Policy-Based QoS Management Architecture in an Integrated UMTS and WLAN Environment

Wei Zhuang, Yung-Sze Gan, Kok-Jeng Loh, and Kee-Chaing Chua, Siemens Singapore

This work was done while the authors were with the ICM Mobile Core R&D Department, Siemens Ptd. Ltd., Singapore.

ABSTRACT

Strong demands for public wireless broadband services will require more capacity than even that can be supplied by advanced mobile cellular systems like the Universal Mobile Telecommunication System. The increasing popularity of WLANs has prompted mobile network operators to consider their deployment in high-density usage areas like indoor/outdoor public hotspots to provide complementary broadband access to their UMTS networks. In order to provide consistent QoS control over an integrated UMTS and WLAN system, a policy-based multidomain QoS management architecture is proposed in this article. Different UMTS–WLAN interworking scenarios are discussed to illustrate the feasibility of the proposed architecture.

INTRODUCTION

The Third-Generation Partnership Project (3GPP) has standardized the Universal Mobile Telecommunications System (UMTS) as the future high-speed mobile telecommunications system that supports real-time and non-real-time multimedia services. However, the costs of acquiring the necessary radio spectrum and the required network equipment upgrades to provide such services via UMTS are very high. Hence, mobile network operators are increasingly interested in the use of 802.11-based wireless technologies (802.11a/b/g) to provide access to these services at so-called hotspots (hotel lobbies, cafes, etc.) [1]. This is because WLANs use license-free radio spectrum to provide low-cost, easily deployable, high-data-rate wireless services. In these hotspots, WLAN technologies provide nomadic high-speed wireless access to existing wired Internet Protocol (IP)-based networks. Mobile network operators envision the selective use of WLAN hotspots to augment areas of high-density usage in their networks where subscribers do not require the wide-area mobility and seamless coverage of UMTS [2].

Mobile network operators can support quality of service (QoS)-sensitive IP applications like voice over IP (VoIP) over the UMTS packet-switched (PS) domain by using the Session Initiation Protocol (SIP)-based IP multimedia subsystem (IMS). Presently, the 3GPP is extending the policy-based QoS control architecture [3–5] for UMTS IMS services to satisfy the end-to-end QoS requirements of other application services in the UMTS PS domain. The policy-based QoS control architecture is based on the concept of policy-based networking (PBN), where service level agreements (SLAs) describing the sets of IP QoS services that network operators have mutually contracted to provide are enforced in their network domains by sets of policy rules. These policy rules describe the amount of network resources required to realize QoS services without going into the details of how to configure the network devices.

The Internet Engineering Task Force (IETF) has defined a policy framework [6] within which sets of policy rules described in the form of policy models [7] are transformed into network device configurations in an administrative domain. The policy rules are stored in the policy repository from which the policy decision point (PDP), alternatively known as the policy decision function (PDF) in the 3GPP context, retrieves the appropriate policy rules in response to policy events that are triggered by the contracted IP QoS services, such as the reception of a Resource Reservation Protocol (RSVP) message by the policy enforcement point (PEP). The PDP translates the acquired policy rules into a set of QoS mechanism configuration actions based on the capabilities of the PEP and the current network conditions. The PEP then executes these PDP-supplied actions to handle the triggering policy events in accordance with the requested IP QoS services.

Thus, PBN provides a high-abstraction view of a network to its operator, and this helps the operator in the deployment of new IP QoS services as it does not need to consider details con-

0163-6804/03/$17.00 © 2003 IEEE

cerning the size or complexity of the network. The use of automated policy translation entities further facilitates the dynamic control of network resources. This is important because the 3GPP desires to reserve QoS resources in the PS domain only in response to the setting up of an IP multimedia session.

End-to-end communications are likely to involve multiple administrative domains controlled by different network operators. To provide a consistent end-to-end service in a multi-operator multidomain environment, the authors previously proposed a hybrid policy architecture for UMTS IMS [8], in which the hierarchical architecture is employed within a single operator's multidomain network, and the peering architecture interconnects multiple operators' networks. A master PDF (MPDF) in an operator's network is peered with MPDFs of adjacent networks through an interdomain policy agent (IPA). After an IPA successfully exchanges updated SLA information with its peering IPAs, the MPDF translates the new SLA into policy rules applicable to its network before updating its policy repository. When a PDF in the network needs to perform local domain policy control for a PS domain session, it just retrieves and enforces the relevant policy rules from the policy repository.

To ensure that consistent IP QoS services can be provided in an integrated UMTS and WLAN environment, the multidomain policy-based QoS control architecture [8] must be extended into the WLAN domain. The objective of this article is to propose how this extension can be realized. A policy-based QoS architecture for the WLAN domain created within the policy framework defined by the IETF is proposed. We discuss how the policy-based QoS architectures of the UMTS PS and WLAN domains can be integrated in different UMTS–WLAN interworking scenarios. We conclude the article and describe some of the outstanding issues in our proposed architecture that will need further study.

Policy-Based QoS Architecture in WLAN Networks

Currently, 3GPP is studying the feasibility of interworking UMTS and WLAN systems. The intent of UMTS–WLAN interworking is to extend UMTS services and functionality to the WLAN access environment so that the WLAN effectively becomes a complementary radio access technology to UMTS. The approach of 3GPP is to design flexible, scalable and general UMTS–WLAN interworking capabilities. There are six scenarios described in [9], which can be implemented in steps from a simple sign-on system for mobile subscribers to a fully seamless intersystem operation. Four of these scenarios allow subscribers of the integrated UMTS–WLAN network to access UMTS services, including IMS-based services. Thus, QoS control is required to handle QoS guaranteed services over the WLAN component of the integrated network. In this section, a policy-based QoS architecture is proposed for WLAN as a prelude to the QoS architecture for an integrated UMTS–WLAN system.

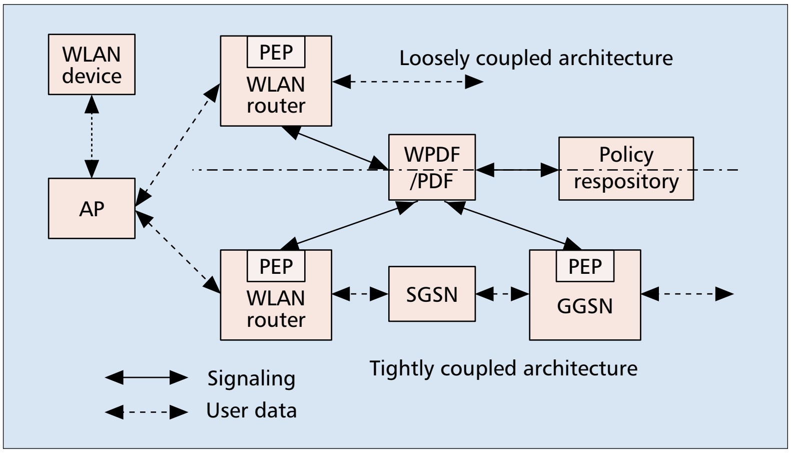

Figure 1. *Policy-based QoS architecture in WLAN.*

Typically, a WLAN network comprises a number of WLAN access points (APs) connected to a WLAN router (WR) that provides access to external networks. QoS mechanisms at the network layer in the form of IP packet header marking, policing and conditioning in the differentiated services (DiffServ) architecture, and at the data link layer in the form of WLAN QoS mechanisms (802.11e), media access control (MAC) frame priority indication (802.1p), and virtual LAN (VLAN) tagging (802.1q) are controllable by the QoS architecture. The MAC QoS mechanisms are provided by the WLAN APs, and the IP DiffServ mechanisms are available at the WR. In the following WLAN policy-based QoS architecture, only the DiffServ mechanisms in the WR are available to the policy control process.

There are two approaches [10] proposed for coupling WLAN networks with UMTS: tightly coupled architecture and loosely coupled architecture (Fig. 1). In a tightly coupled architecture, the WLAN network is connected to the UMTS network as an alternative radio access network. In other words, the WR is connected directly to the serving General Packet Radio Service (GPRS) support node (SGSN) and is treated by the SGSN as a radio network controller (RNC). The data sent by WLAN devices must go through the UMTS PS domain served by the connecting SGSN to reach its destination. To ensure seamless IP QoS services in tightly coupled WLAN–UMTS networks, the UMTS session control entities such as call state control functions (CSCFs) in UMTS IMS are extended into the WLAN network. The session control entities like the CSCFs interact directly with the WLAN devices as if they are normal UMTS user equipment (UE). Thus, a PDF can enforce the network-level policies at the WR directly as if the WLAN network is a part of the UMTS PS domain. The PDF can be an independent entity controlling the WLAN domain only, or an integrated part of the PDF in the UMTS PS domain. In a tightly coupled architecture, the WLAN is an alternative radio access network, so the 3GPP PDF is reused. There is no effect on the 3GPP access control and billing/charging entities.

In a loosely coupled architecture, the WLAN is connected to a gateway GPRS support node

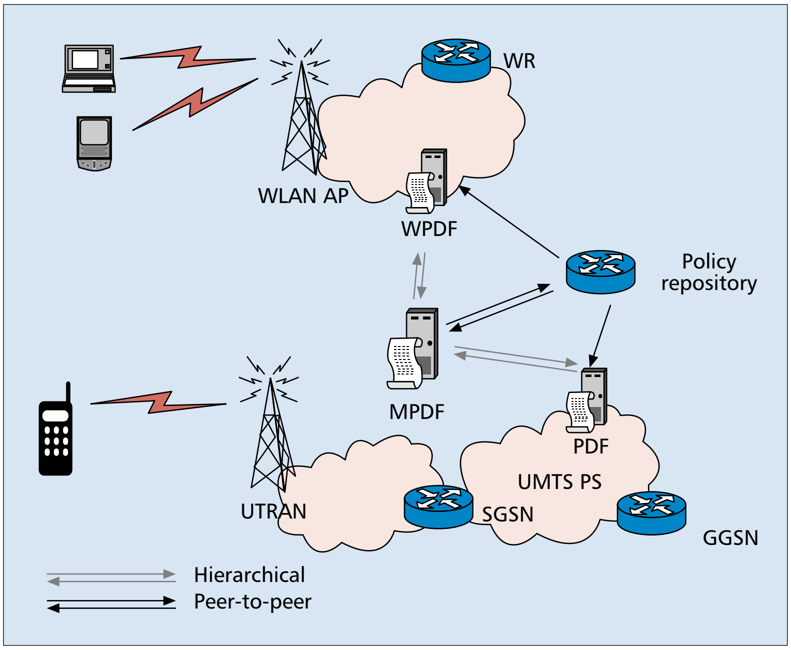

Figure 2. *Scenario 1: WLAN and UMTS networks controlled by one operator.*

(GGSN) of the UMTS network as a separate network. The WR is treated like a GGSN, and the WLAN network is considered a peer UMTS network. This article proposes that the WLAN constitutes a distinct policy domain with its own PDF, called a WLAN PDF (WPDF). The WPDF acts as the PDP in the WLAN domain, and the WR is the PEP that enforces the policy decisions made by the WPDF. In a loosely coupled architecture, the WLAN domain can use session control entities (CSCFs) in the UMTS network. Alternatively, a session control entity can be sited in the WLAN domain to interact with the CSCFs in the UMTS network. This WLAN session control entity is related to the WPDF in the same way as the relationship between the proxy (P)-CSCF and the PDF in UMTS IMS. By adopting an additional session control entity and PDF in the WLAN domain, the distinct interface between the WLAN domain and its interconnected UMTS network is preserved. How the policies implemented in the WLAN domain are related to the policies in the UMTS network is dependent on the interworking scenarios. Consequently, the interaction method between the WPDF and the PDF of the UMTS network is determined by this policy relationship. In a loosely coupled architecture, the WLAN domain behaves as a separate network. The proposed architecture has a standalone WPDF to perform service-level policy control for the WLAN domain, and thus is not expected to affect the access control and billing/charging reference model defined in [10, 11].

The loosely coupled WLAN architecture offers a major advantage over the tightly coupled architecture: integration flexibility offered by the distinct WLAN policy domain. It permits easy integration of the WLAN domain into a multi-operator multidomain environment as either a subordinate policy domain or a peer policy domain. This flexibility permits simple extension of the QoS policy control architecture in [8] into the WLAN domain regardless of the UMTS–WLAN integration scenarios.

Policy-Based QoS Management in an Integrated UMTS and WLAN Environment

Like UMTS networks, WLANs may operate in public, corporate, or residential environments. The different environments may involve different administrative domains and different degrees of network integration [10, 11]. For example, security capabilities and policies may differ between public, corporate, and residential WLANs. In this article a policy-based QoS management architecture extended from the policy control architecture presented in [6] is proposed to support end-to-end QoS in an integrated UMTS–WLAN environment.

Three scenarios are considered in the following subsections to illustrate the feasibility of the proposed architecture:

- One operator controls the UMTS network and WLANs.
- Different UMTS operators share a WLAN.
- An independent WLAN is interconnected to a UMTS operator's network.

These three scenarios are generic enough to cover all cases of an integrated UMTS–WLAN environment. In this article it is assumed that a WLAN and a public land mobile network (PLMN) are the smallest possible units of an administrative domain. In practice, this might not be true due to varying business arrangements. For instance, a café owner may own a hotspot provided within the premise of his/her café. A broadband service provider may coordinate multiple hotspot owners and resell the broadband service to end users. The service provider may also have a roaming agreement with a UMTS operator. In such a scenario, the policy rules may be more complicated than the scenarios presented here; however, the underlying policy architecture should be similar. In extreme cases, only simple variations of the proposed architecture are foreseen.

Scenario 1: UMTS and WLAN Under One Operator

Scenario 1 is a PLMN model where the operator installs and operates an integrated UMTS–WLAN network. The operator fully controls its WLAN sites, and provides normal telecommunications services in addition to WLAN access service. For the integrated UMTS–WLAN environment in a single operator's network, the hierarchical policy architecture is used (Fig. 2). The master policy controller (MPDF) connects to the policy controller of the WLAN network (WPDF) and the UMTS policy controller (PDF). The MPDF of the UMTS network serves as the master policy node of the WPDF so that the policies implemented in the WLAN domain are integrated into the operator's policy hierarchy. The MPDF translates the networkwide policies into domain-specific network-level policies on behalf of the WPDF and PDF, and stores them in the policy repository. The PDFs just retrieve these network-level policies from the repository, translate them into device-level policies, and install these policies in the network devices under their control. Note that these policies are enforced on all new IP multimedia sessions uniformly unless there are policy conflicts regarding the authorization of a new session's QoS requirements. In this case, the policy conflicts must be resolved through the MPDF. The communications protocol between the MPDF and the PDFs is based on the Common Open Policy Services (COPS) [12] protocol.

The policy control process activated during the setting up of an IP multimedia session is illustrated in Fig. 3 and described below:

1. During the session setup period, the session control entity of the WLAN domain will pass the QoS parameters in the session setup signaling to the WPDF.

2. The WPDF will retrieve the relevant network-level policies from the policy repository.

3. The WPDF checks that the requested QoS parameters are permitted by the network-level policies. If the QoS parameters are permitted, the WPDF notifies the WLAN session control entity that policy control has succeeded. If the QoS parameters are explicitly forbidden by the policies, the WPDF notifies the WLAN session control entity that policy control has failed.

4. If there is a conflict in the retrieved network-level policies regarding the authorization of the requested QoS parameters, the WPDF will ask the MPDF to resolve the policy conflict. The WPDF encapsulates the requested QoS parameters in a COPS request message and sends it to the MPDF.

5. The MPDF resolves the policy conflict by creating new network-level policies based on the networkwide policies. Before the new policies are supplied to the WPDF, the MPDF must validate them with the PDF since the UMTS domain is on the data path of the session and thus must be capable of implementing the requested QoS parameters as well.

6–7. Once the policy validation with the PDF succeeds, the MPDF sends the new network-level policies back to the WPDF. At the same time, the MPDF writes the new policies into the policy repository for future retrieval by the PDFs.

Note that a two-level hierarchy is shown in Fig. 2 for illustrative purposes only. The depth of the hierarchy depends on the relationship among the policies that are to be applied to the network. A complex policy relationship is usually represented as a multi-level policy hierarchy. The multi-domain architecture in the integrated UMTS — WLAN network is not restricted to the minimal two-level hierarchy described in the policy control process. The network operator may decide to provide intermediate levels of policy conflict resolution for more granular control over its network.

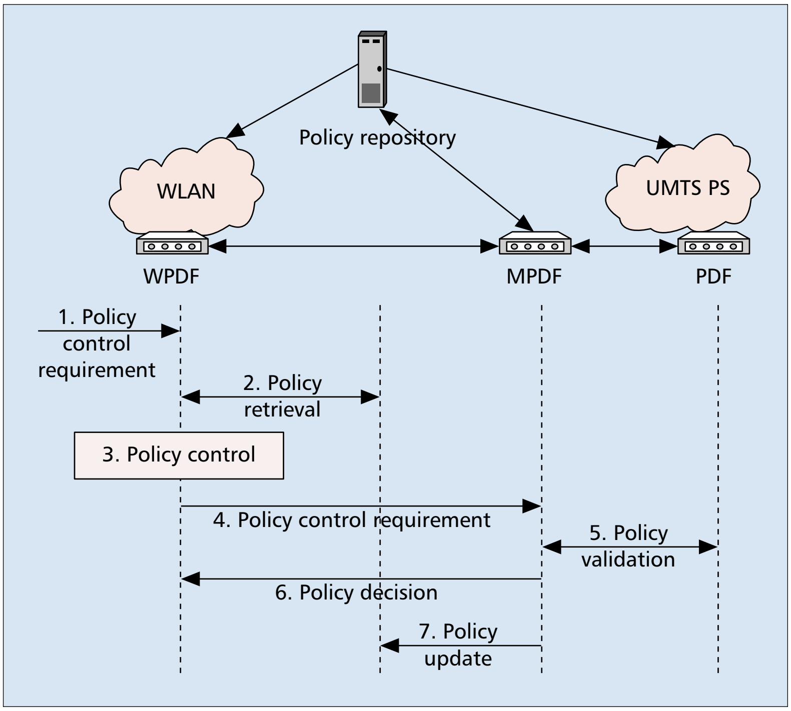

Figure 3. *A procedure for policy control between WLAN and UMTS domains under the control of a single operator.*

Scenario 2: A WLAN Network Shared by Multiple Operators

The normal range of a WLAN AP is less than 150 m in open space and just 50 m in closed environments like within buildings. To provide full WLAN coverage over a built-up area like a city requires deployment of a large number of WLAN APs. As a result, the cost of providing extensive WLAN services is high for a single operator, and it is desirable for operators to share their WLAN infrastructure to reduce upfront costs. Scenario 2 is such a risk-sharing PLMN model to provide different services in an integrated UMTS–WLAN environment. Multiple operators install and operate shared WLAN sites at different locations that are connected to their own UMTS networks. In this scenario, operator A may usually target business subscribers while operator B targets youth subscribers. The shared WLAN sites are configured to provide different amounts of resources at different times of the day. For example, operator A may be allocated more resources to provide business services during daytime, while operator B may get more resources to provide entertainment services after work hours. To provide end-to-end QoS in this environment, a WPDF is deployed in the shared WLAN domain. Scenario 2, shown in Fig. 4, illustrates how the WPDF interacts with the MPDFs of the cooperating operators' networks.

In scenario 2, the QoS policies to be applied in the WLAN domain are subjected to the control of operators A and B. For WLAN traffic going into operator A's network, the WPDF applies the policies supplied by operator A's MPDF. Likewise, WLAN traffic going into operator B's network is subjected to the policies supplied by operator B's MPDF. These policies specify how much resources in the WLAN access router (WR) should be provided to the traffic to be carried by the different operators' networks in order to satisfy the QoS requirements contracted by the operators. In terms of policy rela-

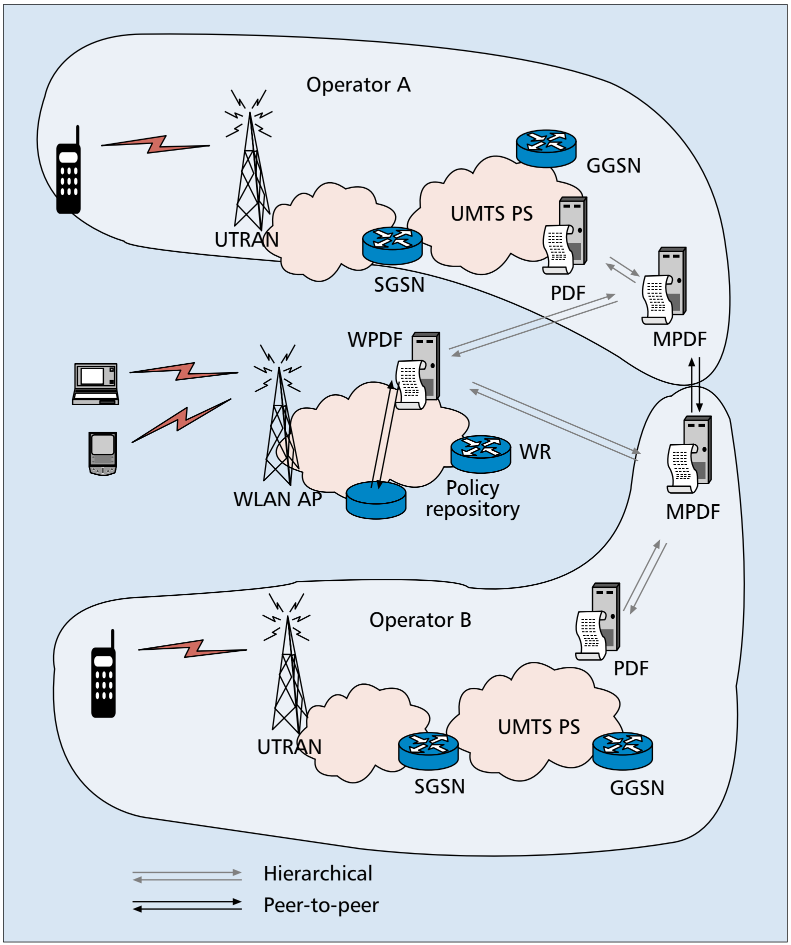

Figure 4. *Scenario 2: WLAN shared by different operators.*

tionship, the WPDF is a child node in the policy hierarchies of operators A's and B's networks, and it is serving two MPDFs that are peers.

In the operational lifetime of the WLAN domain, it is inevitable that the WPDF will encounter traffic conditions that require the application of conflicting policies provided by the different MPDFs. There are two possible ways to resolve the conflicting policies:

- The WPDF communicates information about the conflicting policies to the MPDFs. Then the MPDFs are solely responsible for negotiating new policies that will replace the old policies
- Additional policies known as overriding policies are preset in the WPDF and are used to resolve the conflicting policies. The overriding policies implement the agreement between the cooperating operators on how they share the WLAN resources.

Because the WLAN domain is a shared infrastructure of the operators, it is preferable that the policies that implement the WLAN resource sharing arrangement be solely enforced by the WPDF. This will free the MPDFs of the operators' networks from having to account for the presence of the shared WLAN infrastructure in their policies. In addition, new WLAN domains can easily be added without affecting the policies implemented by the individual MPDFs.

The WPDF of a shared WLAN domain interacts with its parent MPDFs as a node in their two-level policy hierarchies. The policy control process is similar to that illustrated in Fig. 3 except that overriding policies are applied in the WPDF when the policies provided by an MPDF conflict with existing policies governing the traffic destined for the other operator's network. Note that the WPDF maintains its own policy repository. The MPDFs do not have any read or write access to the WLAN policy repository. The policy control process between the WPDF and an MPDF is illustrated in Fig. 5 and described below:

1. During the session setup period, the session control entity of the WLAN domain passes the QoS parameters in the session setup signaling to the WPDF.

2. The WPDF retrieves the relevant network-level policies from the WLAN policy repository.

3. The WPDF checks that the requested QoS parameters are permitted by the network-level policies. If the QoS parameters are permitted, the WPDF notifies the WLAN session control entity that policy control has succeeded. If the QoS parameters are forbidden by the policies, the WPDF notifies the WLAN session control entity that policy control has failed.

4–5. If the retrieved network-level policies are conflicting on the authorization of the requested QoS parameters, the WPDF will ask the MPDF to resolve the policy conflict. The WPDF generates the requested QoS parameters first. Before the WPDF sends out the requested QoS parameters to the MPDF, it must verify that they do not conflict with policies installed by the MPDFs of other cooperating operators. If there is no conflict, the requested QoS parameters are encapsulated in a COPS request message and sent to the MPDF. If there are conflicts between the policies installed by different MPDFs, the WPDF must apply the overriding policies to resolve the conflicts. These overriding policies will modify the requested QoS parameters to conform to the network sharing arrangement that has been agreed for the WLAN domain while ensuring that the requested QoS parameters are satisfied. If the policy modification cannot be performed successfully, the WPDF will reject the QoS request from the session control entity.

6. The MPDF resolves the policy conflict by creating new network-level policies based on its networkwide policies. Before the new policies are supplied to the WPDF, the MPDF must validate them with the PDF of the UMTS domain that is on the data path of the session, which must be capable of implementing the requested QoS parameters.

7–8. Once the policy validation with the PDF succeeds, the MPDF sends the new network-level policies back to the WPDF. At the same time, the MPDF writes the new policies into the UMTS policy repository for future retrieval by the PDFs in the UMTS domain.

9. The WPDF applies the new network-level policies provided by the MPDF and simultaneously updates its policy repository for future retrieval.

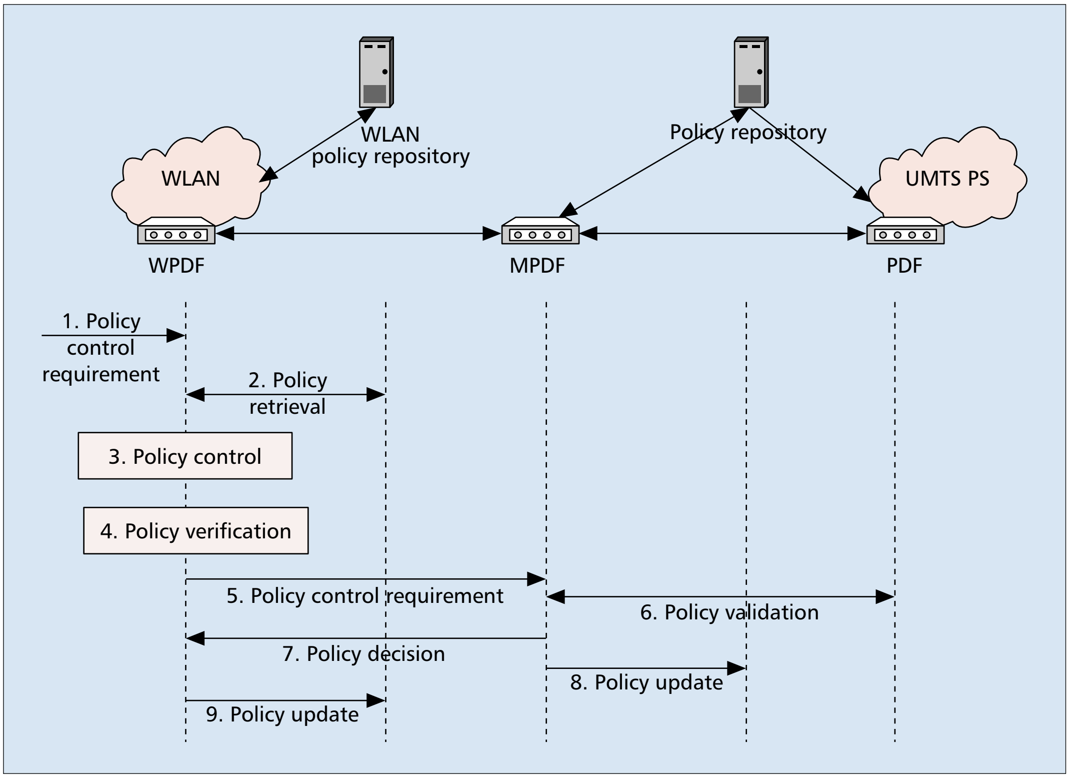

■ **Figure 5.** *A procedure for policy control between a shared WLAN domain and its parent UMTS domain.*

Scenario 3: Customer's WLAN Network Interconnected to an Operator's UMTS Network

This scenario is shown in Fig. 6. Here, the WLAN may belong to an independent Internet service provider (ISP) or an enterprise that is a customer of the UMTS operator. This model allows the UMTS operator to provide wide-area mobile services to customers that have their own WLAN infrastructure. In contrast to scenario 2, the WPDF in the WLAN domain is a peer of the MPDF in the UMTS network. The WPDF has the sole right to update its policy repository. The network-level policies to be employed by interconnecting the UMTS network and the WLAN network are determined by the service level specifications (SLSs) agreed between the peering WLAN and UMTS operators. In these SLSs, there are static and dynamic service requirements. The static service requirements can be directly translated into enforceable network-level policies to be retrieved by the WLAN's PDF and the UMTS' PDFs in the respective networks. However, the dynamic service requirements are dependent on the state of the UMTS/WLAN network (e.g., its resource utilization), and can only be translated into enforceable network-level policies after negotiating with the connecting networks. The purpose of SLS negotiation is to enable the interconnected networks' interdomain policy agents (IPAs) to agree on the specific service requirements that must be supported under the prevailing network states. Once the SLS negotiation is successfully completed, the participating IPAs can translate the agreed on service requirements into enforceable policies in their respective networks. Note that this runtime negotiation may not be initiated on a per-session basis. Instead, SLS negotiation is usually initiated when the IPA detects that the state of its network has changed and the existing policies are no longer enforceable.

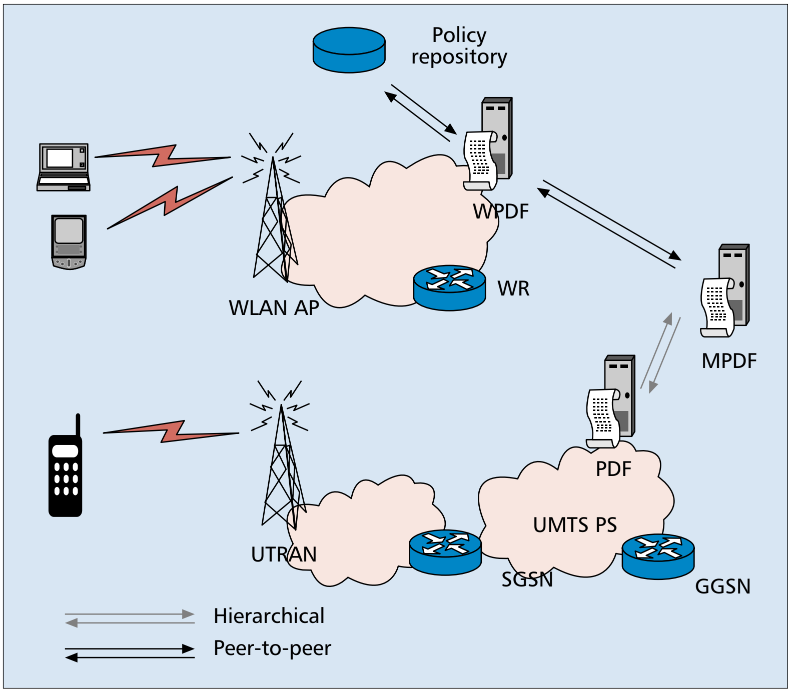

■ **Figure 6.** *Scenario 3: Interworking of a customer's WLAN network and an operator's UMTS network.*

IPAs participating in the SLS negotiation must be interconnected so that SLS information can be exchanged. We propose to use the COPS [13]

IPAs participating in the SLS negotiation must be interconnected so that SLS information can be exchanged. We propose to use the COPS protocol as the communications protocol between peering IPAs (IPA1, IPA2 and IPA3).

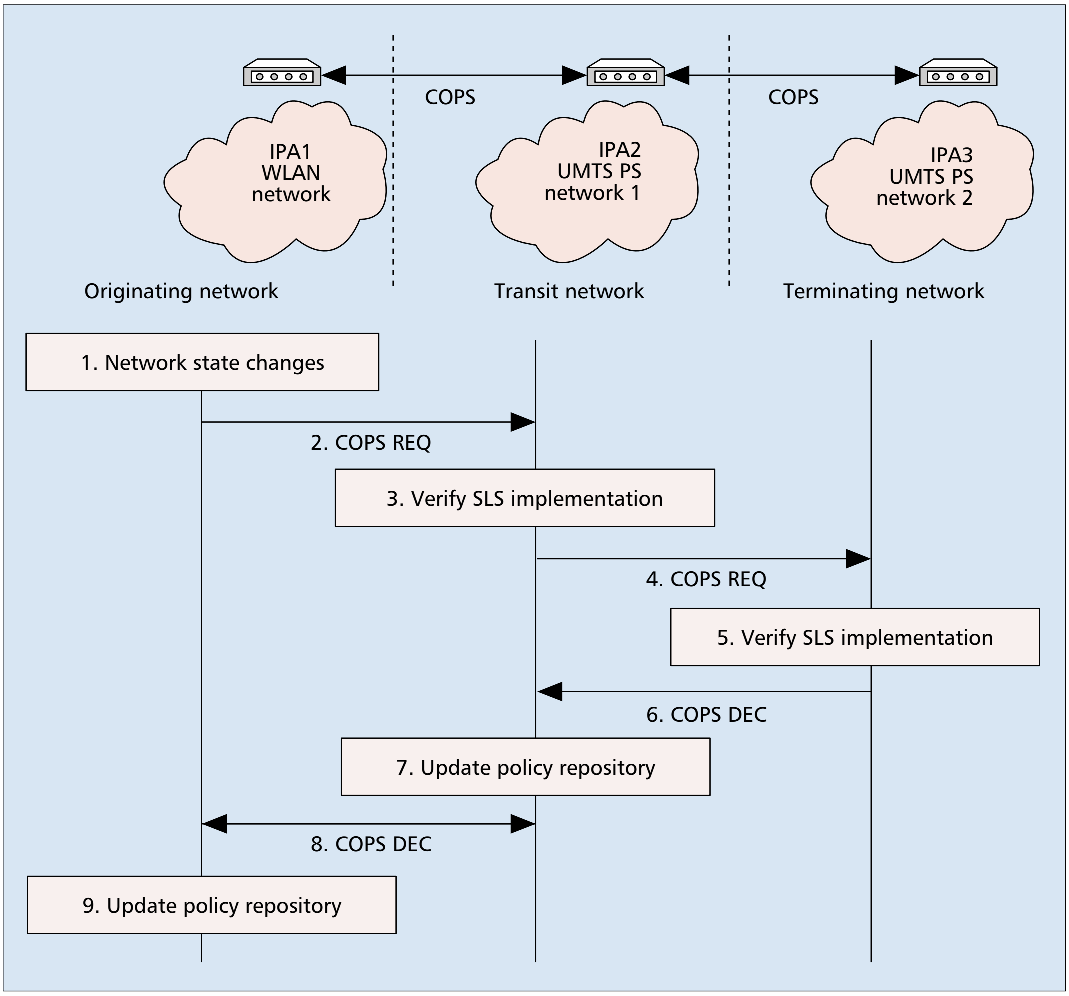

Figure 7. *Policy negotiation between UMTS network and WLAN.*

protocol as the communications protocol between peering IPAs (IPA1, IPA2 and IPA3) (Fig. 7). The UMTS/WLAN operator will configure its IPA with the locations of its peering counterparts. The COPS protocol can be suitably extended with new messages to carry SLS information, as is attempted in the COPS-SLS [13] protocol.

The SLS negotiation process between the IPAs is depicted in Fig. 7:

1. The WPDF detects a change in the network state of the WLAN that invalidates the current network-level policies implementing the dynamic QoS service requirements in the SLS. The WPDF updates the SLS dynamic QoS parameters based on the new network state before translating it into network-level policies.

2. Before the policy repository is updated with the new policies, IPA1 encapsulates the updated SLS parameters in a COPS request message and sends it to IPA2 of the connecting UMTS PS network 1.

3. Once IPA2 receives the SLS information in the COPS request message from its peer IPA1, its MPDF will check whether its current network-level policies can implement the updated SLS parameters requested by UMTS PS network 1. If the current network-level policies can implement the SLS parameters, IPA2 just returns a positive COPS decision message to IPA1. The SLS negotiation terminates at this point because it is assumed that UMTS PS network 1 is guaranteeing that UMTS PS network 2 can support the new QoS service requirements. If the updated SLS parameters cannot be supported, the MPDF of UMTS PS network 1 will translate the updated SLS into network-level policies for verification purposes. If the translation cannot be made, IPA2 will return a negative COPS decision message to IPA1. Otherwise, IPA2 will have to change its network-level policies to meet the new QoS service requirements and consult UMTS PS network 2 about its ability to meet the new QoS service requirements.

4. Before the new network-level policies can be written into the policy repository, IPA2 must check with IPA3 whether UMTS PS network 2 can support the new QoS service requirements by forwarding the COPS request message.

5. IPA3 repeats the SLS implementation verification procedure in step 3.

6. IPA3 returns a positive COPS decision message if the updated SLS parameters can be implemented by current or new network-level policies; it returns a negative COPS decision message otherwise.

7–9. The reception of a COPS decision message from IPA3 will trigger IPA2 to send a matching COPS decision message to IPA1. Thus, IPA1 gets a positive COPS decision message if IPA3 accepts the updated SLS parameters. Once the IPAs receive a positive COPS decision message, they will write the new network-level policies translated from the updated SLS into their policy repositories. If the IPAs receive a negative COPS decision message, the updated SLS cannot be supported, and no update is made to their policy repositories. In that case, the WPDF can modify the SLS dynamic QoS parameters and repeat the SLS negotiation process. The operator of the originating WLAN determines the upper limit on the number of SLS negotiation rounds that is permitted.

Conclusion and Open Issues

A policy-based QoS architecture in the current UMTS PS domain has been described. A simple policy-based architecture to control the QoS mechanisms in a WLAN network has been proposed, along with different ways of integrating the policy-based QoS architectures of the UMTS and WLAN domains for these different interworking scenarios. In these scenarios, the policy architecture deployed ranges from a pure hierarchical architecture (scenario 1) through a mixed peering/hierarchical architecture (scenario 2) to a pure peering architecture (scenario 3). These proposed QoS policy architectures can minimize session setup delay and policy exchange load while maximizing network scalability for the different interworking scenarios. The SLS negotiation and policy conflict resolution mechanisms, which constitute the basic interactions in these QoS policy architectures, have been described.

Several problems, which are the foci of our current and future work, are foreseen in the proposed multidomain QoS policy architecture. These are:

1. To facilitate successful negotiation between IPAs, the parameters of SLS must be standardized to provide the basis for negotiation. This requires analysis of the format of QoS requirements that may be specified in the SLS so that the definition of the QoS resource elements carried in the COPS messages can be determined.

2. The security of the communications channel between the interconnected policy entities is important. This is especially true for the peering architecture used to interconnect different operators' networks. Operators are highly sensitive to the risk of policy leakages through snooping, unauthorized tampering with COPS messages en route between IPAs, and interactions with unauthenticated policy entities. Although the COPS protocol has the ability to secure messages by encapsulating integrity objects, additional mechanisms may have to be deployed to address other security risks.

3. Policy negotiation in a peering architecture is a slow process, especially if the chain of participating networks is long. Proper network design in this case will play an important role in minimizing the delay.

4. The depth of the hierarchical architecture affects the policy provisioning time in an operator's network. PDFs at the lowest level will have to wait longer for a decision from the MPDF when the policy hierarchy is deep. Proper policy hierarchy design will help to reduce the number of policy levels.

Policy negotiation in a peering architecture is a slow process, especially if the chain of participating networks is long. Proper network design in this case will play an important role in minimizing the delay.

References

[1] J. Blau, "Wi-Fi Hotspot Networks Sprout Like Mushrooms," *IEEE Spectrum*, Sept. 2002, pp. 18–20.

[2] T. Bostrom, T. Goldbeck-Lowe, and R. Keller, "Ericsson Mobile Operator WLAN Solution," *Ericsson Rev.*, no. 1, 2002, pp. 36–43.

[3] W. Zhuang et al., "Policy based QoS Architecture in IP Multimedia Subsystem of UMTS," *IEEE Network*, vol. 17, 2003, pp. 51–57.

[4] 3GPP TS 23.228 (v. 5.2.0), "IP Multimedia Subsystem-Stage 2 (Rel. 5)," Oct. 2001.

[5] 3GPP TR 23.917 (v. 0.4.0), "Dynamic Policy Control Enhancements for End-to-end QoS (Rel. 6)," Dec. 2002.

[6] R. Yavatkar, D. Pendarakis, and R. Guerin, "A Framework for Policy-Based Admission Control," RFC 2753, Jan. 2000.

[7] B. Moore *et al.*, "Policy Core Information Model — Version 1 Specifications," RFC 3060, Feb. 2001.

[8] W. Zhuang *et al.*, "Multi-Domain Policy Architecture for IP Multimedia Subsystem in UMTS," *Proc. IFIP/IEEE Net-Con 2002 — Network Control and Eng. for QoS, Security and Mobility with Focus on Policy-Based Net.*, Oct. 2002, pp. 27–38.

[9] 3GPP TR 22.934 (version 6.0.0), "Feasibility Study on 3GPP System to Wireless Local Area Network (WLAN) Interworking (Rel. 6)," May 2002.

[10]. ETSI, "Requirements and Architectures for Inter-working between HIPERLAN/2 and 3rd Generation Cellular Systems," Tech. rep. TR 101 957, Aug. 2001.

[11]. 3GPP TS 23.234 (v. 1.0.0), "WLAN Subsystem — System Description (Rel. 6)," Sept. 2002.

[12]. T. M. T. Nguyen *et al.*, "COPS Usage for SLS Negotiation (COPS-SLS)," work in progress, IETF, Feb. 2002.

[13]. D. Durham *et al.*, "The COPS (Common Open Policy Service) Protocol," RFC 2748, Jan. 2000.

Biographies

WEI ZHUANG [M'97] (zhuangw@ieee.org) received a B.E. degree in electronic engineering from Huazhong University of Science and Technology, Wuhan, China, in 1991, an M.E. degree in communications and electronic systems from Shanghai Jiaotong University, China, in 1994, and a Ph.D. degree from the National University of Singapore in 2000. From 1999 to 2001 he was an executive engineer in the WCDMA group of Fujitsu Singapore Pte. Ltd. In August 2001 he joined Siemens Singapore, where he is currently a senior R&D engineer. His research interests are mainly focused on, but not limited to, 3G WCDMA networks, policy network, and end-to-end QoS control in wired and wireless networks.

YUNG-SZE GAN graduated from the National University of Singapore with Bachelor and Master of engineering degrees in 1998 and 2002, respectively. He joined Siemens Pte Ltd., Mobile Core Research and Development Department, Singapore in 2001 as an R&D engineer. His work is largely concentrated in the UMTS IP multimedia subsystem with some interests in the areas of policy-based network management and multicast networks.

KOK JENG LOH is a technical manager in the ICM Mobile Core R&D Department, Siemens Singapore. He received B.Eng. and M.Eng. degrees from the National University of Singapore in 1995 and 2000, respectively. His current research interests include call state control in UMTS IMS and policy-based QoS in IP-based networks.

KEE-CHAING (K-C) CHUA [M'87] (chuakc@nus.edu.sg) received his Ph.D. degree in electrical engineering from the University of Auckland, New Zealand, in 1990. Following this, he joined the Department of Electrical Engineering at the National University of Singapore as a lecturer, becoming a senior lecturer in 1995 and an associate professor in 1999. From 1995 to 2000 he was seconded to be deputy director of the Center for Wireless Communications (now Institute for Infocomm Research), a national telecommunications R&D institute funded by the Singapore Agency for Science, Technology and Research. From 2001 to 2003 he was on leave of absence from NUS to work at Siemens Singapore where he was the founding head of the ICM Mobile Core R&D department. He has carried out research in various areas of communication networks. His current interests are in ensuring end-to-end QoS in both wired and wireless IP based networks. He is a recipient of an IEEE Third Millennium medal.