PAPER ID # 900827 (Approved for Public Release, Distribution Unlimited)

TIGR IN IRAQ AND AFGHANISTAN: NETWORK-ADAPTIVE DISTRIBUTION OF MEDIA RICH TACTICAL DATA

Benjamin J. Ewy*, Michael T. Swink*, Steven G. Pennington*, Joseph B. Evans†,
Jin M. Kim*, Chun (Steve) Ling*, Samuel L. Earp×, Mari Maeda‡
* Ascend Intelligence, LLC, Arlington, VA
† University of Kansas, Lawrence, KS
× Multisensor Science, LLC, Ellicott City, MD
‡ Defense Advanced Research Projects Agency, Arlington, VA

ABSTRACT

The Tactical Ground Reporting (TIGR) system provides a unique multimedia patrol preparation and debriefing tool tailored to the needs of the patrol leader, and is used by all U.S. Army brigades in Iraq and Afghanistan. The distribution of data from geographically dispersed patrols requires network awareness and policy flexibility. The TIGR system utilizes an application overlay on to the tactical network that is tailored to the distribution of rich media. This architecture results in the theater-wide replication of compressed reports and meta-data, while the storage of full-quality media is distributed to the node closest to where the media was created. Thumbnail and other compressed media representations are replicated according to policy to meet soldier needs. Full-quality media can be accessed on demand from any location in the network. Network utilization is controlled by a policy-based content subscription system, a priority queuing mechanism, and the discrete scheduling of network traffic. In this paper, we focus on the requirements for media rich tactical applications, detail a system that meets those requirements, and discuss experiences using TIGR in theater.

INTRODUCTION

As digital warfare tools progress down to the company and squad levels, there is a need to flexibly tune the distribution of data from geographically dispersed elements using three important criteria: network topology awareness, knowledge of network link capabilities, and policy. TIGR [1] addresses a series of challenging requirements - the ability to efficiently search large amounts of tactical data, support for rich multimedia attachments, data transformation capabilities to provide bandwidth-efficient representations of these media, and the ability to distribute the storage of the full-quality media. To coexist with existing networked applications we must efficiently utilize network links with highly variable properties including delay, loss, link bandwidth, available bandwidth, and reachability, as well as allow for policy based selection of content through subscriptions and the discrete scheduling of media distribution and network utilization.

We have developed and implemented a system that utilizes an application overlay on to the tactical network that is tailored to the distribution and searching of rich media. This system provides a theater-wide distribution of patrol reports, significant actions, and meta-data (location, time, description, author, etc.) for media that is used to build distributed search indices. The system also provides for the storage of full-quality media in a probabilistically uniform manner throughout the network by storing it on the node closest to where the media was created. The full-quality media can be accessed on demand from any location in the network and is then cached locally.

In this paper, we further detail the requirements for media rich tactical applications intended for the "tip of the spear" [2], present an architecture and implementation that meets those requirements, detail models of the system's operation, and discuss experiences based on the system's use in Iraq and Afghanistan.

REQUIREMENTS

We focus this effort by exploring the impact of using TIGR on the tactical network. The two primary use cases for TIGR are After Action Reporting, and Pre-mission Preparation.

- After Action Reporting – Provides a mechanism for company and squad level units to create patrol reports, attach photographs and GPS tracks, and provide responses to priority information requests generated by the reporting chain. This information is then made widely available, both to peer units and higher up the chain of command.
- Pre-mission Preparation – Provides efficient searches – temporal, geographic, and content-based – that allow units to quickly develop a picture of the battlefield

978-1-4244-5239-2/09/$26.00 ©2009 IEEE

prior to patrols. A typical example would be to allow a unit to search for significant actions occurring near a convoy route in the last 2 weeks, and then review pictures and locations associated with these actions. The access to this level of detailed information provides motivation for subsequent rich reporting by the patrols when they return and debrief.

In order to satisfy these use cases TIGR must provide:

- Support for rich media including pictures, PowerPoint, GPS tracks, audio, and video. The ability to create and utilize thumbnail representations of these media types is crucial for efficient utilization of bandwidth, while preserving and providing access to the original full quality media for detailed analysis on demand.
- Distributed storage of the full quality media enables scalability so that giant repositories are not required. The unit that uploads the media is most likely to access it in the future, so storage close to the original server is ideal.
- Rapid and accurate search results. This requires local search indices, and an efficient search mechanism for the tactical data including the meta-data for media.
- The ability to function properly using network links with highly variable properties including delay, loss, link bandwidth, available bandwidth, and reachability.
- Support for policy based scheduling of media distribution to provide network managers the tools to maximize the distribution of information while minimizing the impact of this distribution on other applications and the tactical network in general.
- The ability to function in a disconnected state for long periods, either for units that become cut off from the network due to failures, or pre-planned disconnections where historical data and relevant imagery are loaded prior to going on isolated patrol.

Several observations were also important in guiding the design, namely (1) CPU and disk on servers are inexpensive, (2) LAN bandwidth is inexpensive, (3) WAN bandwidth is expensive and the dominant system constraint, (4) client CPU is expensive, (5) knowing data exists is more important than having everything (i.e., knowing an IED occurred at a location should not be delayed waiting on the high resolution picture of the crater), and (6) new data is more important than old data.

Three key philosophies drove the design process, (1) search must be fast and not dependent on WAN connectivity, (2) WAN activity must be controlled in a very granular way to minimize the impact of this application on the tactical network, and (3) flexible deployment options need to be accommodated due to the wide variety of servers, deployment scenarios, and network capabilities in theater.

ARCHITECTURE

To meet these requirements we have developed an architecture based on servers distributed throughout the tactical network. Each server provides a web based interface and a custom data replication capability for sharing data between the servers. These distributed servers contain imagery, search indices, and a data repository, which maximizes the local experience, while providing a simple to use web interface, even when the wide area network is unavailable. The data replication capability uses an application layer overlay that is tailored to the distribution of rich media. This overlay utilizes the available tactical networks for transport, with a subscription based mechanism for determining which data should be forwarded to a particular server.

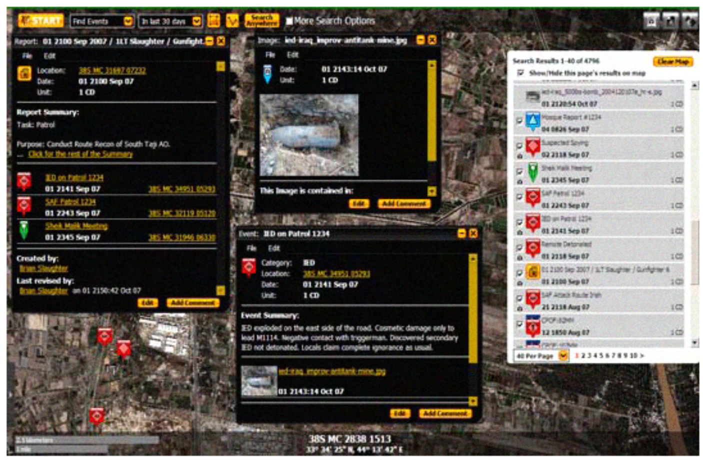

Figure 1 – TIGR User Interface Example

Each server includes the following elements:

- User Interface – The TIGR client is a web application that runs on any computer with a modern web browser. A web based user interface to the tactical data is provided which allows the soldier to easily add new reports and media into the system. The soldier can perform searches for significant activities that occur globally, in regions such as areas of operations, along routes, and constrained by keywords, categories, time, and units. The user interface is provided using Dynamic HTML (DHTML) with JavaScript [3] performing the rendering, and interfaces with the server using AJAX [4] calls, and does not require any special software to be installed on a soldier's computer. The TIGR client communicates with a web server on a TIGR server, typically a server on the same local area network as the client. The three primary interactions that a TIGR client has with a TIGR server are (1) downloading map imagery, (2) downloading search results, individual TIGR content, and media files, and (3) uploading new media files and reports into TIGR. An example user search session is shown in Figure 1, which illustrates the map-based display, search results window, and individual windows for text-based and multimedia content.

- Local Repository – The repository is a combination of an SQL [5] database and the native file system optimized to hold meta-data and media respectively. The repository has four subcomponents:
    - Meta Information Storage – The meta information is stored in the SQL database.
    - Search Index – The search index, mapping words to content including geo-temporal attributes, are also stored in the SQL database. Replicating the search index on each node allows for comprehensive searches without requiring WAN network traffic, and allows accurate searches when the server is disconnected from the WAN either intentionally (a remote patrol) or due to network failure.
    - Imagery Server – Imagery tiles are stored on the file system of the server and mapped to accessible URLs by the web server using the OpenGIS Web Map Service (WMS)[6] specification.
    - Media Storage – Media is stored on the file system of the server, and can be backed up and restored using normal backup software.
- TIGR Multimedia – Enabling wide spread multimedia use on tactical networks requires addressing the needs and concerns of those responsible for maintaining the networks. It is unacceptable to use commercial database replication software to continuously copy an entire theater's worth of data to hundreds of servers. The system needs to be able to tune the data it replicates based on each unit's needs and available system resources.

  As new content enters the server from the user interface, the system decomposes the content and transforms it so that it can be optimized for both local search and remote transmission.

  Optimizing for search first involves extracting meta information from media such as Exif [7] data from digital photographs (location, time, etc.), then associating this meta information hierarchically amongst containing content so that a location specified in a picture that is attached to a report will cause the report to be found when searching that location. Finally, all of the text information and meta information for these content objects is captured in a vector model based [8] search index which has been augmented with geographic, temporal, and unit hierarchy based structures, allowing support for a wide range of tactical searches.

  Optimizing for transmission involves converting media into low bandwidth formats that can be used in addition to the full sized media. For pictures, this involves making a thumbnail, and an 800x600 pixel compressed picture. For PowerPoint files, the thumbnail is a small picture of the first slide, and a series of pictures, one for each slide are provided for a compressed view. Audio and video can be handled similarly, with short snippets provided as a thumbnail, and a compressed form for casual use.
- TIGR Replication – An administrative interface is provided for managing the replication filters by selecting from templates of subscriptions. Subscriptions are used to express the data which a server is interested in receiving (such as all content from a particular brigade or area of operations). Subscriptions can be based on a number of attributes, including geographical location, echelon, content type, freshness (e.g., newest first is most critical), and time constraints (e.g., only newer than a particular date).

  For a given subscription the server maintains a list of content on its peer server matching the subscription. The two servers build this list during initial synchronization and maintain it as new content is received. If an interruption occurs on the connection between the servers, they each revalidate the list of what the remote server has, and then add to the local outgoing queue any content that was acquired during the network outage matching the remote peer's subscription, completing the transfer of any missing content. The creation of the list of content on the remote end is performed in chunks, utilizing checksums of daily portions of the subscription. This overall scheme greatly minimizes the transfer of data between peers, and prevents duplicate transmission of content on the WAN.

  If a soldier wants to retrieve a piece of content that did not match his local servers subscription, a targeted replication is initiated in real time to pull that piece of content from the closest server, and then displays it to the soldier. Once a piece of media or a picture has been viewed by a soldier on a particular server, that media is retained on that server and is then available for immediate viewing by any other soldier using that server, even if network connectivity to the rest of the TIGR network is down. The content is stored locally on each intermediate server that receives it, enforcing a design goal of never transmitting the same content over a particular tactical link more than once.

TIGR SERVERS AND CLIENTS

When deployed to the battlefield, servers are configured based on their role and where they are located in terms of geography, network connectivity, and military organization.

All servers in a TIGR network store the metadata for all TIGR content in a local search index. This helps ensure that when a soldier performs a search within TIGR, they will not be missing critical high-level information even when wide area network connectivity is unavailable. When a search is performed, the metadata in the SQL database is used to retrieve the items best matching that query. The query might be geographically focused, within particular time periods, constrained to particular types of events (e.g., found IED's), and so forth. The retrieved items can then be opened in the user interface to study text and multimedia content. When viewing a report, a media item would typically be viewed at reduced resolution. If there is a need for higher resolution item, it might be already in the local repository, or it might be elsewhere in the network. If it is remote, it can be retrieved through the targeted replication capability. Any media not matching the local subscription will be unavailable if wide area network connectivity is down, and the soldier is notified to try again when connectivity is restored.

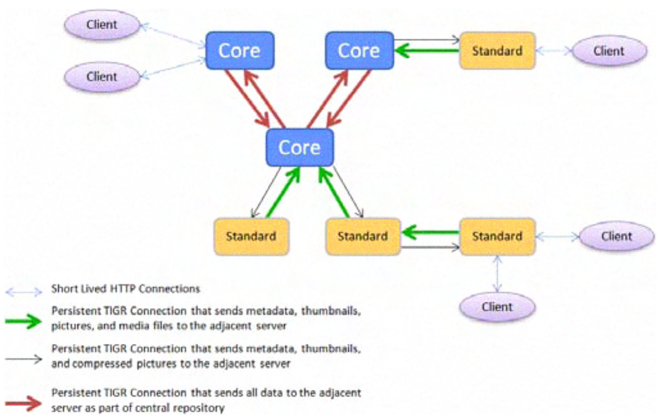

Figure 2 – Sample TIGR Logical Network

An example TIGR logical network is shown in Figure 2, illustrating the data flows between clients and servers of different types, detailed below.

The typical server categories and their respective policies include:

- Core servers – These are typically located in stable locations – large bases with reliable facilities including good network connectivity. These servers provide repositories of content in addition to serving as key participants in the distribution topology. They are home to search information (metadata), thumbnails, pictures, and media files (PowerPoint, videos, etc.) for all content created in a theater of operation. They also have a copy of all map imagery for the entire theater. In some cases, data transfer subscriptions for core servers might include restrictions on large media, but in other cases these sites might provide complete backups of the larger, full quality media.
- Standard servers – These are typically located in smaller forward operating bases and combat outposts, which are relatively fixed but are not as richly equipped in terms of network connectivity and other amenities. These servers provide local content as well as serve as conduit by which content is obtained from the other servers in the network. Standard servers have detailed map imagery for the area of operations of the unit in which they resides, plus any other imagery deemed useful. They contain search information, thumbnails, and compressed versions of pictures for all content created in a theater of operations. Subscriptions typically include restrictions on large media, for example allowing large images to be transferred inward to the core servers for analysis and archiving, but not outward to the edge in order to limit bandwidth impact. A directed replication request is made to the TIGR network to get any non local media or full size pictures requested by a user.
- Clients – Clients can be either web browsers on the same computer where the TIGR server is installed, or a web browser on other computers on the same LAN as one of the servers.

A typical TIGR subscription is to have all data created on the TIGR network flow to core servers to preserve data integrity (backups of all TIGR data can be performed in a single location). It is also policy to have multiple core servers in any deployed TIGR network and to have those servers deployed in separate geographic locations to preserve data integrity.

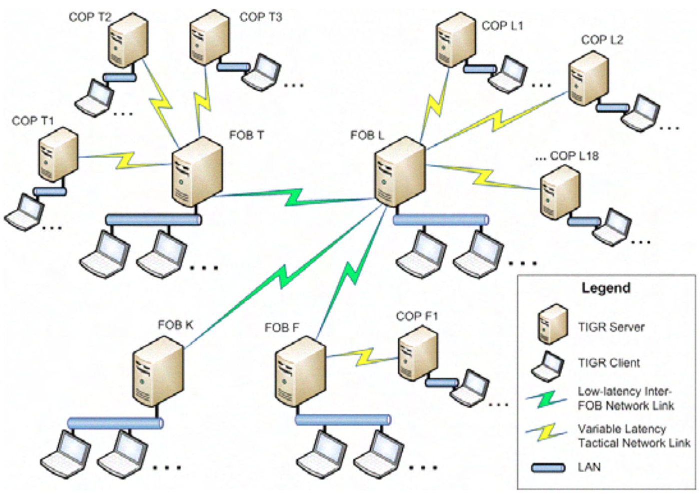

Figure 3 – Notional Physical TIGR Deployment

A notional TIGR server and client deployment is depicted in Figure 3, where FOBs L, F, T, and K have core servers while the remainder of the servers are standard servers. Core servers are typically rack mount 2U servers with significant disk space, while standard servers are often laptops given the increasing capabilities of such platforms.

NETWORK UTILIZATION

The TIGR system was designed to be network-friendly because it must share limited tactical communications resources. Measurements from hundreds of TIGR servers in the Iraq and Afghanistan theaters demonstrate this has been attained.

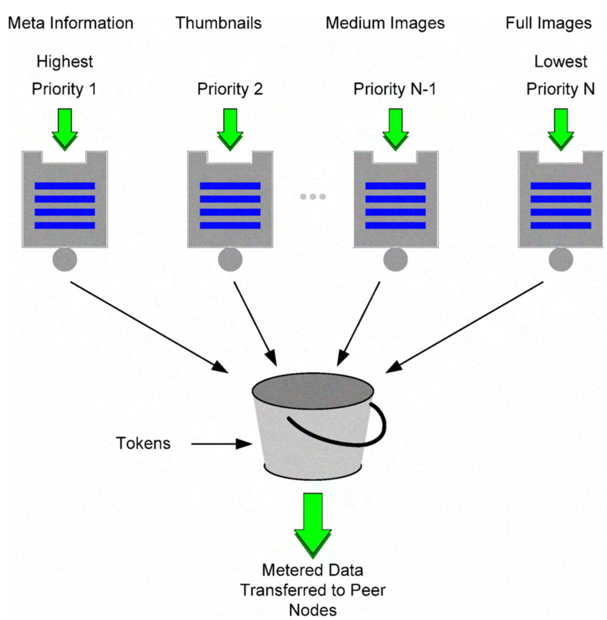

Figure 4 – TIGR Traffic Pacing Architecture

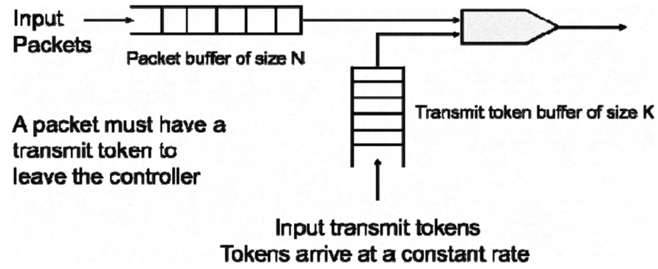

Figure 5 –Application Level Token Bucket

TIGR has incorporated many advanced networking concepts into the application layer, so that these features can be relied on regardless of lower level network transport capabilities. These include both priority and round robin queues [9] and a general open loop token bucket [10] used to control the pace of traffic as illustrated in Figures 4 and 5. The resources can be dynamically modified as a function of policy.

A representative flow of a newly generated report starts with a soldier creating the report and attaching media. The media is then transformed, and stored in the local repository. The meta information is indexed into the local search index. The metadata is sent from the originating server to all peer servers with relevant subscriptions, where they each add it to their respective search indices, along with pointers to the media stored at the originating node. Next thumbnails are transmitted, and finally the full media is replicated in the direction of the nearest core server. Figure 6 shows the average daily data rates from particular edge servers to the core during this same week. The average rate in this case was under 0.3 kbps.

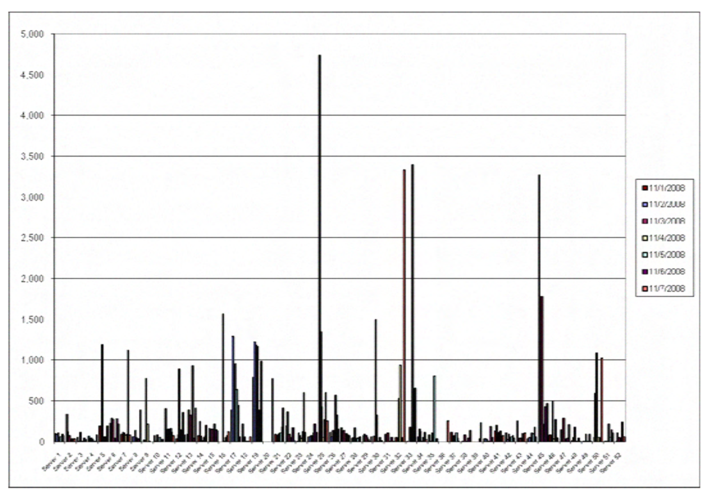

Figure 6 – Data Rates (bits/s) from an Edge to Core

As the data flows to the core servers, the subscriptions for what should be sent out to the other edge servers come into play. In this case content coming into the core from a particular edge server will then be selectively sent out to all of the other edge servers. The meta information will be sent first, and then thumbnails of attached media. Typically a subscription is utilized that also sends compressed media out to all edges, but this is optional based on network bandwidth availability. Figure 7 shows the data rates, averaged over a day, from core to edge TIGR production servers in Iraq during a typical week. The horizontal axis is made up of individual edge servers and the vertical axis is the average data rate in bits per second. The average rate is approximately 1.7 kbps in this sample. The core to edge traffic is much more uniform as expected, with the peaks being transmissions to servers that were offline for some period of time, and are recovering and catching up on historical data during the measurement window.

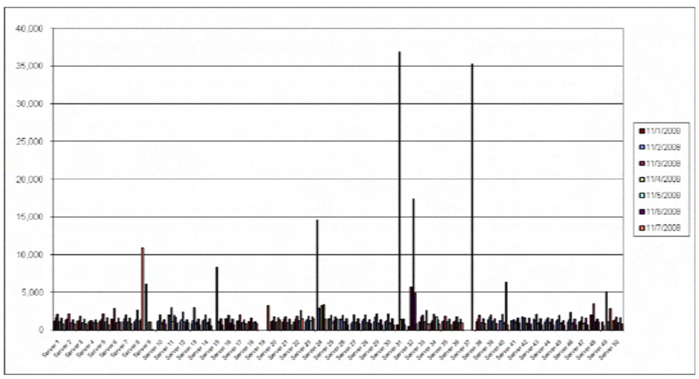

Figure 7 – Data Rates (bits/s) from Core to Edge Servers

These figures illustrate the general characteristics of TIGR data flows. The core to edge traffic is the aggregation of content at all servers in the network, whereas the edge to core traffic is only the content created at that particular edge server. The core to edge traffic is less variable than the edge to core traffic because it is an aggregation of sources.

Table 1 summarizes the average total daily traffic on the WAN being transmitted from an edge to the core. These are averages over the more than fifty TIGR servers in this measurement study.

Table 1 – Average Daily Bytes from an Edge to Core

Meta Info	Thumb-nails	Compressed Media	System Traffic	On Demand	Total Bytes
0.02 MB	0.004 MB	0.968 MB	0.77 MB	0.1 MB	1.85 MB

Table 2 summarizes the average total daily traffic on the WAN being transmitted from the core to the edge for these servers. Two different subscriptions are illustrated, the typical TIGR edge subscription that includes both thumbnails and compressed (800x600 pixels) pictures being sent to the edge, and the limited TIGR edge subscription that only sends thumbnails of images. Note the tradeoff between compressed image bandwidth and the on demand traffic that increases in response to user demand to view the higher quality images. In the case of the limited edge subscription, no compressed media is sent from the core server to the edge (a difference of 8.46 MB per day). This decrease is offset by soldiers retrieving some of this compressed media on demand using the targeted replication mechanism, in this case over 3 MB per day of user requested downloads.

Table 2 – Average Daily Bytes from a Core to Edge

	Meta Info	Thumb nails	Com-pressed Media	System Traffic	On Demand	Total Bytes
Regular Edge	3.4 MB	0.54 MB	8.46 MB	1.67 MB	0.2 MB	14.2 MB
Limited Edge	3.4 MB	0.54 MB	0 MB	1.67 MB	3.19 MB	8.8 MB

Figure 8 provides an example of TIGR traffic pacing using a policy. In this case, a large, unconstrained, full media transfer between core servers was underway between hours 40 and 78, after which time a maximum network utilization of 100 kbps was imposed.

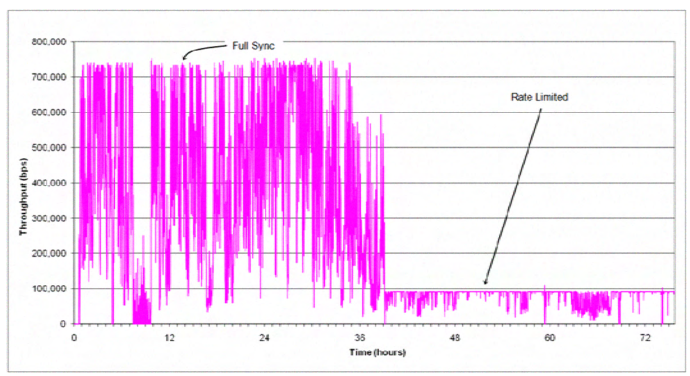

Figure 8 – TIGR Traffic Pacing Example

Another example of applying the bandwidth policy is to only allow media uploads during off usage hours, or cap the traffic at a 64 kbps maximum rate on a low bandwidth connection. A dynamic example can be defined using time such that TIGR will not transmit at all during a regularly scheduled VOIP conference that shares a JNN [11] link with TIGR.

These measurement results demonstrate that TIGR does not stress the WAN links interconnecting outposts and larger bases. The measurements also illustrate that there is flexibility in what is sent (e.g., compressed and uncompressed media) and at what rates, and this provides the ability to control data flows and network utilization.

CONCLUSIONS

The TIGR system has been designed to provide a fast, easy to use search for tactical users. Enabled by an adaptive data distribution architecture, the system performs searches locally to insure rapid response. TIGR supports analysis of data by a wide range of users, including those with high resource demands such as data miners who may require access to substantial amounts of raw, uncompressed media. TIGR is designed to manage communications limitations – its traffic is network friendly, and it is tailored to an environment where unplanned network outages are routine, and must be transparently managed. Deployments in Iraq and Afghanistan have validated the TIGR architectural concepts.

ACKNOWLEDGEMENTS

This work was supported in part by the Defense Advanced Research Projects Agency under contract W15P7T-07-C-P219.

The authors would like to acknowledge the contributions of Bert Hui, Brian Slaughter, Tai Truong, Beau Jepson, Joe Kearney, and Jacqueline Singh.

REFERENCES

[1] David Talbot, "A Technology Surges." Technology Review, vol 111, no 2, pp 70-74, April 2008.

[2] Major General Peter W. Chiarelli and Major Patrick R. Michaelis, "Winning the Peace: The Requirement for Full Spectrum Operations," Military Review (July-August 2005): 16.

[3] E. C. M. A. International, *ECMA-262: ECMAScript Language Specification*, 3rd ed., Geneva, Switzerland: ECMA (European Association for Standardizing Information and Communication Systems), December 1999.

[4] L. D. Paulson, "Building rich web applications with ajax," *Computer*, vol. 38, no. 10, pp. 14-17, 2005.

[5] C. J. Date and H. Darwen, "A Guide to the SQL Standard." Fourth Edition, Addison-Wesley, 1997.

[6] Open Geospatial Consurtium, Inc. *OpenGIS Web Map Service (WMS) Implementation Specification*, v1.3.0, March 2006.

[7] JEITA CP-3451 Exif 2.2.

[8] G. Salton, A. Wong, and C. S. Yang, "A vector space model for automatic indexing," *Commun. ACM*, vol. 18, no. 11, pp. 613-620, November 1975.

[9] L. Kleinrock, *Queueing Systems. Volume 1: Theory*, 1st ed., Wiley-Interscience, January 1975.

[10] J. Turner, "New directions in communications (or which way to the information age?)," *Communications Magazine, IEEE*, vol. 24, no. 10, pp. 8-15, 1986.

[11] "Tactics, Techniques, and Procedures (TTPs) for the Joint Network Node-Network (JNN-N)", *FMI 6-02.60*, September 2006.