CELTIC Initiative Project Madeira: A P2P Approach to Network Management

Martin Zach[1], Daryl Parker[2], Liam Fallon[2], Christian Unfried[1], Miguel Ponce de Leon[3], Sven van der Meer[3], Nektarios Georgalas[4], Johan Nielsen[5]

[1] *Siemens AG Austria*
{martin.zach, christian.unfried}@siemens.com

[2] *Ericsson R&D Ireland*
{daryl.parker, liam.fallon}@ericsson.com

[3] *Telecommunications Software & Systems Group, Waterford Institute of Technology*
{miguelpdl, vdmeer}@tssg.org

[4] *BT Group*
nektarios.georgalas@bt.com

[5] *Ericsson Research*
johan.nielsen@ericsson.com

Abstract

The vision of the Celtic-Initiative project Madeira is to provide novel technologies for a logically meshed Network Management System that facilitates self-management and dynamic behaviour of nodes within networks. These approaches should enable adaptable services and the management of network elements of increasing number, heterogeneity and transience, thereby reducing OPEX. In this paper, we set the scope for investigations within the project and give an outline of our approach. We present a scenario that challenges today's state of the art in Network Management and upon which we are building our case study for a detailed investigation of feasibility. Finally, we describe a preliminary conceptual system architecture and application data model, and give an insight into the expected final project results.

1. Introduction

In today's telecommunications systems, management is mostly achieved through an inflexible, static architecture. As networks are becoming larger and more heterogeneous and exhibit dynamic behaviour, this approach is inadequate. The Madeira project addresses one of the key issues within the telecommunications management sector: the inability of existing Network Management Systems to adapt and evolve towards the service and network system requirements of large-scale ubiquitous networks.

Currently deployed management systems are constrained by rigid interoperability standards such as M.3100 [1] and SNMP [2]. They use static rather than dynamic or self-aware control paradigms and rigid architectures that restrict flexibility and distribution, resulting in implementations with limited capabilities to cater for very large numbers of managed entities and dynamically forming networks of transient elements.

Taking middleware and management as an example, two different aspects can be seen today: Distributed systems access classic management systems and legacy systems are used to manage distributed applications. Here, a few middleware concepts and many different products are used in parallel. Furthermore, management platforms are used as a layer of abstraction for business level management. Middleware concepts, object specifications, protocols and data formats are designed specifically to support distribution, not management functionality. By integrating both aspects, the Madeira project aims to provide a holistic view (framework) following principles e.g. described in [3].

The objective of the Madeira project is to investigate large-scale distribution techniques in Network Management. Our approach is to build a prototype system using the peer-to-peer paradigm [4]. We will show how this architectural technique can be used to solve challenging management problems in next-generation networks and analyse how such an approach, exploiting P2P characteristics such as self

organisation, symmetric communications and distributed control [5][6], will lead to significant reductions in OPEX and a more adaptive network control than in today's systems.

The Madeira project aims at providing an innovative architectural framework, requisite interface protocols, standards, and a reference software implementation. In particular, Madeira will develop a case study focusing on the relationship between fault and configuration management in a dynamically forming network with transient elements, considering the perspectives of the equipment vendor, the service provider, and the network operator.

In the following sections, we first outline the Madeira approach in more detail, then discuss a scenario that challenges current Network Management approaches, and present a high level (preliminary) architecture of our system. Finally we summarise the current status of our work in Madeira so far, and give an insight into the expected final project results.

2. Our Approach

In this section we give a short overview on our approach within the Madeira project and establish the context of the preliminary results presented in this paper.

The starting point for the definition of our system architecture is the analysis of requirements taken from various sources. The top-level requirements originate from the scope as defined in the Madeira project description [7]. These requirements encompass the type and characteristics of the networks, operational goals, and the design paradigms and principles that we wish to evaluate. The requirements are analysed at a high level in an "Architectural Spike" described later in this paper. These sources allow us to derive both architectural and interface requirements.

The above analysis resulted in the selection of the following categories of requirements that are of particular interest to the project and capture the critical issues to be addressed by the Madeira architecture.

Table 1: Important categories of architectural requirements for Madeira

Distributability	Self-Awareness & Observability	Auditability
Scalability	Autonomy	Information Consistency
Interoperability	Security & Trust	Messaging

Guided by this categorisation, lower-level requirements (both functional and non-functional) are derived. The detailed functional requirements are derived from the scenario investigated in the case study. The case study is focused on examining a vertical slice of a selected Network Management area in the context of the Madeira architecture, from the Network Element right through to the Network Management System. This scenario, and the challenges it poses for Network Management, will be described in more detail in the following section.

A core aspect of Madeira is the application of a Model Driven Architecture (MDA) approach [8] to data modelling and manipulation, and showing how Network Management data activities can be captured and represented in a truly distributed, asynchronous system. The objective for modelling is the integration of communications management and software platform middleware concepts to Network Management, while maintaining independence from concrete software technologies. This serves as a basis to develop communication service platforms with integrated management facilities that enable communication applications, services and resources to be used, controlled, operated, administered and maintained in a unified way. Two major activities, information mapping and system management, will offer transformation mechanisms to map Network Management data information across identified levels and manage entities of those levels.

Having the requirements in place – as a result of an ongoing iterative process – and the modelling approach defined, subsequent design iterations are being outlined, following proven architectural strategies [9][10]. The results of the first iteration of our modelling approach and system architecture (named "architectural spike") are presented in section 4.

Finally, a proof-of-concept prototype realising the scenario identified in the case study will develop a vertical slice of Madeira concepts in operation, using a suitable software platform. As already mentioned, the focus of this case study concentrates on the functional areas of configuration and fault management and the relationship between them. Differentiation between faulty network behaviour and normal network

behaviour for transient or moving elements is an area of particular interest. In this context, it is important to investigate how far the traditional functional decomposition within Network Management (FCAPS) can be preserved and where redefinitions will be useful or even required. Note that in this paper we can only introduce the scenario and highlight its challenges for Network Management, as a starting point of the case study; the results of any further investigations based on our implementation will be published in a future paper.

3. The Scenario

Let us highlight the problem domain and the focus of our case study in a concrete example of a heterogeneous, dynamically forming network. For simplicity, a meshed wireless LAN is considered here. Please note that our scenario is equally applicable for any type of network, wireless or wired.

The scenario for the Madeira project must satisfy three criteria: It needs to be

- *challenging*; providing a number of tasks used to exercise the Madeira management approach.
- *grounded*; it must describe identifiable, familiar, and realistic management problems that could, or better, do arise.
- *practical*; it must be possible to demonstrate on a small scale on real, available, and inexpensive network equipment; and it must also be possible to simulate.

The traditional method of deploying a wireless network is to use a wired network for backhaul by connecting wireless base stations at various points. The wireless base stations are deployed independently of each others' backhaul.

In a wireless meshed network, wireless base stations as in a traditional wireless network configuration. Each base station may or may not have a wired connection to a backhaul network. The wireless network sector of each base station co-operates with the surrounding sectors to provide backhaul connectivity to all network elements connected to the meshed network. Only some of the network sectors are connected directly to a wired backhaul network and thus to external networks. Other sectors of the wireless meshed network use adjacent sectors for external connectivity.

Challenges to be resolved in this context include:

- How the management functionality might support deployment of a W-LAN meshed network at a venue where only a subset of the W-LAN access points are connected directly to a wired network.
- How such a meshed network might be configured so that stations can connect externally using any of the W-LAN access points.
- How redundancy can be configured on the W-LAN network for both access points and stations.
- How reconfiguration of the meshed network might be performed when an access point failure occurs.
- How faults are evaluated, weighted and reported given that in some situations access point failure is non-catastrophic, triggering an automatic reconfiguration, and in other cases access point failure requires the intervention of an operator for restoring connectivity.

A set of scenes have been identified, for each of which the challenges are described, and the detailed steps of the scene are given, both from the network's and the operator's point of view.

3.1. Wireless Meshed Network Formation and Re-Formation

After the wireless base stations have been deployed at the venue under consideration (think of a large conference in an exhibition centre), the operator asks the management function of the network to set up a wireless meshed network.

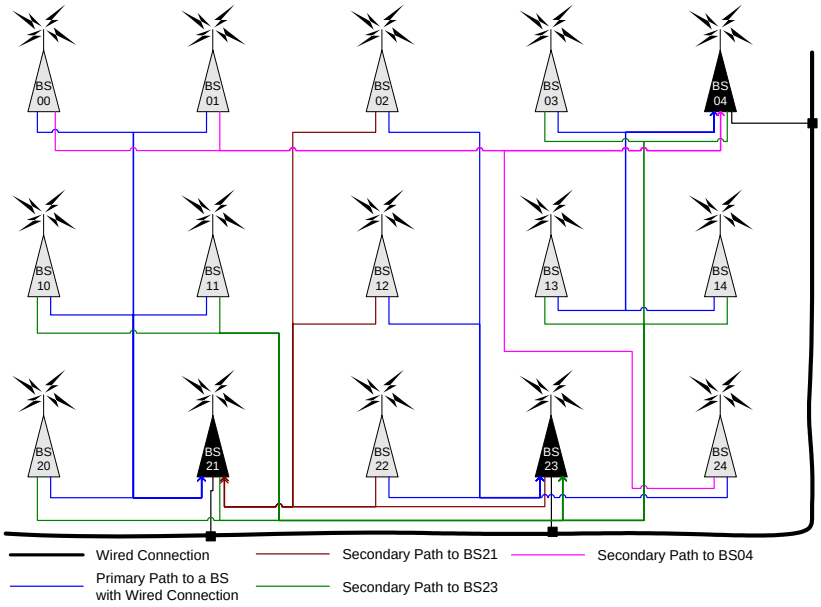

Figure 1: Typical Primary and Secondary Mesh Connectivity

The management function sets up the wireless meshed network. The network depicted in Figure 1 shows a possible resulting connectivity. Each base station has a primary path to one of the available LAN gateway points and a secondary path to another one. The management function ensures that a network element in any exhibition space can connect to at least two base stations. Each base station requires two external connections (Internet) for redundancy purposes.

A base station can have three capabilities:

- Providing wireless connectivity to network elements
- Acting as a gateway for external connectivity
- Bridging connections for other base stations

Figure 1 shows a sample configuration for the meshed network. All the base stations in the network provide wireless connectivity for network elements. BS04, BS21, and BS23 act as gateways for the meshed networks. All base stations can be configured to bridge connections.

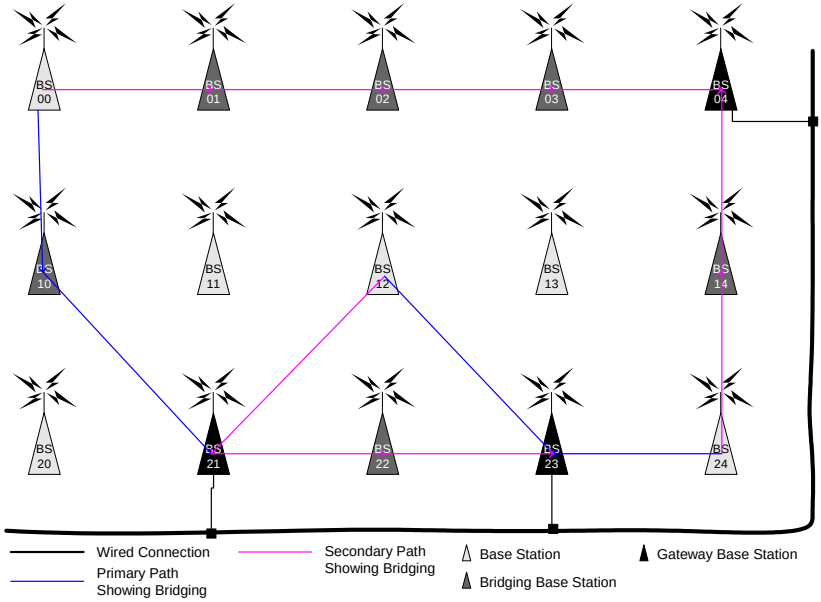

Figure 2: Base Station Capabilities for a Network Subset

Figure 2 shows how the base stations might be configured in the mesh for a subset of the network. The primary connection of BS00 uses BS10 as a bridging node and BS21 as its gateway node; the secondary connection of BS00 uses BS01, BS02, and BS03 as bridging nodes and BS04 as its gateway node. BS12 connects directly to gateway nodes BS23 and BS21 for its primary and secondary connections respectively. BS24 connects directly to BS23 for its primary connection and bridges over BS14 to gateway BS04 for its secondary connection.

Gateway nodes are directly connected to the wired LAN and thus do not bridge for their primary connection. They must bridge for their secondary connections. In Figure 2, BS21 uses BS23 as a gateway for its secondary connection and bridges over BS22 to reach it.

It is important to note that this scene is not only relevant in the beginning (after initial deployment), but applies to all changes to the configuration: The organizers may move some base stations around, add new ones, or face changed performance requirements. Each of these cases would result in a re-formation of the meshed network.

Scene Challenges: Management functionality configures primary and secondary backhaul connectivity for all base stations (Network Formation)

Meshed network adapts to utilize extra resources or if reconfiguration is required (Network Reconfiguration)

Operator Perspective: The operator asks the management functionality to build (re-form) a wireless meshed network. S/he provides the management functionality with some parameters such as the level of redundancy required, the maximum number of bridged base stations to use in a connection, and the identities of the gateway base stations. The management function sets up (adapts) the meshed network and reports back to the operator. The operator can then carry out some tests to ensure that all base stations are connected and that test network elements can connect using the meshed network. The operator then declares the meshed network to be operational.

Management Issues: How are nodes partitioned; what is the best selection process?

What mechanism is used to do resource discovery?

Where is the state information for the mesh stored and cached?

Is there an overseer (Super Peer) to resolve conflicts?

In case of reconfiguration:

- How are existing users affected by this operation?
- What criteria are used to initiate this operation; monitoring or operator action?
- Who/what stores the new topology information or is there any such cache?

Domain Specific Issues: How is IP routing set up in this context?

3.2. Example of a Network Fault: Preferred Base Station Failure

Up to now we have considered only configuration management (CM) related issues. Let us now look at one example of a failure of one component in the network. Although one of the most simple cases to be considered, it will already be evident that the relationship between fault management (FM) and CM becomes non-trivial in a dynamic self-adapting network.

In this scene, the base station being used by wireless equipment fails. The important requirement is that the management functionality re-configures the meshed network, exactly as described in the previous scene.

The wireless equipment switches over to an alternative base station, using that base station’s primary path for external connectivity. All services remain available and are not interrupted.

The meshed network's fault management function issues an X.733 compliant alarm on its northbound interface (towards the operator) with a perceived severity of "warning", since no service for the user is being affected. The raised alarm already comprises the result of the automatic reconfiguration triggered by the management function. Nevertheless the operator has to be informed, in order to get the possibility to repair / replace the base station that is out of order.

Scene Challenge: Meshed network reports fault and adapts its backhaul connectivity to deal with base station failure

Operator Perspective: The operator is informed by an alarm that a preferred base station has failed and that the management functionality has re-formed the wireless meshed network. S/he can view the report generated by the management functionality and carry out some tests to ensure that all base stations are connected and that test network elements can connect using the meshed network. The operator may then repair or replace the failed base station and carry out a wireless network re-formation as described in the previous scene in section 3.1 to bring that base station back into the meshed network. The management functionality will then cease the alarm.

Management Issues: Where is the reported fault sent? How are alarms from different sources correlated?

What distinguishes a fault from a dynamic reconfiguration?

Where is automatic reconfiguration triggered?

Domain Specific Issues: In the recovery scenario, how is IP address allocation resolved?

An illustration of this scene is given in a further simplified picture of the mesh network in Figure 3. It is evident that the base station failure of BS x may be detected from either the connected wireless equipment (as already described) or another BS y that uses BS x for bridging, or another BS z that is used by BS x for bridging. All of these are possible sources of an alarm that has to be distributed in the mesh network, indicated by red arrows. At some nodes these alarms have to be correlated, in order to trigger the required automatic repair action, and to present meaningful information to the operator.

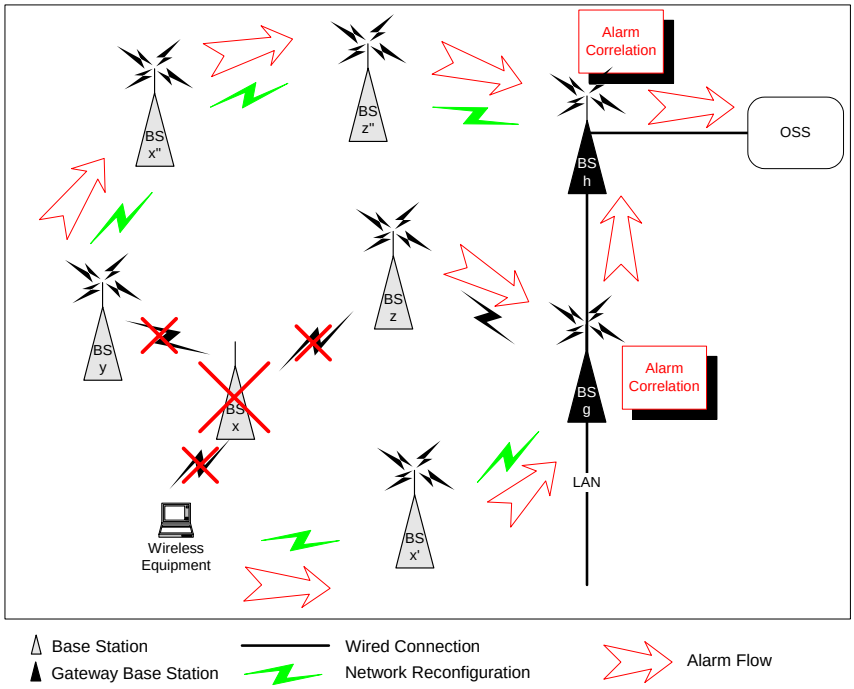

Figure 3: Alarm flows and network reconfiguration in case of a base station failure

This simple example already allows analyzing some general FM issues like alarm/event coordination / correlation between distributed nodes and automatic reconfiguration. Other types of network faults or combinations of faults pose more challenges to Network Management, such as

- correlation between reconfiguration events (CM) and alarms (FM),
- dependency of the "final alarm" shown to the operator on the result of reconfiguration,
- faults resulting in isolated autonomous network parts ("islands"), with different information available, where the work flow for restoring the problem has to be coordinated,
- synchronisation between such temporarily isolated parts, after connectivity has been restored,
- modelling of the state of the network for arbitrary combinations and multiplicities of faults.

4. Madeira High-Level Architecture

Management of the Meshed Network scenario presented in the previous chapter poses many challenges if a traditional approach is adapted. The current generation of management architectures are hierarchical.

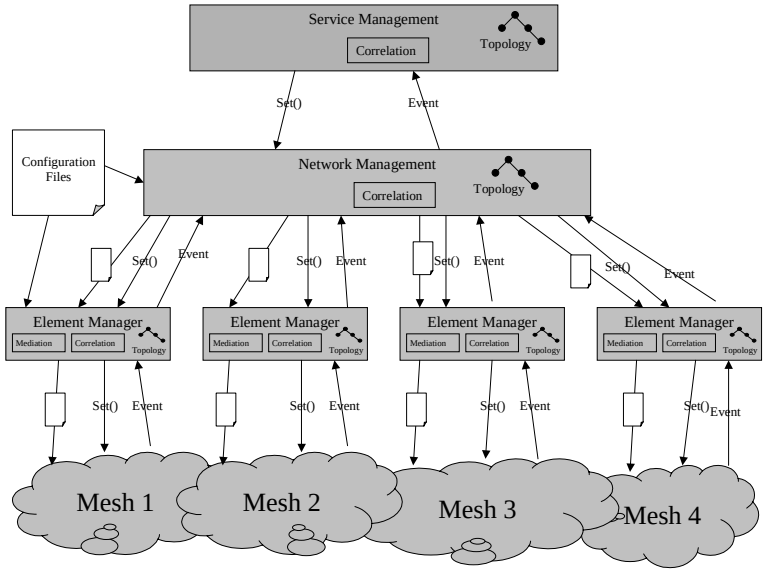

Figure 4: Traditional hierarchical Network Management approach

Fault management is event based using SNMP traps, CORBA notifications and text-based printouts. These events are propagated up through each management layer; with correlation being applied using static topology data (see Figure 4).

Configuration management of connectivity is file based, with data files being generated manually or by off-line planning tools, using static topology data stored in management systems. These files are "run" on network elements, by the element managers or NMSs, thereby applying the configuration changes, and updating the topology data. Configuration management for services is command-response based, with commands propagated down from service management to the network elements and responses going in the opposite direction. The service management system that issues service configuration requests to the network uses static topology data stored in the management systems to implement service provisioning.

The traditional approach as outlined above does not sufficiently support the two new networking concepts of (a) Dynamic Change of Network Topology and (b) Dynamic Change of Network Element Role, as can be seen in the following challenges to traditional systems:

1. Dynamic network reconfiguration difficult if based on off-line file mechanisms.
2. Traditional approaches assume that every network is always connected to the hierarchical management system. Elements in a meshed network may be disconnected from upper management layers during reconfiguration or for extended periods.
3. Dynamic and time-based topology information is hard to handle if alarm correlation assumes that topology information is static.
4. Service provisioning assumes that the network capabilities are static; in a meshed network, the service capabilities of the network can change over time.
5. Current configuration applications assume that the features offered by network elements is static or configurable by management request; in a dynamic network, network element roles can change as the network connectivity changes.

Peer-to-peer systems can help resolve these problems as they are a means of responding to a distributed and dynamic problem with a distributed and dynamic approach. This has been shown widely, e.g. in file sharing environments (Napster, Gnutella). Madeira picks up this idea, stating that the management of distributed and dynamic networks requires a distributed and dynamic management approach.

The management system then turns into a collection of peers, forming an overlay P2P network and exchanging information with each other on an east-west interface (in addition to the standard north-south orientation in today's management approaches). If you look at those interfaces in terms of management views or roles, the following are relevant for each peer, in addition to NMS core internal functions (the Core View, acting as a transaction mediator between all of the other views):

- Manager of network elements, via Southbound interface (NE View)
- Access to an information repository (Database View)
- Exchange of information with other peers, via East-West or P2P interface (Peer View)
- Presentation to a human operator (Operator View)
- Delivering information to a higher, i.e. OSS, layer, via Northbound interface (Service View)

In a similar way of reasoning, this concept can also be applied to network elements: New types of network elements, in addition to being managed by an NMS, may have the capability to exchange management information with their peers autonomously, by means of an overlaid P2P management network. Legacy network elements lacking this capability are integrated by means of proxies, using the peers' southbound interfaces as defined before.

Figure 5 illustrates this concept of views for both "NMS elements" and NEs. In comparing the diagrams the similarities between an NMS and an NE on this level become evident.

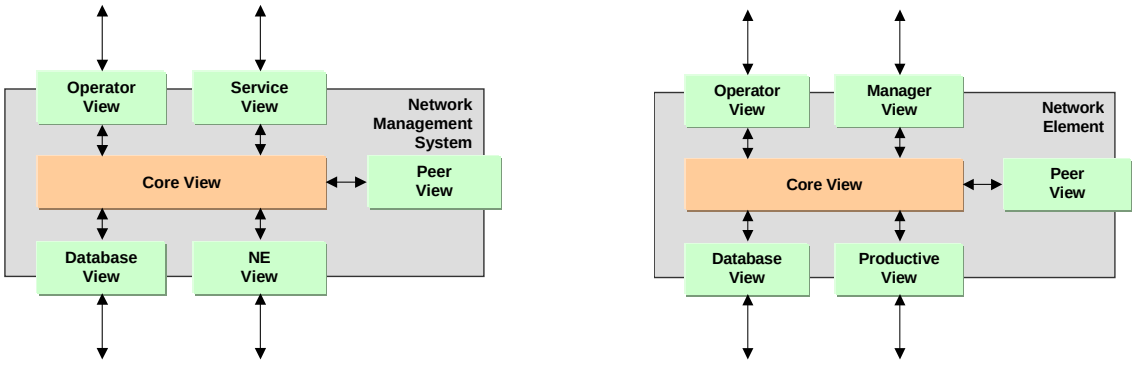

Figure 5: Network management system views (left picture) and Network element views (right picture)

As a further step within this approach, the (new type) network elements and "NMS elements" can be thought on the same layer of the logical P2P network, which merges the southbound interface of an NMS into the P2P interface. In this regard, a management element could be considered as just a special type of network element, or vice versa. Similarly, a northbound interface to a higher layer – that itself supports a P2P interface – could also be replaced. Thus one can draw the following picture of an ideal P2P management system, as shown in Figure 6.

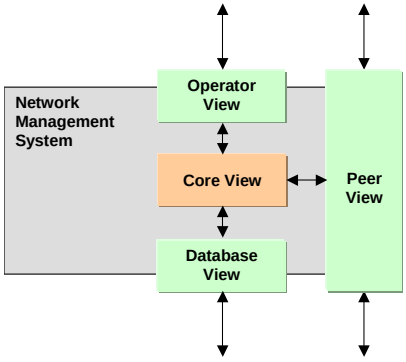

Figure 6: Network management system views in an ideal P2P system

From a modelling point of view, the Madeira system architecture is composed of a set of Adaptive Management Components (AMCs). These AMCs are framework components, which have the ability to exchange and export Network Management information between peer management applications. The AMC is only bound directly to a software platform middleware and Network Management technology when a transformation has been made from a Platform Independent Model to the Platform Specific Model. The AMCs are deployed as an overlay network, communicating using the peer-to-peer paradigm. Relating this concept to the scenario as described before, one could think of deploying a package to each of the involved nodes (both base stations and wireless equipment) that creates and makes use of the overlay P2P management network.

The following steps represent the foundation of this approach:

The first step, the *development step*, supports the creation of with an appropriate formal notation for the specification of an AMC. Defining an AMC in an MDA [8] context requires a mechanism of metamodeling to be employed. A metamodel is "a precise definition of the constructs and rules needed for creating semantic models" [11].

The Madeira Meta Model is the basis of a formal notation language. This Madeira Meta Model decouples the core management component from any concrete middleware, provides a simple, unified interface to the P2P functionality and introduces basic management functionality transparent to the core management component.

The second step is the *execution environment* for an AMC. For this part a heavy lean on Active Network concepts has been made. In the execution environment, active AMC nodes can perform customised operations on payloads contained within the Network Management information. Such operations can be defined by users who inject programs into the network to customise the processing of user and/or application specific data. The execution environment has three layers with each layer providing a unified access to its functionality via one or more APIs.

1. The Application Protocol, transports information between AMC application components. The protocol supports construction of management hierarchies, addressing of hierarchies, scoping, filtering, and transactions.
2. The Application Programming Interface (API) decouples AMC application components from software middleware technology and enables the seamless integration of management functionality into communication applications.
3. The AMC application services realise the naming of objects; enabling the mapping of Meta Schema specified information to directories through the usage of type and data repositories for applications and Madeira architecture components.

The final step is the *deployment* of the components and the distributed system itself. Here, a number of tools are provided for the configuration of the P2P system. These tools are offered in form of an administration application. This application is able to visualize information about the actual state of AMCs, including instantiated objects, request counts, runtime behaviour, monitoring, and log information.

An illustration of this concept is given in Figure 7, showing the context of AMCs, the API, the Application Protocol and various Application Services.

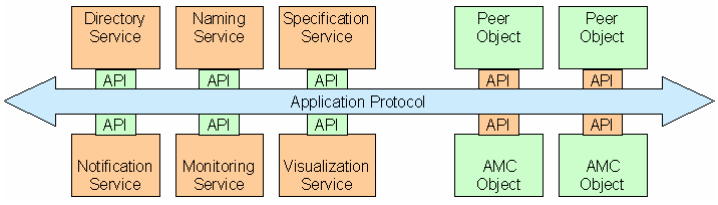

Figure 7: Madeira Modelling Approach

Specification of the components and interfaces introduced in this section is a main goal for the ongoing design activities in Madeira. Further issues to be addressed are the coexistence with legacy Network Management approaches (following an evolutionary strategy) and investigating how this architecture fits into the TMN e-TOM [12] process framework.

5. Conclusions and Future Work

At this early phase of the project, most of the work has been focused on describing key requirements and a challenging scenario for a distributed management system. In this paper we have sketched our approach, presented a scenario for focussing our further investigations, and pointed out some of the key issues that are difficult to address with standard approaches. Further we have presented a high level architecture of a logically meshed Network Management System, serving as an "architectural spike", as the first iteration in an iterative design process.

The ideas presented here give rise to further iterations in specifying requirements, architectural design and modelling. Most of all, the concepts shall now come to life, based on the Madeira software platform and the scenario as described in section 3.

Final project results are expected to be:

- The specification of an architecture for a distributed NMS through an advanced computing framework to support management operations that are massively distributed in nature and across dynamically forming networks.
- A set of standardised vertical (northbound) interface definitions between the NMS and the OSS.
- A set of standardised horizontal interface definitions between NMS elements that allow a high degree of inter-working between the management systems of various network domains.
- Exploration and definition of a new relationship between Configuration and Fault Management for transient dynamic network elements, as a starting point in reconsidering the classical functional areas of Network Management.
- Provide a means of rapidly and efficiently describing and programming management operations that form Network Management applications, through the adaptation of novel application and data modelling techniques.

6. References

[1] ITU-T Recommendation M.3100 (1995), "Generic network information model".

[2] Case, J., Fedor, M., Schoffstall, M., and Davin, J., "The Simple Network Management Protocol", STD 15, RFC 1157, IETF, 1990 (and a huge number of RFCs following RFC 1157).

[3] Hegering, H.-G., Abeck, S., Neumair, B., "Integrated Management of Networked Systems - Concepts, Architectures and their Operational Application", Morgan Kaufman Publishers, 1999.

[4] Risson, J., Moors, T., "Survey of Research towards Robust Peer-to-Peer Networks: Search Methods" Technical report UNSW-EE-P2P-1-1, University of New South Wales, Sydney, Australia, 2004.

[5] Roussopoulos, M., Baker, M., et al, "2 P2P or Not 2 P2P?", The 3'rd International Workshop on Peer to Peer systems, San Diego, USA Feb 26-27, 2004.

[6] Shirky, C., "What is P2P … and what isn't", http://www.openp2p.com/pub/a/p2p/2000/11/24/shirky1-whatisp2p.html

[7] Madeira project website, http://www.celtic-madeira.org/

[8] Object Management Group (OMG) Model Driven Architecture (MDA), http://www.omg.org/mda/

[9] Bass, L., Klein, M., Bachmann, F.: Quality Attribute Primitives and the Attribute Driven Design Method. In: Proceeding of the PFE-4 (2001) 163 – 176.

[10] Hofmeister, C., Nord, R., Soni, D., "Applied Software Architecture" Addison-Wesley, 2000.

[11] Community site for meta-modeling and semantic modelling, http://www.metamodel.com

[12] TeleManagementForum (TMF), Enhanced Telecom Operations Map (eTOM), http://www.tmforum.org/