(Preprint) AAS 23-302

SATELLITE COLLISION AVOIDANCE USING REPEATED GAMES

Sydney Dolan*, Victor Qin*, Geoffrey Ding* and Hamsa Balakrishnan†

The democratization of space and rise of mega-constellations have led to more satellites in orbit than ever before, and thus increasing the number of operator-operator conjunction events. An understanding of maneuver strategies for collision avoidance is crucial to maintaining a safe satellite operating environment. We use game theory to analyze the satellite collision problem and propose autonomous strategies for collision avoidance of different types of satellites by posing it as a repeated game. Our results provide insight into recommended operator decision-making windows and likely strategies of non-cooperative entities.

INTRODUCTION

The number of satellites is expected to increase exponentially over the next decade. With this increased congestion, there is an increased risk of collision between satellites. Even one collision can have significant impacts to the space ecosystem. For example, the 2009 collision of the Iridium-33 and Kosmos-2251 satellites resulted in 823 debris objects that formed two debris clouds in Low Earth Orbit (LEO).[1] This incident resulted in the destruction of the two satellites involved and created long-lasting consequences to others in orbit, as other active satellites had to maneuver to avoid debris from the incident.[2]

Producing a set of coherent regulations for space collisions is difficult, in large part due to the heterogeneity of operators in orbit. Lower launch costs have enabled new entrants to orbit, with over 86 countries now owning and operating spacecraft, a twofold increase from 2010.[3,4] Many of these newcomers have less infrastructure to quickly respond to collision avoidance warnings from screening services. Unfortunately, these screening services often have a high false alarm rate, with a detrimental impact on the trust placed on their alerts by satellite operators.[5] Miscommunications between operators about how each will respond has resulted in numerous near-misses between satellites in orbit.[6,7] Another issue with collision avoidance responses is related to the militarization of space. The presence of different military spacecraft complicates management, as government operators may want to conceal their operations and intentions from others.

To illustrate the complexities of operator behavior for a potential collision, consider a hypothetical high-risk satellite conjunction warning between two operators from different countries that was delivered one day in advance. Due to factors like time zones, staffing, and operator competition, the two operators have not communicated their intended collision avoidance maneuver strategy. Without communication, operators can inadvertently move closer to one another, neglect to move on the assumption the other will, or successfully maneuver away from one another. The collision response

*Graduate Student, Aeronautics and Astronautics, MIT
†William E. Leonhard Professor, Aeronautics and Astronautics, MIT

strategy is determined by factors like mission type, propulsion, maneuverability, and collision avoidance support. This complex decision-trade space motivates the use of game-theoretic methods to study behavior. More specifically, game theory provides a framework to analyze different strategic decisions between operators and determine best practices.

Collision avoidance between non-cooperative operators resembles classical problem in game theory known as the "game of chicken." The game of chicken is a two-player game that considers two drivers who are approaching one another in a head-on collision. If both continue to go straight, they will crash and incur a large negative cost. The only way to avoid a collision is for at least one of them to swerve; however, if one driver swerves and the other goes straight, the one who swerves is considered "chicken" and loses, while the driver that goes straight wins. This game has been used to model interactions in many contexts, such as in the Cuban Missile Crisis,[8,9] climate change,[10] and public goods.[11] In our work, we rely on simulations to obtain a more realistic model of the underlying system than prior studies.

We use two game-theoretic abstractions of the problem: first, we assume probabilistic operator actions and compute mixed Nash equilibria over multiple time steps; and second, we assume a risk threshold-based model of operator actions and determine the thresholds over multiple time steps. The first method finds a probabilistic strategy for each operator that minimizes their expected cost, given knowledge of the other player's risk threshold. The second method finds a series of risk thresholds based on an expectation of the other agent's probability of collision threshold that an operator makes a binary decision at. This parameter can be thought of as a players *risk aversion*. If an agent anticipates a higher threshold value for its opponent, then the expected behavior will be that the opponent is less risky and more likely to move. We then do a parameter study and, by varying the costs of thrust and collision, find how they affect the Nash equilibrium and risk threshold models.

We first describe our collision avoidance model and game, then present some results from tuning short-term and long-term costs in the one-step game. We then solve the multi-step game and present simulated results using a mixed Nash equilibria method and a risk threshold method.

MODEL

In this section, we describe the environment for our model, some preliminaries on game theory, and our game-theoretic framework. First, we provide an overview of our probability of collision calculation and propagation methods. As an introduction to game theory, we provide a brief overview of the game of chicken, which holds some parallels to our problem. Then, we describe our game theoretic framework.

Environment

Before we can calculate the probability of collision, we first propagate each state vector at time step $n-1$ to the time of closest approach t_{ca} using the SGP4 propagator.[12] Then, we transform the state vectors from the Earth-centered inertial frame into the relative motion frame.

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \frac{\vec{r_1}-\vec{r_2}}{|\vec{r_1}-\vec{r_2}|} \\ \frac{\vec{v_1}-\vec{v_2}}{|\vec{v_1}-\vec{v_2}|} \\ \frac{(\vec{r_1}-\vec{r_2})\times(\vec{v_1}-\vec{v_2})}{|(\vec{r_1}-\vec{r_2})\times(\vec{v_1}-\vec{v_2}).|} \end{bmatrix} \tag{1}$$

To calculate the probability of collision at each time-step, we rely on Foster and Estes 2D-P_C method to quantify collision risks.[13] The probability of collision problem is transformed into a 2D collision plane by assuming rectilinear motion of the two objects during encounter plane. This plane is defined as perpendicular to the relative velocities of the two objects, and assumes that the combined uncertainty along the relative velocity vector has no bearing on the calculation of probability of collision.

Assuming a circular cross-sectional area as the hard body radius (HBR), the resultant probability of collision is:

$$P_c = \frac{1}{2\pi\sqrt{|C|}} \int_{-HBR}^{HBR} \int_{-\sqrt{HBR^2-x^2}}^{\sqrt{HBR^2-x^2}} \exp\left(-\frac{1}{2}(\vec{r}-\vec{r_d})^T C^{-1}(\vec{r}-\vec{r_d})\right) d\vec{z}\, d\vec{x} \tag{2}$$

C is the combined position covariance for both objects, $\vec{r}$ is the position on the collision plane, and $\vec{r_d}$ is the debris object's position on the conjunction plane. These positions are given by

$$\vec{r} = \begin{bmatrix} x \\ z \end{bmatrix} \tag{3}$$

and

$$\vec{r_d} = \begin{bmatrix} x_1 - x_2 \\ z_1 - z_2 \end{bmatrix}. \tag{4}$$

Game Theory

Game theory studies interactive decision-making in systems where the choices and outcomes of each agent depend on the actions of others. Games are defined by the number of players N, the actions A_i^t available to each player i at each time t, and the cost to each player i at each time i, u_i^t, which may depend on the actions of other players. For a single stage game, which only contains one time step, each agent receives an expected cost given their action and the action of the opposing agent. The goal of each player is to minimize the cost incurred during the game. Table 1 gives the cost matrix for the game of chicken.

		Agent 2	
		Swerve	Straight
Agent 1	Swerve	$0, 0$	$1, -1$
	Straight	$-1, 1$	$10, 10$

Table 1: Cost matrix for the game of chicken.

Consider the best responses for agent 1 in the game of chicken, which are the actions that minimize agent 1's cost for each action of agent 2. Suppose agent 2 goes straight; then, agent 1 should swerve because $1 < 10$. On the other hand, if agent 2 swerves, agent 1 should go straight because $-1 < 0$. This produces an established phenomenon in game theory known as Nash equilibrium, where everyone is playing the best move they can in the circumstances. As swerving is the best response to going straight, and vice versa, there are two Nash equilibria: one in which agent 1 swerves and agent 2 goes straight, and one in which agent 1 goes straight and agent 2 swerves. These are both pure strategy Nash equilibria.

However, there is actually a third Nash equilibrium, where both agents choose to swerve with probability p; by symmetry, this strategy is identical for both players. To find the p that yields a Nash equilibrium, we can once again analyze best responses. Specifically, we can find the p that causes each agent to be indifferent to swerving or going straight, i.e., that results in both actions leading to identical costs.

The expected cost of swerving in response to the other agent swerving with probability p is $0p + 1(1 - p) = 1 - p$, while the expected cost of going straight is $-1p + 10(1 - p) = 10 - 11p$. Setting these equal to find the p that leads to indifference and thus a Nash equilibrium yields $1-p = 10 - 11p$, so $p = 9/10 = 0.9$. Thus, a strategy for both players of swerving with probability 0.9 and going straight with probability 0.1 is a mixed strategy Nash equilibrium.

We model the satellite collision avoidance problem as a game of chicken extended to multiple time steps $t = 0, 1, \ldots, T - 1$, where each operator decides to move their satellite ('swerve') or wait and continue in their orbit ('straight'). A possible collision occurs at time step T, when the satellites reach their distance of closest approach (DCA).

Another model for the game of chicken over time involves decision-making based on a series of risk thresholds, where an operator will move if the probability of collision is above its risk tolerance at time step t. This model aligns with current operator practice, where satellite operators set collision risk thresholds and move their satellites if the modeled closest approach is too near.

Model Description

Suppose there are two satellites, $i = \{1, 2\}$. We model this system as evolving with discrete time steps $t = \{0, 1, \ldots, T\}$ and $\Delta t = 1$ hr, with $t = 0$ being the first notice of a possible collision event and $t = T$ being the time of closest approach. At each time step, the true DCA is represented by a state $s^t \in \mathbb{R}_{\geq 0}$.

At each time step except $t = T$, each satellite can take an action $a_i^t \in \{0, 1\}$, where $a_i^t = 1$ indicates a maneuver to increase the relative DCA s^{t+1}. We assume that agents will always thrust in the direction that increases the relative DCA and that any thrust by any operator will increase the DCA to above the threshold of danger.

Taking an action incurs a cost for the satellite operator. We model the cost for satellite i as

$$g_i^t(a_i^t) = G^t a_i^t = \frac{G^{\max}}{k^{T-1-t}} a_i^t \ \forall t = \{1, \ldots, T - 1\} \tag{5}$$

for some $G^{\max} > 0, k \geq 1$. If a satellite does not thrust and $a_i^t = 0$, the cost to its operator is 0. If the satellite does thrust, the operator incurs a cost of G^t, with maximum value $G^{\max}$. This cost increases with time, as suggested by the power law term k^{T-1-t}, because a decreased time to closest approach requires more thrust to get a corresponding change in the DCA.

If no action is taken at time step t, i.e., $a_i^t = 0$, then $s^{t+1} = s^t$. A possible collision event is considered to have occurred if neither operator has moved by $t = T$, and both operators suffer the cost of (possible) collision $H \gg G^{\max}$. Given the probabilistic nature of satellite collisions, however, each operator may have a different risk tolerance towards the possible collision. This is represented by their respective types $\theta_i \in \mathbb{R}_{\geq 0}$, and the cost of collision to each operator is $H\theta_i$.

RESULTS

Tuning Short-Term vs Long-Term Costs

The Nash equilibrium strategy of either agent is highly dependent on the values of the short-term cost $G^{\max}$ and long-term cost H.

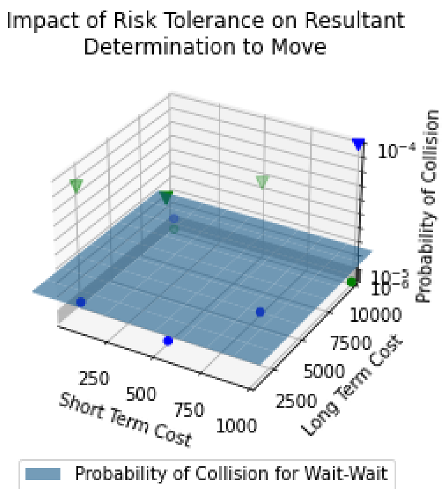

Figure 1: The impact of the short term cost, G, long term cost H, and risk tolerance on resultant action selection. A triangle indicates waiting, while a circle indicates moving. Green and blue differentiate between agents. The 3D plane represents the probability of collision for the wait-wait scenario. This factor has the largest impact on an agent's action for two agents with different risk tolerances.

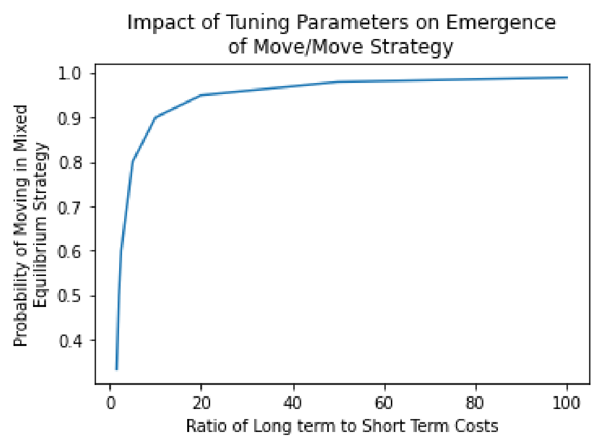

Figure 2: The impact of the ratio of long to short term cost (H/G), on the emergence of a mixed-Nash equilibrium strategy where both agents move with probability p, as shown on the y-axis.

Figure 1 demonstrates that the most significant parameter on the emergence of the move-wait strategy for operators with different risk tolerances is the probability of collision if both agents were to wait. If that probability is above an agent's risk tolerance threshold, then it will move.

Figure 2 compares the impact of tuning parameter ratio value (H/G) on the emergence of a mixed strategy Nash equilibrium, which involves at least one player using a randomized strategy. With a mixed strategy, a player's action is not deterministic, but rather chosen randomly according to a probability distribution over their actions. In the experiments to produce Fig. 2, both agents had a risk tolerance that was less than the probability of collision if both agents waited, $\theta < p_{wait\_wait}$. As both agents have identical risk tolerances and identical tuning parameters, we expect their resultant strategies to be symmetric. It should be noted that these cases also achieved a pure strategies of a_1 = move, a_2 = wait and a_1 = wait, a_2 = move. However, we wanted to investigate how tuning could result in a mixed strategy Nash equilibrium.

Figure 2 demonstrates that as the ratio of long-term to short-term cost increases, the mixed Nash equilibrium strategy distribution favors moving. While the one-step game never produces a pure

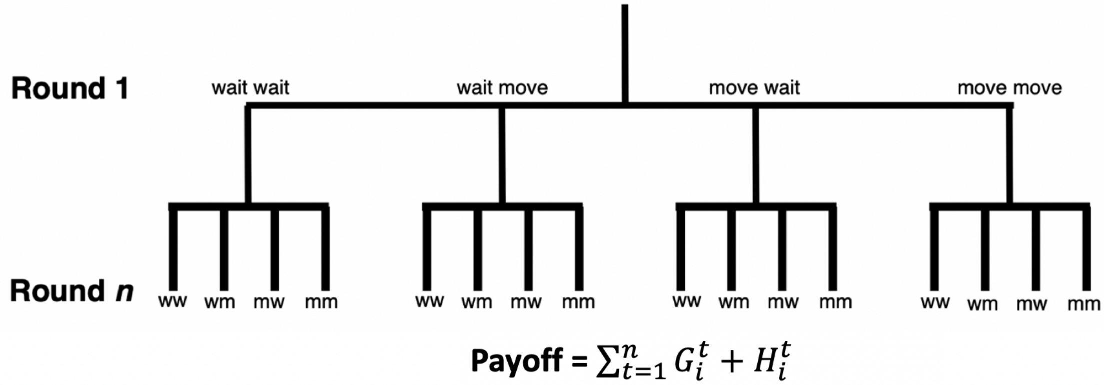

Figure 3: Cost in the multi-step case is a summation of the long-term and short-term costs for each iteration of the T-step game.

strategy where both agents move, when the long-term cost outweighs the short term cost by a factor of about 100, we see that the resultant probability distribution leads to moving over 99% of the time.

MODELING COLLISION AVOIDANCE USING MIXED NASH EQUILIBRIA

In the multi-step analysis, we transform the one-step game to a sequence of T repeated games. At each round, the agents must choose to wait or to move, where moving requires trading off short-term fuel cost and long-term collision cost. First we create and populate a decision tree.

Building the Decision Tree

At each time step, the probability of collision for each strategy is calculated. The risk tolerance of the controlled agent and the covariance (uncertainty) of the system are also updated. A monotonically decreasing risk tolerance as the time to collision gets closer reflects satellite operators taking the warning more seriously and consequently maneuvering.

As the time to collision draws closer, the length of propagation time decreases, and thus accuracy of the probability of collision estimate improves. To account for this engineering trade-off, we assume the covariance associated with the opposite spacecraft decreases at a rate of 10% each time step. This resultant reduction in uncertainty yields a smaller covariance, C, and thus is likely to reduce the probability of collision unless the two satellites are on a collision course with a very small separation.

These two updates create a unique short-term and long-term cost associated with that state for each of the four strategic outcomes (move-move, move-wait, wait-move, and wait-wait); these costs are saved for each round. To create the decision tree, the costs at each round are added together, as shown in Fig. 3. Once the decision tree has been created, we rely on alpha-beta-pruning[14] to determine the optimal solution based on the total reward at each step.

T-Step Results

To determine the impact of repetition on the resultant strategies between operators, we compare strategies between iterations of n-step games. Table 2 lists the recommended strategy for different numbers of iterations and the corresponding time step intervals for decision making. With additional time steps, we find that the satellite operator is more likely to maneuver earlier. With four iterations, we find several solutions, one where the recommended maneuver to move is at time step 2 and one where the recommended maneuver to move is at time step 3.

These results suggest that with additional decision-making intervals, individual satellite operators are more likely to adopt preemptive avoidance measures. This is a result of the risk-tolerance for each operator increasing at each time step. The decreasing risk tolerance means that a collision risk that would have been acceptable initially could become unacceptable at a later time.

Table 2: Recommended Strategies for Different N-step games

Maximum Iterations in Repeated Game	Time Step Interval	Recommended Strategy
2	3 hours	Force opponent to maneuver
3	2 hours	Move at first instance
4	1.5 hours	Maneuver at time steps 2 and 3

MODELING COLLISION AVOIDANCE USING RISK THRESHOLDS

One downside of the tree search analysis is that it searches for the system-optimal solution that maximizes both players rewards by assuming full knowledge of agent costs. Inspired by analysis of the Cuban missile crisis,[9] we can model how operators behave with limited information of the other agent's costs and when forced to make pure move or wait decisions (rather than a probabilistic decision) with respect to risk thresholds.

Suppose two operators $i \in \{1, 2\}$ have satellites that are on a collision course. Each operator can decide whether to wait or move at a time step $t \in \{1, \ldots, T-1\}$; we assume that an operator only requires one maneuver to completely avoid a collision. At time $t = T-1$, the cost matrix is given by Table 3.

		Operator 2	
		Move	Wait
Operator 1	Move	$G^{\max}, G^{\max}$	$G^{\max}, 0$
	Wait	$0, G^{\max}$	$p^{T-1}H\theta_1, p^{T-1}H\theta_2$

Table 3: Last step ($t = T-1$) cost matrix for the satellite collision avoidance game

If operator i decides to move, it incurs a movement cost of $G^{\max}$. If neither operator moves, both incur a catastrophic collision cost of $p^{T-1}H\theta_i$, where the *type* of an operator θ_i is drawn uniformly at random from $[0, \theta^{\max}]$. This type represents an operator's sensitivity to and risk assessment of such a catastrophic collision, where a higher type indicates a greater aversion to risk.

Threshold analysis

An operator with a high type, where $\theta_i \approx \theta^{\max}$, should be very sensitive to the cost of catastrophe and thus act to avoid a collision earlier in the scenario, while an operator with a lower type may be more willing to take risks and see how things develop before moving.

Therefore, we model the likelihood that an agent moves due to excessive risk as follows: We define a sequence of (unknown) thresholds $\theta^0 \geq \theta^1 \geq \dots \theta^{T-1} \geq 0$. If an agent type θ_i is greater than the threshold at the current time step θ^t, it will move. Our goal is to find a relationship backward in time, where θ^{t-1} is related to θ^t. We can find θ^{t-1} from θ^t by comparing the cost of thrusting later to the probability that the other satellite moves instead.

To solve for these thresholds, we assume θ^{T-1} is a constant, solve for the relationship between it and θ^{T-2}, then between θ^{T-2} to θ^{T-3} and so forth until θ^0. We can then tune the value of θ^{T-1} such that $\theta_0 = \theta^{max}$.

Let q^t be the probability that an arbitrary operator moves at time step t. It is operator-agnostic due to our assumption of a uniform distribution of types. If an operator has not moved by time step $t-1$, it must have $\theta^{t-1} \geq \theta_i \geq 0$. Based on the next threshold θ_i, the probability of operator i moving at time step t is the probability that $\theta_i \in [\theta^t, \theta^{t-1}]$. As such,

$$q^t = \frac{\theta^{t-1} - \theta^t}{\theta^{t-1}} \tag{6}$$

A corollary of (6) is

$$1 - q^t = \frac{\theta^t}{\theta^{t-1}} \tag{7}$$

which is the probability that an operator waits.

We first analyze $t = T-1$ from the perspective of operator 1, although the conclusions are the same for operator 2 by symmetry. Operator 1 expects operator 2 to move with probability q^{T-1} and to wait with probability $1 - q^{T-1}$. Then, if

$$G^{\max} < q^{T-1}(0) + (1 - q^{T-1})(p^{T-1} H \theta_1) \tag{8}$$

moving incurs a lower cost than waiting, and operator 1 will move.

We can solve this equation to relate the threshold θ^{T-1} to θ^{T-2}. Assume θ^{T-1} is some constant. First, we substitute thresholds in for q^{T-1} using (7) and rearrange to isolate θ_1.

$$G^{\max} < p^{T-1} H \theta_1 \frac{\theta^{T-1}}{\theta^{T-2}} \tag{9}$$

$$\theta_1 > \frac{G^{\max}}{p^{T-1} H} \frac{\theta^{T-2}}{\theta^{T-1}} \tag{10}$$

When $\theta_1 = \theta^{T-1}$, operator 1 is indifferent between moving and waiting. This is the indifference threshold, where an agent is equally well off by moving and waiting. Then, by replacing (10) with an equality and substituting θ_1 with θ^{T-1}, we can obtain an expression for θ^{T-2} as a function of θ^{T-1}:

$$\theta^{T-1} = \frac{G^{\max}}{p^{T-1} H} \frac{\theta^{T-2}}{\theta^{T-1}} \tag{11}$$

$$\theta^{T-2} = \frac{p^{T-1} H}{G^{\max}} (\theta^{T-1})^2. \tag{12}$$

However, because $\theta^{T-2} \geq \theta^{T-1}$ by definition, we add a max function:

$$\theta^{T-2} = \max\left\{\theta^{T-1}, \frac{H}{G^{\max}}(\theta^{T-1})^2\right\}. \tag{13}$$

The scenario where $\frac{p^{T-1}H}{G^{\max}}(\theta^{T-1})^2 < \theta^{T-1}$ can be interpreted as the cost of thrust $G^{\max}$ being so great that an agent would rather risk collision than move. However, this is unrealistic for this problem, and thus values that fall into this paradigm are filtered out.

This gives us a relationship between the thresholds θ^{T-2} and θ^{T-1}.We continue this process of solving backwards to θ^0 by finding a general relationship between θ^{t-1} and θ^t for $t = T-2, \ldots 2$. First, let us look at the cost matrix for an arbitrary time t, as shown in Table 4. Here, $V_i^{t+1}(\theta_i)$ is the expected cost to operator i of proceeding to the next time step.

		Operator 2	
		Move	Wait
Operator 1	Move	G^t, G^t	$G^t, 0$
	Wait	$0, G^t$	$V_1^{t+1}(\theta_1), V_2^{t+1}(\theta_2)$

Table 4: Time step t cost matrix for the satellite collision avoidance game

We derive the relationship between θ^{t-1} and θ^t by analyzing the threshold where an operator is indifferent to moving or waiting at time t. Consider when $\theta_i = \theta^t$. Then, indifference occurs when:

$$G^t = (1-q^t)V_i^{t+1}(\theta_i) \tag{14}$$

Because $\theta^{t-1} > \theta^t > \theta^{t+1}$, if $\theta_i = \theta^t$ then the agent will concede at the next timestep. Thus then $V_i^{t+1}(\theta_i) = G^{t+1}$, because the agent will move.

By applying (5) and (7) to (14), we obtain the following:

$$G^t = kG^t\frac{\theta^t}{\theta^{t-1}} \tag{15}$$

$$\theta^{t-1} = k\theta^t. \tag{16}$$

This gives us a recursive definition of θ^{t-1} as a function of θ^t. This recursion can get us from θ^{T-2} to θ^1. Then, one more step using (16) gives us θ^0, where we require $\theta^0 = \theta^{\max}$. We tune the value of θ^{T-1} to find where this holds and to derive the set of thresholds an operator should choose to move by. Above the threshold, the agent should move now, while thrusting is cheaper; below the threshold, the agent is willing to take on the risk of moving later.

The operator analysis then proceeds as follows at each time step t:

1. Each operator receives an updated probability of collision p^t. $\theta^{\max}$ is set to the previous threshold of θ^{t-1} because if the game has continued, both parties have $\theta_i \leq \theta^{t-1}$.

2. The updated p^t yields a new final time step cost matrix. Using (13) for the last timestep and (16) for all other timesteps (and tuning to reproduce $\theta^{max} = \theta^{t-1}$), operators follow the threshold analysis to generate thresholds $\theta^{T-1}, \theta^{T-2}, \ldots, \theta^t, \theta^{t-1}$. If $p^t = p^{t-1}$, we should recover the same thresholds as the previous time step.

3. Operators compare θ^t with their own risk threshold and decide whether or not to move. The collision scenario then proceeds to the next timestep.

Results

To study the sensitivity of the thresholds θ to the constants k, T, we implement this recursion for $T = 5$. We hold $\theta^{\max} = 10$ and $p^{T-1} = 1$ constant, and solve for the thresholds with $\theta^0 = \theta^{\max}$ across k and H/G^{max}. We set $k = \{1.01, 1.02, 1.05, 1.1, 1.2, 1.5, 2, 3, 5\}$ and $H/G^{\max} = \{1, 2, 5, 10, 20, 50, 100, 200, 500\}$. The results are given in Fig. 4, 5, and 6. We study θ^{T-2} and θ^{T-1} separately in Figs. 4 and 5 respectively, and the ratio $\theta^{T-1}/\theta^{T-2}$ (the difference in the last two thresholds) in Fig. 6.

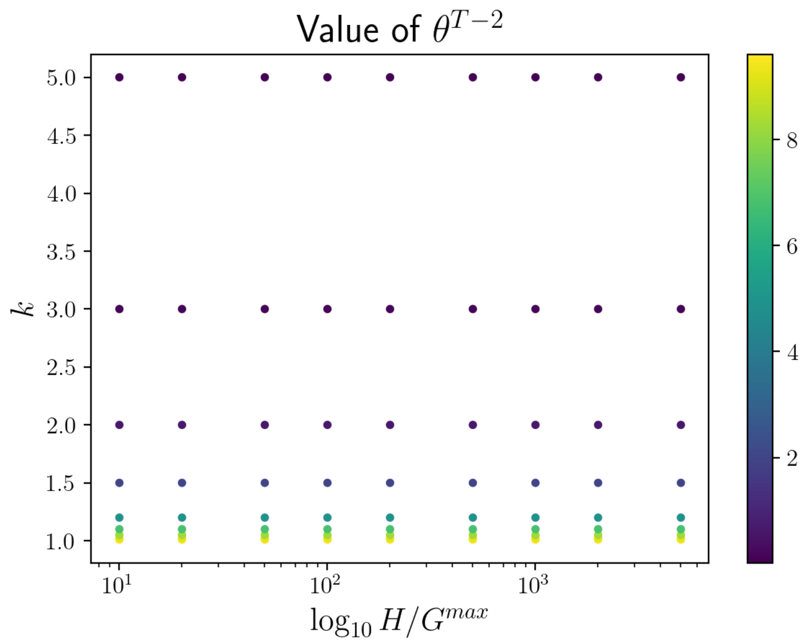

Figure 4: Value of θ^{T-2}

The first Fig. 4 shows that θ^{T-2} is only dependent on k; as k decreases, θ^{T-2} increases. This is because the recursion given in equation (16) only depends on k. A lower value of k means that the cost of moving later does not increase much compared to moving earlier. Then, because θ^{max} is constant, the difference between successive thresholds are low and the value for θ^{T-2} is higher.

In Fig. 5, we see that θ^{T-1} is more complex and depends on both H/G^{max} and k. First, as H/G^{max} decreases (as H reaches parity with G^{max}), the impact of a collision is lessened. Agents are more comfortable with being risky, and the threshold above which operators move on the last time step becomes higher. The dependence on k follows from the above explanation.

There is a large variance in the ratio of $\theta^{T-1}/\theta^{T-2}$ depending on k and H/G^{max}. In Fig. 6, a significant drop in thresholds occurs for high H/G^{max} and low k, showing that many agents will move at the last time step because the cost of collision is very high and the cost of moving later is not too high. On the other hand, for low H/G^{max} and high k, the ratio is low—exactly 1—due to

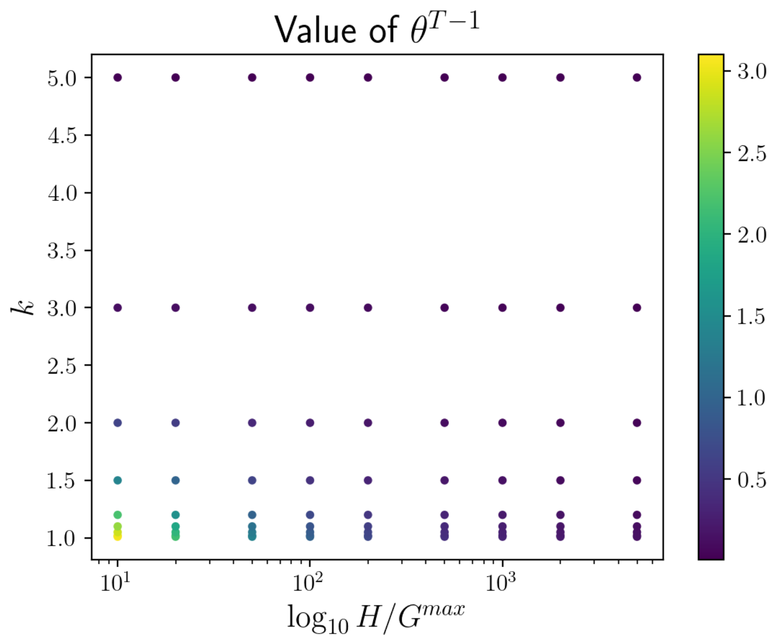

Figure 5: Value of the last threshold θ^{T-1}

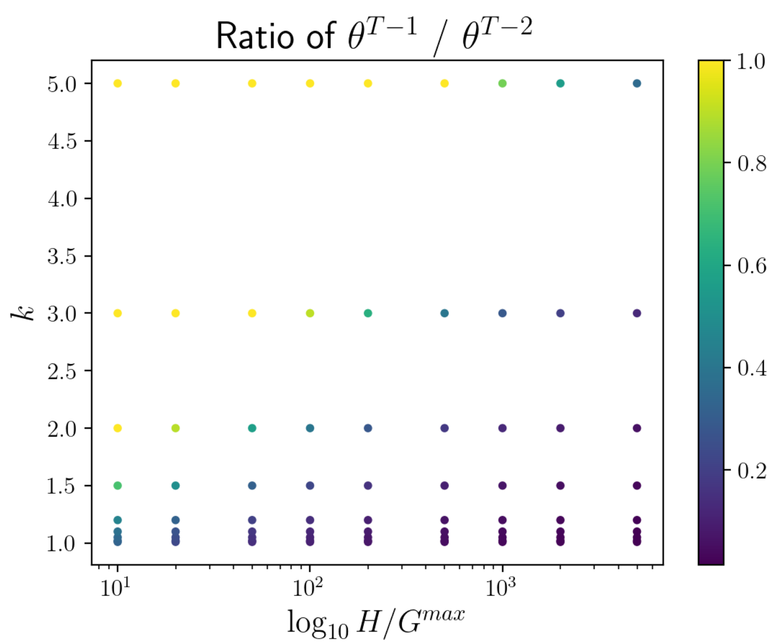

Figure 6: Ratio of $\theta^{T-1}/\theta^{T-2}$

the $\max$ function; if an agent were to move on the last time step, it would have done so at an earlier timestep ($t = T - 2$) because the cost of moving later is very high.

CONCLUSIONS

In this paper, we present a game-theoretic framework to analyze the satellite collision avoidance problem. This framework is designed for collisions between non-cooperative operators, and provides insight into the maneuver timing given a dynamic risk tolerance threshold. We also derive a risk threshold model of satellite collision avoidance that identifies the risk level at which satellite operators should move to avoid a collision. Future work can incorporate satellite position uncertainty into this threshold model to identify an operator's willingness to wait for more accurate position information closer to the time of closest approach.

ACKNOWLEDGMENT

Research was sponsored by the United States Air Force Research Laboratory and the Department of the Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Department of the Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

[1] J. Hecht, "Satellite Collision could have been far worse," *New Scientist*, Apr 2009.

[2] R. S. Jakhu, "Iridium-cosmos collision and its implications for Space Operations," *Yearbook on Space Policy*, 2010, p. 254–275, 10.1007/978-3-7091-0318-0_10.

[3] A. Koop and S. Fortin, "Visualized: Which countries are dominating space?," Jul 2022.

[4] T. Hitchens, "Space traffic management: U.S. military considerations for the future," *Journal of Space Safety Engineering*, Vol. 6, No. 2, 2019, pp. 108–112. Space Traffic Management and Space Situational Awareness, https://doi.org/10.1016/j.jsse.2019.04.003.

[5] R. Hiles, A. Alexander, N. Herzer, D. McKissock, A. Mitchell, and J. Sapinoso, "Report on 2020 Mega-Constellation Deployments and Impacts to Space Domain Awareness," AMOS, 2021.

[6] J. Foust, "SpaceX and OneWeb spar over satellite close approach," *SpaceNews*, Apr 2021.

[7] L. Grush, "China complains to UN after maneuvering its space station away from SpaceX Starlink satellites," *The Verge*, Dec 2021.

[8] F. Zagare, "A game-theoretic history of the Cuban Missile Crisis," *Economies*, Vol. 2, No. 1, 2014, p. 20–44, 10.3390/economies2010020.

[9] A. K. Dixit, D. McAdams, and S. Skeath, "'We Haven't Got But One More Day': The Cuban Missile Crisis as a Dynamic Chicken Game," , 10.2139/ssrn.3406265

[10] K. Pittel and D. T. Rübbelke, "Transitions in the negotiations on climate change: From prisoner's dilemma to chicken and beyond," *International Environmental Agreements: Politics, Law and Economics*, Vol. 12, No. 1, 2010, p. 23–39, 10.1007/s10784-010-9126-6.

[11] B. H. Chong, *Extensions to the Chicken Game in the Context of Public Goods*. PhD thesis, 2020.

[12] D. Vallado and P. Crawford, "SGP4 orbit determination," *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, 2008, 10.2514/6.2008-6770.

[13] J. L. Foster and H. S. Estes, *A parametric analysis of orbital debris collision probability and maneuver rate for space vehicles*. NASA Lyndon B. Johnson Space Center, 1992.

[14] D. E. Knuth and R. W. Moore, "An analysis of alpha-beta pruning," *Artificial Intelligence*, Vol. 6, No. 4, 1975, p. 293–326, 10.1016/0004-3702(75)90019-3.