Predicting Peer Interactions for Opportunistic Information Dissemination Protocols

Marco Marchini[1,2], Mauro Tortonesi[2], Giacomo Benincasa[1], Niranjan Suri[1,3], Cesare Stefanelli[2]
[1]Florida Institute for Human & Machine Cognition (IHMC), Pensacola, FL, USA
[2]Department of Engineering, University of Ferrara, Ferrara, Italy
[3]U.S. Army Research Laboratory, Adelphi, MD, USA

***Abstract*—Tactical edge networks provide one of the most challenging environments for communications, which significantly complicates the development of efficient and robust information dissemination solutions. In our previous work, we found that exploiting highly mobile nodes, such as Unmanned Air Vehicles, with cyclic mobility patterns, as message ferries can significantly improve the performance of information dissemination solutions. However, our experience demonstrated that robust forecasting mechanisms are essential in order to withstand frequent changes in the mobility patterns of message ferrying nodes. This paper presents an extension of the adaptive node presence forecasting component developed for DisService, a Peer-to-peer information dissemination system, that provides estimates of tolerance and accuracy of node mobility forecasts. We tested the extended forecasting mechanism in a simulation environment and found that it can lead to significant improvements in the timeliness and reliability of information dissemination.**

Opportunistic communications; information dissemination; tactical networks, statistical learning.

I. INTRODUCTION

Tactical edge networks provide one of the most challenging environments for communications, with mobile nodes connected via limited bandwidth and highly variable latency wireless ad-hoc links in RF environments. The dynamic behaviour of nodes leads to a frequently changing network topology and widely varying loads being placed on the network by users and applications.

Information dissemination is a critical function that enables tactical network applications, such as sensor data acquisition, to distribute information to multiple peers across the network. Because of the aforementioned challenges, and of the potentially high number of peers that have to be reached, it is extremely difficult to implement efficient dissemination algorithms that provide timely and reliable information delivery in this type of environments.

The peculiar nature of tactical networks calls for ad hoc information dissemination algorithms that can dynamically adapt to the current network topology and conditions. More specifically, the high mobility of some of the nodes that are typically found in tactical networks, such as Unmanned Air Vehicles, suggests the opportunity to use those nodes as message ferries to improve the performance of the dissemination process. This requires explicit support at the middleware level to monitor and process node mobility information, identify common behaviors, and produce reliable forecasts of future contacts with message ferrying nodes.

In our previous work, we have developed a node mobility monitoring and presence forecast solution to detect cyclic node mobility patterns and to leverage that information to support information dissemination [1]. We have integrated the forecast mechanism in DisService, an information dissemination middleware purposely designed for tactical network applications [2] [3].

While the mechanism that we developed proved capable of successfully detecting cyclic movements and predicting the next appearance of the moving node, additional experiments demonstrated that for many applications, this knowledge is not enough and might actually be detrimental. In fact, our experience demonstrated that applications that relied exclusively on future node presence forecasts to schedule their message transmissions to message ferries incurred severe performance losses in case of changes in the mobility patterns. The frequent redeployment and repurposing of highly mobile tactical network nodes, such as UAVs, in response to changing mission objectives required the consideration of more resilient solutions that could withstand changes in the mobility patterns of message ferrying nodes.

To this end, there is the need to provide applications with additional information about the reliability of the forecasts in order to support their decision making. This requires the development of more sophisticated forecasting mechanisms that continuously evaluate their performance and estimate the level of confidence for the forecasts.

This paper presents an extension of the adaptive node presence forecasting component for the DisService project that provides estimates of the forecast *accuracy* and *tolerance* attributes. This is essential to provide applications with contextual information that they can leverage when deciding whether to trust a forecast or not.

We tested the extended forecasting mechanism of DisService in an NS 3 simulation environment and found that providing applications with an estimate of the accuracy of forecasts can lead to a significant increase in the timeliness and reliability of information delivery.

II. TACTICAL EDGE NETWORKS

Tactical edge networks are highly heterogeneous and dynamic environments where both mobile, e.g., soldier platoons and ground vehicles, and stationary, e.g., ground sensor systems and Tactical Operation Centers (TOCs), operate, exchanging data and commands to support the accomplishment of mission objectives. This scenario typically involves many concurrently running applications, such as Blue

Force Tracking (BFT) - applications that provide situational awareness information regarding the presence and location of friendly forces, remote unmanned vehicle control, and sensor data mining/fusion, which run essential tasks and compete for the scarce bandwidth and computational resources.

Unmanned Aerial Vehicles (UAVs), and other airborne assets such as Joint Surveillance Target Attack Radar System (J-STARS), are becoming increasingly prevalent in tactical networks, as they are extremely effective to realize battlefield monitoring, to process information and carry it between disconnected portions of the network, and in general to operate in highly hazardous areas where human presence would be impossible. UAVs can fly autonomously or be piloted remotely, can be expendable or recoverable, and can carry a lethal or nonlethal payload.

Tactical Unattended Ground Sensors (T-UGS) are small ground-based sensors that collect intelligence through seismic, acoustic, radiological, biological, chemical, and electro-optic means. These sensors are networked devices that provide an early warning system to supplement a platoon-sized element and are capable of remote operation. By using P2P models combined with mobile ad-hoc network technologies, sensors can use high-speed short-range radios to exchange and exfiltrate data to other sensors and nearby units. With a high-speed P2P link, a sensor can send high-resolution imagery and motion video that would be impractical to transmit over a low-speed satellite link.

Finally, soldier-carried nodes, such as wearable computers or PDAs, are the last type of nodes found in tactical networks, as they can receive and send information to other entities in the environment. While on a mission, soldiers need to access a variety of information including maps, aerial reconnaissance, various sensor data, intelligence reports, and blue and red force tracking. Some of this data may be preloaded onto the nodes and some may become available later.

The tactical network scenario shares many similarities with non-military applications such as disaster recovery. As UAV technologies improve and become more affordable due to economies of scale, their adoption in civilian applications is expected to increase, spreading well beyond emergency response scenarios. Even now, UAVs are employed for monitoring critical engineering structures (dams, etc.), for search and rescue operations in difficult to reach or hazardous locations, for mail delivery in uninhabitable places (off-shore platforms, polar caps, etc.), for livestock monitoring, etc.

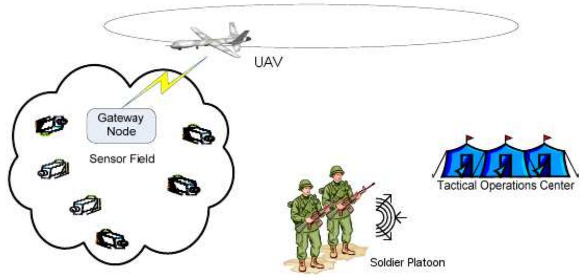

Figure 1. The tactical network scenario.

III. INFORMATION DISSEMINATION IN TACTICAL NETWORKS

Support for information dissemination is essential in tactical networks. The peculiar characteristics of tactical networks call for disruption-tolerant approaches to information dissemination and opportunistic network exploitation. In fact, often information must be delivered to nodes that periodically disconnect from the rest of the network, requiring reliability mechanisms such as caching and periodic retransmission of important (non-obsolete) data. For maximum efficiency, the data dissemination system should not only be capable of withstanding node mobility, but it should also take advantage of it whenever possible.

Tactical applications present multiple patterns of data dissemination. For instance, BFT data is transmitted from each node to every other node in a many-to-many pattern. Sensor fusion data requires many nodes (sensors) to transfer data to one node (gateway sensor or fusion node) and then onto some consumers in a many-to-one and one-to-few pattern. The above requirements typically translate to multiple types of information dissemination services operating concurrently.

There is an opportunity to exploit common patterns in node mobility and information to improve the timeliness and availability of information being disseminated. This requires learning mechanisms that can process node mobility and data/service usage information, identify common behaviors, and produce reliable forecasts that can be used as the basis for decision making in information caching and routing.

More specifically, node mobility forecasts represent an interesting asset to rely upon in order to improve information delivery. In this context, forecasts on future node presence of message ferrying nodes would be important for various reasons. The first one is for adaptive beaconing. Carrying out targeted checks for the presence of a UAV, rather than doing them continuously, allows gateway nodes to remain in a dormant state for a longer amount of time, because they are activated less frequently. This allows resource-constrained nodes publishing information, such as sensor nodes, to save energy they would otherwise spend with frequent network broadcasts.

The second reason is for improved caching support. Leveraging forecast information, nodes can better decide which peers are more likely to be encountered and therefore which among the cached messages are more likely to be delivered. This notion can be exploited to decide the subset of cached messages that should be kept, and the subset that can be dropped. Improved cache management can lead to a significant improvement on information delivery, especially when the storage capabilities of the peers are scarce. In addition, it is possible to further enhance the caching process by cross-correlating node presence forecasts with metadata about message validity time.

Finally, the last reason is to support data transformation. In particular, if the forecasts also include the duration of expected contact, publisher and gateway nodes can predict the overall capacity that will be available to move data. This can be derived as a function of the channel bandwidth and the duration of contact. Nodes can then use this information to prioritize the data that is replicated or to transform the data (for example, by reducing the size of imagery) to make sure that transfers will be successful within the time available.

Being able to take advantage of node presence predictions requires explicit support at the middleware level. The information dissemination middleware on the publishing nodes should continuously check the presence of message ferrying nodes in communications range, extrapolate cyclic mobility patterns, and provide applications with a time estimate for the next contact with message ferries.

IV. Forecasting Node Mobility in DisService

This Section provides an overview of the forecasting mechanism implemented in DisService and discusses the network monitoring phase that collects information to use for predictions, as well as the forecasting model. For additional information, the reader is referred to [1].

A. Network Monitoring

To collect information to feed the forecast model algorithm, each DisService node continuously monitors and records the *contact window* metric, which represents a particularly interesting source of information about node mobility. We define a contact window between nodes N1 and N2 as the tuple containing the start time and the duration of a time interval during which N1 and N2 are in communication range and can be considered neighbors.

Contact window-based metrics can be easily measured by analyzing presence messages that DisService-enabled nodes periodically exchange as they explore new neighborhoods. Each node stores a contact window history table, containing information about its current and previous contact windows with neighbor nodes. Nodes create (or update) contact window as they receive presence messages from their neighbors and close contact windows when they do not receive any presence message for a configured length of time.

By analyzing past contact window information, the forecasting algorithm is capable of detecting cyclic patterns in node mobility and to predict the next time at which a specific mobile node will come back in communication range, as well as the expected duration of that contact.

B. Forecasting Model

DisService adopts a forecasting model based on the Exponentially Weighted Moving Average (EWMA) algorithm. The EWMA algorithm is simple enough that it can be easily implemented in resource constrained nodes such as ground sensors.

The lightweight computational requirements of the EWMA algorithm also enable the building of more sophisticated forecasting solutions on top of it. In fact, the latest version of DisService adopts an adaptive forecasting model that implements 3 different EWMA estimators with different smoothing parameters (0.2, 0.5, and 0.8, respectively) at the same time, and dynamically switches to the most accurate one.

V. Complementing Predictions with Tolerance and Accuracy Estimates

Our experience with the development and deployment of information dissemination services on tactical networks demonstrates that providing applications with forecasts on node presence is not enough to enable them to make good information dissemination decisions.

Given the peculiarity of the tactical network scenario, forecasts that simply produce point estimates on the future presence of given nodes are likely to be of very limited use. In fact, applications that leverage future presence forecasts of message ferrying nodes might base their caching and forwarding decisions upon unreliable information. This is especially likely in case of changes in the mobility patterns of message ferrying nodes.

To address this problem, we have extended the DisService forecasting model discussed in the previous section to complement node presence forecasting information with additional attributes, such as *forecast accuracy* and *tolerance* estimates.

Let us analyze how DisService calculates those estimates on a node A that wants to predict the start time of the next contact with a node B. First, DisService calculates the tolerance estimate, which is obtained by taking a portion, controlled by a system parameter β (with $0 < \beta < 1$) of the average contact window duration, calculated over the last N records:

$$tolerance = \beta \frac{\sum_{i=M-N}^{M} d_i}{N}$$

The calculations of the accuracy estimate are significantly more complicated. Let's call P_S the instant in which the check is scheduled and C_S the start time of the actual contact. Then, we define the displacement Δ as the difference between C_S and P_S:

$$\Delta = C_s - P_s$$

Since Δ is the difference between the actual start time of a contact and the predicted start time of that contact, this quantity can be used to estimate the inaccuracy of the forecasts. In particular, a first estimator could compute the inaccuracy as the arithmetic average of the last N displacements:

$$inaccuracy^{(1)} = \frac{\sum_{i=M-N}^{M} \Delta_i}{N}$$

where M is the number of contacts that node A has registered with node B. However, our extensive experiments demonstrated that the value obtained in this way is a poor estimate of the real inaccuracy, because it considers all the displacements with the same weight. So, if a prediction was well in advance, it could be compensated by a previous prediction that was very late and the final inaccuracy would incorrectly consider the sequence of predictions as accurate.

In order to give more importance to recent displacements, we then devised a second estimator that applies the EWMA algorithm on the last N displacements, with a relatively high value for the smoothing parameter:

$$inaccuracy^{(2)} = EWMA_{\alpha=0.5}\left(\sum_{i=M-N}^{M} \Delta_i\right)$$

However, experiments demonstrated that despite the major importance given to the most recent displacements, this second estimator presents the same problems of the previous one.

We then devised a third estimator that considered the absolute value of displacements. In addition, the estimator checks whether the predicted instant P_S actually lies within the subsequent contact window. Only if a prediction is wrong, the corresponding displacement is evaluated in the inaccuracy formula; otherwise, a null displacement is considered. So we have:

$$inaccuracy = EWMA_{\alpha=0.5}\left(\sum_{i=M-N}^{M} \Delta_i^*\right)$$

$$\Delta_i^* = \begin{cases} |\Delta_i| & if \;\; prediction\, wrong \\ 0 & if \;\; prediction\, correct \end{cases}$$

Our experiments demonstrated that this third algorithm produced good estimates for the inaccuracy of forecasts. We then decided to adopt it in our extended forecasting mechanism.

To produce an estimate of the forecast accuracy, we need to consider the tolerance estimate as well. We then defined the error measure as:

$$error = \max(inaccuracy - tolerance, 0)$$

We use the error measure above to compute the accuracy of predictions, applying an exponential mapping:

$$accuracy = e^{-error}$$

In this way, we have an accuracy that is inversely proportional to the error committed and that conveniently lies within the (0; 1] range.

The current accuracy measures of predictions for every known node are maintained in the World State component of DisService, for each node that has prediction capability. When a node must decide whether predictions-based dissemination strategy is a good choice, it can check if the accuracy measure is greater than a desired threshold, e.g., 0.75.

Once a prediction for the next contact with the node B has been computed by node A, that prediction is stored in the World State component of DisService on node A and the next check for the presence of B is scheduled at the predicted time plus the tolerance estimate.

VI. Adaptive Communication Decisions

DisService provides a customizable controller architecture that allows the dynamic configuration of behaviors such as caching, replication, and forwarding. Leveraging contact window information, the strategy manager component needs to choose whether to discard incoming messages, put them in the message cache, and/or forward them to other neighboring nodes.

The different types of nodes in a tactical network follow different patterns of communication, depending on their purpose, technical features, and mobility. Every node should use a specific information dissemination strategy, which can be based on predictions or not as needed.

That is, computation of forecasts makes sense only to some nodes and only in specific conditions. Two important data in this contest are the availability and the accuracy estimates of forecasts. Thanks to this information, the Strategy Manager component of DisService can decide if it is appropriate to use a prediction-based dissemination strategy or a standard epidemic one.

If predictions are not available or do not have a sufficient level of accuracy, DisService disseminates messages with an epidemic strategy. Otherwise, forecasts can be used in different ways depending on the node that computed them. In fact, predictions are not always available; for instance, when the necessary information for the computation is not available. In this case, the sender needs to do periodic and more closely spaced checks in order to verify the presence of the node to which we are interested in sending messages. Also, predictions can be available, but incorrect. Therefore, we can not blindly rely on the predictions if they are present, but we also need to check their level of accuracy and decide whether it makes sense to use them.

Notice that forecasts on future contacts with other nodes might be interesting not only to drive message caching and processing decisions. In fact, some nodes may be interested in knowing when another node will be in communication range. Then, these nodes can use prediction for the next contact with the entity of interest to know when trying to communicate with that entity. This way of operating avoids continuous probing or monitoring for the presence of the recipient, thus saving energy.

Other nodes may have a lot of data to send, but narrow contact windows, which does not allow the transmission of all this data. In this case, forecasts for the duration of the contact window can be exploited to send selected data, so that they are delivered in time. This way will prevent the transmission of data that would not be received by the addressee because the contact was already over.

VII. Experimental Results

We evaluated the extended DisService node mobility forecasting mechanism by testing it in a simulated environment reproducing a typical battlefield scenario. We used Network Simulator 3 [4], and in particular version 3.11, for all the experiments presented in this paper.

In the reproduced evaluation scenario, two platoons of soldiers (of 18 and 20 units respectively) move in groups across the battlefield. Two groups of stationary T-UGS ground sensor nodes (each composed by 20 units) collect data and disseminate it to subscribers. In each group of nodes, a special node acts as leader and operates as a gateway with respect to other nodes in communications range. Finally a UAV speeds up the data dissemination process by carrying data between the sensor field and the TOC.

All the nodes are connected through wireless 802.11 connections at 6Mbps with a non QoS-enabled MAC layer; the TOC is also connected to each of the soldier patrols through a point-to-point tactical radio link at 1.5Mbps.

In each experiment, the TOC is placed in the same fixed position, while the sensors are positioned randomly in two 300x300-meter areas on the opposite sides of the battlefield.

The UAV follows a modified Random Waypoint Mobility model, with three waypoints set respectively near the TOC and in the middle of each ground sensor field. More specifically, our model allows for configurable-length pauses in the route of the UAV, at selected waypoints. This allow the simulation of particularly long delays, representing unexpected events or changes in the mission objectives, that force the UAV to stop in a certain place, thereby breaking its cycling mobility pattern. After the break, the UAV resumes its journey with the same mobility model that was in place before the stop.

In the first experiment, we tried to use the forecasts computed by DisService without considering tolerance or accuracy estimates. This is to show the limits of using point predictions, without the support of accuracy information. In particular, in these simulations, gateway sensor nodes calculate forecasts for the next contact with the UAV, which has a regular mobility pattern, to detect when to wake up and send data to the airborne vehicle. Each gateway sensor registers the instant predicted and those in which a contact actually occurred.

The results shown in Table I demonstrate that the exact value of the forecast is often in advance of the actual start of the next contact. This is due to small delays that inevitably occur in the motion of the UAV. As a result, there is the need to consider an appropriate level of tolerance for the values produced by the forecasting algorithm, in order to realize a resilient forecasting solution.

TABLE I. EVALUATION OF POINT FORECASTS.

Prediction value (ms)	*Next contact window (ms)*	*Point forecast correct?*	*Error (ms)*
168755	171000 - 193800	No	-2245
306081	304000 - 327200	Yes	
370156	370600 – 392800	No	-444
506460	502600 – 527200	Yes	
569647	569000 - 590800	Yes	
695448	698400 - 720400	No	-2952
763626	766800 - 790200	No	-3174
835000	835200 - 858200	No	-200

However, introducing tolerance estimates is not enough to enable applications to use all the forecast information effectively. In fact, the major disadvantage of forecasts providing point estimates is that, even considering an adequate tolerance level, they could be applied inappropriately, thus leading to performance losses instead of performance gains. For instance, let us consider how gateway sensor nodes could use prediction information about the future presence of UAV. In case of inaccurate forecasts, a gateway would wake up at the (wrong) predicted instant and send its data, which probably will be lost. Augmenting forecasts with information about their accuracy, instead, would allow the sensor to realize when predictions start to be unreliable and, if so, to change its dissemination strategy.

The second experiment aims at demonstrating the robustness of our forecasting mechanism to contingencies that may modify or delay cyclic mobility patterns of message ferrying nodes. For this purpose, we reenabled the calculations of tolerance and accuracy estimates and modified the mobility model for the UAV by inserting a long pause in the middle of the simulation. In particular, we configured the UAV to stop by the TOC for a 60 second interval. After that, the UAV restarts moving with the previous mobility pattern.

Table II presents the values of inaccuracy, error and accuracy estimates that we registered for each future prediction. It can be noted that, for both the gateway sensor nodes, the UAV pause occurs after the sixth prediction performed. The following prediction, which is based on past contact history, is therefore completely wrong, because the UAV returns in communications range after a much larger period of time.

TABLE II. FORECASTS EVALUATION PARAMETERS.

Sensor	*Prediction*	*Inaccuracy*	*Error*	*Accuracy*
1	1	0	0	1
	2	0.198	0	1
	3	0.039	0	1
	4	0.007	0	1
	5	0.356	0	1
	6	0.071	0	1
	7	43.78	41.78	7.2E-19
	8	8.754	6.754	1.17E-3
	9	1.750	0	1
	10	0.349	0	1
	11	0.069	0	1
	12	0.013	0	1
2	1	0	0	1
	2	1.483	0	1
	3	0.296	0	1
	4	0.059	0	1
	5	0.011	0	1
	6	0.002	0	1
	7	39.57	37.57	4.85E-17
	8	7.912	5.912	2.7E-3
	9	1.582	0	1
	10	0.316	0	1
	11	0.063	0	1
	12	0.012	0	1
	13	0.002	0	1

However, the algorithm proves robust to changes in mobility pattern of other nodes, and very reactive as well. In fact, only the first two predictions are wrong; after those, the values of forecasts return to be accurate. In addition, by providing accuracy estimates as well, the extended forecasting mechanism enables the development of adaptive dissemination strategies, that leverage on forecasts only when they are reliable and ignore them during reassessment periods.

The third experiment aims at demonstrating how forecasts could be effectively exploited to save energy in those nodes that have a limited amount available. More specifically, we consider gateway sensor nodes that leverage on predictions for the next contact with UAV to save energy by entering a dormant state and waking up only when necessary.

To evaluate the potential energy savings of adaptive beaconing techniques, we performed several simulations, recording for each one the number of times the gateway sensor wakes up to check for the UAV in a 30-minute interval. Results are shown in Table III. As expected, when increasing the tolerance the reliability of the predictions also increases and the number of checks needed decreases.

Without exploiting forecast information, sensor nodes would have to continuously and frequently probe for the UAV presence, thereby remaining always active. Assuming periodic checks every 200 milliseconds would lead to 9000 checks in the 30-minute interval. Even considering the worst case, that of a null tolerance for the predictions, this leads to a very significant reduction in the number of controls, of approximately 95%.

TABLE III. CHECKS FOR UAV PRESENCE PERFORMED.

Tolerance (ms)	*Number of Checks*
0	459
1000	430
2000	405
3000	403
4000	389
5000	379

VIII. RELATED WORK

Several works have focused on opportunistic networking from the perspective of sharing resources, e.g., internet connectivity, between limited devices [5] [6]. Some of these have also studied social aspects that emerged from the contacts between nodes [7]. Other works have analyzed local link and mobility metrics in order to discover information about the state of the network and take advantage of it [8]. In this research area, several studies have addressed the problem of improving information dissemination on delay-tolerant vehicular networks or MANETs, by analyzing path likelihoods [9] or betweenness centrality [10].

DisService takes a different approach compared to the above mentioned research projects. In fact, it attempts to provide applications with reliable forecasts on future presences of highly mobile nodes in order to enable them to efficiently rely on message ferries to improve information dissemination.

IX. CONCLUSIONS AND FUTURE WORK

The results presented in this paper demonstrate that contextualizing node mobility forecasts with additional information can effectively improve information dissemination. The extended forecast mechanism implemented by DisService effectively provides forecasts of future contacts with message ferrying nodes that are resilient to abrupt changes in mobility patterns.

Future version of DisService will attempt to leverage contact information to try to infer further information about the network state and topology. For instance, the analysis of the number of different nodes encountered, and in particular of those encountered repeatedly, might produce valuable insights to further optimize the information dissemination processes.

X. ACKNOWLEDGEMENTS

This research has been supported by the U.S. Army Research Laboratory, under cooperative agreement W911NF-11-2-0095 and by the Office of Naval Research and the Space and Naval Warface Systems Command, under contract N66001-11-C-0045.

REFERENCES

[1] A. Mazzini, C. Stefanelli, M. Tortonesi, G. Benincasa, N. Suri, “DisService: Network State Monitoring and Prediction for Opportunistic Information Dissemination in Tactical Networks”, in *Proceedings of IEEE MILCOM 2010 Military Communications Conference*, 31 October - 3 November 2010, San Jose, CA, USA.

[2] N. Suri, G. Benincasa, S. Choy, S. Formaggi, M. Gilioli, M. Interlandi, J. Kovach, S. Rota, R. Winkler, “DisService: A Peer-to-peer Disruption Tolerant Dissemination Service”, in *Proceedings of IEEE MILCOM 2009 Military Communications Conference*, 18-21 October 2009, Boston, MA, USA.

[3] N. Suri, G. Benincasa, M. Tortonesi, C.Stefanelli, J. Kovach, R. Winkler, R. Kohler, J. Hanna, L. Pochet, S. Watson, “Peer-to-Peer Communications for Tactical Environments: Observations, Requirements, and Experiences”, IEEE Communications Magazine, Vol. 48, No. 10, October 2010.

[4] NS3 Network Simulator Homepage: http://www.nsnam.org/ (retrieved on January14th, 2012).

[5] M. Conti, S. Giordano, M. May, A. Passarella, “From Opportunistic Networks to Opportunistic Computing”, *IEEE Communications Magazine*, 2010.

[6] M. Conti, M. Kumar, “Opportunities in opportunistic computing”, *IEEE Computer*, 2010.

[7] P. Hui, J. Crowcroft, and E. Yoneki, “BUBBLE Rap: Social-Based Forwarding in Delay Tolerant Networks”, in *Proceedings of 9th ACM Int’l Symp. Mobile Ad Hoc Networking and Computing (MobiHoc 08)*, ACM Press, 2008.

[8] S. Kurkowski, W. Navidi, and T. Camp, “Discovering Variables that Affect MANET Protocol Performance”, in *Proceedings of the 50th IEEE Global Telecommunications Conference (GLOBECOM’07)*, pp. 1237-1242, 2007.

[9] J. Burgess, B. Gallagher, D. Jensen, B. N. Levine, “MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks”, in *Proceedings of the 25th IEEE IEEE International Conference on Computer Communications (INFOCOM ’06)*, pp. 1-11

[10] E Daly, M. Haahr, “Social Network Analysis for Routing in Disconnected Delay-Tolerant MANETs”, in *Proceedings of the 8th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc '07).*