DECENTRALIZED SELF-PROPAGATING GROUND DELAY FOR UTM: CAPITALIZING ON DOMINO EFFECT

Leonid Sedov and Valentin Polishchuk
Communications and Transport Systems, ITN, Linköping University, Norrköping, Sweden
Vishwanath Bulusu
Cal Unmanned Lab, CEE Department, University of California, Berkeley, California, USA

Abstract

We investigate a possible scheme for management of conflicts among autonomous unmanned aerial vehicles (UAVs) in high-density very low level (VLL) uncontrolled airspace. The drones are modeled as disks of a given radius, moving along prescribed trajectories planned without any centralized coordination. Thus, during the motion, the disks may potentially come into contact, which represents loss of separation between the drones. Two overlapping disks get enclosed by a larger disk serving as the protected zone for avoidance maneuvers of all the drones inside it. When the conflict is gone, the disk is deactivated and the UAVs continue towards their destinations. We simulate traffic demand and the evolution of the deconfliction zones over a geographic area and present statistics associated with functioning of the system with and without ground delay to avoid take offs into conflicts. The scheme shows promise and is a good approach to explore further in future work.

Introduction

We investigate a simple ground delay based scheme for autonomous Unmanned Aircraft Systems (UAS) Traffic Management (UTM) in very low level (VLL) uncontrolled airspace. The UAS, referred to synonymously as *drones*, are modeled as disks of a given radius r, analogous to the protected airspace zone around manned aircraft [1]; we color the disks red signifying the fact that a disk is a no-go zone for other disks. We seek to understand whether using *local* information is sufficient to maintain a *global* stability of the system, laying ground for studying decentralized UTM frameworks. The red disks model the extent of local uncertainty, and our control to maintain stability is ground delay. Our premise is that the disks should stay pairwise-disjoint to account for the vehicles' navigation and location uncertainties. Occurrences of disk overlaps therefore become associated to stability of the system.

Ground delay has been extensively researched in manned Air Traffic Management (ATM). It is also being currently explored for UTM [2], albeit with a more centralized approach. The Literature Review section below discusses this in further depth providing the background for choosing a ground delay scheme and trying it with a decentralized approach.

We assume that the drones move along prescribed trajectories, which are planned without any centralized coordination; in particular, the (desired) takeoff times are uncoordinated among the different UAS. Thus, during the motion, the disks may potentially come into contact. This represents a loss of separation between the drones. Our protocol aims at reserving airspace for managing such conflicts and possibly at decreasing the traffic density by disallowing a drone to take off into areas of ongoing deconfliction activities of other drones.

The high-level idea of our scheme is simple: whenever two disks touch, they are enclosed in a larger red disk. The newly formed disk is the conflict resolution zone for the enclosed drones, and so the other disks should stay away from it. Naturally, two de-confliction zones should be disjoint. However when the traffic density is high, it is possible that two red disks start to overlap, either due to a larger red disk being hit by a single original radius-r disk that moves along its projected path, or due to two larger disks covering a common point. This implies, according to our protocol, that the drones in the two overlapping disks get enclosed by a yet larger red disk serving as the protected zone for avoidance maneuvers of all the drones inside it. When the conflict is gone, the disk is deactivated and the UAS continue towards their destinations.

In this paper we study the evolution and characteristics of the above-described system with ground delay as a management tool embedded into it. To implement the ground delay we mandate that no new takeoff is allowed from any point covered by the red disks. We simulate the traffic demand and the evolution of the red zones over Norrköping municipality in Sweden and measure several statistics associated with the functioning of the system for a range of parameters – the disk radius r and the traffic demand for the UAS operations. The algorithm and simulation set up are described in the Simulation section.

The simulation results are presented and discussed under Results section. We give statistics of when and how the conflicts (red zones) move from local spots to becoming a global (area-wide) phenomenon and quantify the experienced ground delay (including comparison of the system performance with and without the ground holding). A short video with excerpt of the simulation is available at http://tiny.cc/tyzfiy (for clarity, only the boundaries of larger disks are shown, i.e., the circles).

In the Conclusion section we discuss how our work contributes to understanding the automation levels feasible for UTM and to quantifying the influence of autonomy on UTM success.

We emphasize that in this work we do not go into specifics of conflict detection and resolution (CDR) approaches. The red disks just signify that space is reserved for it and how it will be done is left as a separate research direction.

Literature Review

The UAS industry envisages services across a broad spectrum of applications such as package delivery, agriculture, sensing and mapping and so on. To enable these operations, autonomous air traffic management strategies have become an important area of research spearheaded by NASA's UTM[3] program.

Inspired by the existing manned ATM architecture, a mainstream UTM designer's dream is some kind of a central system to guarantee global safety. However, due to the highly dynamic and stochastic nature of the above mentioned proposed UAS operations, where "every home will have a drone and every home will serve as an aerodrome" [4], this could be a potential overkill. A more robust distributed approach might be sufficient and better suited.

In fact, the (centralized-by-design) ground delay program (GDP) is being successfully challenged also in ATM – the idea is that airborne delay may be accumulated by speed control so as to make the holding decisions closer to the actual restricted-capacity zone (if the no-go restriction is still there by the time the aircraft reaches it) [5]–[7]: the logic, taken to the extreme, is of the kind "Why waiting on the ground, if the congestion/weather is few hours ahead and it is possible to simply fly with reduced speed, but get back on track if the route clears up".

A centralized ground delay approach for unmanned traffic [2] allows only conflict-free flight plans by checking each plan in space and time. Practically, this depends on advanced planning and approval. In the early days of the industry, this might be sufficient but in the future a large number of flight trips will be generated and conducted in a more real-time manner analogous to today's road vehicular traffic. Hence, a more decentralized plan-as-you-go strategy becomes worth exploring.

De-centralization of UTM may occur on several levels. At the highest level, there stands the business model with several (competing) UTM service providers (as is the case, e.g., with mobile services) [4]. Then, there will be multiple operators, each responsible for a set of drones. Next, on few other possible levels of the hierarchy, there may be many UAS that could be leased or time-shared between the many users. Our focus is on the lowest level: vehicle-to-vehicle (V2V) interaction involving communication and navigation (including sense and avoid).

We consider the UAS ground delay strategy in a more de-centralized fashion than its manned counterpart. Instead of a central authority imposing delays, a UAS delays its own takeoff in real time based on safety information that it can source locally. The deconfliction is also fully distributed – no central system is telling the drones how to avoid each other.

In the system we propose, the lack of safety is conveyed by a red disk overhead. Further, enroute loss of safety and rise of multi-vehicle deconfliction problems are also modeled by red disk overlaps. However, we are not concerned with the specifics of implementation of deconfliction schemes – they may vary depending on the infrastructure in use (cell

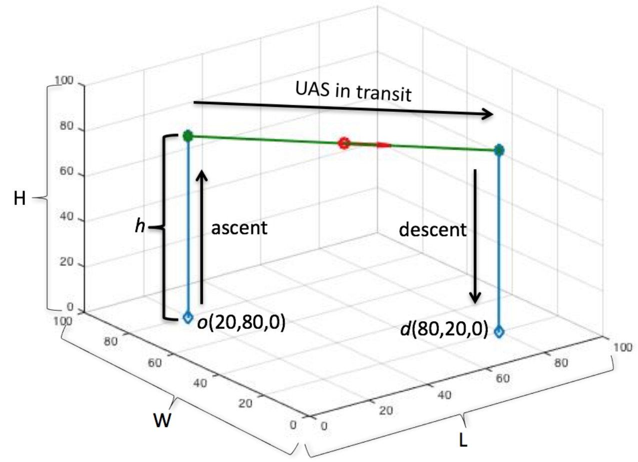

Figure 1. A typical UAS flight path in the Cal model

towers, GPS, WiFi, radio channels, etc.), technology advances (5G, free-space optics) and other details of UTM-to-be.

Simulation

We use the so called Cal model which was introduced in [8]. The airspace is represented with a rectangular volume LWH. The drones takeoff and land strictly vertically, and fly on a fixed level h. Figure 1 illustrates a typical flight. All UAVs are at the same level because with the restrictions on commercial UAS operations [9], there is little room for multiple levels. The setup is hence essentially two-dimensional.

The origins and destinations for the flights are generated randomly based on the population density over the rectangular area. The total number N of flights expected during the day is specified, and the intensity of the traffic starting or ending at a point p of the domain is proportional to the population density at p (the starting times of the flights originating from p form a Poisson process with the rate proportional to the density). The simulations were run for Norrköping municipality in Sweden; the population density is shown in Figure 2. For each setting of the parameters r and N, we simulated 12 hours of traffic.

Algorithm

We maintain the disks (both the original drones and the larger red deconfliction zone disks) in a forest of binary trees. The leaves of the trees are the original radius-r disks identified with the UAVs. Any node of the tree is the disk that contains all the disks in its subtree. The size of the conflict (number of covered UAVs) handled in each node is equal to the number of leaves in its subtree. When two disks touch, they form the larger disk. This disk is the root node in a new tree, and its children are the overlapping disks. We assume, without loss of generality, that the disks start touching each other on-by-one, i.e., that no two pairs of disks start touching at the same time.

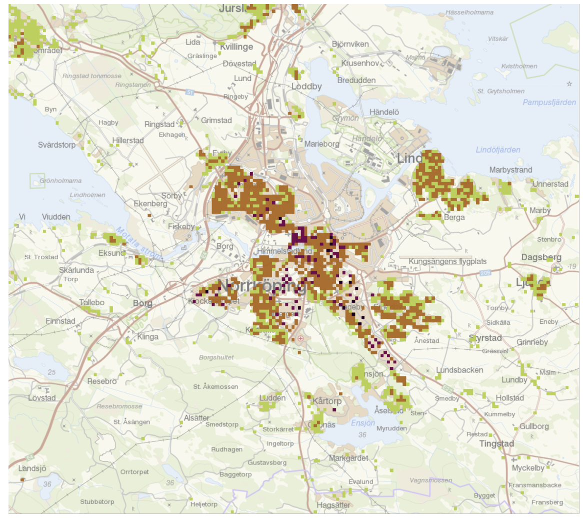

Figure 2. Population Density Map for Norrköping municipality [10].

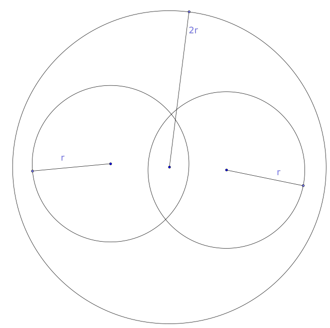

Figure 3. Enclosing disk around 2 UAVs

Figure 3 shows two conflicting UAVs and the enclosing disk for them. Figure 4 shows the compound conflict case.

As soon as the children of a node stop overlapping, the edges from the parent node are removed. Then recursively all parent nodes are removed. The sub-trees of the original tree, which became disjoint after nodes elimination, may intersect each other. We check them for intersection and unite the overlapping disks to form new trees.

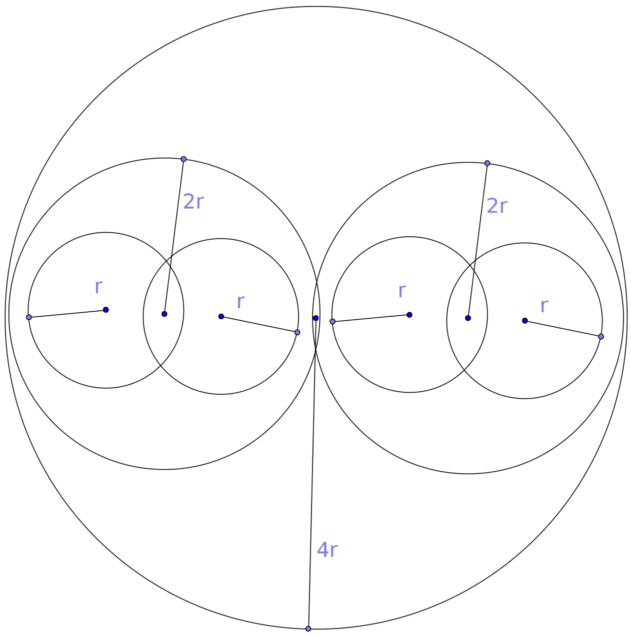

Figure 4. Enclosing disk for the 2 enclosing disks

Figure 5 shows an example of the disks and the forest evolution.

We coded an interactive playground for the disks formation and disappearance; the reader is invited to check it at https://undefiened.github.io/red_prototype/. Initially, the UAV disks are randomly distributed in the area. As per our protocol, if two disks overlap, they are enclosed in a larger red disk; if two red disks overlap, they are enclosed further, and so on. In the playground, the disks *do not move* by themselves; instead, a user can move the UAV disks around and see how the red disks evolve.

Results

The results presented in this section were obtained from simulating 12 hours of traffic for each setting of the parameters N (the number of flights per day) and r. The number N was varied from 10 to 100000, and r – from 2.5 to 25m. We considered two models. In the first one (the model with the ground delay), the drones were not allowed to takeoff if their start location was covered by a red disk; i.e., the drones were not allowed to takeoff into a conflict and had to wait until the red disk above them disappears. In the second model, taking off into a conflict was allowed.

Figure 6 shows average ground delay time (of course, it is applicable only to the model *with* the ground delay; in the second model, the delay is 0). The time is calculated using only the delayed UAVs, in order to report the delay time only for the drones which were actually delayed. As expected, the delay grows with r. Interestingly, N has low influence on the delay. It can be seen that overall, the delay is small.

Figure 7 shows percentage of the UAVs which were delayed (in the model with the ground delay). Equivalently, this is the percentages of UAVs taking off into conflict in the second model. We observe a maximum of 8% delayed traffic at the highest traffic density and largest protected disk radius. Overall, the growth is superlinear in radius for a given density, but approximately at or below linear with respect to N for a given radius.

Figure 8 shows maximum conflict size as a function of N and r during the whole simulation time. With a very slow growth at the low N and r, maximum conflict size swiftly rises when traffic becomes dense. The largest conflict size is similar to [8] but at much lower protected area radius.

Figure 9 shows percentage of area covered by disks on average as a function of N and r.

Next we present few distributions built for the largest parameter values $N = 100000$ and $r = 25$m for the second model (when allowing takeoffs into conflicts).

Figure 10 shows what percent of the UAVs spent different percentage of time in a conflict. It can be seen that two thirds of drones spend less than 10% of their flight time in conflict. Thus, efficiency of the operations is not compromised much.

Figure 11 shows the percent of time spent, on average, in conflicts of different sizes. The drones spend in size 10 or larger conflicts less than 1% of time.

Figure 12 shows percentage of flights taking off into conflict of different sizes. UAVs taking off into free space are not taken into account (this measurement is presented, implicitly, in Figure 7 which shows the complementary measure – UAVs taking off into conflict). Interestingly, most conflicting takeoffs are into small-size conflicts: even though large conflicts cover more area, they are quite rare.

Finally, Figure 13 shows percentage of the time during which different percentage of the area is covered by the red disks. (The average covered area, for

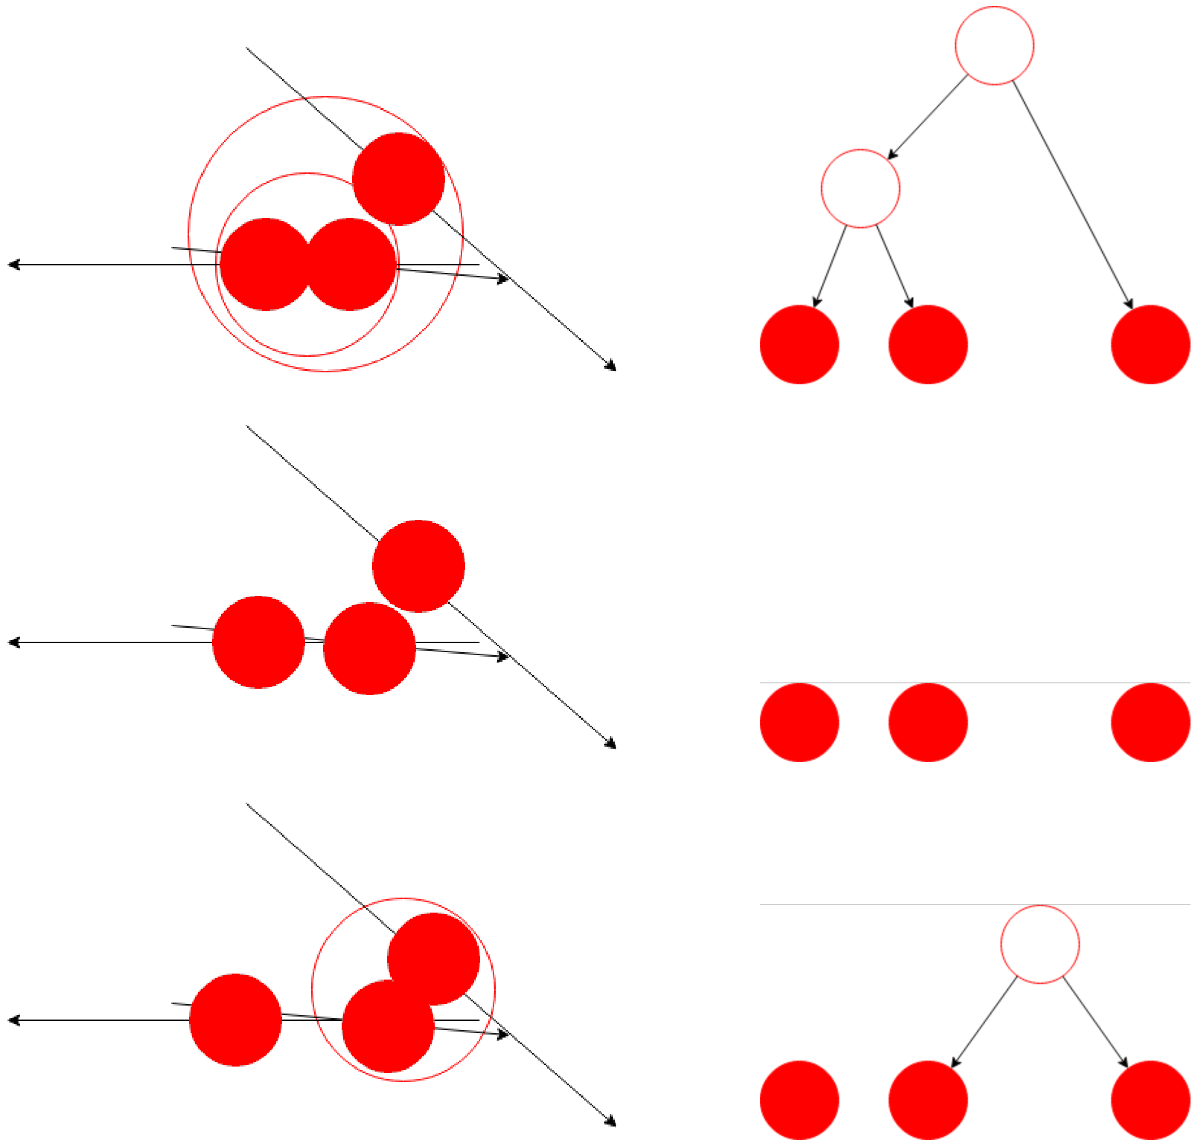

Figure 5. Left: UAVs. Right: the trees

different N and r is shown in Figure 9.)

Conclusion

UTM principles are laid out in NASA ConOps paper [2, Section III.C.2] (which, in turn, resemble Isaac Asimov's [11] Laws of Robotics). Principle 4 essentially states that all actors need to know everything. Still, it may be the case that having only *local* information is sufficient for the traffic management, provided the information spreads around when needed. We believe that more investigation is warranted to estimate how much information should really be handled by the central "know-all" system. Our paper is a step in this direction.

Indeed, on the one hand, in an envisioned centralized UTM each flight plan will be checked for potential conflicts with the already-existing flights – this is similar to the celebrated Ground Delay Program in the conventional ATM (which, by design, requires global knowledge and central authority to impose delays). On the other hand, due to the highly dynamic and stochastic nature of UAS traffic, distributed management is also often mentioned in UTM literature, e.g., in the NASA ConOps paper [2]. If for nothing else, a distributed architecture could work as

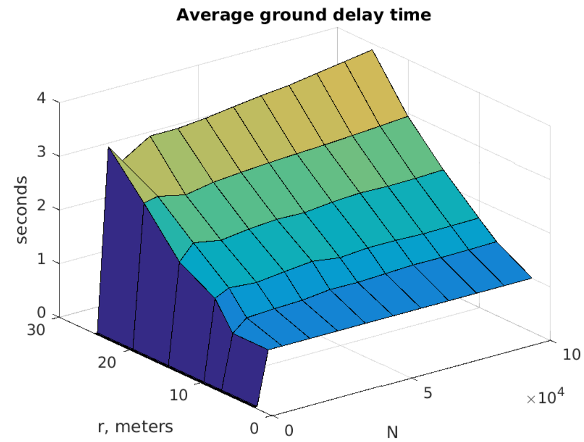

Figure 6. Average ground delay as a function of N and r.

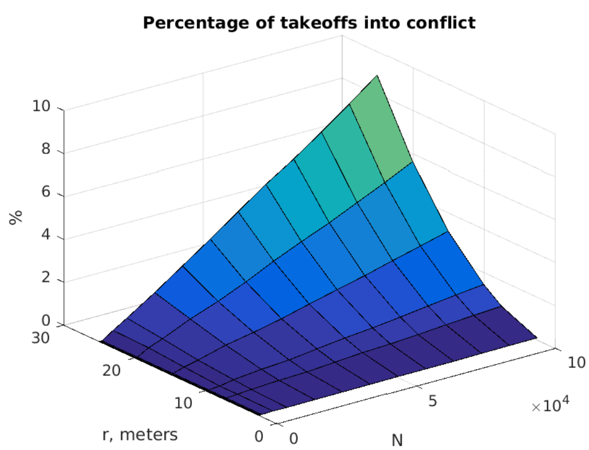

Figure 7. Percentage of UAVs that had a conflict above at the departure time as a function of N and r.

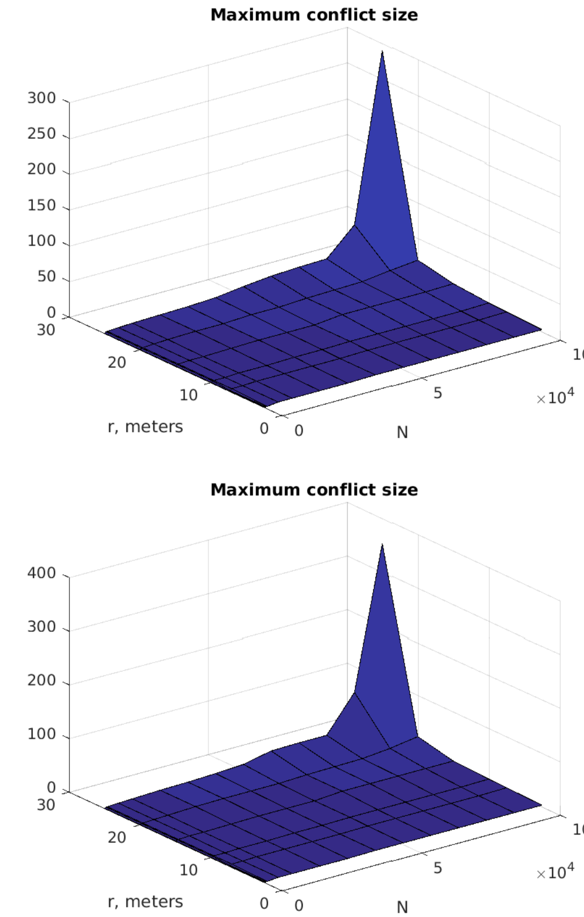

Figure 8. Maximum conflict size as a function of N and r. Top: ground delay model. Bottom: takeoff-into-conflict model

fallback solution for information propagation, providing a contingency plan for the case that the centralized UTM is unavailable or has prohibitively large latency. As one example (Principle 5 from the ConOps paper [2]) consider emergency UAS, prioritized by the UTM over the other drones: a simulation, similar to ours, will help to grasp the ripple effect of public safety UAV plowing through the normally operating drones soup.

More generally, our work fits into the mantras of "Flexibility where possible, structure where necessary" [2] and "When technology is right, regulation is light" [12]. For instance, it may be beneficial to not overload the UTM with local conflict avoidance: if technology is good (V2V communication and onboard CD&R capabilities are strong), there may be less need to provide a centralized mechanism that would precisely tell every drone how to avoid the collisions.

Acknowledgments

The work of LS and VP is part of UTMOK and UTM50 projects supported by the Swedish Transport Administration (Trafikverket) via the Swedish Air Navigation Service provider LFV (Luftfartsverket).

References

[1] J. Krozel, M. Peters, and K. Bilimoria, "A decentralized control strategy for distributed air/ground traffic separation," 2000.

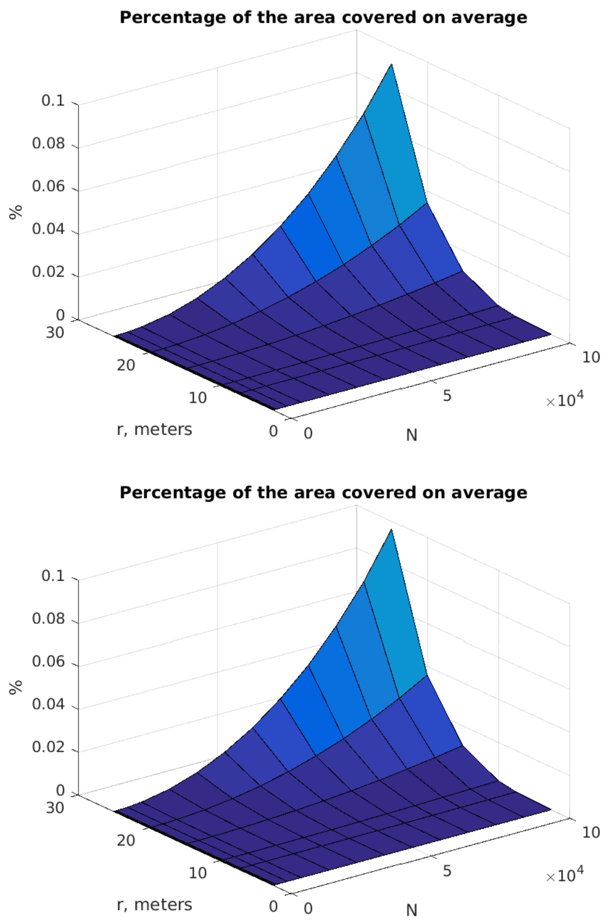

Figure 9. Percentage of the area covered on average as a function of N and r. Top: ground delay model. Bottom: takeoff-into-conflict model

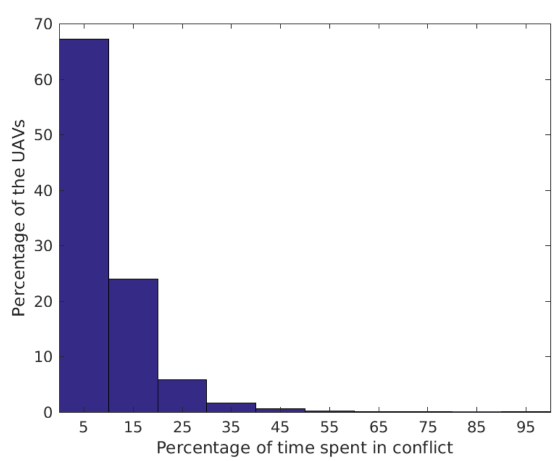

Figure 10. Histogram of the percentage of time spent in conflict

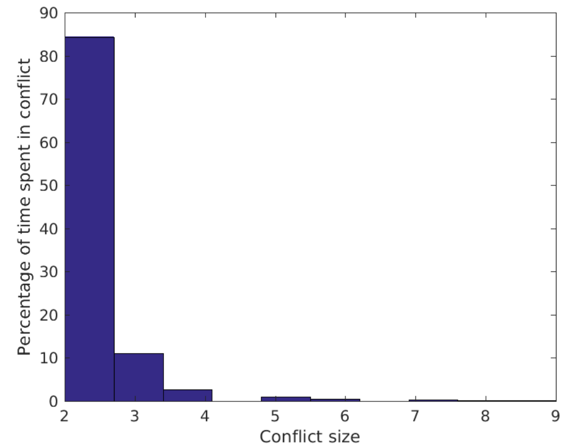

Figure 11. Histogram of percentage of time spent in different-size conflicts

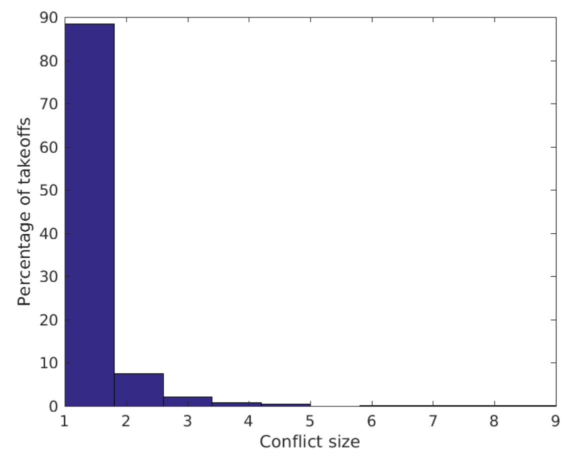

Figure 12. Percentage of takeoffs into different-size conflicts

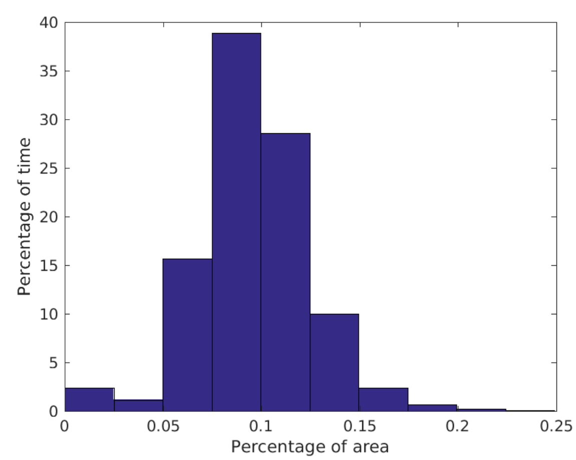

Figure 13. Percentage of time when different percentage of the area is covered by conflicts

[2] T. Prevot, J. Rios, P. Kopardekar, J. E. Robinson III, M. Johnson, and J. Jung, "Uas traffic management (utm) concept of operations to safely enable low altitude flight operations," in *16th AIAA Aviation Technology, Integration, and Operations Conference*, 2016, p. 3292.

[3] P. Kopaderkar, "Safely enabling uas operations in low-altitude airspace," in *IEEE 35TH Digital Avionics Systems Conference DASC*, Plenary talk, 2016.

[4] P. Kopardekar, *Unmanned aerial system traffic management system*, Talks at Google, 2016.

[5] L. Delgado, X. Prats, and B. Sridhar, "Cruise speed reduction for ground delay programs: A case study for san francisco international airport arrivals," *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 83–96, 2013.

[6] L. Delgado and X. Prats, "Effect of radii of exemption on ground delay programs with operating cost based cruise speed reduction," in *Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013)*, 2013.

[7] Y. Xu, R. Dalmau, and X. Prats, "Effects of speed reduction in climb, cruise and descent phases to generate linear holding at no extra fuel cost," in *International Conference on Research in Air Transportation ICRAT*, 2016.

[8] V. Bulusu, R. Sengupta, and Z. Liu, "Unmanned aviation: To be free or not to be free? a complexity based approach," in *7th International Conference on Research in Air Transportation*, 2016.

[9] *Federal Aviation Administration - Unmanned Aircraft Systems*, https://www.faa.gov/uas/, 2016.

[10] *Swedish University of Agricultural Sciences*, http://www.slu.se/en/, 2016.

[11] I. Asimov, *I, robot*. New York: Doubleday & Company, 1950.

[12] *The development of the uas traffic management (utm): An air navigation services perspective*.