Secure Location Verification with a Mobile Receiver

Richard Baker
University of Oxford
richard.baker@cs.ox.ac.uk

Ivan Martinovic
University of Oxford
ivan.martinovic@cs.ox.ac.uk

ABSTRACT

We present a technique for performing secure location verification of position claims by measuring the time-difference of arrival (TDoA) between a fixed receiver node and a mobile one. The mobile node moves randomly in order to substantially increase the difficulty for an attacker to make false messages appear genuine. We explore the performance and requirements of such a system in the context of verifying aircraft position claims made over the Automatic Dependent Surveillance - Broadcast (ADS-B) system through the use of simulation and find that it correctly detects false claims with a peak accuracy of over 97% for the most complex attack modelled; requiring only 75m of deviation between the reported position and the actual position in order for a false claim to be detected. We then report on our design for a mobile receiver and our construction of a prototype using low-cost COTS equipment. We discuss some additional benefits of incorporating a mobile node, examine the difficulties to be overcome and explore the applicability of the approach in other location verification use-cases.

1. INTRODUCTION

Cooperative surveillance systems are extremely widespread. Many planes, trains and automobiles provide an indication of their location and movement to either centralised controllers or to surrounding vehicles. Judicious use of the reported information enables centralised co-ordinators to manage traffic and respond to emergencies, and allows groups of vehicles to avoid collisions and operate more efficiently as a collective. As more and more 'smart vehicles' are produced and integrated with connected infrastructure, the burden on surveillance and tracking technologies looks set only to increase[5].

A concern with cooperative reporting systems is the difficulty in ensuring that reports are genuine, an issue that is far more serious when costly and potentially safety-critical actions are taken in response. However many widely-deployed cooperative surveillance systems have few privacy or security properties, allowing attackers both to learn substantial information about individual vehicles or group behaviours and often to impersonate or invent vehicles, which are then treated as genuine by receivers.

We present a secure location verification (SLV) system that verifies location claims using time-difference-of-arrival (TDoA); wherein the system compares the arrival times of the same message at different receivers to determine whether it has indeed originated from the claimed location. Crucially, at least one receiver is mobile and, by incorporating randomness into its movement patterns, substantially increases the difficulty for an attacker to have a false message incorrectly verified. The approach benefits from requiring little infrastructure investment and having applicability in a range of contexts.

In particular, in this work we:

- propose a lightweight, easily-deployed SLV system that requires no modification to existing systems
- evaluate the effectiveness of the system in a simulation
- describe a suitable architecture and document a prototype implementation
- discuss the system's application in a range of use-cases

2. BACKGROUND

2.1 Cooperative Surveillance

Cooperative reporting systems rely on individual nodes in the system to actively provide an indication of their state, such as their location, to receivers. In contrast with non-cooperative or "primary surveillance" systems such as traditional radar, which actively produce a means of tracking nodes, cooperative systems must rely upon each node to be able to determine its state accurately of its own accord and then report it. For position claims this means the node must be able to localise itself and then provide that location.

Cooperative systems are widely deployed. Air traffic management makes use of Automatic Dependent Surveillance - Broadcast (ADS-B) to track aircraft (discussed in more detail below). The Automatic Identification System (AIS) is a marine tracking system in which vessels determine their location and then broadcast it over a VHF radio link[14]. Sensor networks in which the geographic location of sensors gives meaning to their measurements need the sensors to report that location along with their collected readings. Similarly, knowledge of the locations of nodes is crucial when employing geographic routing or to assess distribution of sensors to ensure appropriate coverage[13]. Connected ve-

CPS-SPC'16, October 28 2016, Vienna, Austria

ISBN 978-1-4503-4568-2/16/10...$15.00
DOI: http://dx.doi.org/10.1145/2994487.2994497

hicles can report status including position via Vehicle-to-Infrastructure (V2I) links in appropriately-equipped cities, these reports can be used to enhance awareness of traffic patterns, as well as be used to implement road pricing and pay-as-you-drive insurance schemes[5][1].

Systems vary widely in their location accuracy, transmission range, frequency of reporting and between those that openly broadcast information in the clear (such as ADS-B or AIS), via those that establish transient links (such as connected vehicles forming platoons) through to those that provide only anonymous reports using secure channels. Where transmission is made by open broadcast, a significant security consideration is the possibility for a malicious party to report false information; to mimic other nodes, to invent fictitious nodes or to alter legitimate reports. Even where the channel is cryptographically secured to ensure message integrity, the problem is not alleviated; knowing that the received message is the same as the one that was sent still does not guarantee that it is an accurate representation of the physical state. Attacks can affect any dependent systems, whether they are centralised control or reporting systems, other nodes in the system or even human operators. Attacks on operators include overloading them with information so that genuine reports go unhandled, or even simply placing them under high levels of stress to increase the risk of human error occurring[18][7]. The open provision of information about individual nodes in the system gives rise to privacy concerns as well; particularly the easy tracking of nodes en masse.

2.2 Secure Location Verification

It is often necessary to have a means of verifying that location claims made by nodes in a cooperative reporting system are genuine. Verification can exploit any property of the transmission that is difficult or expensive for an attacker to influence maliciously; ideally prohibitively so. Common techniques include fingerprinting, distance-bounding, received signal strength measurement, angle-of-arrival and time-of-arrival.

Fingerprinting identifies operating features of the transmitter or protocol implementation that differ between classes of transmitter, or even individual units. If expected values are known for a given transmitter then computed features from a particular transmission can be compared to assess whether it is genuine. Distance-bounding techniques place an upper bound on the distance of a communicating party by using a challenge-response mechanism and measuring the round-trip time. With appropriate hardware the travel time of the signal is the dominant factor and the responder can be constrained to within a certain distance of the challenger. Received signal strength approaches are based on understanding of the attenuation of radio signals during travel. For a given type of environment the expected attenuation over various distances and from various locations can be estimated. With known transmission power the power of the signal at the receiver allows the verifier to identify ranges or locations from which the transmission could have been made. Techniques based on "angle of arrival" can be used to determine a direction to the transmitter and can be used either with ranging techniques or cooperatively in a "triangulation" arrangement to determine the transmitter location. Either mechanically-revolved directional antennae must 'sweep' an area or an antenna array must be employed and arrival times compared at each antenna to compute the angle of arrival. Time-of-arrival techniques measure the time that is taken for a signal to propagate through a medium and use those measurements to constrain possible transmission locations. In a pure time-of-arrival (ToA, or "trilateration") system the time of transmission is known and so the time-of-arrival at a receiver indicates the range from which the transmission was made. Where the time of transmission is not known, instead the can verifier note precisely when a signal arrives at a set of receivers at known locations. These "time-difference of arrival" (TDoA, or "multilateration") readings can then be used by the verifier to identify a set of locations from which the transmission must have come. With more receivers the verifier can constrain the set of locations further, eventually down to a single point[1].

2.3 Unmanned Aerial Vehicles

Unmanned Aerial Vehicles (UAVs or 'drones') have become extremely popular in recent years; employed in military and industrial roles and flown recreationally by individuals. A multitude of aircraft are available, being variously fixed-wing or rotor-wing, electrically or chemically powered and remote-controlled or autonomous. The continuing development of compact, inexpensive control and sensing equipment has made construction of UAVs far more widespread[16]; indeed the low cost of many designs has made UAV usage feasible in situations where an aerial system has previously been unavailable or too expensive to run, such as internal building survey, hobbyist photography or emergency rescue in hazardous areas[2]. Additionally, operating a UAV requires far less skill than a conventional aircraft so usage requires only minimal training and experience. Indeed the prevalence of UAV operators engaging in risky behaviour has prompted regulatory responses. The United Kingdom's Civil Aviation Authority (CAA) mandates that general operations must remain within 500m of the operator, below an altitude of 400 feet and in direct line-of-sight at all times. In December 2015 the U.S. Federal Aviation Administration (FAA) mandated registration of all operators of UAVs with a mass of over 250g[9], to improve operator accountability. The implementation of 'no-fly zones' for UAVs is already widespread, with various technical means such as 'geofences' being employed in an attempt to enforce them, albeit with mixed success. In addition, the CAA specifies tighter controls on "any aircraft which is equipped to undertake any form of surveillance or data acquisition"; mandating a separation of 50m from any person, building, vehicle or vessel not under the operator's control, in an attempt to mitigate the associated privacy risks[6]. Many UAV systems already report telemetry to their controllers that includes position and movement details and it appears likely that future regulations will require this information to be reported to controlling authorities as well.

Many UAVs are also capable of performing autonomous flights and there is considerable work to enhance this capability. Incorporating autonomy can not only reduce the risk of harm from an inexperienced operator, but also supplement operator judgement in hazardous situations (using extra sensors and collision avoidance systems) and reduce operator workload to enable several UAVs to be managed simultaneously by a single human controller. Sufficient technical advances in UAV autonomy and regulatory changes could enable use at greater distances and in far more complex environments, such as for delivery or survey purposes[23][17].

2.4 Air Traffic Management

Air traffic management attempts to ensure safety is maintained during aircraft manoeuvres in the governed airspace, whilst maximising the efficiency of traffic movement. Air traffic controllers maintain communications [1] with aircraft as they operate and provide advisory information or mandatory instructions to pilots. In addition, air traffic control (ATC) operations have long made use of surveillance technologies to assist in tracking aircraft for which they are responsible. Originally this function was performed by non-cooperative means via "primary surveillance radar" and this is still common in military uses. However primary surveillance technologies provide limited information beyond the range and bearing of an object and require considerable infrastructure and expense on behalf of the ATC operators[24]. Practice is moving away from traditional active surveillance and towards a cooperative reporting model in an effort to enhance situational awareness for all parties and accommodate greater traffic; so-called "secondary surveillance radar" (SSR). In this model, aircraft broadcast a variety of status information, which can be received by ground stations and other aircraft. One such SSR, already widely deployed in Europe and the United States (and mandated for use by 2020 by regulators in both jurisdictions) is the Automatic Dependent Surveillance - Broadcast (ADS-B) system[8]. Aircraft are equipped with a transponder that periodically reports details of the aircraft's status in the ADS-B format. ADS-B makes use of one of two possible data links; with the standard for civil aviation being 1090MHz Extended Squitter (1090ES). Earlier SSR systems made use of the 1090ES data link and so its use for ADS-B allows integration with existing transponder equipment. Messages can variously indicate speed, position, callsign, climb rate and emergency status[2], although only position reports are considered in this work. All messages contain an ICAO identifier, a unique 24-bit number identifying the aircraft. The interval between messages varies depending upon the message type; for position messages it is approximately 0.5s[21].

3. RELATED WORK

Location estimation with mobile nodes is described by Luo et al., using a similar approach to that described herein (termed "Mobility-Differentiated Time Difference of Arrival") and applied to the issue of surveying sensor networks for node displacement[13]. However the approach is node-centric (i.e., the node is attempting to estimate its own location) and does not aim to be secure, in that it does not consider adversarial behaviour in the location estimation process.

Perazzo et al. describe a roving verifier that determines whether nodes in a sensor network have been displaced and describe an algorithm to construct a near-optimal route for the verifier to take to conserve fuel. However the localisation approach assumes cooperation from the sensor nodes and hardware to enable a distance-bounding protocol to be used[15]. Čapkun et al. proposed a system that exploits the difficulty for an attacker in claiming a false location when the verifier is moving in an unknown way[3]. However the location verification protocol is again cooperative.

Strohmeier et al. have proposed a location verification system that makes use of widely-distributed, low-cost receivers to improve the coverage of TDoA verification over existing, professionally-deployed systems. They overcome the poorer accuracy of low-cost receivers by collecting multiple messages and testing them together; comparing them statistically to expected values. However, they depend on widespread deployment of receivers and a prior training phase to produce fingerprints that message timings can be compared to. We instead attempt to make the verification process secure solely through the use of mobility; with no prior system training and substantially fewer receivers.

Schäfer, Lenders and Schmitt describe a static multi-receiver ToA approach that considers a 'track' of position claims and uses it to compensate for individual errors to improve verification accuracy. However the model described therein assumes an accurate prover-local timestamp is sent with the message[19] and the authors note that while this capability can be realised from the ADS-B standard, deployment of systems that broadcast on such a dependable timescale is not widespread. Position claims are otherwise sent with much less accurately-measured periodicity.

4. ATTACKER MODEL

Attackers belong to one of three classes, each modelling a different attack case. Attackers broadcast, with an identical capability to legitimate sources, on a defined interval and all messages are assumed to be received successfully. No attacker makes any attempt to determine the location of the mobile receiver. The attacker classes are as follows:

Static attacker.

The simplest attacker class; a static attacker is broadcasts from a fixed location but claims to be in another. This class exemplifies an attacker inventing a 'ghost' vehicle from their home or a parked vehicle nearby.

Mobile attacker.

A mobile attacker is a more complex version of the static attacker. In this case the attacker broadcasts from an airborne mobile platform, modelled as a UAV. This class covers an attacker using a UAV to attempt to fool the verification by legitimately moving the transmitter, but while claiming to be another (perhaps much larger or more important) vehicle, or simply mounting their receiver on a UAV to avoid being caught and the equipment confiscated. It could also model a drone operator flying outside of a permitted area whilst claiming to be inside.

Course deviation attacker.

A course deviation attacker initially broadcasts correct position claims, but then selects a new course and attempts to hide this behaviour by falsely claiming to still be following their old course. For air traffic this attacker class can be seen as an aircraft that has been seized by force or maliciously diverted. Less morbid alternatives exist in other scenarios, such as an attacker in a road vehicle briefly exceeding a speed limit or using a prohibited lane, while continuing to report seemingly law-abiding behaviour.

[1] The primary communication method is currently voice, although this will be replaced by data communication in the near future

[2] Among many other pieces of status information

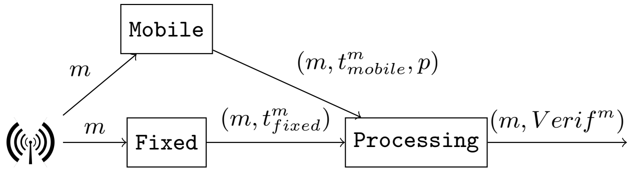

Figure 1: Structure of verification system

5. SYSTEM MODEL

In essence our approach consists of three elements: a mobile verifier node, a fixed verifier node and a processing unit. Both verifier nodes collect the broadcast messages with claimed locations and record them along with the precise time of reception. The mobile node additionally records its location at the time of reception. The recorded (message, time, location) triplets are then sent via an out-of-band channel to the processing unit, which matches claims from each receiver by means of a short-term unique identifier for the message. Upon receipt of matching claims from both nodes, the processing unit performs a TDoA calculation.

Figure 2 demonstrates an example operation of the system in two dimensions as it evolves through time. In this case, the source $S_{0\to 2}$ moves on a linear course, broadcasting messages at a regular interval, while the mobile node $Mobile_{0\to 2}$ moves randomly during the same period. When messages are received by both the fixed and mobile nodes, given that the locations of each node are known, the difference in elapsed time for a message to reach each receiver can be used to compute the possible source locations, which are given by the hyperbolae shown. In each case, only one side of each hyperbola is plotted, as it is determined by which node receives the message first (the fixed node in this example). The expected time difference from the claimed position can also be calculated easily and any discrepancy from the measured difference determined. If the discrepancy is within a defined acceptance threshold α then the message is considered genuine, otherwise it is taken to be false and flagged as such; enabling it to be reported to operators or downstream control systems.

More explicitly, for a position claim m, the distance from it to the two known verifier node positions can be calculated as the Euclidian distance in each case; d_{fixed} and d_{mobile}. Then the expected time difference for the claim $\Delta^m_{expected}$ can be calculated by finding the absolute distance difference and dividing by the propagation speed c (i.e. the speed of light).

$$\Delta^m_{expected} = \frac{|d^m_{fixed} - d^m_{mobile}|}{c}$$

Meanwhile the actual measured time difference Δ_{actual} can be obtained easily from the recorded values at each node.

$$\Delta^m_{actual} = |t^m_{fixed} - t^m_{mobile}|$$

Verification of the claim x is then a matter of comparing the deviation of the actual time difference from the expected time difference against an acceptance threshold α, to produce a result $Verif^x$ that can be output.

$$Verif^m = \begin{cases} Accept & |\Delta^m_{expected} - \Delta^m_{actual}| < \alpha \\ Reject & |\Delta^m_{expected} - \Delta^m_{actual}| \geq \alpha \end{cases}$$

The mobile node is, of course, mobile. It randomly and independently makes changes to its course. These are not pre-determined nor known by any other party. Furthermore, they are not transmitted prior to the course being followed. Upon initialisation, the mobile node begins listening for position claims and commences its movement; it selects a random waypoint and begins moving towards it, upon reaching the waypoint it selects a new one and repeats the process.

It is highly situation-dependent what action should be taken in the case of a claim being falsified. Various reporting, filtering or enforcement options exist but these are beyond the scope of the verification system itself. Some instances are discussed in Section 10.

6. SECURITY ANALYSIS

As with any verification system, an attacker can attempt to fool the verification process, attempt to subvert the verification system or attempt to disable it completely. In this section we consider each in turn and consider the implications for our proposed system.

An attacker attempting to fool the verification process clearly wishes their messages to appear genuine. The capabilities of an attacker are considered in great detail in [18]; noting that in general they can perform message injection, deletion or modification.

In an attempt to have their injected messages accepted as genuine, the attacker can modify their transmission timing. This requires that they know the location of each receiver, such that they can broadcast messages that arrive at each receiver with a time difference consistent with the claimed location. As in Section 5 above, we consider this in two dimensions for ease of discussion, although the generalisation is straightforward. Location verification performed at a single point in time by multiple receivers, wherein the TDoA measurements all refer to the same time of transmission, can detect transmissions that are not made at a single point when there are sufficient ($n \geq 3$ in 2 dimensions, $n \geq 4$ in 3 dimensions) receivers. By contrast, the approach described here only ever uses two nodes for each verification. Therefore in two dimensions it cannot constrain an attacker to less than an entire side of a hyperbola. A message transmitted from any point on that side will appear to be genuine. The security of this approach lies in the fact that as the mobile node moves relative to the fixed one, the hyperbola of possible genuine positions sweeps across a substantial distance, this is particularly noticeable between the second (dashed) and third (dotted) lines in Figure 2. For each message the attacker must accurately determine the location of the mobile node, construct the hyperbola and then transmit from a location on it in order for their message to be accepted as genuine. As the movement of the mobile node is randomised, the attacker cannot predict its location and so must monitor it and move accordingly, with sufficient speed to 'catch' the arc of the hyperbola. Where the attacker is additionally required to transmit regularly, the maximum permissible interval between transmitting claims bounds how quickly the attacker must complete this process. In three dimensions with only two nodes, the location is constrained instead to one sheet of a hyperboloid. The attacker therefore has another degree of freedom for their transmission location, but must still engage in the same reactive behaviour and must now contend with the movement of the mobile node in another axis as well.

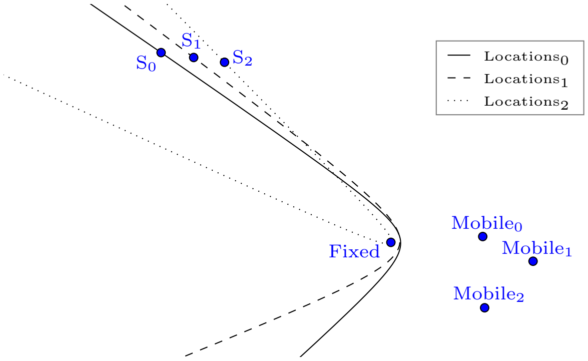

Figure 2: A two-node TDoA system in two dimensions at three successive points in time (as denoted by subscript).

As mentioned above, there are substantial sources of error in the measurements. Even once compensation mechanisms have been applied, there must still be some tolerance for the remaining error. But just as the error tolerance accommodates spurious measurements, it also allows greater accommodation for an attacker to claim a false position. In effect, the error tolerance can be seen as giving the hyperbola or hyperboloid some 'thickness'; expanding the area that an attacker can occupy and still have their claims validated. Figure 3 shows the effect of errors in measuring the transmission time and in localising the receivers. It is already noticeable at S, a relatively short distance from the receivers, but also grows as the distance between the source and the receivers increases.

The problem becomes much more pronounced when the receivers are separated by only a small angle relative to the transmitter. In this case the errors from each receiver overlap, in a phenomenon known as "geometric dilution of precision" (GDOP), resulting in a larger error region. An attacker will benefit from locating themselves in an area of high GDOP, where discrepancies in their positioning will be tolerated the most. However in our approach the areas of high GDOP also change randomly; just as an attacker has difficulty in predicting where they must locate themselves, they also experience difficulty predicting where the areas of high GDOP will be at any point in the future. However, an attacker that is able to make the angle between receivers small relative to themselves can maximise the overlap for a given distance and error tolerance, whilst an attacker that can move further away from the receivers with the angle between them kept small can expand the overlapping region in absolute terms. Depending upon the error tolerances in the system, the space created with such a technique can be large enough to allow a malicious party to perform a useful attack where messages are spoofed for a more distant claimed location. This limitation must be considered carefully when deploying the system. With an airborne mobile node, our system does have a distinct advantage over purely ground-based systems in combating this situation. The greater variation between receiver altitudes reduces the GDOP in the vertical dimension, thus somewhat restricting an attacker's use of altitude to engineer sufficient distance for high GDOP. The vertical range of the mobile node may well be limited, but as long as it is greater than the surrounding terrain elevation, the system benefits from the mobility.

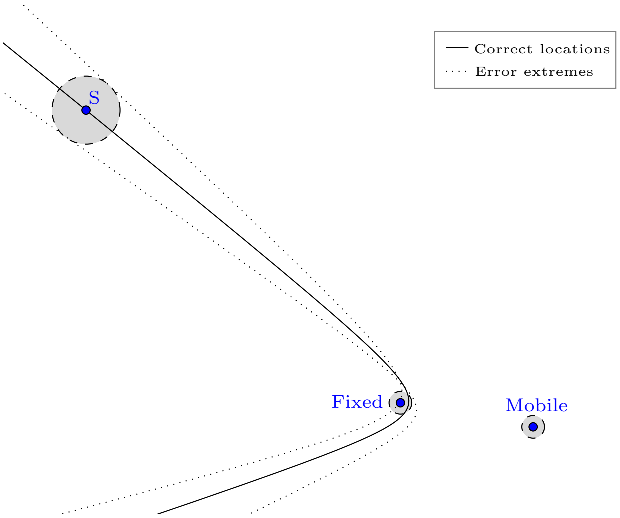

Figure 3: A two-node TDoA system in two dimensions at a single moment in time, with the errors for transmission time measurement and receiver localisation visualised.

Message deletion is possible primarily by selective jamming. As [18] notes, destructive interference (wherein the signal is inverted and superposed onto the original) is extremely difficult to accomplish for a moving aircraft, so we do not consider it here. Constructive interference (wherein a noise 'spike' is injected during message transmission such that the message is corrupted) is feasible and the system presented herein provides no defence against such attacks. Message modification is a derivative of the injection and modification approaches. Again, applying interference to alter a message in-flight is hard to achieve, however a message can be observed, interrupted with constructive interference and re-injected. The normal operation of the verification system for the message injection then applies, with the attacker's timing additionally delayed by having to observe the original message before transmitting.

Aside from attempting to fool the verification system, an attacker can attempt to subvert its operations or render it inoperable entirely. Jamming of the reporting link between the mobile node and the processing unit represents an appealing prospect for an attacker (likewise the reporting link from the fixed node if it is wireless). This would prevent the system from performing verification until either the jamming ceased or the mobile node moved to a location from which it could overpower the jamming signal (e.g. returning to the fixed node, if such functionality were implemented). Moreover the attacker exposes themselves to detection by doing this. Alternative transmission means such as free-space optical (FSO) communication, which are far harder to intercept and jam, could provide a practical countermeasure to this attack in the future. An attacker that is willing and able to take such a blunt approach could instead jam position reporting outright, although this is expensive and risky to do at scale. This system cannot (and indeed makes no attempt to) stymie such an attack.

Spoofing of messages on the reporting link is not a realistic concern however; messages can be encrypted and signed well within the constraints of affordable, portable hardware and the establishment of keys carries little operational cost in a system of only three components. An attacker who spoofs a GPS signal to a legitimate aircraft and causes it to report a false position will be detected indirectly as the aircraft's claims will be falsified, but the ultimate cause of those false claims will not be directly revealed. Detecting GPS spoofing is an area of active research and the likelihood of feasible solutions being found is high[22].

An attractive attack against TDoA systems is the use of directional antennae to break the assumption that each receiver is detecting a message broadcast at the same time. By transmitting each message to only a single receiver, the attacker can apply different time offsets to each identical copy, such that the arrival time differences are consistent with the claimed location. This task is made far more difficult with a mobile node as the attacker must track the mobile node with one antenna and alter their timings based on its position. While logistically this is easier than having to physically move the transmitter for each message, the underlying problem of locating the mobile node remains the same.

A mobile node, of course, needs power and as such its activity will be limited by energy constraints. If the mobile node simply stops while it refuels then the attacker has an open window in which to launch an attack. Some policy must be employed to overcome this, such as using a number of identical mobile nodes (albeit at greater cost).

The length of the mobile node's reporting period determines the maximum time in which an attacker is guaranteed not to have been detected. As such the selection of an appropriate reporting time is crucial in limiting how long an attacker can perpetrate a falsely-claimed track. It also determines how often an attacker can attempt to localise the mobile node using the transmission to assist with a timing attack. A short reporting period keeps the maximum 'guaranteed-undetected' period short, whilst a long reporting period gives the attacker only low-resolution location-estimation capabilities, as well as reducing transmission overhead and associated power consumption. The system is agnostic to any selected reported period and so this factor could be scaled with traffic as necessary and randomised to make the task of the attacker more difficult. The selection of this value represents an important factor in adjusting the performance characteristics of the system to suit the use-case.

7. PRACTICAL CONSIDERATIONS

Typically, the fixed node would be an existing surveillance system receiver, with the processing unit co-located, while the mobile node would be a UAV. There is clearly a requirement that the fixed node and the mobile node should have overlapping reception areas, but there are no other location restrictions on components.

Messages do of course need to be successfully received in order for the system to work. As such it is suitable only for broadcast systems in which the receiver can directly sense messages on the transmission medium. If the report is, for example, relayed en-route to the receivers then the verification results will clearly be incorrect. Similarly while the verification system does not in theory need to know the content of messages in order to be able to record their reception times, in a situation where only a small fraction of traffic is relevant to the verification system (e.g., location claims being made by a mobile device over a 4G connection along with other traffic), the system cannot determine which messages are relevant and must timestamp each one, substantially reducing efficiency.

It must also be possible to match messages using some unique identifier. The parameters of uniqueness in this context are very weak however; the identifier need only be unique for the period between collected claims being reconciled at the processing unit. A target-specific identifier such as a MAC address, callsign or results of transmitter fingerprinting and a message-specific identifier such as a sequence number, hash of message contents or even a sufficiently-precise position claim itself can be appropriate here.

The mobile node will seldom follow the defined course exactly; both due to limitations on the movement accuracy of the mobile platform itself and environment factors such as high winds blowing it off-course. Crucially, tight adherence to the prescribed course is not necessary however, as only the position at message arrival is required for the verification step. The course-following behaviour is simply to create a challenge for an attacker to predict the location of the mobile node and use that information in their attack.

In practice the TDoA calculation is affected by substantial sources of error in measurement and timing, which must be compensated for or tolerated. Errors in the system are modelled following the approaches in [19] and [13]. Two sources of error are considered; clock drift ϵ_{drift} modelled as a linear progression with coefficient t_{drift} applied over any interval between times t_i and t_j and normally-distributed measurement error $\epsilon_{measure}$ that is independent for each measurement.

$$\epsilon_{drift}^{j-i} = (t_j - t_i) \cdot t_{drift}$$

$$\epsilon_{measure} \sim \mathcal{N}(0, \sigma_{measure}^2)$$

So an observation of the current time t_{obs} will include clock drift error applied since the start of the system at $t = 0$:

$$t_{obs} = t + \epsilon_{drift}^{t} + \epsilon_{measure}$$

In contrast with the formulations in [19] and [13] however, we model the clock error differently. The fixed node is assumed to have a high-accuracy clock with negligible drift. The mobile node does not have the luxury of such an accurate internal clock and so experiences noticeable but still linear drift. However this clock is disciplined at a regular interval c_{sync} by a more accurate time synchronisation signal and returns to the correct time, with only some far smaller error in measuring the synchronisation signal $\epsilon_{sync}^{t} \sim \mathcal{N}(0, \sigma_{sync}^2)$. Such a model is consistent with inexpensive GPS-disciplined oscillators that output a pulse-per-second (PPS) output. The clock drift is only in effect during the interval between clock corrections.

As such the clock drift error for a time interval between t_j and t_i is limited to the drift rate applied since the last synchronisation interval plus the synchronisation error:

$$\epsilon_{drift}^{t} = t \bmod c_{sync} \cdot t_{drift} + \epsilon_{sync}^{c_{sync}}$$

Each time the clock is disciplined, the measured drift ϵ_{drift}^{t} is included in a moving average $\overline{\epsilon_{drift}}$. When the mobile node receives a position claim message, it applies a drift compensation based on the average measured drift and the resulting value t_{rec} is recorded.

$$t_{rec} = t_{obs} - t \bmod c_{sync} \cdot \overline{\epsilon_{drift}} + \epsilon_{measure}$$

The mobile receiver also experiences some error in localising itself in order to record its own position when messages are received. In most cases this error is small compared to that introduced due to timing or measurement. As such it is not modelled here.

8. EVALUATION

Primary evaluation of the proposed system was conducted in a simulation, testing verification performance against various attack types and with a selection of property values employed. The use-case was taken to be verification of aircraft position claims broadcast in ADS-B messages to be received by air-traffic control stations.

The simulation modelled the operation of the static and mobile verifiers and a number of attackers. Position claim data obtained from the OpenSky Network[3] [20] were used to provide real flight tracks of legitimate aircraft. The data were captured from a receiver at the University of Oxford Department of Computer Science over a 24-hour period on the 13th June 2012. There were 392,549 position messages in total, covering 1,088 ICAO identifiers with anywhere from a single message to 4,240 messages per identifier. The locations were all treated as genuine and the original transmission times were computed from the reception times by subtracting the propagation delay between the claimed location and the receiver. An attacker was then added, selected from one of three classes. The attacker's false messages are combined with the genuine ones and the resulting dataset fed into the processing unit.

For the purposes of the simulation a static attacker is placed at a random location on the ground, with a randomly-selected flight track that they attempt to mimic. A mobile attacker is airborne, using a rotor-wing UAV, and following their own random flight path. A course deviation attacker is instantiated as an airborne aircraft with a randomly-selected course along with a randomly selected deviation from it up to a maximum of 10% alteration in heading or pitch. The attacker begins by correctly reporting their position but then moves further away from their claimed track as the simulation progresses.

The mobile node was modelled as a rotor-wing UAV and assumed to have already completed its initialisation phase and was instantiated at a distance of 2km from the fixed node, stationary at 250m above the ground. Throughout the simulation the mobile verifier changes course on a randomised interval, selecting a new direction, speed and interval until the next change.

All aircraft were assumed to be in range of the receivers for the mobile node and the fixed node at all times, no reception range limits were applied. Similarly, message loss is substantial for Mode-S transmissions; the ADS-B data used here would suggest only a 7.54% detection rate for position claims. However message loss was not modelled in this simulation. This is not as strong an assumption as it may first appear; each claim is verified individually so a reduction in the number of messages received does not affect the verification itself, only the number of times that it happens.

The acceptance threshold was initially set to allow 100m of deviation; that is $\alpha = \frac{100}{c}$ (for convenience, thresholds are quoted as distance equivalents throughout this paper). As mentioned above, all claims were treated individually; no aggregation of data at a flight level was undertaken.

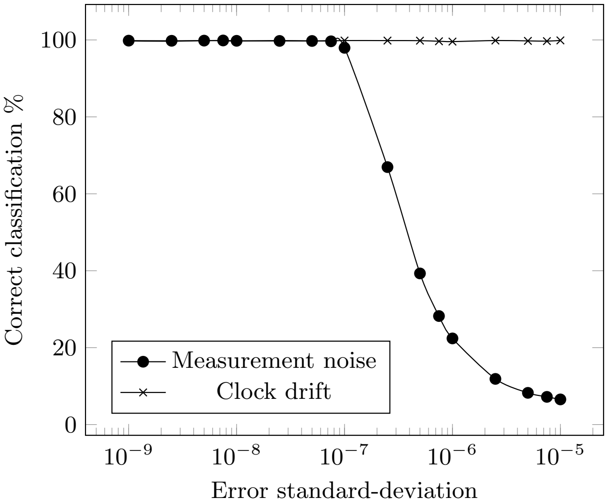

Figure 5: Effect of changing measurement noise ($\epsilon_{measure}$) and clock drift (t_{drift}) standard deviation on detection rates

Figure 4a shows the detection performance of the verification system in simulation against a single instance of each attacker class, as the detection threshold α is varied. For all attacker classes, peak detection occurs at α values equivalent to distances between 50m and 100m. The peak correct classification rates are 99.8% at 100m for static attackers, 99.8% at 100m for mobile attackers and 97.3% at 75m for course-deviation attackers. The most obvious result is that a more permissive detection threshold increases the correct classification rate in almost every case until the threshold becomes unreasonably large. To understand this behaviour one must consider the scale of the air-traffic scenario; ADS-B position claims can be detected hundreds of miles away. Even a threshold on the order of several kilometres still gives an adversary comparatively little area in which to mount their attack compared to that available without the verification system. In this scenario the primary determinant for α is overcoming the sources of error in the system to avoid false negatives. In a deployment around a specific target (such as an airfield), the selection of α would be more constrained. Unsurprisingly, for this reason, performance for static and mobile attackers is near-uniform throughout. While using a mobile platform may make an attacker harder to physically stop, it does little to assist them in making false claims potentially many kilometres away appear genuine. Peak classification of course-deviation attacks is notably lower. This can be attributed to the far smaller variation in position that appears with this attacker class; when the course-change is sufficiently small it will not breach the threshold until the new course has been maintained for some time, hence causing more position claims to be misclassified. In all cases the performance begins to diminish again as α grows to the kilometre scale and beyond. At this level of permissiveness, more false claims are treated as genuine and the higher false positive rate is the cause of the reduction in accuracy.

The relationship between the false positive rate and the true positive rate is shown in Figure 4b, a plot of the receiver operating characteristic (ROC) curve for each attacker class.

[3] A collaborative ADS-B reception and recording initiative, intended to provide data for ATC research

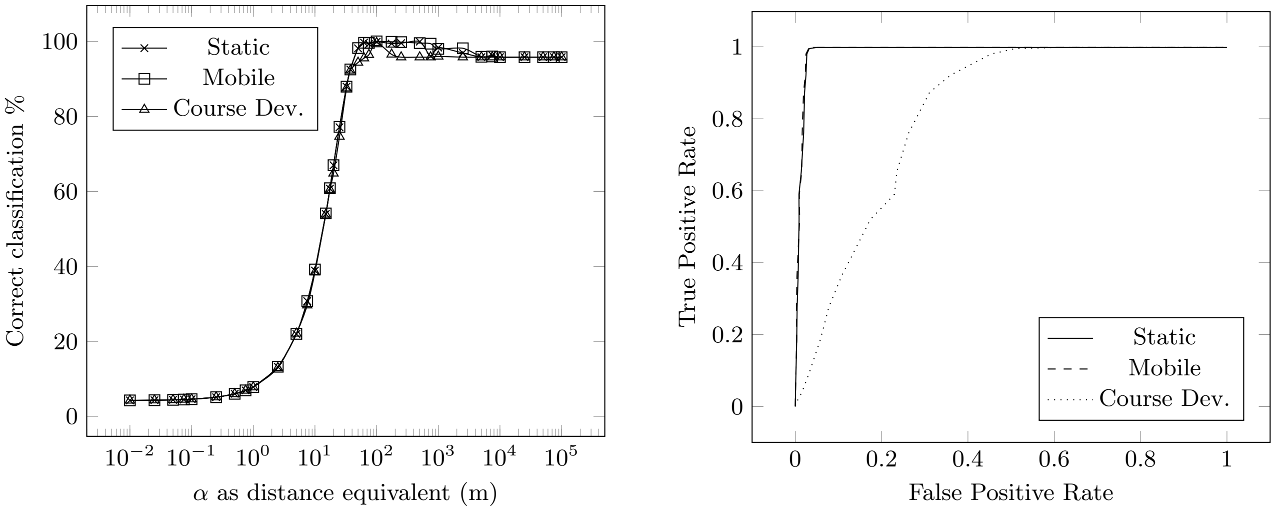

(a) Effect of changing α on detection rates for each attacker class

(b) Receiver operating characteristic (ROC) curve for detector

Figure 4: Classification performance of the verification system. Sub-figure 4a shows detection rates against the three attacker classes, while Sub-figure 4b shows the receiver operating characteristic for the detector.

Here the difference between the static and mobile attackers and the course deviation attacker is particularly noticeable. The potentially enormous difference between a static or mobile attacker's real position and an arbitrarily-chosen one in the sky causes them to be detected reliably until the detection threshold is made unreasonably large. The systematically smaller position difference for course deviation attackers makes avoiding detection more achievable, especially where the selected course change is small.

The effects of the primary sources of error were also explored and are visualised in Figure 5. The measurement error $\epsilon_{measure}$ was varied between 1ns and 10μs, with a single static attacker and an acceptance threshold equivalent to 100m. As would be expected, lower measurement error helps the system to make correct classifications, but only up to a point; further reductions in value below 100ns have little additional effect. Above this level performance falls away substantially, almost entirely due to a sharp increase in the false negative rate as the measurement error approaches the scale of the detection threshold. By contrast, there is almost no effect upon classification performance as the clock drift grows. This is explained by the clock error being sampled only at the start of the simulation and then assumed to drift consistently by the same amount. Under these assumptions the compensation strategy proposed in Section 7 can easily accommodate the drift.

9. PROTOTYPE

We designed a receiver platform for the mobile node and implemented a prototype using extremely inexpensive commercial-off-the-shelf (COTS) equipment. Figure 6 shows the architecture of the receiver platform. There are three main detection subsystems: a Radio receiver to detect position messages from the channel, a Clock to provide a timestamp on each message arrival and a Localisation system to determine the position of the mobile node when a message is received. A central Collector is responsible for recording these values together and storing them until the next reconciliation, at which point they are passed en masse to the processing unit via a Downlink.

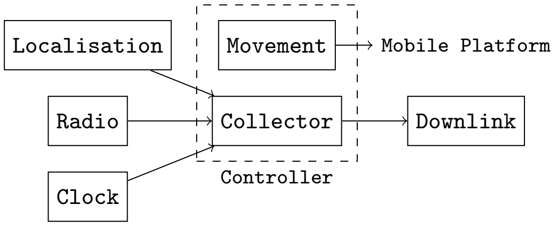

Figure 6: Mobile receiver node architecture diagram

The Radio subsystem could be implemented with off-the-shelf message reception hardware for the protocol being observed (e.g., a Wi-Fi or ZigBee receiver) or with a customisable software-defined radio (e.g., a USRP or bladeRF unit) for increased adaptability. The requirements here are simply that messages are provided to the collector in a timely and predictable fashion in order to avoid introducing additional error and that the detection resolution of the receiver is sufficiently high that the measurement error is kept small. The Clock subsystem can be any sufficiently-precise clock available, either internally or via an external device. The clock should display minimal drift even in the presence of changing environmental factors as a result of movement (such as variations in temperature and pressure), or at least display predictable drift and a means of reporting it. As per the clock error model detailed in Section 7, multi-stage apparatus are a potential choice, such as a local oscillator that is periodically disciplined by a time signal from a global navigation satellite system (GNSS), such as GPS. Any localisation approach that provides sufficient availability and precision is a suitable candidate for the localisation subsystem. Use of a GNSS is the most obvious choice. The sensor subsystems need not be implemented completely independently; for example in the air-traffic management scenario a Radarcape ADS-B detector could be employed as it in-

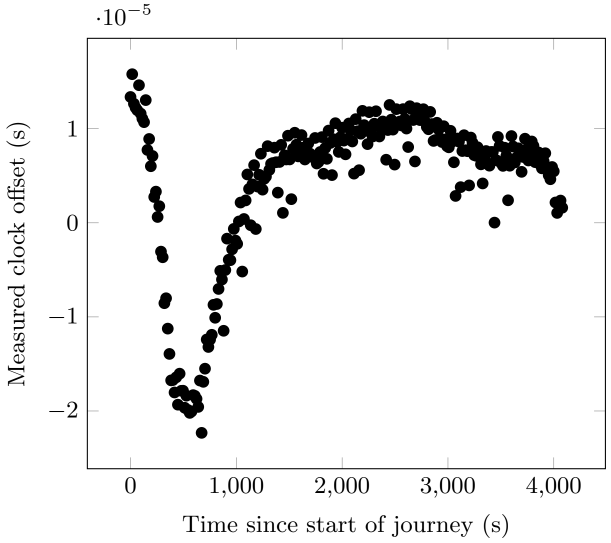

Figure 7: Clock offsets over time

corporates message decoding, high-resolution timestamping and positioning capabilities together, requiring the Collector to simply record the output.

The Collector subsystem runs on an onboard computer attached to the sensor subsystems and to the mobile platform itself, denoted as the Controller. The Controller also runs a movement planning algorithm that is responsible for orchestrating the random movement of the mobile node in conjunction with the underlying mobile platform's control systems. The Downlink can be any suitably long-range, secure connection with the fixed node, such as a common 4G modem.

Our prototype implementation was constructed using off-the-shelf components at a cost of less than £100 ($155). The hardware consisted of a Raspberry Pi 2 Model B acting as Controller and providing the Clock subsystem, a NooElec NESDR Mini 2 dongle (using a Rafael Micro R820T2 tuner and Realtek RTL2832U demodulator chipset) to form the Radio subsystem and an Adafruit Ultimate GPS Breakout v3 (using an MTK3339 GPS module) acting as the Localisation subsystem and assisting the onboard clock. The Raspberry Pi ran the Raspbian Wheezy Linux distribution with a v3.18 kernel and the pps_gpio kernel module to accept a pulse-per-second (PPS) signal. With this capability, the GPS unit not only provided accurate location data for the mobile node, but also acted as a timing source to discipline the Raspberry Pi's internal clock by being configured as a Stratum-0 source for the local Network Time Protocol (NTP) daemon. At the start of every second the PPS signal raises an interrupt to correct the clock if necessary. This corresponds to the time drift model discussed in Section 7, with the manufacturer of the GPS unit quoting a 10ns jitter for the PPS timing signal[11].

Capture of ADS-B messages was performed using a Linux port of the dump1090 utility[4] with extremely minor modifications. Upon receipt of a message the standard behaviour of the utility is to log the raw message, along with the decoded data and the value of a rolling sample count, the utility was modified to also log the system time in this case. The sample count represents an incrementing counter of each sample provided by the DVB-T adaptor. With the sampling frequency set at 2MHz the counter has a nominal resolution of 500ns, dependent upon the accuracy of the oscillator in the DVB-T adaptor and the avoidance of any lost samples.

A Python script was used to monitor the output of the GPS unit and log the position and system time on each update. Statistics were also captured from the NTP daemon to monitor the drift of the system clock against the PPS output provided by the GPS unit. No live downlink was implemented, instead messages along with timestamps and receiver locations were recorded locally and retrieved later.

The prototype mobile node was taken on a representative-scale but ground-based collection route. Figure 8 shows the prototype installed in the vehicle. The equipment as tested weighed 330g, of which 130g was casing. Even with the additional mass of a battery (approximately 300g for a 10,000mAh example at time of writing) mounting the unit on a UAV platform is completely feasible. UAVs with lift capacities in excess of 1kg are widely available at low cost[5].

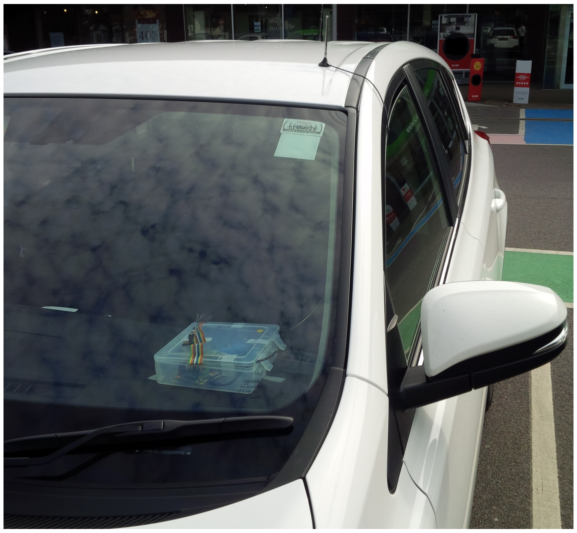

Figure 8: Exterior view of vehicle with prototype installed

The mobile node was mounted on a car and travelled along a journey of approximately 37.5km, as shown in Figure 9. During a 65-minute collection period between 19:00 and 20:05 (BST) on 11th September 2015, 18,464 position claims were received and successfully decoded. A large number of non-ADS-B Mode-S replies were detected but excluded. Similarly, 679 position claims were not decoded correctly and were also omitted. In Figure 10 the maximal detection ranges are overlaid on a map of the southern United Kingdom, demonstrating that the low-cost receiver hardware and off-the-shelf antenna are capable of receiving position messages at considerable distance, even at low elevations. The maximum distance for a received position claim was 207km.

At time of writing, the corresponding data captured at the fixed node during the times in question had not been obtained so the live verification algorithm could not be tested

[4] Originally https://github.com/antirez/dump1090, forked to a Raspberry Pi-compatible Linux version and extended at https://github.com/MalcolmRobb/dump1090

[5] http://www.dji.com/product/spreading-wings-s900/spec

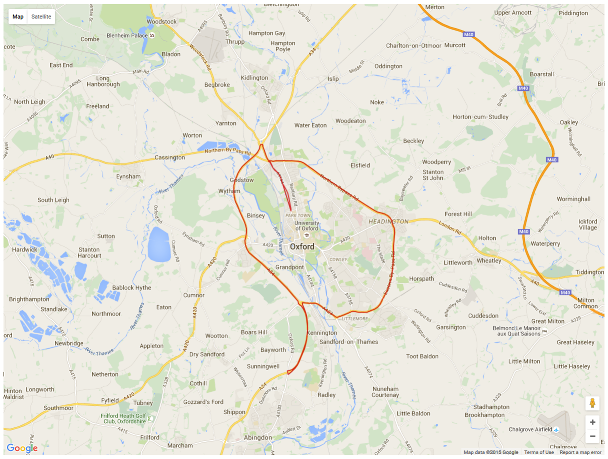

Figure 9: Route taken

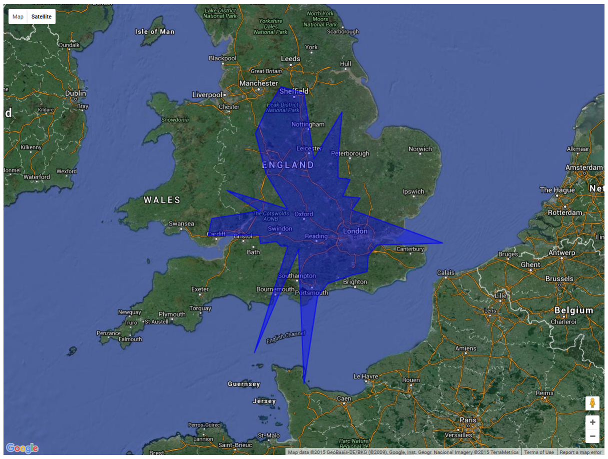

Figure 10: Maximal detection ranges during route

in practice. However, limitations in the timing accuracy suggest that performance would have been poor.

The clock drift was observed at intervals of 16s throughout the capture journey. Clock drift was initially large; in the tens of microseconds as the equipment warmed up. After approximately 1,000 seconds, the conditions had stabilised, enabling the NTP PPS module to compensate for drift more effectively. The clock offset became more predictable; settling to single-digit microsecond values and continued as such for the remainder of the test, as shown in Figure 7. This suggests that clock drift could be compensated for adequately, but also highlights the needs for a mobile verifier to have an 'acclimatisation' period included in its initialisation phase to ensure its clock is stable before it is used to measure message reception times. Furthermore, the effects of pressure on the clock were not demonstrated by a ground-based journey. An airborne mobile node would need to be shielded from the effects of pressure on the clock drift, or a compensation strategy employed.

The measurement error presents the most notable problem however. The maximum stable sampling rate for the RTL2832U chipset has been observed at 2.4 million samples per second (MSPS)[4], so a maximum accuracy of 417ns could be achieved with this hardware, corresponding to approximately 125m of position inaccuracy due to measurement error. This is a lower bound however as the DVB-T dongle passes samples over USB 2.0 with unpredictable delay and local timestamps are only recorded once samples have been transferred and are being processed for decoding by the dump1090 program. Unfortunately the combined effect means that the measurement error in this case is of the order of tens of microseconds and dependent upon the scheduling of the interrupt handler on the Raspberry Pi. As Figure 5 suggests, this level of measurement error would severely impact accuracy. We discuss options for remedying this problem in Section 11 below.

10. DISCUSSION

We have presented a system that can be constructed today using widely-available components. However the capabilities of the system only look set to be enhanced by greater prevalence of cooperative reporting systems, widespread mobile data link provision and progress in mobile platform technology such as UAVs. Advances in flight duration and regulatory developments allowing autonomous operations beyond line-of-sight will enable mobile nodes to provide location verification capabilities over a wider area for a longer time. The equipment required to implement our location verification system could certainly be minimised substantially if intended for widespread deployment, allowing it to be carried by ultra-light UAVs such as those using thermal-hunting, gliding techniques[12].

10.1 Alternative Configurations

We have thus far described a simple configuration of the system, making use of a strictly fixed node and a single mobile node. This approach is suitable for many situations, but is by no means the only configuration. The system could also be implemented with more than two nodes, either to provide tighter verification constraints, or to enable greater coverage. The nodes could be part of a consistent group, or form temporary verification groups from a larger set of mobile nodes, wherever coverage overlaps and the same message is detected by more than one. Similarly the system could be altered to use only mobile nodes and have no fixed node at all. In this case each mobile node would send reports to some remote processing unit to perform the verification, rather than having it co-located at the fixed node. This configuration would greatly enhance the mobility of the verification system, with only the mobile nodes' travel range and the availability of a suitable downlink limiting coverage. Alternatively a hybrid approach is possible, wherein some mobile base station such as a large road vehicle, equipped to act as both fixed node and processing unit, moves to an area and deploys the system ad-hoc to meet a temporary need.

10.2 Potential Applications

The air traffic management scenario modelled in our simulation represents only one use-case for the verification system presented herein. The approach is also applicable to many other scenarios as well, albeit with variations in configuration and different challenges in each case. The system as presented translates easily for other large-scale traffic management instances such as for marine traffic near a busy port. With a suitable VHF radio capable of receiving AIS messages and a long-range airborne mobile receiver (such as an ultralight fixed-wing), the system could be deployed in almost exactly the same configuration. Additionally, the near-constant altitude of marine vessels combined with a comparatively high-altitude mobile receiver substantially reduces the effect of GDOP in the vertical plane. Some other notable scenarios are explored below:

Connected vehicles.

The development of 'connected vehicles' is an active area of research. In this model vehicles report their status, both vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I). Judicious use of this information centrally can allow far more detailed traffic management, while drivers themselves can benefit from information about road conditions provided by vehicles further ahead (one of the claimed advantages of 'platooning' approaches on busy roads, whereby vehicles form an ad-hoc network forwarding relevant information along the column). Such schemes can particularly benefit emergency vehicles; both in obtaining information about routes with lower traffic density and in warning other road users of their approach so traffic can be cleared more quickly. The advent of autonomous vehicles seems likely to only increase the usefulness of all such systems, and the strict adherence to their instructions. In all cases it is vital to detect any false position reports, or indeed any fake claims of the presence of an emergency vehicle in order to clear traffic for a selfish attacker. The fixed node in the verification system proposed here could be co-located with the V2I detection equipment, while the mobile node moves overhead. Alternatively a small number of mobile nodes could be directed reactively to provide location verification in response to unusual or suspicious activity.

Civilian UAV operations.

An attractive use-case is for the policing of civilian UAV operations. Proposals for structured usage of UAVs by commercial operators are often suggesting defined airspace or routes for operations. Inner city 'drone lanes', where vehicles can move across the city in defined airways that keep them away from pedestrians or ground vehicles seem a very distinct possibility[17]. Monitoring traffic is as important in this situation as on road networks; enabling the same reactive management, policing and safety control capabilities. Existing UAV operation proposals note requirements for collision-avoidance systems to be fitted to UAVs, broadcasting the location of the unit publicly so that others can ensure separation. The system then need only detect these position claims. Practical difficulties exist as transmission distances in purely collision-avoidance systems are shorter than those for traffic management. However the low cost of the system would permit widespread deployment of the system in the configuration described here. We discussed above the possibility of replacing the fixed node with another mobile node. If this configuration is viable then it could easily be deployed in a roving model to cover large areas of city, either in a deterministic manner to survey traffic patterns or in an unpredictable way for enforcement purposes.

Sensor network survey.

Security against localisation attacks is a concern in any sensor network where the position of individual sensors is crucial to the validity of the collected information[13]. If individual sensors are programmed to report the location they believe themselves to be at then the verification system herein can be used to detect if any sensor's localisation subsystem is being mislead (e.g., by GPS spoofing). Again, a configuration with only mobile nodes would allow the verification to cover wide areas extremely easily.

11. FURTHER WORK

There are many avenues of future work that would enhance understanding of the verification system.

Firstly, the abstract model could be explored in greater detail. Other use-cases with more stringent constraints on broadcast range and radio propagation could be considered, along with alternative configurations. Localisation error could be included and its effects analysed. Alternative movement patterns could be modelled to include additional practical restrictions or take advantage of other effects. A random movement pattern maximises the difficulty for an attacker in predicting the mobile node's location at a given point in time. However other patterns could be employed in an attempt to improve verification, such as selecting waypoints that move areas of high geometric dilution of precision away from claimed aircraft locations, or that move them by a large amount for each subsequent message. Work to explore the performance characteristics of the system with different routes would help determine whether sacrificing some randomness could be rewarded by a greater increase in detection accuracy. A random movement pattern also ignores limitations that would be placed upon the movements of a real mobile node. For example, a UAV operating in a city or near an active runway would need to contend with prohibited airspace and meet logistical needs such as returning to its base before its fuel is depleted. Studying the effects of respecting these restrictions on the verification performance would be a great step in understanding the practical limitations of a deployment.

Similarly, modelling a more advanced attacker would also enhance the security analysis. A mobile attacker that moves randomly demonstrated very little improvement over a static one. Modelling an attacker that knows the mobile node's movement to some variable accuracy and reacts accordingly, would allow better evaluation of the security level provided by this approach. Better yet, considering a well-equipped attacker with a number of mobile transmitters or with directional transmission equipment would help explore the theoretical boundaries of our approach.

On more practical matters, employing a suitable compensation strategy for measurement error would overcome a critical barrier for the prototype as described here. A radio receiver that delivers messages to the Collector in predictable time would allow the measurement error to move closer to the limit imposed by the receiver hardware itself. One approach with our current prototype is to use the sample counter for samples delivered by the DVB-T dongle. In this manner the variable USB transfer latency, while present, only affects the timeliness of delivery and not the estimation of arrival time as each sample number identifies one capture period on the device, even if that sample does not actually arrive at the Collector until many milliseconds later. This approach is used to implement multilateration techniques in hobbyist air-traffic communities[10]. Only an initial calibration is required to determine the actual timestamp of the first sample, the Clock component is then fulfilled by the Radio receiver's onboard oscillator and a suitable mapping. An different Radio implementation with a higher sample rate would reduce the lower bound on timestamp resolution however; if combined with a more reliable Clock implementation this could substantially decrease the measurement error. Alternatively, considering tracks of claims instead of individual messages would allow the processing unit to filter over the noise parameter with sufficient samples.

12. CONCLUSION

The incorporation of mobility into secure location verification as described herein substantially increases the difficult for an attacker to falsify position claims without being detected; demonstrating >97% accuracy against message injection attacks by the most complex attacker. It is additionally adaptable to many use-cases and is both inexpensive and easily-deployed. As with all TDoA systems, the performance is dependent upon errors being minimised or compensated; the level of timing measurement error is critical to the useful performance of the system and dictates implementation choices. Appropriate selection of a detection threshold depends upon the usage scenario and the level of accuracy provided by the implementation. The system can be practically constructed in a variety of configurations using widely-available equipment and appears to offer an attractive approach to enabling secure location verification in a variety of contexts.

13. ACKNOWLEDGEMENTS

Richard Baker is supported by the UK Engineering and Physical Sciences Research Council (EPSRC). The authors would also like to thank members of the OpenSky Network community for the provision of ADS-B position claim messages and their guidance in working with the data.

14. REFERENCES

[1] A. Bensky. *Wireless Positioning: Technologies and Applications.* Artech House, 2016.

[2] British Broadcasting Corporation. Didcot Power Station collapse: Major search for missing.

[3] S. Capkun, K. B. Rasmussen, M. Cagalj, and M. Srivastava. Secure location verification with hidden and mobile base stations. *Mobile Computing, IEEE Transactions on*, 7(4):470–483, 2008.

[4] Centre Tecnològic de Telecomunicacions de Catalunya. GNSS-SDR operation with a Realtek RTL2832U USB dongle DVB-T receiver.

[5] W. Chen. *Vehicular Communications and Networks: Architectures, Protocols, Operation and Deployment.* Woodhead Publishing, 2015.

[6] Civil Aviation Authority. Regulations for Small Unmanned Aircraft, 11 2015.

[7] A. Costin and A. Francillon. Ghost in the Air (Traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices. *Black Hat USA*, 2012.

[8] Federal Aviation Authority. NextGen Update 2014, 08 2014.

[9] Federal Aviation Authority Unmanned Aircraft Systems Task Force. FAA UAS Task Force Recommendations, 11 2015.

[10] Flightradar24. How We Track Flights with MLAT , 2015.

[11] G. T. Inc. FGPMMOPA6H GPS Standalone Module Data Sheet, 2011.

[12] Jean-Louis Naudin. ThermoPilot - a Thermal Hunter Glider Drone.

[13] J. Luo, H. V. Shukla, and J.-P. Hubaux. Non-Interactive Location Surveying for Sensor Networks with Mobility-Differentiated ToA. In *INFOCOM*, 2006.

[14] F. Papi, D. Tarchi, M. Vespe, F. Oliveri, F. Borghese, G. Aulicino, and A. Vollero. Radiolocation and tracking of automatic identification system signals for maritime situational awareness. *IET Radar, Sonar & Navigation*, 9(5):568–580, 2014.

[15] P. Perazzo, K. Ariyapala, M. Conti, and G. Dini. The verifier bee: A path planner for drone-based secure location verification. In *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2015 IEEE 16th International Symposium on a*, pages 1–9. IEEE, 2015.

[16] P. Pounds, R. Mahony, P. Hynes, and J. M. Roberts. Design of a four-rotor aerial robot. In *Proceedings of the 2002 Australasian Conference on Robotics and Automation (ACRA 2002)*, pages 145–150. Australian Robotics & Automation Association, 2002.

[17] Robert Pearce. Presentation to the UAS COE Public Meeting, 2014.

[18] M. Schäfer, V. Lenders, and I. Martinovic. Experimental analysis of attacks on next generation air traffic communication. In *Applied Cryptography and Network Security*, pages 253–271. Springer, 2013.

[19] M. Schäfer, V. Lenders, and J. Schmitt. Secure Track Verification. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2015.

[20] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm. Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research. In *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*, IPSN '14, pages 83–94, Piscataway, NJ, USA, 2014. IEEE Press.

[21] M. Strohmeier and I. Martinovic. On Passive Data Link Layer Fingerprinting of Aircraft Transponders. In *Proceedings of the First ACM Workshop on Cyber-Physical Systems-Security and/or PrivaCy*, pages 1–9. ACM, 2015.

[22] N. O. Tippenhauer, C. Pöpper, K. B. Rasmussen, and S. Capkun. On the requirements for successful GPS spoofing attacks. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 75–86. ACM, 2011.

[23] Tom Mendelsohn. Amazon to test drone deliveries in UK after government clears runway, 2016.

[24] C. D. Wickens, A. S. Mavor, R. Parasuraman, J. P. McGee, et al. *The future of air traffic control: Human operators and automation.* National Academies Press, 1998.