

RESEARCH CHALLENGES IN MOBILITY AND MOVING NETWORKS: AN AMBIENT NETWORKS VIEW

Ville Typpö¹, Jochen Eisf², Jan Höller³, Ramón Agüero Calvo⁴, and Holger Karl⁵

¹*VTT Technical Research Centre of Finland, Kaitoväylä 1, 90570 Oulu, Finland;* ²*Siemens AG, Otto-Hahn-Ring 6, 81739 Munich, Germany;* ³*Ericsson AB, Torshamnsgatan 23, 16480 Stockholm, Sweden;* ⁴*University of Cantabria, Avda los Castros s/n, 39005 Santander, Spain;* ⁵*Technical University of Berlin, Einsteinufer 25, 10587 Berlin, Germany*

Abstract: This paper aims to address the deficiencies of current internetworking to support mobility and moving networks. We present a novel approach to consider networking in general – the concept of *Ambient Networks*¹ – and discuss related mobility problems. We bring forth several challenges for future research demanding an architectural approach in order to reach a solution in a uniform, efficient, and integrated fashion.

Key words: ambient network; mobility; moving network; network architecture.

1. INTRODUCTION

Mobility is one of the most liberating features, having a profound impact on how communication is evolving into the future. However, a number of challenges remain to be solved, challenges that include revising the architecture principles of the Internet to efficiently and easily cope with new requirements coming from more general application and deployment patterns than those originally perceived for the Internet.

Traditionally, mobility has mainly been about supporting uninterrupted communications to mobile devices connected to a fixed infrastructure. As we will see, there is a new level of mobility support required by the networks supporting emerging ambient and ubiquitous communication needs.

On the one hand, the cellular systems of today, 2G and 3G, have mainly been designed as dedicated and vertical systems, optimised to accommodate a certain set of features including their specialized mobility support. They do provide access to the Internet, but transparency is still far from a reality.

The Internet architecture, on the other hand, was originally designed to be a homogeneous interconnection architecture of networks of stationary computers, providing end-to-end transparency for any type of application. Mobility support was retrofitted later, resulting in a patchwork of specialized solutions at different layers of the system.

At the same time, emerging wireless technologies enable a multitude of new networked devices. For some of these devices, connectivity to a stationary infrastructure is less important than the ability to communicate with and via anything in the surrounding environment in a highly dynamic fashion.

Overall, the environment in which mobility takes place is becoming more heterogeneous. Mobility exposes the users to heterogeneity and dynamics on different levels, e.g. access technologies, networking and trust domains, device capabilities and user contexts. As a consequence, the clear distinction between “network” and “terminal” is increasingly breaking down, as is the distinction between “user” and “operator” with laypersons turning into operators of small personal area or vehicular area networks.

We argue that the challenges faced from mobility include how to cope with this changing communication landscape, how to maintain communications in a heterogeneous networking environment that also includes legacy technologies, and that many of the challenges presented here can only be solved by *addressing the very principles* of the Internet architecture itself. We shall address these challenges in the broader context of the *Ambient Networks*¹ concept.

The remainder of this paper is organized as follows. In Section 2 we review state of the art efforts. Section 3 introduces “Ambient Networks”. Section 4 gives an overview of the diverse aspects of mobility in an Ambient Network; Section 5 provides more details on the various concepts and challenges to overcome. Finally, we conclude the paper in Section 6.

2. RECENT PROPOSALS TO IMPROVE MOBILITY SUPPORT

Mobility solutions can be found by either developing improvements within the current architecture² (as followed within IETF), or by revising the architecture to reflect the changing environment and to comply with the new requirements.

The Internet Indirection Infrastructure (i3)³, e.g., proposes an architecture that offers a communication abstraction based on rendezvous points in an overlay network. These rendezvous points are used to separate packet sending from receiving. Instead of sending a packet to a destination address, the packet is associated with an identifier, which the receiver can use to obtain packet delivery.

Another architectural model, FARA (Forwarding directive, Association, and Rendezvous Architecture)⁴ decouples end system names from network addresses to provide general support for mobility. One advantage of FARA is that neither an entirely new namespace nor a globally unique one is required for the entities. FARA allows several different forwarding mechanisms to co-exist in the network, resulting in variability in the level of characteristics like mobility, identity, and anonymity.

Recent work on HIP (Host Identity Payload)^{5, 6} provides another way of breaking the binding between identities and topological locations of network nodes, a commonly agreed mobility disabler within the current architecture. HIP introduces new cryptographic identities that can be dynamically mapped to IP addresses. While providing support for mobility and multi-homing with a major architectural change in the addressing concept, the solution requires only small changes in current host implementations.

IST BRAIN⁷ defined the architecture and protocols of an IP-based access network with mobility management, in particular, BCMP, the BRAIN Candidate Mobility Protocol. It provides independent, interoperable solutions for local/micro-mobility from global mobility. IST MIND⁸ further developed and validated the concepts and protocols generated in BRAIN by enabling hosts to co-operate with self-organizing wireless ad-hoc networks.

DRiVE⁹ specified a multi-access architecture allowing for seamless intersystem-handover. The concept of a host-controlled flow control was developed to enable parallel usage of different access systems. The architecture is based on Hierarchical Mobile IP, extended by an AAA (Authentication, Authorization and Accounting) component. OverDRiVE¹⁰ extended the scenario with moving networks (e.g. vehicles, trains, etc.) in a multi-radio/multi-access environment, defined a Mobile IP-based solution, and focused on multicast support. The project has strong influence on the ongoing work within the IETF NEMO (Network MObility) group.

3. AMBIENT NETWORKS APPROACH

None of the systems/concepts reviewed in the previous section has been found to provide an all-encompassing solution that can satisfy all requirements for future mobile networking scenarios.

It is one of the basic tenets of the Ambient Networks project¹ that a new architectural approach is necessary. Architectural principles are required that can embrace a multitude of different communication environments, rather than suffer from the heterogeneity. Our approach, which we just started to investigate, is to use *network composition* as such a principle: Instead of terminals, networks as such can form the basic building block of the communication architecture, and the combination of simpler networks into more complicated ones can be a natural operation on these building blocks. One important consequence of this approach is that *everything is a network* – even a single terminal will be regarded as a simple network.

More precisely, everything that supports two interfaces is an Ambient Network (AN). The first one is the *Ambient Network Interface* (ANI), which individual networks can use to negotiate and execute compositions. The second interface is the *Ambient Service Interface* (ASI), which allows applications to access services provided by Ambient Networks (Figure 1).

Such network composition can happen at very different scales, ranging from a PDA composing with a car to two large operator networks composing with each other. The principle and the protocols are the same, the details will differ, e.g., in service level agreements or charging structures.

Network composition is a more powerful concept than the simple internetworking as enabled by the Internet Protocol. While IP can ensure cross-domain/cross-technology transport of *data*, it is not at all clear how to *control* such transmissions. The current Internet assumes homogeneity in the environment in which to provide control. A unifying view like Ambient Networks has the potential to solve this issue of fragmented control.

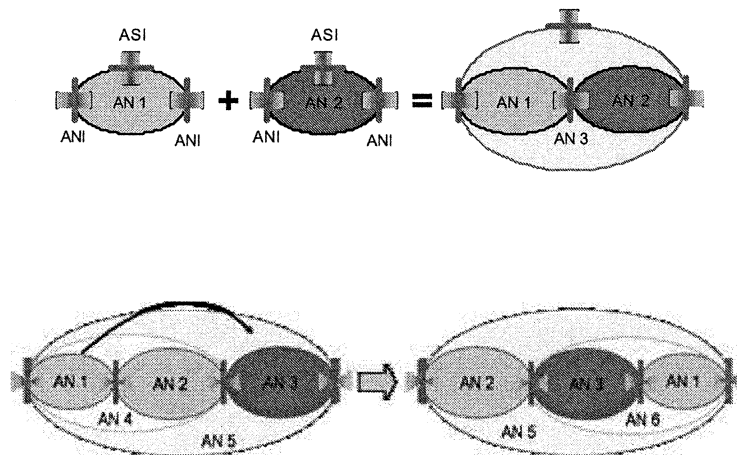


Figure 2. Schematics of a composition-based mobility solution where AN1 moves and triggers the formation of AN6 after dissolving AN4; AN5 is unchanged from the outside.

4. AN ARCHITECTURAL GRIP ON MOBILITY

The vision of composable networks considerably influences mobility support. The primitive mechanisms for composing and separation of networks are closely related to mobility, as mobility can be regarded as a decision on which network should be composed with or separated from which other one and at which point in time. Composition protocols could be simply leveraged for the mechanics of mobility support. Figure 2 outlines such a composition-based mobility support. Moreover, by virtue of the “everything is a network” principle, support for entire mobile networks should not require any special solutions like in the current Internet.

On the other hand, the structure of composed networks is considerably more complex and will have many more opportunities for mobility than a traditional “infrastructure backbone vs. mobile terminal” cellular network. It will make mobility-related decisions much more complex; they also have to be integrated with other considerations as well, e.g., connectivity, application requirements, trust conditions, movement prediction.

When extending networking scenarios from terminals talking to a fixed infrastructure, mobility support can no longer be limited to ensuring a topologically correct IP address after a handover and informing all involved nodes about such an address change. Rather, a broader notion of mobile entities has to be supported with respect to mobility over several dimensions – switching flows between terminals, mobility between trust domains, and mobility across heterogeneous networks with very different capabilities. The current research and standardization efforts, being mostly focused on performance optimisations and scalability problems, are inadequate to solve these fundamental architectural challenges; for example, the 3G mobile network architecture defines clear boundaries between network domains, prescribing a static infrastructure topology.

AN network composition goes beyond what the Internet and mobile networks provide today. Internetworking shall happen not only at the level of basic addressing and routing, but additional functions, for incorporating higher layer support (such as content distribution or context awareness functions) are required.

More precisely, the incorporation of new functions into existing mobility protocols and mechanisms does not appropriately solve the demands of future communication scenarios. We therefore recommend to define a new ‘mobility architecture’, which should be considered as a fundamental component of an overall Ambient Network approach.

Some important architectural principles to be followed include:

- **Diversity:** Supporting and combining several mobility mechanisms and protocols in the very same architecture is considered necessary.

- **Harmonization among layers**: The interaction with mobility support functionalities across layers in the communication stack is necessary.
- **Legacy Awareness**: A migration path from legacy mobility concepts towards Ambient Network is required.
- **Concept of mobile entities**: A mobile entity shall not reflect just the traditional view of mobile terminals, users and networks, but rather smaller communication end-points together with fine grained data flows, which can be moved in Ambient Networks as a result of any communication patterns.
- **Naming and name management**: An appropriate separation of named objects is required together with a dynamic binding process. The nature of the names depends on the involved lifetimes of the objects. For instance, a mobile object's name must be longer-lived than the name of the object that represents its current location.

5. A BROAD APPROACH TO MOBILITY MANAGEMENT

The traditional view of mobility (a host changing its point of attachment to the network) does not suffice. Consider a host which is multi-homed with two active interfaces. Even if the host itself is stationary, an application can at any moment choose to move an ongoing communication flow from one interface to the other. This is also a mobility action concerning a communication end-point which moves among host interfaces. The role of the smallest communicating and mobile entity is crucial to understand. Any other constellation of mobile entities (hosts, networks) may then be viewed as an aggregation or specialization in relation to the smallest mobile entity.

Four aspects of the mobility architecture will be discussed here, with concerns like efficiency and security being understood without saying.

5.1 Mobility in a multi-domain environment

Communication has to be established and maintained while any mobile entity is moving in a multi-domain environment. Multi-domain here implies that there are multiple providers of access and network connectivity, that multiple access technologies may be used, that mobility takes place across multiple trust domains which enforce trust in various ways, and that multiple IP addressing realms are used, e.g. IPv6 as well as public and private IPv4 addressing spaces. Handovers between any such domain combinations must be handled efficiently, avoiding existing “patchwork-type” results.

Trust's importance will grow and handovers will become contingent upon trust. Trust enforcement can, in fact, disable mobility – think of firewalls stopping certain traffic from entering a particular protected domain.

Ambient networking also implies seamless exploitation of multi-homing, multi-access-technology opportunities. Individual data flows of the same device may take different paths across multi-domain networks.

Multi-homing is not limited to physical interfaces only. A host with a single physical interface may internally experience multi-homing by having several logical interfaces. For instance, a user may move an on-going clear communication to an IPSec-encrypted tunnel. Systems commonly treat these capabilities as separate logical interfaces, and this represents a handover.

5.2 Multiple events triggering mobility

Different kinds of events may trigger mobility management actions: Traditional radio link specific conditions, context-dependent, security-related, upper-layer requirements and other system-, application- or user-dependent events.

To cater for all these events, we need general and coherent mechanisms to enable mobility triggering and to identify related events on different protocol layers in a distributed system. Trigger sources and trigger information will have to be included in the mobility architecture. The decision process (arbitration between triggers, policies), particularly the relative roles and cooperation of the involved Ambient Networks regarding mobility triggering and ruling, needs to be handled by a general and uniform framework. Unambiguous decisions on mobility actions are the result.

5.3 The smallest mobile entity

Complementary to treating mobility as a network handover where all traffic is moved from one interface to another at handover, there is a need to define mobility mechanisms for “smaller components” of communication than communication on a per-interface level. The ability to manage flows of related datagrams or sessions, e.g. a streamed media component, between ultimate end-points of the communication is needed. Examples include the multi-homing case where different flows are individually moved between different concurrent active interfaces as well as transferring ongoing communication contexts between different devices. A key is to define this smallest communicating mobile entity which in a protocol stack resides close to where the Socket API lies today. Work on this subject is emerging^{11,4}.

We introduce a session/flow mobility concept as a generic mobility function for any type of payload and transport. This is a more general solution than e.g. an application-layer protocol like the Session Initiation Protocol¹² could achieve. Topics to address include appropriate session management mechanisms, naming of sessions and communication endpoints, and possible state transfer techniques.

Session/flow mobility is also proposed as a key generic function of any overlay doing specific or directed application level processing where there is a need to have dynamic (i.e. mobility) support for splicing flows via intermediary middle-boxes regardless of the type of application level processing these middle-boxes do. One example of such a middle-box is a node performing content adaptation. A user may wish to alter the content format of an already ongoing session because its context is different; for example, streamed visual media might have to be altered to audio as the user is moving into its car. The middle-box will perform this function, but the session will have to be seamlessly moved to go via the middle-box without interrupting the communication.

5.4 Moving networks

Moving networks (persons with a few devices, entire trains) are a prominent component of future networking scenarios. Figure 3 shows a typical example of moving users with several terminals forming temporary moving clusters and network hierarchies while travelling on a train.

A common characteristic for this kind of scenarios is that some mobile entities that are close by move together, forming a *cluster*¹; these devices will often, but not necessarily, be joined together into an Ambient Network.

The formation of these ANs will be highly dynamic, and some kind of hierarchy will be needed in order to integrate them into encapsulating moving networks. These requirements bring about the necessity of new ways of communication between devices as well as protocols to recognize, set-up, and maintain these cluster-triggered Ambient Networks.

Cluster-triggered ANs share many commonalities with traditional ad-hoc networks. While the ad-hoc community has largely focused on the “internal” operation of such a *single* network, we will have to consider how to deal with *many* such networks in close proximity, overlapping each other in space and time, having different trust and connectivity needs, as well as their interaction with the surrounding, more static ambient communications

¹ Not to be confused with the notion of cluster/clusterheads from topology control problems in ad hoc networks.

environment. In the train scenario of Figure 3, many such ANs can be formed, some connected to the train, moving with respect to each other.

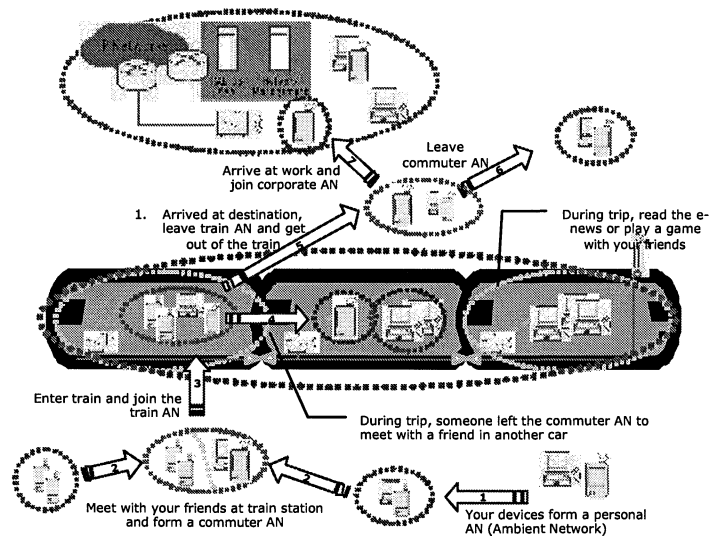


Figure 3. Scenario of moving networks in an Ambient Networks context.

Some issues to consider are:

- **Clusters and network composition:** How can clustered devices turn into a trigger for network composition?
- **Hierarchical ANs:** Network composition naturally results in hierarchies. Clusters, on the other hand, are a concept more tied to physical proximity. How to map these two concepts onto each other?
- **Mobile Routers and Gateways:** To assure continuous reachability, gateways to the outer world are needed. They have to be automatically discovered and re-established after departure/failure of a gateway; exploiting context and access information as well as load balancing between multiple gateways are important.
- **Joint Handovers:** When networks move around, a large number of simultaneous handovers can occur; the signalling traffic should be minimized, even in a hierarchical setting of ANs.

6. CONCLUSIONS

New architectural solutions are needed to comply with the future needs for diverse mobile communication. The novel architectural approach of

Ambient Networks requires new ways of thinking about mobility management as well. We cannot just include new mobility mechanisms without considering the architectural impact. We have presented some architectural principles and concepts, guiding the mobility related research in the Ambient Networks project through its initial phase. We have also discussed several challenges one needs to face when willing to consider mobility in all its richness as a fundamental part of the network architecture.

ACKNOWLEDGEMENTS

The authors would like to thank all our colleagues who contributed to Ambient Networks; Harold Teunissen from Lucent Technologies and Timo Sukuvaara from VTT deserve our special acknowledgements. The Ambient Networks project (IST 507134) is partially funded by the Commission of the European Union. The views expressed in this paper are solely those of the authors and should not be interpreted as necessarily representing the views of their employers or the Ambient Networks project.

REFERENCES

1. N. Niebert, A. Schieder, H. Abramowicz, G. Malmgren, J. Sachs, U. Horn, C. Prehofer, and H. Karl, Ambient Networks - An Architecture for Communication Networks Beyond 3G, *Wireless Communications* **11**(2), 14-22 (IEEE, 2004).
2. B. Carpenter (editor), Architectural Principles of the Internet, RFC 1958 (IETF, 1996).
3. I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, Internet Indirection Infrastructure, *Transactions on Networking* **12**(2), 205-218 (IEEE/ACM, 2004).
4. D. Clark, R. Braden, A. Falk, and V. Pingali, FARA: Reorganizing the Addressing Architecture, *ACM SIGCOMM Comp. Comm. Review* **33**(4), 313-321 (ACM Press, 2003).
5. R. Moskowitz, and P. Nikander, Host Identity Protocol Architecture, Internet Draft draft-moskowitz-hip-arch-05.txt (work in progress), September 2003.
6. P. Nikander, J. Ylitalo, and J. Wall, Integrating Security, Mobility, and Multi-homing in a HIP Way, in: *Proceedings of Network and Distributed Systems Security Symposium (NDSS'03)*; <http://www.isoc.org/isoc/conferences/ndss/03/proceedings>.
7. J. Urban, D. Wisely, E. Bolin, G. Neureiter, M. Liljeberg, T. Robles, BRAIN - an architecture for a broadband radio access network of the next generation, *Wireless Communications and Mobile Computing* **1**(1), 55-75 (Wiley, 2001).
8. P. Eardley, J. Eisl, R. Hancock, D. Higgins, J. Manner, and P. Ruiz, Evolving Beyond UMTS - The MIND Research Project, in: *Proceedings of IEE 3rd Intl. Conf. on Mobile Communication Technologies, Conf. Publ. No. 489* (IEE, 2002), pp. 449-454.
9. T. Paila, S. Alladin, M. Frank, T. Goransson, W. Hansmann, T. Lohmar, R. Toenjes, and L. Xu, Flexible Network Architecture for Future Hybrid Wireless Systems, in: *IST Mobile Summit 2001*, Barcelona, September 2001.
10. M. Ronai, A. Petrescu, R. Tönjes, and M. Wolf, Mobility Issues in OverDRiVE Mobile Networks, *IST Mobile Summit 2003*, Aveiro, Portugal, June 2003.

11. A.C. Snoeren, H. Balakrishnan, and M.F. Kaashoek, Reconsidering Internet Mobility, in: *Proc. of the 8th Workshop on Hot Topics in Operating Systems* (IEEE, 2001), pp. 41-46.
12. J. Rosenberg et al., The Session Initiation Protocol, RFC3261 (IETF, 2002).