



Centrum voor Wiskunde en Informatica

REPORTRAPPORT

SEN

Software Engineering



Software ENgineering

Automated Negotiation and Bundling of Information
Goods

D.J.A. Somefun, E.H. Gerding, S.M. Bohte,
J.A. La Poutré

REPORT SEN-E0407 AUGUST 2004

CWI is the National Research Institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organization for Scientific Research (NWO).

CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Copyright © 2004, Stichting Centrum voor Wiskunde en Informatica

P.O. Box 94079, 1090 GB Amsterdam (NL)

Kruislaan 413, 1098 SJ Amsterdam (NL)

Telephone +31 20 592 9333

Telefax +31 20 592 4199

ISSN 1386-369X

Automated Negotiation and Bundling of Information Goods

ABSTRACT

In this paper, we present a novel system for selling bundles of news items. Through the system, customers bargain with the seller over the price and quality of the delivered goods. The advantage of the developed system is that it allows for a high degree of flexibility in the price, quality, and content of the offered bundles. The price, quality, and content of the delivered goods may, for example, differ based on daily dynamics and personal interests of customers. Autonomous "software agents" execute the negotiation on behalf of the users of the system. To perform the actual negotiation these agents make use of bargaining strategies. We decompose bargaining strategies into concession strategies and Pareto efficient search strategies. Additionally, we introduce the orthogonal and orthogonal-DF strategy: two Pareto search strategies. We show through computer experiments that the use of these Pareto search strategies will result in very efficient bargaining outcomes. Moreover, the system is set up such that it is actually in the best interest of the customer to have their agent adhere to this approach of disentangling the bargaining strategy.

2000 Mathematics Subject Classification: 91B26

1998 ACM Computing Classification System: J4

Keywords and Phrases: negotiation; bundling

Automated Negotiation and Bundling of Information Goods

D J A Somefun¹, E H Gerding¹, S Bohte¹, and J A La Poutré^{1,2}

¹ Center for Mathematics and Computer Science (CWI), P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

² Eindhoven University of Technology, School of Technology Management, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
{koye, Enrico.Gerding, S.M.Bohte, hlp}@cwi.nl

Abstract. In this paper, we present a novel system for selling bundles of news items. Through the system, customers bargain with the seller over the price and quality of the delivered goods. The advantage of the developed system is that it allows for a high degree of flexibility in the price, quality, and content of the offered bundles. The price, quality, and content of the delivered goods may, for example, differ based on daily dynamics and personal interests of customers. Autonomous “software agents” execute the negotiation on behalf of the users of the system. To perform the actual negotiation these agents make use of bargaining strategies. We decompose bargaining strategies into concession strategies and Pareto efficient search strategies. Additionally, we introduce the orthogonal and orthogonal-DF strategy: two Pareto search strategies. We show through computer experiments that the use of these Pareto search strategies will result in very efficient bargaining outcomes. Moreover, the system is set up such that it is actually in the best interest of the customer to have their agent adhere to this approach of disentangling the bargaining strategy.

1 Introduction

Personalization of information goods becomes more and more a key component of a successful electronic business strategy [1]. The challenge is to develop systems that can deliver a high level of personalization combined with, whenever possible, a high adaptability to changing circumstances. In this paper we develop a system which can attain these properties through the manner in which it sells information goods.

We consider the novel approach of selling bundles of news items through a system that allows for bargaining over the price and quality of the delivered goods. Bundling of information goods has many potential benefits including complementarities among the bundle components, and sorting consumers according to their valuation (cf. [2] and the references therein). The advantage of the developed system is that it allows for a high degree of flexibility in the price, quality, and content of the offered bundles. The price, quality, and content of the

delivered goods may, for example, differ based on daily dynamics and personal interest of customers.

The system as developed is capable of taking into account business related constraints. More specifically, it tries to ensure that customers perceive the bargaining outcomes as being “fair” by having customers end up with equivalent offers whenever that seems fair. Partly because of this fairness constraint the actual bargaining process is not really one-to-one bargaining between seller and customer but instead is one-to-many (i.e., between seller and customers).

To accelerate the negotiation process, customers can initiate concurrent negotiation threads for the same bundle with differences in the quality of the delivered bundles. The thread in which the agreement is reached first determines the final bargaining outcome.

In the developed system, autonomous “software agents” perform (part of) the negotiation on behalf of the users of the system. These software agents bargain over a multi-issue price (the price is actually a tariff with a fixed and variable component).

To enable efficient multi-issue bargaining outcomes, we decompose the bargaining strategies into concession strategies and Pareto search strategies. Additionally, we introduce the orthogonal and orthogonal-DF strategy: two Pareto search strategies. We show through computer experiments that the respective use of these two Pareto search strategies by the two bargainers will result in very efficient bargaining outcomes (i.e., these outcomes closely approximate Pareto-efficient bargaining solutions).

In the system the seller agent uses a Pareto search strategy (i.e., the orthogonal-DF) combined with a concession strategy. Although the customer is free to choose other bargaining strategies, the system is set up such that it is actually in the best interest of the customer to have their agent also use a Pareto search strategy (i.e., the orthogonal strategy) combined with a concession strategy.

In Section 2 we discuss the developed system at a more conceptual level. In Section 3 we discuss the customer and seller agent in greater detail. Furthermore, we discuss the type of bargaining strategies these agents use. In Section 4 we study in greater detail the Pareto search aspects of bargaining. Through computer experiments we investigate the efficiency of the introduced bargaining approach. (Note that for this purpose it is not necessary to simulate the entire system as developed, it suffices to consider one-to-one bargaining only.) In Section 5 we discuss the results of the paper and relate the paper to the relevant literature. Conclusions follow in Section 6.

2 A System for Selling Information Goods

2.1 Problem Statement

The goal is to develop a system for the sales of bundles of news items where customers bargain over the price and quality of the delivered goods. The negotiated contract applies to a fixed time interval, which is typically a short period

of time, e.g., a single day. There are roughly three possibilities for implementing the starting time of the negotiation process: customers can negotiate a contract before the news arises, after the news arises, or while the news arises. The system is set up in such a way that all three possibilities can be implemented.

The value customers attach to news items may fluctuate heavily due to daily dynamics. Moreover, there may be wide differences in personal interests of customers. The advantage of the developed system is that it allows for a high degree of flexibility in the price, quality, and content of the offered bundles. The price, quality, and content of the delivered goods may, for example, differ based on daily dynamics and personal interest of customers.

2.2 Bundles of Information Goods

The system sells bundles of news items which become available during a predefined and fixed time interval (e.g., a day). Within the system, prices vary based on the content and “quality of service” of the bundle. A *bundle content* defines which types of news categories the bundle contains. The system distinguishes between k categories. Within a category the system distinguishes between two quality of service levels: i.e., a customer pays a fixed price for either receiving headlines or complete news articles. In the former case we speak of a category with low quality of service, whereas in the latter case we speak of a category with high quality of service. Moreover, with the quality of service (or just quality) of a bundle we actually mean the quality of service specified per category.

A customer bargains with the seller over the bundle tariff. The negotiated tariff is a two-part tariff with a fixed and variable price. The fixed price (p_f) is the price a customer pays for receiving the bundle content with the specified quality of service. Moreover, the variable price (p_v) is the price the customer pays for reading a full article whenever the quality of service only specifies delivery of the article headline.

Consider, for example, the bundle content religion, culture, and politics, where the category religion has a high quality of service and the other two have a low quality of service. Then the customer pays a fixed price for receiving all the full articles in the category religion and only the headlines of all the articles which do not belong to the category religion but do belong to the other two categories. Moreover, the variable price is the price the customer pays whenever she wants to read the full article of a news item that belongs to the categories culture or politics (and does not belong to the category religion).

2.3 Bargaining with Software Agents

We employ the paradigm of “software agents,” where pieces of autonomous software perform (part of) the negotiating on behalf of the users of the system. Customers and seller instruct their agent through a user interface (UI). The agents conduct the actual negotiation. Figure 1 depicts, at a high abstraction level, the bargaining process between a customer and the seller.

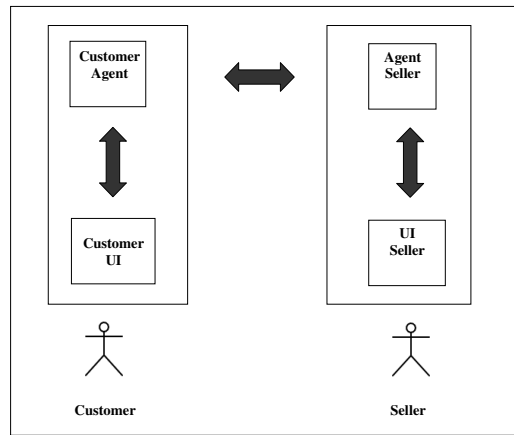


Fig. 1. The one-to-one bargaining process

2.4 Bargaining Process

Bargaining occurs in an alternating exchange of offers and counter offers, typically initiated by the customer. An offer specifies the fixed price, the variable price (uniform for all low quality of service categories), the bundle content, and also the desired quality of service of the information for each category separately. The bargaining process continues until an agreement is reached or one of the bargainers terminates the process. Based on this bargaining process, figure 2 draws the bargaining protocol the customer agents and seller agent use to do the actual bargaining.

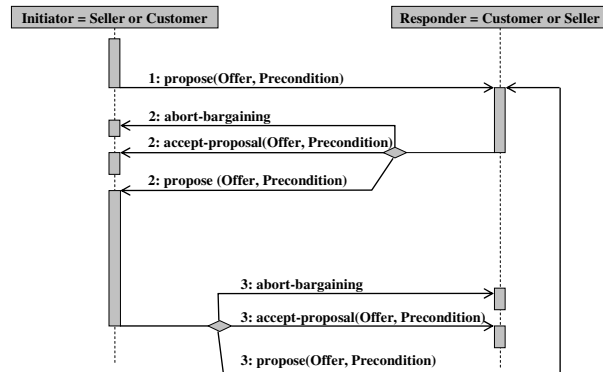


Fig. 2. The agents bargaining protocol

To accelerate the negotiation process, we allow concurrent negotiation threads for the same bundle content with different quality of service. The customer can therefore submit several offers at the same time. In order to discern between threads, each thread must have a different quality configuration for the categories where a quality configuration specifies the quality of service for all the bundle categories. The seller can only respond by varying the fixed and variable price. The thread in which the agreement is reached first determines the tariff and quality configuration for the desired categories.

2.5 Fairness & One-to-Many Bargaining

A possible drawback of bargaining is that two customers may end up paying a substantially different price for very similar bundles. Customers may perceive this as unfair. This is an important concern for the seller, since customers may become dissatisfied or stop using the system altogether.

In the system a notion of fairness is incorporated into the bargaining strategy the seller agent uses. More specifically, within a limited timeframe the seller agent makes equivalent offers to customers who are interested in identical bundles. For offers that specify identical bundles, the actual tariff may still differ from customer to customer. Fairness, however, ensures that the expected revenue of these tariffs is identical for all (counter) offers submitted by the seller agent; the expected revenue (R) of a tariff (p_f, p_v) for a particular bundle is defined as follows:

$$R = p_f + p_v \cdot \rho, \tag{1}$$

where ρ denotes the expected number of articles read in the low quality of service categories (for the average customer). The expected revenue can, however, vary through time. The offer equivalence therefore only holds within a limited time frame. Note that beside “fairness” also other business side-constraints may be implemented. The actual way in which side-constraints, such as fairness, are implemented may be important because it can alter the strategic behavior of customers. It is however beyond the scope of the paper to discuss these issues.

The actual bargaining process between seller and customers is not really one-to-one bargaining between seller and customer but instead is one-to-many. On the one hand, the seller can use his experience in other ongoing bargaining processes between customers to adjust his bargaining strategy; under true one-to-one bargaining the bargaining strategy only depends on the moves of the direct opponent. On the other hand, fairness and/or other side-constraints limit the bargaining options of the seller. These limitations do not apply under true one-to-one bargaining. Figure 3 depicts the one-to-many bargaining process and the possibility of parallel negotiation threads between a customer and the seller.

3 Agents & Bargaining

3.1 Agents

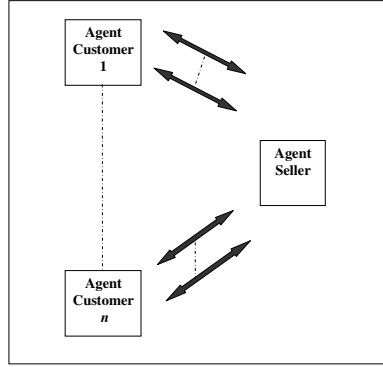


Fig. 3. The one-to-many bargaining with parallel threads

Seller agent Bargaining with a customer is done based on the seller agent's desired aspiration level expressed in expected utils. We define the expected utility u of the seller agent as the expected revenue from selling a bundle given p_f and p_v (see also equation 1), i.e., $u(p_f, p_v) = p_f + p_v \cdot \rho$. We assume that all costs are sunk, i.e., there are no transaction costs. Recall from section 2.5 that ρ denotes the number of expected articles read in the low quality of service categories. The agent can assess the expected number of articles read based on, for example, aggregate past sales data.

Due to the fairness constraint, the seller agent cannot charge different prices to different individuals for identical bundles (within the same time frame). The seller agent can indirectly discriminate between different customers, however, based on differences in the preferred bundle content and/or quality of service. In the system, the seller can discriminate by varying for example the desired expected utility for different combinations of quality of service and bundle content.

Customer Agent The customer agent acts on behalf of the customer. The customer can indicate her preferences by specifying, for each information category she is interested in, the amount of articles she expects to read. The customer can furthermore select between several negotiation strategies to be used by the agent and specify a maximum budget b_{max} . The budget provides the agent with a mandate for the negotiation; the total expected costs should not exceed b_{max} .

Given a tariff (p_f, p_v) for a particular bundle, the customer's expected utility is defined as $u(p_f, p_v) = b_{max} - (p_f + p_v \cdot \rho)$. The second part of the equation is identical to seller's expected revenue (see equation (1)). However, ρ , the number of expected articles read in the low quality of service categories, refers now to the expectation of the individual customer. Given the utility the agent is able to

translate the customer's preferences into offers and also respond to the seller's offers as to maximize the expected utility.

The negotiation protocol allows for multiple negotiation threads for the same bundle content (see Section 2.4). Given a bundle content with n categories, in principle 2^n threads are possible (by varying the quality of service for each category). The customer agent, however, selects only a limited number of combinations based on the customer's preferences, to reduce the amount of communication. In the current system the customer agent initiates $n+1$ threads. In the first thread the quality of service for all categories is set to low. In the second thread, only the quality of service for the category with the highest expected articles read is set to high. In the third thread, this is done for the two categories with the first and second highest expected articles read, and so on. Within a thread, a fixed price and a variable price are negotiated.

3.2 Decomposing Bargaining Strategy

The customer agents and seller agent contain various bargaining strategies to do the actual bargaining over the two-part tariff. These strategies make use of the notion of a utility function to represent the bargainers' preferences for the various tariffs. We decompose bargaining strategies into concession strategies and "Pareto search" strategies.

Concession strategies determine what the desired utility level of an offer will be given a particular sequence of offers and counter offers. Algorithms that implement Pareto search strategies determine, given a particular utility level and a particular history of offers and counter offers, what the multi-issue offer will be, i.e., the fixed price p_f and the variable price p_v of the two part tariff. In terms of a multivariable utility function a counter offer entails both a movement off the iso-utility line and a movement along the iso-utility line. (Given a specified utility level, an iso-utility line specifies all the p_f and p_v points which generate identical utility.) Concession strategies determine the movement of an iso-utility line; Pareto search strategies determine the movement along an iso-utility line.

Pareto search strategies aim at reaching agreement as soon as the respective "concession" strategies permit this. Therefore it may be good for both parties to use it. In more economic terms a negotiated tariff is called Pareto efficient if it is impossible to change the tariff without making one of the bargainers worse off, i.e., one of the bargainers will always attach a lower (or equal) utility to the adjusted tariff. From a system design perspective Pareto efficiency of the negotiated tariffs is clearly desirable.

In Section 3.3 we introduce a particular class of Pareto search algorithms. The experiments in Section 4 show that if the seller agent uses this Pareto search algorithm and customer agents use a similar Pareto search algorithm, then the bargaining outcome will closely approximate a Pareto-efficient solution given a wide variety of concession strategies.

In the system the seller agent uses an instance of the Pareto search algorithms combined with a concession strategy. Although a customer is free to select other bargaining strategies the system is set up such that it is actually in the best

interest of customers to have their agents use Pareto search strategies combined with a concession strategy. We elaborate on this issue in the discussion in Section 5.

3.3 Orthogonal Strategy & DF

Both customer agent and seller agent may use— what we call— an orthogonal strategy as the Pareto search algorithm. This strategy is probably best explained through an example. Suppose, the customer (with whom the seller bargains over the tariff) placed the t^{th} offer of $(p_f(t), p_v(t))$. Moreover, the seller's concession strategy dictates an *aspiration level* of $U_s(t + 1)$: i.e., in utils the (counter) offer should be worth $U_s(t + 1)$. Based on this information the orthogonal strategy determines $(p_f(t + 1), p_v(t + 1))$, the counter offer of the seller, by choosing a (p_f, p_v) combination that generates $U_s(t + 1)$ utils and lies closest (measured in Euclidean distance) to the point $(p_f(t), p_v(t))$. Figure 4 gives the graphical representation of the orthogonal strategy.




Fig. 4. Example of Orthogonal Strategy

The use of the orthogonal strategy by both parties results in a mapping f from a bargainer's aspiration level at t to the aspiration level at $t + 2$. Given convex preferences (cf. [3]) and fixed aspiration levels the mapping f can be shown to satisfy the Lipschitz condition $\|f(x) - f(y)\| \leq \|x - y\|$ for all x and y in the domain of f .³ Thus, given fixed aspiration levels and convex preferences, the orthogonal strategy does imply that consecutive offers do not diverge. (Note, that convex preferences do not rule out changing preferences it only means that at some point in time—i.e., during the bargaining process— the preference relation is convex.) Figure 5 illustrates the use of the orthogonal strategy by both parties

³ The proof is a straightforward application of convex analysis (cf. [4]) given that without loss of generality we can assume that the preferences are bounded. That is, negative and extremely high $(p_f(t), p_v(t))$ combinations can be discarded, without loss of generality.

for the case of tangent iso-utility lines. It draws a sequence of two offers and counter offers with convex preferences and a fixed aspiration level. (U_s and U_c denote the iso-utility lines of the seller and customer, respectively.) The figure illustrates, for instance, how the customer’s offer at time 1 (with aspiration level $U_c(1) = U_c$) is transformed into an offer at time 3 (with aspiration level $U_c(3) = U_c$).




Fig. 5. Sequence of two offers and counter offers with fixed aspiration levels and convex preferences, where (p_f^*, p_v^*) denotes a Pareto-efficient tariff.

The use of just the orthogonal strategy by both parties may lead to very slow convergence to Pareto-efficient bargaining outcomes. To speed up the convergence process we can add an amplifying mechanism to the orthogonal strategy. As the amplifying mechanism we use the derivative follower with adaptive step-size (ADF). (Henceforth we will call this the orthogonal-DF.)

The derivative follower (DF) is a local search algorithm (cf. [5]). It adjust the variable price p_v found by the orthogonal strategy by either subtracting or adding δ to it depending on the result of the previous two adjustments, where δ is called the step-size of the DF. Consequently also the fixed price p_f changes because the adjusted offer still needs to generate the same utility level (specified by the concession strategy). The difference between ADF and DF is that the step-size δ becomes adaptive [6, 7]. We use the ADF proposed by [7]. Intuitively, the idea is to increment the step-size relatively far away from a Pareto-efficient solution and decrement it in the vicinity of a Pareto-efficient solution. Consequently, a quicker and more accurate search of the solution space becomes possible. Algorithm 1 specifies the orthogonal-DF in greater detail and figure 6 illustrates the use of the orthogonal-DF by the seller (the customer uses the orthogonal strategy only).

4 Experimental Setup & Results

The previous sections outlined the general system for selling bundles of news items to several customers through negotiation. As discussed in Section 3.2, negotiation essentially consists of two strategic aspects: the concession of the agents




Fig. 6. Sequence of two offers and counter offers with fixed aspiration levels where the seller uses the orthogonal-DF and the customer only uses the orthogonal strategy.

and the Pareto search method. In this section we focus on the latter aspect of the negotiations. By means of computer experiments we investigate the effectiveness and robustness of the orthogonal and orthogonal-DF approach, to find Pareto-efficient solutions for a wide variety of settings. We evaluate the robustness of the search strategy by experimenting with various concession strategies on the customer side.

Although the system enables customers to initiate several concurrent negotiation threads, within a thread the Pareto search strategy operates independently from the other threads. For researching the efficiency and robustness of Pareto search strategies it therefore suffices to consider only a single negotiation thread in the experiments.

Furthermore, the bundle content in the experimental setup consists of a single category with a low quality of service. The experimental results generalize to negotiations involving multiple categories: Only the shape of the iso-utility curves is affected by the number of categories. In the experiments the shape is varied using different parameter settings.

A general specification of the customer agents and the seller agent was provided in Section 3. Sections 4.1 and 4.2 describe the agent settings which are specifically used within the experimental setup. In particular the agents' preferences and concession strategies are specified in detail in Sections 4.1 and 4.2 respectively. The experimental results are discussed in Section 4.3.

4.1 Agent Preference Settings

We simulate the negotiation with a variety of customer and seller preferences, expressed by the agents' utility functions. The *customer agent's* expected utility depends on ρ the total number of articles the customer expects to read in the low quality of service categories (see also Section 3.1). The value ρ is assumed to be a constant, set randomly between 1 and 20 at the beginning of the experiment. Note that this results in a linear iso-utility curve in the (p_f, p_v) plane (see e.g. Fig. 5). Furthermore, since the purpose is to demonstrate the efficiency of the

Algorithm 1 The orthogonal-DF algorithm

The following is given: (a) the opponent’s last and before last offer $O_1 = (p_f(t), p_v(t))$ and $O_2 = (p_f(t-2), p_v(t-2))$ respectively, (b) bargainer’s utility function $u(p_f, p_v)$ and next aspiration level $U(t+1)$, (c) the step-size δ , and (d) the search direction $s_{dr} \in \{-1, +1\}$. Based on this information the orthogonal-DF computes the next counter offer $O = (p_f(t+1), p_v(t+1))$.

1. Use the orthogonal strategy to compute $O'_1 = (p'_f(t), p'_v(t))$ and $O'_2 = (p'_f(t-2), p'_v(t-2))$, i.e., the points in the (p_f, p_v) plane which generate $U(t+1)$ utils and lie closest to O_1 and O_2 , respectively.
 2. Compute d_1 and d_2 the distance of the opponent’s last two offers, i.e., $d_1 = \|O_1 - O'_1\|$ and $d_2 = \|O_2 - O'_2\|$ ($\|\cdot\|$ denotes Euclidian distance).
 3. Update s_{dr} : whenever $d_1 > d_2$ the orthogonal-DF “turns”, i.e., $s_{dr} = -1 \cdot s_{dr}$, otherwise $s_{dr} = s_{dr}$.
 4. Update δ : decrease δ whenever the orthogonal-DF turns. For a number of periods directly after a turn δ is not increased, and otherwise δ is increased (cf. [7] for the details).
 5. Compute the counter offer $O = (p_f(t+1), p_v(t+1))$: set $p_v(t+1) = p'_v(t) + \delta \cdot s_{dr}$. Next choose $p_f(t+1)$ such that given $p_v(t+1)$ the counter offer generate $U(t+1)$ utils.
-

final deals reached, we set the customer agent’s mandate b_{max} for the bundle such that an agreement is always reached.

The expected utility (i.e., expected revenue) for the *seller agent* is based on ρ , the expected number of articles that the customers will read *on average* in the low quality categories. In contrast to the customer agent, the expectation is not a constant but a function of the variable price p_v . It is assumed that customers who are willing to pay a high variable price are expected to read less than customers with a low variable price (i.e. we assume the law of demand holds cf. [3]). In the experiments we use the linear function $\rho(p_v) = b - a \cdot p_v$ with $b = 20$ and a set randomly between 0.03 and 0.07 at the beginning of an experiment. Note that the seller’s iso-utility curve is now convex (towards the origin).

4.2 Concession Strategies

The customers and the seller can each select their own concession strategies. Although a seller’s concession in the main system can depend on the behavior of all customers (i.e., one-to-many), in the experiments the seller agent’s strategy is simply to decrease the desired utility level or *aspiration level* with a fixed amount each round. The initial aspiration level is randomly varied. Note that the number of customers and their behavior does not affect the seller’s concession when this strategy is used.

On the customer side, on the other hand, we implemented four classes of concession strategies to investigate the robustness of the Pareto search strategy:

1. *Hardhead*. The customer agent does not concede when this strategy is used; the aspiration level remains the same during the negotiations.
2. *Fixed*. A fixed amount c in utils is conceded each round.
3. *Fraction*. The customer concedes the fraction γ of the difference between the current desired utility level and the utility of the opponent’s last offer.
4. *Tit-for-tat*. This strategy mimics the concession behavior of the opponent, based on subjective utility improvement. If the utility level of the seller’s offers increases, the same amount is conceded by the customer. Note that it is the increment in the utility level perceived by the customer. The seller’s actual concession is shielded from the customer agent, as an improvement could also occur when the seller moves along his iso-utility curve. Furthermore, note that the perceived utility improvement could also be negative. To make the concession behavior less chaotic, no negative concessions are made by the customer.

4.3 Results

The seller and the customer in the experiments negotiate in an alternating fashion until an agreement is reached. The efficiency of the agreement is then evaluated based on the distance of the final offer from a Pareto-efficient solution. We measure an offer’s distance from a Pareto-efficient solution as the maximum possible utility improvement for the customer if a Pareto-efficient offer was made, all else remaining equal. To evaluate the quality of the results we compare the out-

Concession strategy	Random search	Orthogonal/DF search	DF/DF search
hardhead	18.92 (± 23.56)	8.03 (± 11.44)	18.63 (± 32.81)
fixed ($c = 20$)	26.52 (± 34.49)	10.43 (± 17.34)	28.82 (± 46.71)
fixed ($c = 40$)	38.91 (± 49.72)	16.21 (± 23.84)	44.29 (± 69.76)
fixed ($c = 80$)	42.12 (± 56.88)	25.61 (± 38.72)	48.84 (± 72.12)
fraction ($\gamma = 0.025$)	30.26 (± 38.37)	10.07 (± 15.03)	32.25 (± 52.81)
fraction ($\gamma = 0.05$)	31.53 (± 40.00)	11.52 (± 16.16)	28.52 (± 52.13)
fraction ($\gamma = 0.1$)	37.81 (± 48.82)	16.91 (± 30.80)	26.28 (± 42.20)
tit-for-tat	72.78 (± 121.35)	59.60 (± 113.27)	56.64 (± 116.82)

Table 1. Average distance from Pareto-efficient solution for various customer concession strategies (rows) and customer/seller search strategies (columns). Results are averaged over 500 experiments with random parameter settings. Standard deviations are indicated between brackets. Best results (see column *Orthogonal/DF*) are obtained if the customer and seller use orthogonal search, and the seller’s search is amplified with a derivative follower.

comes using various search strategies and concession strategies of the customer. Table 1 provides an overview of the results. The row labeled *Random* contains the outcomes when both seller and customer use a random search strategy. This strategy selects a random point on the iso-utility curve⁴. The distance of the

⁴ Only the downward sloping part of the seller’s iso-utility curve is used.

final offer (from the closest Pareto-efficient solution), when random search is used, lies between 1 and 3 percent of the total costs.

Although the inefficiency with random search is only small compared to the total costs, even better results are obtained when one bargainer (typically the customer) uses orthogonal search and the other (the seller) uses orthogonal-DF (i.e., orthogonal search combined with a derivative follower). The results are shown in the column labeled *Orthogonal/DF* of Table 1. The improvements are considerable. The distance of the final offer as a percentage of total costs lies then, for almost all concession strategies, between 0 and 1. Only for the tit-for-tat strategy the distance lies around 1.8 percent. Notice that the Orthogonal/Orthogonal-DF strategy combination is also robust, as best results are obtained using this strategy, relatively independent of the concession strategy selected by the customer.

Table 1 also shows the results if both customer and seller use orthogonal-DF search (column *DF/DF*). These results are very similar to random, however. The derivative follower relies on a consistent response from the opponent to signal the right direction. If both use a derivative follower, this signal is distorted.

Notice that the average distance depends on the concession strategy used by the customer. Although in individual cases Pareto-efficient agreements (with zero distance) are reached using the orthogonal/DF search, the average distance consistently shows some (usually slight) inefficiencies, even when the customer makes no concessions (i.e., the *hardhead* strategy). The reason for this is twofold. Firstly, the DF accelerates finding the efficient solution by making, at times, large steps on the iso-utility curve. At a certain point the algorithm passed the Pareto-efficient point, and then turns. This way the offers keep oscillating around the optimal point. If the concessions are sufficiently large, an agreement can be reached at a point which is less than optimal.

Secondly, the direction and step-size of the DF are based on changes in the Euclidean distance between the seller and customer offers through time. The distance can be influenced by both concessions and movements along the iso-utility curve. As the opponent's iso-utility curve is unknown, the agents are unable to distinguish between the two. This can mislead the DF whenever concessions are very large. Two possible solutions are to make either small concessions, or have intervals with no concessions allowing the search algorithm to find the best deal.

Particularly *tit-for-tat* results in a relatively high inefficiency, because of the reasons described above. Recall that tit-for-tat uses a subjective measure of the opponent's concessions. In practice, the perceived utility increments are sometimes quite large, resulting in bursts of very large concessions. If this occurs near the agreement point this can result in inefficient outcomes.

To conclude, the orthogonal/DF strategy clearly outperforms other combinations of search strategies in the experiments. Inefficiencies still occur, especially if the concessions are large. A trade-off therefore exists between reaching an agreement fast (by making large concessions) and reaching an efficient agreement.

5 Discussion

5.1 The System Revisited

In the paper, we focus on the problem of selling bundles of news items. Clearly, other types of (information) goods can also be sold through the developed system. A key question for extending the use of the system to other application areas is, however, if customers and (to a lesser degree) sellers are willing to have (autonomous) agents do the actual bargaining for them. A prerequisite would be that the traded goods have a relatively low value and transactions are conducted frequently. Consequently, the risks are low and an agent has a lot of opportunities to learn from past experience and gradually improve performance. Note that the negotiation procedure of the system does not require both seller and customer to use the same level of automation. Depending on the particular application of the system, it may be desirable for the customer to rely more or less on the assistant of the software agent.

An additional important aspect of the relevance to other application areas is the potential benefit of using such a system. The developed system appears particularly suitable for selling (bundles) of goods with a high degree of personalization given relatively rapidly changing preferences (as is the case with the news items). More specifically, in the system personalization entails discriminating between customers based on the bundle price and the quality of service. Second-degree price discrimination is the economic term for this type of personalization.

In second-degree price discrimination the price depends on the quantity and/or quality of the purchased good. The distinguishing aspect of second-degree price discrimination is that customers can *self-select* the best purchase. Traditionally, customers are offered a menu of options where a tariff assigns a price to the option in the menu. The work of [8, 9] discusses algorithms which— given a particular tariff structure— learn the best tariffs on-line. They conclude that (especially in a dynamic environment) complex tariffs are generally not the most profitable strategy.

The distinguishing aspect of the developed system is that instead of having explicit tariffs customers can bargain for the most appropriate bundle tariff combination. This can result in a similar (or even higher) degree of discrimination between customers as with explicit complex tariff. In the absence of an explicit tariff structure the seller is, however, more flexible in the degree to which she discriminates. The seller does not have to a priori limit the complexity of the tariff structure. Whenever bundles of (information) goods are being offered, an additional advantage is that, by initiating the negotiation process, customers can explicitly express their interest in a particular bundle of goods. This may facilitate the process of offering customers the appropriate bundles (and consequently it may facilitate the discrimination between customers).

Possibly, bargaining leads to price discrimination based on customers' bargaining skills and not on their preferences. In the developed system this possi-

bility is, however, reduced significantly by the fairness constraint in particular and the fact that bargaining is actually one-to-many in general.

5.2 Bargaining & Pareto Efficiency

In the system the seller agent uses the orthogonal-DF as the Pareto search strategy combined with a concession strategy. The concession strategy determines the next concession relatively independently of the ongoing bargaining process with a particular customer. The idea is that, on the one hand, bargaining with a particular customer should lead to finding the best possible deal for both parties, given the seller’s desired utility level. That is, the bargaining outcome should closely approximate a Pareto-efficient solution. On the other hand, the one-to-many aspect of the bargaining process (i.e., bargaining with more than one customer) should guide the updating of the concession strategy. Thus the seller uses the disentanglement of the bargaining strategy (in a concession and Pareto search strategy) to distinguish explicitly between the one-to-many and one-to-one aspects of the bargaining process.

The experiments in Section 4 show that if a customer agent uses an orthogonal strategy as the Pareto-efficient search strategy then the bargaining outcomes will closely approximate a Pareto-efficient solution. The experiments are conducted for a variety of (customer) concession strategies, customer preferences, and seller preferences. Based on the experimental results we can conclude that any other strategy choice of a customer will probably result in less efficient outcomes. Moreover such a strategy will not influence the concession strategy of the seller (due to the independence of the concession strategy). Consequently, any alternative bargaining strategy of the customer is probably at most as good as the orthogonal strategy combined with a concession strategy that mimics the concessions of the alternative strategy. Thus, given the seller’s choice of the orthogonal-DF combined with a relatively independent concession strategy, it is in a customer’s best interest to choose the orthogonal search strategy combined with a concession strategy. Moreover, this choice results in (a close approximation of) a Pareto-efficient solution.

5.3 Related Work

Related to our work, in [10] a heuristic approach to finding win-win trade-offs between issues is introduced. Contracts which are similar to the opponent’s offer are selected based on fuzzy similarity criteria. Their approach, however, is applied to additive utility functions with independent issues. The orthogonal search method, on the other hand, operates on a more general class of utility functions, which is widely accepted in the economic literature as a reasonable representation for ones preferences.⁵ Additionally, with an orthogonal search method no (domain-specific) similarity function needs to be specified.

⁵ More specifically, with convex preferences the approach works particularly well.

Another approach increasingly used to automate one-to-many negotiations is through auctions. Traditionally, auctions have focused on price as the single dimension of the negotiation. Recently, however, multi-attribute auctions have received increasing attention [11, 12]. Although our system has characteristics similar to those of auctions, bundling and negotiation of information goods have distinct properties which impede the use of current available auction designs. Mainly, a bundle is not sold just once, but can be sold to any number of customers in parallel, with varying contents and configurations. The buyers are also unaware of other buyers in the system and their offers; they perceive the negotiation to be one-on-one. Furthermore, we allow both offers and counter offers, whereas auctions typically have less symmetry.

Rather than via direct negotiation, another way to find (Pareto) efficient solutions for multi-issue problems is through a mediator, see for instance [13–16]. Both parties need to reveal their preferences to the mediator. With a mediator, however, trust becomes an important issue. Furthermore, additional costs are often involved.

6 Concluding Remarks

We introduce a novel system for selling bundles of news items in this paper. Through the system, customers bargain over the price and quality of the delivered goods with the seller. The advantage of the developed system is that it allows for a high degree of flexibility in the price, quality, and content of the offered bundles. The price, quality, and content of the delivered goods may, for example, differ based on daily dynamics and personal interest of customers.

The system as developed here can take into account business related side-constraints, such as “fairness” of the bargaining outcomes. Partly due to these side-constraints (especially fairness) the actual bargaining process between seller and customers is not really one-to-one bargaining between seller and customer but instead is one-to-many (i.e., between seller and customers).

Autonomous software agents perform (part of) the negotiation on behalf of the users of the system. To enable efficient negotiation through these agents we decompose the bargaining strategies into concession strategies and Pareto search strategies. Moreover, we introduce the orthogonal and orthogonal-DF strategy: two Pareto search strategies. We show through computer experiments that the respective use of these two Pareto search strategies by the two bargainers will result in very efficient bargaining outcomes. Furthermore, the system is set up such that it is actually in the best interest of the customer to have their agent adhere to this approach of decomposing the bargaining strategy into a concession strategy and Pareto search strategy.

References

1. Amit, R., Zott, C.: Value creation in ebusiness. *Strategic Management Journal* **22** (2001) 493–520

2. Bakos, Y., Brynjolfsson, E.: Bundling and competition on the internet: Aggregation strategies for information goods. *Marketing Science* **19** (2000) 63–82
3. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Microeconomic Theory*. Oxford University Press (1995)
4. Webster, R.: *Convexity*. Oxford Science Publications (1994)
5. Kephart, J.O., Hanson, J.E., Greenwald, A.R.: Dynamic pricing by software agents. *Computer Networks* **36** (2000) 731–752
6. Dasgupta, P., Das, S.: Dynamic pricing with limited competitor information in a multi-agent economy. In Eztion, O., Scheuermann, P., eds.: *Cooperative Information Systems: 7th International Conference*. Volume 1906 of *Lecture Notes in Computer Science*, Eilat, Israel, BoopIs, Springer (2000) 291–310
7. van Bragt, D.D.B., Somefun, D.J.A., Kutschinski, E., La Poutré, H.: An algorithm for on-line price discrimination. technical report SEN-R0213, CWI (2002)
8. Brooks, C.H., Fay, S., Das, R., MacKie-Mason, J.K., Kephart, J.O., Durfee, E.H.: Automated strategy searches in an electronic goods market: Learning complex price schedules. In: *Proceedings of the ACM Conference on Electronic Commerce*, ACM Press (1999) 31–41
9. Kephart, J.O., Brooks, C.H., Das, R.: Pricing information bundles in a dynamic environment. In: *Proceedings of the 3rd ACM Conference on Electronic Commerce*, ACM Press (2001) 180–190
10. Faratin, P., Sierra, C., Jennings, N.: Using similarity criteria to make negotiation trade-offs. In: *Proc. 4th Int. Conf on Multi-Agent Systems*, Boston, USA, IEEE Computer Society Press (2000) 119–126
11. David, E., Azoulay-Schwartz, R., Kraus, S.: Protocols and strategies for automated multi-attribute auctions. In: *Proc. 1st Int. Joint Conference on Autonomous Agents and MultiAgent Systems*. (2002)
12. Parkes, D., Kalagnanam, J.: Iterative multiattribute Vickrey auctions. Technical report, Harvard (2002) Working paper.
13. Ehtamo, H., Hämäläinen, R.: Interactive multiple-criteria methods for reaching pareto optimal agreements in negotiation. *Group Decision and Negotiation* **10** (2001) 475–491
14. Klein, M., Faratin, P., Sayama, H., Bar-Yam, Y.: Negotiating complex contracts. *Group Decision and Negotiation* **12** (2003) 111–125
15. Raiffa, H.: *The Art and Science of Negotiation*. Harvard University Press, Cambridge, MA (1982)
16. Lin, R.J., cho T. Chou, S.: Bilateral multi-issue negotiations in a dynamic environment. In Faratin, P., Parkes, D., Rodriguez, J., Walsh, W., eds.: *Workshop on Agent-Mediated Electronic Commerce V (AMEC V)*. (2003) 117–123