2014 IEEE Military Communications Conference

Opportunistic Situational Awareness Dissemination at the Tactical Edge

J. David Brown, Mazda Salmanian, and Ming Li
Defence R&D Canada - Ottawa, Ontario, Canada [first.last]@drdc-rddc.gc.ca

***Abstract*—Recent advances in commercial and military technology have contributed to an emerging trend to equip soldiers at the edge of tactical communication networks with wireless computing devices capable of sharing situational awareness (SA) data over mobile ad hoc networks (MANETs). This paper presents an efficient algorithm called "Opportunistic SA Passing" (OSAP) to share short periodic bursts of SA data. OSAP adopts the philosophy that it is not necessary to guarantee delivery of SA about all nodes in the MANET at all times. Rather, OSAP relaxes this requirement, allowing us to explore the network efficiencies gained by distributing SA about "most" of the nodes (typically greater than 90%). We conducted simulations of OSAP in MANETs ranging in size from a light infantry section (or squad) to a platoon. Compared to three other SA distribution schemes—pure flooding, central collection and broadcast, and an efficient broadcast based on Multi-Point Relay—OSAP requires dramatically fewer frames transmitted and forwarded in the MANET to achieve better than 90% coverage of connected nodes.**

Keywords—Mobile ad hoc networks (MANET), Situational Awareness, Tactical Edge

I. INTRODUCTION

Access to up-to-date situational awareness (SA)—including troop location and movement—is critical in tactical operations. Recent technological advances are poised to make commonplace the extension of SA information down to the level of the dismounted solider, as opposed to more traditional techniques of vehicle-gathered SA or SA supplied by a subset of key resources in the network. Furthermore, this SA is not necessarily limited to GPS data, but may include other mission-critical data recorded by infantry-carried sensors. This information must be shared up to commanders and decision-makers and could also be shared among the soldiers themselves.

One technology that may help facilitate periodic broadcasts of SA information to the edge of tactical networks is the mobile ad hoc network (MANET). MANETs offer support for dynamic and beyond-line-of-sight networking in contested environments where fixed communication infrastructures may not exist. In a MANET, allied nodes can act as relays and routers to provide multi-hop connectivity through intermediary nodes to intended destinations. However, as with any wireless communications in a tactical environment, nodes in a MANET must communicate over bandwidth-constrained links. Furthermore, the handheld mobile devices of dismounted soldiers in a MANET are limited in terms of processing power and battery life. Consequently, it is vital to adopt information-exchange protocols that limit processing requirements and channel access frequency when possible.

To this end, this paper proposes an algorithm called Opportunistic Situational Awareness Passing (OSAP): a method for disseminating short bursts of SA information through a MANET such that a high percentage (greater than 90%) of SA information about connected nodes is successfully distributed while using substantially fewer frames than other compared protocols. In OSAP, a randomly selected "initiator" node broadcasts out its own SA data. The nearest-neighbors of the initiator will re-broadcast the received SA, while also appending their own SA to the frame; the re-broadcasting and appending continue until each connected node in the MANET has transmitted the initiator's SA exactly once. After one round of broadcasts has completed, OSAP dictates that a new initiator begin the process all over again. After as few as 3 or 4 rounds (in dense networks), most nodes in the network have an accurate SA picture of most other nodes (i.e., SA concerning 90-95% of their peers).

To evaluate the utility of OSAP in military networks, we present MATLAB simulations of its performance in MANETs ranging in size from 10 nodes (the approximate size of a squad or section) to 30 nodes (the size of a small platoon). We compute the total average SA per node after each round of OSAP controlled broadcasts[1] and compare OSAP against three other existing schemes for disseminating SA information: a "pure flooding" scheme; a scheme that relies on a central node to collect SA information about everyone else and then rebroadcast it; and an efficient broadcast scheme that leverages the multi-point relay (MPR) flooding algorithm from the optimized link state routing (OLSR) protocol [1]. While each of the three schemes for comparison has as an end-goal the complete dissemination of SA (i.e., all nodes receive SA about every other connected node), OSAP can deliver as much or as little SA as is desired, depending upon the number of "rounds" that are executed. We demonstrate that to achieve 90-95% of maximal SA dissemination requires far fewer frames than are required by the three other schemes.

We also present simulations conducted in QualNet to evaluate how OSAP would perform in a more realistic mobile tactical scenario. The QualNet simulations introduce the concept of "SA freshness" and show how the freshness of

[1] where the "average SA per node" is a measure of the average number of nodes about which any particular user has SA information

978-1-4799-6770-4/14 $31.00 © 2014 IEEE
DOI 10.1109/MILCOM.2014.205

IEEE computer society

DRDC-RDDC-2014-P120

information is strongly influenced by network connectivity, mobility, and the frequency of OSAP "rounds".

The rest of the paper is organized as follows. Section II provides an overview of three existing SA dissemination methods—namely, flooding, centralized broadcasts, and Multi-Point Relay- (MPR) based broadcasts—along with specific implementation details for OSAP. In Section III, we present our MATLAB simulations comparing OSAP against the three other SA dissemination techniques, focusing on average SA per node as a function of the number of transmitted frames. The performance of OSAP in realistic mobile tactical scenarios implemented in QualNet is presented in Section IV, where we also discuss how tuning parameters of OSAP can help keep the SA fresh, or up-to-date. Finally, Section V provides a brief conclusion and directions for future research.

II. Assumptions, Definitions and Review of SA Dissemmination

This section introduces some assumptions made in this work, provides definitions of key terms, and describes three existing methods for disseminating situational awareness information in a MANET, along with our proposed OSAP algorithm.

In this paper, we assume that each node in the network is capable of computing certain node-specific information that we call node-specific-SA. This information may be as simple as the node's GPS location or battery life, and we assume that the SA can be compressed to occupy a very small number of bits[2]. Furthermore, we assume that node-specific-SA is always in flux and must be disseminated periodically. If a user in the network is in possession of node-specific-SA for every node in the network, then we say that user has complete network-SA. Similarly, if a user is in possession of node-specific SA for x% of the nodes in the network, that user is defined as having x%-network-SA.

At this point, it is useful to highlight a distinction between the problems of SA dissemination and route discovery. At first glance, the two problems appear to have much in common, with both focused on providing network-wide insight to individual nodes in the network. However, unlike routing information, SA information is not concerned with end-to-end connectivity and does not require an understanding of which node is connected to which. In fact, the three existing SA dissemination schemes discussed in this section do not supplant routing protocols and would not (on their own) provide routing information to the network.

The remainder of this section describes existing and proposed SA dissemination strategies. Throughout the discussions, we adopt the notation that there are N nodes in the network, each node is assigned a unique number from 1 through N, and the node numbered i is referred to as n_i.

A. Simple Flooding

A network flooding strategy (e.g., discussed as one possibility for SA dissemination in [1]) is a straightforward way of disseminating SA in the network and constitutes a simple baseline for comparing other more involved schemes. For node n_i to disseminate its own node-specific-SA using a flood, it proceeds as follows. Node n_i broadcasts a flood frame that uniquely identifies the source (i.e., n_i) and contains a unique sequence number along with its node-specific-SA. All neighbors of n_i receive the frame and check if they have received it before (by consulting the source and sequence number). If they have received the frame before, they simply drop it. Otherwise, if this is the first time they have received the frame they process it and rebroadcast it. In this fashion, once the flood is complete, every (connected) node in the network will have sent the flood frame exactly once, and they will all have node-specific-SA about node n_i.

To achieve network-wide SA, every node in the network initiates an SA flood according to a pre-defined period, denoted here by T_{SA}. In this fashion, every connected node will receive node-specific-SA about all other nodes every T_{SA} seconds. We note that the nodes do not need to be synchronized and can each begin their broadcasts independently[3].

B. Centralized Broadcasts

In a centralized broadcast scheme, one node is designated as a "master". Every T_{SA} seconds, all nodes with connectivity to the master will transmit (unicast) their node-specific-SA to the master node. At this point, the master node has complete network-SA (of the connected nodes). The master node will then initiate a network flood (as described in II.A), where the complete network-SA is included in the flood packet. Thus, every T_{SA} seconds, all nodes connected to the master will receive node-specific-SA about all other nodes connected to the master. Schemes of this type are described in [2] or in the introduction to [3]. This type of scheme is typical of SA systems that rely on the existence of a centralized server or database.

In military networks, a section or platoon Commander is a sensible choice for the master node. However, it is also possible to rotate the identity of the master node in a round-robin fashion.

We note that this scheme requires each node to have a known route to the master node; thus, unlike a pure flooding scheme as in II.A, the centralized broadcast scheme presupposes the existence of a routing protocol *a priori* to the SA dissemination.

[2] Compression strategies for node-specific-SA are beyond the scope of this paper. However, we suggest that for periodically-transmitted SA, an encoding strategy could be adopted whereby the difference in SA is sent relative to a known "base value" or relative to the last known update (similar to a differential encoding scheme). For GPS-coordinates, for instance, this could dramatically reduce the required size of the SA message.

[3] In fact, it is preferable if nodes do not synchronize their broadcasts to avoid causing unnecessary congestion. This could be achieved by introducing a random jitter to the timing.

C. MPR-based Broadcast

An alternative to the simple flooding strategy in II.A is to make use of the concept of multi-point relays (MPRs), as used by the OLSR protocol for the dissemination of topology control messages. The use of MPRs for efficient broadcast in MANETs is discussed in some detail in [4] and the references therein. The OLSR standard [1] contains a heuristic algorithm that describes how each node in a MANET can select its own MPRs among the set of its one-hop neighbors[4]. The MPRs are strategically selected (locally by each node) in an effort to maximize the amount of network coverage that can be achieved based on a minimal set of nodes.

In the MPR-based broadcast strategy, when a node n_i initiates an SA broadcast, only its locally-selected MPRs rebroadcast the message. All other one-hop neighbors will process the message (and record the node-specific-SA) but will not forward it on. The message will continue to be broadcast by subsequent MPRs in the network.

To achieve complete network SA for connected nodes, each node initiates an MPR-based SA broadcast every T_{SA} seconds. This should require fewer packets than the scheme in II.A, as certain (non-MPR) nodes do not rebroadcast the SA.

While the MPR-based scheme does not presuppose the existence of a routing protocol in order to operate, it does necessarily require the use of an MPR-selection scheme to identify the MPRs (thus, suggesting that OLSR be used as a routing protocol in this case).

D. Opportunistic SA Passing (OSAP)

We now consider the new SA dissemination scheme proposed in this paper: Opportunistic SA passing (OSAP).

To begin, we describe the operation of a single OSAP "round". Suppose node n_i is the initiator of a round. Similar to schemes II.A and II.C, node n_i broadcasts a frame containing a source identifier (i.e., n_i is the originator), a unique sequence number, and its node-specific-SA. All one-hop neighbors of n_i check to see if they have already received this SA frame (by examining the source and sequence numbers). If they have not already received it, they append their own node-specific-SA to the frame and rebroadcast it, making sure to leave the source and sequence numbers unchanged. If they have received the SA frame before, they do not rebroadcast it; however, they still inspect the message to extract any additional node-specific-SA that may have been appended since their initial receipt of the frame.

Some thought reveals that after one round, every connected node in the network will have (at a minimum) node-specific-SA about the initiator of the round and node-specific-SA about all of its own one-hop neighbors. Nodes will in all likelihood have additional node-specific-SA about other nodes on the route between themselves and the initiator as well.

Now that we have described a single OSAP round, we propose a simple round-robin method of cycling through nodes in subsequent rounds[5]. Essentially, we want a new node to initiate an OSAP round once every T_{OSAP} seconds, where $T_{OSAP} \neq T_{SA}$ (note that we do not want all nodes initiating rounds simultaneously, but we intentionally want the rounds "staggered"). This could be achieved in any number of ways, where we propose the following method. For N nodes in the network, we let each node in the network have a unique ID from 1 to N. In addition, we assume the nodes are (roughly) synchronized to a common clock. Each node periodically computes the value $n(t) = 1 + \left\lceil \frac{t \bmod (N \cdot T_{OSAP})}{T_{OSAP}} \right\rceil$, where T_{OSAP} is an integer representing the desired interval between OSAP rounds (to make this an integer, an appropriate granularity for the timescale can be selected—e.g., time in milliseconds), t is the current time expressed in the same units as T_{OSAP}, and N is the number of nodes in the network. This computation produces a value for $n(t)$ that cycles through the integers from 1 to N, resulting in a new sequential value every T_{OSAP} units of time. When a node's computation of $n(t)$ is equal to its unique ID, the node assumes the role of initiator and broadcasts out its node-specific-SA to begin a round. Consequently, any particular node will be an initiator once every $N \cdot T_{OSAP}$ units of time.

For OSAP to *guarantee* complete network-SA to all connected nodes, each node must take on the role of initiator, meaning that N rounds are required (resulting in a verbose flooding scheme). However, we will demonstrate in Section III that after only 3 or 4 rounds of OSAP in a dense network, all connected nodes in the network achieve a state of 90- to 95%-network SA, even with N as large as 30. Another notable side benefit of OSAP is that because nodes append node-specific-SA to OSAP messages while propagating through the MANET, this means that having knowledge of any node's SA automatically implies having knowledge of a route to this node as well. OSAP does not pre-suppose the existence of a routing protocol and may (depending upon requirements) provide adequate routing information on its own for many applications.

III. MATLAB Simulation Results

This section presents the results of simulations conducted to evaluate the performance of OSAP in comparison to the three other SA dissemination schemes presented in section II (i.e., flooding, MPR broadcast, and centralized broadcast).

To evaluate the performance of the various schemes, we conducted simulations in MATLAB whereby a square area was populated by nodes randomly distributed according to a uniform distribution. Nodes could communicate directly with

[4] Note that the MPR-selection algorithm pre-supposes that each node is aware of its own one-hop and two-hop neighbors.

[5] In this paper, we assume that the number of deployed nodes in the network is known *a priori*, and that they have been pre-provisioned appropriately with unique source identifiers (IDs) and knowledge of the source IDs of their peers. This is a reasonable assumption in a military deployment with known formations (e.g., squad/section, platoon, company, etc) and constituent units. However, other more complex cycling schemes could be proposed in the absence of this *a priori* information. These more complex schemes are of interest for future research, but are beyond the scope of this paper.

their neighbors if they were separated by a distance less than the transmission range, R. We considered scenarios with

TABLE I
MATLAB SIMULATION PARAMETERS

Sim #	N	A (dist units^2)	R (dist units)	Avg # nodes in R-circle	Density
1	10	80	4	6.28	dense
2	20	160	4	6.28	dense
3	30	240	4	6.28	dense
4	10	80	3	3.53	sparse
5	20	160	3	3.53	sparse
6	30	240	3	3.53	sparse

"dense" networks and scenarios with "sparse" networks, where a dense network was achieved by increasing the transmission range of each node, leading to an increase in the average number of nearest neighbors per node[6]. For both the dense and sparse cases, we examined networks with 10, 20, and 30 nodes (to represent a single squad/section, 2 squads/sections together, and a small platoon). Table I summarizes the six network scenarios that we considered. Table I also includes the values selected for transmission range, R, the size of the simulation area, A, and the density of the network measured as the average number of nodes in a circle of radius R (i.e., the average number of nodes in communication range of a hypothetical node at the centre of this circle). Note that R and A were selected to achieve consistent densities for the dense and sparse scenarios when N was varied.

For each of the six network scenarios, we simulated all four of the SA dissemination schemes presented in section II. For each scheme, we recorded how many frames were transmitted (i.e., the total of all SA frames transmitted by every node in the network). We also recorded a metric we refer to as "average SA per node", denoted as $\overline{SA}$, which we explain as follows. If node n_i possesses node-specific-SA about k nodes in the network (including itself), then n_i has an SA value of k; that is, $SA_i = k$. Since it is assumed that a node always has SA about itself, SA_i has a minimum value of 1. Average SA per node is then simply computed as the average over all SA_i, such that $\overline{SA} = (\sum_i SA_i)/N$.

For the OSAP scheme, we computed the value of $\overline{SA}$ after each round, recording the number of frames required per round—this allowed us to plot $\overline{SA}$ as a function of the number of transmitted frames. The other three schemes were plotted as single points (i.e., the value $\overline{SA}$ and the number of frames that was required to achieve this) since they do not operate on the concept of "rounds".

We repeated our simulations 5,000 times each in order to obtain an average measure of performance on uniformly distributed networks of nodes.

The results of the simulations are shown in Figure 1. The performance of OSAP is shown by the red line in the figures,

[6] Note that increasing the transmission range does not increase the node density in terms of nodes per unit area. It does, however, increase the effective density for our purposes, since we are concerned with the number of nodes within communication range. A dense network has many nodes interconnected, while a sparse network has fewer interconnections.

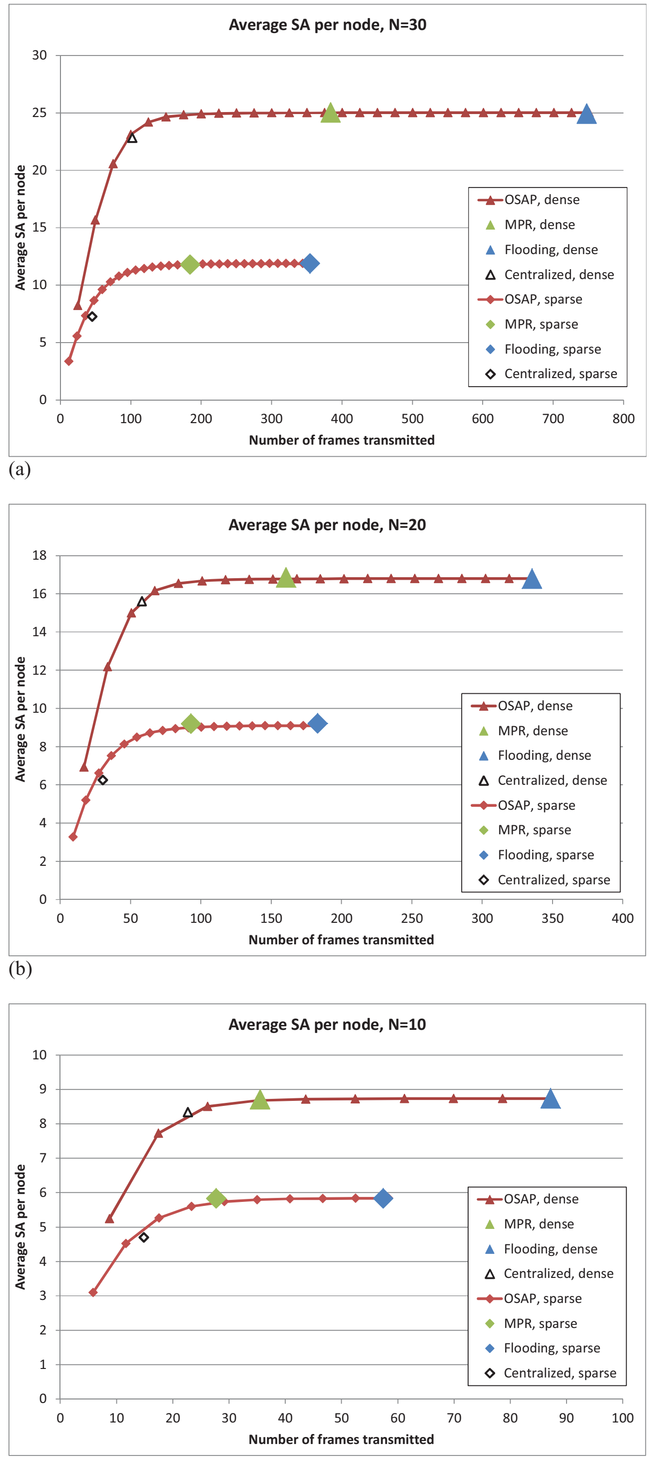

Fig. 1. This depicts the average SA per node, $\overline{SA}$, as a function of the number of network frames transmitted to achieve this level of $\overline{SA}$. Four SA dissemination algorithms are considered for networks with both sparse and dense conditions; (a) has N=10 nodes, (b) has N=20 nodes, (c) has N=30 nodes.

where each marker (triangles for dense networks or diamonds for sparse networks) denotes the measurement of $\overline{SA}$ as a function of transmitted frames after an "OSAP round". In all cases, we note that after only a few rounds, OSAP rapidly approaches an asymptote that represents the maximum average SA per node; this maximum value for $\overline{SA}$ is the same as that achieved by simple flooding and MPR broadcast. We point out here that the maximum value for $\overline{SA}$ does not reach N as one might initially expect; this can be explained by the realization that $\overline{SA}$ will only reach N in a fully connected network (i.e., a network where there exists at least one—potentially multi-hop—path from every node to every other node). In the uniformly distributed networks in our simulations, however, there are many instances of disconnected nodes. We would expect more disconnected nodes in sparse networks and this is precisely what is observed in the figure, with $\overline{SA}$ reaching a lower maximum in sparse networks.

Before discussing OSAP further, we briefly discuss the performance of the other three schemes. We note that MPR achieves maximum $\overline{SA}$ using 40-50% of the frames that would be required by a pure flooding scheme—this efficiency was expected, and is the reason why protocols such as OLSR use MPR to disseminate topology control messages. Additionally, we note that unlike MPR or flooding, the centralized broadcast scheme does not achieve the maximum $\overline{SA}$. This is a result of the potential for disconnected "master" nodes (which can happen easily in our uniformly distributed sceanrios), where the average $\overline{SA}$ in this case will be 1, with each node having node-specific-SA about only iself. The flooding, MPR, and OSAP schemes will still perform well if most nodes are connected; however, the centralized scheme will suffer if the master node happens to be disconnected.

The performance of OSAP when compared to the other three schemes is particularly notable. We observe that OSAP can achieve a value of $\overline{SA}$ that is close to the asymptotic maximum $\overline{SA}$ achieved by other schemes (where "close" is loosely defined here as greater than 90%) using significantly fewer frames than either pure flooding or MPR broadcast. In the "dense" networks, for all values of N (i.e., N = 10, 20, or 30) the OSAP scheme takes 4 or fewer rounds to reach better than 90% of the maximum $\overline{SA}$. In the sparse networks, more rounds are required; however, the number of frames required to reach 90% or better is still far less than the frames required by MPR for complete-SA.

To more clearly demonstrate the efficiency of OSAP (in terms of number of frames transmitted), we directly compare OSAP to the MPR broadcast scheme. As an aside, we note that MPR was selected as a basis for further comparison because MPR clearly outperformed flooding, and because the centralized scheme never reached maximum $\overline{SA}$ (not to mention that the centralized scheme performance plot lies below the OSAP curve in all but one of the scenarios in Figure 1).

For our comparison between OSAP and MPR, we define the term "frame ratio", which we denote by *FR(k)*, where

$$FR(k) = \frac{\text{\# OSAP frames after } k \text{ rounds}}{\text{\# MPR frames}}.$$

The value *FR(k)* denotes the ratio of the number of frames transmitted after k rounds of OSAP to the number of frames transmitted by an MPR broadcast[7]. We also define the terms $\overline{SA}(k)$, which denotes the average SA per node after k OSAP rounds and max($\overline{SA}$), which denotes the maximum asymptotic value of $\overline{SA}$. With these definition in mind we define $\overline{SA}_{\%}(k)$, where $\overline{SA}_{\%}(k) = \overline{SA}(k)/\max(\overline{SA})$; this tells us the fraction of maximum $\overline{SA}$ that is achieved after k OSAP rounds. For instance, in a dense network with N = 30, we find that $\overline{SA}_{\%}(4)$ = 92.4%, meaning that 92.4% of the maximum $\overline{SA}$ is achieved after 4 rounds of OSAP. We note that both *FR(k)* and $\overline{SA}_{\%}(k)$ can be obtained using values available in Figure 1.

Figure 2 shows a plot of $\overline{SA}_{\%}(k)$ versus *FR(k)*. To interpret this figure, consider the meaning of the circled point on the "N=30 dense" curve: this point has a value of *FR(k)* = 0.260 and $\overline{SA}_{\%}(k)$ = 0.924, meaning that for an N=30 dense network, OSAP achieves 92.4% of maximum $\overline{SA}$ using only 26.0% of the number of frames that MPR uses (to achieve maximum $\overline{SA}$).

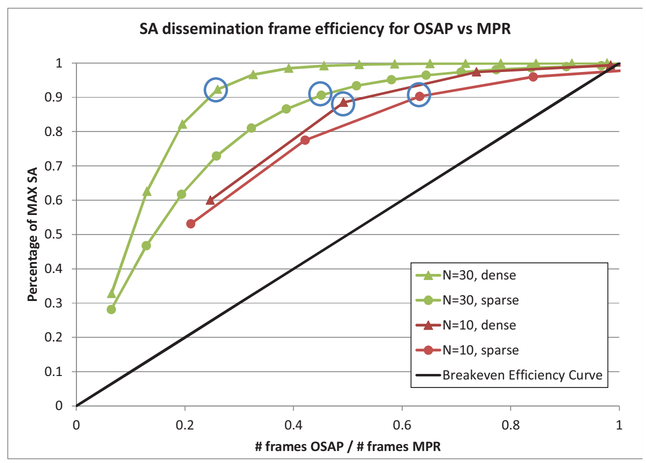

Fig. 2. Frame efficiency of OSAP compared to MPR broadcast scheme. Points lying above the breakeven efficiency curve use relatively fewer OSAP packets to disseminate relatively more SA than an equivalent number of MPR packets.

To further emphasize the frame efficiency of OSAP, we add to Figure 2 a "breakeven efficiency curve", which has a slope of 1. Essentially, for any point above the breakeven efficiency curve with a value of $\overline{SA}_{\%}(k)$, this amount of SA was achievable using less than $\overline{SA}_{\%}(k) \times$*(# frames MPR)*. We note that all the curves shown lie well above the breakeven efficiency curve for all but very high values of $\overline{SA}_{\%}(k)$. If the

[7] To make this clearer with an example, in the dense network scenario with N=30, simulation reveals that *FR(4)* = 26.0% since four rounds of OSAP take (on average) 99.9 frames and a complete MPR broadcast takes (on average) 383.9 frames.

goal is to obtain complete network SA, then an MPR broadcast outperforms OSAP, as seen where the OSAP curves cross the breakeven efficiency curve. However, this occurs only at very high values of $\overline{SA}_{\%}(k)$, where for many applications a lower value is more than adequate. This means that for more modest SA dissemination targets, OSAP is a very frame-efficient protocol compared to MPR (or flooding). To summarize, we note the circled points on all curves:

- For N=30, dense, OSAP achieves better than 92% maximum SA for $FR(k)$ less than 26%;
- For N=30, sparse, OSAP achieves better than 90% maximum SA for $FR(k)$ less than 46%;
- For N=10, dense, OSAP achieves better than 89% maximum SA for $FR(k)$ less than 49%;
- For N=10, sparse, OSAP achieves better than 90% maximum SA for $FR(k)$ less than 63%.

Figure 2 further reveals that OSAP is more efficient in denser networks and in networks with more nodes. The same effect was observed for N=20, but these curves were omitted to avoid cluttering the plot.

At this point, it is useful to revisit the parameters T_{OSAP} and T_{SA} introduced in section II. If OSAP were used to disseminate SA in a network, one way to select T_{OSAP} would be to consider a desirable value for $\overline{SA}_{\%}(k)$ and decide how frequently this amount of SA is required. For instance, assume we are happy obtaining better than 90% SA throughout the network every T_{SA} seconds (where T_{SA} was the period of the flooding, MPR broadcast, and centralized broadcast algorithms). For the "N=30 dense" network, this can be achieved with a value of k = 4 in $\overline{SA}_{\%}(k)$. Thus, we require 4 OSAP rounds every T_{SA} seconds—meaning that $T_{OSAP} = T_{SA}/4$ in this case. These 4 rounds generate a total of approximately 100 frames in T_{SA} seconds—far less than the 384 frames that would be generated by MPR to disseminate maximum $\overline{SA}$ in the same time period T_{SA}. In general $T_{OSAP} = T_{SA}/k$, where k is the number of rounds required to deliver the desired percentage of maximum $\overline{SA}$ in T_{SA} seconds.

As a final note, we point out that the simulations considered here did not include the additional overhead required by the MPR broadcast scheme to compute MPRs in the network, nor did it account for the overhead that the centralized broadcast scheme requires to find unicast routes. This is not a problem for OSAP, which can function in the absence of routes, and as mentioned in II.D, could actually be used to generate (partial) routes depending upon the rigor of the routing requirements.

IV. Application Under Tactical Scenarios

Having demonstrated the effectiveness and frame-efficiency of OSAP in the previous section, in this section we discuss a more "real-world" simulation that investigates how OSAP behaves in a typical mobile tactical deployment. With this second simulation, we were particularly interested in examining the level of "freshness" of the situational awareness in the network (i.e., how recent is the SA we collect). At this point, it is useful to highlight the fact that the MATLAB results in section III described *average* network behavior—for

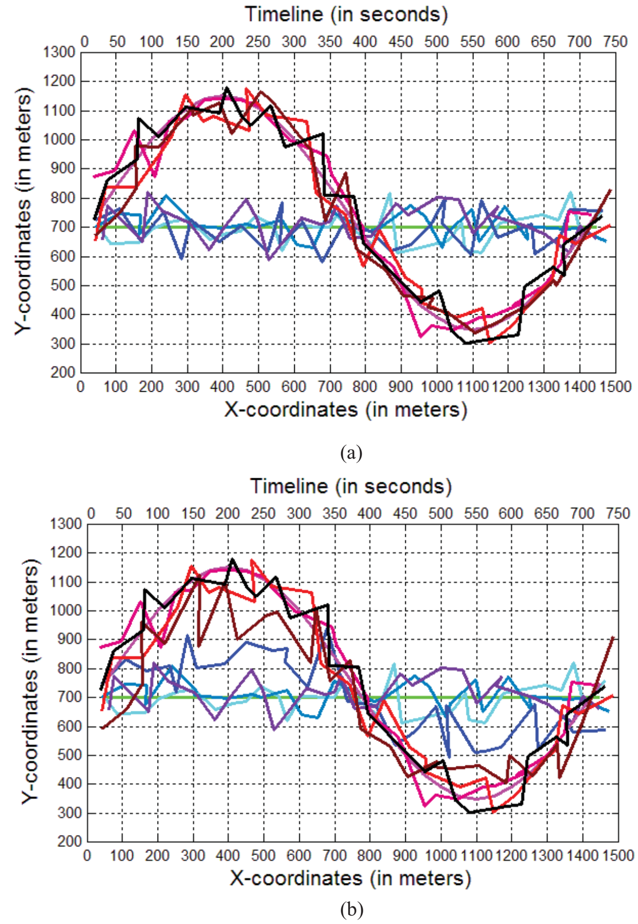

Fig. 3. Paths travelled by N=10 nodes where 5 nodes travel in a straight line and 5 nodes separate from the group; (a) represents tight group mobility, (b) represents loose group mobility.

example, on average an N = 10 dense network will achieve better than 90% $\overline{SA}$ (among connected nodes) in 3 rounds. With the simulations conducted in this section, however, we are interested not in long-term average behavior but in how the SA is distributed in the short-term (in a specific scenario) as time evolves. The freshness of a specific SA picture in the short-term is influenced by the mobility in the network, where some nodes may be out of reach for a period of time and then re-connect later on.

To conduct this simulation, we used the QualNet[8] simulation platform, which is a leading MANET simulator with rich PHY and MAC (physical and medium access control layer) features for high fidelity simulation of real-world wireless networks (for further details see [6]). We implemented OSAP as described in section II.D using a randomized round-robin scheduler to determine the initiating node in every round, where the time interval between rounds was given by T_{OSAP}. The initiating frame in each round was constructed such that it contained a timestamp indicating the

[8] The actual software used was EXata/Cyber 4.1, which is the enhanced version of QualNet with emulation and cyber security features.

time of issue and the node-specific-SA of the initiator; subsequent SA was appended as the frame propagated through the network as described in II.D. The timestamp allowed nodes to compute the age of any node-specific-SA they received in order to evaluate its freshness (as described in more detail later in this section).

We created a mobile tactical scenario consisting of $N = 10$ nodes, where these nodes were divided into two groups, each with 5 nodes. The nodes all travelled at a constant speed following predefined waypoints in an 1000 × 1500 meter area, travelling from one side to the opposite side (from west to east) over a 750 second interval. The paths followed by each of the nodes are depicted (as x-y coordinates over time) in Figure 3. Figure 3(a) depicts nodes travelling in close formation ("tight" mobility), while Figure 3(b) depicts nodes in a looser formation ("loose" mobility), where one node in each group is designated as a "scouting node" with a higher degree of mobility than the other nodes in the group. The transmit power, antenna height and efficiency in the QualNet PHY layer was configured such that the radio range is approximately 85 meters; this results in a complete disconnect between the two groups at certain times. Further details of the simulation implementation with additional experiments and results are given in [7].

As discussed in II.D, in a fully connected network with N nodes, OSAP can *guarantee* that each node will have every other node's SA information after N rounds—i.e., within a time period of $N \times T_{OSAP}$ seconds. In practice, however, OSAP should disseminate SA more quickly than this: as shown in section III, for dense networks we expect that most connected nodes will have SA about most other connected nodes in as little as 4 rounds, or $4 \times T_{OSAP}$ seconds. With this in mind, we select a value of $T_{OSAP} = 2.5$ seconds for our simulation with the aim of disseminating "most" SA among connected nodes within 4 × 2.5s = 10 seconds.

To evaluate the freshness of node-specific-SA information, we define several terms as follows. If any node possesses node-specific-SA about any other node (call this other node n_k), then we use $t_{SA}(k)$ to denote the age of the node-specific-SA for n_k; this value is computed by subtracting the most recent SA timestamp (in a received OSAP frame containing SA about n_k) from the current time. We define two timing thresholds at a receiving node: t_{fresh} and t_{stale}. The SA for n_k is said to be "fresh" if $t_{SA}(k) \le t_{fresh}$; "stale" if t_{fresh} $t_{SA}(k) \le t_{stale}$; and "expired" if $t_{SA}(k) > t_{stale}$. Fresh SA is intended to denote SA that has very recently arrived and is up-to-date; stale SA is intended to denote SA that is still useful, but has not been updated recently; and expired SA is intended to denote SA that is too old to be of use.

In the simulation, we defined a node's SA as "expired" if the node's SA was not updated after 10 rounds of OSAP; thus, we set $t_{stale} = 10 \times T_{OSAP} +$, where is a small interval to account for the processing time of the last OSAP round (i.e., to make sure all frames have been processed). In our scenario we had $T_{OSAP} = 2.5$ s, and = 1.5 s; hence $t_{stale} = 26.5$ seconds. The value of t_{fresh} can be selected to meet whatever freshness criteria is desired; in our scenario, since the intent was to deliver "most" SA within four OSAP rounds, we selected $t_{fresh} = 4 \times T_{OSAP} = 10$ seconds.

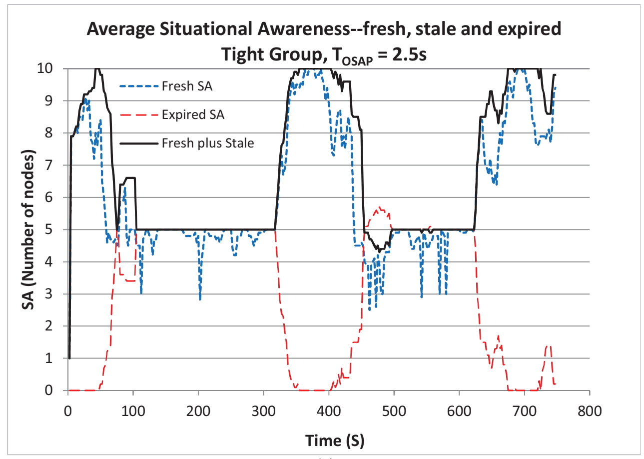

(a)

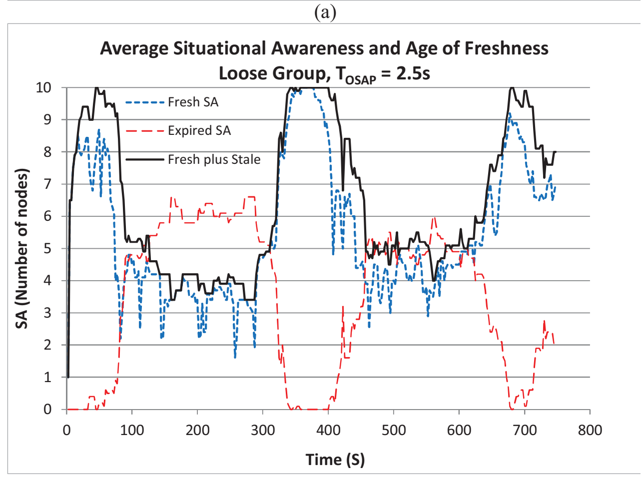

(b)

Fig. 4. Average number of "fresh", "stale", and "expired" SA per node; (a) represents tight group mobility, (b) represents loose group mobility.

After each scenario run in QualNet, we collected the following statistics to verify the effectiveness of the SA passing scheme:

1. Average number of fresh/stale/expired node-specific-SAs at each node as a function of time;
2. Average age of fresh/stale SAs at each node as a function of time;

The statistics above are illustrated in Figures 4 and 5. In each figure, sub-figure (a) represents the tight mobility case and sub-figure (b) represents the loose mobility case.

Figure 4 depicts the average number of fresh, stale, and expired node-specific-SAs recorded per node as a function of time. The figure also depicts the sum of fresh and stale SAs to show "usable" SA, as by our definition a stale SA (less than 26.5 seconds old) is still recent enough to be of use. In Figure 4(a), we can see that during periods of the simulation when the network is fully connected, the average number of fresh SAs rapidly increases and then operates at a level between 8 and 10, meaning nodes have usable information about 80% to 100% of their peers. When the two groups diverge and the network becomes disconnected, the average count of fresh plus stale SA steady-states at about 5, with most of this comprised of fresh SA, which is near 5 for much of the disconnected period; the average count of expired SAs is approximately 5 as well. This is sensible since the two groups of 5 nodes are disconnected, meaning that the nodes have not received updates about the members of the other group. In the loose mobility scenario results shown in Fig. 4(b), we observe more volatility in the results, with generally inferior SA dissemination as compared to the tight group. This is explained by noting that the scouting nodes frequently move so far out of formation that they are disconnected from any other nodes and thus lose opportunities to both receive SA and to append their own SA to OSAP frames, effectively resulting in two groups of 4 nodes and two solitary nodes.

In Figure 5, further insight is provided about the freshness of the node-specific-SA. The red curve in the figure depicts the average age of the fresh SA (i.e., it does not include the age of stale or expired SA), while the dashed blued curve shows the average number of fresh SA. Taken together, these curves tell us how much fresh SA we have and how old it is. Of note is that during the intervals when the two groups are fully connected (i.e., from 0-50 seconds, 350-400 seconds, and 700-750 seconds), the average age of the fresh SA is lower than in the intervals when the two groups are disconnected. This can be explained by realizing that in the connected interval we are guaranteed to have one OSAP round in the network every T_{OSAP} seconds, whereas in the disconnected interval the nodes will see (on average) half as many OSAP rounds due to the randomized round-robin scheduler. Related to this phenomenon is the fact that the average age of the SA in the disconnected interval is "noisier" when compared to the average age of the SA in the connected interval. Again, the smoothness of the average age in the connected interval is a result of the more frequent OSAP rounds. We also note that in the loose mobility scenario, the average age of the fresh SA is noisier, with larger spikes as a result of the irregular connectivity of the scouting nodes.

The results in this section indicate that our SA passing scheme is practical and behaves as predicted in a real-world scenario on a simulated platform. As with any SA dissemination scheme, network coverage depends upon the connectivity and mobility of the tactical network. We refer to [7] for further QualNet simulation scenarios and discussions.

V. Conclusion

We have proposed OSAP (Opportunistic Situational Awareness Passing), a method to disseminate short bursts of situational awareness (SA) information in a mobile ad hoc network (MANET). We showed through simulation that OSAP offers a more efficient method of disseminating SA (in

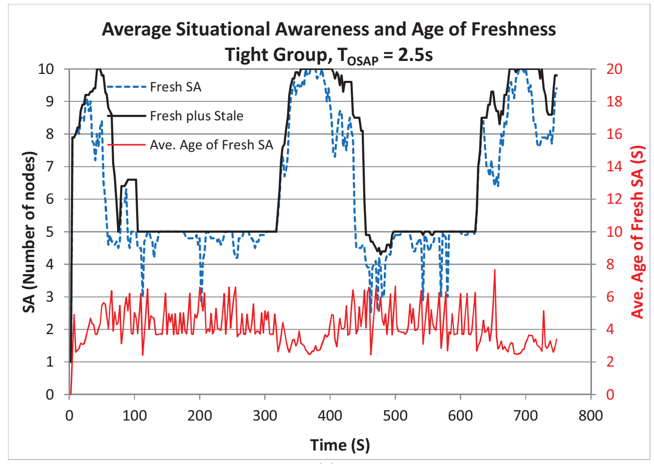

(a)

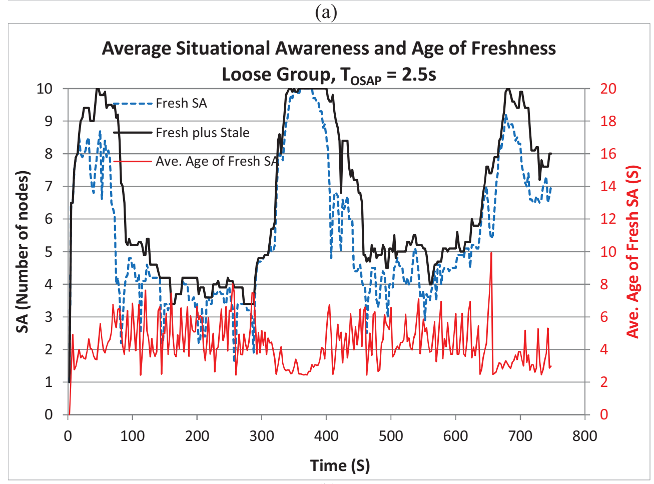

(b)

Fig. 5. Average age of fresh SA per node, measured in seconds is depicted on the secondary axis; the primary axis shows the average number of fresh SA; (a) represents tight group mobility, (b) represents loose group mobility.

terms of number of transmitted frames required) when compared to flooding, centralized broadcasts, and multipoint relay (MPR) broadcasts. These efficiencies are contingent upon relaxing the requirement to deliver SA about all nodes and being satisfied with delivering SA about "most" nodes (where this can typically exceed 90% of connected nodes). OSAP can operate without any pre-existing routing protocols or MPR-discovery schemes; it does, however, require prior knowledge of the number of nodes in the network—though this requirement is likely not difficult to achieve in planned military deployments. Simulations of uniformly distributed networks in MATLAB indicated that for dense MANETs, as few as 3 or 4 OSAP rounds were sufficient to deliver 90-95% of the SA about connected nodes using far fewer frames than competing schemes. Simulations in QualNet demonstrated the feasibility of OSAP in more realistic tactical scenarios and showed the effect of mobility and topology on the freshness of SA in the network.

We are actively researching refinements of OSAP to improve its effectiveness and to further understand how well it operates in real-world tactical scenarios. Of particular interest is the effect of intermittent link quality, where nodes within range will nevertheless have a non-zero probability of dropping frames. In addition, the schedulers for OSAP frames presented here are likely sub-optimal; other "smart" scheduling schemes that adjust when a group of nodes become disconnected are of interest. We note as well that since OSAP appends information to the frame as it propagates through the network, the effect of frame length and maximum frame size must be investigated to understand the size limits of networks in which OSAP can safely operate. Finally, adopting concepts from MPR flooding to OSAP may assist in reducing average OSAP frame size to allow for additional efficiencies.

REFERENCES

[1] Clausen, T. and Jacquet, P., Optimized Link State Routing Protocol (OLSR), IETF RFC 3626, October 2003.

[2] Hammel, T., Farther Out Networking, August 7, 2013. Retrieved from www.darpa.mil/WorkArea/DownloadAsset.aspx?id=2147487187.

[3] Bjorkbom, M., et al., Localization Services for Online Common Operational Picture and Situation Awareness. Access, IEEE, 2013. 1: p. 742-757.

[4] Larsen, E., Landmark, L., Pham, V., Kure, O., and Engelstad, P.E., Optimized group communication for tactical military networks, in *MILCOM 2010*, San Jose, CA, USA, November 2010.

[5] Qayyum, A., Viennot, L., Laouiti A., Multipoint Relaying for Flooding Broadcast Messages in MobileWireless Networks, In Proc. of the 35th Annual Hawaii International Conference on System Sciences (HICSS'2002), Hawaii, USA, January 7-10, 2002.

[6] Scalable Network Technologies Website, Qualnet product description, available at http://web.scalable-networks.com/content/qualnet. Accessed on Feb 20, 2014.

[7] Li, M., Mason, P.C., Salmanian, M., and Brown, J.D., Simulation implementation and performance analysis for situational awareness data dissemination in a tactical MANET, *Proc. Of the 2014 Winter Simulation Conference*, Savannah, GA, USA, 2014 (in press).