(Preprint) AAS 21-756

TOWARD CO-EVOLVING SOLUTIONS OF ADVERSARIAL GROUND STATIONS TRANSIT TIME GAMES FOR P-LEO CONSTELLATION MANAGEMENT

Manuel Indaco,* Sean Harris,† Deacon Seals,† Daniel R. Tauritz,‡ Davide Guzzetti§

As society increasingly relies on space infrastructure for both civil and military applications, the number and the size of satellite constellations in low Earth orbit (LEO) is expected to substantially increase within the next decade; cost reduction combined with increased availability of satellite technology across different functions is granting access to space and space assets to a larger number of parties, potentially including bad actors. Malevolent actors may target both the ground and the space segment on different levels: hacking, spoofing, hijacking and command intrusion are just a few examples of the variety of attacks that can affect the performance of a constellation. In the context of different threats, the impact of an attack is greater the longer it remains undetected. Without satellite interlinks or additional in-orbit supporting networks, orbital motion is one of the primary factors contributing to the detection time for LEO assets. At the minimum, a compromised satellite has to transit from the attack location to the first available defender ground station for the attack to be detected and a mitigation action to occur. In this work, we investigate the particular problem of adversarial ground station transit time by applying coevolutionary algorithms. The solution space of adversarial ground station transit time may rapidly become intractable for analytical or brute force approaches just by the addition of a small number of satellites, orbital planes, or ground stations; thus, after defining a semi-analytical method that is able to describe the problem with one satellite and two ground stations, we apply meta-heuristic algorithms, specifically competitive coevolution, to explore adversarial ground station transit time games with increasingly higher-dimensional solution spaces. For example, coevolutionary algorithms may enable solving adversarial scenarios that include mobile attacking stations, ones that may introduce an additional layer of complexity in determining the origin and attribution of an attack. Results demonstrates how, through evolution, effective strategies for malevolent actors can be found even within scenarios whose mathematical formulation is either too difficult to develop or, if the formulation exists, does not have a closed-form solution.

INTRODUCTION

In the near future, the number of small satellites orbiting the Earth in the form of large constellations is projected to exceed ten thousand. One of the leaders in the field of proliferated low

*PhD Student, Department of Aerospace Engineering, Auburn University, Alabama, 36849, U.S.A..

†PhD Student, BONSAI Lab, Department of Computer Science and Software Engineering, Auburn University, Alabama, 36849, U.S.A..

‡Associate Professor, BONSAI Lab, Department of Computer Science and Software Engineering, Auburn University, Alabama, 36849, U.S.A..

§Assistant Professor, Department of Aerospace Engineering, Auburn University, Alabama, 36849, U.S.A..

Earth orbit (P-LEO) constellations, SpaceX, is currently operating almost thirteen hundred satellites, with the ambitious goal of reaching twelve thousand units (i.e., the number currently approved by the Federal Communication Commission (FCC)). P-LEO constellations enable new forms of omnipresence, from delivering high-speed internet to remote corners of the globe, to continuous Earth monitoring. Security threats to space assets, including to P-LEO constellations, are a relatively new phenomenon, yet they have rapidly come to the forefront of concern as such threats are low-cost, difficult to trace, and, depending on the type of attack, can generate damages of variable extent. Multiple types of attacks can be performed on a variety of targets: malware and command intrusions can be introduced into the space segment, spoofing can affect both the user and link segments, while hacking or hijacking can affect the ground segment.[1] Considering the availability of lightweight commercial technology[2] and the military capabilities of state actors to utilize mobile ground stations or anti-satellite weapons, one of the main characteristics of threats to P-LEO assets is that they can be conducted from fixed unknown ground stations or unknown mobile stations that are mounted on naval, aerial, or terrestrial vehicles, thus becoming difficult to locate. Different authors have initiated effort to prepare for such threats. For example, J.S. Turner[3] casts an Attacker-Defender model for a LEO constellation case into a bi-level integer programming (IP) problem, attempting to identify a scheduling for the possible location and time of an attack; however, given the large dimensionality of the problem, the author limits the application of his framework to GEO constellations with a reduced number of satellites. K. Wagner[4] includes simplified adversarial actions in the maximization of constellation resiliency; in this study, an importance level is associated with each individual satellite, thus hypothesizing that bad actors will prioritize attacks on satellites with the highest importance in order to deliver the largest damage. Adversarial actions are also considered by Zaerpoor et al.[5] while investigating the constellation resilience problem; the authors present a framework to identify actions that yield the highest degradation of system performance for the lowest cost. In all these cases, a deterministic strategy is assumed for the adversary, without considering action-planning for the attack strategy.

Another key element of a P-LEO constellation design to account for while considering adversarial scenarios is the distribution of ground stations across the Earth's surface. The constellation design and the ground station distribution can be simultaneously optimized toward selected objectives, such as to guarantee regional coverage over specific areas. Recently, del Portillo et al.[6] formulated the ground segment distribution problem as a down-selecting optimization problem, with the goal of identifying the minimum number N of ground stations that still offers optimal coverage of the constellation; given the large search space, the solution is obtained via Non-dominated Sorting Genetic Algorithm-II (NSGA-II). In another work by del Portillo et al.,[7] the authors identify availability of the network, latency of the network, and cost as the three drivers to derive Pareto-optimal sites for optical ground stations (OGSs) serving a P-LEO constellation. However, in both cases the authors define the optimal location with the aim of maximizing performance, without including the possibility of adversarial actions that may disrupt/affect service.

Strategic ground station positioning becomes relevant in the context of adversarial ground station transit time games. In fact, the impact of an attack may be proportional to the time such an attack remains undetected: once a satellite is compromised after its passage over an attacker ground station, the threat discovery time will be, at least, equal to the transit time between the attacker and a defender ground station (assuming there are no satellite interlinks nor additional in-orbit supporting networks). As a consequence, malevolent actors may leverage the longest possible transit time interval to select optimal locations to deliver an attack. However, solving the problem of finding optimal locations that would maximize the transit time between ground stations is non-trivial, especially within a large search space which involves multiple ground stations, a large number of satellites, and possibly geopolitical constraints. To cope with such intricate and high-dimensional solution spaces, meta-heuristic algorithms, such as coevolutionary ones, may be exploited to analyze adversarial ground station transit time games.

In this work, we model an attacker-defender adversarial scenario as a ground station transit time game. Within such a game, a malevolent actor selects positions on the globe that maximize detection time. After introducing a semi-analytical method to define the transit time in a simplified game model, we exploit evolutionary algorithms to explore adversarial ground station transit time games with increasingly higher-dimensional solution spaces: at first, we allow only the malevolent actor, i.e., the attacker, to optimize its ground station locations; then, we introduce the possibility for the defender to evolve the location of one defender ground station in addition to the fixed defender ground stations, thus defining a competitive coevolutionary scenario. These experiments set the foundation of a framework to solve transit time games that may be incrementally expanded to solve more realistic scenarios.

SINGLE SATELLITE, 1 ATTACKER VS 1 DEFENDER

Understanding attacker-defender dynamics within ground station transit time problems becomes increasingly difficult with larger numbers of satellites and ground stations. To gain insight on attacker-defender ground station transit time dynamics that may guide more complex investigations, we first consider one of the simplest possible scenarios. Specifically, we assume that the adversarial ground station transit time game only comprises one satellite placed in a polar orbit, and two ground stations, one representing an adversarial actor (i.e., the attacker) and one representing the satellite operator (i.e., the defender). A schematic of the problem is shown in Figure 1: the black antenna is the attacker, while the white antenna is the defender.

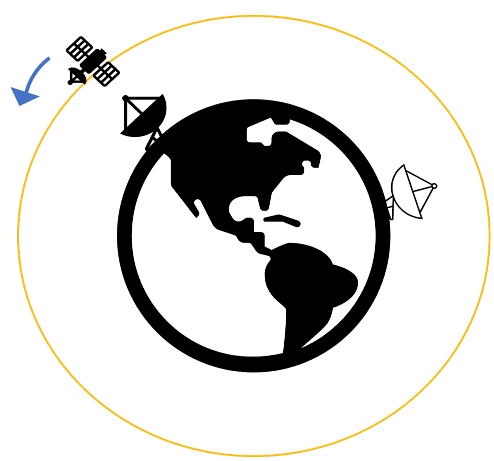

Figure 1: 1 Attacker vs. 1 Defender problem representation. Single satellite on polar orbit.

Assuming a fixed defender location, the attacker's goal is to identify the attack location, in terms of longitude and latitude, that maximizes the defender's time-to-discovery of an interference. We assume that the defender checks the status of the satellite as soon as it enters the corresponding access area, that is, the region where a line of sight exists between the satellite and the defender ground station, as defined by a given maximum central angle, λ_{max} (i.e., the radius of the small circle that defines the access area on an ideally spherical Earth's surface). The epoch when the defender discovers a satellite interference is denoted as detection time, t_d. Additionally, we introduce the following assumptions: 1) the game simulation starts with the satellite at the zenith of the attacker ground station; 2) the attacker compromises the satellite at the initial time, t_0; 3) the satellite travels along a circular orbit in the $x-z$ plane defined within a Cartesian reference system, moving northward with respect to the position of the attacker at t_0; and 4) the Earth is a perfect sphere. This problem is solved by determining the attacker position that maximizes the transit time of the satellite from the attacker to the defender station, $\Delta t = t_d - t_0$. The transit time, Δt, can be found by looking for the condition where the central angle defined by the sub-satellite point relative to the defender ground station is smaller than the maximum given central angle. The geometry of the 1 Attacker vs. 1 Defender problem is reported in Figure 2.

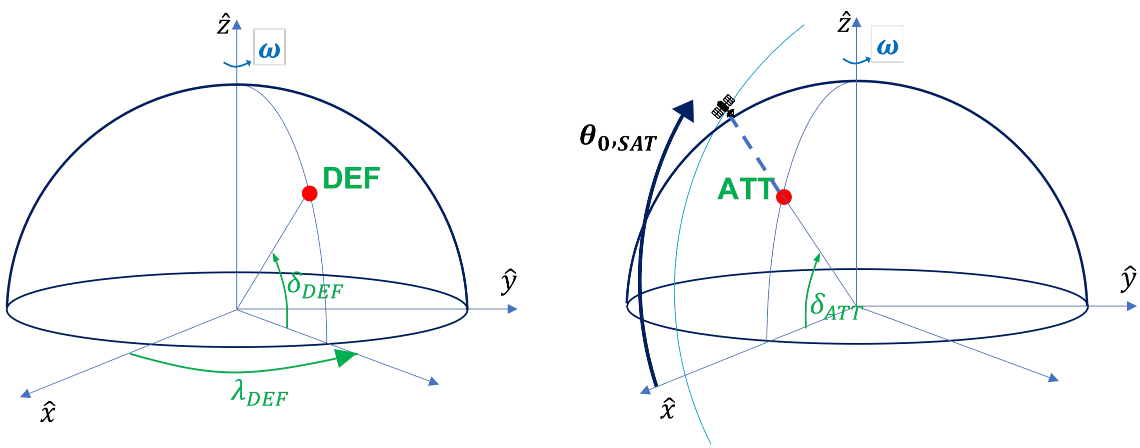

Figure 2: Double ground station problem representation

In Figure 2, *DEF* and *ATT* are, respectively, the defender and attacker ground stations, ω denotes the Earth's rotation rate, and λ_{DEF}, δ_{DEF} and δ_{ATT} are polar coordinate angles through which we define the stations' position; finally, $\theta_{0,SAT}$ represents the true anomaly of the satellite at t_0, assuming that the satellite is on the zenith of the attacker ground station at the initial time. From an orbital mechanics perspective, the satellite position can be described as a function of the Keplerian orbital parameters as:

$$\mathbf{r}_{SAT} = \mathbf{r}_{SAT}(a, e, i, \Omega, \omega, \theta) \tag{1}$$

Considering an ideal two-body motion, the true anomaly evolution along a circular orbit is described by the time law θ(t) = $\theta_{0,SAT}$ + nt, with n denoting the satellite mean motion. Given the spherical geometry of the model, a more suitable representation of satellite position can be expressed through spherical coordinates, these being a function of the polar angles previously introduced:

$$\mathbf{r}_{SAT}(t) = \begin{bmatrix} a\cos(\delta_{ATT} + nt) \\ 0 \\ a\sin(\delta_{ATT} + nt) \end{bmatrix} \tag{2}$$

where δ_{ATT} can be substituted with $\theta_{0,SAT}$ by noting $\theta_{0,SAT}$ = δ_{ATT} at t_0, . Analogously, the defender ground station position can be expressed through:

$$\mathbf{r}_{DEF}(t) = \begin{bmatrix} R_E \cos\delta_{DEF} \cos(\lambda_{DEF} + \omega t) \\ R_E \cos\delta_{DEF} \sin(\lambda_{DEF} + \omega t) \\ R_E \sin\delta_{DEF} \end{bmatrix} \tag{3}$$

with R_E as the Earth's radius. Now, assuming an orbital plane characterized by Ω = 0° and considering the satellite as being at the zenith of the attacker ground station at t_0, Equation (3) can be rearranged introducing the relative position of attacker and defender ground stations:

$$\mathbf{r}_{DEF}(t) = \begin{bmatrix} R_E \cos\delta_{DEF} \cos(\Delta\lambda + \omega t) \\ R_E \cos\delta_{DEF} \sin(\Delta\lambda + \omega t) \\ R_E \sin\delta_{DEF} \end{bmatrix} \tag{4}$$

Introducing a maximum central angle, λ_{max}, the condition for a transit through the defender access area is given by:

$$\cos(\alpha) \geq \cos\lambda_{max} \tag{5}$$

with α being the angle between the spacecraft position and defender's ground station vectors. Extracting α from the definition of scalar product between two vectors, development of Equation (5) leads to the following compact expression:

$$\begin{aligned} &A\cos(\Delta\lambda)\cos(\delta_{ATT}) - B\cos(\Delta\lambda)\sin(\delta_{ATT}) - C\sin(\Delta\lambda)\cos(\delta_{ATT}) + D\sin(\Delta\lambda)\sin(\delta_{ATT}) + .. \\ &.. + E\sin(\delta_{ATT}) + F\cos(\delta_{ATT}) = \cos\lambda_{max} \end{aligned} \tag{6}$$

In Equation (6) the coefficients A, B, C, D, E, and F are dependent on the defender's latitude, Earth's angular rate, satellite mean motion, and transit time, as given in Equation (7):

$$\begin{gathered} A = \cos(\delta_{DEF})\cos(\omega\bar{t})\cos(n\bar{t}) \\ B = \cos(\delta_{DEF})\cos(\omega\bar{t})\sin(n\bar{t}) \\ C = \cos(\delta_{DEF})\sin(\omega\bar{t})\cos(n\bar{t}) \\ D = \cos(\delta_{DEF})\sin(\omega\bar{t})\sin(n\bar{t}) \\ E = \sin(\delta_{DEF})\cos(n\bar{t}) \\ F = \sin(\delta_{DEF})\sin(n\bar{t}) \end{gathered} \tag{7}$$

Please note that the equality sign in Equation (6) is considered as a limit condition (i.e., the original Equation (5) contains an inequality sign, $\geq$). Additionally, it has to be remarked that Equation (6) refers to a polar orbit, $x - z$ plane-confined motion. In the general case of inclined orbit, at most four additional terms should be added to the expression in Equation (6).

Polar orbit scenario

Because Equation (6) is a transcendental function in all its variables, a closed-form solution may be non-trivial, if not impossible. Thus, we start analyzing Equation (6) by considering a limit case where the defender ground station is at one of the poles. When the defender's ground station is at one pole, the coefficients A, B, C, and D in Equation (6) become 0; then, Equation (6) can be explicitly solved for the transit time:

$$\bar{t} = \begin{cases} \frac{\arcsin(\cos\lambda_{MAX}) + \frac{\pi}{2} + \left(\frac{\pi}{2} - \delta_{ATT}\right)}{n} & \delta_{DEF} = -90^\circ \\ \\ \frac{\arcsin(\cos\lambda_{MAX}) - \delta_{ATT}}{n} & \delta_{DEF} = +90^\circ \end{cases} \tag{8}$$

To verify the validity of Equation (8), we performed a grid search for the transit time by discretizing the world map into 1°x1° grid cells using polar angle coordinates. While maintaining the defender ground station at the South Pole, we varied the attacker location across all cells of the grid, propagated the satellite motion (with an orbit altitude fixed at 200 km) until the satellite encountered the defender ground station, and recorded the corresponding transit time. By coloring each grid cell as a function of the corresponding transit time, we produced Figure 3. Figure 3 displays the transit time as a function of the attacker location in polar coordinates, assuming the defender ground station is at the South pole. By sampling attacker locations from the grid search map, we can compare the simulated transit time to the estimate from Equation (8) (which is evaluated for δ_{DEF} = -90°). The results of this comparison are listed in Table 1. The time discrepancy observable in Table 1 derives from the fact that a time step of 60 seconds is adopted in the numerical simulation.

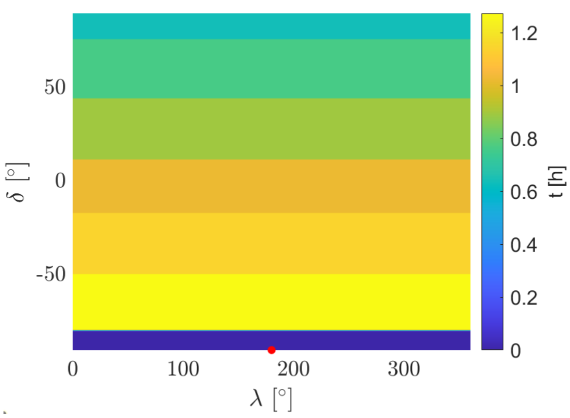

Figure 3: 1 Attacker vs 1 Defender transit time; edge case with defender at the South Pole.

δ_{ATT}	Simulation	Formula
90°	2520 s	2507.3 s
0°	3840 s	3834.7 s
30°	3420 s	3392.3 s
-30°	4320 s	4277.2 s

Table 1: Simulation and analytic formula comparison.

The grid search algorithm developed to validate Equation (6) may be employed to estimate the transit time under conditions for which a closed-form solution of Equation (6) may not be possible. For example, consider a defender's station located at mid-southern latitudes: the transit time is portrayed as a function of the attacker's location (expressed in polar coordinates angles) in Figure 4. In the surface plot in Figure 4, brighter colors indicate a longer transit time, and therefore locations that give advantage to the attacker. Not surprisingly, as the defender ground station is moved south (see Figure 5), the transit time rapidly reduces and the transit time pattern tends to that of the limit case in Figure 3.

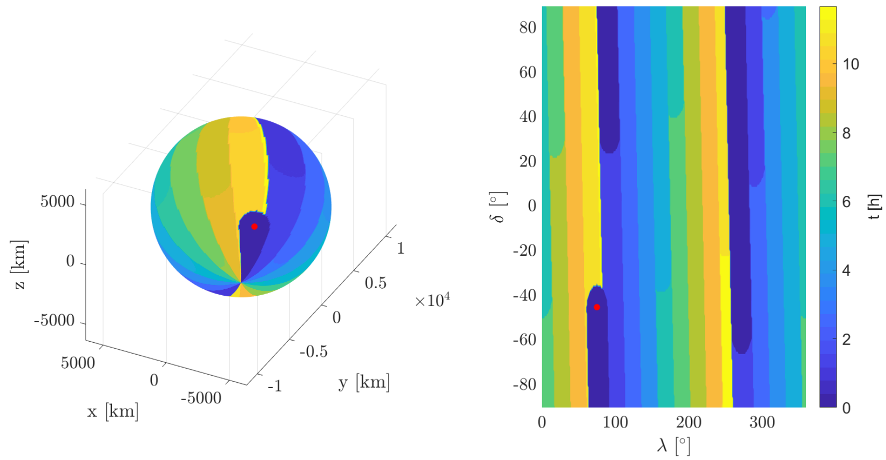

Figure 4: 1 Attacker vs. 1 Defender transit time; the red dot is the position of the defender. (*left*) spherical map, (*right*) equilateral projection.

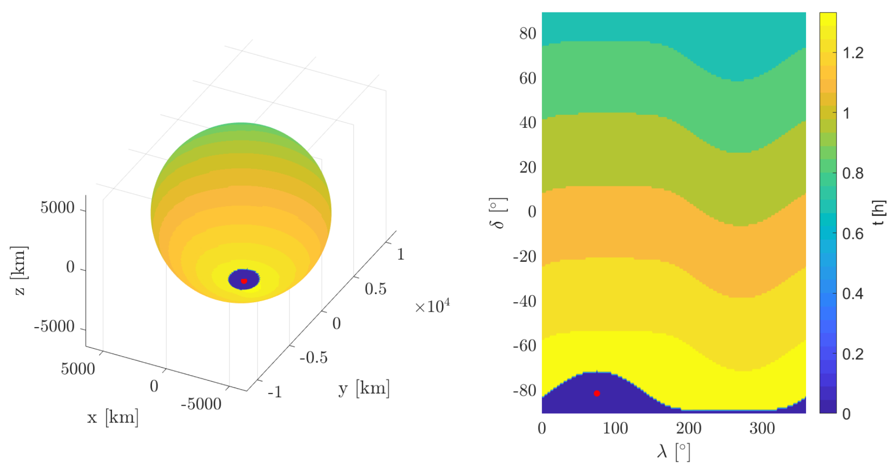

Figure 5: 1 Attacker vs. 1 Defender transit time; defender in polar region. (*left*) spherical map, (*right*) equilateral projection.

Toward problems with multiple ground stations and inclined orbits

Next, we conduct additional studies to understand whether potential, hidden threats exist within ground station transit time games, ones within adversarial games that do not have easy-to-engineer solutions. We start by considering the case of a non-polar orbit with multiple ground stations and a single satellite, which is solved via the grid search method introduced in the previous section. In this analysis, we assume an access area for the attacker with a maximum central angle of 2° and we consider as t_0 the instant when the satellite is in line of sight with the attacker; we no longer assume the satellite is necessarily moving northward with respect to the attacker at t_0. Figure 6 portrays the resulting transit time pattern for a game with two defender ground stations, a single attacker, and a satellite on a non-polar orbit (65° inclination). Looking at Figure 6, one can observe how advantageous locations for the attacker are blended with disadvantageous ones in a non-trivially predictable manner, emphasizing how a minor modification of the scenario parameters may lead to a major difference in the attacker-defender transit time problem solution. The potentially trivial solution shown in Figure 4, Figure 5 quickly becomes less predictable as multiple ground stations and inclined orbits are introduced. Thus, the rising complexity of even simplified problems points to greater challenges in the case of large constellations, necessitating more advanced techniques to understand and solve adversarial ground station transit time games. Evolutionary algorithms are particularly well-suited to find solutions in this type of complex, high-dimensional environment due to their stochastic/heuristic nature.

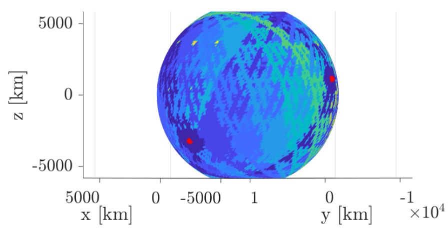

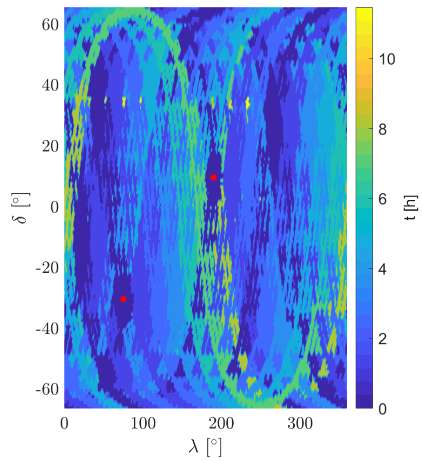

Figure 6: 1 Attacker vs. 2 Defenders; case with an inclined orbit. (*left*) spherical map, (*right*) equilateral projection.

INTRODUCING AN EVOLUTIONARY ALGORITHM APPROACH

Evolutionary algorithms[8] are a class of search algorithms inspired by the process of natural selection, which maintains a population of solution candidates that are iteratively improved through random mutations to their characteristics and sexual recombination of the traits of multiple parent solutions into child solutions. Members of the population in an evolutionary algorithm are evaluated based on a fitness function that assesses the quality of each solution at solving the intended problem. Such a fitness function could be a mathematical function, or a complex simulation that tests the behavior of a solution. Solutions with high fitness are more likely to be chosen to reproduce, and solutions with low fitness are more likely to be eliminated from the population, allowing the average solution quality in the population to increase over time. This allows exploration of the solution space without needing to rely on any analysis of the solution landscape that might be infeasible for complex problems. However, unlike local search algorithms, the genetic recombination operator enables evolutionary algorithms to search in multiple regions at once, and share useful traits throughout the population, which helps to prevent premature convergence to a local optimum.

Before introducing more complex scenarios, here intended as scenarios characterized by a large number of satellites and multiple orbital planes, we directly compare the previous analysis of attacker ground station positioning with the solution provided by our evolutionary algorithm. We evolve a population of latitude-longitude pairs representing attacker locations given a set of fixed defender ground stations, and evaluate these locations using a simulated satellite constellation. The location of each satellite and whether it is controlled by the attacker or defender is tracked at each timestep of the simulation. When a satellite passes within range of an attacker ground station, the attacker compromises it and gains control; when it passes within range of a defender ground station, the defender repairs it and regains control. The attacker's goal is overall degradation of the constellation by controlling as many satellites for as long as possible. The attacker's fitness is the number of satellites controlled per timestep, summed across all timesteps of the simulation, so evolution should enable the attacker to find ground station locations which best enable it to gain control of satellites which the defender cannot quickly retake. This corresponds to maximizing the average attacker-defender transit time for all satellites passing within line of sight of the attacker. Our simulated constellation now consists of 10 evenly-spaced polar orbital planes with 10 evenly-spaced satellites in each orbital plane. The simulation runs for 24 hours with a timestep of one minute; the defender fixed ground station positions (denoted through the letters A and B in Figure 7) are arbitrarily chosen. From the previous analysis, we realized that dominant solutions (for both attacker and defender) may arise in case of ground station located in polar region (for polar orbit); however, such dominant solutions may be unpractical, or, at least, come with high cost, as they would require physical access to polar regions. As a consequence, we impose boundaries on the northernmost and southernmost latitudes. For this example, the imposed constraint is latitude $\in$ [-70°,70°]. The parameters in Table 2 are evolutionary parameters, i.e. parameters commonly used in genetic programming. For a complete explanation and treatment, please refer to this reference.[8]

Initial population (μ)	Children (λ)	# of evaluations	Mutation rate	Mutation amount
200	300	3,200	50%	normal, σ = 1.8°

Table 2: Evolutionary parameters values used for the simulation.

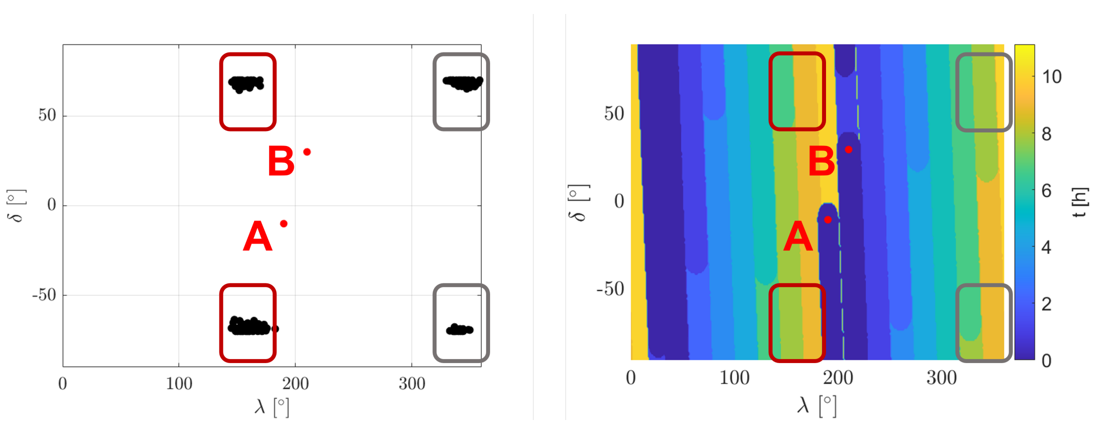

Figure 7: Evolved attacker ground station locations, with defender ground stations in red (*left*) and theoretical grid search results for the single satellite case (*right*). Colormap referred solely to right plot.

Figure 7 shows the evolved attacker's locations with respect to a certain set of defender ground stations; we can observe how the evolutionary algorithm was able to find multiple regions which allowed the attacker to effectively control the constellation. Comparing the two images in Figure 7, it can be observed how the clusters of advantageous regions for an attacker station location depicted in the left picture are in line with the brighter regions in the right plot, indicating that the results obtained by the grid search method and the evolutionary algorithm are consistent.

Competitive coevolution

In the scenarios described so far, only the attacker was allowed to evolve locations for its ground stations, while the defender was passive, i.e. its ground stations were arbitrarily assigned and fixed; however, moving toward more realistic scenarios, it is reasonable to assume some countermeasures from the defender. In our work, we initially model a countermeasure action by granting the defender the ability to evolve the location of an additional ground station (with respect to the already fixed, existing one(s)). It is here relevant to highlight that the additional location is not the result of a reactive behavior of the defender within the simulation time; instead, the additional ground station location is decided before the simulation as a result of evolution across previous generations and then added after 50% of the simulation time has passed, meaning that the defender is evolving a location rather than a strategy. This competitive scenario is game-theoretic in nature, and we solve it with competitive coevolution,[8] a special kind of evolutionary algorithm where the fitness of an individual is dependent on other individuals. Specifically, we employ a two-population competitive coevolutionary algorithm where one population contains attacker solutions and the other contains defender solutions. To assess the effectiveness of the method, we established a set of scenarios with increasing levels of complexity, which is reflected in the number of ground stations (both for the attacker and the defender). Twenty runs are executed for each of these scenarios for a short statistical comparison. A summary of the scenarios is reported in Table 3, while coevolutionary parameters[8] are reported in Table 4.

# of attacker(s)	# of defender(s)	# of orbital planes	# satellites per plane
1	2(+1)	10	10
2	2(+1)	10	10
2	25(+1)	10	10
5	25(+1)	10	10
10	25(+1)	10	10

Table 3: Simulation parameters.

Initial population (μ)	Children (λ)	# of evaluations	Mutation rate	Mutation amount
20	30	2,470	50%	normal, σ = 1.8°

Table 4: Coevolutionary parameters.

In the following, we describe two representative case studies: through the first one, we introduce the tools to interpret more complex scenarios, while in the second one we cast a more realistic application in our framework.

Case study: 2 vs 2(+1) Attacker-Defender

As complexity and dimensionality of the scenario grows, trivial solutions no longer exist and interpretability of the results provided by the algorithm may become challenging; to accommodate this, proper analysis tools must be introduced. In this scenario, we assume the defender to have two fixed ground stations with hand-chosen locations, while the attacker can evolve two different locations, each to place one ground station; additionally, we allow the defender to evolve a location for an additional ground station, thus giving a 2 vs. 2(+1) scenario. Ten equally-spaced, polar planes with ten satellites per orbit are considered in the set up of this experiment; in this case, the simulation duration length is set to 48 hours, with a timestep of 1 minute. Again, constraints are placed on the minimum and maximum latitude (latitude ∈ [-70°,70°]).

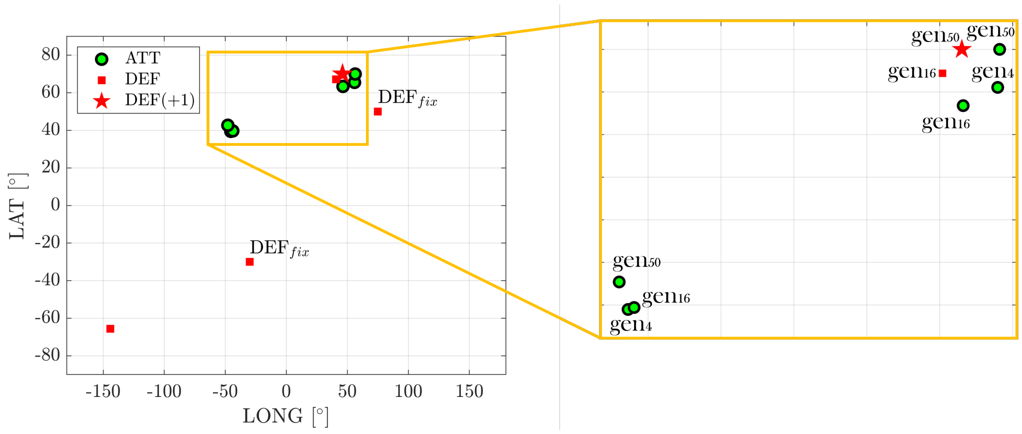

Figure 8: 2 attackers vs. 2(+1) defenders evolved locations. Hand-sampled generations to show how and how quickly evolution progresses toward the final configuration.

Figure 8 portrays the evolution of ground station location sites for both the attacker and defender for a few selected generations, including the last, 50th, generation; more specifically, the locations' coordinates are those associated with the best individual in each population in the corresponding generation. While one may attempt to interpret the figure to determine the quality of the locations discovered (either from the attacker's or the defender's perspective), Figure 8 alone is insufficient to characterize the quality of the solution and does not reveal if those locations are more advantageous for the defender or for the attacker. To this aim, we exploit a tool commonly used in the evolutionary framework: the Current Individual against Ancestral Opponents,[9] or CIAO, plot. Figure 9 displays the CIAO plot resultant from one of the several runs conducted for this scenario, viewed from the defender's perspective. The color characterizing each cell is representative of the fitness value defined in the implementation of the coevolutionary algorithm, i.e., the cumulative satellites uptime; in our case, such uptime represents the summation of the total time for which each satellite belonging to the constellation was under the defender control. Consequently, minimization (or maximization) of the uptime, governed by the evolution of attacker and defender's ground stations, reflect the maximization (or minimization) of satellites transit time among ground stations. More specifically, in the CIAO plot reported in Figure 9, the uptime value represented through the colormap is computed as the average of n best individuals within the corresponding generation (here, $n = 3$). The CIAO plot represents how the coevolution that led to the ground station positions in Figure 8 behaved across generations. Visually, good coevolution should appear in a CIAO plot through a smooth color gradient along diagonal direction and should be am monotonic as possible, without significant (sudden) changes along rows and columns. The actual position represented in Figure 8 is associated to the best individual belonging to the last defender generation. Finally a CIAO plot provides a visual representation of relative fitness, which along is insufficient to determine which player is leading the game; the determination of the leading player requires, instead, a direct comparison between the numerical values of the cumulative uptime and the total theoretical one.

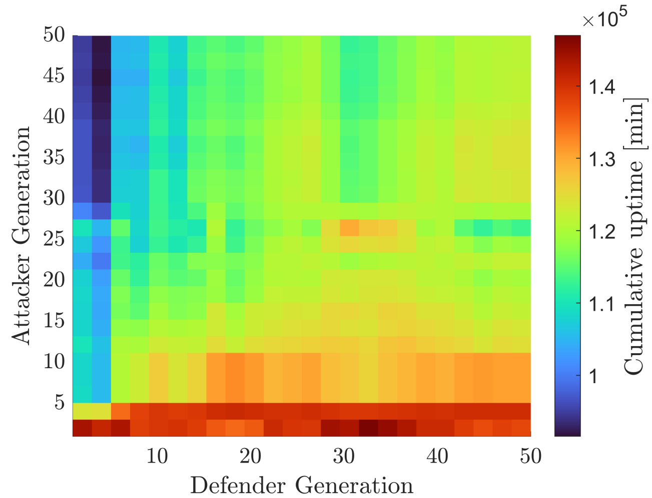

Figure 9: 2 attackers vs 2(+1) CIAO plot.

Grid search simulations are a complementary tool to facilitate a more direct assessment of the solution optimality. For the scenario under consideration, the cumulative uptime map obtained from a grid search simulation run is portrayed in Figure 10.

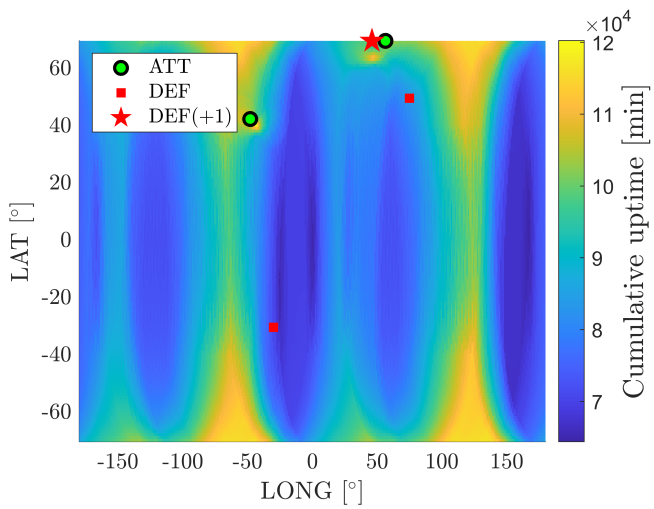

Figure 10: Grid search method visual representation.

For this grid search, longitude and latitude ranges were discretized into a series of 1° intervals; then, fixing the attacker location, a fictitious defender ground station was placed in a single cell and the total cumulative uptime is computed. This process was then applied to all the cells in the grid. Direct comparison of Figure 8 and Figure 10 shows how the defender evolved location falls into one of multiple sub-optimal solution.

Complementary to the ground station geographic location, we can also visualize what consequences those locations have in terms of coverage provided by the satellites. To determine the coverage map, we followed a similar process to the one described for the grid search method: at first, the longitude-latitude ranges were discretized with a 1° resolution; then, a fictitious ground station was placed in each cell and the 48 hours simulation was run. At each step, even if a single satellite was active over a specific cell, it counted as 100% coverage (represented as a 1): at the end, an average across the number of time steps was taken. Figure 11 shows how the addition of the defender ground station in that evolved location counters the solution evolved by the attacker, reducing the attacker effect on the satellites' operation. Please note that the uptime percentage reported in Figure 11 is evaluated with respect to the coverage computed before the presence of the attacker; as a consequence, an uptime of 100% means that the coverage condition is unchanged with respect to the same scenario without the attacker.

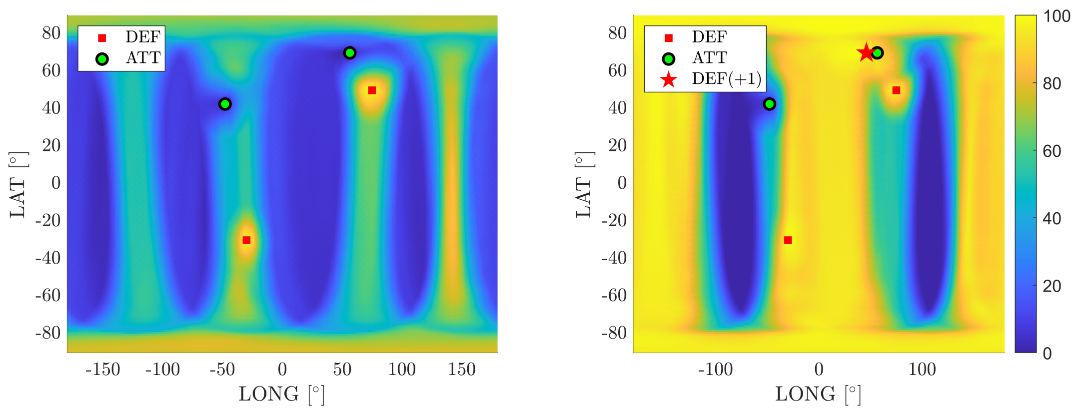

Figure 11: Satellite active regions without (*left*) and with (*right*) the +1 defender ground station.

Inspection of Figure 11 shows how, selecting that specific location, the defender is able to consistently gain back coverage over polar region, at the cost of leaving wide satellites downtime areas at lower latitudes.

Case study: KSAT

So far, we applied our framework to fictitious scenarios, i.e., scenarios with an arbitrary distribution of the ground stations, for both the defender and the attacker, without considering the actual geographic locations and geographic constraints (e.g., regardless of whether the ground stations were located on land or water). Next, we apply our coevolution framework to a more realistic ground station distribution by arbitrarily selecting the KSAT* ground segment. KSAT comprises 25 ground stations distributed around the world, including two polar sites. As the exact position of each ground station was not available, approximated GPS coordinates from the locations names available on the official website are adopted in our experiments. Three different scenarios are explored (see Table 3) with increasing number of attackers; for each case, 20 runs are conducted for statistical purposes. As for the previous examples, we introduce the constraint of latitude ∈ [-70°,70°]; this assumption assumes particular meaning in these scenarios as here we are actually considering an existing ground segment distribution. Additionally, the imposed latitude boundary is in line with

*https://www.ksat.no/ground-network-services/

KSAT ground stations locations, characterized by a ground station distribution within the imposed boundary, with the exception of two ”outlier stations”, respectively at approximately +78° and -72°.

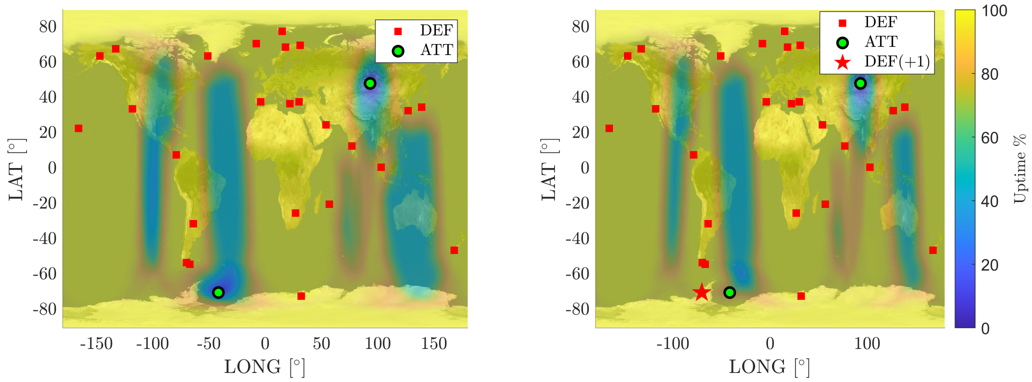

Figure 12: Scenario 2 Attackers vs 25(+1) Defenders. Coverage map before (*left*) and after (*right*) the addition of the defender ground station.

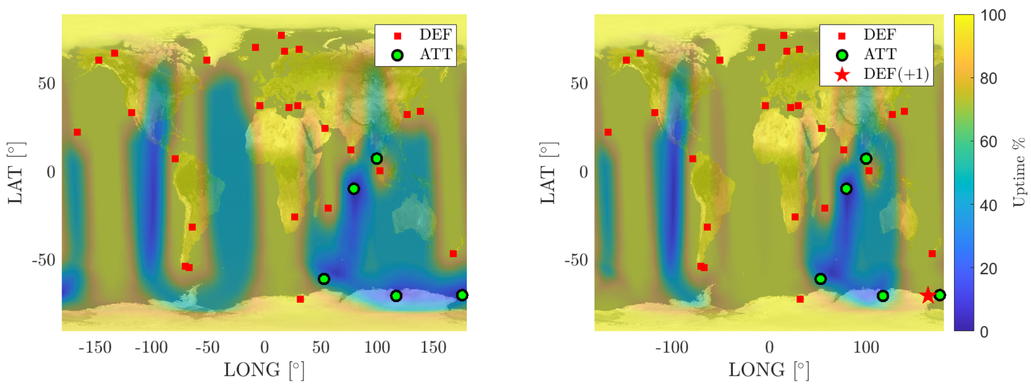

Figure 13: Scenario 5 Attackers vs 25(+1) Defenders. Coverage map before (*left*) and after (*right*) the addition of the defender ground station.

Figures 12,13,14 portray examples of uptime percentage, given relatively to a base uptime map for the constellation. The uptime percentage in Figures 12,13,14 is recorded with no attacker present, for scenarios with respectively 2, 5 and 10 attackers. As expected, the attacker effect on the satellites uptime is greater for a larger number of attackers: while the consequences on constellation performances are modest when solely two attacking ground stations are present, a significant reduction of the original coverage is achieved with five attacking ground stations. Concurrently, Figures 13,14 (right) demonstrate that the defender is able to gain back partial control of the satellites over extended regions with the addition of a single defending ground station situated at the evolved location.

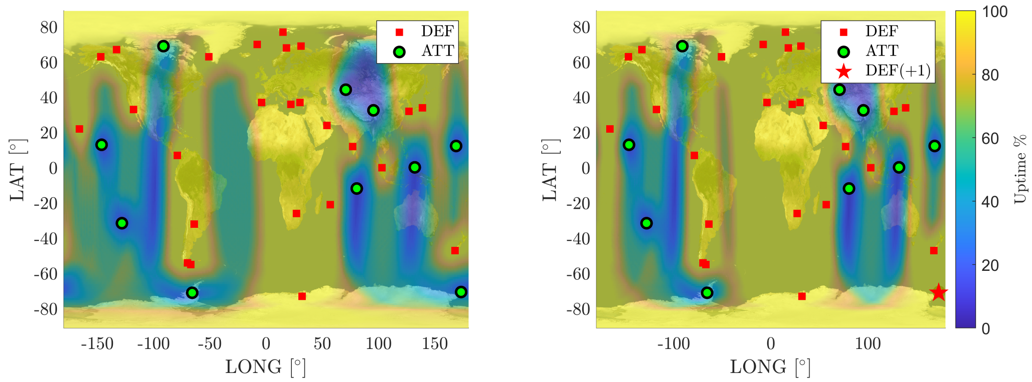

Figure 14: Scenario 10 Attackers vs 25(+1) Defenders. Coverage map before (*left*) and after (*right*) the addition of the defender ground station.

Averaged values across the 20 runs of each case study for the uptime percentage before and after the addition of the evolving defender ground station, together with the average relative improvements, are reported in Table 5. More specifically, the reported average uptime percentage are scaled considering the extension of the area of each 2D cell projected over a 3D sphere. We observe that, in the first two rows of Table 5, that are the scenarios where the number of attackers is comparable to the number of defenders, the presence of an attacker can strongly affect the average base coverage (reduction to an absolute average coverage of 14% for the 2vs2(+1) case, with a base absolute value of 55%). Additionally, the first two scenarios show the best average relative improvements, which was expected when the number of defenders and attackers is comparable. Within games with a comperable number of attacking and defending ground stations, the addition of a single ground station may introduce a important variation of the ultimate game equilibrium solution (in favor of one of the players). Looking at KSAT cases, the average uptime percentage before the addition of the ground station decreases together with the increase of number of attacker (which was, an expected behavior), while the addition of a defender ground station introduces an improvement comparatively smaller with respect to the other two scenarios.

Scenarios	Uptime % (before)	Uptime % (after)	Improvement
1 vs 2(+1)	57%	79%	22%
2 vs 2(+1)	26%	51%	24.8%
2vs25(+1)	85%	90%	4.4%
5vs25(+1)	75%	80%	4.9%
10vs25(+1)	65%	69%	4.9%

Table 5: Experiments results: average percentage uptime across runs for each scenario. Before and after referred to the additional defender ground station. Base absolute coverage is 55%.

DISCUSSION

The experiments presented in Sections and serves to gain the initial insight necessary to develop a robust and capable framework to solve transit time games and gain confidence that such a framework may be expanded to enable early discovery of hidden threats to P-LEO constellations.

From our study emerged that for scenarios involving satellites on polar orbit, the dominant solution for ground station placement tends to the poles or to the highest latitudes available. While this result was expected from simple geometric considerations (in fact, ground stations on the poles are guaranteed to have an access opportunity to each satellite in the constellation at least once every orbit), it is also reassuring that coevolution is able to consistently discover such a known solution. However, it is reasonable to expect that large constellations will comprise a combination of polar and non-polar orbit planes. Engineering considerations such as launch cost and coverage optimization considerations (e.g., optimal coverage of densely populated areas at mid-latitudes) may place a restriction to solely use polar orbits for large constellations. As demonstrated by our analysis, the introduction of inclined orbit may significantly complicate the solution space, which becomes highly tessellated geographically, with advantageous regions blended to poor-performance regions. In other words, less predictable game dynamics, thus threats, may exist when inclined orbit are introduced.

It is also worth noting that our analysis is agnostic to regional coverage and geopolitical constraints; under this assumption, the defender added station will improve coverage regardless of whether the interested areas are ground or water based. Similarly, as visible in KSAT case study, the attacker evolved locations are most likely to be water based or correspond to areas characterized by reduced/ difficult accessibility (where may be impractical to build ground stations). The dominant solution at the northernmost (or southernmost) latitude, also noted in all KSAT experiments for at least one attacking station, may also become suboptimal, or even unfeasible, when accessibility or geopolitical constrains are introduced. While our fitness function is currently based on time, without any spatial attribute, the introduction of geopolitical constraints or cost penalization (e.g., penalizing locations where it would be impractical to deploy a ground station) may yield to significantly different attacking and defending solutions.

Observations were also gathered to develop insight into coevolution. Across all the runs of all the scenarios, coevolution was, on average, well-behaved, meaning that the simulation set up (with all its assumptions) was well-posed; in particular, we noticed that the most compelling driver for a well-behaved coevolution was the maximum latitude range available to evolve a ground station location, which eliminated trivial, dominant solutions. In fact, experiments conducted for the same scenarios without latitude boundaries revealed that the solution for both attacker and defender evolved ground stations tended to rapidly converge toward polar regions.

To conclude, general considerations can be drawn about the computational cost of our experiments; considering that overall the runtime cost scales with the size of the constellation (both in terms of number of orbital satellites and number of ground stations) and with the number of time steps, to conduct our largest simulation (10vs25(+1) scenario) approximately 12 hours were required.

CONCLUDING REMARKS

Technology development, cost reduction, and the constant demand for high-quality service, are pushing the design and development of large satellite constellations; concurrently, the number of threats to ground and space assets is increasing as well, thus requiring the introduction of identification and mitigation strategies. In this work, we study potential threats in the form of a ground station transit time game formulation, which was analyzed at different levels of complexity: at first, we demonstrated with a simple example how the solution to optimal ground station location aimed to the maximization (or minimization, depending from the perspective) of the transit time can be found through a semi-analytical method; then, we underlined how modeling of more complex scenarios requires the introduction of more advanced techniques. Initially, we propose evolutionary algorithms as solution method to tackle this type of high-dimensional, potentially stochastic, environment; then, we modified the base scenario allowing also the defender to evolve a ground station location, thus introducing competitive coevolution as solution method. In the analysis of the solutions, we introduced different visualization tools to aid the interpretation of the results in terms of location, being geographic position alone insufficient to fully comprehend the solution proposed by the algorithm, especially with multiple ground stations, and many satellites configurations. Our current results demonstrate the effectiveness of the proposed strategy and open the possibility to adapt the selected evolutionary techniques to different formulations of the problem; future work will include the introduction of reactive behavior to the defender capabilities, thus allowing the defender to evolve a strategy, and the generalization of the orbital model to non-polar orbit configurations. Additionally, region-based objectives, cost penalization and geopolitical constraints will be introduced in the problem modeling, thus allowing for more realistic solutions.

ACKNOWLEDGMENT

The authors would like to thank Ms. Dhathri Somavarapu for the contribution in the development of the original constellation orbit model.

REFERENCES

[1] "Competing in Space," *NASIC Public Affairs Office, Wright-Patterson AFB, OH*, 2019.

[2] J. Pavur, D. Moser, V. Lenders, and I. Martinovic, "Secrets in the Sky: On Privacy and Infrastructure Security in DVB-S Satellite Broadband," *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*, Miami, Florida, 2019, p. 277–284.

[3] J. S. Turner, "A Methodology for Measuring Resilience in a Satellite-based Communication Network," *Wright-Patterson Air Force Base*, Ohio, 2014.

[4] K. M. Wagner, "Optimization of Disaggregated Space Systems Using the Disaggregated Integral Systems Optimization Technology Methodology," Master's Thesis, Blacksburg, VA, June 2020.

[5] F. Zaerpoor, E. Ghashghai, D. Thai, and J. Kim, "Measuring and Enhancing Resilience of Space Architectures: An Operationally Based Method," *AEROSPACE*, 2013.

[6] E. Crawley, I. d. Portillo, and B. Cameron, "Ground Segment Architectures for Large LEO Constellations with Feeder Links in EHF-bands," *2018 IEEE Aerospace Conference, Big Sky, MT*, March 2018.

[7] I. d. Portillo, M. Sanchez, B. Cameron, and E. Crawley, "Optimal Location of Optical Ground Stations to Serve LEO Spacecraft," *2017 IEEE Aerospace Conference*, , Big Sky, MT, March 2017.

[8] A. E. Eiben and J. E. Smith, "Introduction to Evolutionary Computing," *Springer*, 2003.

[9] D. Cliff and G. Miller, "Visualizing Co-evolution with CIAO Plots," *Artificial life*, Vol. 12, February 2006, pp. 199–202, DOI: 10.1162/106454606776073341.