Mechanism Design for Ontology Alignment

(Extended Abstract)

Piotr Krysta
University of Liverpool, UK
P.Krysta@liverpool.ac.uk

Minming Li
City University of Hong Kong, China
minming.li@cityu.edu.hk

Terry R. Payne
University of Liverpool, UK
T.R.Payne@liverpool.ac.uk

Nan Zhi
University of Liverpool, UK
N.Zhi@liverpool.ac.uk

ABSTRACT

The aim of the ontology alignment problem is to find meaningful correspondences between two ontologies represented as collections of entities. This problem can be modelled as a novel mechanism design problem on an edge-weighted bipartite graph, where each side of the graph holds each agent's private entities, and the objective is to maximise the agents' social welfare. Having studied implementation in dominant strategies with and without payments, we report on findings that for truthful mechanisms, these problems need to be solved optimally. We also study greedy allocation rules with a first-price payment rule, and implementation in pure, mixed & Bayesian Nash equilibria, and have found tight bounds on the price of anarchy and stability.

Keywords

Mechanism Design; Ontology Alignment; Non Cooperative Game Theory

1. INTRODUCTION

In open, distributed computing environments, agents can only communicate if they resolve the differences in the ontological models that they used to represent a common domain. This is typically done by computing an *alignment*: a set of *correspondences* (or *mappings*) stating logical relationships between the entities in the different ontologies [?]. Whilst conventional approaches rely on third parties to compute such an alignment when the ontologies are in the public domain, it is not suitable when the exact composition of each ontology may be private to the owner (an agent, institution, commercial organisation, etc). Furthermore, the ontological model may therefore be partially observable to other agents, as some knowledge encoded within the ontologies may be confidential or commercially sensitive.

As the composition of different subsets of correspondences can result in different alignments, the challenge in negotiating a mutually acceptable alignment is that of selecting and proposing a set of correspondences that result in a preferred alignment for both agents. Whilst various approaches have proposed protocols for exchanging private correspondences with the goal of aligning their respective ontologies, the naive strategies used rely on truthful revelations, and fail to consider how other agents may respond. We model this problem as an edge-weighted bipartite graph where the vertices correspond to entities in the agents' individual ontologies and the edges correspond to candidate correspondences. The objective is to find a *matching* within the graph (i.e. an alignment) that satisfies the aims of both agents, based on a subset of correspondences proposed by each agent. As the ontologies can vary greatly in size, with several in the BioMedical domain possessing tens of thousands of entities, the approach should be computationally tractable.

We explore this problem from a mechanism design perspective, and analyse[1] implementations in *Dominant Strategies* (centralised mechanisms) and in *Nash Equilibria* (decentralised mechanisms). For the implementations in Dominant Strategies, two alternate settings are considered: *with payment*, and *without payment*; where the problem is characterised as a social welfare maximising matching setting, with an additive valuation function.

2. THE ALIGNMENT PROBLEM

We consider a setting in which there are two agents $i \in \{L, R\}$ (the *left* agent and *right* agent), where each agent possesses a private ontology $\mathcal{O}_i$, containing *named concepts*. The alignment is modelled as an edge-weighted bipartite graph $G = (U \cup V, E)$, where the vertices of U and V correspond to these named concepts (i.e. entities) in the agents' individual ontologies respectively, and the edges $e \in E$ correspond to the candidate correspondences. A matching M is a subset of E such that $e \cap e' = \varnothing$ for all $e, e' \in M$ with $e \neq e'$. Each agent $i \in \{L, R\}$ has a non-negative valuation function for different matchings M, denoted $v_i(M)$, where $v_i : M(G) \to \mathbb{R}^+$, which is additive; i.e. $v(S) + v(T) = v(S \cup T)$ such that $S \cap T = \varnothing$ for all $S, T \in M$, and $M(G)$ is the set of all matchings in a graph G. The agents also have the valuation function $v_i : E \to \mathbb{R}^+$ to represent the value $v_i(e)$ the agent i can get from the edge e. The combined value for an edge e is given as $v(e) = v_L(e) + v_R(e)$. Observe that $v_i(M) = \sum_{e \in M} v_i(e)$ for every agent $i \in \{L, R\}$. The goal is to establish an alignment which is equivalent to a matching M that maximises the sum of the combined edges; i.e. $\sum_{e \in M} v(e)$ is maximised. The valuation function v_i rep-

[1]The proofs for the theorems in this extended abstract will appear in the full version of this paper.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

resents the agent's true valuation, or *type* that it attributes to each matching. We use v to represent the combined type profile for both agents, such that $v = \{v_L, v_R\}$, where v_i is the type profile for agent i, and similarly, b denotes the combined bid profile for both agents, such that $b = \{b_L, b_R\}$, where b_i is the bid profile for agent i.

3. AN IMPLEMENTATION IN DOMINANT STRATEGY

To determine the outcome given the bids of the two agents, we consider mechanisms with and without payments. We define a direct revelation mechanism $\mathcal{M}(\mathcal{A}, \mathcal{P})$, which is composed of an allocation rule $\mathcal{A}$ to determine the outcome of the mechanism, and a payment scheme $\mathcal{P}$ which assigns a vector of payments to each declared valuation profile. For the mechanism with payment, the mechanism proceeds by eliciting a bid profile b_i from each agent i, and then applies the allocation and payment rules to the combined bid profiles to obtain an outcome and payment for each agent. As an agent may not want to reveal its type, we assume that b does not need to be equal to v.

For the mechanism without payment, we consider a restricted model of the declaration, whereby each agent's valuation on an edge e, $v_i(e)$ is public (or at least verifiable), and thus $\forall e \in E, b_i(e) = v_i(e)$. What is private is a set of desirable edges $E_i \subset E$ that the agent wants in the outcome. Each agent i therefore declares a boolean value for each edge, denoted $\delta_i(e) \in \{0, 1\}$, such that $\delta_i(e) = 1$ *iff* $e \in E_i$.

3.1 Mechanism design with payment

In this setting, both agents have to pay money to establish a matching. If we are willing to solve the problem optimally (which is possible in polynomial time by simply finding an optimal weighted bipartite matching), then we can use the classic VCG mechanism with Clarke payment. This begs the question: is it possible to have a faster, non-optimal, approximate and truthful mechanism for our problem? We have found that the answer is essentially no.

THEOREM 1. *For the alignment problem with payment, any mechanism not adopting an optimal solution when agents declare their true valuations is either non-truthful, or has an approximation ratio of at least 2.*

THEOREM 2. *For the alignment problem with payment, any deterministic mechanism which does not adopt optimal solution when agents declare true valuation, is either non-truthful, or is maximal-in-range.*

Observe that putting together Theorems **??** and **??** we conclude that for our problem and mechanism design with payment, the only truthful mechanisms are ones that are maximal in range and have approximation ratio at least 2. To complement these lower bound results we have found a very simple truthful mechanism which indeed has an approximation ratio of 2, does not produce an optimal solution and has therefore to be maximal in range.

3.2 Mechanism design without payment

In the mechanism design without payment, if agents can misreport their valuations, no non-trivial truthful mechanism exists, see [**?**]. A natural setting commonly used in previous research assumes that an agent can only declare or hide which edge it wants to match. We will thus adopt this restricted model of the declaration. We assume that agents cannot lie about their valuations, but they may lie about which edge can be used to establish a matching. An instance of the alignment problem on a private bipartite graph is: the valuations of agents on edge e, $v_i(e)$ are public information or verifiable, and agent i's private information is a set of edges $E_i \subset E$ given by $\delta_i(e) \in \{0, 1\}$. An edge e is possible to be accepted in the matching, only if for both agents, $\delta_L(e) = \delta_R(e) = 1$. The agent i will receive value $v_i(e)$ from e if it is matched, and otherwise it receives 0 from edge e. The goal is to maximize the social welfare via a mechanism without money, such that both agents are incentivised to declare their E_i truthfully.

We designed a polynomial time algorithm which determines that, given an instance, if a deterministic truthful mechanism exists with a bounded approximation ratio; if so, then the optimal solution is found. However, if such a mechanism does not exist for the bid, then we show there is no truthful mechanism with a bounded approximation ratio.

THEOREM 3. *There are no randomized mechanisms that are universally truthful and have approximation ratios better than 2 for the setting.*

THEOREM 4. *There are no randomised mechanisms that are truthful in expectation and have approximation ratios better than 1.333 for the setting.*

4. NASH EQUILIBRIA IMPLEMENTATION

Given our results on truthful centralised mechanisms, either the problem should be solved optimally (though costly) or strong lower bounds should be found for the approximation ratios of truthful mechanisms. Thus, we have explored an implementation in Nash equilibria to efficiently approximate mechanisms for matching, using the greedy allocation mechanism. We have found that coupled with the first-price payment scheme, this mechanism implements Nash equilibria which are very close (within a factor of 4) to the optimal matching. Specifically, we have characterised the *Price of Anarchy* of the greedy mechanism completely and have found that it is precisely 4. This bound on the price of anarchy holds even for Bayesian and Mixed Nash equilibria. Furthermore, we have also found that when a pure Nash Equilibrium exists, the *Price of Stability* is at least 2. This provides a complete picture of the complexity of mechanism design for our problem.

5. CONCLUSIONS

The Ontology Alignment negotiation problem was modelled as a mechanism design problem. Our results on truthful (centralised) mechanisms suggest that they require optimal solutions, which are time-costly for large instances. On the other hand, when we change our focus to decentralised mechanisms and we simply let the agents play, then such mechanisms implement Nash equilibria whose social welfare is close to the optimal social welfare.

REFERENCES

[1] J. Euzenat and P. Shvaiko. *Ontology Matching, Second Edition.* Springer, 2013.

[2] S. Dughmi and A. Ghosh, Truthful Assignment without Money. *Proceedings of the 11th ACM Conference on Electronic Commerce*, pages 325–334, 2010.