TAC Market Design: Planning and Specification

Version 1.03
2006-06-15

Contents

1 Introduction

This document is a follow-up on the initial proposal for the TAC Market Design competition, also called CAT[1], and provides a more detailed specification. The initial document is the result of discussions held towards the end of April 2006 in Liverpool with Steve Phelps, Tim Miller, Thierry Moyaux, Peter McBurney and Enrico Gerding, and further discussions during the AAMAS'06 conference with Simon Parsons, Elizabeth Sklar and people from the TAC community. It is intended as a basis for more discussion and will be updated over time to include changes and improvement resulting from these discussions.

2 Change Log

- Version 1.00 : first version by Enrico Gerding
- Version 1.01 : (Enrico Gerding) Incorporated timetable from original proposal into section 4.2. Aims of various components and who designs what more explicit in section 5.
- Version 1.02 : (Simon Parsons) Revised timetable from Peter McBurney, changed terminology from "broker" to "specialist", and added other comments from the CUNY team.
- Version 1.03 : (Peter McBurney). Added list of issues arising from MBC Project Meeting on 2006-06-06, in section 3.

3 Current Issues

This section presents issues for debate and resolution.

3.1 Soton June 2006

The first list of issues here arise from comments and questions made by members of of the MBC project team at the meeting in Southampton on 2006-06-06.

1. (Dave Cliff) The CAT game structure was viewed as a good match with typical market structures in national stock markets, with the national exchange (LSE, NYSE, etc) mapping to our meta-exchange. One difference is that, in real markets, traders may interact directly with the meta-exchange. Should we enable this? If not directly, should we enable it via a dummy specialist who simply passes bids and asks onto the meta-exchange without attempting to match them?

[1]The reverse TAC, and also an abbreviation of "catallactics", which is the science of economic exchanges.

- Enabling this would set the meta-exchange in direct competition with the specialists. Do we want this?
- DC thought that enabling this would increase the attractiveness of the CAT game to exchanges such as the LSE.

2. (Dave Cliff) In real equity markets, the reasons for the existence of specialists are:

   - Anonymity of trading (ie, trading through intermediaries)
   - Potentially faster trades than on the meta-exchange
   - Greater liquidity (since specialists can buy and sell on their own account, they bring their own money to the market; in addition, being smaller and more focused, they may be better at marketing to particular segments of traders than a monolithic exchange).
   - Price competition in the form of commission rates and structures.

3. (Dave Cliff) What model of time will we use in CAT? For example, we could allow each round to have an elapsed time of (say) 1 hour, but to represent a day of trading, with specialists only interacting with the meta-exchange at the end of each such day.

   - PM view: We will have to articulate the precise sequence of interactions between traders, specialists, meta-exchange and registry.

4. (Robin Mason) Should we make the traders (or, equivalently, their strategies) available to the entrants? In real markets, trader strategies are not known to specialists. On the other hand, if any of the MBC partners enter the competition we will need to ensure all entrants have the same information about trader strategies. Entrants may also wish to be able to play with a version of the game as part of their preparation, so that we would then need to provide potential traders to entrants.

   - PM view: I believe our original idea was that we would release publicly the finite set of strategies which our traders could adopt, but not their relative proportions across the population of traders. This proportion would, in any case, change as the game progress, since it provides us with a leaver to create market dynamism.
   - Someone suggested that we could provide a finite super-set of strategies, and say that our traders would be drawn from this set. Thus, no particular strategy would necessarily be adopted.

5. (Robin Mason) How will specialists communicate their offers to traders? If they are allowed to advertise, do we insist they signal truthfully? If not, there may be a mis-match between advertisement and reality (eg, with regard to average settlement time), and so traders would need to factor this into their decision-making.

- PM view: I think *caveat emptor* is best here (let specialists lie if they wish), although our traders may require more sophisticated learning mechanisms to cope with this.

6. (Robin Mason *et al*) If we permit specialists to lie, then the game may also be used to study issues of trust and reputation. It may even be possible to run a trust-based game as an overlay on CAT.

   - It was decided by the meeting to park this issue at least for the first year of CAT.

7. (Robin Mason) What happens to specialists if they do not clear?

8. It was suggested that we learn more about the 2006 AAAI game (which may be similar), and also contact the providers of the two TAC games for their experiences.

9. It was proposed to solicit feedback on our game design from experts in economic mechanism design. We have the funds to pay for this advice. The people suggested were:

   - Steve Gerstad, Arizona
   - Tom Stenhouse, London Stock Exchange (formerly HP).
   - Mark Armstrong, UCL/ELSE
   - TAC people.

   Robin Mason offered to suggest other possible reviewers from economics.

4 Resources and Planning

4.1 Resources

Steve Phelps will leave the project full time from the 20th of May, but may continue working on the competition part time. Simon Parsons, Elizabeth Sklar, Jinzhong Niu and Kai Cai, are willing to contribute to the competition, and it is expected that they will lead the implementation.

4.2 Planning

A prototype of the competition should be implemented as soon as possible to identify any issues with the proposed design (some of the expected issues are detailed in this document, but others are likely to emerge). The prototype will be an extension of the existing JASA framework. Work on the prototype is expected to start early June. Others, e.g. Southampton, can then test the system by designing mechanisms for the prototype.

Our provisional timetable is as follows.

Now - end June 2006 Outline specification of prototype, cost and timing estimates and approvals

Begin June - mid July 2006 Development of initial prototype platform by CUNY, scoping and design of trader strategy engine

Mid July 2006 - end August 2006 Assessment of lessons of prototype, creation of trading platform (hopefully, building on work of prototype), development of trader strategy engine.

September - October 2006 Integration and testing of trader strategy engine with trading platform, porting and testing of combined system onto fast hardware platform (this may be straightforward, or maybe not), refinement of API and rules.

November 2006 Release of detailed rules, API, platform and trading agents for potential entrants.

April 2007 Competition entrants register initial expression of interest.

May 2007 Test Run of Competition (allow some or all entrants to submit their markets, as a trial of our system)

Summer 2007 Operation of TACMarket Design, alongside TAC Classic/TAC SCM.

5 Market Architecture

This section describes the various components of the TAC market design game, their objectives, and how these components interact. The market consists of specialists[2], traders, and a meta-exchange (see figure 1). Traders buy or sell a certain quantity of a homogenous good by placing bids and asks at a particular specialist (see section 5.1 for details). The goal of each trader is to maximise a given utility function based on the demand and supply function. The specialist is a local exchange that matches the bids and asks and sets parameters such as participation costs and costs of information. In addition, the specialist can trade with other specialists via the meta-exchange. A specialist can thus decide to trade on the global market through the meta-exchange rather than clear the auction within its own local market. The goal of a specialist is to maximise a given scoring function based on profits and possibly other factors such as trading throughput (section 5.2 describes the way in which specialists can make a profit,

[2]Earlier versions of this document used the term "broker" for the artificat that CAT entrants will develop. This earlier usage seems to run against the meaning of the term — a broker is usually taken to be someone who trades on behalf of someone else. What we are conered with here is closer to the "specialist" or "market maker" of the New York Stock Exchange, who provides a markey in a given kind of share. The difference, of course, is that here all "specialists" provide markets in the same commodity, so the term isn't a perfect fit.

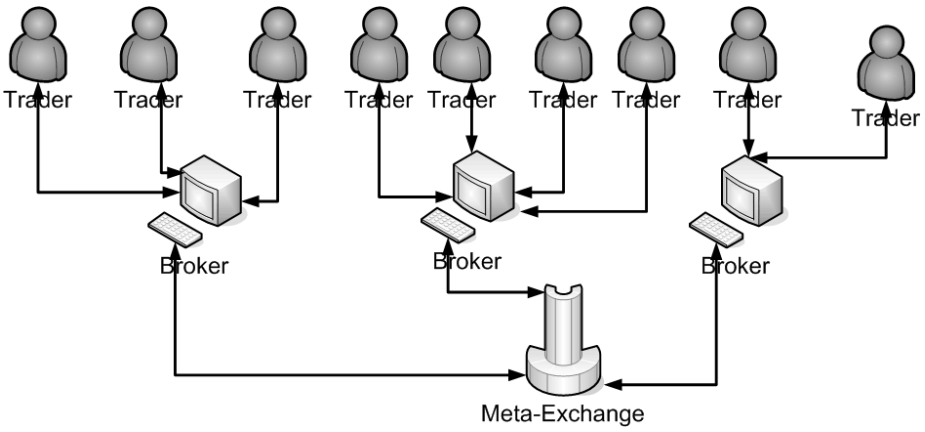

Figure 1: Market architecture.

section 6.1 describes the scoring in more detail). Finally, the objective of the meta-exchange is simply to facilitate trade between specialists.

The competition designers provide the traders and the meta-exchange, described in sections 5.1 and 5.3 respectively. Given these components, each entrant to the competition is required to design a single specialist, i.e., the rules with which the specialist performs the matching, set the parameters such as costs (these settings may change during the course of a game), and decide when and how to trade on the meta-exchange. The specialist is further detailed in section 5.2.

NOTE: Although a specialist is allowed to place a trader's bids and asks on the meta-exchange, currently the specialist is not allowed to own any of the traded goods. As a result, all bids and asks placed on the meta-exchange have to match with trader's bids and asks. This places the emphasis on the mechanism and away from designing trading strategies. In future implementations, however, the specialist could also be allowed to own goods.

5.1 Traders

Each trader has a demand and/or supply function, a bidding strategy, and a specialist selection strategy. In addition, a trader may have some initial endowments and a set of goods which can be traded on the market. Traders can be either buyers or sellers, depending on the demand and supply function, the current market price, and whether a trader has resources to offer.

5.1.1 Bidding Strategy

As identified in the proposal, one of the main hurdles is probably the design of the strategies for the space of mechanisms that the entrants can implement.

The idea is to use the existing set of strategies in JASA and variants of these strategies (e.g. with different parameter settings) for the initial prototype.

Most of the bidding strategies are *adaptive* and learn from past transactions. If the competition is run online, then the learning needs to be fast, both computationally and needing few iterations before having fairly good solutions. In addition, the traders could be initialised by learning with a default mechanism before trading in the actual competition.

5.1.2 Specialist Selection Strategy

Traders that can switch between specialists are called *mobile*, whereas non-switching traders are called *stationary*. The advantage of mobile traders is that it creates direct competition amongst specialists. We have not yet decided in detail how mobile traders choose a specialist, but we are leaning towards a simple n-armed bandit approach in which traders maintain an estimate of the expected value of trading in each specialist's market, and make their choice on this basis. A more advanced specialist selection strategy could utilise the information of the specialists directly to make a decision. This information includes the price of participating and of information, and other information such as market price and past interactions.

5.2 Specialists

The specialist is designed by the entrants. The entrants need to do the following:

- Implement a clearing and pricing policy by matching bids and asks from traders.
- Set fees: e.g. price of information, participation costs, price for placing a bid or ask. It should be possible for a specialist to price discriminate between buyers and sellers, and possibly charge for different quantities.
- A specialist can also provide information, and possibly request information from other specialists. Information may include transaction history or the order book, as specified in the proposal.
- A specialist should also be able to transfer existing bids and asks to the meta-exchange. Although the specialist can also be viewed as a trader as such, it is probably not desirable that the specialist is allowed to own any traded goods as well.

  NOTE: It is not yet clear whether a specialist should be allowed to change the quantity and/or price of the bids and asks.

Specialist profit/loss The specialist can make a profit by setting appropriate fees. The designer of the specialist has to carefully balance the fees since traders could go to another specialist if fees are too high. In addition to the fees, the

specialist can make a profit by taking a share of the surplus (the difference between ask and bid price) according to its pricing policy. Note that it may also be possible for a specialist to (temporarily) make a loss by matching a bid and an ask where the bid price is lower than the ask price (this may be done e.g. to attract buyers and sellers, but it could also be used to increase the amount of trade).

NOTE: We may need to detail the extent to which a specialist can run a deficit, whether it has to "borrow" money to support a "deficit" and so on.

Announcing the policies Whether or not the policies of the specialist need to be announced depends on the complexity of the trading agents. For example, if a trading agent calculates its expected utility based on the fees, then it makes sense to announce the fees publicly before any bids and asks are placed. It will be more complex, however, to also announce the clearing and pricing policies since (1) a language needs to be specified which might restrict the freedom of the specialist designer, and (2) traders will have to be sufficiently intelligent to parse and understand the implications of different policies.

Bookkeeping Bookkeeping is required in order to assure that specialists will not perform illegal actions, such as matching non-existent traders, claiming they have more profit than they actually have, or give out erroneous information about the market to the traders. Bookkeeping is closely related to system architecture and the communcation language and is discussed in more detail in section 7.

5.3 Meta-Exchange

The meta-exchange allows specialists to interact with one another by placing bids and asks to balance excess demand and supply at a local specialist. The pricing and clearing mechanisms need to be pre-defined by the competition designers. A standard continuous double auction could be used. There may also be costs involved in participating through the meta-exchange in order to discourage the global interaction and stimulate local interactions within a specialist. However, this may not be necessary as local interactions are probably more profitable for a specialist anyway.

NOTE: Who will run the meta-exchange? Would it be better to allow specialists to go into other markets as traders?

6 The Game: Scoring, Rounds and Moves

A game consists of a number of trading days, and each trading day consists of a number of rounds. Broadly spekaing, in each round, every trader has the chance to make an offer, and every specialist has the chance to clear their market. (It

Agent Type	Action	Interact with
Trader	Select specialist	-
	Place bid or ask	Specialist
	Request information	Specialist
Specialist	Set fee parameters	Trader
	Provide information	Trader
	Match bids and asks	Trader
	Set trading price	Trader
	Place bids and asks	Meta-Exchange

Table 1: Summary of possible agent moves.

is not decided whether the meta-exchange will clear every round, or whether, for example, it will only clear at the end of a trading day.) Table 1 presents an overview of the possible actions that the agents can take during the competition.

6.1 Scoring

In the initial proposal the scoring is based on a weighted combination of measures such as profit, trade volume, etc. Since any weighted function is rather arbitrary, an idea is to announce the weights only just before the game starts, and possibly change from one day to the next. This way, the specialists need to handle a a wide variety of goals. During the discussions at AAMAS, however, some concerns were raised as to having such a function as opposed to using a single "clean" scoring based only on the specialist's profits.

7 System Architecture

Whereas the market architecture describes the competition on a conceptual level, this section is concerned with implementation issues. To this end, this section describes the various components of the actual system, how these interact, and the communication language used. Furthermore, implementation issues such as security and robustness are also discussed here. More specifically, these include issues such as how to ensure that errors in a specialist implementation do not affect the entire system (robustness), and how to prevent illegal actions such as setting buy prices above the bid or sell prices below the ask (security).

The system architecture is depicted in figure 2. As shown in this figure, all communication between specialists and clients and between specialists and meta-exchange goes via the server. The server passes the messages to the registry which keeps a record of the transactions. The registry also plays an essential role in security and bookkeeping and is further explained below.

NOTE: The precise role of the registry is an open question. For example, the specialists could also directly communicate with the registry to request informa-

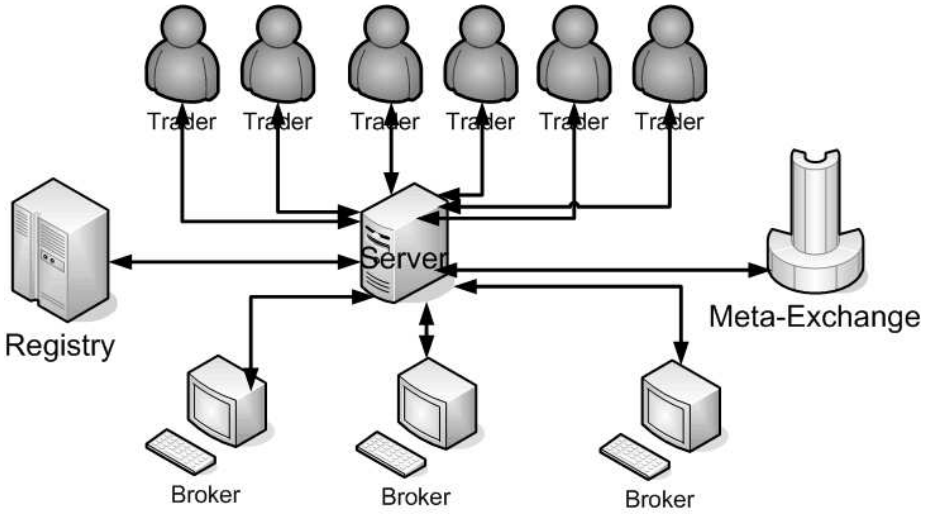

Figure 2: System architecture.

tion regarding other specialist's transactions (possibly for a fee). In that case the registry is also part of the market architecture.

7.1 Registry

A main task of the registry is to keep track of all the costs and profits made by the specialists (bookkeeping) and to intercept all the messages in order to verify that the actions are legal (security).

7.1.1 Bookkeeping

In order for the system to work properly, the registry must calculate all the profits/losses resulting from matching, request for information, etc (since the specialists have an incentive to misreport their profits). As a result, the registry needs to be informed about all the fees and keep track of the specialist/trader actions. This is feasible using the current architecture since all the messages pass through the central server and are passed on to the registry.

7.1.2 Security

The registry must also ensure that no illegal actions can be executed. For example, the buying price is never allowed to exceed the bid and the selling price is never allowed to go below the ask when a buyer and seller are matched. Note, however, that a matching may take place even if the bid price is below the ask price. In that case, the specialist makes a loss. It is also important that the trading price for a matched buyer and seller may differ, thereby allowing the specialist to obtain a share of the surplus.

Furthermore, in the current proposal the specialist is not allowed to actually own any of the traded goods; the sole purpose is to match buyers and sellers. As a result, specialist is only allowed to match a buyer if a seller exists, and visa versa. This can be easily implemented by restricting the communication language. However, the registry still needs to verify the matched bids and asks actually exist, and that they are not expired. In addition, the interaction with the meta-exchange may complicate matters.

7.2 Communication Language

7.2.1 Specialist-Trader Interaction

- Bids and Asks. A trader can place bids and asks containing price-quantity pairs, and a deadline after which it will expire. The system could also be extended to allow a trader to submit multiple XOR asks and bids to reflect the demand/supply curve. The server attaches an ID to each submitted bid and ask.

  - *ask(quantity,price,deadline)*

  - *bid(quantity,price,deadline)*

- Matching. The specialist can match a buyer's bid and a seller's ask using the bid and ask IDs. The trading price for the buyer and the seller can differ, enabling the specialist to take a cut of the profit. In addition, the quantity need not match with any of the quantities specified in the bids and ask. As a result, the syntax is as follows:

  - *match(bid-id,ask-id,buy-price,sell-price,quantity)*

7.2.2 Specialist-Meta Exchange Interaction

TBD

8 Miscellaneous

8.1 Running the Competition

Current TAC competitions (TAC classic and TAC supply chain) are running online over a number of consecutive days. Furthermore, the code of each entrant runs on a local machine, and communicates remotely with the server where the actual competition is running, allowing entrants to change their code during the course of the competition. Ideally, a similar approach should be used for the new competition. However, communication and computational issues may prevent such an approach to work in practice (e.g. we expect the mediator to have frequent communications with clients). As a result, although the distributed approach is more desirable, we may need to consider alternative approaches such as running the entire competition centrally, possible using a cluster. The

prototype can be used to determine whether or not the distributed approach is feasible.

Whereas initially the idea was to run the competition offline and using empirical game theory to determine the result of the competition, currently the idea is to have the competition run online and use less computationally intensive methods to evaluate the mechanisms submitted by the entrants.

8.2 Deadline Effect

Entrants may exploit the fact that the game has a (artificial but necessary) deadline. For example, since traders need a few iterations to respond to changes (due to the learning mechanism), a specialist could suddenly charge high costs and/or claim the entire surplus in the final rounds of the game. One way of dealing with this is to introduce uncertainty when the deadline occurs (this approach is used in the ART trust and reputation game). Alternatively, the deadline effect need not be avoided at all, and could be considered as a feature of the game (this approach is used in the TAC supply chain). A third possibility is to have a known deadline, but to rate the specialists upon their profits (and whatever else is chosen as a metric) at sme point before the deadline (some point after the traders have carried out some learning and before any deadline effect takes place).

The choice comes down to what we want to measure, and what kind of specialist strategies we want people to create. If we want entrabts to create strategies that could be useful in practice in commodity markets, we need to assess the steady state performance of the specialists, and so want to exclude or reduce any deadline effects. However, from the point of creating an interesting competition, allowing deadline effects gives entrants another aspect of strategic behavior to manipulate.