Re-Evaluating the Hider-Finder Problem

Regan Copple
Harrison Schramm
UNITED STATES OF AMERICA

rcopple@groupw.com, hschramm@groupw.com

ABSTRACT

The United States, her Allies, and adversaries are embracing advances in computing environments and technology (henceforth, AI) for both civil and military applications. Our work proposes to explore a central tension in autonomous and semi-autonomous systems, which is the fundamental tradeoff between the length of an unmanned system's deployment (sojourn time) and the uncertainty between individual system failures due to unscheduled maintenance and those due to adversary action. This paper is unique in that it will explore the problem from both a policy lens as well as an applied statistics point of view and provide insights for both the procurement and employment of unmanned systems more broadly.

INTRODUCTION

What does it mean to trust an unmanned system? This seemingly innocuous question is a central problem in a variety of both civil and military applications of unmanned technology. As the world becomes more automated and less supervised by human caretakers, autonomous systems will be left to endure for extended periods of time with limited monitoring. While autonomous systems with prolonged endurance could be a valuable asset for civilian and military maritime monitoring purposes, the decay of constellations of autonomous systems is an inevitable reality. But the discrepancy between the expected rate of degradation and actual rate could provide an early indicator of malign interference with a constellation of unmanned systems.

The aforementioned analysis raises an important issue, which is the attribution of failure when no response is detected and a given system is presumed destroyed. If nonresponse is the sole indicator of system failure, it could be caused by a variety of factors including the environment, mechanical failure, or most worryingly, attrition, which we define as malign adversary action intended to disable and destroy unmanned systems. Using this definition, attrition is separate from failure because failure results from scenarios not involving purposeful adversary action. Attrition is the most serious scenario for military platforms, which could be carrying classified sensors and payloads onboard that the operator does not want falling into the hands of an adversary's military.

Similarly, understanding natural degradation patterns could enable one state to degrade an adversary's systems in a manner that is difficult to distinguish from random failures. Such an approach would alternatively enable an ability to engage in offensive operations or limit Blue's detection of indications and warnings (I&W) before Red engages in a military operation.

Our contribution in this note is to consider both the reliability and game aspects of unmanned systems degradation in a unified structure. Our paper is structured as follows. In Section Two, we review literature from both the applied mathematical and policy viewpoints. In Section Three, we lay the preliminary, mathematical unifying aspects of both reliability and game theory. In Section Four, we explore specific scenarios before finally summarizing conclusions and providing recommendations for future analysis.

LITERATURE REVIEW

This paper ties together thoughts from both applied probability theory, reliability, and some aspects of Game Theory. For a treatment of reliability as a Markov Process, see Grimmet and Stirzaker (2001). For an application of reliability to similar military problems, see (Gaver et al, 1984 and similar). For an exploration of strategically choosing a so-called 'strike time' in a two-player competition with varying payoff functions, see Schramm et al (2014).

From a statistical point of view, determining when a process is maliciously subverted is a classic problem, touching on issues of both reliability as well as fraud detection. From a game theoretic point of view, two 'games' are being simultaneously played in a unified construct; first, a game is played against an unthinking adversary – unscheduled failures – where Nature is the adversary. Second, a game is being played by a thinking adversary who may have partial information about the state of Nature against a thinking adversary. We recommend interested readers to see recent papers by (GG Brown) about the difference between these games.

Analytically, we make a strong presumption of stationarity for the reliability process, by which we mean that the hazard experienced by any remote, autonomous entity is solely a function of the amount of exposure to the hazard it experiences. This leads to a fixed hazard rate. It can be shown (see Ross or similar) that the presumption of a fixed hazard rate leads directly to the Poisson Process. Nuanced readers may object to the use of the Poisson Process for a mechanical system. However, we justify it on two grounds. First and perhaps most importantly for our purposes, it allows us to frame both the 'games' as a single Markov Chain. Secondly, we assert that for the systems we consider, the hazard time is sufficiently long that even if the process is not truly stationary, it is at least locally stationary for some insights.

As improvements in fuel economy and battery technology enable deployment of autonomous systems for months at a time, longer deployments will create more opportunities for adversaries to tamper with unmanned surveillance systems (Sliverajan et al, 2018). The U.S. is already adapting tactics and strategies including Distributed Maritime Operations (DMO) that are specifically designed to leverage long-term deployment of unmanned systems to tax adversary ISR systems by providing adversary ISR nodes with more targets than each can simultaneously track (Popa et al, 2018).

From a policy perspective, determining attrition of unmanned systems can influence the employment of autonomous systems for surveillance and reconnaissance purposes. One of the principal values of unmanned systems is their ability to operate farther forward than manned platforms while taking on less risk due to the lack of personnel onboard. Tampering or destroying an unmanned system does not lead to a loss of life, which creates fewer consequences for Blue and additional incentive to use unmanned systems more frequently if there is insufficient will among national security decision-makers to deploy manned personnel into risky areas of action.

A second policy issue worth considering is the attribution problem. Classifying a system's nonresponse as failure or attrition is complex due to the relative anonymity of the sea and air, and a lack of ex post facto forensic evidence that would indicate the cause of an unmanned system's nonresponse. In this respect, attributing the failure of unmanned systems is similar to the attribution problem present in cyberspace (see Finlay and Payne, 2019). Altogether, the political inclination to use unmanned platforms for ISR in forward operating areas coupled with an attribution problem for long-endurance systems creates the need to better understand how attrition may occur and what steps could be taken to mitigate Red attempts at attrition against Blue's unmanned ecosystems.

ANALYSIS

Mathematical Underpinnings

We begin by defining the state space, which differs depending on the observer. The fundamental matrix of transitions contains omniscient knowledge of ground truth, where the Blue player (who owns the unmanned systems) has partial information on the state of deployed systems as a function of their interrogation / communication policy. Without loss of generality, we call the Blue player the maximizing player in the sense that they want to maximize the number of units operating either at any given time (the general problem) or at a given strike time (which we call the culmination problem). Conversely, the Red player is the minimizing player and either seeks to minimize Blue’s force generally or at a given time of their choosing. Thus, there are four possible instantiations of this game:

Our first line of inquiry is as follows: We consider a system where unmanned vehicles are deployed remotely from their home station to a distant location. After deployment, these vehicles are subject to two types of 'loss':

a. Mechanical malfunction (failure)

b. Enemy action (attrition)

From the home station's point of view, both types of losses are instantiated as a 'failure to report', and therefore are identical.

Our approach is to create a continuous time Markov Chain, where the state space represents the number of entities that are still operational per unit time. The side deploying the entities, which will refer to as the 'Blue' side without loss of generality, will know which entities are still operational, but information as to the cause of failure will be 'hidden'. The Red side seeks to minimize the number of operational units; but does so with the constraint that they want to do so in such a way that hides in the 'noise floor.’ If the Blue side sees too many inoperative systems, they may infer that Red is actively interfering.

This construct is oddly similar to cheating at cards; one wants to be 'luckier than average', but not so lucky as to arouse the suspicions of the house (or other players).

First, we consider the case where there is no interference from Red at all. In this instance, where there are 10 remotely deployed units, we can see the rate of attrition rises the longer the ecosystem of unmanned units operates. At the same time, the rate of failure falls because there are fewer units available to fail independent of any tampering or malign influence.

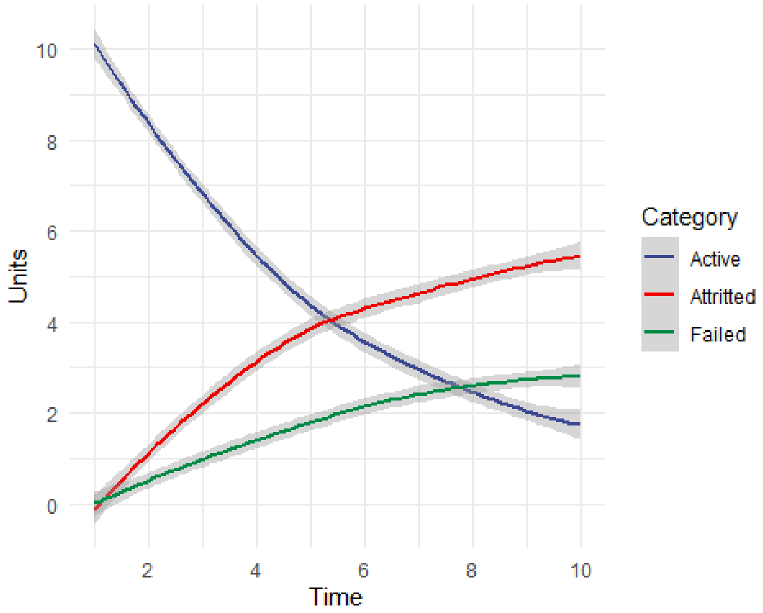

Figure 1: System Attrition from Opportunistic Red Strike Times

Figure 1 displays the scenario where Red attacks and destroys any Blue unmanned system they may find. This opportunistic pattern is not likely in reality as Red will likely attempt to avoid creating Blue suspicion that their systems are being attrited rather than failing due to mechanical malfunction and natural causes. In this case, we would rate risk as 1, or 100 percent, since every system found by Red will be destroyed.

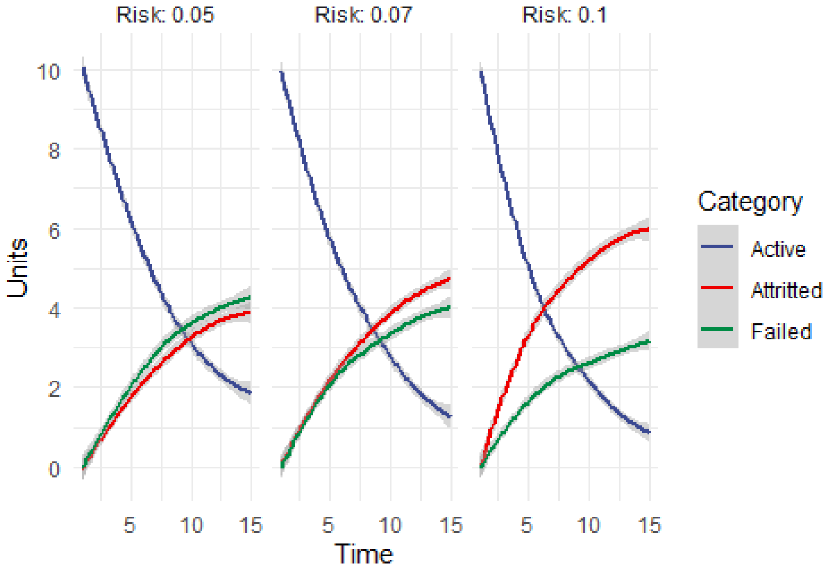

Figure 2: Health of Fleet as a Function of Risk with Perfect Knowledge

In Figure 2, risk is defined as a subjective assessment the Blue commander makes regarding the likelihood that any one individual system will be attrited in a specific period of time. In this case, the time period under examination is ten months. The graph on the left demonstrates a risk-averse commander's rate of attrition and failure, while the graph on the right demonstrates a risk acceptant Blue commander. This, of course, is more information than the Blue commander will see. From Blue's perspective, the difference between the red and green lines in the above graphs is not observed because Blue can only determine if a signal is lost, not why an unmanned ship's signal has been lost. Seen through this lens, the difference between the two graphs above is not nearly as stark. However, it is important to note that the difference between five and ten percent attrition is relatively minimal, because a ten percent risk of attrition is still relatively low when considering Blue would likely be deploying unmanned assets forward of manned assets, which would mean unmanned ships and aerial vehicles would likely operate close to an adversary's territorial waters (Talley, 2017). Operating close to shore or in contested areas would increase the probability of attrition to nearly half, which would yield far more drastic attrition curves than those shown below.

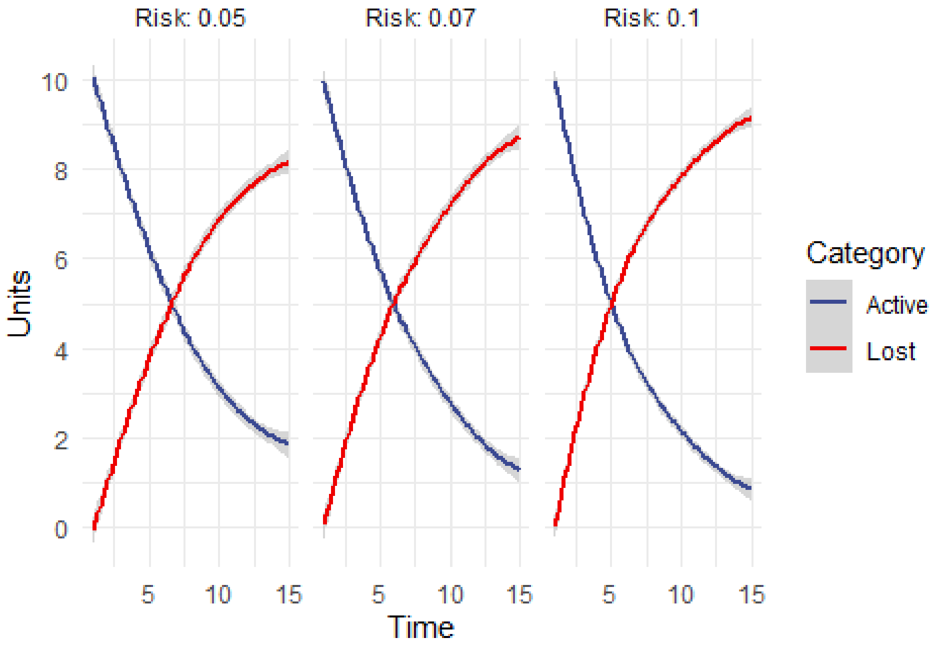

Figure 3: Health of Fleet as a Function of Risk from Blue's Perspective

Figure 3 shows while the per-unit time risk from the Red force has risen significantly, the overall impact from Blue's perspective is minimal. Blue commanders may not immediately recognize one additional asset lost as a sign of increased attrition but may instead assume the higher-than-expected attrition rate is due to an unforeseen mechanical issue rather than malign interference. As a result, Red can conduct significant attrition without creating suspicion among Blue commanders.

Red's Opportunity

Red's opportunity to find Blue targets is a function of both the number of Red assets and their *radius of action*. While detailed, statistical models of particles in motion exist, simulation suffices. We model the Blue fleet as randomly distributed particles in a 2-dimensional field, with reflecting barriers.[1] A single Red searcher is also present, and when the distance between the Red searcher and a Blue target is within the

[1] By which we mean that when a particle hits a barrier, it is reflected into the field.

search distance, the target is considered to be 'acquired'. As the number of times a Red searcher closes within search distance of a Blue target, the fewer undetected Blue units exist. This finding is important because the greater the number of units Red detects, the greater number of units Red is capable of destroying or influencing.

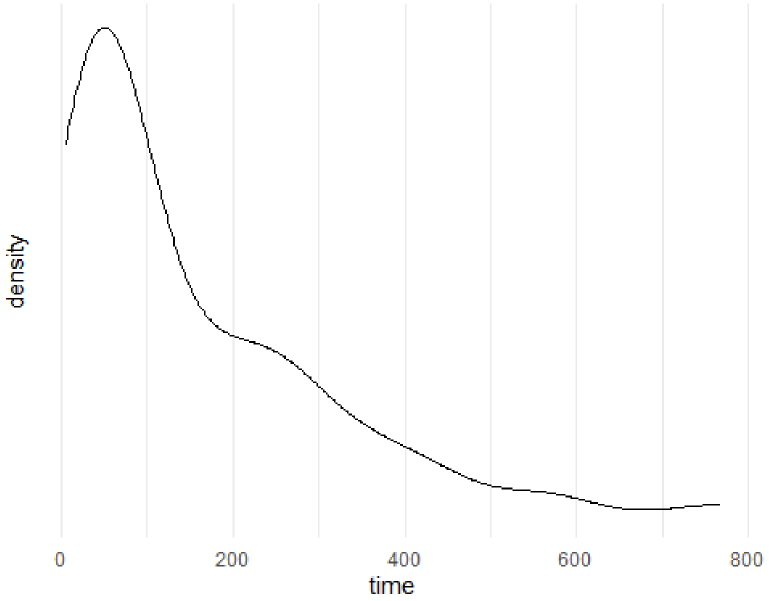

Figure 4: Number of Blue Units Remaining Hidden Over Time

Shortly after the deployment of a constellation of unmanned systems, Red will have the most potential targets to search for because not only are most if not all Blue platforms still operational, but they have also not been discovered. Therefore, the number of operational assets in a given search area will be at its highest and the grouping of assets will be at its most “dense.” However, as time progresses, the more units that have been discovered through a search increase, and the time between the discovery of each additional unmanned asset lengthens. The curve above demonstrates the number of Blue units that Red has yet to discover and the time between each discovery. As more Blue units are discovered and some Blue units fail, the number of undiscovered Blue platforms in the search area will diminish, and the time between discovering new Blue units will increase.

Strategic Extensions

Our plot above brings into focus the notion that a minimizing player (Red) has *agency* in the sense that they may be able decide upon a so-called 'strike time'.[2] Red's strategy is to maximally degrade Blue's force over time without arousing Blue's suspicions. We assume that the Red player is indifferent to whether Blue's assets are attritted by Red, or if they simply fail. Part of Red’s ability to degrade the Blue fleet depends on Red’s ability to detect the location of Blue unmanned units, which is depicted in the plot above. We explore three cases below:

[2] Here, 'strike time' takes on both a literal and mathematical meaning. The strike time is when the Red player decides to exercise its *option*, which in this case is an actual strike.

1. Red has no information about Blue's fleet status

From a technical perspective, the case where Red has no information or knowledge as to the disposition of Blue's fleet is the most likely scenario. In this case, Red is faced with the dual problems of not knowing the true reliability of Blue's assets (failure due to unscheduled maintenance) or the effectiveness of their own countering efforts. In this instance, there is unlikely to be a nuanced solution that will give Red an advantage without tipping their hand to Blue. A follow-on question for further consideration is whether or not this case implies that Blue themselves does not have a good estimate of their own capability.

2. Red has incomplete information

In the case where Red has partial (in the sense of not being completely reliable) information on the state of Blue's force, they may be able to take a hedging strategy. This strategy relies on the two-sided, joint probability functions,

$$E[ObservedLoss|Loss]_{red},\ E[ObservedLoss|Loss]_{blue}$$

These functions are clearly not independent, and are expressions of each side's risk tolerance, strategy, and to some degree - fear - as they are statistical reasoning.

3. Red has complete information

In the case where Red has complete information, they may pursue a strategy where commanders know the number of Blue assets, they would like to attrit, have some sense of the amount of time (availability) to be attritted by Red, and the *Strike Time*.

Reported Versus Expected Reliability

Perhaps one of the best strategies for Red is to take advantage of is the fact Blue deploys its' forces with an *expected reliability*, which is determined using various techniques of Operations Research and Applied Statistics. However, as highlighted in a number of recent articles on readiness (see McLemore et al, 2021 and similar), expected readiness is generally not seen in the 'wild'. Red can take advantage of this by intentionally causing failures in Blue's force early. These losses in force from Red action are not easily differentiated from mechanical failures, with Blue assuming that the excess losses are from inaccurate estimates of reliability instead of adversary action. This problem is therefore similar to an issue seen in the commercial satellite sector (see Gures et al, 2019).

CONCLUSION

In this short note, we have highlighted the following tensions: First, the need for users of fleets with remote units that are unable to return to base for maintenance to consider the tension between failure and adversary action. Second, from an adversary perspective, the strategic tradeoff between minimizing Blue's overall fleet in time vice minimizing its capability at a particular moment in time (strike time). Understanding this tension and optimal Red strategies will - for the Blue player - lead to better technical and policy decisions.

Further work along these lines of inquiry will involve replacing the notions about failure and attrition with actual distributions, distilled from data. An additional area for exploration will be to consider how these distributions will be adjusted in real time during deployment.

Declaration: The authors declare no conflict of interest.

REFERENCES

[1] Fearon, James. 1995. “Rationalist Explanations for War.” *International Organization* 49 (3): 379–414.

[2] Feller, William. 2009. *An Introduction to Probability Theory and Its Applications. Vol. 1*. 3. ed., rev. print., [Nachdr.]. Vol. 1. Wiley Series in Probability and Mathematical Statistics. S.l.: Wiley.

[3] Finlay, Lorraine and Christian Payne. 2019. “The Attribution Problem and Armed Cyber Attacks,” *American Journal of International Law* 113 (1): 202-206.

[4] Gaver, D. P., P. A. Jacobs, and G. Latouche. 1984. “Finite Birth-and-Death Models in Randomly Changing Environments.” *Advances in Applied Probability* 16 (4): 715–31. doi:10.2307/1427338.

[5] Grimmett, Geoffrey, and David Stirzaker. 2001. *Probability and Random Processes*. 3rd ed. Oxford; New York: Oxford University Press.

[6] Gures, Seda Demirbas, Ilkay Ulusoy, and Burak Durmaz. 2019. “Satellite Failure Estimation Vs. Reliability Prediction Analysis.” In *2019 Annual Reliability and Maintainability Symposium (RAMS)*, 1–5. Orlando, FL, USA: IEEE. doi:10.1109/RAMS.2019.8769031.

[7] McLemore, Connor, Shaun Doheney, Sam Savage, and Phillip Fahringer. 2021. “Military Readiness Modeling: Changing the Question from ‘Ready or Not?’ to ‘How Ready for What’.” *Military Operations Research* 26 (1).

[8] Schramm, Harrison, David Alderson, W. Matthew Carlyle, and Nedialko Dimitrov. 2014. “A Game Theoretic Model of Strategic Conflict in Cyberspace.” *Military Operations Research* 19 (1): 5–17.

[9] Silverejan, Bilhanan, Ocak, Mertand Nagel,Benjamin. "Cybersecurity Attacks and Defences for Unmanned Smart Ships," *2018 IEEE International Conference on Internet of Things and IEEE Green Computing and Communications and IEEE Cyber, Physical and Social Computing and IEEE Smart Data*, 2018, pp. 15-20, doi:10.1109/Cybermatics_2018.2018.00037.

[10] Talley, Edward W. “Defending an Area with Autonomy: Autonomous Intelligence, Surveillance, and Reconnaissance Capabilities Leveraging Unmanned Aerial Systems for Defending Forward Operations Locations,” *United States Air Command and Staff College*. Montgomery, Alabama: Air University Press, p. 12-15.

[11] Thie, Paul. 1983. *Markov Decision Processes*. The UMAP Expository Monograph Series. Lexington: COMAP.