Capítulo 3

Optimización de las comunicaciones entre redes móviles vehiculares

En el capítulo anterior se han presentando algunos escenarios que podrían beneficiarse de utilizar un enfoque basado en una solución de movilidad de redes. Claramente, parece que la provisión de acceso a Internet desde plataformas móviles (como trenes, aviones o autobuses) parece el escenario más relevante. Además, el escenario de comunicaciones vehiculares está recibiendo gran cantidad de atención por parte de la investigación académica e industrial.

El escenario particular de las comunicaciones vehiculares es cada vez más popular debido a que existe un gran número de aplicaciones potenciales que podrían beneficiarse de disponer de la capacidad de comunicarse a través de Internet. Principalmente, existen 2 problemas a tratar: el acceso a Internet desde coches (el denominado *car-to-Internet scenario*) y las comunicaciones entre vehículos (*car-to-car scenario*). Dada la naturaleza de las comunicaciones vehiculares y su relevancia, resulta necesario estudiar la aplicabilidad de una aproximación basada en movilidad de redes.

Este capítulo introduce primero el escenario vehicular, presentando los retos particulares derivados del mismo y analizando los diferentes enfoques que están siendo propuestos para soportar dicho escenario.

3.1. Introducción

En la sociedad moderna actual, mucha gente pasa una gran cantidad de tiempo en sus coches. En el pasado, las posibilidades de comunicación pasaban mayoritariamente por las redes celulares. Posibilitar las comunicaciones de banda ancha en coches [KBS+01] es una contribución muy importante en la tendencia global hacia las comunicaciones ubicuas. Los coches deben proporcionar acceso a Internet y deben ser capaces de establecer comunicaciones entre ellos, soportando nuevos servicios y aplicaciones.

Hay una gran cantidad de aplicaciones y servicios potenciales que son de gran interés para los usuarios de automóviles. En la Figura 3.1, se muestran algunos ejemplos representativos, clasificados en cinco categorías diferentes, pero con cierto solapamiento entre ellas:

- **Servicios de comunicaciones personales.** Las aplicaciones clásicas de telecomuni-

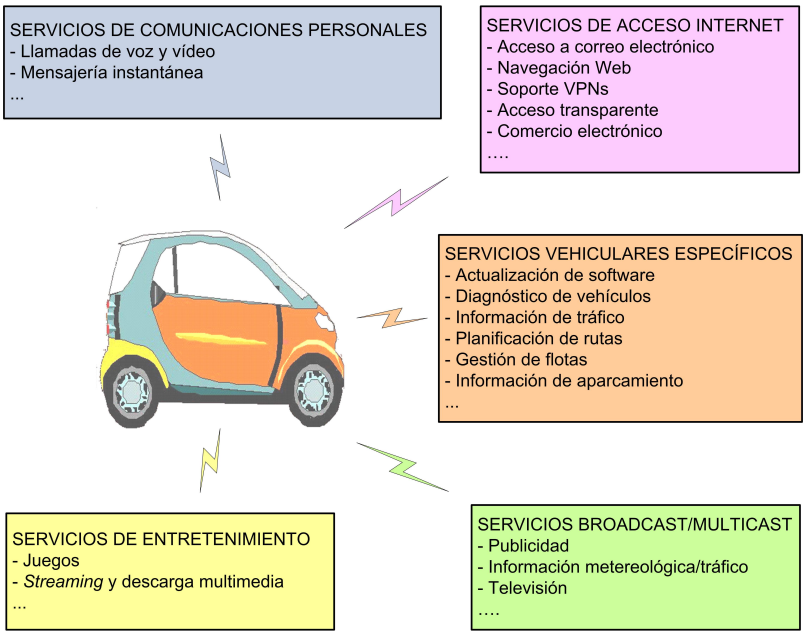

Figura 3.1: Algunos ejemplos de aplicaciones y servicios en un escenario vehicular.

caciones, tales como las comunicaciones de voz, tienen que ser integradas para su uso en los coches. De hecho, algunas de ellas están disponibles en los coches actuales (p.e., comunicaciones manos-libres utilizando un sistema celular integrado). Sin embargo, se espera que en el futuro aplicaciones más complejas estén disponibles en los coches, aprovechando las mayores capacidades que se espera que tengan los vehículos – comparada con las de los terminales de comunicaciones portátiles actuales.

- **Servicios de acceso a Internet.** Los vehículos, especialmente los servicios de transporte público, como trenes y autobuses, deben facilitar el uso a bordo de las aplicaciones típicas de trabajo (p.e., correo electrónico, software de VPN, etc.), mediante la provisión de acceso transparente a Internet, ya sea usando dispositivos embebidos en el vehículo o los terminales de los propios pasajeros.

- **Servicios vehiculares específicos.** Existen algunas aplicaciones que son específicas del escenario vehicular, como por ejemplo la descarga de información relativa al tráfico, la monitorización y diagnóstico de vehículos, y el control y la actualización del software instalado en los vehículos. En general, la seguridad es un aspecto crítico en este tipo de aplicaciones (p.e., en la diagnosis o la actualización de software).

- **Servicios de entretenimiento.** Los juegos multi-jugador y el *streaming* multimedia son aplicaciones ampliamente extendidas hoy en día, que muy probablemente serán de gran importancia en escenarios vehiculares (p.e., niños en los asientos traseros del coche, o personas yendo a su lugar de trabajo, jugando mientras que se desplazan). Además, estos servicios pueden beneficiarse de información de localización.

- **Servicios broadcast/multicast.** El envío de contenidos a grupos de receptores es un

servicio de interés en el entorno vehicular. Esta clase de servicio será probablemente proporcionado empleando tecnologías de acceso específicas, como DVB, por lo que tendrán que tenerse en cuanta además consideraciones adicionales.

Por todo lo anterior, parece que los coches dejarán de ser sistemas aislados dentro de poco [KBS+01], surgirán nuevos servicios y aplicaciones cuando los coches tengan la capacidad de conectarse a Internet y de comunicarse entre ellos [Ern06]. Estos nuevos escenarios supondrán nuevos retos que tendrán que ser resueltos, principalmente relacionados con la gestión de la movilidad, pero también con la provisión de calidad de servicio y la seguridad. Algunos de estos problemas están siendo estudiados por proyectos e iniciativas de investigación conjuntos, como los siguientes:

- El proyecto europeo DRiVE[1] (1999) y su continuación OverDriVE[2] (2001), que se centraron en facilitar la entrega de servicios multimedia a vehículos y el desarrollo de un router vehicular que proporcionara acceso, mediante múltiples tecnologías de radio, a una red intra-vehicular (intra-vehicular network, IVAN) móvil [LJP03], [LN03], [WS03].

- El proyecto InternetCAR[3] (1996), investigó cómo podría facilitarse la conexión transparente de vehículos a Internet. En algunas fases del proyecto se llevaron a cabo experimentos reales (con un número de vehículos que alcanzaba hasta 1640). Algunos resultados de estos experimentos pueden encontrarse en [EMU03], [EU02], [USM03], [WYT+05] y [KLE05].

- El proyecto *Red sobre Ruedas*, "Network On Wheels" (NOW[4]) (2004) se centra en IPv6 y la tecnología IEEE 802.11 para desarrollar comunicaciones entre vehículos basadas en conceptos de redes ad-hoc. Esencialmente, este proyecto está explorando maneras de que vehículos en movimiento puedan establecer dinámicamente enlaces con otros coches, motos y camiones en la vecindad, para compartir información de tráfico.

- El proyecto FleetNet ("Internet en la carretera[5]") (2000) fue formado por un consorcio de seis compañías y tres universidades con objeto de promover el desarrollo de sistemas de comunicaciones entre vehículos.

- El proyecto Daidalos[6] (2002) es un Proyecto Integrado del Sexto Programa Marco de la Unión Europea, actualmente en su segunda fase. Uno de sus objetivos es la integración óptima de tecnologías de acceso heterogéneas para permitir a los operadores de red y proveedores de servicio ofrecer nuevos y más rentables servicios. Las redes móviles son una de las tecnologías de acceso consideradas por el proyecto. Hasta el momento, Daidalos ha trabajado en tres aspectos dentro de la movilidad de redes [BSC+05b], [BSC+05a]: el desarrollo de una implementación del protocolo de

[1] `http://www.ist-drive.org/`
[2] `http://www.ist-overdrive.org/`
[3] `http://www.sfc.wide.ad.jp/InternetCAR/`
[4] `http://www.network-on-wheels.de/`
[5] `http://www.et2.tu-harburg.de/fleetnet/index.html`
[6] `http://www.ist-daidalos.org/`

Soporte Básico de Movilidad de Redes [dlOBC05], la extensión del protocolo básico para soportar tráfico multicast [vHKBC06] y el diseño de una solución de optimización de rutas para redes móviles [BBC04].

Los proyectos anteriormente descritos son solamente algunos de los más relevantes. Existen muchos otros esfuerzos de investigación en esta línea, como el proyecto InternetITS[7] [MUM03], el consorcio Car2Car Communication[8] o el proyecto CarTALK 2000[9]. Dada la gran cantidad de esfuerzos de investigación relacionados con las comunicaciones vehiculares, queda patente que el escenario vehicular es un tema de investigación de actualidad. La mayoría de estos esfuerzos están dirigidos a proporcionar soluciones para los dos escenarios principales considerados en las comunicaciones vehiculares:

- **Comunicaciones *Car-to-Internet*.** Este es un escenario muy común ya que muchas de las aplicaciones que se espera se necesiten en un vehículo, implican comunicaciones entre un nodo dentro de un coche y otro extremo en Internet (p.e., navegación web, correo electrónico, etc.). Inicialmente sólo se empleaban las comunicaciones celulares en este tipo de escenarios [AVN00]. Recientemente, con el éxito de la tecnología inalámbrica IEEE 802.11, otras tecnologías están siendo consideradas. Se está investigando como solventar las limitaciones (p.e., costes, bajos anchos de banda, altos retardos, etc.) de las tecnologías celulares existentes hoy en día, mediante el uso de WLAN 802.11 ([LG04] presenta un estudio sobre la posibilidad o no de utilizar WLAN 802.11 para conectar trenes a Internet) y WiMAX.

- **Comunicaciones *car-to-car*.** Existen diversas aplicaciones vehiculares, como los juegos en red, la mensajería instantánea, la información de tráfico o los servicios de emergencia, que pueden implicar comunicaciones entre vehículos que se encuentran relativamente cercanos entre sí, y que incluso pueden moverse juntos (p.e., convoys militares). Además, hay algunas aplicaciones emergentes que son exclusivas del entorno vehicular. Por ejemplo, los servicios de información al conductor podrían informar de forma inteligente acerca de atascos, negocios y servicios que se encuentren en las cercanías del vehículo, u otro tipo de noticias. Estos servicios emergentes no están bien soportados en la actualidad. Numerosos retos tecnológicos han de ser solventados antes de que las comunicaciones inter-vehiculares puedan llegar a ser ampliamente desplegadas. Estos escenarios están siendo investigados mayoritariamente por la comunidad ad-hoc, debido a que los protocolos de encaminamiento ad-hoc resultan muy apropiados para este tipo de problema (es decir, topologías que cambian rápidamente a medida que los coches se desplazan, carencia de infraestructura previa, etc.).

La conectividad puede proporcionarse en ambos escenarios empleando una solución genérica de movilidad de redes (p.e., el protocolo de Soporte Básico de Movilidad de Redes [DWPT05]). Sin embargo, tal y como se describirá más tarde, el caso vehicular presenta

[7] http://www.internetits.org/
[8] http://www.car-2-car.org/
[9] http://www.cartalk2000.net/

algunas particularidades que hacen que el rendimiento cuando se emplean soluciones genéricas de movilidad de redes y de optimización de rutas sea muy bajo, requiriendo por lo tanto el estudio de nuevos tipos de soluciones.

3.2. Haciendo posibles las comunicaciones vehiculares

En esta sección se presenta una visión panorámica del estado del arte en comunicaciones vehiculares, clasificando las propuestas existentes en tres categorías diferentes.

3.2.1. Soluciones basadas principalmente en ad-hoc

Hay una gran cantidad de trabajo de investigación en el campo de las redes ad-hoc. Algunos de los mecanismos desarrollados por la comunidad ad-hoc parecen ser apropiados para el escenario vehicular, al menos como punto de partida. Por lo tanto, en los últimos años se han propuesto muchas soluciones para permitir las comunicaciones vehiculares basadas en el concepto de redes ad-hoc vehiculares (Vehicular Ad-hoc Networks, VANETs). Dentro de esta categoría, a la que llamamos *soluciones basadas principalmente en ad-hoc*, incluimos a todos aquellos mecanismos que afrontan el problema de las comunicaciones vehiculares utilizando soluciones ad-hoc exclusivamente, sin emplear IP móvil.

3.2.1.1. Redes ad-hoc vehiculares

Las redes ad-hoc surgen como alternativa a las redes basadas en infraestructura, debido a las demandas de movilidad y al reto que supone desplegar redes de acceso inalámbricas sin zonas *muertas* (sin cobertura). En particular, una red ad-hoc móvil (Mobile Ad-hoc Network, MANET [CM99]) es un grupo de dispositivos móviles inalámbricos que cooperan para formar una red IP. Esta red no necesita ningún tipo de infraestructura para trabajar, ya que los dispositivos de los usuarios de una red MANET son la propia red, por lo que un nodo no sólo se comunica directamente con los dispositivos que tiene dentro de su radio de alcance, sino también con otros utilizando rutas multi-salto a través de otros nodos de la MANET.

Una red ad-hoc vehicular (Vehicular Ad-hoc Network, VANET) es un tipo particular de red ad-hoc en la que los nodos se encuentran en vehículos [FTMT$^+$05]. Mediante la configuración de una VANET, los vehículos pueden comunicarse localmente sin necesitar ninguna infraestructura (ver Figura 3.2).

El escenario vehicular tiene características que lo diferencian de otros escenarios de comunicaciones en red. Por ejemplo, por un lado tiene similitudes con los escenarios clásicos ad-hoc, debido a que presenta una topología que cambia rápidamente a medida que los coches se mueven. Sin embargo, por otro lado, las limitaciones y optimizaciones son diferentes. Primero, la eficiencia energética no es tan importante en las comunicaciones inter-vehiculares como lo es en las redes ad-hoc tradicionales, debido a que los vehículos disponen de una potente fuente de energía recargable. Segundo, los vehículos en general se mueven en carreteras (y dentro de un mismo carril la mayoría del tiempo).

De cara a hacer que el escenario de la Figura 3.2 funcione correctamente, hay algunos aspectos que deben resolverse:

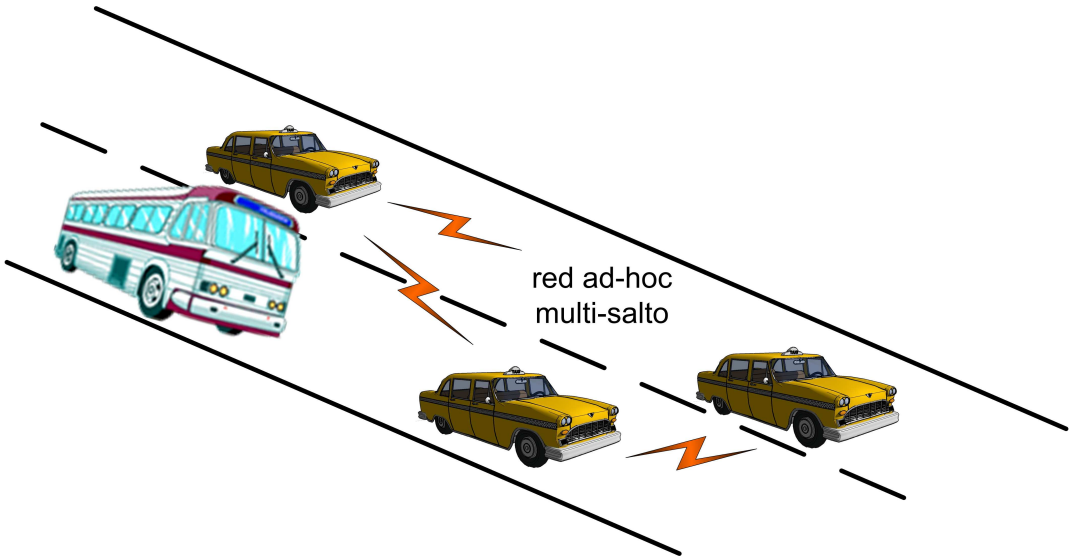

Figura 3.2: Red ad-hoc vehicular (VANET).

- *Encaminamiento.* En una red ad-hoc, no hay ninguna infraestructura de encaminamiento pre-establecida, por lo que los nodos tienen que colaborar en la configuración y mantenimiento de rutas multi-salto. Por lo tanto, se necesitan protocolos de encaminamiento específicos para escenarios ad-hoc.

- *Seguridad.* Debido a la falta de gestión que caracteriza a las redes ad-hoc, la seguridad es un aspecto crítico. Los protocolos dirigidos a trabajar en redes ad-hoc deben diseñarse prestando una especial atención a sus posibles debilidades de seguridad.

- *Autoconfiguración de direcciones IP.* Los protocolos existentes para la autoconfiguración de direcciones IP (en redes con infraestructura) no funcionan en las redes ad-hoc, por lo que tienen que definirse nuevos mecanismos que soporten la autoconfiguración IP de los nodos ad-hoc.

Si, además de las comunicaciones entre vehículos (car-to-car communications), se quiere proporcionar conectividad a Internet a los nodos de una VANET (car-to-Internet scenario), entonces debe resolverse también el siguiente aspecto:

- *Descubrimiento de una pasarela a Internet.* Se necesita un nodo especial, llamado *Internet Gateway* (pasarela a Internet), que conecta la red ad-hoc con la infraestructura. Permitir que los nodos ad-hoc descubran de forma eficiente las pasarelas a Internet supone ciertas dificultades, debido a la naturaleza de las redes MANET.

A continuación analizamos brevemente cada uno de los aspectos enumerados anteriormente.

3.2.1.2. Encaminamiento ad-hoc

Las redes ad-hoc han recibido una gran atención en los últimos años [CCL03], [AWW05], [FJL00]. Debido a su naturaleza inalámbrica, multi-salto y a su alta movilidad, los protocolos de encaminamiento tradicionales (utilizados en redes cableadas) no funcionan correctamente, y por lo tanto no pueden ser empleados, en redes MANET. Una gran cantidad de protocolos de encaminamiento han sido propuestos, la mayoría de ellos dentro del IETF. Algunos de ellos reciben el nombre de *reactivos*, porque se lanza el proceso de encontrar y establecer una ruta hacia un destino sólo cuando hay paquetes que tienen que ser enviados hacia dicho destino (como por ejemplo Ad-hoc On-Demand Distance Vector – AODV [PBRD03] – y Dynamic Source Routing – DSR [JMH04], [JMB01]).

Existen también protocolos conocidos como *proactivos*, porque los nodos proactivamente mantienen una entrada en su tabla de encaminamiento para todos los destinos alcanzables (como Optimized Link State Routing – OLSR [CJ03]), reduciendo de esta forma el tiempo que se necesita para establecer una ruta hacia un destino, aunque ello incrementa la complejidad del protocolo. Más información sobre protocolos de encaminamiento ad-hoc puede encontrarse en [AWD04], [RT99] y [CCL03].

El rendimiento de las redes ad-hoc depende en gran medida del protocolo de encaminamiento empleado y de la tecnología de radio empleada. La mayoría de los trabajos de investigación realizados hasta el momento en temáticas ad-hoc han sido realizados mediante simulación [KCC05], aunque también existen algunos trabajos experimentales, que estudian el rendimiento real de prototipos de redes ad-hoc [MBJ01]. Algunos de estos trabajos se centran en escenarios vehiculares [SBSC02], [SBS$^+$05], demostrando que es factible desplegar redes ad-hoc utilizando equipamiento IEEE 802.11b. Por otro lado, algunos autores afirman que es muy complicado conseguir que una red ad-hoc funcione con más de 10 nodos y 3 saltos intermedios [TLN03]. Se necesitan más trabajos de investigación que analicen el rendimiento de redes ad-hoc reales. Además, la disponibilidad de nuevas tecnologías DSRC (Dedicated Short Range Communication) [ZR03] impulsarán aún más las redes ad-hoc.

3.2.1.3. Seguridad

La seguridad es un aspecto crítico en las redes ad-hoc. Dada la naturaleza inalámbrica, el dinamismo de las redes MANET, y la falta de una infraestructura pre-establecida y de mecanismos de control, proporcionar a esta clase de redes un nivel de seguridad similar al de la Internet clásica (basada en infraestructura) es realmente complejo. Todos las funcionalidades enumeradas anteriormente (encaminamiento, autoconfiguración IP y descubrimiento de pasarelas de Internet) comparten este problema. Existen muchas publicaciones al respecto de la seguridad en ad-hoc, algunos de ellos analizando las amenazas, como [ZH99] y [SA99], y otros proponiendo soluciones a problemas específicos.

Aunque existen otros aspectos de seguridad que han sido tratados, tal y como el estimulo de la cooperación entre nodos, el problema del descarte de paquetes por parte de nodos maliciosos [SBR03], o la provisión de una autoridad de certificación fiable y segura en redes ad-hoc [HBC01], [CBH03], el problema del encaminamiento seguro es el que ha recibido más atención.

Algunos de los protocolos de encaminamiento propuestos actualmente, como AODV [PBRD03], DSDV [PB94] y DSR [JMH04], tienen vulnerabilidades de seguridad que per-

miten realizar ataques fácilmente. Debido a las importantes diferencias entre las redes IP basadas en infraestructura y las redes ad-hoc, es necesario desarrollar nuevos mecanismos de seguridad

Existen varios tipos de ataques de seguridad que pueden realizarse contra los protocolos de encaminamiento ad-hoc [SDL+02]. A continuación resumimos los más relevantes:

- *Ataques de modificación.* Un nodo malicioso puede causar la redirección de tráfico de datos o ataques de denegación de servicio (Denial-of-Service, DoS) introduciendo cambios en los paquetes de control de encaminamiento o reenviando mensajes de encaminamiento con valores falsos.

- *Ataques de suplantación.* Un nodo malicioso puede suplantar la dirección IP de un nodo legítimo y, por lo tanto *robarle* su identidad y realizar este tipo de ataque combinado con un ataque de modificación. El mayor problema de este tipo de ataques es que es difícil trazar quién es el nodo malicioso.

- *Ataques de fabricación.* Un nodo malicioso puede crear y enviar mensajes de encaminamiento falsos. Este tipo de ataque es difícil de detectar, ya que no es sencillo verificar que un mensaje de encaminamiento en particular es inválido, especialmente cuando indica que un vecinos no puede ser alcanzado.

Los autores de [SDL+02], [SLD+05] proporcionan los siguientes requisitos como aquellos que debe cumplir un protocolo de encaminamiento ad-hoc seguro:

1. La señalización de encaminamiento no puede ser suplantada.
2. No pueden inyectarse mensajes de encaminamiento fabricados en la red.
3. Los mensajes de encaminamiento no pueden ser alterados en tránsito (salvo acorde al funcionamiento normal del protocolo de encaminamiento).
4. No pueden formarse bucles en el encaminamiento como resultado de una acción maliciosa.
5. Las rutas más cortas no pueden ser sustituidas por otras como resultado de una acción maliciosa.

La comunidad investigadora ha tratado los anteriores problemas de seguridad en los protocolos de encaminamiento ad-hoc, intentando proponer mecanismos que cumplan algunos, si no todos, de los requisitos mencionados anteriormente. Un gran número de soluciones ha sido propuesto. A continuación describimos brevemente algunas soluciones representativas:

Ref. [HPJ05] propone una versión segura de DSR (llamada ARIADNE), utilizando claves simétricamente pre-distribuidas o criptografía simétrica pre-desplegada para la autenticación.

SEAD [HJP02] es un protocolo de encaminamiento proactivo seguro, basado en DSDV [PB94], que utiliza cadenas de funciones hash (*hash-chains*).

SAODV [ZA02] es una propuesta para proporcionar seguridad a AODV [PBRD03]. Se utilizan dos mecanismos para asegurar AODV: firmas digitales para autenticar la información no mutable de los mensajes y *hash chains* para asegurar la información mutable (es

decir, el número de saltos). Para la información no mutable, la autenticación se realiza extremo a extremo. Sin embargo, la misma técnica no puede aplicarse a la información que puede cambiar en tránsito, porque sería posible que un nodo intermedio suplantara la identidad de un nodo legítimo y modificara el campo que indica el número de saltos en un paquete de petición de ruta (*route request*). Para evitar esto, se utilizan cadenas de hash para proteger la información mutable.

En SRP [PH02], se asume que existe una asociación de seguridad para cada par origen-destino, para poder crear una ruta multi-salto. Este protocolo es vulnerable a algunos ataques, como la fabricación de mensajes de error en una ruta.

Una propuesta muy interesante es ARAN [SDL+02], [SLD+05]. Esta solución utiliza mecanismos de criptografía de clave pública para evitar todos los ataques enumerados con anterioridad. Sin embargo, presenta la desventaja de requerir certificados emitidos por una tercera parte. Este requisito puede afectar al despliegue de la solución, especialmente en entornos vehiculares.

En resumen, podemos concluir que cualquier mecanismo dirigido a trabajar en un escenario ad-hoc debe tener muy en cuenta aspectos de seguridad, si bien muchos protocolos actuales de MANET no lo hacen.

3.2.1.4. Autoconfiguración de direcciones IP

De cara a permitir que las redes MANET puedan soportar servicios IP, todos los nodos de la red deben configurar al menos una dirección IP. Sin embargo, no hay ningún mecanismo estándar que proporcione información de configuración IP a nodos de una red MANET, por lo que es necesario configurar los nodos a priori, lo que evita la formación espontánea de redes ad-hoc.

Los protocolos de configuración IP existentes [TN98] para redes tradicionales con infraestructura asumen la existencia de un único enlace con capacidad de transmisión multipunto para la señalización. Tal enlace no existe en las redes multi-salto sin infraestructura, por lo que es necesario diseñar nuevos mecanismos que permitan la autoconfiguración de direcciones IP en una red MANET [SKP+06], [RRGS05].

De cara a tratar del tema de la autoconfiguración IPv6 en redes MANET, se creó un nuevo grupo de trabajo dentro del IETF, denominado AUTOCONF. Este grupo ha identificado dos posibles escenarios principales [RSCS06] dónde es necesaria la autoconfiguración de direcciones IP para redes MANET:

- Red ad-hoc *aislada* (Stand-alone): una red ad-hoc que no está conectada a ninguna red externa, como por ejemplo las redes en conferencias, las redes en campos de batalla, las redes de vigilancia, etc. Lo más probable es que en estos casos no exista ningún tipo de entidad para la delegación de direcciones o prefijos preestablecida en la red. En este escenario, las direcciones IPv6 no tienen porqué ser globales.

- Redes ad-hoc *híbridas* (en el extremo de una red con infraestructura): una red aislada conectada a Internet. Los nodos de una red híbrida deben obtener direcciones IPv6 globales, para que puedan comunicarse con cualquier otro nodo de la Internet. Esto típicamente requiere descubrir el prefijo IPv6 disponible en la red MANET y configurar una dirección única (no utilizada) a partir de este prefijo [BC06].

Aunque el grupo de trabajo AUTOCONF está aún trabajando en la definición del protocolo, ya existen muchas soluciones en la actualidad. Una clasificación de las más relevantes puede encontrarse en [BC05].

3.2.1.5. Descubrimiento de una pasarela a Internet

Para proporcionar conectividad a una red MANET híbrida [RRGS05], además de un direccionamiento IPv6 global, se necesita de un tipo de nodo especial en la red ad-hoc. La pasarela a Internet (Internet Gateway, IGW) es un nodo que tiene conexión tanto a una red de acceso con infraestructura como a la red ad-hoc, y que proporciona conectividad a los nodos conectados a esta última. Un IGW puede ser móvil o fijo y es de vital importancia para proporcionar conectividad a los nodos que están del lado MANET. Debido a las características de las redes MANET, es deseable que se desplieguen múltiples IGWs (por ejemplo, para mitigar problemas relativos a congestión).

Actualmente, una propuesta muy común para proporcionar conectividad a Internet en los vehículos consiste en desplegar IGWs a los lados de las carreteras, de forma que los vehículos que pasen puedan usarlos para acceder a Internet. Uno de los retos que supone esta arquitectura es cómo descubrir eficientemente los IGWs disponibles [BWSF03], ya que uno de los componentes clave que afectan al rendimiento global es el algoritmo empleado par descubrir y seleccionar IWGs [RGS04].

Desplegar una infraestructura de red consistente en varios IGWs en las carreteras y confiar en el encaminamiento multi-salto en las redes ad-hoc vehiculares formadas no es suficiente. Podrían existir ”agujeros” en la conectividad, que podrían evitar que los vehículos se comunicaran (no sólo entre sí, sino también con Internet). Además, los IGWs podrían no pertenecer todos al mismo proveedor y por lo tanto, no sería posible que un vehículo pudiera mantener la misma dirección IPv6 al moverse de un IGW a otro. Aunque existen soluciones que mitigan el efecto de la conectividad intermitente, como la descrita en [OK04] para una red no ad-hoc (basada en pasarelas de aplicación y proxies), debería permitirse la posibilidad de conmutar a otra interfaz de red (p.e., celular, como GPRS o UMTS) manteniendo de forma transparente las sesiones existentes (es decir, soporte de movilidad transparente real).

Las soluciones basadas principalmente en ad-hoc presentan algunos inconvenientes. Por ejemplo, hay algunos aspectos de seguridad que no están resueltos aún y no proporcionan un soporte de movilidad global. Por lo tanto, este tipo de solución no es válido para satisfacer todos los requisitos del escenario vehicular.

3.2.2. Soluciones basadas en movilidad de terminal

Una forma diferente de soportar comunicaciones vehiculares consiste en considerar cada coche como un terminal individual y emplear técnicas que utilizan IP móvil par soportar la movilidad del terminal. A este tipo de soluciones es a las que denominamos *soluciones basadas en movilidad de terminal*.

Este tipo de soluciones están basadas en aprovechar los protocolos inalámbricos y de movilidad existentes, haciendo los cambios necesarios para incrementar el rendimiento en

entornos vehiculares. Un ejemplo muy sencillo consiste en utilizar soluciones basadas en redes celulares 2.5G/3G [AVN00].

Otro ejemplo es la arquitectura definida en el proyecto Drive-thru [OK04]. Está basada en proporcionar algunos servicios de Internet útiles en entornos con conectividad intermitente. Los coches obtienen esta conectividad intermitente conectándose a puntos de acceso WLAN desplegados en la carretera.

En algunos escenarios resulta interesante combinar un mecanismos de IP móvil con soluciones ad-hoc, de cara a soportar el movimiento de los vehículos entre redes ad-hoc y redes con infraestructura. Esto requiere permitir la *movilidad global* entre diferentes tipos de redes de acceso (ad-hoc o con infraestructura) para preservar de forma transparente la conectividad de los vehículos. La mayoría de las propuestas para la gestión global de la movilidad en redes ad-hoc están basadas en adaptar los mecanismos de IP móvil existentes para ser utilizados con protocolos de encaminamiento ad-hoc particulares.

Una de las propuestas más conocidas es MIPMANET [JAL+00], que básicamente propone una solución basada en IPv4 Móvil y AODV. Para combinar la naturaleza reactiva de AODV con la proactiva de IPv4 Móvil, los Agentes Foráneos (Foreign Agents, FAs) se anuncian periódicamente en la red ad-hoc. Los Agentes Foráneos son utilizados como pasarelas a Internet (IGWs), de cara a mantener información sobre en qué red ad-hoc se encuentra localizado un nodo y para encaminar los paquetes hacia el borde de dicha red ad-hoc. Se utiliza AODV para entregar los paquetes entre el FA y el nodo móvil. Se emplea un enfoque en capas y túneles para el tráfico saliente de la red ad-hoc, de forma tal que se separa la funcionalidad de IPv4 Móvil del protocolo de encaminamiento ad-hoc. Se propone un mecanismo similar en MEWLANA [EP02], pero adecuado para el protocolo de encaminamiento Destination-Sequenced Distance Vector (DSDV) [PB94]. En [RK03], se combinan técnicas tales como limitar la inundación de los anuncios de los Agentes Foráneos a un vecindario de n-saltos de radio – utilizando para ello un campo de tiempo de vida (TTL) en los mensajes de anuncio –, y espiar y cachear mensajes de agentes, para mejorar el rendimiento. De manera similar, un mecanismo que integra IPv4 Móvil y OLSR se propone en [BMA+04].

En relación al soporte de IPv6, [PMW+02] describe cómo proporcionar conectividad a Internet con soporte de IPv6 Móvil a redes ad-hoc. IPv6 Móvil utiliza el protocolo de *Neighbour Discovery* como parte de su mecanismo de detección de movimiento, con la adquisición de una dirección IP globalmente encaminable. Este mecanismo de detección de movimiento es modificado en las redes ad-hoc, en las que el IGW juega el papel de router local y los *Router Advertisements* son reemplazados por *Gateway Advertisements*. La dirección IPv6 configurada del prefijo de la red MANET que se incluye en los anuncios emitidos por el IGW es utilizada como la CoA del MN. Esta manera de realizar la detección de movimiento tiene el inconveniente de que requiere más tiempo que descubrir el movimiento entre dos puntos de conexión a la Internet fija, debido a que los Gateway Advertisements no son enviados con tanta periodicidad como los Router Advertisements (para evitar consumir recursos radio en exceso). Otros mecanismos propuestos para el soporte de movilidad global IPv6 en redes ad-hoc son [HSFN04] – que adopta una arquitectura jerárquica (basada en HMIPv6) para permitir que los nodos ad-hoc se registren en más de un AR/IGW simultáneamente – y [HLWC05] – que propone un protocolo que automáticamente organiza la red ad-hoc en una arquitectura en árbol para facilitar el direccionamiento y encaminamiento dentro de la red MANET.

También existen algunas soluciones propuestas que tratan específicamente con el escenario *car-to-car*. Un ejemplo es [BW05], que es similar a MIPMANET [JAL+00], en el sentido de que reutiliza el concepto de Agente Foráneo de IPv4 Móvil – colocado en el IGW – para gestionar la movilidad global de nodos ad-hoc. Se sigue empleando comunicación IPv4 entre el HA y el FA (utilizando túneles IPv6-en-IPv4), ya que la solución asume una Internet basada en IPv4 (los autores también proponen el uso de una arquitectura basada en proxies para soportar que vehículos que soporten sólo IPv6 puedan comunicarse con CNs IPv4 en la Internet). Tal y como se hace en [RK03], los FAs anuncian activamente su servicio, pero limitado a áreas locales, para evitar inundar toda la red vehicular.

Las soluciones basadas en movilidad de terminal tienen un inconveniente principal, consistente en que no tienen en cuenta que un vehículo muy probablemente contendrá más de un dispositivo que podría beneficiarse de disponer de acceso a Internet. Las soluciones basadas en movilidad de terminal requieren que todos los dispositivos gestionen su propia conectividad y movilidad (aunque se muevan todos juntos a la vez) y por lo tanto evita que nodos sin soporte de movilidad puedan ser desplegados en coches.

3.2.3. Soluciones basadas en movilidad de redes

Debido a que el escenario vehicular involucra a un grupo de dispositivos que se mueven juntos, tanto el caso de comunicaciones *car-to-Internet*, como el caso *car-to-car*, pueden ser abordados asumiendo que hay un router móvil desplegado en cada vehículo, encargado de gestionar la movilidad del grupo de dispositivos que se encuentra dentro del vehículo (a este tipo de soluciones, las denominamos *soluciones basadas en movilidad de redes*). Sin embargo, cabe destacar que después de estudiar la literatura relacionada, hemos encontrado muy pocas propuestas que consideran enfoques de movilidad de redes para escenarios vehiculares, puesto que la mayoría de los mecanismos consideran los coches como terminales individuales.

El escenario particular de comunicaciones *car-to-Internet* encaja muy bien en el paradigma de la Movilidad de Redes. Por lo tanto, la aplicación de un conjunto de soluciones genérico de movilidad de redes debe ser estudiada para el caso de comunicaciones ente vehículos e Internet. De hecho, es un buen ejemplo de escenario dónde se necesitan soluciones de optimización de rutas para movilidad de redes.

El escenario *car-to-car* también puede ser abordado utilizando una solución genérica NEMO. Sin embargo, ese tipo de solución no ofrece un rendimiento demasiado bueno en una comunicación entre vehículos, incluso aún cuando se utiliza alguna solución genérica de optimización de rutas para NEMO. Las razones para este subóptimo rendimiento son las siguientes:

- Las Redes Hogar de dos coches que se están comunicando pueden no ser la misma o estar localizadas lejos una de la otra. Esto hace muy necesaria la implantación de una solución de optimización de rutas, para evitar un incremento en el retardo debido a la utilización del protocolo de Soporte Básico de NEMO.

- Los coches muy probablemente obtendrán conectividad a Internet mediante una red celular 2.5/3G. Esta clase de redes típicamente presenta unos retardos muy elevados y

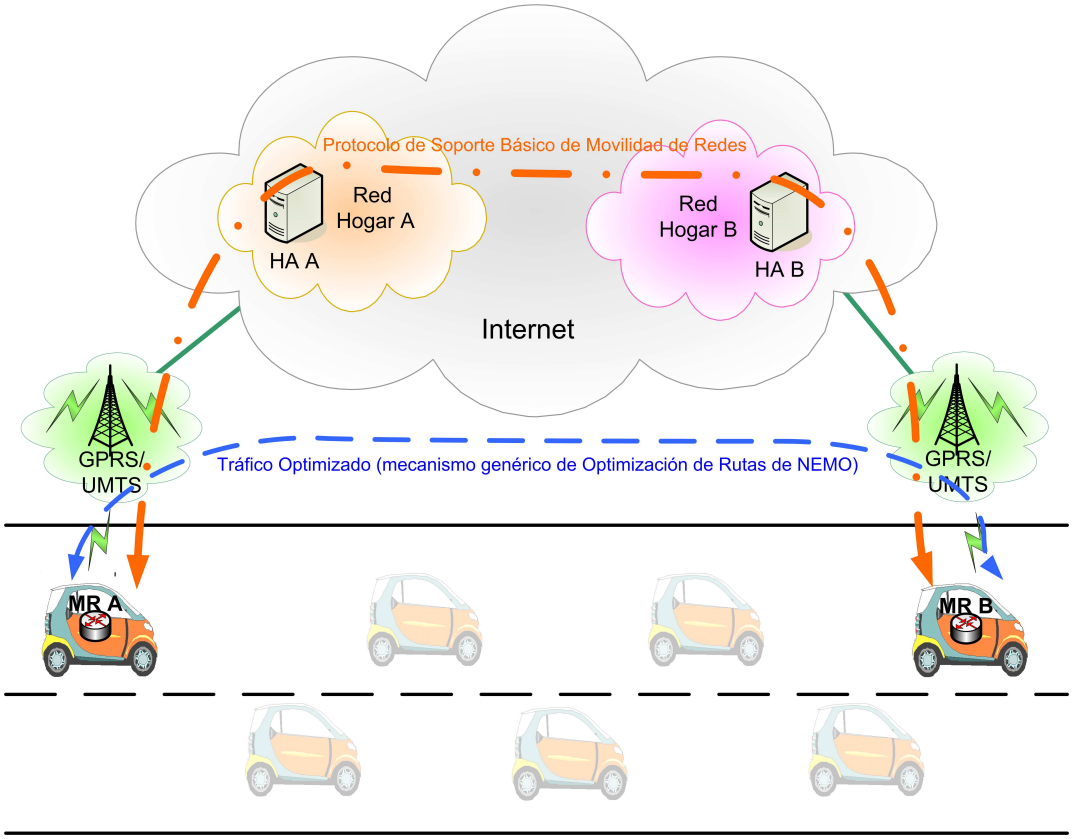

Figura 3.3: Funcionamiento de una solución genérica de Movilidad de Redes en el escenario *car-to-car*.

unos anchos de banda reducidos [VBM+05], [VBS+06]. Esto tiene un gran impacto en las comunicaciones inter-vehiculares cuando se usa una solución NEMO genérica. Un ejemplo de escenario *car-to-car* se muestra en la Figura 3.3. Si se utiliza el protocolo de Soporte Básico de Movilidad de Redes [DWPT05], el tráfico de datos fluye desde el Router Móvil de un coche (MR A) hacia su Red Hogar (Red Hogar A), dónde los paquetes son reenviados hacia la Red Hogar del otro coche (Red Hogar B) y finalmente entregados al Router Móvil B (MR B). Ésta es claramente una ruta subóptima. Si se utiliza una solución genérica de optimización de rutas, los paquetes no atraviesan las diferentes Redes Hogar de las redes móviles involucradas, pero siguen teniendo que pasar por la infraestructura para ser enviadas de un coche a otro. Esto puede significar un elevado retardo (las redes GPRS presentan retardos de unos 500 ms sólo en un sentido [VBM+05], [VBS+06], mientras que las redes UMTS tienen unos 150 ms [MdlOS+06], [OMV+06]) y un ancho de banda reducido. Por lo tanto, deben explorarse diferentes esquemas de optimización de rutas.

Aunque los coches pueden comunicarse entre sí a través de la infraestructura – empleando el protocolo de Soporte Básico de NEMO –, una conclusión que se puede obtener de la anterior discusión, es que los esquemas clásicos de optimización de rutas NEMO no ofrecen

un buen rendimiento en las comunicaciones inter-vehiculares. Sin embargo, hay una oportunidad de optimización que está siendo actualmente estudiada en el campo de las comunicaciones inter-vehiculares (inter-vehicular communications, IVC). Esta optimización está basada en la utilización de redes ad-hoc vehiculares (VANET), para explotar la conectividad entre coches vecinos, y el establecimiento una ruta multi-salto para soportar los servicios inter-vehiculares. La aplicación de este enfoque a una solución vehicular basada en NEMO es uno de los objetivos de esta Tesis Doctoral. En ella exploramos como diseñar un mecanismo basado en movilidad de redes que optimiza las comunicaciones inter-vehiculares, aprovechando que los MRs pueden establecer redes ad-hoc para comunicarse directamente (evitando pasar por la infraestructura).

De acuerdo a nuestro conocimiento, sólo hay una propuesta combinando los enfoques de NEMO y ad-hoc [WOKN05], [WMK+05], [WMK+04], [OWUM04]. La solución (definida en [WOKN05], [WMK+05], [WMK+04]) básicamente considera redes MANET que se mueven juntas (por ejemplo, dentro de un coche) e integra MANET y NEMO, colocando las funcionalidades del IGW y el MR en lo que denominan *Pasarela Móvil* (Mobile Gateway, MG). El protocolo de Soporte Básico de Movilidad de Redes [DWPT05] es responsable de proporcionar conectividad a Internet a la MANET móvil (por lo tanto, no es necesario que los nodos de dicha MANET soporten el protocolo IPv6 Móvil), mientras que un protocolo de encaminamiento ad-hoc adicional es ejecutado ente las Pasarelas Móviles, creando una red MANET superpuesta para comunicaciones entre diferentes redes móviles. Este esquema permite la comunicación directa entre nodos de MANETs móviles que pertenecen a la misma MANET superpuesta (usando la denominada *ruta directa*), mientras que el protocolo de Soporte Básico de Movilidad de Redes se utiliza para el resto de los casos (*ruta nemo*). Además, el mecanismo permite que un MG pueda tomar prestada la conectividad a Internet de un MG adyacente (*ruta detour*). Se propone utilizar un protocolo de encaminamiento ad-hoc proactivo para la MANET superpuesta, en particular, OLSR es considerado en [OWUM04].

La solución descrita anteriormente es una primera aproximación para combinar de manera óptima NEMO y ad-hoc para soportar comunicaciones vehiculares. Los autores han dejado como trabajo futuro el análisis de seguridad, por lo que en la arquitectura que ellos proponen, por ejemplo, nada previene a nodos malintencionados robar tráfico o realizar un ataque de inundación a la Red Hogar (Return-to-Home Flooding [NAA+05] attack). Esta falta de seguridad es un aspecto crítico, especialmente en entornos inter-vehiculares.

El diseño de un mecanismo basado en una solución de movilidad de redes combinado con soporte ad-hoc, de una forma *segura*, para permitir comunicaciones vehiculares óptimas es uno de los objetivos clave de esta Tesis Doctoral.

Part II

Route Optimisation for Mobile Networks in IPv6 Heterogeneous Environments

Optimización de Rutas para Redes Móviles en Entornos IPv6 Heterogéneos

Chapter 4

Goals and Design considerations

4.1. Introduction

This chapter enumerates the main goals of this PhD thesis and provides some design considerations that have been followed in the fulfilment of these goals.

4.2. Goals

The main goal of this PhD thesis is to design a set of solutions providing Route Optimisation support for Network Mobility, in such a way that the solutions are secure and easily deployable. A second objective of this PhD thesis is to develop a Route Optimisation solution for vehicular environments, by combining in a secure way Network Mobility and Ad-hoc concepts. Next, we elaborate more on the specific goals and requirements that the designed solutions solution should meet.

There are many possible scenarios where Network Mobility will play a key role, some of them related to the provision of Internet access from mobile platforms, such as cars or buses. We believe that there is a severe drawback in the existing standardised solution supporting Network Mobility [DWPT05], [dlOBC06], that is the sub-optimal packet routing that this solution forces in order to preserve mobility transparency. This sub-optimality can lead even to prevent communications from taking place, and therefore should be tackled if it is desired to deploy moving networks in practice.

A general overview of the existing proposals on Route Optimisation for NEMO[1] has been provided in Section 2.4, highlighting their problems and those issues that are still unsolved. Taking this into account, a brief summary of the requirements that we consider that a generic Route Optimisation (RO) for NEMO solution should fulfil (in order to be rapidly deployed in current scenarios) is provided next:

- The NEMO Route Optimisation solution should provide **support for legacy nodes**. In order to facilitate the practical deployment of the solution today, it should not require changes on the operation of any node but the Mobile Router and maybe the

[1]The interested reader may refer to [NZWT06], [PSS04b] and [LLKC05] for published surveys on existing Route Optimisation solutions for NEMO.

Home Agent of the mobile network. Therefore, Correspondent Nodes, Local Fixed Nodes and Mobile IPv6 hosts connected to a Mobile Network should not require any modification to be compatible with the NEMO RO mechanism. This does not necessarily imply that the mechanism should be able to optimise the traffic of any node in any scenario, but that at least the use of the RO mechanism on a network should not prevent any node from being able to communicate through the mobile network.

- The NEMO Route Optimisation solution should support the optimisation of communications of **Local Fixed Nodes** attached to a NEMO, that is, the mechanism should enable direct path routing between a CN and an LFN, by providing Angular Route Optimisation. Local Fixed Nodes will represent a large number of the nodes attached to a NEMO, for example in a moving vehicle, such a car, where it is likely to expect that many of the nodes that will require Internet connectivity, such as sensors, on-board computer, infotainment devices, etc, will have no mobility support at all. Therefore, providing them with NEMO RO is critical to exploit the benefits of a network mobility solution. To achieve that, mobility transparency (i.e. LFNs not being aware of the mobility of the NEMO) should be preserved.

- The NEMO Route Optimisation solution should support the optimisation of communications of **Visiting Mobile Nodes** attached to a NEMO. This means that the mechanisms should enable direct path communication between a VMN and its HA – when the VMN is operating in Bidirectional Tunnel (BT) mode – or between a VMN and a CN – when the VMN is operating in Route Optimisation (RO) mode.

- The NEMO Route Optimisation solution should support the optimisation of communications involving **nested configurations of Mobile Routers**, that is the number of tunnels traversed by packets belonging to a communication between a MNN attached to a sub-MR and a CN should be eliminated (or at least minimised), by providing Multi-angular Route Optimisation. At least, 3 levels of nesting should be supported, since nowadays, a higher number of nesting levels is not foreseen.

- The NEMO Route Optimisation solution should be **secure**, at least as secure as current NEMO Basic Support protocol and Mobile IPv6 are. Therefore, the solution should provide a similar level of security than the Route Optimisation solution of Mobile IPv6 [JPA04], [NAA+05].

- If the specific mechanism designed to provide Angular Route Optimisation differs from those used to cope with the different configurations of Multi-angular Route Optimisation, they should be **interoperable** in such a way that an optimal route results from the combined operation of both.

- The NEMO Route Optimisation solution should be **scalable**, so that in case of large mobile networks in terms of number of attached nodes, the designed RO solution should allow many sessions to be optimised, minimising the additional resources needed in any node of the network, compared to the resources required by plain NEMO Basic Support protocol.

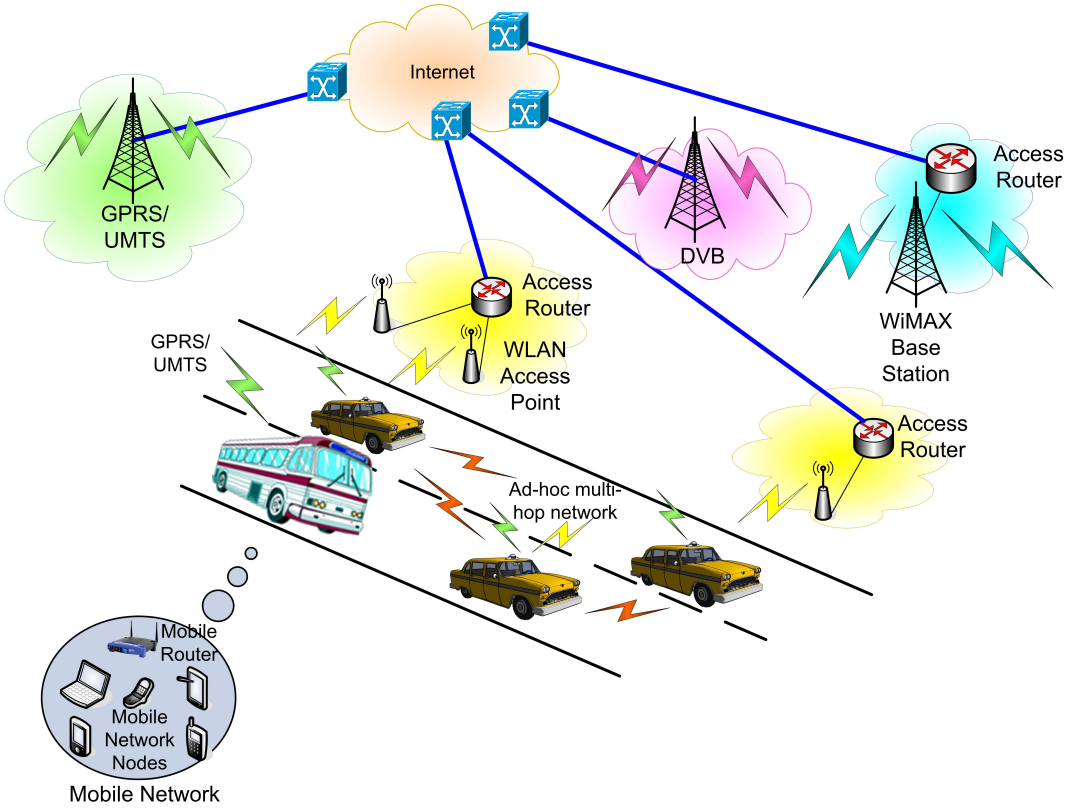

Figure 4.1: Vehicular communications scenario.

- The NEMO Route Optimisation solution should **minimise additional protocol complexity, processing load and signalling overhead** in order to facilitate the deployment of the solution in real scenarios.

- The NEMO Route Optimisation solution should **minimise the effect of increased handover delay** that it may have, when compared to the plain NEMO Basic Support protocol handover.

These are very **strong requirements** that, according to our analysis of the current state of the art (summarised in Chapter 2), are not fulfilled by any existing solution. This PhD thesis proposes a generic Route Optimisation for NEMO solution (described in detail in Chapter 5) that take these requirements as the basic design criteria.

Chapter 3 introduced the issue of vehicular communications. Vehicular scenarios (an example is shown in Figure 4.1) are becoming very important nowadays, since there exists a number of new services and applications that will likely be deployed in cars when they are provided with communication capabilities. Because of the importance of this particular scenario, it is needed to provide the tools and mechanisms that enable the optimisation of vehicular communications.

An objective of this PhD thesis is the study of an existing opportunity for **optimisation** of local communications in vehicular environments by using an **ad-hoc network** formed by

the vehicles involved in the communication and perhaps other vehicles in their surroundings. The design of a solution that optimally combines a Network Mobility approach with multi-hop ad-hoc communications in a secure way, and the evaluation of its performance in car-to-car communications, is the second main goal of this PhD work.

4.3. Design considerations

In this section we briefly summarise the considerations that we have adopted for the design of the mechanisms proposed in this PhD thesis. The rationale behind these considerations is provided in the PhD thesis when the specific mechanisms are described.

Since we aim at designing a generic Route Optimisation solution for NEMO and a mechanism that addresses the specific issue of vehicular communications, being both suitable for existing legacy nodes and not requiring changes on the operation of any node but the Mobile Router, there are several design aspects that should be addressed:

- *Layer of the protocol stack.* It seems clear that if it is required to avoid changes on the nodes, IP it is the only possible choice, since IP already has some mobility support, provided by the Mobile IPv6 [JPA04] protocol. It may be argued that the mobility can also be handled at other layers, but as it was briefly discussed in Section 2.1 (when talking about how Network Mobility could be provided), making it at the application [HLZ06] or transport [CAI06] layers requires the modification of all the applications or transport protocols to support mobility, which is redundant. On the other hand, making it at the link layer is not feasible if mobility across different technologies is required. Therefore, the option chosen is to **implement the designed mechanisms at the IP layer, using/modifying the Mobile IPv6 protocol.**

- *Entities involved on the optimisation.* Since one of our requirements is not to change any node but the MR (and maybe also the HA), it is clear that neither VMNs, not LFNs nor CNs can be changed. However, there would be the possibility of involving some nodes on the routing infrastructure in the optimisation, such as proposed in [WKUM03] and [WW04]. Nevertheless, as the deployment and scalability of the solution are also very important design considerations, the requirement of involving external nodes to perform the optimisation is not feasible. Therefore, the option chosen is to **put all the specific new functionalities required to perform the Route Optimisation on the Mobile Router only.** This does not mean that the designed mechanisms cannot make use of available mobility capabilities that certain Mobile Network Nodes (e.g., VMNs) may have.

Capítulo 4

Objetivos y Consideraciones de Diseño

4.1. Introducción

Este capítulo enumera los principales objetivos de esta Tesis Doctoral y proporciona algunas consideraciones de diseño que han sido seguidas para el cumplimiento de los objetivos perseguidos.

4.2. Objetivos

El objetivo principal de esta Tesis Doctoral es el diseño de un conjunto de soluciones que proporcionen soporte de optimización de rutas para redes móviles, de tal forma que las soluciones propuestas sean seguras y fácilmente desplegables. Un segundo objetivo de esta Tesis es el desarrollo de una solución de optimización de rutas para entornos vehiculares, mediante la combinación de forma segura de los conceptos de movilidad de redes y ad-hoc. A continuación se desarrollan un poco más las metas específicas y los requisitos que las soluciones diseñadas deben cumplir.

Existen muchos posibles escenarios donde la movilidad de redes jugará un papel clave, algunos de ellos relacionados con la provisión de acceso a Internet en plataformas móviles, como coches o autobuses. Creemos que hay un severo problema en la solución estandarizada existente para el soporte de la movilidad de redes [DWPT05], [dlOBC06], consistente en el encaminamiento subóptimo de paquetes que fuerza dicha solución de cara a preservar la transparencia de la movilidad. Este efecto subóptimo puede llegar incluso a impedir que las comunicaciones lleguen a efectuarse, y por lo tanto debe ser resuelto si se desea que las redes móviles puedan llegar a desplegarse en la práctica.

Una panorámica general de las propuestas existentes en relación a la optimización de rutas en redes móviles[1], ha sido proporcionada en la Sección 2.4, resaltando sus problemas y aquellos puntos que están todavía sin resolver. Teniendo esto en cuenta, a continuación se incluye un breve resumen de los requisitos que consideramos que una solución genérica

[1]El lector interesado puede acudir a [NZWT06], [PSS04b] y [LLKC05], para encontrar publicaciones que clasifican soluciones existentes de optimización de rutas para redes móviles.

de optimización de rutas para redes móviles debe cumplir (de cara a facilitar un rápido despliegue de la solución en los escenarios existentes):

- La solución de optimización de rutas para redes móviles debe proporcionar **soporte para nodos legados**. De cara a facilitar un rápido despliegue de la solución, ésta no debe requerir cambios en el funcionamiento de ningún nodo salvo el Router Móvil y quizás el Agente Local de la red móvil. Por lo tanto, ni los Nodos Corresponsales, ni los Nodos Locales Fijos ni los Nodos Móviles (que implementen IPv6 Móvil) conectados a una red móvil deben precisar cambios para ser compatibles con el mecanismo de optimización de rutas. Esto no implica necesariamente que el mecanismo deba ser capaz de optimizar el tráfico de todos los nodos en cualquier escenario, pero sí que al menos, la utilización de un mecanismo de optimización de rutas en una red no impida que ningún nodo sea capaz de comunicarse a través de la red móvil.

- La solución de optimización de rutas para redes móviles debe soportar la optimización de comunicaciones de **Nodos Locales Fijos** conectados a la red móvil, es decir, el mecanismo debe habilitar la comunicación directa entre un CN y un LFN, proporcionando optimización de rutas angular. Los Nodos Locales Fijos representarán un gran porcentaje de los nodos conectados a una red móvil, por ejemplo en un vehículo, como un coche, dónde es probable esperar que muchos de los nodos que requieran conectividad a Internet, como sensores, ordenadores de abordo, dispositivos de entretenimiento, etc., no tendrán soporte de movilidad alguno. Por lo tanto, es vital proporcionar a esta clase de nodos soporte de optimización de rutas, de cara a explotar los beneficios que brinda una solución de movilidad de redes. Para conseguir esto, la transparencia de la movilidad (los LFNs no deben ser conscientes de la movilidad de la red a la que están conectados) debe preservarse.

- La solución de optimización de rutas para redes móviles debe soportar la optimización de comunicaciones de **Nodos Móviles Visitantes** conectados a una red móvil. Esto significa que los mecanismos deben habilitar la comunicación directa entre un VMN y su HA – cuando el VMN está operando en modo de Túnel Bidireccional – o entre un VMN y un CN – cuando el VMN está operando en modo de Optimización de Rutas.

- La solución de optimización de rutas para redes móviles debe soportar la optimización de comunicaciones que involucren **configuraciones anidadas de routers móviles**, es decir, debe eliminarse el número de túneles atravesados por un paquete perteneciente a una comunicación entre un MNN conectado a un sub-MR y un CN (o al menos minimizar dicho número), proporcionando optimización de rutas multi-angular. Al menos deben soportarse 3 niveles de anidamiento, ya que actualmente no se prevé un número mayor de niveles de anidamiento.

- La solución de optimización de rutas para redes móviles debe ser **segura**, al menos tan segura como el protocolo de Soporte Básico de Movilidad de Redes e IPv6 Móvil. Por lo tanto, la solución debe proporcionar un nivel de seguridad similar al que presenta la solución de optimización de rutas de IPv6 Móvil [JPA04], [NAA+05].

- Si el mecanismo específico diseñado para proporcionar optimización de rutas angular es diferente de aquel utilizado para soportar las diferentes configuraciones de optimi-

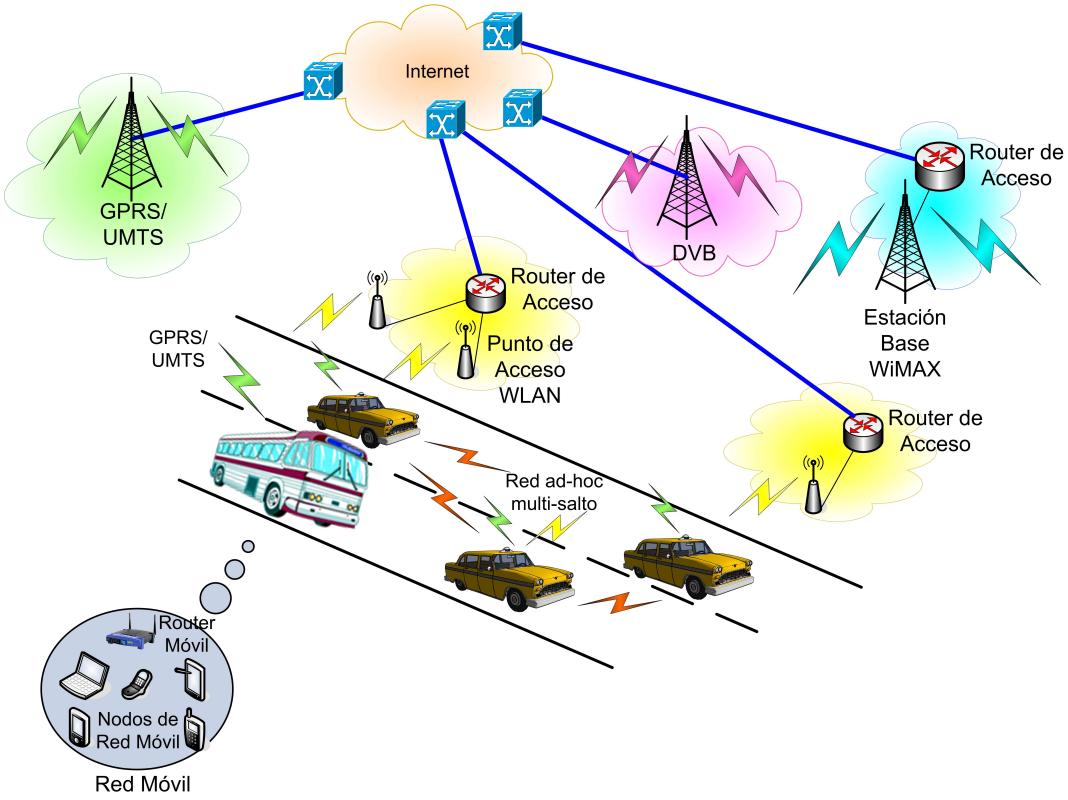

Figura 4.1: Escenario de comunicaciones vehiculares.

zación de rutas multi-angular, los dos mecanismos deben ser **interoperables**, de una forma tal que resulte una ruta óptima del funcionamiento combinado de ambos.

- La solución de optimización de rutas para redes móviles debe ser **escalable**, de forma tal que la solución diseñada permita optimizar un gran número de comunicaciones en caso de redes móviles muy grandes, minimizando los recursos adicionales requeridos en cualquier nodo de la red, comparado con los recursos que requiere el protocolo de Soporte Básico de Movilidad de Redes.

- La solución de optimización de rutas para redes móviles debe **minimizar la complejidad adicional del protocolo, la carga de procesamiento y la sobrecarga de señalización**, de cara a facilitar el despliegue de la solución en escenarios reales.

- La solución de optimización de rutas para redes móviles debe **minimizar el efecto de incremento en el retardo del traspaso**, que pudiera tener, comparado con el retardo de un traspaso utilizando el protocolo de Soporte Básico de Movilidad de Redes.

Los anteriores, son **requisitos muy fuertes** que, de acuerdo a nuestro análisis del estado del arte actual (resumido en el Capítulo 2), ninguna solución existente cumple. Esta Tesis Doctoral propone una solución genérica de optimización de rutas para redes móviles (descrita en detalle en el Capítulo 5) que adopta los requisitos anteriores como criterio básico.

El Capítulo 5 introdujo la problemática de las comunicaciones entre vehículos. Los escenarios vehiculares (un ejemplo se muestra en la Figura 4.1) están adquiriendo gran importancia actualmente, dado que existe un número de nuevos servicios y aplicaciones que muy probablemente serán desplegados en los coches cuando éstos dispongan de capacidades de comunicación en red. Debido a la importancia de este escenario particular, es necesario proporcionar las herramientas y mecanismos que faciliten la optimización de las comunicaciones vehiculares.

Un objetivo de la presente Tesis Doctoral es el estudio de la oportunidad de **optimización** que existe en comunicaciones locales entre vehículos, por medio de una **red ad-hoc** formada por los vehículos involucrados en la comunicación y quizás otros vehículos cercanos. El diseño de una solución que combine de forma segura el enfoque de movilidad de redes con el de las redes ad-hoc multi-salto, y la evaluación de su rendimiento en comunicaciones inter-vehiculares, es el segundo objetivo principal de esta Tesis.

4.3. Consideraciones de Diseño

En esta sección resumimos brevemente las consideraciones que hemos adoptado en el diseño de los mecanismos propuestos en esta Tesis Doctoral. El razonamiento subyacente se proporciona en la Tesis cuando los mecanismos específicos son descritos.

Dado que queremos diseñar una solución genérica de optimización de rutas para redes móviles y un mecanismo dirigido a la problemática específica de las comunicaciones vehiculares, siendo ambos aptos para su utilización con nodos legados y sin que se precisen cambios en el funcionamiento de ningún nodo, salvo el router móvil, hay algunos aspectos de diseño que deben ser considerados:

- *Capa de la pila de protocolos.* Parece claro que si se quiere evitar cambiar el funcionamiento de los nodos, IP es la única posibilidad, dado que IP ya cuenta con cierto soporte de movilidad, proporcionado por el protocolo IPv6 móvil [JPA04]. Podría argumentarse que la movilidad se puede gestionar en otras capas de la pila de protocolos, pero como ya fue brevemente discutido en la Sección 2.1 (cuando se hablaba sobre cómo se podía proporcionar soporte de movilidad de redes); hacerlo en las capas de aplicación [HLZ06] o transporte [CAI06] requeriría la modificación de todas las aplicaciones o protocolos de transporte para soportar la movilidad, lo cual es redundante. Por otro lado, no es posible hacerlo en el nivel de enlace si se quiere soportar movilidad entre diferentes tecnologías. Por lo tanto, la opción escogida es **implementar los mecanismos diseñados en la capa IP, usando/modificando el protocolo IPv6 Móvil.**

- *Entidades involucradas en la optimización.* Dado que uno de requisitos es no cambiar ningún nodo salvo el MR (y quizás también el HA), parece claro que ni los VMNs, ni los LFNs ni los CNs pueden ser cambiados. No obstante, existiría la posibilidad de involucrar a algunos nodos de la infraestructura de encaminamiento en la optimización, como se propone en [WKUM03] y [WW04]. Sin embargo, como la desplegabilidad y escalabilidad de la solución son también consideraciones de diseño muy importantes, no parece adecuado involucrar a nodos externos en la optimización. Por lo tanto,

la opción elegida es **poner todas las funcionalidades requeridas para realizar la optimización de rutas en el router móvil exclusivamente.** Esto no significa que los mecanismos desarrollados no puedan hacer uso de las capacidades de movilidad que determinados Nodos de Red Móvil (p.e., MNNs) puedan tener.

Chapter 5

Generic Route Optimisation solution for Network Mobility

5.1. Introduction

The Network Mobility (NEMO) Basic Support protocol [DWPT05] enables complete networks to roam among different access networks, without disrupting network nodes' ongoing sessions and without requiring any specific mobility capability in the hosts. Nevertheless, it has some important limitations in terms of performance (see section 2.3), due to the increased path length and the packet overhead that this solution introduces. Such limitations triggered the need for what has been called Route Optimisation (RO) for NEMO. Although there exist a number of proposed solutions that try to overcome the suboptimal routing problems that arise due to use the NEMO Basic Support protocol (see section 2.4), there is no solution that addresses the multifold problem of Route Optimisation in such a way that it solves all the main limitations of the NEMO Basic Support protocol under the most important deployment scenarios (e.g., nested and non-nested NEMOs, non-mobile and mobile capable nodes attached to the NEMO, etc). An additional requirement that current proposed mechanisms do not meet, is not to put any additional strong requisite on the operation of the nodes of the Mobile Network in order to be able to benefit from Route Optimisation.

This chapter describes a generic Route Optimisation solution for Network Mobility, designed in this PhD thesis, called Mobile IPv6 Route Optimisation for NEMO (MIRON) [BBC04], [BBCS05], [CBB+06]. MIRON is composed of two main modes:

- For those nodes of the mobile network that do not have any mobility capability, the Mobile Router (MR) performs all the Route Optimisation and mobility tasks on their behalf (what some authors [NZWT06] have called *Proxy MR*).
- For those nodes and (mobile) routers with standard Mobile IPv6 support, an address delegation mechanism, based on PANA (Protocol for Carrying Authentication for Network Access) [JLO+06] and DHCP [DBV+03], provides these nodes with topologically meaningful addresses (i.e. addresses that are directly reachable without requiring special rendezvous points, such as Home Agents, to be deployed to re-route any packet towards the actual location of the node). This enables these nodes to manage their own mobility and to perform the Route Optimisation by themselves.

These two different key modes of operation of MIRON combined give as a result a complete Route Optimisation solution for mobile networks, enabling traffic from any kind of node (with and without mobility support) and network configuration (including nesting) to be optimised. This is achieved without requiring changes on the operation of any node except Mobile Routers.

This chapter is organised as follows. An overview of the protocol is first provided in Section 5.2, before describing in detail how the proposed solution works. This description is presented next, divided into two sections, in order to deal with the Angular (Section 5.3) and Multi-angular (Section 5.4) Route Optimisation issues separately. A validation and evaluation of the protocol, taking into consideration security and scalability concerns, is provided in Section 5.5. In Section 5.6, the solution is compared with some existing RO proposals. After that, Section 5.7 explores an additional long term topic: the provision of Route Optimisation for NEMO by using secure delegation of signalling rights based approaches. Finally, some conclusions are provided in Section 5.8.

5.2. Protocol Overview

MIRON aims at improving the overall performance of communications involving nodes within a NEMO, by both avoiding data packets passing through the MR's HA and reducing the packet overhead due to the additional IPv6 headers introduced by the NEMO Basic Support protocol. MIRON does not introduce any change on the operation of the Correspondent Nodes and the Mobile Network Nodes, but only of the Mobile Routers.

Figure 5.1 shows a possible Route Optimisation target scenario for MIRON. It considers a mobile network deployed in a train, consisting of different types of MNNs:

- *Fixed nodes* in the train without mobility support (i.e. LFNs), such as internal servers or passengers' laptops.
- *Mobile devices* (i.e. VMNs), such as passengers' laptops, running Mobile IPv6, that keep using their Home IPv6 Addresses.
- Nested mobile networks, such as Personal Area Networks (PANs), e.g., a passenger's laptop, acts as a MR of his devices and is connected to the train's MR.

All of these devices access the Internet through the train's MR. This scenario includes almost every possible mobile network communication, involving LFNs, VMNs and nested NEMOs. Figure 5.1 also shows the different components every entity is composed of. Both the components and the way they work together to construct a complete Route Optimisation solution will be described in detail later in this chapter.

MIRON addresses two different Route Optimisation aspects:

- *Angular routing*. Angular routing is caused by the MRHA bidirectional tunnel introduced by the NEMO Basic Support protocol, since packets of a communication involving a MNN have to be forwarded through the HA of the NEMO. MIRON addresses this problem in two different ways, depending on whether the MNN that is communicating with a CN has mobility support or not. If the MNN has no Mobile

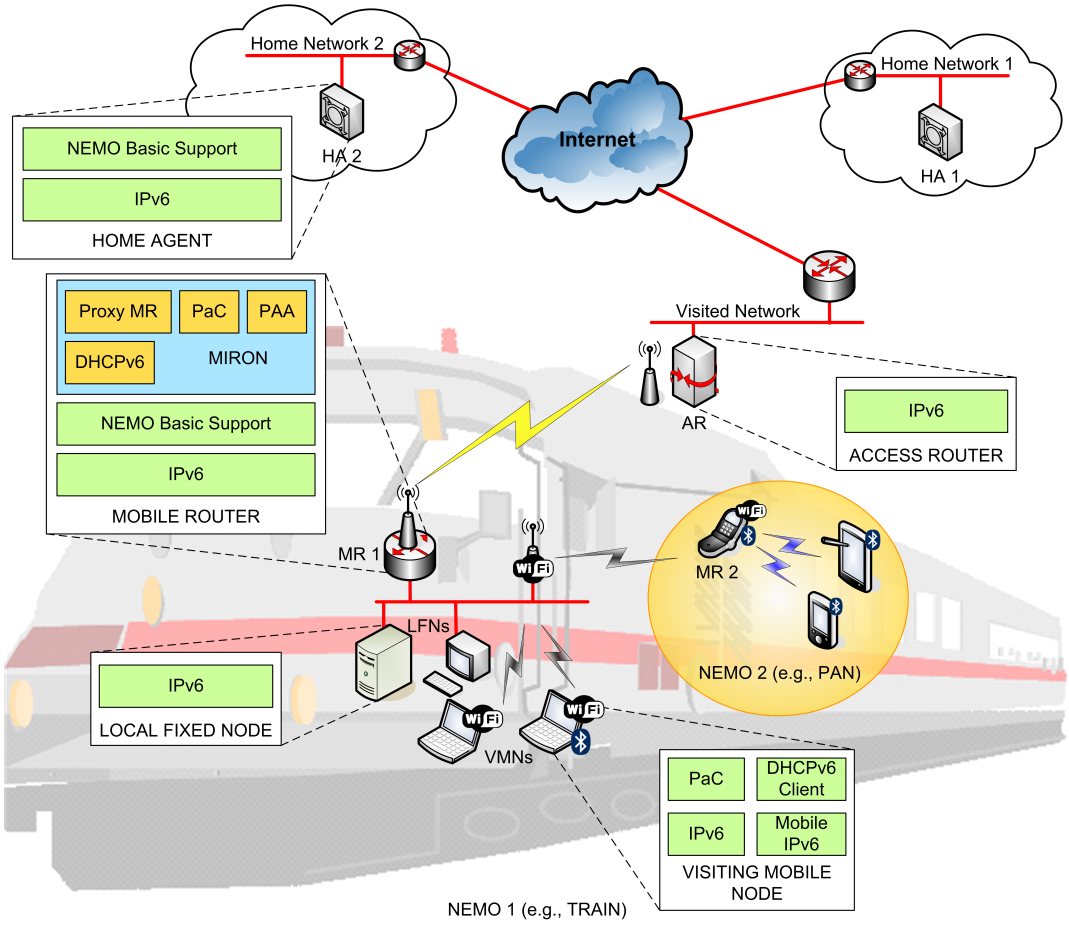

Figure 5.1: Overview of the MIRON architecture in a practical scenario.

IPv6 capabilities (i.e. an LFN), the approach followed by MIRON consists in delegating the Route Optimisation functionality to the MR, that performs all the RO signalling and packet handling on behalf of the LFNs. Therefore, the MR is a kind-of "Proxy MR" [NZWT06] for the LFNs of the NEMO. On the other hand, if the MNN is a Mobile IPv6 MN (i.e. a VMN) that is visiting the mobile network, MIRON takes advantage of the already available mobility support that the MN has. In this case, by using PANA and DHCPv6, the MR provides a topologically meaningful IPv6 address (that is, an address belonging to the network that the MR is visiting) to every VMN attached to the NEMO and updates it every time the NEMO moves. This, in addition to a routing mechanism that enables these addresses to be routed inside the NEMO, allows the VMN to make use of its own Mobile IPv6 Route Optimisation functionality, therefore avoiding traversing the MR's HA and reducing the packet overhead.

- *Multi-angular routing.* Multi-angular routing is caused in nested NEMOs by the chain of nested MRHA bidirectional tunnels that packets should traverse. MIRON addresses this problem by using PANA and DHCPv6 to provide topologically meaningful IPv6 addresses to every MR in the nested NEMO hierarchy. In this way, every MR has an

IPv6 address belonging to the network that the root-MR (that is, the MR of the NEMO at the top of the hierarchy) is visiting. This, in addition to a routing mechanism that enables these addresses to be reachable, makes it possible to avoid traversing any HA.

The set of mechanisms of MIRON enables direct path communication between a MNN (LFN or VMN) and a CN, avoiding the suboptimal MR-HA path. The recursive tunnelling due to nesting is also eliminated, therefore optimising the traffic in every possible configuration of a mobile network. MIRON only introduces changes in the MR (see Figure 5.1), while MNNs, HAs and CNs remain unchanged, thus facilitating the deployment of the solution. The next two sections provide a detailed protocol walk-through of MIRON.

5.3. Angular Route Optimisation

If no Route Optimisation mechanism is used, all the traffic sent/directed to a MNN goes through the bidirectional tunnel set up between the MR and its HA. MIRON enables direct communication – without traversing the MR's HA – by following one of the next approaches, depending on the type of MNN:

- Local Fixed Node (LFN). LFNs do not have mobility support, so any mechanism that attempts to optimise their traffic should be implemented without requiring support from the LFN itself. The MIRON mechanism for LFNs is basically a proxy-MR approach, in which the MR performs the Mobile IPv6 Route Optimisation [JPA04] on behalf of the LFN.

- Visiting Mobile Node (VMN). VMNs are Mobile Nodes that are visiting the mobile network, managing their own mobility. By default, the Care-of Address obtained and used by a VMN attached to a NEMO belongs to the Mobile Network Prefix of that NEMO, so although these mobile nodes may be performing Route Optimisation with the CNs they are communicating to, there still exists a tunnel – between the NEMO's MR and the MR's HA – introduced by the NEMO Basic Support protocol. In this case, our proxy-MR approach is not feasible, therefore a different mechanism is used. MIRON takes advantage of the mobility support that VMNs already have. Basically, we propose a mechanism, using PANA and DHCPv6, that enables the VMNs to configure topologically valid IPv6 addresses (i.e. those addresses that belong to the address space of the foreign network the NEMO is visiting) as CoAs, and letting the VMNs manage their mobility and perform their Route Optimisation tasks.

5.3.1. Detection of the type of node

In order to apply the appropriate Route Optimisation mechanism, the MR should first be able to determine which kind of node (LFN or VMN) every node that is communicating is. The MR performs such a task by looking for Binding Update messages received at its ingress interfaces, since an MN right after gaining connectivity to a foreign network and configuring a new CoA (from the MNP), has to send a Binding Update to its HA to inform it about its new location (i.e. MN's CoA).

5.3.2. Route Optimisation mechanism for LFNs

Local Fixed Nodes are nodes without any mobility support running, therefore a mechanism that optimises their traffic cannot relay on any mobility function implemented by them. MIRON puts this LFN mobility support into the MR, that performs all the required mobility and Route Optimisation tasks on behalf of the LFNs attached to it.

The mechanism basically consists in enabling a MR to behave as a proxy for the LFN, performing the Mobile IPv6 Route Optimisation signalling and packet handling [JPA04] on behalf of the LFN. In order to do that, the MR first tracks the different communications that LFNs have established and decides which of those will be optimised, since optimising a traffic flow involves a cost – in terms of signalling and computation resources at the MR – that may not be worth for some kinds of flows (e.g., DNS queries). This decision (that is, whether to perform Route Optimisation for each flow or not) is out of the scope of this PhD thesis.

For those LFN-CN pairs whose traffic is to be optimised, the MR starts to send the RO signalling described for standard Mobile IPv6 in [JPA04]:

- The BU is sent by the MR.
- The BU contains the LFN's address as the Home Address (HoA) and the MR's CoA as the CoA (since the MR's CoA is the only topologically meaningful address available).

The Route Optimisation mechanism defined by Mobile IPv6 [JPA04] requires an additional procedure to be performed before sending the BU message, in order to mitigate possible attacks [NAA+05]. Basically, this mechanism, called Return Routability (RR), verifies that the node that is reachable at the HoA is able to respond to packets sent to a given CoA (different to the HoA of the node). This mechanism can be deceived only if the routing infrastructure is compromised or if there is an attacker between the verifier and the addresses (that is, HoA and CoA) to be verified. With these exceptions, the test is used to ensure that the MN's Home Address (HoA) and MN's Care-of Address (CoA) are collocated.

In our solution we adopt the procedure described above. For this purpose, the MR has to perform the Mobile IPv6 Return Routability procedure [JPA04] on behalf of the LFN. Such a procedure involves sending the Home Test Init (HoTI) and Care-of Test Init (CoTI) messages to the CN and processing the replies (Home Test message – HoT – and Care-of Test message – CoT). These messages are sent as specified in [JPA04], using the LFN's address as the source address in the HoTI message – which is sent encapsulated through the MR's HA –, and the MR's CoA as the source address in the CoTI message. With the information contained in the HoT and CoT messages, sent by the CN in response to the HoTI and CoTI messages respectively, the MR is able to build a BU message to be sent to the CN on behalf of the LFN. This message is sent using the MR's CoA as the packet source address and carries a Home Address destination option set to the LFN's address.

Besides performing the Route Optimisation signalling on behalf of the LFN, the MR has to process the packets sent by and directed to the LFN. Packets sent by the CN follow a direct path to the MR, not traversing the HA, as a result of the Route Optimisation. These packets carry the MR's CoA as destination address, and also carry a Type 2 Routing Header with the LFN's address as next hop. The MR processes and removes the Routing Header of the packet, checking if the next hop address belongs to one of its LFNs and, if so, delivering

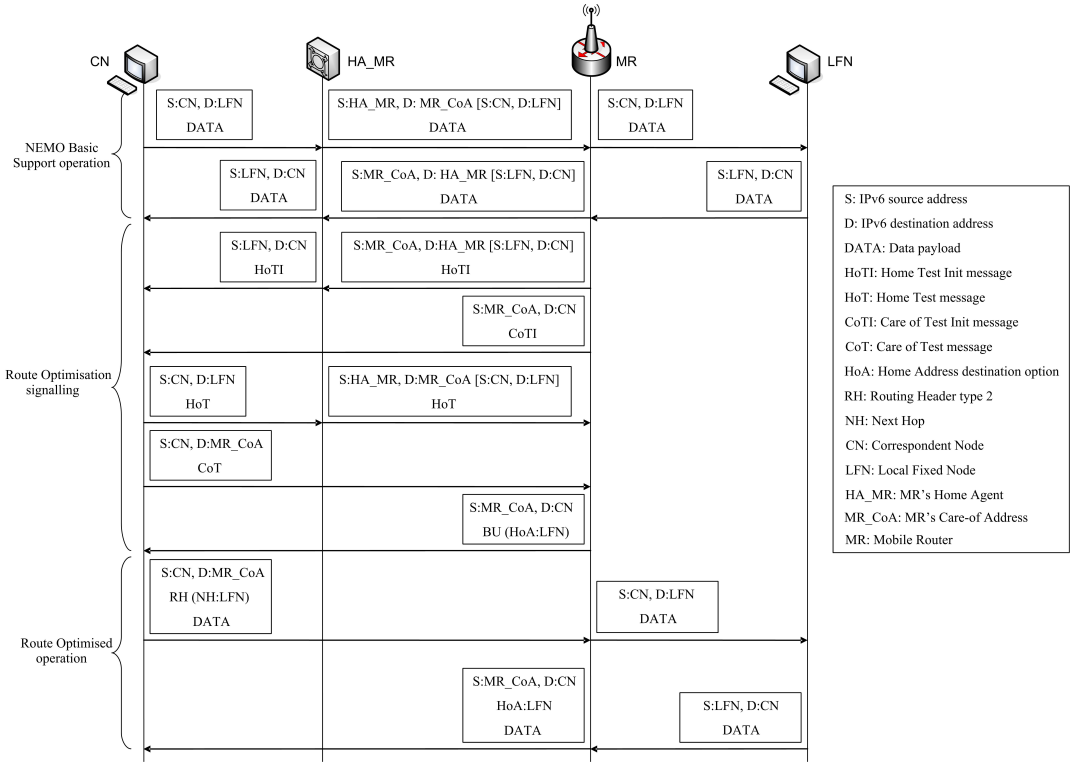

Figure 5.2: Route Optimisation mechanism for LFNs: Proxy-MR operation.

the packet to the LFN. Current Mobile IPv6 specification [JPA04] defines that IPv6 nodes which process a Type 2 Routing Header must verify that the address contained within is the node's own Home Address. This is done in order to prevent packets from being forwarded outside the node. In MIRON this has been changed and the MR verifies that the address contained in the Routing Header is the address of one of the LFNs that the MR is acting as Proxy-MR. In the opposite direction, the MR receives the packets sent by the LFN and performs the following actions on every packet:

- Set the MR's CoA as IPv6 source address.
- Insert an IPv6 Home Address destination option, carrying the address of the LFN.

Figure 5.2 shows the signalling and data flows of the proposed Route Optimisation mechanism for LFNs, including at the top of the figure the NEMO Basic Support protocol data flow for comparison purposes.

5.3.3. Route Optimisation mechanism for VMNs

Visiting Mobile Nodes are nodes that support mobility (that is, nodes running Mobile IPv6 [JPA04]) and are visiting a mobile network. Therefore the VMN is attached to an Access Router that is the NEMO's MR, and the address that the VMN obtains and configures as CoA belongs to the Mobile Network Prefix. In this case, our proxy-MR mechanism used

for LFNs cannot be used, as the VMN itself may generate Route Optimisation signalling with its CNs. Besides, the MR cannot modify the RR and RO signalling sent by the VMN in order to make the MR's CoA the CoA that the CN uses to send the packets to the VMN, because part of the RR signalling is protected by IPsec (the HoTI message is sent through the VMN's HA protected by IPsec ESP).

In this thesis, two different approaches were explored to provide VMNs with Route Optimisation, although only one was finally adopted by MIRON. The first mechanism is based on linked Mobile IPv6 Binding Cache entries, while the second is based on the use of PANA and DHCPv6 to provide topologically meaningful IPv6 addresses to VMNs. A detailed description of these two mechanisms is included next.

5.3.3.1. Linked Mobile IPv6 Binding Cache Entries

As we have previously explained, a MR cannot modify the Route Optimisation signalling generated by an attached VMN in order to make the MR's CoA be the address that the CN uses as the VMN's CoA. However, Mobile IPv6 specification [JPA04] does not prevent from having linked Binding Cache (BC) entries (as long as a circular reference – a loop – is not created). A linked BC entry exists when the Care-of Address of an entry appears as the Home Address of a different entry.

A VMN attached to a NEMO configures a CoA that belongs to the Mobile Network Prefix of the moving network. This CoA is used by the VMN in the Home Registration to its HA, but also in the Route Optimisation signalling sent to its CNs. Therefore, a CN communicating with a VMN performing RO, would have an entry in its BC as follows:

Home Address	Care-of Address
VMN's HoA	VMN's CoA ($\in$ MNP)

A MR may help a VMN by sending Binding Updates to the CN on its behalf. This BUs would bind the VMN's CoA to the MR's CoA, so an additional entry would be added to the CN's BC:

Home Address	Care-of Address
VMN's HoA	VMN's CoA ($\in$ MNP)
VMN's CoA ($\in$ MNP)	MR's CoA

The Mobile IPv6 specification (RFC 3775 [JPA04]) does not clearly define what should be the behaviour of a CN with linked BC entries (such as the previous one) when it has to send a packet to VMN's HoA. The logical processing would be to the following:

- The CN examines its Binding Cache for an entry for the destination address (VMN's HoA) to which the packet is being sent. Since the CN has a BC entry for this address, it adds a type 2 Routing Header to route the packet to the VMN (the destination node) by way of its Care-of Address (that is, the destination address of the packet is set to the VMN's CoA).

- The CN examines again its BC for an entry for the new destination address (VMN's CoA) to which the packet is now being sent. Again, the CN has a BC entry for this

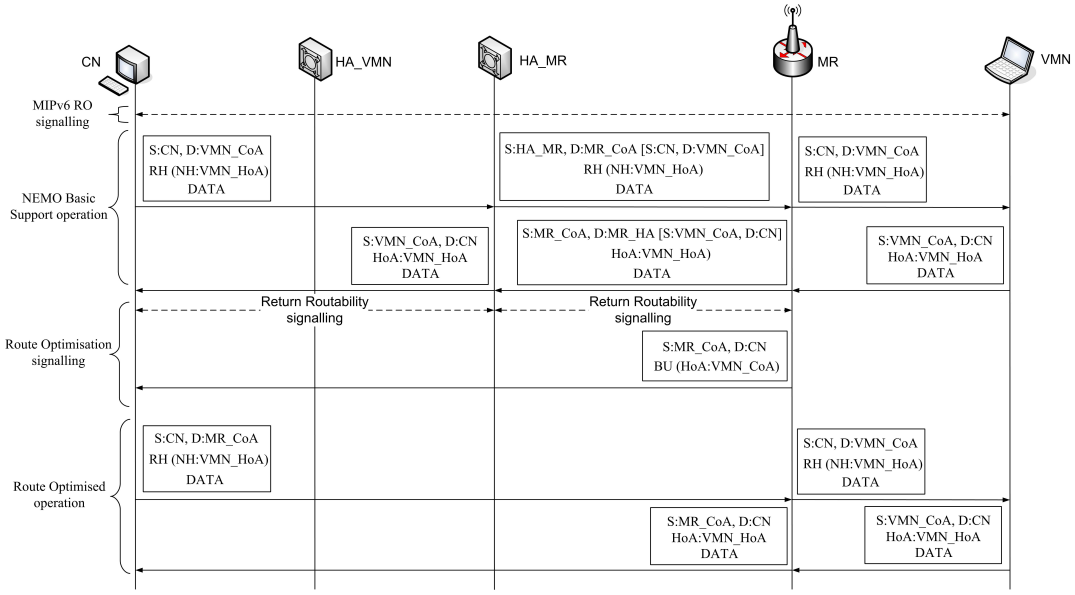

Figure 5.3: Route Optimisation mechanism for VMNs: Linked Mobile IPv6 Binding Cache Entries.

address, so it adds a new type 2 Routing Header to the packet, and sets the destination address to the Care-of Address of this BC entry (i.e. MR's CoA).

However, IPv6 specification (RFC 2460 [DH98]) states that a Extension Header (such a Routing Header) can occur at most once in a packet (except for the Destination Options Header, that can occur at most twice). Besides, RFC 3775 [JPA04] restricts the type 2 Routing Header to carry only one address, so the above behaviour cannot happen.

On the other hand, it is not clear what a Correspondent Node Mobile IPv6 implementation would do if it receives Binding Updates that creates linked BC entries. A Mobile Router would still benefit from sending a BU, binding the VMN's CoA to the MR's CoA, if the CN performed the following processing when sending a packet:

- The CN examines its Binding Cache for an entry for the destination address (VMN's HoA) to which the packet is being sent. Since the CN node has a BC entry for this address, it adds a type 2 Routing Header to route the packet to the VMN (the destination node) by way of its Care-of Address (that is, the destination address of the packet is set to the VMN's CoA).

- The CN examines again its BC for an entry for the new destination address (VMN's CoA) to which the packet is now being sent. The CN has a BC entry for this address, and it sets the destination address of the packet to the Care-of Address of the last BC linked entry (i.e. MR's CoA).

If a CN behaved like described above, a MR would be able to enable direct path communication between a CN and an attached VMN by performing the following operations (see Figure 5.3):

- Once the MR has detected that a VMN has attached to the NEMO, and that it is doing Route Optimisation with a CN, the MR sends a BU to the CN binding the VMN's CoA (that belongs to the NEMO's MNP) and the MR's CoA. This BU creates a linked BC entry in the CN.

- The MR processes packets sent by the CN to the VMN, since the CN – because of the linked BC entry – now is sending packets to the VMN with the MR's CoA as destination address. The MR replaces this destination IPv6 address with the VMN's CoA.

- The MR processes packets in the reverse direction as well (that is, packets sent by the VMN to the CN), replacing the source IPv6 address with the MR's CoA (the VMN originally sets it to the VMN's CoA).

Although this Route Optimisation mechanism is allowed by current Mobile IPv6 specification (RFC 3775 [JPA04]), a crucial point was to check how existing MIPv6 implementations behaved about linked BC entries. This was the next step after designing the mechanism presented above (that was actually a natural extension of the Proxy-MR operation defined for the RO support of LFN communications). In order to investigate the operation of MIPv6 CN implementations, MIPL[1] was chosen as a candidate for analysis, because it is one of the most used available MIPv6 implementation. MIPL is an open source implementation of Mobile IPv6 fully compliant with RFC 3775 [JPA04]. It was checked whether a CN running MIPL 2.0 RC2 would support our mechanism, and the result was that it was not possible without modification of MIPL (although it would be possible to easily modify it to support linked BC entries).

Based on our analysis of a representative Mobile IPv6 implementation, it seems that linked BC are not well supported by existing MIPv6 implementations, which means that a modification of CNs would be required to enable our proposed mechanism to work. Therefore, we decided to seek for a different approach to address our requirement of enabling NEMO RO for VMNs. The mechanism that we finally designed and adopted is described in the next section.

5.3.3.2. PANA-based Address Delegation

The Route Optimisation approach that MIRON defines for Visiting Mobile Nodes attached to a NEMO is based on taking advantage of the mobility support that these nodes already have, providing the means to the VMN to perform the RO. In order to allow the VMN to manage its own mobility and enable it to perform Route Optimisation with the CNs (in a way that avoids the MRHA bidirectional tunnel), we propose the following:

- Provide a topologically meaningful IPv6 address to the VMN. These addresses are those that belong to the network that the root-MR is visiting.

- Enable this address to be routable inside the NEMO, as it only has topological meaning in the visited network. The MR has to perform proxy neighbour discovery for this

[1] Mobile IPv6 for Linux `http://www,mobile-ipv6.org/`

address in the egress interface that is attached to the network to which the address belongs. Besides, the MR has to insert a host route for this address to be able to route packets destined to it.

- Perform source address routing in the MR in order to send directly (that is, avoiding the bidirectional MR-HA tunnel that still exists and is used for non-optimised traffic) packets sent by the VMN.
- Update the address of the VMN when the NEMO moves.

A mechanism that fulfils the previous goals should be able to allow VMNs, and only the VMNs – the mechanism must not affect other type of nodes –, to obtain a new IPv6 address to be used as the CoA, whenever the MR wants to, and in a secure way that does not introduce any new security threat.

The Route Optimisation mechanism for VMNs that we propose in this section uses a particular functionality that is included in the PANA protocol[2], namely, the capability of telling a node that it must change its IPv6 address and how to get a new one.

This imposes the requirement that PANA client (PaC) software must be available in VMNs for providing them Route Optimisation, and PaC and PANA agent (PAA) software must be available in MRs. The PaC software in the MRs is needed to optimise nested NEMOs as it will be described in the next section. The PAA software in MRs is needed to support the Route Optimisation for VMNs visiting the NEMO, and to support the Route Optimisation of nested NEMOs. PANA support is not required by MIRON in the access network that the root-MR is visiting (in the infrastructure access network) nor in the LFNs attached to the NEMO.

The assumption of availability of PANA software in the MRs is not a problem, because MIRON is based on modifications in the MRs software, PANA is just an additional software to have. The assumption of PANA in VMNs can be more restrictive. The idea of the solution is that MIPv6 compliant Mobile Nodes can visit the NEMO and optimise its routing just like when they visit an infrastructure access network. This will not be true if they do not have a PANA client installed.

Most of current access networks (such as hotspots deployed in airports and cafeterias) require users to authenticate to the network before gaining Internet access. As the number of hotspots continues growing in the coming years, authentication mechanisms will be more and more important in order to avoid non-authorised users using and wasting the network resources. Using a standard protocol to perform such authorisation and authentication tasks would help in the deployment of ubiquitous access “anytime anywhere” networks. Our Route Optimisation protocol, MIRON, assumes that (i) an authentication protocol will be used in public heterogeneous access networks and that (ii) PANA [YOP+05], [JLO+06], [FOP+06] will be a standard protocol widely deployed and used, so PANA support will be available in VMNs.

We argue that assuming that VMNs will have PaC software does not limit the practical usability of MIRON for Route Optimisation in VMNs, since, on one hand, it is not realistic to assume public access networks to be open and not to require any kind of authentication. On the other hand, we assume that PANA support will be available on VMN, because it is

[2] A brief summary of PANA is provided in Appendix A.

expected that PANA will become a standard authentication protocol once its specification is concluded within the IETF, finishing with the current status of multiple possible authentication mechanisms (e.g., IEEE 802.1X, proprietary web-based systems, etc).

Even if that is not finally the case, and PANA does not turn to be the standard authentication protocol in heterogeneous networks, a different protocol that is able to provide IPv6 routable addresses to arriving VMNs and change them every time the NEMO moves, could be alternatively used. One example could be the use of the DHCPv6 *Reconfigure* mechanism [DBV+03], using some authentication information between MR and VMN obtained from any other means to authorise the MR changing the IPv6 address used by the VMN.

Anyway, if neither PANA nor an alternative protocol is available in a VMN, this VMN – attached to a MIRON MR – will just not benefit from the Route Optimisation mechanism provided by MIRON, and its traffic will follow the suboptimal path provided by the NEMO Basic Support protocol.

The mechanism to provide an IPv6 address to the VMN using PANA works as follows (see Figure 5.4): when a VMN attaches to a NEMO, it initiates the PANA session (PANA discovery and handshake phases). Immediately after that, the actual authentication and authorisation phase (by executing EAP between the PAA and PaC) takes place. Then, the VMN is authorised to access the network and it has an IPv6 address. This address is obtained by using the address autoconfiguration mechanism available at the NEMO. Initially, we assume that we are using stateless address configuration for addresses of the Mobile Network Prefix, but later we will see that we can also use stateful address configuration within the NEMO. The VMN then sends a Binding Update message to its Home Agent, informing about its current location. Once this BU is received at the MR, it becomes aware that a new VMN is now attached to the NEMO. The MR discards this BU message and starts a PANA re-authentication phase.

During the PANA re-authentication phase, the PAA located in the MR tells the PaC located in the VMN that it should obtain and configure a new IPv6 address (Post-PANA address, POPA) and how to obtain it, by including available configuration methods in a Post-PANA-Address-Configuration (PPAC) AVP contained in a PANA message (PANA-Bind-Request). DHCPv6 is the only available configuration mechanism listed in the message, and upon the reception of that, the VMN requests an address using DHCPv6. There is a DHCPv6 component located at the MR that receives the DHCPv6 requests from the VMN and then obtains (using one of the available autoconfiguration mechanisms at the foreign network) an IPv6 address. The DHCPv6 component generates a DHCPv6 reply – including this address – that is delivered to the requesting VMN. This DHCPv6 component implements the client part of DHCPv6 and also some reduced functionalities of the server part (e.g., the generation of DHCPv6 responses), but it is not a DHCPv6 server (for example, the DHCPv6 component does not have a pool of available addresses, each time an address is needed, it obtains it from the foreign network), although the implementation of this DHCPv6 component can be performed very easily from the code of a normal DHCPv6 client and server implementation. Once the MR has sent the DHCPv6 reply – including the (/128) delegated address – to the VMN, the PaC in the VMN conveys this newly configured IPv6 address to the PAA in the MR by sending the PANA-Update-Request message.

The use of stateful address configuration (DHCPv6) within the NEMO (to configure addresses from the MNP) is also possible, but it requires the DHCPv6 component at the MR

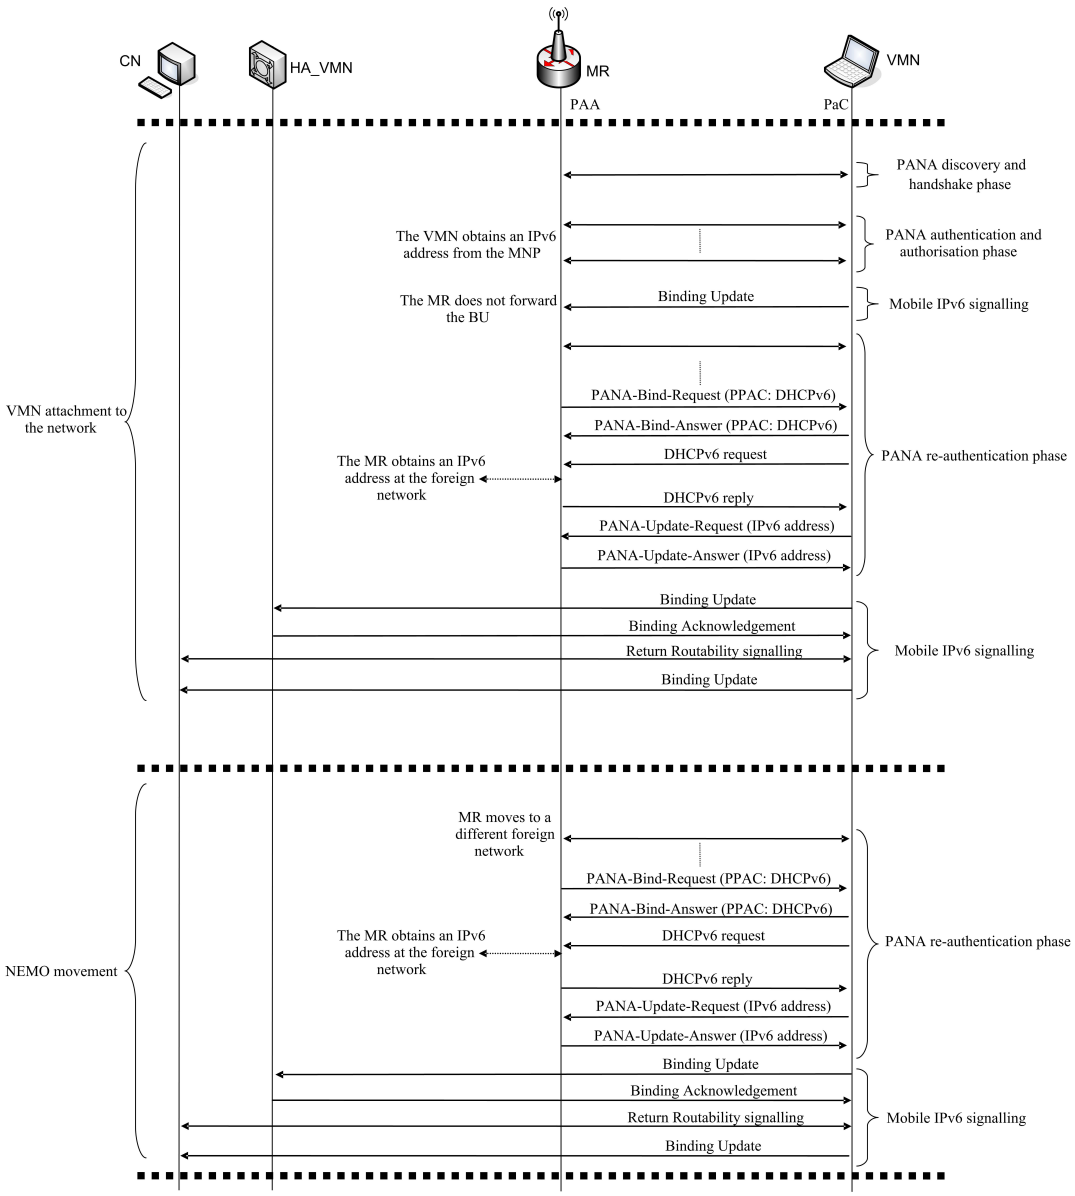

Figure 5.4: Route Optimisation mechanism for VMNs.

to implement the complete server functionality and to check, before providing an address, whether or not the requesting node is an identified VMN, to know if this address should belong to the MNP or to the visited network address space. Nodes are identified as VMNs by the MR according to the procedure described above.

In order to enable the VMNs' addresses reachability inside the NEMO, the MR has to add a host route for each VMN's address and perform proxy neighbour discovery on its egress interface (the interface that is connected to the link where the address has topological meaning), allowing the MR to forward packets to their final destinations. Both the delegated IPv6 addresses and the host routes have a lifetime that prevents this state to remain in the

network after a sub-NEMO or a node leaves a parent-NEMO (the value depends on the lifetime of the address obtained by the root-MR).

The VMN, triggered by the change of address, starts the Mobile IPv6 location update process, sending first a Binding Update (BU) message to its HA. The VMN may then update the location information in the CNs it is communicating with (if the VMN is running Route Optimisation with its CNs). This process consists of the VMN performing the Return Routability process [JPA04] and sending a BU to every CN whose traffic is to be route optimised.

When the NEMO moves to a different foreign network, the MR requests new IPv6 addresses and provides them to the VMNs attached to the NEMO by starting a new PANA re-authentication phase. The MR requests VMNs to configure a new IPv6 address using DHCPv6.

Due to the PANA and DHCPv6 signalling, MIRON takes much longer to finish its handover than the NEMO Basic Support. Similarly to the case of Mobile IPv6, micromobility solutions such as Fast Handovers for Mobile IPv6 [Koo05], may be designed/adapted to MIRON to alleviate the increase in the handover delay [BSM+05].

5.4. Multi-angular Route Optimisation

The routing inefficiencies due to the MRHA bidirectional tunnel are exacerbated when NEMOs are attached to other NEMOs, forming a nested NEMO. Packets belonging to a communication between a MNN of a nested NEMO and a CN have an additional IPv6 header per nesting level and traverse the HAs of every MR of the nested NEMO.

The problem of enabling RO for nested NEMOs (i.e. MRs visiting a NEMO) is very similar to that of VMNs (i.e. MNs visiting a NEMO). Both VMNs and MRs are nodes that are mobile-capable and can manage their own mobility. Routing inefficiencies arise from the fact of not using topologically meaningful addresses (i.e. addresses belonging to the NEMO MNP) as CoAs. Section 5.3.3 describes an address delegation mechanism with a built-in routing system that is able to provide IPv6 addresses – belonging to the foreign network that the MR is visiting – to a VMN in a secure way, by using PANA facilities.

MIRON extends that solution, used for providing Angular RO for VMNs, to enable Multi-angular RO in nested NEMOs. Basically, the solution consists in providing topologically meaningful addresses – that is, those that belong to the foreign network that the root-MR is visiting – to every MR in the nested NEMO. The same PANA-with-DHCPv6-based mechanism is used to provide an IPv6 address to a MR that attaches to a NEMO (and to change it when one of the parent NEMOs moves). MRs have both a PAA and a PaC component and also a DHCPv6 component, so when a MR connects to a mobile network, they are able to get and configure a new IPv6 address.

Providing topologically meaningful addresses is not the only required step to avoid the suboptimal multi angular routing in nested networks. Another requirement that has to be met is that these addresses are globally reachable. To enable that, every MR in the nested NEMO keeps track of the address of the node requesting an IPv6 address using DHCPv6, so when the delegated address is received, it can insert a host route entry in its routing table that allows it to route packets destined to that address afterwards. This information is also used to perform source address based routing for the packets generated inside the NEMO, as every

MR should know for each packet if it has to be sent directly to the router it is connected to (in this way, avoiding the tunnel), or it has to be sent towards the HA, through the bidirectional tunnel (for traffic that is not being optimised).

This address delegation mechanism with built-in routing avoids the multi encapsulation and multi angular routing in nested networks. Besides, it enables angular MIRON route optimisations to work when applied to a NEMO located at any level of a nested NEMO.

5.5. Validation and evaluation of the proposed solution

This section provides both an experimental and analytical evaluation of MIRON. The main aim of this evaluation is to study the performance of MIRON and compare it with the NEMO Basic Support protocol. A security and scalability analysis is also provided, in order to demonstrate that the designed mechanism has no critical security or scalability issues that may have an impact on the deployment of the solution.

5.5.1. Experimental evaluation

5.5.1.1. MIRON implementation

In order to be able to conduct real experiments that allowed us to evaluate the performance of the NEMO Basic Support protocol and the improvements provided by MIRON, we first implemented the NEMO Basic Support protocol [dlOBC06]. A first prototype of MIRON was also implemented, providing all the Route Optimisation mechanisms [BOC+06].

Packets belonging to a communication flow optimised by MIRON must not traverse the bidirectional tunnel. Therefore, for outgoing traffic, a host route towards the CN of the flow should be inserted at the MR, to avoid the default route through the tunnel interface. Besides, there may be simultaneously communications in a NEMO – from different MNNs – with the same destination CN that are not all being optimised, thus source address based routing is necessary.

The required additional protocols and procedures (such as the Return Routability and DHCPv6) were completely implemented, with the exception of PANA, that is currently being implemented and integrated. The fact that PANA is not implemented does not have an impact on the results obtained in the tests, as we have focused on the TCP throughput, and the PANA signalling is generated during handovers (and also periodically to renew the lifetime). In this PhD thesis, we have not been concerned about the performance of our solution during handovers as we address the problem of Route Optimisation, just in the same way that the Route Optimisation solution for Mobile IPv6 does. Improvements in the handover latency (like the ones designed for Mobile IPv6 [Koo05], [SCMB05], [BSM+05]) requires further study and will be addressed in future works.

The NEMO Basic Support protocol [dlOBC06] and MIRON were mostly implemented in user space, because in this way the development was easier and quicker than doing that in the kernel. The main software characteristics are: a Linux machine with kernel linux-2.6.x (tested with linux-2.6.8.1) with support for IPv6-in-IPv6 tunnels (used for the MRHA bidirectional tunnel) and Netlink sockets, and the *pcap* library (used for the capture and processing of the mobility related signalling).

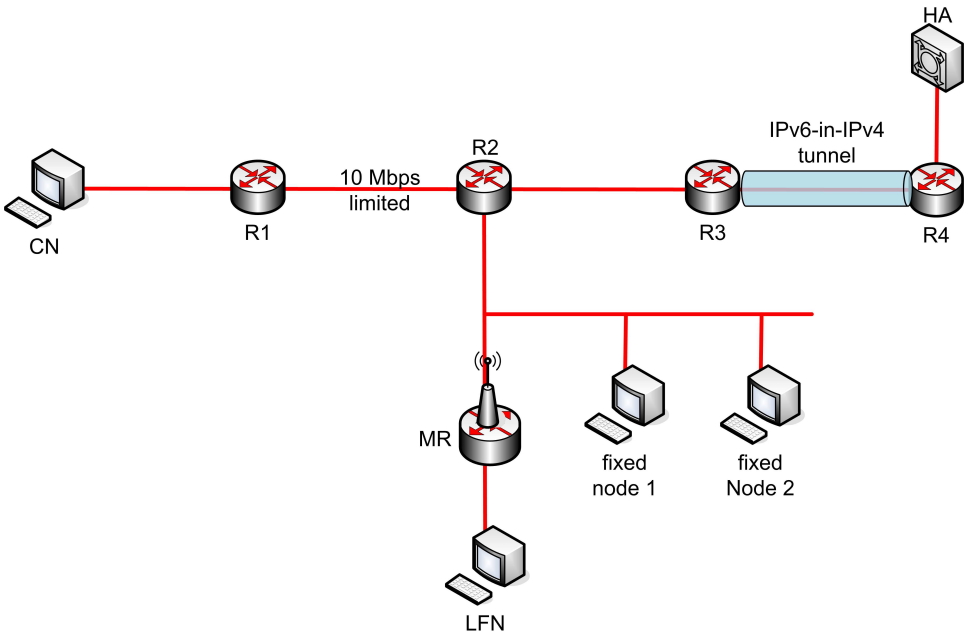

(a) Non-nested scenario.

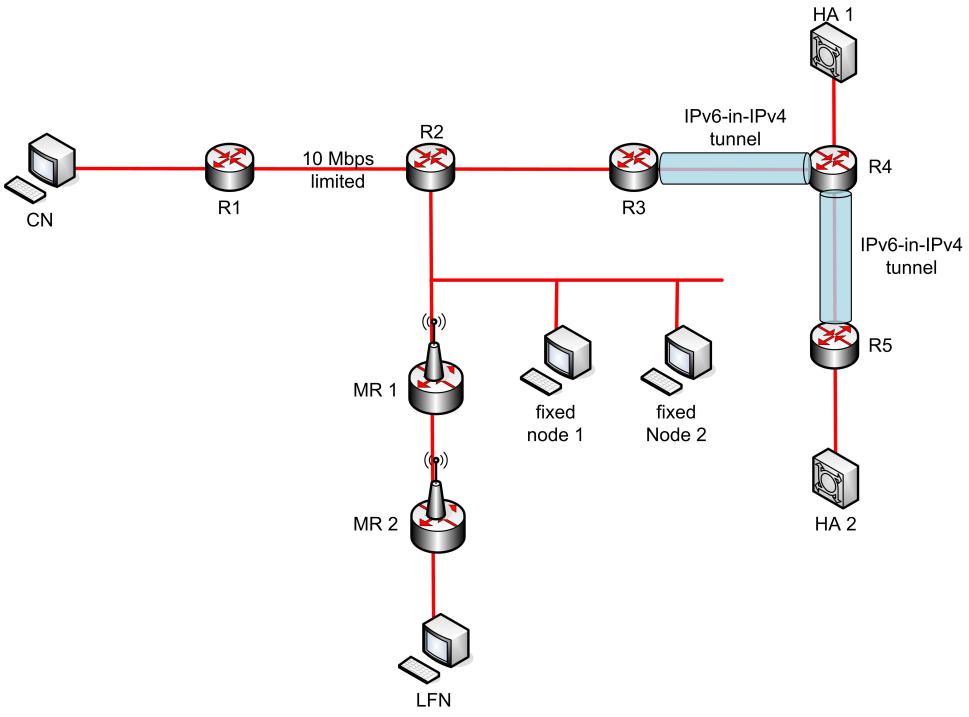

(b) Nested scenario.

Figure 5.5: Network mobility testbed employed during the experimental evaluation.

5.5.1.2. Studied scenarios

Two different scenarios (Figure 5.5) were deployed to allow us to experimentally test the performance of MIRON and compare it with the NEMO Basic Support solution. The first

one (Figure 5.5(a)) was used to evaluate the performance in a non-nested case, whereas the second (Figure 5.5(b)) is an extension of the former to include nesting.

Next, we describe the second scenario, as it is an extension of the first one. This scenario (Figure 5.5(b)) consists of thirteen Mandrake 10.0 Linux machines (all with linux-2.6.8.1 kernels, except 3 routers that run linux-2.4.22). Five of them act as *fixed* (i.e. non-mobile) routers (R1 to R5), two as Home Agents (HA1 and HA2), two as Mobile Routers (MR1 and MR2), one as Correspondent Node (CN), one as Local Fixed Node (LFN) and two as *fixed* nodes (Fixed nodes 1 and 2). This is part of the IST Daidalos[3] project testbed at the Universidad Carlos III de Madrid.

All mobility-aware nodes run network mobility software, that is, the NEMO Basic Support protocol (at the HA and MR) and MIRON (at the MR only) developed by us. The CN runs MIPL 2.0 RC2, with the support of Route Optimisation enabled.

We need to be able to modify the delay in the path followed by packets of a communication between a CN and a MNN (that is, the path between the CN and the MNN's HA and/or the path between the MNN's HA and the foreign network the MR is currently attached to). This allows us evaluate how the performance of a particular network mobility solution is affected by network characteristics, such as the particular location of mobile networks, home networks and correspondent nodes. For this purpose, we used the NIST Net emulator[4]. NIST Net allows a single Linux PC, set up as a router, to emulate a wide variety of network conditions (e.g., latency, jitter, packet loss, ...).

We were interested in studying how the delay (and also the packet overhead) introduced by the MRHA bidirectional tunnel affects the performance of applications. TCP performance is heavily dependent on the round trip time (RTT) between the communication peers. Taking this into consideration and the fact that 85% of the traffic in the Internet is generated by TCP connections [MC00b], the TCP study case becomes very interesting to be performed and analysed. Therefore, we set up an scenario that allowed us to modify the delay in the CN-HA-MR path.

Other network characteristics, besides the delay, that do not have an special effect in the TCP performance and that are also present in non-mobile networks, were not modified.

NIST Net software runs only in IPv4 and with linux-2.4.x kernels (at the moment we performed these tests). Therefore, in order to use it in our testbed, we had to set up an IPv6-in-IPv4 tunnel – between R3 and R4 and between R3 and R5 – using it in our IPv6 scenario. In the first scenario, the non-nested one (Figure 5.5(a)), every packet in the CN-HA-MR path traverses the IPv6-in-IPv4 tunnel, which allows us to modify the network behaviour by changing the parameters of the NIST Net emulator running in R3 and R4. In the rest of the path followed by packets, native IPv6 is used, so the tunnel inclusion does not affect the overall test performance except for the small added delay due to IPv6-in-IPv4 tunnelling and the reduction of the PMTU (the situation is not different from having a change of the transport link technology in the path and it is transparent to the IPv6 behaviour). Actually, the IPv4 tunnel clearly shows the current status of IPv6 networks in the Internet, with lot of IPv4 clouds connecting IPv6 native networks. In the nested scenario, a second IPv6-in-IPv4 tunnel was set up – between R4 and R5 – to allow us to modify the network delay between the two different home networks (i.e. between HA1 and HA2).

[3] `http://www.ist-daidalos.org/`
[4] `http://www-x.antd.nist.gov/nistnet/`

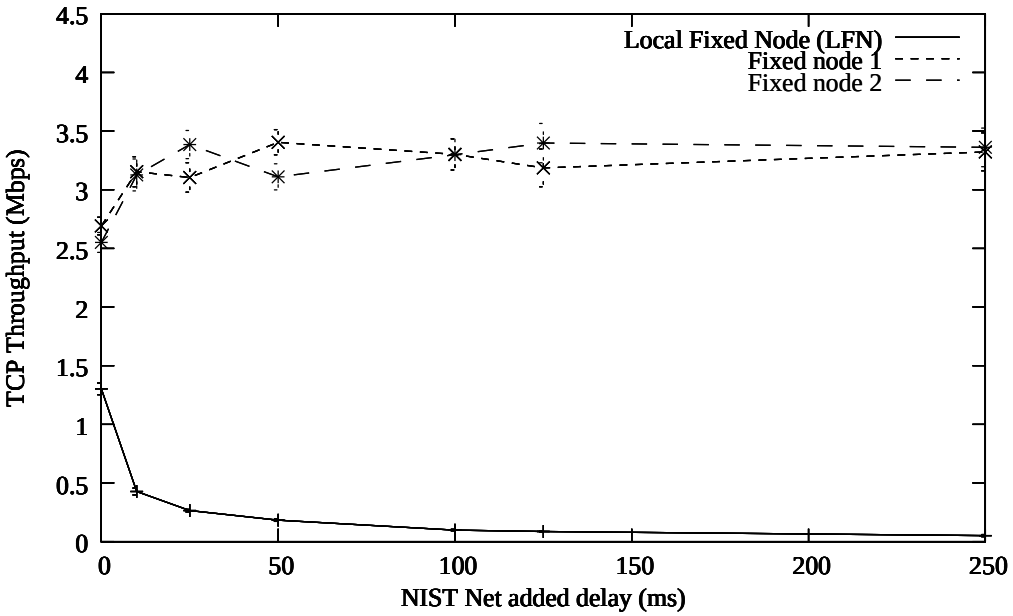

Figure 5.6: Impact of NEMO Basic Support protocol on the TCP throughput.

To avoid the influence of the wireless media characteristics and interferences from other neighbouring wireless networks, the performance tests were conducted using wired mobile routers, although experiments using wireless mobile networks were also performed to check the correctness of our solution.

5.5.1.3. Impact of network mobility on the TCP performance

The suboptimal routing introduced by the NEMO Basic Support protocol [DWPT05] makes packets not follow the direct CN-MR-MNN path, but the usually longer CN-HA-MR-MNN path. This adds a delay in the packet delivery that can significantly reduce the performance of certain applications. Furthermore, packets are encapsulated between the HA and the MR, thus reducing the PMTU. Both effects, increased delay and reduced PMTU, have an impact in the performance of applications.

The test consists in measuring the average TCP throughput of an MNN (an LFN in the tests) downloading a file from a CN, while two other non-mobile network hosts (Fixed nodes 1 and 2), attached to the same network the NEMO is visiting, simultaneously download the same file, both in a non-nested and in a nested scenario (see Figures 5.5(a) and 5.5(b)). The available bandwidth between the CN and the network that the mobile network is visiting was limited to 10 Mbps, by setting the R1-R2 link to 10 Mbps Half-Duplex. The tool used for the download was *scp* (secure copy) and the size of the file was 50 MBytes.

Each average TCP throughput sample was calculated over a 20 seconds independent interval of download and at least 30 samples were obtained for each test (to guarantee the statistical validity of the measurements).

For the non-nested scenario, the unidirectional NIST Net added delay of the link R3-R4 – *delay1* – was varied between 0ms (i.e. home network, visited network and CN locate very

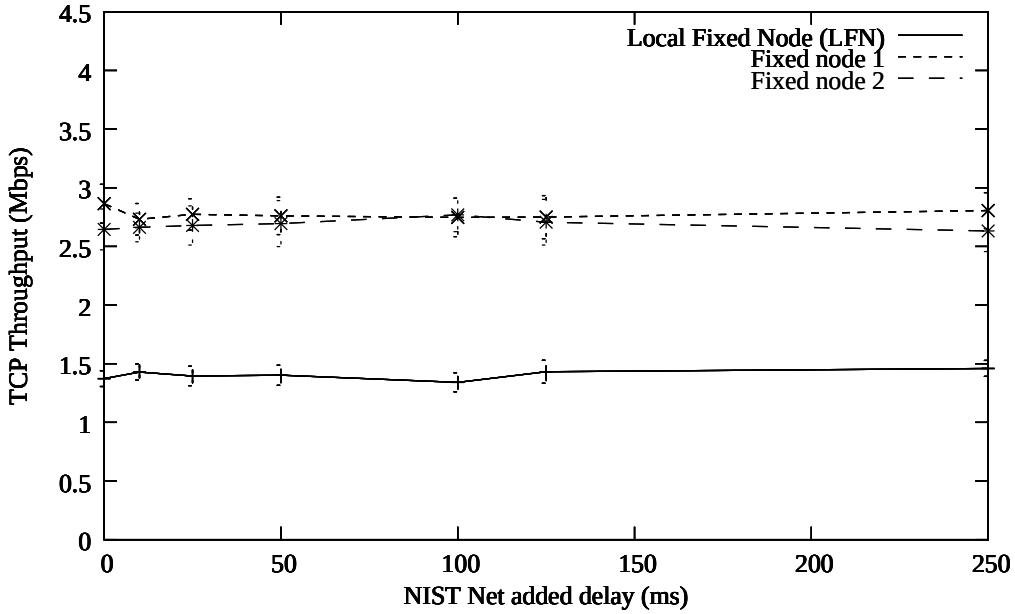

Figure 5.7: Impact of MIRON on the TCP throughput.

close each other) and 250ms (this value represents a high, but still common RTT value of 500ms in the Internet nowadays). *Delay1* is part of the CN-HA and HA-MR delays, thus affects the overall delay in the CN-HA-MR-LFN path followed by packets of the CN-LFN communication. Results for the case of using the NEMO Basic Support protocol are shown in Figure 5.6. Results for the case of using MIRON are shown in Figure 5.7. Confidence limits (95%) are also shown in both figures.

When the NEMO Basic Support protocol is used, the effect of a higher value of *delay1* in the performance of TCP application is clear: the effective throughput decreases as the delay increases (Fig 5.6). The LFN would obtain a much higher effective throughput if it was connected directly to the foreign network instead of the NEMO. This difference in the throughput increases with the delay in the CN-HA-MR-LFN path. Therefore, nodes of a mobile network located way from its home network and/or from the CN they are communicating with, would obtain extremely low TCP throughput when competing with other TCP flows, because of the suboptimal path introduced by the NEMO Basic Support protocol. Even for a value of *delay1* equal to 0ms the throughput obtained by the MNN is almost a half of the one obtained by the non-mobile nodes. Although *delay1* is 0ms, the RTT between CN and MNN is bigger than the RTT between CN and fixed nodes, because the path is not direct and there are more hops, and this difference, even though small, has an important effect on the TCP fairness. Moreover, there exists a difference in the PMTU because of the overhead that also has an influence in the TCP performance.

If MIRON is used, the performance improvement is substantial (see Figure 5.7). The TCP throughput remains constant despite the value of *delay1*. This result is as expected, because with MIRON data packets do not follow the CN-HA-MR-LFN sub-optimal path, but the direct CN-MR-LFN path. Part of the difference in the TCP throughput of the fixed nodes and the LFN is due to the packet overhead (MIRON introduces a 24-byte per packet

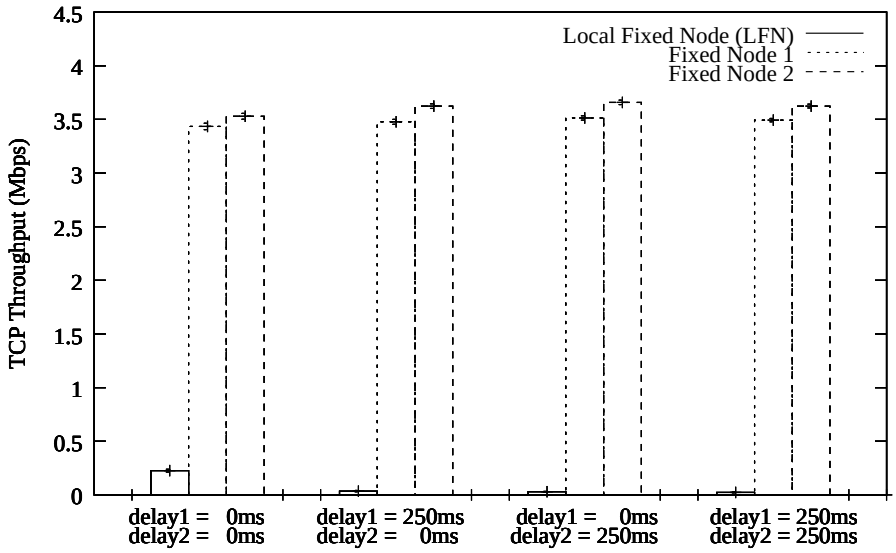

Figure 5.8: Impact of NEMO Basic Support protocol on the TCP throughput in a 2-level nested Mobile Network.

overhead, because of the Routing Header type 2 and the Home Address destination option). The performance of the MIRON prototype used during the tests (completely implemented in user space) may also have something to do with the obtained difference, although this difference could be reduced by improving the implementation (e.g., by implementing it in kernel space, or at least those tasks that have an strong impact in the overall performance).

For the nested scenario (Figure 5.5(b)), besides evaluating the effect of the varying *delay1*, that is, the delay of the path CN-HA-MR-LFN, a second adjustable delay – *delay2* – was introduced between R4 and R5, allowing us to evaluate also the effect of the distance between the home networks of two different mobile networks that are nested. Figure 5.8 shows the obtained throughput results for the NEMO Basic Support protocol and Figure 5.9 for MIRON.

As in the non-nested test (see Figure5.9), the improvement achieved by MIRON is clear. The NEMO Basic Support protocol performs worse than in the non-nested scenario, even for the null added delay case. This is because the actual RTT is bigger for the LFN than for the fixed nodes due to the longer path that packets have to traverse (CN-HA2-HA1-MR1-MR2-LFN) and the reduced PMTU. On the other hand, the performance obtained with MIRON is the same as in the non-nested scenario, as packets follow the optimal direct path and the overhead remains the same, no matter what number of nesting levels the mobile network has. As in the non-nested scenario, the TCP throughput of the LFN is lower than the one achieved by the fixed nodes because of higher RTT (packets go through more intermediate hops – MR1 and MR2) and the impact of implementing MIRON completely in user space.

5.5.2. Analytical evaluation

We have analysed how the added delay due to the suboptimal CN-HA-MR-MNN path introduced by the use of the NEMO Basic Support protocol affects the performance of TCP

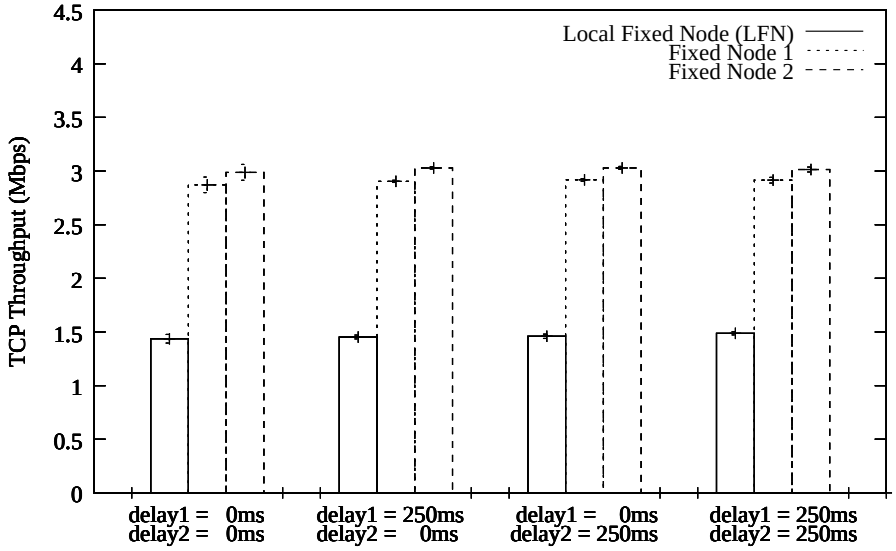

Figure 5.9: Impact of MIRON on the TCP throughput in a 2-level nested Mobile Network.

applications. In addition to the severe effect that the RTT has in the TCP performance, and the obvious effect that the delay itself has on real time applications[5], there is another effect that impacts performance: the packet overhead (and the associated PMTU reduction).

A 40-byte IPv6 header is added to every packet in the MR-HA bidirectional path due to the NEMO Basic Support protocol. Moreover, an IPv6 additional header is added per nesting level. The effect of this overhead can be negligible for non real time applications, but it can be very important for real time ones, such as VoIP applications. In order to quantitatively evaluate this effect, we analyse next the effects of the NEMO Basic Support protocol and MIRON, comparing it with plain IPv4 and IPv6, in a VoIP communication using the widely utilised Skype[6] application. Skype [BS04] uses the iLBC (internet Low Bitrate Codec) [ADA+04] codec, which is a free speech codec suitable for robust voice communication over IP. The codec is designed for narrow band speech and results in a payload bit rate of 13.33 kbps with an encoding frame length of 30 ms and 15.20 kbps with an encoding length of 20 ms.

Table 5.1 shows the packet overhead and the bandwidth consumed by a VoIP communication using UDP/RTP and the iLBC codec, for plain IPv4, plain IPv6, the NEMO Basic Support protocol and MIRON. The overhead of MIRON is less than the one introduced by the NEMO Basic Support protocol and remains constant though the number of nesting levels. The reader should notice that a nested mobile network connected to the Internet through a 64 kbps connection would not be able to support this kind of VoIP traffic (VoIP applications are expected to be very important in forthcoming 4G networks). In [dlOBC06] an additional analysis of the packet overhead in network mobility environments is presented.

[5] There are analytical studies [KT01] that state that the maximum tolerable delay in a voice communication is about 50 ms.

[6] `http://www.skype.com/`

Protocol	Bitrate (kbps)	Packet Overhead (%)
IPv4	31.2	51.28
IPv6	39.2	61.22
NEMO (without nesting)	55.2	72.46
NEMO (2 nesting levels)	71.2	78.65
NEMO (3 nesting levels)	87.2	82.57
MIRON (without nesting)	48.8	68.85
MIRON (2 nesting levels)	48.8	68.85
MIRON (3 nesting levels)	48.8	68.85

Table 5.1: iLBC bitrates and packet overhead (20ms encoding length).

5.5.3. Security considerations

From the security point of view, allowing the MR to perform some operations on behalf of the LFNs attached to it (i.e. Proxy-MR operation) does not introduce any threat, because LFNs trust their MR for the routing of all their traffic. From the architectural point of view, the solution is also natural, as the Route Optimisation support defined by Mobile IPv6 [JPA04] conceptually could be implemented in multiple boxes. MIRON just applies this mechanism, by dividing the functionalities among two different physical boxes, but actually the conceptual basis of the solution is the same as the one defined in RFC 3775 [JPA04].

It may be argued that an attacker may induce the MR to initiate the Route Optimisation procedure with a large number of CNs at the same time, by sending to an LFN of the NEMO a spoofed IP packet (e.g., ping or TCP SYN packet) that appears to come from a new CN. MIRON shares this and others vulnerabilities of Mobile IPv6 [NAA+05], but the solutions proposed to mitigate these attacks in [NAA+05] are also applicable to MIRON. For example, to avoid bringing down the MR by making it send unnecessary Binding Updates (after performing the complete Return Routability procedure), the Mobile Router should apply some local policies [NAA+05], such as:

- Setting a limit on the amount of resources (i.e. processing time, memory, and communications bandwidth) that it uses for the *Proxy-MR* functionality. In this way, when the limit is exceeded, the MR may decide to stop initiating the RO procedure for new CN-LFN communications, following the plain NEMO Basic Support protocol operation for these ones.
- The MR may also recognise addresses with which an LFN had meaningful communication in the past and only start the RO procedure for those addresses.

[NAA+05] proposes additional mechanisms for a Mobile Node to avoid attacks regarding to Route Optimisation. Most of them may also be considered by a Mobile Router implementation that provides MIRON capabilities.

5.5.4. Scalability considerations

MIRON requires some additional operations to be performed in the MR. This section briefly analyses the scalability of MIRON and provides some implementation considerations

to ensure a scalable deployment.

Basically, there are three different aspects that may affect to the scalability of MIRON:

- *Signalling load.* In order to optimise a CN-LFN flow, the MR has to perform the MIPv6 RO signalling with the CN on behalf of the LFN. This signalling grows linearly with the number of CN-LFN pairs being optimised. Similarly, to optimise the traffic of a VMN or a nested NEMO, the PANA and DHCPv6 signalling also grow linearly with the number of VMNs/MRs. This linearity is important, since it makes the required resources in a MR proportional to the size of the NEMO and it seems natural to expect MRs of large mobile networks (such the ones deployed in trains) to be powerful enough and not be resource-constrained. On the other hand, resource-limited devices, such as cellular phones and PDAs are not expected to be the MR of networks with more than a few attached nodes.

- *Memory consumption at the MR.* MIRON needs some additional information to be stored at the MR, such as host routes, extended Binding Cache entries (since state information regarding each LFN-CN optimised pair is required), and information about delegated addresses. The required memory to store a host route, a binding entry or the information about a delegated address is relatively small and grows linearly with the number of mobile nodes (i.e. VMNs and MRs) being optimised and LFN-CN route optimised pairs.

- *Processing load at the MR.* MIRON requires the MR to perform some additional operations: inspection of every packet, special handling (that is, removal of the Routing Header in the CN to LFN direction and addition of the Home Address Destination Option in the LFN to CN direction) of route optimised packets and source routing. Regarding packet inspection, MIRON just needs to look at the source and destination addresses of every packet to track LFN-CN flows and also to certain IPv6 headers to detect new arrived VMNs/MRs attached to the NEMO, so this inspection is quite similar to the normal inspection that a router does. Even if some local policies are implemented at the MR to enable smarter decisions about whether a certain flow should be optimised or not, requiring the MR to look also at other fields in a packet (such as transport headers), this inspection is not much different than the inspection than typical firewall software does in an border (access) router. Besides, the amount of traffic being processed by a MR is in general related to the size of the NEMO, so the same reasoning about the size of the NEMO and the resources of its MR also applies here.

  The special packet handling is performed by MIRON only to packets that belong to an LFN-CN communication that is being route optimised. Therefore, neither the optimised packets from VMNs or MRs, nor the packets of communications that are not being optimised, require such special packet handling. This special packet handling adds some delay in the packet processing time that depends on the MR capabilities and how this processing is implemented.

  Finally, source routing at the MR is needed to avoid route optimised packets to be forwarded through the MRHA bidirectional tunnel (instead of following the optimised direct path). Therefore, MIRON requires a different routing table per LFN that has traffic being optimised. Each of these routing tables has an entry per each CN the LFN

is communicating with. Therefore, the amount of routing entries grows linearly with the number of different LFN-CN pairs being route optimised and it is independent of the nesting level.

We can conclude that MIRON required resources grow linearly with the number of optimisations being performed, independently of the nesting level. This allows practical deployments, since it is natural to expect that the capabilities and resources of a MR to be proportional to the size of the managed NEMO. Besides, a limit on the amount of resources (memory, processing power, etc) used by MIRON can always be set, so the MR may stop starting new RO operations when that limit is exceeded.

5.6. Comparison with previous work

This section compares two of the most well-known Route Optimisation for NEMO proposals with MIRON, in terms of performance, signalling load and complexity.

As we described in section 2.4.1.1, authors of [LJP03] propose to allow the Mobile Router directly to inform the CN about the location of the MNP (using the Prefix Scope Binding Update, PSBU) [EMU03]. We will next compare this proposal (hereafter BU for Network Prefixes) with MIRON. In particular, we will consider the benefits and the costs associated with each one of them. With respect to the costs, the main difference concerns the deployment effort associated with the different proposals. MIRON, as we have already mentioned, uses the existent MIPv6 protocol unchanged. This means that the deployment of MIRON only implies modifications to the MRs. CNs do not need any upgrade since they do not require any MIRON-specific mechanism. On the other hand, BU for Network Prefixes requires not only upgrading the MRs but also upgrading all the potential correspondent nodes, i.e. all the nodes in the Internet. This is a huge deployment cost, which may not be worth depending on the resulting benefits, which will be considered next.

The benefit resulting from the adoption of any of the proposals is the optimised path through which packets are routed between the MR and the CN. However, the approach based on BU for Network Prefixes requires less signalling than MIRON. We will next quantify the difference in order to evaluate if this overhead reduction can justify the deployment cost previously identified. Consider a moving network with N MNNs. Suppose that each MNN communicates simultaneously with M CNs in average. This means that with MIRON, $N \times M$ Binding Updates messages will be required to optimise these communications. On the other hand, if the approach based on BU for network prefixes is used, the number of BU required depends only of the number of different CNs that are communicating with at least one MNN. This is so, because the BU message refers to the whole MNP, implying that if two or more MNNs are communicating with the same CN, only one BU message is needed. The net benefit resulting from the adoption of BU for Network Prefixes with respect to MIRON is a reduction in the amount of BU messages required proportional to the number of MNNs that are simultaneously communicating with a common CN. It should be noted that this only applies for those CNs that do not belong to the Home Network, since those nodes residing in the Home Network already benefit from a direct routing with the mobile network thanks to NEMO Basic Support protocol. So, the benefits provided by an approach based on BU for Network Prefixes heavily depend on the expected number of MNNs that will communicate

with a common CN outside the Home Network. The costs, on the other hand, are objective and account for the upgrading of all the nodes of the Internet to support the new option. MIRON, on the other hand, is compatible with standard Mobile IPv6 CNs.

The NEMO Basic Support protocol when applied to the case of nested mobile networks is quite inefficient as was mentioned in section 2.4.2. The Reverse Routing Header (RRH) mechanism [TM04a] proposes a solution to alleviate these inefficiencies. The proposal requires modifications in MRs and HAs, but not in LFNs, VMNs, or CNs. Besides, this proposal requires the use of Tree Discovery [TM04b] to allow the MRs to find out the level of hierarchy in the nesting.

RRH introduces an overhead (see section 2.4.2.1 for details of the mechanism operation) that can be quantified in one IPv6 header plus one routing header plus one IPv6 address per level of nesting of the Mobile Network, i.e. $(40 + 8 + n \times 16)$ bytes $= (48 + n \times 16)$ bytes, where n is the number of levels in the nesting (at least 2). This overhead is required in all the packets that go to and from the mobile network. It could be eliminated from some packets in the way out of the mobile network only at some cost in functionality (ability to detect changes in the nesting) and security. Notice that the solution of MIRON for nested mobile networks only requires the 40 bytes of the tunnelling and even that is avoided when an end-to-end optimisation of the path between the mobile network and the CN is used.

The additional need for using Tree Discovery [TM04b] implies changes in MRs and routers included in the nesting, because Router Advertisements must support the functionality of Tree Discovery. This also implies an overhead in signalling because Router Advertisements in the nesting must have a minimum of 32 bytes more than normal Router Advertisements. This must be compared with the signalling load required to distribute topological valid addresses to MRs in MIRON.

Based on the previous analysis, we can conclude that MIRON provides better Route Optimisation support than the two chosen proposals for comparison: PSBU and RRH. Besides, MIRON does not require to change any node but the MR, and in most cases requires less signalling load to optimise traffic.

5.7. A long term approach: secure delegation-based RO mechanisms

MIRON is a Route Optimisation solution for NEMO that has been designed with a very strong requirement in mind: not to impose any change on the operation of any node of the Internet (i.e. CNs) or any node attached to the Mobile Network (that is, MNNs). This is so in order to enable an easy deployment of the solution. However, it may be argued that there are certain scenarios that could benefit from different NEMO RO mechanisms (e.g., those scenarios that require stronger security guarantees or need to limit even more the signalling load required by the solution). Because of that, in this section a brief discussion about other alternative approaches is provided, based on the secure delegation of the signalling rights [NA03] to the Mobile Router.

There are several good reasons to let the Mobile Router in a NEMO send the signalling on behalf of the MNNs belonging to that NEMO (that is, whenever it is needed to optimise a MNN-CN communication flow, the MR sends a Binding Update to the CN, binding the

MNN's HoA to the MR's CoA). One of these reasons is the reduction of the signalling overhead within the NEMO, since it is the MR the one that manages the Route Optimisation and therefore MNNs no longer send any signalling regarding NEMO RO.

Although MIRON partially follows this kind of approach, there are several scenarios in which a different mechanism may be needed. For instance, it is known that in different contexts there have been doubts about the goodness of the MIPv6 Return Routability mechanism[7], and therefore it may be necessary to think of different approaches to the NEMO RO problem (compared to the MIRON solution). Future Route Optimisation mechanisms may take advantage from introducing changes on the operation of Correspondent Nodes and/or Mobile Network Nodes. For instance, this may enable the use of strong cryptography mechanisms to provide Route Optimisation support for NEMO.

A strong cryptography approach to protect Binding Updates must be based on a security association between the two nodes participating in the communication (i.e. MNN and CN). When signalling messages (e.g., Binding Updates) are sent, the problem is then how to efficiently create a security association between these nodes. Some solutions have already been proposed to solve that in Mobile IPv6 for host mobility scenarios. We consider the following important solutions: solutions based on the availability of a Public Key Infrastructure (PKI), solutions based on the use of Cryptographically Generated Addresses (CGAs) [Aur05], [Aur03], [AVH06] and solutions based on Crypto Based Host Identifiers (CBHIs) [vB04].

Letting the MR send the location update signalling on behalf of the MNNs has some advantages (such as a reduction of the signalling overhead). In MIRON, the MR behaves as a Proxy-MR for the RO signalling of the LFNs. However, in order to enable the MR to send also the signalling on behalf of LMNs and VMNs, a delegation of the signalling rights to the MR is needed. That is, some procedure must be carried out to allow the MR to send signalling messages on behalf of MNNs, in a way that enables the CN to verify that the MR is actually allowed to send this signalling.

Next, different approaches to the secure delegation of the signalling rights are explored.

5.7.1. Delegation based on PKI certificates

As a first approach, the delegation may be expressed in the form of certificates generated by a PKI. This general concept can be easily adapted to be used in NEMO. Basically, the PKI assigns prefix certificates to MRs, binding a MR public key to a NEMO Mobile Network Prefix.

$$CERT = [MNP, K_{MR}+]_{K_{CA}-}$$

Basically, the certificate states that the MR owning the public key K_{MR+} is authorised to bind a CoA to a HoA with network prefix MNP. This certificate is signed by a Certification Authority (CA).

[7] Many mobile operators seem to be reluctant to use a solution based on Return Routability as compared to "strong cryptography" to protect the location information updates (i.e. Binding Updates sent to Correspondent Nodes) in their Mobile IPv6 deployments. Essentially RR is considered a "weak security mechanism" and it is accused of introducing a non-negligible burden of signalling in the network, which is a relevant handicap in links where resources are scarce (i.e. the wireless access link from a NEMO to the infrastructure).

5.7.1.1. Procedure of operation

- The MR obtains a certificate from the PKI, containing the Mobile Network Prefixes associated to the MR.
- Each Binding Update sent by the MR to a CN on behalf of a MNN is signed with the MR's private key. The message also contains the MR's prefix certificate.
- The Correspondent Node, when receiving a Binding Update, obtains the prefix certificate associated with the HoA contained in the BU, and verifies it. If the Binding Update is valid, the CN adds an entry in its Binding Cache.

5.7.1.2. Analysis of the solution

In this approach, a high protection against identity attacks is provided, but the major drawback of this solution is the requirement of a global key infrastructure, which is an unrealistic requirement for the whole Internet nowadays (although it is a solution feasible in more restricted environments).

Using prefix certificates introduces the non-trivial issue of the Prefix Ownership and this problem is much more complex than the basic Address Ownership issue that arises with Mobile IP.

5.7.2. Delegation based on self-signed certificates

In this case, the MNN is assumed to have a Cryptographically Generated Address (CGA) as its HoA. As described in [Aur05], a CGA[8] is an IPv6 address, which contains a set of bits generated by hashing the IPv6 address owner's public key. This property allows the user to provide a "proof of ownership" of its IPv6 address. On the other hand the MR (i.e. delegate) has a certificate as follows:

$$CERT = [CGA, K_{MR+}]_{K_{MNN-}}$$

Basically, the certificate states that the MR owning the public key K_{MR+} is authorised to bind a CoA to the CGA (MNN's HoA) included in the certificate. In other words, the MNN, identified by the CGA, delegates the right to send Binding Updates (location update messages) to a trusted node, the delegate, identified by K_{MR+}. This certificate is signed with the MNN's private key associated to the CGA (K_{MNN-}).

5.7.2.1. Procedure of operation

In this scenario, whenever a MNN-CN RO is needed, the MR performs it on behalf of the MNN and sends to the CN a location update message (BU) linking the MNN's HoA to a CoA. The process is the following: the Binding Update is signed with the MR's private key and it includes the certificate. When the CN receives this location update message, it first verifies the certificate using the MNN's public key associated to the CGA (HoA of the BU) and then it verifies the received message using the MR's public key (K_{MR+}), included in the certificate.

[8] A more elaborated description of CGAs is provided in Section 6.3.2.1.

5.7.2.2. Analysis of the solution

The main advantages of this approach are the following:

- It does not require the deployment of a PKI infrastructure. This is a crucial point because assuming the availability of a global PKI infrastructure is not very realistic in large networks (e.g., Internet), at least nowadays.
- On the other hand, it would be potentially compatible with SEND [AKZN05].

Also, some drawbacks can be pointed out:

- The solution is not transparent for the CN, since any CN must understand the address format and the procedures involved, which requires changes on the software of the CN.

5.7.3. Implicit Delegation

In this approach to the delegation of signalling rights, there is not explicit delegation from the Mobile Network Node to the Mobile Router. Instead, the MNN gives to the MR the right to send signalling on its behalf by accepting the use of an address with a particular structure (the address format is proposed in [vB04]).

5.7.3.1. Address format

This address (an IPv6 address) is composed of the network prefix (64 bits) and the Interface Identifier (64 bits). The network prefix is simply the Mobile Network Prefix. The Interface Identifier (IID) is called a Crypto Based Host Identifier [vB04] and is created in the following way:

$$IID = [4\ control\ bits,\ 48\ bit\ site\ identifier,\ 12\ host\ bits]$$

The format for this IID is proposed and described in [vB04]. The 4 control bits are: one reserved, one to distinguish between 80 bit identifiers and 64 bit identifiers (in this application we are only interested in 64 bit identifiers), and the usual universal/local bit and group bit. To ensure EUI-64 compatibility, [vB04] proposes to set the u/l bit to ”universal” and the group bit to indicate a group address. Because we have 12 host bits, we will be able to address $2^{12} = 4096$ hosts, which seems to be large enough for a NEMO.

The site identifier contains cryptographic information that allows Correspondent Nodes to verify that the address is used legitimately. The site identifier must contain the following information (this is different from what is proposed in [vB04] because of the reasons explained in next section):

$$NEMO\ Site\ identifier = Hash(MNP, K_{MR+})$$

5.7.3.2. Procedure of operation

A MR willing to serve a NEMO by sending the signalling on behalf of its MNNs, must generate a pair of keys: public/private. Then, it generates and provides addresses (Home Addresses or HoAs) to the MNNs. The addresses have the format explained in the previous section.

If the MR wants to send a Binding Update on behalf of a MNN of its NEMO to a Correspondent Node, the MR signs the BU with its own private key. The MR also informs the Correspondent Node of its public key and Mobile Network Prefix (that must match the one of the HoA included in the BU).

The CN can verify the address by re-calculating the Site Identifier (it has the MR public key and the NEMO Mobile Network Prefix) and checking that it matches that of the HoA. Using the MR public key, the CN can also verify the authenticity of the BU.

An attacker cannot generate a fake BU that binds a certain HoA to a CoA. To be able to do that, the attacker would need to authenticate the BU with a private key that corresponds to the public key used to create the Site Identifier of the HoA.

An attacker can also try to generate pairs of public/private keys and create a dictionary of 2^{48} different Site Identifiers. Then, if the attacker detects a particular HoA that she wants to attack, she only has to look up the public/private key corresponding to the Site Identifier of the HoA in the dictionary. Using the Mobile Network Prefix in the calculation of the Site Identifier makes this attack much more difficult, because the dictionary must include not only Site Identifiers but also network prefixes: $2^{48} \times 2^{64}$ entries.

Notice that in this section we focused on the conceptual ideas of this mechanism, a practical solution would use some improvements, for example a symmetric key could be generated from the public/private key for doing authentications less computationally costly. Also, the particular hash algorithm or public key cipher method are not analysed.

5.7.3.3. Analysis of the solution

The main advantage of this delegation solution is that is very simple. Nevertheless some disadvantages can be pointed out:

1. It is incompatible with stateless address autoconfiguration and other solutions that work with the IID as CGAs (what can have a negative effect in SEND for example).

2. The solution is not transparent for the CN, i.e. any CN must understand the address format and the procedures involved, which requires changes to the software of the CN.

3. It imposes a limit of 2^{12} to the number of hosts in a NEMO. This does not seem to be a great problem for a NEMO. A solution for this limitation would be to use more than one prefix in the NEMO (this, of course, uses address space).

5.7.4. Secure-delegation of signalling rights: summary and final remarks

In this section, we have analysed the need of a delegation of the signalling rights in those environments in which a Route Optimisation for NEMO using strong cryptography is required.

The delegation of signalling rights can be done in an explicit way, by means of authorisation certificates, or, as it has been devised here, in an implicit way, accepting the use of an address with some particular characteristics.

Likely, the simplest solution, implicit delegation, has also several limitations (as incompatibilities with other mechanisms like SEND or stateless address auto-configuration). The most flexible solution, the one based on PKI certificates, requires an important infrastructure. The solution based on CGAs can be a good compromise between complexity and flexibility.

5.8. Conclusions

The NEMO Basic Support protocol [DWPT05] enables whole networks to move and change their point of attachment, transparently to the nodes of the network. This solution introduces some limitations and problems in terms of performance (increased delay in packet delivery and packet overhead, decrease in available PMTU, the HA becoming a bottleneck, etc). To overcome these limitations we have designed and implemented a Route Optimisation solution: MIRON, that enables direct path communication between a node of the mobile network – supporting any kind of node, with and without mobility capabilities – and a Correspondent Node.

MIRON has two modes of operation: the MR performing all the Route Optimisation tasks on behalf of those nodes that are not mobility capable – thus working as *Proxy MR* [NZWT06] – and an additional mechanism, based on PANA and DHCP, enabling mobility-capable nodes (i.e. Mobile Nodes attached to a NEMO) and routers (i.e. nested Mobile Routers) that actually have mobility and Route Optimisation capabilities to manage their own Route Optimisation.

To validate the design of the solution and evaluate the actual performance of it, a prototype of MIRON was implemented in Linux. The NEMO Basic Support was also implemented so we could compare the results obtained with MIRON with the basic solution for network mobility. Tests involving TCP applications showed that the increased RTT perceived by the nodes of a NEMO (due to the suboptimal path followed by packets) has a severe impact on the performance (in terms of effective throughput, when sharing some link with traffic from other active non-mobile TCP nodes). This effect is exacerbated when NEMOs are nested. On the other hand, the same tests conducted with MIRON showed a better performance, by obtaining much higher effective TCP throughput than in the case of the NEMO Basic Support, also in the case of nested networks.

The effect of packet overhead was described by means of a quantitative analytical study of the overhead that several protocols add to packets belonging to a VoIP application, such as Skype. These results show that the packet overhead introduced by the NEMO Basic Support protocol is significant for this kind of application, specially when there is nesting.

We could also think about a long term NEMO Route Optimisation solutions that could be developed without the constraints of keeping CNs (i.e. any potential peer that a node in the mobile network may have) unmodified. An interesting approach along this line is the secure delegation of the signalling rights to the Mobile Router. Three different solutions based on this scheme have been proposed. We think that there is the need of working in the design of NEMO RO solutions that, by taking advantage of introducing some changes on Correspondent Nodes and/or Mobile Network Nodes, could be more efficient than MIRON.

But both types of solutions will coexist, because a large installed base of legacy nodes will require a solution like MIRON.

In conclusion, this chapter proposes a Route Optimisation for NEMO solution (MIRON), that provides significant performance improvements over the NEMO Basic Support protocol and that is implemented only modifying the software in the MRs. LFNs, VMNs, or CNs do not need to be modified for MIRON to work, which facilitates the deployment of the solution. The validity of the solution has been proven by making experiments and tests with an implementation for Linux.

We want to highlight that MIRON is cited in a document of the IETF NEMO Working Group that analyses the NEMO Route Optimisation solution space [NZWT06], considering MIRON as one of the reference solutions.