

Observations on GeoNetworking

Document HTG1&3-3

Version: 2012-11-12

EU-US ITS Task Force
Standards Harmonization Working Group
Harmonization Task Groups 1 & 3



U.S. Department of Transportation



Table of Contents

1 INTRODUCTION..... 3

2 GEOGRAPHICALLY SCOPED INFORMATION DISSEMINATION 3

3 OVERVIEW OF THE CURRENT GEONETWORKING PROTOCOL 5

4 RF CHANNEL CONGESTION ISSUES 7

5 SECURITY-RELATED ISSUES 8

5.1 Denial of service attacks.....8

5.2 Cryptographic processing9

5.3 Layering issues and security10

5.4 Authorization10

5.5 Privacy11

6 POSSIBLE IMPROVEMENTS..... 11

6.1 Congestion mitigation11

6.2 Privacy, authentication and other security related issues12

7 APPLICABILITY OF GEONETWORKING 14

8 CONCLUSIONS 15

9 REFERENCES..... 16

1 Introduction

This document is intended as an introduction to the concept of geodissemination of information and a preliminary analysis of an implementation thereof. Geodissemination of information was investigated and initial concepts were developed in the EU research project GeoNet (<http://www.geonet-project.eu>). This document contains analysis of technical features of geodissemination, including the GeoNetworking protocol being standardized at ETSI TC ITS (in WG3). Since the US is currently not investigating geodissemination, a divergence analysis was not possible. Consequently, the output of this effort was not consistent with the contents of the other HTG documents which focus on divergences between and gaps among existing standards. Thus, the results and conclusions of this analysis were produced as a separate informative report. This report also contains recommendations for how to achieve harmonized implementations of geodissemination of information in the future.

Geodissemination of information is not currently under consideration in the US; however, within ETSI (cf. [1],[2],[3],[4],[5],[6],[7]), GeoNetworking is a candidate for mandatory implementation. This is potentially a significant obstacle to future harmonization of ITS communication standards.

2 Geographically scoped information dissemination

A number of use cases in ITS communications involve the dissemination of information in a particular geographic region. This ITS requirement has led to the development of the concepts of location-based addressing and routing of data (packets). In this document, "geodissemination" is used to refer to the basic functionality of dissemination of information within a prescribed geographic area, while the term "GeoNetworking" is used to refer to the ITS station networking and transport layer protocols specified in the standards currently being developed at ETSI (cf. [1],[2],[3],[4],[5],[6],[7]).

As an example of geodissemination, Figure 1 (cf. D1.2 "Final GeoNet Architecture Design," June 2010, of the GeoNet project, cf. <http://www.geonet-project.eu>) illustrates a scenario where road traffic hazard information (black ice) is being transmitted to all vehicles located in targeted geographic areas.

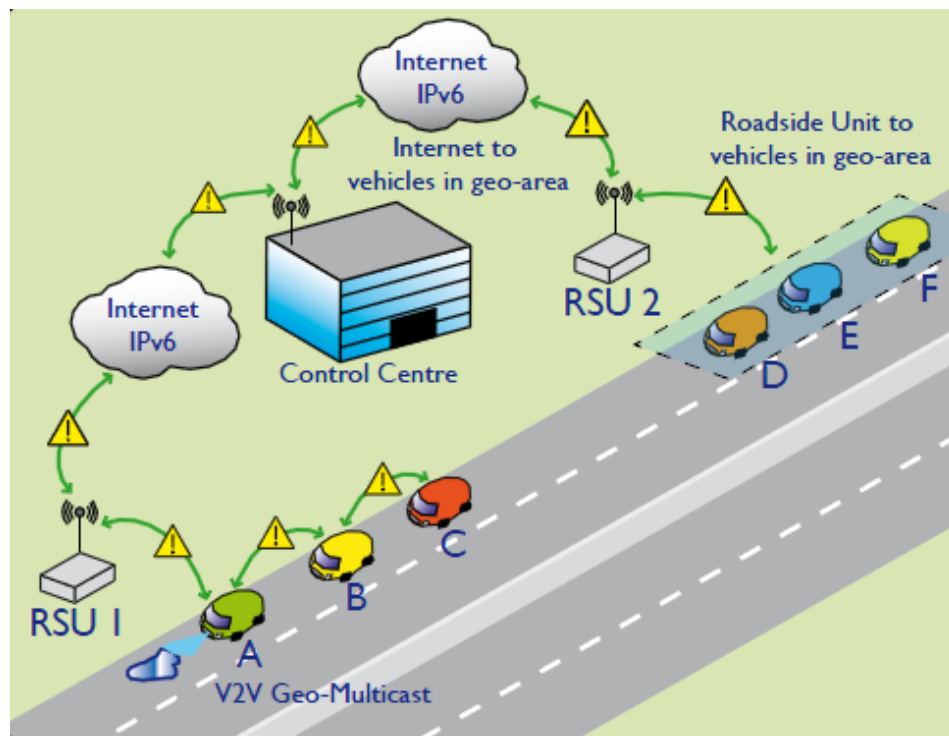


Figure 1: Geodissemination of road traffic hazard information

In this example, there are two targeted areas, one being the section of roadway in the immediate vicinity of the hazard (black ice) and the second being a section of roadway used by vehicles approaching the hazard (e.g. to allow vehicles the opportunity to choose another road). Local dissemination of data is generally performed over short-range wireless technologies (e.g. 5.9 GHz), and internet and/or cellular communications may be used for longer range dissemination. In Figure 1, RSU1 connects to a Control Centre and subsequently to RSU2 using fixed infrastructure (e.g., the Internet). RSU2 is chosen in this example because it is in the appropriate location to disseminate the hazard warning to vehicles within its RF coverage and thereby topologically scope the distribution of the information to only potentially interested recipients. Note that geodissemination of information could also be performed using other wireless technologies (e.g. cellular).

In the example of Figure 1, vehicle A detects black ice and notes its position. It sends a geographically scoped warning message destined for any and all vehicles in the area nearby. The message is received by vehicle B which then forwards the warning to vehicle C, and so on, until either the message reaches the boundary of the specified target geographic area, or there is no vehicle present to forward it further. The message is also given an expiry time by the originator, and is forwarded within the target area until it expires or is cancelled, or no vehicle remains within the target area. Vehicle B will therefore forward the message back towards vehicle A as well as towards vehicle C. If another car comes into range, vehicle A will resend the message such that the new vehicle can receive it.

In the figure, vehicle A also sends a notification message to a nearby roadside unit (RSU 1), which forwards the message to the traffic hazard control center (other application models are, of course, possible). The hazard control center determines an appropriate geographic area for dissemination of a warning, and periodically dispatches the warning to RSUs serving that

area (in this case, RSU 2). RSU2 then transmits the warning (using its 5.9 GHz medium) to all cars located in the target geographic area (vehicles D, E, F). While not shown in Figure 1, message dissemination in the area near RSU 2 might also involve forwarding of the warning by the vehicles in that area. Transmission of messages from an RSU to any vehicle within range without further retransmissions is an example of MAC layer single-hop communications. Note that herein the terms MAC layer single-hop applies only to the wireless segments of a communication path.

Figure 1 illustrates three different geodissemination scenarios:

- Information sent by a vehicle to all vehicles in an immediate geographic area around the vehicle (GeoBroadcast).
- Information sent by a vehicle to a server in the Internet (IPv6 unicast/GeoUnicast) for subsequent dissemination in a geographic area.
- Information sent by a server in the Internet to all vehicles in a given geographic area (IPv6 multicast/GeoBroadcast).

Other possible communication scenarios that are not illustrated include:

- Information sent to any one vehicle located in a target area (IPv6 anycast/GeoAnycast);
- Information sent to one specific vehicle located in a given target area (IPv6 unicast/GeoUnicast)

In these scenarios, the sender and the receiver(s) could either be a vehicle, a roadside unit, or any node in the Internet, and the target area could be local (immediately surrounding the sender) or remote (reachable via other vehicles, roadside, or the Internet).

Note that it is also possible to eliminate the link between RSU1 and RSU2 and attempt to replace it with a time-varying ad hoc network of vehicles extending from vehicle A to vehicle D, thereby eliminating the IPv6 backhaul. While the probability of successful dissemination of information to the intended geographic area would thereby be (dramatically) decreased, it is nonetheless possible.

3 Overview of the Current GeoNetworking Protocol

A geodissemination protocol is currently being specified inside ETSI TC ITS WG3 as a network layer protocol ("GeoNetworking") for ITS stations that specifies mechanisms for packet forwarding in an ad hoc collection of ITS stations. Geographic locations of ITS stations are used as "addresses" for packet destinations (cf. [4],[5],[11],[12],[13]). The current set of ETSI standards ([1],[2],[3],[4],[5],[6],[7]) mandate GeoNetworking implementations in all ITS stations, and mandate its use for communications over 5.9 GHz in Europe including for the periodic transmission of safety-related (CAM/DENM) messages.

While not discussed in detail herein, GeoNetworking is also intended to support point-to-point ("unicast") communication between pairs of ITS stations based on geographical locations as well as the dissemination of packets in geographical areas.

Figure 2 illustrates the location of GeoNetworking headers within a MAC sub-layer protocol data unit (MPDU). MAC addresses are used to address peer stations either in broadcast mode or in unicast mode and an LLC header is used to direct the network layer protocol data unit (NPDU – the PDU exchanged between the LLC sub-layer and the higher layer, in this case the GeoNetworking layer) to the appropriate networking protocol. In Figure 2, FPDU stands for “facilities layer protocol data unit” and NSDU stands for “network layer service data unit”.

MAC Header	LLC Header	GeoNetworking Header	GeoNetworking Security Header (optional)	FPDU (=NSDU) (optional)
-------------------	-------------------	-----------------------------	---	--------------------------------

Figure 2: MAC/LLC/GeoNetworking frame structure

The GeoNetworking header is comprised of a common header and an extended header as shown in Figure 3 (cf. ETSI draft TS / EN 302 636-4-1).

Common Header	Extended Header
----------------------	------------------------

Figure 3: GeoNetworking header structure

The common header is 36 octets in length (cf. clause 8 in [4]) and contains the geographical location (24 octets) of the sender of the packet. The contents of the extended header depends on the functionality (i.e., GeoUnicast, GeoAnycast, GeoBroadcast, etc.) and as an example, the GeoUnicast extended header is 52 octets in length, which results in a GeoNetworking header that is 88 octets in length. The GeoNetworking standards also distinguish between "single-hop" and "multi-hop" packets; single-hop packets include the BEACON and SHB, and multi-hop packets include GeoUnicast, TSB, GeoBroadcast, GeoAnycast, LS Request, and LS Reply. Herein, single-hop and multi-hop refer to the sending of these packets.

The optional GeoNetworking security header is currently unspecified, and as such, its precise impact on the overhead and security processing required are difficult to assess. Issues in security are discussed later in this document.

Each GeoNetworking router maintains a neighbor location table containing time-stamped address, position, and speed for ITS stations (routers) in its vicinity. GeoNetworking forwarding algorithms use this neighbor location table to make forwarding decisions. Note that this table contains information about nearby ITS stations that is also found in the Local Dynamic Map (LDM) and its associated neighbor location table built from CAM/DENM messages received from nearby ITS stations.

Since all datagrams sent from upper layers have the common header, all received datagrams can be used to build the neighbor location tables. According to the GeoNetworking protocol, when there is no higher layer activity, datagrams containing only the common header (network-layer beacons) must be sent "periodically."¹

All GeoBroadcast messages that are received by a router are cached (unless the maximum hop count is exceeded) and must be forwarded to any and all new routers that show up in the neighbor location table. Cached messages are purged when their expiration time is exceeded or when they are canceled. If a message has not expired or been canceled, a router that has a cached copy must forward it to new routers. Current ETSI draft standards propose mechanisms to reduce the potential for flooding due to this basic algorithm.

4 RF Channel Congestion Issues

When considering safety of life and property applications that exchange information over capacity-constrained RF channels, channel congestion can become a critical issue. The current draft standards for cooperative awareness applications specify that each vehicular ITS station broadcasts a basic safety/cooperative awareness message (BSM/CAM) on a 10 MHz safety channel on the order of 10 times per second. However, preliminary field tests in the US, the results of which are intended to be made publically available in the near future, indicate that this 10 MHz RF channel has insufficient capacity to serve a significant number of such vehicular ITS stations within several hundred meters of each other. While this problem could be mitigated somewhat by increasing the channel capacity (e.g. by changing to 20 MHz channel bandwidths), current tests are being conducted only with the 10 MHz channels currently specified in EU and US regulations.

As currently specified (cf.[1],[2],[3],[4],[5],[6],[7]), GeoNetworking introduces the following channel congestion concerns.

- 1) The GeoNetworking header contains source position information which is also a key data element of the BSM/CAM message itself. Since BSM/CAM messages are single-hop broadcast messages, there is no networking (forwarding or routing) to be performed on the BSM/CAM packets (FPDUs). This means that the entire contents of GeoNetworking header are superfluous. Elimination of the GeoNetworking would result in a savings of 36 to 88 octets depending on the header type used. Such a reduction is significant considering the BSM/CAM message itself is on the order of a few hundred octets.²
- 2) Forwarded messages run the risk of flooding the channel with information that is redundant from the point of view of the receiver. Flooding may occur because a copy of a message is forwarded by two or more separate nodes, potentially more than

¹ In many cases, it will not be necessary to send beacons as units will in general send CAMs; however, units that do not send CAMs are required to send beacons.

² Tests currently being conducted in the US involve an approximately 285-octet BSM, roughly half of which is security overhead.

doubling the channel capacity consumed by that message. This may of course be mitigated by duplicate detection mechanisms at the network layer as currently specified within the GeoNetworking protocols. However, network layer mechanisms cannot prevent duplication and flooding if two vehicles both report the same incident, for example the same patch of black ice. In this case only mechanisms at the application or facilities layer can detect whether a message contains duplication of already known information. Without this application-layer involvement, redundant messages will continue to be forwarded, reducing channel capacity for messages that contain new information.

5 Security-Related Issues

5.1 Denial of service attacks

It is well known that there are many DoS attacks for MAC-layer single-hop communications, from simple jamming to more sophisticated MAC attacks, including overloading the crypto-processing capacity of nodes with fake packets. Since these are attacks below the networking layer, protection of the GeoNetworking protocol does not mitigate any of these attacks. However, an unprotected GeoNetworking protocol could lead to severe channel flooding attacks because of the store-and-forward feature of GeoNetworking. A single attacker can cause congestion and therefore denial of service in larger areas, including areas beyond its own radio range.

Possible channel-flooding attacks include:

- 1) **Large-area, long-lifetime forwarding:** In large-area forwarding, a sender generates a message with a target zone that is extremely large, extremely far away, or both, and the lifetime of the message is very long. This causes the message to continue to be forwarded in a large area for a long time, consuming channel capacity.
- 2) **Many-message forwarding:** In many-message forwarding, an attacker takes advantage of the fact that all cached messages must be forwarded to any new node that appears in the neighbor location table. An attacker can simply broadcast a large number of messages which are received and must be rebroadcast by a number of intermediate nodes. These may be vehicles approaching an intersection. When a new node appears (such as if a new vehicle approaches the intersection on a cross-street), all of the intermediate nodes that see the new vehicle forward their duplicates of the original messages, causing significant channel congestion. The attacker may potentially also masquerade as one or more new vehicles to prompt unnecessarily frequent retransmission. Since new vehicles are identified by new Link Layer addresses (LL_ADDR, in practice, typically the MAC address) that are not in the current location table, masquerading as a new vehicle is straightforward.

Large-area, long-lifetime forwarding attacks can be mitigated, though not eliminated entirely, by the use of communication security services. If GeoNetworking target information (target location, target area, message lifetime) is required to be signed and the signature authorized by a certificate, the permissions in the certificate can be set so as to limit the distance a

message can be sent, the area it can be sent to, or the lifetime it can have. Note that in such cases, the forwarding nodes would have to verify the originator's signature on every received message.

Large-area, long-lifetime forwarding attacks could potentially also be addressed by non-cryptographic means, such as assigning each message a probability of being discarded without forwarding such that the probability increases with the age of the message and/or the distance from the origin.

Many-message forwarding attacks may be mitigated by general congestion control mechanisms at the network layer and lower to address channel congestion (e.g., duplicate packet detection), and by access control and contention management mechanisms for the internal buffer that stores messages to be forwarded. Cryptographic mechanisms also mitigate this attack, but to a limited extent: an attacker can still broadcast a large number of initial messages, but if messages are signed, the signature is checked by the GeoNetworking processing, and a new vehicle is identified by a new certificate rather than a new LL_ADDR,³ the attacker will need multiple distinct certificates to masquerade as multiple distinct vehicles.

An additional method to mitigate both of these attacks would be to allow intermediate nodes to use application-layer logic to determine whether messages should be forwarded because they are relevant and contain fresh information. Some context sensitive forwarding decisions may also be made at the network layer based on GeoNetworking headers. There is work currently ongoing on incorporating such techniques within the bodies developing the GeoNetworking standards (ETSI TC ITS WG3, et al). However, more sophisticated forwarding decisions are possible using higher layer information. For example, higher layer logic can detect messages that are no longer relevant even if going strictly by the GeoNetworking headers the message has not yet expired and should be forwarded. Higher layer logic can also be used to detect the presence of redundant information that need not be forwarded such as the same incident reported by multiple observers.

5.2 Cryptographic processing

As noted in the previous section, if GeoNetworking messages are signed by the originator with appropriate permissions given in a certificate and have their signature verified whenever they are forwarded, this can mitigate large-area long-lifetime forwarding attacks. However, there are the following concerns:

- 1) Signatures add overhead, both in terms of processing and packet size. It would be desirable to have each packet signed only once by the transmitter. In this case, the question arises of where in the stack the packet should be signed if it is to be signed only once. Section 5.3 considers issues that arise when attempting to protect network-layer information and application-layer information with a single signature in an architecture where there may be multiple network stacks and multiple applications.

³ Or if the LL_ADDR is contained in the signed payload.

- 2) Signature verification creates a performance burden on the forwarding nodes which, it can be argued, they should not have to bear, especially if they are not drawing benefit directly from the message itself. This burden could be mitigated by allowing application-level logic to make forwarding decisions based on the contents of the message. By making use of higher layer information, higher layer logic may be able to decide not to forward a message without the need to verify the signature thereon. This results in a computational savings not just for the forwarding node, but also for all "downstream" nodes. Forwarding decisions made at lower layers do not have the same ability to reduce processing loads.

5.3 Layering issues and security

Some observations with respect to the standard layered communication model and GeoNetworking include:

- 1) If application PDUs are signed at the network layer, and if different applications over GeoNetworking use different certificates (for privacy or robustness reasons), the networking layer needs to distinguish between the higher layer entities that originate packets in order to use the correct certificate to sign the packets. This is a layering violation.
- 2) If the GeoNetworking protocol signs its payloads (NSDUs) and other networking protocols do not, then to prevent signatures from being created at multiple layers, higher layers must be aware of which network stack they are using so that they *can sign* packets that are not going to be transmitted via GeoNetworking and *not sign* packets that are going to be transmitted via GeoNetworking. Any application PDU that is transmitted over multiple media and/or different networking and transport layer protocols may need to be signed multiple times within the station for every transmission.
- 3) In the case where the implementation of the ITS station is distributed and contains more than one router, application PDUs may be sent from any or all of the routers. If signing is done at the network layer (inside each router), then a signing certificate and private key need to be maintained in each router. This imposes management complexity and creates a security risk.

All of these considerations suggest that signing should be applied as high in the stack as possible. This may not be consistent with the GeoNetworking protocol's current location at the network layer because the GeoNetworking information would not be protected.

5.4 Authorization

By originating a GeoBroadcast message, a unit is instructing intermediate units to forward it to the target area. This causes the intermediate units to incur a computational burden. It is not

clear that all units should have the authority to generate a GeoBroadcast message and require other units to incur that burden. This aspect should be investigated properly.⁴

5.5 Privacy

In many instances, an ITS station (ITS-S) will be running multiple applications. In such situations, the combination of applications being run by the ITS-S could be considered Personal Identifiable Information (PII) (see HTG1-1 [15] for a further discussion of why this is the case). As such, ITS standards should provide protection against an eavesdropper discovering that two different applications are being run on the same ITS-S. The proposed GeoNetworking protocol poses a significant challenge to providing such protection.

All applications in an ITS-S using the GeoNetworking stack can be associated with that station if the ITS-S does not move a great distance between transmissions from the different applications. An eavesdropper can use the kinematic information in the GeoNetworking headers to determine that these messages have come from the same vehicle. If these messages come from different applications, the eavesdropper will have determined that the applications are running on the same vehicle. With GeoNetworking at the network layer, the only mitigation of this attack is to encrypt all PDUs at a layer that protects all application identifiers, which is impractical. Users who run multiple applications over the GeoNetworking stack will thus likely suffer a loss of privacy.⁵

6 Possible improvements

6.1 Congestion mitigation

As noted in 5.1, channel flooding attacks based on GeoNetworking may be addressed by internal and external congestion control mechanisms based on information available at the network layer and below.

Concerns about this mitigation include the following:

- 1) This mitigation is restricted to information available at the network layer and below due to GeoNetworking being located at the network layer, and it generally affects all packets being transmitted on the congested channel, not just those using GeoNetworking. If the congestion control mechanisms were able to take higher-layer information into account, more sophisticated decisions could be made.
- 2) Congestion control on its own may be used to address the large-area, long-lifetime attacks by devising a method for discarding messages. However, without cryptographic mechanisms to link a message to a sender, the same discard algorithm would have to be applied to all messages. The use of cryptographic mechanisms

⁴ Note the cost-benefit considerations here. Incorporating GeoNetworking headers and the required processing thereof causes a burden to the system compared to a system without GeoNetworking. The benefit of GeoNetworking should be commensurate with this cost, noting that reducing the number of messages that can be sent lowers the benefit.

⁵ This loss of privacy might be acceptable if the users were to affirmatively opt in to it. However, it is not clear that users understand multi-application privacy considerations with enough clarity to make this decision.

allows different senders to have different capabilities. For example, a highly trusted sender could be allowed to send a message twice as far as a less trusted sender. This suggests that it could be useful to sign the origin and destination information in a GeoNetworking packet as well as using congestion control. The next section considers how these cryptographic mechanisms could be applied.

This suggests that:

- 1) Congestion control mechanisms should be developed (this is an active work area within several SDOs and industry consortia).
- 2) Geodissemination at the facilities layer (information centric forwarding) may be a more effective way of implementing store-and-forward messaging to geographic locations.
- 3) Cryptographic mechanisms should be used to permit how far, and for how long, ITS-Ss may request messages to be sent. Different ITS-S could have different permissions.

6.2 Privacy, authentication and other security related issues

Section 5 has identified three different components of a message sent over GeoNetworking that may need to be cryptographically protected:

- Application payload
- Extended header
- Common header

The first two are generated by the originator of the message, and the last is changed at every intermediate hop of a multi-hop dissemination.

Risks associated with these components are as follows:

- There is a low risk to the system from attacks based on false data in the common header, because (in the GeoBroadcast mode) the content of the common header is generally not used to make forwarding decisions. The residual risk arises from an attacker who is able to provoke retransmissions of stored messages by appearing to be a new node.
- There is significant risk to the system if the extended headers are not authenticated, as an attacker may be able to launch a large-area, long-lifetime forwarding attack as described in section 5.
- There is significant risk to the system if the application payload is not authenticated.

These risks could obviously be addressed by signing each component individually. However, this would result in increased processing costs and packet size. It is greatly preferable to sign a packet only once before transmission to keep these costs to a minimum.

Based on this analysis, options to address security concerns include the following:

- a) Originator signs the facilities layer PDUs, GeoNetworking headers are not signed.
- b) Originator signs the facilities layer PDUs and includes the geodissemination information in the signature (cf. signing pseudoheaders).

- c) Originator signs the payload plus extended header (not including the common header) in the case of multi-hop, and signs the payload plus common header in the case of single-hop.
- d) Originator signs the payload plus extended header; all forwarders sign the common header and the immutable fields (payload plus extended header). This results in attaching two separate signatures, one for the payload and the other for the combined header and signed payload.
- e) Originator only signs the facility/application layer PDU (payload) which contains all relevant geodissemination information, and forwarding decisions are made within the facilities layer based on that relevant information and information contained in the Local Dynamic Map (LDM). Geodissemination information is not inserted at the network layer or below.

Option a) does not protect the GeoNetworking headers and therefore is unacceptable.

Option b) has the drawback of being cross-layer and having different treatment of packets for different communication media as discussed in section 5.3.

Option c) has the drawback noted above of increased complexity due to the necessity of distinguishing between single-hop and multi-hop. Additionally, it has the cross-layer drawbacks of option b).

Option d) has the drawback of increased packet sizes which is particularly problematic when this protocol is to be used over capacity-constrained media such as 10 MHz safety channels at 5.9 GHz, and especially so when these channels are used for safety of life and property.

Option e) moves the processing of geodissemination packets (and the maintenance of neighbor location tables) to the facilities layer, and in doing so allows the use of any available networking protocol for packet delivery. Depending on the details of the geodissemination protocol, there may be a need to ensure consistency of neighbor location tables on multiple hosts in distributed ITS-S implementations.

Many of the security-related issues discussed above can be mitigated by moving the geodissemination functionality/protocol higher up in the protocol stack to a place at the bottom of the ITS station facilities layer (cf. option (e) above). Herein, the implementation of geodissemination functionality in the facilities layer is referred to as "geoforwarding." The use of geoforwarding for geographic location dependent message distribution creates a single location for signing messages thereby protecting all sensitive geographic information introduced at the facilities layer and above.

An additional important benefit of moving geodissemination functionality to the facilities layer is that different network layer protocols and/or communication media could be used to transport the message to its intended destination, and it becomes possible to choose an optimum networking and access technology for each link between two stations. This includes the possibility of using legacy networks and devices that are not geoforwarding-aware, as well as using IPv6 multicasting mechanisms to accomplish more efficient and selective dissemination of messages in geographic areas. This would also hold true for relay stations

which translate messages from one set of RF parameters to another to support applications at a border crossing.

GeoNetworking, as currently specified, will increase the load on capacity limited wireless links in the communication system, and could lead to severe RF channel congestion.

In addition to reducing the sizes of packets (PPDUs) sent over the air and allowing the use of legacy networks and devices, moving geodissemination functionality to the facilities layer allows more informed decisions as to whether the message should be forwarded or not based on information available to the local application, leading to a reduction of RF channel congestion resulting from large packet sizes and multiple redundant messages.

7 Applicability of GeoNetworking

In conventional networking, an application generally neither knows nor cares about the geographic location of the physical devices with which it intends to communicate, and assumes that the network to which it is attached will handle whatever is necessary to deliver the information to the intended destination(s). On the other hand, when addressing (a set of devices in) a specific geographical area in an ad hoc network of mobile devices which is constantly changing, an application is more concerned with the location rather than the identity of the physical devices for which the information is intended. Generally, in such time-varying distributions of ad hoc nodes there can be no presumption that the information will reach its destination, nor that there will be any devices in that area to receive the information if it should ever arrive somewhere nearby.

In such dynamic situations, the conventional networking concept of preconfiguring a network path or set of paths from source to destination(s) and subsequently transmitting and attempting to forward packets along the path(s) will not succeed. GeoNetworking replaces this concept with one of broadcast forwarding of packets hoping they will ultimately find a path to the intended destination. Since it is very hard to make guarantees about network service qualities with this approach, it means that GeoNetworking is best suited to applications of a particular type: ones where messages are not time-critical, do not rely on maintaining a stateful communications session, are not unicast, and refer to data that will be valid for some time. Given the overhead added by GeoNetworking to packets for all applications, it is not clear that the cost of implementing GeoNetworking in the network stack, where the cost is borne by all applications, outweighs the benefit for the single limited class of applications for which GeoNetworking is more valuable. This particularly applies to bandwidth-limited channels that will carry time-critical safety-of-life information. This also suggests that the GeoNetworking function of geodissemination of packets might best be implemented as a series of application-layer forwarding decisions rather than as mandatory network-layer functionality.

8 Conclusions

Preliminary analysis of geodissemination of information and ETSI's GeoNetworking protocol performed by HTG1 and HTG3 has led to the following conclusions:

- Significant GeoNetworking protocol overhead and redundancy have the potential to cause congestion in capacity-constrained (RF) channels such as the 10MHz channels at 5.9GHz.
- Maintaining a given level of security and privacy in GeoNetworking without opening up the system to various attacks is problematic.
- Moving geodissemination functionality above the network layer addresses many of the security concerns with network layer implementations.
- Moving geodissemination functionality above the network layer allows legacy networks and devices to participate without modification.
- Moving geodissemination functionality above the network layer allows this functionality to be introduced seamlessly in the US when and if needed.
- GeoNetworking is only useful for a small subset of all ITS applications and should not be mandated for use by applications of other types.

9 References

- [1] ETSI TS 102 636-1 V1.1.1 (2010-03), Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 1: Requirements
- [2] ETSI TS 102 636-2 V1.1.1 (2010-03), Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 2: Scenarios
- [3] ETSI TS 102 636-3 V1.1.1 (2010-03), Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 3: Network architecture
- [4] ETSI TS 102 636-4-1 V1.1.1 (2011-06), Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 4: Geographical addressing and forwarding for point-to-point and point-to-multipoint communications - Sub-part 1: Media-Independent Functionality
- [5] ETSI TS 102 636-4-2, Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 4: Geographical addressing and forwarding for point-to-point and point-to-multipoint communications - Sub-part 2: Media dependent functionalities for ITS-G5A media (draft)
- [6] ETSI TS 102 636-5-1 V1.1.1 (2011-02), Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 5: Transport Protocols; Sub-part 1: Basic Transport Protocol
- [7] ETSI TS 102 636-6-1 V1.1.1 (2011-03), Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 6: Internet Integration; Sub-part 1: Transmission of IPv6 Packets over GeoNetworking Protocols
- [8] ETSI ES 202 663 V1.1.0 (2010-01), Intelligent Transport Systems (ITS); European profile standard for the physical and medium access control layer of Intelligent Transport Systems operating in the 5 GHz frequency band
- [9] ETSI TS 102 687 V1.1.1 (2011-07), Intelligent Transport Systems (ITS); Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access layer part
- [10] ETSI TS 102 724 V1.1.1 (2012-10), Intelligent Transport Systems (ITS); Harmonized Channel Specifications for Intelligent Transport Systems operating in the 5 GHz frequency band, Channel specifications 5 GHz
- [11] Offenlegungsschrift DE 10 2008 026 183 A1, 2009.12.03, Bundesrepublik Deutschland, Deutsches Patent-und Markenamt, "Verfahren und Kommunikationssystem zur Übertragung von Nachrichten", Brakemeier, Achim; Seeberger, Dieter; Weiss, Christian
- [12] WIPO International Publication Number WO 2008/092475 A1, 1 February 2007, "Method for Information Dissemination in a Communication Network", Applicant: NEC Deutschland GMBH; Applicants: Torrent Moreno, Marc; Hartenstein, Hannes; Festag, Andreas

- [13] WIPO International Publication Number WO 2009/018835 A1, 6 August 2007, "Method for Automatic Address Configuration in a Mobile Ad Hoc Network", Applicant: NEC Europe Ltd.; Applicants: Baldessari, Roberto; Festag, Andreas
- [14] European GeoNet Project - URL: <http://www.geonet-project.eu>
- [15] EU-US ITS Task Force, Standards Harmonization Working Group, Harmonization Task Group 1, HTG1-1:2012, Status of Security Standards