

Resilient internetwork routing over heterogeneous mobile military networks

Lars Landmark, Erlend Larsen,
Mariann Hauge

Norwegian Defence Research Establishment (FFI)
P.O. Box 25, NO-2027 Kjeller, Norway
{Lars.Landmark, Erlend.Larsen, Mariann.Hauge}@ffi.no

Øivind Kure

Department of Telematics
Norwegian University of Science and Technology
Norway
okure@item.ntnu.no

Abstract—Mobile networks in the military tactical domain, include a range of radio networks with very diverse characteristics and which may be employed differently from operation to operation. When interconnecting networks with dissimilar characteristics (e.g. capacity, range, mobility) a difficult trade-off is to fully utilize the diverse network characteristics while minimizing the cost. To support the ever increasing requirements for future operations it is necessary to provide tools to quickly alter the rule-set during an ongoing operation, due to a change in operation and/or to support different needs.

Our contribution is a routing protocol which targets these challenges. We propose an architecture to connect networks with different characteristics. One key point is that low capacity links/networks segments can be included in the heterogeneous network, these segments are protected from overload by controlling where and when signaling/data traffic is sent. The protocol supports traffic policing, including resource reservation. The other key point is the ability to quickly alter the network policy (rules-set) including QoS support during an operation or from operation to operation.

Keywords—*Heterogeneous Networks; Policy; QoS; Ad-hoc;*

I. INTRODUCTION

In the near term, mobile networks in the military domain will continue to consist of a set of networks dedicated for specific operational needs. These networks differ in their characteristics in terms of capacity, transmission range, transmission delay, rate of topology change, power consumption etc. Interconnecting available networks in an operation into one common network can increase the robustness, enhance the availability, and improve the overall network capacity. The drawbacks are increased risk of inconsistent routes, potential congestion of low capacity networks and hence an unstable network.

In our design we create an overlay network, and use depth-first searches for route discovery. A depth first search provides strict control of the search paths. It is easy to implement rules to protect low capacity networks from excess signaling traffic and low priority traffic. Additional route discovery methods might be utilized in network segments with sufficient capacity.

Depending on the network capacity, an overlay node may only be aware of its neighbors or of a larger scope of overlay nodes. Only neighbors would be the typical in case of only connected low capacity networks. Each of the overlay nodes has a policy rule-set per interface used to control signal and

data traffic. The policy set can be known by other nodes, but this is not a requirement. Our design shares similarities with pathlet routing[12], but in contrast to pathlet routing, our design does not require pathlets to be distributed. Instead it is based on depth first searches.

Our main contribution is to take control of the signaling traffic over connected low capacity networks and further to direct the traffic load according to local network policy. For this purpose, proactive routing is not feasible, unless the protocol use dissimilar signaling timers in different sub-networks. A weakness of such approach is high convergence time, and thus high risk of inconsistent routes [7]. Reactive routing and breadth-first search is also costly. In case of many low capacity networks, a breadth search will flood each low capacity network for each search. This method is also vulnerable for repeated flooding of the low capacity networks due to different transmission delays and search timeouts.

A great risk in heterogeneous networks is overloaded low capacity network segments. The worst case scenario is that a low capacity link or network ends up as the only interconnection between two high capacity islands. In such a scenario, the routing mechanism must be able to limit both control traffic and data traffic based on network policy and available resources. Hence, automatic mechanisms must be in place to route traffic with different requirements according to local policies and available network resources. The second contribution is the protocol's ability to perform some of these automated mechanisms. Hence, our focus is on signal forwarding, policy control and network resource utilization to make sure the networks are used as expected by the mission planners while providing necessary flexibility. Our cost is potential slightly longer paths

The rest of the paper is structured as follows. In Section II we discuss related work. The protocol design is described in Section III and the simulation environment is shown in Section IV. Simulation results are given in Section V, and Section VI concludes the discussion.

II. RELATED WORK

Little research has been done to study efficient IP connectivity between MANETs with different characteristics (e.g., capacities and delay). The work that we are aware of, can be grouped in three groups: (1) Proposals for new inter-domain protocols suitable for mobile environment [2][4]. (2) Proposals for modifications to make BGP [1] better suitable




Figure 1: Physical links within different networks (A) mapped to their respective network (B). (C) These networks are connected by virtual links/physical links used by the overlay Depth-First-Search (DFS) routing protocol.

for mobile networks [5][6]. (3) Proposals of hybrid, hierarchical, and composite routing that handle both internal and external routing [7][8]. A recent survey of inter-domain routing protocols is given in [3]. There are also some proposals that aim to improve BGP for a mobile environment. Arguments that favor BGP is that it is the most used inter-domain routing, and it is a well-studied protocol. It is potentially easier to support connectivity between a modified BGP protocol and the BGP protocol running in the backbone. In [6], the problem of AS-split is tackled. A dynamic AS reconfiguration is proposed. The paper also raises the discussion of whether the original policy should follow a deployed (and split) AS, or not. In [5], the two problems of dynamic BGP-peer discovery and slow convergence time of BGP are treated. A distributed peering broker service is implemented, and the BGP peers announce their mobility (stationary, low, medium, high) in order to select more stable paths. However, it is hard for BGP to provide multiple paths towards a destination while conforming to the policies of individual networks. In heterogeneous networks multiple paths are crucial as a consequence of capacity ranging from kbps to Gbps. In [12], it is referred to BGP and its tradeoffs when it comes to providing multipath. Various problems are listed, such as the problem of implementations not supporting all of BGP's routing policies. This may result in BGP exposing only a limited set of additional paths, making it difficult to know which path is being used. A consequence is that BGP mainly selects a single path for each destination, which is installed and further announced to its neighbors.

All the work presented above allows each MANET to preserve its internal legacy routing protocol. Some work is also done to study the use of one protocol in a heterogeneous environment. In these proposals, a scenario with many small MANETs of different characteristics is assumed. Instead of requiring a protocol for inter-domain connection, the same IP layer protocol (but different MAC and PHY protocols) is tried utilized in all nodes. In [7], it is shown that different protocol timers and varying transmission delays can increase the number of routing loops in a heterogeneous MANET. In [8], OSPF-MT is used in such an environment to provide some admission control and QoS. In [16] OLSR and OSPF-MDR are combined in composite routing. OSPF-MDR is a mobility

extension to OSPFv3 that interoperates natively with OSPF. In this architecture, OLSR is used in each MANET, and can be tuned to the MANET's properties. OLSR and OSPF-MDR are merged and use OSPF-MDR's functionality to interact with OSPF. OSPF ties the different MANETs together.

Few, of the mentioned protocols can adapt its signaling rate to also include low capacity networks in the common network, while providing consistent routes, network policy and multiple paths. In our work, we include low capacity networks at the cost of potential longer traffic paths, while providing support for QoS, policy and multipath.

III. PROTOCOL DESIGN

The proposed protocol aims to interconnect diverse wireless networks while protecting the low capacity networks from congestion. The protocol also supports resource reservation and policing. A network in this context may be a network segment of varying size or of only a point to point link.

Our protocol is designed to run as a network overlay as shown in figure 1. One benefit of using an overlay is that the protocol must be installed only on a subset of the nodes in the network of networks. The distance between any two neighboring nodes can be large. Routing to set up the link between the overlay neighbors can be based on local interior protocols or even BGP in case of large distance. The minimum requirement is that each node knows its overlay neighbor(s). This can be solved either by pre-configuration, or with a neighbor discovery protocol.

Depth-First Search (DFS) is used as the search technique. DFS has the property that a search can be steered and well controlled. It is easy to implement rules to control signaling traffic and to steer data traffic. The drawbacks of a DFS approach are an increased path setup delay, and a likelihood of discovering slightly longer paths than the optimal. In addition, reactive in contrast to proactive protocols, are normally associated with an overhead that increases with the number of flows. In our work, these disadvantages are mitigated by the use of search/policy/QoS history and other ongoing traffic as hints for new searches.

A hint is used to advise searches and can be acquired from any source, such as other routing protocols, management traffic, intercepted data traffic etc. In our implementation, we collect hints from ongoing searches and store them as *hints*. A node might end its search with a *notfound* message in case of dead end. This status is returned, either when there are no more links to try, or due to implemented policy. Either way, the information is sent back along the path to the source and stored as hints for future searches. The *hint repository* is used by all searches. A search based on incorrect hints does not lead to an erroneous search but will have an increased route setup delay, and a higher overhead.

One advantage with the DFS strategy is that traffic can be easily controlled using different policies. Our protocol is designed to discover a path conforming to the search requirements and the current network policy. The policy property of a network can be learned either from other protocols, or by the search protocol itself. In the latter case, each search carries information about the resource request, address information, priority etc. A network owner might block a search. In this case, the search bounces back to the previous hop and the result is stored. With this design, we support the possibility to block a link or network segment for all external signaling and data traffic when this best suites the operation. This will often be the default setting for narrowband network segments.




Figure 2: Protocol state machine

The DFS operation is depicted in figures 2 and 3. A new search is initiated when no route information is available. Similar AODV2 [10], packets are stored, and a *search* message is initiated towards the destination while the route back to source is stored. To improve connectivity, the addresses of the intermediate nodes are enclosed in the *search* message. Each *search* packet contains $\langle source, destination, sequence\ number, priority, QoS, flag, nextHop\ and\ the\ route\ table\ of\ the\ previous\ hop \rangle$. The route table is updated at each intermediate node, and all route entries are checked against the current policy before they are added to the search message. Hence, the search information is stored in each visited node, and the route table is updated regarded to the previous hop's announced route table, similar to BGP.

Searches can match the route table on variables in the triplet $\langle source, destination, TOS \rangle$. Route entries are added to the node's route table, based on which of the three variables that are filled. First method writes all, and is used by individual searches. Second writes $\langle -1, destination, TOS \rangle$ and is filled by visited nodes. Last writes $\langle -1, destination, -1 \rangle$ and is used for route exchange between neighbors.

The main reason for this separation is to differentiate on the accuracy of the hints. We trust the route entries by the number of match variables in the route entry. Hence, a match on all

three route tuples is trusted, while only destination is less trusted. Depending on the policy and flow requirement, data packets can be dispatched, or a new search is needed.

In case a new search is necessary, a *search* message is issued from the source. Next, the source consults its hint database for the next hop. The hint database will take into account its topology view, available resources and learned policy rules in the networks towards the destination. If the hint database fails to deliver a next hop towards the destination, the search randomly select a neighbor. The *search* message is then dispatched with the FLAG set to SEARCH and the node STATE set to SEARCHING. The corresponding link is marked as tested. Each link is only tested once per search. A consequence of operating over unreliable networks is the need for the next hop node to acknowledge the received message. If no acknowledge is received before a *timeout* has expired, the search is retransmitted.

A *search* message received by intermediate nodes can result in one of two actions; it is either forwarded or returned to previous hop. In case of a new search, the search variables are stored, and the state is set to SEARCHING before a new next hop is tried. On the other hand, if the search is already processed and the state is SEARCHING, and if all available links are tried, the intermediate node replies with the status set to *NOTFOUND*.

When a *search* message is successfully received at the destination, the destination sets the state to FOUND and replies with a *found* message back to the source. Each intermediate node and the source receiving the *found* message changes its state from SEARCHING to FOUND. A *found* message might be lost resulting in a terminated search. The source will then issue a new *search* after a *source-search-timeout*.

Similarly to AODV2, the protocol responds to a link break for ongoing flows. Each link, for ongoing flows, is monitored by the use of *hello* messages. If a node does not receive a *hello* messages within a predefined *hello-timeout* window, the neighbor is set to lost, and the neighbor adjacency is removed from the hint database. After setting the link to lost, an *error* message is sent to all sources affected by this link.

Our DFS protocol supports both admission control and resource reservation. A requirement is the ability to reserve resources on the physical bearers between overlay nodes. Otherwise, it is used for resource management within the overlay network. It supports resources reservation similar to




Figure 3: Protocol signaling for node A to node B

A. Network Resource Management

RSVP [15], for a specific $\langle source, destination, TOS \rangle$ triplet, where source and destination could be unique nodes or source/destination networks. The request for resource

reservation is carried by the *search* packet, and each node along the path consults its *available resources* database on the calculated next hop link. Resources are either manually preconfigured, or acquired by measurements. In our work the *available resources* were preconfigured with the link capacity and further altered by each flow's QoS request. The used resource is soft stated, meaning that the resource reservation times out, unless a continuous flow of data packets keep the timer updated. The resource timeout is a function of priority. Hence, a high priority flow holds its resources longer than a low priority flow.

B. Network Policy

In this work, we assume all subnetworks conforms to a common set of policy rules. However, the different network owners are free to choose which of the rules in the common set to enforce at a given time. Hence, each network preserves its own autonomy both in time and space. With this architecture the network owners can reserve a fraction of the resources for e.g., coalition use and apply national QoS and policy rules on the remainder of the resources. The resources reserved for external use can then be policed according to the needs of that network (e.g., a Protected Core coalition Network (PCN) [11]).

Each node will have policy implemented per neighbor adjacency. The policy might be static or dynamic, or a mixture of static and dynamic. A static policy will typically be long lived, e.g., traffic from operator A is not allowed into the network, while a dynamic policy might be short term, e.g., current capacity.

A flow is admitted or rejected, based on network connectivity and policy setting. Each *search* message contains a header declaring the type of traffic in terms of *priority*, *QoS* etc. In case a flow is blocked, the policy that led to the rejection is reported in the *notfound* message to the previous router and stored in the previous router's *hint repository* with a timer. The timer is dictated by the router rejecting the search request.

The DFS protocol allows the use of local policies, and hence does not require distributed consistent policy setting for all nodes in order to obtain consistent forwarding tables. With this scheme, the policy can be determined, set and altered locally, without the need for global negotiations to assure consistent policies. In case of a local change in policy, the sources of the affected traffic are notified.

IV. SIMULATION SETUP

The protocol has been implemented in the ns-3 framework [13] using point-to-point links to represent the overlay network. To test the effect of hints, we first evaluated the protocol on static 8x5 grid topology. We collected statistics between two sources and destinations with a nominal distance of 7 hops. Background traffic was generated randomly between any source and destination every 2 seconds and lasted 5 seconds, and hints were updated if a shorter path was found.

Next, we ran simulations to evaluate DFS's ability to protect low capacity networks from overload and prioritize mission critical traffic, using policy rules and resource reservation. This behavior was tested over a static topology referred to as the islands topology (figure 4). Two high




Figure 4: Simulated islands topology

capacity network islands are connected via network with different capacities to illustrate the key problem of keeping control of signaling traffic and traffic management.

The two high capacity islands are connected via three different link capacities; one link representing a typical VHF network (between nodes 2 and 12), one representing SATCOM (between nodes 4 and 14) and one representing UHF (between nodes 0 and 10, through 5 and 20). The VHF link capacity is set to 64 kbps, the UHF links are 1 Mbps and the SATCOM link provides 640 kbps. The transmission delay for VHF and UHF is 1 ms, while the SATCOM link has a 600 ms delay. All local links within each islands are also UHF, but with a higher capacity of 11 Mbps. In a normal operation, the majority of the traffic will be generated and terminated within the network of the source, but an increasing amount of traffic needs to traverse several networks. In our test, we evaluated the protocol by sending traffic between the islands.

The islands topology was used for three different tests; 1) the no policing case, 2) the policing case, and 3) with resource reservation. In the first simulation, we did not use policies and treated all traffic as best effort. All traffic was allowed over all link types. In the second simulation, the high priority traffic was allowed over all link types, the medium priority was allowed over SATCOM and UHF, and the best effort class was only allowed over UHF. In the third simulation, we ran the same traffic classification as in the second case, but with resource reservation.

The following flows were generated: Two high priority CBR flows of 30 kbps each, with a duration of 200 seconds initiated at 90 and 100 seconds. Three medium priorities flows, with duration of 200 seconds, each of 150 kbps and being initiated at 50, 60 and 70 seconds. The low-priority/best-

effort flows were generated every 2 seconds with duration of 5 seconds. The background best-effort data rate is variable and shown on the x-axis in the figures 7 to 12. The y-axis shows the number of packets being dropped. All traffic classes have source and destination located on different islands. Hence, the available crossover paths were VHF between node 2 and 12, SATCOM between node 4 and 14, and a number of UHF links between node 1 and 10. Each presented result is the average of 5 simulation runs with different seeds.

A node's route entry was timed out and further rewritten as $\langle -1, destination, -1 \rangle$ after 6 seconds. The change in route entry made it useful as hints for further searches, but not for data forwarding.

In these simulations we only represented the overlay network, thus all nodes participated in the routing. An overlay node discovered its neighbor using a neighbor discovery protocol. Each node was aware of its link capacity, but no policy information was pre-shared for the neighboring node(s). Signaling packets were prioritized in the simulations by being inserted in the front for the queue instead of appended to the back. Hence, signaling packets were not lost due to tail drop, and had low delay.

V. SIMULATION RESULTS

In this section, we present simulation results. Using the resulting packet loss and delay, the protocol is evaluated with regards to the effect of hints, and the ability to do network policing and resource reservation.

A. Hints versus no hints

This section evaluates the benefit of hints with regards to search time and path length. The Depth First Search (DFS) algorithm is evaluated based on the number of hops to the destination. The path length was measured as a snapshot at time 50 s and time 550 s.

At time 50, the nodes only know about their overlay neighbors and will do a random search, while at time 550 the hint database is filled based on earlier searches. As more knowledge and hints are acquired based on ongoing searches, the search time and distance is reduced. Given a uniform $\langle source, destination, tos \rangle$ search pattern, it is possible to build a route table close to the optimal over time. However, a route hint is used only as long as it conforms to the search requirements and network policy.

In terms of experienced hop count (figure 5), there was a reduction in the number of hops by 50% from time 50 to time 550. The time points 50 and 550 were only selected as snapshots, and the convergence time depends on the amount of cross searches and stale route entries. In this simulation, there was no mobility. The advantage of hints on the total number of routing transmissions is shown in figure 6. Over the course of the simulation, searches that go in random directions, and does not learn from previous searches, generate three times as many routing transmissions.

DFS is evaluated by the number of hops. However, the shortest path is not a suitable indicator of the actual length or cost of the path. The reason is that the underlying overlay links may differentiate in terms of capacities and number of physical hops. A better option is to optimize on resources, where resources include e.g. delay, bandwidth and energy.




Figure 5: Hop count for two flows with different starting times.

B. DFS without policing and resource reservation

In the first simulation using the islands topology (figure 4), OLSR [9] and Depth First Search (DFS) were run without any priority/policing or QoS for comparison. OLSR is included for illustration purpose, showing the consequence of shortest path routing. More comparisons with OLSR are shown for a previous version of our protocol in [17].

As expected, OLSR congest low capacity links, due to the shortest path routing (figure 7). DFS is not a shortest path routing algorithm, but starts randomly without any hints. As a result, flows will initially rarely follow the shortest path, but be spread over alternative paths. Similar observation is made for breadth search [14], as searches tend to reach their destinations first over less loaded paths. DFS will approach shortest path by time, if shortest path is the chosen optimization.

The end-to-end delay results (figure 8) tell an interesting tale. The DFSP delay varies between 7 and 28 s, while OLSR begins with a 70 s delay, which decreases with increased traffic. The reason is high queuing delay over VHF. VHF has less capacity, and thus contributes to the delay through queueing. Our simulation results show that more packets are transmitted over the SATCOM link with increased data rate. Although SATCOM has higher transmission delay, it has more capacity to transmit packets, causing the average overall queueing delay to be reduced with increasing traffic rate. At low traffic rate, the VHF queue is filled, but with few drops. At high traffic rate, the VHF queue starts to drop packets, but at the same time more packets are sent over SATCOM without being dropped. Hence, the delay statistic is skewed towards the SATCOM transmission delay with increasing data rate.

C. DFS with policing

With neither policing nor resource reservation, important traffic is not prioritized. In the following we enforce a policy that reserves the low capacity networks (SATCOM and VHF) for the high and medium priority. Thus, a policy was set up, where the high priority traffic was allowed over all link types, the medium priority traffic was allowed over SATCOM and UHF, and the best effort traffic was only allowed over UHF.

These simulations gave an interesting result for the situation with no packet loss on any traffic type (figure 9) and low priority data rate below 300 kbps. The main effect of




Figure 6: Routing transmissions on grid topology for DFSP.




Figure 7: Packet loss, number of lost packets in terms of offered traffic load.

reserving the low capacity networks for high priority traffic was to force the low priority traffic over the higher capacity links between the islands. Investigating the packet loss for the high priority traffic (figure 10) reveals the same behavior as for the low priority traffic, starting when each low priority flow reaches a data rate of 300 kbps, leading to congested UHF links. The total traffic load for prioritized traffic is less than the capacity of VHF and SATCOM, but still these flows experience packet loss. The reason is that prioritized traffic is allowed over UHF, and not forced over SATCOM and VHF. Hence, SATCOM and VHF is less utilized, while the remaining prioritized traffic contests for resources over UHF, with similar likelihood of packet loss as the low priority traffic. A quick fix to this is to implement differentiated service and priority scheduling.

D. DFS with policing and resource reservation

Protecting low capacity networks using policy is not enough. Resource reservation is required to protect prioritized traffic while not starving low capacity flows. In this work, the resource that is reserved is capacity, but other types of resources could be just as relevant. By searching for a path with available resources, the traffic is routed over links with available capacity.

The total packet loss (figure 11) increases when the route search includes resource reservation. The reason is




Figure 8: End-to-end delay.




Figure 9: Packet loss with policing, number of lost packets in terms of offered traffic load.

twofold. First, as more links need to be searched before a valid path is found, more packets are buffered resulting in higher likelihood of buffer overflow. Second, more packets are dropped due to no valid path found and thus more flows are denied admittance to the network. The packet loss for high priority data reveals that although the total packet loss now is higher and begins earlier than without resource reservation, the high priority traffic is now protected.

E. Lesson learned

The main purpose of resource reservation is to ensure reduced loss for prioritized traffic. With both policing and resource reservation, we experience no loss for prioritized traffic. However, the loss and channel throughput is sensitive to the soft-state timers. A long timer results in low throughput since resources are held, but not used. On the other hand, with a short timer, link entries may be timed out, due to irregular inter-packet time, resulting in increased no-route loss.

VI. CONCLUSION

In this work we have showed that a routing protocol based on the depth-first search principle can be steered to protect low capacity networks from both unwanted signaling traffic and data traffic in a heterogeneous network of many networks.

Depth First Search is associated with long search delays and the risk of not finding an optimal path. We have shown that these problems can be mitigated by using hints. Routing hints also reduce the signaling traffic. By collecting results from previous searches and/or ongoing management traffic, experience is gathered and used in subsequent searches. Policy is implemented for better control of both signaling and data traffic. Each link is configured with a policy based on operation. As a consequence, searches are only forwarded over links that conform to the current policy.

The protocol supports resource reservation. When resource




Figure 10: Packet loss for high priority traffic with policing, number of lost packets in terms of offered traffic load.




Figure 11: Total packet loss when applying policing and resource reservation, number of lost packets in terms of offered traffic load.

reservation is used, neither signaling traffic nor data traffic is forwarded over networks without sufficient resources.

In our example, we differentiate traffic by TOS, but we could also police traffic based on e.g., source, destination, protocol type, and visited intermediate nodes. The idea is to ensure that networks forward traffic as intended. In these simulations we have used a static network topology to represent an environment of mobile networks. The overlay network will not experience the same mobility as a MANET since it is the job of the underlying network to reroute the link between overlay nodes, however there will be some long time link breaks and lots of short breaks that we will introduce in future simulations.

REFERENCES

- [1] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," *RFC 4271*, Jan. 2006, www.ietf.org.
- [2] C.-K. Chau, J. Crowcroft, K.-W. Lee, and S. H. Y. Wong., "Inter-domain routing for mobile ad hoc networks," in *proceedings MobiArch*, Seattle, WA, USA, Aug. 2008, pp. 61-66
- [3] T. Gibbons, J. Van Hook, W. Na, T. Shake, D. Street, and V. Ramachandran, "A Survey of Tactically Suitable Exterior Gateway Protocols," in *proceedings MILCOM*, Nov. 2013, pp. 487-493
- [4] S.-H. Lee, S. H. Y. Wong, C.-K. Chau, K.-W. Lee, J. Crowcroft, and M. Gerla, "InterMR: Inter-MANET routing in heterogeneous MANETs," in *proceedings MASS*, San Francisco, CA, USA, Nov. 2010, pp. 372-381.
- [5] M. Kaddoura, B. Trent, R. Ramanujan, and G. Hadynski, "BGP-MX: Border Gateway Protocol with Mobility Extensions," in *proceedings MILCOM*, Baltimore, MD, USA, Nov. 2011, pp. 687-692
- [6] S. Hares and R. White, "BGP Dynamic AS Reconfiguration," in *proceedings MILCOM*, Oct. 2007, pp. 1-7.
- [7] L. Landmark, M. Hauge, and O. Kure, "Routing Loops in Mobile Heterogeneous Ad Hoc Networks," in *proceedings MILCOM*, Nov. 2013, pp. 112-118
- [8] M. Hauge, J. Andersson, M. A. Brose, and J. Sander, "Multi-Topology Routing for QoS Support in the CoNSIS Convoy MANET," in *proceedings MCC*, Gdansk, Polen, Oct. 2012, pp. 179-197
- [9] T. Clausen and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)", RFC 3626, Oct. 2003, www.ietf.org.
- [10] C. Perkins, S. Ratliff, and J. Dowdell, "Dynamic MANET On-demand (AODVv2) Routing", *draft-ietf-manet-aodvv2-03*(work in progress), Feb. 2014, www.ietf.org.
- [11] R. Schutz, S. McLaughlin, T. Daeleman, M. Luoma, M. Peuhkuri, P. Carlen, and J. Haines, "Protected Core Networking (PCN): PCN QoS and SLA definition," in *proceedings MCC*, St.-Malo, Oct. 2013
- [12] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," in *proceedings ACM SIGCOMM*, Barcelona, Spain, Aug. 2009, pp. 111-122.
- [13] ns-3 network simulator, <http://www.nsnam.org/>, last accessed January. 2015.
- [14] L. Landmark, Ø. Kure, Knut Øvsthus, "Performance analysis of the AODV ad hoc routing protocol in a dual radio network", in *proceedings WMuNeP'05*, pages 106-112, 2005
- [15] Braden, R., et al, "Resource ReSerVation Protocol (RSVP)", RFC 2205, September 1997
- [16] F. Jin, T. Goff, and P. Guangyu, "Comparison studies of OSPF-MDR, OLSR and Composite Routing," in *proceedings MILCOM*, San Jose, CA, USA, Oct. 2010, pp. 989-994.
- [17] M. Hauge, L. Landmark, E. Larsen, P. E. Engelstad, and Ø. Kure, "Resilient internetwork routing with QoS support over heterogeneous mobile military networks", *STO-MP-IST-123 Symposium on Cognitive Radio and Future Networks*, The Hague, The Netherlands, May 2014.