Formal Verification for Fostering the Design of Mechanisms for Social Good

Munyque Mittelmann
University Federico II
Naples, Italy
munyque.mittelmann@unina.it

Aniello Murano
University Federico II
Naples, Italy
aniello.murano@unina.it

Laurent Perrussel
Université Toulouse Capitole, IRIT
Toulouse, France
laurent.perrussel@irit.fr

ABSTRACT

Social Good requires to be innovative while designing new resource allocation protocols. New key properties have to be addressed such as fairness, diversity, or equity on top of the classical properties such as strategyproofness and individual rationality. In this paper, we show how we can take advantage of the recent progress in formal verification to fully revisit the automation of Mechanism Design. The contribution will be at first the technical progress on logics for strategic reasoning and then, more importantly, a road map for (i) representing key Social Good properties in Social and (ii) synthesising mechanisms handling these properties.

KEYWORDS

Formal Verification, Strategic Reasoning, Mechanism Design

1 INTRODUCTION

Mechanism Design (MD) is a central problem in our digital society and it consists of designing new games (a.k.a. mechanisms) for aggregating preferences in multi-agent settings [11]. Such mechanisms may be used for addressing inequality issues, climate change, or civic involvement. The key MD challenge is to design a mechanism that chooses a good outcome, with respect to the designing criteria, even though agents may be self-interested and may lie about their preferences. The designing criteria may specify a preferable behavior of the players (e.g. truthfulness) as well as desirable features of the outcome[24] (e.g. social welfare maximization). In principle, almost any kind of market-based institution or organization can be viewed as a mechanism[23]. Some classical types of mechanisms are voting systems, household allocation, ride-sharing platforms, auctions, and fair division protocols. Moreover, in recent years, there has been a growing effort at designing novel mechanisms for a wide variety of problems and settings, including school choice programs with quotas and matching problems with individual and regional quotas (eg. [15]). Daily decisions based on algorithms impact the whole society and Social Good (SG) can't be ignored: at first, by offering tools for automating and verifying the design stage and second by allowing non-MD experts to be involved in the definition of future mechanisms and assessment in terms of social good.

Although logic-based languages have been widely used for verification [10] and synthesis [12] of Multi-Agent Systems (MAS), the use of formal methods for reasoning about strategic behavior in mechanism design has not been much explored yet. An advantage in adopting such a perspective lies in the high expressivity and generality of logics for strategic reasoning [25]. Moreover, by relying on precise semantics, formal methods provide tools for rigorously analyzing the correctness of systems, which is important to improve trust in mechanisms generated by machines. The problem of formally reasoning about mechanisms is, however, nontrivial: it requires considering quantitative information (e.g., utilities and payments), private information about the participant's preferences, and complex solution concepts (such as strategy dominance). Pauly and Wooldridge in [25] first argued that strategic logics developed for the formal verification of MAS could be good candidates as formal frameworks to reason about mechanisms. They considered Alternating-time Temporal Logic (ATL) [4] and showed with two case studies based on voting systems that some relevant properties for the verification of such systems can be expressed in this logic. However, ATL is not expressive enough to represent complex solutions concepts, a key condition for representing mechanisms and there has been no significant progress for years. Three key contributions, Strategy Logic [9] and its quantitative and probabilistic extensions [5, 8] were the breakthrough for revisiting Pauly and Woodridge's initial intuition. Recently, Maubert *et al.* [19] show that such language is expressive enough to tackle the main challenges raised by Mechanism Design. They first show that assessing a mechanism, i.e., evaluating it with its target properties, boils down to encoding a model checking problem [10] and second, designing a mechanism, with respect to a set of expected properties, can be rephrased as a synthesis problem [21]. Finally, Mittelmann *et al.* [22] has shown how to lift the approach to handle the verification of Bayesian and stochastic mechanisms.

In the following sections, we detail the main milestones for addressing this challenge: how to fully revisit the traditional Mechanism Design process for addressing the huge need for mechanisms with a Social Good dimension. We need automation for scaling up and Formal Verification is a promising tool for this problem.

2 A ROADMAP

The goal of this position paper is to sketch out the key steps for taking advantage of the recent progress in formal verification for revisiting the automation of Mechanism Design and fulfilling the need for new mechanisms addressing Social Good problems. Hereafter, we detail a two-step method: first, going further on logics for strategic reasoning and, more importantly, an extensive characterization and representation of the key properties considered in Mechanism Design for Social Good. The results will enable us to move to the second step: fully revisiting the typical design process of mechanisms. By taking advantage of tools and methods for automation, MD can then be rephrased as a human-AI collaboration process where logics for Multi-Agent Systems are the cornerstone. Model-checkers can be extended to synthesize different potential mechanisms based on the formal representation of the requirements; these latter will be elicited in cooperation with stakeholders and Mechanism Designers to collect key properties of an intended mechanism (e.g. social welfare). These properties will then be guaranteed by design as the synthesis procedure is correct by construction and will only explore the set of possible mechanisms: in other words, trust will also be guaranteed by design. If synthesis is not possible, model-checking will explain why: what are the conflicting requirements and, still based on model-checking techniques, candidate mechanisms may be generated for approximating the intended requirements. Still, this new perspective on designing mechanisms for Social Good is nontrivial and requires addressing two interwoven and challenging questions:

- *Going further in terms of Mechanism Design for Social Good*: typical SG problems are related to education, health, environmental sustainability, or security to name a few and there are strongly related to the popular Sustainable Development Goals[1]. There is a recent growing interest in bringing AI-based solutions to SG Problems [27]. Several classical MD properties are usually relevant when building SG mechanisms: e.g., Pareto optimality and truthfulness. However, some specific SG properties should also be addressed: social welfare and fairness are classical ones, but we may face ethical considerations in problems such as kidney exchange [13], vaccine distribution, diversity in school choice [3], and refugees' relocation [14]. All these less-standard properties should be explicitly defined and discussed. It will help to characterize different families of mechanisms.
- *Going further in terms of expressiveness and computation*: addressing a large scope of mechanisms is a key success factor. The specification Language should be able to consider qualitative and quantitative information, tolerate to represent different notions of approximation [17], and consider imperfect information, uncertainty, and high-order beliefs. As we go further on the specification language side, the computational cost should also be considered. During these last years, even if model-checkers and SAT tools have made huge progress, the computational complexity of ATL and Strategy Logic is still an obstacle. Eliciting fragments leading to the improvement of the complexity is a key issue as it will pave the way for using model-checkers and SAT-based software at scale.

By addressing these challenges, we will be in a position where humans may be helped by machines for assessing and validating potential mechanisms. This new perspective will help to close the gap between stakeholders and mechanism designers. This is a key issue when Social Good is at the heart of the definition of a mechanism. The Formal Verification toolbox will help to assess different promising variants of mechanisms. It will also provide techniques for explaining properties inferred by those variants. Finally, Formal Verification brings a trustworthy dimension: all variants will be well-defined and explainable by design.

3 AN ILLUSTRATION

Mechanism Design is, up to now, a non-trivial task as it requires manually proving each property. As mentioned earlier, tools are there for automating the design stage, allowing us to give more importance to the definition of the goals of the mechanism and its evaluation. Let us illustrate by first revisiting some recent results on the synthesis of classical allocation mechanisms, namely auctions. In [19], we argue that *Quantified Strategy Logic* ($\mathsf{SL}[\mathcal{F}]$) [8] is a powerful tool for reasoning about resource allocation protocols where:

(1) there is a strategic dimension: agents may not be cooperative and compete to gain access to some scarce resources;
(2) resources and utility functions that encode rationales for decision-making may be quantitative.

3.1 Representing Properties

The syntax of $\mathsf{SL}[\mathcal{F}]$ is defined as follows:

$$\phi ::= p \mid \exists s.\phi \mid (a,s)\phi \mid f(\phi_1, ..., \phi_n) \mid \mathsf{Next}\phi \mid \mathsf{Always}\phi \mid \mathsf{Future}\phi$$

where p is an atomic proposition, s is a variable representing a strategy, a is an agent symbol (Ag is the set of agent symbols), and f is a function symbol. The intuitive reading of the operators is as follows: $\exists s.\phi$ means that there exists a strategy such that ϕ holds; $(a,s)\phi$ means that when strategy s is assigned to agent a, ϕ holds; Next, Always, and Future are the usual temporal operators "next", "always" (universal temporal quantification), and "future" (existential temporal quantification). The meaning of $f(\phi_1, ..., \phi_n)$ depends on the function f. Boolean operators may be encoded as functions. The detailed semantics is described in [19, 21] but the underlying model is a state-transition model with two main characteristics:

- fluents have values in $[-1, 1]$ instead of Boolean ones;
- models are finite.

About strategies, we adopt the classical definition for memoryless strategies in logics for MAS: a function stating for each agent what actions it performs in each possible state.

Main requirement when considering Mechanism Design is the ability of expressing that a mechanism satisfies a property. To do so, the prerequisite is to capture the notion of equilibrium. The type θ_a of an agent a determines how she values each choice, and write $\boldsymbol{\theta}$ to denote a type profile describing a type for each agent in Ag; we assume some valuation function and thus an associated utility function which will be encoded by function symbol $util_a$. A strategy profile $\mathbf{s} = \bigcup_a s_a$ is a Nash equilibrium (NE) if no agent a can increase her utility (represented by function $util_a$) a unilateral change of strategy. The following $\mathsf{SL}[\mathcal{F}]$-formula characterizes Nash equilibria:

$$NE(\mathbf{s}, \boldsymbol{\theta}) := \bigwedge_a \forall t[(Ag_{-a}, s_{-a})(a,t)\mathsf{Future}(term \wedge util_a(\theta_a)) \\ \leq (Ag, s)\mathsf{Future}(term \wedge util_a(\theta_a))]$$

The above encoding means that verifying whether the strategy associated to variable $\mathbf{s}$ is an equilibrium w.r.t. $\boldsymbol{\theta}$ boils down to a model checking problem. Notice that fluent $term$ encodes the assumption that the mechanism is finite, i.e. there is a decision state. Transforming the existence of a Nash Equilibrium problem into a model checking problem has for primary effect to define this concept as a plain and general computational object. In other words, high level formal languages enable to represent classical Mechanism

[1] https://sdgs.un.org/

Design properties. As an example, we are able to represent what *Individual Rationality* as an SL[$\mathcal{F}$]-statement [19]:

$$\bigwedge_a util_a(\theta_a) \geq 0$$

Verifying that a mechanism satisfies such constraint consists in its encoding as a state-transition model and to verify if the property holds in all states which is an equilibrium (e.g., Nash equilibrium). In terms of computational complexity, checking whether a property holds in a model with imperfect information is PSPACE-Complete.

Social Good Mechanisms. Let us consider the allocation protocol detailed in [18]. In this work, the problem of allocating home to homeless people is encoded as a Temporal Resource Allocation Problem (TRAP). Clearly, such problem may be encoded as a state-transition model. The key issue is the proposed protocol is the fairness dimension. The authors suggests that there is no standard definition of fairness and adopt later on a group fairness definition; this is clearly related to the notion of coalition of agents which is a built-in concept in ATL and SL[$\mathcal{F}$]. Fairness is then rephrased in terms of statistical parity: each group of agents should of a metric z is similar for every group; let $\bar{z}$ the expected value of a metric over all groups: at some stage in the future, for all groups, the gap between the metric value and the expected one is no worse than ε:

$$\exists \mathbf{s}(\bigwedge_g (s_g, g)\mathsf{Future}\ (z(g) - \bar{z} \leq \varepsilon))$$

Notice that, in [22], we present a probabilistic variant of Strategy Logic that is able to encode Bayesian Mechanisms. Still, the variant combining quantity and probability has to be defined but such language will be expressive enough for checking formulas representing nontrivial concepts such as Group Fairness.

3.2 Synthesis of Mechanism

The next question is whether we can go further by providing a partial specification of an intended mechanism and then create a model that will encode a potential candidate mechanism. In [21], we show that it is actually feasible. Model checkers may be extended to output a model instead of a decision about the satisfaction of a formula. The method is the following:

(1) specify the overall behavior and key properties;
(2) with the help of a meta-algorithm, create a model and check if the specification is satisfied.

Let us illustrate this 2-step meth with the design of an Auction mechanism [21]. We first define the overall behavior: the auction should have one winner, be sequential protocol, consider a specific number of goods... For instance, the following statement specifies that the targeting auction protocol should guarantee that the winner (fluent *choice*) of the auction should pay the current price:

$$\mathsf{Always} \bigwedge_a (choice = win_a \rightarrow pay_a = price)$$

while the "non-winners" should not pay anything:

$$\mathsf{Always} \bigwedge_a (choice \neq win_a \rightarrow pay_a = 0)$$

Market Description Language [26] or Auction Description Language [20] are at first interest for such specification but are unable to address the strategic dimension. Hence, the specification should be completed by formulas like the following one stating that in any equilibrium, the protocol should be individual rational

$$\bigwedge_\theta \forall \mathbf{s}\, NE(\mathbf{s}, \theta) \rightarrow \bigwedge_a util_a(\theta_a) \geq 0$$

The second step consists of designing an algorithm for synthesizing a state-transition model encoding a possible mechanism. In [21], we detail such algorithm and, more importantly, some conditions describing whether a solution may be founded; in other words, the conditions that guarantee the decidability. A key dimension is the "price to pay": all synthesis algorithms are exponential, typically k-EXPTIME where k represents the number of alternations in the quantification of strategies.

Social Good Mechanisms. Let us revisit the homeless above example. In the paper, the authors assume that the protocol is fully designed and check its properties first and second, provide some experimental results. But, a different perspective may be considered:

- should the allocation procedure be like a repeated game?
- what are the key properties? Especially parity, fairness...
- what if one property prevents finding an eligible mechanism: can we consider an alternative definition?

Several mechanisms may then be created and, consequently, experiments may be run on different configurations and parameters. This is a key progress offered by the automation of MD: we not only give more importance to the properties but also create the conditions for more diverse experiments. In [18], the authors stress that the diversity dimension may be considered with a horizon: after n steps, parity should be reached. The built-in temporal dimension of the state transition model will help to compute the impact of the hypothesis about n or the computation of the horizon for n.

4 AUTOMATING MD FOR SOCIAL GOOD

As stressed by the above illustration, the very first challenge is to identify the key concepts and properties in mechanisms impacting Social Good and rephrased them as general properties. A promising starting point is the general definitions of mechanisms' properties in the specific context of auctions [19]. This characterization for Social Good will raise numerous challenging questions in terms of expressiveness as its associated properties require handling complex notions (the previous Parity example is a good illustration).

4.1 Going Further on Expressiveness

Three key challenges in MD and SG should be addressed for going further on automation; they are key as they condition the success of using formal verification for contributing to SG.

(1) What kind of memory the mechanism should assume about the stakeholders?
(2) Is there a stochastic dimension (uncertainty about information or resources)?
(3) How far should we go about the resources used for reasoning (resource-bounded reasoning)?

4.1.1 Memory impact. Up to now, key results related to auctions capture mechanisms where memoryless strategies are sufficient to represent the agents' behavior, that is, mechanisms where previous experience has no influence. An extensive range of mechanisms may assume that configuration, going from household allocation to school selection. However, when participating in sequential mechanisms, agents could gather information from other agents' behavior and act based on what happened in previous steps of the mechanisms (e.g., in an iterative voting procedure). For such situations, a promising alternative technique is to assume partial recall instead of perfect recall as it leads to undecidability. Ågotnes and Walther [1] investigate strategic abilities of agents with bounded memory, while Belardinelli *et al.* [7] consider bounded memory as an approximation of perfect recall. Both frameworks prevent undecidability issues and pave the way for considering formal verification techniques for SG mechanisms: vaccination or rent division mechanisms may assume less standard memory configurations.

4.1.2 Stochastic Mechanisms. Generalizing from the deterministic to the probabilistic setting is challenging: the wide and heterogeneous range of settings considered in the literature obscures the path for a general and formal approach to verification. The setting may consider deterministic or randomized mechanisms, incomplete information about agents' types (Bayesian mechanisms), mixed or pure strategies, and direct or iterative mechanisms. Moreover, considering Bayesian mechanisms brings out different methods for evaluating a mechanism according to the timeline for revealing the information as the mechanism is executed. The Main Challenge is to investigate how the probabilistic setting for synthesizing SG mechanisms introduces additional and useful criteria: expected optimality and approximation. Numerous SG mechanisms consider a stochastic dimension such as Vaccination or Refugee Relocation.

4.1.3 Bounded Resources Impact. Numerous variants of logics for strategic reasoning considered agents with bounded resources have been proposed (e.g.[2]). Those languages enable us to verify that to achieve some goals, an agent will only use a limited quantity of resources (including time). This is at first interest if we want to verify that an agent participating in a mechanism does not use unconsidered resources for her profit. The main challenge consists of assessing the impact of restricting resources on the behavior of a mechanism. Typically, Mechanism Design assumes that agents are strategic with an endless power for reasoning. If this assumption is no longer true, then one should consider what will be the impact on classical properties such as incentive compatibility or individual rationality and then on SG properties. For instance, fairness or equity may assume that all agents should have the same access to a limited number of resources.

4.2 Going Further on Scalability

Progress should not be restricted to the expressiveness dimension. The long-term goal is to impact social good and formal verification tools are one of the pillars for building the bridge between stakeholders, AI researchers, and MD specialists. To do so, we should take advantage of the existing model checkers but deeply improve their scalability. To do so, two inter-related questions on computational complexity and techniques should be addressed:

(1) Is there any fragment of logics for strategic reasoning offering a benefit in terms of computational complexity?
(2) How to take advantage of the recent progress of model-checkers and SAT solvers for the synthesis of mechanisms?

4.2.1 Identifying relevant fragments. In [21], we show that the synthesis of a deterministic mechanism is k-EXPTIME. Recent results show that checking the existence of a Nash Equilibrium in a probabilistic context is 2- EXPTIME [16]. This high cost is usually caused by the assumption that the full expressiveness of the language is considered. However, the high expressiveness of languages for strategic reasoning may not always be needed for simple classes of mechanisms. Therefore, a potentially fruitful direction for future work is to study the complexity of synthesis for those specific fragments [6]. The main challenge consists of identifying what fragments are relevant for representing the core concepts of typical SG properties and then identify the associated minimal fragments.

4.2.2 Synthesis of Mechanism. The overall goal is to go further by using model checkers to verify whether some subsets of properties related to social good (fairness, no discrimination...) and synthesizing SG mechanisms. An open question is the scalability: once the limits are identified, how can we speed up the synthesis stage? Do we still go for a variant of existing tools or do we fully change the perspective: going from SAT to Linear Temporal Logic (LTL) model-checking and from LTL to ATL. As we can see, the questions go far beyond the SG dimension but are a condition of success.

5 CONCLUSION

As we sketched out in this position paper, formal verification is a promising answer to the question of building resource allocation mechanisms for going further on Social Good. Clearly, such a question is of utmost importance: Social Good boosts the need for mechanisms and (partially) automating the design stage is of primary interest. The main output of this research question will be at first *the general specification of properties usually associated with mechanisms impacting social good.* That will help to design and test different variants of mechanisms. The second key output is the *foundation of a new pipeline for automating the design of mechanisms*; this foundation relies on the characterization of different specification languages and tools allowing a white-box perspective on the synthesis of mechanisms. It will directly impact trustworthiness and explainability which are key success conditions.

Acknowledgments. This research has been supported by the ANR project AGAPE ANR-18-CE23-0013, and the EU Horizon 2020 Marie Skłodowska-Curie project with grant agreement No 101105549.

REFERENCES

[1] Thomas Ågotnes and Dirk Walther. 2009. A logic of strategic ability under bounded memory. *Journal of Logic, Language and Information* 18, 1 (2009), 55–77.
[2] Natasha Alechina, Brian Logan, Nguyen Hoang Nga, and Abdur Rakib. 2010. Resource-bounded alternating-time temporal logic. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1.* Citeseer, 481–488.
[3] Maxwell Allman, Itai Ashlagi, Irene Lo, Juliette Love, Katherine Mentzer, Lulabel Ruiz-Setz, and Henry O'Connell. 2022. Designing School Choice for Diversity in the San Francisco Unified School District. In *Proceedings of the 23rd ACM Conference on Economics and Computation* (Boulder, CO, USA) *(EC '22)*. Association for Computing Machinery, New York, NY, USA, 290–291. https://doi.org/10.1145/3490486.3538271

[4] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (sep 2002), 672–713. https://doi.org/10.1145/585265.585270

[5] Benjamin Aminof, Marta Kwiatkowska, Bastien Maubert, Aniello Murano, and Sasha Rubin. 2019. Probabilistic Strategy Logic. In *IJCAI*. ijcai.org, 32–38.

[6] Francesco Belardinelli, Wojciech Jamroga, Vadim Malvone, and Aniello Murano. 2019. Strategy logic with simple goals: Tractable reasoning about strategies. In *28th International Joint Conference on Artificial Intelligence (IJCAI 2019)*. 88–94.

[7] Francesco Belardinelli, Alessio Lomuscio, Vadim Malvone, and Emily Yu. 2022. Approximating perfect recall when model checking strategic abilities: theory and applications. *Journal of Artificial Intelligence Research* 73 (2022), 897–932.

[8] Patricia Bouyer, Orna Kupferman, Nicolas Markey, Bastien Maubert, Aniello Murano, and Giuseppe Perelli. 2023. Reasoning about Quality and Fuzziness of Strategic Behaviors. *ACM Trans. Comput. Logic* 24, 3, Article 21 (apr 2023), 38 pages. https://doi.org/10.1145/3582498

[9] K. Chatterjee, T.A. Henzinger, and N. Piterman. 2010. Strategy Logic. *Information and Computation* 208, 6 (June 2010), 677–693. 2010/CHP10.html

[10] Edmund M. Clarke, Orna Grumberg, Daniel Kroening, Doron A. Peled, and Helmut Veith. 2018. *Model checking, 2nd Edition*. MIT Press. https://mitpress.mit.edu/books/model-checking-second-edition

[11] Vincent Conitzer and Tuomas Sandholm. 2002. Complexity of mechanism design. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence* (Alberta, Canada) *(UAI'02)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 103–110.

[12] Cristina David and Daniel Kroening. 2017. Program synthesis: challenges and opportunities. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375, 2104 (2017), 20150403. https://doi.org/10.1098/rsta.2015.0403 arXiv:https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2015.0403

[13] Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John P. Dickerson, and Vincent Conitzer. 2020. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence* 283 (2020), 103261. https://doi.org/10.1016/j.artint.2020.103261

[14] Daniel Freund, Thodoris Lykouris, Elisabeth Paulson, Bradley Sturt, and Wentao Weng. 2023. Group fairness in dynamic refugee assignment. In *Proceedings of the 24th ACM Conference on Economics and Computation* (London, United Kingdom) *(EC '23)*. Association for Computing Machinery, New York, NY, USA, 701. https://doi.org/10.1145/3580507.3597758

[15] Masahiro Goto, Atsushi Iwasaki, Yujiro Kawasaki, Ryoji Kurata, Yosuke Yasuda, and Makoto Yokoo. 2016. Strategyproof matching with regional minimum and maximum quotas. *Artificial Intelligence* 235 (2016), 40–57. https://doi.org/10.1016/j.artint.2016.02.002

[16] Julian Gutierrez, Lewis Hammond, Anthony W Lin, Muhammad Najib, and Michael Wooldridge. 2021. Rational verification for probabilistic systems. *KR' 21* (2021).

[17] Jason Hartline. 2010. Approximation in mechanism design. In *Proceedings of the Behavioral and Quantitative Game Theory: Conference on Future Directions* (Newport Beach, California) *(BQGT '10)*. Association for Computing Machinery, New York, NY, USA, Article 35, 1 pages. https://doi.org/10.1145/1807406.1807441

[18] Ashwin Kumar and William Yeoh. 2023. Fairness in Scarce Societal Resource Allocation: A Case Study in Homelessness Applications. In *Proceedings of the workshop "Autonomous Agents for Social Good"*. online.

[19] Bastien Maubert, Munyque Mittelmann, Aniello Murano, and Laurent Perrussel. 2021. Strategic Reasoning in Automated Mechanism Design. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning*. 487–496. https://doi.org/10.24963/kr.2021/46

[20] Munyque Mittelmann, Sylvain Bouveret, and Laurent Perrussel. 2022. Representing and reasoning about auctions. *Auton. Agents Multi Agent Syst.* 36, 1 (2022), 20. https://doi.org/10.1007/S10458-022-09547-9

[21] Munyque Mittelmann, Bastien Maubert, Aniello Murano, and Laurent Perrussel. 2022. Automated Synthesis of Mechanisms. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, Lud De Raedt (Ed.). International Joint Conferences on Artificial Intelligence Organization, 426–432. https://doi.org/10.24963/ijcai.2022/61 Main Track.

[22] Munyque Mittelmann, Bastien Maubert, Aniello Murano, and Laurent Perrussel. 2023. Formal Verification of Bayesian Mechanisms. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 10 (Jun. 2023), 11621–11629. https://doi.org/10.1609/aaai.v37i10.26373

[23] Roger B. Myerson. 1989. *Mechanism Design*. Palgrave Macmillan UK, London, 191–206. https://doi.org/10.1007/978-1-349-20215-7_20

[24] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.

[25] Marc Pauly and Mike Wooldridge. 2003. Logic for mechanism design—a manifesto. In *Proceedings of the 2003 Workshop on Game Theory and Decision Theory in Agent Systems (GTDT-2003), Melbourne, Australia*. Citeseer.

[26] Michael Thielscher and Dongmo Zhang. 2010. From General Game Descriptions to a Market Specification Language for General Trading Agents. In *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*, Esther David, Enrico Gerding, David Sarne, and Onn Shehory (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 259–274.

[27] N Tomašev, J Cornebise, F Hutter, S Mohamed, A Picciariello, B Connelly, DCM Belgrave, D Ezer, FCVD Haert, F Mugisha, G Abila, H Arai, H Almiraat, J Proskurnia, K Snyder, M Otake-Matsuura, M Othman, T Glasmachers, WD Wever, YW Teh, ME Khan, RD Winne, T Schaul, and C Clopath. 2020. AI for social good: unlocking the opportunity for positive impact. *Nature Communications* 11, 1 (2020).