

Autonomous Bidding & Coordinated Acceptance in One-to-Many Negotiations

Thimjo Koça

*Intelligent & Autonomous Systems Group
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
thimjo.koca@cw.nl*

Tim Baarslag

*Intelligent & Autonomous Systems Group
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
T.Baarslag@cw.nl*

Catholijn M. Jonker

*Interactive Intelligence Group
TU Delft
Delft, The Netherlands
c.m.jonker@tudelft.nl*

Abstract—This work presents the Autonomous Bidding & Coordinated Acceptance framework (ABCA): an agent-team design that allows general bilateral agents to engage in one-to-many negotiations in a setting where (possibly overlapping) deals with multiple opponents are desirable. We propose also a coordinated acceptance strategy that uses the estimated outcomes of its bilateral negotiations while deciding to accept a deal.

Index Terms—one-to-many negotiations, multi-deal, composite negotiations, concurrent negotiations, procurement

I. INTRODUCTION

Procurement is the process businesses use to agree to terms for acquiring goods or services, often via competitive bidding [1]. Setting up a purchase agreement is a time-consuming process, with an average cycle time of 12 weeks, most of which is spent on negotiations [2]. As more and more businesses are moving their purchases online and the volume of online transactions is increasing, it is paramount to implement new technologies, such as automated negotiations, to reduce procurement duration and complexity.

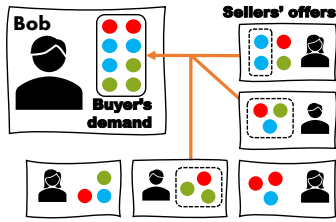


Fig. 1. An example of the multi-deal setting.

In a common procurement setting, known as *one-to-many multi-deal setting*, a buyer has a demand list that contains various items¹ which can usually only be satisfied by reaching a deal with multiple sellers at once. For instance, suppose a buyer named Bob needs to buy two red items, three blue items, and three green ones in a scenario as depicted in Fig. 1. None of the sellers can satisfy Bob's demand fully and therefore he has to strike multiple agreements with each seller individually. Typically, Bob may prefer to conduct these negotiations concurrently with all of them so that he is able to

compare prices and regulate his demand for each seller based on the supply from the other sellers. Hence, Bob faces two types of challenges: (1) negotiating bilaterally with each seller to achieve an agreement that (partially) satisfies his demand, and (2) coordinating among bilateral negotiations to reach a satisfactory overall outcome.

The negotiation and coordination tasks cannot be separated fully, since to support coordination, the bilateral negotiations need to respond to input and provide information, and to reach bilateral outcomes that fit the overall goal, bilateral negotiation should receive updates regarding the status of the whole one-to-many negotiation. Yet, to negotiate bilaterally in a proficient way, a one-to-many negotiator can benefit from the vast literature on bilateral negotiation strategies, ranging from simple concession-based strategies [3], to more refined heuristics [4], [5], with new bilateral strategies being proposed on a regular basis (for instance, new strategies are presented annually in Automated Negotiating Agents Competition [6]).

To effectively use general bilateral strategies in a one-to-many setting, we need to integrate coordination duties either while taking the decision of acceptance, or during the bilateral bidding process. Furthermore, to do this without the need to make any context-dependent changes, the bilateral agent's decision-making process must not be altered.

Focusing on the acceptance strategy of bilateral agents, it is not straightforward how to integrate coordination duties, without altering the bilateral acceptance process. A fundamental difficulty comes from the fact that bilateral agents in literature operate under single-acceptance protocols, typically the Alternating Offers Protocol (AOP) [7], according to which parties exchange offers until one accepts the offer of the opponent, walks away, or the negotiation deadline is reached. As a consequence, whenever a seller accepts an offer proposed by the general bilateral agent representing the buyer, an agreement will be reached without considering any coordination. Furthermore, a coordinated acceptance done well, must take into account the likely outcome of each other ongoing bilateral negotiation the buyer is involved in, yet general bilateral agents are not able to provide any information about their ongoing negotiations, or use any information about the possible outcomes of the other negotiations the buyer engages in.

¹such as products, services, etc.

To the best of our knowledge, there is no previous work that proposes a way to enable a systematic reuse of general bilateral agents in the context of one-to-many multi-deal negotiations (by "systematic reuse" we mean the use of general bilateral agents in a one-to-many multi-deal negotiations without the need to make any context-dependent changes). Rahwan et al. [2] proposed an approach to tackle a simplified version of our setting, in particular that of a buyer trying to reach a unique deal with only one of the available sellers (i.e. a *single-deal* setting). However, their focus was on the agent's overall architecture and the identification of few coordination heuristics. The same holds true for follow up works that used Rahwan's agent design in a single-deal setting [8], [9], or a multi-deal one [10], and for works that proposed agents the coordination and negotiation strategies of which are not decoupled [11], [12]. Moreover, there is no previous work that proposes some coordination strategy that uses outcome estimations in its decision-making.

In this work, we propose the Autonomous Bidding & Coordinated Acceptance framework (ABCA): an agent team design which allows the integration of general bilateral agents in a one-to-many multi-deal setting. ABCA is composed of a coordinator and several Oblivious Bilateral Agents (OBAs), each using a general bilateral strategy to negotiate with its seller. While an OBA can bid autonomously, the coordinator, using estimates about the outcomes the rest of OBAs can reach, must grant its permission to accept a deal. OBAs in ABCA communicate solely with the coordinator and using the AOP, unaware of the setting they operate in or the existence of a coordinating agent, while the coordinator is in charge of the communication with the sellers, so that the acceptance of each bilateral deal can be coordinated.

II. RELATED WORK

In the context of one-to-many negotiations, the majority of works focus on the proposal of coordination heuristics, or the identification of adequate one-to-many bidding, or acceptance strategies. To the best of our knowledge, no previous work has provided a systematic way of reusing general bilateral agents in the one-to-many context.

Rahwan et al. [2] were the first to introduce a modular architecture composed of a coordinator and several sub-negotiators (each responsible to bid with a unique seller) and propose some basic coordination strategies for the single-deal setting. In some follow up works Nguyen & Jennings [8] enhanced their coordinator with the ability to decommit from a settled agreement paying a penalty fee, An et al. [13] proposed a coordination strategy that controls when its sub-negotiators would bid, and Williams et al. [9] proposed a coordinator which can estimate the utility that each bilateral thread could bring and use that to tune its sub-negotiators' parameters. Mansour & Kowalczyk [14], in a context where multiple, non-overlapping deals over single-issue items can be reached, proposed three coordination strategies (or meta-strategies as they call them), through which the reservation values and/or the concession rate of its sub-negotiators could

be controlled. Najjar et al. [10] operated in a one-to-many multi-deal setting to investigate the problem of quality-of-experience management. They propose a coordinator that can update its sub-negotiators' parameters, as well as take control of the bidding if a deal has not been reached near the deadline. Note that, none of these works had at its focus integrating general bilateral strategies. Hence, their authors either adapted some simple time-based concession strategies for their sub-negotiators [8], [10], [14], or picked bilateral strategies that could be adapted to fit the type of coordination they proposed [2], [9], [13]. Moreover, Williams et al. [9] are the only to propose a coordinator that uses estimates, yet, their estimates are utility estimates and not outcome estimates as in our case.

Few authors have proposed monolithic agents to tackle one-to-many negotiations and are therefore unable to reuse general bilateral agents [11], [12], [15]. On some other works, Beam et al. [16] provided the earliest discussion of the challenges that one-to-many multi-deal negotiation might raise in the context of procurement; Mohammad [17] provided a methodology to approximate optimal policies when the acceptance strategies of the opponents are known in the context of supply chain management; Niu et al. [18] proposed a method to group bilateral threads in batches that will be executed sequentially, in an environment where multiple deals can be reached, but not all bilateral negotiations can be executed concurrently; Amini et al. [19] presented a BOA agent that operates in a multi-party negotiation; Sanchez-Aguix et al. [20] treated the problem of how a team of independent agents can negotiate with a single opponent; and Yokoo & Hirayama made a survey of algorithms that solve the distributed Constraint Satisfaction Problem [21].

III. PROBLEM SETTING

In our problem setting, a buyer attempts to satisfy part of their demand by striking deals with multiple sellers concurrently over (possibly overlapping) bundles of items.

A. Domain Model

A buyer in our setting seeks to acquire a number of indivisible items among the set of all possible items \mathcal{I} , which can be characterized by various attributes (quantity, price, expiration date, etc). Different sellers supply items in *bundles*. Formally, a bundle b is a tuple $(I, p_1, p_2, \dots, p_n)$, where $I \subseteq \mathcal{I}$, and $p_k : I \rightarrow \mathbb{P}_k$ assigns a value for the attribute p_k to each element in I .

Example 3.1 (Bundle): In a scenario where a bundle $b = (I, p_1, p_2)$ is composed of items $I = \{\text{Gin}, \text{Tonic}\}$ characterized by quantity $p_1 : I \rightarrow \mathbb{N}$ and unit price $p_2 : I \rightarrow \mathbb{R}^{\geq 0}$, a seller offers a bundle composed of 3 Gin items for €4 each and 5 Tonic items for €6.

The set Ω of all bundles considered by the agents is called the *outcome space*. The empty bundle is denoted by \emptyset .

The buyer possesses a utility function $u : \mathcal{M}_\Omega \rightarrow [0, 1]$ defined over \mathcal{M}_Ω , the set of possible multi-sets² over Ω .

²We define u over multi-sets rather than sets to be able to evaluate several identical bundles put together.

Example 3.2 (Utility over multi-sets): Given a multi-set of bundles $B \in \mathcal{M}_\Omega$, characterized by the same attributes as in Example 3.1, a typical utility function can be constructed as follows: given $U_i^{p_1} : \mathbb{N} \rightarrow [0, 1]$ and $U_i^{p_2} : \mathbb{R}^{\geq 0} \rightarrow [0, 1]$, two triangular utility functions [5] over quantity and unit price of an item $i \in \mathcal{I}$ and B_i , a multi-set that contains all bundles of B in which item i is present, we can construct utility functions over the quantity and unit price of an item i :

$$\begin{cases} u_i^{p_1}(B_i) = \begin{cases} U_i^{p_1}[\sum_{b \in B_i} p_1^b(i)], & B_i \neq \emptyset \\ 0, & \text{otherwise} \end{cases} \\ u_i^{p_2}(B_i) = \begin{cases} U_i^{p_2}[\sum_{b \in B_i} \frac{p_1^b(i) \cdot p_2^b(i)}{p_1^b(i)}], & B_i \neq \emptyset \\ 0, & \text{otherwise} \end{cases} \end{cases}$$

where $p_1^b : \Omega \rightarrow \mathbb{N}$ and $p_2^b : \Omega \rightarrow \mathbb{R}^{\geq 0}$ are the quantity and unit price of items of bundle b . The additive utility is then:

$$u(B) = \sum_{i \in I_B} [w_i^{p_1} \cdot u_i^{p_1}(B_i) + w_i^{p_2} \cdot u_i^{p_2}(B_i)]$$

where $w_i^{p_1}, w_i^{p_2} \in [0, 1] \wedge \sum_{i \in I_B} [w_i^{p_1} + w_i^{p_2}] = 1$, and $I_B = \{i \in \mathcal{I} | \exists b = (I_b, p_1^b, p_2^b) \in B : i \in I_b\}$.

B. One-to-Many Multi-Deal Negotiation

A one-to-many negotiation N between a buyer and competing sellers $\sigma_1, \dots, \sigma_n$ consists of n concurrent bilateral negotiations between the buyer and sellers σ_k (see Fig. 2), all of which take place within a predefined time interval \mathfrak{T} . We will refer to each of these bilateral interactions as *negotiation threads*.

A negotiation thread is *concluded* if: (1) the parties reach an agreement within \mathfrak{T} ; or (2) one of the parties decides to terminate the negotiation without an agreement before the end of \mathfrak{T} ; or (3) no decision has been made by the end of \mathfrak{T} . A non-concluded thread is *ongoing*. Moreover, at each time $t \in \mathfrak{T}$ we will denote ongoing threads by $\mathcal{G}(t)$. For a concluded thread T , we denote the outcome $O(T) = b$ in case the parties of T reach an agreement over a bundle b , and by $O(T) = \emptyset$ in case the parties fail to reach an agreement.

We are interested in *multi-deal* one-to-many negotiations, i.e. one-to-many negotiations in which the buyer has a demand list that contains various items which can usually only be satisfied by reaching a deal with multiple sellers at once. Consequently, the buyer's possible bilateral outcomes may overlap and therefore the buyer needs to coordinate its efforts among the different threads.

The outcome $O(N)$ of the one-to-many multi-deal negotiation N is defined as the multi-set composed of the bilateral outcomes $\{O(T_1), \dots, O(T_n)\}$. at each time $t \in \mathfrak{T}$ the buyer has reached a partial (possibly empty) outcome \mathcal{PO}_t through its concluded threads $\{T_k\}_{k=1}^n \setminus \mathcal{G}(t)$.

The buyer's goal is to maximize utility $u(O(N))$.

C. Oblivious Bilateral Agents

Our aim in this work is to provide a modular design that enables the buyer to employ a general bilateral agent in each of the threads it engages. Since the general bilateral agents

are not aware of the one-to-many context, we will refer to them as *Oblivious Bilateral Agents (OBAs)*.

An OBA in a bilateral negotiation is assigned some initial preferences (most frequently in the form of a utility function) which cannot be altered, and exchanges offers with its opponent until the negotiation concludes, unable to receive or send any information regarding the progress of their interaction.

A typical bilateral negotiation obeys the rules of Alternating Offers Protocol (AOP) [7]: a protocol that allows parties to exchange offers in a turn-based manner within a predefined time interval. Each offer can be a bundle that one party proposes to its opponent, an *accept* message signaling agreement over the opponent's latest proposed bundle, or a *walk-away* message signaling the sender is terminating the negotiation without reaching an agreement.

D. Problem Statement

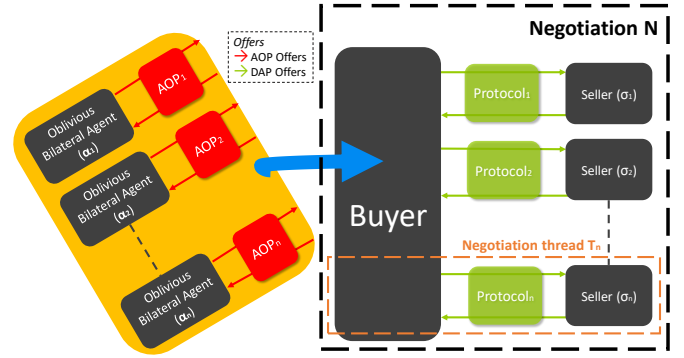


Fig. 2. The problem of fitting general bilateral AOP agents in a one-to-many multi-deal setting.

A design that enables the use of general bilateral agents in a multi-deal one-to-many negotiation must fulfill some requirements. More specifically, the design must allow for a coordination among threads even when a general bilateral agent is unable to: (1) share any information about the status of its ongoing negotiation; (2) generate information regarding the progress of the other threads of the buyer; (3) use information from the other threads in its decision-making.

IV. ABCA AGENT TEAM-DESIGN

We propose the Autonomous Bidding & Coordinated Acceptance framework (ABCA): a design of agent teams that enables the integration of general bilateral agents in a one-to-many multi-deal setting. In ABCA, the bidding is done autonomously by the OBAs, while acceptance is determined by a coordinator agent (see Fig. 3) with three responsibilities (each addressing one of the problem requirements discussed in section III-D). More specifically, the coordinator: (1) acts as a mediator that forwards and keeps track of the bids received by an OBA/seller to the corresponding seller/OBA; (2) estimates the outcomes of the buyer's ongoing threads; and (3) intervenes whenever an acceptance decision must be taken.

Moreover, to engage in the decision to accept even when the acceptance comes from a seller, the coordinator communicates with each seller through the double-acceptance AOP (DAOP) [22], an extension of AOP that requires a confirmation of acceptance to reach an agreement.

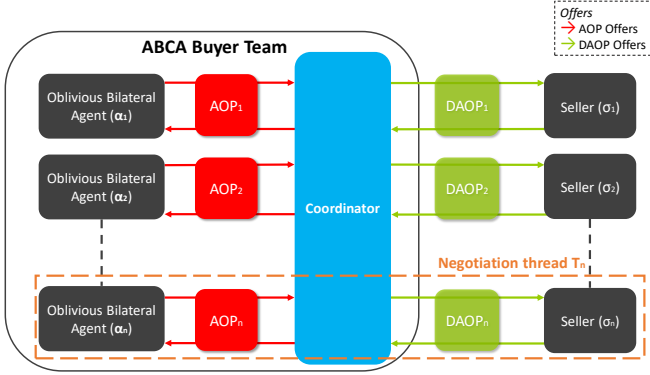


Fig. 3. The ABCA design and communication with the sellers.

Our design differs from the architecture of Rahwan et al. [2] in that it enables multi-deal negotiations, and it allows any general bilateral strategy to be used as it is, without requiring any adjustments to enable communication with the coordinator.

A. Autonomous Bidding

In the ABCA framework, the coordinator sets up a thread with each seller and assigns to an OBA α_k with a utility function u_k ; and the responsibility to bilaterally interact in thread T_k . Whenever an offer o is sent by seller σ_k to the buyer, the offer is received by the coordinator. In case o is a bid, the coordinator forwards it to α_k . At receipt of o , α_k uses its general bilateral strategy to generate a response offer o' and sends it to the coordinator. If o' is a bid or a walk-away, the coordinator forwards it to the seller σ_k . The cycle repeats until α_k or σ_k send an accept to the coordinator, or the negotiation thread fails (i.e. a walk-away message is sent or the negotiation time is over).

In this way, ABCA can incorporate any general bilateral agent and allows the agent to bid autonomously, while being unaware of the fact that it operates in a one-to-many context and that its negotiation with the corresponding seller is mediated by the coordinator.

B. Coordinated Acceptance

While the OBAs are responsible for the bidding in their threads, the accept and confirmation messages are tackled by the coordinator, i.e. whenever an OBA α_k (or seller σ_k) informs the coordinator that it wants to accept a bundle b , the coordinator determines whether to forward the acceptance (or send a *confirm* message) to the seller σ_k .

We propose the Coordinated Acceptance Algorithm (Algorithm 1) by which the coordinator accepts/confirm b whenever it is part of the expected negotiation outcome that

maximizes its utility or walks away from the corresponding thread otherwise. In detail, when receiving an accept message in T_k , the coordinator uses its outcome estimator function $\mathcal{EO} : \{T_k\}_{k=1}^n \rightarrow \Omega$ to generate $\mathcal{EO}(T)$ for all its other ongoing threads $T \in \mathcal{G}(\text{now}) \setminus \{T_k\}$. Next, assuming that the estimated outcomes will be eventually reached by the OBAs, the coordinator combines them with its current partial outcome $\mathcal{PO}_{\text{now}}$ to calculate the maximum utility u_b that can be achieved if b is accepted, and compares it to the maximum utility $u_{\bar{b}}$ that can be achieved in case b is not accepted. Whenever $u_b > u_{\bar{b}}$, the coordinator decides to accept b and informs the seller about its decision. Otherwise, the coordinator sends a walk-away message to σ_k and the negotiation in T_k concludes without an agreement.

In this way, ABCA controls the bilateral deals that it strikes and considers information from the whole one-to-many negotiation when deciding to reach an agreement.

Note that we choose to send a walk-away message since after α_k sends an accept message, it assumes T_k has concluded with success and terminates its execution, hence, the buyer has no negotiation strategy in T_k from that moment. Moreover, if the coordinator refuses to confirm the acceptance or walk away when a seller σ_k sends an accept message, the coordinator must generate an offer to forward to α_k , distorting the negotiation history in T_k and OBA's perception of σ_k 's strategy.

Algorithm 1 Coordinated Acceptance Algorithm

Input: received *Accept* by α_k or σ_k for bundle b

- 1: $u_b = \max_{\mathcal{X} \subset \mathcal{G}(\text{now}) \setminus \{T_k\}} u(\{b, \mathcal{PO}_{\text{now}}\} \cup \{\mathcal{EO}(T)\}_{T \in \mathcal{X}})$
 - 2: $u_{\bar{b}} = \max_{\mathcal{X} \subset \mathcal{G}(\text{now}) \setminus \{T_k\}} u(\{\mathcal{AO}_{\text{now}}\} \cup \{\mathcal{EO}(T)\}_{T \in \mathcal{X}})$
 - 3: **if** $u_b > u_{\bar{b}}$ **then**
 - 4: **if** isSentBySeller(*Accept*) **then**
 - 5: **return** *Confirm*
 - 6: **else**
 - 7: **return** *Accept*
 - 8: **end if**
 - 9: **else**
 - 10: **return** *Walk-Away*
 - 11: **end if**
-

V. EXPERIMENTS

The ABCA framework enables general bilateral agents to represent a party in a one-to-many multi-deal setting, in a way that allows reusability and is more modular compared to the use of a single monolithic agent. However, the modularity and reusability of ABCA may come at a cost since the bilateral agents are not using any information from the other threads of the one-to-many negotiation. Therefore, we choose to study the theoretical limit of the performance of an ABCA team by comparing it to the performance of a monolithic Oracle in ideal circumstances in Experiment 1, and under more realistic circumstances, where imperfect OBAs are used, in Experiment 2.

A. Setup

For each experiment, we perform 1000 simulations for scenarios that involve a buyer seeking 10 different types of items (characterized by quantity and unit price) offered by 10 sellers and measure the average buyer's utility. We assign to the buyer a goal bundle $g = (G, p_1, p_2)$ composed of all items, with a goal quantity $p_1 = 100$ and a unit goal price $p_2 = 1$ each. The preferences of the coordinator are expressed through a utility function as introduced in example 3.2, where the triangular utilities around quantities and unit prices peak at $p_1^g(i)$ and $p_2^g(i)$ respectively and have their end points specified by multiplying $p_1^g(i)$ or $p_2^g(i)$ with the amount over which the buyer is willing to deviate from its goal quantity or goal unit price, called the tolerance parameter (lower quantity tolerance $t_l^{p_1} \in [0, 0.5]$, upper quantity tolerance $t_u^{p_1} \in [1, 1.5]$, lower price tolerance $t_l^{p_2} = 1$, and upper price tolerance $t_u^{p_2} = 100$). For the OBAs we constrain the coordinator's goal bundle to the items that their corresponding seller can provide and construct their utility functions in the same way.

In each scenario, a seller offers some random items out of the buyer's goal bundle. Moreover, each seller possesses 1000 acceptable bundles composed of the items the seller can offer, with quantities that do not exceed the buyer's quantities, and prices $p_1^s \in \{2, \dots, 100\}$.

B. Experiment 1 - ABCA validation

In the first experiment we investigate the performance cost of ABCA framework. For this purpose, we abstract away bilateral negotiations and compare the performance of an idealized ABCA team with the performance of a monolithic Oracle, which solves approximately the scenario as a centralized constrained optimization problem that maximizes the buyer's utility given the possible multi-set of bundles acceptable to each seller. The Idealized ABCA Team is composed of a coordinator that possesses a perfect outcome estimator, and several OBAs, each of them able to secure the deal that maximizes its utility among the acceptable bundles of their corresponding seller. Since all ABCA components are optimized, every difference in performance can be attributed to ABCA's structure, which allows to assess the performance cost of the ABCA framework.

According to our results (see Fig. 4), the Idealized ABCA Team is able to secure an average utility of 0.68, or 87% of the utility that the monolithic Oracle could achieve, which indicates that a buyer represented by an ABCA Team is capable of achieving high utilities in a one-to-many negotiation even though its bilateral bidding strategies are not coordinated, and its bilateral goals cannot be updated.

In Experiment 1, each OBA always secures the best possible deal. However, bilateral strategies in practice are not able to systematically maximize the OBA's utility. We relax this assumption in the second experiment.

C. Experiment 2 - Using imperfect bilateral strategies

In the second experiment we evaluate the performance of an ABCA team when its OBAs are imperfect. We evaluate

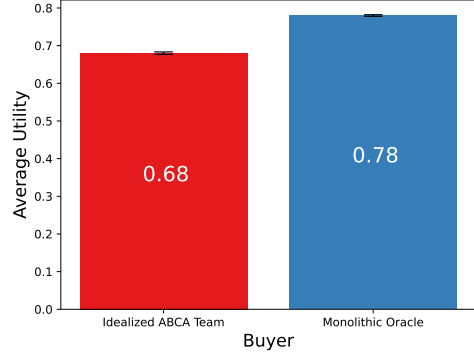


Fig. 4. Average utility obtained by the Perfect ABCA Team and the Monolithic Oracle. The error bars display the confidence intervals.

the *Estimating Team*, a team composed of a coordinator whose outcome estimator is perfect, and 10 imperfect abstract OBAs. The OBAs are modeled such that the best among them always secures a deal that lies in the highest 10% (with respect to the OBA's utility) of the corresponding seller's acceptable bundles, the second best secures deals that lie in the second highest 10%, and so on. Furthermore, to investigate if the qualitative impact of bilateral strategies on the overall buyer's performance is affected by the presence of outcome estimation, we include the *Desperate Team*, composed of Rahwan's Desperate coordinator [2] adapted for the multi-deal setting and the same abstract OBAs as the Estimating Team. Lastly, we use as a baseline the *Baseline Team*, composed of the Desperate coordinator and 10 Median OBAs, i.e. OBAs that secure as an agreement the median bundle (with respect to the OBA's utility) among the set of acceptable bundles of the corresponding seller. Since the Estimating Team differs from the Idealized ABCA of Experiment 1 only by the performance of its OBAs, this setup allows to study ABCA's performance drop caused by imperfect OBAs. Moreover, the inclusion of the Desperate Team enriches the study and allows us to explore ABCA's performance drop caused by imperfect OBAs in the presence or not of outcome estimations.

As we increase the performance of the OBAs (see Fig. 5), the coordinators of both the Desperate Team and the Estimating Team secure overall outcomes with progressively higher utilities ($p < 0.01$) since their coordinators have to pick among better bilateral deals, indicating that the capacity of the OBAs to reach profitable agreements has a decisive impact on the performance of an ABCA Team, which is not qualitatively affected by the absence of an outcome estimator. Furthermore, the Estimating Team always has a significant advantage ($p < 0.01$) compared to the Desperate Team, since the Estimating Team's coordinator benefits from taking a more informed decision, indicating that information about future outcomes improves the performance of an ABCA team. Lastly, the fact that the Estimating Team performs better than the Baseline Team for the OBAs of intervals 20 – 30% and

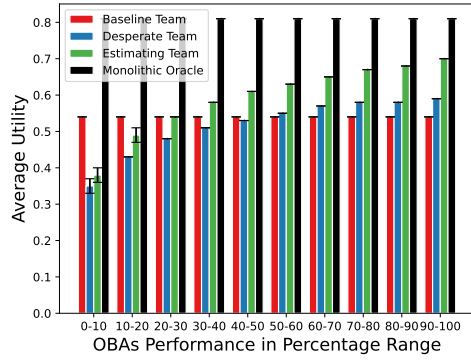


Fig. 5. Average buyer utility as we vary the performance of the OBAs of the Estimating and Desperate teams.

30–40%, while the Desperate Team performs worst, indicates that a coordinator that can estimate bilateral outcomes well can recover some of the performance losses caused by bad bilateral strategies.

VI. CONCLUSIONS & FUTURE WORK

This work introduces the Autonomous Bidding & Coordinated Acceptance framework (ABCA), an agent-team design that enables the integration of general bilateral agents in a one-to-many multi-deal setting and Coordinated Acceptance Algorithm that uses bilateral outcome estimates while deciding to accept a possible deal. We find that that ABCA in principle can achieve results matching those of a monolithic design, while having a more modular design and enabling the use of general bilateral strategies proposed in literature. The ABCA framework may thus help to bridge the gap between bilateral negotiation research and the needs of procurement.

Future work may extend the Coordinated Acceptance Algorithm to account for distributions over possible outcomes which can lead to more robust strategies, or build on ABCA to propose architectures that permit coordination of general bilateral bidding strategies which can result in more effective one-to-many negotiations. Moreover, the existing experimental setup can be extended to account for imperfect outcome estimation (and as a consequence possible unpredictability of bilateral strategies) so that its effect on ABCA's overall performance can be investigated.

ACKNOWLEDGMENT

We are indebted to Flavio Gaier and Niklas Hall from Acumex ApS for our discussions on real-life procurement settings. The research reported in this article is part of Vidi research project VI.Vidi.203.044, which is financed by the Dutch Research Council (NWO).

REFERENCES

[1] J.-J. Laffont and J. Tirole, *A theory of incentives in procurement and regulation*. MIT press, 1993.

[2] I. Rahwan, R. Kowalczyk, and H. H. Pham, "Intelligent agents for automated one-to-many e-commerce negotiation," in *Proceedings of the twenty-fifth Australasian conference on Computer science-Volume 4*. Australian Computer Society, 2002, pp. 197–204.

[3] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation decision functions for autonomous agents," *Robotics and Autonomous Systems*, vol. 24, no. 3–4, pp. 159–182, 1998.

[4] S. Chen and G. Weiss, "OMAC: a discrete wavelet transformation based negotiation agent," in *Novel Insights in Agent-based Complex Automated Negotiation*. Springer, 2014, pp. 187–196.

[5] C. R. Williams, V. Robu, E. H. Gerding, and N. R. Jennings, "IAMmhaggl: A negotiation agent for complex environments," in *New Trends in Agent-based Complex Automated Negotiations*. Springer, 2012, pp. 151–158.

[6] T. Baarslag, R. Aydoğan, K. V. Hindriks, K. Fujita, T. Ito, and C. M. Jonker, "The automated negotiating agents competition, 2010–2015," *AI Magazine*, vol. 36, no. 4, pp. 115–118, 2015.

[7] M. J. Osborne and A. Rubinstein, *A course in game theory*. MIT press, 1994.

[8] T. D. Nguyen and N. R. Jennings, "Managing commitments in multiple concurrent negotiations," *Electronic Commerce Research and Applications*, vol. 4, no. 4, pp. 362–376, 2005.

[9] C. R. Williams, V. Robu, E. H. Gerding, and N. R. Jennings, "Negotiating concurrently with unknown opponents in complex, real-time domains," in *20th European Conference on Artificial Intelligence*. IOS Press BV, 2012, pp. 834–839.

[10] A. Najjar, Y. Mualla, K. D. Singh, G. Picard, D. Calvaresi, A. Malhi, S. Galland, and K. Främling, "One-to-many negotiation qoe management mechanism for end-user satisfaction," *IEEE Access*, vol. 9, pp. 59 231–59 243, 2021.

[11] A. Byde, M. Yearworth, K.-Y. Chen, and C. Bartolini, "AutONA: A system for automated multiple 1-1 negotiation," in *EEE International Conference on E-Commerce*, 2003. IEEE, 2003, pp. 59–67.

[12] B. An, V. R. Lesser, D. E. Irwin, and M. Zink, "Automated negotiation with decommitment for dynamic resource allocation in cloud computing," in *Proceedings of AAMAS 2010*. Toronto, Canada: IFAAMS, 2010, pp. 981–988.

[13] B. An, K. M. Sim, L. G. Tang, S. Q. Li, and D. J. Cheng, "Continuous-time negotiation mechanism for software agents," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 36, no. 6, pp. 1261–1272, 2006.

[14] K. Mansour and R. Kowalczyk, "An approach to one-to-many concurrent negotiation," *Group Decision and Negotiation*, vol. 24, pp. 45–66, 2015.

[15] P. Bagga, N. Paoletti, B. Alrayes, and K. Stathis, "ANEGMA: an automated negotiation model for e-markets," *Autonomous Agents and Multi-Agent Systems*, vol. 35, no. 2, pp. 1–28, 2021.

[16] C. Beam, A. Segev, M. Bichler, and R. Krishnan, "On negotiations and deal making in electronic markets," *Information Systems Frontiers*, vol. 1, no. 3, pp. 241–258, 1999.

[17] Y. Mohammad, "Concurrent local negotiations with a global utility function: a greedy approach," *Autonomous Agents and Multi-Agent Systems*, vol. 35, no. 2, p. 28, 2021.

[18] L. Niu, F. Ren, and M. Zhang, "Feasible negotiation procedures for multiple interdependent negotiations," in *Proceedings of AAMAS 2018*. Stockholm, Sweden: IFAAMS, 2018, pp. 641–649.

[19] M. Amini, M. Fathian, and M. Ghazanfari, "A BOA-based adaptive strategy with multi-party perspective for automated multilateral negotiations," *Applied Intelligence*, vol. 50, no. 9, p. 2718–2748, 2020.

[20] V. Sanchez-Anguix, V. Julian, V. Botti, and A. Garcia-Fornes, "Tasks for agent-based negotiation teams: Analysis, review, and challenges," *Engineering Applications of Artificial Intelligence*, vol. 26, no. 10, pp. 2480–2494, 2013.

[21] M. Yokoo and K. Hirayama, "Algorithms for distributed constraint satisfaction: A review," *Autonomous Agents and Multi-Agent Systems*, vol. 3, no. 2, pp. 185–207, 2000.

[22] B. Alrayes, Ö. Kafalı, and K. Stathis, "Concurrent bilateral negotiation for open e-markets: the conan strategy," *Knowledge and Information Systems*, vol. 56, no. 2, pp. 463–501, 2018.