

Architecture of a Mobile Internet

Dick Schefström

Center for Distance-spanning Technology, 97187 Luleå, Sweden

dick@cdt.luth.se

Abstract

In this paper, we outline the next major evolution of the Internet architecture, where support for mobility is not the exception but the typical case. We show which parts of the Internet architecture needs to be extended, how they interact, and why they are needed.

We point out the suitability of especially license free radio access equipment not only technically, but as a well harmonizing piece in the non-centralized and distributed way that Internet grows.

We show that a wired Internet supporting mobility drastically changes the market situation in the mobile and wireless sector, possibly to a level where special mobile operators are not needed. We show that this leads to a situation where wireless and mobile access is just as inexpensive as its wired counterpart

1. Mobile Internet: What is it?

The answer is not self-evident, and will differ significantly whether you ask someone from the Internet or from the mobile phone community. This is understandable, since the differences in view between those communities is wide, and is not only a question of protocols and technologies, (like the classic argument between circuit and packet orientation), but also of the basic principles of systems construction: distributed and non-hierarchical or centralized and hierarchical, equal access open services or closed islands, (sometimes called walled gardens). The question even reaches on to regulation and the role of government, and is thereby also a political one. The difference in culture between original Internet community and traditional operator approach is real and important, and can be characterized as in the following table:

Internet wireless mobile?	- but Mobile telephony - but and with services like on the Internet?
----------------------------------	---

Same services and applications as on the Internet	New separate services special to the mobile customer.
---	---

Computer companies and universities	Telecom operators
-------------------------------------	-------------------

Openness and freedom emphasized	Control, closedness and restrictions
---------------------------------	--------------------------------------

Distributed "democratic".	and Centralized and hierarchical.
---------------------------	-----------------------------------

Starts from computers as terminals	Starts from telephones as terminals
------------------------------------	-------------------------------------

Simple payment model - usually flat rate.	Complicated payment model - time, volume, person, place, time of day, etc, etc.
---	---

Internet Culture Telephony Culture

Table 1: Internet versus Telecom culture

The importance of cultural aspects of Internet, and its various roles, has been well described in for example (Castells 2001, p 36-63).

When now combining those concepts, of Mobility and Internet, we therefore stand before a decisive crossroad of great significance: Do we want an Internet extended to support mobility, or a cellular/mobile world extended with Internet-like contents? The question is of much greater significance than has been observed this far, and some of us remember a similar situation in the early 90's, but then the issue was fixed networks.

To some people it is almost self-evident that mobile Internet should be taken literally, that is making the same services and applications available "mobile" as are already done using fixed access. There is, however, a strong alternative view, much driven from the mobile telephony world, according to which "mobile Internet" is a new concept, largely independent of what we now call Internet. The reasoning here is largely driven by an assumption that the limited capacity of terminals, the limited bandwidth and different pricing of the mobile network, makes it necessary to develop new services tailored to this context. Those may, of course, be built on the corresponding services of the fixed Internet but will, anyway, be different. Suspicious thinkers have even raised the view that the arguments for separate services may have something to do with a desire to create walled gardens of services with closer ties to the operators. This is, however, just a rumour.

On the other hand, creating new services, or at least special variants of existing ones, for the mobile case, is of

course an extra burden. If it could be avoided there would be great benefits, simplifications and savings. A single set of services, and no variants, for all of Internet, be it fixed or mobile. The obvious alternative idea is therefore to *not* accept the lower performance and higher pricing as a law of nature for mobile systems, but try to change this. Actually, it can be done. Now, to understand the issue better, let us take a look at some background.

2. The Expectations

We believe most IT professionals would agree on summarizing the 90's in two points:

- The extreme growth of the Internet.
- The extreme growth of mobile telephony.

At the start of the decade, Internet was a strange corner of the communications world, of concern only to obscure academics and not considered of wider interest. This was at least the unanimous vote of telecom operators. At the end of the same decade, Internet was the self evident platform for all sorts of beyond-voice communication applications, and an essential component of modern society. Typically, well over 50 percent of households in US and western Europe has Internet access, and a quickly growing number of those have broadband connections.

The characteristics of usage are adapted to "flat-rate", that is a fixed cost independent of time and volume transferred. This is very important, since it stimulates a natural relaxed usage pattern and the application of sound, images and video. That is, much of what makes life in cyberspace fun and valuable. This approach also has the positive effect of making charging and customer management subsystems simple and inexpensive.

At the start of the decade, mobile phones were considered exotic high profile gadgets for the executive, if present at all, and were mostly used to impress ladies at expensive restaurants. At the end of the same decade, 95 percent of teenagers in Scandinavia have a mobile phone, and not having one may be seriously hurting your social life and status.

The charging culture is that of traditional telephony, that is depending on distance and time, but typically orders of magnitudes more expensive than wired calls. The author suspects that the great growth of SMS usage, especially among teenagers, is to a large extent driven by the pricing factor. Teenagers often have their determined phone budget and prefer to send a fixed price SMS rather than risking an expensive phone call.

Anyway, all this fed a great growth in IT industry, and has sure been good for wealth, society in general and quality of life for the individual. Then, at the start of a new millennium, the telecom industry naturally expected a next boom, now to the power of two, by multiplying

Internet with Mobile Telephony. For some reason still to be understood, much of the expectations got attached to the UMTS, often called 3G, standard, and a wild bid for radio spectrum licence rights started, ending with a situation where European operators have paid well over 140 Billion Euro just for the right of being alone in a piece of radio spectrum. Still, investments in hundreds of thousands of radio base stations and upgraded special transport networks have to be financed because of the funny tradition of building a separate wired network to connect the radio base stations. All this is even more strange as UMTS was not originally designed for carrying Internet. That is as afterthought caused by the wired Internet's overwhelming success.

It is therefore not surprising that the pricing suggestions - shown among other things by the so called GPRS, or generation 2.5 - are causing worries for anyone concerned about natural Internet usage. While charging by time unit and distance seems less likely, most operator charge by downloaded volume, and a typical cost seems to be in the order of 5 Euro per megabyte. Now think about what this means. It means that downloading an mp3 tune would cost as much as buying the whole CD in a store, just for the bit transport. It means that all large data structures, like pictures, audio, and especially video, would be severely punished and therefore probably be removed. It would mean that you have to think very carefully before surfing in on a site or following a link. It means that most of what makes the user interface fun and rewarding will be removed, throwing us back into bit-preserving but boring and unfriendly user interfaces. It means that a decade of improvements and creation of a relaxed and natural style of investigating resources on the Internet will be more or less killed.

Now view this, again, in contrast to the wired Internet access, where a typical broadband home customer has 10 mbit/s bi-directional for a flat rate cost in the order of 40-50 Euro per month. In the worst cases, it is ADSL like asymmetric, but anyway. Besides the pricing issue, the planned UMTS based mobile Internet is not even "broadband", since it provides at best a few hundred kbit/s to a single user.

Will people accept such a huge difference in price/performance in wired and unwired access to the same thing, the Internet? Is it reasonable that an IP package, jumping hop after hop through the network, then, at its final hop, gets a 1000 times higher price-tag because this last hop was wireless?

No, we do not believe this is reasonable. We believe that a Mobile Internet should, and could, be built so that mobile access, wired or unwired, has close to similar price/performance characteristics as the fixed and wired network.

We believe that this is best done by building a single IP network to which radio base stations and computers are

connected on equal footing. This single IP broadband network that is anyway built for dozens of other reasons and whose deployment has started well in some countries but has wide support at all levels. At the EU Summit Meeting in Barcelona 13-16 March, there were a few main decision reported, among them the “Widespread availability of broadband in EU by 2005”. This network is an IP network. It is part of the Internet. It must be built to support mobility.

3. The Radio Interface

Until now, mobile and cellular systems have been designed largely driven by its radio interface. The first, second, and the third generation mobile telephony system all started there. The reason for this might be that the radio resource indeed is a limiting factor, and that the very concept of radio seems to be so intuitively connected to that of mobility. The focusing on radio may also have its roots in that this medium has been under heavy regulatory control and therefore attracts a certain administrative and legal attention.

However, it does not need to be that way, and a holistic architectural analysis can just as well show that the radio interface is in fact *not* the appropriate starting point. Actually, the physical mobility of terminals is only at a very first level a question of radio, when being confined to the coverage area of a single base station. As we move outside this, it is the wired network part that reroutes traffic to another, better suited, connection point. When picking up your device after a flight, turning it on, it is the wired system behind the radio interface that checks your identity, decides to let you in, and makes it known to potential callers that you are now at a completely different place. All mobile/cellular networks are therefore mostly wired and fixed networks from an architectural point of view.

It should therefore be natural to see that the complexity of mobility, like identification, “roaming”, authentication, handover, etc, is handled by the “new mobile Internet”, while radio base stations are to be treated more like the final hop providers that they are. More like Ethernet interfaces that happen to be wireless. Simple devices with a localized specific role. Radio is a necessary, but not the essential and determining, aspect of a mobile system.

Radio base stations have however not traditionally had a reputation for being inexpensive. On the contrary, they have been very expensive, complex, and not something you buy in the corner electronics shop. A UMTS basestation, for example, cost at least 50.000 Euro. Then, you get 2 mbit/s to share in coverage area. And you should not buy one anyway, since you are not allowed to use it unless you have a license for its use of the radio

spectrum. And that is troublesome and very expensive to get.

On the other hand, part of the radio spectrum was opened, first in the 2.4 GHz and now on the 5 GHz range, for use by anyone without license. And suddenly, like some kind of magic, radio base stations were neither expensive nor complex anymore. Actually, they became very inexpensive consumer-electronics like pieces, costing about 250 Euro, being easy to install and delivering 11 mbit/s capacity, (IEEE 802.11b). As this is written the IEEE802.11a and c are being deployed providing 20 and 55 mbit/s raw capacity, with European HiperLAN as a challenger in the 55 mbit/s class. In any case, RadioLAN basestations are now at a price-level where it is considered consumer electronics. I have it at home, and soon it will be available in every gas-station, at every hotel, at every supermarket, at every airport, in every living-room. It will be built in as standard equipment in stationary computers just like DVD-players are today.

4. The New Approach

The goal of the RadioSphere project, (<http://www.radiosphere.nu>), run by CDT and its associated partners, is to build a “Mobile Internet”, but with a price/performance ratio that revolutionizes usage, and thereby allowing for applications otherwise seldom thought of in this context. The research plan is derived directly from a consistent idea and represents a well thought out puzzle rather than spurious research items. This holistic approach, the architects view, is an important quality.

The basic idea is radically different and innovative, and while being basically simple it has very far reaching consequences: *build the Mobile Internet by extending the wired broadband Internet by inexpensive and largely consumer owned radio basestations*, when appropriate combined with existing wireless wide area networks. This leverages on the ongoing broadband deployment, (where some countries have reached further, among them parts of Scandinavia), and extends that directly into wireless use and mobility. It increases the value of the current broadband investment and makes wireless bandwidth a people’s commodity. It can deliver a price/performance ratio which makes it possible to run the same applications in the mobile and fixed contexts. This represents huge savings and simplifications.

Philosophically, we believe that Mobile Internet should be built in the “Internet Way”: Build a single wired network, where radio base stations and regular computers are on equal footing. This provides a number of advantages, recognizable from the wired Internet culture:

- Organic growth. Anyone can add to the complete network.
- Little central control, activating grassroots in a distributed system.
- Use of standard off-the shelf components.

The core of this mobile internet is then in the wired infrastructure, (as it should be), which must support true Mobile IP, extremely fast authentication and encryption, mobile QoS, integration with DVB, portable/mobile UI's, and a number of other features as identified below.

We then need a way to automatically manage this system and make it attractive to share bandwidth. We need society-wide, neutral and hustle-free ways of authenticating people. We need easy-to-manage ways of handling access-rights to basestations. We might need ways to automatically, by distributed negotiation, adjust the emitted power and spectrum usage of basestations. For many applications we need to deploy a mobile IP allowing us to address a terminal under a static ID while its anchor point in the network changes. In addressing the traditional realms of UMTS and GSM, we need to deploy IP-telephony with the special problems it means to ensure QoS in a mobile IP and radio access situation. We need to exploit the client-to-client and efficient multicasting potential of the new RadioLAN technology. It leads naturally the way to ad-hoc networking.

Then, we can expect every gas station to have a radio basestation, and all those who bothered to get a bonus card get free access to the network. Shopping centers will offer free parking and free mobile Internet. Airlines offer the gold-card owners free access. Well, you will probably have it as soon as you get the paper-slip for having applied for the bonus programme.

At CDT we run the Arena project, (<http://www.cdt.luth.se/projects/summary/arena.html>), where we extend the ice-hockey match experience by putting a camera in the referees helmet, pulse-sensors on hockey players, info from other matches, and much more, all distributed using inexpensive radio basestations to PDA-owning people in the audience. Audience get many replays and flashy pictures to choose among. You get access as you buy the ticket. Automatic expiry at midnight the same day. Obviously, you can buy an electronic-only ticket as well, sitting in a restaurant downtown.

And so we take the final step and establish the People's Movement for Free Mobile Internet and let each other in to our basestations. I have a base station at home already, of course. As you come to my home, I let you in. Probably.

Installation and operation of basestations is then distributed, and stops being a bottleneck. No more need for serious men in white coats nursing equipment, since any computer interested teenager will install the basestation just as if it was a new graphics card or a joystick. In general, size and operation of this Mobile Internet can grow in the same efficient and robust way as the young wired Internet. The principle is the same. It will work again.

The far reaching potential consequences of this should thrill anyone: Inexpensive operator free mobile Internet as a side-effect of the flat rate wired internet to homes and gas-stations. It may drastically change business. It is good, however, very good, for the general public.

As is further developed below, we have identified about two dozen research items that needs to be addressed, and each of those will be explicitly addressed, while possibly with different emphasis. While the basic building blocks are there, a challenging technology and integration adventure remains, but which can now be fast and efficient given this focused but still visionary goal. There are several problems to solve and many research issues to address and the following is a partial list of those:

1. *Authentication*. Needed for many reasons, ranging from network access and implementation of mobility to securing of business transactions. PKI needed.
2. *Encryption and Privacy*. Can be arbitrarily sophisticated. Care to be tracked?
3. *Customer and group membership management*. Light-weight and partner neutral.
4. *Electronic economic transactions*. In case you have to pay for access it should be electronic.
5. *The Business model*. New ways of securing economies of mobile access must be advised. But, the goal is less money for bit transport, anyway.
6. *User Interface technology for mobile devices*. You must have it with you.
7. *Mobility of IP-addresses*. MobileIP and associated questions. Triangle routing is not an option. Multicasting must be efficiently supported.
8. *Adaptive and self-regulating cell-coverage*. Necessary for consumer electronics scenario.
9. *Mobile Subnetworks*. As many clients move at the same time, important and necessary optimizations can be done. One Foreign Agent for the whole ship?
10. *Advanced Applications*. The world beyond the web and beyond the phone. Handheld electronic meetings and interactive radio and televisions.
11. *Location-aware applications*. Among few genuinely mobility oriented services.
12. *Integration with broadcast technologies and DVB-T in particular*. There is a fruitful marriage possible here, from transport technology to contents.
13. *Application Architectures for Resource Limited Small Computers*. Such as ASP-style, where large storage and applications must be hosted by the "network" rather than the client.

14. *Service discovery*. Where is a Printer? A CD-Burner?
15. *Context mobility*. Keeping all state, security associations, open files, etc, while moving.
16. *Ad-hoc networking*, peer-to-peer and multicast optimization.
17. *Quality of Service in Mobile and IP networks*. Maintaining reservations while path changes.
18. *Handover management*. Fast and smart handover. IP level probable bottleneck.
19. *Mobile IP telephony*. In addition to QoS, header compression is an issue.
20. *Smart antennas*. Mobile networks, such as cars, trains, busses, boats, can carry them.
21. *Short range gigabit/s radio*. Some call it "Infostations".
22. *Ipv6*. Is it needed before Mobile Internet can be real? Must Mobile IPv6 be deployed at same time?
23. *Interworking with UMTS*. Attaching any kind of radio basestation to the RadioSphere network is a core goal. This includes, of course, UMTS.
24. *Legal Issues*. Can lawyers sabotage bandwidth sharing? Is FCC/Government regulation really needed long term?

Each of those is a fruitful area as such, but with RadioSphere they get a meaningful common structure leading to a fascinating goal.

5. Research Items

Each of the 24 points above has its own in depth discussion and plan within the RadioSphere project, but due to space limitations in this paper we can only select of few of them and then provide a brief overview.

5.1. Authentication

An obvious prerequisite for mobility is the change of network connection point, which in the general case then also implies change of access point owner. That owner most likely wants to know which people utilizes his equipment, or at least be sure that the user is member of a group to which this right is granted.

This is an essential function in our vision, since it is the mechanism that lubricates the free exchange of bandwidth, which in turn is the key to the extraordinary efficiency of the approach. This is the basic mechanism that quickly ensures that you have the "Premium Card" of the gas-station chain or the shopping-mall; that you have paid the membership to the "peoples network"; that you have checked in with the airline that now offers you wireless Internet access while waiting.

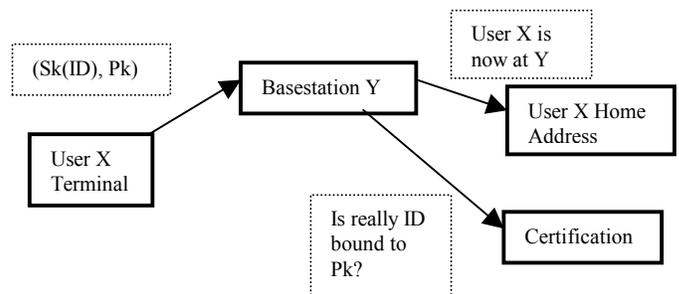
Authentication also turns out to be a prerequisite for addressing other problems like packet forwarding and route optimization in a mobile network. The incoming IP-telephony calls to your home IP-address should be rerouted to your current mobile connection point, and be so along the nearest path between the caller and your current location. To allow this, we must be sure that it is actually *You* who asks for all your traffic to be rerouted to this new place. Otherwise, the network cannot trust the rerouting request.

There is a well-known suite of proposed standard protocols addressing this problem proposed by the Internet Engineering Task Force – IETF – under the name "IPSec",

<http://ietf.org/html.charters/ipsec-charter.html>.

Anyone should however feel a bit worried about the amount of documents and complexity that has been accumulated in this area, and possibly there is reason for a fresh simplifying view on the subject.

In principle, the authentication can be done securely by assuming the existence of public and secret key pairs, (Pk, Sk), for each individual according the RSA asymmetric encryption principles.



As User X gets in contact with Basestation Y he has to prove that he *is* X. He then uses his secret key, Sk, to encrypt a piece of data, (say his ID – such as social security number), and sends that to the basestation together with his public key. The basestation decrypts the ID using the public key, Pk, getting the ID. He now knows that this must be user X, since no one else could have encrypted ID so that Pk decrypted it. However, he must also be sure that ID is actually bound to Pk – that the public key offered actually is in fact corresponding to User X. To do this he must go to a "trusted third party", usually called a certification instance, to validate this information. This machinery has been called a Public Key Infrastructure, or "PKI" for short, and is needed for all sorts of reasons.

As we now know that it is actually User X that has shown up at Basestation Y, we can store this information as valid at User X's permanent home address site. The simple idea is then that anyone who wants to reach User

X goes to her well known home site, and there gets the information that X is now reachable under Basestation Y, so that packets would be directly addressed to this place.

Anyone knowledgeable in the security area knows that many issues and problems remain, to some extent depending on how paranoid we feel. Anyway, the point is that reasonably secure identifications might be possible with simplified and efficient protocols. A fresh and critical view should be taken on the IPSec suite of solutions.

So, we have at least two different authentications: one for getting access to the basestation, and one for updating the permanent home address with information about current location. The former may or may not require updates to basestation software. In any case, we already see the need for an open software platform with standard OS in this basestation, or just as likely, the "basestation" being a card with associated drivers and software hosted by any computer.

5.2. Customer and group membership management

It must be *extremely* easy to authenticate and get access to basestations as we move around. It should, in fact, involve zero manual intervention whenever theoretically possible. The current early and somewhat primitive market, (such as pushed in Sweden by Telia HomeRun), requires that you manually log in and submit a password which you have gotten from this particular operator. This has to be done at each connection event. If you move to another point, and maybe another operator, you have to repeat the process.

This is, of course, not acceptable nor realistic in a wider perspective. You would be very irritated if, when using your mobile phone, you had to log in and submit password to the network as you want to use the phone.

Rather, your terminal, (computer), should authenticate you, as described above, and your identity be used to check a possibly very wide set of reasons why you should be let in, ranging from being a trusted customer to the owner of the access point to various memberships in clubs, societies, bonus programs or the famous "Peoples Network Movement". One could compare this to the quick electronic execution of the browsing through your pile of club-and-bonus cards which many of us have too many of in our wallets.

The management of such a (distributed) database is however a real challenge given the scalability issues of world wide usage and potentially a million membership options. Its design must for obvious reasons not rely on central servers, but be truly distributed to allow for fault tolerance, robustness, efficiency and organic growth.

An even more important aspect of this issue is that of "Always Best Connected", ABC, which is a one of those buzzwords that in the last year has reached almost shamanistic status in some contexts. Yes, we all know it: It has mostly been used as a way to make it like different radio-access technologies could complement each other rather than compete. And yes, it has mostly been a question of handover between GPTR/UMTS and WLAN.

While we have always felt that there is a strong element of self-deception and tranquilizer in this approach there are deeper, more relevant and much more interesting aspects of the concept of Always Best Connected. Unexpected to some, but nevertheless... When I now want to make a phone call, a mobile or GPRS one, there are several operators in the air and I would obviously like to choose the one who currently gives me the best price. No, I don't want to be confined to a single one, but I would like them to compete for this customer and bid under each other! Comviq, Cingular, Telia, Europolitan/Vodafone or whatever, depending on who happen to give me the best price/performance at this particular time.

That, would be Always Best Connected. That, would be real market economy. Designing that system, would be really important research, serving the only viable long term objective: society at large.

However, I now guess that a certain angst appear in the minds of many current actors when such a scenario is proposed. This, the most self-evident and positive interpretation of "ABC" is probably forbidden for those who want to lock in customers and preferably like to have a secure monopoly. Nice and cozy for those, but not in the general good for society at large. But many ingratiating researchers follow the orders and instead use their time and ingenuity for something "less dangerous".

5.3. Mobility of IP Addresses

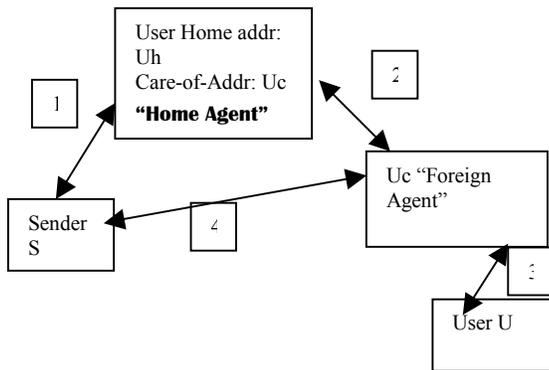
Being mobile means connecting to the Internet at different places from time to time. This is a special challenge given the unfortunately dual role of the IP-address: it is *both* an identification and a location. The IP-address is used as the basis for routing in the reasonably stable, or slowly changing, fixed network, and the efficiency of routers and routing tables is depending on that addresses with the same prefix, (up to some number of bits), are located in the same part of the network.

Since "any problem of computer science can be solved by inserting yet another level of indirection", one can use two addresses: a *home-address*, which is fixed and works as your identity, and a *care-of address*, which varies depending on where you are connected in the network.

Packets sent to the home-address are then rerouted to the care-of address.

This is the basic principle of Mobile IP, (Perkins 1998, Solomon 1998), which spends most of its contents on solving the various complexities and inefficiencies that may appear when using this quite basic idea.

We call the users home address, his static identity, “Uh”, and his Care-of Address “Uc”:



Mobile IP introduces the concept of a “Home Agent”, which in its simplest form catches packets destined to home-address Uh, wraps them in another message and forwards them to the care-of address Uc, where the “Foreign Agent” unwraps it and delivers the original packet to U, (3).

An extensive report could be written about the effects of this machinery, various optimizations, and the effects it has on the TCP/IP suite as well as potentially the application level. It turns out rather complex after a while.

Let us just conclude that we believe there is plenty of room for new solutions and improvements – there must be simpler approaches – and it can be observed that Mobile IP is not in large scale operational use anywhere. Some starting points are the following:

The simple forwarding solution, called “triangular routing” following paths 1 and 2 in the figure above, is unrealistic in a general scenario. Making an IP-call to your Luleå based colleague as you both are in Paris makes packets travel over half of Europe while you’re in the same city.

All communication should go directly between correspondents, following path 4 above. Sender should only pick up care-of address Uc at User Home, and then send actual messages directly to Uc.

The Session Initiation Protocol, SIP, is designed to do many of the things we need for Mobile IP, such as forwarding and registering actual position. In addition it allows for media selection and negotiation. We believe a Mobile IP solution must architecturally be integrated with SIP, (RFC2543, [ftp://ftp.isi.edu/in-notes/rfc2543.txt](http://ftp.isi.edu/in-notes/rfc2543.txt)).

An important reason Mobile IP has not taken off is that it relies on that User repeatedly reports home her new

care-of address, (along 3 and 2 in figure above). This cannot be done safely unless authentication is in place.

The User U must itself be given an IP address at the care-of address Uc, which may be for a example a radio basestation. In most scenarios this implies a mapping between Uc and U, since U may not make sense outside the local subnet. This issue is also related to the question of availability of IP-addresses, (where Ipv6 provides many of them), and the current Ipv4 solution to this problem called Network Address Translation, NAT, (RFC3022, [ftp://ftp.isi.edu/in-notes/rfc3022.txt](http://ftp.isi.edu/in-notes/rfc3022.txt)). We believe integration of NAT in the Mobile IP architecture, versus the use of Ipv6, is important to study.

While viewing the over all architectural picture it seems natural to note that users should not be known by IP-numbers but by “mail-addresses” which in turn are mapped to IP-numbers without user intervention. This is done by the Domain Name System, DNS, (RFC1024, [ftp://ftp.isi.edu/in-notes/std/std13.txt](http://ftp.isi.edu/in-notes/std/std13.txt)). The relationship between this, SIP and the Mobile IP home-address/care-of address must be studied and optimized.

We could continue this list by discussing hand-over, roaming, session mobility, mobile IP and multicasting, Quality of Service and service differentiation, handling of packets-in-transfer, transfer to IPv6, (which helps!), etc, etc, but leaves the reader with the hopefully strong impression that important inventions are still to be made. Providing smooth handover of IP streams with small cells and fast movement will be a really challenging problem to tackle for some time to come, and could easily deserve a whole point/subarea in itself.

The problem of IP and mobility is also addressed in the context of Ipv6, and a special addition called “Mobile Ipv6” is described in the Internet draft “Mobility Support in IPv6”, (<http://search.ietf.org/internet-drafts/draft-ietf-mobileip-ipv6-13.txt>).

It advices a way in which mobile nodes could move between different connection points while still maintaining a single address. That is, mobility transparent to transport level protocols and applications. This, obviously, what we want.

On the other hand, we should never forget that in many situations this complex machinery is not needed at all! Given that we accept the following restrictions:

1. No change of connection point without restarting session.
2. Very limited ability to accept incoming “calls”.

Then, we do not any Mobile IP at all. Given that we accept the above two points everything works immediately.

And finally, maybe most important, we must remember that the mobility of IP addresses, that is what

is addressed by for example Mobile IPv6, concerns only a limited part of what the concept of "mobility" is about. The ability to "roam", to have reasons and right to attach to the Internet wherever you are, is the most important aspect of mobility.

5.4. Adaptive and self-regulating cell-coverage

Current cellular networks are built under the assumption that there is a central operator planning cell coverage, placement of stations, usage of channels and avoidance of interference. The RadioSphere cannot assume that, since it builds on a non-centralized model where the networks grow organically *without* central planning and control. RadioSphere is a truly distributed system, and is therefore more efficient and robust since it avoids the bottlenecks and failure-sensitivity of central control.

So, RadioSphere basestations must adapt themselves. There are at least two parameters to play with: emission power and channel usage. Standard cellular network procedure is to decrease cell size as we need to increase capacity by increasing the number of basestation in the same area. To avoid interference radio emission power should be adjusted downwards. In the same way we should choose "channels" so that we maximize the distance to neighbouring stations. For example IEEE802.11b has 13 different channels, where about three can be used at the same time without interference. In a mosaic of stations over a given area there is an optimal allocation of channels to stations which minimizes interference and maximizes capacity.

It is the mission of the Distributed Adaption Protocol to find this allocation and suitable power level. We foresee that this is done by stations listening to each other, exchanging information, and negotiating a common solution.

HIPERLAN/2, (<http://www.etsi.org/technicalactiv/hiperlan2.htm>), provides facilities in this direction, in addition to very high speed transport at 54 mbit/s.

It should also be possible to build the network so that radio coverage information is automatically collected, just as a side effect of users moving around. Assuming that at least some clients knows two things:

- its position, by basestation intersection or through GPS,
- its experienced signal quality at that position,

Then, we can let each of them send the data pair, (Position, Signal Strength) with certain intervals, and automatically build up a map of the radio coverage situation. This can then be used in several ways:

1. Just providing the user with that info, probably painted on a suitable map, as to know what to expect. Quality of Service is to some extent a question of expectations. If the user knows about variations, he will feel like the overall quality is better than if he didn't know.
2. Directly hint the user of "moving 20 meters north" to improve connection.
3. Increase precision of signal-strength based positioning by comparing actual signal strength with previously measured.

The approach taken here is also principally interesting, since it meets the goal of the "self-organizing system". No central authority is needed, but information is gathered automatically just as an invisible side effect of the little user-ants working around. Statistics say that quality only improves as times go by.

5.5. Ad-hoc Networking

The concept of "Ad-hoc networking" has been an increasingly popular research subject for a few years, and numerous are the grad students who have been sent out to penetrate this area by trend-sensitive supervisors. However, as usual when a concept attracts attention, its meaning starts to evolve and bend into all sorts of directions, its context is lost, and unfortunately often things turn into a search for complexity rather than solutions. Has this happened here? Judge yourself:

The attraction we see in Ad-hoc networking is directly derivable from its most fundamental definition:

Ad hoc networking definition 1: *Communication without fixed infrastructure.*

OK, we may use an "ether" to transfer radiowaves, but that's all. The most obvious and rewarding example of this is of course the nice optimisation gained by sending messages directly between terminals, if they are within mutual radio coverage, rather than taking the way over a (wire-connected) basestation and then out again.

However, I would expect operators not to like this, because this makes it very hard for them to charge for the message, since you do not utilize their equipment. But from all other perspectives this is good, since it saves resources and lets people spend their money on even more creative things. All the millions of daily SMS's - and mobile multimedia mails in the future - which are sent between youngsters in schoolyards then become free! Further, think about the easily available mass distribution that naturally comes with this approach. Easily addressing all girls in the class in a single message, implemented in a very efficient way.

An addressing technology and protocol which makes this optimization whenever possible would therefore be a real step forward for mankind. Simple as it may sound, it is not done yet. Now some of you may feel that this sounds far to simple, and may look for more complexity, and then come up with a definition that more explicitly points to trouble:

Ad hoc networking definition 2: Routing through a population of non-fixed, mobile, nodes, usually with the purpose of finding a path to fixed infrastructure..

This certainly creates some more trouble in the general case, since part of the traditional routing table is now a) distributed into every node and b) changing often. While there are already several proposed protocols for managing this, I hesitate since we now seem to somewhat lose contact with purpose and look for complexity for its own sake. The objections to the most complex interpretation of Ad Hoc networking is that a) it seems applicable only in rare and unusual situations and b) it relies on really extremely cooperative and altruistic clients.

The context for application of the multi-hop ad-hoc network approach is cases where a client is not within radio reach with a wired basestation, but with other clients, which in turn may have contact with some client who finally ends up at some wired place. Think about this again. Think about it in the context of the world as you know it or expect it to be. You probably end up finding this care very rare, because if you have contact with some client who have contact with a basestation, it is overwhelmingly likely that you already had that yourself. At least in a civil society where fixed infrastructure is continuously rolled out and is growing every day. Actually, the scenario becomes less and less likely as time goes by.

Yes, there are certainly situations, special situations, where general ad-hoc networking is needed. Such as in case much of the infrastructure we usually rely on become destroyed, or in cases where we temporarily need to set up networks in seldom visited places. I get to think about center of Jenin or the Tora Bora complex, and seriously such scenarios were the original reason why US DoD started to fund ad-hoc networking projects. And still do.

However, if your objective is to “create a mobile and wireless Internet with a price/performance close to the fixed network”, this is probably not a prioritised task. But don't be sad, there is anyway enough complexity further down the road!

The second reason for hesitation is the reliance on client terminals as intermediate nodes in the ad hoc network. In the special case of emergency situations as exemplified above this is easy to mandate, but in a free market of independent terminal owners it is much less

likely. This especially as the mobile terminal is likely to be a resource limited piece of equipment with little free resources to give away.

While we certainly *could* imagine such an altruistic and cooperative population, it reminds us of the much less intruding RadioSphere idea of that people would let others use, borrow or buy their fixed basestation capacity to get out on the Internet. The “People Movement” idea of RadioSphere then seems much more likely to be acceptable than the general ad-hoc idea since it is using a resource which is likely to be less limited. Anyone who have had a problem with the RadioSphere approach must then feel *very* uncomfortable with multi-hop ad-hoc networking.

6. Related work

While the general mobility debate today is dominated by 3:rd generation UMTS efforts, there are still many projects developing, or exploiting aspects of, Radio LAN technology. The currently most popular standard, IEEE 802.11b, was standardized in 1999, (<http://grouper.ieee.org/groups/802/11/index.html>). See also (O'Hara & Petrick 2000).

Work within Mobile IP, (Perkins 1998, Solomon 1998), is active and important, and is often assuming Radio LAN as the context. It does not rely on it in any way, however. Related to this, much authentication and security is done, for example IETF Diameter, (<http://search.ietf.org/internet-drafts/draft-calhoun-diameter-18.txt>). There are many drafts available in this area and can they be found starting at <http://www.ietf.org/ID.html>.

UMTS is the much talked about 3:rd generation mobile telephony, or really Mobile Internet, (<http://www.umts-forum.org/>). In many countries, licences for UMTS have been sold for surprisingly high fees. Many now start to doubt the value of those licenses. The much cited “i-mode”, (<http://www.nttdocomo.com/i/index.html>), in Japan, is a kind of mobile Internet, but not at all the full uncompromised case we are looking for. It has been much more successful than the European WAP initiative, (<http://www.wapforum.com/>), because of much better “telephones”, (Java, larger colour screens, etc), and an ability to stimulate i-mode contents production. The business model of letting contents provider keep most of the revenue may have something to do with this.

Many companies offer RadioLAN access at hot spots like airports, (for hot prices), like Telia HomeRun in Sweden, (<http://www.homerun.telia.com/eng/start.asp>). The US coffee shop chain Starbucks offer RadioLAN access at their 3000 shops, and many airlines offer it in cabin during flights, (such as SAS, <http://www.sas.se>).

Many smaller companies start to offer limited coverage Mobile Internet using this technology, or provide support for its installation. A couple of Swedish examples are PowerNet AB, (<http://www.powernet.se/>), and Åkerströms Björbo AB, (<http://www.akerstroms.se/>). The Swedish companies Columbitech, (<http://www.columbitech.com/>), and “A Brand New World”, (<http://www.abnw.com>), are further examples of new brave startups taking this direction.

We could just as well add Service Factory, (<http://www.servicefactory.com>), supplying for example authentication and user management software, and Excilan, offering roaming even internationally for RadioLAN's, (<http://www.excilan.com/>). Recent startups such a Joltage, (<http://www.joltage.com>), do even offer products to support small franchising business based on the idea of selling off part of your bandwidth to visiting people using RadioLAN technology.

There is however still far too little work on the path shown in this paper, maybe because of its “too” far-reaching effects. Work here seems this far largely to be guerrilla-like loosely connected communities like Elektrosmog, (www.elektrosmog.nu), and Rooftops, (ref). Many more exist, and the “links” page of Elektrosmog, (<http://www.elektrosmog.nu/links>), contains some of those. The MIT professor David Reed argues that the technical reasons for spectrum regulations are obsolete and maintains a website, (<http://www.reed.com/OpenSpectrum/>), with presentations and material supporting and discussing this view. It is interesting since he is in direct close interaction with FCC as a one of its expert advisors.

The Stockholm Open initiative, (<http://www.stockholmopen.net/>), is a recent and very interesting initiative in the RadioSphere spirit, which tries to stimulate people to share their radio bandwidth in an ad hoc way. The project is started by the Royal Institute of Technology in Stockholm, (<http://www.kth.se/>), which also runs the wireless center [Wireless@KTH](http://www.wireless.kth.se/), (see http://www.it.kth.se/tillfalle/wireless_invigning.htm. No English website available, however).

Concerning scenarios for future wireless developments, there is an excellent work called Wireless Foresight available at <http://www.wireless.kth.se/foresight/>, published June 2002.

7. Conclusion

What we want is Internet, but without a wire. We do not want any “special services” or other things. We want *Internet*, and we want it wireless. How can we get that?

The core idea is a somewhat ruthless exploitation of two facts: 1) that wired broadband Internet is becoming ubiquitous, inexpensive and flat rate, and b) the availability of inexpensive radio basestations of high capacity.

Combining those gives us a Mobile Internet while essentially keeping the price/performance characteristics of the wired Internet. This may make 3:rd generation mobile systems, UMTS, obsolete before being deployed. It may make the role of mobile operations redundant. It is stunning.

If this scenario is right, which we believe, the following conclusions should be drawn for telecom equipment suppliers:

1. The terminal issue must be approached from the pocket computer side rather than from the mobile phone side.
2. Radio base stations become consumer electronics, just as already is the case for computers and phones.
3. Equipment for building the open network of “broadband to homes and gas-stations” become much more important than the closed transport network behind traditional cellular mobile systems.

This all harmonizes wired and non-wired systems, with great efficiency gains for society at large. It is certainly an improvement for all of us.

10. References

Network references have been included directly in the text and are not repeated here.

- [1] (Perkins 1998), Mobile IP, Design Principles and Practices, Charles Perkins, Addison-Wesley 1998.
- [2] (IEEE Wireless Communications 2002), Special issue on Technologies for 4G Mobile, April 2002, Vol 9 No. 2.
- [3] (Solomon 1998), Mobile IP, The Internet Unplugged, James Solomon, Prentice Hall 1998.
- [4] (O'Hara & Petrick 2000), The IEEE 802.11 Handbook: A Designer's Companion. IEEE, Bob O'Hara and Al Petrick.
- [5] (Muratore, ed. 2001), UMTS – Mobile Communications for the Future, Wiley 2001.
- [7] (Castells 2001) The Internet Galaxy, Manuel Castells, Oxford University Press 2001
- [8] (A. G. Valko 1999). Cellular IP - A New Approach to Internet Host Mobility, *ACM Computer Communication Review*, January, 1999.