



AUTOCONF - Stating the Problem

Thomas Heide Clausen, Ulrich Herberg

► To cite this version:

Thomas Heide Clausen, Ulrich Herberg. AUTOCONF - Stating the Problem. [Research Report] 2007. inria-00192959v2

HAL Id: inria-00192959

<https://inria.hal.science/inria-00192959v2>

Submitted on 30 Nov 2007 (v2), last revised 4 Dec 2007 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

AUTOCONF - Stating the Problem

Thomas Heide Clausen — Ulrich Herberg

N° ????

November 2007

Thème COM

A large blue rectangle occupies the lower half of the page. Overlaid on the left side of this rectangle is a large, light grey stylized letter 'R'. To the right of the 'R', the words 'Rapport de recherche' are written in a white serif font, with 'Rapport' on the top line and 'de recherche' on the bottom line. A horizontal grey brushstroke underline is positioned below the text.

*Rapport
de recherche*



AUTOCONF - Stating the Problem

Thomas Heide Clausen, Ulrich Herberg

Thème COM — Systèmes communicants
Projets Hipercom

Rapport de recherche n° ???? — November 2007 — 32 pages

Abstract:

This memorandum outlines the goals for and constraints on an IP address and prefix autoconfiguration mechanism for mobile ad hoc networks.

Key-words: mobile network, ad hoc network, network architecture, address configuration, routing, IP networks, MANET

AUTOCONF - Stating the Problem

Résumé : Ce mémorandum détaille les objectifs et contraintes pour un mécanisme d'autoconfiguration d'adresse IP et préfix pour des reseaux ad hoc.

Mots-clés : réseaux mobiles, réseaux ad-hoc, architecture de réseau, configuration d'adresses, protocole de routage, réseaux IP, MANET

1 Introduction

A Mobile Ad hoc NETwork (MANET) is, essentially, a collection of mobile nodes, communicating among themselves over wireless links and thereby forming a dynamic, arbitrary and often multi-hop graph. The core tenants of MANETs are their autonomy (*ad-hoc-ness*), *i.e.* that they can be deployed with no a priori configuration, and their ability to cope with a high degree of network topology dynamism (the term *mobile* is indicative of this, although a dynamic network topology may be due to factors other than mobility).

MANETs are largely made possible due to the emergence of commodity wireless radio interfaces, and protocols maintaining connectivity in a highly dynamic, multi-hop wireless network have been developed and standardised by the IETF and by others in form of routing protocols such as [15].

Such MANET routing protocols operate with the assumption that MANET routers (MRs) already have IP addresses and, if applicable, prefixes assigned¹ – however no dedicated mechanisms for assigning IP addresses or distributing prefixes to MRs have been standardised by the IETF. It is worth noting that academic literature has presented numerous such mechanisms, however that these may not necessarily have been developed with an explicit requirement of seamless integration with the Internet architecture and other IETF-related protocols.

1.1 Memorandum Outline

The remainder of this memorandum is organised as follows: section 2 outlines the notion of IP addresses and prefixes, as is currently employed in the Internet, with the purpose of retaining the basic notions and properties of addresses and prefixes. Section 3 briefly recapitulates the architecture of MANETs as seen by the IETF and described in [9] and [10]. Specific effort is made to include those architectural issues which are pertinent to understanding the requirements and constraints that an IP autoconfiguration mechanism for MANETs must satisfy, as well as to define a terminology used in later sections.

Two common terms are used when describing MANET deployments: "*connected MANETs*" and "*disconnected MANETs*", with the latter also sometimes referred to as "*stand-alone MANETs*". These terms are **wrong, misleading, unfortunate** and **unhelpful** since they do not provide an useful classification of MANETs – in fact, these terms do not provide any classification of MANETs. Section 4 explains why in detail, and proposes an alternative taxonomy – which, contrary to "*connected MANETs*" and "*disconnected MANETs*", accurately classifies MANET deployment scenarios with respect to autoconfiguration issues.

Section 5 explicitly, and with reference to section 2 and section 3, tries to answer the simple question of "what are the entities that are to be configured in a MANET, and what does it mean to configure those?". In other words, which of interface(s), router(s) are to get which of addresses and prefixes. In a sense, this section is stating the problem that is to be solved.

¹At this point, the term "assigned" is used for both prefixes and addresses being "given" to a given MR. A distinction will be made in section 2.

In the Internet today, mechanism for automatically configuring network interfaces with IP addresses already exists. Similarly, mechanism for automatically delegating a prefix from a given router to another given router exists. Section 6 briefly summarises those, as well as outlines the assumptions that these mechanisms are designed to work under. While these mechanisms can be said to "solve the problem" stated in section 5 for the Internet, the assumptions under which they do "solve the problem" are closely related to the classic model of an IP link.

MANET interfaces and the logical structure of a MANET differs in part from the classic model of an IP link, as described in section 3. Verbatim application of the mechanisms for automatically configuring network interfaces with IP addresses described in section 6 may therefore not be possible, or may not satisfy the goals stated in section 5 – and this is described in section 7 in some detail.

Section 8 recapitulates the key points brought forward in these discussions, and thus concludes this memorandum.

2 What is an IP Address?

Before elaborating on the specific issues regarding IP autoconfiguration of MANETs, as well as on existing IP autoconfiguration mechanisms, it is important to retain the basic notions of addresses – and their complements, prefixes. This section will therefore briefly describe the nature of addresses – which includes introducing the terms “address assignment” and “prefix delegation”. The discussions in this section are intended to provide an overview and motivation for the IP addressing architecture in the Internet, and not to be an exhaustive architectural discussion of the IP addressing architecture.

2.1 The Basics

In the most basic form, an IP address can be considered simply as follows:

- *an IP address is a unique identifier of a network interface.*

I.e., an IP address allows, uniquely, to address IP datagrams to one network interface among the many network interfaces present in the Internet².

According to this, the important property of an IP address is its network-wide uniqueness – and absent any other constraints, in order to uniquely identify any network interface on the Internet, this would be a sufficient property to ensure.

2.2 Networks

The Internet is big — very big — so a large number of IP addresses are required in order to ensure that each network interface can have a unique IP address. IPv4 offers 2^{32} unique addresses, and IPv6 offers 2^{128} unique addresses. If, as indicated in section 2.1 no other constraints than uniqueness were imposed, this would imply that each router in the Internet would have to maintain a routing table explicitly containing 2^{32} or 2^{128} entries – and, in order to construct its routing table, would have to exchange this much information with other routers. In order to reduce the amount of state information that each router would have to maintain (the size of the routing table) and exchange, constraints other than uniqueness are imposed on IP addresses as assigned throughout the Internet.

The basic idea in reducing the state that routers are to maintain and exchange is simple: assign IP addresses in an organised fashion, as indicated in the example in figure 1. Thus, all network interfaces with IP addresses from within a given interval would be attached to the same router – and the routers would each have to maintain and exchange state for each interval, rather than for each explicit IP address.

²A note may be in order to point out that an IP address is not on its own considered a unique identifier of an “end-point of communication”, since that requires to not only identify the network interface of a networked entity, but also the process running on that entity that should receive a given IP datagram. In the Internet, an IP address AND a process identifier (a port number or a protocol number) serves as the unique identification of an “end-point of communication”. Incidentally, port numbers and protocol numbers are, therefore, often also termed “code points”.

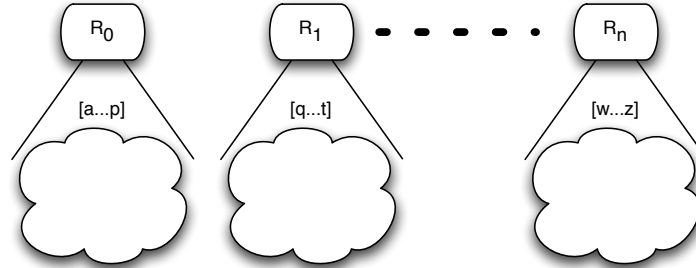


Figure 1: **Flat topology of routers:** each router having an address interval, with all network interfaces configured with addresses from within that interval being directly connected to that router.

With 2^{32} or 2^{128} IP addresses, if the Internet was constructed uniquely as in the simple example in figure 1 illustrates, it would imply either a large number of routers, each with a still large number of network interfaces attached – or a small number of routers, each with even more network interfaces attached. A large number of routers still entails a large amount of state for each router to maintain and exchange – and a small number of routers each with a very large amount of network interfaces attached would concentrate traffic with ensuring capacity and reliability issues. Thus, and with respect to IP addresses, each of the clouds from figure 1 may itself contain routers, as illustrated in figure 2.

IP addresses within a cloud in figure 2 are assigned in the same organised fashion as in the example in figure 1, although now hierarchical: all addresses from within the interval $[a-p]$ are in the cloud “under” router R_0 ; all the addresses from within the interval $[a-c]$ are also “under” router R_1 etc.

The term used for giving a router an address interval and allowing that IP addresses from within that interval are attached to only *that* router is called *delegation*. Thus, an IP address assigned to a network interface indicates *where* in the hierarchy this network interface is located.

This implies two things, which are important to retain:

- a router is *delegated* an IP address interval from someone, who has authority to do so
 - i.e. from a router has had (at least) that IP address interval delegated itself;
- two properties apply to IP addresses:
 - an IP address is a *unique identifier of a network interface* (as in section 2.1);
 - an IP address implies a *precise topological location in the addressing hierarchy*.

The last item is of special importance, since it implies that an address is not *just* an identifier, but that there’s a semantics associated to an address – and that this semantics is important for the organisation of the Internet.

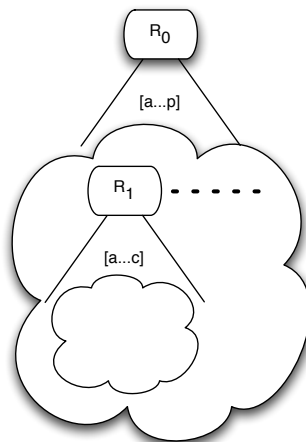


Figure 2: **Hierarchical topology of routers:** router R_0 delegates a subinterval $[a...c]$ from its interval $[a...p]$ to the subordinate router R_1

2.3 Prefixes – the Internet Term for Intervals

Syntactically, an IP address is a simple sequence of bits – 32 bits for IPv4 and 128 bits for IPv6. In order to achieve an easy way of designating "IP address intervals" and "sub intervals" of such an IP address interval, prefixes are used. Structurally, a prefix is the n leftmost bits of an IP address, which are common to all IP addresses from within that interval, and n is said to be the *prefix length*. A classic example of a prefix and a prefix length from IPv4 would be 130.225.194/24, where 130.225.194 is the prefix and 24 subsequently is the prefix length – indicating an IP address interval of [130.225.194.0–130.225.194.255].

The general notation for a prefix in IPv6 is $p::/n$, implying that p is the prefix, the length of this prefix is n , and the $::$ indicating that everything following the prefix (*i.e.* the bits other than the n leftmost bits of the address) is a sequence of zeros³.

Thus, and with reference to figure 3, if the router R has the prefix $p::/32$ assigned, subordinate routers R_1 , R_2 and R_3 can be delegated sub-intervals by simply delegating these the prefixes $p:1::/64$, $p:2::/64$ and $p:3::/64$ respectively.

It shall be noted that:

- the prefixes thus delegated to these three routers are non-overlapping – *i.e.* the address ranges from which the routers R_1 , R_2 and R_3 may respectively assign addresses are also non-overlapping;

³Although not relevant for the present discussion in this memorandum, note that $::$ in general in IPv6 indicates any number of consecutive four-digit groups (in hexadecimal) whose value is zero. Thus, $p::1$ would be a valid way of representing an IPv6 address where the first bits are p , followed by one or more four-digit groups of zeros, followed by 1

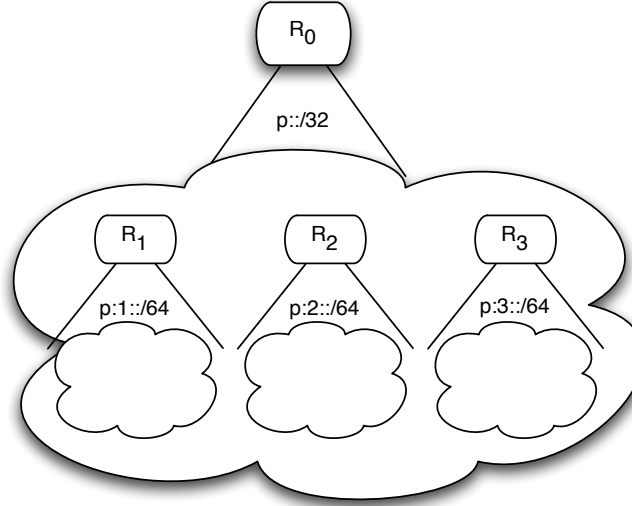


Figure 3: **Prefix delegation:** router R_0 has delegated non-overlapping subprefixes $p:1::/64$, $p:2::/64$ and $p:3::/64$ of its prefix $p::/32$ to its subordinate routers.

- prefixes delegated to the routers R_1 , R_2 and R_3 overlap with the prefix delegated to R_0 , and thus the precise topological location of the routers R_1 , R_2 and R_3 is subordinate to R_0 .

Commonly, and for the remainder of this memorandum, prefixes will simply be written as $p::$ or $p:1::$, i.e. the prefix length will be omitted when written. The prefix length is in this case simply assumed equivalent to the length of the leftmost and explicitly written bits, i.e. that which appears to the left of the $::$.

2.4 Delegation and Assignment

Prefix delegation is, thus, the way in which a router acquires an IP address interval, which is topologically correct with respect to its location in the addressing hierarchy. From within this IP address interval, the router may – as detailed in section 2.3 – either delegate parts to subordinate routers, or assign addresses to network interfaces attached to it, as illustrated in figure 4. The important distinction in this is, that a delegation of a prefix does not provide any network interface with a unique identity but rather allocates part of the IP address interval for exclusive use by a subordinate router – whereas address assignment implies that a network interface is configured with a specific unique and topologically correct IP address.

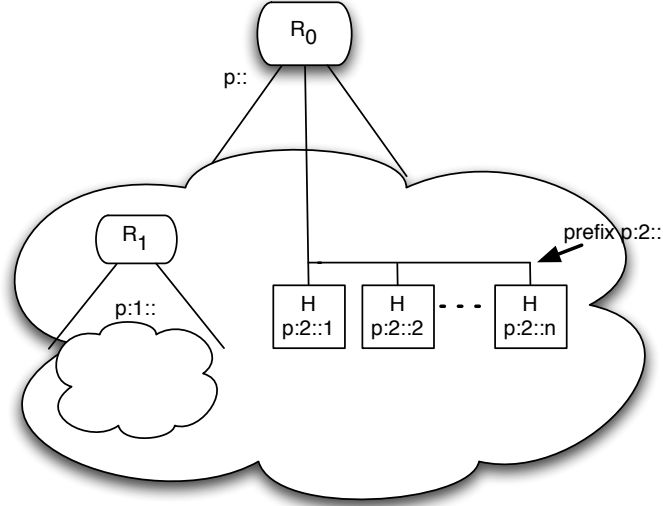


Figure 4: **Prefix Delegation and Address Assignment:** a router assigning IP addresses from part of its prefix to network interfaces directly attached, and delegating another part of its prefix to a subordinate router.

2.5 Summary

In this section, basic notions of IP addresses and prefixes, as is currently employed in the Internet, have been outlined. An IP address can be assigned to a network interface as a unique identifier in the network. In addition, it is also a precise topological location in the addressing hierarchy. Prefixes are intervals of addresses that can be delegated to routers for the purpose of either redelegating parts to subordinate routers, or assigning addresses from the interval to network interfaces attached to them.

3 The MANET Architecture

[9] and [10] describe, in detail, an architectural view of MANETs, as when seen as part of the Internet and thus integrated in the Internet architecture. In this section, the basic concepts relevant to a discussion on IP autoconfiguration for MANETs will be briefly reiterated. In particular, the notion of a "link" as well as the morphology of a MANET router (MR) is important for the discussion in this memorandum and are, therefore, presented in section 3.2 and section 3.3, respectively. In these sections, reference will be made to what can be called the *Classic IP Link Model*, which is therefore introduced in section 3.1.

3.1 The Classic IP Link Model

A core notion in the Internet is that of a *link*. A link, illustrated in figure 5, is assumed to have the following properties:

- IP datagrams are not forwarded at the network layer when communicating between network interfaces which are on the same link; hence
- TTL/hop-limit in IP datagrams are not decremented when communicating between network interfaces on the same link;
- IP datagrams with a TTL/hop-limit of 1 are (modulo data loss) delivered to all network interfaces on the same link and;
- link-local multicasts and broadcasts are received by all network interfaces on the same link – without forwarding.

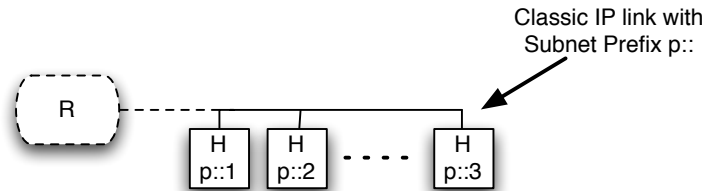


Figure 5: **Classic IP Link Model:** hosts (H) connected to the same link have assigned IP addresses from a common prefix, possibly assigned by a router (R).

When assigning an IP address to a network interface, this network interface is also configured with a prefix, such that the following constraint is respected:

- all network interfaces configured with addresses from within the same prefix $p::$, and with the same prefix $p::$ assigned to the interfaces, can communicate directly with one another.

It follows from the above that the notion of "IP link" is tied with the notion of an "IP Subnet" (IPv4) or a prefix (IPv6), in that all network interfaces which are configured with the same subnet address or prefix are considered to be on the same IP link and thus that for communication between nodes within the same subnet, no forwarding is required and no decrement of TTL/hop-limit is performed.

Network interfaces within the same prefix or, for IPv4, within the same subnet, are within the classic IP link model assumed to also be attached to the same classic IP link as described above. For completeness, it should be mentioned that the inverse is not necessarily true: in some network configurations, interfaces connected to the same classic IP link may be configured within different prefixes or subnets.

3.2 The Missing Link

MANET interfaces are a specific class of network interfaces in that they exhibit what is commonly denoted as *semi-broadcast characteristics*. With reference to [10], this roughly implies that otherwise neighbouring nodes may experience distinctly different local connectivity. As indicated in figure 6, the MANET interface of N2 is able to directly communicate with the MANET interface of N1 and N3 (*i.e.*, no forwarding, TTL/hop-limit is not decremented, a transmission from the MANET interface N2 is received by the MANET interfaces of both N1 and N3, including link-local multicasts and broadcasts).

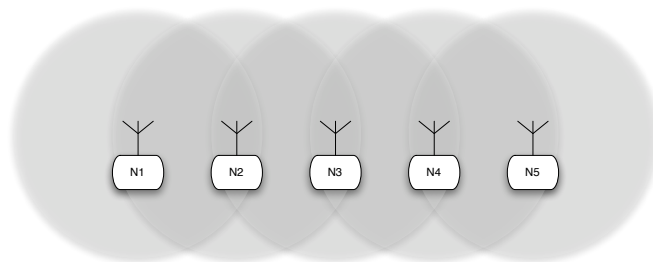


Figure 6: **MANET**: nodes (N) with MANET interfaces. The light grey area indicates the **coverage area** of each MANET interface.

Considering the properties listed for an IP link in section 3.1, this might imply that the MANET interfaces of N1, N2 and N3 would be connected to the same link. However the semi-broadcast nature of MANET interfaces implies that this is not so: transmissions from N1 will not reach N3 without forwarding – which entails that TTL/hop-limit is decremented, and that link-local multicast and broadcast from the MANET interface of N1 will not reach the MANET interface of N3.

The semi-broadcast nature of MANET interfaces implies that:

- any two MANET interfaces can not be assumed to be connected to the same link; thus

- the IP address / prefix configuration of MANET interfaces must be such that their configuration (as per the constraint in section 3.1) does not indicate that they can be assumed to be connected to the same link.

3.3 The Morphology of a MANET Router

A MANET router is a router having, at least, one MANET interface – *i.e.* an interface exhibiting semi-broadcast characteristics as indicated in section 3.2 – and otherwise retaining the usual characteristics of a router. In particular, a MANET router:

- may have a prefix delegated to it;
- may delegate (part of) that prefix to other subordinate routers;
- may have networks attached to it, in particular these networks may be of any type, (including, but not limited to, other MANETs), thus respecting usual routing and addressing hierarchies.

3.4 'M' for Mobility – 'A' for Ad hoc

As indicated in section 1, MANETs are characterised by being mobile – *i.e.* that the network topology, in particular the set of participating MANET Routers as well the ability for two MANET routers to communicate among themselves may change over time. In particular, this entails that a formerly single MANET may separate into two independent MANET, and that two formerly independent MANETs may merge into one single larger MANET.

Additionally, the 'A', for ad hoc, implies that the network may emerge spontaneously and, potentially, without any pre-determined infrastructure elements (or without any such elements altogether) such as a router providing access to an outside network such as the Internet. This implies that a MANET may appear either in isolation, where the *uniqueness property* of IP addresses is the sole priority, configuration-wise – or in a situation where it may be integrated/attached to an outside network where, in addition to the *uniqueness property*, IP addresses / prefixes should correspond to the topologic location of the MANET within the addressing hierarchy.

This latter point is somewhat analogous to other IP networks, in which a collection of hosts may connect with each other in a private network without connection to outside networks and wish to acquire IP addresses where the *uniqueness property* of IP addresses / prefixes is the sole priority (link-local addresses) – as well as be part of a network where in addition to the *uniqueness property*, IP addresses / prefixes should correspond to the topologic location of the network interfaces within the addressing hierarchy (global addresses). As will be detailed in section 6, these two are both supported in the Internet via IPv6 Stateless Autoconfiguration – and, an analogy to MANETs exist, as indicated in this section.

3.5 Summary

A MANET router is a router which has at least one MANET interface, *i.e.* a network interface exhibiting semi-broadcast characteristics, but which otherwise retains the usual characteristics of a router. Thus, the singular issue with respect to MANET routers is, that the constraints regarding IP address and prefix configuration of network interfaces must be respected, also on the MANET interface(s) – which in particular entails that:

- any two MANET interfaces can not be assumed to be connected to the same link; thus
- the IP address / prefix configuration of MANET interfaces must be such that their configuration (as per the constraint in section 3.1) does not indicate that they can be assumed to be connected to the same link.

4 Connected vs. Disconnected MANETs ???

Occasionally, the terms "*connected MANET*" and "*disconnected MANET*" (or, for the latter also "*stand-alone MANET*") are employed when describing different deployments scenarios (for example in [12, 13, 14, ?, 16]). These terms are both **unfortunate** and **wrong** – and their use should be **discouraged** and **discontinued**. This section will show why they're wrong in section 4.1, why they're unfortunate in section 4.2 – and propose an alternative and correct way of describing different MANET deployment scenarios in section 4.3.

4.1 On why these terms are wrong

In order to correctly analyse why the terms "*connected MANET*", "*disconnected MANET*" and "*stand-alone MANET*" are wrong, consider first the scenario depicted in figure 7. This MANET contains four MANET routers (MRs), each of which having (of course) one MANET interface and one additional network interface towards a classic IP link. MR₁ has, on this classic IP link, a (non-MANET) router, R, which has another network interface towards another classic IP link. There are no components (other than, perhaps, hosts) in this network, other than those depicted in figure 7.

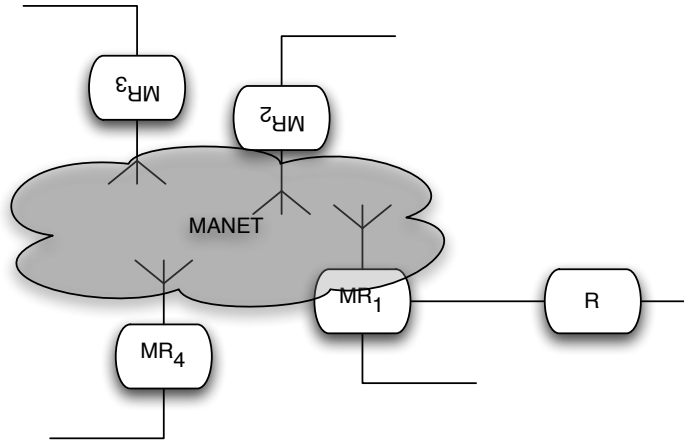


Figure 7: **A MANET:** four MANET routers (MRs), with MR₁ connected via a classical IP link to a non-MANET router R. Is this a "*connected MANET*"?

The obvious question to ask, then, if this MANET in figure 7 is a "*connected*" MANET or if it is a "*disconnected MANET*"?

4.1.1 Graph Theory?

Being "connected" may refer to a property of a graph, recalling that a *connected graph is an undirected graph that has a path between every pair of vertices*.

A MANET is intended – even expected – to possess this property, which the MANET in figure 7 certainly does possess.

- Thus, if the intended semantics of a "*connected MANET*" is with reference to graph theory, then the MANET in figure 7 indisputably is a "*connected MANET*".

4.1.2 Outside Networks?

The common use of the term "*connected MANET*", however, seems intended to carry a semantics that go beyond that of graph theory. Assuming that MANETs are supposed to be composed of a connected graph of MANET routers, a MANET being "*connected*" may be in reference to the existence of links from MANET routers over non-MANET interfaces to other networks. Considering the scenario in figure 7, this MANET indisputably is connected to another network over a non-MANET interface – the easiest example hereof is the existence of a connection from the MANET via MR1 to R.

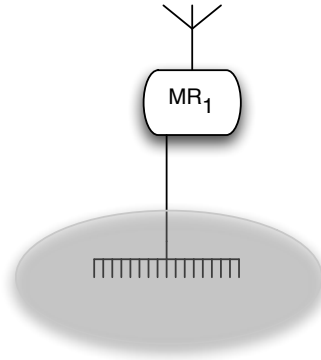


Figure 8: **A Minimal "*connected*" MANET:** the mobile router MR₁ has a single MANET interface and a single network interface towards a classic IP link.

Considering an even simpler MANET, such as the one depicted in figure 8 where a single MANET router, MR₁, has a single MANET interface and a single network interface towards a classic IP link, the MANET can be said to be composed of the single MANET interface – which, indeed is connected to another network, indicated by the grey oval, via MR₁.

- Thus, if the intended semantics of a "*connected MANET*" is with reference to the existence links connecting the MANET to other networks, then the MANET in figure 7 indisputably is a "*connected MANET*". Moreover, a single MANET router with

a single MANET interface and a single interface connected to a classic IP link is indisputably a *connected MANET* – and, indeed, so is every possible MANET.

4.1.3 The Internet?

It may be that the intention of the term "*connected MANET*" is to indicate if a MANET is connected to "The Internet", with the term "*disconnected MANET*" being used otherwise. There are a number of problems inherent in this usage, since:

- "The Internet" as such does not exist as an unique and identifiable entity, but is rather a concept of "a network of networks" [11]. Thus, what does it mean to be connected to the Internet? That one can access a particularly common www search-engine? Are "Internets" (other than the one where this particularly common www search-engine exists) then not "Internets"? If not, then what, other than arbitrary existence of a given service, distinguishes the "real" and the "other" Internets.
- Even if an unique "Internet" could be identified and defined, this would not suffice, since:
 - a router is a router is a router, thus a router which connects to another router which is part of this "Internet" does not become "connected to" the Internet – but becomes "part of" the internet. A MANET router connecting via a network interface to a classic IP link and thereby to another router which is "part of" the "Internet" therefore also becomes "part of" this "Internet".
 - specifically, routers in the "Internet" are not segregated into classes of "Internet routers" and "Internet routers, connected to the Internet";
 - as such, routers do not signal if they are indeed "Internet routers" or not – implying that the scenario in figure 7 would be indistinguishable from the scenario where the outside network was this "Internet".
- Thus, if the intended semantics of a "*connected MANET*" is with reference to the existence of links connecting the MANET to the "Internet", then if a given MANET is a "*connected MANET*" or not is indeterminable.

4.2 On why these terms are Unfortunate

The use of the terms "*connected MANET*" and "*disconnected MANET*" indicates a desire for classifying MANETs into two categories, with networks of either category having a clearly defined and easily identifiable set of properties, distinct from networks of the other category.

The fact that since a single MANET router with a single MANET interface and a single network interface connected to a classic IP link, as well as more complex MANETs can be said to be a "*connected MANET*", as per both section 4.1.1 and section 4.1.2, indicates, that all MANETs are, indeed, in the "*connected MANET*" category. Or, in other words, the clearly defined and easily identifiable set of properties that would set a "*connected MANET*"

and a "*disconnected MANET*" apart are not to be found in neither graph theory nor in the existence of links from the MANET to outside networks.

Since, as per section 4.1.3, if a MANET, or more generally "any router" is "connected to the Internet" or not, is indeterminable, using this as a classifying metric is unfortunate for the same reasons: it does not provide a well-defined set of criteria for determining if a MANET is a "*connected MANET*" or not.

All of the above illustrates that these proposed classifications and terms are unfortunate, or perhaps unnecessary, since the terms "*connected MANET*", "*disconnected MANET*" and "*stand-alone MANET*" do not clearly specify scenarios in which distinct problems are present and therefore where distinct considerations are required.

4.3 A Better Taxonomy

While the terms "*connected MANET*" and "*disconnected MANET*", as described in section 4.1 and section 4.2 are **wrong** and **unfortunate** and therefore should be discouraged and discontinued, they have been proposed for a reason – a reason which is valid enough, but which needs to be clearly described and, consequently, a valid taxonomy needs be defined.

In order to understand the issue, it is important to recapitulate the properties of an IP address, as described in section 2 and repeated below:

- *an IP address is a unique identifier of a network interface* (as in section 2.1);
- *an IP address implies a precise topological location in the addressing hierarchy.*

Of particular importance for the discussion in the following is the second of these two properties. Considering the two scenarios in figure 9 for illustrating the importance this makes.

4.3.1 Autonomous MANETs

In figure 9a, the MANET exists, with a number of connections to other networks. The particularity of this scenario is, that none of these networks imposes an addressing hierarchy on the MANET. To make reference to the tree-like figures in section 2, the MANET is the "root" of the tree, with routers in the other networks being subordinate to the MANET. As such, these non-MANET routers should have prefixes delegated to them from the MANET, such that these prefixes reflect the non-MANET routers topological location below their upper MANET routers.

- in the network illustrated in figure 9a, the MANET is not subordinate to any "upper router" which imposes an addressing hierarchy – and, in particular, the MANET is the "root" of the addressing hierarchy.

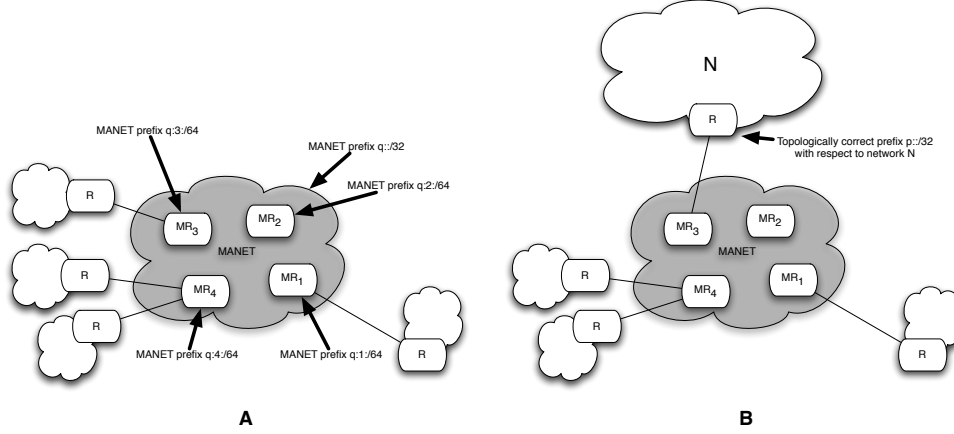


Figure 9: **Autonomous and Subordinate:** On the left side, the MANET is the "root" of the addressing hierarchy. On the right side, the MANET is imposed an addressing hierarchy by a superordinate router.

4.3.2 Subordinate MANETs

In figure 9b, the MANET exists, with a number of connections to other networks. The particularity of this scenario is, that one of these networks, the one indicated with a bold-face "A", imposes an addressing hierarchy on the MANET. To make reference to the tree-like figures in section 2, the MANET is an "internal node" of the tree, with the router R_A in the network A. Thus, to ensure that the second property of IP Addresses (that they be correct with respect to a topological location in the addressing hierarchy imposed by the network A), IP addresses within the MANET must be assigned from within a prefix delegated to R_A – possibly by R_A delegating part of its prefix for use within the MANET.

As with the "Autonomous MANETS", described above, non-MANET routers in other networks, subordinate to the MANET, should have prefixes delegated to them from the MANET, such that these prefixes reflect the non-MANET routers topological location below their upper MANET routers.

- in the network illustrated in figure 9b, the MANET is subordinate to any "upper router" which imposes an addressing hierarchy – and, in particular, addresses within the MANET must respect this addressing hierarchy.

On the topic of subordinate MANETs, such may be *multi-homed*, i.e. be subordinate to two or more routers, each of which impose their own topological location in an addressing hierarchy. This situation is no different than that of any other multi-homing situation, in that the MANET may acquire a prefix from either or both of these, and assign addresses from either or both of these – so long as it is ensured that traffic from the MANET which

is generated with a given source IP address is transited through the appropriate router (i.e. the one from whose' prefix that source IP address is acquired).

4.4 Summary

This section takes exception to the usage of the terms "*connected*" and "*disconnected*" MANETs, arguing that these are both wrong, unfortunate and unhelpful – in particular since it is not possible to uniquely classify MANETs into either "connected" or "disconnected", and since they therefore do not assist in understanding if different considerations apply for different MANET deployments.

With reference to an understanding of address hierarchies, it is understood that there exists two distinct deployment scenarios for MANETs – with the distinction being if the MANET is subjected to an address hierarchy, imposed from an outside network to which it is connected, or not. The particularities of these two scenarios are discussed.

The **Autonomous MANETs** is proposed, classifying MANETs which are **NOT** subordinate to another network in an addressing hierarchy. The term **Subordinate MANETs** is proposed, classifying MANETs for which another network (or set of networks) impose an addressing hierarchy onto the MANETs.

The distinction of "Autonomous MANETs" and "Subordinate MANETs" allows, uniquely, to classify MANETs into two disjoint classes.

5 AUTOCONF – The Goals

With the general considerations on IP addresses and prefixes in section 2, and the taxonomy in section ?? identifying two distinct MANET deployments, the goals that a MANET AUTOCONF solution should satisfy can be summarised into a simple set of goals:

- provide each MANET router with a unique prefix;
- respecting addressing hierarchies and the Internet addressing architecture;
- contain any MANET-specific behaviours to the MANET interfaces of MANET routers.

This section will detail these goals a bit further, starting with general architectural considerations in section 5.1, followed by specific considerations for Autonomous MANETs and Subordinate MANETs, in section 5.2 and section 5.3, respectively. Each of these sections will specify the precise goal for an AUTOCONF solution for each of these two MANET classes.

Section 5.4 will reflect on the architectural consequences of this approach, specifically with respect to that which applies to the MANET interfaces.

5.1 Architectural Considerations

The overall goal of AUTOCONF – AUTOMATIC CONFIGURATION for MANETs – is to ensure that network interfaces can be uniquely identified, respecting the addressing architecture in section 2 and the constraint outlined in section 3.1. In other words, that network interfaces are automatically configured, respecting the Internet and the IP architecture.

In a MANET router, this entails that the MANET interface is to be configured with an IP address and prefix, as detailed in [9] and [10] – and that network interfaces attached to a given MANET router, as well as routers subordinate to a given MANET router are to be, respectively, assigned IP addresses and delegated prefixes from within a unique prefix over which that MANET router is authoritative. This implies that:

- each MANET router must acquire a unique prefix, over which it is authoritative.

Under which constraints this prefix is to be acquired will be detailed in subsequent sections. With each MANET router having acquired such a unique prefix, the following considerations apply:

- a MANET router behaves, with respect to subordinate routers and networks, identical to any other router, in particular in respecting the addressing hierarchies;
- given that the MANET router has acquired a unique prefix, the MANET interface can be configured with an address from within this prefix and a prefix length of /128 (for IPV6) or /32 (for IPV4) – or use the MANET interface as an unnumbered interface [1].

5.2 Acquiring a Prefix in an Autonomous MANET

The particularity of an Autonomous MANET is the absence of an externally imposed addressing hierarchy – and thereby the absence of an entity which can provide unique and topologically correct prefixes to MANET routers.

In order for each MANET router to acquire a unique prefix, an AUTOCONF solution must therefore be developed that:

- allows MANET routers to, from within a well-known prefix similar to the well-known Link Local Prefix, and in a distributed fashion, each acquire a prefix which is unique within this MANET;
- addresses and prefixes from within this well-known prefix are characterised by:
 - being unique only within one MANET.

5.3 Acquiring a Prefix in a Subordinate MANET

The particularity of a Subordinate MANET is the presence of one or more externally imposed addressing hierarchies – and thereby the presence of one or more entities which can provide unique and topologically correct prefixes to MANET routers.

In a Subordinate MANET, the solution detailed for Autonomous MANETs may be applied in order to acquire MANET-wide unique and local prefixes for each MANET router. In addition, an AUTOCONF solution must be developed that:

- allows MANET routers to request delegation of a prefix from one or more or all of the externally imposed addressing hierarchies – a prefix or prefixes which will be correct with respect to the MANETs topological location within that/these addressing hierarchies.

5.4 Consequences

Recalling the tree-like addressing hierarchies from section 2, the set of MANET routers in a MANET appear as a single "internal node" in the tree. This means that:

- "downwards" in the tree, *i.e.* to connected networks and subordinate routers, the addressing hierarchy is respected with respect to each MANET router;
- "upwards" in the tree, if the MANET is a Subordinate MANET, the MANET routers will respect the addressing hierarchy with respect to the external entity(ies) from which prefixes have been acquired;
- "horizontally" within the MANET, and due to the potentially dynamic topology of the MANET, no permanent relationship – and, in particular, no hierarchy – is ensured between MANET routers. Negative consequences hereof are prevented by configuring

the MANET interfaces such that their configurations do not allow assumptions that they all be connected to the same "link", as detailed in section 3. This particularity is exposed **ONLY** horizontally between MANET routers, and is neither propagated upwards or downwards in the tree.

5.5 Summary

In this section, the goals of MANET autoconfiguration based on the architectural considerations and the difference between Autonomous and Subordinate MANETs have been classified. Each MANET router must be provided with a unique prefix. If the MANET is a Subordinate MANET, prefixes have to be topologically correct with respect to the imposed addressing hierarchies. In addition, only the MANET interface of a router has to be exposed to MANET-specific autoconfiguration requirements; classical IP networks attached to a non-MANET interface are not affected.

6 Automatic Configuration in the Internet

Other than assigning IP addresses and prefixes to network interfaces in a static manner, there are several mechanisms of automatic configuration. Multiple solutions have been standardised by the IETF and are wide-spread in implementations. In this section, an overview of the most common autoconfiguration mechanisms is presented. Among them are DHCP (presented in section 6.1), and Stateless Autoconfiguration (refer to section 6.2). These mechanisms rely on certain assumptions of the Internet, some of them being mentioned in section 6.3.

6.1 Dynamic Host Configuration Protocol (DHCP)

DHCP ([17] for IPv4, [18] for IPv6) enables automatic allocation of an IP address to a host by a DHCP server. A host requiring an IP address contacts a DHCP server and requests a new address. The DHCP server will dynamically⁴ assign an address from a certain interval of addresses, and allocate a so called “lease” of that address to the client. The client can then use the address for a certain time. If it wants to keep the address for a longer time, it has to prolong the lease. If the DHCP server is not in the same subnet as the DHCP client, it is possible to use a DHCP relay agent which can forward the messages to a different IP network.

For automatically delegating IPv6 prefixes to routers, there exists an option in a DHCP message called DHCP-PD (as specified in [19]). A router may demand a prefix from another router by sending a DHCP request including the DHCP-PD option. The DHCP server may then delegate a subordinate prefix of its own prefix to the client. The prefix delegation option may also be relayed through a DHCP relay agent.

6.2 Stateless Autoconfiguration (SLAAC) / Neighbour Discovery Protocol (NDP)

Stateless Autoconfiguration (SLAAC) (as specified in [21]) is another way of assigning IP addresses to hosts which has been introduced with IPv6 (and cannot be applied in IPv4). In contrast to DHCP, a host requiring an address does not need to contact any kind of server. Instead, the protocol is completely distributed.

A host first selects a tentative IPv6 address by attaching its host identifier (in most cases its MAC address) to the well-known link-local prefix. It then has to verify that no other host in the same subnet has the same address by broadcasting NDP messages [20] to the link. If the address is not unique, the autoconfiguration process will be aborted. Upon a successful address uniqueness test, a host may request a prefix from any router on the link by an exchange of NDP messages. It will again attach its host identifier to that router prefix and possibly repeat the address uniqueness test.

⁴For details of manual and automatic allocation refer to [17], [18]

6.3 IP architecture assumptions

The two mechanisms presented in the sections 6.1 and 6.2, are based on certain assumptions which hold in the Internet. Most importantly, these assumptions are:

1. Messages (e.g. DHCP Request, NDP) are link-locally broadcasted with a TTL of 1, and may thus never be forwarded by a router.
2. The local topology is relatively stable.
3. Routers acquiring a prefix from a delegating router are in a subordinate topology with respect to that router.

An example for assumption 1 are NDP messages. They have a hop limit of 255, and whenever a node receives an NDP message with a decremented hop limit, it will reject the message, and the uniqueness test will fail. Thus, NDP messages may never be forwarded by a router.

Assumption 2 means that in classical IP networks, links will not break down very frequently, and the topology of the LAN will not change in an ad-hoc manner. This assumption is for example important for DHCP where all clients are required to reach the DHCP server to acquire or prolong their lease. If the link is broken, the allocated address will be invalidated after the lease has expired.

Assumption 3 applies for instance for DHCP-PD. The requesting router is topologically subordinate to the delegating router, as depicted in figure 2.

6.4 Summary

In this section, the currently most-used autoconfiguration mechanisms DHCP and SLAAC/NDP have been described in their basic functionality. They do solve the autoconfiguration “problem” in the Internet, relying on certain assumptions which hold in the Internet. Important assumptions among them are that link-locally broadcasted messages will never pass any router, that the local topology is relatively stable, and that routers acquiring a prefix from a delegating router are in a subordinate topology with respect to that router.

7 Internet vs MANETs

As presented in section 6, there are several well-working standards for autoconfiguring IP addresses and prefixes in the Internet. These are all based on certain assumptions that are valid in the architecture of the Internet, some of them having been presented in section 6.3. However, as shown in section 3, MANETs have some different architectural properties in comparison to classical IP networks. These special properties avoid the verbatim usage of the presented autoconfiguration mechanisms, which will be demonstrated in this section.

7.1 Different link-model of MANETs

One reason why the autoconfiguration methods presented in section 6 cannot directly work on MANETs is the different link model of MANETs in comparison to the classic IP link model (refer to section 3.1 and 3.2 for a more detailed description of the link model in the MANET architectural model). In the classic IP link model, a link-locally broadcasted packet will never be forwarded by a router. However, to assure that MANET router prefixes are non-overlapping within a certain scope, uniqueness of router prefixes must be assured over several hops. This invalidates the usage of NDP for the prefix uniqueness test, as the NS and NA messages would have to be forwarded by routers.

7.2 Routers vs. Hosts

Several proposals for MANET autoconfiguration in literature suggest to simply use NDP messages to configure MANET nodes. However, this contradicts the MANET architecture as proposed in section 3 where MANET nodes are described as routers and possibly attached nodes. RFC 4862[21] specifically states: “*The autoconfiguration process specified in this document applies only to hosts and not routers*”. The reason for this is that one main goal of autoconfiguration is to allocate prefixes to MANET routers (refer to section 5) and not IP addresses to hosts. It has to be assured that these prefixes are non-overlapping so that hosts attached to that router can then be autoconfigured using classical autoconfiguration mechanisms as presented in section 6.

7.3 Mobility in MANETs

One main difference of MANETs in comparison to classical wired networks is the mobility of the nodes and thus a dynamic network topology (refer to section 1). Routes between MANET nodes may change very frequently and may also break down. When using DHCP for prefix delegation, it is important that all clients reach the DHCP server which is not guaranteed in a MANET. For example, the node running a DHCP server can be cut off the MANET, and thus no new node would ever get a prefix, and as soon as the leases of prefixes expire, participating nodes in the MANET will lose their prefixes. In addition, neither DHCP servers nor NDP can handle merging and partitioning of MANETs as described in section 1.

7.4 Relationship between MANET routers

As seen in the example in section 6.3, in the Internet, routers acquiring a prefix from a delegating router are compulsorily in a subordinate topology with respect to that router. This is not always true in a MANET. While MANET routers may allocate a subprefix from a superordinate router, they are not forced to. This case is depicted in figure 10, where two routers in a MANET with a prefix $p:1::$ and $p:2::$, respectively, are on the same hierarchical level rather than in a subordinated topology. This is the reason why delegating prefixes using DHCP-PD will not directly work because $p:2::$ is not subordinate prefix of $p:1::$ and can thus not be delegated by MR_1 .

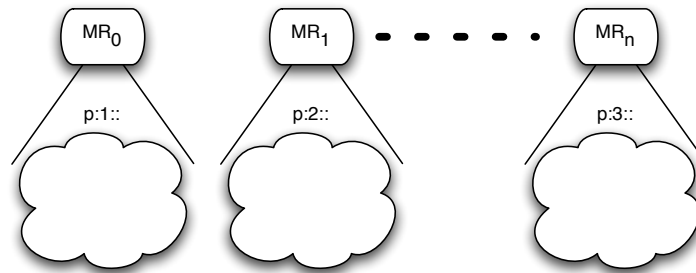


Figure 10: **MANET routers:** which are on the same hierarchical level thus having non-overlapping prefixes.

7.5 Summary

MANETs have some different architectural properties in comparison to classical IP networks. Verbatim application of classical autoconfiguration mechanisms is not possible due to different underlying assumptions. First, MANET nodes have a different link-model as classical IP links thus inhibiting verbatim application of these mechanisms in MANETs. In addition, NDP cannot be applied in MANETs as it only applies to configure hosts, not routers. Mobility in MANETs avoids the usage of DHCP, as the network topology may frequently change and the connection to the DHCP server might not be constantly available. At last, MANET routers are in contrast to classical routers not compulsorily in a subordinate topology with respect to some other router. Hence, they can be in the same hierarchical level thus requiring some other mechanism as DHCP-PD.

8 Recapitulation

This memorandum has described goals for and constraints on IP address and prefix auto-configuration mechanisms for mobile ad hoc networks, which are necessary for a seamless integration of MANETs with the classical IP architecture.

In doing so, it has been illustrated that the unfortunately commonly used terms "*connected MANET*", "*disconnected MANET*", and "*stand-alone MANET*" are unsuitable – in particular since these terms do (i) not allow classification of MANETs with different properties and (ii) do not allow this since the terms are not descriptive of relevant properties in MANETs. Indeed, it has been argued that any and all MANETs can be classified as being "*connected MANETs*", rendering the terminology and the classification that the terminology aims at providing useless.

Considering the Internet addressing architecture, a taxonomy has been proposed which, contrary to the terms indicated above, allows to classify MANETs into two categories, describing properties relevant for each category with respect to autoconfiguration of MANETs. This taxonomy aids in simplifying both the description of the goals of autoconfiguration in MANETs, as well as the architectural considerations that apply for how IP addresses and prefixes are assigned and delegated to interfaces and routers.

While the goals for autoconfiguration in MANETs are not that different from those for autoconfiguration in other network types, the characteristics of the MANET interfaces, the requirement that each MANET router shall acquire a prefix and the possible deployment of MANETs as Autonomous MANETs render verbatim application of existing autoconfiguration mechanisms impossible: the fundamental assumptions behind such existing mechanisms are distinctly different from those of MANETs – core examples hereof have been described in this memorandum.

Acknowledgements

The architectural view presented in this memorandum is the result of exhaustive discussions within the IETF, in particular with the participants in the MANET and AUTOCONF working groups. Special thanks to Jari Arkko (Ericsson, Finland) and Dave Thaler (Microsoft, USA) for sharing their insights and for reviewing various pieces of this memorandum, and Francois Morain (Ecole Polytechnique, France), for providing clarifying questions to parts of this memorandum.

References

- [1] F. Baker, "RFC1812: Requirements for IP Version 4 Routers", Standards Track, <http://www.ietf.org/rfc/rfc1812.txt>
- [2] T. Narten, S. Thomson, "RFC2462: IPv6 Stateless Address Autoconfiguration", Standards Track, <http://www.ietf.org/rfc/rfc2462.txt>
- [3] V. Devarapalli, R. Wakikawa, A. Petrescu, P. Thubert, "RFC3963: Network Mobility (NEMO) Basic Support Protocol", Standards Track, <http://www.ietf.org/rfc/rfc3963.txt>
- [4] D. Thaler, "Multilink Subnet Issues", Internet-Draft (Work in Progress), <http://www.ietf.org/internet-drafts/draft-iab-multilink-subnet-issues-02.txt>
- [5] T. Clausen, E. Baccelli, R. Wakikawa, "NEMO Route Optimisation Problem Statement", Internet-Draft (Work in Progress), <http://www.ietf.org/internet-drafts/draft-clausen-nemo-ro-problem-statement-01.txt>
- [6] T. Clausen, E. Baccelli, R. Wakikawa, "Route Optimisation in Nested Mobile Networks (NEMO) using OLSR", Proceedings of the IASTED International Conference on Networks and Communications Systems (NCS), April, 2005
- [7] IETF MANET Working Group Charter, <http://www.ietf.org/html.charters/manet-charter.html>
- [8] IETF AUTOCONF Working Group Charter, <http://www.ietf.org/html.charters/autoconf-charter.html>
- [9] I. Chakeres, J. Macker, T. Clausen, "Mobile Ad hoc Network Architecture", Internet-Draft (Work in Progress), <http://www.ietf.org/internet-drafts/draft-ietf-autoconf-manetarch-07.txt>
- [10] T. Clausen, "A MANET Architectural Model", INRIA Research Report RR-6145, January, 2007
- [11] J. F. Kurose, K. W. Ross, "Computer Networking: a Top-Down Approach Featuring the Internet Package", 1st edition, Addison-Wesley Longman Publishing Co., Inc., 2000
- [12] T. Boot, "Analysis of MANET and NEMO", Internet-Draft (Work in Progress), <http://www.ietf.org/internet-drafts/draft-boot-manet-nemo-analysis-01.txt>
- [13] K. Mase, K. Akima, "Prefix Distribution Framework for Connected MANETs", Internet-Draft (Work in Progress), <http://www.ietf.org/internet-drafts/draft-mase-autoconf-prefix-framework-00.txt>

- [14] S. Ruffino, P. Stupar, T. Clausen, S. Singh, "Connectivity Scenarios for MANET", Internet-Draft (Work in Progress), <http://www.ietf.org/internet-drafts/draft-ruffino-autoconf-conn-scenarios-00.txt>
- [15] T. Clausen, P. Jacquet, "RFC3626: Optimized Link State Routing Protocol (OLSR)", Experimental, <http://www.ietf.org/rfc/rfc3626.txt>
- [16] N. Shenoy, Y. Pan, V. G. Reddy, "Bandwidth Reservation and QoS in Internet MANETs", Proceedings of the 14th ICCCN International Conference on Computer Communications and Networks, pp. 353-358, October, 2005
- [17] R. Droms, "RFC2131: Dynamic Host Configuration Protocol", Standards Track, <http://www.ietf.org/rfc/rfc2131.txt>
- [18] R. Droms, J. Bound, B. Volz, T. Lemon, C. Perkins, M. Carney, "RFC3315: Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", Standards Track, <http://www.ietf.org/rfc/rfc3315.txt>
- [19] O. Troan, R. Droms, "RFC3633: IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", Standards Track, <http://www.ietf.org/rfc/rfc3633.txt>
- [20] T. Narten, E. Nordmark, W. Simpson, H. Soliman, "RFC4861: Neighbor Discovery for IP version 6 (IPv6)", Standards Track, <http://www.ietf.org/rfc/rfc4861.txt>
- [21] S. Thomson, T. Narten, T. Jinmei, "RFC4862: IPv6 Stateless Address Autoconfiguration", Standards Track, <http://www.ietf.org/rfc/rfc4862.txt>

Contents

1	Introduction	3
1.1	Memorandum Outline	3
2	What is an IP Address?	5
2.1	The Basics	5
2.2	Networks	5
2.3	Prefixes – the Internet Term for Intervals	7
2.4	Delegation and Assignment	8
2.5	Summary	9
3	The MANET Architecture	10
3.1	The Classic IP Link Model	10
3.2	The Missing Link	11
3.3	The Morphology of a MANET Router	12
3.4	'M' for Mobility – 'A' for Ad hoc	12
3.5	Summary	13
4	Connected vs. Disconnected MANETs ???	14
4.1	On why these terms are wrong	14
4.1.1	Graph Theory?	15
4.1.2	Outside Networks?	15
4.1.3	The Internet?	16
4.2	On why these terms are Unfortunate	16
4.3	A Better Taxonomy	17
4.3.1	Autonomous MANETs	17
4.3.2	Subordinate MANETs	18
4.4	Summary	19
5	AUTOCONF – The Goals	20
5.1	Architectural Considerations	20
5.2	Acquiring a Prefix in an Autonomous MANET	21
5.3	Acquiring a Prefix in a Subordinate MANET	21
5.4	Consequences	21
5.5	Summary	22
6	Automatic Configuration in the Internet	23
6.1	Dynamic Host Configuration Protocol (DHCP)	23
6.2	Stateless Autoconfiguration (SLAAC) / Neighbour Discovery Protocol (NDP)	23
6.3	IP architecture assumptions	24
6.4	Summary	24

7	Internet vs MANETs	25
7.1	Different link-model of MANETs	25
7.2	Routers vs. Hosts	25
7.3	Mobility in MANETs	25
7.4	Relationship between MANET routers	26
7.5	Summary	26
8	Recapitulation	27



Unité de recherche INRIA Futurs
Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399