Authenticating a Mobile Device's Location Using Voice Signatures

Jack Brassil, Ravi Netravali†, Stuart Haber, Pratyusa Manadhata, and Prasad Rao
HP Laboratories, Columbia University†

***Abstract*—Providers of location-based services seek new methods to authenticate the location of their clients. We propose a novel infrastructure-based solution that provides spontaneous and transaction-oriented mobile device location authentication via an integrated 802.11x wireless access point and 3G femtocell access system. By simply making a voice call while remotely monitoring femtocell activity, a calling party can verify a (cooperating) called party's location even when the participants have no pre-existing relationship. We show how such a traffic signature can be reliably detected even in the presence of heavy cross-traffic introduced by other femtocell users. We describe how the verification proceeds without revealing details of the authentication – or even the parties involved – to the location provider.**

Keywords: location tracking, GPS, security and privacy, macrocells, distance bounding, detection theory, proximity testing, E-911

I. Introduction

Users of mobile devices, e.g., phones and netbooks, increasingly perform day to day tasks such as banking and shopping on their devices by utilizing third party services. The service providers, especially location based application providers (LAP), can offer better services to the users based on the users' locations. For example, a bank may be able to authenticate a customer's Automated Teller Machine (ATM) transactions from the customer's location. If the bank can authenticate the customer's location and conclude that the customer is near the ATM, then the bank may infer that the transactions are legitimate. If, however, the customer is not near the ATM, then the transaction is suspicious, e.g., a thief might be using a stolen ATM card to withdraw cash from the customer's account.

LAPs, ranging from discount distributors such as *LivingSocial* and *GroupOn* to geo-social services including *Foursquare*, seek to stimulate and direct consumer purchase decisions based in part on the distribution of targeted retail-oriented 'flash' deals and discounts. Many LAPs not only seek to locate clients, but also authenticate those client locations. In many cases, those clients are new users of the LAP's service with whom they have no pre-existing relationship, such as a consumer entering a shopping mall.

Despite nearly 2 decades of research, authenticating a mobile device's location remains difficult [1], [2], [3], [4]. Though mobile operators have ubiquitously deployed location services for subscribers (e.g., Verizon's VZ Navigator), these services are often unavailable or poorly suited for third parties including LAPs, with perhaps the notable exception being mandated E-911 service. Hence inexpensive and widely deployed GPS receivers have made handset-based location service the preferred choice of LAPs. Existing services generally rely on a user's assertion of location (e.g., via an application uploading GPS coordinates). As users benefit from location authentication, e.g., location based discount coupons in Foursquare, the economic incentives to provide false location information are growing. We unsurprisingly find many location spoofing applications on the Android market. Hence we anticipate that authenticating client location will become increasingly important as emerging location-driven ecosystems evolve, and that some LAPs will demand to authenticate clients to both enhance and measure service delivery quality. For example, a LAP could report to its advertisers that 95% of coupons were distributed to mobile consumers whose locations were authenticated.

In this paper, we propose a novel approach to device location authentication using off-the-shelf femtocells [5]. The short wireless range of these basestations permits us to locate associated mobile devices to within tens of meters. By impressing a unique *voice* traffic signature while remotely monitoring femtocell ingress link activity, a remote calling party can verify *any* called party's location. A particular advantage of this scheme is that location could be authenticated for any mobile device, including voice-only mobile phones. Our lightweight and non-cryptographic approach requires no modification to existing infrastructure and avoids the complexities of managing trusted infrastructure such as public key infrastructure (PKI) and trusted platform modules (TPM). Hence our approach is vulnerable to challenging and difficult to mount attacks, e.g., most location authentication systems – including ours – are vulnerable to *wormhole attacks*, where a remote party colludes with an on-site associate to fake one's presence. Our authentication strength, however, is consistent with the commercial needs of existing LAPs.

The rest of the paper is organized as follows. Section II describes our design goals. The next section provides a brief review of femtocell technology, then outlines our proposed authentication system architecture and operation. Section IV examines the problem of designing and detecting traffic signatures in the presence of interfering cross-traffic including voice calls, text messages and data transfers by other parties sharing the femtocell. We then describe a prototype system we constructed. We introduce a general analytical model of noise in

978-1-4673-1430-5/12/$31.00

Section VI and show that we can evaluate the probability that our voice signal can be correctly detected in the presence of multiple types of simultaneously occuring cross traffic types. We present a basic security analysis in Section VII, where we describe plausible attacks and attempts to defeat location authentication. In the final sections we review related work, summarize our contributions, and identify several envisioned enhancements of our authentication approach.

II. Design Goals

LAPs seek a relationship with their clients that is consistent with the dynamic, transient nature of their services, such as the immediate delivery of a discount coupon for clients entering a shopping mall. In constructing a suitable location service, we have identified the following desired system properties:

- *Transaction-oriented* The service should provide a one-time location check, not a continuous location 'tracking' service. Users should opt-in with each location check.
- *Spontaneous, fast, and easy* Authenticating a location should not require scheduling or planning, and must execute quickly with little or no mobile user burden.
- *Mobile device-independent* Any mobile 3G or 4G device carried by users should permit location checks, including phones with no data services.
- *Correct and trusted location service* Both LAPs and their clients must trust that the location service provider reliably delivers correct location information.
- *Private* Only a single, specific LAP should be authorized to perform each location check. Even the location service provider should be unaware of the transaction, with no records kept.
- *Location accuracy both indoors and outdoors* A location check should provide fine-grain location information – such as GPS does – even if the client is indoors.
- *Authentication Strength* The authentication service's strength should be consistent with the LAPs' needs. For example, LAPs must be confident that multiple parties (e.g., the location service provider and the located party) can not easily collude to verify a false location. The service, however, may be vulnerable to difficult to mount attacks where the cost of mounting an attack exceeds the attacker's benefit (Section VII).

In contrast, it is instructive to look at the characteristics of currently available location services as offered by mobile network operators. These services have been tremendously successful in providing infrastructure-based and device-based location service. Operators are generally trusted. The services are cheap and ubiquitous. Nonetheless, there are several reasons why these services are often a poor fit for LAPs. First, each service is carrier-specific, locating only subscribers' devices. Continuous location tracking is often emphasized, rather than one-time location checks. Handset-based location (i.e., GPS) is often unavailable indoors. With the exception of E-911 service, operators have been reluctant to open the service to 3^{rd} parties. Hence, we seek an alternate infrastructure-based location service that is designed specifically for LAPs.

III. System Operation

A. Femtocells

Prior to discussing the operation of our system we pause to highlight key technical features of femtocells that we will exploit; please see Chandrasekhar et al. for more details on femtocell technology [6]. Femtocells are low-power wireless access points that operate in licensed spectrum to connect standard mobile devices to a mobile operator's network, typically using wired public internet access. Despite their limited wireless range (e.g., tens of meters), femtocells meet the various regulatory, compliance and spectrum use requirements of cellular base stations, including supporting location service. Though generally intended to improve cellular coverage inside buildings and areas with relatively poor cell tower coverage, we will use various femtocell properties (e.g., limited transmission range, exposed uplink, private ownership, integrated GPS) to authenticate the location of a femtocell-associated mobile device, without requiring mobile operator involvement or any modifications to operator infrastructure or services.

Residential femtocells typically support only 2-8 active mobile device associations (i.e., users), though such limits are frequently dictated by an assumption about the necessary available uplink bandwidth to ensure adequate quality-of-service for multiple active voice calls. Each call might consume roughly a continuous 50 kbs duplex rate, depending on the coding mechanism employed. Enterprise femtocells supporting 8-32 active users are rapidly emerging, with interconnection technologies and interference management in large-scale deployments being topics of considerable current research interest [7].

Voice calls may originate on femtocells, and subsequently be handed over to cell towers as callers move, however active calls originating elsewhere may not be handed to a femtocell. Femtocell owners may specify access control lists (e.g., family members only, any subscriber). In most ways a femtocell is best viewed as remotely managed and largely closed infrastructure that happens to reside on customer premises.

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 mobile operator's core network and femtocell is tunneled and encrypted with protocols such as the Encapsulated Security Payload (ESP) protocol [8]. Hence, confidentiality is assured against exactly the passive monitoring that we will describe in the next section.

B. System Architecture

Participating in a location authentication are:

- *Bob* is a mobile device user whose location is to be authenticated. He is willing to cooperate with the authentication, but we can not trust his assertion of his location. To be located Bob requires a voice-only phone (or smartphone) capable of associating with a femtocell at his current location.

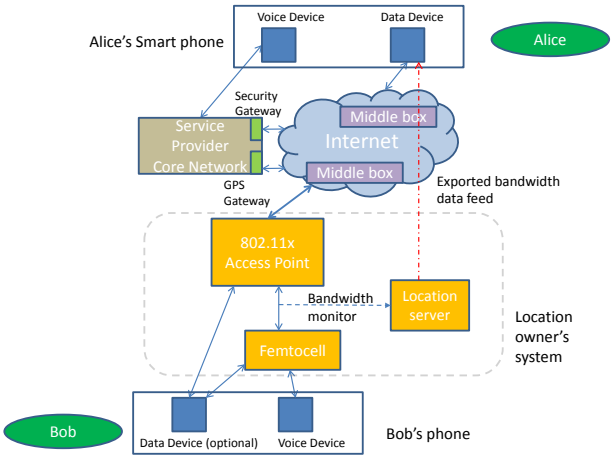

Fig. 1. Architecture of a single-carrier location authentication system.

- *Alice* seeks to verify Bob's location (with his explicit approval). Alice and Bob do not need to have any pre-existing relationship. However, in some applications Alice might have a relationship with Bob (e.g., his family member or employer) that compels his cooperation. Alternately, Alice might be a previously unknown LAP willing to extend Bob a benefit (e.g., discount) upon a location verification. Alice must have the equivalent capability of a smart phone, or more precisely a (mobile or wired) voice-only phone plus some minimal compute and display capability; a web browser suffices.
- The *Location Service Provider (LSP)* seeks to provide a public-access location authentication service. The location itself – say a coffee shop – might already offer a public WiFi service. The LSP is incented to provide location service to realize some benefit (e.g., to be known as a discount coupon distributor). 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, who can remain permanently anonymous to the LSP.

Figure 1 depicts the 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. Note that the computational burden of the location server is sufficiently small that in practice it can be run directly on either the AP or the femtocell.

Depending on his mobile device capabilities, Bob may optionally send or receive data traffic using either the femtocell (i.e., 3G or 4G) or the 802.11x AP. Note, of course, that though only Bob's device is depicted, other parties on site might be sharing these channels.

Consider the following basic authentication process:

1) Bob binds to femtocell (e.g., he calls Alice) so he is within several meters of the device.
2) Bob provides Alice with the LSP's URL, which offers complete site location information and access to the exported data stream.
3) The location server continuously monitors the (encrypted) AP-femto downstream link's average bandwidth each second and exports a stream of these values.
4) Alice calls Bob, impressing a voice traffic signature on the AP-femto link; the timing of Alice's call is determined entirely by her.
5) Alice monitors the exported bandwidth feed for characteristics of her own traffic.
6) If the behavior of the bandwidth feed convinces Alice that she is observing her own voice traffic traverse the AP-femtocell link, Alice confirms Bob's phone's association with the femtocell, and concludes that Bob is present at the specified location.

By calling Bob, Alice impresses a distinct traffic envelope on the AP-femtocell downlink, whose timing she controls. 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 system, this is roughly 50 kbs. She expects a similar decrease within a few seconds of hanging up. Though Alice can choose to visually display and study the bandwidth stream, in the next section we will describe an automatic detection algorithm she can run to reliably detect the presence or absence of her call, even when competing with significant cross-traffic from other users of the AP-femto link.

Though we will discuss attempts to defeat the location system in Section VII, we pause for a moment to consider two questions which arise immediately.

- *Can't the location server provide false location information to support Bob?* Yes. But the location owner doesn't know Alice. Alice can query the URL again at any future time, or have other parties query for her. If Alice or her proxy does not receive the same location information for each inquiry, she can invalidate the previous confirmation of Bob's location. Because the LSP does not know either Alice or Bob's identity, or even that an authentication took place, the LSP is unable to consistently provide transaction-specific false location information over time.
- *Can the location owner transmit a phony bandwidth stream to trick Alice into believing she is using this femtocell (when she is communicating with Bob at another location)?* Only Alice knows when and how she impresses a signal on the channel. The location owner is unable to guess when and how Alice (or her proxy) communicates with Bob to perform a verification.

IV. Signal Design and Detection

We next consider the challenging problems of 1) the design of the traffic signal Alice chooses to use to serve as her

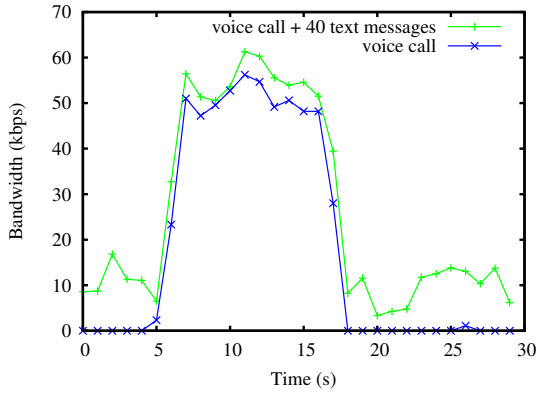

Fig. 2. Average bandwidth measurements captured from our prototype for a 30 second epoch containing an 11 second voice call (blue), and the same call with 40 randomly time-offset text messages (green).

fingerprint that she is indeed using the link, and 2) extracting that signal from other traffic generated by femtocell users on site (e.g., voice calls, text messages, web accesses), and 3) evaluating the probability that Alice herself is using the link, and consequently authenticating Bob's location.

We will limit our attention to traffic signals Alice can send with no change to existing mobile handsets or infrastructure; we will revisit this assumption later. Fig. 2 illustrates the captured bandwidth samples of a typical inbound 11 second voice call (blue). To represent a heavily used link, the figure also shows the aggregate bandwidth of Alice's call occurring concurrently with cross-traffic we constructed by adding 40 randomly time-offset copies of a captured text message in a 30 second interval (green). This 'noisy' signal is returned to Alice (typically after a few seconds delay) for her to evaluate the presence or absence of her call.

Suppose Alice places a *single* voice call to authenticate Bob's location. Though she initiates the call, Alice has imprecise control over both call establishment timing and the shape of the bandwidth envelope associated with the call's packets arriving to the monitored link. Prior to initiating a call test, Alice defines an observation window (or epoch) of duration T, taken to be sufficiently long to complete her call test and observe its effect on the return channel (e.g., $T = 30$ sec.). Alice records an estimate of call start time $\hat{t}_{start}$ and stop time $\hat{t}_{stop}$, and calculates an estimated call duration $\hat{D} = \hat{t}_{stop} - \hat{t}_{start}$.

The signal observed by Alice on the return channel in each epoch is $r[i], \quad i = 0, 1, \ldots, T-1$. Alice executes a detection algorithm on the received stream to choose between the hypotheses

$$r[i] = \begin{cases} s[i] + n[i] & H_1 : Alice's\ call\ present \\ n[i] & H_0 : Alice's\ call\ not\ present. \end{cases} \quad (1)$$

The received signal is modeled as the sum of two components: a signal $s[i]$ of duration D corresponding to the transmitted call, and a noise signal $n[i]$ which captures the bandwidth contribution of any cross-traffic on the link.

Our detection algorithm comprises 3 heuristics informed in part by the maximum-likelihood detection of signals in classical digital communication systems such as pulse-width and pulse-position demodulation. However, unlike a conventional communication system, the transmitted signal is not completely known to the sender, but can be constructed approximately. The 3 components of our detection algorithm are:

- *Signal amplitude* The amplitude test ensures that there is a contiguous set of bandwidth measurements in the epoch of sufficient magnitude to indicate the presence of at least one call.
  Since the envelope of the transmitted signal is only roughly known, we construct a signal $\tilde{s}[i]$ as an estimate of $s[i]$. Roughly speaking, this estimate is a sequence of $\hat{D}$ values of the nominal voice call bandwidth magnitude of 50 kbs. Observing the received bandwidth stream, Alice expects her call to begin at time $\hat{t}_{start} + R\hat{T}T$, where $R\hat{T}T$ is the estimated round trip time between the start of her transmission and its appearance on the return channel.
  Rather than use the values of the magnitude of the received signal directly, the amplitude detector calculates the convolution of the received signal with the estimated signal, i.e.,
  $$c[j] = r[j] * \hat{s}[j], \quad j = 0, 1, \ldots, 2T-1. \quad (2)$$
  If a signal is present on the link, we would expect the convolution to have a well defined peak; the estimated maximum value of $c[j]$ is calculated to be $M = \max_{j \in [0, 2T-1]} \hat{s}[j] * \hat{s}[j]$. The amplitude test checks if a sequence of k contiguous values of $c[j]$ (with $k < \hat{D}$) exceed some threshold αM, where α is a constant satisfying $0 < \alpha < 1$; if so, there is sufficient 'continuous' traffic on the link to indicate that a call may have occurred.
- *Signal edges* The edge test ensures that there is an observed measured bandwidth increase and decrease in the received signal near in time to the expected call start and stop time. Suppose the nominal measured call bandwidth magnitude is B, and a bound on the error in our RTT estimate is 2δ seconds. Observing the received signal, at the call start (stop) we would anticipate the measured bandwidth to rise (fall) by a value within some constant bandwidth $B \pm \epsilon$. Since we are averaging bandwidth over 1 second intervals, the expected rise or fall is not instantaneous but observed over measurements 2 seconds apart, i.e., for $i \in [\hat{t}_{start} - \delta, \hat{t}_{start} + \delta]$ we expect
  $$B - \epsilon < r[i+1] - r[i-1] < B + \epsilon, \quad (3)$$
  and for $i \in [\hat{t}_{stop} - \delta, \hat{t}_{stop} + \delta]$ we expect
  $$B - \epsilon < r[i-1] - r[i+1] < B + \epsilon. \quad (4)$$
  Both the up and down edges must be successfully detected for the edge test to indicate the presence of a voice call.

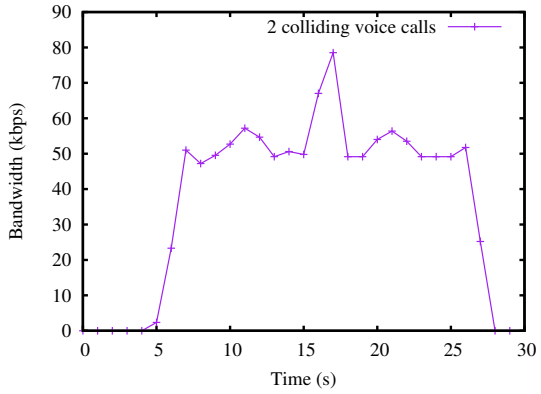

Fig. 3. Average bandwidth measurements when a second voice call slightly overlaps in time with the call we seek to detect.

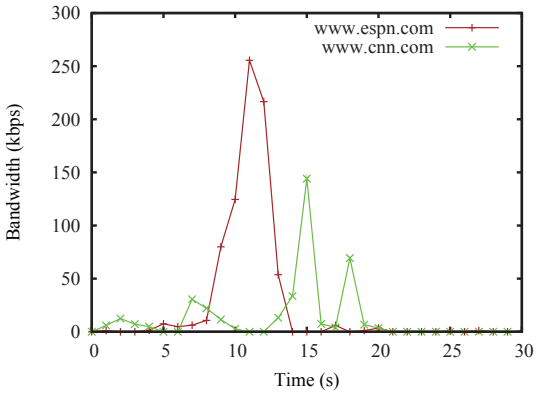

Fig. 4. Average bandwidth measurements for two web page downloads. While some downloads consume high bandwidth and are of short duration (ESPN), others may span a large fraction of the observation interval (CNN).

- *Convolution MSE* The convolution error test finds the minimum sum of the squared difference of the convolution of the received signal and estimated signal, and the estimated signal with itself, i.e.

$$min \sum c[j] - \hat{s}[j] * \hat{s}[j], \tag{5}$$

over all possible start times of the estimated signal. Minimizing this error helps determine the timing and duration of the transmitted signal embedded in the received signal.

The interfering cross traffic types we face are text messages, data transfers (primarily web downloads), voice calls, and control traffic. We will not consider control traffic here since 1) it consumes negligible bandwidth in the femtocell's 'operational' state, and 2) we have no control over its transmission.

Text messages are typically low bandwidth (e.g., 1 or 2 kbs) transfers of only a few seconds duration. To study a large number of text messages arriving independently of each other in an epoch, we sampled the bandwidth of an actual arriving text message, then summed copies of the bandwidth samples with random time offsets across an epoch. Figure 2 shows that the aggregated messages form smooth, time-homogeneous traffic; these message have virtually no impact on our ability to detect the presence of a voice call. Note, however, it is possible that several text messages could be queued in the mobile operator's network due to network congestion or temporary transmission failure, and suddenly be released in a burst. In such a case, even a modest number of arriving text messages (e.g., 10) could interfere with voice call detection.

It is the timing rather than the magnitude or number of interfering voice calls that cause a voice call detection failure. For example, one or more existing calls that outlast a test call look like time-homogeneous background traffic that do not inhibit call detection. But interfering calls that either start of stop near in time to the test call can disrupt its detection. Figure 3 shows the received signal when the voice call of Fig. 2 'collides' with a second voice call that begins at time $t = 15$ secs. Though this interference appears disruptive, note that the trailing edge of Alice's call is intact, and our algorithm successfully detects the presence of the call.

Data traffic (e.g., file transfers, web pages, streaming media) are typically high bandwidth (e.g., $100 - 400$ kbs) bursts of several to tens of seconds duration; Figure 4 shows typical examples. Our tests show that the transmission of even a single data transfer near in time to the start or stop of a test call is nearly certain to disrupt detection. In general, a voice call whose duration exceeds that of interfering data traffic promises to be most easily identified. Of course, Alice is always at liberty to issue a sequence of multiple test calls or data transfers of varying duration if she is uncertain that she is observing her own traffic.

V. Prototype System

We have constructed a single-carrier femtocell-based location authentication system prototype. When not performing tests, we observe typical system behavior by allowing the femtocell to serve occupants of our small office. Our prototype uses the Verizon 3G Network Extender (Samsung 2CS-2U01) femtocell; the bandwidth measurements reported here are representative of codecs deployed by Verizon. An x86-based commodity PC with multiple ethernet NICs running a standard Linux 2.6.34 kernel serves as the location server. In contrast to Figure 1, the server is located inline between the AP and femtocell, and traffic is forwarded between NICs via a standard network bridge. Bandwidth measurements are taken by reading a bridged interface directly with one of various, widely available tools such as *ifstat v.1.1*. The upstream link from the wireless AP is a shared DSL connection of approximately 1.5 Mbs upstream and 15 Mbs downstream.

An Apache web server offers users a static page with detailed site location information, including GPS coordinates, and a URL to access online bandwidth measurements. Real-time measurements are initiated on-demand, and exported via *netcat* on a separate interface to not impact bandwidth measurements. Verifiers are also able to request graphical views of a bandwidth measurements for an epoch; compact *sparklines* are generated with javascript for remote parties who are display-limited (e.g., smartphones). Our offline detection algorithm is compactly implemented in Python.

VI. An Analytical Model of Detection

We next develop an analytical model to determine the probability that we will be able to detect a voice signal in the presence of interfering cross-traffic. As long as the channel bandwidth is not fully utilized (i.e., saturated), a mix of interfering voice, text and data traffic bandwidth is additive. In general this noise forms a non-stationary process, though to begin we consider a stationary noise model.

Suppose that in each epoch we take the arrival times of individual interfering traffic bursts to be a Poisson point process with rate λ. The expected number of arrivals in an interval of duration t is then λt. Each interfering traffic burst has a variable bandwidth, and has a duration lasting several seconds, and hence interfering traffic bursts can overlap in time. These overlaps can cause the instantaneous bandwidth consumed by interfering traffic to vary greatly, in some cases exceeding the level of the voice signal we seek to detect.

Suppose we let $h[n], n = 0, 1, ..., T-1$, be a sequence corresponding to the bandwidth consumed each second (kbs) by a noise burst of maximum duration T, with $h[n] = 0$ for $n < 0$ and $n > T$. Each such burst might represent a single interfering message, such as a text message, voice call or data transfer, or an aggregate of several such messages arriving in an interval of duration T. The resulting noise model is similar to the well-studied, continuous-time 'shot noise' in electronic circuits, where poisson impulses arrive at random time instances t_i to a circuit with impulse response $h(t)$ to produce a noise process

$$\hat{s}(t) = \sum_i h(t - \hat{t}_i). \tag{6}$$

The probability density function of the shot noise $\hat{s}(t)$ ([9], p. 565) is

$$f(s) = e^{-\lambda T} \sum_{k=0} g_k(s)(\lambda T)^k / k!, \tag{7}$$

where $g_0(s) = \delta(s)$, $g_1(s) = g(s) * \delta(s) = g(s)$, ... , $g_k(s) = g_{k-1}(s) * g(s)$, ... , where the density function $g(s)$ satisfies

$$\int_0^s g(x)dx = Pr[h(t) \leq s]. \tag{8}$$

Informally, in our setting we seek the probability that sum of a sufficient number of interfering noise instances – each defined by a continuous-valued, discrete-time function $h[n]$ – arriving at an average rate λ will exceed some threshold value s (i.e., Pr$[\hat{s}(t) > s]$ and hence interfere with our detection. Suppose we model the bandwidth of each instance of noise by the sequence of bandwidths $h[n]$ depicted in Figure 5. In this example, the average duration of the noise instance is 10 seconds, and the magnitude is initially 14 kbs trailing off to 5 kbs. The corresponding probability density function $g(s)$ is shown in Figure 6.

Figure 7 depicts the cumulative distribution function of the aggregated noise as we vary the arrival rate of this noise. Recall that the voice signal we seek to detect has bandwidth of roughly 50 kbs. Hence, an aggregate noise bandwidth nearing that value near the start or end of a voice call will likely disrupt

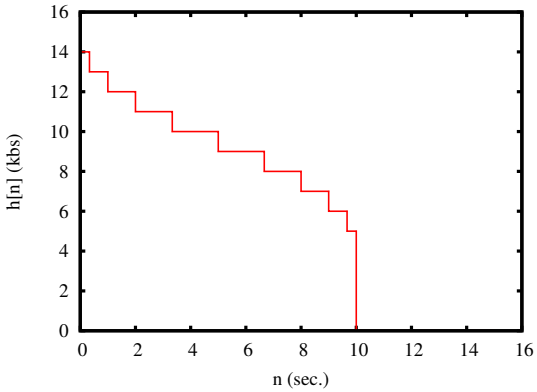

Fig. 5. An example of the bandwidth consumed by a single instance of interfering cross traffic. The duration of a noise burst is 10 seconds.

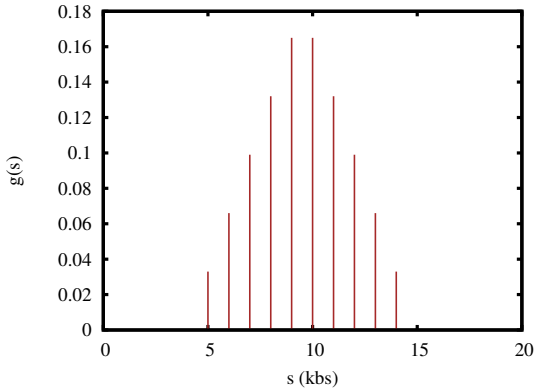

Fig. 6. The probability density function of a noise instance.

signal detection; in our implementation noise exceeding 40 kbs near the signal edge would be disruptive. The figure shows that the probability that our aggregate noise exceeds 40 kbs is $1 - .991 = .009$ when the noise arrives at rate $\lambda = 1$, but increases to $1.0 - 0.385 = 0.615$ when the arrival rate increases to 5. This latter result is consistent with our intuition; each arriving interfering signal has bandwidth of roughly 10 kbs, and the expected number of (overlapping) noise signals is 5, then we would expect the aggregated noise to exceed 40 kbs with probability of more than half.

VII. Security Analysis

We next conduct a security analysis of our system and identify plausible attacks. A plausible attack does not imply that the attack will be carried out, especially if the "cost" of mounting the attack exceeds the "benefit." For example, if an attack requires extensive planning and execution, but only results in the attacker obtaining a retail discount, then the attack may not happen.

A. Fooling the System

Bob may *fool* our system, i.e., he may prove to Alice that he is near a Location Server (LS) even though he is not. First, he may be able to *modify* the bandwidth feed observed by Alice and may impress a fingerprint on the feed so as to indicate his presence near the LS. For example, Bob may

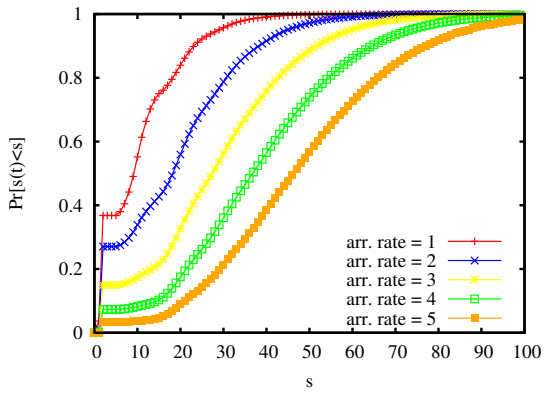

Fig. 7. Cumulative Distribution Function of aggregated noise as arrival rate λ varies.

compromise the LS and modify the bandwidth feed exported by the LS. Similarly, Bob may modify the bandwidth feed during its transmission from the LS to Alice, and may also be able to compromise Alice's device to present an appropriate bandwidth feed to Alice. Each of these attacks must be carried out in real-time when Alice calls Bob.

Second, when Alice calls Bob, he may be able to have an on-site colluder initiate a phone call through the location's femtocell, and hence may be able to impress a fingerprint on the bandwidth feed returned to Alice. For example, upon receiving Alice's call Bob could immediately call his associate from another phone and hang up as soon as Alice hangs up her phone. The bandwidth feed Alice monitors will now have a fingerprint similar to what she expects. In one variation, Bob may collude with the location owner to take a call when Alice calls Bob. All these attacks, however, may be detected due to the time lag between when Alice calls Bob and when Bob calls his onsite colluder; in particular, Alice's hanging up is particularly difficult to determine and respond to quickly. Moreover, Alice may use multiple, complex signals as her fingerprint, e.g., the bandwidth required to transfer an image of her choosing, that Bob can not easily mimic.

B. Disrupting System Operation

Our system is prone to *denial of service* attacks, i.e., Bob may not be able to prove his location to Alice even if he wanted to. First, an attacker may modify the bandwidth feed observed by Alice to indicate that Bob is not near the LS. For example, the attacker may compromise the LS and modify the exported bandwidth. Similarly, as discussed in the previous section, the attacker may tamper with data transmission or may compromise Alice's device to modify the observed bandwidth feed. Again, these attacks must be carried out in real time when Alice calls Bob.

Second, an attacker may perform a network DDoS attack on either Alice's network or the LS's network, and hence may prevent Alice from accessing the exported bandwidth data. Alice, however, can easily detect these types of attacks as she won't be able to connect to the LS.

Third, an attacker may prevent Bob's phone from being associated with the site's femtocell. For example, the attacker may simultaneously use many phones at the location and may saturate the AP-femtocell link. Then Bob will be blocked from using the femtocell to receive calls and Alice won't be able to reach Bob.

Fourth, the attacker may compromise the location server and may prevent the server from measuring the bandwidth and/or exporting the bandwidth. For example, the location owner may disable the location server and may deny all location authentication at her location.

In the first attack, Alice concludes that Bob is not near the LS even though Bob is. Alice cannot conclude anything about Bob's location in the last three attacks; Bob may or may not be near the LS. Hence the first attack is more severe than the last three, but is more difficult to mount.

C. Privacy Concerns

An attacker may be able to learn Bob's location by attacking our system. For example, if the attacker compromises Alice's device and gets access to her phone records and web access logs, then the attacker can learn Bob's location. Similarly, if the location server stores information about authentications and gets compromised, then an attacker may learn Bob's location [10] .

Our system assumes that the location of the femtocell and an associated site description are openly published. Further, the LSP's exported data feed is available to any remote party, and the feed's (in)activity might provide some general indication about the presence/absence of people on site.

VIII. Related Work

Despite nearly 2 decades of research, authenticating a mobile client's location remains difficult. Classical authentication system proposals often relied on distance bounding [11], [12]. Location proof architectures almost invariably rely on deploying trusted infrastructure, often distributing trust across multiple system elements in a complex authentication overlay. Such systems typically strive to achieve a high degree of confidence in verification, frequently using cryptographic protocols to bind devices and identities. In contrast, our system places no trust in infrastructure beyond their normal operation, and aims for a simple architecture that avoids the complexities of trusted infrastructure management, but provides authentication strength consistent with the commercial needs of existing LAPs.

Our approach is similar to related work in two aspects. First, in principle, we assume that we trust an entity's location and then prove that a mobile device is near the entity; the entity could be a femtocell or an 802.11x AP [10], [13]. Second, in implementation, we extend existing infrastructure by adding femtocells and location servers. In comparison, prior work requires certification authorities [14], APs capable of issuing cryptographic location proofs [10], [13], and trusted platform modules (TPM) [15], [16], [17]. Our approach, however, differs in one key aspect: we don't use any cryptographic primitives and rely on lightweight traffic signals for authentication.

Researchers have recently turned to physical characteristics of broadcast wireless channels (e.g., broadcast packets [18], [19], modulated power [20]) to form shared secrets based on channel observations to provide mutual co-location. To avoid trusted infrastructure and address collusion, other proposals rely on the presence of corroborators – sometimes establishing trust through reputation systems [21]. Bertino and Kirkpatrick explore Near-Field Communications and dedicated location devices to create an access control scheme [22].

Relatively little research has focused on the role femtocell technology can play in providing location services. Borgaonkar et al. describe how the lack of physical security makes femtocell location reporting an appealing target for hackers [23]. Indeed, it is precisely this lack of physical security – femtocells are located on customer premises – that permits us to construct an authentication service.

Despite the proposed location proof systems' broad diversity, most systems – including ours – remain vulnerable to collusive 'wormhole' attacks where a remote party colludes with an on-site associate to fake one's presence. Though distance bounding techniques may be a practical solution to these threats [24], it too suffers from weaknesses [25].

Despite these vulnerabilities, LBS systems have enjoyed tremendous success in practice. WiFi Positioning Systems (WPS) and hybrid WPS/GPS systems (e.g., Skyhook Wireless [26]) are the most popular location determination systems in use today for indoor/outdoor applications, despite locating only smartphones and other 802.11 equipped devices and being vulnerable to location-spoofing and denial-of-service attacks [27].

More recently, *location-as-a-service* or *Where 2.0* companies (e.g., LOC-AID [28] and Veriplace [29]) have begun to serve as intermediaries between mobile operators and third parties seeking client location. While promising, bootstrapping these services is challenging; each client and third party must proactively establish a relationship with each aggregator.

IX. Conclusion

To address the rapidly changing requirements of internet-based location application providers, we have proposed and demonstrated a new infrastructure-based location authentication solution that operates in a spontaneous, transaction-oriented, collusion-resistant fashion. Using instrumented femtocells, a simple voice call is enough for a calling party to verify a (cooperating) called party's location. We believe that this feature is unique among non-operator based location authentication system proposals, and crucial to reach the $2/3$rds of consumers that use voice-only phones.

Though we have focused on a system which does not require changes to mobile infrastructure or handsets, we believe that more robust authentications can be performed with the involvement of mobile operators. In particular, operators control the infrastructure, have preferential network vantage points, and can create easily discernible authentication fingerprints using simple techniques such as manipulation of packet headers (e.g., DiffServ Code Points).

References

[1] R. Want, A. Hopper, V. Falco, J. Gibbons, "The Active Badge Location System", *ACM Transactions on Information Systems*, Vol. 10, No. 1, January 1992, pp. 91-102.

[2] N. Priyanatha, A. Chakraborty, and H. Balakrishnan, "The Cricket Location Support System," *Proceedings of 6th Int. Conf. on Mobile Computing and Networking (MobiCom'00)*, Aug. 2000, pp. 32-43.

[3] D. E. Denning, P. F. MacDoran, "Location-Based Authentication: Grounding Cyberspace for Better Security," *Computer Fraud & Security*, Feb. 1996.

[4] T. Kindberg, K. Zhang, N. Shankar, "Context Authentication Using Constrained Channels", *Proceedings of the Fourth IEEE Workshop on Mobile Computing Systems and Applications*, 2002, pp. 14-21.

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

[6] V. Chandrasekhar, J. Andrews, A. Gatherer, "Femtocell Networks: A Survey", *IEEE Communications Magazine*, Vol. 46, No. 9, September 2008, pps. 59-67.

[7] Qualcomm, http://www.qualcomm.com/product-services/wireless-networks/femtocells

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

[9] A. Papoulis *Probability, Random Variables and Stochastic Processes*, McGraw-Hill, 1965.

[10] W. Luo, U. Hengartner, "VeriPlace: A Privacy-Aware Location Proof Architecture," *Proc. of 18th ACM SIGSPATIAL GIS 2010*, 2010, pp. 23-32.

[11] S. Brands, D. Chaum, "Distance-Bounding Protocols," *Advances in Cryptology - EuroCrypt*, Lecture Notes in Computer Science, 1994, vol. 765/1994, pp. 344-359.

[12] N. Sastry, U. Shankar, D. Wagner, "Secure Verification of Location Claims", *Proceedings of the 2nd ACM workshop on Wireless security (WiSe '03)*, 2003.

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

[14] 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.

[15] A. Dua, N. Bulusu, W. Hu, W. Feng, "Towards Trustworthy Participatory Sensing," *Proc. of USENIX HotSec*, August 2009.

[16] S. Saroiu, A. Wolman, "I Am a Sensor, and I Approve This Message," *Proc. of HotMobile 2010*, pages 37-42.

[17] P. Gilbert, J. Jung, K. Lee, H. Qin, D. Sharkey, A. Sheth, L. Cox, "YouProve: Authenticity and Fidelity in Mobile Sensing," *ACM SenSys*, 2011.

[18] Y. Wei, K. Zeng, P. Mohapatra, "Adaptive Wireless Channel Probing for Shared Key Generation," *Proceedings of IEEE Infocom 2011*, 2011.

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

[20] Y. Zhang, Z. Li, W. Trappe, "Power-Modulated Challenge-Response Schemes for Verifying Location Claims," *Proceedings of IEEE GlobeCom 2007*, 2007.

[21] M. Talasila, R. Curtmola, C. Borcea, "Location Verification through Immediate Neighbors Knowledge," *Proc. of the 7th International ICST Conference on Mobile and Ubiquituous Systems (MobiQuitous '10)*, Sydney, Australia, December 2010

[22] M. Kirkpatrick, E. Bertino, "Enforcing Spatial Constraints for Mobile RBAC Systems," *Proc. SACMAT'10*, 2010, pp. 99-108.

[23] R. Borgaonkar, K. Redon, J.-P. Seifert, "Experimental Analysis of the Femtocell Location Verification Techniques," *Proceedings of the 15th Nordic Conference in Secure IT Systems (NordSec)*, Helsinki, Finland, 2010.

[24] K.B. Rasmussen, S. Capkun "Realization of RF distance bounding," *Proc. of 19th USENIX Security Symposium*, 2010.

[25] C. Cremers, K. B. Rasmussen, S. Capkun. "Distance Hijacking Attacks on Distance Bounding Protocols," *Cryptology ePrint Archive: Report 2011/129*, 2011.

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

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

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

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