Securing a Femtocell-based Location Service

Jack Brassil, Pratyusa K. Manadhata
HP Laboratories

***Abstract*—Mobile device users are increasingly incented to falsify their locations to retain location privacy while capturing economic benefits such as location-based retail discounts. Location spoofing is easily achieved with several widely-used location services that rely on smartphone applications to convey GPS coordinates, IP addresses, or WiFi Positioning System radio environment data. In earlier work we introduced a network infrastructure-based system that provides spontaneous, rapid, and robust mobile device location authentication by supplementing existing 802.11x APs with off-the-shelf femtocells. The proposed system has the property of leveraging mobile operator infrastructure, without requiring operator participation in either providing or authenticating location. In this paper we present a security analysis of the location authentication system. We assess its resistance to DoS attacks, identify various approaches for a mobile user to deceive a location verifier with and without the assistance of a colluder, and explore the tradeoffs between cost and complexity in mounting such attacks. Finally, we identify a collection of system modifications and countermeasures to anticipated attacks designed to decrease location authentication system vulnerabilities and increase privacy protection.**

Keywords: location privacy, GPS, WPS, distance bounding, proximity testing, E-911, indoor positioning, covert channels

I. Introduction

Internet location-based application providers such as *EverSave* and *foursquare* potentially stand to benefit by verifying the location of their mobile users. Such services ideally seek to spontaneously authenticate mobile device location without the need to have a pre-existing relationship with each user. But few options are available for location services to rapidly authenticate a new client. While mobile operators provide ubiquitously available network-based location services, this service is generally accessible only to subscribers, not third parties. Mobile operators currently have no straightforward means of authorizing and sharing subscriber location information with third-party location applications, while also ensuring that subscriber security and privacy are protected, even if all parties agree to such sharing.

As a result smartphone applications relaying GPS coordinates or WiFi Positioning System (WPS) radio environments observations have become the preferred choice of location service for internet application providers. Existing services generally rely on a user's assertion of location (e.g., via an application uploading GPS coordinates). But the economic incentives for users to provide false location information are growing. As incentives such as these location-based retail discount coupon distribution have grown, we unsurprisingly find a surging number of location spoofing apps available on the smartphone application marketplace. Hence we anticipate that systems that *authenticate* client location will become increasingly important as emerging location-driven ecosystems evolve.

To address this demand we have proposed authenticating a mobile device's location by placing femtocells at existing public WiFi sites [1]. The short wireless range of these basestations permits us to locate User Equipment (UE) to within tens of meters, and indoor operation is supported. By sending either distinctive *voice* or *data* traffic while remotely monitoring femtocell ingress link activity, a remote calling party can verify a called party's location by analyzing a reverse communication channel characterizing femtocell activity.

In this paper we focus on the security and privacy characteristics of the femtocell-based Location Authentication (*LocAuth*) system. We argue that several properties of a femtocell-based approach – short-range wireless, managed infrastructure, encrypted uphaul, etc. – make the solution surprisingly difficult to defeat relative to comparable GPS and WPS approaches that simply forward unverifiable observations and are accepted uncritically. The remainder of the paper is organized as follows. Section II begins with a quick review of femtocell properties that are exploited by our system. Next we introduce the participants in a location verification, and sketch our proposed authentication system architecture and operation. Readers can find more detailed information about system performance and a prototype implementation in [1], [2]. We then present a security analysis in Section III, where we categorize and describe plausible attacks and attempts to disrupt system operation, deceive the verifier, and obtain private information about the participants. Resistance to collusive attacks – often a key weakness in previous proposals – is considered, as are countermeasures to limit security vulnerabilities. In the final sections we review the vulnerabilities of some widely deployed location systems, summarize our major contributions, and identify several envisioned enhancements of our authentication approach.

II. Location System

A. Key Femtocell Properties

We now briefly highlight a few key properties of existing 3G femtocell technology that are critical to understanding the operation of our location service and its potential security weaknesses; additional femtocell details are well described in survey articles including [3]. Femtocells are low-power, limited range (e.g., tens of meters) wireless access points that operate in licensed spectrum to connect subscriber's mobile devices to their mobile operator's network. The principal application of residential femtocells today is to provide wireless coverage in areas not well served by cell towers. Femtocells typically use wired public internet access as backhaul. They

978-1-4673-0937-0/12/$31.00 ©2012 IEEE

satisfy the various regulatory, compliance and spectrum use requirements of macrocells, including supporting location service.

Residential femtocells typically support only 2-8 active mobile device associations (i.e., users). A voice call consumes roughly a continuous 50 kbs duplex rate, depending on the coding mechanism employed; data calls approach sustained download speeds of 2 Mbs. Voice calls can originate on residential femtocells, and subsequently be handed over to cell towers as callers move, however active calls originating elsewhere may not be handed to a femtocell. Equipment owners may specify access control lists (e.g., family members only, any subscriber). GPS signal availability is typically required, and can be achieved indoors through cabled remote antennas.

Voice and data traffic to and from the femtocell are directed to a Security Gateway (SG) at the edge of the operator's core network. Some control traffic may also be directed to other service points, such as a GPS Gateway. Voice, data and control traffic between the SG and femtocell is tunneled and encrypted with the Encapsulated Security Payload (ESP) protocol (tunnel mode) [4], and transported over UDP. Hence, confidentiality is assured against snooping. ESP provides integrity checking for the packet payload and protects against replay attacks. ESP uses DES or 3DES to provide data confidentiality by encrypting the packet's contents. Note that an observer of the SG-femtocell channel will see packets that appear to originate at those endpoints, and not at UE or their communicating party.

B. Participants in Location Authentication

Bob is a mobile device user whose location is to be authenticated. He is willing to cooperate with the authentication to realize some benefit but we can not trust his assertion of his location. Though not mandatory, for our purposes we will assume that Bob carries a smartphone. *Alice* seeks to verify Bob's present location (with his explicit approval). Alice and Bob do not need to have any pre-existing relationship. Alice must have the equivalent capability of a smart phone, or more precisely a (mobile or landline) voice-only phone plus minimal compute and display capability; a web browser suffices. The *Location Service Provider* (LSP) provides a public-access location authentication service. The location itself – say a coffee shop – might already offer a public WiFi service. The site location is assumed to be fixed over time. The LSP – the coffee shop owner – has no prior relationship with either Alice or Bob, each of who can remain permanently anonymous to the LSP.

C. System Architecture and Operation

Figure 1 depicts a basic single-carrier authentication system architecture. To an existing 802.11x access point with an internet connection, an LSP minimally adds 1) a femtocell, and 2) a computer operating as a *location server*. The location server hosts a web server, and offers a public page with detailed site location information (e.g., GPS, postal address, contact information, etc.) The location server also continuously monitors the average bandwidth on the (encrypted) downlink between the AP and femtocell; an average bandwidth for each 1 second interval is measured, and these values form a data stream that is publicly exported. Other communication channel metadata might also be measured and exported, such as the numbers of packets of each length l observed in the interval. We refer to this reverse communication channel as the *Return Data Feed* (RDF).

The figure also depicts Bob's mobile Service Provider's core network. Alice need not share a common operator network with Bob, nor even know Bob's operator. Assume Bob is in range of the femtocell; any voice or data communication from Alice to Bob will ultimately traverse Bob's operator's network and be forwarded to the femtocell on route to Bob. Note, of course, that other subscribers of Bob's mobile operator might be present at the location, be associated with the femtocell, and also might be receiving voice and data traffic through the femtocell.

Alice can communicate to Bob by initiating a voice call, or performing a data transfer. When Alice communicates with Bob and simultaneously observes the RDF, she expects the bandwidth measured on the femtocell ingress to increase and expects the bandwidth to fall when she terminates communication. Suppose Alice initiates a voice call. Alice's call to Bob impresses a distinct traffic envelope on the AP-femtocell downlink. Within a few seconds of Bob's off-hook, Alice expects to observe the measured average bandwidth values increase by the bandwidth consumed by her call; in our prototype system this is roughly 50 kbs. She expects a similar decrease within a few seconds of hanging up.

Alternately, Alice can transfer data to Bob. Alice must be capable of controlling a data transfer to impress a unique traffic signal on the femtocell ingress, such as by performing rate-control or manipulating packet lengths. The data can be pushed or pulled, and the underlying transfer protocol is unrestricted. One simple approach is for Alice to provide Bob the URL of a data file on a web server she controls, and allow Bob to initiate the data transfer. Note that *http* transfers potentially avoid the need for Bob to have a special-purpose application to receive the transfer.

Consider the following basic authentication process used by Alice to confirm Bob's association with the femtocell, and hence his location:

1) Bob successfully binds to the femtocell.
2) Bob messages Alice, and provides her with the LSP's location URL.
3) The location server continuously monitors the (encrypted) AP-femto downstream link and exports an RDF with two (logical) component streams: 1) the average bandwidth over each one second interval, and 2) the number of packets received in the previous second of each observed packet length.
4) Alice transfers data to Bob and controls either the transfer rate or packet lengths (or both) to impress a data traffic signature on the AP-femto link.
5) Alice monitors the RDF for characteristics of her data transfer.

If the behavior of the RDF convinces Alice that she is observing her own voice traffic traverse the AP-femtocell link,

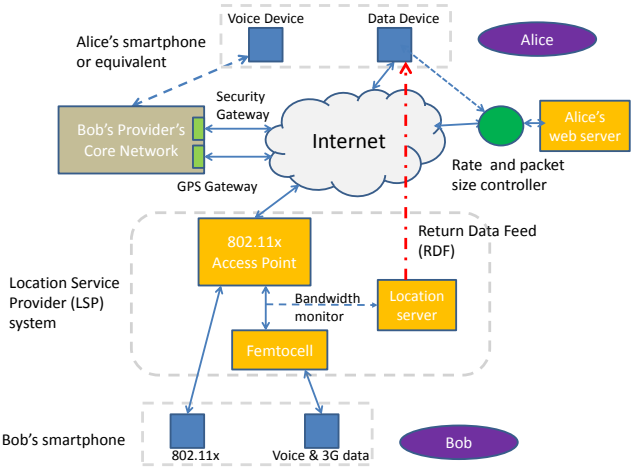

Fig. 1. Architecture of a single-carrier location authentication system.

Alice confirms Bob's phone's association with the femtocell, and concludes that Bob is present at the specified location. If the observed RDF does not reflect Alice's communications, she can not conclude that Bob is on-site. Alice can elect to retry her call at a later time to confirm Bob's presence.

We have constructed a prototype of this system and have successfully demonstrated its operation. Observe that if Bob is carrying a voice-only device, an analogous verification scheme can be performed by Alice by simply making a voice call to Bob. More detail on the design of signals that Alice send to Bob, the robustness of the system, and the speed at which Alice can perform an authentication is available in [2]. In the remainder of this paper we will focus on the security and privacy vulnerabilities of the proposed system, and discuss how these weaknesses can be overcome.

III. Security Analysis

We next examine the security and privacy properties of the proposed location system, and discuss its resistance to some frequently suggested attacks. While many attacks are easily conceived, they can be deceptively complicated or costly to successfully implement. We do not strive to exhaustively consider all possible variants, rather we highlight those that we consider likely to be most effective or difficult to prevent. Finally, we describe several simple enhancements to improve overall system robustness. We also introduce countermeasures for specific attacks, some of which are simple to implement yet can render attacks ineffective, more easily detected, or more expensive to mount.

The location system's attack surface is defined by the three principals (or actors) – Bob, Alice, and the LSP – and five critical system components, namely the AP, femtocell, location server, web server, and UE (i.e., primarily smartphones). Of the principals, either Bob or the LSP might interfere with a verification. Recall that we consider Bob to be cooperative but untrustworthy. More precisely, Bob may simply seek to appear to be cooperative, even if misrepresenting his location and working to interfere with a verification. In general, the LSP stands to gain by hosting a verification venue (e.g., perhaps by receiving direct or indirect compensation from Alice), and any of his current and future gain is put at risk if he is caught interfering with authentications. External parties might also attack the system; we refer to these as malicious non-principals (MNPs).

A. Disrupting Service

An attacker can prevent a LocAuth from succeeding either by compromising one or more system components, or by disrupting operation of the network(s) interconnecting those components. Like any internet-attached service, a LocAuth system is subject to network-based DoS attacks. Under an attack, Alice might be unable to verify Bob's location, or Bob might be unable to establish his location. Denials-of-Service (DoS) can be executed by MNPs or the system principals themselves. As an example, an LSP can simply interrupt its service at any time (e.g., denying any parties within range from being authenticated by remote parties).

A DoS attack on an authentication site can focus on either the femtocell or location server. A location server under attack might be unable to respond to a web request for the location URL, or continuously transmit the exported feed. An attack on the femtocell or AP can saturate the downstream bandwidth to the device such that either incoming or outgoing calls can not be forwarded by the femtocell. Interestingly, attacks on the femtocell or AP can be initiated locally by an onsite MNP, either by consuming all available bandwidth (preventing Bob from communicating) or by broadcasting a radio jamming signal. A network DoS attack launched against Alice's network can also impede her ability to authenticate any users. Of course, Alice can mitigate such attacks by initiating authentications from multiple locations, or having proxies perform authentications on her behalf.

Observe that the authentication system comprises a network of decentralized, independent, geographically distributed verification sites with no centralized component. This offers considerable protection against DoS attacks; while individual LocAuth venues can be attacked, the effort (e.g., bandwidth) required in a multi-location attack grows with the number of attacked authentication locations.

An attacker can alternately subvert a verification by *compromising* any system component (e.g., establishing supervisory control of a component). AP and web server compromises are both achievable and well-studied, but are outside the scope of this paper. Lack of physical security has also been raised as a vulnerability potentially facilitating femtocell compromises [5], [6]. Finally, smartphone users' willingness to download non-certified applications with little reservation remains a compromise threat whose extent has yet to be fully understood [7].

B. Deceiving the Verifier

We next consider how Bob can act to deceive Alice by attempting to convince her that he is at the claimed authentication site when he is elsewhere. Deception attacks invariably take one (or both) of the following forms; either Bob attempts to deceive Alice that she is communicating with the claimed location rather than his actual present location, or Bob deceives

Alice by manipulating the RDF exported from the claimed location to indicate his presence.

Deception attacks can be mounted individually by Bob, or with the assistance of a colluder. Let's first consider the former. Suppose Alice transfers network traffic (either voice or data) through the femtocell by initiating a communication to Bob. For Bob to deceive Alice, she must observe behavior on the RDF that closely resembles what she expects – an increase in traffic of approximately 50 kbs shortly after a voice call initiates, and a similar decrease when she terminates the call, or a sequence of packet sizes consistent with her data transmission. To accomplish this, Bob must either 1) ensure that a call (or data transfer) with timing, bandwidth usage, and packet sizes consistent with Alice's expectation is received by the femtocell, or 2) modify or substitute the RDF with a counterfeit feed consistent with her expectations.

To achieve the former, Bob can remotely 'forward' a (logical) copy of Alice's transmitted packet stream to the femtocell ingress at the claimed location. Note that any traffic forwarded to the femtocell does not necessarily need a recipient (or receiving application); even if dropped by the femtocell the bandwidth appearing on the femtocell ingress is sufficient for deception. Such a forwarding action must be performed quickly or Alice might detect the delay in the appearance of her traffic to the femtocell. Forwarding traffic to the femtocell through the femtocell's associated mobile operator would take several seconds, and likely be detected. Hence, Bob's preferred approach would be to direct a data stream mimicking Alice's transfer directly to the femtocell's IP address.

To achieve the latter, Bob can alternately send a *modified* exported bandwidth stream to Alice. For example, he can insert himself 'in-the-middle' between the (claimed) location server and Alice, and forward a modified version of the location server's RDF, enhanced to falsely include the channel characteristics associated with Alice's transfer.

As an alternative, Bob can send a substitute stream to Alice by providing her a false location URL, pointing to a web site he operates from which he can also export an RDF he controls. A particularly elaborate version of this attack is as follows. Suppose Bob operates a Private Location Authentication system (PLA) at his current location, effectively impersonating the claimed location's LocAuth system. Bob provides Alice with a location URL that mimics the claimed location's, with his own RDF. In the absence of a central database of valid authentication sites, Alice places her trust in a network of unverifiable LocAuth system operators. Without taking additional steps to verify the legitimacy of the PLA site, it is possible for Alice to be deceived.

But note how difficult it would be for Bob to sustain this deception over time if he attempts to use the same PLA system to support multiple deceptions. LocAuth sites are more-or-less permanent and fixed; that is, a location page and exported bandwidth feed are expected to be unchanged over long time periods. Hence, Alice expects to be able to reach these resources at *any* time in the future. As are result, Bob is obligated to keep these services running indefinitely. If Alice revisits the location URL, and finds it unavailable or changed, she can invalidate any previous confirmation of Bob's location.

Further, Alice can maintain a list of all URLs and RDFs previously provided by Bob. Suppose Bob attempts to use a single PLA to deceive Alice about his location in a sequence of deceits over time. When performing an authentication, Alice expects that the only RDF indicating her call is that of Bob's present location; Alice can monitor all past feeds to check for any activity her current call generates at Bob's purported previous (other) locations.

An example demonstrates the effort required by Bob. Suppose Bob operates a single PLA. In deceit A, he claims to be at Location A, and creates a location web page that mimics that of the claimed location. For a later deceit B, he creates a second location web page mimicking location B. For Alice not to detect the deception, Bob must ensure that the IP addresses of the (supposedly different) web servers differs. More significantly, both location URLs would export the same bandwidth feed from the single femtocell in his PLA. If Bob claims to be a location B, Alice could monitor the previously provided feed for location A and detect Bob's deception. Hence, Bob must not only execute his current deception, but ensure that all previous deceptions remain active and do not raise suspicion.

C. Collusive Deceptions

Let's next turn to collusive attacks where Bob has assistance from a confederate (i.e., the colluder). Such an assistant is usually equipped with Bob's phone and positioned at Bob's claimed location. Of course, if Bob is unknown to Alice and passes his smartphone to an on-site colluder, we will be unable to distinguish him from another; Bob's private key on his smartphone is his identity.

Recall that if Bob possesses a voice-only phone, the best Alice can do is locate Bob's phone, and establish that Bob is in possession of his phone by speaking with him (if he is known to her). Hence a commonly proposed attack is for an on-site colluder to 'forward' Alice's voice conversation to Bob.

A common misconception is that this attack can be executed with mobile operator-based 'call forwarding' service. In many cases, however, this service is network-based redirection, and the incoming call would not reach the femtocell targeted by Alice, and she would not observe expected activity on the femtocell ingress. It is possible, however, for the colluder to implement UE-based forwarding. Forwarding via a mobile network is likely to result in a call-initiation delay detectable by Alice. Yet a call-setup delay can be eliminated if the colluder keeps a pre-established connection to Bob in anticipation of Alice's call. A preferred attack is for the colluder to convert the received voice signal to VoIP, and send to Bob over an internet connection. Such an attack would require modest technical sophistication to prevent Alice from hearing audible indications of forwarding (e.g., echoes and dial tone).

Collusive attacks require the existence of a relatively low-latency communication channel between Bob and the colluder to support a coordinated, timely deception. The low delay requirement generally rules out 3G/4G communication channels, where end-to-end delays can be significant. Bob can communicate with the colluder using IP over the claimed location's

802.11x AP, or a separate IP communication channel 'carried in' by the colluder that does not rely on LSP infrastructure.

A challenging collusive attack to detect has the LSP operating as Bob's colluder. In this attack, Bob signals the LSP to modify the bandwidth of the exported data stream by simply indicating the call initiation and termination times. The incentive for an LSP to collude with Bob would necessarily have to outweigh the risk that the deceit is detected, and the LSP's service is flagged as untrustworthy; a loss of all future revenue for the LSP could be the result.

Next we consider a collection of minor system modifications and *countermeasures* that make deceiving the verifier more difficult. A first tool to detect deception lies in the amount of information that is returned to Alice. In general, more information can assist Alice in both verifying location and identifying suspicious behavior, at the expense of consuming additional traffic on the exported feed. In some cases, even a small amount of additional information can be valuable. For example, if the femtocell is receiving and dropping an excessive amount of incoming traffic, the femtocell could indicate a 'health' status indicating that the system might be under attack.

As a second example, the system can supplement the RDF with measurements of the femtocell *egress* link characteristics (e.g., average bandwidth measurements). Alice can observe the egress data to detect attempts to manipulate traffic on the femtocell ingress. As an example, if Alice is speaking with Bob she would expect the femtocell egress to behave in a fashion similar to the ingress. If Bob is simply redirecting a data stream to the femtocell from a remote location, the RDF would not exhibit the expected behavior. Rather, depending on system implementation Alice might see a small increase in traffic associated ICMP redirects in response to data sent to an inactive TCP or UDP port.

Another means of protecting the system from external manipulation is to have the bandwidth monitor on the femtocell ingress report only the amount of bandwidth traffic whose source IP address corresponds to the femtocell's mobile operator security gateway (or domain). Any attack traffic originating from other IP sources would not be observed by Alice, forcing Bob to spoof IP source addresses.

Where possible, Alice should control the timing of her communication with Bob; she should 'push' data transfers to Bob, rather than let Bob 'pull' data from her web server. Otherwise, Bob can quickly send a URL provided by Alice to a colluder at the claimed location, who can then pull the data himself. This attack is particularly threatening since with no additional challenges a colluder would not need to carry Bob's phone to mount the attack.

Of course Alice can also perform a variety of actions to confirm a verification if she suspects deceit. Alice can execute multiple authentication transfers to Bob. Another approach is for Alice to add a challenge such as a request that Bob call her. If an exported feed also reports bandwidth on the femtocell egress, than Bob would need to have a colluder support a forwarding of Bob's call to Alice (opposite in direction from the earlier attack where Alice calls Bob and the colluder forwards).

D. Privacy

We next examine the information exchange – and hence the potential information loss – between the system participants and/or external parties. To begin note that Bob realizes his principal privacy requirement, the ability to *opt-in* to an authentication on a per-transaction basis. His opt-in action takes the form of informing Alice of his location and the site URL.

In most applications Bob reveals his identity to Alice. If Bob has a voice-only phone, his speaking voice can serve as a personal identifier if he is known to Alice. In the case where Bob has a smartphone and the authentication is entirely automated, Alice can confirm Bob's identity by asking for part of his exchange to include a digital signature. Alice need not be known to Bob, though she too can sign an exchange to reassure Bob that he is revealing his location information to the intended party. This action can help Bob detect a malicious party seeking to track his location.

The LocAuth system has the unusual property that the location service provider – namely the site operator – need not have any knowledge of Alice nor Bob, nor the fact that they engage in a transaction. Both Alice and Bob are protected from revealing their identity (or relationship) to the LSP. Further, since transactions are encrypted and usually not known to the LSP, no records are maintained that might later be revealed in a compromise of the system. In particular, an LSP eavesdropping on the femtocell ingress does not see Alice's IP address; all traffic appears to be to/from the operator's security gateway. Note, however, that Alice should anonymize her network address to minimize her risk that her identity can be determined by the LSP when she accesses the site URL. The LSP is also in a position to monitor unencrypted traffic through the public AP (i.e., that traffic not associated with the femtocell), making femtocell traffic analysis a relatively unattractive target to an eavesdropping LSP.

Consider the amount of information revealed to an external party eavesdropping on an RDF. Though the data stream appears to contain little valuable information, it forms a *covert* channel that can provide a remote party with an indication of the site occupancy. Such information is of potential value to a burglar waiting for an empty store to rob. In another example, a business analyst could examine the overall network utilization of all RDFs of every location of a certain business (e.g., a coffee shop chain) as an indication of store visit trends, and perhaps infer business activity. Of course a location can be densely occupied, but if none of the occupants are using the femtocell than this information is not revealed to an observer of the feed. Similarly, even a single occupant using the femtocell can download enough data to nearly fully utilize the femtocell downlink. Note, however, that traffic analysis by the eavesdropper might be able to distinguish between a single user, and multiple users, of a femtocell.

Finally, we note that Alice's data transfers to Bob exiting her web server may not be encrypted until reaching the SG and are subject to eavesdropping. But any traffic sent from Alice to Bob appears to be destined to a web proxy in Bob's operator's network domain. Hence an observer eavesdropping on Alice's

server will not be aware when and if she is communicating with Bob.

IV. Vulnerabilities of Alternative Location Authentication Services

Having considered the security provided by our proposed system, we next briefly review the security vulnerabilities of some widely deployed location systems. We note that a large number of compelling LocAuth architectures have been introduced [8], [9], [10], [11], but have as yet attained relatively limited deployment.

- *Mobile Operator Location Service:* The network-based location service used by mobile operators offers relatively strong protection against many classes of attacks by MNPs, due in part to system properties including physically secure cellular infrastructure and hardware-based UE identifiers (e.g., with SIM cards). This level of protection comes at the expense of limiting location service offerings to subscribers, and avoiding release of location information to third parties (with the exception of E-911 service). To address these limitations, 'location-as-a-service' providers [12], [13] are developing mechanisms to securely share operator location information with third parties, though the vulnerabilities of these newly emerging services has yet to be established.
- *WPS:* While WiFi Positioning Systems (WPS) such as offered by Skyhook Wireless [14] offer location service, these services do not provide authentication. Location information is easily spoofed by tools designed to effectively mimic the observable APs known to exist at a claimed, remote location, and also by spoofing one's current IP source address [15].
- *Handset-based GPS:* Handset-based GPS receivers remain a preferred location technology for smartphone applications operating in environments where GPS signals are available. Like WPS, the GPS coordinates imported to these applications are easily falsified.

V. Conclusion

We have examined the security and privacy properties of a new approach to location authentication that operates by supplementing existing WiFi hotspots with an unmodified femtocell AP. The system exploits mobile operator technology without directly involving the operator in a verification. By supporting spontaneous, transaction-oriented verifications, the approach is well aligned with the evolving needs of internet location-based application providers, and particularly their desire to authenticate new users, such as a customer arriving to a retail shopping mall. In contrast to earlier work, we do not strive to achieve a high degree of system security by constructing trusted location infrastructure. Instead, we explore how cooperating third parties – a verifier and the operator of a site offering authentication service – can jointly provide a service without the direct participation of mobile operators.

Femtocells are, of course, not widely deployed today, and widespread deployment would be required to build a network of authentication sites. But, apart from enabling new services, the basic advantages of wider deployment of femtocell technology – both to operators and consumers – remain plentiful. Our proposed system requires no changes to operator infrastructure or mobile user equipment. Hence, the technology required to deploy a large-scale location authentication system exists, is inexpensive, operates off-the-shelf, and can be deployed easily. While future large-scale deployment of femtocells is uncertain, we do envision the integration of femtocell and 802.11x radios in a single multi-access unit as being a potential catalyst for wider-scale deployment. We believe that support for femtocell-based applications service can be a key driver of future growth in the small cell market.

Mobile device users are naturally concerned about protecting their location privacy. Our system let's users opt-in to each transaction, mitigating some fears about location tracking. We have demonstrated that the system has the unusual property of the location service not participating in transactions directly, nor even being necessarily aware of their existence. Hence mobile users have one less party they must trust. We have presented a high-level system security analysis, and have argued that deceiving a location verifier is a non-trivial technical challenge. We have shown that a decentralized authentication system offers protection from DoS attacks. Finally, we have introduced a number of enhancements to the originally proposed system that reduce its attack surface, and have explored various countermeasures that increase resistance to commonly suggested system attacks.

References

[1] R. Netravali, J. Brassil, "Femtocell-assisted Location Authentication (poster/extended abstract)," *IEEE LANMAN 2011*, Oct. 2011.

[2] J. Brassil, P.K. Manadhata, "Robust Location Authentication with Femtocells," *submitted for publication*, 2011.

[3] V. Chandrasekhar, J. Andrews, A. Gatherer, "Femtocell Networks: A Survey", *IEEE Communications Magazine*, Vol. 46, No. 9, Sept. 2008.

[4] S. Kent, "IP Encapsulating Security Payload (ESP)," *IETF RFC 4303*, December 2005.

[5] R. Borgaonkar, K. Redon, J.-P. Seifert, "Experimental Analysis of the Femtocell Location Verification Techniques," *Proc. of 15th NordSec*, 2010.

[6] M. Gorlatova, R. Aiello, S. Mangold, "Managing location privacy in cellular networks with femtocell deployments," *Proc. of WiOpt'11*, 2011.

[7] W. Enck, M. Ongtang, P. McDaniel, "On Lightweight Mobile Phone Application Certification", *Proc. of 16th ACM CCS*, 2009.

[8] S. Saroiu, A. Wolman, "Enabling New Mobile Applications with Location Proofs," *Proc. of HotMobile 2009*, pp. 1-6.

[9] V. Lenders, E. Koukoumidis, P. Zhang, M. Martonosi, "Location-based Trust for Mobile User-Generated Contents: Applications, Challenges and Implementations," *Proc. of Hotmobile 2008*, 2008.

[10] A. Narayanan, N. Thiagarajan, M. Lakhani, M. Hamburg, D. Boneh, "Location Privacy via Private Proximity Testing," *Proc. of NDSS 2011*, 2011.

[11] M. Talasila, R. Curtmola, C. Borcea, "Location Verification through Immediate Neighbors Knowledge," *Proc. of Mobiquitous'10*, 2010.

[12] Veriplace, Inc., http://veriplace.com.

[13] LOC-AID, Inc., http://www.loc-aid.com.

[14] Skyhook Wireless, http://www.skyhookwireless.com/

[15] N. Tippenhauer, K. Rasmussen, C. Ppper, S. Capkun, "Attacks on Public WLAN-based Positioning," *Proc. of MobiSys'09*, 2009.