

# An Extensible Network Resource Abstraction for Applications on Mobile Devices

Arjan Peddemors<sup>♦\*</sup>, Ignas Niemegeers<sup>\*</sup>, and Henk Eertink<sup>♦</sup>

<sup>♦</sup>INCA group  
Telematica Instituut  
Enschede, The Netherlands  
{Arjan.Peddemors, Henk.Eertink}@telin.nl

<sup>\*</sup>WMC group  
TU Delft  
Delft, The Netherlands  
I.Niemegeers@ewi.tudelft.nl

**Abstract**— Mobile devices are increasingly equipped with multiple network interfaces that can dynamically support access to different kinds of wireless and fixed networks. Dealing with these dynamics and heterogeneity is challenging for many types of applications running on mobile devices and is hindered by the lack of a consistent overview of available network resources.

In this paper we present a system service which offers an extensible network resource abstraction to applications running on a mobile device. It can be used in conjunction with well-known APIs that provide access to these resources. We argue that the cross-layer and mobility reflecting transparency introduced by the model underpinning this abstraction is essential for applications that need to dynamically adapt to changes in available network resources. Ease of adaptation may reduce the reliance on dedicated mobility management protocols. Furthermore, we show that different types of mobile applications may use this abstraction at various levels of detail, to match their specific decision making needs. We introduce the architecture of the service, discuss experiences with its reference implementation, and show how the system can be extended to accommodate new network technologies. Experiments with a real-world application indicate that uncomplicated application logic suffices for a dynamic network environment.

**Keywords:** *network adaptation; network heterogeneity; mobility management; communication middleware; mobile application; multi-homing; host mobility*

## I. INTRODUCTION

With the establishment of the mobile device as a major computing platform, users become accustomed to mobile applications that take into account the differentiating characteristics of this platform. Common aspects such as the personal link between user and device, its continuous physical closeness and availability, and the battery powered operation, require applications to be designed differently than those on more traditional platforms such as the Personal Computer (PC). As a consequence, this requirement translates into the need to support these applications with platform specific facilities at the operating system level.

One such aspect is the availability of dynamic and heterogeneous network resources. The primary difference between applications that are involved in data communication on mobile devices and on more traditional computers such as desktop PCs or servers, is the difference between a dynamic and a relatively static network configuration. The layering as seen in the protocol stack has been a good match with this static configuration. The elementary abstraction offered on more traditional hosts is the abstraction presented by the transport layer – setting up and maintaining connections with other end-points – combined with a general notion of a binary IP connectivity state: there is connectivity or there is none. This also matches well with the network hardware most often found on traditional computers: the fixed Ethernet interface with a, generally speaking, binary connectivity state. The application deals with setting up an end-to-end connection with another host and does not need to care about settings and configurations of entities at lower layers (network, link, physical) that take care of establishing and maintaining this connection. This is the power of the layered approach and it has served us well over the years.

The situation is different in a mobile setting: the wireless network hardware can often be in one of multiple modes such as ‘off’, ‘sleep’, ‘power safe’, or ‘fully on’, and detect multiple networks with different and fluctuating network capacity – inherent to mobility. Additionally, network hardware may be able to sustain connectivity to more than one network and mobile devices often get equipped with multiple network interfaces: currently a configuration with 3 or 4 interfaces is not exceptional (e.g. Bluetooth, 802.11, cellular, and fixed serial such as USB). This opens up the possibility to choose an available or potentially available data network to carry the traffic for a transport layer connection. There are many reasons to influence this decision: capacity, cost, security, power usage, etc. Furthermore, different applications (such as email, web browsing, file download, streaming media, p2p clients) may have different needs, perhaps even varying in time. With this in mind, the simple transport layer abstraction we used before does not suffice anymore for all cases. We need to support applications with information about the state of entities at the network layer and below in order to be able to indicate, for example, that connection traffic must go over wireless network A and not B. However, we also need to make sure that this

---

This research has been supported by the Dutch Freeband Communication Research Programme (AWARENESS project) under contract BSIK 03025.

information is offered with the right level of detail: just enough to make decisions, providing a high level view that does not reveal all information at all layers to the application, i.e. keeps the layered model intact.

In this paper, we propose an extensible network resource abstraction that supports applications on mobile devices with their decisions on the usage and manipulation of network resources. We describe a network resource model that constitutes the foundation for the abstraction offered to applications. Furthermore, this paper introduces an Application Programming Interface (API) and Windows CE system service implementation that provides the network resource description. This API can be used in conjunction with well-known other APIs that provide the actual access to the network resources: it does not, for instance, replace the Socket API. The main contribution of this paper is that it indicates the necessity for balanced cross-layer network resource transparency for mobile applications and shows this can be realized without breaking the layered model. We show that, perhaps contrary to intuition, the design of applications that handle cross-layer information does not need to be complex or laborious given the right level of abstraction. This focus is on dealing with characteristics of network resources that can be readily observed from the mobile host: i.e. without relying on information or measurements gathered from within the network infrastructure.

Although the main focus of this work is on supplying information to applications, many aspects are also useful to system oriented functionality. For instance, mobility management functionality such as Mobile IP may show default behavior on how packets are tunneled and uses information on available 'real' paths to handover the tunnel from one network to another. Additionally, a system power manager may use information on the power consumption of network interfaces in their various modes to make decisions on activation and mode changes of these interfaces.

The work presented here assumes an existing technology infrastructure consisting of common – commercially available – mobile devices with multiple network interfaces, general purpose operating systems available on these devices, a set of basic IP technologies as well as a mixture of various mobility management protocols. This focus is motivated by the market uptake of devices, primarily 'Smartphones' and devices with integrated PDA and phone functionality, that have multiple network interfaces and an open system configuration, allowing for additional application to be installed. The combination of growth of this platform with a lack of deployment of mobility management functionality such as Mobile IP results in a shift of adaptation responsibility from the system towards the application. Hence, applications need better information on available network resources.

This paper is organized as follows. Section II discusses common network operations and the necessary level of awareness for mobile applications; section III introduces the model that delivers an abstraction of these resources. Section IV describes the architecture of the system service we use to offer this abstraction to applications on common mobile device operating systems, which is followed, in section V, by a discussion on the experiments with its implementation and an

application. Section VI discusses the results. This is followed by an overview on related work and we wrap up with conclusions.

## II. APPLICATIONS: NETWORK AWARENESS AND NETWORK OPERATIONS

This section provides an analysis of the level of network resource awareness shown by a number of example mobile applications. Additionally, its objective is to illustrate and better understand the kind of inter-layer information exchange taking place on a mobile device. It illustrates that network resource awareness is not something for 'performance oriented' applications only, but is relevant to many types of networked applications. Furthermore, we show that the states of mobility management mechanisms – such as Mobile IP or the Host Identifier Protocol (HIP) – are an inherent part of the overall network resource state in a mobile environment.

Modern operating systems for popular mobile devices allow us to use various means to obtain an indication of currently available network resources. Information can be retrieved through different kinds of mechanisms, each targeting their own specific information range, e.g. system APIs to obtain settings and notification of change in network interfaces, detected wireless networks, or IP configuration, as well as APIs – sometimes hardware vendor specific – to provide information on network hardware settings. Although useful by themselves, this scattering of APIs or similar mechanisms makes it hard for applications running on the device to get a complete overview of network resources. But more importantly, the APIs commonly do not adhere to the same underlying model or semantics to describe these resources. This obstacle needs to be overcome in order to give applications a clear and complete image of communication capabilities.

### A. Examples of Network Operations

The application examples given here provide an indication of the possible network resource aspects that must be dealt with. The following types of network operations are considered:

- *Detection and Scanning.* Deals with the configuration of detection and scanning settings for the available network interfaces. Enabling this is generally costly in terms of power consumption and can interfere with data transfer. Controlling this configuration may require network technology specific knowledge.
- *Activation.* Takes care of activating one or – if supported by the technology – multiple links over a network interface as well as establish an IP configuration over these links. This operation requires understanding of how to identify networks, where network identifiers may be network technology specific. Link activation credentials as well as IP configuration parameters may be cached by the OS (is not necessarily an application concern).
- *Sending and Receiving Data.* Involves the selection of an available path to carry the data.

TABLE I: INDICATION OF NETWORK OPERATIONS FOR THREE EXAMPLE APPLICATIONS

Network Operation	Remote Control	Video Newsfeed Player	Media Stream Player
	Application for remote control over equipment and appliances in the user's surroundings. Appliances are IP enabled and reachable over local wireless networks (802.11 / Bluetooth (BT) / Ultra-WideBand (UWB) in both infrastructure and ad-hoc modes).	Video Newsfeed Player with integrated download functionality, which tries to obtain a new video news item at regular intervals. Providing relatively fresh news also when temporarily deprived from network connectivity. Uses Mobile IP to deal with handovers.	Media Stream Player which adjusts the stream parameters given the available network resources at a certain time, in particular the available throughput capacity. Objective is to supply the user with high quality and continuous streams.
<b>Detection/Scanning:</b> configure relevant interfaces for (occasional) detection of available networks	Enable detection of those networks that offer access to the remote control interfaces of appliances. See <i>activation</i> below.	It is only necessary to connect when trying to download new items. For power efficiency, only discover when necessary, i.e. just before and during downloads.	Detect when the user moves out of range of the current network or detect that a better than current network is available.
<b>Activation:</b> activate available link over a network interface – possibly more than one depending on network technology – and activate IP settings over link	Possible strategies could be a) activate only those links that are known to provide access to known appliances or b) try to opportunistically connect to any device with a remote control (IP) interface available over any link.	Active the (physical) link with the best characteristics for high-volume downloads: fast and inexpensive. Reactivate – if not activated already – other (next-best) links when currently used link becomes unavailable.	Activate the link with the best characteristics for continuous and high quality streaming. Reactivate – if not activated already – other (next-best) links when currently used link becomes unavailable.
<b>Data Traffic:</b> sending and receiving IP data, typically using transport layer functionality	This involves controlling over which outgoing path the connection will be established: especially an issue when the appliance is more than one network hop away.	Controlling over which outgoing path Mobile IP will establish its tunnel, so that traffic generated over the Mobile IP virtual link will go over the right physical link	Controlling over which outgoing path the connection will be established.
<b>Monitoring:</b> follow the characteristics of network resources in use – as seen from the perspective of the mobile host	Monitor existing connections to managed appliances. For a), monitor availability of links that may provide access to known remote appliances. Helps e.g. to notify user that access to one or more appliances has been lost.	Because the connection is sustained by the mobility management protocol, it is useful to know which real network eventually carries the data packets, for instance to control costs.	Reconnect the control and stream connection and change the stream parameters accordingly. It is assumed that a large enough playout buffer exists that can absorb these temporary stream interrupts.

- *Monitoring.* Status information used as input to the decision process that initiates other network operations.

See table I for a high-level overview of the network operations relevant for three examples applications all running on a mobile device. Many other applications could be listed here. We have left out, on purpose, such applications as Voice over IP (VoIP), because cross-layer issues are clearly important and often focus on performance aspects such as fast handovers. The examples above show that cross-layer interaction is equally important to a wide range of less demanding applications. In fact, much of the interaction is more based on the existence of relationships between the various entities within layers (is available network X activated through interface Y?) than on the characteristics of these entities (e.g. maximum throughput for a link).

#### B. Contours for Cross-Layer Network Resource Awareness and Control

Considering the examples above, it is clear that mobile applications need to operate on network resources at different layers. This goes against the principle of the layered stack model, where an application uses only the interfaces offered by the transport layer functionality. Still, many choices can be enforced using the transport layer interface, typically the Socket API. Most network stack implementations can be configured such that binding a socket to a local IP address determines over which link the outgoing traffic will be

directed<sup>1</sup>. This means that *local address binding* allows an application to choose, for example, that packets for a TCP connection are sent out over a link available through a physical interface with certain desired maximum throughput characteristics, or over a Mobile IP virtual link and interface which will sustain the TCP connection even in the event of alternating availability in networks. An application, however, can only make this kind of decisions when it has received information that associates a local IP address with a link and a link with an interface, and characteristics such as maximum link throughput: essentially cross-layer information. The following observation can be made: for the *enforcement of the usage of network resources* the current interface between the application and transport layer offers substantial possibilities, while the cross-layer information used for the *decision on which network resources to use* cannot be obtained through this interface. This observation has formed a leading motivation for the creation of the network resource abstraction presented in this paper: we have existing means to enforce the usage of network resources but do not have proper input to the decision process. As the examples show, not all enforcement can go through the Socket API. Lower layer (network, link) operations, however, seem to have a general form, e.g. operations such as ‘enable scanning’ or ‘activate link’, that can be defined as abstract operations. Furthermore, applications

<sup>1</sup> And, for most practical configurations, determines that the incoming traffic for that socket will be received over the same link, although this decision is not taken by the mobile host but by the routing infrastructure.

that need highly specific lower-layer control over network resources can use additional, more specific, APIs such as those available to detect fixed Ethernet link characteristics, or to manipulate specific 802.11 settings, or to control Bluetooth behavior.

Given the analysis of the examples and the considerations above, the following list of requirements for the network resource abstraction can be formulated:

- Cross-layer network resources information should be presented ‘as is’: being true to what’s generally understood as basic elements and concepts inside the layered stack model, as well as inside network technologies. This enables straightforward interoperability with existing well-known APIs.
- Being capable of reflecting what is typically available on personal mobile devices such as Smartphones, from operating system and network hardware point of view.
- Offer flexible levels of abstraction to provide applications with a view that matches with their decision making needs.
- Simple and light-weight interaction: mobile devices are battery powered and have relatively few resources and are normally only capable of accommodating light additions.

### C. Interaction with Default System Behavior

So far, we considered the influence of applications on network resource activation and usage. Obviously, a part of the decision making can be the responsibility of the operating system. The interaction between application and system when it comes to the assignment of available resources has been an extensive topic of study [8]. The target platform for this paper is a single-user mobile device with a general purpose operating system that simultaneously runs multiple applications. We take a straightforward approach by assuming that the operating system implements default behavior that can be overridden by the applications. We do not consider mechanisms that may help to resolve conflicts between applications: we think of this as a responsibility of the user.

We adhere to the following principle. The up front default operating system configuration for the activation and assignment of network resources is determined by a rough estimate of the needs of the applications that will run on the devices. On the other hand, the applications are designed with a general notion of this default system behavior in mind. For instance, the OS assumes that at least one network interface must be activated to offer internet connectivity to applications. A network resource unaware application assumes IP connectivity to the internet, while an aware application may activate and use other networks than the system activated network to better match application needs. When looking at the examples, certain actions such as network discovery settings and network activation decisions may therefore already be realized by the system: in that case the application only needs to select from what’s available. This can be executed by obtaining a snapshot of the currently available network

resources, decide which outgoing path to use, and bind to this path – using the Socket API – during connection initiation.

## III. NETWORK RESOURCE MODEL

To identify the basic elements needed to describe network resources, we go back to the concepts that are fundamental for each of the layers [6]. They are widely known, although the inter-relationships between them are often not explicitly expressed. The Network Resource Model (NRM) considers entities at three different layers of the protocol stack: the link layer including everything below, the network layer and the transport layer. At the (physical/) link layer, bits and frames are transferred over *links*, where a link is maintained by a network interface. An example is the link setup between a host’s Ethernet interface and an Ethernet switch. At the network layer, packets are sent over *paths*, where a path consists of a concatenation of links. At the transport layer, data is sent over end-to-end *transports*, where a transport uses one or multiple paths.

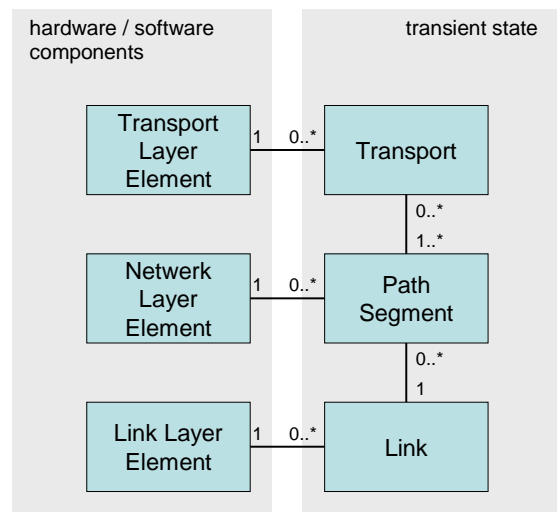


Figure 1: Basic (abstract) NRM entities and their relationships

The lifetime of each of these fundamental entities is managed by the hardware and software functionality at the respective layers. Links are maintained by network interfaces, paths by IP stacks, transports by transport layer stacks (UDP, TCP, SCTP, DCCP). These pieces of functionality we respectively call *Link Layer Elements*, *Network Layer Elements*, and *Transport Layer Elements*, and are defined as basic entities in the model. All NRM entities are described by their properties and their relationships with other entities. At the transport layer, and especially the network layer, the available entity types are relatively few, but at the link layer they can be numerous, depending on the number of network technologies that need to be incorporated. The types are defined following an object-oriented paradigm resulting in type hierarchies: basic and abstract types define generic characteristics, while deriving types define more specific characteristics. Note that the NRM does not exclude ad-hoc networking: these can simply be modeled with links and paths.

We define a number of interactions that can be used to manipulate or monitor resources with our model. These currently are simple and entail retrieval of a resource description as well as means to subscribe to changes on the whole or parts. For instance, a subscription can be made to changes in a layer, a type, or an instance of a type. Operations could also be defined for influencing network resources, such as the selection and activation of 802.11 networks on a specific 802.11 interface. These interactions are discussed further in section IV with the Network Abstraction Layer interface.

The NRM is notated in XML schema [14]. Network resource descriptions that adhere to the model may be expressed as XML documents that have a structure dictated by the model’s XML schema. We use various XML schema mechanisms to describe the model: abstract types, deriving (concrete) types, keys and key references to enforce uniqueness of identifiers, etc. Unfortunately, the model’s integral schema cannot be included here because of its length, but can be obtained from [3].

Alternatively, network resource descriptions may be provided through other means than an XML document, as XML is not the preferred format for all types of applications. In fact, the most basic interface to obtain network resource info is through the C API of the Network Abstraction Layer described in the next section. This interface describes the same NRM elements in terms of C data structures. We merely use the NRM schema as a formal notation to define the model: it can have different ‘bindings’ in the form of XML, C, scripting, CORBA, etc.

Considering the discussion above, we now have a model with abstract types as depicted in figure I, which indicates the general cardinality of the relationship between the entities (not the properties). A *Transport Layer Element* may maintain between zero and more transports, but a *Transport* is always associated with a single Transport Layer Element. Likewise, a Transport uses one or more *Path Segments* and a Path Segment supports between zero and more Transports, etc.

The abstract model covers only those segments of paths (single hop) that are directly connected to the host, not full end-to-end paths, because a multi-hop path is not completely visible from the hosts’ standpoint. Therefore, the host cannot uniquely identify an entire multi-hop path. Furthermore, paths may be highly dynamic as each IP packet is individually routed, although in practice subsequent packets most likely will travel the same way. Thus, strictly speaking, gathering characteristics such as delay and jitter is not possible for paths but is feasible for end-to-end connections. The host is capable, however, to provide a good description of the first segment of a path. For mobile hosts, the first hop in the path is often highly relevant, as this is mostly over a wireless link, which is in many cases the segment of the path with the most restricted resources and also the segment that has the highest influence on cost of usage – from the user’s perspective. So, information about the available outgoing paths may in many situations provide valuable information to applications generating network traffic. Additionally, the path segment is the only part of the entire path that is selected by the mobile host. Typically, path

segments coincide with the entries found in the operating system’s routing table.

When replacing these abstract entity types with concrete types, we can get a view of the NRM as provided in figure II. Note that the cardinality of the relationship between the TCP connection and the IPv4 path segment has changed compared to the abstract case: a TCP connection uses a single path, but in general, transports may use multiple paths. Furthermore, this figure indicates that the 802.11 interface has two relationships with, possibly, multiple 802.11 networks: one to indicate the 802.11 networks in range (zero or more) and one to indicate the currently active network (zero or one). So we see that the nature of the relationships is often protocol and technology specific and a simple model is capable of capturing major characteristics. Similarly, the IPv4 stack has a relationship that reflects all IPv4 paths segments and a relationship with one particular path segment (if available) that marks the default route, i.e. the path segment that is selected when a connection is not explicitly tied to a specific path segment.




Figure II: Example NRM view for an IPv4 enabled host with a WLAN interface

Naturally, in case the host has a more complex network configuration, the NRM description also becomes more complex. This is depicted in figure III. Note the following aspects: 1) an SCTP association (connection) may use multiple path segments<sup>2</sup>, either IPv4 or IPv6, 2) path segments may use an Ethernet link, an 802.11 link, or a Mobile IP tunnel virtual link, 3) the Mobile IP tunnel virtual link is established using an IPv4 or an IPv6 path.

#### A. Type Hierarchies

The possibility of defining type hierarchies is a major feature of the NRM: it allows applications to choose the level

<sup>2</sup> Actually, an SCTP association uses one or more *paths segments* (abstract type) where a path segment may be an *IPv4 path* or *IPv6 path* (concrete type): for the sake of readability, both figure II and figure III do not show the abstract types and the relationship with their concrete subtypes. Likewise, the Mobile IP tunnel uses one *path segment*, which can be an *IPv4* or an *IPv6 path*, and a path segment uses one *link*, which can be an *Ethernet*, an *802.11 network*, or a *Mobile IP Tunnel Virtual Link*.

of abstraction of the network resource description offered to them. As shown in figure III, a real-world mobile device configuration may involve many different types which not all may be of interest to the application. Suppose an application is only interested in choosing the outgoing path which is best suited to carry a certain data stream. It uses the ‘maxSpeed’ characteristics defined for the abstract type ‘Link’ to select suitable link entities and uses the local IP address of the available gateway path segment realized over this link to setup the data stream connection. This application is only interested in entities that are instances of two specific types and uses – for the link information – only generic characteristics, not technology specific characteristics.

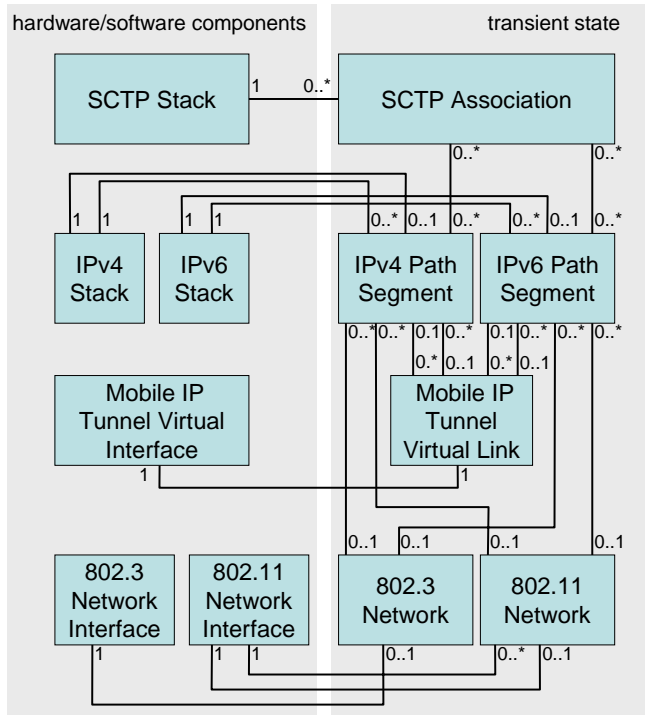


Figure IV: More complex NRM example view: an IPv4 / IPv6 enabled host with two network interfaces, and with Mobile IP and SCTP functionality

Furthermore, type hierarchies allow the model to be extended. When a new link layer technology becomes available it can ‘hook’ into the model by deriving from the available abstract types that best match the technology. For instance, a GPRS cellular interface may be defined that derives from the more generic CellularInterface. This more specific interface may, as an extension, provide means to read out all GPRS-specific identifiers of base station that are in range.

### B. Type Overview

All types depicted in figure IV are defined in the current version of the NRM. These figures, however, do not provide a full overview of the available types: some type hierarchy intermediate types are not visible. The current version of our model defines a number of simple characteristics for basic concrete types such as the TCP and UDP connection type, the IPv4 and IPv6 stack types, as well as for a number of network interface technology types most often found on mobile devices

(USB serial line, 802.11 Wireless LAN, Bluetooth, cellular, Ethernet).

## IV. SYSTEM ARCHITECTURE

This section describes the system architecture we use for experiments with the applications using the network resource abstraction. Its main objective is to define a system service for operating systems such as Symbian, Windows CE or Linux that run on common mobile devices with multiple network interfaces. All these operating systems provide means to obtain information on network resources, with varying degrees of detail, and allow for manipulation of these resources, mostly through a set of different APIs. Additionally, these OSs are general purpose and therefore suitable to support many different types of applications.

The system service focuses at implementing a Network Abstraction Layer (NAL) API for usage by applications or other system services. Obviously, such an API may eventually be implemented as a core operating system facility – our objective here is to experiment with our network abstraction, which is most conveniently implemented as an add-on component to an existing environment.

The NAL API is dedicated to provide network resource information in a light-weight form and adheres to simple application interaction often found in system APIs. Its basic subscribe/notify mechanism suffices for most applications because the number of entities that describe the full network resource state is limited: a more elaborate query mechanism is unnecessary and adds complexity.

This section starts with the description of the NAL interface – including the basic C API – followed by the architecture of the NAL system service.

### A. Network Abstraction Layer Interface

The NAL offers the application a view on the available network resources structured according to the model presented in the previous section. The NAL interface towards applications consists of a C API.

The means of interaction between an application and the NAL consists of messages send back and forth over one or more *full duplex pipes*. After its creation, a pipe can be used to start two types of interaction: it is used to subscribe to notifications on creation, updates, or deletions of entities in one of the layers, or it is used to send a command (`ioctl`) to one of these entities. See table II for an overview of available messages.

Although the interaction between application and NAL is a form of Inter-Process Communication (IPC), and therefore may consist of unmarshalled, machine dependent C structs, all message payloads must be *packed* before sending or *unpacked* after receiving, because many entities are represented by variable length data structures. In C, variable length data structures are often implemented as chained elements connected by pointers, which do not have a meaning in another process space.

The subscribe message `NAL_MSG_SUBSCRIBE` consists of an array of subscription elements. A subscription element indicates a subscription to all changes in a specific layer (`NAL_SUBSCRIBE_LAYER`), all changes of entities of a specific type (`NAL_SUBSCRIBE_TYPE`), or a change in one specific entity (one instance of a type: `NAL_SUBSCRIBE_INSTANCE`). The subscription elements together are a mask interpreted in an ‘AND’ like fashion, e.g. making it possible to subscribe to change in state of the entire network layer (IP stack, outgoing paths) in combination with a subscription to change in all 802.11 networks (instances of type ‘802.11 network’).

When the NAL receives a subscribe message, it replies with a dump of the currently existing entities that match the mask. After the dump, the pipe remains open and is used by the NAL to send changes in the network resource state. These changes are the creation of an entity (`NAL_MSG_CREATE`), the update of an entity (`NAL_MSG_UPDATE`), and the deletion of an entity (`NAL_MSG_DELETE`). A create and update are followed by the complete description of the entity, while a delete only indicates the entity’s identifier. All notifications take place in the context of a transaction: a `NAL_MSG_TRANS_BEGIN` indicates the begin and `NAL_MSG_TRANS_END` indicates the end of a transaction. This enables the application to construct at all times a consistent view on the current network resources.

TABLE II: NAL MESSAGES

Message ID	Description / Parameters
<code>TRANS_BEGIN</code>	Indicate start of transaction. No parameters
<code>TRANS_END</code>	Indicate end of transaction. No parameters
<code>CREATE</code>	Data of newly created entity
<code>UPDATE</code>	Data of updated entity
<code>DELETE</code>	Identifier of deleted entity
<code>SUBSCRIBE</code>	Subscribe to parts of the resource description. Takes a subscribe array as parameter.
<code>IOCTL_CALL</code>	Execute an ioctl command. Parameters: call id, entity id, ioctl number, ioctl parameters
<code>IOCTL_RESULT</code>	Result of an ioctl call. Parameters: call id, ioctl result value

One important feature of the NAL interface is the possibility for an application to unpack an entity as a type it was packed, or *any of its super types*. This means that an application can choose *the most appropriate level of detail*, by unpacking an incoming entity to the type with just enough information.

The command message `NAL_MSG_IOCTL` is used to send a command plus parameters to one of the existing entities. The ioctls are defined per type and are typically used to control an entity of this type. For instance, an `ioctl` may be used to activate a detected network through one interface, change the default route (default outgoing path), or to change characteristics of the TCP stack. `Ioctls` may be defined for generic types to allow for generic operations, or more specific

types to support for instance highly technology dependent operations.

This C API offers the basic NAL functionality. It is possible to provide additional NAL bindings in different programming languages, such as Java, C++, or scripting languages. Actually, due to its object-oriented character, the NAL interface may be most conveniently described using an object-oriented language. We provide a C binding here as the most basic binding, as C is supported on all operating systems and is the most common language for system oriented APIs.

### B. NAL System Service

The NAL system service is responsible for supporting the features offered through the NAL interface. Its architecture is setup such that it can be realized on any of the common operating systems for more advanced mobile devices.

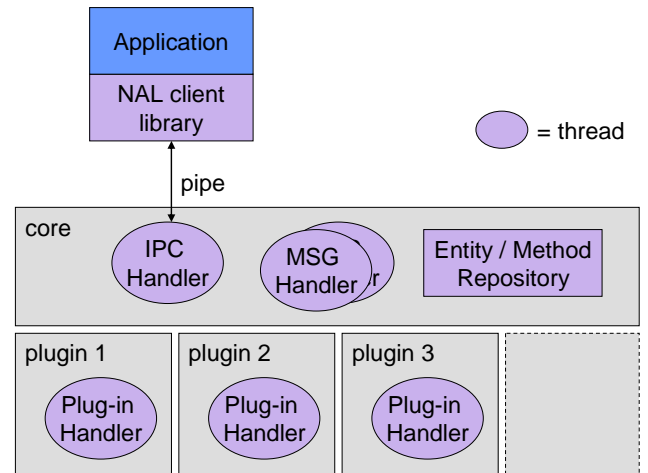


Figure IV: The NAL System Service

As depicted in figure IV, it consists of the following components:

- One IPC handler (Thread)
- Message Handlers (Thread)
- An Entity and Method Repository
- Plug-in Handlers (Threads)
- Client library (exposing the NAL API)

As part of the core NAL functionality, the IPC listener thread waits for incoming full-duplex pipe establishment by the application. A new pipe is immediately passed on to a message handler thread; the IPC handler goes back to listening for incoming pipe initiation. The message handler checks the incoming message: if it is a `NAL_MSG_IOCTL` message, a lookup of the method is performed using the ‘Entity / Method Repository’, the method is executed, and results are passed back to the client. If the incoming message is of type `NAL_MSG_SUBSCRIBE`, the message handler registers the pipe to the Repository, so that later changes to the entities in the repository can be passed back to the client, and dumps the current repository content.

The current state of the available network resources is stored in the Repository. The NAL follows a fully ‘pluggable’ architecture: the content of the Repository is kept up to date through plug-ins. This makes the NAL extensible and easy to configure for different kinds of mobile devices. Each plug-in is responsible for its own set of NRM types. At startup, the NAL looks for plug-ins, typically implemented as shared libraries, and starts the ‘threadproc’ plug-in entry point function in its own thread. Every change to the Repository by plug-ins is executed in the context of a simple transaction. Every transaction consists of three phases: a *main* phase where the major changes to the repository are made, a *hook* phase where all other plug-ins have a chance to react to these changes, and a *fix* phase that gives the transaction initiator the opportunity to make final fixes before the end of the transaction. The fix phase calls the ‘transhook’ functions of all available plug-ins. This phased transaction is convenient when references between entities managed by different plug-ins must be made. At transaction end, the Repository will notify all subscribers – if necessary – of the changes.

## V. EXPERIMENTS

We developed a NAL reference implementation for the Windows CE operating system with plug-ins for abstract types and types for the network and transport layer (nal\_base), 802.11 network types (nal\_ndis), cellular network types (nal\_cell), and USB fixed serial types (nal\_ras). Additionally, we created a NAL command line client that reflects all notifications generated by the NAL. The NAL implementation consists of almost 11000 lines of C code (.c and .h files) and has an executable size (.exe and .dll-s) of 90kB. The 802.11 plug-in gathers info on available networks and associating access points and received signal strength indications (rssi). Per default, it obtains a list of access points every 5 seconds. Similarly, the cellular network plug-in retrieves info on in-range GSM/GPRS networks and operators and refreshes the cellid and rssi of the currently used base station every 5 seconds. The operator list is renewed every 5 minutes.

When running the NAL in combination with the command line client on a Qtek 9090 Windows CE device in a constant environment, where the device is not moved, the cellular and 802.11 interfaces are activated, and no further applications run (‘NAL-on’), it consumes little extra energy compared to the same configuration that does not run the NAL (‘802.11-on’). As shown in figure V, both lines coincide: running the device without NAL and the 802.11 interface switched off, however, shows decreased power consumption (‘802.11-off’). In this two hour period, the NAL produce around 40 messages at startup, to provide a full snapshot of the available network resources, and slightly more than one update per minute for the rest of the run.

To validate the abstraction offered by the NRM as well as to test the NAL reference implementation, we have executed a number of experiments with an implementation of the streaming media example application of section 2.1. At the server side, the test configuration is based on a modified streaming audio application called SlimServer [12] which allows a user to stream audio content from his personal collection to a node on the internet. The streaming client itself

is a modified GSPlayer [5] for the Windows CE operating system. Neither the SlimServer nor the GSPlayer software is capable of maintaining streams in case of network handovers.




Figure V: Power usage for device running with and without the NAL service in combination with client

The test environment consists of a server machine storing a music collection and running the SlimServer software and a GSPlayer on a Qtek 9090 Windows CE device with, amongst others, a cellular GPRS interface and an 802.11b network interface. The device can reach the server either through the cellular IP network or through a local 802.11 network with a single access point (AP). No mobility management protocols are available: in case of network resource changes the application is responsible for re-initiating streams between server and client.

We want to realize ‘reasonable’ stream handover from one network to another, meaning that handover may be audible but the music must continue from where it left off. Furthermore, different networks may support different bitrates with different audio qualities: when possible offer the best quality, otherwise adjust the quality to the available capacity. This functionality is realized by monitoring the amount of data in the stream buffer and monitor the notifications offered through the NAL. When the streaming buffer size is below a certain threshold, the currently used link is not considered for some time and the stream is reestablished over another link (if available). When a NAL notification indicates that the current link is not available anymore or a better link is available, the stream is reestablished over the newly preferred link, without waiting for the buffer to reach the low watermark. This straightforward application logic delivers continuous streams over GPRS and 802.11 and switches between them when the 802.11 network is switched off and on. The switch is audible due to the difference in quality as well as a little loss of sync of less than 1 second. Synchronization can most likely be improved when the start time indication is expressed in milliseconds instead of seconds. The client only uses Link and Gateway Path Segment (to get the local IP address to bind to the stream socket) notifications received from the NAL.

The server had to be adapted to accommodate IP address independent session identification. Additionally, although capable of transcoding MP3 files through a piped process to a lower bitrate stream suitable for transfer over the cellular network (32 kbit/sec), it was altered to transcode high quality MP3 files to lower quality files on a disk cache. This enables streaming songs starting at the indicated time.

## VI. DISCUSSION

The experiments show that a real-world application can be made network resource aware using the network resource abstraction offered through the NAL. They show that, although offering a full and detailed view on the available network resources, the resource abstraction can be applied in a selective way, using specific information when needed and more generic information when highly specific details are not necessary. Note that often a lot of functionality is already available; the application logic is just not adapted to dynamic network resources. Additionally, the experiments show that the abstraction can be offered without a high power consumption penalty.

When considering the possible usage of the resource abstraction by the remote control example application as described in section II, a number of other advantages can be illustrated. Assuming a configuration with multiple appliances being approached using IP over 802.11 infrastructure and Bluetooth point-to-point links. Establishing connections with these appliances may require – without our facility – the use of various system APIs, as listed in table III.

TABLE III: WINDOWS CE SYSTEM APIS AS USED BY THE REMOTE CONTROL EXAMPLE IN ABSENCE OF NRM/NAL

API	Description
ndisuiio	Enable 802.11 scanning, query 802.11 settings
wzc	Activate / deactivate 802.11 links
iphlpapi	Lookup IP settings, mapping from IP to interface identifiers, interface identifiers to interface names, route table information
ws2bth bthutil winsock2	Control Bluetooth discovery, query Bluetooth settings and activate/deactive Bluetooth links (Microsoft stack), or
Broadcom SDK	Bluetooth functionality as above (Broadcom stack)

The use of the different APIs and libraries requires the developer to manually associate the entities between these APIs which includes dealing with differences in semantics. The definition of links between entities is an explicit part of the NRM. Additionally, certain pieces of functionality are vendor specific: in this case, the same Bluetooth functionality may be accessible through different APIs, depending on the vendor and chipset of the currently used device. Furthermore, the available APIs do not allow selecting a ‘level of abstraction’ as offered through NRM / NAL: they require the developer to descend to a highly technology specific level to obtain such information as network identifiers. Our facility offers the necessary functionality through a single API, in a consistent manner with the right level of abstraction.

The subscribe/notify mechanism used for the interaction between applications and the NAL in combination with `ioctl` style commands is sufficient to obtain the right level of network resource information and control. The number of information entities that describe the current status are few (typically below 100) which makes advanced query mechanisms unnecessary.

One important limitation of our approach is the lack of features to add characteristics to types in a dynamic manner. This could be added to the expense of introducing more complexity and more run-time demands, and is most likely hard to implement as a simple C system interface.

## VII. RELATED WORK

The work related to the topic of this paper covers a number of different areas: models for cross-layer information interchange, abstractions for network resources, mechanisms for application and network adaptation, mobility management in heterogeneous network environments, and existing system facilities. This section describes a non-exhaustive selection of related work coming from these areas.

The framework presented in [10] introduces aspects of application awareness towards the mobility management process on a mobile device. It defines triggers relevant to a mixed set of mobility management protocols and shows how applications may interact with these. As such, the work presented here has a broader scope because it looks beyond application awareness in relation to mobility management protocols only. The scope becomes broader in [9] with the definition of ‘communication context’ that takes into account both application notification and application control, but it does not elaborate on implementations and experiments.

The MobileMan project [2] defines a cross-layer reference architecture for the exchange of data between layers in a conventional protocol stack (network, transport, application, etc.). It uses a ‘Network Status’ vertical element to keep each layer informed about the state of other layers, with the explicit goal to support the dynamics in mobile ad-hoc networks and provide performance optimizations. Our work incorporates cross-layer information exchange with particular focus on the application layer for application adaptation, and does look at cross-layer performance enhancements. It is not clear, however, how the cross-layer information is represented.

The work presented in [11] provides a formal definition of network context for ad-hoc mobile nodes. It primarily looks into the specification of the ad-hoc network topology within the vicinity of a mobile host and defines the cost for paths taken through this topology, but does not elaborate, as we do, on other aspects such as mobility management facilities.

Many approaches exist that focus on ways mobile applications can adapt to changing resources. Although the work presented here proposes a means to supply information on network resources so that applications may be adaptive and does not specify a mechanism for application layer adaptation itself, the actual strategies for adaptation or mechanisms forcing the application in a certain form of adaptation are important because they provide hints on what kind of information is relevant. The Odyssey platform for adaptive mobile data access models the adjustment of applications around the high level concepts of agility and fidelity. The adaptation functionality is divided between the application and the operating system. Experience with this system shows that adaptation must balance between agility and stability [8]; rapid changes in network resources must, for a more stable user perception, not always result in swift adaptation.

The AwareCon architecture described in [1] seeks to incorporate context information in the network itself, and supports adaptation at various layers in the network stack. The context information is distributed as a part of the payload of the communication packets set between the nodes that make up a network. It primarily focuses on low-power nodes and applications mostly operating in an ad-hoc fashion. Our focus is more on the adaptation of applications given a certain set of network resources, assuming standard network hardware and network stacks.

In our definition of the network resource model, a mobile host may provide multiple facilities for mobility management to handle the dynamically changing point(s) of attachment to the Internet. The work described in [13] introduces a mechanism that adapts the kind of mobility management depending on the behavior of the application, i.e. the kind of connection initiated by the application. The application itself is not adaptive.

The API presented in [4] shows similarities with the NAL interface. It has a modular architecture and defines a query interface for retrieving network resource info. It, however, primarily focuses on the link layer – not providing a cross layer interaction as we do – and does not elaborate on how the different needs of different applications are handled.

The MIB types [7] share some aspects with the NRM entities defined here, although the scope and objectives are quite different. Many MIB types are not relevant from a (mobile) host perspective and vice versa: MIB does not seem to support a wide variety of wireless networks.

The topics covered by the IETF Detecting Network Access (DNA) workgroup as well as the IEEE 802.21 standardization effort on handover between heterogeneous network look at mechanisms to transfer information from one layer to another. However, these have a more narrow scope, focusing on interaction between link and network layer.

## VIII. CONCLUSIONS

We have introduced a network resource abstraction suitable for application on mobile devices. It is extensible from two perspectives: new types of network resources can be added to the network resource model by inheriting from existing, more basic type, and the implementation offers the abstraction to application can be extended with plug-ins.

The application analysis shows that our approach addresses a real problem in an operating environment that is used by millions on a daily basis: we recognize the mobile host as a major computing platform. We do stick to existing entities available at the various layers of the protocol stack and defined as part of existing network technologies, but provide levels of abstraction by means of type hierarchies.

The experiments show that the NAL implementation can support real-world applications and that it does not consume substantial amounts of energy.

## IX. AVAILABILITY

The browsable and annotated NRM XML schema definition, the NAL reference implementation source code in C for Windows CE, as well as various Windows CE NAL binaries can be obtained from [3].

## X. ACKNOWLEDGMENTS

We would like to thank Jacco Brok for his contributions to our discussion on the network resource model. Furthermore, we would like to thank Stanislav Pokraev for suggesting SlimServer.

## XI. REFERENCES

- [1] M. Beigl, A. Krohn, T. Zimmer, C. Decker, and P. Robinson, "AwareCon: Situation Aware Context Communication", In *Proceedings of the 5th International Conference on Ubiquitous Computing (UbiComp 2003)*, Oct. 2003
- [2] M. Conti, G. Maselli, G. Turi, and S. Giordano, "Cross-Layering in Mobile Ad Hoc Network Design", In *IEEE Computer*, Feb. 2004
- [3] CoSphere NAL: Network Abstraction Layer Software for Windows CE, Retrieved on 18/8/2006 from <http://cosphere.telin.nl/nal/>
- [4] M. Bandholz, A. Gefflaut, J. Riihijärvi, M. Wellens, and P. Mähönen, "Unified Link-Layer API Enabling Wireless-Aware Applications", In *Proceedings of the 17th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC'06)*, Sept. 2006
- [5] GSPlayer software, Retrieved on 18/8/2006 from <http://hp.vector.co.jp/authors/VA032810/>
- [6] V. Kawadia and P. Kumar, "A Cautionary Perspective on Cross Layer Design", *IEEE Wireless Communication*, Feb. 2005
- [7] K. McCloghrie and M. Rose (Ed.), "Management Information Base for Network Management of TCP/IP-based internets: MIB-II", *IETF RFC 1213*, March 1991
- [8] B. Noble and M. Satyanarayanan, "Experience with adaptive mobile applications in Odyssey", In *Mobile Networks and Applications 4*, pp. 245-254, 1999
- [9] A. Peddemors, H. Eertink, and I. Niemegeers, "Communication Context for Adaptive Mobile Applications", In *Workshop Proceedings of the 3rd Conference on Pervasive Computing - Middleware Support for Pervasive Computing Workshop (PerWare'05)*, March 2005
- [10] A. Peddemors, H. Zandbelt, and M. Bargh, "A Mechanism for Host Mobility Management supporting Application Awareness", In *Proceedings of the Second International Conference on Mobile Systems, Applications, and Services (MobiSys'04)*, June 2004
- [11] G.-C. Roman, C. Julien, and Q. Huang, "Network Abstractions for Context-Aware Mobile Computing", In *Proceedings of the 24th International Conference on Software Engineering*, pp. 363-373, May 2002
- [12] SlimServer software, Retrieved on 18/8/2006 from <http://www.slimdevices.com/>
- [13] X. Zhao, C. Castelluccia, and M. Baker, "Flexible Network Support for Mobility", In *Proceedings of the Fourth Annual International Conference on Mobile Computing and Networking (MobiCom'98)*, Oct. 1998
- [14] W3C XML Schema, <http://www.w3.org/XML/Schema>