Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics
San Antonio, TX, USA - October 2009

Transporting Sensor Data over Radio Networks in a System of Systems

Network and Bandwidth Planning

Robert E. Donnelly, Ph.D.
US Army, RDECOM, CERDEC, PM C4ISR On-The-Move
CSC
Fort Monmouth, NJ, USA

***Abstract*—Individual optimization of sensor data transport over a fixed network or designing a network around fixed sensor data distribution requirements is less efficient than joint optimization. Simplified rules and methods of estimating bandwidth and planning networks are needed to tackle the more complex problem of joint optimization. A process of estimating and combining steady-state and burst data flows and using this information in optimizing network design and sensor data transport is presented. This includes using rules for data dissemination to bridge distinct networks, sensors, and users in a tiered network of networks environment. Overall, a system of systems approach balancing user needs and technical solutions produces optimal C4ISR technology deployment.**

***Keywords*—communications, system of systems, sensor networks**

I. INTRODUCTION

Abraham Lincoln once said, “The right information at the right time is nine-tenths of any battle.” Geographically and organizationally distributed sensors can provide not only an abundance of data but also the challenge to get it to the right people. This paper gives approaches for using radio networks to get dispersed sensor data into the hands of those who need it.

The overall problem of distributing sensor data in a networked system of systems has many facets, including:

- Estimating both steady-state and burst bandwidth
- Network planning, both homogenous and tiered
- Bandwidth shaping based on these estimates and plans
- Handling and accounting for radio network loss
- Interfacing to sensors locally before distribution
- Managing, advertising, and geo-locating sensor data
- Distribution via push and/or pull mechanisms
- Shaping/sizing still images, video clips and streams
- Packet-level reliable transport over lossy radios

Over the past five years, Product Manager Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (C4ISR) On-The-Move (PM C4ISR OTM) has worked on these problems in the context of a C4ISR System of Systems (SoS) environment, combining live systems in a realistic field environment [1] and simulation in a high-performance computing environment [2]. An integration approach that bridges the live and simulation domains [3] brings live and simulated sensor data to consumers. Of the facets above, several such as local sensor interfaces and managing, geo-locating, and distributing sensor data have been previously explained [3] and others such as transport may be described in future work. The key for such interface and transport approaches to be successful is network planning that accounts for the high bandwidth needs of sensor data and bandwidth estimates and shaping that maximize use of the network plan, as described here.

II. PROBLEM SPACE

Just within the planning phase, there are many aspects to the problem of integrating sensor data and networks. Within a C4ISR SoS, sensor data must coexist with other network users: command and control (C2) such as email-like person-to-person messages, situational awareness (SA) such as positions of various systems and people, other collaborative services such as white boarding and chat, and network management and quality of service (QoS) overhead. Other domains will likely have similar messaging, e.g. sensor position, status, and management. Besides overall average traffic loads, two useful characterizations are 1) steady state data, e.g. a GPS-based position that is sent on a regular basis, and 2) burst data, e.g. imagery from a sensor. In any non-trivial context, planning should account for the coexistence of bursty sensor data and other data sharing the network.

A. Sensors

The type and connectivity of sensors impact the type of data. Fig. 1 shows three typical sensors in a C4ISR SoS.

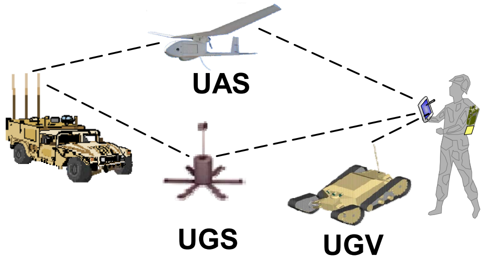

Figure 1. Sensor Types and Initial Injection Points

An Unattended Ground Sensor (UGS) often has acoustic and seismic sensors that trigger optical and infrared sensors that

978-1-4244-2794-9/09/$25.00 ©2009 IEEE

capture still imagery or possibly a short series of frames. Individual imagery is small, but either artificial intelligence or a soldier-in-the-loop is needed to limit the number of images that are injected from the initial contact, dismount (right), vehicle (left), or fixed location (not shown). An Unmanned Ground Vehicle (UGV) is usually under direct control of a dismount or, at times, a manned vehicle. Generated imagery is usually stills or at times a short video clip. An Unmanned Aerial System (UAS) streams data to a dismount, a vehicle, and/or a fixed location. Its altitude means it may be able to stream directly to several users at once, and its viewpoint means that sending its stream, clips, or stills to other locations may all be of value. Note that a lower-echelon UAS will fly at a lower altitude and have less range and time in air, impacting both the available data and the value of further distributing this data versus a more capable UAS assigned to higher echelons.

B. Networks

Sensor data usually uses dedicated communication assets to the initial viewer of its data, be it the controller of the sensor or possibly a remote video terminal (RVT) in the case of a UAS. Distributing RVTs to soldiers to provide them direct feeds to any UAS to which they can connect requires balancing the benefits of the feeds versus the drawbacks: cost, weight, power, setup time, security concerns while focusing on the RVT, etc. The primary viewpoint of this study is further distribution of sensor data after its initial contact to a user, accounting for the differences in data just noted. Given that context, these tactical networks are now considered.

Most tactical networks share several characteristics: varied and shared use, loss and outages, an ad hoc and/or dynamic nature, and limited bandwidth compared to potential demand. A C4ISR SoS will likely have a network of networks involving tiers as shown in Fig. 2.

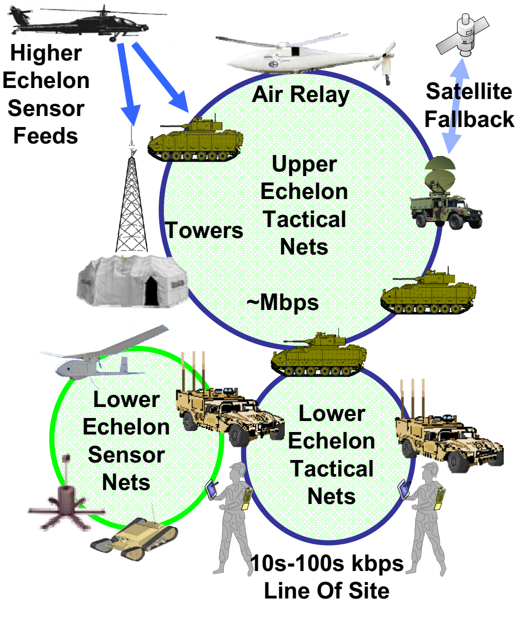

Figure 2. Tiered Tactical Networks

Lower echelon tactical networks (company and below) are mostly line-of-sight (LoS) based and have limited bandwidth per node. Higher echelon tactical networks (battalion and above) include higher bandwidth LoS networks plus potential air and/or tower relays and fallbacks to satellite links. Just as the distinct sensor types generate distinct types and scales of data, the distinct networks, especially between lower versus upper echelons as noted here, have distinct characteristics and capacities. Beyond bandwidth, range, and relay capabilities, there may be distinctions on per-node bandwidth allocation between the two types of networks. A carrier sense multiple access (CSMA), lower-echelon radio is able to allocate bandwidth to high-need users on demand. A high-bandwidth, LoS, upper-echelon network may need to pre-allocate bandwidth to advantaged nodes such as an air relay. Satellite uplink bandwidths also are often pre-allocated, e.g. giving a fixed location more bandwidth than a mobile node. Distributing sensor data over the networks efficiently requires considering these distinctions.

C. Sensors, Networks, Users, and Data Flows

Understanding the types and characteristics of the varied sensors and networks, the problem is now to optimize the data flows available to the users. Lower echelon sensor data is shared on the lower echelon networks and sent to higher when warranted. Upper echelon sensor data is shared to high-value nodes (using high-value communication assets on the upper echelon networks) but may also be needed, at times in formats requiring less bandwidth, by soldiers at the lower echelons. The desire is to meet the distinct needs of the intelligence analyst and the forward soldier by jointly optimizing the sensor data transport and the networks that carry it.

III. APPROACH

Separate design of application use of bandwidth and the networks that provide bandwidth leads to non-optimal overall design. Especially given the challenges of mobile, tactical wireless networking compared to wired or cellular networks, cross-layer optimization, covering application, middleware, and networks, is needed [4]. Security concerns must also be addressed, and architectural approaches such as protected core networking [5] balance security and availability. Tools for mobile ad hoc network (MANET) design consider the number and location of nodes, mobility, and transmit power [6]. In a system of systems environment with the ability to set the rules for how an integration framework puts sensor and other data on the network [3], it is possible to jointly design sensor feeds and networks. Since the tactical networks must carry multiple types of data as noted earlier, the joint design must consider multiple applications and the networks they use.

This complex problem is made tractable by separating the calculation of steady-state and burst data rates. Steady-state data includes most smaller-sized messages, e.g. C2, SA, and collaborative services as defined above. Sensor imagery, be it still images, video clips, or video streams, is counted as burst data. Voice call groups, network management, and QoS overhead may be allocated as fixed percentages of network capacity, with the remaining available for user data in the steady-state and burst categories. Optimizing voice call group overhead is possible but mostly by the allocation of network nodes to distinct frequencies, as discussed below.

A. Steady-State Message Calculations

The frequency of message rates per node, e.g. frequency of position reports and C2 messages, leads to rates per node. C2 messages sent by a user to addressees may be hard to estimate, as are other user-driven messages. Data capture and analysis of realistic tests provides the best estimate of this traffic [7]. As automatically generated messages often form the majority of the steady state messages, estimates for the less frequently sent user-driven messages are less critical. For a homogenous network segment, the steady-state rate estimate may be reduced to four components:

- Internal human-generated
- External human-generated
- Internal auto-generated
- External auto-generated

Considering the first two, often lower bandwidth, categories as fixed leaves optimization to the latter two. Internal auto-generated bandwidth is determined by setting a frequency, e.g. 30 nodes that each generate a position report every 2 minutes yields 15 messages a minute. External auto-generated bandwidth is determined by a frequency and information management rules. For example, of a potential total of 3000 positions, perhaps 200 would be relevant to the 30 nodes of this network segment, based on geography and/or organizational relationships. Injecting these into the network every 5 minutes leads to 40 messages a minute. If there is a single injection point, e.g. a vehicle bridging the two tactical networks shown in Fig. 2, gains from aggregating external position reports may be possible, e.g. 4 (larger) messages with 10 positions each every minute. The steady-state message calculation is now tractable with just a few variables to consider: frequency of internal and external auto-generated messages and rules for injecting auto-generated messages external to a given network segment.

B. Burst Message Calculations

An average over time of burst-based traffic may be of little value. The goal is to handle the expected bursts patterns so the user has the right information at the right time. While the military environment has challenges such as mobility and a lack of infrastructure such as cellular towers, there are several ways to simplify the problem of optimizing burst, high-bandwidth traffic in this setting versus the commercial Internet.

Consumer video viewing over broadband has or will bypass peer-to-peer as the dominant traffic on the Internet [8-9]. This has led to proposals to use peer-to-peer methods for video streaming [10] and video on demand [11] that reduce demands on the original servers and balance network loading. User Datagram Protocol (UDP) with multicast addressing is not frequently used for consumer video. Internet Service Providers (ISPs) may not support it, and it requires prior agreement on multicast groups and appropriate forwarding in routers. A military environment allows determining the sources and viewers of video based on operational needs. With proper planning of video feeds, networks, and multicast routing, the demand on the original servers is reduced. In a MANET environment, UDP multicast to multiple users uses limited additional bandwidth over UDP unicast to a single user, especially on a given network segment. These distinctions between consumer and military use of video are important to understand, and our approach may apply to other fields such as homeland security or industrial monitoring.

Information management rules that determine the sensor data available to each node, the benefits of multicasting, and operational and command-enforced limits on network use combine to make burst data management tractable. Restrictions on the sizes of downloadable files limit the download times so simultaneous downloads are limited. Multicasting limits the simultaneous video streams. The net result is that the following two components of burst sensor data may be estimated for a given network segment:

- Streaming data
- Still image and video clip downloads

The number of streams is limited via planning of who sees what video and using multicast to share a single stream to multiple viewers. The size of a given stream is settable by adjusting frame size, frame rates, and quality. For the most part, we have implemented a pull mechanism for stills and clips. The number of simultaneous downloads on a given network segment may be assumed to be relatively small. By appropriately sizing the downloads, the duration of a single pull is limited. While Transmission Control Protocol (TCP) attempts to use all of the available bandwidth, somewhat by trial and error, we have used UDP with application-layer reliability and bandwidth shaping to pre-determine the bandwidth used by any given file (image or clip) transfer.

C. Variables, Optimization, and Estimates

The variables in generating a combined bandwidth estimate have been reduced to a relatively small set: number of nodes per network segment and frequency of reports, number of external nodes reported upon and frequency, size and number of streams, still and clip file sizes, and bandwidth (BW) used for each file transfer.

The combination of steady-state and burst traffic estimates may be best illustrated by a simplified example. Consider a network segment that has 100 kbps aggregate bandwidth. Estimates for human-generated traffic and rules for auto-generated may be used to find the total steady-state bandwidth.

TABLE I. STEADY-STATE BANDWIDTH

	Nodes	**Freq. per min**	**Total per min**	**Size**	**BW kbps**
Internal, human			25	512	1.67
External, human			25	512	1.67
Internal, auto	200	0.5	100	512	6.67
External, auto	300	0.5	150	512	10
Total					**20**

With these assumptions, the 20 kbps steady-state data leaves 80 kbps for burst data. A steady-state percentage of total bandwidth of 25-50% is typical for our scenarios. Assuming up to 2 simultaneous downloads, a per-download rate of 40 kbps is appropriate. Consider a 50 kiloByte image in 36 packets with overhead. At 40 kbps, this takes 10 seconds to transmit initially. This assumes no packet loss and does not account for retries and acknowledgements. To complete an entire image, a minimum number of retries per packet, R, for total packets, P, and per-packet loss, L (as a fraction), is:

$$R = \log(P) / \log(1/L)$$

To have success at up to 50% packet loss, each packet in a 36-packet image needs to be retried up to 5-6 times or transmitted a total of 6-7 times. A reasonable retry mechanism is then initial transmission over 10 seconds and up to 9 retries every 10 seconds. This provides a likelihood of success in the order of a minute even with 50% packet loss, and often 20-40 second latency with less loss. On a limited bandwidth, lower-echelon network, this is a reasonable time to wait.

Taking this last calculation in reverse, we can use the available, per-download burst bandwidth (40 kbps) and a desired maximum transmit time (about a minute) to obtain a reasonable image size (50 kB). The retry mechanism timing and timeout will enable success even if 3-5 multiple downloads are attempted simultaneously. Field testing is critical and allows tuning of all of the parameters mentioned (image size, time between retries, total time-out period) due to changes in loss, steady-state bandwidth, and other overhead.[1] These rather simple calculations do provide a mechanism to obtain a reasonable estimate of target image sizes that has been proven workable in field experiments.

The application of available burst rate to sensor video streams is simpler. The above steady-state calculation leaves 80 kbps for a single video stream. This is a low bit rate for video, but is workable by reducing frame sizes and rates. Attempting video streams over a lower-echelon network will only be workable when loss is low. Also, operational control may be needed to stop other imagery downloads at the same time. By taking appropriate measures, video streaming over such low-bandwidth, lossy wireless networks is possible. Note that similar calculations are possible for Mbps-capable upper-echelon networks resulting in the ability to handle much larger still images and video clips and larger streams, noting that the steady-state bandwidth needs will also be much higher.

IV. APPLICATION

The above bandwidth calculations started with fixed assumptions about a network. To optimize network design, these calculations are tools to evaluate various scenarios, which are then varied to achieve optimal results. Several examples of this are given in this section.

A. *Optimized Network Design of a Lower-Echelon Tier*

In a MANET, network management overhead increases with the number of nodes. In order to handle a larger number of nodes, tiering of otherwise homogenous networks may be needed. Optimal design requires considering nodes per frequency, the total bandwidth per frequency, bridging capabilities, and segmenting both data flows and voice call groups. An example network design is given in Fig. 3.

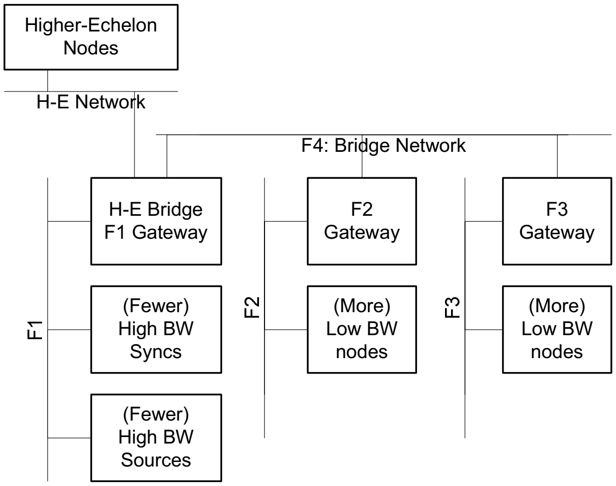

Figure 3. Tiered Tactical Networks

Collecting operationally associated nodes on the networks with frequencies F2 and F3 allowed keeping certain data to these networks. Certain voice call groups and data multicast groups could remain strictly on these frequencies. The nodes on F2 and F3 had sensors with limited bandwidth needs (e.g. stills only), if any. The users on the nodes also had limited need to pull large amounts of imagery, except for perhaps the gateway to the F4 bridge network, assigned to a leader position. This allowed using a relatively large number of nodes on each of these networks. Key leaders, who view more sensor data (syncs), and nodes connected to high-bandwidth sensors (sources), such as a ground control station for a UAS, were collected on a separate network with frequency F1. One vehicle on this network also served as a bridge to a more capable, higher-echelon (H-E) network.

Within this network design, the above approach of determining steady-state bandwidth, bandwidth available for bursts, and thus burst bandwidth shaping could be applied. Alternative network lay downs can be compared in a series of “what if” scenarios to find the optimum combined sensor and network plan.

B. *Combining Upper- and Lower-Echelon Tiers*

In designing a more complex, multi-tier, non-homogeneous network, the same methods are applied, with a given calculation as above for each network segment and consideration of the inter-segment traffic, both steady-state and burst. While it may be possible to send the position reports of thousands of entities (both live and simulated) to all nodes directly on a higher-echelon network, bridge nodes to the lower-echelon networks would only forward hundreds. Operational analysis provides estimates for human-generated traffic within each and between the various network segments. So while there are a greater number of domains to consider

---

[1] The results of field testing and parameter tuning for specific waveforms are not releasable to the general public. [1]

plus determining appropriate rules for bridging domains, the bandwidth estimation methods are basically the same.

With an estimate for available burst bandwidth, file-based images may be appropriately sized. The file advertisement mechanism may use a filter so a lower-echelon node does not attempt to pull larger images. Note that the bandwidth limiter would prevent flooding the network in any case, but there would just be a low likelihood of completion for large images.

Consider balancing the needs of an intelligence analyst who may make use of information that is lost when video streams are set to low bitrates, a commander who uses this information to make deployment decisions, and a forward soldier who may need a timely but less detailed feed. Each of these individuals will likely have distinct network support and access. Transcoding video at network boundaries is a way to meet the needs of each of these. An analyst in a Brigade (BDE) Tactical Operations Center (TOC) may have access to a standard television quality 5 Mbps video stream from a UAS. This may be sent over an upper echelon network at ~0.5 Mbps by using an efficient codec and reduction in frame sizes and rates. Similarly, a UAS feed available at a mobile node that bridges upper and lower echelon networks could transcode this feed appropriately for either network. Fig. 4 illustrates this concept, where the maximum bitrates are based on the per-network analysis calculations noted above. So the three sets of users see streams appropriate to both their needs and networks.

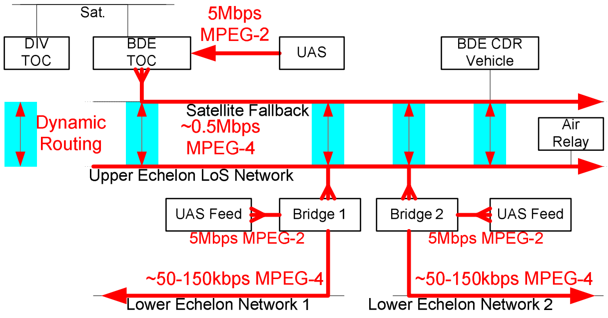

Figure 4. Tiered Networks with Video Transcoding

V. DISCUSSION

The prior section applies an approach for bandwidth estimation to two examples where networks and sensor data flows are jointly optimized. Competing network uses such as C2, SA, collaboration, and network overhead are also considered and included in bandwidth estimates. The result is optimization of all facets of networks and network users in a C4ISR SoS. Two mechanisms (bandwidth-shaped, reliable transport of files and transcoding of streams) were briefly discussed, but the focus of this paper is the planning phase. These mechanisms were used to enable calculations of optimal image size and allow matching streams to user needs and network capacities.

These examples were somewhat simplified for illustration. Additional considerations are used in the planning phase. One is the impact of message size on available throughput rates. A message that does not fill a complete packet frame may mean that the remaining capacity is lost. Reaching theoretical bandwidth limits would require that all packets are the size of the maximum transmission unit (MTU). The Ethernet MTU is typically 1500 bytes, but tactical radios may user smaller values for over-the-air transmission. Further, a radio network may require request to send (RTS), clear to send (CTS), and acknowledge packets that may add distinct overhead rates to small versus large packets. In general, network and sensor data optimization should account for the characteristics of the waveform of the network segment under study, and this was done for several waveforms.

The distinct types of sensors were mentioned early but limited attention was then paid to these differences. Just as UAS video streams were diagramed in Fig. 4, the mostly file-based data flows of UGS and UGV sensors are diagramed from source to user over alternative network designs. These may be used to estimate input to bandwidth calculations such as the likely number of simultaneous image pulls over a given network segment and output such as optimal file size.

Network conditions frequently change in a tactical setting. Redundancy in networks may result in varying available bandwidth over time. While this paper emphasizes pre-planning versus response to changing conditions, sensor integration may account for known likely scenarios. For example, fallback from a higher bandwidth LoS radio to a satellite link could be recognized via router polling with a response of adapting file download rates and stream parameters.

One untried extension would be to vary streams based on feedback. The Real-time Transport Protocol (RTP) has been used to send MPEG-4 H.264 video streams on our networks, as is done for video over third generation (3G) phones. The associated Real-time Transport Control Protocol (RTCP) may be used for client feedback on loss and delay [12]. The server could adapt to these reports, reducing bitrates if loss and delay increase. This may be preferable to fixed reservations of video stream bandwidth via QoS mechanisms that could crowd out other uses while still degrading to be unwatchable.

Unexpected consequences may result from integrating networks and sensors. An example is that when a vehicle was parked in an UGS field, repeated reports flooded the network. The next planning phase included looking at appropriate controls, human in the loop or improved reporting mechanisms, to prevent this.

Operational control, be it simple in-the-loop forwarding or authorizing and limiting sensor imagery and stream viewing, is an example of human response to a given situation. Emergent behavior in a complex system of systems may require human adaptation that complements the planning stage.

VI. SUMMARY

Optimizing the transport of sensor data over tactical radio networks has been shown to be complex problem. A system of systems environment includes distinct user needs, multiple mechanisms (files and streams), and requirements to share networks with other uses. The presented approach segments the data types, reduces the problem to a few parameters, and then looks at the resulting capabilities given these assumptions. This method may be iterated for comparisons of distinct assumptions and designs. Both homogeneous and non-homogenous networks of networks are supported by applying the method to each network segment and determining appropriate rules for bridging these segments. The net result is an approach that moves sensor data from a varied set of sensors over a tiered network to a set of users with distinct interests.

A key assumption and implementation is shaping bandwidth based on the network and sensor plan versus adapting to ongoing guesses of network status. As TCP's connection orientation and congestion-control mechanisms were developed for wired networks and tend to misinterpret loss [4], bandwidth shaping and UDP are used in the transport mechanism. This also allows efficiency via multicast often not seen in consumer-oriented video distribution.

Network status information may be used to adapt sensor flows, but via major status changes, e.g. failover from a network of one capacity and set of characteristics to another, rather than in short-lived changes. An analogy may be basing investment strategies on long-term trends and the available investment vehicles versus the prior hour of trading. In both cases, the short-term past may not predict the short-term future, and a solid plan based on trends is preferable.

Finding reasonable deployments of sensors, sensor controllers and terminals, and network assets involves additional technical (e.g. weight and power) and non-technical (e.g. cost, logistics, and operational use and need) components. Our methods have been used to set up various scenarios for testing in a combined high-performance simulation and realistic field environment over several years. Parameters may be adjusted during an exercise and deployments varied each year, but the bandwidth estimation and network planning methods presented here for handling sensor data flow provide a reasonable starting point. The end result is C4ISR SoS investigations that are relevant and enable future deployment of such systems. Joint bandwidth estimation and network planning are the keys that put the right information in the right hands in the time frame needed to make a difference.

ACKNOWLEDGMENT

The work described in this paper is the result of many organizations working together to inform on a SoS problem that none of them individually "owns". PM C4ISR On-The-Move provides a Live/Virtual/Constructive venue that enables investigation of this SoS problem space, and its staff reviewed the content of this paper. The government programs and technology providers that participate in PM C4ISR On-The-Move's events enable the analysis presented here. The support and funding of the Army/Land Division of CSC and the CSC Defense Group Fellows Program enabled this paper to be written and presented. The author would like to acknowledge all of these contributions.

REFERENCES

[1] Requests for reports of PM C4ISR On-The-Move's events should be made through: US Army, CERDEC HQ, Ft. Monmouth, NJ, PdMC4ISROTM@conus.army.mil.

[2] J. Santiago, "Stimulation of live C4ISR experimentation environments by using HPC simulation," Proc. of the 2009 Department of Defense High Performance Computing Modernization Program Users Group Conference, San Diego, CA.

[3] R. Donnelly, "Bridging live and simulated domains with a common integration approach," Proc. of the 2009 Spring Simulation Multiconference, San Diego, California, pp. 355-364.

[4] North Atlantic Treaty Organisation, "Information management over disadvantaged grids," RTO-TR-IST-030, 2007, http://www.rto.nato.int/pubs/rdp.asp?RDP=RTO-TR-IST-030 (accessed November 3, 2008).

[5] G. Hallingstad and S. Oudkerk, "Protected core networking: an architectural approach to secure and flexible communications," IEEE Communications, vol. 46, no. 11, pp. 35-41.

[6] L. Kant, et al., "Network science based approaches to design and analyze MANETs for military applications," IEEE Communications, vol. 46, no. 11, pp. 55-61.

[7] M. Amabile, J. Sypniewski, and M. Griffis, "Advances in mechanisms supporting data collection on future force networks: Product Manager C4ISR On-The-Move," Proc. of the 26th Army Science Conference, 2008, Orlando, FL, http://www.asc2008.com/manuscripts/B/BP-01.pdf (accessed 23 March 2009).

[8] S. Hansell, New York Times, October 31, 2008, http://bits.blogs.nytimes.com/2008/10/31/video-is-dominating-internet-traffic-pushing-prices-up/ (accessed March 19, 2009).

[9] G. Kim, IP Business News, November 15, 2008, http://www.ipbusinessmag.com/articles.php?issue_id=73&article_id=469 (accessed March 19, 2009).

[10] E. Mykoniati, et al., "Scalable peer-to-peer streaming for live video content," IEEE Communications, vol. 46, no. 12, pp. 40-46.

[11] A. Nafaa, S. Murphy, and L. Murphy, "Analysis of a large-scale VOD architecture for broadband operators: a p2p-based solution," IEEE Communications, vol. 46, no. 12, pp. 47-55.

[12] H. Schulzrinne, et al, Network Working Group, Request for Comments 3550.