The Lemonade Stand Game Competition: Solving Unsolvable Games

MARTIN A. ZINKEVICH
Yahoo! Research
and
MICHAEL BOWLING
University of Alberta
and
MICHAEL WUNDER
Rutgers University

In December 2009 and November 2010, the first and second Lemonade Stand game competitions were held. In each competition, 9 teams competed, from University of Southampton, University College London, Yahoo!, Rutgers, Carnegie Mellon, Brown, Princeton, et cetera. The competition, in the spirit of Axelrod's iterated prisoner's dilemma competition, which addressed whether or not you should cooperate, asks the questions, "how should you cooperate, and with whom?" The third competition (whose results will be announced at IJCAI 2011) is open for submissions until July 1st, 2011.

Categories and Subject Descriptors: B.6.1 [**Logic Design**]: Design Styles—*Logic Arrays*

General Terms: Economics, Experimentation

Additional Key Words and Phrases: Templates, Skeletons, Things

1. INTRODUCTION

The Lemonade Stand Game was introduced on the Yahoo! Group lemonadegame:

> It is summer on Lemonade Island, and you need to make some cash. You decide to set up a lemonade stand on the beach (which goes all around the island), as do two others. There are twelve places to set up around the island like the numbers on a clock. Your price is fixed, and all people go to the nearest lemonade stand. The game is repeated. Every night, everyone moves under cover of darkness (simultaneously). There is no cost to move. After 100 days of summer, the game is over. The utility of the repeated game is the sum of the utilities of the single-shot games.

If all the lemonade stands are at different spots, then your utility is the distance to the person clockwise you plus the distance to the person counterclockwise you, measured in spots. For example, if Alice sets up at the 3 o'clock location, Bob sets up at 10 o'clock, and Candy sets up at 6 o'clock, then first we arrange them clockwise from 1 o'clock (Alice, Candy, then Bob): there are 3 spots clockwise between Alice and Candy, 4 spots clockwise between Candy and Bob, and 5 spots

Authors' addresses: `maz@yahoo-inc.com`, `bowling@cs.ualberta.ca`,`mwunder@cs.rutgers.edu`
For more on the competition, see `http://martin.zinkevich.org/lemonade/`.

clockwise between Bob and Alice. Therefore, Alice gets \$8, Bob gets \$9, and Candy gets \$7. If all the lemonade stands are located at the same spot, everybody gets \$8. If exactly two lemonade stands are located at the same spot, the two collocated stands get \$6 each and the loner gets \$12. So, the total utility is always \$24.

Given this call, nine teams competed each year, from University of Southampton, University College London [Sykulski et al. 2010; de Cote et al. 2010], Yahoo!, Rutgers [Wunder et al. 2010], Carnegie Mellon [Reitter et al. 2010], Brown, Princeton, et cetera.

2. OBJECTIVES OF RUNNING THE TOURNAMENT

In competitions in two-player zero-sum games, a conventional approach is to attempt to approximate the equilibrium, via methods such as minimax search or abstraction and equilibrium computation. However, this method assumes that the game is solvable (that equilibrium strategies are interchangeable [Nash 1951]) or at least that combining strategies from different equilibria yields reasonable approximations of equilibria. In the lemonade stand game, combining equilibria can yield highly suboptimal (even worst-case) strategy profiles, and therefore, this competition forces players to focus not on computing equilibria, but on selecting equilibria, and "convincing" others to play their equilibria.

The competition was modeled in part on Axelrod's famous iterated prisoner's dilemma competition [Axelrod 1980; 1984]. In the prisoner's dilemma, how to cooperate is clear: there is an action labeled "cooperate". Also, who to cooperate with is clear: one plays with one opponent at a time. However, in the lemonade stand game (which is a type of location game), there are many ways to cooperate. The simplest and most used was to play on opposite sides of the circle. However, there are 12 such configurations for 2 players, and 36 such configurations overall. So which cooperation specifically and with whom is critical.

Before running the competition, I ran preliminary experiments where a constant strategy (which played a single action the whole game) won a tournament against a variety of sophisticated AI programs. Thus, I knew that traditional methods would not fare well. This information was shared with the competitors.

3. STABLE AND UNSTABLE COOPERATION

Before the first competition, we wanted to see if players could collaborate: we had two types of collaboration in mind. For the sake of illustration, assume Alice and Bob are collaborating against Candy.

(1) **stable:** Could Alice and Bob collaborate by playing opposite each other, forcing Candy to get 6 (the safe value)?

(2) **unstable:** Could Alice and Bob collaborate such that Candy got less than 6 utility (the safe value)? E.g., a "sandwich": Candy is at 2 o'clock and not moving; Alice moves to 1 o'clock and Bob moves to 3 o'clock.

One can think of these as two variants of Stackelberg equilibria for the game where Alice and Bob play as one in a zero-sum game against Candy: in the first, Alice and Bob are the leaders; in the second, Candy is the leader.

In each year, nine teams submitted bots.[1] Let us begin by counting the percentage of rounds where one player got below the safe value (unstable cooperation). In general, for three players playing uniformly at random, there is a $\frac{5}{12} \approx 41.7\%$ chance that someone will get below the safe value on a given round. In the first competition, this occured in 14.8% of the rounds. In the second competition, this occurred in 1.6% of the rounds. Thus, even if such opportunities exist, players are not exploiting them. It is also possible that the opportunity cost (e.g., not pursuing a stable cooperation) outweighs the short-term gains of an unstable cooperation.

The second kind of cooperation is where two of the bots play opposite one another (stable cooperation). By this definition, everyone playing uniformly at random will generate this cooperation approximately 22.9% of the rounds. In the first competition, there was cooperation 67.6% of the rounds. In the second competition, there was cooperation 96.7% of the rounds. Moreover, if you ignore the utilities for all rounds where no cooperation occurred, the ranking remains the same in both competitions. Thus the competitions were decided on two factors: amount of cooperation and the amount of utility received during cooperation.

4. CONCLUSION

Can a bot agree upon an equilibrium in an unsolvable game with another bot designed by someone else? In this game we have empirically demonstrated that this is possible, which raises the question, what about other unsolvable games? Right now the rules of the 2011 competition are being finalized. The deadline for submitting bots will be July 1, 2011. In the new competition, we will specify a distribution over location games on a circle instead of a single game. The results will be presented at the TADA workshop at IJCAI 2011.

The iterated prisoner's dilemma competition of Axelrod was interesting and powerful because it was *simple*. The lemonade stand game was a natural next step. Further steps will involve distributions over more complex games, communication, and lifelong learning; moving toward answering a fundamental question in multi-agent learning, "Machines can work *for* us, but can they work *with* us?"

REFERENCES

Axelrod, R. 1980. More effective choice in the prisoner's dilemma. *The Journal of Conflict Resolution 24,* 3 (Sep.), 379–403.

Axelrod, R. 1984. *The Evolution of Cooperation.* Basic Books, United States of America.

de Cote, E. M., Chapman, A., Sykulski, A., and Jennings, N. 2010. Automated planning in repeated adversarial games. In *26th Conference on Uncertainty in Artificial Intelligence (UAI 2010).* Association for Uncertainty in Artificial Intelligence, Avalon, California.

Nash, J. 1951. Non-cooperative games. *Annals of Mathematics Journal 54*, 286–295.

Reitter, D., Juvina, I., Stocco, A., and Lebiere, C. 2010. Resistance is futile: Winning lemonade market share through metacognitive reasoning in a three-agent cooperative game. In *Proceedings of the 19th Behavior Representation in Modeling and Simulation (BRIMS).*

[1] In the first year, all but one program (Brown) were really fast with the competition finishing in a few minutes (Brown's bot was fast enough for the rules, but took two weeks to play against the other programs). Thus, as many of the results that followed involved re-running the competition, we left Brown 2009 out of the analysis of the results.

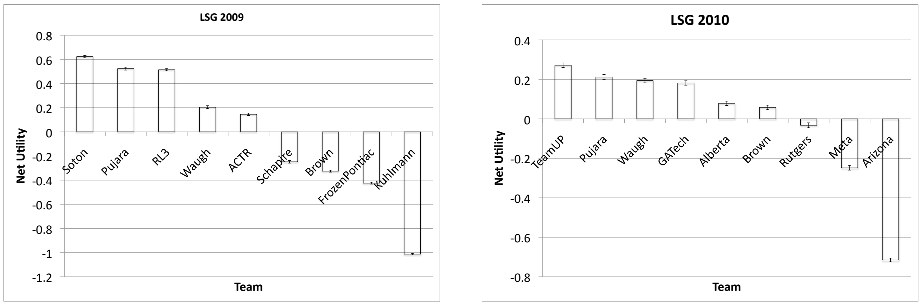

(a) (b)

Fig. 1. The performance of the teams during the 2009 (a) and the 2010 (b) competitions. The net utility is the average utility per round minus 8, so that the average performance is zero.

Sykulski, A., Cote, A. C. E. M. D., and Jennings, N. 2010. EA^2: The winning strategy for the inaugural lemonade stand game tournament. In *19th European Conference on Artificial Intelligence*.

Wunder, M., Kaisers, M., Littman, M., and Yaros, J. 2010. A cognitive hierarchy model applied to the lemonade game. In *AAAI Workshop on Interactive Decision Theory and Game Theory (IDTGT)*.