BEYOND RAW IP: MUOS ARCHITECTURE AS A PLATFORM FOR FUTURE SERVICES

Adam Bawor
General Dynamics C4 Systems
Scottsdale, AZ

ABSTRACT

The widespread migration of communication solutions away from circuit-switched architectures to "IP" represents a profound paradigm shift in the way that features and services are conceived, developed and deployed. The basic Internet Protocol (IP) connectivity unifies the method in which various components communicate with one another irrespective of their lower layer transport technologies or higher layer applications. Mobile User Objective System (MUOS), which is a global, satellite-based communication network replacing the legacy UHF Follow-On (UFO) system, features an all-IP architecture where properly layered protocols at given node implement a feature or service. The precise set and configuration of end-user features is dependent on the chosen implementation of the terminal, which allows tailoring of services to the application. The implementation of the underlying radio network leverages proven commercial products – specifically, the 3rd generation (3G) Wide Code Division Multiple Access (WCDMA) radio interface which has been customized to meet the specific system requirements, as well as the constraints of satellite radio links. Though global in reach, MUOS comprises only a handful of facilities and utilizes Defense Information Systems Agency (DISA) services for backbone IP communication. Leveraging the latest commercial technologies allows MUOS to adapt existing products and services to speed up implementation of a design that not only meets the current set of requirements, but also provides a platform for the future. Commonality of design with the commercial world will result in the possibility to both provide new services and enhance the existing ones with emerging technologies. In this paper, we discuss the layered, all-IP commercially-based architecture of MUOS, and how its design can keep up with the evolving communication needs of the warfighter.

INTRODUCTION

The continued evolution of the military strategy for conducting business operations, warfare, and enterprise management towards the Net-Centric Operations and Warfare (NCOW) model is resulting in increased reliance on the availability of communication links and ever more complex applications, which offer services beyond basic voice and IP (e.g. Quality of Service (QoS) controls, resource management, support for mobility, multihoming, etc.) MUOS is likely to play a preeminent role in this emerging NCOW vision because it is a system founded on enabling technologies: it offers global reach of its communication paths and the potential to exploit service enhancements in line with commercial trends.

MUOS services include group connectivity, dial-up-like IP connectivity and call-like IP connectivity between two end-points. In this paper, we discuss only the last of the these – the End-to-End (E2E) services – which comprise voice services (G.729a,b and enhanced Mixed Excitation Linear Prediction (MELPe)) and data services (with Burst, Stream or Flow QoS controlled through appropriate configuration of the Over-the-Air (OTA) link and Differentiated Services (DiffServ) architecture on the ground). Both IPv4 and IPv6 are supported to the User Entry (UE) Segment, as is mobility (including inter-satelite handoffs) and selection of the Access Point Names (APNs) (e.g. addresses published or unpublished in the Domain Name Server (DNS)). Since all bearer data is transported over IP between the endpoints, MUOS is a type of a Voice over IP (VoIP) network. However, MUOS does not connect voice services directly to any external IP-based voice network (Defense Switched Network (DSN) is an SS7 network).

Figure 1 shows the overall physical view of the MUOS system.

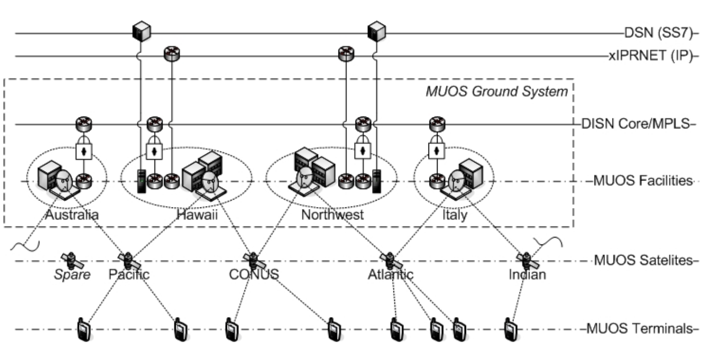

Figure 1 MUOS – Global View

Four MUOS satelites provide footprint coverage of the entire globe (save for Polar Regions) and relay signals between MUOS terminals and four communication sites distributed around the world, all of which are inter-connected via leased Defense Integrated Switched Network (DISN) lines. All sites maintain connectivity with two satelites but only the facilities located on the U.S. soil interconnect to external military networks: NIRPNET, SIPRNET, and (SS7-based) DSN.[1]

ADAPTATION OF COMMERCIAL PRODUCTS

The basic idea behind MUOS is adaptation of a commercial WCDMA system to provide access service for globally-distributed subscribers. This can be visualized as a simple replacement of a ground antenna with a geostationary satelite, as shown in Figure 2.

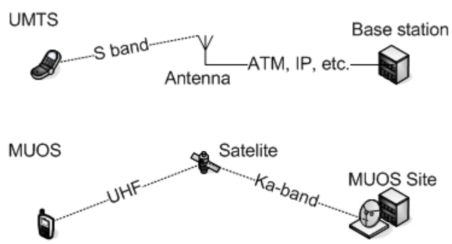

Figure 2 Adaptation of UMTS Architecture in MUOS

The most important change from a commercial system is adaptation of the OTA interface to the radio interface and the requirements of the inter-satelite links so that IP connectivity can be provided from the terminal to the Ground System (GS). Once the IP connectivity is provided, additional features and services can be built on top of it. In MUOS, the E2E services (such as locating the terminating party) are built around the IP Multimedia Subsystem (IMS) architecture, which is a Session Initiation Protocol (SIP) based architecture defined for the cellular industry. Although the baseline set of services and features targets requirements, the system can accommodate a wide range of additional services through incremental upgrades. These are discussed later in this paper.

The interconnectivity between the sites is provided via leased DISN lines, which utilize a MultiProtocol Label Switching (MPLS) network. A Service-Level Agreement (SLA) is put in place to insure appropriate level of QoS (including bandwidth requirements) associated with each link, which means that the only links that can get congested are the OTA links (the number of DS0s out to DSN can be increased to any number).

The MUOS terminal is designed around the Joint Tactical Radio System (JTRS) platform such that, in theory, only a Common Air Interface (CAI) Waveform may need to be developed for MUOS. If successfully pursued in practice, this model will result in reuse of existing technology, reduction of risk, cost and time required to supply a working terminal to the customer.

To make MUOS a secure VoIP system, MUOS features a custom solution for transfer of signaling messages – SIP messages are sent on a dedicated Packet Data Protocol (PDP) Context (including the IP address) and over a MUOS-internal subnetwork, which is not shared with user bearer data. Furthermore, the SIP-based call control logic in the UE is isolated from the MUOS Compatible Terminal (MCT), which means that no application running on the MCT may gain connectivity to the MUOS SIP infrastructure.

The voice codecs used in MUOS (MELPe and G.729ab) are both widely used (in, respectively, military and commercial applications) and will be available in the MCT. Data connectivity to DSN is based on commercial standards, too. The V.150.1 Modem Relay technology is used to facilitate transmission of data (including secure voice over devices such as the Secure Communications Interoperability Protocol (SCIP) phone) between the MCT and DSN users (clear, secure and clear-before-secure modes are all supported by MUOS, though not necessarily by all types of MCTs). A modified Commercial Off-The-Shelf (COTS) gateway (GW) is used to interwork MUOS-internal SIP-based network with the SS7-based DSN.

Adaptation of the commercial WCDMA system where only the Packet-Switched (PS) domain is used, required the support for aggressive header compression. MUOS provides a “zero-overhead” compression mechanism for all voice services (even with Discontinuous Transmission (DTX) turned on). “Zero-overhead” compression is also available for data services, whenever possible (this depends on the characteristics of the transmitted data). Finally, a custom SIP compression mechanism is employed to dramatically reduce the size of the signaling messages that are otherwise known for their immoderate consumption of bandwidth.

[1] The figure shows only the communication equipment because the scope of the paper excludes topics related to satelite control, geolocation and network management.

ALL-IP LAYERED ARCHITECTURE

The telephony industry (including cellular 3rd Generation Partnership Project (3GPP) standards) is undergoing a transition from the circuit-switched (CS) technology to a mix of CS and PS networks, with the goal of eventually utilizing PS domains only. The transition, however, is a slow process and will not complete for many years. In that regard, the MUOS system has leapfrogged the industry by implementing all of its services over the PS network. To accomplish this, the service architecture had to be designed around this new technology.

MUOS' all-IP architecture is strictly layered, meaning that IP is the converging protocol for all service nodes, and that features are designed by aggregating various layers which, when combined, make the node exhibit the desired behavior.

The layering principle, especially with regard to what's "below IP" and "above IP," is seen particularly well in the design of the OTA mechanism, which comprises a number of various protocols (RRC, RLC, MAC, etc.) but, in essence, provides the physical and link **(**L1/L2) layers such that IP packets can be sent from the UE to RAN where they are relayed via the GPRS Tunneling Protocol (GTP Tunnel) to the UE's serving Gateway GPRS Serving Node (GGSN). The aggregate behavior of these protocols, that accounts for the characteristics of the link between the UE and GS, can be configured for a variety of applications (in the baseline there are two basic types of configuration: one optimized for streaming services and one optimized for bursty data transmission).

The OTA interface is the most complex part of the system that implements protocol modifications, includes support for spectrum adaptation, dynamic application of Dovetail Interleaver (DTI) in order to increase capacity in times of congestion, a SuperRAB in the User-to-Base (U2B) direction, aggressive header compression (HC) with support for DTX, progressive management of OTA resources, Power Saving Mode (which preserves the IP address while the UE is idle), support for Multi-RABs (multiple, concurrent and independent IP services to a single UE), handling of mobility scenarios, custom rest state (CELL_FACH) and channel switching for dynamic adjustment of bandwidth allocation. Although designing this 'bridge over the air' was a tough task in itself, what compounded the difficulty was the desire to utilize an unmodified COTS Core Network (CN). This cost-reduction measure resulted in a reduced solution space for a number of problems and consequently challenged the talents of the MUOS engineering team. Needless to say, this interface is the 'crown-jewel' of the MUOS technology and the foundation for any services offered at the terminal. Figure 3 shows a simplified view of the virtual L1/L2 mechanism underlying the IP connectivity between the UE and GGSN, along with the identified components which were modified and which were not.

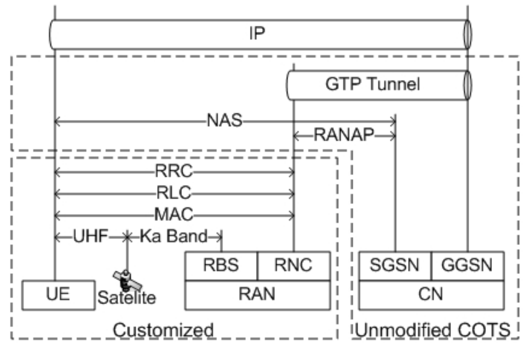

Figure 3 Design of the OTA Link

The layering of the architectural components is also visible in the design of the terminal, which is separated into logical layers, each providing a different set of functionality. The CAI Waveform for MUOS was designed to be loaded onto a JTRS terminal and, therefore, it is decoupled from the platform on which it runs. This separation of the UE from the MCT allows great flexibility in that different terminals may be designed for different purposes, such as different form factors (handheld vs. man pack), different uses (voice-only vs. data-only), or different end-user interfaces (auto-answer vs. non-auto-answer mode), etc. New end-user services may be added at the application layer, at the MCT layer, inside MUOS (UE and/or GS), or through a combination of any of these three approaches. Understanding of the users' needs is key in assessing the layer modifications that may be required to implement the new features.

The Modem Relay technique used in MUOS is likewise an example of the power of layering. What from the UE's perspective looks like a G.729ab voice transmission, may in fact constitute transfer of data via the v.150.1 protocol, which is implemented in the MCT and the DSN Gateway.

Finally, transmissions between the sites across the DISN Core occur via tunnels which render encrypted MUOS data transparent to the rest of the DISN Core and, at the same time, hide from MUOS the underlying MPLS mechanism utilized by the DISA's network.

Having the mechanism of transporting IP packets between the UE and its GGSN, additional system features are provided to meet communication needs of the end user. While the basic connectionless IP service works much like a "dial-up" service in that the user has control over the configuration of its leg over the air, E2E services are built by inserting a call control logic between this raw IP service and the Application Programming Interface (API) to the MCT. Call control is provided with the help of SIP, which has become a *de facto* standard signaling protocol for IP-telephony in North America. More specifically, the call control mechanism is based on the IMS Architecture, which is part of the 3GPP standards and which is becoming the foundation of IP-based service platforms of major telecommunication operators. The main components of the architecture are the Call Session Control Function (CSCF), Home Subscriber Server (HSS) and the Application Server (AS). The CSCF and AS were defined to work in tandem such that the AS implements custom behavior desired by the operator (for example, in MUOS AS prevents certain users from accessing DSN) while the CSCF performs the generic functions which are common across the majority of IP-telephony systems (e.g. location services). MUOS takes advantage of this paradigm, which allows use of an unmodified COTS CSCF with a customized MUOS Application Server (MAS). Figure 4 shows an abstracted view which presents how each terminal is provided with IP services, that connect it to the MUOS IP Backbone. The SIP protocol operates above the IP layer and provides the location services, as well as coordination of configuring OTA links at the two sides involved in a single call. The "SIP Server" denotes the CSCF, MAS and associated equipment of the IMS Core. The DSN GW also "talks" SIP, as well as it translates and transcodes, respectively, between SIP and SS7 and between G.729/RTP media stream and G.711/PCM.

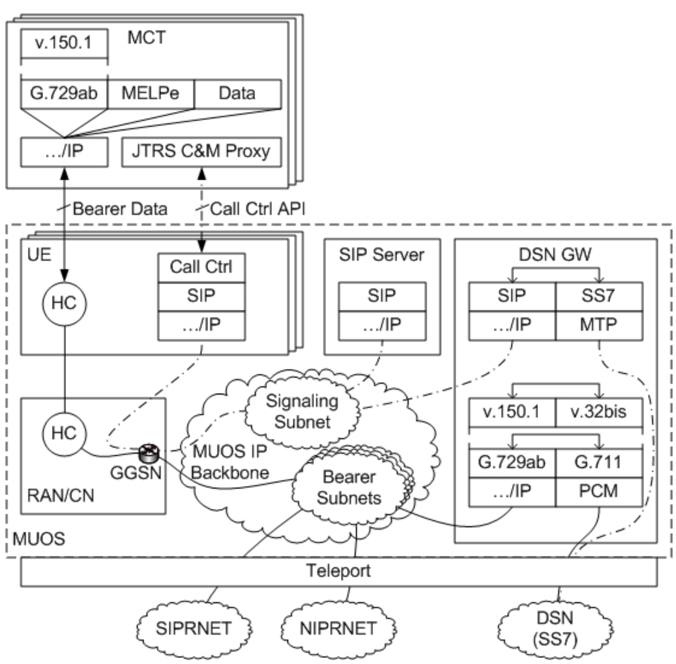

Figure 4 IP-Based Services

Note that connections to NIPRNET and SIPRNET are purely IP without any use of SIP, which is used only for E2E type of calls between two UEs and between UEs and DSN GWs acting on behalf of DSN users. The subnetting is designed for security reasons to separate various types of traffic (signaling, data with or without access to NIPRNET, etc.).

Finally, each of the two IMS Cores maintains registration information about all the users who are logged-in. Load Balancers (not shown in Figure 4) front the IMS Core such that call requests are always forwarded to the closest working IMS Core. This mechanism increases the availability posture of the MUOS system and reduces the call setup time.

In summary, the IP protocol forms the foundation upon which the services are provided by stacking up additional layers of protocols. MUOS IP Backbone, which internally utilizes DISN lines to interconnect GGSNs at the four MUOS sites, allows IP-enabled nodes to communicate with one another irrespective of their location, underlying link layers, or the applications running above IP. The simplicity and power of the IP paradigm is precisely this independence of the mechanisms below IP and 'applications' above IP.

MUOS is aligned with commercial trends in terms of the overall architecture and the specific components and protocols. As such, it is a great platform for introduction of new services – some of which have already become the staple of cellular operators – that could take military communications to an entirely new level of usability.

PLATFORM FOR FUTURE SERVICES

MUOS is oftentimes described as a replacement for the legacy UFO system. This is true in that both systems are satelite-based and both offer global mobility of the users. The similarities, however, end there. The difference between the UFO system and MUOS, in simple terms, can be characterized as that between "a CB radio" and "an IP phone system." The UFO system offers only a bent pipe for transmitting at pre-assigned frequency ranges. Therefore, management of the access method is provided by a combination of planning, discipline and the communication equipment relying on custom features developed by various user organizations. As such, UFO does not offer any services. MUOS, on the other hand, is a complete system: it provides location services, it actively manages radio resources and facilitates the setup of radio bearers to match the specific needs of the application type. MUOS even supports exchange of MCT-MCT passthrough information (for example, signaling to setup crypto devices for the session).

In the baseline release of MUOS, the set of features is limited to those necessary to meet service requirements with emphasis placed on the adaptation of the air interface. This phase of the program may be thought of as building a platform for services to be added in the future, as needed by the customer. It should be noted that even though some of the advanced features may not be useful today, the opportunities offered by MUOS may drive development of new end-user applications, which in turn may necessitate upgrades of the MUOS system.

The following is a list of enhancements that may be added to MUOS (note that the analysis of these potential features, including any feasibility studies, has not yet been conducted; nevertheless, no major obstacles are anticipated with respect to most, if not all, of these future services):

- Utilization of existing features supported by SIP, such as subject line, different types of alerts, information about the caller (e.g. vCard, unit info), information about the call to which this one is responding ("Re:"), etc.
- Standard supplementary telephony features such as Call Hold, Call Park, Call Transfer, Call Forward, Call Forward – No Reply, etc.
- Voice mail (with different greetings when the user doesn't answer, when the terminal is not connected, or when the call cannot be established due to congestion, etc.)
- Conferencing (*ad-hoc* conferencing as a useful way to create 'group communication' without prior planning)
- Presence feature showing real-time status of select users
- Agency partitioning to dedicate MUOS resources on a per agency basis
- Establishment of recurring sessions without explicit signaling (end-to-end signaling is used only once to setup, for example, 10 second feeds from sensors to a logger, recurring every 10 minutes for the next 20 days)
- Directory services facilitating dialing by name, rank, function, etc. (e.g. in a unit with leaders rotating every 8 hours, a soldier could simply dial 'leader' and be connected with the one currently on duty)
- Sequential or parallel search of users with multiple phone numbers, including follow-me service out to DSN
- Ringback when free (or a flavor thereof) allowing notification when the busy number is ready to accept a call (e.g. when the user is done with the previous service or when the congestion conditions do not exist any more)
- Group Ringing and Call Pickup (forking of call announcements to a group of users and establishment of the call with the first member to answer)
- Enhanced dialing features (e.g. four phone numbers assigned to a single terminal, where one is active at all times, one is active only on Sundays, one is active only in the afternoons, and one is active only when the congestion conditions are not grave in the beam that is serving the UE – each of these could then be published in different directories)
- Call blocking based on the caller or type of call (e.g. a voice-only terminal would not receive any invitations to data sessions)
- Enhanced Management of PDP Contexts (support for secondary PDP Contexts with varying QoS or more than two primary PDP Contexts)
- Enhanced Power Saving Mode (preservation of all PDP Contexts, not only the signaling one)
- Call logging (e.g. for training purposes)
- Asymmetrical calls (client vs. server) or one-way calls (sensors/telemetry vs. control).
- Additional Multi-RABs (i.e. beyond the configurations offered in the baseline) that would extend the set of possible configurations of the radio links in support of multiple concurrent services
- Additional APNs for fine-tuning of access rights management
- Vocoder negotiation allowing selection of the best-quality or the most-efficient vocoder, or any common vocoder (for example, during the time of congestion allow only the most-efficient vocoder)
- Additional audio (e.g. wideband codecs for superior voice quality or MELPe 600bps for conservation of capacity) and video codecs in the MCT with support for appropriately configured OTA links and aggressive header compression in the UE
- Defaulting to the best-quality vocoder whenever possible and, during congestion conditions, downgrading calls (G.729=>MELPe) instead of preempting them
- Support for banks of transcoders on the ground in order to stitch individual call legs with different vocoders
- OTA header compression for the Red-Side headers
- Distinguishing between Call Precedence and Resource Precedence, where the former is set by the end-user while the latter is set by the system (this feature could become attractive especially in complex scenarios involving heterogeneous systems, that can make this distinction, and multihomed terminals with the purpose of preserving scarce resources, balancing of

loads across various domains, routing of important call through more secure or reliable networks, etc.)

- Enhanced compression of SIP messages
- Mobility at the IP layer (IPv6)

Finally, there is a handful of opportunities that warrant a more detailed coverage.

Interconnection with other SIP systems or user agents – One of the most obvious and desirable features that could be added in the future, is an upgrade to make MUOS communicate directly with other SIP-based systems and/or clients, within Information Assurance (IA) limitations. This service could be implemented in MUOS by addition of a fronting mechanism in the IMS Core for the purposes of performing SIP authentication of non-MUOS servers/clients. There might also be a desire to implement a mechanism ensuring that appropriate level of QoS (especially, the bandwidth) is available across the whole path of transmission (i.e. including the section outside of MUOS).

MUOS Services via Non-MUOS IP Service – MUOS services could be made available at any IP-enabled node as one could envision a use of a MUOS-compatible IP Terminal on xIPRNET such that end-to-end services (not just the raw IP connection) could be established between the UE and a xIPRNET-based node (presently, a UE can communicate with xIPRNET nodes via connectionless – i.e. "dial-up" like – service only). Also, a new type of a terminal could be devised with support for multihoming such that MUOS services were to consume the OTA resources only when not in the proximity to other wireless or wireline networks (similar in principle to the concept behind the Unlicensed Mobile Access (UMA) technology in cellular).

Sharing of MCTs – Presently, there is a large base of legacy equipment which was designed around minimal bandwidth requirements (e.g. 75 bps teletype machines). In MUOS, the nominal data service is 9.6 kbps (which is 128 times more than the 75 bps teletype machine expects). An analysis of the options could be conducted to determine the best way of accommodating those small-bandwidth applications. For example, a mechanism of 'timeslotting' over a single service leg could allow multiple low-bandwidth applications to share a single connection. Although this could be done at the application layer, providing this type of support inside MUOS would allow for each of the terminal extension to dial other terminal's 'extensions' (a single MCT could register multiple directory numbers). This way, a connection could be established between multiple nodes "behind" the MCT with support for header compression. Careful accommodation of these legacy applications would preclude unnecessary expansion of the applications to consume all the bandwidth that is available. Moreover, extending the reach of MUOS services past the MCT itself could greatly reduce the number of needed terminals.

Subscription – a very interesting class of applications can be developed around SIP subscription services. Multiple applications could then be developed in association with subscription based on time, rank, danger level, or even geography (for terminals equipped with Global Positioning System (GPS) capabilities). In short, MUOS could automatically manage sources of messages (e.g. weather alerts, danger alerts, biological alerts etc.) and distribute them to select subscribers (based on the criteria selected during subscription). The subscription services could be controlled (e.g. sourcing of alerts, redefinition of the zones, etc.) either from a MUOS terminal or a node on the xIPRNET.

Dispatch Console for Situational Awareness – Equipping MCTs with GPS could bring situational awareness to a whole new level. For example, each terminal could upload (to a server) its GPS location periodically (the interval depending on the speed of the terminal). This information about the location of each terminal (and perhaps other telemetry data) could be collected in a 'dispatch console' (accessible either from the xIPRNET or any MUOS terminal), where it could be visually overlaid over the map of the area along with other information about the location (e.g. the weather). Subscription zones could be changed from the console, buddy lists could be displayed on a side with information about each of the subscriber in the area, and some devices could be controlled directly from the console (for example, a picture could be taken by sending a signal to a camera equipped with an MCT and then uploaded directly to an *ad-hoc* created group of users who appear to be approaching the intersection where the camera is located). The usability of such application integrated with MUOS could appeal to a wide range of users.

Interoperability with Iridium – Given the widespread adoption of Iridium services within the Department of Defense (DoD), interoperability with the next generation Iridium handsets (i.e. beyond the services offered via DSN) seems worth considering. Presently, a phone connection between MUOS and Iridium terminals requires a two-stage dialing and double transcoding (since Iridium uses an ultra-small bandwidth codec, the quality of voice service between a MUOS and Iridium terminals can be expected to be poor). Iridium also offers a "Direct Data" service with data rates of up to 10 kbps, where the modem is run over the air to the Iridium gateway, which terminates the modem and serves as an Internet Service Provider (ISP) by providing a direct connection to an IP cloud (the Internet). It would be worthwhile to consider any options with regard to bridging MUOS and Iridium services – both voice and data – such that mobile users with either terminal could communicate beyond what is possible via a DSN line.

Still other possibilities exist, such as increasing the capacity through addition of new satelites, creation of multi-UE terminals for handling of data rates in excess of 384 kbps, etc. The intent of this paper is not to present an exhaustive list of possible upgrades but to convey the unusual flexibility of the MUOS architecture. All-in-all, there is a plethora of opportunities for MUOS to evolve such that it enables deployment and use of new communication services (these will have to take into account the needs of the end user (soldier, operator, application, etc.), and the appropriate layer at which they should be implemented).

CONCLUSION

MUOS is an exemplary model of how systems should be designed such that they not only meet the current requirements but allow for natural evolution of the system which is expected to accommodate proliferation of ever new and more powerful communication applications. MUOS exhibits the chief characteristics of a robust platform upon which to build. Instead of a radio where the user is scanning the frequencies to latch on to a channel, MUOS offers a handful of very useful services provided in a cleanly engineered, integrated, and standardized package. With some upgrades, MUOS can offer to the end-user a Blackberry™-type of a terminal with support of various voice, video, messaging, and web applications. Such options invite probing of opportunities in terms of enabling technologies so that instead of the people working around the equipment, the system works for people. In short, MUOS architecture encourages pursuit of opportunities that will appeal to its user community in the years to come.