Toward decentralised consensus and offloading for area coverage in a fleet of drones

Hanna Kavalionak, Emanuele Carlini, Pietro Cassarà, and Carlo Meghini

Istituto di Scienza e Tecnologie dell'Informazione (ISTI)
Consiglio Nazionale delle Ricerche (CNR)
<name.surname>@isti.cnr.it

Abstract. A precise and dynamic visual coverage of a given area is an essential task in many smart contexts, ranging from civil communities to military applications. Due to the last years advancement in hardware miniaturization and efficiency, area coverage is often performed with a combination of static and moving devices, such as unmanned aerial vehicles (drones). Drones are useful to cope with the highly unpredictability and dynamicity of environments, but require specific and efficient solutions toward and efficient area coverage. In this paper we proposes an initial work toward a drone-based approach for the task of area coverage. In particular, we focus our analysis on the following points: (i) decentralized consensus for movement planning, and (ii) the integration of cloud computing infrastructures and technologies for computation offloading, both for image analysis and movement planning.

Key words: Decentralized consensus, drones, distributed tracking, dynamic environments

1 Introduction

The active monitoring of a geographical area through sensors is a fundamental and widespread aspect for a wide range of applications including private surveillance, crowd tracking, and public security. Active monitoring can be exploited in a number of both military and civil applications, such as surveillance of national borders to control immigration/emigration or controlling the flow of tourists in large cities. A core aspect of active monitoring is represented by the task of *area coverage*, i.e. the ability to place sensors in order cover the considered area in an optimal way. Since in many applications the condition of the monitored area can change abruptly, the best way to place sensors also change over time. An infrastructure made of static sensors can be not enough to cope with unexpected events that can result from the inherent unpredictability of crowd behaviour and the environment, such as for example a broken camera or unexpected visual obstacles. In addition, ground sensors takes time to be installed, and therefore cannot be deployed in an unexpected situation if not foreseen in advance. Also, the monetary investment for the monitoring of a single event can not be justified in certain scenarios (e.g. research activities).

Therefore, a crucial aspect is the degree of adaptability that sensors (e.g. cameras, temperature, sound, etc..) are able to exploit. In order to cope with highly unpredictability and dynamicity of coverage activities, recently several approaches exploits Unmanned Aerial Vehicles (UAVs, informally known as *drones*) as a valid option to carry out many different kind of monitoring [15]. The technological advancements of UAVs rapidly increased in the recent years mostly due to military reasons, nevertheless most of the technology is also available to civil and research purposes. Drones can be deployed to different locations on demand, with a very short notice and without requiring a dedicated static infrastructure placed beforehand. Their behaviour can be reprogrammed while in mission, making them suitable to adapt to fast and unpredictable events within the same mission.

In this paper we consider the challenge of an *active area coverage* by means of a fleet of UAVs. For the purpose of this paper, we assumes UAVs are equipped with means to communicate to each other and with the sensors necessary for their mission. The usage of a fleet of drones presents multiple benefit when compared to a single UAV in terms of: (i) size of the coverable area, (ii) duration of coverage, (iii) coverage redundancy. However, a careful orchestration is required in order to accurately plan the movements of UAVs in order to obtain the best coverage possible. A straightforward way to organize the movement of UAVs is to employ a centralized entity (e.g. a server) that continuously collects their position and generate the new movement plan. However, the effectiveness of this solution is limited, as it suffers in terms of robustness (what if the server crashes?) and scalability (the frequency and the number of communication can saturate the server, which is not able to produce the new movement plans in time). Therefore, in we advocate a decentralized and distributed approach, in which the drones self-organize their movement toward an effective active area coverage.

Besides the self organization of the fleet, we also draw several considerations in relation to the utilization of cloud computing technologies and mechanisms for the image analysis activity of the drones. In particular, we consider the case in which specific image analysis techniques are too computationally heavy to be executed by drones (e.g due to battery constraint) and it would be more effective to communicate the images to be analyzed to a cloud server which can return back the results of the computation. We analyze this approach according to the work already done in the field of mobile cloud, in which mobile devices (typically smarthphones) offload their computation to nearby cloud computation units (cloudlet).

The paper is organized as the following. Section 3 presents the reference architecture and the envisioned scenario considered in the paper. Section 4 discusses a preliminary model for the area coverage and its exploitation in a best-of-n problem formulation. Section 5 elaborates the possible integration of cloud computing technologies for the offloading of computation in the considered scenario. Finally, Section 6 concludes the paper.

2 Related Work

The task of self-organize the movements of a number of entities in a decentralise fashion is not new, and it might be referred to as *flocking* [11]. Generally, such task can be abstract as a specialized version of distributed consensus, and it has been tackled in many research fields, although often with different nomenclature and purpose. For example, in the field of multi-agent systems, holonic systems define an organizational model of agents based on self-similarity into "super-agents" that are seen as single agents from the outside [7]. In peer-to-peer many approaches rely on self-organization techniques to organize the peers of the network in overlay for multiple purposes, such as area coverage [3] or in order to estimate a distribution of network parameters [12].

Recently, decentralized and peer-to-peer flocking algorithms have been applied to drone networks with the aim of self-organizing a fleet of drone toward the completion of complex task. The literature about this topic is very vast; hereby we provide several pointers to recent results. For example, Vásárhelyi et al. [17] provides an algorithm solution based on short-term repulsion and long-range attraction of drones, and it is validated via a numerical simulation. Another recent approach, Yuan et al. [18] proposed a decentralized model predictive control (DMPC) flocking algorithm to self-organize the movement planning of drones, by using the XBee communication technology.

Several recent works specifically dealt with some version of the problem of area coverage and in specific application domains. For example, [2] propose the usage of drones for area coverage in agricultural applications. Rosalie et al. [13] proposed an ant-colony algorithm paired with a way-points based mobility model to improve the area coverage of drones. The approach of Schleich et al. [15] is to maintain a connected network among the drone by exploiting a tree-based overlay network, as well a mechanism that allows drone to predict positions of one-hop neighbours in the tree.

Most of the above works do not consider the remaining level of the battery for the drones in the system when planning for movement or actions. However, there are few approaches that take batter in consideration for different purposes. For example, Messous et al. [10] propose an approach that tries to keep network connectivity taking into account the battery level of the drones in a fleet. Although in a preliminary shape, our approach differentiates from the above ones as it embeds the following features: (i) consideration of battery consumption in the decentralized modeling for area coverage, and (ii) integration of cloud computing technologies for the offloading of computation devoted to area coverage activity.

3 Reference Scenario and Architecture

The envisioned scenario is depicted in Figure 1. In such scenario, groups of persons move across the considered area. The area is potentially large and can contains obstacles such as trees or buildings. The area can already be equipped

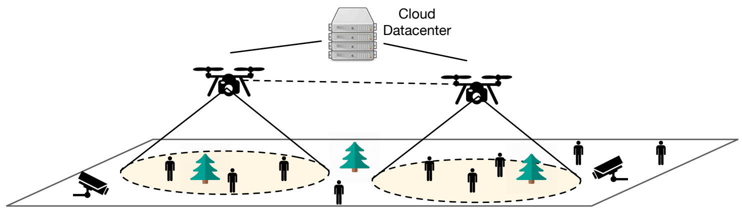

Fig. 1. The envisioned drone-assisted cloud-based crowd control scenario

with static ground sensors devoted to area coverage that are interconnected via a wireless or wired network. The typical operations conducted by these sensors include the estimation of the crowd density, motion and behaviour. However, an infrastructure made solely of static sensors can be not enough to cope with unexpected events that can result from the inherent unpredictability of crowd behaviour and the environment, such as for example a broken camera or unexpected visual obstacles. In addition, ground sensors takes time to be installed, and therefore cannot be deployed in an unexpected situation if not foreseen in advance. Also, the monetary investment for the monitoring of a single event can not be justified in certain scenarios (e.g. research activities). Therefore we advocate a scenario in which sensor-equipped drones complement with the ground sensor network in order to resolve many of the aforementioned issues. In such scenario a fleet of drones flies above the area, each drone connected to each others and with the network of ground sensors. Drones can be used as highly-moving computational and storage units, allowing for a dynamic access point toward remote cloud datacenters. They can be deployed to different locations on demand, with a very short notice and without requiring a dedicated static infrastructure placed beforehand. Their behaviour can be reprogrammed while in mission, making them suitable to adapt to fast and unpredictable events within the same mission.

In the light of aforementioned vision, the project focuses on two tightly connected aspects:

- a scalable and decentralized support for drones-to-drones and drone-to-ground communication, with the aim of disseminate information about both the state of the sensor and behaviour of the crowd in the drone-assisted area coverage network.
- an effective and QoS-aware orchestration of the computation related to area coverage in terms of computational resource selection, task management, and offloading to remote computational resources, organized by means of the Cloud Computing paradigm.

An high level overview of a reference architecture for the internal software of the drone is depicted in Figure 2. On the bottom level of the architecture lies the drone hardware. We assume drones to be equipped with sensor for manoeuvrability (e.g. GPS, rotors controller, etc..) and image acquisition (e.g. cameras).

We also assume they are equipped with relatively high battery capacity and computational power.

The *communication manager* module will take into account the management of the drone-to-drone and drone-to-ground communications. Since communication is a costly operation, a particular care will be taken such that information dissemination will be done in an effective way, maximising the usefulness of information sent. The component will also take into account the unreliability aspects of the communication channels. The information obtained by means of the communication module will feed the *local context manager*. The context models the view of a drone about its surrounding, and contains information of other drones, ground sensor and about the crowd. The information of drones range from their positions, direction and speed, to battery level and computational capacity. An important features of the context is the exploitation of prediction algorithms to predict ahead the context, which will allow the drones to plan in advance their behaviour so to possibly anticipate or avoid critical situations. On the top of the stack, the *application manager* orchestrates the computational aspect of the drones. The computational tasks can be related to the decentralized organization of the fleet and area coverage activities. The *movement planner* decides the trajectory of the drone considering the local context, and in such a way to globally optimize the area covered by the fleet. The *crowd tracking* module will employ image recognition algorithms already existing in literature in order to acquire information and build models of crowd behaviour. The application manager will coordinate the computation underlying these modules by deciding whether to execute the related tasks locally or remotely according to the local context.

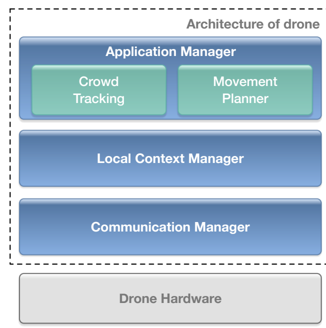

Fig. 2. High-level architectural view

4 A Model for area coverage

In general terms, the targeted objective is to *cover* as much area as possible by exploiting a fleet of drones, while minimizing battery consumption. The exact meaning of cover depends on the specific application scenario; In this context we consider an area coverage devoted to crowd control, in which the main task carried out by drones drones acquire images of a certain area and perform some analysis with the aim of recognize certain items (i.e. crowds, people, objects). Therefore, our problem can be formulated as: for a given accuracy associated with the recognition/classification of the items, we have to find/compute a combination of states for the drones that maximizes the surveillance area coverage while minimizing their battery energy consumption.

We consider each drone having a state composed by three parameters: (i) its geographical (GPS) position (x, y), where x and y are linear coordinates; (ii) the altitude h (distance from the ground), which also specify the area coverage S; and (iii) energy level of the battery E. The task is then to find the proper combination of the states that would maximize the area coverage for the given accuracy while minimizing the battery level consumption.

The proposed approach considers to divide the behaviour of the drones in two phases:

1. A local adjustment phase in which drones applies the necessary changes in order to respect the given accuracy;
2. A phase in which (i) drones exchange each other possible option plans for their positioning, and (2) reach a consensus on which plan to apply. This phase exploit a best-of-n formulation to reach the consensus.

In the following we provide a preliminary analysis and modelling of the two phases.

4.1 Local adjustments for accuracy

The given recognition accuracy is reached by the drones by operating locally on two parameters:

– The altitude of the drone h (Figure 3). According to its altitude, a drone can cover a different portion of an area with a different level of details. A lower altitude allows to receive more details of the items of interest and, presumably, to perform a more detailed and precise recognition/analysis of such items. Nevertheless, in this case the coverage area of the drone decreases together with the attitude $S = \pi r\hat{}2$, where r is the radius of the coverage circle. Considering the ϕ as angle of horizontal field of view of the drone to be fixed, while h changes in time we can derive the area coverage radius as:

$$r(t) = h(t) * \tan(\frac{\phi}{2}) \tag{1}$$

– The specific algorithm used for image recognition. This choice has an impact on battery consumption due to different complexity of the computation. Normally, the more accurate algorithm is for the recognition, the more energy consumption computation has to be executed. Hence decreasing the accuracy of the recognition can decrease the required energy consumption.

Therefore, we can derive the accuracy of the recognition by the drone as a function of its attitude and energy consumption. In other words given the required accuracy and the local state of the battery the drone can compute in what way to satisfy the accuracy, by: (i) decreasing the covered area; (ii) increasing the computation complexity of the algorithm increasing battery consumption; or (iii) applying a combination of both these approaches.

4.2 Best-of-n formulation

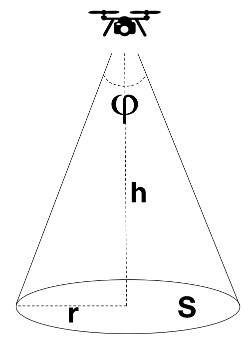

Fig. 3. Area coverage for a drone according to its height

Self-organization is a popular research topic in robot swarm, especially in its Best-of-n Problem formulation [16]. In particular, the best-of-n problem refers to problem of collective decision making done by a set of agents. According to Valentini et al. [16] *"The best-of-n problem requires a swarm of robots to make a collective decision over which option, out of n available options, offers the best alternative to satisfy the current needs of the swarm"*. As a consequence, a decision among the options is taken according to the concept of majority (i.e. when a sufficient number of agents favour a specific option), which generally depends on the specific application.

Valentini et al. also categorize the best-of-n problems according to two specific characteristics of an options, namely *quality* and *cost*. Both these characteristics depend on the application scenario considered. A best-of-n problem is then categorized according to the symmetry or asymmetry of both quality and cost. If in a given problem all the options have the same quality, the n-problem is symmetric with respect to quality. Otherwise if at least two options have different quality, then the problem is asymmetric. The same reasoning goes for the cost.

Few assumptions are necessary to frame the active area coverage defined in this paper into a discrete best-of-n problem. The first assumption is done by making discrete the problem of coordinate a set of drones (flocking). This can be done by considering the following two factors: (i) limit the area of actions of drones and (ii) divide this area into a grid of tiles, and the movement of the drones are defined as movement from one tile to another. With these two assumptions, the flocking problem goes to continuous to discrete, as the possible actions (movements) of the drones are finite. The second assumption is that the drones have in place the proper protocol an technology (such as XBee [18]) to communicate to each other in order to exchange the computed plans and to select the one that best satisfy the area coverage problem.

By applying the same criteria used by Valentini et al. we characterize the area coverage as a best-of-n problem, by associating the quality of a solution with the amount of area covered (at the given quality) by the fleet, and the cost with the amount of energy spent by drones in order to provide such coverage. According to this formulation, both quality and cost are asymmetric. The interaction between cost and quality can be defined as synergic or agnostic: in the former case the best option has the best quality with the minimum cost, in the latter the best option results in a tradeoff between cost and quality. In our case the interaction between area coverage and energy consumption is synergic, as the best solution would be the one that maximizes the area coverage while minimize the energy consumption.

5 Offloading to cloud computing

In the last years, many approaches have dealt with scenario in which computation is offloaded from mobile devices to cloud datacenters [6]. The benefit of such offloading is to improve the capacity of mobile and thin devices, usually limited in terms of CPU, memory and battery life, so that even simple devices can run complex and demanding applications. Among the many proposals, MAUI [5] and CloneCloud [4] are based on virtual machine migration and focuses on offloading of computation from mobile devices to remote servers at execution time, allowing the developers of applications to decide which computation can be offloaded.

In terms of computational resources, Cloud computing could represent an ideal back-end solution to manage the computation related to crowd tracking and image processing [9, 8]. However, due to the large amount of data collected, which needs to be transferred to the cloud, and the inherent dispersion of entities that performs data collection, it can be infeasible or inconvenient to transfer the computation toward a large remote datacenter. This is specially true in out envisioned scenario, as the behaviour of the crowd for the purpose of area coverage shall be identified fast such to allow the drone fleet to adjust their position.

This scenario points toward the case in which several ground sensors, or some powerful drone, assume the role of *cloudlet* [14], while normal drones the role of mobile devices. In the cloudlet model, drones would offload their computation to cloudlets, which are relatively small computational units connected with the full blown remote cloud server. Cloudlets are deployed locally to the area of interest and often placed in common and crowded areas to achieve physical proximity with mobile devices. This aspect provides devices with low latency and high bandwidth connections, thereby allowing an interactive response for demanding applications.

The approaches defined for cloudlet currently developed target mobile devices like smart-phones or laptops. The difference with respect to our scenario is the fact that offloading from smart-phones does not affect the context of the cloudlets or the devices. Instead, in our scenario the offloading also affects the behavior of a drone, which in turns can affect the whole fleet. In other terms, the decision whether to offload is not only affecting the quality of the application but potentially affects the area coverage scenario as a whole, for example by modifying the behavior of the other drones in the fleet. Therefore, we plan to adapt existing or design new distributed algorithms that: (1) orchestrate the computation also considering the effect that offloading can have in all the entities related to the crowd tracking, and (2) perform fast and effective brokering of cloudlet resource [1], in order to guarantee the quality of service demands from the crowd tracking tasks.

6 Conclusion

The organization of the activities of a fleet of drones is a relevant task in many of today's smart environment. In this paper, we present several initial considerations about the area coverage, i.e. the activity devoted to the analysis of an area through image analysis. Specifically, we analysed the problem of decentralized consensus for movement plan in a best-of-n problem formulation, and we reviewed the current trends and approaches for computation offloading, specifically for image analysis, in the frame of cloud computing technologies. As future work, we plan to integrate the analysis provided in the paper into a concrete proposal, both from a technological and algorithmic viewpoints.

References

1. G. F. Anastasi, E. Carlini, M. Coppola, and P. Dazzi. Qbrokage: A genetic approach for qos cloud brokering. In *2014 IEEE 7th International Conference on Cloud Computing*, pages 304–311. IEEE, 2014.
2. A. Barrientos, J. Colorado, J. d. Cerro, A. Martinez, C. Rossi, D. Sanz, and J. Valente. Aerial remote sensing in agriculture: A practical approach to area coverage and path planning for fleets of mini aerial robots. *Journal of Field Robotics*, 28(5):667–689, 2011.
3. E. Carlini, L. Ricci, and M. Coppola. Integrating centralized and peer-to-peer architectures to support interest management in massively multiplayer on-line games. *Concurrency and Computation: Practice and Experience*, 27(13):3362–3382, 2015.
4. B.-G. Chun, S. Ihm, P. Maniatis, M. Naik, and A. Patti. Clonecloud: elastic execution between mobile device and cloud. In *Proceedings of the sixth conference on Computer systems*, pages 301–314. ACM, 2011.
5. E. Cuervo, A. Balasubramanian, D.-k. Cho, A. Wolman, S. Saroiu, R. Chandra, and P. Bahl. Maui: making smartphones last longer with code offload. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, pages 49–62. ACM, 2010.
6. H. T. Dinh, C. Lee, D. Niyato, and P. Wang. A survey of mobile cloud computing: architecture, applications, and approaches. *Wireless communications and mobile computing*, 13(18):1587–1611, 2013.
7. A. Esmaeili, N. Mozayani, M. R. J. Motlagh, and E. T. Matson. A socially-based distributed self-organizing algorithm for holonic multi-agent systems: Case study in a task environment. *Cognitive Systems Research*, 43:21–44, 2017.
8. H. Kavalionak, E. Carlini, A. Lulli, C. Gennaro, G. Amato, C. Meghini, and L. Ricci. A prediction-based distributed tracking protocol for video surveillance. In *2017 IEEE 14th International Conference on Networking, Sensing and Control (ICNSC)*, pages 140–145, May 2017.
9. H. Kavalionak, C. Gennaro, G. Amato, and C. Meghini. Dice: A distributed protocol for camera-aided video surveillance. In *Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM), 2015 IEEE International Conference on*, pages 477–484. IEEE, 2015.

10. M.-A. Messous, S.-M. Senouci, and H. Sedjelmaci. Network connectivity and area coverage for uav fleet mobility model with energy constraint. In *Wireless Communications and Networking Conference (WCNC), 2016 IEEE*, pages 1–6. IEEE, 2016.
11. R. Olfati-Saber. Flocking for multi-agent dynamic systems: algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, 2006.
12. A. H. Payberah, H. Kavalionak, A. Montresor, J. Dowling, and S. Haridi. Lightweight gossip-based distribution estimation. In *2013 IEEE International Conference on Communications (ICC)*, pages 3439–3443. IEEE, 2013.
13. M. Rosalie, J. E. Dentler, G. Danoy, P. Bouvry, S. Kannan, M. A. Olivares-Mendez, and H. Voos. Area exploration with a swarm of uavs combining deterministic chaotic ant colony mobility with position mpc. In *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 1392–1397, June 2017.
14. M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies. The case for vm-based cloudlets in mobile computing. *Pervasive Computing, IEEE*, 8(4):14–23, 2009.
15. J. Schleich, A. Panchapakesan, G. Danoy, and P. Bouvry. Uav fleet area coverage with network connectivity constraint. In *Proceedings of the 11th ACM international symposium on Mobility management and wireless access*, pages 131–138. ACM, 2013.
16. G. Valentini, E. Ferrante, and M. Dorigo. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:9, 2017.
17. G. Vásárhelyi, C. Virágh, G. Somorjai, N. Tarcai, T. Szörényi, T. Nepusz, and T. Vicsek. Outdoor flocking and formation flight with autonomous aerial robots. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3866–3873, Sept 2014.
18. Q. Yuan, J. Zhan, and X. Li. Outdoor flocking of quadcopter drones with decentralized model predictive control. *ISA Transactions*, 2017.