Literature Review

Submitted in partial fulfilment

of the requirements of the degree of

Bachelor of Science (Honours)

of Rhodes University

Grahamstown, South Africa
17th June 2014

Contents

Chapter 1

What is a drone

1 Drone definitions and classification

A drone is an unmanned vehicle that is controlled from a distance by a pilot or driver. The term drone has been popularised in recent times due to the high profile of Unmanned Aerial Vehicles (UAVs) commonly used in modern warfare and also popularised in popular culture. They have become synonymous with the US led War on Terror in the Middle East. Despite the prevalent image of a drone being a small airplane, equipped with a large array of sensors and weaponry, and controlled by a person sitting in an air-conditioned room far away, this represents only a very small and specialised application for drones. The term drone is the banner label for all unmanned vehicles. There are then many subsets of drones such as UAVs and Unmanned Ground Vehicles (UGVs). These can then be split up further depending on classifying features.

UAVs have been around, in a much simpler form, since the early 20th century. These simple devices were intended to perform basic reconnaissance or to drop off munitions and were either balloons or kites with no remote control capabilities. They relied on timers and tethers to perform their intended purpose. (Krock, 2002)

The development of UAVs continued over the years and by the 1930s they were capable of being controlled remotely. The Queen Bee, created in the UK, was the first reusable UAV. It was used by the Royal Air force and Navy as a target for shooting practice. (Krock, 2002)

Since then the industry has taken off. Military applications for UAVs were seen as indispensable in warfare as they could be made smaller than conventional aircraft and

therefore avoid detection. With their reduced radar signature and lack of risk to a pilot, UAVs quickly become popular for their reconnaissance capabilities in dangerous areas.

Recently UAVs have taken on the additional role of remote strike tools. The United States first made use of the Predator drone armed with a Hellfire missile to remotely eliminate a target in late 2001. (staff, 2005) This marked a turning point for larger UAVs as they took on the dual role of reconnaissance and strike vehicle. While larger UAVs worked on increasing airtime and armament size, smaller UAVs began to be employed as close reconnaissance support for soldiers in the field.

In the last few years the creation of Micro Aerial Vehicles (MAVs) has become a huge industry in both the military and the public sector. MAVs are now being used as research tools, for disaster and agricultural surveillance as well as for recreational use.

1.1 Drone classification categories

Drones come in many different shapes, sizes and capabilities. However, they are usually all grouped together under the general heading 'drones'. Dr. Arjomandi (Arjomandi *et al.*, 2007) suggests that there are six features that should be used in the classification of drones into sub categories. These 6 categories are weight range and endurance, maximum altitude, wing load, engine type and mission aspect. The first five categories are outlined in tables 1.1 to 1.5 and the mission aspect is discussed afterwards.

Classification by weight		
Designation	**Weight Range**	**Example**
Super Heavy	>2000 kg	Global Hawk
Heavy	200 - 2000 kg	A-160
Medium	50 - 200 kg	Raven
Light	5 – 50 kg	RPO Midget
Micro	<5 kg	Dragon Eye

Table 1.1: Classification by weight

Range and endurance			
Category	**Endurance**	**Range**	**Example**
High	>24 hours	>1500km	Predator B
Medium	5 − 24 hours	100 − 400 km	Silver Fox
Low	<5 hours	<100 km	Pointer

Table 1.2: Range and endurance

Maximum altitude		
Category	**Max Altitude**	**Example**
Low	<1000 m	Pointer
Medium	1000 − 10000 m	Finder
High	> 10000m	Darkstar

Table 1.3: Maximum altitude

Wing loading		
Category	**Wing loading kg/m2**	**Example**
Low	<50	Seeker
Medium	50 − 100	X-45
High	>100	Global Hawk

Table 1.4: Wing loading

Engine type	
Engine	**Example**
UEL Rotary	Outrider
Turbofan	Global Hawk
Two-stroke	Pioneer
Piston	Predator
Turboprop	Predator B
Electric	Dragon Eye
Push & Pull	Hunter
Prop	LEWK

Table 1.5: Engine type

Dr. Arjomandi (Arjomandi *et al.*, 2007) identifies six mission areas in which a drone may be used.

- Intelligence, Surveillance, Target Acquisition and Reconnaissance (ISTAR)

- Combat (UCAV)

- Multi-Purpose

- Vertical Take-Off and Landing (VTOL)

- Radar and Communication Relay

- Aerial Delivery and Resupply

A drone may fulfil more than one of these tasks to varying degrees. To this end Dr Arjomandi (Arjomandi *et al.*, 2007) devised a ranking system to decide how well a drone fulfils a given task. These ratings are shown below:

1. Does not perform mission aspect at all.

2. Performs mission aspect with the lowest performance characteristics.

3. Performs mission aspect with medium performance characteristics.

4. Performs mission aspect with high performance characteristics.

5. Performs mission aspect with the greatest performance characteristic capabilities.

1.1.1   Conclusion on classifications

These classifications allow for a diverse range of drones to be differentiated. Although the intent is mainly for the classification of military drones, there is a large range of smaller drones that would either be difficult to classify or would all be grouped together despite some quite important differences in design and intended purpose. Two examples of this are the Parrot AR Drone and the Black Hornet Nano. While the AR Drone is primarily designed for gaming use, the Black Hornet is designed for military reconnaissance. Both fall into the micro UAV category despite the Black Hornet being substantially smaller. They also fall into the same categories for all of the other classifiers except for engine type where neither can be classified. These classifiers need to be broadened in order to allow for the vast number of micro drones that are now becoming available.

Chapter 2

The Parrot A.R. Drone 2

The A.R. Drone is a small remote controlled quadcopter developed by French company Parrot. It was released in 2010 (Cooper, 2010) at the Consumer Electronics Show in Las Vegas. It has quickly become of one of the most popular recreational drones due to its low cost and ease of control. (Selden, 2013)

A quadcopter is a form of rotary wing aircraft that makes use of four individually controlled rotors on the four corners of the device. Using four rotors allows for a lot more control over the device's movement at the expense of having to individually manage these rotors. When controlling a device there are 3 axes on which the device can be manoeuvred. These axes are yaw, pitch and roll (see figure 2.1).
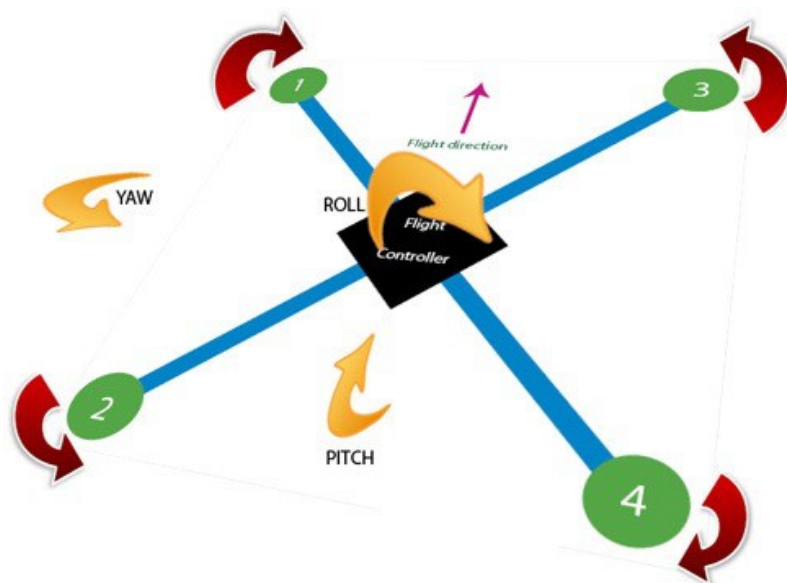
Figure 2.1: Yaw, pitch and roll of a quadcopter
(GAJENDRAN, 2012)

1 Aerodynamic effects on the quadcopter

There are many aerodynamic effects that need to be considered for a quadcopter to either maintain a hover position or to move in a specified direction. The quadcopters four rotors prove to have some unique advantages and challenges when compared to the traditional rotary wing aircraft, the helicopter.

1.1 Thrust

The primary function of all the rotors is to provide thrust. Thrust is used to push the quadcopter into the air and counter the effects of gravity on its body (Hoffmann *et al.*, 2007). All four rotors provide thrust which is an advantage over the helicopter where the rear rotor is only used for stabilization.

1.2 Rotational Torque

Having four rotors does not mean that the exertion of the rotor's torque on the body of the quadcopter is automatically negated. The blades spin in opposite directions to the two on

either side of them and in the same direction to the one opposite. Note in figure 2.1 that rotors 1 and 4 spin in the same direction and rotors 3 and 2 spin in the opposite direction. This has the effect of countering out the rotational torque of the rotors. Without this the quadcopter would yaw about its centre axis. (Hoffmann *et al.*, 2007)

1.3   Blade-flapping

Blade flapping is the phenomenon seen when two blades spin around the same axis. The leading blade will see more resistance than the trailing blade. This increased resistance on the leading blade causes it to rise upward, while the reduced resistance on the trailing blade causes it to drop. The blades then swap positions as they rotate about an axis causing a flapping effect. If the blades are of exactly the same weight and length, the effect is cancelled out and there is no resultant instability. (Hoffmann *et al.*, 2007)

1.4   Pitch and roll

Pitch and roll on the quadcopter are controlled in the same way with different combinations of rotors. By increasing the speed of the two rotors at one end of the quadcopter so that they are running faster than the two rotors at the opposite end, it will cause the faster running end to lift up and the quadcopter to move in the direction opposite to the lifted rotors (see diagrams 2.2 and 2.3).
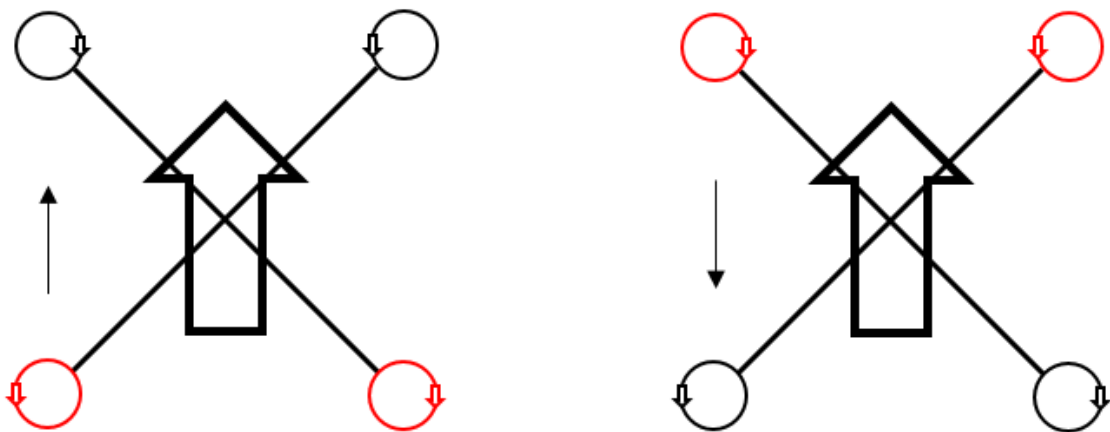
Figure 2.2: Pitch forward and backward
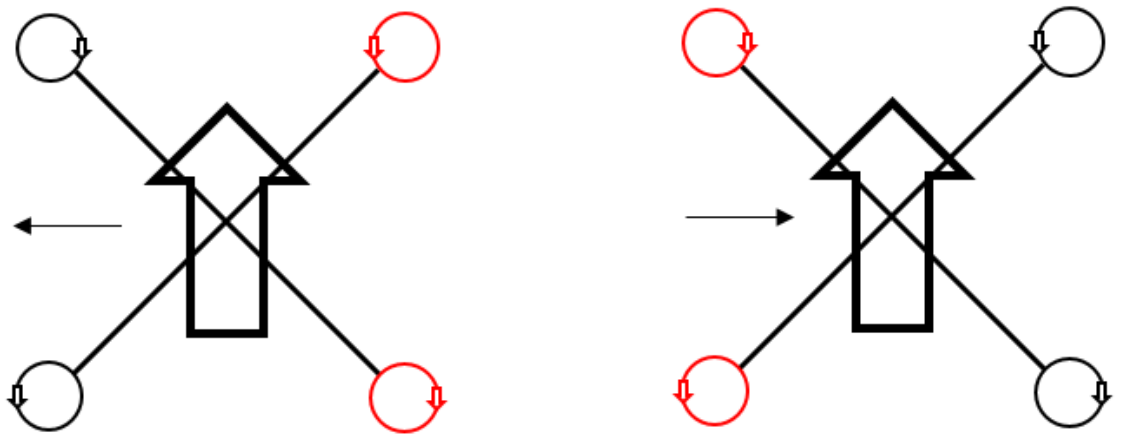
Figure 2.3: Roll left and right

1.5 Yaw

Yaw is the rotation about the devices vertical axis. By rotating two opposite rotors at a higher speed than the other two, the induced torque will cause the device to rotate. If the speed up of two of the blades is proportional to the slowdown of the other two then the device will stay at the same height and rotate on itself. (Hoffmann *et al.*, 2007)

Figure 2.4: Yaw clockwise and counter clockwise

1.6 Conclusion on aerodynamics

Analysis of the aerodynamics of the A.R. Drone highlights the complexity of the on-board control system that enables stable flight. It also highlight the need to have very accurately moulded parts for the A.R. Drone. Given the phenomena of blade flapping, the rotor blades n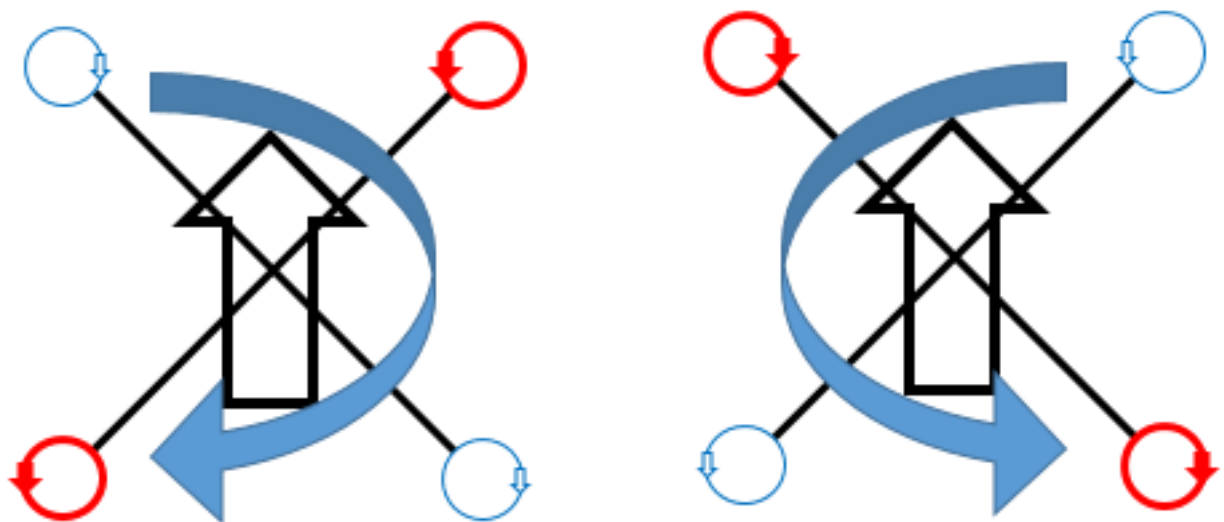eed to be exactly the same length and weight to minimise the effect. If a blade is damaged, it will lead to unstable flight. Unstable flight due to impaired blades would also manifest itself in the yaw torque of the drone. Lastly, uneven blades on one of the rotors could lead to the drone's thrust being impaired on that rotor. This would result in the on-board control system having to apply excessive amounts of correction to maintain stable flight.

2 A.R. Drone systems

2.1 Body

The A.R. Drone is designed for recreational use and is therefore designed with a cushioning Expanded Polly Propylene body. The body encases the internal control circuitry and battery. The four rotors for the drone are then affixed to a carbon fibre cross that is attached to the body. The A.R. Drone comes with two hulls, one for indoor and one for outdoor use. The outdoor hull simply covers the battery and USB port while the indoor hull includes a shroud that protects the rotors in the event of the drone striking an obstacle (see figure 2.5). (Bristeau *et al.*, 2011)
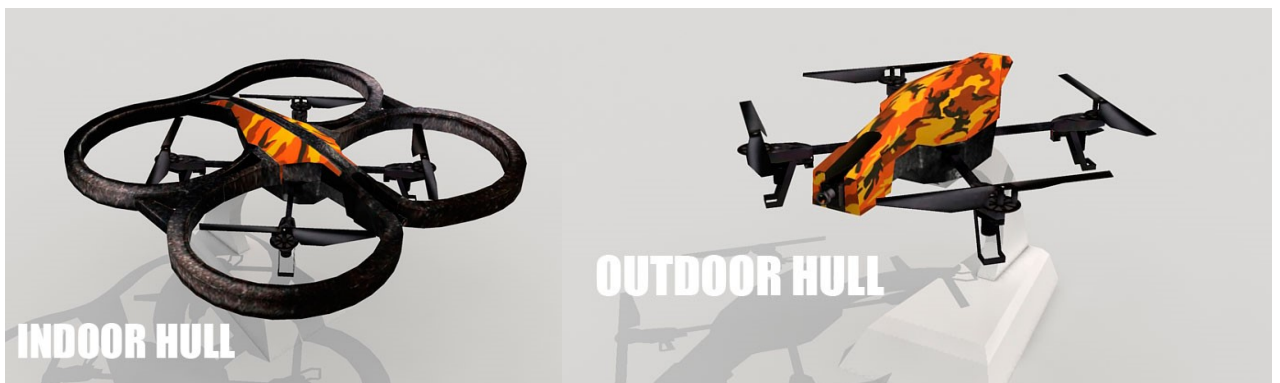

Figure 2.5: AR Drone with indoor and outdoor hull

(Perdana, 2012)

2.2 Engines

The A.R. Drone uses four independent rotors in order to generate lift. These rotors are directly driven by individually controlled brushless motors. Each motor is controlled by an ATMEGA8L processor allowing obstacles to be individually detected and the motors power to be cut-off. This feature helps prevent damage to the motors and rotor blades. (Bristeau *et al.*, 2011)

2.3 Lithium-Polymer Batteries

The A.R. Drone uses a single lithium-polymer battery to sustain flight for 10 to 15 minutes. The battery comprises three, series-connected cells, each with a nominal voltage (fully charged) of 4.2V. The overall battery voltage is 12.6V with a capacity of 1000mAh. The battery is housed in the centre of the A.R. Drone in a foam basket to cushion it. (Bristeau *et al.*, 2011)

2.4 Forward facing camera

Another of the A.R. Drones sensors is a forward facing camera. The output of the camera is in VGA format with a resolution of 640x480 pixels and a frame rate of 30 frames per second and a 92° field of view. (Parrot, n.d.b) The camera is not used by the drone's operating system. It is an auxiliary device that provides additional functionality to the end user. (Bristeau *et al.*, 2011)

2.5 Operating System

The A.R. Drone makes use of a custom, real time Linux operating system that handles all of the input and output interfaces that control the drone. The operating system is responsible for the following tasks:

- maintaining threads on Wi-Fi communications
- video data sampling from the on-board camera

- video compression (for efficient wireless transmission of the on-board camera's video data back to the operator)

- image processing of the downward facing camera for ground speed calculations

- sensor data acquisition

- pose estimation (the orientation of the drone in 3D space)

- closed-loop control

(Bristeau *et al.*, 2011)

2.6 On-board sensors

The A.R. Drone makes use of six on-board sensors in order to help determine its direction of travel, its speed and also its height above ground level. This helps the on-board system maintain stable flight as well as providing feedback to the user.

2.6.1 Ultrasound sensor

The ultrasound sensor is used for close range altitude sensing. It can therefore detect when the drone moves over a new object. The ultrasound sensor is capable of accurately measuring the altitude of the drone accurately up to a range of 6m. The ultrasound sensor thus enables the drone to fly over objects without getting too close to them. It also prevents the user from lowering the drone too rapidly onto the ground. (Parrot, 2012)

2.6.2 Pressure sensor

The pressure sensor is also used by the drone to determine its altitude at heights greater than 6m above ground level i.e. when the drone is beyond the 6m operating range of the ultrasound sensor. When the drone is switched on, the pressure sensor calibrates itself to ground level. As the drone rises and moves beyond the 6m range of the ultrasound sensor, the pressure sensor becomes the primary method of measuring the drone's altitude. (Parrot, 2012)

2.6.3   Downward facing camera

Besides the front-facing camera, the A.R. Drone has a secondary, downward facing camera mounted at the base of the main body. The video data from this camera can also be accessed by the user (in the same way as is done for the primary, front-facing camera). It's primary function is to enable the drone's speed relative to the ground to be determined. (Parrot, 2012)

2.6.4   3 axis gyroscope

A three axis gyroscope allows for the direction of motion to be detected along the X, Y and Z axis (InvenSense, n.d.). The gyroscopes are small solid state microchips on the motherboard of the drone that provide feedback to the drones on-board operating system enabling it to stabilise itself. (Parrot, 2012)

2.6.5   3-axis accelerometer

The 3-axis accelerometer works in a similar fashion to the gyroscope except that it has magnitude rather than a simple binary value. It is also a small silicon microchip on the drone's motherboard and is responsible for returning a value describing the drone's acceleration vector. (Parrot, 2012)

2.6.6   3-axis magnetometer

The magnetometer provides a reference to the earths magnetic north pole. It allows for a finer degree of control over the drone flight path. By having a point of reference, the drone is able to maintain a fixed direction and not yaw about its own axis. (Parrot, 2012)

Chapter 3

A.R. Drone SDK

On releasing the A.R. Drone, Parrot also released a free-to-use System Development Kit (SDK) so that independent developers could create their own games for the drone. The SDK allows the programmer to control the drone directly as well as access all information from the on-board sensors.

1 Communication with the A.R. Drone

The A.R. Drone communicates with a client device through an ad-hoc Wi-Fi network. On starting, the drone creates a Wi-Fi network (usually called "adrone_xxx"). (Piskorski *et al.*, 2012) The client then requests a connection to the drone. The drone allocates the user a free IP address. The client and the drone are then paired and the client can start issuing instructions while the drone will start streaming information from it's sensors to the client. (Piskorski *et al.*, 2012)

1.1 Communication channels

Due to the large volume of data being streamed to and from the drone, there are three separate channels used for communication. The first channel uses UDP port 5556 and is responsible for communicating the client's instructions to the drone. These "AT commands" or control commands are usually sent at a rate of 30 times per second, allowing the user to direct the drone's functionality and configuration. (Piskorski *et al.*, 2012)

The second line of communication is done over the UDP port 5554. This line is used to send data from the drone's on-board sensors. This data is referred to as "navdata" and is also transmitted around 30 times a second. The navdata reports information such as the status of the drone, information on the rotors as well as information from the other sensors. (Piskorski *et al.*, 2012)

The third line communicates over UDP port 5555 and is used for sending the information captured by the drone's cameras. The data is encoded and then sent to the client who can decode the information using a built in decoder in the SDK. (Piskorski *et al.*, 2012)

An optional fourth line of communication can be established on TCP port 5559, the control port. This is used when the data that is being sent is essential. This is compared to the data sent over the UDP ports that is non-essential. It can be used to send critical data or confirm that configuration data has been sent. (Piskorski *et al.*, 2012)

1.2 AT Commands

AT commands are issued to the drone by the client in order to direct the drone's functionality. The AT commands are sent in a formatted string that allows the drone to interpret the instructions. The instructions are in the format:

AT*	*An instruction*	=	*Sequence number*	*Comma separated arguments*	Line Feed

Each string that is sent to the drone can contain more than one instruction. The instruction field houses the command that will be sent. The sequence number tells the drone which order to execute the instructions (begins at 1). Lastly, the comma separated arguments are used to give additional information to the instruction. (Piskorski *et al.*, 2012)

1.2.1 AT command summary

AT command	Arguments	Description
AT*REF	Input	Takeoff/Landing/Emergency stop command
AT*PCMD	Flag, roll, pitch, gaz, yaw	Move the drone
AT*FRTIM	-	Sets the reference for the horizontal plane
AT*CONFIG	Key, value	Configuration of the A.R. Drone
AT*CONFIG_IDS	Session, user, application ids	Identifiers for AT*CONFIG commands
AT*LED	Animation, frequency, duration	Set a led animation on the A.R. Drone
AT*ANIM	Animation, duration	Set a flight animation on the A.R. Drone
AT*COMWDG	-	Reset the communication watchdog

Table 3.1: AR Drone AT commands

1.3 Navdata

In order to receive navdata from the drone the user needs to instruct the drone on which data needs to be sent. The AT Command CONFIG is used to do this and allows the user to receive either a small set of the navdata or all of it. Once the AT command has been sent the drone will begin sending the information to the client. The client needs to create handles to manage the data. (Piskorski *et al.*, 2012)

1.4   Video stream

The video steam from the drone is managed by the video pipeline implemented in the SDK. The SDK takes the raw output from the drone, decodes the video and gets it ready to be displayed. The user can then add their own functionality into this pipeline so that the cameras can be used as an additional sensor. (Piskorski *et al.*, 2012)

2   SDK design

When an application is made using the A.R. Drone SDK, it will be created using a layered architecture approach. This allows for a separation of concerns and makes programming simpler for the client. This separation makes it easier for inexperienced programmers to interface with a more abstract code base, while more advanced users can access the controls at a lower level and send the AT commands directly. (Piskorski *et al.*, 2012)
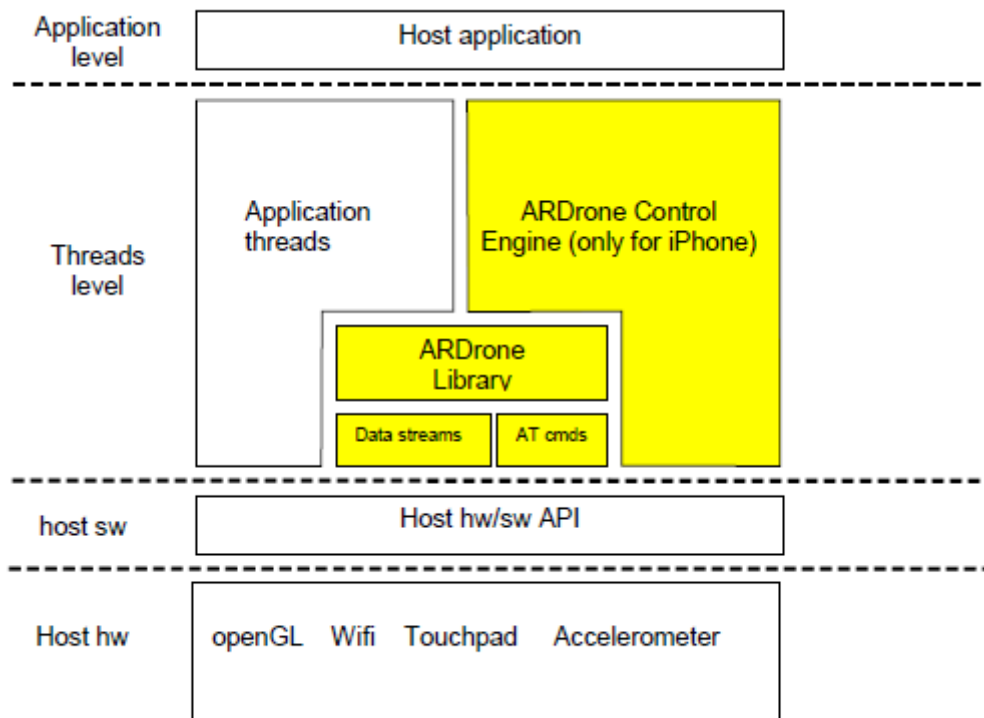
Figure 3.1: SDK layered architecture

(Piskorski *et al.*, 2012)

3 Limitations of the SDK

While the SDK allows for the majority of control to be conducted by the user, it does not allow for the on-board software of the Drone to be modified. This limitation is mainly in place to prevent the user from damaging the drone by disabling the flight stabilisation and other such functions that are core to the drone's ability to manage itself in the air. One disadvantage of this is that the drone's built-in height controller cannot be disabled. While at heights bellow six meters, low enough for the sonar to work, the drone automatically compensates to try and keep its height constant. This feature is intended to allow the drone to 'climb' up inclined surfaces. While this feature can prove useful when flying up hills, it can also prove problematic if a fixed height is desired, such as the recording of a relief map. (Piskorski *et al.*, 2012)

Chapter 4

Analysis of visual search algorithms

Search patterns are a strategic way of covering a given area in order to extract certain information. In order for the desired information to be extracted from an environment, search patterns need to be customised for the search area, the search vehicle as well as the type of search that is being conducted. There are many different methods of conducting a search depending on what the aim of the search is. Searches can be conducted on foot, from ground based vehicles or from aircraft, each of which have their own advantages and disadvantages. Given the context of an A.R. Drone, aerial search patterns will need to be analysed in order for the drone to perform a search pattern from the air.

1   Aerial searching

Aerial searches, as the name suggests, are conducted from above the ground utilising various aircraft. The most common aircraft used are either airplanes or helicopters. Airplanes allow for a large area to be covered from both low and high altitudes although not to the level of detail that can be achieved with a helicopter which is able to stop and hover over places of interest. With drones becoming ever more popular, their use in search and rescue has just begun to be realised. Due to the large array of drones available for use they can be customised to work in almost all environments. If an in-depth search is desired, a small drone can be used to fly over the ground at a very low altitude. If, on the other hand, a broad search for a large target is desired then a large drone with high altitude capability can be used.

2 Drone selection for search

When choosing a drone to conduct a search there are many questions that need to be asked in order for the best results to be achieved. Choosing the wrong drone for a job could result in the search turning up nothing or irrelevant data being returned.

A study by Waharte et al. investigated the possibility of utilising drones for search and rescue operations in an attempt to speed up the response time of rescue crews. Waharte identified four areas that may challenge drone operations for search and rescue, the quality of the on-board sensors, energy and endurance of small drones, navigation of unknown areas and the use of multiple drones. These four problems are discussed bellow:

Sensor quality - The higher the quality of the sensors, the faster the drone can cover an area. An example of this is if a drone has a high resolution camera, the drone will be able to give the same level of detail as a drone with a low quality camera from a much higher altitude. (Waharte & Trigoni, 2010)

Energy limitations - Most recreationally available drones have a very short operation time (usually in the order of a 10 to 30 minutes). This limits the range in which they can operate as well as the amount of sensors that can be attached before the power drain impedes the drone's ability to perform its tasks. (Waharte & Trigoni, 2010)

Navigating unknown environments - There are many environmental factors that could affect the drone's performance. Probably the most severe of these environmental factors is the wind that could blow a drone off course as well as into objects. There are also objects that could impede the drone's path such as trees and large rock outcrops. (Waharte & Trigoni, 2010)

Multi drone searches - If the area being covered is large then making use of several drones makes covering the area very efficient. However, the drones need to be aware of each other's locations, as well as the data that they are collecting. This induces the problem of whether the data should be correlated at a central server or whether the drones should share this data autonomously with each other. (Waharte & Trigoni, 2010)

Waharte also mentions several other constraints that could be a challenge to deal with but which would separate drones from traditional search and rescue vehicles. One of these constraints is that of vegetation. A large tree will cast a large footprint on the ground that an aerial search will not be able to see under. While this can be mitigated to some degree by an angled search, a large cluster of trees will still yield a blind spot. Using a

small drone here could yield a great advantage as, with the correct camera and sensor setup, the drone could be used to fly in between the trees leaving no ground uncovered. (Waharte & Trigoni, 2010)

An area that is not mentioned by Waharte is how the drones know where they are. A drone could make use of relative positioning given a known start location. This could be done based on a known speed of travel or by tracking how much ground is covered. However, given environmental factors such as a wind shifting the drone off its designated path or cambered ground these methods would not yield accurate results. A more elegant solution would be to use a GPS module so that the drone's location could be monitored precisely. Parrot has recently released a GPS module for the A.R. Drone that has an accuracy rating of $+/-$ 2 meters. (Parrot, n.d.a) With the accuracy that the GPS module combined with the height sensor the A.R. Drone is a candidate for high accuracy searching.

3 Search patterns

The objective of a search pattern is to minimise the amount of time it takes to find a given objective. Search patterns need to vary depending on what is being looked for, the location that is being searched, as well as the amount of information available on the location of the objective. In search and rescue there are six primary methods of search that are used depending on the goal of the search. (of Transportation, 2008)

While some searches are for a passive object, others are for objects such as an airplane or a boat which can make use of broadcasting tracking information on emergency channels. Not only do these broadcasts alert emergency personal that there is an emergency but they also assist in finding the target. For the purposes of using an A.R. Drone, the only method of tracking will be done visually utilising the on-board cameras. (of Transportation, 2008)

3.1 Route search

A route search is used to track something travelling in a given direction, for example, a vehicle travelling down a road or an airplane travelling in a known direction. The objective of this type of search is to cover a long thin area. Figure 4.1 shows how the search path runs along the most likely path of the target.(of Transportation, 2008)
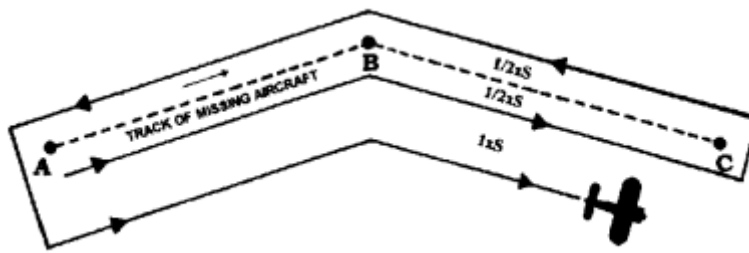
Figure 4.1: Route search

(of Transportation, 2008)

3.2   Parallel search

A parallel search is conducted if the area that is being searched is also rectangular, but the location of the target is unknown in the given area; for instance to detect a target in an open field where there is no locational information. (of Transportation, 2008)
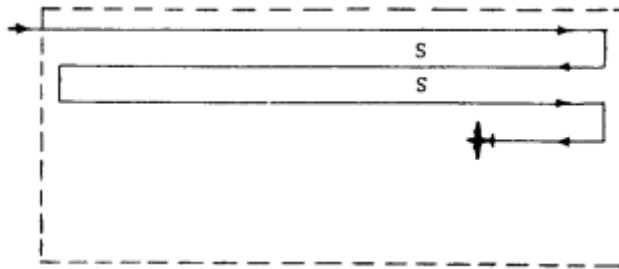
Figure 4.2: Parallel search

(of Transportation, 2008)

3.3   Creeping line search

The creeping search is similar to the parallel search except that it works on flying the shorter line of the rectangle first and working along the long side. This could be used in the instance where the general direction travel for the target is known. This search puts an emphasis on searching one side first and is usually performed when there is information about the likely location of the target. (of Transportation, 2008)

Figure 4.3: Creeping line search

(of Transportation, 2008)

3.4   Expanding square search

The expanding square search works on searching from a known start point and expanding outward. This search places an emphasis on finding the target close to the centre and less likely as the search expands. This search also allows for a second search to easily overlap the first to give a more in-depth search. (of Transportation, 2008)

Figure 4.4: Expanding square search

(of Transportation, 2008)

3.5   Sector search

The sector search is similar to the expanding square search but with a greater emphasis on finding the target at the centre. The sector search uses quite a complex pattern that requires preplanning on order for it to be executed correctly. This search pattern can be used when there are problems with viewing the target such as dense brush or low visibility as it allows for multiple passes to maximise the chance of finding the target. (of Transportation, 2008)
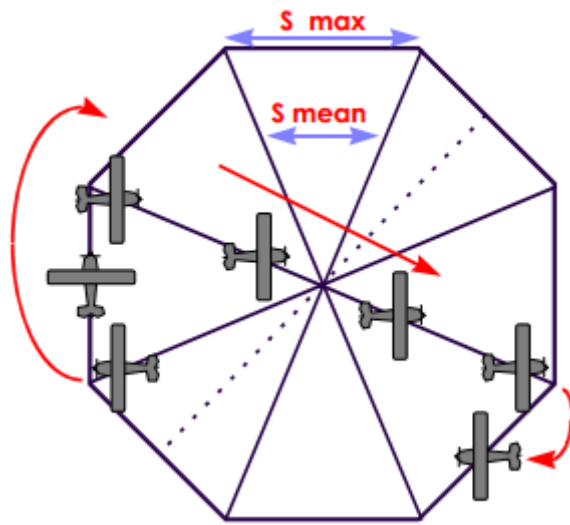
Figure 4.5: Sector search

(of Transportation, 2008)

3.6 Basic contour search

Lastly the contour search is used when the area being searched has a high degree of elevation such as a mountain range. Simply flying over a mountain range will increase the likelihood of the target not being sighted. With the contour search method, the search pattern is flown around the elevated land at increasing or decreasing altitudes. (of Transportation, 2008)
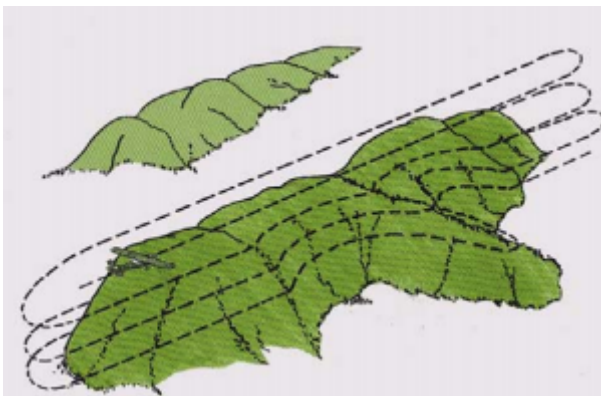

Figure 4.6: Basic contour search

(of Transportation, 2008)

4   Group searches

One of the ways in which a search can be improved is to divide the work among multiple drones. By making the area of search smaller for each individual drone, the speed at which the search will be conducted is increased as well as dealing with the problem of a small drone's limited energy. On a purely practical level, a single drone searching a large area is highly problematic. The drone can only fly at a given speed and so will take linearly longer as the size of the area increases. This, coupled with the drones limited energy and need to continuously come back to be recharged, could result in a search taking much longer than is necessary.

4.1   Search area division

The first problem that needs to be addressed when using multiple drones to search an area is how to divide up the area. Maza et al. performed research on dividing a polygon, given the co-ordinates of the points as well as the locations of the drones. Their work took into account the abilities of each of the drones with regards to range as well as the total area to be covered. Their algorithm was designed to solve the problem of a given area and the need to find the most suitable way of dividing up the work of the drones. (Maza & Ollero, 2007)

A study by Kolling and Kleiner took a second approach to this study and worked on the premise that the targets were aiming to avoid detection and could move at any speed. This type of search also requires the use of multiple drones. However, these drones will need to work together so that once an area is cleared, the target is unable to re-enter the area. The drones essentially need to block off the cleared area while continuing to cover the uncleared area. Kolling and Kleiner describe an algorithm for determining the minimum number of drones required for searching an area in which no target can escape. (Kolling & Kleiner, 2013)

4.2   Positional correlation and avoidance

One of the problems with having multiple drones flying in a given area is the problem of correlating their positions with each other so that a search is conducted thoroughly and they are able to avoid each other. The first problem can be easily solved by preprograming

the given areas of interest into the drones from a central control point. An article by Edmondson et al. suggest that this solution may not be efficient and that the drones should rather communicate and decide on their search areas collaboratively. With this method avoidance is implicitly controlled as the drones can constantly update nearby drones to their location. (Edmondson *et al.*, n.d.)

4.3   Information sharing and collection

While a search is being conducted the controller, as well as other drones in the area, need to be aware of what each other is doing. These updates could be that a target has been potentially sighted or that a drone is no longer operational. These updates need to be propagated between drones as well as back to the controller. Edmondson et al. made use of the drone's on-board wireless capabilities to allow for communication as well as allowing the drones to create a path that could link the furthest drone back to the controller. (Edmondson *et al.*, n.d.)

5   Mapping a search area

When flying a search pattern, it is important to record data on the area that is covered so that the same areas are not covered and so that data on possible target acquisition can be extracted. In the context of a drone that may be flying out of communication range of a human controller who is able to identify the target, it is important that the drone record images and scans of the area it has covered. If a drone is used to take pictures of the ground that it is searching, these pictures can be collated to form a map of the area that the drone has searched. The complete map can then be scanned for the target object after the drone has returned to a base station.

The principle of aerial cartography is to take a set of overlapping photographs of an area and then place them together to form a complete map. With the use of computers, this process can be automated to speed up the process of joining the images. The basic principle of joining images is to identify a set of matching pixels in two images so that the overlapping areas can be identified. This approach, however, is very resource intensive and too slow to be used reliably. A method of identifying key points of interest in an image was developed by David Lowe in 1999 and is called Scale Invariant Feature Transform or SIFT. SIFT is a lot less resource intensive and due to the fact that it only checks points of interest

where the pixel make up changes drastically, it has fewer points that it needs to correlate. These points are determined by using a Gaussian pyramid by repeatedly smoothing and subsampling the image. By matching only key points in an image SIFT offers a drastic improvement in speed over traditional methods of image matching. (Lindeberg, 2012)

A study done by Brown and Lowe into the efficacy of using SIFT to stitch together panoramic images attempted to create a reliably merged panoramic image given rotation, changes in light intensity and change in shooting angle. The first step in the process was to use SIFT to identify the points of interest in all the images. Brown and Lowe then utilised a random sample consensus (RANSAC) method to match the portions of the images that overlapped. The images were then adjusted for changes in rotation and camera angle changes. Lastly, the images were adjusted so that the lighting in the images was as consistent as possible. This method has proved to be highly efficient as well as reliable for stitching images together. (Brown & Lowe, 2007)

6 Target detection

Target detection is the process of sifting through the collated data from a search and detecting the desired target. There are many different methods that can be used to detect the target, from the controller visually identifying it, to running visual analysis techniques. Two methods of searching through the data automatically involve looking for a specific colour or, to improve the accuracy, a pattern. These methods can be used together or separately depending on the desired output required.

6.1 Colour matching

If object being searched for has a colour that is distinctly different from the environment that it is, for example, a red ball in a white room, the simplest way of identifying it is by its colour makeup. When an image is captured, it is captured as a matrix of pixels, each with an associated red, green and blue value that make up the displayed colour. By analysing these three values, the colour that is being searched for can be identified. In the instance of the red ball in the white room, every pixel should have the values (255,255,255) which correspond to white. When the red ball is found in the picture, the values will change to (255,0,0). This makes visual analysis by colour a simple yet powerful tool.

6.2   Template matching with glyphs

A more advanced method of identifying an object in an image is template matching. Template matching is the process of identifying a given pattern in a larger dataset. In image processing, the given template is what the user is looking for, while the dataset is the complete image that is to be searched. While it is simple for the human brain to adapt a template to apply to a new scenario, such as a model car and the real thing, this process is incredibly complicated for a computer. For the scope of this research, complete pattern matching in the interest of search and rescue has been omitted in favour of glyph recognition.

"Optical glyphs are two-dimensional graphic signs with square grid structure" (Redlarski *et al.*, 2013). Glyphs use a uniform grid pattern that is characterised by a white background on which sits a black bordered grid of black or white tiles (see figure 4.7). In order for the dimensions of a glyph to be calculated each row and column needs to contain, at least, one white square. While the patterns on a glyph can be random, they are often used to store data that is inferred from the shape of the glyph. (Redlarski *et al.*, 2013)

Figure 4.7: Glyph examples

(Kirillov, 2010)

The first step in utilising a glyph is to isolate it in an image. Redlarski et al. uses a set process to take a given image and extract a glyph image from it. The first stage, or pre-processing, is made up of three subprocesses. First the image is greyscaled, then edges are detected using differential edge detection and lastly Otsu thresholding is used to increase the contrast of the image. Having reduced the amount of information in the image, Redlarski et al. then identify potential glyphs by looking for white bounded quadrilaterals. These images are then extracted as potential glyphs and tested to ensure that they are glyphs. (Redlarski *et al.*, 2013)

Once a glyph has been identified and extracted, it can be compared to a database of know

glyphs. This can either be used as a final confirmation process that the given image is a glyph or to retrieve the associated meaning of the glyph.

Work done by Dementhon and Davis into the use of the Pose from Orthography and Scaling with ITerations (POSIT) algorithm proves that the orientation, or pose, of an object in an image can be calculated if the dimensions of the object are known. The algorithm relies on calculating an approximate pose of the object and using that to correct the perspective of the image. By calculating the pose, correcting the image and then recalculating the pose, the perspective of the image is eliminated and a scaled orthographic image is calculated. From this the pose can now be accurately calculated. (Dementhon & Davis, 1995)

This work has been extended by Oberkampf et al. to work with coplanar objects such as glyphs. In instances when all the points that are used are coplanar, such as glyphs, the POSIT algorithm returns two possible solutions for the pose of the object. Oberkampf et al. extend the POSIT algorithm to account for coplanar points by checking both possible solutions and calculating which is the most likely. The process used for checking the algorithm is done by creating a projection from the given plane and comparing it relative to other points in the image. In this way the correct pose can be calculated in instances when more than four points can be identified. (Oberkampf *et al.*, 1996)

Bibliography

Arjomandi, Maziar, Agostino, S, Mammone, M, Nelson, M, & Zhou, T. 2007. Classification of Unmanned Aerial Vehicles. *The University of Adelaide. Australia.[En linea] Disponible en: http://personal. mecheng. adelaide. edu. au/maziar. arjomandi/Aeronautical,* **20**.

Bristeau, Pierre-Jean, Callou, François, Vissière, David, Petit, Nicolas, *et al.* 2011. The navigation and control technology inside the ar. drone micro uav. *Pages 1477–1484 of: 18th IFAC World Congress,* vol. 18.

Brown, Matthew, & Lowe, David G. 2007. Automatic panoramic image stitching using invariant features. *International Journal of Computer Vision,* **74**(1), 59–73.

Cooper, Sean. 2010. Parrot AR.Drone hands-on: a quadricopter for the rest of us. Online. Available from: `http://www.engadget.com/2010/01/07/parrot-ar-drone-hands-on-a-quadricopter-for-the-rest-of-us/`. Accessed[14/04/2014].

Dementhon, Daniel F, & Davis, Larry S. 1995. Model-based object pose in 25 lines of code. *International journal of computer vision,* **15**(1-2), 123–141.

Edmondson, James, Cahill, Gene, & Rowe, Anthony. On Developing User Interfaces for Piloting Unmanned Systems.

GAJENDRAN, SRINIVASAN. 2012. Quadcopter Basics. Online. Available from: `http://creativetechno.wordpress.com/2012/06/13/quadcopter-basics/`. Accessed[16/04/2014].

Hoffmann, Gabriel M, Huang, Haomiao, Waslander, Steven L, & Tomlin, Claire J. 2007. Quadrotor helicopter flight dynamics and control: Theory and experiment. *In: Proc. of the AIAA Guidance, Navigation, and Control Conference,* vol. 2.

InvenSense. Triple Axis (X/Y/Z) MEMS Gyroscopes. Online. Available from: `http://www.invensense.com/mems/gyro/tripleaxis.html`. Accessed[17/04/2014].

Kirillov, Andrew. 2010. Glyphs' recognition. Online. Available from: `http://www.aforgenet.com/articles/glyph_recognition/`. Accessed[19/05/2014].

Kolling, Andreas, & Kleiner, Alexander. 2013. Multi-UAV motion planning for guaranteed search. *Pages 79–86 of: Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems.* International Foundation for Autonomous Agents and Multiagent Systems.

Krock, Lexi. 2002. Time Line of UAVs. Online. Available from: `http://www.pbs.org/wgbh/nova/spiesfly/uavs.html`. Accessed[07/04/2014].

Lindeberg, Tony. 2012. Scale invariant feature transform. *Scholarpedia*, **7**(5), 10491.

Maza, Ivan, & Ollero, Anibal. 2007. Multiple UAV cooperative searching operation using polygon area decomposition and efficient coverage algorithms. *Pages 221–230 of: Distributed Autonomous Robotic Systems 6.* Springer.

Oberkampf, Denis, DeMenthon, Daniel F, & Davis, Larry S. 1996. Iterative pose estimation using coplanar feature points. *Computer Vision and Image Understanding*, **63**(3), 495–511.

of Transportation, Washington State Department. 2008 (March). *WSDOT Aircrew Training Reference Text.* Washington State Department of Transport.

Parrot. Parrot AR Drone 2. Online. Available from: `http://ardrone2.parrot.com/apps/flight-recorder/`. Accessed[13/05/2014].

Parrot. Technical Specifications. Online. Available from: `http://ardrone2.parrot.com/`. Accessed[15/05/2014].

Parrot. 2012. TECHNICAL SPECIFICATIONS. Online. Available from: `http://ardrone2.parrot.com/ardrone-2/specifications/`. Accessed[17/04/2014].

Perdana, Julius. 2012. Win A Real AR Drone 2.0 - Paper Model Contest. Online. Available from: `http://www.paper-replika.com/index.php?option=com_content&view=article&id=7748:win-a-real-ar-drone-20-paper-model-contest&catid=144:ar-drone-2-&Itemid=200144`. Accessed[19/04/2014].

Piskorski, Stephane, Brulez, Nicolas, & Eline, Pierre. 2012. *AR Drone Developer Guide SDK 2.0.* 1.7 edn. Parrot.

Redlarski, Grzegorz, Pałkowski, Aleksander, & Ambroziak, Dominik. 2013. Optical glyphs based localization and identification system. *Pomiary, Automatyka, Robotyka*, **17**, 266–271.

Selden, Anthony. 2013. These Are The High-Tech Drones That You Can Actually Buy. Online. Available from: `http://elitedaily.com/envision/these-are-the-high-tech-drones-that-you-can-actually-buy/`. Accessed[14/04/2014].

staff, Defense Industry Daily. 2005. Predator Kills Important al-Qaeda Leader in Pakistan. Online. Available from: `http://www.defenseindustrydaily.com/predator-kills-important-alqaeda-leader-in-pakistan-0547/`. Accessed[07/04/2014].

Waharte, Sonia, & Trigoni, Niki. 2010. Supporting search and rescue operations with UAVs. *Pages 142–147 of: Emerging Security Technologies (EST), 2010 International Conference on.* IEEE.