

# TIGR: The Tactical Ground Reporting System

*Joseph B. Evans, University of Kansas*

*Benjamin J. Ewy, Michael T. Swink, and Steven G. Pennington, General Dynamics C4 Systems*

*Debra J. Siqueros, PEO C3T*

*Samuel L. Earp, Multisensor Science, LLC*

## ABSTRACT

The Tactical Ground Reporting System (TIGR) is a cloud-based application that facilitates collaboration and information sharing at the patrol level. DARPA began development efforts in 2005, and introduced the system to a single brigade in Iraq at the beginning of 2007. Its use spread to all Army units in Iraq and Afghanistan, as well as Kuwait, the Philippines, South Korea, and the Horn of Africa, and it has proven to be an indispensable tool for counterinsurgency operations. The system enables collection and dissemination of fine-grained information on people, places, and events, and offers a media-rich view of the battlefield with digital photos, videos, and high-resolution geospatial imagery. In 2011 TIGR transitioned to the Army and is now managed by the Project Manager for the Joint Battle Command-Platform (PM JBC-P), where it has expanded to include vehicle-based systems and other on-the-move operational scenarios. Despite the richness of its multimedia content, TIGR has successfully operated for years in multiple theaters of war with minimal impact on low-bandwidth tactical networks. This article describes the TIGR system and how it has evolved, as well as the challenges of designing and deploying a distributed system for disadvantaged networks, and the unique data dissemination architecture that has made TIGR such a success.

## INTRODUCTION

The Tactical Ground Reporting System (TIGR) [1] was designed by the Defense Advanced Research Projects Agency (DARPA) for use by patrol soldiers at the company echelon and below. These soldiers have adopted the system as their primary tool for capturing and sharing combat patrol information, and there are now more than 95,000 users of the system. At the time of writing, TIGR is actively used at hundreds of sites by thousands of deployed U.S. Army and British Army soldiers in Afghanistan and throughout the world, as well as all conti-

mental U.S. (CONUS) Army units undergoing pre-deployment training.

TIGR addresses a series of challenging requirements: the need to efficiently search large amounts of tactical data and the need to support rich multimedia attachments, using bandwidth-efficient thumbnails and compressed versions of the multimedia. TIGR was developed as a system tailored to the distribution and searching of rich media, and fills a gap between operations and intelligence by providing theater-wide distribution of patrol reports, significant actions, and other tactical information.

This article examines key design aspects of the system that are responsible for its widespread adoption: a web browser client that is intuitive, lightweight, and scalable, and a cloud-based architecture that makes extensive use of caching and compression to minimize network utilization and optimize the user experience over tactical networks with frequent outages. The article also provides a description of the user base and use cases of TIGR that drive the utilization of the system.

## TIGR USERS

TIGR was designed originally for patrol leaders and other operational users at the company echelon and below, and focused on post-patrol debriefs and pre-patrol mission planning. Rather than solely relying on higher echelon intelligence staff for relevant information, a platoon leader or company commander could search TIGR for all information on the area of interest, and also directly enter reports and pictures into TIGR after a mission. TIGR breaks from the traditional hierarchical, bottom-up filtered information flow of reporting, and instead builds on the successes of direct peer-to-peer collaboration that originated in efforts such as CompanyCommand [2] and CAVNET [3].

Since its initial fielding, TIGR has not only strengthened information sharing among patrol leaders but has allowed the common operational picture to be shared among disparate groups on the ground. For example, TIGR allows “land-owning” units to collaborate with other units

Portions of this work appeared at MILCOM 2009.

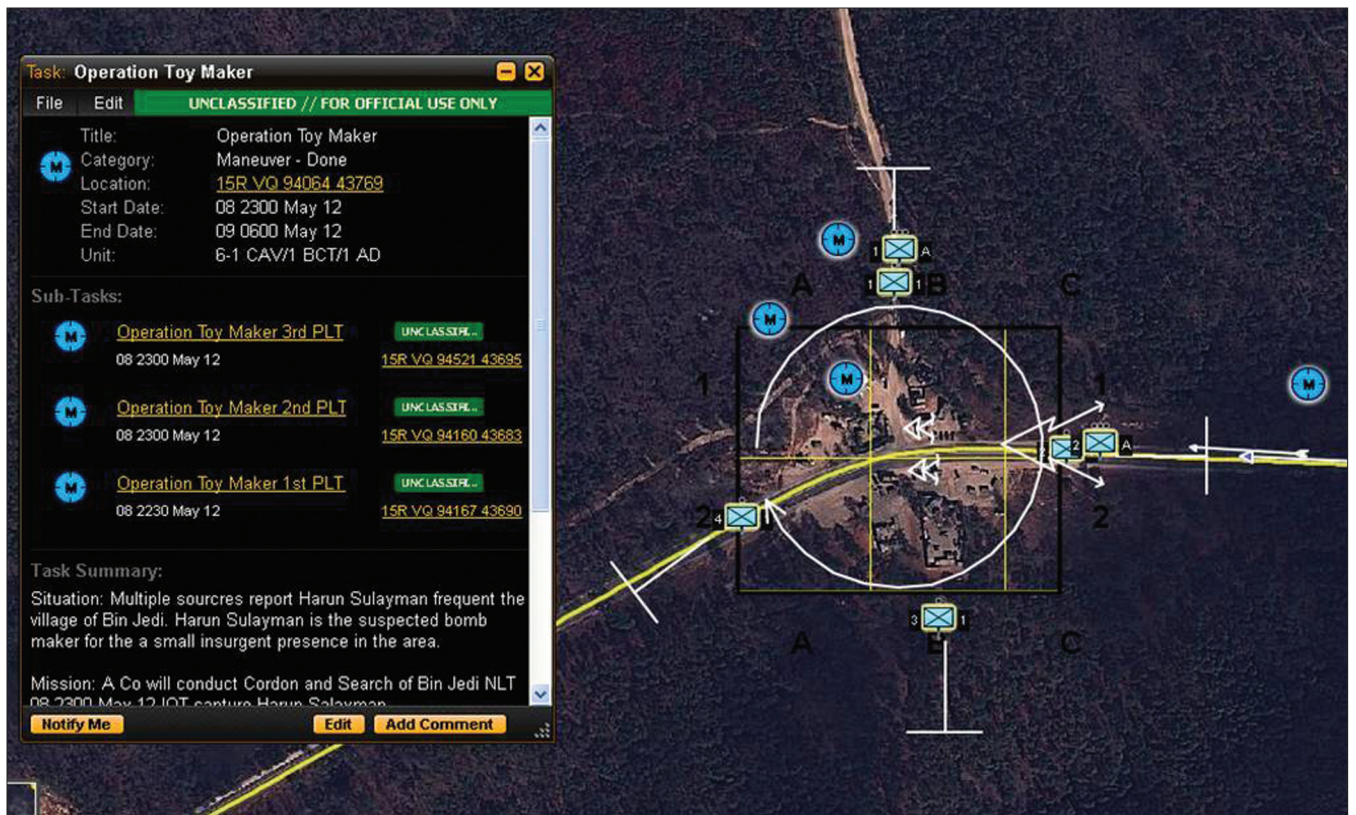


Figure 1. Example task window, and the corresponding icons and operational graphics on the map.

that work in or travel through their area of responsibility (AOR) such as the Military Police (MP), Combat Service Support (CSS) and logistics units, and Human Terrain Team (HTT) members. All these groups can share and utilize the information they are collecting and tracking. Also, notably, TIGR facilitates information push downward to the “tip of the spear” so that patrol leaders now have access to detailed improvised explosive device (IED) threat analysis and ultra-high-resolution map imagery previously only available at the higher echelons. Patrol leaders can view geographically specific warnings posted by the intelligence teams of impending attacks as well as possible IED plantings along their patrol route. Traditionally, information was vetted and filtered as it went up the chain of command; TIGR promotes peer-to-peer information sharing at every echelon. TIGR was originally focused on planning and debriefing activities, but has been extended to provide real-time situational awareness information such as displaying position location information. Beginning in 2012, TIGR systems began being deployed as part of the Joint Capabilities Release (JCR) software update [4] in some vehicles, and in 2013 TIGR systems are deployed in Warfighter Information Network — Tactical (WIN-T) [5] commanders’ vehicles allowing for use on the move, and during the execution phase of missions.

## TIGR INFORMATION MODEL

One requirement of designing an intuitive user interface is the need for an information model that is both simple and general. Traditionally,

military reporting consists of filling out forms with many obligatory fields that reflect the information needs of higher echelons. The structured fields often become irrelevant or obsolete as the fight evolves. The TIGR data model comprises just five types of data objects — **Events**, **Places**, **Reports**, **Tasks**, and **Collections** — and their content is unobtrusively structured. The number of mandatory fields is kept to the minimum.

**Events** are represented as icons on the map and can describe non-kinetic activities (e.g., meeting a local leader) as well as enemy activity (e.g., small arms attack). TIGR Events are created by selecting the Event Category from a drop-down menu and dragging the icons onto the map. An Event’s key fields are the time of occurrence and location. Additional structured data can include its title, creator name, creator unit, drawings on the map, narrative summary, and any relevant media. Media commonly consist of photos, but may also include video, audio, Microsoft® Office documents and other files.

**Places** are similar to Events but only have a location, not a time of occurrence. Places are things such as a power plant, a mosque, or other infrastructure.

**Reports** are summaries that pull together Events and Places with additional narrative context that may describe an overall mission and include the route or track taken if appropriate. Reports are created when a soldier returns from a patrol and describes the details of his/her mission, including the Places visited and the Events that occurred while on patrol.

**Tasks** are used for mission planning and execution purposes. Using Tasks, a user can place a



Figure 2. High-resolution aerial imagery with ground-level view perspective.

Task on the map, assign a unit to perform the desired task, assign a date/time, attach relevant media such as orders, and create and associate sub-tasks (i.e., for platoon- or squadron-level sub tasks).

**Collections** are a user-defined group of Events, Places, Tasks, and Reports. It is similar to a folder; examples include all detainees for the month of April, high value target reporting for a particular area of operations, or a group of all reconstruction projects in a particular AOR.

## TIGR CLIENT

The TIGR client is browser-based, and has a simple interface that supports the full range of operations for a lower echelon unit. TIGR was designed using the look and feel of popular social media websites, making its usage very intuitive for younger generations of soldiers. There are specially optimized interfaces available for smaller screened laptops, as well as touch screen panels found in vehicle installations. The client has a presentation based on maps, and makes use of menus and tabs. There are five main tabs: Search, Create, Forums, T-Mail, and Profile. The Search for content is done with the map interface. The Create tab allows content such as Events and drawings to be added to the map. The Forums tab provides access to a forum where users in the TIGR community exchange information of tactical relevance as well as postings related to desired feature enhancements in TIGR or technical problems encountered. The T-Mail tab accesses the TIGR internal messaging capability. The Profile tab accesses the user profile including the contact information for a user, which is especially valuable for users to collaborate outside of TIGR. TIGR makes it easy to find and contact

the soldier on the ground with direct knowledge of the people, places, and events of interest. This direct tie between the data in TIGR and real people provides a level of trust and believability that is unique in tactical systems.

Information is presented in a spatial context, using icons or drawings as appropriate. The data in TIGR is searchable by keyword, date time group range, category, and unit. Most important, data is searchable using the map — a soldier can draw a box, polygon or route on the map and constrain the search to the designated area. Searches can be modified, for example, by resizing the geospatial search or changing the hostile event categories of interest. When search criteria are updated, the results automatically appear both geospatially on the map as well as in a list view.

The map icons are active elements: a mouse rollover of an icon will show thumbnails of attached pictures, and a mouse click will bring up a window with detailed information about the icon. For example, a click on an Event icon will reveal the detailed narrative of the event that occurred, and all of the attached media and photo thumbnails. A mouse click on one of the thumbnails will bring up a larger view of the photo.

The amount of content in each Report, Event, and Place varies greatly in terms of textual detail and the number of media added. A user in Afghanistan, for example, combined the use of TIGR with a GPS-enabled camera to capture a significant number of photos during a fight. The resulting report, with over 100 Mbytes of photos, would have been impossibly hard to share using traditional means, but TIGR makes widespread dissemination possible and information discovery easy as each photo is represented on the TIGR map with an icon, rather than being just a file attachment in a large assortment of files.

High-quality terrain imagery is of key importance to tactical users who rely on maps for mission planning and analysis. TIGR currently provides the best tactical map imagery to users by tapping into a variety of different sources. Commercial geospatial imagery from a variety of sources as well as the National Geospatial-Intelligence Agency (NGA) provides 0.6 to 1 m per pixel resolution imagery covering hot spots such as Afghanistan and the Horn of Africa. TIGR utilizes a forward deployed team to process and distribute the very latest aerial imagery collected in theater (e.g., Buckeye [6] and PeARL™ [7]) to provide even higher-resolution views. TIGR utilizes elevation data from satellite and Light Detection and Ranging (LIDAR) to provide terrain information to enhance mission planning using the built-in line of sight and intervisibility tools. All of this imagery is provided by an OpenGIS standards compliant Web Map Service (WMS) [8] server, which allows other systems to leverage the local imagery repositories. Imagery is stored locally on servers throughout the cloud to reduce WAN bandwidth requirements. Due to constraints imposed by available storage space, a typical server is configured with a base set of NGA imagery for the whole country, enhanced with a higher-resolution imagery set that covers the particular region the server is supporting.

TIGR defaults to automatically selecting the best imagery as a user pans or zooms the map using an algorithm that takes into account the resolution and collection date of the imagery. The user can also choose among all available imagery choices for a particular area so that they can see a topological map, an older map showing a region prior to a reconstruction project, and so on. In addition to elevation, satellite, aerial, and topological map imagery, TIGR offers a 360° ground-level view of approximately 9000 street miles in Iraq and 1000 street miles in Afghanistan. This data was originally collected as part of the Rapid Equipping Force [9] sponsored TOURIST project, but PM JBC-P has begun integrating ground-level perspective imagery from other sources, and this is expected to increase significantly as the typical vehicle sensor packages expand.

## TIGR SERVERS

TIGR is a distributed cloud-based system designed to operate on disadvantaged networks. TIGR uses persistent TCP connections over existing network assets — tactical network assets under unit control and the strategic network — as available. The TIGR cloud consists of servers located at large forward operating bases (FOBs), as well as remote outposts such as joint security stations (JSSs) and coalition outposts (COPs).

Server hardware varies from multi-CPU rack-mounted servers with uninterruptible power supplies to laptop computers. The laptop servers are deployed in large numbers at smaller outposts. A snapshot of the Iraq server topology when at its most distributed level is depicted in Fig. 3.

The server distribution is critical to providing reliable functionality on tactical networks.

Tactical networks are unreliable for many reasons; for example, power disruptions or difficult terrain can cause outages. Hence, link outages and disruptions are the norm rather than the exception. Multiple primary replication connections are utilized at key servers, as well as secondary connections configured to be utilized as alternate server options, particularly for mobile servers. By placing servers close to the users at the edge, we can ensure that soldiers can continue to exercise TIGR functionality locally even when all of these network and server redundancy mechanisms are insufficient. Once the WAN connectivity is restored, TIGR's synchronization mechanism will automatically exchange the data generated during an outage.

The TIGR architecture [10] consists of clients, mobile servers, edge servers, and core servers. Functionally, the difference between core servers and the smaller edge and mobile servers is simply the amount and type of information that is synchronized with its peers. 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 WAN connectivity is unavailable. When a search is performed, the metadata in the local database is used to retrieve the items best matching that query.

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 a higher-resolution item, it might already be in the local repository, or it might be elsewhere in the network. If it is remote, it can be retrieved through a targeted on-demand replication capability. Any media not in the local repository will be unavailable if WAN connectivity is down, and the soldier is notified to try again when connectivity is restored.

An example TIGR logical network is shown in Fig. 4, illustrating the data flows and policies between clients and servers of different types:

- Core servers: These are typically located in stable locations — large bases with reliable facilities including good network connectivity. These servers provide comprehensive 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.
- Edge servers: These are typically located in smaller forward operating bases and combat outposts, which are relatively fixed but not as richly equipped in terms of network connectivity and other amenities. These

*TIGR is a distributed cloud-based system designed to operate on disadvantaged networks. TIGR uses persistent TCP connections over existing network assets — tactical network assets under unit control and the strategic network — as available.*

servers provide local content and also serve as a conduit through which content is obtained from the other servers in the network. Edge servers have detailed map imagery for the area of operations of the unit in which they reside, 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, plus full media for the AOR served by the edge server. If a user accesses media that is not local to the edge, a directed media request is sent to the edge server's peers to retrieve the data, and it is then cached locally for future reuse.

- Mobile servers: The touchscreen optimized TIGR servers may operate disconnected on low-bandwidth vehicles such as JCR systems, or connected as mobile edge servers

in vehicles with higher bandwidths such as commanders' vehicles in units with WIN-T Increment 2 [5].

- Clients: Clients can be either web browsers on the same computer on which the TIGR server is installed or a web browser on other computers on the same LAN as one of the servers.

The typical TIGR data policy is to have all data created on the TIGR network sent from the edges to the core servers to preserve data integrity and reliability (backups of all TIGR data can be performed at any single geographically dispersed core server location). The core servers will propagate back down to the edge servers only the metadata, thumbnails, and compressed versions of media to reduce bandwidth used while allowing for a complete local search repository.




Figure 3. TIGR cloud topology in Iraq at its largest.

## TIGR INTEROPERABILITY

Interoperability and data exchange is an important element for any tactical system. There are other databases that hold data of interest to TIGR users, and users of other systems that are interested in analyzing TIGR data. TIGR is capable of doing data exchange manually by exporting and importing data in formats used by Microsoft Office, FalconView™, Google Earth™, and ESRI ArcGIS®. TIGR has used manual imports to bring legacy repositories such as large spreadsheets of places into TIGR. TIGR users have manually exported TIGR data for visualization in tools such as Analyst's Notebook®, Axis-Pro®, and Google Earth.

TIGR provides a SOAP-based web service to allow external systems to access and manipulate TIGR data in an automated fashion. This capability is used to pull TIGR data into other databases including Combined Information Data Network Exchange (CIDNE®), Command Post of the Future (CPOF), Palantir, and Distributed Common Ground Station — Army (DCGS-A). TIGR automatically imports significant activities (SIGACTs) and a number of other reports from CIDNE, CPOF, and MarineLink, among others.

## SERVER MANAGEMENT

Because deployed TIGR systems are composed of servers utilizing disadvantaged tactical communications networks, a unique management system is required to overcome connectivity, delay, and bandwidth issues. The TIGR monitor provides four capabilities: overall TIGR network situational awareness, detailed interserver traffic monitoring, information on server configuration and software status, and remote patching. The management system utilizes monitoring processes that run on each TIGR server. These monitoring processes collect information from the various TIGR processes as well as the host operating system, and then summarize this information for periodic updates to a centralized management server. The central management server provides web-based graphical visualization of the cloud status, including server status and connectivity, synchronization policy settings, and server software revision level (the topology diagram in Fig. 3 is taken directly from this graphical interface). The central monitor allows for remote network monitoring, the collection of soldier usage statistics, and configuration updates. The central management server also has the ability to reach out to remote servers to selectively update portions of the TIGR software in a bandwidth-efficient yet persistent manner.

## NETWORK UTILIZATION

The tactical networks deployed in theaters like Iraq and Afghanistan are often taxed, and the network utilization of an application is a critical parameter in successfully coexisting with the other applications on the tactical WAN.

A representative flow for a newly authored report starts with a user creating the report and attaching media. Thumbnails and compressed versions of the media are created and stored in

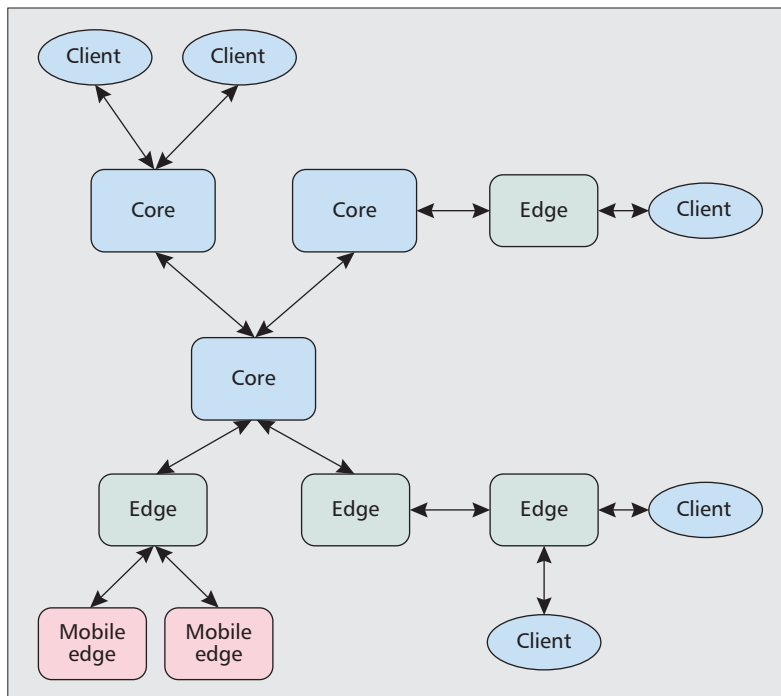


Figure 4. Sample TIGR logical network.

the local repository. The metadata is indexed into the local search index. Next, a priority queuing system first sends the metadata toward the core of the network. Thumbnails are the next queue transmitted, then the compressed versions, and finally, the full size media is sent to the core. On average, in Afghanistan we see edge servers transmit approximately 2 Mbytes/day [10] up toward their closest core server.

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 metadata 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. Usually, in Afghanistan a core server will transmit 15 Mbytes/day of data down to each edge server [10].

The general characteristics of TIGR data flows is such that 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 and subject to more variability due to the changing nature (in particular, media attachments) of any given day's reporting requirements. The core to edge traffic is less variable than the edge to core traffic because it is an aggregation of sources.

Traffic collected from Afghanistan provides an opportunity to explore these traffic characteristics in more depth. Figure 5 plots the daily bytes sent from an edge to the core for 100 servers for seven days each. The resulting histogram has a Poisson distribution that is characteristic of networking systems describing the

traffic being offered into the system. The samples in the tail of the distribution are typically either a large media attachment that dominates a particular sample or the result of a long outage at the tactical edge if WAN connectivity is lost for multiple days; when the edge server comes online, it has a larger than normal daily transmission as it must send multiple days worth of reporting to the core.

Figure 6 shows the distribution of daily bytes sent from the core servers back out to the edge. This is the histogram of received bytes for 100 servers, one sample for each of seven days. Here the traffic distribution is Gaussian, as might be expected by the central limit theorem, because this traffic is the aggregate of a large number of independent data sources. This data shows the effect of multiple-day outages on data sent from the core to the edge. The average number of bytes sent was approximately 15 Mbytes/day, but also discernible in the tail of the distribution are harmonics at the two-day outage and three-day outage values of 30 and 45 Mbytes/day.

This data illustrates that TIGR traffic is well behaved and does not stress the WAN links

interconnecting outposts and larger bases, and that the system has flexibility in what is sent and at what rates [11, 12]. This provides the ability to control data flows and network utilization to support a wide variety of network deployments.

## TIGR DATA

TIGR has become a very rich historical data repository filled with multimedia information. There is over 170 Gbytes of data in the Iraq repository, and the biggest single weekly growth was over 2 Gbytes in the week of December 21, 2008. The combined repositories in Afghanistan contain almost 300 Gbytes of soldier-generated reports and media. The single biggest week of data being added in Afghanistan was the week of June 12, 2011, when over 3.7 Gbytes of data was added, including more than 5000 photographs.

The consolidated TIGR repositories as of January 2013 have more than 618,000 Events, 343,000 Places, 172,000 Reports, 29,000 Collections, and 293,000 photographs along with another 281,000 media files (PowerPoint presentations, Word documents, etc.).

In 2012, in Afghanistan alone, 9000 PowerPoint attachments taking up 24.5 Gbytes of disk space were added, along with 20,750 photographs taking up 55.7 Gbytes of disk space. All of this data is completely searchable using keywords, units, categories, geographic regions, and time. In 2012 soldiers performed almost 3 million searches in Afghanistan. In a typical week in Afghanistan during 2012, 3000 soldiers logged into TIGR more than 9000 times to perform 56,000 searches, and create 3600 Events, 600 Places, and 900 multimedia documents adding 2 Gbytes of data to the repository, all while operating in a low-bandwidth network environment with high latency, power outages, and difficult environmental conditions.

## CONCLUSIONS

As the Army moves toward its Common Operating Environment, the TIGR model (web-based, WSDL interfaces, use of standard open technologies, cloud-like data provider and syncing mechanism) easily supports the new requirements across multiple computing environments to include mounted, handheld, sensor, and command post. TIGR is unique because of its combination of collaborative, intuitive products that have enabled soldiers to find emergent usages for the product itself, ranging from using TIGR as a “yellow pages” and a logistics tracking tool, for security surveillance, and more.

## ACKNOWLEDGMENTS

The authors would like to thank Dr. Mari Maeda for her contributions to this work. The technology described in this article was developed under US Government contracts W15P7T-07-C-P219 and HR0011-10-C-0030.

## REFERENCES

- [1] D. Talbot, “A Technology Surges,” *Technology Rev.*, vol. 111, no. 2, Apr. 2008, pp. 70–74.

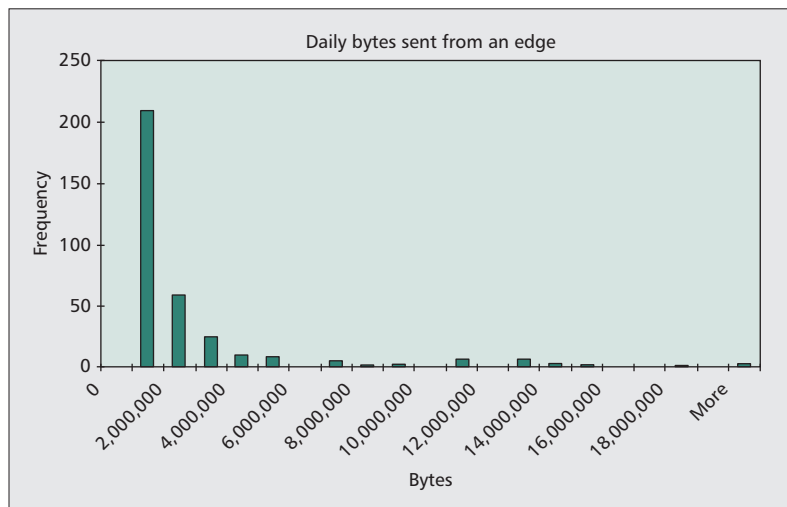


Figure 5. Traffic histogram, edge to core.

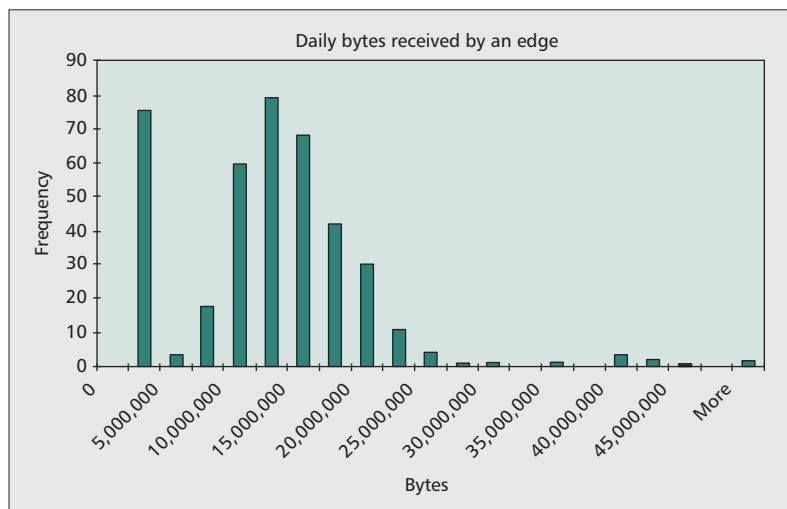


Figure 6. Traffic histogram, core to edge.

- [2] N. M. Dixon *et al.*, *CompanyCommand: Unleashing the Power of the Army Profession*, Center for the Advancement of Leader Development & Organizational Learning, 2005; [CompanyCommand.army.mil](http://CompanyCommand.army.mil).
- [3] Maj. P. Michaelis, speaking on *FRONTLINE: A Company of Soldiers: Innovating & Improvising*, WGBH and PBS, Feb. 22, 2005, <http://www.pbs.org/wgbh/pages/frontline/shows/company/lessons/>.
- [4] C. Schwerin, *Army Fields Next-Generation Blue Force Tracking System*, <http://www.army.mil/article/61624/>
- [5] K. ain, *Army Rolls Out Networked Capability Set 13*, <http://www.army.mil/article/87085/>
- [6] D. Miles, *Buckeye System Brings New Digital Capability to Warfighters*, American Forces Press Service, June 13, 2006.
- [7] Urban Robotics, Inc., *PeARL Digital Imaging Sensors*, [http://www.urbanrobotics.net/home.php?page=products\\_pearl](http://www.urbanrobotics.net/home.php?page=products_pearl).
- [8] Open Geospatial Consortium, Inc., "OpenGIS Web Map Service (WMS) Implementation Specification," v1.3.0, Mar. 2006.
- [9] D. Miles, *Rapid Equipping Force Speeds New Technology to Front Lines*, American Forces Press Service, Aug. 12, 2005.
- [10] B. J. Ewy *et al.*, "TIGR in Iraq and Afghanistan: Network-Adaptive Distribution of Media Rich Tactical Data," *IEEE MILCOM '09*, Oct. 2009.
- [11] *Tactical Ground Reporting Network Bandwidth Analysis Report*, PEOC3T Systems Integration and Engineering Division Futures Office, 8 Mar. 2007.
- [12] S. West *et al.*, *Tactical Ground Reporting System In Theater Assessment (FOUO)*, Fort Monroe, VA, 24 Jan. 2008.

## BIOGRAPHIES

BENJAMIN EWY ([benjamin.ewy@gdc4s.com](mailto:benjamin.ewy@gdc4s.com)) is a member of technical staff for General Dynamics C4 Systems. He has been involved in design, implementation, and deployment of the TIGR program since 2006. He holds an M.S. in electrical engineering and a B.S. in communications engineering, both from the University of Kansas. His research interests include probabilistic network management and cognitive networks.

JOSEPH B. EVANS ([evans@ku.edu](mailto:evans@ku.edu)) is the Deane E. Ackers Distinguished Professor of EECS at the University of Kansas (KU). He has served as director of the Information & Telecommunication Technology Center and director of Research IT at KU, and as a program director at the National Science Foundation. He has co-founded startups including a network gaming company acquired by Microsoft in 2000 and the defense-oriented venture acquired by General Dynamics in 2010 that developed TIGR. He received a Ph.D. from Princeton University in 1989.

DEBRA SIQUIEROS ([debra.siquieros@gmail.com](mailto:debra.siquieros@gmail.com)) was the chief engineer of the TIGR System for PM FCB2 and worked for the Army for over 30 years. She is in the systems engineering Ph.D. program at Stevens Institute of Technology, from which she also holds a Master's of Engineering in network information systems. Her research interests include patterns for security systems of systems design. She also writes and lectures for her consulting firm, Red Hot Security.

SAMUEL EARP ([searp@multisensorscience.com](mailto:searp@multisensorscience.com)) received his Ph.D. in electrical engineering from Duke University. He has held a wide variety of academic, technical, and managerial positions in the defense industry over the last 30 years. For the last decade, he has been a principal of Multisensor Science, LLC, and a consultant to DARPA.

MICHAEL SWINK ([michael.swink@gdc4s.com](mailto:michael.swink@gdc4s.com)) is a software engineer with General Dynamics C4 Systems. He is the TIGR systems architect and has been responsible for the TIGR data storage and replication system since 2006. He holds an M.S. in electrical engineering and a B.S. in communications engineering, both from the University of Kansas.

STEVEN PENNINGTON ([steven.pennington@gdc4s.com](mailto:steven.pennington@gdc4s.com)) is a member of technical staff with General Dynamics C4 Systems. He has been a systems architect responsible for the geospatial capabilities of the TIGR program since 2006. He holds an M.S. in electrical engineering and a B.S. in communications engineering, both from the University of Kansas. His research interests include geospatial analysis, software defined radios, and RF spectrum reuse.