

Flexible Resource Allocation and Composition Across GSM/3G Networks and WLANs

Mohammad Al-Fares
Comp. Science and Eng.
UC, San Diego
La Jolla, CA 92093
malfares@cs.ucsd.edu

Martin Johnsson
Ericsson Research
Isafjordsgatan 14
164 80 Stockholm, Sweden
martin.johnsson@ericsson.com

Per Johansson
Calit2
UC, San Diego
La Jolla, CA 92093
pjohansson@soe.ucsd.edu

Amin Vahdat
Comp. Science and Eng.
UC, San Diego
La Jolla, CA 92093
vahdat@cs.ucsd.edu

ABSTRACT

A significant challenge facing wireless and mobile networking is seamless global network connectivity. As expected, no single network operator and no single technology dominates the wireless access landscape. Instead, heterogeneity and change characterize both operators and access technologies. On the face of it, such a competitive landscape should support rapid innovation and more rapid rollout of global connectivity. Unfortunately, end users today are still unable to seamlessly leverage the broad array of available networks. Thus, while coverage may be nearly ubiquitous, access is not. One reason for this difficulty is that wireless providers cannot quickly and efficiently enter into the equivalent of roaming agreements. In this paper, we present a system architecture to enable seamless composition of wireless network access across a range of technologies. Importantly, we do not require pre-existing agreements on the part of operators or active involvement on the part of end users. Rather, operators advertise network capabilities and price while users have built-in preferences for cost, performance, battery life, etc. to allow end devices to both choose an appropriate network and to ensure that end-to-end billing takes place appropriately. We have completed an initial system prototype and our performance evaluation is promising for potential future low-overhead deployment.

Categories and Subject Descriptors

C.2.4 [Distributed Systems]: Distributed applications; C.3.9 [Special-purpose and Application-based Systems]: Process control systems

General Terms

Design, Economics, Standardization

Keywords

Resource allocation, mobile roaming, network composition

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiArch'08, August 22, 2008, Seattle, Washington, USA.
Copyright 2008 ACM 978-1-60558-178-1/08/08 ...\$5.00.

1. INTRODUCTION AND MOTIVATION

The current standards and systems for cellular and other types of wireless networks are challenged by recent networking trends. The business and technology environment is becoming both more heterogeneous and more dynamic. The cellular and wireless world is facing a rapid development and planning of new radio access technologies such as Wi-Fi, WiMAX and LTE. At the same time, we also see new types of actors emerging, such as virtual operators, “white label” service providers (e.g. the Cloud [8]), and aggregators/brokers (e.g. Boingo [4]). Avoiding further market fragmentation, providing an integrated view from a user perspective, and coping with the general high level of dynamicity in both technology and business require new concepts and mechanisms to accommodate roaming and trading of services and resources ubiquitously across various networks. Agreements between users and networks and also between networks must be possible to establish on-the-fly connectivity with little or no user intervention, as users like to simply connect and use the services of whatever wireless networks are available.

Network Composition is a concept that was first studied and developed in the Ambient Networks project [17, 9]. Network Composition defines one uniform procedure, where also user devices are treated as individual networks capable of establishing so-called *Composition Agreements* dynamically at runtime without any prior agreements between the owners of the different interacting networks. As an example, users would be able to roam into a visited network and use the services and resources of that network without the need for any prior arrangement between the operator of the visited network and the user’s home operator. By executing the composition procedure, composition agreements would be settled both between the user and the visited network, as well as between the visited network and the home operator. Such agreements will control what resources and services would be made available for each of the composing networks, and will regulate compensation for the usage of these services and resources.

The primary goal of this paper is to create a resource peering framework that implements this seamless and transparent network composition between arbitrary voice/data wireless network operators, regardless of the access technology (large GSM/UMTS networks, small Wi-Fi hotspots, etc.), or type of service required (voice or data). In addition, the control plane of this composition should be a single global network that all operators can access, namely, the Internet. Towards this end, this paper presents resource allocation mechanisms to allow arbitrary peering between multiple providers and access technologies. These mechanisms are described in Section 2.

Furthermore, since today’s handheld mobile devices and laptop computers incorporate multiple wireless access technologies

(GSM, Wi-Fi, Bluetooth, etc.), when a device is in the coverage area of multiple network operators simultaneously, users should be able to automatically switch to the best access method depending on several user-defined criteria (e.g., service type, price, power consumption, available bandwidth, signal strength, etc.). Among other reasons, the compensation heterogeneity of those methods is a central real-world barrier that makes such a dynamic switching capability difficult, if not impossible. However, a unified resource reservation and allocation system across multiple access technologies should eliminate this barrier.

2. BACKGROUND

SHARP [11] is a distributed framework for secure resource federation and peering through bartering. Initially developed and tested for PlanetLab [7], a global-scale networking research grid, the SHARP system is based on cryptographically signed capability objects termed *claims*. These in turn are separated into *tickets* and *leases*, which represent promises and rights, respectively, to control specific resources at specific times.

GENI [18] is a network infrastructure designed to support experimentation with a variety of novel protocols and applications. Its proposed resource allocation system incorporates many of the ideas and philosophy of SHARP, but transforms it into an full-fledged economy with the addition of *tokens*. While the value of a token is abstract and open to interpretation by design, tokens may be thought of as the fundamental units of value in the system, and thus can provide a common ground to value and trade resources (similar to digital cash). Tokens are signed by the issuer and canceled on expenditure, and are atomic, unique and not reused. Tokens can be issued and grouped in *token sets* to aggregate their transfer, management, and verification.

A resource in GENI is an abstract object that is fully described by its associated *RSpec* document. An instance of such a resource could be a generic ‘machine-hour’ for example, in which case the RSpec describes the physical resources of that machine (CPU, RAM, HDD, etc.), its network connectivity (1Gbit Ethernet), the type of service (dedicated vs. shared), its geographic location, the start time, etc. Resource discovery mechanisms, such as SWORD [19], would be used to try matching a user’s query for a given resource to previously-donated resources that match a subset of the queried fields in the RSpec.

2.1 Resource Allocation Components

We give a brief overview of GENI resource reservation elements, as described in the GENI specification [2]. Interested readers are referred to the specification for more details.

Site Manager.

The site manager (SM) is in charge of user-account creation and maintenance. It donates a fraction of the resources under its control for tokens and distributes those tokens among its users as it sees fit. There is one site manager per ‘institution,’ and every site manager has a parallel component manager that keeps track of previously donated resources.

Component Manager.

The component manager (CM) honors tickets to previously donated resources and issues the leases to those resources and gives the user access (e.g. instantiates a virtual machine and returns login credentials, modifies an access control list, etc.).

Resource Broker.

A resource broker (RB) acts as a clearinghouse for donated resources from several site managers, and matches resource requests with previously donated resources. When a resource is donated a number of tokens is returned in exchange according to a predefined policy or going rate for that resource. The resource broker accepts tokens in exchange for resource tickets, which the user presents to that resource’s component manager in order to gain access. In a decentralized system, there would be many resource brokers; however, a broker is responsible only for resources donated directly to itself (i.e. RBs share no state, and therefore no coordination or synchronization is necessary between them). Of course, to which resource brokers an SM donates resources depends on which other institutions the SM expects to barter with, and the distribution for each one.

User Interface.

The user interface (UI) allows the user to: (a) discover what resources are available, (b) request a reservation from an RB, and (c) present tickets to CMs to gain access to resources. The UI keeps track of a user’s token balance and outstanding tickets, and also provides an interface to the home provider from which the user can buy/refill tokens.




Figure 1: Resource reservation steps. Notice that the site and component managers are paired, as they share state regarding previously donated resources.

2.2 Resource Reservation Steps

To reserve and consume a given resource, the following steps must be taken. Figure 1 illustrates this reservation protocol.

1. The site manager responsible for a resource makes a donation to a resource broker of its choosing for an agreed upon number of tokens.
2. When a user wants to ‘buy’ this resource, it contacts the responsible resource broker and ‘pays’ the required number of tokens and when the tokens’ validity is checked, gets a ticket for that resource in return.
3. The ticket is presented to the component manager responsible for the resource, which verifies its authenticity. The component manager then gives the user access credentials for that resource.

3. ARCHITECTURE

To demonstrate the flexibility of our approach, we show as examples mapping the aforementioned components onto two types of service providers: a 3rd generation cellular network, UMTS, and a Wi-Fi network. Note that the UMTS network is backwards compatible with GSM/GPRS¹ (often referred to as 3GSM) and that our proposed approach will also work with GPRS capable handsets since they offer Internet access.²

3.1 UMTS Network Architecture

The UMTS network architecture is divided into two domains; a circuit-switched (CS) for voice services and a packet-switched (PS) domain for data services. Furthermore, the network hierarchy has three main hierarchical components, the:

1. User Equipment (UE); the
2. Radio Access Network (RAN); and the
3. Core Network (CN).

The RAN and CN both contain functionality to be backward compatible with GSM/GPRS devices (MS).

The UE contains the *UMTS Subscriber Identity Module* (USIM), or just SIM in the case of a GSM/GPRS MS, that stores the user's credentials. The USIM/SIM chip can be transferred between different UMTS devices, i.e. it identifies the user and not the device.

The RAN manages the radio resources and contains functions for handover, call admission control, etc., for both the CS and PS domains. Radio network specific nodes for GSM/GPRS are also hosted in the RAN to allow backward compatibility.

In UMTS, the radio base station is called *Node B*, and a *Radio Network Controller* (RNC) manages the radio resources of several Node B units. Together they form a *Radio Network Subsystem* (RNS). Typically, an RNC serves several hundreds of Node B units, each of which in turn serves an area roughly of on the order of two thousand subscribers.

For GSM/GPRS part of the system, the radio base station is called a *Base Transceiver Station* (BTS), and the radio resource management node is called a *Base Station Controller* (BSC). Together they form a *Base Station Subsystem* (BSS).

Both the RNC and BSC nodes divide the traffic from the mobile devices (UE or MS) into circuit switched (voice) and packet switched (data). In turn, voice and data traffic is separated and passed on to the circuit switched and packet switched domains of the core network, respectively.

In the CS domain of the CN, a *Mobile Switching Center* (MSC) routes voice calls from both UMTS and GSM devices to other mobile networks, or to the fixed landline network (PSTN) via a *Gateway MSC* (GMSC). The MSC also a central role in the authentication procedures of a device's USIM/SIM when it is attaching to the network. The *Home Location Register* (HLR) contains the user credentials in his home network, and is queried by the MSC to verify the user's authenticity. A user is then authenticated via a challenge-response exchange sequence.

In addition to user authentication, the HLR also supports user location management together with a *Visitor Location Register* (VLR). The VLR, which is collocated with the MSC, keeps track of all foreign roaming users on a network and informs the HLR of

¹GPRS (General Packet Radio Service) augmented the voice only (2G) GSM network with IP network access capability.

²The main difference between GSM/GPRS and UMTS is the introduction of a new radio interface (WCDMA) that offers much higher data rates than the TDMA based GSM system.




Figure 2: Location of reservation components in the UMTS network architecture.

their home networks of their current Location Area (LA) so that incoming calls can be routed to the visited network.

In the PS domain of the CN, two types of nodes handle the data traffic: the *Serving GPRS Support Node* (SGSN) and the *Gateway GPRS Support Node* (GGSN). The tasks of the SGSN include logical link management, packet routing to/from the UE/MS, mobility management, location management; a location register in the SGSN stores user profiles and locations of all users registered with the SGSN. The GGSN serves as the UE/MS's gateway router to the Internet and also provides the UE/MS with an application-level IP address.

User mobility in the PS domain is maintained by the *GPRS Tunneling Protocol* (GTP), where GTP tunnels are set up between SGSN and GGSN. As user devices move in and out of the area serviced by different SGSNs, the traffic is moved between GTP tunnels while still being anchored to the same GGSN (similar to how IP tunnels are moved between Home and Foreign Agents in Mobile IP). This way the user device will always maintain the same IP address while moving.

We refer the reader to [13, 14] for more details about architecture, addressing and mobility management in UMTS and GSM/GPRS networks.

3.2 UMTS Network Mapping

The user interface is implemented as an application on the mobile station's USIM (the UMTS SIM card). This application manages the resource discovery queries, and keeps track of the current token balance and outstanding unredeemed tickets.

We co-locate a site manager and component manager at every RNC/BSC, with their associated resource broker located in the Core Network, common for both the CS and PS domain (Figure 2). The site manager dynamically monitors load on the RNS/BSS

and donates slices of its unused capacity (voice and/or data) at small regular intervals to the RB. Note that it is possible to issue multiple tickets for the same donation if it was specified as a shared resource. The donation and reservation in this case is a purely local decision, as no coordination or synchronization with other brokers is necessary.

The control plane between these components may be network dependent, but the only requirement here is that the resource brokers have an interface to the Internet to bill the operators of roaming users. We briefly discuss the scalability of the system in Section 3.6.

3.3 Wi-Fi Hotspot Mapping

The mapping for an 802.11 wireless network is logically simpler; the resource broker, component and site managers all functionally reside at the access point router (Figure 3). The router is configured as an open network, such that all new users are assigned an IP address using DHCP, and the control over access is achieved through dynamic management of firewall rules. The AP router adds a firewall rule to allow traffic to a defined control port and specifically addressed to the router to allow communication with the local reservation components. In addition, the AP router keeps a list of active users, and blocks traffic that is not explicitly addressed to or from an active user, or not matching the control port.




Figure 3: Location of reservation components in a 802.11 hotspot. The users include both laptop owners and wifi-enabled smart-phones.

The resources an SM donates can be shared best-effort (e.g., the common ‘access-hour’ at airports and cafés), a guaranteed block of specified capacity (using traffic shaping), or a cap on a certain data size. In this case, the UI of a new user communicates with the RB using the control channel and requests the ticket of a certain option and presents its tokens in exchange. The RB must verify the tokens’ authenticity before granting access, and therefore must contact the user’s original service provider, whose address is given in the token, to cancel and redeem these tokens. If the transaction is accepted, the user is added to the firewall list of active users and his traffic is unblocked.

3.4 Resource Discovery

The problem of resource discovery is simplified by the nature of proximity to wireless networks. As mobile users are interested only in resources of networks within radio range, the user interface

is expanded to include a broadcast request that includes the RSpec of the required resource and the service class (e.g. voice-minute, data-hour, guaranteed vs. shared, etc.), which can be interpreted as a ‘seeking service’ broadcast request.

The RANs and APs in range delegate this request to their associated brokers, which respond with a subset of available resources matching the request, along with their exact specs and rates. Of the several responses to the request, which resource to choose is determined automatically by the user interface using a pre-defined set of preferences (e.g. “lowest rate with acceptable signal strength”, “highest available bandwidth”, etc.). The UI then initiates the ticket request as discussed in Section 2.2.

3.5 Reachability and Roaming

The preceding sections describe how users can *request* resources (call-out), but not how idle users can be reached. To overcome this problem, we propose using an analog of the traditional HLR/VLR registration system that tracks the location and roaming network of idle users. Mobile IP [20] employs a somewhat analogous system for hosts roaming across different subnets with the use of a *home* IP address and a roaming location-dependent *care-of* IP address, to which a home-agent tunnels packets.

When idle, the user interface ‘registers’ with a reachable network of its choice (this step is also done when starting a new lease). This registration itself is a service, that of reachability, and therefore follows the same discovery and reservation steps mentioned above (by broadcasting a ‘seeking registration’-type RSpec). The local visitor location register (GENI-VLR) then advertises the presence of this MS on its network by sending a notification to the home location register on that MSs home network (GENI-HLR).³ Finally, any attempt to reach that MS is preceded by contacting the GENI-HLR in the MSs home network to find its current care-of IP/MSRN/etc. In this sense, the GENI-HLR combines the functions of a DNS and a traditional HLR.

The GENI-VLR must also distinguish between active and idle users, as it has the extra responsibility of informing the MS about an incoming voice call or packet flow in order for the MS to initiate reservation steps to ‘answer’ it. The network drops traffic to MSs that refuse to respond to such notification.

Seamless roaming from one operator to another is an important architectural consideration, but beyond the scope of this paper. As future work, we intend to investigate incorporating related techniques (such as UMTS inter-operator and inter-technology handover [15]) into our architecture.

3.6 Scalability

Apart from token verification, the resource donation and ticketing steps are entirely local. In addition, we anticipate that for both mappings described, the RBs administer relatively small geographic domains. Therefore the broker’s computational load would be limited to the number of active users in a small area. Also, the capacity of the required link to the Internet need not be substantial, as it is used for token validation only and not data traffic.

In essence, this system is a distributed Internet service with no shared state. The only possible source of contention is token validation traffic at home operators. But as this is a logically independent operation, this load could easily be distributed over many servers and geographic areas. For these reasons, it is not expected that the scalability of the system would be negatively impacted as it expands to global scales.

³This notification must involve a recent signature by that MS, otherwise trust is required among the operators, as false positive and false negative notifications could lead to MS unreachability.

3.7 Policy Issues

3.7.1 Token Creation and Billing

Each network operator creates its own tokens. These tokens become the currency of that operator and are pegged to a real-world currency using a set exchange-rate. Tokens are given to users in exchange for paying their monthly bill, or prepaid by request. A user's token balance must be positive at all times, as resources cannot be bought on credit.

When users spend their tokens on other operators' networks, the spent tokens are presented to the issuers for reimbursement. That is, a network provider 'bills' the home-operator of a user that has roamed onto its network for the usage time spent, with the signed tokens as proof. Billing is done in tandem with resource consumption, because tokens must be canceled immediately on usage lest they be copied and reused elsewhere by dishonest users. The token billing interface is implemented as a uniform web interface by all operators in the system, regardless of access method (tokens fully specify their issuer and the address of that interface). Users of delinquent/defaulting service operators are banned.

Optionally, operators may create online exchanges where tokens of different operators are bought and sold, similar to the foreign exchange marketplace. This phenomenon already exists for virtual economies such as SecondLife [16]. Such exchanges could also determine the floating token rates for donations and tickets, based on dynamic supply and demand, and the RSpec parameters of the resource in question (geographic location, service class, etc.).

3.7.2 Virtual Operators

Nothing in the system design specifies that operators need to be brick-and-mortar and operate their own networks. Virtual operators can exist, transparently leveraging other providers' networks, as long as they can profit from selling tokens to their users. Furthermore, virtual operators could provide public web portals for users not subscribed to a home provider (as in the pay-as-you-go model), where anyone could purchase that operator's tokens using a credit card. Several such operators would emerge to compete for this class of sporadic user.

3.7.3 Rate Spread

In the general case, the token rate for the donation of a resource and the rate for that resource's ticket need not be equal, and therefore the brokers need not act as neutral agents. As described above, brokers handle donated resources from the same interest domain (i.e. both the broker and site managers belong to the same operator), therefore the donation rate is meaningless in this context and may be set to zero.

3.7.4 Security and Arbitration

Donations, tickets, and tokens must be cryptographically signed with the private key of their issuer for verification of authenticity. Further, all control communication between operators needs to be authenticated and encrypted in order to minimize or make more difficult man-in-the-middle, replay, and spoofing attacks. The problem of public key distribution is well-known, but orthogonal to this project. However, in this instance, it might be possible to devise a chain-of-trust bootstrapping protocol to make incrementally-trusted key queries in a manner similar to the DNS hierarchy.

In our system, we consider token and ticket signing as a legally-binding guarantee for real-world compensation for token billing and ticket adjudication, as it provides non-repudiation as well. And while we aim to have operator membership in system be open to arbitrary operators world-wide, we must rely in part on a reputa-

tion system to ensure that operators do not over-promise and under-deliver resources. In particular, the barriers to entry to the system (registration, due diligence, credit analysis) must be high enough to discourage small spurious providers, but low enough to ensure wide accessibility. The main reason being that larger dishonest providers are easier to blacklist, avoid, and hold accountable.

4. IMPLEMENTATION

We developed a proof-of-concept Internet-based implementation of the resource allocation components in Java, and tried to follow the API and data structure XML formats described in the GENI specification as closely as possible. The components are run as web services, using the Apache implementation of the XML-RPC standard (ver. 3.0 [3]), which provides a simple server/client remote procedure call interface. The locations of these services (IP addresses and ports) are known at initialization time.⁴

Donations, tickets, and tokens are signed using public-key cryptography with the SHA1-DSA standard with 1024-bit keys (the assumption here is that all components are initialized properly with a global public-key library). These signatures are verified each time a donation is made, a token is accepted, or a ticket is issued.

5. EVALUATION

It is vital that this resource allocation system not adversely affect the user experience. In particular, the steps of the reservation protocol must be unnoticeable to the user and as close to instantaneous as possible.

5.1 Test Description

To measure the impact of the reservation protocol on the call-establishment time, the following tests attempt to measure the computational overhead for a generic resource request (a 'voice-minute' for example). We measure the resource donation, ticket request, and ticket redemption times for the Java prototype described above. These tests were run on a 2.33 MHz dual-core Intel machine with 2GB of RAM, running the 1.5.0_13 JVM release on Mac OS X 10.5.

Described briefly, a site manager donates 1000 generic resources in succession to a resource broker. A thousand different users serially request tickets to those resources, and then redeem them at the their component manager. The average donation, ticket request, and ticket redemption times are measured. Note that these tests measure only the processing overhead of resource reservation, and do not include any link latencies between components that might be incurred if a deployment required their physical separation. We run the tests with and without digital signatures to measure the impact of signing and verification on performance.

5.2 Results

Table 1 shows the results of the aforementioned tests. As expected, the signature signing and verification are the most computationally expensive components of the reservation process. For both donations and ticket requests, two signature verifications are performed (on both sides of the RPC); while redemption only checks the ticket signature, hence the reduced redemption time in comparison. Of course, requests for subsequent/continued resource reservations would be processed ahead of time and overlaid on current resource leases by the UI for uninterrupted service.

For this reason, the expected major contributor to the call-setup time is not these local reservation transactions, but the token billing

⁴Note that for a distributed deployment, these addresses would be explicitly specified in the RSpec of the resource.

Step	w/ signing (ms)	w/o signing (ms)
Resource Donation	29.818	7.852
Ticket Request	24.597	2.235
Ticket Redemption	10.86	1.857

Table 1: Latency overhead of resource donation, ticket request, and ticket redemption.

step, which requires communication with the user's home service provider. However, since the common assumption is that most users are honest, network operators could choose not to stall new callers while their tokens are being checked (and bear the small risk of non-payment), in order to provide a more streamlined service.

In addition, the digital signing and verification is performed in software here. And while we acknowledge that the specialized hardware elements under consideration (small Wi-Fi routers, BTSs, etc.) might be ill-equipped to handle such computationally expensive cryptographic operations, we argue that if this scheme is to be widespread at the consumer level, we expect that introduction of commodity hardware cryptographic acceleration would reduce this overhead much further.

Overall, the evaluation of this resource reservation system is still in the preliminary stages. As future work, we also aim to test the latency, load, and scalability characteristics of the individual components as they are subjected to anonymized real-world traffic traces for both types of network deployments.

6. RELATED WORK

Traditionally, network composition in cellular networks translates to inter-operator roaming. When two PLMNs want their subscribers to be able to use each other's networks, an explicit roaming agreement and a dedicated link are required between the operators. As n , the number of operators globally grows rapidly, this required n^2 such agreements, and this quickly becomes inefficient. Projects such as GPRS Roaming Exchanges (GRXs) [12] try to provide centralized IP backbone networks that attempt to aggregate GPRS data peering points between multiple operators, however these exchanges are currently relatively localized between neighboring networks and fail to achieve global scales. Current research efforts into future 4G tracks (such as 3GPP LTE) focus on all-IP cellular infrastructures, and aim to push this functionality onto the open Internet.

The idea of tokens can be seen as a simplified form of digital money [6, 5], and faces many of the same social and technical challenges (e.g. adoption, privacy, auditing, policing, etc.). However, the use of online money exchanges ([1, 21]) and private currencies ([21, 10]) has slowly been gaining widespread acceptance across the globe. Specifically, some practical uses of digital currency in real-world applications similar to our proposed system have been widely successful; the public transportation systems in Hong Kong (Octopus card), Singapore (EX-Link), and Sweden (M-ticket) are notable examples. It would not be surprising to see the trend of online micro-payments for connectivity and services in general slowly overtaking traditional billing methods in the near future.

7. CONCLUSIONS

By adapting GENI's generalized capability-based resource claim bartering and trade model, we have presented in this paper a flexible resource allocation mechanism that allows for global, transparent network composition between wireless service providers. We then showed as examples how this mechanism can be applied to UMTS

and Wi-Fi networks. In addition, we provide an initial prototype implementation that shows the resource reservation time overhead to be negligible from a usability standpoint, which is encouraging. Future work includes investigation of real-world deployments simulations with realistic traffic workloads, as well as further investigation of the feasibility of incorporating our architecture into deployed systems.

Acknowledgments

We are grateful to the anonymous reviewers for their valuable comments on previous drafts of this paper.

8. REFERENCES

- [1] PayPal. <http://www.paypal.com/>.
- [2] T. Anderson and A. Vahdat. GENI Distributed Services. <http://www.geni.net/GDD/GDD-06-24.pdf>.
- [3] Apache Software Foundation. Apache XML-RPC. <http://ws.apache.org/xmlrpc/>.
- [4] Boingo Wireless. Boingo. <http://boingo.com/>.
- [5] D. Chaum. Security without identification: transaction systems to make big brother obsolete. *Communications of the ACM*, 28(10), 1985.
- [6] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In *CRYPTO '88: Proceedings on Advances in cryptology*, 1990.
- [7] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. PlanetLab: An Overlay Testbed for Broad-coverage Services. *SIGCOMM Computer Communication Review*, 2003.
- [8] Cloud Networks, The. The Cloud. <http://www.thecloud.net/>.
- [9] D. Zhou, et al. Validated Composition and Compensation Architecture (D26-G.2). <http://www.ambient-networks.org/deliverables.html>.
- [10] e-gold Ltd. e-gold: Internet Payments. <http://www.e-gold.com/>.
- [11] Y. Fu, J. Chase, B. Chun, S. Schwab, and A. Vahdat. SHARP: An Architecture for Secure Resource Peering. In *ACM symposium on Operating systems principles*, 2003.
- [12] GSM Association. GPRS Roaming Exchange (GRX). http://www.gsmworld.com/technology/gprs/grx_prov.shtml.
- [13] J. Eberspacher, H. Vogel, C. Bettstetter. *GSM Switching, Services and Protocols*. John Wiley & Sons, Inc., New York, NY, USA, 2001.
- [14] H. Kaaranen, S. Naghian, L. Laitinen, A. Ahtainen, and V. Niemi. *UMTS Networks: Architecture, Mobility and Services*. 2001.
- [15] H. Kwon, K.-R. Jung, A. Park, and J.-C. Ryou. Consideration of UMTS-WLAN Seamless Handover. In *IEEE International Symposium on Multimedia*, 2005.
- [16] Linden Lab. LindeX. <http://secondlife.com/currency/>.
- [17] M. Johnsson, et al. Final System Description (D18-A.4). <http://www.ambient-networks.org/deliverables.html>.
- [18] National Science Foundation. GENI: Global Environment for Network Innovations. <http://www.geni.net/>.
- [19] D. Oppenheimer, J. Albrecht, D. Patterson, and A. Vahdat. Design and Implementation Tradeoffs for Wide-area Resource Discovery. In *High Performance Distributed Computing*, 2005.
- [20] C. Perkins. IP Mobility Support for IPv4. RFC 3344, Internet Engineering Task Force, 2002.
- [21] WebMoney. WebMoney: On-line Payment Solution. <http://www.wmtransfer.com/eng/>.