

COOPERATION IN SOCIAL NETWORKS: COMMUNICATION, TRUST, AND SELFLESSNESS

David C. Arney

Army Research Office
Research Triangle Park, NC 27709

Elisha Peterson

Department of Mathematical Sciences
United States Military Academy
West Point, NY 10996

ABSTRACT

The Army’s myriad social networks connect not only humans and soldiers, but also machines, computers, and robots. And most of the social/biological/physical/informational connections in these networks are in the form of cooperation – entities working together to achieve a common goal. So how does this complex web of relationships, collaborations, and communities of diverse entities work? This paper introduces a new framework in which measures of cooperation can be precisely defined. We demonstrate how the framework can be applied to social networks, and examine the dynamic networks arising in the case of pursuit and evasion games. Finally, we relate the concepts of competitive and altruistic cooperation to trust and the nature of communication within a team.

1. INTRODUCTION

The Army’s myriad social networks connect operational battlespace entities (not only humans and soldiers, but also machines, computers, and robots). And many of the social/biological/physical/informational connections in these networks are in the form of cooperation – entities working together to achieve a common mission or goal. A major component of this net-centric warfare is to “interact and collaborate in the virtual (informational) domain” (Alberts et. al., 1999). At a basic level, such cooperation can be traced to a fundamental synchronization of entities, which is also being studied mathematically (Strogatz, 2003). At the highest levels of application, businesses are embracing net-centric and collaboration concepts with the hope of inventing “organizations for the twenty-first century that will not only be more economically productive but also more humanly desirable” (Malone et al., 2003; Tapscott and Williams, 2006).

How does this complex web of connections, relationships, collaborations, and communities of diverse entities work? What are the most useful and appropriate metrics for a cooperative network and its social and informational value? What is the essence of cooperation and what makes a cooperative system, a collaborative network, or a flat organization effective? What forms of communication and operations within a network enhance cooperation? How do we measure trust and selflessness? Can our network structures, processes, tools, connections, communications, and languages enhance cooperation, achieve synergy, and optimize our networks?

This paper reveals the role of mathematics as a valuable tool to study these questions about networks and cooperative systems and provides the first steps in a theoretical formulation of the fundamental principles, relationships, and metrics of these phenomena. The unifying concept throughout the paper is a *subset team game*, a framework that assumes the existence of a function describing how each subset of players within a team value the possible outcomes of a simulation. This concept is a slight generalization of ideas in classical cooperative game theory, and can be used to measure an *altruistic contribution* and a *competitive contribution* for each subset of players.

We demonstrate several examples in which this framework can be applied to draw conclusions regarding the nature of cooperation within a scenario, focusing primarily on network science and pursuit and evasion games. In particular, we discuss how one can decide which behaviors or algorithms are the most altruistic, how to tell whether a teammate is trustworthy, and how one can determine whether a “cooperative system” (or network or organization) is truly cooperative in a mathematical sense.

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE DEC 2008	2. REPORT TYPE N/A	3. DATES COVERED -	
4. TITLE AND SUBTITLE Cooperation In Social Networks: Communication, Trust, And Selflessness		5a. CONTRACT NUMBER	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Army Research Office Research Triangle Park, NC 27709		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited			
13. SUPPLEMENTARY NOTES See also ADM002187. Proceedings of the Army Science Conference (26th) Held in Orlando, Florida on 1-4 December 2008, The original document contains color images.			
14. ABSTRACT			
15. SUBJECT TERMS			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	UU
			18. NUMBER OF PAGES 8
			19a. NAME OF RESPONSIBLE PERSON

2. A FRAMEWORK FOR COOPERATION

Currently, cooperation is not well-understood from a mathematical point-of-view. The most classical mathematical construct in this area is found in John von Neumann’s “cooperative game theory”, introduced in the early 20th century (Neumann, 1928). However, this concept applies primarily to a group of entities working together for *selfish* reasons. In contrast, most forms of cooperation in social and military networks involve a team working together for a common good. In reality, von Neumann’s theory only applies to a kind of cooperation that is mostly economic in nature.

In this section, we describe a framework that unifies von Neumann’s form of cooperation, which we call *selfish* or *competitive cooperation*, with a more team-oriented or *altruistic cooperation*. This framework uses the notion of a *subset team game*, meaning a situation or scenario in which the value of a given outcome, as perceived by a certain team, can be measured. Moreover, each subset of players on the team may have a different perception of the value of an outcome.

Within this broadly applicable framework, there exist clearly defined metrics for both selfish cooperation and altruistic cooperation. In later sections, we will exhibit several sample applications of this cooperation framework, and discuss how it aids an understanding of communications, trust, and altruism within social networks.

2.1. Brief History of Mathematical Cooperation

Our framework is an extension of von Neumann’s *cooperative game theory*, which analyzes what we call competitive cooperation. The *oligopoly* is one real-life example. If a market is dominated by a small number of firms, they may, without directly communicating, take actions which improve their own earnings. They may also directly cooperate, colluding to fix high prices for example. In such a situation, no firm actually acts altruistically, but rather cooperates solely to advance its own agenda. Cooperative game theory was developed to understand and analyze the economics of such situations, providing a means to compensate players based on their *marginal* contributions to a larger group.

A *cooperative game with transferable utility* consists of a situation involving (i) a set of players T called the *coalition*, and (ii) a *payoff function* (or *utility function*) $v: 2^T \rightarrow \mathcal{R}$ associating a particular *value* or *utility* to each subset of the coalition (Burger, 1963; Osborne, 2003). Note that 2^T is the collection of all subsets of T .

In these games, the *marginal contribution of a player A to the coalition T* is defined to be

$$m_A(S) = v(S) - v(S \setminus A). \quad (1)$$

This is interpreted as the additional value of the outcome, through the utility function, when the players in A participate. The term “transferable” indicates that all players enjoy the same payoff. In this classical theory, the marginal contribution can be used to determine appropriate compensations to each player in such a coalition.

A related notion is the *cooperative game with non-transferable utility*, which involves (i) a coalition of players T , (ii) a set X of possible outcomes, (iii) a *consequence function* $V: 2^T \rightarrow X$ mapping each subset $S \subseteq T$ to an outcome, and (iv) a payoff function $u_A: X \rightarrow \mathcal{R}$ defined for each player $A \in T$ associating a value to each possible outcome (Burger, 1963; Osborne, 2003).

2.2. The Framework: Subset Team Games

Unfortunately, this previous theory has limits and cannot explain many forms of cooperation. The growth of the user-written encyclopedia *Wikipedia* provides one example of unexplained cooperation. Thousands of contributors made this the largest encyclopedia in the world within just a few years, despite little recognition for their work and no monetary gain. In addition, human subjects often cooperate in situations where it is not *rational* to do so, and there is little reason to work with another player. In the game of *Prisoner’s Dilemma*, two players are pitted against each other in a situation where the rational choice, in the sense of Nash equilibrium, is to turn on the other player. However, when the game is simulated in real life most human beings do not follow the rules of economics. Rather, they frequently choose to cooperate (Axelrod, 1984, 1999) and often people are legitimately concerned with both the organization’s effectiveness and other people’s payoffs. (Taylor, 1987)

Additionally, many social networks are driven more by altruistic cooperation than by competitive cooperation. Within sports teams, the best players often subject their own interests to that of the team. The same definitely can be said of Army organizations. In short, cooperative game theory cannot adequately explain many instances of cooperation.

We can, however, extend the existing theory to capture both altruistic and competitive cooperation. The idea is to generalize the notion of cooperative games so

that the value of an outcome can be assessed not just for teams or individual players, but also for any subset of a team. The precise definition follows.

Definition. A *subset team game* consists of a situation involving (i) a coalition of players T , (ii) a set X of possible outcomes, (iii) a *consequence function* $V : 2^T \rightarrow X$ mapping each subset $S \subseteq T$ to an outcome, and (iv) a *payoff* or *utility function* $u_S : X \rightarrow \mathbb{R}$ for each subset $S \subseteq T$ taking an outcome to a real number representing the value.

Note that, for a given subset $S \subseteq T$, the term $V(S)$ is interpreted as the outcome produced when only the players in S participate in the game.

Given subsets $A \subseteq B \subseteq T$, define the *subset payoff function* $u_A(B) = u_A(V(B))$. This represents the payoff due to the players in A , which we call the *assessing subset*, in the outcome produced when only the players in B , called the *coalition* participate. Fig. 1. Illustrates the roles of sets A , B , and B^c in this scenario.

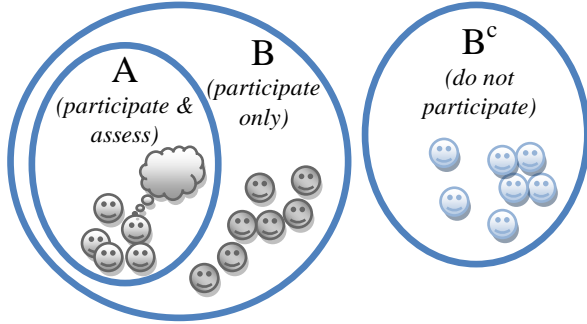


Figure 1. The subset payoff function $u_A(B)$ represents the value perceived by the players in A when the players in B participate.

We would generally like to limit our discussion to situations where more players leads to a more “successful” outcome. In other words, (i) adding more players to the game should never reduce the value of the outcome, and (ii) the value of an outcome perceived by a group of players should not be smaller than the value as perceived by a subset of that group. This idea is illustrated in Fig. 2, and motivates the following definitions:

Definition. A subset team game is *fully-cooperative* if $u_A(C) \leq u_A(D)$ whenever $A \subseteq C \subseteq D \subseteq T$. A subset team game is *cohesive* if $u_A(C) \leq u_B(C)$ whenever $A \subseteq B \subseteq C \subseteq T$.

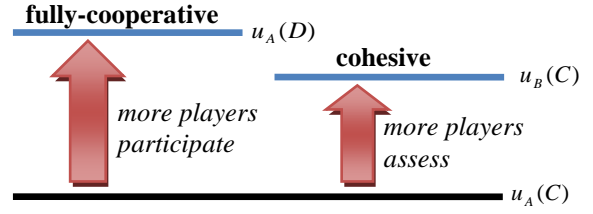


Figure 2. The fully-cooperative and cohesive conditions describe situations in which the value of an outcome increases when either more players participate or more players assess (respectively).

2.3. Metrics of Cooperation

We can define metrics of competitive and altruistic cooperation within this framework, provided a game is both cohesive and fully-cooperative. The metrics demonstrate how the classical marginal contribution can be subdivided into a “competitive contribution” and an “altruistic contribution”.

Definition. Given a payoff function $u_A(B)$ in a subset team game, the *total marginal contribution* of a subset A to the team T is

$$m(A) = u_T(T) - u_{A^c}(A^c). \quad (2)$$

If the game is both cohesive and fully-cooperative, then the *competitive contribution* of A is

$$c(A) = u_T(T) - u_{A^c}(T) \quad (3)$$

and the *altruistic contribution* of A is

$$a(A) = u_{A^c}(T) - u_{A^c}(A^c). \quad (4)$$

These metrics are the centerpiece of the framework, since they permit an analysis of the types of contribution made by any subset of players. Note that the altruistic term disregards the perception of value by any player in A . Note also that the total marginal contribution decomposes as

$$m(A) = c(A) + a(A). \quad (5)$$

Fig. 3 illustrates this breakdown.

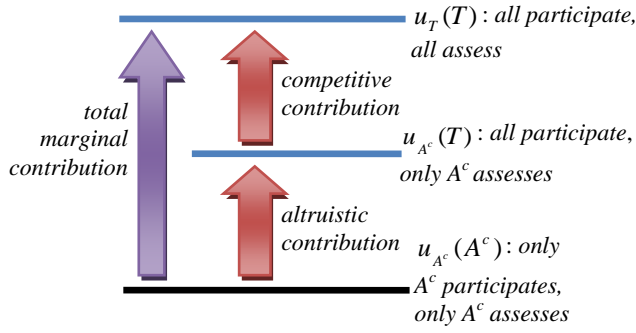


Figure 3. The altruistic, competitive, and total marginal contributions of a subset.

3. USING THE FRAMEWORK

In the onslaught of equations in the previous section, it is possible to lose sight of the general applicability and adaptability of this new framework. In this section, we demonstrate how the framework can be used in practice.

3.1. Setting up and Utilizing the Framework

The steps required to set up and analyze a scenario or simulation within this framework are:

1. Assume we have a game with a particular set of players T .
2. Determine how “value” should be interpreted within the scenario. Choose something that makes sense for individuals as well as subsets of players.
3. Construct the corresponding *payoff function* $u_A(B)$, which describes how each subset of players A assesses the value of an outcome in which only the players in B participate.
4. Given a particular player or subset of players A to be analyzed, use the payoff function to compute $u_{A^c}(A^c)$, $u_{A^c}(T)$, and $u_T(T)$.
5. Use these values to compute the competitive contribution (eq. 3) and altruistic contribution (eq. 4) of a particular player or subset of interest.
6. Adjust the behavior of players based on these metrics.

One must be careful in choosing the payoff function. The assignment of value as perceived by a subset of a team can be a somewhat arbitrary process. If not all outcomes are known, it may be a considerable challenge or may not even be possible. Moreover, the cohesive and fully-cooperative conditions are often not met in practice. As is always the case, the mathematics is not a perfect model for reality, yet nevertheless offers substantial

insights. The true power of this framework lies in the ability to conduct analysis relating the competitive and altruistic measures back to the issues of trust and communication.

3.2. Example: Wikipedia

Our first example is a qualitative analysis of a user’s contribution to *Wikipedia*. The value of this website to mankind is immense, yet it depends fundamentally upon the selfless actions of tens of thousands of contributors.

First, let the “value” represent the average value of the information on *Wikipedia* to an assessing subset. Consider a single individual A that contributes to the website. Then $u_{A^c}(A^c)$ represents the value of the encyclopedia *without* the contributions of A . The term $u_{A^c}(T)$ represents the value (to everyone except for A) of the encyclopedia *with* the contributions of A . Finally, the term $u_T(T)$ represents the value of the encyclopedia *with* the contributions of A , as perceived by everyone.

While all three terms should have high values, note that the competitive contribution $u_T(T) - u_{A^c}(T)$ is highest when the contributions of A contain information valuable to A but few others, while the altruistic contribution $u_{A^c}(T) - u_{A^c}(A^c)$ is highest when the contribution contains information valuable to nearly everyone. Thus, our common-sense interpretation of selflessness matches the mathematical notion of altruistic cooperation.

3.3. Example: Basketball

Our second example involves a basketball team. We define the “value” of a completed game to a subset of players to be the total points per game (ppg) scored by those players. Suppose that a team involves seven players whose scoring is as shown in Table 1.

Table 1. Hypothetical basketball statistics.

Player	Player’s PPG	Points scored when the player is absent
A	25	84
B	20	90
C	20	80
D	15	70
E	10	70
F	10	99
G	10	90
<i>Total</i>	110	

Note that the average score of the team is 110, which is the sum of each player’s PPG. Player A’s selfish contribution is simply his/her PPG, or 25. Player A’s altruistic contribution is $(110 - 25) - 84 = 1$, which is the difference between the number of points scored by the other players when A does and does not participate. Thus, player A could be said to be a very “selfish” player.

Player E, on the other hand, makes a selfish contribution of 10 and an altruistic contribution of 30. So this player is a highly altruistic player. If all the players are compared, players D and E both make the highest total contribution to the team’s success, since without them the team scored on average 40 points less. Availability of new sports statistics (e.g., plus/minus statistics in hockey and basketball) are making the thorough analysis of player’s contributions and their cooperative nature more transparent to coaches and fans.

3.4. Limitations

In practice, it can be difficult to find situations in which the cohesive and fully-cooperative conditions are met. However, even in such situations the metrics as defined above provide useful information about behaviors and algorithms. The pursuit and evasion games discussed later provide a more interesting and concrete example of this.

Another difficulty arises in assessing the value of multiple outcomes. Frequently, it is difficult to determine the value of all possible outcomes of a scenario or simulation. A lot of data is required to assess the various metrics of cooperation within our framework. For example, the basketball scenario above requires data for games that a particular player does not participate. In many cases, there are not enough of these games to make the results statistically significant. Thus, while the metrics from these data may indicate that players may be somewhat “selfish” or “altruistic”, the results may be far from conclusive.

4. COOPERATION IN SOCIAL NETWORKS

One can argue that the essence of social networks is cooperation, as even the links between agents in such networks usually indicate cooperation of one form or another. Networks of all kinds are especially well-suited to our framework. Indeed, in this context the mathematical language of graph theory becomes an abundant source of payoff functions for such networks.

Note that the vertices of a social network are themselves a subset of players, offering the first step in constructing the subset team game framework. Moreover, many graph invariants satisfy the conditions required to construct a value function that is both cohesive and fully-cooperative. The remainder of this situation demonstrates how even the simplest graph invariants can be useful within this framework.

Let $G = (V, E)$ be a graph with vertex set V and edge set E . Since the vertices represent the “players” in this case, the framework requires an “outcome” when only a subset $B \subseteq V$ participates. We let the outcome be the subgraph G_B induced by the vertices in B . Then, any function defined for all subsets of a graph will induce a “payoff function” of the form $f_A(B)$.

For example, suppose that $f_A(B)$ represents the maximum size of a connected component of G_B containing a vertex in A . This very simple and straightforward “connecting” function satisfies both the cohesive and semi-cooperative conditions, permitting us to measure both the “altruistic” and “competitive” contributions of each collection of vertices.

Fig. 4 demonstrates this calculation for a set of four players $A = \{1,2,3,4\}$. The complement of the set is a single vertex $\{5\}$. The size of the component containing vertex 5 is 1 in Part (a) and 2 in Part (b) of Fig. 4, so the *selfless contribution* of the set is +1. The size of the largest component in the entire graph is 3, so the *competitive contribution* of the set is also +1.

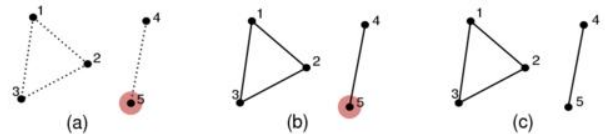


Figure 4. Schematic for finding the marginal contributions of vertices $\{1,2,3,4\}$ to the network.

In the example shown in Fig. 5, the contribution of node 2 is entirely “selfless” (+3), since it increases the maximum component size of the complement by 3.

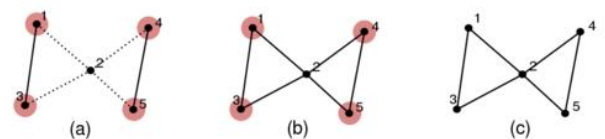


Figure 5. Schematic for finding the marginal contributions of vertex 2 to the overall graph.

To see what the numbers above mean, suppose the graphs indicate a communications network within an organization. In the second situation (Fig. 5), agent 2 makes a strong altruistic contribution to the organization, since without him/her the groups $\{1,3\}$ and $\{4,5\}$ cannot communicate. Mathematically, this contribution is entirely altruistic. The contribution of $A = \{1,2,3,4\}$ in the first case (Fig. 4) is a mixture of altruistic and competitive cooperation since the remaining vertex $\{5\}$ is not as well-connected as the other vertices. Even this simple “connecting” function demonstrates the nature of cooperation and its utility.

5. COOPERATION IN PURSUIT AND EVASION GAMES

In this section, we describe how the framework applies to pursuit and evasion games. These seemingly simple games provide a “toy model” simulation for demonstrating that our metrics of altruistic and competitive cooperation can be an effective way to evaluate player behaviors.

A *pursuit and evasion game* is a game involving a team of *pursuers* and a team of *evaders*. The objective of the pursuers is to “capture” the evaders, while the objective of the evaders is to avoid capture (Isaacs, 1965). These simple games are played out endlessly on battlefields, on playgrounds in the form of *tag* and *capture the flag*, on the floor of multimillion-dollar stadiums in the form of football, on the silver screen in the form of car chases, and in the natural world between predator and prey.

There are countless variations of this game. One can adjust the number of players and their properties, the environment in which the game is played out, even the number of teams involved. The goals of the players also vary from game to game. It may be that the pursuers need to capture all of the opposing team, or they may just need to capture a single player. The evaders may be trying to reach a particular location or just to run away and hide. Mathematical solutions exist for just a few of these variations. In any situation with these kinds of complexities, no “exact” solution is known (Isaacs, 1965; Nahin, 2007).

5.1. A Java Platform for Analysis of Pursuit and Evasion Games

In order to analyze these games, we developed a Java platform capable of simulating a wide variety of pursuit and evasion games. The platform enables the user to

adjust the number of players and goals on each team, as well as the starting location, speed, behavior, and other characteristics of each individual agent. The number of teams can also be adjusted if desired.

The platform is highly visual and dynamic, allowing the user to alter starting positions or other parameters and immediately see the impact on the computed paths; the computed paths move continuously as the user drags a starting position along the screen. Fig. 6 shows a screenshot of one simulation.

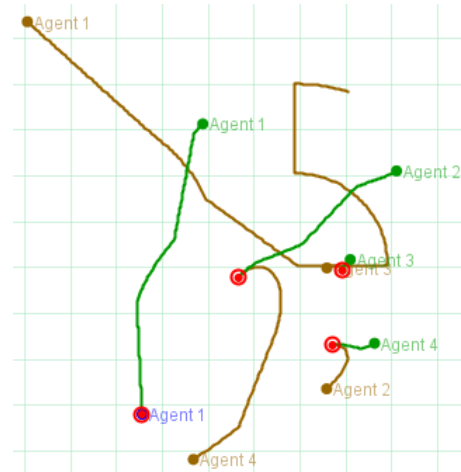


Figure 6. Paths of players in a two-team pursuit and evasion game.

We also implemented a feature to compute arbitrary metrics of success, and the corresponding metrics of cooperation. This feature was also dynamic; a user can adjust the visibility distance of several players and immediately see the impact on the level of altruistic cooperation. Finally, the platform also collects statistics for large numbers of simulations.

The highly visual and dynamic nature of the platform, coupled with the ability to create highly complex scenarios, enabled us to explore a wide variety of pursuit-evasion games and quickly gain intuition regarding how certain parameters impacted metrics of success and cooperation.

5.2. Initial Results

Using the Java platform, we applied the framework to a simple pursuit and evasion game involving players with spatially limited visibility and communications. Each player constantly passed information about the positions of players within their line-of-sight to other teammates within a certain range. Our goal was to assess the cooperative value of various behaviors in terms of altruism and competitiveness. We compared several autonomous behaviors, in which no central entity

controlled behavior, with a baseline fully-controlled system. The first payoff function used was the average distance between players on the two teams.

Our initial observations in this scenario follow:

- Control algorithms were the most successful and, in the sense of our definitions, the most highly cooperative. This is not surprising because the system control dictated the cooperation of the system.
- The most altruistic players were slow-moving, but could see and communicate over long distances.
- The most selfish players were quick, and relied on communications from other players.

To summarize, the numeric data underlying these observations matched our expectations of altruism and competition.

The first step in further validating the subset team game framework in this context is to choose a better payoff function. The average distance is generally not a fully-cooperative function, since a subset of players may have a smaller average distance than the entire set of players. Despite this limitation, the relative values of the altruistic and competitive metrics still offered insight into the nature of the players' behaviors.

5.3. Networks in Pursuit and Evasion Games

Pursuit and evasion games typically involve highly dynamic social networks. The players in the game represent the vertices, and edges may represent either two players that can see each other or two players that are able to communicate with each other. This gives rise to a *communications network* for each team and a *visibility graph* between the two teams. As time passes, the players continually change positions, and so the communications and visibility graphs continually change as well.

A natural question to ask is how the structure of these networks can be exploited to design more effective algorithms for pursuit and evasion. Can we develop more efficient communications and cooperation strategies? How does this dynamic network structure inform our knowledge of the game?

6. COMMUNICATION, TRUST, AND ALTRUISM

With the assumption that a system is completely autonomous, *trust* becomes a very important issue. In real life, opposing teams in a scenario are almost always made up of specialists, that is, players with particular skills (speed or maneuverability) or patterns of behavior

(aggressive or passive). In situations with this variety of behaviors or algorithms, trust can play a central role in determining the proper way for such teams to work together. A high level of trust reduces the need for communication among players, freeing up the communications infrastructure or cognitive load on each player for other things. In this sense, one can view trust as a sort of "implied communication" between players.

Our framework is directly applicable to the issue of trust. We assume that player A *trusts* player B when *player A has a high confidence that player B will behave in a way that benefits the team*. This is precisely the notion of altruistic cooperation developed earlier. Thus, if an appropriate payoff function can be determined, our framework can be used to determine which players are most deserving of trust. Players with high altruistic contributions ought to be given a high level of trust, while those with high competitive contributions may not deserve much trust. Basketball is a good example here. A "ball hog" that never passes the ball and always takes the shot should not be trusted by teammates as much as a skilled team player who always passes the ball to the player most likely to score. A basketball team consisting of players who earn the trust of one another through altruistic cooperation often overachieves and succeeds despite their individual skills and thereby clearly demonstrate the power of this form of cooperation.

A related notion is *predictability*, in which one player has a high confidence in the behavior of another player and can use this information to adjust its own behavior. This concept is outside the scope of the current cooperative framework, but also comes into play since a high level of predictability also reduces the need for communications. Further study and extensions to the framework may help illuminate this aspect of cooperation

7. CONCLUSIONS

The ideas presented in this paper are simply a beginning. In the many situations where the framework can be applied, one can define metrics of altruism and competitive cooperation. These metrics make it possible to alter behaviors or algorithms in order to increase the level of altruism within a team. However, many questions regarding this concept of altruism remain. Is a more altruistic team always better, or are there situations in which some degree of competitive cooperation is beneficial? Does optimizing altruism correspond to optimizing a team's payoff? What additional conclusions can be drawn in the case of social networks?

We can say that the metric of altruism should provide a means of establishing trust within a network or team of players. In highly autonomous systems, this mathematical route to trust may prove to play a fundamental role in determining how agents interrelate. This higher level of trust may lead to a reduced communications load within the system, without sacrificing the systems efficacy.

On a higher level, this may also offer some insight into the total level of trust within a network, by evaluating the altruistic cooperation metric on a global level. One can also evaluate whether a “cooperative system” is truly cooperative, or which form of cooperation that system possesses. A cooperative system that is primarily competitive in nature is very different from one that is primarily altruistic in nature.

The US Army has many networks and organizations that operate primarily on the basis of cooperation. The Army’s culture is based on highly technical, cooperative teams (e.g., teams of teams, systems of systems). And many of the new emerging technologies of net-centric warfare involve using cooperation in the form of hybrid systems – specialized teams of people, machines, computers, and robots (Alberts, et. al., 1999). At the next layer of detail, the major elements of cooperation are found in the mix of trust and autonomy of the agents (i.e., lack of strict control). The American Army is known for its trust and autonomy in its missions and tasks, and this collaborative culture will continue in future doctrine. Therefore, understanding, designing for, and implementing cooperation in our new systems are critical elements in meeting the goals of the future highly-networked Army. The fundamental research presented in this paper on the basic mathematical principles of cooperation, especially the new framework introduced, can contribute greatly to that effort and to the fulfillment of these important Army goals.

ACKNOWLEDGMENTS

This research was completed while the second author was under the support of an NRC/Davies Postdoctoral Teaching/Research Fellowship. The authors are grateful to former USMA cadet Andrew Plucker and current cadet Lucas Gebhart for their insights into pursuit and evasion games. In their senior projects and summer internships at the Army Research Office, they developed productive simulation tools and conducted many of the simulations discussed in section 5.2, and their work helped to formulate the initial observations listed there.

REFERENCES

- Alberts, D. S., J. J. Garstka, and F. P. Stein, 1999: *Network Centric Warfare: Developing and Leveraging Information Superiority*. CCRP, 256 pp.
- Axelrod, R. and M. D. Cohen, 1999: *Harnessing Complexity*. Basic Books, 206 pp.
- _____, 1984: *The Evolution of Cooperation*. Basic Books, 241 pp.
- Burger, E., 1963: *Introduction to the Theory of Games*. Prentice-Hall, 202 pp.
- Isaacs, R., 1965: *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control, and Optimization*. Wiley, 420 pp.
- Malone, T. W., R. Laubacher, and M. S. S. Morton, 2003: *Inventing the Organizations of the 21st Century*. MIT Press, 433 pp.
- Nadeau, E. G. and D. J. Thompson, 1996: *Cooperation Works! How People are Using Cooperative Action to Rebuild Communities and Revitalize the Economy*. Loan Oak Press, 205 pp.
- Nahin, P., 2007: *Chases and Escapes*. Princeton University Press, 270 pp.
- Neumann, J. v., 1928: Zur Theorie der Gesellschaftsspiele. *Math Annalen*, **100**, 295-320.
- Osborne, M. J., 2003: *An Introduction to Game Theory*. Oxford University Press, 560 pp.
- Strogatz, S., 2003: *Sync: The Emerging Science of Spontaneous Order*. Hyperion, 338 pp.
- Tapscott, D., and A. D. Williams, 2006: *Wikinomics: How Mass Collaboration Changes Everything*. Portfolio, 324 pp.
- Taylor, M., 1987: *The Possibility of Cooperation*. Cambridge University Press, 209 pp.