Flying Smartphones: When Portable Computing Sprouts Wings

Ross Allen, Marco Pavone*, Mac Schwager
Department of Aeronautics and Astronautics, Stanford University
{rallen10, pavone, schwager}@stanford.edu
*Corresponding author

Introduction

For decades pop culture has imagined for us a future filled with robotic companions that attend to our daily chores. While often featured in sci-fi, this vision of the future maybe more accurate, and more near-term, than expected. What the movies and TV shows may have gotten wrong, however, is the form of our future robotic companions. Instead of humanoids, aerial drones seem to be rapidly approaching adoption for everyday tasks.

From personal computers to smartphones to (smart)drones

At first glance aerial drones may seem a non sequitur in the list of personal computers (PCs) and smartphones, however they may indeed represent the next step-change in technology that connects the physical and digital worlds. Personal computers were the first technology to provide digital processing power to the average person. Smartphones brought the next step-change in technology, not only because they are mobile, but also because they integrated a basic set of sensors to a processing platform. Fusing a processor with a GPS receiver, accelerometers, magnetometers, and internet connectivity has enabled so many unique applications that app developers will be exploring the design space for decades to come. Drones mark the next leap in this progression. Along with a processor and sensor suite, drones incorporate actuators---propellers to move themselves around, and potentially grippers to manipulate objects in the world. The fusion of these three elements (computation, sensing, and actuation) along with developments in the theory of robot autonomy allow drones to *actively* engage with the world around them; this is in contrast to the relatively passive interactions between humans and PCs, and humans and smartphones.

Smartdrones

Not all unmanned aerial vehicles are consistent with a comparison to smartphones. Remote controlled aircraft have existed for decades but, lacking any form of autonomy, can not be considered "drones" and are likely to remain, strictly, a hobbyist's pursuit. On the other end of the spectrum lie military drones, which are typically very expensive, complex, and highly task-specific. The type of drone for which we draw parallels to smartphone technology, arguably the type that will have the most impact on the average person's daily life, are drones we will refer to as *smartdrones*. Smartdrones have several defining features:

- **Affordability**: to achieve wide-scale use, smartdrones will likely fall in the same price bracket as smartphones and modern laptops, i.e. $500 - $2000, which makes them affordable by a common household.
- **Lightweight structure**: many potential applications for smartdrones will directly or indirectly involve operations in proximity to human subjects. This immediately opens the issue of safety. Safety not only depends on a robust autonomy/software architecture, but also on the drone's

physical design. Lighter weight, slower drones are inherently safer, and thus will be the platforms of choice for operations in human-centric environments.

- **Autonomy**: beyond just being robust and reliable, smartdrones must also be intuitive to use. In the same way that a smartphone can be used by anyone, regardless of prior computer knowledge, smartdrones must be easily usable by those with no technical background. To achieve this, drones will have to be highly autonomous; relying on the onboard processor for all low-level control, and leaving only application selection and a few input options to the user.
- **Standardization**: perhaps the essential, defining characteristic of a smartdrone is its flexibility to a wide range of applications. To achieve this wide range, smartdrones will share a quasi-standardized set of hardware and a unified control/autonomy structure. Hardware will range from components typical of smartphones (e.g., accelerometers, gyros, barometers, cameras, thermometers, even microphones), to propulsive systems and manipulator/grasping mechanisms for payloads. On the software side, with a standardized GPS module and communication protocol, smartdrones will need to have identical behavior when avoiding restricted airspaces.

So far, we have made no distinction between quadrotors and fixed wing aircraft when referring to smartdrones. This is because either of these platforms is capable of meeting our definition, so either or both may be adopted. It is worth noting, however, that the hover capabilities of quadrotors tends to add an additional safety layer over fixed-wing craft, therefore making them the more attractive of the two for applications in human proximity. Fixed-wing aircraft, on the other hand, offer much greater range and endurance. Hybrid craft, such as tiltrotor craft, would also fit our definition of smartdrones.

In the remainder of this article we focus our attention on the autonomy feature (arguably, one of the most critical), from the technology itself, to its safety aspects, and, finally, to the range of applications that it enables.

Control/autonomy structure

A unified control/autonomy structure is key for the smartdrone concept so that app developers will know that the software they develop will interact with firmware and hardware in much the same way, regardless of smartdrone model or manufacturer (similarly to how an app can be released on android and iOS with little additional work). The unified control/autonomy structure will likely mirror the structure that has been developed for many research-based quadrotors; a high-level representation of such a control structure is given in Figure 1. The control structure is composed of a set of nested loops. Outer loops, responsible for more abstract decisions, feed information down to inner loops, usually in the form of setpoints or reference targets, which drive the direct control of the smartdrone hardware. Sensors feedback information about the state of the smartdrone to the relevant control layer.

Specifically, the user selects an application for the smartdrone and the app produces a set of high-level objectives. The motion planner fuses these objectives with information about the world---such as obstacle locations, no-fly zones, or speed restrictions---to come up with a feasible plan for achieving the objectives. The position controller is tasked with executing the plan by comparing the desired position from the motion planner with the actual position read by the sensors and performing feedback control. The attitude controller is tasked with stabilizing the aircraft along with executing the positioning commands from the position controller. Since most drone platforms are underactuated, the attitude controller is a “slave” to the position controller in that arbitrary positions and velocity cannot be achieved independent of attitude, so the attitude controller accommodates the desired positions and velocity. While the outer loops of the control structure will employ sophisticated optimization, control, and decision-making techniques, the inner loops will likely apply simple, robust proportional-integral-derivative controllers.

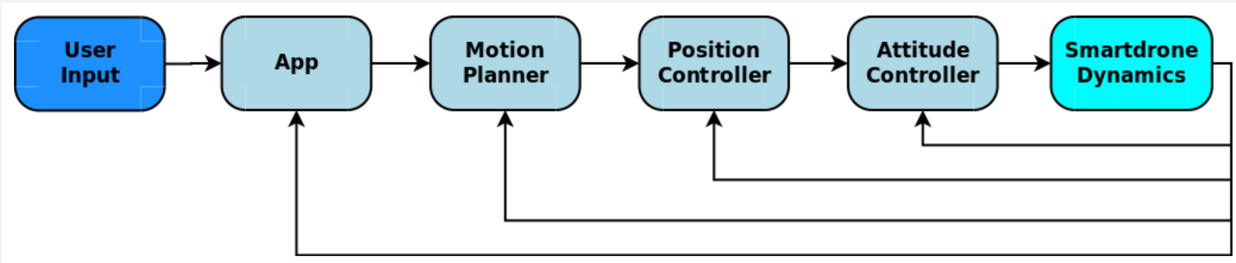

Figure 1: High-level control/autonomy structure for smartdrone platforms.

Safety

As PCs, smartphones, and smartdrones introduce progressively more powerful technological applications, they also carry an ever increasing burden of risk---an example of the proverbial double-edged-sword. For example, PCs allowed the average person to digitize most of their personal credentials and financial information. This greatly simplified tedious tasks such as filing taxes, but also opened the door to risks such as identity fraud. The primary safety issues that are being addressed as smartdrones are adopted in wide-scale use fall in the categories of: sensing, planning, verification, and system-level integration.

Sensing: Each layer of the smartdrone control/autonomy structure, as given in Figure 1, requires its own sensing hardware that is used in estimating the current state of the craft. The innermost layer, representing attitude estimation, is for the most part a solved problem. Even inexpensive, off-the-shelf IMUs are sufficient to estimate and control attitude. This is why one can purchase a quadrotor “toy” for less than $50 and have it hover and perform basic motions. Such remote-controlled toy quadcopters are often well trimmed, so they are capable of hovering in place fairly reliably. Being well-balanced to avoid drifting during hover is, however, very different from autonomously controlling the position of a drone.

Position control, indicated by the second innermost layer in Figure 1, is a greater challenge because position estimation requires considerably more sophisticated hardware than that of attitude control. For absolute position, a smartdrone would require a GPS module. GPS modules are relatively expensive, running at $80 for a hobby-grade component. Furthermore, GPS alone may be insufficient to guarantee safe operation. GPS relies on line-of-sight to GPS satellites, making it unreliable in environments with partial or full obstruction of the sky (e.g. canyons, forests, near tall buildings, and indoors). Thus, smartdrones will likely have to supplement GPS information with localized position information to provide terrain-relative position data. Several sensors types are capable of achieving this, including sonar, lidar, and vision. In the end, position estimation will be achieved by a mixture of these technologies.

Planning: Even with perfect and complete sensor data and an infallible controller, a major issue exists in how to decide what trajectory a smartdrone should take through a complex, dynamic world. These questions have been at the center of the field of robot motion planning for years. Smartdrones present a particularly challenging form of robot motion planning because they require the consideration of a high-speed robot in a changing environment. This form of motion planning, termed real-time kinodynamic planning, is an active field of research. Recent work at the Autonomous Systems Laboratory at Stanford University has developed a framework for solving such problems in real-time [Allen and Pavone, 2016]. The framework operates on an offline-online computation paradigm, whereby a library of trajectories is precomputed offline and then efficiently pruned online when environment data becomes available. Machine learning and optimal control techniques make such a procedure fast and accurate, in the sense that near-optimal trajectories are repeatedly computed every few milliseconds. Figure 2 shows the application of such a framework to the control of a quadrotor that dodges a fencing blade.

Figure 2: Demonstrating real-time kinodynamic planning on a quadrotor dodging a fencing blade.

Verification: Recent work has sought to verify the safety of smartdrone systems by embedding verification directly into the design of the control/autonomy module. The field of formal methods, which was traditionally developed to verify the correctness of computer programs, has now been applied to design drone control systems that are correct "by construction." For example, work at the Stanford Multi-robot Systems Lab, and Boston University Robotics Lab has led to formal methods algorithms that construct provably safe trajectories for multiple smartdrones to perpetually monitor an environment, while scheduling sufficient time to recharge their own batteries [Leahy et Al., 2015].

System-level integration: The US Federal Aviation Administration (FAA) has recently overhauled its regulations regarding the use of Unmanned Aircraft Systems (UAS), now requiring private operators to register their drones in a national database, and prohibit flying of recreational drones near airports and other areas with sensitive airspace. As drone capabilities grow and as autonomous features find their way into commercialized drone technology, the FAA is incrementally taking steps to integrate drones into the already complex US airspace (https://www.faa.gov/uas/). Likewise, the European Aviation Safety Agency (EASA) is taking precautions to integrate private drone usage safely into the European airspace (https://www.easa.europa.eu/easa-and-you/civil-drones-rpas). Nonetheless, significant challenges remain to safely integrate drones, whether commercial or private, into the airspace and the regulations governing drone usage are expected to evolve considerably in the coming years.

Smartdrone Apps

Some of the first applications for quadrotors (and, more in general, smartdrones) were realized in the laboratory. Due to their ease of control and robustness to changing configurations, quadrotors became an excellent demonstration platform for navigation, planning, and network algorithms [Mellinger et Al. 2012]. These research demonstrations paved the way for many of the commercial and military applications we see being developed today.

Perhaps the application that has received the most attention by the public is the proposed use of unmanned aerial delivery platforms such as Amazon Prime Air. Drone delivery has the potential to radically change the way we access consumer products as it would lower the delivery time for online purchases from days to minutes. Amazon's delivery system does not quite match our description of smartdrones, however, since it involves expensive, large, task-specific aerial robots; not flexible, inexpensive platforms usable by the public. Restaurants, on the other hand, could utilize the more universal smartdrone concept for delivery of small food items to local communities. Similarly, medicine and first aid supplies could be delivered to remote or hazardous areas during disaster events.

Another drone application that has made its way into mainstream media is that of recreation use, specifically for action/adventure sports. Established drone companies, such as DJI, and startups, such as Lily Robotics, are planning to offer multirotor aircraft that are designed to autonomously follow a user and shoot video. While these products are not yet on the market, the significant number of companies and startups pursuing this concept gives credence to the idea that we will soon see quadcopters chasing skiers down mountains. Currently the most economically viable, albeit lesser known, application for drone technology lies in agriculture. Companies such as 3DR are providing autonomous multirotor craft that can survey crops by recording multispectral images of farmland.

The power of the smartdrone concept becomes even more apparent when one imagines multi-drone collectives acting collaboratively to carry out large-scale tasks. Just as the benefits of smartphones have exploded with the advent of mobile apps for social networking that cull data from a collective of users, the capabilities of smartdrones will explode as the interconnectedness of the drone network increases. Today, mobile apps that mine data from hundreds of thousands of daily users, such as Waze and Tealeaf, are able to effectively predict phenomena as diverse as traffic and stock prices. Tomorrow, smartdrones will leverage the perpetual networked aerial drone presence to give rich, real-time data about agricultural crops, traffic, weather, the movement of wildlife, the activities of suspected criminals, and give early warnings for everything from wildfires to freeway pileups. Furthermore, many of the deficits of the small size of smartdrones, including limited flight time, limited range, and limited payload, can be alleviated when one considers the coordinated actions of large groups of drones. A thriving research community in multi-robot systems and multi-agent control is currently devoted to solving problems of large-scale coordinated autonomy. New decentralized algorithms are emerging for control, perception, and trajectory planning over a wireless network to enable multi-drones systems: groups of drones that reach collective inferences about the world and make collective decisions about what actions to take in the world to accomplish a task.

The potential applications of this smartdrone collective are vast. Perhaps the first capability that will be realized from smartdrone collectives will be large-scale distributed perception. Drones will provide us with a perpetual sensor network in the sky, to sense diverse forms of data for diverse purposes [Schwager et Al. 2011], as illustrated in Figure 3. As mentioned, farmers are already using individual drones for crop sensing, to see daily or weekly detailed snapshots of crop health. These snapshots then inform decisions about watering, fertilizing, and applying pesticide to specific areas of the crops where they are most needed. With the advent of smartdrone networks, farmers could have an on-demand updated computer model of the health of their crops for crop management decisions. Smartdrone networks will also help search and rescue teams find lost hikers in the wilderness, or victims of boating accidents lost at sea. The key is the ability of a smartdrone network to parallelize the task of gathering information over a large area. The larger the area, the more drones one can deploy to search it efficiently. Construction sites, which are frequently targeted for theft, and large-scale infrastructure, which requires frequent inspection, could employ smartdrones for persistent surveillance. High-tech border security could implement a fleet of smartdrones that could monitor large stretches of remote terrain.

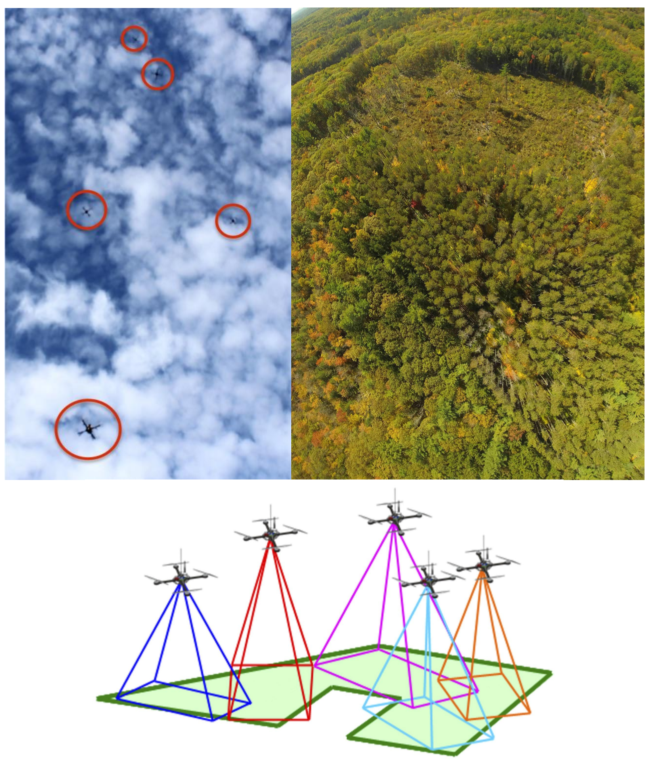

Figure 3: Five smartdrones cooperatively surveying a research forest.

Beyond merely sensing the environment, smartdrones *interacting* with the environment (for example with grippers, display lights, and other actuators) will open up a new range of exciting applications. For example, a group of smartdrones with colored LEDs can form a massive 3-D display, creating a new medium for art, entertainment, communication, and marketing. Researchers at ETH Zurich have already shown the promise of such drone displays [Alonso-Mora et Al. 2012]. When equipped with grippers, a smartdrone collective might soon replace cranes in construction sites, collaborating to hoist heavy beams into place to build buildings and bridges alongside human construction workers [Lindsey et Al. 2012]. One day, national forests may employ groups for autonomous smartdrones, to not only monitor for forest fires, but to fight them with the targeted application of fire retardant; and farmers may use smartdrones, not only to monitor crop health, but to actively manage crops by applying water, fertilizer, and pesticide with surgical precision. Indeed, the most transformative applications for smartdrones are most likely still waiting to be discovered by the app developers and drone users of the future.

Conclusion:

Due to their ability to actively, autonomously interact with the world, lightweight, highly-autonomous drones are emerging as the next step-change in consumer electronic technology, much in the same way that smartphones revolutionized personal computing. While research is ongoing to ensure safe, autonomous operation, smartdrone systems are already being utilized in several applications, with many more applications soon to emerge. After two decades of research and development, portable computing has finally sprouted wings!

References:

R. Allen and M. Pavone. "A Real-Time Framework for Kinodynamic Planning with Application to Quadrotor Obstacle Avoidance". In: AIAA Conf. on Guidance, Navigation and Control. San Diego, CA, Jan. 2016

K. Leahy, D. Zhou, C.-I. Vasile, K. Oikonomopoulos, M. Schwager, and C. Belta. "Persistent surveillance for unmanned aerial vehicles subject to charging and temporal logic constraints." Autonomous Robots, 1-16, 2015.

D. Mellinger, N. Michael, and V. Kumar. "Trajectory generation and control for precise aggressive maneuvers with quadrotors." International Journal of Robotics Research, 31(5):664-674, 2012.

M. Schwager, B. J. Julian, M. Angermann, and D. Rus. "Eyes in the sky: Decentralized control for the deployment of robotic camera networks." Proceedings of the IEEE 99, no. 9 (2011): 1541-1561.

J. Alonso-Mora, M. Schoch, A. Breitenmoser, R. Siegwart, and P. Beardsley. "Object and animation display with multiple aerial vehicles." IEEE/RSJ International Conference on Intelligent Robots and Systems, 1078-1083, 2012.

Q. Lindsey, D. Mellinger, and V. Kumar. "Construction with quadrotor teams." Autonomous Robots 33(3):323-336, 2012.