

# Authentication on the Edge: Distributed Authentication for a Global Open Wi-Fi Network

Paper Id: 1569031687

## Abstract

A global-scale, outdoor Internet access infrastructure with low infrastructure costs is finally attainable. Emerging projects are leveraging the proliferation of private Wi-Fi networks to build an access infrastructure from autonomous, independently owned Internet connections. These open Wi-Fi networks aggregate the broadband Internet installations at homes and other private properties into a global-scale shared infrastructure for ubiquitous Internet access. To ensure the traceability and accountability required by the broadband ISPs and private owners of these Wi-Fi networks authentication and authorization are needed. This paper describes Authentication on the Edge, a localized and distributed authentication method for open Wi-Fi networks. Three main ideas are used to adapt to the variability and unreliability of these networks: the use of certificate-based authentication, the distribution of certificate revocation list segments to all entities, and the self organization of access points into a social look-up network. Authentication on the edge combines centralized administration and operator assistance with distributed algorithms to confine the authentication to the edge of the network. These methods achieve the scalability needed for the overwhelming size and volume of a global network and increase resiliency against temporary failures of the infrastructure. Overall authentication delays are reduced by as much as 71% compared to existing authentication schemes applied in the same scenario. Analysis of SMS traces from a large cellular provider show that the social network backend will satisfy all authentication requests as a fallback mechanism.

## 1 Introduction

Ubiquitous Internet access at high speed and low cost has been a long-standing vision for years, attracting major effort from both academia and industry. As shown in recent studies [5] [7], indoor Internet connectivity is becoming pervasive in the U.S. as a result of the steady growth of broadband penetration to the home and the proliferation of 802.11-based wireless

local area networks (WLANs) deployed in households and other popular indoor locations. On the other hand, outdoor coverage for broadband Internet access is seriously lagging behind despite the explosion of wireless network technologies and deployments over the last decade.

We see a fundamentally new approach emerging to bridge this gap which has recently been conceived in both academia (PERM [28], P2PWNC [15], OBAN [14]) and industry (Fon [4], ABitCool [1]). Following the peer-to-peer spirit individual users share their subscribed broadband access over their private wireless routers under the supervision of a single authoritative party. The approach is born with unprecedented *scalability* and *cost-effectiveness* when compared to the other alternatives because each new user adds the necessary infrastructure to support himself. This method to construction and maintenance of a global-scale Internet access infrastructure based on *privately* owned wireless Internet access points (either residential or small business), henceforth called GIANT, is simple: a user opens up her own Internet connection to other participating users who are nearby the user's private Internet point of attachment. In return, when she is away from her own fixed Internet connection, the user is granted Internet access through other compliant users' wireless routers in range. A trusted third party is responsible for managing user credentials and handling billing. To ensure operation of the GIANT network on the *global scale*, management cannot be left to the individual owners of access points. For the majority of end users the task is technically too challenging and time consuming. More importantly, legal and regulatory requirements posed by Internet Service Providers (ISPs) for traceability and accountability further necessitate a management system that handles authentication and authorization among other services. For fair and controllable sharing at the global scale some *operator-assistance* is required.

In this paper we present our design and implementation of a scalable and resilient service for *authentication on the edge* (AGE) of GIANT networks that is used for access control of the nomadic

users. A set of *semi-distributed* algorithms and protocols operate under light-weight centralized coordination to authenticate users for controlled access. Our design of AGE seeks to strike the optimal balance between fully centralized approaches and fully distributed approaches in order to capitalize on the simplicity and ease of management of the former and scalability and ideal robustness of the latter.

Existing authentication mechanisms are unfit in GIANT for three main reasons: its vast physical size and large number of nodes, the variance and unreliability it suffers running over unmanaged Internet paths and untrusted hardware, and its security demands. To overcome these issues AGE localizes and completely decentralizes the authentication process itself while relying on a central server to manage, maintain, administer and disseminate updates of authentication material as a task separate from authentication itself.

AGE's mechanisms make it well suited to GIANT networks. AGE supports a single authentication authority allowing clients to access the service anywhere in the world with the same user id and authentication credentials. Authentication in AGE proceeds with as little user interaction as possible (the user only has to select the GIANT SSID for association) and AGE is resilient to the variable network conditions in GIANT including potential loss of connectivity to the central server.

At its foundation AGE is based on EAP-TLS authentication [9], an EAP method using certificates and private keys. The TLS authentication method avoids the use of passwords and allows mutual authentication of the authenticator (access point) and supplicant (mobile client). A central server operates the AGE certificate authority (CA) which manages the certificates for all GIANT users. The central server pushes updates to all relevant parties when authentication material changes. To continue operation in the face of server failure and avoid delays caused from accessing an authentication server in the Internet, each AGE access point (AP) runs a self contained authenticator, confining the authentication process to the wireless link only. The CA root certificate is embedded at every entity allowing clients and access points to verify each other's certificates locally. AGE uses Certificate Revocation Lists (CRLs) to inform AGE entities when a certificate has been revoked before its expiration. The CRL is also maintained by the AGE central server. Rather than querying the central server during each authentication, in AGE each entity (AP or client) carries the most recent CRL along with its certificate and exchanges the CRL during authentication. Because of

GIANT's large size with millions or tens of millions of global users at 28 bytes per entry the size of the CRL could strain the storage capacity of AGE entities. Therefore the AGE server divides the CRL into segments. Each CRL segment states the validity of a set of certificates, in essence providing a more up to date validation of a certificate. Even after presenting the CRL segment the situation might arise that the timestamp on a CRL segment violates the freshness constraints of an AP's security policy. In this case AGE maintains a peer-to-peer social overlay network for satisfying trusted freshness queries. The social overlay is built through friendships with other GIANT users and is maintained by the AP owner himself. AGE makes no stipulation on the freshness constraints imposed by any access point.

We have implemented AGE as a new EAP module for the `FreeRADIUS` server and `wpa_supplicant` Linux software packages. We ported the software to the `OpenWRT` open source router firmware for the Linksys WRT home wireless router. We wanted to show that AGE is immune from Internet delays and achieves low authentication times, continues correct operation in the face of failure of the central server, and is secure. Measurement results comparing EAP-AGE to EAP-TLS in the GIANT scenario show that AGE satisfies requests with between 49.7% and 71.6% lower delay, around 490 msec and 1614 msec, providing a faster and more predictable authentication. Using SMS traces from a large nationwide cellular provider we also show that with only 5 social connections as routing paths AGE's social network will provide a nearly 100% success rate of lookups. Finally we analyze AGE's security features and show that AGE achieves the same level of security against known threats as private wireless access points.

The rest of the paper is structured as follows: Section 2 reviews related work. Section 3 outlines the challenges for authentication and access control in GIANT networks and the reasons why current approaches are not appropriate. Section 4 describes the design of AGE and Section 5 covers the implementation details of the AGE framework. The evaluation and security analysis of AGE are found in Section 6. Future work is covered in Section 7 and the paper concludes in Section 8.

## 2 Related Work

The first subsection describes existing open Wi-Fi networks similar in concept to GIANT and the authentication methods in place for those networks. The second section covers projects which share ideas common to the approaches used in AGE.

## 2.1 Authentication in Open Wi-Fi Networks

The concept of opening privately owned access points for public usage has been suggested in both academia and industry. The PERM project [28] enables neighbors to share their broadband connections with each other for improved overall performance and reliability. PERM uses predictive mechanisms to multihome each participant's traffic across the available links. Although PERM uses certificates to verify neighbors' identities the sharing incentives and access control are not designed for mobile users. PERM instead targets long lasting mutual sharing relationships which do not necessarily exist in GIANT networks.

FON [3] and aBitCool [1] are communities for sharing wireless access points. A captive portal is used to redirect new users to a centralized authentication page. Each access point advertises an open SSID that uses no wireless encryption and requires a user to type his password whenever associating to a new access point. The user database is maintained at the central server which is queried for each network authentication, incurring high network delays and which is vulnerable to outages and DoS attacks making it unsuitable for GIANT.

Several companies provide access to thousands of hotspot access points at many locations in the world [2, 6]. Hotspot deployment takes advantage of the popularity of certain locations to for targeted deployment. Users pay a monthly fee to offset the cost of deployment. Authentication in hotspots is very similar to that used in FON and aBitCool where users typically authenticate through a captive portal. In general hotspots will suffer the same drawbacks when deployed in GIANT as FON and aBitCool.

The OBAN [14] project focuses on improving inter-AP handoff in open Wi-Fi networks by dividing the community into cells similar to cellular phone networks. The cell manager serves as an authentication proxy for each cell. OBAN requires the placement of servers in every cell, or one per 50,000 users as specified by the designers. Given GIANT's targeted global deployment the OBAN approach requires high maintenance overhead and might struggle to find an optimal deployment of manager servers to cover GIANT's unplanned growth.

The MoB project [12] provides an eBay like trading mechanism for sharing wireless access between individual users in public areas. MoB focuses on forming trustworthy relationships between unknown identities and accurate reporting of usage. MoB does not try to provide a static network and therefore does not address authentication and access control. P2PWNC [15] similarly focuses on ensuring fair trading rather than authentication in an open Wi-Fi network by us-

ing centralized servers to establish paths of trust for exchange of bandwidth credit. Both of these schemes could benefit from the accountability provisioned by AGE. MoB or P2PWNC mechanisms could be used on top of AGE to encourage participation in the GIANT network.

## 2.2 Other Related Work

The idea of localizing authentication to remote parts of the network has been used for network file systems [18]. Authentication information is prefetched and cached to enable fully local decisions when authenticating users. The caching of the filesystem user and group information results in some inconsistencies between the cached and central filesystem, but the delay is deemed acceptable. AGE follows a similar approach by having access points and clients store the CRL segments. The CRL segments at a given point in time might be older than the newest version at the central server. The benefit is that AGE authentication can operate without central server intervention.

AGE makes a tradeoff of distributing authentication information in the background instead of transferring it during the authentication. This tradeoff places a computation and communication cost on the directory server to generate and publish authentication material. The same tradeoff has been exploited in CRS [24] and CRT [20], where higher off-line computation is invested at the directory server to speed up on-line lookup for revoked certificates. We move one step further by pushing the CRL database away from the directory server towards the Internet edge where the authentications actually take place, therefore saving not only on-line computation at centralized servers but also on-line communications with those central authorities over the Internet.

The SPROUT project [23] uses social connections from the Friendster online social web site to augment a traditional DHT. The social links provide increased trust in the authenticity of the routing path. The results show improved query hit ratio and reduced delay in the face of malicious nodes. AGE also leverages the increased trustworthiness of social links to provide DHT-like timestamp queries for verifying CRL segments.

## 3 Challenges for Authentication and Access Control in GIANT

This section defines the characteristics of GIANT networks, the challenges GIANT poses on authentication methods and the failures of existing authentication schemes in the GIANT setting. Many efforts have been made to provide outdoor ubiquitous wire-




Figure 1: IEEE 802.1x LAN port authentication framework. The client (1) associates to an access point through which (2) it is tunneled to an authentication server on the LAN which finally (3) grants permission for outgoing connections.

less Internet access. Each of these solutions suffers either performance problems, expensive infrastructure, or low scalability.

In response to the shortcomings of existing solutions various projects have been started which attempt to leverage the popularity of residential broadband and the coverage of private wireless access points to create a global Internet access infrastructure. Given the steadily increasing residential broadband Internet penetration in the U.S., the growing density of Wi-Fi signals on the street, and the minimal cost of sharing Internet access through those Wi-Fi networks, we believe that this approach will lead to a ubiquitous Internet access infrastructure in populated areas. In essence, the GIANT approach aggregates the peer-to-peer installations of individual Wi-Fi Internet connections into a large-scale shared Internet access infrastructure. From an individual user's perspective the bandwidth of her fixed broadband Internet connection is made *portable* with a very small amount of one-time investment in the installation of a software-upgraded Wi-Fi access point (a.k.a. wireless router). Although these networks are easy for the operator to deploy due to their self expanding nature they are difficult to manage because of the vast size, unmanaged hardware, unreliability of the interconnection and private and untrusted control of the hardware.

In order for GIANT networks to see wide adoption there must be strict access control to ensure traceability and accountability, required for legal and regulatory reasons by the ISPs over which GIANT will run. Leaving the management of GIANT systems to individual end users as many community projects do is not viable as most end users are not willing or technically able to handle such tasks. Authentication and access control overseen by a central authority is a must.

There are two predominant approaches currently

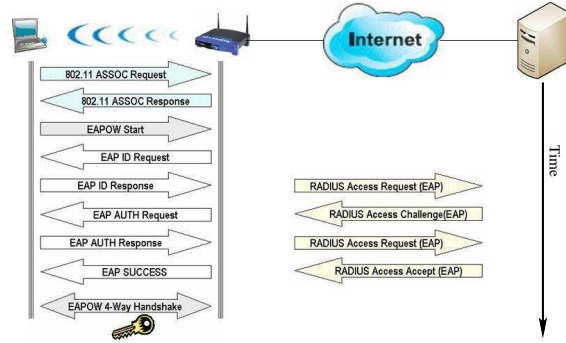


Figure 2: Direct application of IEEE 802.1x and EAP in GIANT. The authentication server is located in the Internet requiring RADIUS tunneled EAP messages to cross the WAN.

in use to manage access to wireless networks: captive portals and LAN port authentication schemes. Captive portals are deployed in hotspots and several open Wi-Fi projects. A captive portal is a firewall application which blocks all traffic going from the client to the Internet until the user has authenticated through a custom web page<sup>1</sup>. Once the user attempts to make a web connection her browser is directed to the authentication page where she enters her access credentials. The request is forwarded to a central database, usually through the Internet, which confirms or denies the request. Authentication is done at the application layer so there is no possibility for layer 2 wireless encryption. Because captive portals are password based, provide no wireless encryption and suffer long authentication delays when communicating to the central server they are ill-suited for GIANT networks.

Existing LAN port authentication schemes, such as WPA and 802.11i standards, are based on the IEEE 802.1x framework and the extensible authentication protocol (EAP) as illustrated in Figure 1. These schemes provide wireless encryption and support a wide range of different authentication methods. A supplicant, the client side access control entity, associates to an authenticator located in a LAN access point such as an Ethernet switch or Wi-Fi access point. The supplicant is authenticated through a closed port on the access point against an authentication server, for example RADIUS or MS Active Directory server. If the authentication is successful the authentication server signals the authenticator to open the port to allow authorized traffic from the supplicant, e.g. access to servers in the LAN or Internet. Although the standardized EAP methods work well in enterprise networks, GIANT changes many

<sup>1</sup>see NoCatAuth (<http://nocat.net/>), Chill-iSpot (<http://www.chillispot.org/>), WifiDog (<http://dev.wifidog.org/>)

assumptions about the operating environment which reduce the effectiveness of existing methods (as illustrated in Figure 2).

Firstly, the existing methods are not designed to scale to the size of GIANT. Intended for deployment in enterprise LANs with hundreds or maybe thousands of clients the authentication server will be seriously challenged in GIANT with potentially millions of active clients.

Secondly, because the existing frameworks are designed for LAN authentication it is assumed that all entities are deployed in the local enterprise network or even the same LAN segment. The GIANT system acts as an overlay operator (or an overlay wireless ISP) on an access infrastructure that is collectively owned by its actual users. Network management systems currently used in enterprise access network deployments [30] are not suitable in this context due to the lack of control over the interconnection medium in GIANT. Centralized authentication in GIANT is vulnerable to general Internet outages and distributed DoS attacks. The access points should be able to operate reliably and autonomously, even when access to network functions provided by the back-end management system is temporarily unavailable or inaccessible.

Thirdly, GIANT operates over the wild Internet spacing an authentication server far away from both clients and access points. Network anomalies, latency, and intermittent loss of connectivity are frequently encountered in wide area networks at the Internet scale. The entire authentication process using EAP methods involves several rounds of communication between the supplicant and authentication server. Traversing the Internet leads to high and variable delay in the completion time and in the case of packet loss it becomes difficult to gauge an appropriate timeout before restarting the authentication.

The unreliable wide-area connection has considerable impact on the existing IEEE 802.1x authentication process. Tunneling EAP messages through RADIUS does not use reliable transport meaning it is vulnerable to any lost packet. To quantify the effect on authentication success rates we measured WPA EAP-TLS authentications in a GIANT-like deployment. A RADIUS authentication server configured for EAP-TLS authentication was installed in the United States. A client supplicant in Europe attempted an EAP-TLS authentication through a local wireless access point once every minute. Over a seven day period on average approximately 1% of all authentications experienced a timeout. A maximum of 2.1% of one day's authentication attempts failed. Compared to local deployments with close to

100% success rate the penalty is high. Note that the impact of every lost packet is considerable given the default 30-second timeout in most implementations. Note also that these measurements were also taken under favorable conditions: the authenticator contacted an unloaded authentication server over a high speed LAN to the Abilene Internet2 network. Real deployments of AGE in residential areas connect over considerably slower and less reliable DSL or Cable connections. The capacity of the server and the network that hosts the server will be seriously challenged when handling the large-volume of authentication requests from the entire GIANT network. Also the impact of server failure will be much more severe when serving requests from millions of users. Although the overloading problem can be alleviated by re-directing the request to other authorized delegates, it will significantly increase the operation and maintenance cost, requiring constant adaption of the backend infrastructure to respond to GIANT's unplanned growth.

Finally, as access points in GIANT multiplex wireless connectivity of their owners and other nomadic visitors, multiple wireless networks with different security profiles, traffic shaping and prioritization policies need to be advertised and served by each of these access points. Virtualization of the access points [8] is a key requirement towards the wide acceptance of GIANT among performance and security concerned users. GIANT should provide equivalent security provisioning for both the public and the private networks. GIANT's independently operated access points introduce a number of security challenges. Clients must consider any access point to be malicious at all times. The lack of centralized control of the underlying infrastructure means that individual users can do as they please with their router. Also, authentication crosses untrusted paths opening the door for man-in-the-middle, wormhole, and spoofing attacks. Malicious clients and access points have the ability to collaborate with each other and perform coordinated attacks. These attacks could result in forged authentication material, replay of old material or the capture of authentication credentials. In the case of centralized authentication the well known authentication server is vulnerable to denial of service attacks through the Internet. The involvement of a third party is necessary to overcome these issues.

## 4 Authentication on the Edge

To address the challenges of scale, outages, delay and security we propose *Authentication on the Edge* (AGE) for GIANT networks, which localizes the authentication process on the access points and avoids



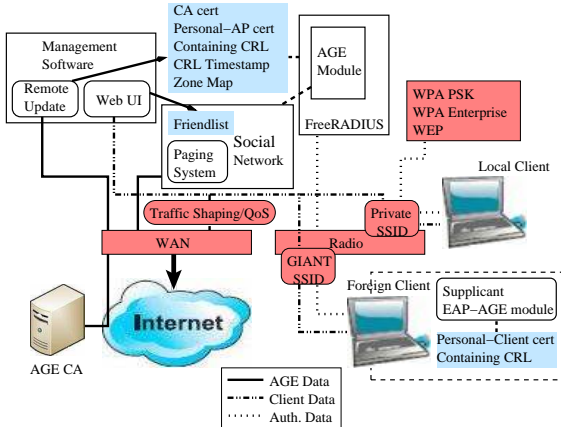


Figure 4: The entire AGE framework. Connecting lines show the flow of data through AGE. Pre-existing components are in pink (darkest color) and configuration information is in light blue (medium color).

ing the CRL to fit within EAP’s maximum packet size. Once the CRL transfer is complete the authenticator optionally begins verification of the CRL segment freshness by initiating a query to the AGE social network. While the query is completing the normal EAP-TLS handshake is resumed. When the TLS handshake reaches the point to verify the supplicant certificate it waits for the timestamp query to complete, if not already. If the presented CRL is valid the TLS verification continues as normal using the supplicant’s CRL. Otherwise the authentication fails.

## 5 AGE Framework

The AGE framework is shown in Figure 4. In this section we describe the design and function of each of the pieces in more detail.

### 5.1 CRL Segments

Assigning the containing CRL segment for a given node is an important task in AGE. If an authenticating client has the same containing CRL as the AP (a local authentication) then no CRL timestamp lookup is required. In order to reduce potential network overhead the assignment of clients to CRL segments should be done such that the number of foreign authentications required is minimized. Users typically spend most of their time in a small number of physical locations (in the same town, at work, at home, at a favorite cafe). Therefore the majority of authentications can be optimized if entities in the same physical locations share the same containing CRL segment. In this way users living in the same area, most likely to be authenticating each other, share the same CRL segment and therefore do not need to perform any remote CRL verification, saving both in authentica-

tion latency and network bandwidth. To perform the actual mapping of locations to CRL segments we define zones which cover a well defined physical area. Each zone is covered by one CRL segment which is assigned a unique CRL id. User certificates are generated with the CRL segment id embedded in the certificate. Users are bound to the same zone for the lifetime of the certificate. When the certificate expires the user can be reassigned to a different zone. To reason about CRL zones each access point keeps a listing signed by the AGE CA of all known zones and their locations - the AGE map.

### 5.2 Social P2P Network

As previously described, an authenticator with strict freshness requirements will need to validate each foreign CRL segment by looking up the segment timestamp. To ensure AGE’s operation when the central AGE service is unavailable the AGE access points arrange themselves into a p2p network. The network needs to support basic DHT functionality: mapping the CRL segment id to timestamps. However, traditional DHTs are untrustworthy making them vulnerable to replay attacks in AGE where malicious nodes can intercept requests for a particular key and respond with stale timestamps on behalf of their collaborating clients. This would allow an invalid client to continue accessing the network after it has been banned. Rather than using a traditional DHT, AGE access points organize themselves into a social overlay. Overlay links are maintained between entities based on the friendship of the owners for trusted routing.

On top of the trusted social network overlay there needs to be some type of routing scheme. The timestamp query is unique in that the query needs to only be forwarded to an area, not a specific location. Any access point with a containing CRL that matches the request can reply. Routing through the social network then is a matter of finding a friend or a friend of a friend, etc. who is in the target zone. Using the intuition that a friend closer to the target zone is more likely to know someone in the target zone. Greedily forwarding the request to the friend closest to the target should minimize the number of routing hops.

The AGE social network uses the location of each zone, specified in the AGE map, to forward messages between friends. A node’s location is the same as its containing zone. The problem of forwarding messages in the AGE social network then becomes the same as greedy forwarding in ad hoc and sensor wireless networks. Many solutions to this problem have been proposed in previous research [10, 19, 21, 22]. In each

of these protocols a node uses greedy forwarding until a void is reached in which the current node is closest amongst its neighbors to the destination. The protocols then switch to some variant of face routing to traverse the edge of the void and continue progress. Face routing protocols are well known to succeed in planar unit disk graphs. The AGE social network, based on friendship connections over the AGE zone map is clearly not planar nor unit disk. Using localized Delaunay triangulation [17] a node can convert his graph of friends into a planar graph for face routing. However, recent work has shown that not all variants of face routing succeed when routing in arbitrary undirected planar network graphs [16]. Based on this result we use the GFG algorithm [10] to perform routing in the AGE social network.

The basic GFG routing will still fail in the situation that a node has no social links outside of its own zone. Because every node in the same zone has the same location the GFG algorithm will consider all nodes in the same zone as a single node. To make progress within the same zone AGE uses *intra-zone* connectivity to select the best next hop. Selecting the friend with the most intra-zone connections ensures face routing can resume and increases the likelihood of finding a long distance link toward the destination. In the extreme case that a node has no neighbors with intra-zone links the node randomly selects one of its neighbors to forward the message. In our trace analysis only 6% of nodes did not have a direct connection to another zone. Our analysis in Section 6.3 of a nationwide cellular provider’s SMS traces shows that with at least 5 friends the network is fully connected.

### 5.2.1 Bootstrap Paging System

AGE supports a paging system to enable bootstrap of the social network. The paging system maps user ids with their current IP address. The paging system is built using a regular DHT because correctness of returned results is not important as will be shown. Any of the popular DHTs could be used [25, 26, 27]. Nodes insert their IP address into the DHT under their username. The values are signed with the node’s public key from its certificate. To find the current address of a friend nodes query the paging system using the friend’s username. Using the friend’s certificate if cached, the digital signature can easily be checked on the returned result. Next the querying node performs TLS authentication with the node at the given IP address to confirm that the friend is still executing at the given address. If at any point some authentication fails the friend will be considered unavailable.

The very first time a node joins the GIANT network it registers itself with the central AGE server.

At this time the central server feeds the new node with the IP addresses of the last  $n$  entities to register. The new node uses these addresses to bootstrap the paging system. On subsequent reboots the node uses the cached addresses of its friends. If a friend cannot be found in the paging system it is considered to be inactive. Each entity is responsible for updating its location information in the paging DHT.

### 5.3 AGE Insertion Tree

The central server maintains all of the authentication information including certificates, CRL segments and CRL timestamps. When some information changes the server must distribute the update to all affected entities. The central server has to store the username to zone mappings for each user in order to generate new certificates and maintain the CRL segments. When a new CRL segment update is ready the central server knows all of the users that should receive that update. The central server uses the above described paging system to find the current IP address of each access point and then pushes the update to each user in turn. As access points receive updated CRL segments they begin forwarding the new CRL to their friends who are in the same zone. This helps cut the latency of the update and the load on the central server. CRL timestamps are always distributed along with the CRL segments to all access points. Mobile clients synchronize with their home access point whenever they are associated updating the CRL segment as necessary. In addition, if a client roams to a foreign access point in its same zone it can compare the time stamp of the CRL segment from the access point to its own and keeps any newer CRL segments. New certificates can be pushed to individual entities following the same methods.

## 6 Evaluation

Our primary goals for our evaluation were to measure AGE’s performance, confirm its correctness and verify that AGE is secure against attacks. We compared the authentication times of our AGE implementation against industry standard implementations of other similar authentication methods in a realistic GIANT deployment. To confirm that AGE’s social network will sufficiently satisfy freshness queries we analyzed the SMS traces from a large nationwide cellular provider. We examined the relationships between users and the connectivity of users after dividing the users into AGE zones. The evaluation finishes with a thorough analysis of potential attacks against AGE and the mechanisms AGE uses to thwart those attacks.

	EAP-TLS/LAN	EAP-TLS/GIANT	EAP-AGE/GIANT
Client Processing Time	0.021556 [1.00x]	0.021171 [0.98x]	0.028369 [1.32x]
Server Processing Time	0.033561 [1.00x]	0.013828 [0.41x]	0.485165 [14.46x]
Network Delay	0.065948 [1.00x]	0.964469 [14.62x]	0.007132 [0.11x]
Total Auth. Delay	0.121065 [1.00x]	0.999468 [8.26x]	0.520666 [4.30x]

Table 1: Latency breakdown in seconds for EAP-TLS over LAN (EAP-TLS/LAN), EAP-TLS over the Internet (EAP-TLS/GIANT), and AGE over the Internet (AGE/GIANT). The increase compared to the baseline (EAP-TLS/LAN) is shown in brackets for each measurement.

## 6.1 Implementation of AGE Framework

We have implemented AGE as a new EAP type. New authentication modules for EAP-AGE were added to the popular Linux client supplicant software `wpa_supplicant`<sup>2</sup> and the open source RADIUS authentication server `FreeRADIUS`<sup>3</sup>. Our implementation modifies the existing EAP-TLS methods by inserting the CRL exchange immediately following the EAP-Identity exchange. A background thread runs to validate CRLs performing a lookup if necessary. The server configuration specifies the CRL tolerance to determine the acceptable age of a CRL. Using the `OpenSSL` library the CRL timestamp is checked. The validation thread also inserts a hook into the `OpenSSL` TLS engine for retrieving the CRL. When the EAP-TLS authentication reaches the point of CRL verification it calls the hook. The validating thread blocks until a reply is received from the social network. If the returned timestamp matches the timestamp of the client’s CRL then the TLS handshake ensues. Otherwise the validation thread returns an error and the TLS handshake fails. The `FreeRADIUS` AGE module is configured by a new `age` section in the `eap.conf` configuration file.

The AGE social network is implemented as a library exporting the single function `lookup` which takes as an argument a CRL segment id and returns the CRL timestamp. The AGE map used in routing is a mapping between zone id and GPS coordinates for the zone given in decimal format. The node pre-processes the file to determine which zones are its adjacent zones and which are long distant zones. After the friendlist table has been completed the node also precomputes the friend closest to every zone. When a request arrives the node checks the precomputed list to make the routing decision using either greedy or face routing. In total the access point software is around 3500 lines of code and the client `wpa_supplicant` module is only 500 lines of code. We ported all of the access point software to the `OpenWRT` Linux distribution for the Linksys WRT54G AP. The AGE firmware image is 2.8MB which fits within the

<sup>2</sup>[http://hostap.epitest.fi/wpa\\_supplicant/](http://hostap.epitest.fi/wpa_supplicant/)

<sup>3</sup><http://www.freeradius.org>

	Avg.	Min.	Max.
EAP-TLS GIANT	0.998509	0.987507	2.253010
EAP-AGE GIANT	0.557370	0.497009	0.639851

Table 2: Variation in authentication times using normal EAP-TLS authentication compared to EAP-AGE.

4MB WRT54Gv4 onboard flash memory.

## 6.2 Performance Evaluation

We use our AGE implementation to show how AGE reduces the typical authentication delay. We measure both network and processing latencies at all stages of the authentication process were measured. As a baseline we compare AGE’s performance to a typical IEEE 802.1x EAP-TLS deployment where the client connects to a powerful authentication server over fast and reliable LAN links, as shown in Figure 1. For performance measurements we also compare AGE to a direct application of IEEE 802.1x using EAP-TLS in a GIANT network as illustrated in Figure 2, where the authentication server is reachable over the potentially slow and unreliable Internet. An access point was installed on a campus in Europe configured to run either EAP-TLS or EAP-AGE. A Pentium 4 laptop equipped with an internal 802.11b/g card was used to perform a series of authentications using `wpa_supplicant`. In the EAP-TLS over LAN case the access point forwarded authentication requests to a second laptop in the same campus, configured to run the `FreeRADIUS` server. In the EAP-TLS over GIANT scenario the requests were forwarded to a Pentium 4 workstation in the United States to perform RADIUS authentication. Finally the local access point was configured to perform EAP-AGE as in Figure 3. The CRL tolerance was set low so that no social network lookup was required.

Table 1 shows the breakdown of authentication costs for each measured deployment, using EAP-TLS over LAN, GIANT and EAP-AGE over GIANT. The server side processing latency for AGE at the access point is significantly higher than at the Pentium-4 PC because the access point is equipped with a 200MHz microcontroller. There is also a slight increase in the client side processing time even though the client machine remained the same in all measurements. This

Zone	Members	Zone	Members	Zone	Members
1	11579	9	102333	16	10318
2	8635	10	41310	17	2718
3	39127	11	12236	18	131364
4	25348	12	71980	19	106329
5	420962	13	265402	20	13154
6	238867	14	562028	21	47395
7	102302	15	730577	22	264036
8	96880				

Table 3: Members per zone in SMS traces.

comes from the extra overhead of transferring and processing the CRL segment. However, compared to EAP-TLS over GIANT, AGE cuts the overall total authentication latency by half. The savings comes from the reduced network overhead, which is an order of magnitude lower than even the EAP-TLS over LAN configuration. Note that 94.3% of the AGE latency actually comes from the processing latency at the access point, which can be easily further reduced by a faster microcontroller. From Table 2 it is clear that the EAP-AGE protocol reduces jitter in authentication times offering more stability for sensitive applications like VoIP.

### 6.3 SMS Trace Analysis

AGE’s ability to correctly authenticate users depends on the success rate of the social network look-up. If the AP is not connected to the target zone then it will not be able to verify the client and will reject a valid client. On the other hand if a path exists the facing routing algorithm will succeed. To verify that social networks are well connected we analyzed SMS traces from a large nationwide cellular provider. The traces cover the individual SMS messages from over 6 million users. In the traces we consider two users friends if they have exchanged at least one SMS message. Cellular networks provide a natural method for dividing users into zones. The processing of SMS messages is handled by the mobile switching center (MSC) in a hierarchical fashion so that all messages sent from a group of cell towers in a given area are forwarded to a single MSC. Overall in our trace there are 22 MSCs. The home MSC of any user can be extracted from his telephone number prefix. Each user is assigned to a zone based on his MSC. Using this breakdown we have the zones with populations as shown in Table 3.

We first analyzed the friendships of each user. One third of the users have only one friend, likely infrequent SMS users. In contrast, some users have tens of thousands of friends. These users likely are bots or messaging services from the cellular provider. Over the entire data set the mean, min and max number of friends is 2.3, 1, and 46789. Over 99.99% of users

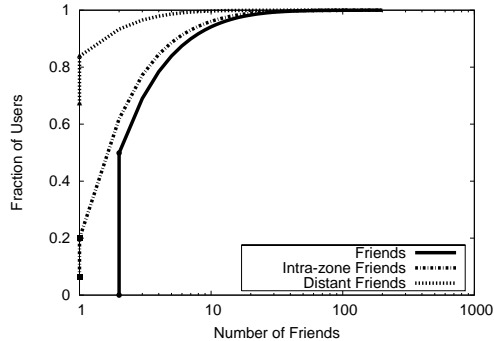


Figure 5: For each user the total number of friends, the number of intra-zone friends and the number of distant friends. The fraction of users with 0 of each type of friend is marked by the lowest points on the line. Only users with between 2 and 200 friends were considered.

have fewer than 200 friends. Filtering out the outliers with fewer than 2 friends (infrequent users) and more than 200 friends (bots) the mean, min and max change to 4.2, 2, and 200 friends. Figure 5 shows the number of friends, number of intra-zone friends and number of long distance friends using the filtered user set. Intra-zone friends are those in different zones and long distant friends are those not in immediately adjacent zones. We use the GPS coordinates of each MSC to determine which zones are adjacent. Using the filtered user set, overall 94% of all users had a friend in a different zone. 33% of users have at least one long distance link.

To decide if these connectivity properties could support the social network we evaluated the reachability of every user. We used the filtered set of users from above and followed the graph formed by friendship links of each user, searching for each zone. Figure 6 shows how many zones were reachable by individual users. It is clearly a bimodal distribution. The users in the left hand side of the figure, those with incomplete reachability, had on average between 1.71 and 3.17 friends. Those on the right hand side, the well connected users, had on average 5.02 friends. This finding suggests that having at least 5 friends is sufficient to maintain the reachability to all zones in AGE’s the social network.

### 6.4 Security Analysis

In our analysis we assume private keys have not been compromised and that the CA is trustworthy. Any router and client can be compromised and can collaborate with other compromised entities to carry out attacks. The AGE verification is based on TLS, which is known to be secure, with the additional new step for exchanging the CRL segment. This step does not change the functioning of the TLS protocol itself,

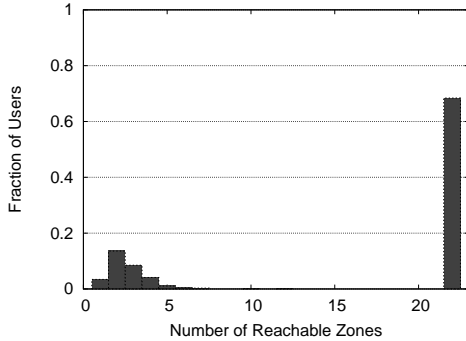


Figure 6: The maximum number of zones reachable by fractions of the population following only social links in SMS traces. The total number of zones is 22.

so TLS remains secure. Because the CRL and certificates are signed and verifiable by client and AP there is no threat of forgery. The certificate certifies the containing CRL id of the client so the client must present the correct CRL segment.

Private keys are never transferred, and therefore never vulnerable, except during the initial installation of the AGE software. We assume that the software (with embedded private keys) is distributed via some secure means, either from a trusted web server over SSL connections or physically installed by a trusted operator.

The biggest threat to AGE security comes from its social network backend. P2P overlay networks are very untrustworthy, threatened by Sybil attacks, routing misbehavior, forgery, and man in the middle attacks [11, 13]. The design of AGE’s social overlay overcomes most of these challenges. The identity of each AGE member can be verified using the member’s certificate thus overcoming the Sybil attack. Because the AGE social overlay only stores digitally signed timestamp files which can be verified with the CA public key, there is no threat of forgery of the timestamps. Malicious nodes can drop messages which will cause the current authentication to fail. By building the overlay based on social connections to friends, the routing paths are considered more trustworthy and therefore less likely to contain malicious nodes. In addition, timestamp requests can be fulfilled by any member of the target zone, so if one route fails there is likely another route available.

A malicious node in the social overlay may also collaborate with a client. The only attack the malicious nodes can perform is to replay old timestamps of a CRL matching the last state in which the client was not revoked. The only way to avoid this attack is to gather multiple timestamp files from different hosts. This technique is not guaranteed to succeed but it increases the likelihood of finding the correct

result. Only one copy of the newest timestamp is needed from amongst all samples. Because a CRL has an expiration date the timestamp replay attack will only have a limited window of success.

## 7 Future Work

We have not addressed the issue of how to effectively distribute CRL segments to AGE clients. There may be long periods of time when the client does not associate to the home AP for example during a business trip, or holiday. If the client has invalid authentication material then it cannot authenticate and therefore cannot download the most recent material. One possibility is to open a port on the AGE access point which is used solely for updating authentication material. Another option is to allow clients a grace period after a failed authentication during which the client has the opportunity to update its credentials.

We leave billing, fraudulence detection and certificate revocation policies undefined in this work. Each of these issues, while important for operation of GIANT networks is highly dependent on operator preferences and policies.

An important optimization for GIANT is support for fast hand-offs between GIANT access points. Most existing solutions rely on a proxy authenticator to handle fast rekeying or pre-authentication by the client at neighboring access points. GIANT’s unmanaged infrastructure makes such approaches difficult or impossible. A future line of research will examine how to best support fast re-authentication in GIANT networks.

An optimization for AGE’s social p2p network is centralized guidance for maintaining friendships. AGE perimeter routing will benefit if every user has connections to every adjacent zone. Although the lookup in the AGE overlay is fully distributed, the overlay construction and maintenance can be managed and assisted by centralized management systems such as delegates of the back-end authority, similar to the design of HOURS [29].

One vulnerability inherent in GIANT networks is the threat of snooping. Although AGE protects the over-the-air transmissions from snooping attacks using the latest WPA encryption, a malicious AP owner can easily install a sniffer inbetween the router and the broadband link. The only solution to this problem is to employ some form of end-to-end encryption or to tunnel all sensitive traffic through an encrypted proxy like VPN.

## 8 Concluding Remarks

There exists a hot on-going debate on the social and legal impacts of sharing Internet access through Wi-Fi. Residential Internet users typically understand

how to enable encryption and MAC filtering to prevent intruders and free-riders into their wireless networks. As GIANT networks expand in popularity one major concern from the perspective of ISPs is that GIANT potentially increases the load on the local/regional ISP networks. For similar reasons very few ISPs allow a user to share the Internet connections, except for a few such as Speakeasy, Stephouse Networks, and EasyStreet. We believe that GIANT will benefit not only GIANT users, but also broadband ISPs in attracting more broadband subscribers.

GIANT networks offer a promising new approach for outdoor Internet access. AGE running in GIANT networks will pave the way for wide scale adoption of GIANT as users can be confident in the security and performance enabled by AGE and operators can be confident in having the traceability and accountability they require. AGE's mechanisms of localizing and distributing authentication will enable security and accountability to easily scale with the enormous growth potential inherent in GIANT networks.

## References

- [1] Abitcool wi-fi community. <http://www.abitcool.com/>.
- [2] Boingo wireless service. <http://boingo.com/>.
- [3] Fon. <http://www.fon.com/>.
- [4] Global wi-fi community. <http://www.fon.com/>.
- [5] June 2006 bandwidth report. <http://www.websiteoptimization.com/bw/0606/>.
- [6] T-mobile hotspots. <http://hotspot.t-mobile.com/>.
- [7] Wi-Fi Surpasses Ethernet for Home Networking. [http://www.parksassociates.com/press/press\\_releases/2005/gdl2.html](http://www.parksassociates.com/press/press_releases/2005/gdl2.html).
- [8] B. Aboba. Virtual access points, May 2003. IEEE P802.11 Wireless LANs.
- [9] B. Aboba and D. Simon. RFC 2716: PPP EAP TLS authentication protocol, Oct. 1999.
- [10] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.
- [11] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach. Secure routing for structured peer-to-peer overlay networks. *SIGOPS Oper. Syst. Rev.*, pages 299–314, 2002.
- [12] R. Chakravorty, S. Agarwal, S. Banerjee, and I. Pratt. Mob: a mobile bazaar for wide-area wireless services. In *Proceedings of ACM MobiCom '05*, pages 228–242, 2005.
- [13] J. R. Douceur. The sybil attack. In *Proceedings of International Workshop on Peer-to-Peer Systems*, 2002.
- [14] E. Edvardsen. Fixed and mobile convergence. In *Proceedings of BroadBand Europe 2004*, 2004.
- [15] E. Efstathiou, P. Frangoudis, and G. Polyzos. Stimulating participation in wireless community networks. In *Proceedings of IEEE INFOCOM*, 2006.
- [16] H. Frey and I. Stojmenovic. On delivery guarantees of face and combined greedy-face routing in ad hoc and sensor networks. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, 2006.
- [17] J. Gao, L. J. Guibas, J. Hershberger, L. Zhang, and A. Zhu. Geometric spanner for routing in mobile networks. In *MobiHoc '01: Proceedings of the 2nd ACM international symposium on Mobile ad hoc networking & computing*, 2001.
- [18] M. Kaminsky, G. Savvides, D. Mazieres, and M. F. Kaashoek. Decentralized user authentication in a global file system. In *Proceedings of ACM SOSP*, 2003.
- [19] B. Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proceedings of ACM MobiCom*, 2000.
- [20] P. Kocher. A quick introduction to certificate revocation trees (crt). In *Technical report, ValiCert*, 1990.
- [21] F. Kuhn, R. Wattenhofer, and A. Zollinger. Worst-case optimal and average-case efficient geometric ad-hoc routing. In *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, 2003.
- [22] B. Leong, S. Mitra, and B. Liskov. Path vector face routing: Geographic routing with local face information. In *Proceedings of ICNP 2005*, 2005.
- [23] S. Marti, P. Ganesan, and H. Garcia-Molina. Sprout: P2p routing with social networks. In *Proceedings of First International Workshop on Peer-to-Peer and Databases (P2PDB 2004)*, 2004.
- [24] S. Micali. Efficient certificate revocation. In *Technical Memo MIT/LCS/TM-542b*, 1996.
- [25] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker. A scalable content-addressable network. In *SIGCOMM '01: Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*, 2001.
- [26] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location and routing for largescale peer-to-peer systems. In *Proceedings of Conference on Distributed Systems Platforms, Heidelberg*, 2001.
- [27] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup protocol for Internet applications. *IEEE/ACM Transactions on Networking*, 11(1):17–32, February 2003.
- [28] N. Thompson, G. He, and H. Luo. Flow scheduling for end-host multihoming. In *Proceedings of IEEE INFOCOM*, 2006.
- [29] H. Yang, H. Luo, Y. Yi, S. Lu, and L. Zhang. HOURS: Achieving DoS resilience in an open service hierarchy. In *Proceedings of DSN*, 2004.
- [30] L. Yang, P. Zerfos, and E. Sadot. Architecture taxonomy for control and provisioning of wireless access points (CAPWAP). IETF Request For Comments (RFC4118), June 2005.