A Symbolic Sequential Equilibria Solver for Game Theory Explorer

Demonstration Track

Moritz Graf
University of Freiburg
Freiburg, Germany
grafm@cs.uni-freiburg.de

Thorsten Engesser
IRIT
Toulouse, France
thorsten.engesser@irit.fr

Bernhard Nebel
University of Freiburg
Freiburg, Germany
nebel@uni-freiburg.de

ABSTRACT

We present the first implemented symbolic solver for sequential equilibria in general finite imperfect information games.[1]

KEYWORDS

game theory; extensive-form games; sequential equilibrium

ACM Reference Format:
Moritz Graf, Thorsten Engesser, and Bernhard Nebel. 2024. A Symbolic Sequential Equilibria Solver for Game Theory Explorer: Demonstration Track. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

1 INTRODUCTION

The sequential equilibrium is a standard solution concept for extensive-form games with imperfect information, introduced by Kreps and Wilson [4] in 1982. Sequential equilibria are so-called *assessments*, consisting of a *strategy profile* β and a *system of beliefs* μ. The strategy profile specifies for each action a and information set I the probability $\beta(I)(a)$ with which the acting player plays a at I. The belief system specifies for each history h and information set I the probability $\mu(I)(h)$ that the acting player attributes to being in history h, given that they know they are in information set I. An assessment (β, μ) is a sequential equilibrium if it satisfies the two properties (1) *sequential rationality* and (2) *consistency*.

Intuitively, sequential rationality ensures that strategies are rational given players' beliefs. That is, for any information set I in which player i acts, the *believed utility* of playing β is at least as good as the believed utility of playing any strategy profile β' that differs from β only in the action probabilities for player i. Formally, for all information sets I of player i and strategy profiles $\beta' = (\alpha_i, \beta_{-i})$,

$$U_i^B(\beta', \mu|I) \leq U_i^B(\beta, \mu|I).$$

Conversely, consistency ensures that beliefs are sound given the players' strategies. It requires the existence of a series of fully mixed assessments (β^n, μ^n) where $\lim_{n\to\infty}(\beta^n, \mu^n) = (\beta, \mu)$, and

$$P_{\beta^n}(h) > 0, \quad \text{and} \quad \mu^n(I)(h) = \frac{P_{\beta^n}(h)}{P_{\beta^n}(I)}$$

for all information sets I and histories $h \in I$. Here $P_{\beta^n}(\cdot)$ is the probability of history h or information set I being reached, assuming all players act according to the strategy profile β^n.

The complicated definition of consistent assessments as the limit of an infinite series is necessary because some information sets may not be reached, in which case $P_\beta(I) = 0$. It is also one of the main reasons why finding sequential equilibria is a challenging problem, both conceptually and in terms of computational complexity.

In our technical paper [1] we describe an algorithm for symbolically solving sequential equilibria. In this demonstration paper, we present our implementation, which we have integrated into the well-established open-source software *Game Theory Explorer* [6].

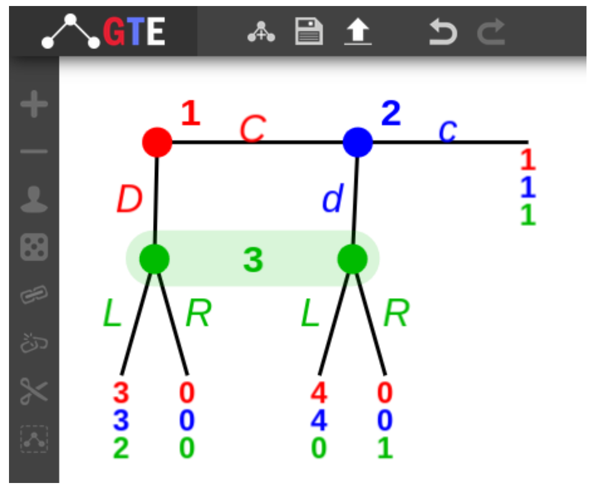

Figure 1: Game Theory Explorer input for Selten's horse.

2 EXAMPLE: SELTEN'S HORSE

Figure 1 shows *Selten's horse*, a well-known example of imperfect information games, first introduced by Selten [7] in 1975.

Note that there is only one information set that contains more than one history. If this information set can be reached (i.e., if player 1 plays D with non-zero probability, or if player 2 plays d with non-zero probability), consistency requires that the beliefs be the conditional probabilities of each of the histories being reached. If the information set cannot be reached, the beliefs can be arbitrary as long as the sequential rationality property is still satisfied.

Some more complex examples where arbitrarily chosen beliefs in one information set affect the choice of beliefs in other parts of the game tree are discussed in our technical paper [1].

[1] See https://github.com/tengesser/GTE-sequential for the implementation and our demonstration video.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

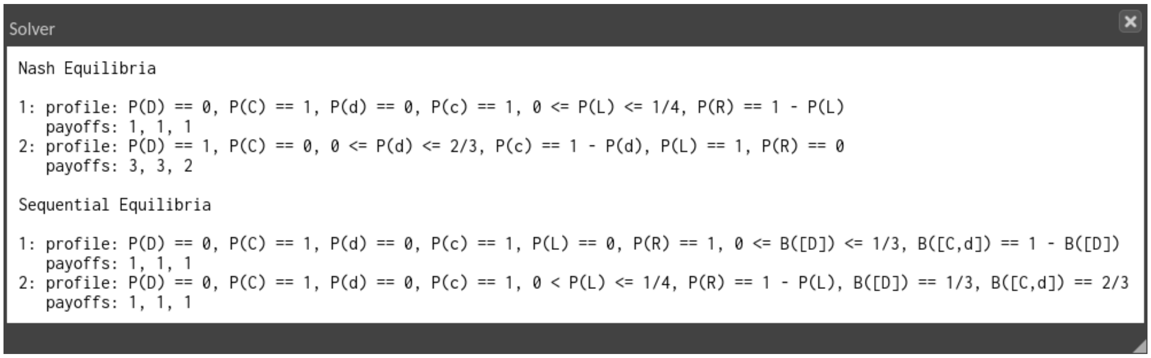

Figure 2: Output of the solver. $P(\cdot)$ are the action probabilities and $B(\cdot)$ are the beliefs.

3 SYSTEM OF EQUATIONS AND INEQUALITIES

Our algorithm works by generating (and then solving) a system of polynomial equations and inequalities characterizing the set of all sequential equilibria. The variables are the probabilities $\beta(I)(a)$ for each action a to be played at its information set I, and the beliefs $\mu(I)(h)$ that the players assign to each history h at I. When the information set is clear, we often write $\beta(a)$ and $\mu(h)$. We obtain the following system of polynomial equations and inequalities:

$$\beta(I)(a) \geq 0, \qquad (1a) \qquad\qquad \mu(I)(h) \geq 0, \qquad (2a)$$

$$\sum_{a \in A(I)} \beta(I)(a) = 1, \qquad (1b) \qquad\qquad \sum_{h \in I} \mu(I)(h) = 1, \qquad (2b)$$

$$\left(\sum_{h \in I} \mu(I)(h) U_i^E(\beta | \langle h, a \rangle)\right) - U_i^B(\beta, \mu | I) \quad \leq 0, \qquad (3)$$

$$\beta(I)(a) \cdot \left(\left(\sum_{h \in I} \mu(I)(h) U_i^E(\beta | \langle h, a \rangle)\right) - U_i^B(\beta, \mu | I)\right) = 0, \qquad (4)$$

$$\prod_{p_i > 0} \alpha_i^{p_i} \prod_{p_i < 0} \gamma_i^{-p_i} = \prod_{p_i > 0} \gamma_i^{p_i} \prod_{p_i < 0} \alpha_i^{-p_i}. \qquad (5)$$

The equations of type (1a-2b) ensure that strategies $\beta(I)$ and beliefs $\mu(I)$ are probability distributions. The equations of type (3-4) are based on the one-shot deviation principle for imperfect information games [2] and ensure sequential rationality. They are quantified over all information sets I and actions a applicable in I. Note that the expected and believed utilities $U_i^E(\cdot)$ and $U_i^B(\cdot)$ are polynomial functions of β and μ. The equations of type (5) provide consistency. Obtaining the exponent and coefficient vectors p, α, and γ is not straightforward and requires us to compute the extreme directions of a set of polyhedral cones. The idea was proposed by Kohlberg and Reny [3] and elaborated in our technical paper [1] in sufficient detail to be easily implemented.

For Selten's horse, the only such equation is $\mu(\langle D \rangle)\beta(C)\beta(d) = \mu(\langle C, d \rangle)\beta(D)$. The equation corresponds to our intuition that, assuming the information set for player 3 can be reached, the belief assigned to each history must be the conditional probability of the history being reached. If both $\beta(d) = 0$ and $\beta(D) = 0$, then both sides of the equation are zero and the belief can be arbitrary.

Our implementation allows restricting solutions to pure strategies or beliefs, which often improves the time to solve the system. This is done by adding additional equations to the system.

4 SOLVING THE POLYNOMIAL SYSTEM

Our implementation uses the cylindrical algebraic decomposition algorithm implemented in Mathematica [8]. It partitions the solution space of the system into connected components, which are represented in a stratified way (i.e., without cyclic dependencies).

Consider the output of our solver for Selten's horse in Figure 2. The instance was solved in about three seconds on an Intel i7-1195G7 CPU. Both sequential and Nash equilibria (of which the sequential equilibrium is a refinement) are shown. There are two connected components of Nash equilibria, with only the first component having corresponding sequential equilibria.

The set of sequential equilibria is further divided into two connected components. The sequential equilibria of Selten's horse are also discussed by Osborne and Rubinstein [5]. Interestingly, they only consider the case where $\mu(\langle D \rangle) = \frac{1}{3}$, for which both components can be combined into one. They seem to be either unaware of the remaining equilibria, or at least do not discuss them. This illustrates the difficulty of manually finding all sequential equilibria, which requires proving that no other equilibria exist.

5 CONCLUSION

We have presented the first implementation of a symbolic sequential equilibrium solver for general finite imperfect information games.

A drawback of our approach is its high algorithmic complexity: Cylindrical algebraic decomposition is double exponential in the number of variables in the worst case. Nevertheless, we argue that our solver is a significant improvement over the state of the art, which is to find sequential equilibria by hand. Besides analyzing small games, we see a practical use for teaching the concept of sequential equilibria in game theory courses.

ACKNOWLEDGMENTS

This work was supported by the ANR LabEx CIMI (grant ANR-11-LABX-0040) within the French State Programme "Investissements d'Avenir", and the EU ICT-48 2020 project TAILOR (No. 952215). We thank Robert Mattmüller, Umberto Grandi, Bernhard von Stengel, and Martin Antonov for their helpful feedback.

REFERENCES

[1] Moritz Graf, Thorsten Engesser, and Bernhard Nebel. 2024. Symbolic Computation of Sequential Equilibria. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6-10, 2024*.

[2] Ebbe Hendon, Hans J. Jacobsen, and Birgitte Sloth. 1996. The One-Shot-Deviation Principle for Sequential Rationality. *Games and Economic Behavior* 12, 2 (February 1996), 274–282.

[3] Elon Kohlberg and Philip J. Reny. 1997. Independence on Relative Probability Spaces and Consistent Assessments in Game Trees. *Journal of Economic Theory* 75, 2 (August 1997), 280–313.

[4] David M. Kreps and Robert Wilson. 1982. Sequential Equilibria. *Econometrica* 50, 4 (July 1982), 863–894.

[5] Martin J. Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory*. The MIT Press.

[6] Rahul Savani and Bernhard von Stengel. 2015. Game Theory Explorer: Software for the Applied Game Theorist. *Computational Management Science* 12, 1 (2015), 5–33.

[7] Reinhard Selten. 1975. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory* 4 (1975), 25–55.

[8] Adam Strzebonski. 2023. Cylindrical Algebraic Decomposition. https://mathworld.wolfram.com/CylindricalAlgebraicDecomposition.html. Accessed on 2023-10-09.