This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the ICC 2007 proceedings.

A policy-based approach for managing ubiquitous networks in urban spaces

Antonis M. Hadjiantonis, Marinos Charalambides, George Pavlou
Centre for Communication Systems Research
Dept. of Electronic Engineering, University of Surrey
Guildford, UK
{initial.surname}@surrey.ac.uk

Abstract— **Modern cities are becoming increasingly networked environments and a plethora of computing equipment interacts with today's urban citizens. We refer to these ubiquitous networked environments as "urban spaces" and attempt to manage these under a unified framework based on policies. Since users actively participate in urban spaces with their owned devices and demand more control and privacy, we introduce a scheme to protect user privacy and respect their preferences. We adopt a multiple manager paradigm to enable more entities to offer their services and cooperatively shape management logic based on their objectives. Detailed policy examples illustrate the concepts and simulation results measure the effect of policies on network performance.**

Keywords: ubiquitous computing, policy-based management framework, privacy protection, policies

I. Introduction

Computing devices are everywhere and our everyday life is undeniably linked to several of these. Mobile phones, PDAs, media players or laptops are the indispensable companion of the urban dweller. Beyond our controlled gadgets, myriads of devices require and expect our interaction in an increasingly networked urban environment. In order to describe these complex networked environments we use the notion of "urban spaces" and illustrate the concept in Fig.1. The plethora of computing equipment that needs to communicate and provide seamless assistance to the modern citizen of urban spaces motivates our research in an effort to provide a framework to manage these devices and utilize their capabilities.

We focus on a specific case study which is a subset of the general case of ubiquitous computing. Consider a network formed by the infrastructure of a Network Operator and the devices of individual users. The plethora of wireless devices differentiates such networks from the traditional Internet concepts. The operator's infrastructure includes for example media servers, information kiosks, traffic cameras etc. Users' devices may include mobile phones, laptops, PDAs, as well as home network devices like TVs or media players. The Network Operator has agreements with independent Service Providers, who can use the network infrastructure to offer different services to the users. We propose a policy-based approach to manage the whole network and allow more than one entities to cooperatively perform management tasks. This is possible with the adoption of a multiple manager paradigm where both the Network Operator and the Service Providers can introduce their own policies, while a conflict detection and resolution mechanism is in place. Users participating in the network are willing to share resources, but at the same time they demand more control over their devices and protection of their personal data. The proposed framework integrates the users' preferences and caters for their data protection by using regulatory policies.

The paper is organized as follows: Section II provides the background and our previous related work, while Section III presents the policy-based management framework for ubiquitous networking in urban environments. Section IV provides details on the system architecture with detailed examples on the framework realization. A simulation based evaluation of the network performance with and without policies is given in Section V and we conclude in Section VI.

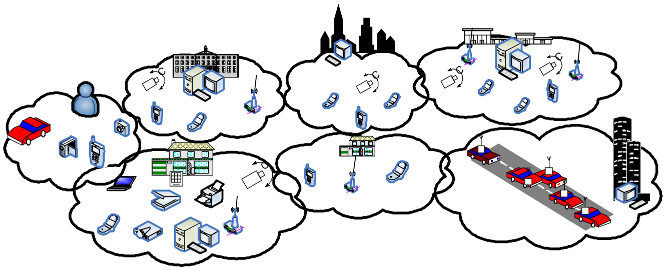

Figure 1. Ubiquitous Urban Space

II. Related Work

A. Literature Review

Ubiquitous networking has received both academic and commercial interest. In [1] a detailed description of the challenges for ubiquitous computing is presented from different perspectives. With the proliferation of wireless networks and increasingly networked environments different approaches have been adopted. In [2], ubiquitous computing is proposed for home networks And in [3][4] spontaneous approaches to networking are presented, focusing on users' interaction and services. Different enabling technologies have been considered as the basis for ubiquitous communication. Mobile Ad Hoc Networks (MANETs) offer fast and cheap deployment without the need of existing infrastructure while emerging Mesh technologies attempt to combine the benefits of MANETs with the support of wired access points [5]. These

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) and the EPSRC PAQMAN project on Policy Analysis (GR/S79992/01).

1-4244-0353-7/07/$25.00 ©2007 IEEE

networks require different management paradigms due to their inherent dynamicity and fluidity. Policy-based approaches have been proposed in [6,7].

Policy-based management simplifies the complex management tasks of large scale systems, since policies monitor the network and automatically enforce appropriate actions in the system [8,9]. In an environment where a number of policies need to coexist, there is always the likelihood that several policies will be in conflict, either because of a specification error or because of application-specific constraints. It is therefore important to provide the means of detecting conflicts in the policy specification [10,11]. Considering the different conflict types, it is possible to define rules that can be used to recognise conflicting situations in the policy specification. These rules usually come in the form of logic predicates and encapsulate application-specific data and/or policy information as constraints. Examples on how these rules can be used as part of a detection process can be found in [12,13]. Another issue regarding policy-based systems is whether the policies should apply to all users and how their preferences are respected. In [14] the authors consider cases where no absolute control from an authority is accepted, while in [15] a "promise theory" attempts to provide "political autonomy" to entities and decentralize policy management.

B. Our Previous Work

The work presented in this paper is based on our previous work on the management of mobile ad-hoc networks [6]. A hybrid organizational model introduced a two tier hierarchical structure and distributed management operations among top level nodes which form the "hypercluster". More than one manager can cooperatively introduce policies to the system using the policy-based functionalities, thus implementing a "multi-manager" paradigm. In addition, context collection and processing functionalities complement the system and provide a feedback mechanism to the PBM system. Three roles are defined for management purposes, namely Manager Node (MN), Cluster Head (CH) and Cluster Node (CN), one of these is assigned to each node. A distributed algorithmic process assigns a role to all devices, depending on a Capability Function which expresses their current status and connectivity parameters. Policies are stored in the Distributed Policy Repository (DPR) which is a set of repository replicas located always on Manager Nodes (MN) and on selected Cluster Heads (CH).

The "hypercluster" notion refers to a set of nodes that are assigned the MN or CH roles, based on the actual device capabilities and utilizing available context information. The algorithmic construction of the hyper-cluster was also presented and evaluated for the case of MANETs. The nodes of the "hypercluster" have a Policy Decision Point (PDP) that evaluates policy conditions and enforces action to the managed cluster. We refer the reader to [6] for a more detailed presentation of the management framework including the context-aware components.

III. POLICY-BASED MANAGEMENT FRAMEWORK FOR UBIQUITOUS NETWORKING IN URBAN SPACES

The objective is to provide a unified and simplified management of networked urban spaces. The nature of these networks sets different requirements, compared to traditional management of fixed ones. More than one manager may have different management objectives. It is essential to detect and resolve conflicts among managers in order to avoid inconsistencies. In addition, users' devices participate in the network but the users require respect for their privacy and preferences. The high mobility environment, in which a user interacts, has an inherent ad hoc element, since he/she intermingles with users and services on the move. The following sections describe the key aspects of our work, in an effort to design a policy-based management framework for networking in urban spaces.

A. Organisational model and multi-manager paradigm

We apply the organizational model described earlier to a specific case-study, i.e. the management of urban spaces under a unified policy-based system. The multi-manager paradigm is ideal for the described case study. The assignment of nodes to the role of a Manager Node (MN) needs to be static, in order to ensure that the "eligible entities" are always selected. An "eligible entity" is a public or commercial organization which has some interest in the management of the ubiquitous network. This interest can be either commercial exploitation of the network by providing value added services or regulatory safeguard of the data and functionality of networked devices. The proposed multi-manager paradigm enables the coexistence of more than one "eligible entities" as Manager Nodes (MNs). Each MN can introduce policies in the system to express its high level goals and these policies are interpreted in the management logic of the network. The distribution of the policies among the hypercluster nodes helps on one hand to distribute the load of management and decision making and on the other hand gives localized control to Cluster Heads. However, the coexistence of distinct administrative authorities raises the issues of conflict detection and resolution. As it will be discussed later, we incorporate a conflict analysis mechanism to our system to alleviate the problem.

For the examined case study of urban space networking, "eligible entities" can be Network Operators (e.g. Mobile Networks Operators), Service Providers (e.g. Multimedia providers), Local Authorities (e.g. Tourism Office), Data Protection Agencies (e.g. Information Commissioner's Office-ICO). To demonstrate our ideas, we choose three entities with competing interests in managing the network: a Network Operator, a Service Provider and a Data Protection agency. Fig.2 displays the deployment of the proposed organizational model in the urban space depicted in Fig.1. Each cloud of devices from Fig.1 forms a cluster. The three Manager Nodes (MN) and the Cluster Heads (CH) form the "hypercluster". The rest of the devices take the role of Cluster Nodes (CN). Table I describes the PBM components consisting each node. The interactions in this multi-manager scenario will be explained in detail in Section IV.C.

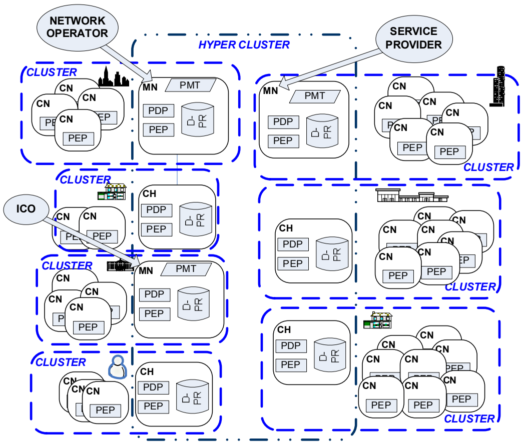

Figure 2. Organizational model in an urban space

Policy-based management can be seen as a way to implement partial autonomy of clusters, by providing them with the management logic (expressed within system policies) and let clusters decide, based on their preferences and local conditions. At the top hierarchy level, network managers need only high-level information and do not need to know about the specifics within each cluster.

B. *Policy-based management architecture*

The principles of a policy-based management (PBM) framework have been introduced in [6]. In this paper we extend the PBM framework to accommodate the needs of urban space networks. Table I summarizes the components of PBM systems. DPR is a distributed version of a traditional Policy Repository.

TABLE I. POLICY-BASED MANAGEMENT COMPONENTS

Components		*Active for*	*PBM functions*
PMT	Policy Management Tool	MN	introduce, edit
DPR	Distrib. Policy Repository	MN, CH	store, distribute
PDP	Policy Decision Point	MN, CH	monitor, decide
PEP	Policy Enforcement Point	MN, CH, CN	enforce, report

The complexity of the environment and the vast numbers of devices provide a challenging environment where the deployment of a policy-based system can significantly simplify management tasks and accelerate devices' configuration. There are several policy types necessary in order to effectively manage urban networks:

1) *Location-Based Services (LBS) policies*
2) *Content delivery policies*
3) *Network-wide Preferences policies*
4) *Charging policies*
5) *Security policies*

We choose a subset of the above in the effort to demonstrate the applicability of PBM through examples applying to the management of networked urban spaces. Location-Based Services (LBS) policies can provide a rich and customizable experience to a mobile user, depending on his/her physical location as well as his/her privacy settings. Content delivery policies can control the information that a user receives while at home or on the move. Network-wide Preferences policies can provide users with the recommended settings and the parameterization of their controlled devices.

For simplicity and clarity, we use a restricted notion for policy specification. Policies follow the established event-condition-action (ECA) specification and can be easily adapted to a complete policy language (e.g. Ponder). For the examples of the defined policies, events are omitted since policies are grouped under the same triggering events. A description of the event is provided for better understanding. To complement the design of our PBM architecture, we employ a mechanism for the detection and resolution of policy conflicts. A number of conflicts may arise in the policy specification, like modality and mutual exclusion conflicts, conflicts of duty and multiple manager conflicts. This work focuses on the last type and addresses the inconsistencies that can occur within the adopted multi-manager paradigm. A detailed conflict detection and resolution example is presented in Section IV.C.

C. *End-user privacy protection*

When it comes to managing a network where the networked devices belong to individuals rather than organizations, issues like privacy and data protection should be considered. In European Union for example, strict legislation by the European Data Protection Supervisor (EDPS, Directive 95/46/EC, http://www.edps.europa.eu) mandates the processing and acquisition of personal data and national authorities have been established to monitor their enforcement (e.g. ICO in UK). Different regulations apply in the US, where a territorial approach is adopted. It is evident that the management of a network consisting of individuals' devices should or is legally obliged to respect the directives regarding the collection and processing of personal data. In order to tackle this issue a twofold protection mechanism is incorporated in the proposed policy-based management framework:

1) *User-centric control:* Individuals can set their privacy preferences to their controlled networked devices and explicitly restrict access to their personal data, regardless of the network policies.

2) *Policy-based regulation scheme*: The national or regional data protection authority has the ability to introduce appropriate policies to the managed system that will ensure users' personal data are not collected or exploited.

The described case study refers to a trusted ubiquitous environment and we assume that the network is always managed by trusted entities. The requirement is to respect users' preferences and safeguard the unfair use of their personal data; therefore we propose a scheme that prevents manager entities to acquire information against the users' will. The case of non-trusted environments poses the requirement of rigorous security schemes and malicious node detection which are out of the scope of this paper. The next section introduces a differentiation between managed objects to accommodate the needs of user-centric control. We further elaborate on technical details for both protection mechanism in the Section IV.A and IV.B.

D. Policy Free and Policy Conforming Objects

We establish the definition and differentiation between Policy Free Objects (PFO) and Policy Conforming Objects (PCO) by indicating the benefits and complications imposed to the system. The motivation behind this differentiation is presented here.

Network management can be seen as a set of operations on managed objects in order to achieve effective FCAPS management, as defined by ISO. Traditionally, a human network manager can control every MO in the system by setting or retrieving values, monitoring the status and reacting to reported events. In other words, a central administrative authority owns and controls the managed network. But as previously explained the case of ubiquitous networking in urban spaces is fundamentally different from traditional networks. The individual users are reluctant to entrust the management of their devices to a central authority and demand more control over their owned devices. This contradiction has motivated our idea to differentiate MOs and introduce Policy Free Objects (PFO) and Policy Conforming Objects (PCO).

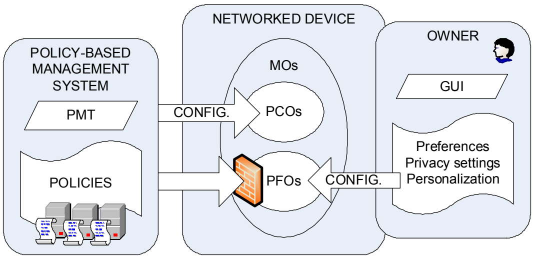

Figure 3. Policy Free and Policy Conforming Objects

A policy-based management system automates the control of network devices, by enforcing policies over their managed objects (MOs). We define Policy Free Objects (PFO) as the MOs of a networked device which are directly controlled by the device's owner and their values and/or status are not influenced by policy decisions. Policy Conforming Objects (PCO), similarly to traditional MO, are controlled by the PBM system, i.e. their values and/or status are influenced by policy decisions. Fig.3 presents conceptually the above definitions. The realization of our ideas is presented in Section IV.A.

IV. SYSTEM ARCHITECTURE - CASE STUDY INVESTIGATION

This section provides technical details on the implementation of the proposed framework. The twofold protection mechanism of user's privacy and preferences is described. First, the user-centric control scheme employs the defined Policy Free and Policy Conforming Objects with example policies. Next, the details of the policy-based regulation scheme are presented with applicability examples. Finally, the conflict detection and resolution mechanism is defined and demonstrated with another example.

A. User-centric control

As outlined earlier, individual users are reluctant to grant complete control of their devices to a central authority and demand more influence on their behavior and data disclosure. The presented idea of Policy Free and Policy Conforming Objects (PFO/PCO) can accommodate these demands and offer a way for users to set their privacy preferences and explicitly restrict access to their personal data, regardless of the network policies. As proof of concept we present an example.

In this example, we define a limited set of Managed Objects (MO) and allow the devices' owners to set their preferences using a user friendly interface. Depending on the users' input, the MOs are classified as Policy Free (PFO) and Policy Conforming Objects (PCO). The mapping is straightforward and the devices automatically carry out the classification. As a result, read/write permissions are set by the user for the information he/she considers sensitive, as well as preferred values for device settings. Table II lists the managed objects and their set of values. The bold values are the ones selected by the user of this specific example.

TABLE II. LIST OF USER MANAGEABLE OBJECTS

Managed Object		***Values***	***Read Access (RA)/ Write Access(WA)***	***PFO PCO***
DST	Device Status	**on**, off, auto	RA Allowed WA Restricted	PFO
PWU	PowerUsage	normal, low, sleep, **auto**	RA Allowed WA Allowed	PCO
SBW	SharedBandwidth	[0-100]%, **auto**	RA Allowed WA Allowed	PCO
SMR	SharedMemory	[0-100]% (**30%**), auto	RA Allowed WA Restricted	PFO
Access Control Object for external Data		***Values***	***Read Access (RA)***	***PFO PCO***
SL	ShowLocation	yes, **no**	RA Restricted for Location data	PFO
SB	ShowBattery	**yes**, no	RA Allowed for Battery status	PFO

The MOs that had their values explicitly set by the user are classified as PFO and they will not be affected by network policies (DST, SMR, SL, SB). The ones with values equal to "auto" are classified as PCO and the PBM system can access and modify them (PWU, SBW). The management system can operate, regardless of the users' selection but cannot override their preferences. Table III contains system policies and based on the user's preferences, policies P3, P4 will not affect the particular user, while policies P1 and P2 will.

TABLE III. NETWORK OPERATOR POLICY EXAMPLES

P#	***Policy***	***affects***
P1	if (SB=yes)^(Battery>30%) then setPWU(normal)	yes
P2	if (SB=yes)^(avgFreeBW>60%)^(Battery>80%) then setSBW(40%)	yes
P3	if (time=[2:00..4:00])^(avgFreeBW>90%) then setDST(off)	no
P4	if (Battery>50%)^(PWU:=normal)^(avgFreeMR>60%) then setSMR(50%)	no

For simplicity, the example policies are not overly complex, yet useful enough to demonstrate the proposed concepts. The case study assumes a network consisting of personal users' devices (mobile phones, PDAs etc), as well as devices controlled by the network managers (information kiosks, wireless traffic cameras, etc). Some of the networked devices may operate unsupervised and the management system must ensure their proper operation. The Network Operator introduces the above policies (Table III) to the system with the purpose of conserving the battery of managed devices (P1,P3) and to allocate shared resources according to device statistics and remaining battery (P2,P4). Statistics such as the average free bandwidth (avgFreeBW) and memory (avgFreeMR) are recorded by the devices and can be used in policy conditions. The user of the example defines his/her preferences for the owned devices, by explicitly setting the device status to on and the shared memory to 30%. Also, the user restricts access to the device's location data but allows the system to read the battery status. As a result, policies P3 and P4 do not apply to the user's device, while policies P1 and P2 do apply and configure the PCO objects, i.e. the shared bandwidth and the power usage profile. Regarding data protection, the disclosure of the user's current position is protected but he/she may not benefit from Location-Based Services (LBS) that utilize his/her position details. The same set of policies affects all networked devices. However, devices that are controlled by the NO operate as normal policy controlled devices, i.e. have all their objects in PCO status. This allows their full configuration by the network manager.

B. Policy-based regulation scheme

In addition to the explicit user defined preferences, the PBM system has the ability to control unfair exploitation of user data by deploying a regulation scheme with appropriate policies. Having explained the rationale for multiple managers and the notion of "eligible entities", we explain how the regulations of data protection can be enforced in the system and more importantly not overridden. In our multi-manager case study, we consider a data protection agency (e.g. ICO, Information Commissioner's Office for UK) as an "eligible entity" that has the control of one Manager Node. Using the PMT (Policy Management Tool) interface, ICO has the ability to manage the lifecycle of policies and introduce appropriate policies to the managed system according to current regulations. In addition, it can review, edit or disable existing policies so at to ensure users' personal data are not collected or exploited by other "eligible entities"; in this case study, by the Network Operator or a Service Provider.

For example, users who are willing to reveal their location data (SL=yes) should be protected from services that can continually track their position. Tracking is possible by frequently polling the user location and comparing consecutive measurements, depending on the accuracy of the available positioning method and the users' speed. With the increased penetration in the consumer market of high accuracy GPS-enabled devices and improvement of indoor positioning methods, this issue is becoming quite important. Let us assume that current regulations state that tracking the position of civilians is allowed within a circular area of uncertainty that has a defined minimum radius, e.g. minimum radius for pedestrians (min_rad) of 100m. The polling interval of location data must have a minimum value (Min_poll_int) so that between consecutive polls, the user can be found in an area with high uncertainty, i.e. uncertainty radius > min_rad.

$$\textit{uncertainty radius} = \textit{accuracy} + \textit{speed}*\textit{polling interval} \quad (1)$$

Using a simple equation (1) the ICO can formulate an appropriate policy that will enforce the described regulation:

if (SL=yes)^(0<Loc.speed<1.5m/s)^(Loc.accuracy<min_rad) then set_Min_poll_int((min_rad-Loc.accuracy)/Loc.speed)

Further than configuration policies, a regulatory body can use the policy-based system to monitor the collection of user data and gather information for offline processing. Simple policies can periodically log information about the services that retrieve user data. The logged details can be reviewed and analyzed statistically to extract information about how Service Providers use the location data of users and investigate their unfair exploitation.

The flexibility of a PBM system allows complex policies to be formulated during runtime and be introduced to the system without disruption. This allows managing entities to adapt to changes and simplifies the complex task of configuring a large scale network as in the examined case study. A change in regulations can be applied by editing existing policies or introducing new ones, without disrupting the operation of the network and affecting the users. From a business point of view that means less cost for software maintenance and less effort for manual configuration and updates of devices. However, from an administrative point of view, the system should incorporate sophisticated mechanisms to resolve policy conflicts in the described multi-manager environment

C. Conflict detection and resolution

Every policy-based system inevitably needs to deal with arising policy conflicts. The proposed PBM system is no exception and this section attempts to enhance the system with a conflict detection and resolution (CDR) mechanism. Although several conflict types can be identified with regard to our application domain, our interest focuses on conflicts arising between policies originating from different managing entities (MNs) as these are closely related to the adopted multi-manager paradigm. We refer to these conflicts as *inter-manager conflicts.*

The proposed CDR mechanism is part of a protocol for the communication of manager nodes (MNs). The protocol defines the procedure for policy updates with conflict detection and resolution and ensures the consistency of the Distributed Policy Repository. This is presented by the sequence diagram in Fig.4. In this case study three "eligible entities" cooperatively manage the network: the Service Provider (MN1), the Network Operator (MN2) and ICO (MN3). The procedure is the same for any number of manager nodes.

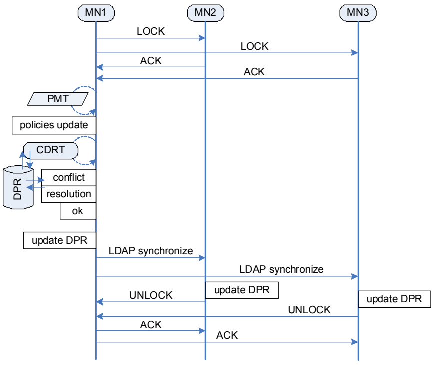

Figure 4. Sequence diagram for policy updates

For the introduction or editing of policies in the system, a MN must send a LOCK message to all other MNs to ensure that no concurrent policy changes occur and ensure the consistency of the Distributed Policy Repository. Once confirmations (ACK) are received the initiating manager can use its Policy Management Tool (PMT).Using the CDR Tool, all new or changed policies are analyzed locally for conflicts based on a set of global detection rules that the eligible entities have agreed upon and specified a priori. In the event of a conflict, resolution can be achieved in different ways depending on the conflict type, the entities involved and any prior agreements between management entities as we demonstrate further below. Once CDRT has verified the consistency of all policies, the initiating MN can update the Distributed Policy Repository, which will automatically propagate changes to other MNs. Once the Manager Nodes have updated their DPR, they reply with an UNLOCK message to the first MN to confirm changes. The MN that initiated the changes sends ACKs to all MNs which release all PMTs for further policy updates.

The occurrence of inter-manager conflicts lies in the fact that each manager has its own high level objectives which are expressed by different policies. Inevitably, these policies may contradict because of incompatible management interests. We provide below an illustrative example which describes such situations and serves as proof of concept for our proposed method of conflict detection and resolution.

In our case study, a Service Provider (MN1) specializing in media delivery wants to maximize profit by providing media to as many users as possible. The Network Operator (MN2) on the other hand, monitors the network to discover bottlenecks and ensures its stable operation by configuring controlled and user devices. Consider a simple scenario where both managers want to configure the shared bandwidth (SBW) of the devices that are located in a specific area with low bandwidth availability and high user density, e.g. a stadium. SBW value is divided in bandwidth for management (mngBW) and bandwidth for forwarded traffic (p2pBW). Both managers want to achieve their objectives by configuring system devices (access points, information kiosks) as well as user devices (mobile phones, PDAs) that allow the configuration of SBW (i.e. SBW is PFO). The Network Operator's policy is to use most of the shared bandwidth for management purposes because a stable network is more important than forwarding p2p data and user traffic. Using the PMT at MN2 the following policy (p1) is composed that sets SBW to 40% of which 30% will be used for management traffic and routing data and 10% for peer-to-peer and forwarded traffic:

```
if (SBW=auto )^(SL=yes)^(locateUser(Stadium))
then setBW((SBW:=40%),(mngBW:=30%),(p2pBW:=10%))
```

The Service Provider on the other hand wants to utilize the users' shared bandwidth for distributing media and content (e.g. advertisements, video replays) among customers and needs more bandwidth for traffic forwarding over multiple hops. To realize these goals, the PMT at MN1 is used to formulate the following policy (p2) that sets SBW to 60% of which only 20% will be used for management traffic and routing data and 40% for forwarded traffic:

```
if (SBW=auto) ^ (SL=yes)^(locateUser(Stadium))
then setBW((SBW:=60%),(mngBW:=20%),(p2pBW:=40%))
```

Assuming that the above policies are triggered by the same event, i.e. the entrance of a user to the stadium area, the two policies above are conflicting since they both aim at configuring the same resource with inconsistent parameters. This is a specialization of an inter-manager conflict and can be detected with a rule of the following form:

```
if [p1.setBW(SBW1, mngBW1, p2pBW1) ^
  p2.setBW(SBW2, mngBW2, p2pBW2)] ^
  [(SBW1 != SBW2) v (mngBW1 != mngBW2) v
  (p2pBW1 != p2pBW2)] ^
  p1.locateUser(_) == p2.locateUser(_)
then signalConflict(BWAlloc(p1, p2))
```

The resolution process proposed here is automated, once a resolution action for each conflict type is agreed upon and pre-specified by the manager entities. This action is triggered when the conflict has been detected and, in this scenario, acts as a mediator between the managers objectives, i.e. allocates a weighted average based on the values provided by the two policies:

```
if signalConflict(BWAlloc(p1, p2))
then setBW((SBW:= p1.getSBW * 0.6 + p2.getSBW * 0.4),
     (mngBW:= p1.getmngBW * 0.6 + p2.getmngBW * 0.4),
     (p2pBW:= p1.getp2pBW * 0.6 + p2.getp2pBW * 0.4))
```

The value of the weights used in the averaging process depends on the contractual agreement between management entities and/or the business model of the managed network. In this example the Network Operator policy values have a weight of 0.6 while the weight is 0.4 for Service Provider policies. This agreement reflects that is more important to maintain a stable network in such an area and give more bandwidth to management traffic. The resulting policy that will be propagated to the other managers is the following:

```
if (SBW=auto )^(SL=yes)^(locateUser(Stadium))
then setBW((SBW:=48%),(mngBW:=26%),(p2pBW:=22%))
```

V. Performance Evaluation

In order to demonstrate how policies can improve network performance, we present an example to evaluate their effect to a wireless network. Consider the described model where clustering is used for management purposes. For policies with cluster-wide scope, a cluster head makes decisions based on local events and conditions. In this example, the aim is to transfer media files between two devices within a cluster and the management system uses policies to examine local conditions and decide the best way to transfer a file, i.e. whether to download the file locally or stream it from the source. In our case study, we focus on wireless networks where clusters can be formed, for example, within a house or among users visiting an attraction.

The Service Provider defines a set of policies that are enforced whenever a media file transfer is requested within a cluster (Table IV). The conditions of these policies use two new metrics that express the current conditions in the cluster:

1) network utilization (NU): expresses the average bandwidth utilization between the source and destination based on the maximum real bandwith of each device

$$NU=(1/2)*[avgBW_s/maxBW_s + avgBW_d/maxBW_d]$$

2) media capacity (MC): provides a metric of how the minimum free bandwidth between source and destination devices compares to the bitrate of the requested media. A bigger MC shows better bandwidth availability for media streaming

$$MC=[min(maxBW_s - avgBW_s\,,\ maxBW_d - avgBW_d)]/mbr$$

where *avgBW* is the average value of a device's utilized bandwidth over time, *maxBW* is the maximum real bandwidth of a device and *mbr* is the requested media bitrate. Subscript s and d refer to source and destination devices respectively.

TABLE IV. MEDIA TRANSFER POLICIES

P#	*Policy*
P1	*if (NU<0.3) then download(file)*
P2	*if(NU>0.3)^(MC>1) then stream(file)*
P3	*if(NU>0.3)^(MC<1) then stream_reduced(file)*

The action of P1 is to download the file if the conditions between source and destination are good (NU<0.3). When NU>0.3, i.e. the average availability of bandwidth is reduced, policies P2 and P3 decide on the action by evaluating MC. If media capacity is sufficient (MC>1) the file is streamed to the user (P2). However, when MC<1 streaming the file at the original bitrate would cause bad media quality as well as further network congestion. Therefore, the action of P3 is to reduce the bitrate of the file before streaming. Bitrate reduction may be achieved by providing an alternative medium format with lower bitrate so as to avoid resource-consuming transncoding.

The defined metrics offer a comparable way to describe the local conditions between source and destination devices. The cluster head evaluates the policy conditions by calculating NU and MC in order to enforce the appropriate action. Although these metrics take into consideration the conditions only at source and destination, we argue that this is sufficient for our proposed management model since the created clusters are relatively small [6]. This is necessary in order to avoid the severe bandwidth reduction over multiple hops in MANETs.

In order to evaluate the effect of the above policies to the network performance, we used the ns2 simulator. The purpose of the simulations was to measure the performance of a wireless ad hoc network based on 802.11 with or without the presence of the aforementioned policies. We setup transfers over a multi-hop MANET cluster and emulated file download with a FTP traffic generator and media streaming with a UDP generator. Additional TCP traffic flows were created to emulate the avgBW values. The effective bandwidth of 802.11 based networks is much less than the theoretic maximum of 11Mbps, therefore we set maxBW to 1Mbps for our calculations.

The simulation scenario included the transfer of different file types (Table V) between two users under various network conditions. We performed several tests for each file type and measured performance characteristics for downloading or streaming the same media file. The chosen media had the same duration, so as to illustrate the option of streaming different versions of the same file. For each test, the values of NU and MC were calculated and the policies decided which action to enforce.

TABLE V. MEDIA TABLE

	Size(s)	*Bitrate(Kbps)*	*Dur.(s)*	*Popular Formats*
M1	2880	96	240	MP3 podcasts, 3GPP video
M2	24000	800	240	MPEG4 video

Fig.5 shows the downloading throughput from source to destination with respect to the network utilization. We can see that the enforcement of P1 ensures that when downloading (NU<0.3) the throughput is sufficient.

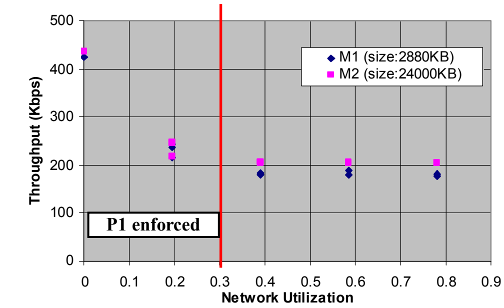

Figure 5. Throughput for downloading between source and destination

In addition, the download time remains reasonable as demonstrated in Fig.6, where the download time ratio (the time of each test over the minimum download time for NU=0) is low for NU<0.3. A ratio=2 means the user has to wait twice as much if the same file was downloaded for NU=0.

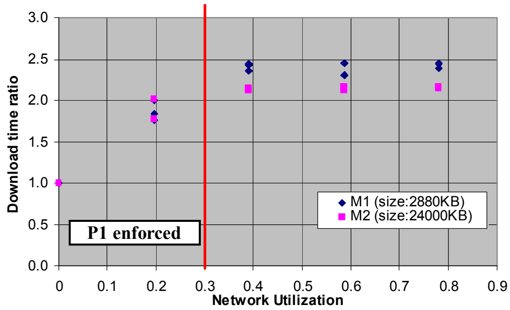

Figure 6. Download time ratio

For NU>0.3 the PBM system decides to stream media, in order to avoid excessive download times. Based on media capacity value (MC), policy P2 or P3 is enforced. Streaming tests were performed for the same network conditions as in the previous simulations and the same media were used. For streaming media, a representative metric of the quality is the end to end delay of the received packages. As expected, the smaller the MC the bigger the delays observed. The long delays while streaming M2 (bitrate 800Kbps) can be avoided with the enforcement of P3, since in those cases MC<1. By streaming the alternative version M1 (bitrate 96Kbps) the delays are significantly reduced and MC remains above 1.

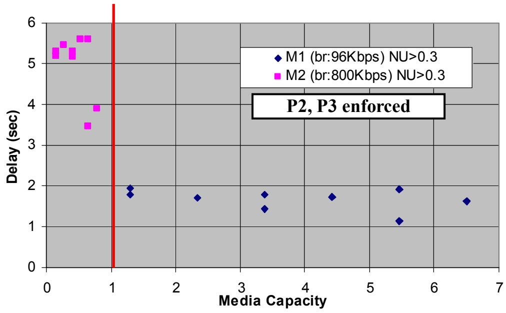

Figure 7. Received packet delays for streaming media

In addition, we calculate the throughput ratio as the transmitted throughput over the actual media bitrate. The measurements presented in Fig.8, indicate that high bitrate media (M2) cannot be transmitted under the current conditions and the degraded ratio translates to bad media quality. Streaming low bitrate media (M1) is possible and the ratio is near 1, demonstrating excellent media quality. Again the value of MC reflects the local conditions and the enforcement of policies P2 and P3 prevents the initiation of a high bitrate transmission when the conditions do not allow for satisfactory media transfer rates.

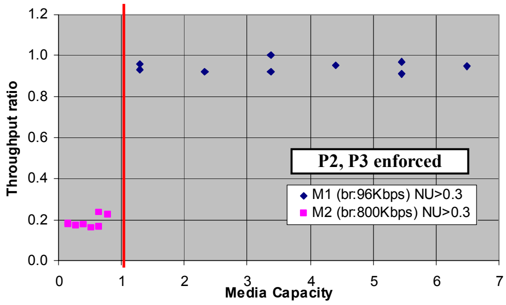

Figure 8. Throughput ratio for streaming media

Based on the presented results, we argue that the proposed PBM solution can provide tangible benefits to the network performance. A significant improvement can be achieved since policies control the creation of media traffic flows and prevent further congestion. From the users' point of view, the experience in sharing media is improved. Although user's experience is subjective, measurements of packet delays and download times offer an objective metric to evaluate the quality of media delivery. These metrics show reduced packet delays with the deployment of appropriate policies and improved quality of delivered media.

VI. Conclusions

We have presented a novel policy-based framework for the management of the increasingly networked urban spaces. The framework offers each user the opportunity to set his/her individual privacy settings and preferences, regardless of the network policies, while regulatory bodies can monitor the acquisition of user data and investigate unfair exploitation. By adopting a multi-manager scheme we allow more entities to offer different services to the users, without violating their privacy concerns and preferences. The framework integrates an automated conflict detection and resolution mechanism that prevents policy inconsistencies among different managers.

The detailed examples illustrate the potential of our framework while simulation results show improved network performance. In our future work we plan to introduce more advanced policies that cover complex scenarios and enrich the functionality of the framework. A wider range of policy detection and resolution cases needs to be investigated, to enhance the stability of the PBM system. In addition, we plan to investigate the cooperation of hyper-cluster nodes for sharing alternative media formats between clusters. Our aim is to provide a complete management framework to facilitate rich user experience and interaction in a seamlessly networked urban environment.

References

[1] D. Chalmers et al, "Ubiquitous Computing: Experience, Design and Science", Ver.4, accessed September 2006, http://www-dse.doc.ic.ac.uk/Projects/UbiNet/GC/Manifesto/manifesto.pdf

[2] H. Schulzrinne et al, "Ubiquitous computing in home networks", IEEE Communications Magazine, Vol.41/Iss.11, Nov.2003

[3] L.M. Feeney, B. Ahlgren, A.Westerlund, "Spontaneous networking: an application oriented approach to ad hoc networking", IEEE Communications Magazine, Vol.39/Iss. 6, Jun.2001

[4] J. Latvakoski, D. Pakkala, P. Paakkonen, "A communication architecture for spontaneous systems", IEEE Wireless Communications, Vol.11/Iss.3, Jun.2004

[5] R.Bruno,M.Conti,E.Gregori, "Mesh networks:commodity multihop ad-hoc networks", IEEE Communications Magazine, Vol.43/I.3, Mar.2005

[6] A. Hadjiantonis, A.Malatras, 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)

[7] R. Chadha et al, "Policy Based Mobile Ad hoc Network Management" 5th IEEE International Workshop on Policies for Distributed Systems and Networks (Policy 2004)

[8] M.Sloman, E. Lupu, "Policy Specification for Programmable Networks", Proceedings of First International Working Conference on Active Networks (IWAN'99), Berlin, June 1999

[9] D.C. Verma, "Simplifying network administration using policy-based management" , IEEE Network, Vol.16,Iss.2, Mar-Apr.2002

[10] E.C. Lupu, M.S. Sloman, "Conflicts in policy-based distributed systems management," IEEE Transactions on Software Engineering, v.25, 1999

[11] J.D. Moffett, M.S. Sloman, "Policy conflict analysis in distributed system management," Journal of Organisational Computing, v.4, 1994.

[12] M. Charalambides et al., "Policy conflict analysis for quality of service management," 6th IEEE Workshop on Policies for Networks and Distributed Systems (Policy 2005)

[13] 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)

[14] M.Burgess, G.Canright, "Scalability of peer configuration management in logically ad hoc networks", eTransactions on Network and Service Management, Volume 1/ No.1 Second Quarter 2004

[15] M. Burgess, "An approach to understanding policy based on autonomy and voluntary cooperation", 16th IFIP/IEEE Distributed Systems: Operations and Management Workshop (DSOM2005), LNCS 3775