A context-aware, policy-based framework for the management of MANETs

Antonis M. Hadjiantonis, Apostolos Malatras, George Pavlou
Centre for Communications Systems Research, Department of Electronic Engineering,
University of Surrey, UK
{a.hadjiantonis, a.malatras, g.pavlou}@surrey.ac.uk

Abstract

Mobile ad hoc networks (MANETs) are an emerging paradigm in wireless communications that has recently attracted a lot of attention. Their inherent benefits such as unrestrained computing, lack of centralization and ease of deployment at low costs are tightly bound with relevant deficiencies such as limited resources and management difficulty. There is a need for new management approaches to handle the requirements and specific characteristics of these networks. We propose a hybrid approach, employing a hierarchical and distributed organizational model for MANET management. We adopt a policy-based network management (PBNM) approach together with context awareness and we present our system architecture that is capable of effectively managing a MANET. We present and evaluate our approach under various application scenarios.

1. Introduction

There exists an increasing interest towards wireless and particularly ad hoc networking. In ad hoc networks, mobile nodes move randomly and organize themselves arbitrarily; thus, the network's wireless topology may change rapidly and unpredictably. Conventional wireless networks require some form of fixed network infrastructure and centralized administration for their operation. In contrast, since MANETs are self-creating, individual nodes are responsible for dynamically discovering other nodes they can communicate with. This way of dynamically creating a network often requires an equally dynamic ability to manage the network and supported services according to higher-level management goals (i.e. policies), taking also into account the surrounding conditions (i.e. context).

We believe that this highly dynamic environment can benefit from a Policy Based Network Management (PBNM) approach and the emerging context-driven autonomic communications trend. One of the major advantages of adopting a policy-based approach is the relevant "controlled programmability" that can offer an efficient and balanced solution between strict hard-wired management logic and unrestricted mobile code migration and deployment. While there has been previous research on deploying PBNM solutions for MANETs, our work introduces a novel organizational model specifically targeted to the needs of MANETs by incorporating context awareness to dynamically adapt to the continuously changing conditions. Context information can be used to trigger cross-layer changes (network and application configurations) according to policies, leading to a degree of autonomic decision-making. To the best of our knowledge the proposed context-aware PBNM framework is the first to consider exploiting context information in conjunction with a PBNM system for MANET management. Our organizational model is evaluated based on realistic assumptions, taking into account network topology and device capabilities.

The rest of this paper is organized as follows. After this brief introduction, Section 2 reviews related work in the area. In Section 3 we present our proposed management framework, focusing on issues such as the management model and hyper-cluster formation. Our system architecture is the focus of Section 4, with emphasis being placed on PBNM specifications and context management. Section 5 evaluates our work using both simulation and testbed measurements, while Section 6 concludes the paper and gives pointers to ongoing and future work.

2. Related Work

Related literature in mobile ad hoc networks management is limited and proposed solutions attempt to only partly solve relevant issues. Existing approaches vary regarding the adopted organizational model. Recently there has been a shift towards PBNM systems through a hierarchical approach, which is though not well suited for a MANET environment. Also the proposed deployment of mobile agents

0-7695-2598-9/06 $20.00 © 2006 IEEE

amplifies their inherent security implications. The policy-based paradigm [1,2] offers a promising solution since it allows dynamic alteration and controlled programmability of management logic based on the supported policy types.

	Tiers	Hierarchical	Distributed	Policy-based	Agent-based	Modules	Storage	Managers
ANMP [3]	2	+	-	-	-	1	anmpMIB	1
Guerilla [4]	2	+	+	-	+	4	MIB	1
R.Chadha [5]	2	+	-	+	+	1	mySQL	1
K.Phanse [6]	2	+	-	+	-	2	PIB	1
Our Proposal	**2**	**+**	**+**	**+**	**-**	**2**	**LDAP**	**≥1**

Table 1. Taxonomy of related work on MANET management

The first efforts to tackle MANET management were presented in [3]. The suggested Ad hoc network management protocol (ANMP) was based on hierarchical clustering of nodes in a three level architecture. The two proposed clustering algorithms limit severely its applicability due to their centralization. The “Guerilla” architecture [4] adopts an agent-based two-tier distributed approach where at the higher level “nomadic managers” make decisions and launch active probes to fulfill management objectives. Mobile agents exploit a utility function to decide their migration and probe deployment.

In [5] a PBNM system using intelligent agents is proposed. Policy agents are deployed and manage the network through a two tier hierarchical architecture. Policy definitions follow the principles of IETF but the use of several proprietary protocols (YAP, AMPS, DRCP/DCDP) restricts its wider adoption. Another PBNM approach is presented in [6] in order to provide QoS in MANETs. The proposed k-hop clustering scheme and extensions to COPS for policy provisioning (COPS-PR) protocol add policy server delegation and redirection capabilities. Although in RFC status, COPS and COPS-PR have found little acceptance and their relatively heavyweight nature may limit their applicability to MANETs. Table 1 presents a taxonomy and summarizes related work in the area, also including our work.

The exploitation of context in network management has been addressed before since the potential benefits can be tangible. Relevant adaptive systems can be deployed, interacting with the surrounding environment and functioning according to relevant conditions [7-9]. The main drawback of all these approaches, including our previous work [7][13], is the static evaluation of context against predefined rules. The use of network policies to achieve a more dynamic nature at a higher conceptual level has not been considered in the past and this is one of the particular innovative aspects of our approach.

3. MANET Management Framework

The proposed framework introduces a novel policy-based organizational model which combined with context information processing can effectively manage a MANET. The proposed model has a number of novel features such as its hybrid organizational approach and hyper-cluster formation, a distributed and replicated policy repository and context-driven policy enforcement.

3.1. Distributed and hierarchical model

We adopt a hybrid approach by proposing a distributed and hierarchical model. Using the policy-based paradigm and context awareness we aim to provide a suitable organizational model for MANET management.

Before analyzing the proposed model we first explain the differentiation between node “modules” and “roles” (Figure 1). The three “roles” refer to the MN (Manager Node), CH (Cluster Head) and CN (Cluster Node) roles, as used in clustering schemes. Beyond the traditional duties for these roles, their behavior and mission is guided by the defined policies. A “module” is the preinstalled software of a node. Our design has two modules: CM (Cluster Manager) and TN (Terminal Node).

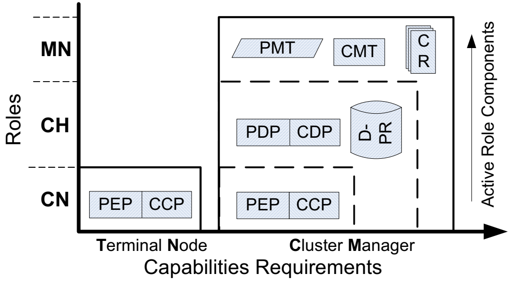

Figure 1. Node roles and modules

This module separation was deemed necessary to accommodate a wider range of node capabilities in the MANET. The TN is the simplest module and is lightweight to make it suitable for limited capabilities devices, e.g. smart mobile phones. Thus TNs can only be assigned to the CN role. On the other hand CM modules have full PBNM system functionality and context processing capability. Therefore CMs are collaboratively responsible for MANET management and can be assigned to all three roles. The selection of the appropriate module for each network device depends mainly on its capabilities. A set of minimum requirements offers a prescribed guideline and indicates whether a device can efficiently host the CM module.

Effectively the differentiation between “roles” and “modules” refers to the differentiation of the organizational role of an entity in the network as opposed to the actual software capabilities it carries. The node roles and modules are depicted in Figure 1 where their respective policy and context related components are also shown. Depending on the assigned role of a cluster manager (CM), the respective components are either active or dormant. During initial network setup, default policies stored in CM are loaded in order to guide the model’s deployment.

Based on the above classification, we present our policy-based organizational model. The multi-manager paradigm and the hyper-cluster formation are introduced, aiming to offer a balance between the strictness of hierarchical models and the fully-fledged freedom of distributed ones. At the same time our model embraces both as it can be deployed as either of these. In addition, policy-based management offers controlled programmability to the system to suit the MANET dynamic environment. The components for each of the three roles are depicted in Figure 2 as well as the information flows between them.

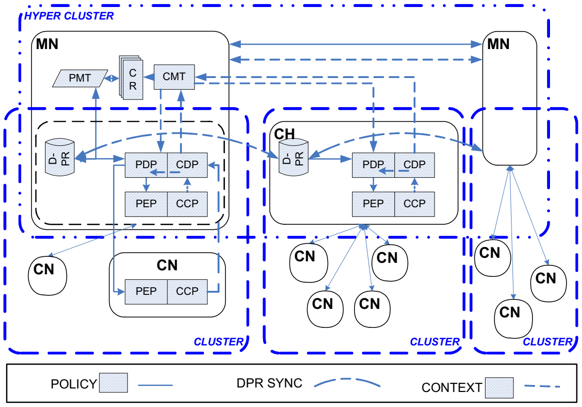

Figure 2. Organizational model and node roles

The idea behind the multi-manager paradigm lays in the nature of ad hoc networks and the purpose of their formation. Having more than one manager gives the flexibility to form networks between distinct trusted administrative authorities. This is performed without any of these being forced to forfeit its management privileges. Instead managers cooperatively introduce policies which guide the overall network’s behavior. For example, a MANET can be setup for a corporate meeting between two companies’ representatives. The multi-manager paradigm treats the companies’ managers as equals and allows both to affect network behavior by introducing policies. In addition, from a functional point of view, in large scale ad hoc networks scalability issues demand more than one manager in order to control and administer effectively the numerous cluster heads. A deployment example with two managers is depicted in Figure 3, mirroring the organizational model in Figure 2.

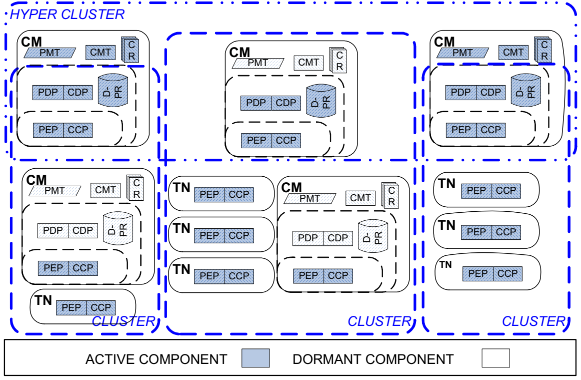

Figure 3. Example deployment and node modules

3.2. Hyper-cluster formation

We introduce the hyper-cluster notion referring to a set of Cluster Manager (CM) nodes that are assigned the MN or CH roles, utilizing available context information. The construction of the set is discussed in detail later. Nodes that hold the MN role encapsulate the CH role as well. The assignment of a node to the MN role depends on the MANET formation purpose and its application use. In the corporate meeting example, the two devices which the companies’ managers use are assigned the MN roles. In a military oriented scenario, e.g. platoon leaders would become MNs. If there is no apparent specification for the MN assignment or in dynamic conditions where MN re-assignment is necessary, then this occurs in an algorithmic fashion as described later. Future work will consider the utilization of context information or a reputation-based election procedure for the selection of MNs.

Of fundamental importance to our architecture is the Capability Function (CF) of a node. It denotes its current ability to host resource-consuming software modules. The CF reflects two aspects of the nodes’ capabilities, one referring to their computing attributes and another to their mobility (Mobility Ratio - MR). The driving concept is that if a node moves constantly and is responsible for link breaks, then it should not be deemed as capable. In our approach the CF considers specific computing attributes, namely memory (MEM), processing power (PP), battery power (BP) and computing load (CL). MEM, PP and BP have a proportional relationship with the CF and MR and CL an inversely proportional one. By assigning weights to these variables based on their significance and incorporating time we come up with Equation 1.

$$CF(t) = \frac{w_1 \times MEM(t) + w_2 \times PP(t) + w_3 \times BP(t)}{w_4 \times MR(t) + w_5 \times CL(t)} \quad (1)$$

For the CF to be comparable, the variables are normalized to a value range of [0, 1] by dividing each one with its maximum counterpart as defined by us (i.e. we consider a maximum PP of 3 GHz). The sum of the weights w_1, w_2 and w_3 should be 1 while the same holds for w_4 and w_5. This ensures a bounded CF within the range (0, 1]. It is possible for a variable to exceed the maximum defined value; the normalization process will still yield a value of 1. The default policies loaded at startup initialize the weight values.

MEM, PP and BP are obtained from the system profile. For CL the product of the percentage of CPU utilization and the PP value is used. Regarding the MR calculation, we cannot simply consider the frequency of node movements, since this is not indicative of topology changes. All nodes i.e. might constantly move towards the same destination. In addition a non-mobile node might not be deemed as capable since all other nodes might be moving away from it. Considering these, we associate the MR of a node with the frequency of link breaks with its neighbors as these have occurred up until the time of measurement. This information is obtained from the network layer (i.e. routing table). In the future we plan to use context-driven mobility predictions to enhance the MR proactively. Equation 2 yields the MR from the time the MANET was setup up until time t with the value being normalized in the range of [0, 1] (the number of link breaks cannot exceed that of neighbors).

$$MR\ (t) = \frac{link\ \_\ breaks\ \ ([\,0\,,t\,])}{neighbors\ \ ([\,0\,,\ t\,])} \quad (2)$$

The proposed algorithm (Figure 4) receives input regarding the module installed on every node, its Capability Function and the network topology, yielding the role assignment to each node. When the network is initially setup every node calculates its CF and periodically tracks it. Some, or even only one, of the nodes carrying CM will eventually be marked as MNs. This can be statically configured at the MANET setup as discussed.

After the initial phase, a selection process commences to establish the role of the nodes. All nodes carrying the TN module acquire the CN role. Our framework is generic and allows for nodes having the CM functionality to be assigned any of the three roles. This has the benefit of forming the hyper-cluster from a restricted set of CHs and MNs amongst all possible candidates. For reliability reasons backup CHs can be enabled in case a CH fails due to insufficient resources or mobility. The hyper-cluster will be consisted of nodes that form the dominating set (DS) of the graph of nodes with CM modules, thus ensuring one-hop accessibility for the remaining nodes with CM module, which are assigned the CN role. The idea is borrowed from backbone overlay networks used for routing in MANETs where the use of DS is prominent. Based on the extensive work on the area [10-12] we propose exploiting a distributed algorithm for the DS formation using as an optimization heuristic the defined CF.

```
Hyper Cluster Formation Algorithm

Assumptions
- N is the set of nodes in the MANET
- Nb(x) is the neighboring set of a node x
- MN, CH, CN boolean flags denote node role
- CM, TN boolean flags denote node module functionality

Construction Rules
1. $\forall\ x \in N, if\ TN(x)=1 \rightarrow CN(x)=1$
2. $\forall\ x \in N,\ calculate\ \ CF(x)$
3. $\exists\ x \in N : MN\ (x)=1\ by\ default$
4. $\neg \exists\ y \in Nb(x):CM\ (y)=1\ AND\ CM\ (x)=1 \rightarrow CH\ (x)=1$
5. $\forall\ y\ \in Nb(x):CM(x)=CM(y)=1 \rightarrow (x \xrightarrow{Nb(x)} y) \wedge (y \xrightarrow{Nb(y)} x)$
6. $\exists\ y,w \in Nb(x):CM\ (x)=CM\ (y)=CM\ (w)=1,$
   $y \notin Nb(w)\ AND\ w \notin Nb(y) \rightarrow CH\ (x)=1$

Optimization Rules
1. $x, y \in N, Nb(y) \subset Nb(x), CF(x) \geq CF(y),$
   $CH(x)=CH(y)=1 \rightarrow CH(y)=0 \wedge CN(y)=1$
2. $x, y, w \in N, Nb(w) \subset Nb(x)\ Y\ Nb(y),$
   $\min\{\ CF(x),\ CF(y), CF(w)\} = CF(w),$
   $CH(x)=CH(y)=CH(w)=1 \rightarrow CH(w)=0 \wedge CN(w)=1$
```

Figure 4. Hyper-cluster formation algorithm

If a node with the CM module has no other neighbors with the CM module then it marks itself as CH. The remaining nodes with CM module exchange neighboring tables (only entries including nodes with CM functionality) with their neighboring nodes with CM module. If such a node has two unconnected neighbors with CM module it acquires the CH role. To ensure optimized construction of the dominating set, a trimming process takes place that prunes the CH set from redundant CHs with small CF value.

Effectively, the CHs together with the MNs form the hyper-cluster and collectively manage the MANET. Every node with CM module is registered to it and can access it in one hop. Every CN registers itself to its CH neighbor with the highest CF value, while those that do not have such neighbors acquire a route to one of them through their CN neighbors. Depending on the application use of the MANET, the MNs are either dynamically introduced or statically configured upon the initial construction of the MANET. In the latter case our algorithm takes this into account by assigning explicitly these nodes to the MN role and thus to the hyper-cluster. In the first case the aforementioned algorithm is executed once again upon the selected set of CHs to deduce the set of MNs of the MANET. The result is a clustered MANET with nodes in all three of the defined roles. Space limitations avert us from describing the algorithm used to maintain the hyper-cluster when node movements affect network topology.

3.3. Distributed & replicated PR

The policy repository (PR) is a critical component in every PBNM system and we cannot rely on a single node to store it. The idea of storage replication is not new and is widely used in fixed networks as a backup in case of failures. In ad hoc networks however due to the intermittent nature of wireless links, it is expected that nodes will become disconnected frequently. Thus access to a central repository cannot be guaranteed depending on the networks volatility and mobility. In order to tackle this deficiency we propose the DPR (Distributed Policy Repository) component.

DPR is an enhanced version of the Policy Repository [14] and is consisted of repository replicas distributed among hyper-cluster's nodes. Instead of simply replicating the PR among the nodes, we incorporate a sophisticated policy-based replication scheme. We further elaborate on the policy related issues in Section 4. By utilizing context information (i.e. the Mobility Ratio) and based on the introduced policies, the system automatically enforces the appropriate replication state among hyper-cluster nodes, depending on how volatile the MANET is. In this way, we provide alternative access options in case a repository is corrupted or disconnected and distribute traffic load and processing overhead among nodes.

The DPR component lies within the CM module and it can effectively store network policies in a directory server. In order to balance resource consumption and policy accessibility, a selection of the hyper-cluster nodes activate their DPR component and carry a replica of network policies. The DPR state of each node is imposed by the network policies which define the overall policy replication state. Management objectives reflected in policies and MANET volatility, influence the DPR replication degree to conserve resources and ensure maximum repository availability. Specifically, when network mobility is high and links are exceedingly intermittent, reliable access to a remote Policy Repository may be impossible. In this case, policy objects (PO) monitoring network mobility detect the high volatility and proactively increase the replication degree of DPR. Effectively the network will respond with increased decentralization of the policy repository, pushing the storage points (DPRs) closer to the decision points (PDPs). Each manager node (MN) or cluster head (CH) with an active DPR accommodate a replica of the repository and serve as access points for repository requests within their cluster, balancing this way processing load and traffic in the network. A CH with a dormant DPR can access policies from a list of neighboring CHs or MNs with an active DPR.

In effect policies and context guide the DPR behavior and replicas' distribution, ensuring on one hand maximum repository availability (distributed copies) and on the other hand a single logical view of the stored policies (replicated content). Thus, efficient management of clusters can be achieved even when temporarily disconnected from the network manager.

3.4. Context driven policy enforcement

Apart from the policy related components as these are identified by the IETF policy framework we introduce a group of new entities related to context collection and processing. These are necessary for our system being capable of sensing, communicating with its surrounding environment and adapting to changing conditions. Incorporating context awareness into our policy-based management framework makes it flexible and dynamic in response to the inherently unstable MANET domain, allowing a degree of autonomy to be reached.

The key point is to ensure uniform network management. This is achieved by the presence of an identical and synchronized Policy Repository at the MNs, which in turn guarantees that all PDPs behave in the same way and enforce the same actions in a MANET-wide fashion. Because of the unified manner aggregated context is presented to the MNs the conditions' evaluation is the same.

4. System Architecture

Based on the requirements expressed in previous Sections and the proposed organizational model, we present the design principles for our system architecture. It is an extension to the traditional PBNM system that takes into account the inherent characteristics of ad hoc networks. By incorporating context information this leads to a certain degree of dynamic adaptation and allows for self-management. We provide more details on these aspects in the following sections.

4.1. Policy-based Management

In order to apply the PBNM paradigm to our system we adopt the standardized by IETF/DMTF information model for policy representation. Defined policies are represented according to PCIM/PCIMe [15-16]. Based on this representation we map policies to the standardized LDAP data model [17-19].

Our initial efforts focus on the definition of the necessary policies for MANET management, rather than the formal definition of a policy language. We believe this representation is both effective and lightweight so as to cater for the policy needs in the resource-poor MANET environment:

{Roles} [TimePeriod] if {conditions} then {actions}

We define the policy's "enforcement scope" as the set of nodes where actions need to be enforced, when the policy is triggered by this set's collected context. Regarding the actual policy design, we model realistic examples of policy types in order to illustrate our concepts. These policies are a first step towards a flexible and adaptable management framework specifically designed for the needs of a MANET. Based on the above definition, three enforcement scopes are realized for the needs of our design:

a. MANET-wide - Routing adaptation

Policies can be triggered at the Manager Nodes by the context collected and aggregated from all network nodes. MN's PDP decide and enforce the actions network wide to all MANET nodes' PEPs. These policies are identified by their assignment to the MN node role.

Routing adaptation example policy: a plethora of protocols has been proposed to solve the multihop routing problem in MANETs, each based on different assumptions. A generic classification can distinguish them into proactive and reactive regarding the strategy used to establish routes between nodes.

Based on the above, we decided to model a policy which would enable dynamic on the fly adaptation of the routing protocol. Network conditions and context on one hand and administrator defined management goals on the other, can both be expressed by this type of policy which effectively alters routing strategy and increases network performance:

{MN}[T] if {RM=(n..m)} then {RoutProt:=k}

This policy type is used to adapt network behavior by switching the routing protocol (RoutProt) according to the network's relative mobility (RM). RM is aggregated context information extracted from the network-wide knowledge of node movements, GPS positioning data, mobility ratio and other context. The simple condition monitors if RM value lies within the range (n..m) in order to enforce the associated simple action that activates the appropriate routing protocol. In our implementation the idea is to use a proactive routing protocol (OLSR, k=1) when relative mobility is low and a reactive (AODV, k=2) when high. Depending on management goals and on network-wide aggregated context, compound conditions and actions can be introduced in all the proposed policy types, in order to take more parameters into account.

The network-wide enforcement scope of this policy implies that the condition variables used should have an aggregated network-wide value. The RM value for example is extracted from the aggregation and processing of node context such as the Mobility Ratio (MR) and it is collected at the Context Management Tool (CMT) components of the Manager Nodes. This higher level context information will drive the triggering of actions that should be enforced globally. Each CMT forwards this value to the local PDP and to the PDPs of all CH's under their MN. Each PDP enforces the triggered action to all cluster nodes, including itself. Finally all CHs report successful execution to their MNs. More information about context related components is given in Section 4.2

b. Hyper-cluster wide - Repository replication

Policies can be triggered at all hyper-cluster nodes by the context aggregated within the hyper-cluster. Decisions are enforced only at the hyper-cluster nodes. These policies are identified by their assignment to MN and CH node roles only.

Repository replication example policy: the need for repository replication has already been explained in detail in Section 3.3. For this purpose we model a policy type to guide the replication degree of the Distributed Policy Repository. A manager node has the ability to dynamically define the behavior and the replication degree of the DPR by introducing related policies on the fly and without shutting down the system or the DPR component.

{MN,CH}[T] if {FM=(n..m)} then {ReplDegState:=k}

This policy type is used to guide the replication degree (ReplDegState) of the Distributed PR (DPR) component. The fluidity metric (FM) is a cluster-wide aggregated value similar to the RM metric. We implement three states of replication, namely k=1:Single, k=2:Selective and k=3 Full. These states reflect the need for PR replicas within the hyper-cluster nodes and adapt according to the volatility of the MANET. As mentioned earlier, the idea is to increase the DPR replication degree when network fluidity increases (Figure 5).

Based on cluster-wide information, such as the FM, the cluster's Context Decision Point (CDP) informs the collocated PDP and policies of this type may be triggered for hyper-cluster wide enforcement. Each action {*ReplDegState:=k*} implies a state transition to State k and is enforced differently.

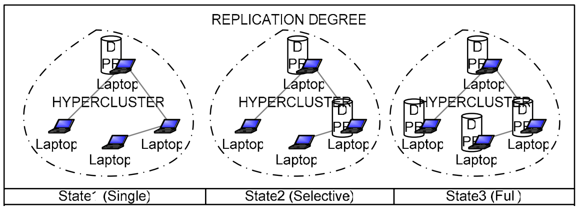

Figure 5. Replication Degree States

A transition to full replication (State 3) is succeeded by activating the DPR components at all nodes participating in the hyper-cluster. In order to transit to selective replication (State 2) a selection algorithm, similar to the roles' selection one is executed, this time only among the hyper-cluster nodes. A dominating set is selected to host the active DPR components. Finally the transition to single replication is achieved by keeping active the DPR of the MN with the highest capability function (CF), while maintaining backup files. So, policy instances of this type implement the adaptation of the DPR replication degree according to network fluidity changes.

c. Cluster wide - Energy conservation

Policies can be triggered at all hyper-cluster nodes by the context aggregated within their cluster. Decisions are enforced only at the cluster nodes that triggered the policy. These policies are identified by their assignment to all three roles (MN, CH, CN).

Energy conservation example policy: a major issue in MANETs is the conservation of device resources. We tackle this by introducing a policy type that adaptively configures energy consumption according to their current state and environment as well as the overall management objectives:

{MN,CH,CN}[T] if {BP=(n..m)} then {TransPow:=k}

This policy type is used to manage effectively the device resources by influencing relevant configuration parameters. The Battery Power (BP) context is used here to affect the node's transmission power (TransPow). In our implementation k ={1,2}, where 1=Normal Power and 2=Low Power. The idea is to use a threshold battery level in order to reduce transmission power and conserve remaining battery power. Policies of this type only need cluster-wide context knowledge since their enforcement is independent among clusters. The PDP of every Cluster Head receives context information for the registered variables and enforces the actions to the PEP of the reporting CN. In these cases context information is withheld within the cluster, thus reducing overall traffic load and processing resources.

4.2. Context Management

Context information collected from all the nodes forming the MANET refers to their computational and physical environment and is tightly coupled with the policy-based management system since it is this information being monitored that may trigger a certain policy. Every node collects its own context information based on its available sensors. The term sensor is generic since it can refer to a battery monitor, a GPS receiver etc. A context model is needed to represent the collected information efficiently and accurately. Based on these requirements we propose exploiting Unified Modeling Language (UML) design principles for our context model, which is not presented in detail due to space limitations.

The model incorporates the notion of semantics to describe the context information, the sensors and their relationships. The general context of a node consists of higher level contexts that have been deduced from simpler ones i.e. mobility prediction of nodes deduced from device capabilities, GPS readings, personal diaries etc. The UML model is inherently associated with the data representation of the collected context. It allows for expressiveness and can be easily mapped to an XML document for interoperable storage.

Our system involves four components that deal with the context awareness requirements, the Context Collection Point (CCP), the Context Decision Point (CDP), the Context Management Tool (CMT) and the Context Repository (CR). There is an obvious matching of these components to the four elemental components of a PBNM system. The major design difference is that the flow of information is reverse to the one in PBNM systems, where a top-down approach is adopted. Here, context is collected and processed at the lower layers of the architecture and is passed to the higher layers for management decisions to be taken.

The CCP is installed on every node and its responsibilities include monitoring the environment through the available sensors and collecting the relevant information. After some basic pre-processing, context information is stored using the proposed model. This information is passed on from every CCP to their respective Cluster Head's CDP. The CDP aggregates the context information in order to gain cluster-wide knowledge of context of particular interest. It communicates with the local PDP entity to examine if context conditions are met so that the actions of a certain policy are triggered for the cluster.

The aggregated context information is also sent by the CHs to the MNs. The Context Management Tool (CMT) collects and aggregates the CH-aggregated context information and exchanges it with the other MNs. This ensures a MANET-wide common knowledge regarding context information. It is this context information that is passed from the CMT to the hyper-cluster PDPs and may trigger the appropriate policies. The previous two conditions ensure MANET-wide concurrent triggering of policies and thus network-wide adaptation.

4.3. Configuration Enforcement

In our PBNM system, PDPs monitor and evaluate their conditions against the context information they have received, in order to trigger the actions of the corresponding policies. These actions are enforced to the Managed Objects at the PEPs under the PDP's control. Our proposed architecture is built upon a programmable platform which is extended to cater for efficient provisioning of policy decisions to the PEPs. For more details regarding our programmable platform we refer the reader to our previous work [7][13]. The actions of a policy include altering dynamically the software currently installed on a node or activating/ deactivating software modules on-the-fly. Enforcement is achieved using the lightweight programmable middleware platform that uses autonomous software modules. These are installed in mobile nodes to extend or alter their functionality. Examples of such software modules include routing protocols, clustering algorithm implementations or any other service/ protocol the MANET requires. Within the examined trusted environment, module control is dictated by the introduced policies, thus ensuring controlled programmability and reconfiguration. In not trusted environments however security issues are raised, which currently lie outside the scope of our work.

5. Evaluation

We first evaluate our proposed management model in terms of performance and scalability. We then proceed in presenting our initial evaluation regarding our policy-based approach based on results obtained from our prototype implementation. The implementation process has not yet been completed but initial results are encouraging.

We used the ns2 simulator to evaluate the proposed hyper-cluster formation algorithm. We experimented with large node populations and topologies so as to assess in more detail the scalability of our approach. The transmission range of each node is set to 100 m, and the link capacity takes a value of 2 Mbps (worst-case scenario). The simulations were performed for a stationary MANET. In order to assess the effect of increasing network size on our dominating set clustering scheme, the terrain-area is accordingly increased, so that the average node-density is kept constant during simulations. The number of nodes in this case is varied from 25, 100, 225, 400 to 625 and the terrain-area size from 200x200 m^2, 400x400 m^2, 600x600 m^2, 800x800 m^2 to 1000x1000 m^2 respectively. Figure 6 shows the average hyper-cluster size during the clustering as a function of node population. As it can be deduced the increase in hyper-cluster size is almost linear to the increase in MANET size, which confirms the scalability of our approach.

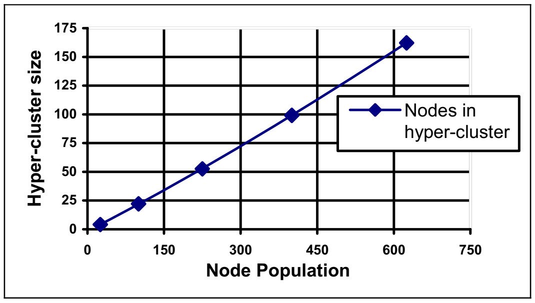

Figure 6. Hyper-cluster size as a function of MANET size

We then examine the time required for our distributed algorithm to lead to a stable hyper-cluster. Figure 7 presents the average time required for nodes to deduce their roles in the clustering process.

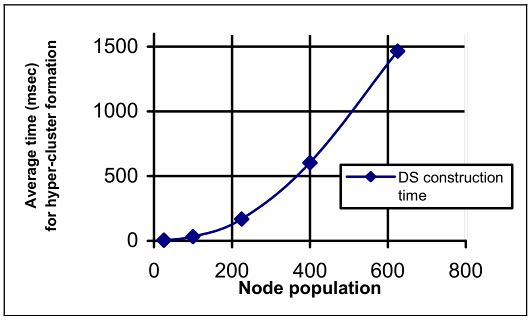

Figure 7. Average hyper-cluster formation time as a function of MANET size

Although the increase appears to be almost exponential, the time required for each node is in the range of 1465 msec when 625 nodes are considered which is acceptable. This is the time required for the whole clustering process to be completed since this is completely distributed. The sharp increase in time is attributed to the increased node connectivity when the node population increases. This leads to delays when neighborhood information is exchanged.

Another parameter we examined was the effect on the hyper-cluster size incurred by the TN to CM ratio. As is evident from Figure 8, the larger this ratio the less the number of nodes that form the hyper-cluster. This was to be expected since under these conditions only a small fraction of all nodes have the capabilities to participate in the hyper-cluster.

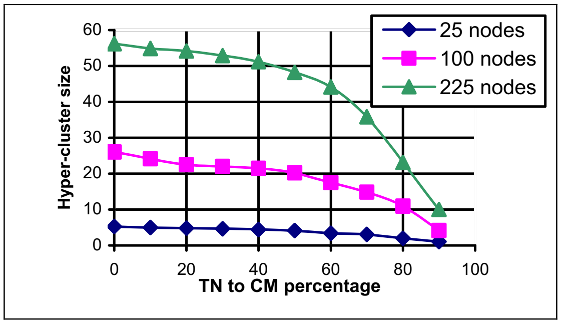

Figure 8. Hyper-cluster size as a function of TN / CM ratio

In addition to the simulation experiments, we performed traffic measurements related to the Distributed Policy Repository (DPR) deployment, which is one of the core and innovative parts of our proposed model. The purpose of these measurements was to evaluate the replication cost of the DPR component as well as the cost of synchronization maintenance. We used an openLDAP directory server on a Pentium M 1.4GHz laptop. Two instances of openLDAP emulate the DPR components of a Manager Node (MN) and a Cluster Head (CH).

First we assume that the MN possesses all MANET policies and a CH needs to be informed of the policies in order to create and load the respective Policy Objects. In traditional PBNM systems, a policy client at the CH would perform an LDAP search operation [19] to the MN's policy repository. Our scheme would instead automatically replicate the contents of the MN's DPR to the CH's DPR using the openLDAP replication engine (sync operation) [20], making the policies available at the CH for local access. The measurements are performed for different directory sizes. Our LDAP implementation for the presented policies requires 4 entries per simple policy rule instance while compound policy rules may require more entries. Therefore directory size in entries instead of rules is the independent variable depicted.

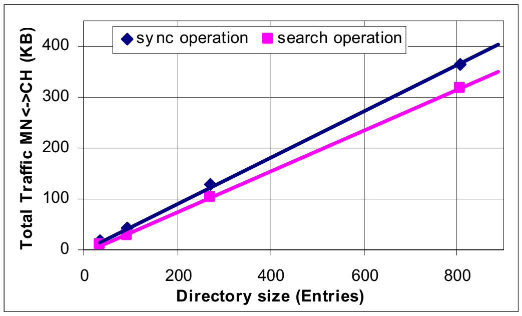

Figure 9. Generated traffic for full repository retrieval

Figure 9 shows the measured traffic between a MN and a CH for full repository retrieval. From the graph we observe an expected linear increase of traffic with respect to repository size for both methods, since all entries have approximately equal storage size. What is worth noticing is the slight traffic increase incurred from the sync operation compared to the search operation for the same directory size. There is a 12.2% increase in total traffic when using sync operation. Therefore we argue that the proposed replication scheme does not significantly increase generated traffic and the achieved benefits of distributing the repository among MANET nodes make this solution attractive.

A second case is examined, where new policies are added to the Manager Node's repository and the Cluster Heads should be informed. For this reason we experimented by introducing new rules and thus new entries to the active MN repository (Figure 10).

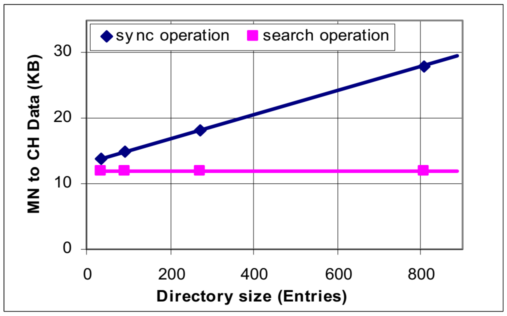

Figure 10. Generated synchronization traffic

Using the search operation the traffic cost of retrieval would be independent of repository size but the CH must be informed first about the location of the new policies, i.e. the rules' distinguished names [19]. Using the replication engine, the sync operation automatically disseminates all changes and new rules to the CH's repository. This comes at an additional traffic cost which was found to be linearly dependent on the directory size, since the sync operation carries additional presence information for all replicated entries. After measuring this traffic, we argue that is acceptable for a MANET environment as for 800 directory entries it would not exceed 30KB. The traffic difference of the two approaches is 0.018*DS (KB) where DS is the directory size in entries. Finally we measure the traffic cost of keeping the replicated directory in synchronization for different directory sizes. This cost under the examined conditions is 854 bytes and is required at configurable periodic intervals. Given that no change has occurred, this cost is independent of the directory size.

Based on the above measurements we argue that the replication scheme needed for the realization of the proposed DPR incurs small traffic overhead to the MANET and is a viable solution with added benefits.

6. Conclusions and Future Work

We have presented a novel policy-based management framework, specifically targeted at mobile ad hoc networks (MANETs). Based on the introduced distributed and hierarchical organizational model we deploy a replicated Distributed Policy Repository (DPR) among the hyper-cluster's nodes. By exploiting context information we can achieve effective evaluation of policy conditions that trigger actions with specified enforcement scope.

In order to assess the applicability of our concepts we have performed simulations and measurements. From the initial simulation's results we conclude that the proposed organizational model can scale well for increased node populations while initialization time is kept low. Based on traffic measurements, we evaluate the applicability of the replication scheme, needed to implement the DPR concept. The results indicate that the scheme can be deployed to resource-constrained MANETs since the traffic overhead is minimal.

Having implemented and tested key features of our proposed framework, we set forth our future objectives aiming to complete implementation and validation of our design. We will need to address policy conflict analysis and resolution, a critical issue in every PBNM system which is further amplified with the introduction of a multi-manager paradigm. Further improvements and evaluations will be made for our clustering algorithm. Additionally a dynamic election mechanism for the Manager Nodes will be investigated, in the light of novel reputation-based schemes. Finally we cannot overlook the inherent security implications and we will focus on securing the policy repository's content which essentially encapsulates the management logic of every policy-based system.

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

8. References

[1] Sloman M.,"Policy Driven Management For Distributed Systems", Journal of Network and Systems Management, Vol 2(4), Dec. 1994

[2] Strassner J., "Policy-Based Network Management, Solutions for the next generation", Morgan Kaufmann, ISBN: 1558608591, 2003

[3] Chen W., Jain N., Singh, S., "ANMP Ad hoc network management protocol", IEEE Journal on Selected Areas in Communications, Vol 17(8), Aug.1999

[4] Shen C., Srisathapornphat C., Jaikaeo C., "An adaptive management architecture for ad hoc networks", IEEE Communication Magazine, Vol 41(2), Feb.2003

[5] Chadha, R., H. Cheng, Y.-H. Cheng, Chiang, J., Ghetie, A., Levin, G., Tanna, H., "Policy Based Mobile Ad hoc Network Management", 5th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY2004), June 2004

[6] Phanse, K.S., DaSilva, L.A., "Protocol support for policy-based management of mobile ad hoc networks", IEEE/IFIP Network Operations and Management Symposium (NOMS), 2004

[7] A. Malatras, G. Pavlou, **"**Context-driven Self-Configuration of Mobile Ad hoc Networks"**,** IFIP International Workshop on Autonomic Communications (WAC 2005), October 2005

[8] Bellavista, P., Corradi, A., Montanari, R., Stefanelli, C., "Context-aware middleware for resource management in the wireless Internet", IEEE Transactions on Software Engineering, Vol 29(12), December 2003

[9] Yau, S.S., Karim, F., Wang, Y., Wang, B., Gupta, S.K.S., "Reconfigurable Context-Sensitive Middleware for Pervasive Computing", IEEE Pervasive Computing, July-September 2002

[10] P.-J. Wan, K. M. Alzoubi and O. Frieder, "Distributed construction of connected dominating set in wireless ad hoc networks", IEEE Infocom 2002

[11] R. Friedman, M. Gradinariu and G. Simon, "Locating cache proxies in MANETs", ACM MobiHoc 2004

[12] U. Kozat and L. Tassiulas, "Network layer support for service discovery in mobile ad hoc networks", IEEE Infocom 2003

[13] S. Gouveris, S. Sivavakeesar, G. Pavlou and A. Malatras, "Programmable Middleware for the Dynamic Deployment of Services and Protocols in Ad Hoc Networks", IEEE/IFIP Integrated Management Symposium (IM 2005), May 2005

[14] Westerinen A. et al, "Terminology for Policy-Based Management", RFC 3198, Informational, Nov.2001

[15] Moore, B., Elleson, E., Strassner J., Westerinen, A, "Policy Core Information Model-Version 1 Specification", RFC 3060, Feb.2001

[16] Moore B., "Policy Core Information Model (PCIM) Extensions", RFC 3460, Jan.2003

[17] Strassner J., Moore, B., Moats, R., Elleson, E., "Policy Core Lightweight Directory Access Protocol (LDAP) Schema", RFC3703, Feb.2004

[18] Pana M., Reyes, A. Barba, A., Moron, D., Brunner, M., "Policy Core Extension Lightweight Directory Access Protocol Schema (PCELS)", RFC 4104, Jun.2005

[19] Wahl, M., T. Howes and S. Kille, "Lightweight Directory Access Protocol (v3)", RFC 2251, Dec. 1997

[20] Zeilenga K.D, Choi J.H., "The LDAP Content Synchronization Operation", Internet-Draft, Sep. 2004