

Concurrent Negotiations with Global Utility Functions

JAAMAS Track

Yasser Mohammad^{1,2}¹NEC Corporation

²National Institute of Advanced Industrial Science and
Technology
Japan
y.mohammad@nec.com

Shinji Nakadai^{1,2}¹NEC Corporation

²National Institute of Advanced Industrial Science and
Technology
Japan
nakadai@nec.com

ABSTRACT

Automated Negotiation is attracting more attention from researchers recently as it is becoming more relevant to industrial and business applications with increased reliance on automated systems. Most research in this area assumes either a single negotiation thread with a well-defined utility function for each agent involved or a set of concurrent negotiations with an ordering of outcomes in each local negotiation. In this paper, we consider an agent engaged in a set of concurrent negotiations with a utility function defined only for the *complete set of agreements* in all of them and no locally defined ordering of outcomes in any negotiation independent from what happens in the others. We argue that this problem setting is interesting both from the academic and the industrial points of view. The paper then presents an algorithm that allows such agent to maximize its expected global utility by orchestrating its behavior in all negotiation threads. The performance of the proposed method is analyzed theoretically and empirically using simulation.

KEYWORDS

Automated Negotiation; Concurrent Negotiation; Agreement Technologies

ACM Reference Format:

Yasser Mohammad^{1,2} and Shinji Nakadai^{1,2}. 2022. Concurrent Negotiations with Global Utility Functions: JAAMAS Track. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 3 pages.

1 INTRODUCTION

Most research in *automated negotiation* focuses on single-threaded negotiation in which two or more agents are engaged in a common negotiation process [2, 4]. In many practical scenarios, the agent needs to negotiate *simultaneously* with multiple other agents with its utility function defined only for complete sets of agreements in all of these negotiation threads (*concurrent negotiation* hereafter).

Concurrent negotiation is attracting more interest recently with variants of the problem being studied by different groups. Nguyen and Jennings [8] proposed a heuristic model which assumes that partners are either a conceiver or a non-conceiver agent utilizing a time-based strategy. Alrayes et al. [1] recently proposed a method that uses a new protocol allowing for reservation of goods in a buyer-seller negotiation with a penalty. The ability to *reserve* goods

allows agents to make tentative commitments that they can decommit from later with a penalty. Other methods allowing decommitment have been investigated (e.g. [9]). Li et al. [3] proposed a method based on approximating negotiation with outside options by an auction and utilized insights from auction theory to propose three heuristics for *concurrent negotiation* with a single price-like issue. This method relies on the ability to model the effect of any negotiation thread on the others as a change in *reservation value* (the value expected on failure) which makes it difficult to extend to more general settings.

A more important problem in most cases is the — sometimes implicit — assumption that an ordering can be defined for outcomes in one negotiation thread without referencing other threads. This is appropriate only in some limited scenarios including resource allocation with additive utility [10]. This paper has four main contributions. Firstly, a rigorous formulation of the concurrent local negotiations with global utility function problem is articulated. Secondly, an offline algorithm for solving this problem efficiently with guaranteed convergence near a local maximum of the global utility function with arbitrary confidence is given. Thirdly, an online *heuristic* adaptation of the offline algorithm for realistic time-dependent acceptance models is provided. Finally, an empirical evaluation of the proposed method against state-of-the-art methods is reported showing the viability of the proposed approach.




Figure 1: Concurrent Negotiations.

Fig. 1 shows an instantiation of the problem discussed in this paper with a single center C engaged in four closed bilateral negotiations (negotiation threads) with four *edge* agents. Each edge agent has its own utility function which is used to evaluate outcomes it receives (and plan to offer) to the center. The center agent has a single *global* utility function that is defined only for a complete set of agreements/ disagreements in the four negotiation threads.

The set of E simultaneous negotiation threads (a *concurrent negotiation* hereafter) Υ is defined – from the point of view of the agent C engaged in them – as a tuple $(E, u, \Upsilon^e \forall e \in [1, E])$ with E negotiations and a single utility function u defined for the set of negotiations Υ^e each of which is defined by an outcome-space Ω^e , a round limit T^e and an acceptance model a^e representing all the information the agent have about the corresponding edge agent.

The problem that this paper tries to solve is the following: Given a *concurrent negotiation* Υ , find the *multithreaded-policy* π that maximizes the expected utility for agent C . Formally, find π^* such that:

$$\pi^* = \arg \max_{\pi} \mathcal{EU}(\pi) \equiv \arg \max_{\pi} \int_{\Omega} p(\omega|\pi) u(\omega) d\omega, \quad (1)$$

where ω is a complete set of outcomes (one for each thread) of dimensionality E , and $p(\omega|\pi)$ is the probability of reaching agreement ω given the policy π ¹. The outcome ω can always be decomposed into ω^e representing the outcome of thread e , and ω^{-e} representing the outcomes on all the others.

2 PROPOSED APPROACH

The main idea of the proposed method is to utilize the following theorem to disentangle the computations required for each negotiation thread.

THEOREM 2.1. *The problem of finding the optimal policy π^{e*} given the policies of all other threads π^{-e} is equivalent to solving the following problem:*

$$\pi^{e*}|\pi^{-e} = \arg \max_{\pi^e} \int_{\Omega^e} p^e(\omega^e|\pi^e) \mathcal{U}^e(\omega^e|\pi) d\omega^e, \quad (2)$$

where $\mathcal{U}^e(\omega^e|\pi) \equiv \int_{\Omega^{-e}} p^{-e}(\omega^{-e}|\pi^{-e}) u(\omega^e, \omega^{-e}) d\omega^{-e}$, and $p^{-e}(\omega^{-e}|\pi^{-e})$ is the joint probability of $\omega^{-e} \in \Omega^{-e}$.

The main insight of Theorem 2.1 is that finding the optimal policy assuming optimal behavior in all other threads *given the policies on all other threads* is equivalent to finding the optimal policy in a single thread negotiation in which the agent has only probabilistic information about the utility function that is summarized by the expectation \mathcal{U}^e .

To solve Equation 2, we need to calculate p^e and \mathcal{U}^e . p^e can easily be found [5]. This leaves the calculation of \mathcal{U}^e as the core problem.

Exact evaluation of \mathcal{U}^e is $O(K^{TE})$. To find an efficient approximation, we rewrite \mathcal{U}^e as

$$\mathcal{U}^e(\omega^e|\pi) \equiv \int_{\Omega^{-e}} p^{-e}(\omega^{-e}|\pi^{-e}) u(\omega) d\omega^{-e} = \mathbb{E}_{p^{-e}}[u(\omega)] \quad (3)$$

This means that \mathcal{U}^e is an expectation over some probability distribution p^{-e} and can be approximated by sampling.

So far we showed that solving for one thread given the solution for all other threads can be done using sampling. This process is then repeated iteratively until convergence which constitutes the proposed algorithm called Iterative Thread Policy Enhancement (ITPE). Convergence is guaranteed by the following theorem:

¹ ω is a tuple of outcomes on all edges. The special outcome ϕ represents disagreement.




Figure 2: Utilities against state-of-the-art suppliers.

THEOREM 2.2. *Given a confidence level c where $0 < c < 1$, using $\frac{-\ln(1-c)}{2\epsilon^2}$ samples for approximating \mathcal{U}^e can guarantee – with confidence c – the convergence of ITPE to a multithreaded-policy π where $\mathcal{EU}(\pi^+) - \mathcal{EU}(\pi) < 2\epsilon$ and π^+ is an 2ϵ -Nash-Equilibrium for the induced negotiators’ game (i.e. \mathcal{EU} cannot be improved more than 2ϵ by changing a single policy).*

The subproblem of finding the optimal policy is solved using the Quick Greedy Concession Algorithm (QGCA) which has a time complexity of $O(KT)$ for static models. Extensions to more general acceptance models are known [6] and can be employed here.

To adapt ITPE to non-stationary acceptance models, the following simple online approach can be used. Given a change δ in the acceptance probability of outcome ω^e for thread e , it is trivial to show that \mathcal{U}^e and p^x for $x \neq e$ do not change (because they do not depend on $a^e(\omega^e)$). The main idea of the Predictive Policy Enhancement Algorithm (PPE) is to avoid re-sampling when changes in the acceptance model are detected by reusing existing samples weighing them according to the change in the acceptance model.

3 EVALUATION

We compared the performance of ITPE to two state-of-the-art groups of *concurrent negotiation* agents [7, 11] negotiating with state-of-the-art edge agents in a series of experiments. Fig. 2 shows the results of one of them. PPE outperformed all other algorithms, while ITPE behaved roughly at the same level as the method proposed by Williams et al. [11] while being able to handle more general scenarios and having theoretical guarantees in special cases.

Despite these encouraging results, the proposed system has some limitations that we provide pointers to possible directions of future research to handle these limitations. The most obvious limitation of the proposed system is that it does not scale well to negotiations with numerous outcomes or with a continuous outcome-space.

CONCLUSION

This paper introduced the optimal policy discovery problem in the context of concurrent negotiations with a global utility function. It provided the first approximate offline solution to the problem when acceptance models of different partners are known and stationary. An online greedy version is then introduced for non-stationary acceptance models. Evaluation experiments show that the proposed method can outperform state-of-the-art methods and is robust to errors in acceptance models. The proposed method suffers from scalability issues as it is linear in the size of the outcome-space which itself is exponential in the number of negotiation issues.

REFERENCES

- [1] Bedour Alrayes, Özgür Kafalı, and Kostas Stathis. 2018. Concurrent bilateral negotiation for open e-markets: the CONAN strategy. *Knowledge and Information Systems* 56, 2 (2018), 463–501.
- [2] Reyhan Aydoğan, David Festen, Koen V Hindriks, and Catholijn M Jonker. 2017. Alternating offers protocols for multilateral negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*. Springer, 153–167.
- [3] Cuihong Li, Joseph Giampapa, and Katia Sycara. 2006. Bilateral negotiation decisions with uncertain dynamic outside options. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 36, 1 (2006), 31–44.
- [4] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. 2014. Genius: An Integrated Environment for Supporting the Design of Generic Automated Negotiators. *Computational Intelligence* 30, 1 (2014), 48–70.
- [5] Yasser Mohammad. 2021. Concurrent local negotiations with a global utility function: a greedy approach. *Autonomous Agents and Multi-Agent Systems* 35, 2 (2021), 1–31.
- [6] Yasser Mohammad. 2021. Optimal Deterministic Time-Based Policy in Automated Negotiation. In *PRIMA 2020: Principles and Practice of Multi-Agent Systems: 23rd International Conference, Nagoya, Japan, November 18–20, 2020, Proceedings* 23. Springer International Publishing, 68–83.
- [7] Yasser Mohammad, Enrique Areyan Viqueira, Nahum Alvarez Ayerza, Amy Greenwald, Shinji Nakadai, and Satoshi Morinaga. 2019. Supply chain management world. In *International Conference on Principles and Practice of Multi-Agent Systems*. Springer, 153–169.
- [8] Thuc Duong Nguyen and Nicholas R Jennings. 2004. Coordinating multiple concurrent negotiations. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent System*, Vol. 3. IEEE Computer Society, 1064–1071.
- [9] Thuc Duong Nguyen and Nicholas R Jennings. 2005. Managing commitments in multiple concurrent negotiations. *Electronic Commerce Research and Applications* 4, 4 (2005), 362–376.
- [10] Kwang Mong Sim and Benyun Shi. 2009. Concurrent negotiation and coordination for grid resource coallocation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 40, 3 (2009), 753–766.
- [11] Colin R Williams, Valentin Robu, Enrico H Gerding, and Nicholas R Jennings. 2012. Negotiating concurrently with unknown opponents in complex, real-time domains. In *Proc. of the Twentieth European Conference on Artificial Intelligence*. 834–839.