AutONA: A System for Automated Multiple 1-1 Negotiation

Andrew Byde
Hewlett-Packard Laboratories,
Filton Rd., Stoke Gifford,
Bristol BS34 8QZ, UK.
Andrew.Byde@hp.com

Kay-Yut Chen
Hewlett-Packard Laboratories,
1501 Page Mill Rd., Palo Alto,
California 94304, US.
Kay-Yut.Chen@hp.com

Claudio Bartolini
Hewlett-Packard Laboratories,
1501 Page Mill Rd., Palo Alto,
California 94304, US.
Claudio.Bartolini@hp.com

1. INTRODUCTION

For any organisation, saving on procurement costs has an impact on profitability that is multiplied by gross margin. Although much research focus has been placed on achieving the lowest possible cost for the goods/services purchased, we concern ourselves here with the operational procurement costs of the organisation.

The AutONA (Automated One-to-one Negotiation Agent) system was conceived as a means of reducing these operational procurement costs, enabling procurement departments to automate as much price negotiation as possible, thus creating the option of reducing direct costs and/or redeployment of operational effort into strategic procurement requiring high human involvement. The problem domain has been limited to the automation of multiple 1:1 negotiations over price for quantities of a substitutable good subject to the organisations procurement constraints of target quantity, price ceiling and deadline.

We present the design of the core reasoning system and preliminary results obtained from a number of experiments conducted in HPs Experimental Economics Lab. The architecture of AutONA is that there is a central **reasoner** that sets goals, targets and price-caps for each seller-quantity pair, and for each seller there is a **bidding agent** that interprets these control parameters, and handles negotiation with a given seller, maintaining a record of the history of interaction, and acting accordingly. Thus the reasoner is responsible for assessing the merit of each seller's position relative to the others, while the bidding agent chooses good local policy accordingly.

Our main conclusion is that AutONA could reasonably be deployed for automated negotiation, having shown no evidence for being identified as an automated system by suppliers, and having demonstrated comparable gains from trade.

In the next section we review previous work in the automated negotiation domain. In §3 we specify the reasoner, in §4 we specify the bidding agents. In §5 we describe the experimental setup for the human-based evaluation of AutONA; in §7 we summarize the results of these experiments; and §8 we conclude.

2. BACKGROUND

The problem of bargaining is an old one. When studied from a Game-Theoretic perspective (see, e.g. [3]) the problem of what offers to make (in a single 1-1 negotiation context) almost always reduces to a calculation of the correct *first* offer to make, which is immediately accepted by one's opponent - so that no actual negotiation occurs! This reduction is due to assumptions about the nature of rewards (in particular, there are no hard deadlines), information, and rationality that simply do not hold in the real world.

In [1], Faratin et al. present a pragmatic approach to the replication of reasonable human 1-1 negotiation strategies in machines. The usage of the term **tactic** here (in §4.2) mirrors the use there. They consider three main types of negotiation tactic: time-based, resource-based, and imitative. An important class of behaviour missing from this taxonomy is that of *competition*-based behaviour, which is only feasible in a context where there are several parallel streams of negotiation being conducted.

As far as we are aware, there are no attempts in the literature to address this problem. Instead, researchers have investigated methods for conducting multi-variable 1-1 negotiation, e.g. [2], [5], or have focused on the application of negotiation technology to various distributed computing problems, such as resource allocation (e.g. [4]).

3. REASONER SPECIFICATION

3.1 Assumptions

The beginning of any procurement process is a **purchase request**, specifying the quantity desired Q and the maximum price acceptable for the full quantity, P. This quantity Q can be bought from one or more **seller**s, each of which has a minimum quantity they will consider selling, a maximum quantity they will consider selling, and whose potential sale quantities jump in some specified minimum increments. These parameters are specified (for a seller S) as q^S_{min}, q^S_{max} and q^S_{step} respectively.

The reasoning about how much to offer for each quantity centers around **options**, where an option o is defined by a seller and a quantity. For each option, the system forms a series of estimates regarding the likely price of purchasing the specified quantity from the specified seller. These estimates are parametrized by a **risk parameter**. The possible values for the option risk parameter are **best**, **expected** and **no-risk**, giving rise to three prices for each option o, $p_-(o)$, $p_e(o)$ and $p_+(o)$ respectively.

3.2 Price Estimates

The price estimates for an option are calculated using estimates of the distribution of the lowest price a seller will accept for the option's quantity. A **belief** is a probability distribution over prices *per unit*, parametrized by the properties that an option may have.

There are many possible ways to represent beliefs; observation of frequencies in historical data can be used to build non-parametric models, but when the amount of data is small, these methods are not suitable. We choose to assume a log-normal distribution on prices, and select the mean and standard deviation to minimize squared-error with respect to observed closing prices in prior negotiations. The observed closing prices in previous negotiations are normalized with respect to a benchmark price that carries information on the market price on the date that the negotiation was concluded. By doing so, we reduce the impact of the variation of market prices over time. To this effect we also introduce a customizable scale factor that gives exponentially less weight to older data.

To each seller S we associate a belief function $b_S(p,q)$, with the interpretation that the probability of the price for the option $o \in \mathit{option}(S)$ closing between prices p_1 and p_2 (per unit) is believed, prior to the start of negotiations, to be

$$\mathit{Prob}(p(o) \in [p_1, p_2]) = \int_{p_1}^{p_2} b_S(x, q_o)dx. \tag{1}$$

The price estimates for a given option are generated as follows:

1. The **best** price $p_-(o)$ of an option o is defined to be the current highest offer that AutONA has made for the specified option, or some fixed minimum, $p_{min}(o)$ otherwise (i.e. if no offer has yet been made).
2. The **no-risk** price $p_+(o)$ of an option o is the larger of the best price and the largest number p such that for all $p' < p$, $\mathit{Prob}(p(o) \in [p', \infty]) > 0$. Informally, it is the highest price to which AutONA should attatch non-zero probability (via the belief).
3. Given an option o with quantity q_o and seller S, associated belief $b_S(p,q)$, best price $p_-(o)$, and no-risk price $p_+(o)$, the **expected** price of the option o is given as

$$p_e(o) = \frac{\int_{p_-(o)}^{p_+(o)} x\, b_S(x/q_o, q_o)\, dx}{\mathit{Prob}(p(o) \in [p_-(o), p_+(o)])}. \tag{2}$$

3.3 Spreads

Although it may be that negotiations will be for the full quantity with each seller, it is also quite possible that due to quantity constraints, it will be necessary to divide the full purchase quantity Q between several sellers; the trade-offs that AutONA then makes will be between alternative ways of dividing the quantity up between the available sellers. We call such a "dividing-up" a **spread**. Formally, a spread is a set of options.

Just like an option, associated to a spread are a quantity, and a range of estimates of the price at which it can be obtained, where the prices are parameterized by risk. The quantity of a spread $\sigma = \{o_1, \ldots, o_k\}$ is just the sum of the quantities of its corresponding options, and the prices are defined likewise:

$$\begin{aligned} \mathit{quantity}(\sigma) &= \textstyle\sum_{i=1}^{k} \mathit{quantity}(o_i) \\ p_*(\sigma) &= \textstyle\sum_{i=1}^{k} \mathit{quantity}(o_i) p_*(o_i), \end{aligned} \tag{3}$$

where, as for options, $*$ can be $-$, e, or $+$. (Recall that option prices are *per-unit* – spread prices are for the full quantity $\mathit{quantity}(\sigma)$).

3.4 Targets

In order to determine how hard to bargain for each option under consideration, the reasoner sets **targets** for each option, which are calculated with reference to the other sellers and the options they offer. The target of option o belonging to seller S is intuitively understood to be the maximum price per unit likely to be acceptable for o, and is calculated via a sort of "credible threat" reasoning: It is worth considering o at price p only if there is a completion of o to a spread no more expensive than the best spread available *not* including options belonging to S. This understanding is modified by risk-parameters "-", "e", "+", that capture best-case, average-case and worst-case qualifications of the above clauses.

Formally, for each option o, and some set of potential purchase spreads $\mathcal{M}$, we make the following definitions.

1. The set of **alternatives** to o in $\mathcal{M}$ is the set of those purchase spreads in $\mathcal{M}$ which do not contain any options belonging to the seller S of o:

$$\mathit{alt}_{\mathcal{M}}(o) = \{\sigma \in \mathcal{M} | \, \mathit{options}(\mathit{seller}(o)) \pitchfork \sigma\}.$$

2. The set of **completions** to o in $\mathcal{M}$ is the set of spreads that, with o added, become an acceptable purchase spread:

$$\mathit{comp}_{\mathcal{M}}(o) = \{\sigma | \, \sigma \cup \{o\} \in \mathcal{M}, \mathit{options}(\mathit{seller}(o)) \pitchfork \sigma\}.$$

3. The **target** of o relative to the set of purchase spreads $\mathcal{M}$ is defined for any pair of spread-risk preferences, $r_1, r_2 \in \{-, e, a, +\}$, as

$$\begin{aligned} t^{\mathcal{M}}_{r_1,r_2}(o) = {} & \min(p_{r_1}(\sigma) | \, \sigma \in \mathit{alt}_{\mathcal{M}}(o)) \\ & - \min(p_{r_2}(\sigma) | \, \sigma \in \mathit{comp}_{\mathcal{M}}(o)). \end{aligned} \tag{4}$$

EXAMPLE 1. *Suppose that the options are* $\{a, b, c, d, e, f\}$ *with associated quantities 1,3,2,4,6,2, and suppose that each is associated to a unique seller. If the purchase request quantity* Q *is 5, then the set of all acceptable purchase bundles* $\mathcal{M}$ *is*

$$\mathcal{M} = \{\{a, d\}, \{b, c\}, \{b, f\}, \{a, c, f\}\}.$$

Selecting option a *to calculate targets for, we have*

$$alt_{\mathcal{M}}(a) = \{\{b, c\}, \{b, f\}\},$$

$$comp_{\mathcal{M}}(a) = \{\{d\}, \{c, f\}\}.$$

The target for a *relative to* $\mathcal{M}$*, with expected prices for alternatives and no-risk prices for completions is therefore*

$$\begin{aligned} t^{\mathcal{M}}_{e,+}(a) &= \min(p_e(\sigma) |\, \sigma \in \{\{b, c\}, \{b, f\}\}) \\ &\quad - \min(p_+(\sigma) |\, \sigma \in \{\{d\}, \{c, f\}\}). \\ &= \min\Big(3p_e(b) + 2p_e(c), 3p_e(b) + 2p_e(f))\Big) \\ &\quad - \min\Big(4p_+(d), 2p_+(c) + 2p_+(f)\Big). \end{aligned}$$

3.5 Acceptable Purchase Spreads

The set of acceptable purchase bundles $\mathcal{M}$ in (4) should ideally be the set of *all* possible spreads consistent with the purchase request, i.e.

$$\mathcal{M}_{Q,P} := \{\sigma |\, p_-(\sigma) \le P, quantity(\sigma) = Q\}. \quad (5)$$

Notice that we require the spread's total quantity to be *exactly* Q, so that $\mathcal{M}_{Q,P}$ may be empty. Future implementations may allow flexibility in the purchase request, and hence the set of all acceptable purchase spreads.

When there are several sellers with small feasible quantity steps q_{step}, the set $\mathcal{M}_{Q,P}$ may be too large to reason over, in which case it is necessary to restrict attention to some sub-collection of spreads.

It can be shown that if the price per additional unit is non-increasing with quantity for each seller, then the set of spreads that can minimize total price is given by the extreme points of the convex hull (in quantity space) of the set of all acceptable purchase spreads, $\mathcal{M}_{Q,P}$. This fact, and the intuition that at any given time there will be some seller that is "favourite", and from whom we should like to buy as much of the quantity Q as possible subject to quantity constraints, informed our choice of algorihtm for restricting the set of spreads under consideration.

3.6 Starting and Ending Negotiations

3.6.1 Starting

We assumed that the procurement process begins with the buyer sending out a request for quotes to each seller, in response to which they will each quote an ask for the requested quantity (which is, of course, not always Q, depending on seller constraints). AutONA then has to make a counter offer; the seller counter offers again, and from then on the tactics selected by the reasoner will specify counter offers. This process requires us to specify how AutONA's first bid is generated.

We chose the first bid on an option o to be $0.94p_-(o) + 0.06p_+(o)$, i.e. close to the best one could expect. This choice was made on the assumption that our first bid would almost certainly not succeed, but that a successful transaction would be concluded only after negotiations. If the initial bid were set too high, it would almost certainly be accepted, which could lead (via the construction of beliefs on the basis of historical trade information) to inflation in the price that AutONA would consider reasonable.

3.6.2 Ending

The reasoner controls completion of individual negotiations: AutONA continues trading until the difference between the worst-case and expected case prices is less than a pre-defined (small) proportion, ϵ of the worst-case price:

$$best_+(\mathcal{M}) - best_e(\mathcal{M}) < \epsilon \cdot best_+(\mathcal{M}), \quad (6)$$

where

$$best_r(\mathcal{M}) := \min(p_r(\sigma) |\, \sigma \in \mathcal{M}). \quad (7)$$

4. BIDDING AGENT SPECIFICATION

4.1 Option choice

When negotiating with a direct seller S, there may be many options with respect to which negotiations could proceed. We choose to order the options according to the best expected price amongst acceptable purchase spreads containing them.

1. The **best spread** with respect to risk option r, $\mathbf{B}_r$ is any spread in the maximization set $\mathcal{M}$ such that $p_r(\mathbf{B}_r) = best_r(\mathcal{M})$. We assume that there is an implicit total ordering on spreads which allows us to select $\mathbf{B}_x$ consistently and un-ambiguously.

2. If $\mathbf{B}_e \cap options(S) \neq \emptyset$, then the option o which forms the intersection is the most favoured option for seller S.

3. Otherwise, o is the smallest-quantity option which minimizes the expected price function over spreads containing an option from the given seller:

$$q_o p_e(o) + best_e(\mathcal{M}^c_o) = best_e(\mathcal{M} \setminus \mathcal{M}^a_o). \quad (8)$$

4.2 Tactics

A **tactic** is a rule specifying a new value to offer in response to the thread of negotiation that has so far taken place with a given seller. The tactics used by AutONA are all **Alpha-Beta** tactics, which are specified by two numbers, α and β. A new bid is given with respect to the preceding one, the last ask, and the most recent change to the ask, as

$$\begin{aligned} \text{new bid} = \min\big(\text{old ask},\ &\text{old bid } + \\ &+\ \alpha \times (\text{change in ask}) \\ &+\ \beta \times (\text{ask} - \text{bid})\big). \end{aligned} \quad (9)$$

More specifically, we use two sub-families: pure alpha and pure beta tactics:

- the **fixed alpha tactics** A_j, $j = 0,1,2,3,4$ are the 5 alpha-beta tactics with $\beta = 0$, $\alpha = \{\alpha_0, \frac{1}{2}(1+\alpha_0), \frac{1}{2}, 1, 0\}$ respectively; and
- the **fixed beta tactics** B_j, $j = 0,1,2$ are the 3 alpha-beta tactics with $\alpha = 0$, $\beta = \{0, \beta_{small}, \beta_{big}\}$ respectively;

here $\alpha_0 > 1$, $0 < \beta_{small} < \beta_{big} < 1$ are constants for which the values chosen were 2, $\frac{1}{5}$ and $\frac{1}{2}$, respectively. Note that $A_4 = B_0$.

4.3 Tactic Selection

The choice of which tactic to use with each option o depends on the relative standing of that seller (for that quantity) with respect to the others.

The intuition behind tactic selection is that the value of the expected price relative to the expected-price-alternatives, govern the use of the α parameter; the β parameter is determined by "how far the seller has to go": the normalized difference between the current ask and the expected price.

If the change between the previous and current ask is non-zero, i.e. if the seller has conceded at all since his previous offer, we choose the tactic for option o to be the fixed alpha tactic A_j, with j selected according to the following algorithm:

1. Define[1]

$$\begin{aligned} t_0 &= t_{-,e}(o), \\ t_1 &= \tfrac{1}{2}(t_{-,e}(o) + t_{e,e}(o)), \\ t_2 &= t_{e,e}(o), \\ t_3 &= \tfrac{1}{2}(t_{e,e}(o) + t_{+,e}(o)), \\ t_4 &= t_{+,e}(o). \end{aligned}$$

2. Choose j such that $|t_j - p_e(o)|$ is minimized.

The intuition is that if the expected price of the option $p_e(o)$ is close to t_0, for example, then it is expected to be comparable to the *best-case* for its best possible alternative, and hence is valuable, so that we should concede in order to keep the seller happy; if $p_e(o)$ is close to t_4 then we expect o to be comparable (when completed) to the *worst-case* alternative: hence it is the seller's responsibility to concede towards us if he wants to be considered seriously.

If the change between the previous and current ask is zero, the current tactic for option o is chosen to be the fixed beta tactic B_j according to the following algorithm:

1. Let

$$s = \begin{cases} \dfrac{p_+(o) - p_e(o)}{p_+(o) - p_-(o)} & \textit{if } p_+(o) > p_-(o) \\ 0 & \textit{otherwise} \end{cases}$$

2. If $s < \frac{1}{4}$, choose $j = 0$;
3. if $\frac{1}{4} \le s < \frac{3}{4}$, choose $j = 1$; and
4. if $\frac{3}{4} \le s$, choose $j = 2$.

[1]Recall that the most suitable option o to negotiate over is chosen using (8)

5. EXPERIMENTS

5.1 Overview

Since AutONA is designed for real procurement applications, it is essential to understand its performance before any deployment in real business environments. More specifically, there are three key questions that we seek to answer:

1. Are the negotiation algorithms on which AutONA is based exploitable by clever sellers – is it possible for sellers to detect that they are bidding against a "machine" when negotiating with AutONA?
2. How well does AutONA perform in different trading environments. The goal here is to identify, as much as possible, a relationship between specific features of the purchasing environment and the performance of AutONA.
3. How well does AutONA perform compared to human traders in similar circumstances.

A sequence of laboratory experiments was conducted to perform the tests, following standard experimental economics methodology. The subjects were given accurate information about the game, in particular, how their actual monetary rewards depended on their aggregate performance over the course of the session. Experimental anonymity with respect to roles and payment was preserved and no deception was used. Experiments with all human subjects were conducted, to serve as benchmarks to measurements of AutONA's effectiveness. The same experiments were then run again with AutONA replacing one of the human buyers.

5.2 The Experimental Model

The goal of the experimental design phase was to capture important aspects of the true procurement environments in which AutONA is intended to participate. To remove any conscious or unconscious biases in the experimental design, very little information about how AutONA works was provided to the experimenter who designed the experiments. The primary information used to construct the experiments came from the HP procurement organization, which provided detailed descriptions of, and data from, their procurement operations.

Due to business and scientific considerations, we chose to examine a scenario similar to that of DRAM procurement. Important aspects of this scenario, such as the small numbers of buyers and sellers, their relative market power, the inflexibility of short-term capability and the possibilities of shortages, were included in the design of the experiments. Some complications, such as inventory carry-over and timing of delivery, were ignored.

The experimental model has three central components: the buyers, the sellers, and the negotiation process.

5.2.1 The buyers

Each buyer's objective is to procure a certain amount Q, which will be referred to as the **target quantity**, of a single homogenous commodity. Buyers are rewarded according to the following formula:has a linear download sloping demand function with a cut-off point at Q and an additional bonus if he procures an amount not less than Q. Thus his demand function is

$$Demand(q) = \begin{cases} a - b \times q & \text{if } q < Q, \\ a - b \times Q + bonus & \text{if } q = Q, \\ 0 & \text{if } q > Q, \end{cases} \tag{10}$$

where a and b are positive constants obeying the constraint $a - b.Q > 0$, so that buyers are always incented to buy no less than Q goods. This demand function gives rise to the reward function,

$$Reward(q) = \begin{cases} a.q - b.q^2 & \text{if } q < Q, \\ a.Q - b.Q^2 + bonus & \text{if } q \geq Q. \end{cases} \tag{11}$$

A player's total payoff for purchasing quantity q is given by $Reward(q) - C(q)$, where $C(q)$ is what the buyers pay for the goods. This payoff function provides no incentive to procure any amount more than Q, which is similar to the situation in which a buyer is trying to procure enough DRAM to manufacture computers for a specific fixed quantity contract with a downstream reseller.

5.2.2 The sellers

Each seller has a cost function $K(q)$ where q is the quantity they sell. Their payoff function is $C(q) - K(q)$, where $C(q)$ is what the buyer(s) pay him. The cost function $K(q)$ is assumed to have a fixed cost (F), a variable cost (c) and a capacity (k):

$$K(q) = \begin{cases} F + c \times q & \text{if } q \leq k \\ F + c \times k + 10c \times (q - k) & \text{if } q > k \end{cases} \tag{12}$$

It is assumed that when a seller tries to sell above capacity, he has to incur 10 times the normal costs. This is probably more realistic than assuming that it is impossible to sell more than capacity, since sellers can, if they wish, always procure goods on the spot market to cover short-falls in supply. The net result of the extra factor of 10 is to make production beyond capacity expensive but not impossible, which is realistic in the DRAM environment.

Sellers were always played by human subjects.

5.2.3 Supply & Demand Calibration

There are only a few major players in the DRAM market: Four major suppliers cover roughly 70–80% of the market. The market is a bit more fragmented on the buyer side, but there are only a few players (such as HP, IBM and Dell) that have the market power to negotiate substantial deals with the major sellers.

The experiment was set up with four homogenous buyers and four heterogeneous sellers. The sellers' capabilities reflected true market share in the DRAM market. The total market capacity is normalized to 1000. Both capacities and cost functions were fixed throughout the experiment, so that the only uncertainty exists on the demand side. The demand parameters were set up so that the market equilibrium quantity was the smaller of either the total capacity or the totals of all the buyer's target quantities, Q. This allowed us to measure the effectiveness of a buyer by simply looking at the amount he has procured.

Buyers' target quantities were generated by a random process consistent with actual demand fluctuations. The HP Procurement Risk Organization has been analyzing the distribution of DRAM demand over the years. A normalized form of this distribution was used in the experiment.

Two supply and demand scenarios were considered. In the first scenario, the average total target quantity was slightly higher than the total capacity. However, demand was generated according to a log-normal distribution, so the chance of a shortage (total target more than total capacity) was roughly 50%. In the second scenario, the total target quantity was **always** greater than the total capacity. Thus, every trading period is in shortage, although it is uncertain of the degree of the shortage.

5.2.4 The Negotiation Process

The negotiation process was modeled as a round-based multiple 1-1 negotiation game. In each round, buyers and sellers take turns to make offers consisting of a price and a quantity, with no requirement to improve on previous offers. Each offer is directed at only **one** player on the other side of the market, and are private information between the buyer-seller pair. In each round, a player can make a new offer, accept the offer on the table, or stop the negotiation.

A limited form of cheap talk was allowed: A player could send a message consisting of a price and a quantity to anyone on the other side of the market, with no commitments: There were no consequences of this communication other than information exchanges.

A time cost was introduced to provide incentives for timely negotiation. The first 8 rounds of negotiation were free, but after that each round cost a fixed amount to any player who has an active offer on the table. The trading period terminates if either side of the market (buyers or sellers) has no active offers. This process does not guarantee termination, but in practice negotiations usually terminated in about 10-14 rounds.

6. CUSTOMIZATION

AutONA was designed before the experiments were. The design criteria behind AutONA were for it to be applicable to a wide range of procurement situations and exhibit flexibility through customization. To play the game, AutONA needed to be customized; this section covers some of the customization choices that we made, and discusses the impact they had on the experiments results.

Customization can be seen as consisting of two components: a set of parameter values for certain control parameters; and heuristics and rules relating to the way in which data are fed to and from the system by an operator.

6.1 Customization Parameters

6.1.1 Termination condition

The parmaeter ϵ (see §3.6) sets the point at which AutONA will recommend to the buyer that a price is accepted and that negotiation with the seller over a particular quantity should be concluded. We decided to set epsilon to 5%, meaning that AutONA will recommend to close a deal when the price that the seller offers is within 5% of the price it expects for that seller at that quantity.

6.1.2 History scale factor

AutONA was pre-loaded with a history of previous negotiation with the various sellers and with the market price for previous rounds. Because of the accelerated time in the experiments, we had the freedom of placing the actual periods in time at our will. We decided that the history scale factor would be set so that the all the data of the previous experiments would count for about a half of the data of the current experiment. Between periods in the same experiment the time difference was considered to be negligible.

6.2 Heuristics and rules

6.2.1 Deal definition

During the deal definition phase, the user operating AutONA sets values for parameters such as quantity required and price ceiling. The obvious choice to make was to define the quantity to procure as the target quantity of the game (section 5.2.1). For price ceiling, we use a value that is equivalent to the reward that AutONA would receive for procuring the target quantity, as defined in (12). With these settings, we ensure that AutONA will not form deals that will encur greater costs than its maximum reward.

6.2.2 Seller selection

In the seller selection phase, the quantities q^S_{min}, q^S_{max} and q^S_{step} are defined for each of the sellers. q^S_{max} is one of the most important parameters of the game, as it represents the capacity that sellers have available. But that piece of information was not available to AutONA (nor to any other buyer-side player). Nor had AutONA been designed to elicit that knowledge as the game progressed. The values of q^S_{max} for the sellers determine how AutONA builds its spreads. For the first experiment we used a heuristic that would have AutONA build spreads that divide the required quantity nearly equally among the four sellers. The rule was to set q^S_{max} for each seller at 27% of the target quantity. Having observed that AutONA was not so successful in procuring the target quantity (see discussion of the second result), later on we decided to have AutONA build spreads where one of the sellers was getting the biggest share of its target quantity. We did that by setting the maximum quantity available from each of the seller to be 75% of the target quantity.

6.2.3 Negotiation

The protocol used in the experiments prescribed that the buyer put in the first offer, whereas the protocol that AutONA had been designed to play a game where the seller would submit the first offer, for a quantity requested by AutONA (see 3.6.1). To comply with the rules of the game, we had to define a heuristic for the first offer that wouldn't suggested through the AutONA user interface. To play fairly, we needed to bind the heuristic to information that was available to AutONA. Our decision was that the first offer would be submitted as a percentage of the price that AutONA expected for negotiation from a given seller ($p_e(S)$). In the first experiment we guessed that 90% might be a fair value. Having observed that in the second experiment AutONA procured prices with a spread of 93.8% to 105.8% on the mean, we set it to be 94% for the fourth experiment. In both cases AutONA exhibited a less than brilliant performance in procuring the target quantity (see discussion on the second result). To improve things in the last experiment, we decided that the first bid was to be submitted at exactly $p_e(S)$, resulting in a better performance of AutONA quantity-wise.

6.2.4 Recomputing spreads

AutONA was designed to attempt to impose quantity on the suppliers, through the RFQ process. The game would go smoothly if suppliers did accept the quantities by responding with a counteroffer on the same quantity. We observed that this was not the case during the experiments. Whenever the seller proposes a different negotiation quantity, the AutONA operator faces a decision on whether to proceed negotiating over quantities appearing as options in AutONA spreads or restarting AutONA to recompute the spreads. In the spirit of making the experiments as repeatable as possible, we needed to put the operator in condition to use deliberation as little as possible. So we defined a rule that if none of the sellers responded to the quantity suggested, AutONA should be restarted by the seventh round. Likewise, AutONA needed to be restarted if sellers would not respond even after a deal has just been struck. In that case, the operator shoud restart AutONA subtracting the deal quantities achieved so far from the game target. Restarting AutONA is not ideal, but in both cases gives us the advantage that q^S_{max} can be set using information taken from offers that sellers have made. This tactic is useful in reducing the number of rounds required to achieve deals, thus avoiding round costs. More importantly it is useful to actually secure the quantity that was needed especially in cases of supply shortage. To respond to the problem that AutONA was having in procuring target quantity, in the fifth experiment we modified the rules so as to restart and recompute the spreads after the fifth round, using seller information on quantities to set q^S_{max}. A further rule was that after bundle recomputation, AutONA operator would accept standing seller offers that would fall within the percentage of $p_e(o)$ that was set to determine the first offer.

7. RESULTS

A total of five experiments were conducted, each with eight players (four buyers and four sellers). Two of the five experiments were all human experiment. In the rest, AutONA played the role of **one** buyer. In the fifth experiment, a modified version of AutONA was used, to counteract behavioural traits discovered in the first four experiments, discussed below. In each case, AutONA was provided with data from previous experiments as simulations of market inputs.

Experiment	Supply/Demand treatment	Players
1	Random shortage	All Human
2	Random shortage	AutONA
3	All shortage	All Human
4	All shortage	AutONA
5	All shortage	Modified AutONA

Table 1: Summary of the experiments

Our first result addresses the first experimental question.

RESULT 1. *AutONA passed a limited version of the Turing test. There is no obvious method for the human subjects to exploit AutONA.*

Support: In the beginning of each experiment that involved AutONA we announced that one of the players would be played by a robot. At the end of the experiment, we informally quizzed all the subjects as to the identity of the robot. The answers we obtained were random. There is no evidence that human subjects can identify which player was played by AutONA. Furthermore, we can conclude that no subjects have found and used any logical loop-holes in AutONA's algorithms.

The other two experimental questions are concerned with performance. There are two primary measures we use to benchmark the performance of a buyer: price, and quantity with respect to target. Payoff is not relevant for the reason that AutONA is not designed to optimize the experimental payoff, and indeed is not even aware of the existence of a payoff function.

All things being equal, quantity with respect to the target is the most important measure. Table 2 summarizes buyers' performance as measured by the quantity they procured as a percentage of their targets.

Since the buyers are homogenous, their performance should be roughly the same if all of them are playing rationally. In all the experiments with two exceptions, human subjects have procured rough a similar amount (compared within experiment) with respect to their targets. The two exceptions are buyer 4 in experiment 1 and buyer 3 in experiment 5. Some variations are expected since humans do not negotiate equally. Experiment 1 seems to show a larger variation, which can probably be explained by inexperienced subjects.

The "Market" column in Table 2 lists the total quantities procured in the market (by the 4 buyers) as a percentage of the total capacity. Experiments 1 and 2 have the same supply and demand parameters, while experiments 3, 4 and 5 have another set of parameters. It is clear that aggregate results are consistent across experiments. The percentage bought with respect to target quantity is within 2 percentage points across buyers for each experiment. This is strong evidence that experimental results were repeatable and human subjects understood their instructions and responded well to monetary incentives.

This brings us to our second major result.

RESULT 2. *The original AutONA was procuring substantially less, relative to its target quantity, than human buyers. This is particularly significant when there is a shortage.*

Support: As can be seen from Table 2, the quantity procured by AutONA is substantially lower than that of human players in experiment 2 and 4. Table 2 also reports a summary of experiment 2 with only the periods in shortage. In those periods, AutONA was procuring even less, at 53% of target, which is consistent with the results in the "all shortage" experiment (experiment 4).

On the basis of experiments 1 through 4, it is clear that AutONA has a severe behavioral bias. Roughly speaking, it is not aggressive enough in completing negotiations with successful transactions: it spends too long negotiating, and sellers go elsewhere. This problem is exacerbated by a shortage. Human buyers seem to be able to recognize the importance of grabbing supplies as fast and as aggressively as they can, while AutONA does not.

RESULT 3. *AutONA received lower prices than the human players.*

Support: From Table 3, we see in experiment 2 and 4 that AutONA has the lowest average price.

Experiment	Buyer 1	Buyer 2	Buyer 3	Buyer 4 / AutONA
1	\$167	\$161	\$183	\$167
2	\$183	\$172	\$172	**\$163**
2 (shortage)	\$192	\$174	\$170	**\$163**
3	\$191	\$191	\$181	\$189
4	\$191	\$210	\$182	**\$182**
5	\$184	\$185	\$184	**\$185**

Table 3: Summary of buyers' performance as measured by the average price of transactions.

It seems that AutONA was trading off prices with quantities: one reason why AutONA procured significantly less than expected was its strong stance on price. This bias towards aggressive price negotiation is due to a design assumption: that there was no competition against other buyers, and hence that time is much less of an issue. When time is not an important issue, there is no reason to negotiate speedily, except to meet purchasing deadlines, and so it is advisable to bargain hard. The DRAM procurement game, especially when there was a shortage, definitely involved competition between buyers, and although AutONA does well on price, it does poorly on quantity.

Bearing in mind the relative importance of the two performance measures, we decided to modify the behaviour of AutONA. To begin with, instead of opening on option o with a bid of $0.94p_-(o) + 0.06p_+(o)$ (see §3.6.1), we re-configured AutONA to open with $p_e(o)$. This (unsurprisingly) led to many negotiations concluding immediately, and on average reduced the duration of negotiations considerably, at the expense of leading to more expensive trades. In addition, an adjustment of the seller constraints (encoded in q^S_{min} and q^S_{max}) such that the largest component in each spread took up about 75% of Q, seemed to result in superior performance. Both of these modifications were at the configuration level. We anticipate that each specific negotiating

Experiment	Buyer 1	Buyer 2	Buyer 3	Buyer 4 / AutONA	Market
1	100%	93%	87%	72%	87%
2	89%	92%	94%	**67%**	85%
2 (shortage)	96%	91%	91%	**53%**	82%
3	71%	85%	71%	75%	76%
4	89%	82%	87%	**52%**	77%
5	83%	80%	65%	**80%**	77%

Table 2: Summary of buyers' performance as measured by the percentage of the target quantity purchased. Results in bold are for AutONA.

environment will place different requirements on AutONA and hence will lead to different configurations. The fifth experiment was run with this modified version.

RESULT 4. *A modified version of AutONA performed significantly better on quantity, and not as well on price – its payoff of price and quantity was similar to that of humans.*

Support: See experiment 5 in Tables 2 and 3.

8. CONCLUSIONS

In this paper we have presented a system, AutONA, for conducting multiple simultaneous 1-1 negotiations over price and quantity. The use of competition between sellers to guide negotiation tactics is key.

We have implemented this system, and conducted human trials to evaluate it on the basis of its ability to negotiate "reasonably", and on its performance with respect to a trading game that was designed independently of the system itself.

We find that AutONA passes a limited version of the Turing test: The experiments did not reveal any obvious exploitation that a human trade can use against AutONA. On the other hand, AutONA in its original configuration exhibited significantly different aggregate behavior from human traders; it was less aggressive on quantity and more aggressive on price – a behavioral bias that is non-desirable in the HP DRAM procurement context in which it was evaluated. Subsequently, AutONA was modified, and the modified version behaved more in line with human traders in the experiments, but does not exhibit any significant advantage over human traders.

9. ADDITIONAL AUTHORS

Mike Yearworth (HP Labs, email: `Mike.Yearworth@hp.com`), and Nir Vulkan (The Saïd Business School, Oxford, email: `Nir.Vulkan@sbs.ox.ac.uk`).

10. REFERENCES

[1] P. Faratin, C. Sierra, and N. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 3-4(24):159–182, 1998.

[2] P. Faratin, C. Sierra, and N. Jennings. Using similarity criteria to make trade-offs. In *Proc. 4th Int. Conf on Multi-Agent Systems*, pages 119 – 126, 2000.

[3] Abhinay Muthoo. *Bargaining Theory with Applications.* Cambridge University Press, 1999.

[4] A. Sathi. *Constraint Directed Negotiation of resource allocations*, pages 163–195. Morgan Kaufman, 1989.

[5] K. Sycara. *Multi-agent compromise via negotiation*, pages 119–139. Morgan Kaufman, 1989.