Budget-Feasible Mechanism Design for Cost-Benefit Optimization in Gradual Service Procurement

Extended Abstract

Farzaneh Farhadi
Aston University
Birmingham, UK
f.farhadi@aston.ac.uk

Maria Chli
Aston University
Birmingham, UK
m.chli@aston.ac.uk

Nicholas. R. Jennings
Loughborough University
Loughborough, UK
n.r.jennings@lboro.ac.uk

ABSTRACT

We consider a procurement problem where a software agent procures multiple services from self-interested providers with private costs and uncertain reliabilities to complete a budget-limited task before a strict deadline. Over the last decade, several truthful budget-feasible procurement mechanisms have been developed to extract the true cost information from strategic providers. Most of these mechanisms have focused on maximizing the procurer's value (e.g., the task's success probability), and hence procuring as many services as the budget allows, even if the returned benefit is lower than the incurred cost. In this paper, however, we focus on the more realistic objective of balancing the cost-benefit tradeoff and propose a novel approach for designing budget-feasible mechanisms that invoke services gradually over time and whenever they are cost-optimal. A major barrier to achieving this goal was the strong dependencies among the decision variables caused by budget constraints. We overcome this barrier by proposing a conservative decomposable approximation to budget constraints. This is the first such approximation technique, which opens a path toward designing budget-feasible mechanisms for contingent planning problems.

KEYWORDS

Budget-Feasible Mechanism Design; Self-interested Agents; Cost-Benefit Tradeoff; Service Procurement

ACM Reference Format:
Farzaneh Farhadi, Maria Chli, and Nicholas. R. Jennings. 2023. Budget-Feasible Mechanism Design for Cost-Benefit Optimization in Gradual Service Procurement: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

The complex dynamic structure of today's markets increases the need for just-in-time service procurement where the necessary elements of a service are only marshaled together when requested by a customer [1]. Composing a service depends on procuring a number of subtask services that are often outsourced to multiple self-interested service providers, for example, providers at the next echelon in the supply chain. These service providers may offer uncertain delivery times while setting a wide range of prices. At the same time, the procurer's request for each subtask is characterized by a limited, fixed budget and a strict deadline, imposed by the market constraints. This mismatch makes for a very challenging problem to resolve. In such settings, the need for a procurement mechanism that can manage the limited budget to achieve the optimal cost-benefit tradeoff, whilst remaining agile, is essential.

The required mechanism belongs to the class of budget-feasible mechanisms, where the payments used for either service procurement or supporting truthfulness should satisfy the budget constraint [2, 4, 5, 7, 13]. The introduction of budget constraint, which applies not to the costs but to the payments, reveals a new dimension of difficulty to mechanism design, as it introduces strong inter-dependencies among decision variables. This challenge comes from the fact that the truth-inducing payments not only depend on the providers' actual bids, but also on the set of all possible bids that could be submitted by the providers [9].

Most of the available techniques for budget-feasible mechanism design have concentrated on problems with discrete decision spaces, as the procurements are often assumed to be made simultaneously at time 0 or at a few predetermined points in time (so the problem is equivalent to selecting one or more subsets of the available services) [6, 8, 11, 15]. These methods are also often aimed at maximizing a value function that depends only on procurements and not on payments [3, 5, 8, 11, 15]. These two assumptions reduce the budget-feasible mechanism design problem to choosing a set(s) of services that brings the highest possible value to the procurer while satisfying the budget constraint.

Our work is distinct from this literature in two aspects:

(1) **Payment-dependent objective:** The goal of our problem is to balance the cost-benefit tradeoff, which requires including payments in the objective function. This feature, which prevents the procurer from wasting the budget to procure low-quality services, is valuable as it allows the procurer to save the budget for future use. However, it brings new challenges, as it violates the monotonicity of the objective function. This violation is due to the fact that procuring a new service may lead to a decrease in the value function if the returned benefit does not compensate the costs.
(2) **Continuous decision space:** Gradual service procurement, when each service is procured at an optimal time, has shown to be a powerful tool for optimizing cost-benefit tradeoff [14]. However, it makes the problem extremely challenging as it introduces a continuous dimension to the mechanism designer's decision making.

Although there have been significant advances in the field of budget-feasible auction mechanisms, the existing mechanisms are not powerful enough to handle our problem in the setting above.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

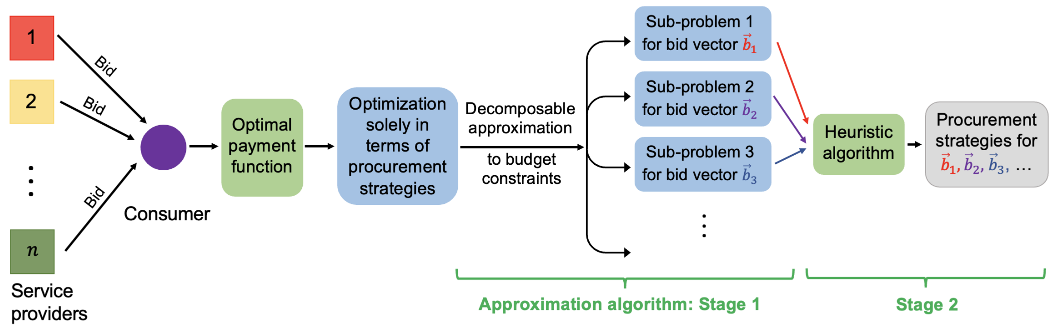

Figure 1: The key elements of our proposed procurement mechanism

In this paper, we summarize our budget-feasible mechanism design problem and outline a gradual procurement mechanism that overcomes the aforementioned challenges. The idea behind this mechanism is not restricted to procurement problems, but can be widely used for designing cost-efficient budget-feasible mechanisms for contingent-planning problems [10, 12]. The proposed mechanism is built upon our previous work presented in [9].

2 PROBLEM SPECIFICATION

A consumer C with a budget B would like a task to be completed before a deadline D. The task has a value V for the consumer if and only if it is executed before the deadline. There are n service providers, given by the set $\mathcal{N} = \{1, \ldots, n\}$, that can perform the task for the consumer. The consumer can invoke any number of providers at arbitrary times to increase the chance of success. However, it must compensate providers for their efforts with the payments that at least cover their costs. This creates a tradeoff for the consumer between success probability and invocation cost. However, the consumer does not have full information to optimally balance this tradeoff. This is mainly because the cost incurred by each provider for performing the task is its own private information and the provider may misrepresent this information, if it promises to increase its profit. In such settings, the consumer should design the procurement strategies and the payments, which together are called a procurement mechanism, such that the following conditions are fulfilled: 1) The providers voluntarily participate in the procurement process and do not regret their participation (ex-post individual rationality (ex-post IR)), 2) Irrespective of what others do, each provider finds truthful disclosure of its private information to be optimal (Dominant strategy incentive compatibility (DSIC); 3) The payments do not exceed the consumer's budget limit (Budget-feasibility); and 4) The outcome strikes a balance between success value and invocation cost (Cost-benefit efficiency).

There are four main challenges in designing such a procurement mechanism:

(i) The decision space is infinite-dimensional as the procurement strategies and payments need to be designed for every possible cost vector that could be declared by the providers.
(ii) Budget-feasibility constraints induce strong inter-dependencies among decision variables. Therefore, the problem is not decomposable to simpler finite-dimensional problems.
(iii) Cost-benefit optimization requires including payments into the objective function, which violates its monotonicity. This prevents the consumer from using standard techniques developed to handle the budget-feasibility constraints for problems with monotone objective functions [7, 13].
(iv) Invoking providers at optimal and not pre-determined times adds a continuous dimension to the problem, which has not been investigated before.

In the next section, we briefly discuss our design technique and how it helps overcome the above-mentioned challenges.

3 OUR DESIGN TECHNIQUE

We design a *direct* procurement mechanism (see Fig. 1), where the providers are asked to reveal their private cost information directly in terms of bids, and the procurement strategy and the payments are determined based on these bids. Our first step towards this goal is to design a payment function that can guarantee DSIC and ex-post IR of any gradual procurement mechanism at the minimum possible cost, and hence is proved to be optimal for our problem.

Designing such a payment function enables us to transform our problem into a simpler optimization problem, which aims to design the optimal procurement strategy for every possible bid vector. This optimization problem has a mixed continuous/discrete infinite-dimensional search space, a non-monotone objective function, and a set of budget constraints that makes it non-decomposable by intertwining the decision variables. Non-decomposability of this problem makes it extremely difficult, if not impossible, to solve, without effective approximations. Therefore, we propose a two-stage approximation algorithm to tackle the problem. The principal idea behind our approximation algorithm is to build a set of decomposable constraints that are stronger than the budget constraints and can be substituted for them as an approximation (Stage 1). This approximation enables us to decompose the problem into an infinite number of finite-dimensional sub-problems and hence overcome Challenges (i) and (ii). We then in Stage 2, propose a low-complexity heuristic algorithm to handle Challenges (iii) and (iv). Our heuristic algorithm searches the mixed continuous/discrete search space of each decomposed sub-problem and finds an approximated cost-benefit optimal gradual procurement strategy for each possible bid vector. Through both game-theoretical and empirical analysis, we prove that our proposed procurement mechanism satisfies the goals we set out in Section 2.

REFERENCES

[1] Aslı Aksoy, Eric Sucky, and Nursel Öztürk. 2014. Dynamic strategic supplier selection system with fuzzy logic. *Procedia-Social and Behavioral Sciences* 109 (2014), 1059–1063.

[2] Nima Anari, Gagan Goel, and Afshin Nikzad. 2014. Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. 266–275.

[3] Eric Balkanski, Pranav Garimidi, Vasilis Gkatzelis, Daniel Schoepflin, and Xizhi Tan. 2022. Deterministic Budget-Feasible Clock Auctions. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2940–2963.

[4] Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. 2012. Budget feasible mechanism design: from prior-free to bayesian. In *Proceedings of the 44th annual ACM symposium on Theory of computing (STOC)*. 449–458.

[5] Arpita Biswas, Shweta Jain, Debmalya Mandal, and Y Narahari. 2015. A Truthful Budget Feasible Multi-Armed Bandit Mechanism for Crowdsourcing Time Critical Tasks. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1101–1109.

[6] Hau Chan and Jing Chen. 2016. Budget Feasible Mechanisms for Dealers. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 113–122.

[7] Ning Chen, Nick Gravin, and Pinyan Lu. 2011. On the Approximability of Budget Feasible Mechanisms. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 685–699.

[8] Ludwig Ensthaler and Thomas Giebe. 2014. A dynamic auction for multi-object procurement under a hard budget constraint. *Research Policy* 43, 1 (2014), 179–189.

[9] Farzaneh Farhadi, Maria Chli, and Nicholas R. Jennings. 2023. Optimal and Efficient Auctions for the Gradual Procurement of Strategic Service Provider Agents. *Journal of Artificial Intelligence Research (JAIR)* (2023).

[10] Jörg Hoffmann and Ronen Brafman. 2005. Contingent planning via heuristic forward search with implicit belief states. In *Proceedings of the 15th International Conference on Automated Planning and Scheduling (ICAPS)*. 71–80.

[11] Minming Li, Chenhao Wang, and Mengqi Zhang. 2022. Budget feasible mechanisms for facility location games with strategic facilities. *Autonomous Agents and Multi-Agent Systems* 36, 2 (2022), 1–22.

[12] Dorin Shmaryahu, Guy Shani, and Jörg Hoffmann. 2019. Comparative criteria for partially observable contingent planning. *Autonomous Agents and Multi-Agent Systems* 33, 5 (2019), 481–517.

[13] Y. Singer. 2010. Budget Feasible Mechanisms. In *Proceedings of the 51st IEEE Symposium on Foundations of Computer Science (FOCS)*. 765–774.

[14] S. Stein, E.H. Gerding, A.C. Rogers, K. Larson, and N.R. Jennings. 2011. Algorithms and mechanisms for procuring services with uncertain durations using redundancy. *Artificial Intelligence* 175, 14 (2011), 2021 – 2060.

[15] Jun Wu, Yuan Zhang, Yu Qiao, Lei Zhang, Chongjun Wang, and Junyuan Xie. 2019. Multi-Unit Budget Feasible Mechanisms for Cellular Traffic Offloading. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*.