Supporting VHO by a Self-Organizing Multidimensional P2P Overlay

Simon Oechsner,Tobias Hoßfeld, Kurt Tutschku
Department of Distributed Systems
University of Würzburg, Germany

Frank-Uwe Andersen
Siemens AG,
Siemensdamm 62, Berlin, Germany

***Abstract*— Vertical handovers (VHO) are expected to be a key feature in Beyond 3G (B3G) networks. This paper presents a CAN-based P2P overlay network for supporting vertical handover in B3G networks. The P2P overlay is used to quickly locate attachment points (APs) for mobile entities and to retrieve rapidly the configuration and coverage information of these APs. The advantage of the P2P-based solution is its distributed nature, its scalability, and its self-organizing capability. We show by means of simulation the efficiency and the scalability of our approach in comparison to a standard CAN implementation.**

I. Introduction

For mobile communications beyond 3G, heterogenous access systems are expected. A mobile terminal moving through a landscape and variety of access systems such as WLAN (802.11), WIMAX (802.16), 2.5G or 3G, needs to perform frequently *vertical handovers (VHO)*, i.e. pass the ongoing connection from one access system to another, as well as from one operator to another. All of these access systems consist of numerous *attachment points (APs)*.

The execution of VHO can be accelerated and made more secure if the involved entities are aware of the attachment points and their coverage areas, cf. [1], [2]. This concept saves considerable time which would otherwise be spent in scanning the environment, e.g. field strength. The identification of candidate APs is achieved by comparing the position of the mobile device and the attachment points. An entity controlling the execution of the VHO can query an *information database* (DB) for AP candidates, their configurations, and their coverage areas and thus decide where to hand-off to.

Typically, the information will be stored in a central entity. Central databases, however, are vulnerable to system failures and overload. A central database is restricted if time-limited coverage area information, e.g. measurements by mobiles, needs to be stored and retrieved quickly.

This work describes the application of a distributed, self-organizing *peer-to-peer (P2P)* mechanism to overcome the aforementioned disadvantages. The suggested P2P-based vertical handover support architecture (P2P-VHOSA) employs a modified Content Addressable Network (CAN) [3] to distribute the DB containing the coverage information. Besides solving the overload problem, the application of a self-organizing P2P mechanism permits new ways of operating mobile communications systems. APs can be inserted easily with no or less manual interference, such as configuration tasks, due to the self-organization of CAN. Hence, APs of different technologies, and even of different operators, can easily be included into the system.

The rest of this work is organized as follows. Section II gives a brief overview on work dealing with VHO and introduces the basic information system for supporting VHO. The enhanced architecture of the distributed information system based on the CAN P2P overlay is described in Section III. Section IV presents numerical results on the efficiency and the scalability of our solution. Finally, Section V concludes this paper and gives an outlook on future work.

II. State of the Art

Vertical handover related issues embrace many different problems. The two mainly investigated areas are the one dealing with protocols to support VHO and the one addressing the resource management to keep QoS requirements. For example, [4] proposes a full-featured protocol suite by integrating Mobile IP, while Balasubramaniam et al. [5] focuses on the decision process adopted for VHO for multimedia applications in pervasive systems and wireless networks. The proposed decision process guarantees also certain QoS requirements. Gurtov and Korhonen [6] offer a complementary perspective of the problem, analyzing effects of the VHO on TCP-friendly protocols.

In [1], Siebert, Lott et al. describe a *Hybrid Information System* (HIS) considering the resource management of VHO. The HIS aids inter-system handover or vertical handover by creating a database where information gathered from the different Radio Access Technologies (RATs) is made available.

The database is filled with measurements of mobile subscribers or attachment points like received signal strength indicator (RSSI) or block error rate (BLER). These measurements are associated with the locations that they correspond to, so that the database resembles a map of the covered area with measured values for points on this map. This reduces scanning effort and provides information for a fast handover decision.

When a handover decision has to be done for a mobile device, the database is queried with the location of that mobile. Then, an Intelligent Service Control (ISC) decides which measurement reports are of interest by computing a so-called Decision Area (DA) and merges all reports covered by that area, cf. [1]. This processed information is then used by the entity deciding about the handover, which may be both network-initiated/device assisted or device initiated/network assisted.

However, an infrastructure is required that enables this resource management. Being a centralized component, the database used by HIS

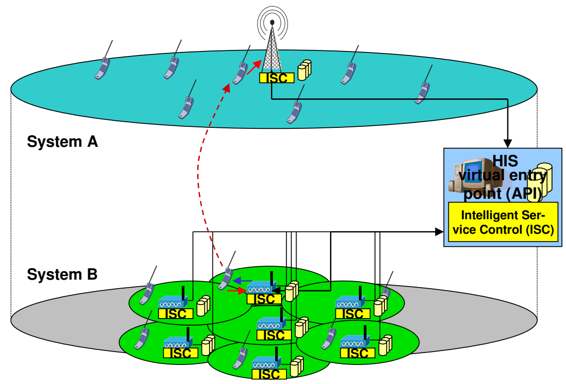

Fig. 1. Distributed information system: the database consisting of measurements is distributed among APs

requires a high performance element resulting in high costs. Additionally, the signaling delay is dominated by the propagation delay from any entity in the network to the server. This means that a high load at the server is observed although the information is only required locally, since a mobile asks for measurement around its current position.

III. VHO Support Architecture

To solve these problems, we propose to distribute the database, along with the functionality of the ISC, to the attachment points themselves. The measurement information is stored locally near the potential requestors, which offers two advantages. First, we reduce response times to queries if we manage to route efficiently to the AP storing the requested information. Second, we segment the database in a way that allows us to put only a low load on each of the devices where pieces of this database are stored, achieving advantages like cost-reduction and resilience. A schematic for this new architecture can be seen in Figure 1.

The key to achieve this distribution and to storing the information near to the requestors is a geographical layout of the distributed database. Each part of the database stores only the measurements that are done in a two-dimensional geographical zone around it.

In order to manage these zones, and to enable communication between the nodes holding the database, we will use the basic overlay architecture of CAN [3] and adapt it to our needs. CAN is a peer-to-peer network that uses a d-dimensional torus as a basic structure. Each peer is responsible for a zone and holds a routing table consisting of the nodes that own the neighboring zones. Routing uses a proximity metric in the Cartesian space to forward messages.

Overlay layout The overlay we construct consists of all attachment points of the radio access networks (e.g., UMTS nodeBs, WLAN access points) as peers. The peers are positioned on the overlay according to their physical position, thus achieving a mapping of reality. This positioning is done utilizing the coordinates of the geo-location of the peers as their ID in the overlay, which is also a coordinate in the two-dimensional overlay network. The physical coordinates can be retrieved using e.g. the global positioning system (GPS).

To properly map the different access technologies, we create one 'reality' for each technology that is considered. This means there is one 'UMTS reality', one 'WLAN reality', and so on. These realities are similar to the realities of CAN in the sense that they are basically separate overlays that contain the same set of peers, but have different layouts. All nodes that are part of the overlay are part in every one of these 'technology realities'.

However, there is an essential difference to standard CAN. The attachment points are only responsible for their own zone in the technology reality that matches their own technology (i.e., a WLAN access point only has its own zone in the WLAN reality). In all other technology realities, they are just put into the zone befitting their position, but the zone is not split or altered (similar to the zone overloading mechanism described for CAN). Through this mechanism, we introduce a logical separation of different technology types into the network. The resulting layout for two technologies is visualized in Figure 2.

When a new node joins the overlay, it uses its position to be inserted in the right zones in every reality. In the reality for its own technology, the CAN insertion using overloading is used, i.e., the nodes from different technologies that also lie in the zone are divided up to the two newly created zones (by splitting the old zone) according to their position.

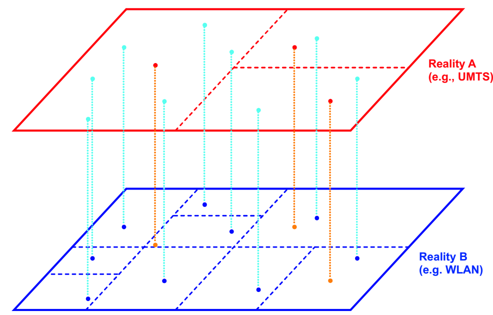

Fig. 2. An architecture with two technology realities

In the other realities, the new node is just inserted in the according zone, informing the node responsible for that zone of its presence. This node is also kept as a separate entry for the differing technology in the local nodes routing table, for a purpose we will explain later. We will refer to these entries as *bridging links*. With this mechanism, each node has one bridging link to every technology reality differing from its own. In case the corresponding nodes in other realities hand over their zone, or the part of the zone containing the local node, the according bridging link also has to be updated.

If a peer wants to leave the network, it has to hand over the zone it is responsible for. This is described for the normal CAN architecture and works as well for our problem. Additionally, it has to inform the nodes in other realities which are responsible for the zones containing the leaving node.

Searching for information The measurements are treated as documents with the measurement location as an ID. They are stored on the node that is responsible for this ID, i.e., the node that is responsible for the zone containing this location.

Requests for reports are assumed to originate in a reality different from the one where the requested information is stored, since VHO means a change of access technology. To cope with this, the request is transferred to the target technology reality via the bridging links. The peer where the request originated can use his bridging link to the target technology reality to forward the query to that reality. There, it is routed with the standard CAN routing algorithm, until it reaches the node that stores the requested reports. This node is then responsible to compile the reports and to send them back to the requestor via a direct link.

IV. Numerical Results

In this section, we describe our simulation scenarios and the results of our experiments. All simulations were conducted with a simulator developed in Java for this purpose. We take a look at the search speed of the proposed P2P-based VHO support architecture (P2P-VHOSA) in comparison to an implementation which is based on the standard CAN algorithms, denoted as standard CAN. In this standard CAN, nodes and measurements get a random ID in the 2-dimensional overlay which is independent of the underlying geographical layout. All APs are in one single reality. Additionally, we compare the effect of the network composition on the search process for P2P-VHOSA and standard CAN.

The search speed for measurement reports is one of the most important performance indicators for ensuring a seamless VHO. We used a network setup of 60,000 UMTS nodeBs and a number of WLAN access points varying from 40,000 to 180,000 which were uniformly randomly distributed in a 40,000 km^2 area, roughly the country size of Germany. All of these attachment points were assumed to be connected with a capacity of 1 mbit. The size of a typical search message was estimated to be 204 byte.

We derived the transmission delay between two nodes from their physical distance which allows us to assess the locality of a search. The processing time of packets at nodes was neglected, however, the hop distribution allows to approximate the total processing time for a search query.

To gather data about search requests, we started close to one million searches at nodes chosen randomly from the network. In our approach, the location for which information was requested was close to the according peer, simulating a mobile in the coverage area of that attachment point. The requested measurements were always from the technology differing from the technology of the starting node, i.e., if the search started at a WLAN AP, UMTS measurements were requested. For the standard CAN network, random keys were searched, starting from random peers. For each request, we recorded the time it took to reach the destination node, as well as the number of hops taken.

Figure 3 shows the distributions for the number of hops search messages had to take for the P2P-VHOSA and the standard CAN, each with three different network sizes.

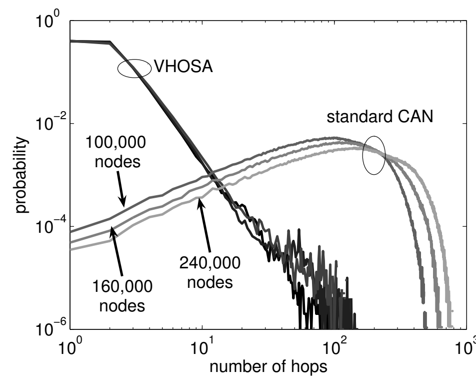

Fig. 3. Hop distribution for a standard CAN network and our modified CAN

We can observe that the number of hops needed to reach the search destination grows with a larger number of peers in the standard CAN. The path length is in the expected range, O($n^{\frac{1}{d}}$), for the number of dimensions $d = 2$ and the number of nodes n between 100,000 and 240,000 [3].

In contrast, the P2P-VHOSA shows a significantly lower number of hops for search requests. Close to 80% of the searches need one or two hops, while the rest rarely exceeds 10 hops. This can be explained by the structure of our network. Measurement reports are requested for a small area surrounding the mobile for which handover is considered. Since this mobile is in the coverage area of the access node that originates the query in the overlay, the location of these reports is close to the location of the access node. Due to the layout of the network, the bridging links also lead to an access node that is in that same geographical area. Moreover, the reports are stored on a peer that is in the same range as the mobile. All of this combined leads to short routing paths, i.e. a low number of hops. Moreover, the network size has no negative effect on the hop distribution, since a higher network density does not change the local nature of a search.

The effect of the locality of searches can be seen even better when we take a look at the search speed. Figure 4 shows the cumulative density functions (CDF) for the search times, again for the same types and sizes of networks.

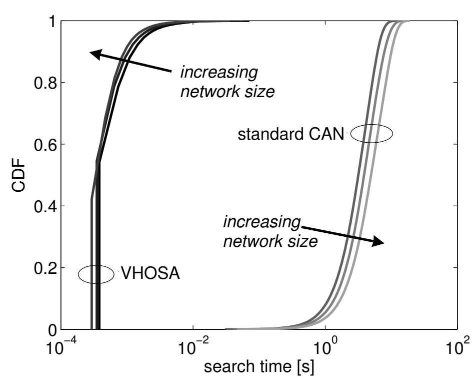

Fig. 4. CDF of the search time for a standard CAN network and our modified CAN

We can observe that requests in our modified CAN are answered faster than in the standard algorithm. This can in part be attributed to the lower number of hops needed for a search. Moreover, under our assumptions these hops are also shorter than the medium hop length in the standard CAN, which spans half of the network. The same effect leads to the faster search times in larger P2P-VHOSAs, since a higher node density leads to closer nodes, while the number of needed hops is about equal.

V. Conclusion

This contribution outlines a CAN-based vertical handover support architecture which stores and locates radio measurements is a fully distributed way. The measurements can be used to increase the reliability of VHOs. The advantages of the proposed architecture are *a)* the avoidance of expensive central databases, *b)* the fast retrieval of information due to short overlay paths, and *c)* the self-organization capability of CAN which permits a highly variable topology of the radio access network, thus enabling a new way of adding APs and of operating mobile systems.

We showed that our approach offers an efficient search for information while it is still scalable. Future work will focus on the resilience of the proposed architecture and will compare the chosen network model to other setups.

Acknowlegment

The authors would like to thank Melanie Brotzeller for her help conducting the simulation and her assistance to the authors.

References

[1] Matthias Siebert, Matthias Lott, Marc Schinnenburg, and Stephan Göbbels. Hybrid information system. In *Proceedings of IEEE Semiannual Vehicular Technology Conference (VTC2004-Spring)*, Milan, Italy, May 2004.

[2] Matthias Siebert, Daniel Bültmann, and Matthias Lott. Inter-system handover and coverage detection for 3g/wlan cooperation (invited). In *Proceedings of 11th European Wireless Conference 2005*, pages 328–336, Nicosia, Cyprus, Apr 2005.

[3] Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard M. Karp, and Scott Shenker. A scalable content-addressable network. In *SIGCOMM*, pages 161–172, 2001.

[4] Y. Zhang G.G. Richard III S. Helal, C. Lee. An architecture for wireless lan/wan integration. In *Proceedings of the IEEE Wireless Communication and Networking Conference (WCNC)*, 2000.

[5] Indulska Balasubramaniam, S. Vertical handovers as adaptation methods in pervasive systems. In *Proceedings of the 11th IEEE International Conference on Networks*, pages 705–710, 2003.

[6] J. Korhonen A. Gurtov. Effect of vertical handovers on performance of tcp-friendly rate control. *ACM Mobile Computing and Communications Review*, 8:73–87, 2004.