An Adaptive Service Management Framework for Wireless Networks

Antonis M. Hadjiantonis, Marinos Charalambides,
and George Pavlou, *University of Surrey*

Abstract: *Nowadays, services targeting home and business users, such as Internet access, digital television or online entertainment are taken for granted. However, our lifestyle creates the need for extending the reach of such services but also the need for new ones, targeted to people on the move. With complex wireless networks forming around us and a new range of services becoming available, an enhanced service management paradigm is required. This work introduces a novel framework for the management of services in wireless environments, considering the diverse requirements and conditions. Using a wireless Media Service as a case study, simulation results indicate that our policy-based approach offers a flexible framework where both business and user expectations are met. User preferences guide service customization while statistical processing of user requests directs service adaptation.*

Keywords: *policy-based management, adaptive framework, customization, wireless services, media delivery*

Digital Object Identifier 10.1109/MVT.2008.915322

1. Introduction

Our need for accessing any kind of information is constant. Access to corporate networks, e-mails or simply entertainment are new necessities posed on an increasingly networked wireless world. In the era of mobility and connectivity, a multitude of wireless devices surround us in our everyday life. Wireless digital assistants such as mobile phones, laptops or PDAs must be able to cope and offer the desired services at anyplace and at anytime.

Inevitably, a framework for Service Management is required, taking into account the diverse conditions and requirements of wireless networks. Specialized solutions that can adapt to changes fast and reliably, while ensuring security over the wireless interfaces are needed. One has to consider that wireless

1556-6072/07/$25.00©2007IEEE

networked devices are not under the strict control of a Network Operator (NO) as in traditional infrastructure-based networks e.g., fixed IP, cellular. A critical requirement is to respect the owner relationship between the users and the devices. Individual users are reluctant to entrust the command of their devices to the operator and demand more control. These issues motivate our research efforts, with the goal of providing an adaptive Service Management framework where user preferences are respected and the consumed services are tailored to their needs. Industrial predictions mention that "extending the service portfolio is one of the best options for growth" [1], fueling more research interest in novel solutions for mobile users.

Policy-based management (PBM) has been traditionally used for the management of large-scale complex IP networks [2], [3]. Policies capture the high-level management objectives and are automatically enforced to devices, thus simplifying and automating compound and time-consuming configuration tasks. Service Management is becoming increasingly complex, and research efforts point out that it could benefit from a policy-based approach [4]. By adopting a PBM paradigm and extending a framework for wireless ad hoc networks, we propose a novel Service Management framework. Our framework incorporates features such as full service customization based on user preferences and adaptive behavior of services based on statistical observations. Policies orchestrate the operations of the Service Adaptation Logic (SEAL) entity that encapsulates the core functionality of our framework. As proof of concept, a detailed specification of a Media Service illustrates our proposal, while simulation of service adaptation demonstrates the benefits of such an approach.

The rest of the paper is organized as follows. After a brief background overview, Section 3 introduces the proposed Service Management framework, detailing the architecture and the SEAL entity. In Section 4, we present a detailed case study based on the realization of a Media Service. Section 5 presents the simulation results of the Media Service adaptation process.

2. Background

We choose to build our proposed policy-based service management architecture on an existing hybrid organizational model for Mobile Ad Hoc Networks (MANETs) [5]. This model was chosen due to its inherent policy-based attributes and the ability to anticipate the dynamic and unpredictable condition changes of wireless networks.

By combining the benefits of hierarchical and distributed management schemes, the model offers the desired properties of policy-based management through Policy Management Tools (PMTs), distributed decision making by cooperating Policy Decision Points (PDPs) and distributed policy storage. An algorithmic process organizes the wireless network in clusters, where assigned Cluster Heads (CHs) perform local management tasks. The rest of the nodes become Cluster Nodes (CNs), register to their nearest CH and remain under its supervision. The correspondence of physical devices to roles depends on context-aware metrics (e.g., current mobility, device capabilities); generally speaking, lightweight devices like cell phones become CNs while more powerful devices like laptops or access points can become CHs. Depending on the formation purpose of the wireless network and the business rules of the Network Operator, one or more privileged nodes are assigned the Manager Node (MN) role. Together the MNs and the CHs constitute the *hypercluster*, which performs the management tasks in a distributed and cooperative manner.

A policy-based framework had been integrated in the aforementioned model, adding the desired self-managing capability and controlled programmability to ubiquitous wireless networks [5], [6]. A critical issue for every PBM system is the analysis of policies in order to detect and resolve conflicts and inconsistencies [7], [6]. In order to employ a policy-based paradigm to wireless networks and service management, one has to consider that the networked devices are not under the strict control of the Network Operator (NO). The individual users are reluctant to entrust the management of their devices to a central authority and demand more control over their owned devices [8]. In [6] we have introduced the differentiation of managed objects into *Policy Free Objects (PFO)* and *Policy Conforming Objects (PCO)*, thus offering the ability to users to define the desired access rights to their preferences.

3. Framework

In order to manage a complex set of services and offer the expected Quality of Service (QoS) to users, a Service Provider (SP) has to take into account several parameters

WE CHOOSE TO BUILD OUR PROPOSED POLICY-BASED SERVICE MANAGEMENT ARCHITECTURE ON AN EXISTING HYBRID ORGANIZATIONAL MODEL FOR MOBILE AD HOC NETWORKS DUE TO ITS POLICY-BASED ATTRIBUTES AND ABILITY TO ANTICIPATE THE DYNAMIC CHANGES OF WIRELESS NETWORK.

and constraints. But for a service to be successful, a certain degree of control must be given to the end-user. Preferences offer some control to users and allow for the customization of available services. A user's preferences may express general device settings or access to integrated hardware (e.g., power profile, GPS receiver). We refer to these as *basic preferences* so as to differentiate from *service-specific preferences*. The latter refer to user options aiming to customize a specific service. Both preferences and device capabilities affect the adaptation process of deployed services.

3.1 System Architecture

The architecture presented in Figure 1 is based on the aforementioned Policy-Based Network Management system, which is extended and customized by introducing the Service Adaptation Logic (SEAL) and User Preferences Control (UPRC) components. The novel features introduced, together with detailed policy design, facilitate a flexible and extensible Service Management framework.

The Service Adaptation Logic (SEAL) component accepts users' requests and provides a customizable and adaptive service management framework by taking into account device capabilities and service-specific preferences. SEAL interacts with the User Preferences Control (UPRC) on a user's device, aiming both at the enhancement of users' experience and the optimization of offered services.

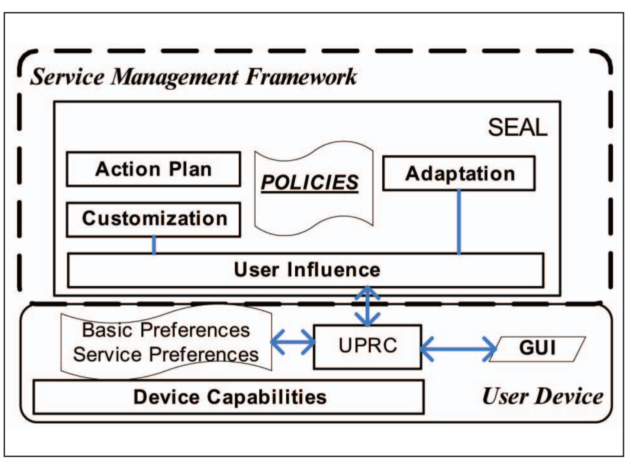

FIGURE 1 System Architecture.

3.2 Service Adaptation Logic

The Service Adaptation Logic (SEAL) component is a network-side entity responsible on one hand for adapting offered services according to specific user's preferences and on the other hand for influencing these preferences in order to optimize service utilization. These tasks are policy-driven, enabling a flexible and extendable service creation and execution environment. The detailed case study scenario in the next section demonstrates the functionality of SEAL, while simulation results follow.

The tasks of service customization and adaptation are directed by user service requests. Each request contains necessary information for the operation of the service, such as device capabilities and service-specific preferences. Before a service is offered, SEAL performs a three-level customization procedure. The first level is based on the requesting device capabilities. In addition, two extra levels of customization are introduced, which depend on the users' preferences, differentiating between basic and service-specific preferences. These parameters are examined by relevant policies and result in device and service-specific configuration. With the aim of service provisioning optimization, SEAL may attempt to influence user's preferences. This task can be executed directly by the Service Provider (*proactive influence*) or can be triggered during the Service customization task (*reactive influence*). The latter refers to the notification of a user during service initiation with the purpose of improving the requested service. The user is informed about the improvements and prerequisites, i.e., which preferences should be changed to allow the SP to offer the improved service.

While users' preferences need to be respected at all times, a Service Provider may need to proactively influence them for certain services to operate smoothly. For example, a file sharing service cannot operate, if all users choose not to share any files in their sharing preferences. In these cases, the SP needs to influence users (proactive influence) to change their preferences [9].

Service adaptation can be achieved by statistical analysis of the service-specific users' preferences and device capabilities. By analyzing these data, SEAL may identify current trends in service requests and profile the capabilities of users' devices. Based on the extracted information, SEAL dynamically changes the provisioned service aiming to satisfy more users' requests with less overhead. Section 5 demonstrates the adaptation process by simulation, based on the case study presented in the following Section.

On the client-side, the User Preferences Control (UPRC) entity communicates with SEAL, in order to visually notify the user and handle necessary device configuration changes. This lightweight entity manages all preferences and based on user input replies to the influence notifications from SEAL.

4. Case Study

To demonstrate the introduced ideas we present a detailed case study and elaborate on policy definitions implementing the service management framework.

4.1 Media Service Scenario

The increasing popularity of music downloads and video sharing activities among internet and mobile users has motivated our selection of a Media Service for experimentation. In this scenario, as depicted in Figure 2, users can have access to media services (audio, video, picture) while traveling on trains. This scenario is particularly attractive in the case of Underground Train networks, where user connectivity is limited.

A Network Operator (NO) deploys Cluster Heads (CHs) onboard trains and offers the infrastructure to different Service Providers (SPs). A multiple manager (MNs) environment is possible, where policies orchestrate manager interaction [6]. CHs are wireless access points with processing and caching capabilities. Depending on physical dimensions and passenger density, each train carriage can be considered as a separate cluster managed by a Cluster Head (CH). Economic considerations affect the hardware specification of CHs where trade-offs between cost and user coverage need to be made.

The CHs are interconnected (forming the *hypercluster*) and share a common media database which is physically located in the middle of the train. CHs also interact with the Manager Node (MN) controlled by the SP to update policies and report critical events. Users are able to request media items available on this database as well as items shared by other users. All service requests are made to the CH and the latter maintains a list of all available media items either on the network-wide Media Database (MD) or the cluster-wide Shared Media (SM) table. Apart from identification keywords and source location, this list describes items in terms of media/content type, quality and operational requirements.

To access the media service a user presents the CH with a request of the following format:

The above request consists of three main attributes: the *device capabilities*, user *media* and *basic preferences*,

mediaRequest(userID,
devCabs[codecs, freeSpc],
mediaPrefs[type{audio, video, picture, any},
quality{high, medium, low, any},
content{news,sports,entertain, any},
source{MD, SM}],
basicPrefs[connect{bluetooth, wifi},
battStrgy{norm, pwrSave}]);

WHILE USERS' PREFERENCES NEED TO BE RESPECTED AT ALL TIMES, A SERVICE PROVIDER MAY NEED TO PROACTIVELY INFLUENCE THEM FOR CERTAIN SERVICES TO OPERATE SMOOTHLY—FILE SHARING SERVICES REQUIRE FILES TO BE SHARED, FOR EXAMPLE.

reflecting the three-level customization process (Figure 3). For this case study, the procedure can be viewed as a filtering process on matching media items where policies are used to guide the selection decisions of the CH.

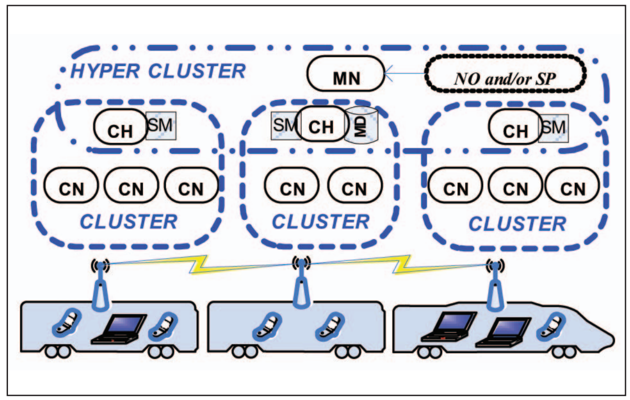

FIGURE 2 Case study scenario.

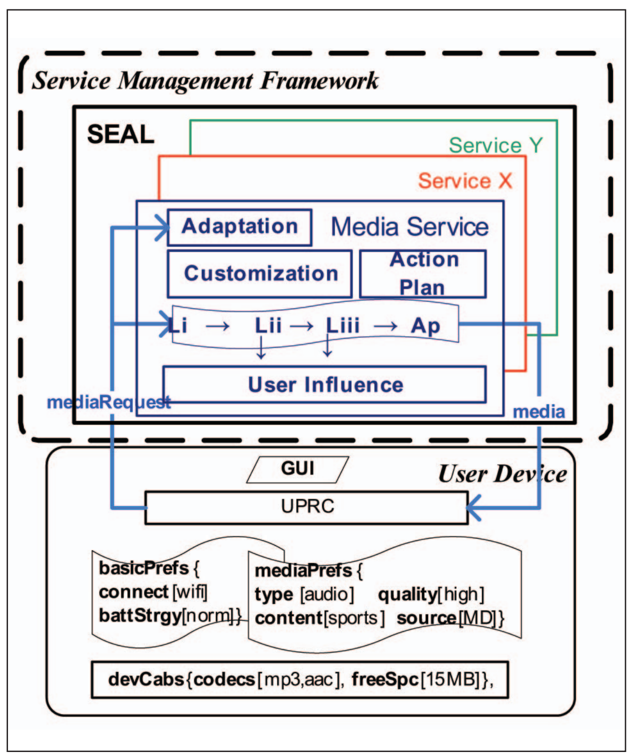

FIGURE 3 Media Service and Service Adaptation Logic (SEAL).

POLICIES GUIDING THE CUSTOMIZATION BASED ON USER'S REQUESTED MEDIA SERVICE PREFERENCES AIM TO DETERMINE MEDIA ITEMS THAT FIT TO THE QUALITY AND SOURCE PREFERENCES.

Media and *basic preferences* are optional and depend on the user's demand for personalized media delivery. In the subsequent sections we describe these policies, along with their specification and usage.

4.2 Policy-Based Service Customization and Adaptation

Once a request is received by a Cluster Head, a three level customization process is initiated aiming to satisfy the user's request. It should be noted that the service is able to recommend changes to user preferences in order to provide alternative media when current settings fail to return any results. The process allows a fully customized and tailored Media Service delivery to the user. At the same time, the service adapts to the users' demands.

TABLE 1 ***First customization level policies.***

Name	Policy
LiP1	**if** supportCodecs(devCabs[codec,_], mediaList[name,codec,_]) **then** selectCodec(mediaList[name, _])
LiP2	**if** supportSize(devCabs[_, freeSpc], mediaList[name, size, _]) **and** selectedCodec(mediaList[name, codec, _]) **then** selectItem(mediaList[name, _])
LiP3	**if** source(mediaList[name, _]) = = MD **and** mediaType(mediaList[name, _]) == (audio **or** video) **then** selectStream(mediaList[name, _]) **and** selectItem(mediaList[name, _])

TABLE 2 ***Second customization level policies.***

Name	Policy
LiiP1	**if** supportQuality(mediaPrefs[_, quality, _], mediaList[name, quality, _]) **and** supportSource(mediaPrefs[_, source], mediaList[name, source, _]) **then** selectItem(mediaList[name, _])
LiiP2	**if** usrFlag(mediaPrefs, not_informed) **then** informUsr(options[]) **and** setUsrFlag(mediaPrefs, informed)
LiiP3	**if** timeout = = FALSE **then** chkServPrefs(quality, source)

4.2.1 Capabilities and Preferences Customization

Initially, the CH searches the Media Database and Shared Media list for media items matching the criteria by keyword, content type and media type. Besides the usual media selection based on device capabilities (Li), two additional levels of customization (Lii,Liii) are introduced, which make use of the basic and service user preferences respectively. The initial generated list (*mediaList*) contains all matching media along with their metadata and triggers the first level of customization according to the requesting device capabilities. Three sequential policies (Table 1) apply here, aiming to determine media items on the generated *mediaList* with matching codecs and free memory space. These policies are initiated (triggered) by a *chkDevCabs* event, signaling the first filtering level.

Policies LiP1 and LiP2 check for media in the list that match the supported codecs of the user device and additionally satisfy free space requirements. Policy LiP3 applies only to audio and video media (event: *chkStream(mediaList[name, _])* and is triggered only if there is a match for codecs but the available space does not satisfy the requirements of that media.

The output of the first filtering level is an updated *mediaList* of items matching the requesting user's device capabilities. It should be noted that this list also includes items that the user cannot download because of limited free space and are marked as possible streaming media.

Policies guiding the customization based on user's requested media service preferences are shown in Table 2. These policies aim to determine media items that fit to the quality and source preferences. If a match is not found, then the user is informed to change the media preferences so as to result in alternative options.

Policy LiiP1 will be invoked at the second customization level, with triggering event *chkServPrefs*. Its action selects media items from the list, if matching quality and source are found. If this is not the case (event *noMatch ((quality, source), mediaList[_]))*, then policy LiiP2 notifies the user about failing to match his/her media service preferences and suggests changes to these preferences aiming to provide alternatives. The next policy (LiiP3) processes the user's reply (event *usrReply (mediaPrefs[quality, source])*) and checks the new preferences. Note that the action of this policy acts as a trigger for the first (LiiP1) indicating that the process starts again with alternative user preferences. The condition of the second policy (LiiP2) checks if the user has already been informed once, so as to avoid looping when he/she does not change any preferences or the notification expires.

Similarly, a third customization level aims to satisfy the basic user preferences. The policies of Table 3 select a media item if matching connectivity between the user and the media source is found and notify the user about failing to match his/hers connectivity preferences. For example, when a WiFi user requests media found on another user who uses only Bluetooth connectivity, then the system suggests a change to the first user's connection preferences to allow him/her to receive the desired media. The initiating event for this customization level is *chkBasicPrefs(connection)*.

TABLE 3 *Third customization level policies.*

Name	Policy	
LiiiP1	**if**	supportConnect(mediaPrefs[_, connection, _], mediaList[name,connection, _])
	then	selectItem(mediaList[name, _])
LiiiP2	**if**	usrFlag(connection, not_informed)
	then	informUsr(options[]) **and** setUsrFlag(connection, informed)
LiiiP3	**if**	timeout = = FALSE
	then	chkBasicPrefs(connection)

TABLE 4 *Action Plan policies.*

Name	Policy	
ApP1	**if**	streamSelected(mediaList[*name*,_])==TRUE
	then	setupStream(mediaList[*name*,_]) **and** streamTo(*userID*)
ApP2	**if**	streamSelected(mediaList[*name*,_])==FALSE **and** source(mediaList[*name*, _]) = = MD
	then	setupFileTransfer(mediaList[*name*,_]) **and** downloadTo(*userID*)
ApP3	**if**	streamSelected(mediaList[*name*,_])==FALSE **and** source(mediaList[*name*, _]) = = SM
	then	setupFileTransfer(mediaList[*name*,_], *sourceUserID*) **and** downloadTo(*userID*)

4.2.2 Action Plan

After the customization process, the user will be presented with a list of media items to choose from. The user's reply will serve as the trigger for the action plan policies (event is *userSelect(medialist[name,_], userID)*). Based on these policies (Table 4), the Cluster Head decides whether to stream the selected media to the user, provided the first customization level had marked that media for streaming. Otherwise, depending on its source, the media is downloaded on the user's device from the CH's database or from another cluster user (*sourceUserID*).

For clarity, the above policies are simple; however the Service Provider has the ability to change the action plan by editing existing policies or introducing new ones, taking into account more parameters or operational conditions. For example, ApP1 could include conditions like link quality or utilization between the user and the Cluster Head, in order to avoid significant packet losses that would degrade streaming media quality [6]. In addition, as technology evolves, the option of P2P streaming media between users can be easily integrated to the PBM system with the introduction of a few new policies, instead of fully upgrading the Media Service software.

4.2.3 Service Adaptation

An important task of SEAL is to adapt existing Services according to statistical analysis of users' prevailing service-specific preferences and device capabilities. This adaptation improves both service performance as well as users' experience. For this case study and the simulation presented in the next section, SEAL monitors the media-specific preferences for requested quality and device capabilities for codecs availability. By calculating the Weighted Moving Averages (WMA) of certain request parameters, SEAL can identify the trends in media requests and device capabilities among the served users. Using the flexibility of the underlying policy-based system, the Service Provider can anticipate users' demands and accelerate the processing of their requests. The following policy example illustrates the benefits of our design.

The adaptation process takes place at the Cluster Heads (CHs) using the aggregated parameters of their cluster requests. A periodic event (*calculateWMA(qualityCnt[ ],codecCnt[ ])*) triggers the above adaptation policy. The Weighted Moving Average is a statistical formula used to analyze time series data in order to smooth out short-term fluctuations, thus highlighting longer-term trends. By counting the occurrences of *low (L)*, *medium (M)* and *high (H)* for the media *quality* preference, the highest WMA value (*popQualityWMA*) identifies the most popular quality (*popQuality*) request. In the same way, the most popular *codec* (*popCodec*) can be identified, i.e., the one available on the majority of the devices during the examined period.

If the average occurrences of the popular formats exceed the ones defined by thresholds (thr1,thr2) and the Cluster Head processing load (*chLoad*) is below 25%, then the CH begins the adaptation action, i.e., transcodes the most requested (*mostReq[ ]*) media files within its cluster using these two parameters (quality *q*, codec *c*). As a result, available media options can be significantly increased for the majority of users. In addition, conditions prevent CHs to start the resource consuming transcoding process, if they are already busy serving users' request (higher *chLoad*).

HIGHER MEDIA AVAILABILITY—THE RATIO OF AVAILABLE MEDIA OF A SPECIFIED PREFERENCES OVER THE TOTAL NUMBER OF AVAILABLE MEDIA—REFLECT A HIGHER PROBABILITY OF A USER'S REQUEST BEING SATISFIED AND A WIDER RANGE OF MEDIA OPTIONS FOR THAT COMBINATION.

5. Evaluation

We have simulated the adaptation process with the enforcement of policy SaP1, measuring its effect over time on the described Media Service. The metrics used are the number of requests per interval for a combination of quality and codec, and the estimated difference of media availability. Media availability is the ratio of available media of a specified preferences and/or device capabilities combination over the total number of available media. We measure the relative media availability difference, when compared with the media availability of the same service without adaptation. Higher ratios reflect a higher probability of a user's request being satisfied and a wider range of media options for that combination. The availability improvement is depicted as a positive relative difference.

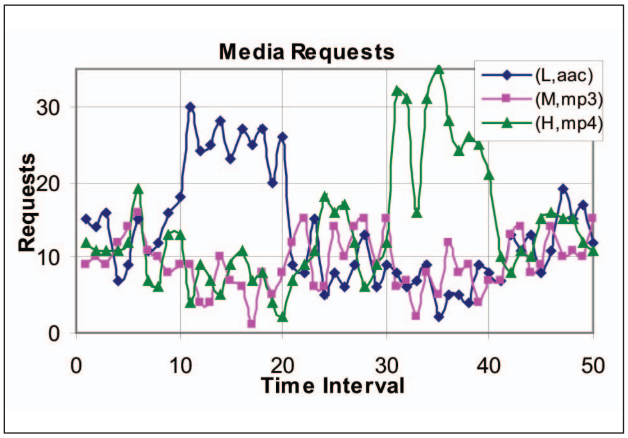

FIGURE 4 Media requests over time.

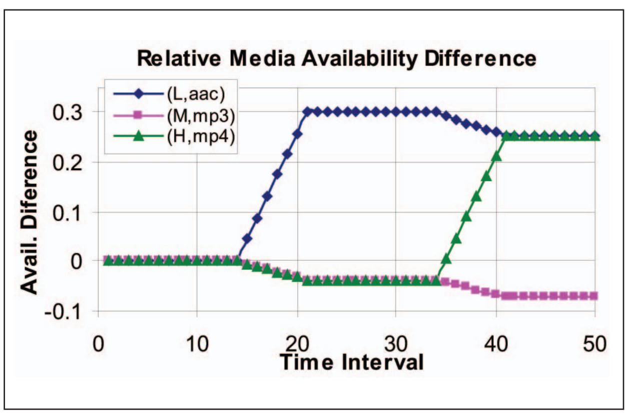

FIGURE 5 Adapting behavior of media availability.

A request generator is programmed to send 100 requests per time interval for a total of 50 intervals. This generator produces random requests except during specified intervals where we deliberately bias the request parameters. The purpose of the bias is to simulate the increase in media requests for a specific combination of quality and codec (q, c). In real life, this would happen when the passengers of a carriage have a common behavior that differs from the average user of the service. For example, a group of students using their mobile phones try to download low quality tracks while on the train, resulting in increased (L,aac) requests. Or commuters of a first class carriage try to access high quality video news on their laptops during peak hours, resulting in increased (H,mp4) requests. These behaviors are simulated with a bias in the request generator during interval periods 11–20 and 31 to 40 respectively. For the purpose of the simulation, we choose three codec formats, namely aac, mp3 for audio and mp4 for video, while three options for quality can be available (Low, Medium, High). Figure 4 shows part of the simulation results for the number of requests for the mentioned combinations ((L,aac),(H,mp4)), plus an additional random one (M,mp3) for comparison. The peaks on the graph are the result of the generator bias.

For every time interval, policy SaP1 is triggered and the popularity thresholds are checked. If both thresholds are exceeded, indicating a very popular quality and codec combination, then the action of transcoding is enforced. This results in an increased number of available media for that combination, thus resulting in increasing media availability during those periods (Figure 5). The use of a weighted moving average ensures that adaptation is not triggered for a sporadic increase in requests. This is also reflected in the delayed triggering of the adaptation process (after interval 14 and 34), ensuring that a trend in users' requests has been established. Effectively, a 12.2% average increase in requests for low quality tracks (L,aac) results in a 3% increase in media availability for that combination (period 11 to 20). Similarly a 15.2% average increase in requests for high quality video (H,mp4) results in a 2.9% increase in media

TABLE 5 *Service Adaptation policy.*

Name	Policy	
SaP1	**if**	(popQualityWMA > thr1) **and** (popCodecWMA > thr2) **and** (chLoad < 25%)
	then	transcodeItems(mostReq[],popQuality,popCodec)

availability (period 31 to 40). For the total period, the effect on media availability for a random combination of media requests (M,mp3) is minimal (−0.7%).

6. Conclusions

In this paper we have presented an adaptive policy-based service management framework for wireless networked environments. The framework accommodates a level of control from the end-user through generic and service-specific preferences. While these can guide the provider towards a fully customized service, they can also be influenced to achieve optimized service utilization. Another important feature of the framework is the support of service adaptation. This functionality is based on statistical and contextual information and as demonstrated through simulation it can potentially enhance service performance and user experience. The overall concept of adaptive and customized service provisioning is driven by policies, which facilitate a flexible and extendable service creation, enhancement and deployment environment.

The various components and functionality of our framework are demonstrated through an extended case study involving a Media Service scenario and by simulating the adaptation procedure. Service management is supported with the specification and description of policies influencing the different levels of processing required, from service creation to service delivery. The examined scenario and simulation results validate the applicability and potential of our approach, despite the relative simplicity of the introduced policies. For more advanced services involving intelligent algorithms and concurrent service/resource utilization, policies tend to increase in numbers and complexity. As such, part of our future work will involve the integration of our ongoing work on policy conflict analysis to support the needs of a complete Service Management framework.

Acknowledgments

The research work in this paper was partly supported by the EU EMANICS Network of Excellence on the Management of Next Generation Networks (IST-026854).

Author Information

Antonis M. Hadjiantonis (A.Hadjiantonis@surrey.ac.uk) is a researcher and Ph.D. candidate at the Center for Communication Systems Research, Department of Electronic Engineering, University of Surrey, United Kingdom, where he is an active member of the Networks Research Group. He holds a Diploma in Electrical and Computer Engineering from the National Technical University of Athens, Greece. His research interests focus on policy-based management, ad hoc networks, management technologies for wireless networks and ubiquitous computing.

Marinos Charalambides (M.Charalambides@surrey.ac.uk) is a Research Fellow and Ph.D. candidate at the Center for Communication Systems Research, Department of Electronic Engineering, University of Surrey, United Kingdom, where he is an active member of the Networks Research Group. He holds a B.Eng. in Electronic and Electrical Engineering and an M.Sc. in Communications Networks and Software, both from the University of Surrey. His research interests focus on policy-based management, policy analysis, and IP quality of service.

George Pavlou (G.Pavlou@surrey.ac.uk) is a professor of communication and information systems at the Center for Communication Systems Research, Department of Electronic Engineering, University of Surrey, United Kingdom, where he leads the activities of the Networks Research Group. He holds a Diploma in engineering from the National Technical University of Athens, Greece, and M.Sc. and Ph.D. degrees in computer science from University College London, United Kingdom. His research interests focus on network management, networking, and service engineering, covering aspects such protocol performance evaluation, traffic engineering, quality of service management, policy-based systems, multimedia service control, programmable networks, and communications middleware.

SERVICE MANAGEMENT IS SUPPORTED WITH THE SPECIFICATION AND DESCRIPTION OF POLICIES INFLUENCING THE DIFFERENT LEVELS OF PROCESSING REQUIRED, FROM SERVICE CREATION TO SERVICE DELIVERY.

References

[1] Deloitte TMT Industry Group, "Technology, Media & Telecom. Trends: Predictions 2007," online report, accessed in Jan. 2007 http://www.deloitte.com/dtt/research/0,1015,sid%3D1012%26cid%3D108298,00.html?theme=tmt1

[2] P. Flegkas, P. Trimintzios, and G. Pavlou, "A policy-based quality of service management system for IP DiffServ networks," *IEEE Network*, vol. 16, no. 2, pp. 50–56., Mar.–Apr. 2002

[3] K. Chan et al, COPS Usage for Policy Provisioning (COPS-PR), RFC 3084, Standards Track, Mar. 2001.

[4] A. Yew, A. Liotta, and G. Pavlou, "Applying a policy-based framework to manage quality of service requirements in the virtual home environment," *IEEE Intl. Conf. on Communications 2002 (ICC2002)*.

[5] A.M. Hadjiantonis, A. Malatras, and G. Pavlou, "A context-aware, policy-based framework for the management of MANETs," *7th IEEE Intl. Workshop on Policies for Distributed Systems and Networks (Policy 2006)*.

[6] A.M. Hadjiantonis, M. Charalambides, and G. Pavlou, "A policy-based approach for managing ubiquitous networks in urban spaces," *IEEE Intl. Conf. on Communications 2007*, Glasgow (ICC2007).

[7] M. Charalambides et al., "Dynamic policy analysis and conflict resolution for DiffServ quality of service management," *IEEE/IFIP Network Operations and Management Symposium (NOMS 2006)*.

[8] M. Burgess and G. Canright, "Scalability of peer configuration management in logically ad hoc networks," *eTransactions on Network and Service Management*, vol. 1, no. 1, Second Quarter 2004.

[9] P. Obreiter and J. Nimis, "A Taxonomy of Incentive Patterns—The Design Space of Incentives for Cooperation," Technical Report No. 2003-9, Faculty of Informatics, Univ. of Karlsruhe, (2003).

VT