

Data Collection for Heterogeneous Handover Decisions in Beyond 3G Networks

Marc Fouquet¹, Christian Hoene², Morten Schläger³, and Georg Carle¹

¹ Technical University of Munich

² University of Tübingen

³ Nokia Siemens Networks, Berlin

Abstract. Future mobile networks will be increasingly heterogeneous. Already today, wireless LAN is used by many mobile network operators as an addition to traditional technologies like GSM and UMTS; WiMax and 3GPP Long Term Evolution (LTE) will be added. Having heterogeneous wireless networks, one challenging research question needs to be answered: Which user should be served by which access network at which time and when to conduct a handover? For such decisions, information on the state of networks and terminals is required.

In this publication we simulate mobile networks in which a central entity called Network Resource Management (N-RM) gives handover recommendations to mobile terminals. Based on these recommendations and local knowledge on link qualities, the terminals choose the cell to switch to.

The N-RM should have a global view on the networks to give best recommendations. We designed the Generic Metering Infrastructure (GMI), a publish/subscribe system to collect information about access networks and terminals efficiently.

We investigate the tradeoff between the signalling overhead caused by data collection and the quality of the handover decisions and show, how smart monitoring can reduce the amount of measurement data while ensuring the efficient use of heterogeneous networks.

1 Introduction

Future mobile networks will be heterogeneous, i.e. consisting of GSM, UMTS, WLAN, WiMax [1], LTE [2], and other radio access technologies. 3GPP is currently standardizing IP-based mobility solutions that will allow seamless handovers between such technologies [3].

Having a heterogeneous network, we need to decide, which user should be served by which access technology at which time. As different user and operator preferences must be considered, a policy-driven decision engine has to make the handover decisions [4]. It is probably impossible to develop one “perfect” algorithm that fits all needs.

Deciding handovers locally in the mobile terminal is sub-optimal, as the global situation in the network, i.e. the load in the different cells, can not be taken into account. It would be advantageous to have a “wizard of oz” view on the access networks and make network-side handover decisions with perfect information, but, of course, it is impossible to collect all data and to make decisions on time. In addition flat hierarchies in future mobile networks make the process of data collection increasingly difficult.

In LTE networks there will no longer be a node like the UMTS RNC, which already has load- and radio data of hundreds of cells, but only evolved nodeBs located much closer to the antennas. With wireless LAN access points, the situation is similar. Collecting management data is expensive as the base stations are spread in the countryside, with costly rented or wireless “backhaul”-links connecting them to the operator’s core network. Thus, any central view on the network will be inaccurate and delayed.

To get an efficient view on the state of the network, we have designed the **Generic Metering Infrastructure (GMI)**[5], a publish/subscribe system for collecting and distributing measurement data in an operator’s network. It uses various compression techniques for efficient data transport. The GMI is intended for all kind of management applications, i.e. for fault management, security management, and resource management. In this publication we concentrate on the usage of GMI for making handover decision.

We developed a network-side decision engine called **Network Resource Management (N-RM)**, which resides in the core network (Figure 1). It implements algorithms, which give the terminals handover recommendations. The mobile terminals make the final handover decisions based on their local knowledge on the radio conditions and the N-RM recommendations. We base this exemplary system on research results by Fan et al. [6] achieved in the BmBF project “ScaleNet”.

In this work we do not aim at increasing the performance of the handover system because it could follow diverse goals and any arbitrary policies. Instead, we take a few algorithms as examples and study, on how the accuracy of information influences the quality of central decisions. More precisely, we primarily focus on the effect of different strategies for data collection and the resulting quality of mobility decisions.

We describe the GMI in Section 2 and the setup for our experiments in Section 3. Results are presented in Section 4 and interpreted in terms of data volume in Section 5. We discuss related work in Section 6 and finally conclude the paper in Section 7.

2 The Generic Metering Infrastructure

In this section we describe the Generic Metering Infrastructure (GMI), our publish/subscribe system for future mobile networks. It allows for efficient data distribution by enabling the clients to selectively subscribe for information they need, by distributing the data to interested parties in a multicast-like fashion and by compressing data.

Like all publish/subscribe systems, the GMI offers an event service. Clients can subscribe for “types” of information they are interested in. Whenever new information is created at a producer, it is sent to the event service which then takes care of distributing the new data to all interested clients.

2.1 GMI Design

The GMI’s event brokers are called **Metering Management and Collection Entities (MMCEs)**. These nodes manage the subscriptions and create distribution trees if multiple receivers are interested in the same data.

Our generic metering infrastructure is a subject-based publish/subscribe system. This means that the “types” of information as described above are organized in a tree, an individual leaf of the tree is identified by a DNS-like address.

The GMI supports 3 different types of requests:

Periodic measurements: A metering client can subscribe to periodic metering tasks. In this case the subscriber specifies a desired report period and the subject it wants to stay informed about.

Triggers: It is also possible to set triggers for measurements. For such a subscription the metering client specifies one or multiple thresholds for a metered value. If the value rises above or falls below the given threshold the client is informed immediately.

Request/Reply: The last type of reporting is an immediate response to a request of a metering client for a certain value of data. This notification is not an event in the classical sense of a p/s-system. In this case the metering client simply sends a request for a subject (which is a message similar to a subscription) and receives a reply containing the value (which is handled like a notification).

In each case the GMI uses appropriate methods to optimize the data transfer, i.e. by merging similar tasks and building distribution trees.

2.2 Granularity Periods and Report Periods

Monitoring in 3GPP networks usually does not directly deal with raw data from the network elements, but with derived “key performance indicators” (KPI).

A network element collects raw data for an interval called **granularity period (GP)**. After the GP is finished, this data is used to calculate the KPIs [7]. With today’s network elements like RNCs, the minimum GP is 5 minutes, while 30 minutes are a far more typical value. Of course monitoring with such a granularity barely helps when building a heterogeneous resource management system.

For the GMI we expect the meters to work the same way as they do today, but with shorter granularity periods. Each meter has a specific minimum GP. New data for the GMI’s publish/subscribe system is only available when a GP has ended.

This means that with *periodic measurements*, the minimum report interval is one GP. With *triggers*, the system can only evaluate at the end of each GP, whether or not the trigger has fired. *Request/reply-style queries* are answered after the current GP has ended, so they may get delayed for one GP.

For periodic measurement jobs, the client specifies the reporting interval in multiples of the granularity period. We call this property of measurement jobs the **report period (RP)**. So a report period of 10 GPs means that each 10th metered value is actually reported to the client.

2.3 GMI in future mobile networks

Figure 1 shows, how the GMI could be deployed in a 3G-beyond network. For UMTS, meters would be placed at the Radio Network Controller (RNC), a central node that controls hundreds of cells. However in other radio access networks, the necessary data has to be collected at much more distributed locations.

The links between the cells (that may be located somewhere in the countryside) and the packet core network are a scarce and expensive resource for mobile network operators. So any management task has to be careful to save bandwidth here.

The interface between MMCEs and meters depends on the type of meter and the link between meter and MMCE. Legacy meters could use well-known protocols like SNMP here, while GMI-enabled meters would use protocols that are optimized to save data.

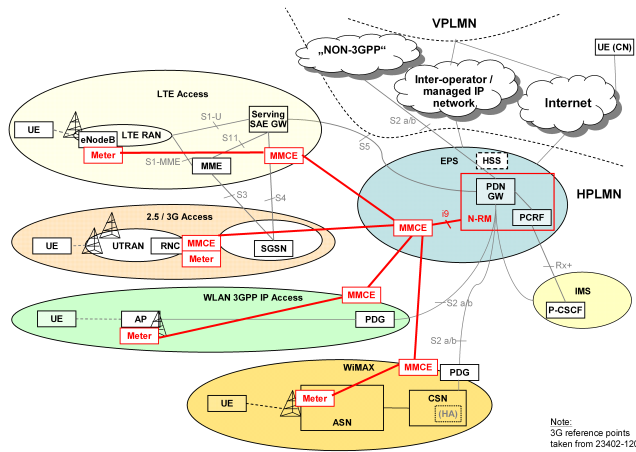


Fig. 1. Mapping of the GMI to the SAE network architecture.




Fig. 2. Overview of the experiment setup.

3 Simulation Setup

In this section we describe the setup of our experiments. All components of our simulations can be seen in Figure 2. They were run on a single machine, but as separate applications using TCP/IP communications.

3.1 The User- and Radio Network Simulator

The simulator is an application that simulates the mobile networks and their users. Several cells belonging to different **radio access networks (RANs)** are placed on a map. The users move around on this map and start and stop sessions. We use an “accelerated real-time” simulation, events in the simulator happen roughly 15 times faster than they would in reality.

The map is also shown in Figure 2. It contains one LTE cell, five UMTS cells and three WLAN hotspots. Signal quality is based on the distance between the user and the

center of the cell, so our UMTS, LTE and WLAN cells only differ in bandwidth and range.

The movement model of the users in our simulations is a modified random waypoint pattern. The users select a new waypoint somewhere on the map with a probability of 40%. With 10% probability they will move to one of the three interesting locations that are covered by the WLAN hotspots. In 50% of the cases the user will stay at his current position for an exponentially distributed random time. Session duration is exponentially distributed as well, between the sessions there is an also exponentially distributed waiting time with same expectation, therefore each user has 0.5 sessions on average.

Our users produce actual load on the air interface by opening sessions, while the GMI produces signalling load which only occurs on the backhaul links between the base stations and the core network. Our assumption here is that there is no direct interaction between these two kinds of traffic, even though they have to share the backhaul link. The rationale for this assumption is the idea that the end-users' traffic will primarily be limited by the available bandwidth on the air interface. The backhaul-link should be dimensioned as small as possible as it produces costs for the operator, but this will be done according to estimates for the sum of user and signalling load.

On startup, users will connect to the strongest RAN. There is a network-internal resource management build into the simulator, which allows handovers between cells of the same RAN, the handover logic for this case is given in Algorithm 1 which is executed for each user periodically.

There are no handovers between different RANs as long as no global resource management application (N-RM) is running. When users leave the coverage area of the current RAN they will lose their sessions and scan for a new RAN as soon as they notice that the old one is gone.

The simulator is able to accept GMI subscriptions for a large number of parameters regarding cells and users. The minimum granularity period of all meters is one second, which would be 15 seconds in the real system.

3.2 N-RM

The network resource management (N-RM) is our global application for handover management. It gets data from the simulator via GMI, its decisions are passed back to the simulator and influence the mobile terminal's cell selection (see Figure 2).

We assume that N-RM does not know the exact location of the user — it only knows the user's location on a "per cell" granularity. Therefore it can not tell exactly, which cells are available at the user's current position, it only has a rough idea, which cells overlap.

The mobile terminal initially also does not know about cells of other RANs. In real life it would have to scan different frequencies to find alternative cells — and this constant scanning would waste battery power. So in our case, N-RM will give the mobile terminal hints, which cells to search for. In reality these hints would also contain radio parameters. The mobile terminal will only scan for cells of different RANs after receiving such a hint.

We are testing five different setups:

Algorithm 1: Mobile terminal- and RAN-internal RM decisions

```
input: users, cells
foreach  $u \in users$  do
  // mobile terminal-side logic. If N-RM recommended cells,
  // the mobile terminal tries to connect there.
  foreach  $c \in cells\_recommended\_by\_global\_rm(u)$  do
    // check signal quality
    if  $sq(c, u) > th_{minimum\_sq}$  then
      |  $u.handoverTo(c)$ ;
      |  $continue\_with\_next\_user$ ;
  // This is the "local RM" logic. It only makes handovers
  // within cells of the same RAN.
  foreach  $c \in nearby\_cells\_of\_current\_ran$  do
    // check minimum requirements for load and signal
    if  $sq(c, u) > th_{minimum\_sq}$  and  $l(c) < th_{critical\_load}$  then
      // find best available cell
      |  $score(c) \leftarrow calculate\_score(l(c), sq(c, u))$ ;
      | if  $score(c) > score(former\_best\_cell)$  then
      | |  $new\_best\_cell \leftarrow c$ ;
  if  $new\_best\_cell$  found then
  |  $u.handoverTo(new\_best\_cell)$ ;
```

Algorithm 2: N-RM decision logic on bad signal quality

```
precondition:  $u \in users, sq(u) < th_{minimum\_sq}$ 
input : users, cells
 $cells\_recommended \leftarrow \emptyset$ ;
foreach  $n \in neighbour\_cells(cell(u))$  do
  | if  $l(n) < th_{acceptable\_load}$  then
  | |  $cells\_recommended \leftarrow cells\_recommended + \{n\}$ ;
sort\_by\_load( $cells\_recommended$ );
send\_recommendation( $u, cells\_recommended$ );
```

Algorithm 3: N-RM decision logic on cell overload

```
precondition:  $c \in cells, l(c) > th_{critical\_load}$ 
input : users, cells,  $user\_list(c)$ 
 $cells\_recommended \leftarrow \emptyset$ ;
foreach  $n \in neighbour\_cells(c)$  do
  | if  $l(n) < th_{acceptable\_load}$  then
  | |  $cells\_recommended \leftarrow cells\_recommended + \{n\}$ ;
sort\_by\_load( $cells\_recommended$ );
 $receivers = random\_subset(users(c))$ ;
foreach  $r \in receivers$  do
  |  $send\_recommendation(r, cells\_recommended)$ ;
```

- *Experiments without N-RM*
In these runs, the users will stay in their RAN until they lose connectivity.
- *Experiments with a purely trigger-based N-RM*
The N-RM subscribes for the load in each cell and the signal quality of each user. The subscriptions are trigger-based, which means that N-RM is notified whenever a threshold-value is crossed.
In case of a user with bad signal quality, N-RM will request load information from surrounding cells using ONCE requests. Based on the results it will recommend cells to the user (see Algorithm 2). This recommendation will have an impact on the next run of the local RM (Algorithm 1) inside the simulator.
In case of an overloaded cell, N-RM will again request load information from the neighbour cells. It will choose a subset of the current users in the overloaded cell and send its recommendations to this subset (see Algorithm 3).
With this basic triggered N-RM, the resource management will only be informed when a critical situation occurs, there is no feedback whether the situation persists despite the countermeasures. As our triggers are defined with hysteresis, N-RM will only take action again when i.e. the load in a cell crosses the upper threshold again after it has crossed the lower threshold.
- *Experiments with an N-RM based on periodic reports*
N-RM subscribes to load information of each cell and signal quality information of each user on a periodic basis. This means that there is a constant flow of reports which does not change during the simulation. N-RM does not request current information when it needs specific data, but it uses the last reported value.
The basic decision algorithms are still Algorithm 2 and Algorithm 3. These are run periodically whenever new data has arrived.
- *Combination of triggers and periodic reports*
This variant works with triggers again, but after a trigger has fired and N-RM has taken action, it enters a success control loop. In this state the trigger will be turned off and replaced by a periodic measurement job which continuously monitors the critical value.
With each arrival of a current value, N-RM will check if the system still is in a critical state. If this is the case, it will again request additional data which is needed for the decision and take action accordingly. There is a different (much lower) threshold for cancelling the periodic measurement job and returning to the triggers when the situation has been resolved.
- *N-RM with full information*
For comparison all simulations have also been run using a N-RM which subscribes for periodic reports of all values in the simulator with a report period of 1 GP. This can be seen as the theoretical maximum amount of data that N-RM could possibly subscribe to. The decision algorithms are still the same.

4 Results

Figure 3 shows the amount of traffic which is produced by GMI reports. On the horizontal axis we have increased the number of users in the system and therefore the offered load.

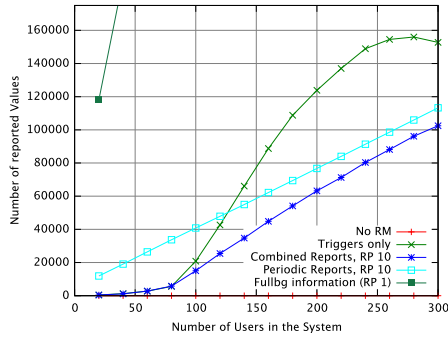


Fig. 3. Number of transmitted values.

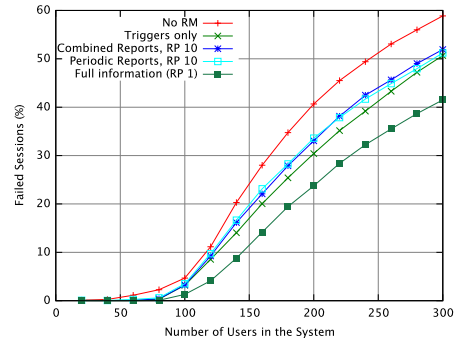


Fig. 4. Failed sessions.

One can see that the curve for periodic reports with a report period of 10 GP is linear in the number of users as expected. The curve for a report period of 1 GP is also linear but comparably high.

When the system load is low, triggers produce only very few messages as almost no critical events happen. However, when the load increases, the number of trigger-messages explodes. As one can see, the number of messages decreases again when the load is extremely high, as the parameters always stay above their thresholds.

The “combined reports” curve shows the result of an N-RM algorithm which switches to periodic reports after receiving a trigger for a subject. As one can see this curve always produces less messages than the triggered or periodic curves.

The quality of the decisions is evaluated using Figure 4, which shows the percentage of failed sessions. One astonishing fact is that the trigger curve is better than periodic reports and combined reports. This was not expected as the used trigger algorithm is quite primitive, it only acts when it gets an ascending trigger and does not control the success of its actions. However the variance in the measured parameters is high, so the resource management is still triggered very often — which causes N-RM to send more recommendations to the users.

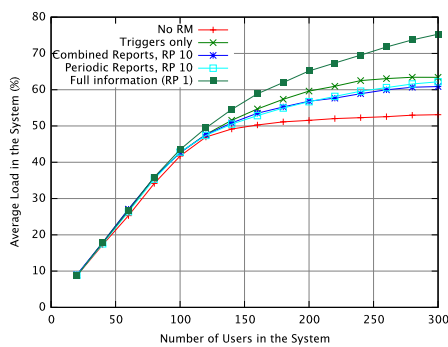


Fig. 5. Average load.

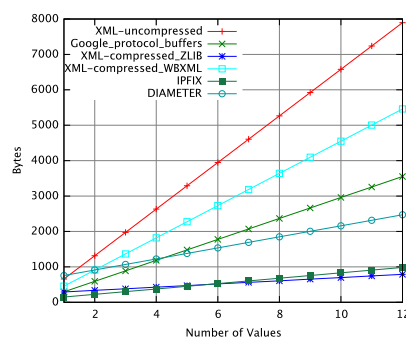


Fig. 6. Data volume when transporting measurement data using different protocols.

Figure 5 shows the average load in the system. Here the number of users can be seen as a measure of the offered load, while the vertical axis is the actual load which could be handled by the system. The main room for improvement through a network resource management are the WLAN cells here. Without resource management, the users will never switch to WLAN but stay in UMTS or LTE even when passing by a WLAN cell. However with RM, some users can be moved to WLAN, therefore the available bandwidth of the system can be used better.

5 Data Volume

Resource Management	Number of Users	Failed Sessions	Data Volume per Cell
None	80	2.28%	0.00 kBit/s
Combined Reports (RP 10)	80	0.40%	0.01 kBit/s
Maximum Information (RP 1)	80	0.09%	0.64 kBit/s
None	160	28.01%	0.00 kBit/s
Combined Reports (RP 10)	160	22.09%	0.09 kBit/s
Maximum Information (RP 1)	160	14.16%	1.18 kBit/s

Table 1. Data volume vs. success

So far, we have only been talking about the number of values that were sent through the GMI. In this section we will discuss the volume of management data which has to be sent over the backhaul links for heterogeneous resource management.

In our implementation, the GMI uses an uncompressed XML data format. However in a production environment it would be favourable to use more bandwidth-efficient protocols. Basically any protocol which transmits key-value-pairs is suitable for transporting the GMI messages.

Figure 6 shows the performance of six different candidate protocols when transmitting a representatively chosen GMI dataset. Note that this is not a hard comparison of these protocols which is simply not possible in such a short analysis. This section is rather meant to give a hint on how to estimate the practically required data volume. With each of the candidates there are almost unlimited possibilities to encode the data differently which of course affects the volume. It should also be noted that so far, only application layer data volume has been considered, there are no IP-headers and also no headers of the layer 4 protocols that might be used in combination with our six candidates (TCP, UDP, SCTP).

- Uncompressed XML has advantages in terms of transparency as the data is transmitted in plain text. However it wastes bandwidth on the expensive backhaul links.
- WBXML [8] is a standard by the Open Mobile Alliance, which replaces XML tags by shorter binary strings, but leaves the actual content of the tags unchanged. In absence of a DTD, it builds a string table from the tag names and uses references to this table afterwards, which was the case here. In our example the data has been compressed to 69% of the size of the original XML.

- Google Protocol Buffers [9] is another representation which preserves the hierarchical structure of XML. In the example the data was compressed to 45% of the original size. The direct comparison with WBXML may be a bit unfair, as the Protocol Buffers encoder was able to use meta-information about the structure of the document. With a DTD, we would expect WBXML to perform roughly equivalent to Google Protocol Buffers.
- Diameter [10] is shown in this comparison as it is a common accounting protocol in 3GPP networks. It requires the GMI messages to be mapped to flat key-value pairs. With our example data, the data volume used by Diameter was 25% of the original XML.
- IPFIX [11] is based on Cisco’s Netflow v9 protocol. It is basically meant to export traffic flow data, but it is flexible enough for our purpose. IPFIX also works on flat attribute-value-pairs, but on the link it separates the attributes from the values. The attributes are sent once in so-called template records in the beginning of the transmission, while the values are sent separately in data records. Therefore IPFIX can save bandwidth compared to Diameter, our example-data was compressed to 12% of the original size.
- The last candidate protocol is a simple LZ77-zipped [12] version of the original XML data. The messages have not been compressed one by one, but the state on both sides is held during the transmission of multiple messages. In the long term, after 50 messages, the data volume could be reduced to 7% of the original XML size. However this advantage comes with increased costs in terms of memory- and CPU-usage at sender and receiver.

For Table 5 the results from Section 4 have been combined with knowledge about message sizes. Here we assume data transport by IPFIX, which was the second-best solution in the comparison above, as we want to avoid the computational effort of compressing the data using LZ77. We also assume that each value is sent in a separate packet — which is a worst-case scenario — and add TCP and IP headers. As mentioned in Section 3.1, our granularity period is 15 seconds.

As we can see, heterogeneous access management gives a high benefit while using very little bandwidth. With 80 users in the system the result with our combined periodic and triggered reporting is almost as good as the result obtained with the maximum available information.

With 160 users the radio network is already overloaded as we can see from the success rate of the sessions. However the number of messages we need for N-RM still remains negligible compared to the amount of user data transferred through HSDPA or LTE cells. Here one could even consider to send all available data (GP 1) to the N-RM, which produces 14 times the data volume of the combined reports method, but still stays around 1 kBit/s per cell.

6 Related Work

In recent years, there has been a lot of research on handovers in heterogeneous networks. The approaches in [13], [4] and [14] leave the decision which network to choose

to the mobile terminal. This is a reasonable design concept since the information about signal quality of the surrounding base stations is available there, while with network-centric decision engines this information must be transported to the core network. Transporting this data also introduces undesired delay that leads to potentially imprecise values.

On the other hand a management facility inside the network is able to take global information, i.e. on the load situation into account and therefore is able to make better decisions. Additionally, it can help the mobile terminal to find adjacent networks without forcing it to scan for available access points which would deplete its battery power.

The authors of [15] use a network-centric approach which basically integrates WLAN into an UMTS UTRAN network. Our approach attempts to be more general by abstracting the handover-logic from details of the RANs. [16] adjusts load triggers on the network-side to optimize handover performance. This work is about the actual handover decision process, i.e. in a combined GERAN/UTRAN network, while we primarily focus on data collection and intentionally keep the decision process as simple as possible.

The authors of [6] use a network-assisted policy-based approach. They're using two decision engines, one of which is located in the core network while the other resides on the mobile terminal. This scenario is also the base of our work, the GMI could be used here to provide the network- and RAN-related information which is required by the decision engine on the network side.

A related standard is IEEE 802.21 [17], which specifies an information service, an event service and a command service to support heterogeneous access decisions and therefore consists of several building blocks that were similarly realized for our simulations. However IEEE 802.21 leaves the actual question of data transport open.

7 Conclusion

Today, users of mobile networks increasingly demand data services and voice-calls for flat prices. This makes business difficult for network operators, there is hard competition on the market and revenues are shrinking. Operators have to cut costs and one way to do so is increasing network efficiency. This requires heterogeneous networks with smart management, as different access technologies have different strengths and weaknesses, which leads to a need for new methods of data collection.

We have shown that the flexibility that our GMI provides gives advantages when making heterogeneous handover decisions. It is possible to make good decisions with fewer data by switching between periodic reporting and setting triggers. Our simulation results show that customer experience could be enhanced significantly, while the cost in terms of produced overhead was negligible.

Our concept of the GMI is not only suitable for heterogeneous access management, but also for general management tasks or for security (i.e. distributed intrusion detection). If the GMI would be used for those purposes as well, the gains could be even bigger because of the late duplication property of the GMI's publish/subscribe system.

8 Acknowledgements

The authors would like to thank Mark Schmidt and Vladimir Maldybaev for their help regarding the evaluation and Andreas Monger for his conceptual GMI work and the initial implementation of the MMCEs.

References

1. IEEE802.16, "Air Interface for Fixed Broadband Wireless Access Systems," IEEE Standard for Local and Metropolitan Area Networks, Oct. 2004.
2. 3GPP, "TS 23.401 v8.4.1: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access," 2008.
3. ———, "TS 23.402 v8.4.1: Architecture enhancements for non-3GPP accesses," 2008.
4. H. J. Wang, R. H. Katz, and J. Giese, "Policy-Enabled Handoffs Across Heterogeneous Wireless Networks," in *WMCSA '99: Proceedings of the Second IEEE Workshop on Mobile Computer Systems and Applications*. Washington, DC, USA: IEEE Computer Society, 1999, p. 51.
5. A. Monger, M. Fouquet, C. Hoene, G. Carle, and M. Schläger, "A metering infrastructure for heterogeneous mobile networks," in *First International Conference on COMMunication Systems and NETWORKS (COMSNETS)*, Bangalore, India, Jan. 2009.
6. C. Fan, M. Schläger, A. Udugama, V. Pangboonyanon, A. C. Toker, and G. Coskun, "Managing Heterogeneous Access Networks. Coordinated policy based decision engines for mobility management," in *LCN '07: Proceedings of the 32nd IEEE Conference on Local Computer Networks*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 651–660.
7. 3GPP, "TS 32.401 V8.0.0: Telecommunication management; Performance Management (PM); Concept and requirements (Release 8)," 2008.
8. B. Martin and B. Jano, "WAP Binary XML Content Format," <http://www.w3.org/TR/wbxml/>, 1999.
9. Google Inc., "Protocol Buffers," <http://code.google.com/apis/protocolbuffers/>, 2008.
10. P. Calhoun, J. Loughney, E. Guttman, G. Zorn, and J. Arkko, "Diameter Base Protocol," RFC 3588 (Proposed Standard), Sept. 2003. [Online]. Available: <http://www.ietf.org/rfc/rfc3588.txt>
11. B. Claise, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information," RFC 5101 (Proposed Standard), Jan. 2008. [Online]. Available: <http://www.ietf.org/rfc/rfc5101.txt>
12. J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," *IEEE Transactions on Information Theory*, vol. 23, pp. 337–343, 1977.
13. V. Gazis, N. Alonistioti, and L. Merakos, "Toward a generic "always best connected" capability in integrated WLAN/UMTS cellular mobile networks (and beyond)," *Wireless Communications, IEEE*, vol. 12, pp. 20–29, 2005.
14. E. Stevens-Navarro and V. Wong, "Comparison between vertical handoff decision algorithms for heterogeneous wireless networks," *Vehicular Technology Conference, 2006. VTC 2006-Spring. IEEE 63rd*, vol. 2, pp. 947–951, May 2006.
15. R. Pries, A. Mäder, and D. Staehle, "A Network Architecture for a Policy-Based Handover Across Heterogeneous Networks," in *OPNETWORK 2006*, Washington D.C., USA, Aug 2006.
16. A. Tölli and P. Hakalin, "Adaptive load balancing between multiple cell layers," *Vehicular Technology Conference, 2002. Proceedings. VTC 2002-Fall. 2002 IEEE 56th*, vol. 3, pp. 1691–1695 vol.3, 2002.
17. IEEE802.21, "Media Independent Handover Services," <http://www.ieee802.org/21/>, 2007.