



Dengji Zhao <jason.dengji.zhao@gmail.com>

[AAMAS2016] Your Paper #699

AAMAS2016 <aamas2016@confmaster.net>

25 January 2016 at 11:24

To: Dengji Zhao <dengji.zhao@gmail.com>

Cc: Sarvapali Ramchurn <sdr@ecs.soton.ac.uk>, Nick Jennings <nj@ecs.soton.ac.uk>, AAMAS2016 <aamas2016@confmaster.net>

Dear Dengji Zhao, Sarvapali Ramchurn, Nick Jennings,

We are pleased to inform you that your paper #699 with title: Fault Tolerant Mechanism Design for General Task Allocation has been accepted for publication as a *Full Paper* in the proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS2016).

The review process was extremely selective. Out of the 550 submissions that were reviewed, the programme committee selected 137 papers as Full Papers (8 pages or 4 pages for Blue Sky papers) and 131 papers as Extended Abstracts (2 pages). These totals include the Innovative Applications, Robotics, Embodied Virtual Agents and Human-Agent Interaction, and Blue Sky Ideas special tracks.

The reviews of your paper are included below. In addition to the text of the (possibly revised) initial reviews that you have already seen, the reviews include a numerical score (from 1 to 10) and, in some cases, the review of the SPC member and any additional reviews that were solicited during the discussion phase. Please be aware that the numerical ratings are not always indicative of the comparative quality of your paper. As you know, each PC and SPC member only saw a small subset of the overall submissions. In selecting the final programme, the programme chairs read through the detailed recommendations from the PC and SPC and also had a short live call with almost all of the SPCs to discuss the potential outcome of each individual paper. The decisions were based on the overall set of reviewed papers (even though we were able to follow the SPC recommendations in almost all cases).

Though there will be no further review of your paper, in some cases the acceptance decision was influenced by the assumption that you would improve your paper according to the reviewer comments, or your rebuttal. As you prepare the final version of your paper, please take into account the reviewers' suggestions.

Full Papers are allowed 8 pages in the proceedings plus any bibliographic references. The final versions of Full Papers are due by the *19th of February 2016*. In a later message we will provide you with the copyright and formatting instructions so that you can prepare and submit the final camera-ready version.

For your paper to be published, at least one author must register for the conference and attend to present it. We strongly advise that at least one of the authors registers by the earlier registration deadline (to be announced). There will be a final deadline for an author to register to secure inclusion of the paper in the proceedings. We will inform you about this deadline when we send the camera-ready instructions.

All accepted papers (both full papers and extended abstracts) will also be presented as posters in several poster sessions. As a result, the poster sessions will be a very important part of the overall programme and we look forward to having your poster there.

The conference programme will be available on the conference website in a few weeks.

Congratulations on having your paper accepted and we look forward to seeing you at AAMAS2016 this coming May in Singapore.

Best regards,

John Thangarajah and Karl Tuyls
AAMAS-2016 Programme Chairs----- Review -----

Relevance : 9
 Originality and novelty : 7
 Significance : 8
 Readability and organization : 9
 Technical quality and soundness : 8
 Overall recommendation (Full Paper) : 8
 Overall recommendation (Extended Abstract) : 8
 Reviewer Expertise : 9
 Reviewer Thoroughness : 7

-- Comments to the author(s):

Summary of the Paper

=====

The paper considers a task allocation setting with uncertainty that is induced by failures. The authors provide an ex-post truthful mechanism that incentivizes participation of agents and conclude by examining an application of this in ride sharing.

Justification & Comments

=====

The paper is easily readable and well organized. The setting being explored is relatively clean.

I thought that the "application to ride sharing" is not much more than an example, on a toy problem so I do not think the term "Application" should have been used here.

I think that the discussions in sec 6 are more interesting than 5, so I would reorder or feature them more prominently.

I wonder: are there (perhaps restricted) settings in which agents are not risk-neutral where a mechanism can still be found?

Questions for Rebuttal

=====

Does the allocation considered include redundancy? i.e., will several paths to completion be considered, or is only a single path selected that has the highest expected utility?

If it is considered, how hard is it in general to compute the optimal allocation, and also to apply the mechanism?

If it is not, then is there a way to extend the results to such settings?

-- Summary:
[Seen by the authors and (S)PC members]

----- End of Review -----

----- Review -----

Relevance	: 7
Originality and novelty	: 6
Significance	: 6
Readability and organization	: 7
Technical quality and soundness	: 7
Overall recommendation (Full Paper)	: 6
Overall recommendation (Extended Abstract)	: 8
Reviewer Expertise	: 6
Reviewer Thoroughness	: 7

-- Comments to the author(s):
Summary of the Paper

=====

The paper studies truthfulness and individual rationality of an allocation mechanism that explicitly takes into account execution uncertainty. It generalizes previously known condition on the truthfulness of the mechanism and explains how the mechanism could work in a ride sharing domain.

Justification for the Scores

=====

Relevance

The studied setting is relevant to AAMAS and practical applications. Mechanism design in task allocation is an important problem and the aspect of execution uncertainty is crucial for some applications.

Originality

The work is rather incremental. The main contribution is in generalizing some of the assumptions of the previous works and defining the sufficient and necessary condition for truthfulness of a specific class of mechanisms.

Significance

The presented work is nice from surely theoretical perspective, because it provides a complete characterization of the conditions that make POV mechanism truthful. However, the paper could have made a better job in explaining the practical implications of the presented generalizations.

What kind of practical problems can be modelled using the generalized multi-linear valuations, but could not be modelled using the existing assumptions? It would be better to make more explicit what parts of the example in Section 5 make the truthfulness guarantees in the existing work insufficient.

Still I believe that the paper presents a sufficiently interesting extension of understanding of a relevant class of mechanisms.

Readability and Organization

The paper is very well written and easy to understand especially at the beginning. The later sections feel a little detached and verbose. For example, the example in Section 5 works as demonstration of the concepts used in the paper, but not really as validation of the proposed method or a source of deeper understanding of applicability of the method. The example can be substantially shortened or better, the analysis should be deeper. It is not worth the space in the current form. The argument of Section 6 is also quite unclear. I understood the problem stated and the method that should be used to make the claim, but the actual claim in the last paragraph is quite unclear to me. It would be helpful to elaborate a little more on this part.

Technical Quality and Soundness

I did not find any technical problems in the paper. The theorems seem to be correct and the relation to the precious world clearly explained.

A minor issue with the model I have is that the success probability is defined for the whole set of tasks allocated to an agent. In real world examples, the agent may often fail to execute only a subset of the assigned tasks. Are there any problems in generalizing the framework in this direction?

Furthermore, the trust part is a rather strange. The framework assumes that there is an aggregation function that will aggregate the individual agent's trust to actual true success probability some other agent. This does not sound realistic at all. Trust is generally quite an imprecise value and it may be rational not to be truthful about it if an agent has some believe about the precision of other agents' trust estimates. The mapping should be better argued for to make it believable.

Questions for Rebuttal

=====

1) Are the conditions from truthfulness of the PEV mechanism from previous works insufficient to assure truthfulness in the example in Section 5?

-- Summary:

I like the theoretical contribution of the paper that completes the characterization of situations in which the studied

mechanism is truthful. The paper could have better explain practical implications of the generalization.

----- End of Review -----

----- Review -----

Relevance	: 8
Originality and novelty	: 6
Significance	: 6
Readability and organization	: 5
Technical quality and soundness	: 8
Overall recommendation (Full Paper)	: 7
Overall recommendation (Extended Abstract)	: 7
Reviewer Expertise	: 5
Reviewer Thoroughness	: 3

-- Comments to the author(s):

Summary of the Paper

=====

This paper studies a particular mechanism termed "the PEV-based mechanism." Like VCG it chooses an efficient allocation, but unlike VCG it is designed for settings where agents may be assigned to carry out actions that may fail. Thus payments differ, and in particular depend on whether the actions succeeded (i.e. Post-Execution Verification). This mechanism has been applied in specific settings, but this paper proved a general result characterizing when it is ex-post incentive compatible, and illustrates an application to ride sharing.

Justification for the Scores

=====

As far as I can tell this paper is correct, but I found it very dense and tough going. Even in places where intuition is provided, I often found the intuition too technical to really help (e.g. after Definition 4). I am not familiar with the techniques used in [14] or [15], so I am uncertain how much of a delta this technical development represents, but I am willing to give it the benefit of the doubt. Similarly, the example of ridesharing involves very rich preferences---I would say probably too rich for a real setting. So for me this isn't really compelling as an example of something this extension allows us to do.

That said, the results seem potentially interesting. While they are restricted to a particular mechanism, it is a natural one, and this may be a benefit as it points to future opportunities to explore mechanisms that do not necessarily allocate efficiently, as in work on partial verification. So as before I am inclined to give this paper the benefit of the doubt.

Constructive Comments to Improve the Paper

=====

The mechanism could really use a better name than "the PEV-based mechanism"

Questions for Rebuttal

=====

Why is it without loss of generality to assume that Bob, Tom, and Alice all have $p = 1$?

-- Summary:

A very dense technical presentation with a not-entirely-convincing example application, but still potentially interesting.

----- End of Review -----

////////////////////////////////////
Powered by ConfMaster.net
////////////////////////////////////