Manager Selection over a Hierarchical/Distributed Management Architecture for Personal Networks

Jose A. Irastorza, Ramón Agüero, and Luis Muñoz

Department of Communications Engineering – University of Cantabria, Spain
{angel,ramon,luis}@tlmat.unican.es

Abstract. In spite of having been the focus of several works, traditional management architectures, usually based on a centralized model, are not suitable for the particular characteristics of personal networks and their underlying multi-hop topologies. A hierarchical/distributed approach is proposed in this work, which also analyzes different strategies to optimally select the nodes taking the manager role. In order to assess the benefits and drawbacks of these mechanisms, a proprietary simulator was developed, and different metrics were studied (probability for a node to take part on the management architecture, number of hops needed to reach a manager, and fairness of the distribution of the management burden). A novel heuristic is proposed to enhance one of the analyzed strategies, and it is shown to outperform the rest of algorithms.

Keywords: Personal Networks, Management Organization Model, Distributed-Hierarchical Models, Algorithmic assessment.

1 Introduction

The evolution and proliferation of wireless devices and peripherals have been as remarkable as the aim of applying network technologies to interconnect them. The resulting communication architectures are conditioned by some particular characteristics of wireless environments: dynamicity of nodes, heterogeneity of the devices and involved technologies, as well as energy and bandwidth constraints. A typical scenario can embrace multiple kinds of devices, ranging from modern laptops to low cost, low capacity sensors and actuators, interconnected via e.g. a wireless multi-hop network, which might be spontaneously deployed over a geographically limited area, which can be referred to as personal surface. This illustrative scenario has been given the name of personal network (PN) [1]. Multi-hop (or mesh) network topologies are usually adopted to implement the subjacent connectivity over these network deployments.

One key issue of any type of network, and especially a personal network, comes from its management. Hence, the management tasks associated to personal networks can be seen as a challenge that must be tackled in order to facilitate the operability of these deployments. Although different network management architectures and models have been widely studied over fixed networks, it must be considered that managing a personal network poses several new difficulties, taking into account the specific characteristics of such deployments. First, communication links within wireless multi-hop networks are, intrinsically, unreliable, dynamic (due to node movements) and

K. Pentikousis et al. (Eds.): MONAMI 2010, LNICST 68, pp. 210–224, 2011.

showing varying capacity. Furthermore, nodes have several constraints, like limited battery and data storage capacity. The main consequence is that the resulting topology is unpredictable, and thus automatic on-demand reconfiguration procedures are usually required. These relevant facts pose some requirements to the network management task and, generally, to any service to be implemented over these networks. Questions like discovery protocols/procedures, autonomous topology and node reconfiguration, security and signaling overhead, just to name a few, must be resolved in order to efficiently manage a personal network.

The work presented in this paper can be focused on providing some answers to the "how to manage personal networks" question, by proposing an organizational model with an optimal distribution of the manager role to seek that a maximum number of nodes are managed, as well as balancing the management burden (fair distribution of agents between the available managers), so as to minimize the overhead introduced by the corresponding management traffic.

There are some other works which have analyzed the implications of multi-hop topologies over network management; two of the most referenced ones are the GUERRILLA framework [2] and the Ad Hoc Network Management Protocol [3], which propose a hierarchical approach. The former uses a clustering division among nodes, while the second one makes use of active probes which introduce some intelligence within the network. The work presented in [3] should be also mentioned, since it defines a complete information model, describing and implementing a prototype of a probe-based management architecture. One common aspect of these three previous works is that they analyzed both the information and the organization management models. There exist other works focusing their contribution on the organization model, paying special attention to management architectures based on a combination of both distributed and hierarchical approaches. Among these, some relevant works to be mentioned are [5], [6], [7], which are mainly based on specific clustering techniques, or [8], which analyzes how to optimally distribute management operations within the network.

The paper is structured as follows: Section II introduces relevant aspects about the organizational model for a management architecture, tailored to the specific characteristics of personal network environments. Section III discusses how manager and agent roles might be distributed between the network entities, identifying the parameters which should be taken into account so as to assess the appropriateness of the selection. In addition, different manager selection strategies are presented, ranging from rather pessimistic to almost optimum situations. Since the performance of those different strategies heavily depends on the particular characteristics of the network topologies, Section IV discusses the connectivity of random network deployments, paying special attention on the number of connected components (subgraphs) which might be expected for a particular network configuration. Afterwards, Section V, using an extensive analysis conducted over a proprietary simulator, evaluates the aforementioned strategies, discussing their benefits and disadvantages. Finally, Section VI concludes the paper, advocating some items which are left for future works.

2 Distributed/Hierarchical Management of Personal Networks

It is now believed that future personal networks will bring about some new requirements, including a much more relevant dynamicity, regarding their spontaneous configuration and self-managed topology.

Hence, in order to ensure an effective management of future personal networks, usually comprising multi-hop topologies and having a significant number of heterogeneous and mobile nodes, we propose using the adaptive and decentralized management paradigm. In this sense, the drawbacks of traditional management systems, characterized by centralized and static organizational models, are overcome.

Almost all current management frameworks are based on the traditional manager/agent model, which usually assumes a centralized control scheme, clearly inadequate to manage multi-hop or mesh topologies, since relying on a unique central entity to manage the whole network is not a sensible alternative, provided that nodes are intermittently connected. In such type of configurations, the manager would represent a single point of failure, and this is not admissible in future personal networks.

Another well known limitation of this centralized approach is its almost non-existing scalability; in this sense, having a single manager station might lead to heavy management traffic exchanged between this and the corresponding agents, as well as a high processing load on this particular node, which could lead to long execution times for management operations. This becomes even worse over personal networks; as it is assumed that the relevance of management operations (accounting, security, etc) would be quite high over such networks, and thus, the resulting overhead might heavily augment, decreasing the network performance. To sum up, we may say that the intrinsic characteristics of personal networks and multi-hop/mesh topologies, such as temporary links, sparse bandwidth and limited resources, impose a different approach rather than a traditional centralized architecture for the management framework.

There have been already some proposals aimed at overcoming these drawbacks, proposing some refinements to the basic centralized scheme. On of the most relevant ones is the so-called Management by Delegation (MbD) approach [9], which fosters the delegation of some tasks from the management station to the corresponding agents by means of dynamic downloadable scripts. Other extensions of this centralized scheme are based on the establishment of certain agent hierarchies, enabling direct interactions between agents. Nevertheless, these refinements (which are still based on a centralized approach) do not efficiently address all the previously enumerated shortcomings; on the contrary, the use of management architectures based on a decentralized operation and relying on self-management mechanisms could become a more sensible choice. Following this idea, in this work we present an organizational model based on a distributed and hierarchical model.

The proposed management framework is logically structured according to a three-level hierarchy, composed by a top level manager, which could be selected from a number of second-level managers. These take a localized manager role, controlling a set of nodes which could be seen as a cluster (characterized by some sort of connectivity between its components). The agents are thus the third level of the hierarchy. In spite of the three defined levels, there are only two management communication

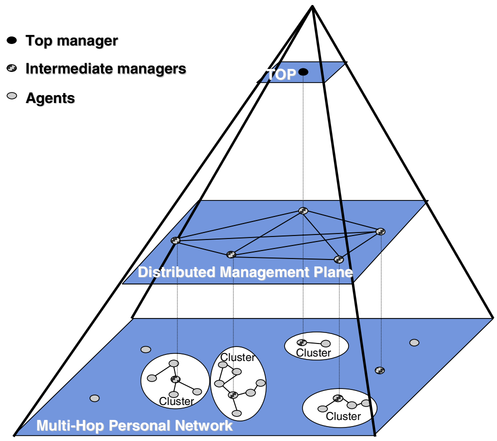

Fig. 1. Distributed/hierarchical management organization model for personal networks

planes: one established between the agents and their corresponding manager (second level); and another one which interconnect all the second level managers between them and with the overall manager (first level). As can be seen on Figure 1, this manager plane creates an overlay network on top of the PN which interconnects the second level managers by means of virtual links, each of which might correspond to a particular path, which could comprise several physical links, either over the underlying network or using another parallel communication facility.

The proposed architecture entails a distributed management plane (overlay network) with a number of nodes which take a manager role, each of them would control a subnetwork (or network component), communicating with the rest of managers in a peer to peer and collaborative manner. Using this distributed approach, the network management subsystem would achieve higher reliability and efficiency, as well as less overhead, on both communication and system resources.

In addition to that distributed plane, the management architecture also presents a hierarchical approach, since the manager role is distributed between two different levels: the top one represents overall network manager, while the 2nd level managers can be referred to as intermediate managers. Each of them controls its own domain (network component or cluster), collecting and processing information from the corresponding agents and forwarding such data to the upper level manager, if necessary. It also delivers management information from the top manager to its own domain nodes. As can be seen, the depicted management framework follows a distributed/hierarchical organization model.

3 Problem Statement

For the rest of this section we will assume that the network comprises a set of N nodes, M of which take the manager role while the resulting A (N-M) are agents. In addition, we will assume that the number of covered (able to access, at least, one manager) agents is A_C.

We have seen in the previous section that most of the sensible choices for the management architecture to be employed over a wireless multi-hop topology (as the one which we could map on a personal network) promote the balancing of the corresponding manager load between a set of nodes, thus bringing about hierarchical/distributed management architectures. Hence, the problem to be addressed is how to select this optimum set of managers. In order to be able to assess the suitability of the selection, we need to establish a number of aspects of merit. Their combination is what would yield the optimum selection strategy. Below we discuss three of the most relevant ones, which are the ones that will be later evaluated.

- **Covered probability.** this is the most obvious one, and refers to the percentage of nodes which are able to access a manager and, thus, can participate in the management architecture. The goal would be to reach a full coverage, meaning that all nodes are either managers or are covered by, at least, one manager.
- **Average number of hops.** one of the classical problems associated with multi-hop networks is the interference that communications can cause. In order to avoid a serious increase of the overhead brought about by the management traffic, it would be interesting having a small number of hops between each node and its corresponding manager.
- **Agent distribution.** the main goal of balancing the management burden is to avoid concentrating too much traffic into a single node; in this sense, the selection of managers should try to provide a fair distribution of the agents between them. In order to measure this aspect, we introduce the following parameter:

$$\beta = \frac{1}{M}\sum_{m}\frac{\left|A_m - \frac{A_C}{M}\right|}{\frac{A_C}{M}} = \frac{1}{A_C}\sum_{m}\left|A_m - \frac{A_C}{M}\right| \tag{1}$$

In the previous expression, AC accounts, as was already said, for the overall number of covered agents, while Am is the number of agents which are covered by the particular manager ‘m’. Therefore the β parameter can be described as the relative difference between the optimum distribution (in which all managers have the same number of associated agents, AC/M) and the current one. The lower this parameter is, the closer it is with the fairest distribution.

In addition to the three parameters described before, we will also look into the combination of the first two, by studying the probability of being managed when a limit on the number of hops to be used to reach a manager is assumed.

The aforementioned figures of merit will be evaluated and studied using different manager selection strategies. Below, a description of each of them is provided.

- **Strategy 1 Ramdom manager deployment.** In this case, we assume that the M managers are randomly selected, without any kind of previous planning. This reflects a quite unlikeable situation since, depending on the network characteristics, there might be managers without any node within their coverage area. From an implementation point of view, this option poses no major difficulties.
- **Strategy 2: Topology-agnostic optimal manager deployment.** In this case, we assume that managers are placed in those points which ensure a maximum

"geographical" coverage of the whole area. This does not mean that the number of covered agents is maximized, since this would depend on their particular position in any particular network instantiation.

- **Strategy 3: Topology-aware optimal manager deployment.** In this case, the current topology of the network is used to optimally assign the managers. In order to accomplish this, we solve the p-median problem [10]. The p-median aims at finding a set of p managers from the overall deployed nodes which minimizes:

$$\sum_{j \in N} \{\min_{i \in M} d_{ij}\} \tag{2}$$

In our case, and without loss of generality, we have used a constant cost per link, so the overall cost between nodes i and j can be seen as the number of hops between them (provided that a path exists between them). Hence, the p-median problem would aim at minimizing the overall distances needed by all nodes to reach a manager; provided that all agents are covered (all the demand is satisfied).

- **Strategy 4: Topology-aware sub-optimal manager deployment.** As briefly introduced before, one of the main drawbacks which is usually attributed to the p-median problem is that its first goal is to cover all nodes [10] (i.e. all the demand should be covered). Depending on the network topology, there might be situations in which it could be more appropriate not managing some nodes, if that jeopardizes the rest of the assignment strategy. In order to better resolve this issue, we propose a simple heuristic, in which we do not account for those subgraphs having a reduced number of nodes, resolving the p-median problem only over the rest of the network. In this case, an additional design parameter is the size of the subgraphs to be discarded. A trade-off must be done between the nodes which are lost and the additional benefit which can be obtained by letting them out of the management architecture. Clearly, deleting isolated nodes is a sensible option, since those nodes are not able to communicate with the rest anyway. However, it has to be discussed whether it is sensible deleting subgraphs with a greater number of nodes.

Figure 2 aims at providing an illustrative example of the behavior of the four strategies. In this case, the network consists of 50 nodes, and 7 managers are deployed. Furthermore, Table 1 presents the particular figures which were obtained for each of the parameters which are being analyzed. The two first strategies (top part of the figure) are rather straightforward. First, the random deployment leads, in this particular case, to a relevant number of nodes which are left uncovered, while there are some managers which are unnecessarily close between them.

On the other hand, the second strategy, characterized by an optimal geographical placement of the 7 managers, ensures a high coverage, although it might lead to having some uncovered nodes, as those which are between the two lowest managers. In order to discuss the other two strategies, it is worth noting that the particular network example which has been used yields seven subgraphs, two of them consisting of isolated nodes. In this case, the advantages and drawbacks of the two latter manager selection approaches are clear: strategy 3 offers (for this particular network deployment) full coverage, since it devotes a manager to each of the existing subgraphs; however, the distribution is not very fair, since there is one manager which needs to handle a large number of nodes, which require, in some cases, a larger

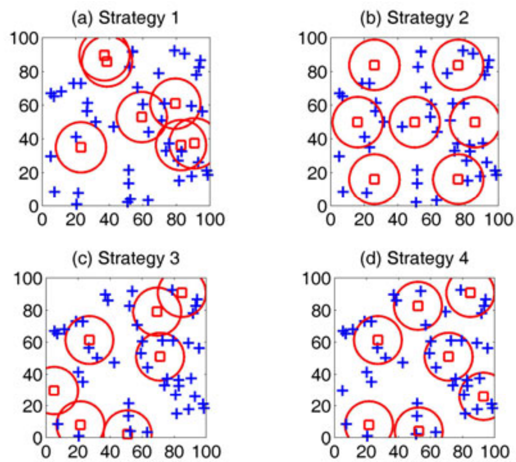

Fig. 2. Manager selection strategies. Squares represent managers (M=7) and crosses represent agent nodes.

Table 1. Summary of manager selection strategies results

Strategy	β	Avg_hops	Disconnected	2hops_covered
1	0.661	2.143	0.349	0.465
2	0.336	1.579	0.116	0.837
3	0.963	1.977	0.000	0.721
4	0.445	1.341	0.047	0.907

number of hops to reach the corresponding manager. Besides, strategy 4, in which the two isolated nodes are not considered, fails to provide a full coverage, but, on the other hand, it ensures a fairer distribution of the nodes between the available managers, thus leading to routes with a fewer number of hops. As can be seen on Table I, the connectivity is higher for strategy 3, but if we limited the maximum number of hops to reach a manager to 2, then the covered probability of strategy 4 clearly outperforms the p-median based one.

4 Connectivity of Random Network Deployments

Since one of the main goals of the analysis is to establish the number of managers which should be deployed (for the different analyzed strategies) in order to guarantee a certain connectivity degree, it is worth analyzing the connectivity of the subjacent network scenario. This would provide the information required to better understand the corresponding results. In this sense, we introduce a connectivity degree parameter (ζ), which accounts for the number of subgraphs (SG) of a particular deployment of N nodes.

$$\xi = \frac{N - SG}{N - 1} \tag{3}$$

From the above expression, it is straightforward that in a fully connected network ζ equals 1, while for a network in which there is not any connection amongst the nodes, ζ equals 0. E.g. the connectivity of the network ranges from 0 (all nodes are isolated) to 1 (there is, at least, one possible path between any pair of nodes within the network). Furthermore, and in order to be as generic as possible, we define the normalized coverage (ρ), as the ratio between the particular coverage of the Radio Access Technology and the side of the area under analysis (this allows us carrying out the most generic analysis as possible). In that sense, if we maintain the number of nodes (N) and ρ constant, the connectivity (as it has been defined) should not change, no matter the actual dimension of the area under analysis is. This means that it is possible establishing a function g so that ζ = g(N, ρ) (without depending on the particular area dimension). However, in many cases, network deployments are characterized by node density (D) and the coverage (R) of the subjacent technology. It would be also appropriate being able to establish a similar function, taking D and R as its arguments. Let's assume that we have a squared area (side Li), then:

$$\begin{aligned} D_i &= \frac{N}{L_i^{\,2}} \\ R_i &= \rho \cdot L_i \end{aligned} \tag{4}$$

By inspection, we can see that, provided we maintain N and ρ constants, the product of D and R^2 is also kept constant. In fact:

$$D_i R_i^{\,2} = \frac{N}{L_i^{\,2}} \rho^2 L_i^{\,2} = N\rho^2 = \Lambda \tag{5}$$

Hence, if we define the Λ parameter, as the product of N and ρ^2, or, equivalently, as the product of D and R^2, we should seek for a function f such that ζ = f(Λ).This result would be of outer relevance, since it would allow a quick estimation of the particular network configuration in order to reach a required connectivity degree; this might be quite beneficial, e.g. when deploying wireless sensor networks to cover a particular area.

Figure 3 shows the connectivity degrees which were obtained for the different (N, ρ^2) pairs, after a simulation analysis which consisted of 400 different scenarios, and 1000 independent runs per scenario. As it was expected, it is possible establishing the following relationship between the two parameters:

$$N\rho^2 = \Lambda \;\rightarrow\; \log(N) + \log(\rho^2) = K \tag{6}$$

Where Λ and K are constants and must be functions of ζ, which also linearly depends on log (N) + log (ρ^2). Hence, the problem is to find the relationship between Λ and ζ.

Before that, Figure 4 shows the standard deviation which was observed for the 1000 independent runs which were executed per scenario. As can be seen, the variation of the results is higher (around 0.15) for network deployments having fewer nodes; however, the standard deviation quickly decreases and the behavior for relatively large values of Λ is pretty much predictable.

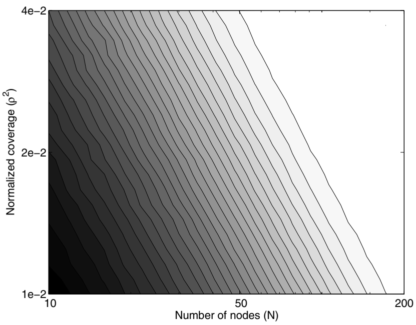

Fig. 3. Connectivity degree (ζ) for various (N,ρ^2) pairs. ζ ranges from 0.13 (darker) to 1.00 (lighter).

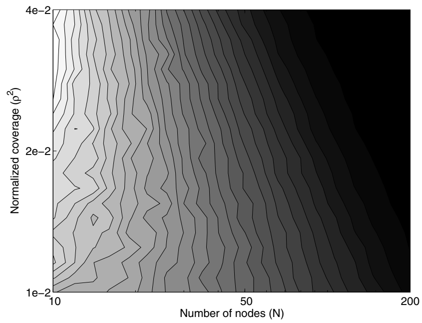

Fig. 4. Standard deviation of the connectivity degree (ζ) for various (N,ρ2) pairs. STD(ζ) ranges from 0.00 (darker) to 0.15 (lighter).

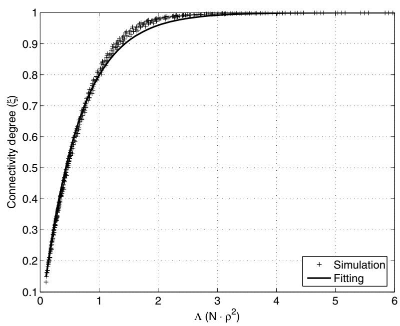

Fig. 5. Relationship between the network topology (Λ parameter) and its connectivity (ζ)

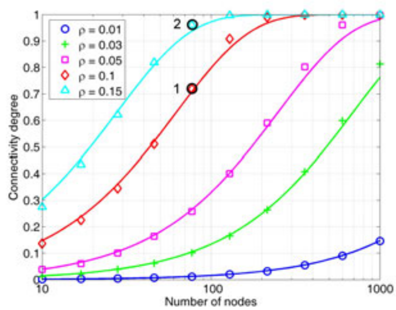

Fig. 6. Connectivity degree (ζ) for various node densities and normalized coverages (ρ)

The conclusion is that the estimation of the connectivity degree can be accurately estimated based on the particular characteristics of the network deployment (number of nodes, coverage of the radio access technology and dimensions of the area under analysis).Now the question is to find the expression that better fits the results which were obtained in the different simulation runs. Figure 5 shows the different (Λ,ζ) pairs which were observed during the Montecarlo procedure, as well as the exponential function which better fits them (LMSE is less than 0.02). The resulting parameter of such function is 1.6, so we can write:

$$\xi \approx 1 - e^{-1.6\Lambda} = 1 - e^{-1.6N\rho^2} = 1 - e^{-1.6DR^2} \tag{7}$$

Figure 6 shows the connectivity degree for various values of the normalized coverage. The lines have been obtained with Eq. [7], while the markers are empirically acquired with a Montecarlo simulation comprising 1000 runs per scenario. As can be seen the values which are provided by the proposed expression are rather close to the real ones, assessing the validity of such expression to estimate

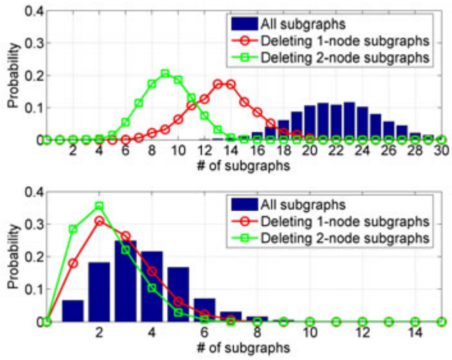

Fig. 7. pdf of the number of subgraphs for a network of 80 nodes over the two selected scenarios (top figure for ζ=0.7 and bottom graph for ζ=0.9)

the connectivity degree. This allows us to better select the parameters of the scenarios over which we will analyze the manager selection strategies. In this sense, in order to have a more thorough comparison of the different manager selection strategies, we select two illustrative network deployments (see Figure 6): (1) a sparse network, in which ζ is around 0.7, and (2) a high-connected scenario, in which ζ is higher than 0.95. We ensure these two scenarios by fixing N to 80 and using two normalized coverages: 0.10 (Λ = 0.8) and 0.15 (Λ = 1.8), respectively.

The last step before discussing the obtained results is to establish the design parameter of the fourth algorithm. In that sense, Figure 7 shows the probability distribution function (pdf) of the number of subgraphs which were encountered over the two scenarios which will be used during the evaluation of the different alternatives.

The number of subgraphs which are reflected on Figure 7 corresponds to those which could have been expected by using Eqs. [3] and [7]. First, in the case of the connected network, Λ equals 1.8, which yields a connectivity of ζ = 0.95, thus resulting in an average 4.95 subgraphs. On the other hand, for the sparse network, in which Λ equals 0.8, the connectivity is, according to Eq. [7], around 0.72, which yields 23.1 subgraphs, corresponding with the histogram which is represented in Figure 7. It has to be mentioned that, according to the differences between simulated and analytical values, the approximation is slightly better in the case of the sparse network.

Besides, the figure also permits inferring the potential advantages when discarding components of a particular size. As can be seen, by cutting components of 1 and 2 nodes, the additional benefit is rather notable, since with 5 managers the whole network will be covered in more than 90% of the cases for the high connected scenario. On the other hand, for the sparse case, since it is quite likely to have subgraphs of reduced size, there is a clear benefit when leaving out of the analysis those subgraphs (as can be yielded from the graph, the decrease on the number of resulting components is much more relevant than in the previous scenario). Furthermore, by cutting those network chunks, we ensure that the size of the remaining subgraphs would be much bigger and the p-median problem might provide better results, especially in terms of the number of hops which are required to reach a manager, since more managers will be available for the remaining subgraphs (this is true for the two scenarios). Hence, we will assume that if the number of nodes which are reachable is fewer than 3, then a node would not take the manager role (in fact, it will not belong to the managed network).

5 Discussion of Results

This section discusses the results which were obtained when applying the previously defined algorithms to assign the manager role, based on an extensive simulation campaign. A Montecarlo method was applied, and the metrics which were previously described were analyzed.

The simulation setup was based on a proprietary simulator, while the p-median problem was solved with the popstar tool [9]. Since the main goal of this work is to analyze the manager selection strategies, an ideal propagation channel was assumed. Incorporating more realistic models would not add too much complexity to the simulator, since an average coverage might be used in such cases; in addition, by pursuing this approach, we ensure that the analysis could be easily extended to other type of scenarios, e.g. wireless sensor networks. We assume two simulation areas, over which we randomly (Poisson Point Process) deploy N=80 nodes. The first scenario (sparse network – $\zeta \approx 0.7$) is characterized by a normalized coverage of $\rho=0.1$, while in the second one (high connected network – $\zeta > 0.95$) the normalized coverage is $\rho=0.15$.

The next step was to select the managers, following the four previously discussed algorithms; afterwards, an analysis of how the agents would be distributed between them was performed. In this sense, we will study the additional benefit of increasing the number of managers for all cases, i.e. M will be increased from 1 to 10. Each individual setup was simulated 200 times, in order to get tight confidence intervals.

Figure 8 shows the probability for an agent to be covered by, at least, one manager. As can be seen, the worst manager deployment strategy leads to the smallest coverage probability. On the other hand, there is not a remarkable difference between strategies 3 (legacy p-median) and 4 (cutting the 1-node and 2-node network subgraphs), staying below 3% for all cases. The effect of deleting such nodes appears in the high connected scenario, as the coverage probability asymptotically leads to a value slightly smaller than 1 for the fourth strategy, while a full coverage is almost reached with 5 managers for the legacy p-median case. The results also show that the topology-agnostic approach does not reach the same coverage as the one which can be achieved by means of the p-median based algorithm. Last, it is worth referring to the benefit which is achieved over the sparse network with strategies 3 and 4; in this case, the difference between these two manager selection algorithms is even less relevant than the one which is observed for the high connected network.

One of the benefits that the proposed heuristic should have over the legacy p-median is that it should achieve a fairer distribution of the agents between the different managers. Figure 9 shows how the different strategies distribute the management burden between the selected managers, representing the β parameter for

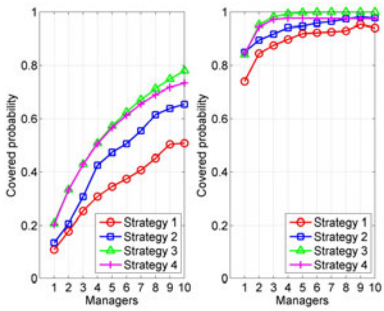

Fig. 8. Coverage probability for the four manager deployment strategies and the two connectivity scenarios (left ζ=0.7 and right for ζ=0.9)

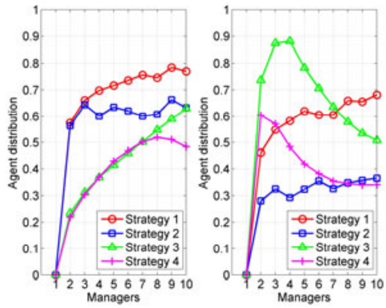

Fig. 9. Distribution of agents between managers for the four manager deployment strategies and the two connectivity scenarios (left ζ=0.7 and right for ζ=0.9)

the four algorithms. We shall observe an important result for the high connected network, since the third strategy (applying the p-median over the whole network) provides similar performance than the worst one (in terms of coverage probability, i.e. strategy 1), at least when the number of managers is reasonably low. This is the consequence of reserving managers to cover all subgraphs, no matter their size is. However, with the proposed modification (strategy 4), the distribution is strongly improved, and it almost reaches the behavior of the best solution (topology-agnostic, strategy 2) when the number of managers is greater than 4. For the sparse network, there is not a clear difference between strategies 3 and 4 (providing, both of them, a better performance than the others), although the graph shows that the tendency of the measured parameter tends to decrease for $M > 8$ for strategy 4, while for the legacy p-median approach the parameter still shows a growing trend.

The third aspect which was identified as a parameter to be optimized is the average number of hops which are required to reach a manager. The fewer the needed hops, the less overhead which would be caused by the management traffic. As can be seen in Figure 10, the heuristic presented in this paper again shows a very similar behavior than the second strategy for the high connected scenario, being slightly better for more than 4 managers. Obviously, the random selection approach provides the worst results, while for the p-median case, the graphs yields that for a small number of managers, the analyzed parameter is rather high (almost reaching the values obtained for the worst case); this is due to the fact that the p-median main goal is to cover all nodes (as mentioned earlier) and thus, it establishes managers in all subgraphs, with the consequence that in those ones in which the number of nodes is greater, more hops are needed to reach the selected manager, since there are fewer managers to serve those network components. For the sparse network, the proposed heuristic also shows the best performance, although in this case it is comparable with the legacy p-median. However, when the number of managers is higher than 8, there is a change in the corresponding trends (the benefit of the novel heuristic starts to be more relevant), due to the fact that the average number of resulting subgraphs is between 8 and 9, as was shown in Figure 7.

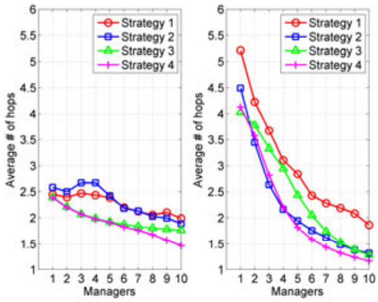

Fig. 10. Average number of hops required to reach a manager for the four manager deployment strategies and the two connectivity scenarios (left ζ=0.7 and right for ζ=0.9)

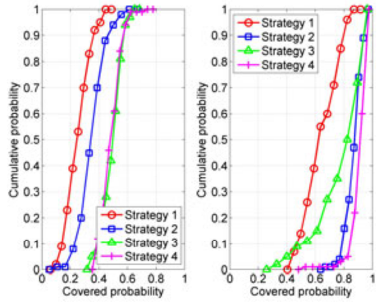

Fig. 11. cdf of the probability to reach, at least, one manager when the maximum number of hops is set to 3. The number of managers to deploy (M) is set to 5. Left figure: sparse network (ζ=0.7) and right figure (high connected network (ζ=0.9).

One conclusion which could be derived by analyzing the results of both Figures 8 and 10 is that, if a bound was set on the maximum number of hops which could be used to reach a manager, the behavior of the proposed scheme might be even much better than the rest of analyzed strategies. As can be seen in Figure 11, which shows the cumulative distribution function (cdf) of the probability of being connected to a manager using a maximum of three-hop routes (and fixing the number of managers to deploy, M, to 5), the proposed solution is the one which offers the best performance for the high connected network, since the probability of being connected with routes of either one or two hops is much higher with this strategy. In this case it is worth highlighting the improvement compared to the traditional p-median. Although this solution was able to provide a higher coverage (since it did not cut any node from the network), this benefit would not be perceived if a limit was established on the maximum number of hops which could be used to reach a manager. For the sparse network scenario, the difference is not that remarkable, since in this case we might need to deploy more managers to start noticing the improvements brought about by the proposed approach, as it has been discussed above.

The previous results show that the proposed heuristic provides a quite good performance for selecting managers on a hierarchical/distributed management architecture which could be used over forthcoming personal network scenarios. The tests carried out with two illustrative network deployments prove that this behavior can be extrapolated to other scenarios.

Another important aspect to be mentioned is that, although the analysis which has been performed is fundamental, we have kept an implementation perspective in mind, and e.g. the size of the subgraphs to be cut from the network would allow to bring the algorithm to a straightforward implementation, since the required information may be acquired by means of neighboring discovery processes, which are expected to be part of any personal network system.

6 Conclusions

This paper has tackled the management of personal networks. Opposed to the traditional centralized model, we have proposed a distributed/hierarchical architecture, in which the manager burden is shared between a number of nodes, thus making the whole system more scalable and alleviating the overhead associated to management traffic. These managers are together joined in an overlay network, used to interchange management related information. Furthermore, an optional higher level may be added, since it may be interesting that some individual nodes are able to gather all the information to be managed.

One of the most challenging issues which arise in the previously depicted architectures is how to carry out an appropriate selection of the manager roles. This work has specifically looked into this problem. We have identified different metrics which should be analyzed to assess the appropriateness of the selection, accounting for how a fair sharing of the management burden is fulfilled, the probability for any node to participate within the management subsystem and the resulting overhead (based on the number of hops which are required to communicate with the respective manager). Based on these metrics, different strategies have been analyzed: a rather pessimistic algorithm, in which the managers are randomly deployed, and two optimum approaches, based on the geographic position of the manager nodes and on the particular network topology. For this latter case, an enhancement, based on a novel heuristic, has been proposed, by eliminating those network components (subgraphs) with a relatively small number of nodes.

The results have yielded two main conclusions: on the one hand, it is of outer relevance to facilitate an appropriate selection of the managers, since there might be relatively large differences depending on the particular selection strategy; in addition, the proposed heuristic provides very interesting results, since it brings about much better performance in terms of agent distribution between the selected managers, as well as considering the number of hops which are used to reach the corresponding manager, while it does not severely degrade the coverage probability. In order to account for a wide range of potential scenarios, the analysis has been conducted using two illustrative network deployments.

The research which has been conducted in this work has been based on network graphs, but a major consideration has always been the possibility of mapping the proposed algorithms on top of real protocols and network deployments. In this sense, and thanks to the management simulation framework which was presented in [12], future work will analyze the implications of the proposed management architecture as well as the manager selection architectures on the communication and protocol performance. Regarding this future line of research, there are some interesting works available in the literature, mostly dealing with wireless sensor networks [13-14].

Acknowledgments

The authors would like to express their gratitude to the Spanish government for its funding in the following two projects: Mobilia - CELTIC Program (Avanza I+D TSI-020400-2008-82) and "Cognitive, Cooperative Communications and autonomous SErvice Management", C3SEM (TEC2009-14598-C02-01).

References

1. Lu, W., Gu, Y., Prasad, R.V., Lo, A., Niemegeers, I.: A Self-organized Personal Network Architecture. In: Third International Conference on Networking and Services, ICNS 2007, June 19-25, pp. 36–36 (2007)
2. Shen, C.-C., Srisathapornphat, C., Jaikaeo, C.: An Adaptive Management Architecture for Ad Hoc Networks. IEEE Communications Magazine 41(2), 108–115 (2003)
3. Chen, W., Jain, N., Singh, S.: ANMP: Ad Hoc Network Management Protocol. IEEE Journal on Selected Areas in Communications 17(8), 1506–1531 (1999)
4. Badonnel, R., State, R., Festor, O.: Management of Mobile Ad Hoc Networks: information model and probe-based architecture. International Journal of Network Management 15(5), 335–347 (2005)
5. Sivavakeesar, S., Pavlou, G., Liotta, A.: Stable Clustering Through Mobility Prediction for Large-Scale Multihop Intelligent Ad Hoc Networks. In: Proc. of WCNC 2004, Atlanta, USA (March 2004)
6. Fallon, L., Parker, D., Zach, M., Leitner, M., Collins, S.: Self-forming Network Management Topologies in the Madeira Management System. In: Bandara, A.K., Burgess, M. (eds.) AIMS 2007. LNCS, vol. 4543, pp. 61–72. Springer, Heidelberg (2007)
7. Badonnel, R., State, R., Festor, O.: A Probabilistic Approach for Managing Mobile Ad Hoc Networks. Transactions on Network and Service Management 4(1), 39–50 (2007)
8. Lim, K.-S., Adam, C., Stadler, R.: Decentralizing Network Management. KTH Technical Report (December 2005)
9. Yemini, Y., Goldszmidt, G., Yemini, S.: Network Management by Delegation. In: Second International Symposium on Integrated Network Management IM 1991, Washington, D.C., pp. 95–107 (April 1991)
10. Resende, M.G.C., Werneck, R.F.: A hybrid heuristic for the p-median problem. Journal of Heuristics 10(1), 59–88 (2004)
11. Canós, M.J., Ivorra, C., Liern, V.: Fuzzy p-median problem: A global analysis of the solutions. European Journal of Operational Research 130(2), 430–436 (2001)
12. Irastorza, J.A., Agüero, R., Muñoz, L.: Fostering the simulation-based evaluation of management architectures over multi-hop topologies. In: IEEE/IFIP Network Operation and Management Symposium (NOMS 2008), Salvador do Bahia, Brazil (April 2008)
13. Ruiz, L.B., Silva, F.A., Braga, T.R.M., Nogueira, J.M.S., Loureiro, A.A.F.: On impact of management in wireless sensors networks. In: IEEE/IFIP Network Operations and Management Symposium, NOMS 2004, April 19-23, vol. 1, pp. 657–670 (2004)
14. Ruiz, L.B., Nogueira, J.M., Loureiro, A.A.F.: MANNA: a management architecture for wireless sensor networks. Communications Magazine 41(2), 116–125 (2003)