Exploring Value of Information-based Approaches to Support Effective Communications in Tactical Networks

Niranjan Suri[1,2], Giacomo Benincasa[1], Rita Lenzi[1], Mauro Tortonesi[3], Cesare Stefanelli[3], and Laurel Sadler[2]

[1]Florida Institute for Human and Machine Cognition, Pensacola, FL, USA
[2]US Army Research Laboratory, Adelphi, MD, USA
[3]University of Ferrara, Ferrara, Italy
{nsuri,gbenincasa,rlenzi}@ihmc.us
{mauro.tortonesi,cesare.stefanelli}@unife.it
{niranjan.suri.ctr,laurel.c.sadler.civ}@mail.mil

Abstract
Tactical networking environments present many challenges in terms of bandwidth, latency, reliability, stability, and connectivity. Sensors increasingly generate very large data sets that exceed the ability of tactical networks to transfer and disseminate them in a timely manner. Furthermore, a rapid decrease in the cost of sensors, combined with the desire to cover larger areas with persistent sensing capabilities, have resulted in a wide-scale deployment of sensors in increasing numbers, further widening the gap between the volume of information that is generated and the subset that can be successfully delivered to consumers. This paper explores the notion of determining the value of information in order to prioritize and filter information that is disseminated over these tactical networks, focusing on the dissemination of information to and from dismounted soldiers in a battlefield environment. This is a promising approach to mitigate the constraints of tactical networks and to reduce information overload on soldiers.

1. Introduction
Tactical networking environments present many challenges from the communications perspective in terms of bandwidth, latency, reliability, stability, and connectivity [1]. While sensing, computation, and storage capabilities have advanced rapidly, communication capabilities in tactical edge networks have not been able to achieve a similar growth rate. Sensors increasingly generate very large data sets that exceed the ability of tactical networks to transfer and disseminate them in a timely manner. Furthermore, a rapid decrease in the cost of sensors, combined with the desire to cover larger areas with persistent sensing capabilities, have resulted in a wide-scale deployment of sensors in increasing numbers, further widening the gap between the volume of information that is generated and the subset of that information that can be successfully delivered to consumers. Soldiers as sensors, equipped with smartphones or other portable computing devices, will place a further load on the already congested networks.

These trends have motivated researchers to increasingly focus on the challenging problem of filtering information and of prioritizing and transmitting only those subsets that would be useful to consumers. In fact, recent research in multiple disciplines has raised the question of determining the Value of Information (VoI) as an enabler for effective decision-making [2], thus enabling the filtering and prioritization of information according to the corresponding value perceived by the consumer on an individual basis [3].

Solutions that can analyze information and infer its value represent a natural complement for tactical communications middleware. In fact, the latter were designed to withstand node mobility and communication path disruptions and to exploit the scarce communication resources in the most efficient way, typically by implementing smart and reliable message prioritization mechanisms [1] and data fusion [4]. VoI-based solutions help by further reducing the bandwidth requirements and improving the communication latency, essentially trading off the delivery of non-critical information to ensure that important and high-priority information can reach consumers that need it in a timely manner.

An equally important motivation for filtering information based on value to the consumer is to reduce information overload. Delivering and presenting unnecessary information to soldiers actively performing a task at the very least results in an unnecessary increase in their cognitive workload. In the worst case, it could become a distraction and cause them to make mistakes.

This paper provides a working definition of VoI and a short survey of how other researchers and systems have applied this concept. Then, the paper explores VoI-based concepts for the purpose of timely dissemination of essential information to and from dismounted soldiers in a battlefield environment. The approach we describe is generalizable to the dissemination of information to other platforms and vehicles that are also interconnected via tactical networks, and is an extension of an earlier realization of information selection and prioritization based on relevance, which we described in [5].

2. Tactical Edge Networks

The reader is likely to be familiar with the characteristics of Tactical Edge Networks (TENs), but we summarize them in this section for completeness. Many types of nodes typically operate in the tactical environment. Some are mobile, such as manned and unmanned ground and air vehicles and portable devices carried by dismounted soldiers. Other nodes are stationary, such as Tactical Operation Centers and unattended sensors (often grouped into wireless sensor networks to perform coordinated information gathering and object tracking tasks). Most of the nodes communicate through wireless links of various types (satellite, cellular, and ad-hoc), usually in a hostile RF environment. As a result, in the tactical environment, severely constrained bandwidth, highly varying communication latencies, disconnected nodes, and network partitions are more the norm than the exception.

From the information-centric perspective, the objective of communications middleware operating in TENs is to manage discretized units of information content (henceforth referred to as Information Objects or IOs) and deliver them to consumers in the most effective way. IOs may be as simple as an assertion of some facts, location data of a friendly or enemy unit, a graphic such as a map, a picture, a document such as an intelligence report, or a full motion video clip. In a military context, information may be deemed to be of value if it increases the Situational Awareness (SA) of the consumers and/or causes them to alter their course of action for a better outcome. The overall objective is to convey useful information to support decision making while reducing bandwidth consumption, delivery latency, and cognitive workload. Any reduction in bandwidth utilization alleviates the constraints of TENs and is beneficial. Furthermore, the sorting and delivery of IOs in priority order reduces the latency of delivery of important information, which is an added benefit.

3. Defining Value of Information

Research in sensor networks, where strict constraints on computation, energy, and channel access make communications particularly expensive, has recently identified two interesting metrics for ranking Information Objects (IOs) with filtering and prioritization purposes: Quality of Information (QoI) and

Value of Information (VoI) [6] [3]. In this paper, we generalize these metrics to other IOs, such as tracks and documents, in addition to sensor reports.

More specifically, QoI represents an *"internal" and objective metric* that considers the *intrinsic characteristics* of an IO. For example, an IO containing a photograph or an image might have a QoI value defined by aspects such as level of detail (or resolution), clarity, contrast, exposure, etc. In some cases, it is possible to devise automated systems that determine the QoI of an IO by analyzing its contents. For example, IOs generated by infrared sensors will typically have higher QoIs than those generated by visual sensors at night, and vice-versa by day. In more complex cases, such as IOs containing documents/reports, the determination of the QoI may have to rely on a human operator.

On the other hand, VoI represents an *"external" and subjective metric* that classifies an IO according to the *utility it provides to its consumer* – that is, to its ability to support the consumer in more effective decision making. In fact, the VoI concept originated in economic and decision making research communities, with the purpose of investigating the advantages that additional information provided to decision makers [7] [8]. VoIs are dynamic values that change according to many factors, such as a consumer's needs and information availability. Also note that the same IO may have different VoI values for different consumers and that an IO may have a very high QoI (for example, it may be a very accurate, recent location report for an enemy unit) but it may have a very low VoI for a particular consumer, for whom it represents irrelevant information. In other words, an IO's VoI is a function of its QoI, of its suitability or applicability to the consumer given the consumer's current context, as well as of the previous history of IOs sent to the consumer.

While QoI and VoI are novel concepts, researchers have already started investigating their adoption in mostly-static / steady-state operational conditions, either through the application of multiple-criteria decision making solutions such as the Analytic Hierarchy Process [3], or of Von Neumann-Morgenstern utility functions [9]. These earlier works focus on static VoI values and essentially adopt the congestion control objective. More sophisticated solutions applied the VoI concept to optimize the scheduling of message transmissions [10] or the traveling path of unmanned data harvesters [11] in underwater wireless sensor networks, considering time-varying but system-wide, i.e., non consumer-specific, VoI measures.

However, the highly varying characteristics of TENs, as well as the ever changing contexts and mission objectives of all the personnel and devices operating in those environments, call for significantly more intelligent middleware solutions. In particular, they must be capable of considering multiple and dynamic VoI values for each IO, according to the corresponding consumers and their current contexts, and of integrating with tactical communication solutions [1] to implement dynamic IO filtering and prioritization policies. Since the contextual information about consumers and their relative interests might be incomplete or out of date, the reasoning component of the middleware should be capable of dealing with missing data and uncertainty and of predicting personnel's future context and mission objectives – and consequently of forecasting the future VoI of a given IO for all the potentially interested consumers it knows of. Finally, the middleware should allow VoI determination to dynamically adapt based on feedback from the consumers in the field.

Calculating the actual VoI of an IO for a particular consumer is challenging, as it requires the system to model each consumer in terms of their existing knowledge, their objectives, their information needs, and their decision-making strategy. A general solution that comprehensively addresses this problem does not exist to the best of our knowledge. This paper describes a specific implementation, DSPro, which realizes VoI-based filtering for a small set of critically important IO types in tactical environments, such as tracks, sensor reports, and other documents with metadata that supports such evaluation.

4. Value-based Information Dissemination Patterns

As part of our ongoing research efforts in the application of VoI concepts within tactical environments, we have identified three different deployment patterns for exploiting information value-based dissemination, which are shown in Fig. 1. The first deployment pattern applies to information from a command / operations center flowing to dismounted soldiers using a tactical network. The intuition here is that the operations center is essentially an enterprise network node, with very few constraints on computation, storage, and network resources. However, not all of the information available at the operations center can be transmitted to the tactical edge given the capacity limitations. Therefore, each of the dismounted soldiers pushes their current context (denoted as the user context) to the operations center where an information valuator component examines the available IOs (and new incoming IOs – for example from deployed sensor networks) to determine their value, filter IOs that do not satisfy a relevance threshold, prioritize the IOs that are selected, and transmit those to dismounted soldiers. In this deployment pattern, data from sensor networks is first exfiltrated back to the operations center from one or more sensor network gateways (via a number of possible network links) before being evaluated and transmitted to dismounted users.

The second deployment pattern applies to information from the tactical network, including dismounted soldiers and sensor networks, being transmitted back to an operations center, where the consumer may be an analyst, or other tactical network users that are also connected to the operations center. In this case, the IOs are generated by sensors as well as soldiers and include tracks, detections, pictures, reports, and other potentially large objects. As discussed before, the tactical network does not have the capacity to transfer all of this data back to the operations center. Therefore, in this second deployment pattern, IOs increasingly stay where they are gathered, and are pushed out of the tactical network to the operations center based on demand. For example, an analyst may express his or her interest in different types of IOs, which would represent their user context. These user contexts would be pushed out to sensor network gateway nodes, where valuator components would match locally generated IOs to the consumers. Again, IOs that satisfy a relevance threshold would be selected, prioritized based on their relevance, and transmitted to the consumers.

The third deployment pattern applies to information sharing directly at the tactical edge – for example from one soldier to another or from a sensor network to a soldier. Unlike the first two cases, these peer-to-peer exchanges occur on an ad-hoc basis based on potentially opportunistic contacts between the edge users. Hence, the user contexts are not pre-shared but exchanged upon contact, at which point a node with relevant IOs would push those to the other node. While not shown in the diagram, user nodes and sensor network gateway nodes contain information valuators as well as the node contexts for connected peers/users (we will use the word peer and user interchangeably – a user is represented by a node that is a peer on the network).

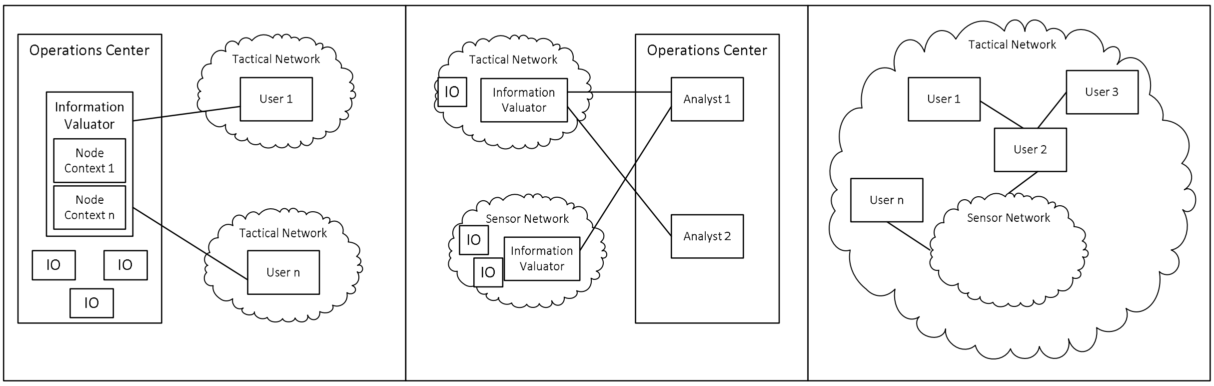

(a) From OC to Edge Network (b) From Edge Network to OC (c) At the Edge Network

Fig. 1. Deployment Patterns for Information Value Driven Dissemination.

5. DSPro Middleware

The DSPro middleware is a concrete realization of a VoI-based dissemination system that currently supports deployment patterns 1(a) and 1(c) (it is being extended to handle pattern 1 (b) as well). DSPro implements a peer-to-peer architecture, where the same node can act as a consumer and provider of information to other nodes, and builds on top of the experience we developed with information dissemination in tactical environments in the context of the DisService research project [12].

Fig. 2 shows the high-level architecture of DSPro and sketches the key components in the system and the interactions with information producers and consumers. In this middleware architecture, consumer nodes generate and push their context to other nodes that are potential sources or repositories of information. The source nodes have a matchmaker component that evaluates the value of each IO against a consumer node's context. IOs that have a value higher than a chosen or pre-determined relevance threshold are then transmitted to the appropriate consumer. IOs that have been sent to a consumer are tracked in a history store, so that they would not be considered again in the future. The matchmaker is triggered by two events – either the arrival of new IOs or changes to a consumer node's context. When a new IO arrives, it is persisted in the local data store and evaluated by the matchmaker for all known consumers. On the other hand, if a peer node context changes, all IOs in the data store that were not already transmitted are examined to see if any of them would be of value to the consumer given the updated node context. If so, those are transmitted to the consumer and tracked in the history store.

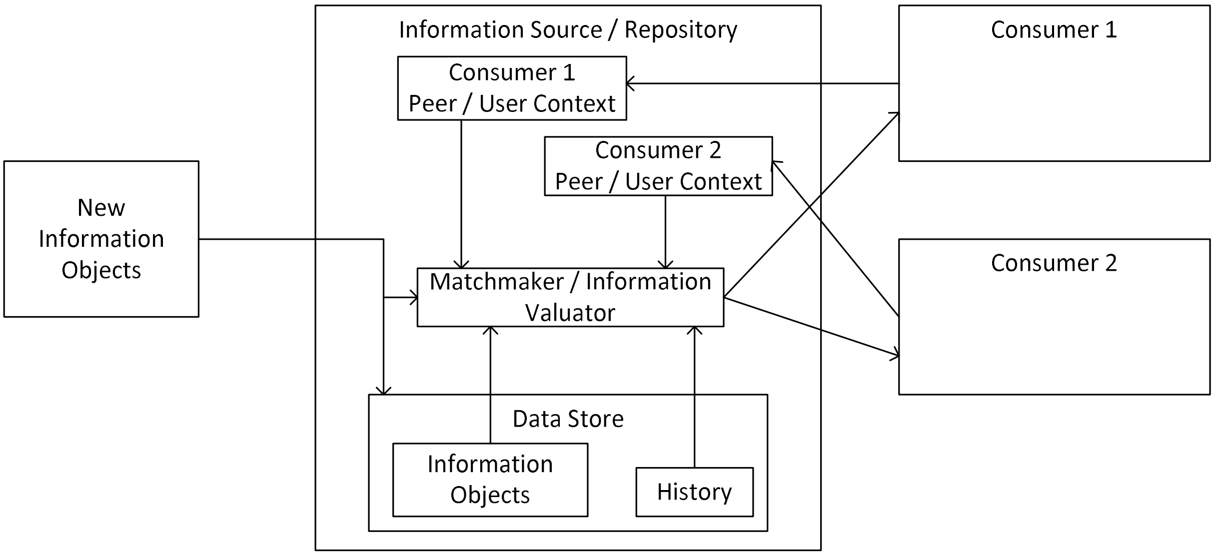

Fig. 2. High-level Component Architecture of the DSPro Middleware.

DSPro adopts a pragmatic, computationally efficient approach to determining the VoI of an IO for each user/consumer. The evaluation function depends on the data type of the IO. The user node context representation in DSPro consists of a user's current location, a projected route (or routes), along with any temporal information (usually the result of a route planning tool), type of mission, role of the consumer, current activity, as well as policies about the value to be assigned to an IO based on the metadata. Given that a vast majority of tactical information is geographically driven, extra flexibility was provided in controlling the valuation based on geographical proximity. In particular, the user context can specify the notion of a "useful distance" for different IO types, which is taken into account when evaluating the VoI based on geographical proximity. Note that the evaluation is more than a simple geographical filter. In particular, the evaluation considers future planned positions (based on planned routes) in addition to the current position. Therefore, DSPro will match IOs that may not be currently, relevant to a consumer based on the consumer's current position, but might be relevant to where the consumer is expected to be in the future, with the value being inversely proportional to the expected time at that position.

Another enhancement to DSPro's evaluation of the value of information is to consider the previous history of information that has been matched and delivered to a consumer. For example, with tracks (position updates about entities), DSPro will compute the degree of change of the track from the last update that was sent to a consumer as part of determining the value of that track. This also takes into account the distance of the track from the consumer. So, a track that moves 100 m, but is 1 km away, is assigned a higher value than a track that moves 100 m, but is 10 km away.

Proximity-based valuation also considers the range of influence of the entity represented by the IO. For example, an airborne platform, given its speed, has a much wider range of influence than a dismounted soldier walking on the ground. The range of influence policies in DSPro are specified based on the MIL-STD-2525 symbol code for the entity, which is part of the metadata of the IO. Using the MIL-STD-2525 symbol code makes it very flexible to change the valuation policies in DSPro. For example, it is simple to express a policy in DSPro that assigns a high value to airborne elements versus ground elements. Future implementations will also take into account the lines of bearing / direction of motion (moving towards a consumer versus moving away) when computing the value of that particular track for a given consumer.

The user context also specifies the weights for a ranking function, which can adjust the relative importance across the different parameters of geographical proximity, temporal proximity, mission relevance, role relevance, and activity relevance. DSPro also supports custom policies for evaluating VoI on other metadata attributes of IOs. For example, Fig. 3 shows a simple policy that computes value based on the Affiliation attribute in the metadata. This particular example is a static policy (i.e., does not change based on other information in the user context) with a weight of 4.0 (out of a maximum of 10.0). If the attribute matches the word Coalition, the value is determined to be 5.0. On the other hand, if the attribute matches the word Hostile, the value is determined to be 9.0.

```
<?xml version="1.0"?>
<RankerPolicy>
	<Type>Static</Type>
	<Attribute>Affiliation</Attribute>
	<Weight>4.0</Weight>
	<Alternative>
		<Match>Coalition</Match>
		<Value>5.0</Value>
	</Alternative>
	<Alternative>
		<Match>US</Match>
		<Value>6.0</Value>
	</Alternative>
	<Alternative>
		<Match>Hostile</Match>
		<Value>9.0</Value>
	</Alternative>
</RankerPolicy>
```

Fig. 3. Custom Ranking Policy in DSPro for Metadata

Determining the VoI of an IO for a user then consists of evaluating the metadata of the IO against that user node context. In addition to the above parameters, other attributes of the IO, such as the source (e.g., the commander of the mission), pedigree, and designated importance level are taken into consideration. As mentioned earlier, once the VoI is calculated for an IO, if the VoI falls below a configured “minimum value/worth” threshold (specified as part of the consumer node context), the IO is not transmitted to the consumer. If the VoI is higher than the relevance threshold, the IOs are sorted in priority order, based on their VoI, and transmitted to the consumer accordingly. Additional details about the matchmaking process are described in [5].

The context that is pushed to other nodes can change dynamically over time to reflect changes in the nature of information that is desired (or would be of value) to the consumer. Changes can include updates to policies (such as the one shown in Fig. 3), adjustments to the weights assigned to the different factors of geographical and temporal proximity, range of relevance based on the MIL-STD-2525 symbol code, changes to planned routes, current position, etc.) This allows a consumer to specify and control the nature of information that is desired. DSPro has been integrated with tactical applications such as ATAK (Android Tactical Assault Kit)[1] that provide the user interface to the consumer. Due to space considerations, we do not address the user interface in this paper.

Since the matchmaker may run on resource-constrained nodes, the computational cost should be considered. In DSPro, the computational complexity is O (n), where n is the number of IOs in the data store. If a peer node context changes, all n IOs in the data store have to be evaluated against the updated peer node context. On the other hand, when a new IO arrives, it has to be evaluated against each peer (so, if there are m peers, there are m evaluations). The computational cost is kept low by not considering interactions between IOs. For example, DSPro does not consider that sending a report from Sensor 1 may make it unnecessary to send a different report from Sensor 2 since both the reports were covering the same target.

IOs are also immutable, so an update to an existing IO is handled as a new IO through the system. Furthermore, DSPro linearizes the IOs so that they can be handled in some sequential order. The simplest approach to linearization is to use a metric such as the creation time or arrival time of the IO at a node in the network (for example, a node that has to process the IOs).

Because military operations are inherently group efforts, it is often the case that multiple consumers have similar node contexts, and therefore have interest in receiving similar subsets of IOs. For instance, all the members of a team may be interested in receiving situational awareness data from the area of deployment. A naïve implementation of Value-Based Information Dissemination that transmits the selected IOs objects via unicast may even under-perform a simpler IP multicast-based implementation that blindly transmits all the IOs to every consumer, if the overlap among the matched subsets is large. In DSPro, while the IO selection is performed by matching a single consumer node context against a single IO metadata, the actual transmission of the IOs is performed by taking the network topology into account. The matchmaker component aggregates IOs that need to be sent to multiple consumers in the same subnetwork and then transmits a single copy of the message via multicast.

DSPro relies on a reliable multicast capability provided by DisService [12] (which, in turn, is based on hop-by-hop UDP multicast). Because IOs are multicast, nodes may receive irrelevant IOs, which are not processed, but are locally cached for possible later use. If a previously cached IO becomes relevant in the future, the node is simplify notified that the IO is now relevant to the consumer, without having to transmit the IO again, thereby saving on bandwidth. DisService goes even further with its implementation of opportunistic listening [13], where each network packet is self-contained and self-describing, which allows intermediate nodes that listen in on this traffic to be able to cache the packets, and potentially make them available to other peers at a later point in time. This approach increases the availability of the information in the face of disconnections and network partitioning.

It is important to note that false positives (where the information is not relevant but is not filtered) can be tolerated because while the system may not save as much bandwidth as possible, it would still be an improvement over the baseline. However, false negatives (where information is relevant but is still filtered) are detrimental as it might cause the consumer to lose SA and/or make an incorrect decision.

6. Experimental Results

We performed an experimental evaluation of DSPro using a reference scenario based on Agile Bloodhound - an annual US Department of Defense (DoD) Office of Naval Research (ONR) sponsored

[1] https://atakmap.com/

technology demonstration, where DSPro has been deployed over the last four consecutive years. This scenario realizes the deployment pattern from Fig. 1(a). As shown in Fig. 4, the scenario consists of multiple military hub vehicles connected via SATCOM to a command center. Each hub vehicle supports multiple dismounted soldiers, either on foot or in vehicles of their own, using a tablet/smartphone and communicating through a MANET. The scenario emulates a typical operation that involves a variety of information flows, including friendly (blue) force tracks, enemy (red) force tracks, sensor reports, documents such as intelligence reports, logistics reports, and messaging. For the purpose of this experiment, we consider only two data types – tracks and sensor reports. The metadata and data format for tracks and sensor reports were XML messages as defined by the Marine Air-Ground Task Force (MAGTF) Command and Control (C2) Tactical Service-oriented Architecture (TSOA). However, it should be relatively easy to adapt the system to use other standardized metadata and data formats, particularly if they are XML-based.

Tracks were fed into DSPro at the operations center from JTCW (the Joint Tactical Common Operating Picture Workstation), but could come from any number of track management systems such as DCGS (Distributed Common Ground System). Tracks are also generated by the vehicles on the move and by the dismounted soldiers at the edge. Note that there are a wide variety of tracks, representing platforms ranging from ships to airborne vehicles to ground units, that are part of multiple, ongoing, sometimes unrelated missions. Therefore, not all tracks are relevant to every consumer.

Sensor reports are generated from a variety of sensors deployed throughout the area of operations. A sensor report typically contains metadata and a payload – typically an image. Each sensor report affects or covers a geographic sub-region. Again, not every sensor report is relevant to every consumer.

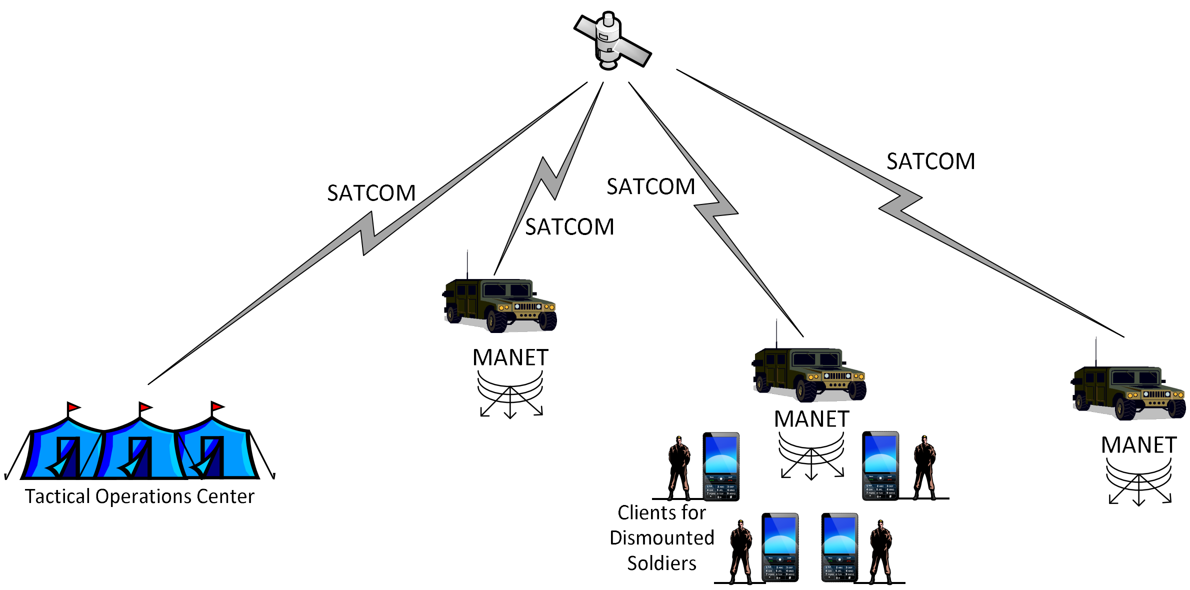

Fig. 4. DSPro Experiment Scenario.

We compared three different information dissemination strategies: a baseline strategy, in which all tracks and sensor reports are delivered via DisService's reliable multicast communication function to each of the four clients; and the naïve and DSPro strategies, which are VoI-based approaches that independently evaluate and select the IOs to deliver to each consumer. The naïve approach to implementing such a system is to independently transmit the relevant IOs to each consumer. But this approach could be highly inefficient in a scenario where there is a large degree of overlap between the

IOs that are selected for a set of consumers in physical proximity, since the overlapping IOs would be duplicated on the MANET network as they are transmitted independently to each consumer. As mentioned earlier, the DSPro implementation aggregates data and avoids multiple transmission of the same IOs on the same network links, thereby saving bandwidth.

It is important to note that the results could vary widely based on the policy selected for the VoI-based filtering and prioritization, and on the actual scenario itself (in terms of the number of tracks, their positions, the number of sensor reports, their sizes, and their coverage). In fact, DSPro provides many customizable options for determining the VoI as well as setting the filtering relevance thresholds. As a result, the number of IOs transmitted by the three strategies to consumers are significantly different.

Table 1 shows the results we obtained for the subset of soldiers supported by one of the hub vehicles. In total, there were 3445 tracks and 510 sensor reports generated during the mission execution phase. Results for both the Naïve implementation and the DSPro optimized implementation are shown for the three different selectivity thresholds (high, medium, and low) defined by setting the corresponding useful distance to 0.5 kilometer, 1 kilometer, and 2 kilometers for sensor reports. In all cases, the geographic proximity filter for tracks was set to be 1000 kilometers, a large value to ensure that all nodes have the same common operating picture regarding tracks in their area of operations. This setting still resulted in tracks for unrelated missions being filtered out, but each of the four clients that are part of the same mission receive the same set of tracks (i.e., an overlap of 100%).

With regard to sensor data, the number of reports filtered out for each client varies according to selected dissemination strategy and, for the naïve and DSPro strategies, to the configured relevance threshold. As can be seen in the results, applying a VoI-based filtering mechanism in this scenario is very effective, reducing SATCOM bandwidth utilization between 26.85% and 45.39% and MANET bandwidth utilization between 48.60% and 64.86%. Note that there is a small difference between the Naïve and DSPro implementation with the same selectivity, with delivery count being higher with DSPro. This is caused by DSPro being more efficient, and being able to send more reports to clients before their position (and consequently their node context) changes. Once the node context changes, any unsent reports that are no longer relevant are simply dropped.

Table 1. DSPro Performance Results.

		Tracks				Sensor Reports				Bandwidth (Kb/sec)		Performance Improvement	
	Selectivity	Total Number	Delivered Per Client	Filtered Out	Average Overlap	Total Number	Delivered Per Client	Filtered Out	Average Overlap	SATCOM	MANET	SATCOM	MANET
Baseline	N/A	3445	3445.0	0	100.0%	510	510.0	0	100.0%	642.1	936.2		
Naïve	High	3445	997.0	2448	100.0%	510	105.5	315	75.9%	361.9	466.8	43.64%	50.14%
DSPro		3445	997.0	2448	100.0%	510	107.8	315	76.9%	350.7	328.9	45.39%	64.86%
Naïve	Medium	3445	997.0	2448	100.0%	510	124.2	303	82.1%	369.4	515.7	42.47%	44.92%
DSPro		3445	997.0	2448	100.0%	510	125.0	303	82.6%	361.5	338.9	43.71%	63.81%
Naïve	Low	3445	997.0	2448	100.0%	510	173.5	223	80.8%	436.7	644.5	32.00%	31.16%
DSPro		3445	997.0	2448	100.0%	510	176.5	223	83.6%	469.7	481.2	26.85%	48.60%

7. Conclusions

Value of Information-based approaches to information management and dissemination are a particularly promising direction for tactical network communications. They are an effective mechanism to counter the increasing disparity between the volume of data gathered/generated and the bandwidth available in the network to move all of that data. VoI based approaches also have the potential of reducing operator overload by filtering out unnecessary information that can be distracting. Determining the VoI in a generic, open-ended system is a difficult, unsolved problem. However, this paper has described the DSPro middleware that exploits VoI in a tactical information management context. DSPro has been applied to the problem of disseminating information from an operations center to dismounted soldiers, and also between soldiers and sensor networks at the edge. Initial results are promising in terms of the bandwidth reduction. Future efforts will focus on more comprehensive representations of users' contextual information that is used to evaluate VoI, and on more flexible VoI evaluation mechanisms that can accommodate multiple deployment scenarios.

References

[1] N. Suri, E. Benvegnù, M. Tortonesi, C.Stefanelli, J. Kovach, J. Hanna, "Communications Middleware for Tactical Environments: Observations, Experiences, and Lessons Learned", IEEE Communications Magazine, Vol. 47, No. 10 (Special Feature on Military Communications), pp. 56-63, October 2009.

[2] R. Laxminarayan and M.K. Macauley (Eds.), "The Value of Information: Methodological Frontiers and New Applications in Environment and Health", Springer, 2012.

[3] C. Bisdikian, L. Kaplan, M. Srivastava, "On the quality and value of information in sensor networks", *ACM Transactions on Sensor Networks*, Vol. 9, No. 4, Article 48, pp. 48:1-48:26, July 2013.

[4] H. Mitchell, "Data Fusion: Concepts and Ideas", Springer, 2012.

[5] S. Rota, G. Benincasa, M. Interlandi, N. Suri, B. Bonnlander, J. Bradshaw, M. Tortonesi, S. Watson, K. Boner, "Supporting Information on Demand with the DisServicePro Proactive Peer-to-peer Information Dissemination System", in *Proceedings of the 2010 Military Communications Conference (MILCOM 2010), pp. 984-991, San Jose, CA, USA, 31 October – 3 November 2010.*

[6] I. Todoran, L. Lecornu, J. Le Caillec, A. Khenchaf, "Assessing Information Quality in Fusion Systems", NATO SAS-106 Symposium on Analysis Support to Decision Making in Cyber Defence & Security, 09-10 June 2014, Tallinn, Estonia, 2014.

[7] S. Galanis, "The Value of Information under Unawareness", Journal of Economic Theory, Vol. 157, pp-384-396, May 2015.

[8] J. Quiggin, "The Value of Information and the Value of Awareness", Theory and Decision, 2015, DOI: 10.1007/s11238-015-9496-x.

[9] D. Cansever, "Value of Information", in *Proceedings of 2013 Military Communications Conference (MILCOM 2013)*, pp. 1105-1108, San Diego, CA, USA, 18-20 November 2013.

[10] L. Bölöni, D. Turgut, S. Basagni, C. Petrioli, "Scheduling Data Transmissions of Underwater Sensor Nodes for Maximizing Value of Information", in *Proceedings of the 2013 IEEE Global Communications Conference (GLOBECOM 2013),* pp. 460-465, 9-13 December 2013.

[11] S. Basagni, L. Bölöni, P. Gjanci, C. Petrioli, C.A. Phillips, D. Turgut, "Maximizing the Value of Sensed Information in Underwater Wireless Sensor Networks via an Autonomous Underwater Vehicle", in *Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM 2014)*, Toronto, Canasa, 27 April – 2 May, 2014.

[12] 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 (Special Feature on Military Communications), pp. 60-69, October 2010.

[13] N. Suri and G. Benincasa, “Opportunistic Listening System and Method”, US Patent Number 8,493,902 issued July 23, 2013.