knowledge.

Trading Price The **trading price** (tp) of a product is a weighted average of its past prices, which weighs newer contract prices more heavily than older ones. The trading price is used by the simulator to set the price range of all negotiations, and for calculating penalties.

Balances Factories have an associated balance—seeded at the start of the game with some finite amount—from which they withdraw to pay for supplies, etc., and into which their sales revenue is deposited.

Bulletin Board The SCM world contains a world-readable **bulletin board** that conveys both static and dynamic information about the game environment and all factories over the course of the simulation.

The static information includes the simulator settings (e.g., number of simulated days), and product information, namely a list of the consumers and producers of all products (i.e., all factory’s positions in the production graph), and the initial trading prices (called catalog prices).

The dynamic information includes a trading price list (per product), which reports a weighted average of each product's past prices; and a financial reports section (also per agent), which is updated only periodically, that summarizes the financial standing of all factories (e.g., their balances).

Finally, the bulletin board also contains an **exogenous contract summary**, which reports the total quantity and average unit price of exogenous contracts each day.

The Simulation Each simulation of the SCM world runs for multiple (say, 100) days. Before the first day, each agent is assigned a production cost. During each day:

1. The world generates exogenous contracts, and samples storage costs and shortfall penalties for all agents from their corresponding distributions.
2. Agents engage in multiple (say, 20) rounds of negotiations with their negotiating partners. They can also read the bulletin board.
3. All contracts are executed: i.e., products are moved from the seller's inventory to the buyer's, and money is moved from the buyer's account to the seller's.¹
4. The bulletin-board is updated, most notably to reflect new trading prices, updated financial reports, and the day's exogenous contract summaries.

Differences from SCML 2023 The Standard 2024 game is a simplification of the standard 2023 game, designed to simplify the API. Notably, it replaces the spot-market mechanism with the shortfall-penalty and simplifies contract processing by avoiding complicated breach processing and bankruptcy rules replacing all of this with independent execution of contracts.

The Standard 2024 game is an extension of the OneShot 2023 game designed to take into account inventory management as well as negotiation while keeping the simple interface of the OneShot game. The specific changes from OneShot 2024 and their implications are:

1. Products are not perishable in the Standard game allowing factories to accumulate stock of their input product. This means that profits calculated every day are lower limits as they do not take into account the value of accumulated stock which can be manufactured and sold in the future. This implies that there is no disposal of products at the end of each day, instead there is a small cost associated with carrying stock (storage cost).
2. Agents can negotiate future contracts instead of only being allowed to negotiate about deliveries on the same day.
3. Negotiated quantities have a larger range to allow for proactive stockpiling of the input product.
4. Negotiated prices have a larger range which necessitates reasoning about prices as well as quantity matching.
5. Production graphs can be deeper which means that some agents will negotiate with both their suppliers and consumers at the same time.

How to Compete To participate in the Supply Chain Management League (SCML), you should write and submit code for an autonomous agent that acts as a factory manager.

In the Standard track, at most one instantiation of each agent will run in each simulation, together with an unknown mix of additional agents prepared by other participants and by the organizing committee. An

¹Contracts are executed in order, starting from the raw material and ending with the final product. Production is run in the same order, which guarantees the accuracy of the utility functions.

agent's performance will be measured by its score, which will be computed as the truncated mean² of the utilities (i.e., profits) accrued by all the factories it is assigned to manage across all simulations.

All tournaments will be conducted in two rounds, a qualifying round and a final round. All entrants that are not judged to break any of the SCML and ANAC submission rules will be entered into the qualifying rounds. Top-scoring agents in the qualifying round will then be entered into the final round.

The final results will be announced at AAMAS 2024. It is expected that finalists will send a representative to the ANAC session (at AAMAS 2024), where they will have the opportunity to present their agent.

²An agent's truncated mean will be calculated by first sorting that agent's scores in all the simulations, and then removing the top and bottom x_t and x_b scores from that agent's sorted list, where x_t and x_b are values selected by the organizing committee to balance test efficiency (taking into account scores from as many simulations as possible) and robustness (insensitivity to outliers, or to a few simulations in which the agent realizes extremely high or low profits).