M2M: From Mobile to Embedded Internet

Geng Wu, Shilpa Talwar, Kerstin Johnsson, Nageen Himayat, and Kevin D. Johnson, Intel

ABSTRACT

Is M2M hype or the future of our information society? What does it take to turn the M2M vision into reality? In this article we discuss the business motivations and technology challenges for machine-to-machine communications. We highlight key M2M application requirements and major technology gaps. We analyze the future directions of air interface technology improvements and network architectures evolution to enable the mass deployment of M2M services. In particular, we consider the salient features of M2M traffic that may not be supported efficiently by present standards, and provide an overview of potential enhancements. Finally, we discuss standards development for M2M.

INTRODUCTION

Machine-to-machine (M2M) communications in the context of the mobile Internet has been a subject of intense discussions over the past two years. Some see it as the next technology revolution after the computer and Internet. Some consider it simply hype. Others are cautious with a wait-and-see attitude.

Part of the confusion has been that M2M is not something completely new. For those familiar with embedded control, M2M is a natural extension of their existing business. They fail to see the explosive growth that others are excited about. Other people also remember the high-tech bubbles in recent history, and question the practical future of M2M.

Intel recently completed an extensive study on the issues critical to the M2M industry. We exchanged views with leading equipment manufacturers, software vendors, and service providers. In this article, we share some of our learning.

We begin with our vision of the future embedded mobile Internet. Then we look at several M2M use cases that offer significant market potential. We discuss the requirements and challenges associated with mass-scale M2M networks, and describe potential system architectures and deployment options that can enable the connectivity of billions of low-cost devices. We describe the salient features of M2M traffic that may not be supported efficiently by current standards and provides an overview of potential enhancements. Finally, we summarize the progress of standards development for M2M.

THE FUTURE OF EMBEDDED INTERNET

Mobile Internet is at a turning point. In this section, we discuss what motivates the evolution and share our vision of M2M for the future embedded Internet.

THE TECHNOLOGY AND ECONOMIC MOTIVATIONS FOR M2M

The proliferation of mobile Internet provides nationwide ubiquitous coverage and mobility support. Today's advanced wireless networks are ready to deliver broadband data service at a significantly lower cost than in the past, thanks to extensive standardization [1]. These networks offer many of the features necessary to enable M2M services in the future embedded Internet.

Technology is one of the main drivers of M2M. The semiconductor industry's shrinking lithography and improved yield continue to reduce chipset cost and power consumption. Carrier WiFi, small cells, relay, and peer-to-peer communication further extend the coverage of wireless networks while dramatically reducing cost per bit transferred.

There are also profound economic motivations for the wireless industry to aggressively pursue M2M. As voice revenue continues to deteriorate, operators are under tremendous pressure to introduce new services that will fill their revenue gap. M2M, cloud computing, and application stores top the list of potential revenue-generating services.

THE VISION OF INTERNET OF THINGS

To the authors, M2M represents a future where billions to trillions of everyday objects and the surrounding environment are connected and managed through a range of devices, communication networks, and cloud-based servers.

There are three essential components to this "Internet of Things" vision.

A continuum of devices from low-cost/low-power to compute-rich/high-performance: In the M2M market, large numbers of devices are expected to be embedded, requiring extremely low price points and low power consumption. However, higher-end devices such as gateways, machine control modules, intelligent vision systems, and even consumer electronic products are also important growth segments.

Ultra scalable connectivity: This is arguably the most important component of the M2M vision. A device that is not connected cannot easily be managed or work in concert with other devices. Our most critical challenge, therefore, is to enable low-cost connectivity that addresses not only the massive network scale but also the vastly diverse requirements dictated by the device continuum.

Cloud-based mass device management and services: The vision of the future is no longer one device acting alone, but many devices acting together. Thus, although distributed processing is critical to address the complexity of M2M applications, centralized decision making and management of billions of devices within the cloud will become an essential value of the Internet of Things vision.

Although this vision is not new, it is only now gathering momentum as ubiquitous connectivity is finally becoming a reality, and Moore's Law has driven device cost and size low enough to justify "smart devices everywhere."

The Essential Elements of M2M Solutions

Third-/fourth-generation (3G/4G) wireless technologies will play a central role in the future of M2M. Its high data rate enables high value services. In markets where 2G is reaching the end of its life cycle, 3G/4G is the only option.

As the M2M market expands, operators will encounter significant technical challenges. *Security* will be of paramount concern. A major security breach in a network connecting billions of devices is unthinkable. It is expected that advanced solutions including "security-on-chip" will be developed.

As the M2M market expands, "zero-touch" *manageability* and *information overload* will pose significant challenges to the network. We have observed in large-scale surveillance networks that the number of video feeds overwhelms human operators. Thus, there will be an urgent need for video analytics at the installed end devices.

Similarly, as M2M solutions evolve, *optimum distribution of device and cloud intelligence* will become critical. With increased intelligence at the device, *augmented sensing* will be practical for innovative value-added services.

Scaling smart device installations and supporting future technology is a crucial smart system architecture concern. We envision standardized M2M plug-n-play capability to be essential to overall acceptance of M2M technologies.

Finally, future M2M solutions will need to support *a mix of legacy and new services and devices*. We expect M2M gateway/aggregation points to play a key role in bringing the installed short-range sensors online and providing interworking with different wireless technologies. These gateway/aggregation points can also become a platform for value-added services to enable an explosive growth of short-range smart sensors, fully managing the scale

Toward the Future of Embedded Internet

Future M2M ecosystems will be complex and span many industries, including telecom and electronics. Unlike current M2M markets, which are highly segmented and often rely on proprietary solutions, future M2M markets will need to be based on industry standards to achieve explosive growth. This standards process will be much broader than writing a specification, as it involves not only interfaces, but also platforms and services. The M2M industry needs to leverage existing vertical market solutions, design platforms that horizontalize the market, and avoid the narrow solutions that come from chasing "killer apps." The industry also needs to ramp up efforts to develop critical technologies for an optimized air interface, device manageability, network architecture, and security in order to enable future mass deployment of embedded devices.

Future M2M ecosystems will be complex and span many industries, including telecom and electronics. Unlike current M2M markets, which are highly segmented and often rely on proprietary solutions, future M2M markets will need to be based on industry standards to achieve explosive growth.

Representative M2M Usage Models

One can envision creating an immensely rich set of applications when thousands of objects surrounding us are connected. Some examples are smart homes, where intelligent appliances autonomously minimize energy use and cost; "connected cars" that react in real time to prevent accidents; and body area networks that track vital signs and trigger emergency response when life is at risk. In this section, we study three M2M applications and provide a brief description of other applications to demonstrate the broad market potential of M2M.

Utilities (Smart Grid)

Smart grid integrates communication capabilities with utility generation (e.g., electric power, gas, water) and delivery infrastructure to automate monitoring and control. Sgnificant savings in resource consumption is also possible when utility supply is dynamically matched with demand. Key smart grid applications are smart metering, distribution network automation, demand response, equipment diagnostics, as well as wide area monitoring and control. An example of smart grid network architecture is described in [2].

An M2M-enabled smart meter collects utility usage information from home appliances via short-range radio or a home area network and sends the information to the M2M server by communicating directly through the 3G/4G network or via an M2M aggregation device and then to the 3G/4G network. The M2M aggregation device collects information from many smart meters in the area and sends the aggregated information to the M2M application server. The home area network interface to the smart-meter can be based on several short-range wireless technologies such

Security and public safety	Surveillance systems, control of physical access (e.g., buildings), environmental monitoring (e.g., for natural disasters), backup for landlines
Smart grid	Electricity, gas, water, heating, grid control, industrial metering, demand response
Tracking and tracing	Order management, asset tracking, human monitoring
Vehicular telematics	Fleet management, car/driver security, enhanced navigation, traffic info, tolls, pay as you drive, remote vehicle diagnostics
Payment	Point of sale, ATM, vending machines, gaming machines
Healthcare	Monitoring vital signs, supporting the aged or handicapped, web access telemedicine points, remote diagnostics
Remote maintenance and control	Industrial automation, sensors, lighting, pumps, vending machine control
Consumer devices	Digital photo frame, digital camera, ebook, home management hubs.

Table 1. *Example M2M use cases with wireless WAN coverage and mobility support.*

as WiFi, Home-Plug, ZigBee, or even 3G/4G if used with femto-like capability.

This hierarchical network architecture is actually typical in many M2M applications.

Vehicular Telematics

Most vehicular M2M applications can be categorized into one of the following: safety and security, information and navigation, diagnostics, or entertainment.

An example of a safety and security service is Automatic Crash Notification. This service utilizes various crash sensors on the vehicle to report the location and extent of damage to the vehicle in the event of a crash. It also initiates a voice call to facilitate reporting of the crash to Emergency Services.

Information and navigation services provide access for the vehicle occupant to a variety of location sensitive information and content, similar to what we have access to today via Google map search.

Diagnostic services enable the occupant and/or vehicle maintenance/repair centers to collect data from a multitude of sensors located throughout the vehicle in order make maintenance and/or repair recommendations.

Many vehicular M2M applications require a combination of short-range low-power low-throughput wireless access (e.g., Zigbee) for sensing processes and local connectivity within the vehicle, and long-range low-latency high-throughput wireless access such as 3G/4G for reporting functions and Internet access for media content services.

Healthcare (M-Health)

M-Health is a nascent market aimed at improving the quality of patient care and reducing healthcare costs. Services include telemedicine to improve patient care by virtue of more accurate and faster reporting of changes in the patient's physical condition, automated connectivity of medical devices to the hospital network and remote management of these devices, and electronic representation and exchange of medical data between hospitals and medical groups such as laboratories or pharmacies to lower transaction costs.

To date, the healthcare industry has spent significant resources on telemedicine. One of the primary services is remote patient monitoring and care, wherein a patient wears bio-sensors to record health and fitness indicators such as blood pressure, body temperature, heart rate, and weight. These sensors forward their collected data to an M2M device (e.g., a patient's cell phone) that acts as an information aggregator and forwards the data to the M2M application server in the cloud. The M2M server responds to the collected data by sending alerts and appropriate medical records to medical providers. In emergency situations, an M2M device can directly provide the medical status of a patient en route to the hospital (e.g., in the ambulance), allowing physicians to prepare for treatment in advance of the patient's arrival. This is a scenario where reliable high-speed connectivity such as 4G cellular is required.

Other Applications and General Usage Models

Many other applications can benefit from wireless WAN capability, such as new generations of consumer devices that incorporate personal navigation, e-readers, remote digital frame functions, public surveillance systems, environment monitoring, remote maintenance and control, e-payment systems, and tracking/tracing.

Table 1 shows example M2M use cases that require or benefit from wireless WAN coverage. The table is not exhaustive, but it gives an idea of the range of M2M applications that have been discussed in the Third Generation Partnership Project (3GPP) and IEEE standards [2–4].

M2M Connectivity and Network Architectures

As seen from the wide range of usages, the industry needs a cost-effective, scalable M2M solution that will support a variety of applica-

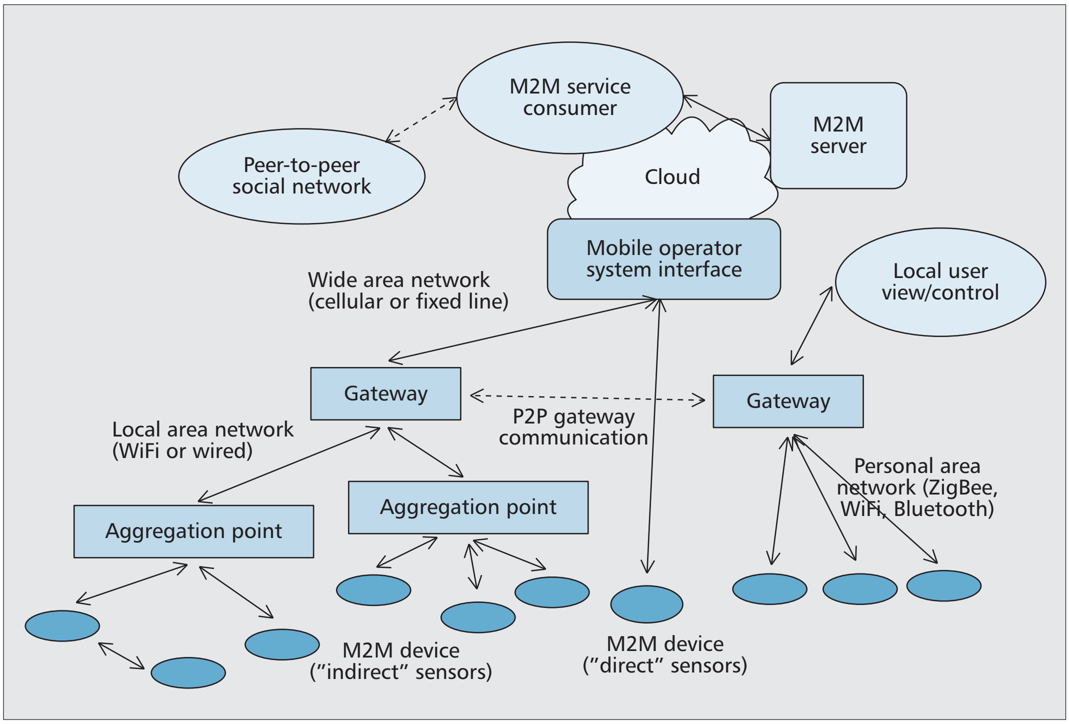

Figure 1. *A high-level M2M system architecture.*

When the numbers of devices explodes, network access deteriorates rapidly. The expected increase in M2M devices also poses a network capacity concern. Hierarchical network architectures offer an effective solution.

tions and M2M devices. When the numbers of devices explodes, network access deteriorates rapidly. The expected increase in M2M devices also poses a network capacity concern. Hierarchical network architectures offer an effective solution.

Hierarchical Network Architectures for Scalable Connectivity

Multiple connectivity options are available today to connect M2M devices to a server and each other. However, to do this on the scale of billions to trillions, particularly when many devices are limited in range due to cost/size/power constraints, hierarchical deployments that provide reliable, efficient interworking between multiple communication protocols (PAN/LAN/WAN) will be needed. Figure 1 captures a high-level view of a hierarchical M2M system architecture.

The M2M device can connect to the M2M server directly through a WAN connection (e.g., cellular 3G/4G) or an M2M gateway (aggregation point). The gateway is a smart M2M device that collects and processes data from simpler M2M devices and manages their operation. Typically, connecting through a gateway is preferred when devices are sensitive to cost, power, or location. There are several lower-cost radio protocols, such as IEEE 802.11, IEEE 802.15, and power line communications, through which these devices can communicate.

Many M2M applications will require connectivity between end devices. Peer-to-peer (P2P) connectivity can be supported in this architecture at various levels of hierarchy depending on latency requirements and the type of information exchanged.

Hierarchical Networks for High Capacity

Recent studies have pointed to the explosive growth in mobile data demand driven by popular devices and video services [5, 6]. M2M will further add to the pressure.

The most cost-effective solution to this explosive increase in traffic demand lies, once again, in hierarchical network architectures, comprising both multiple tiers and multiple radios.

Multitier: In the multitier hierarchy shown in Fig. 2, large cells provide ubiquitous coverage to M2M devices and support high mobility; while smaller network elements such as relays and pico/femto access points (APs) bring connectivity closer to the devices, improving link reliability and increasing system capacity [7]. The lower cost of smaller APs makes them an attractive method of adding capacity, since it is done at a lower cost per bit [7, 8]. Devices (mobile stations) can also serve as a tier in the network hierarchy by creating P2P nanocells.

Multiradio: Figure 2 also shows multiple access networks being integrated and managed as part of a single hierarchical network. Here, the additional spectrum and connectivity available across different networks (e.g., WiFi and cellular) may be exploited synergistically to further improve system capacity and device quality of service. The cost associated with this additional capacity may be very low since the alternate spectrum could essentially be free (e.g., unlicensed spectrum). New network devices, such as the integrated femtocell access point (AP), with licensed and unlicensed capabilities, can implement tighter coupling across these radio technologies, efficiently utilizing the available spectrum.

M2M systems need to be able to detect unusual events (such as changed device location, device damage) and support appropriate levels of authentication for M2M devices and gateways. Enhanced monitoring and security may require changes to the network entry/re-entry procedure.

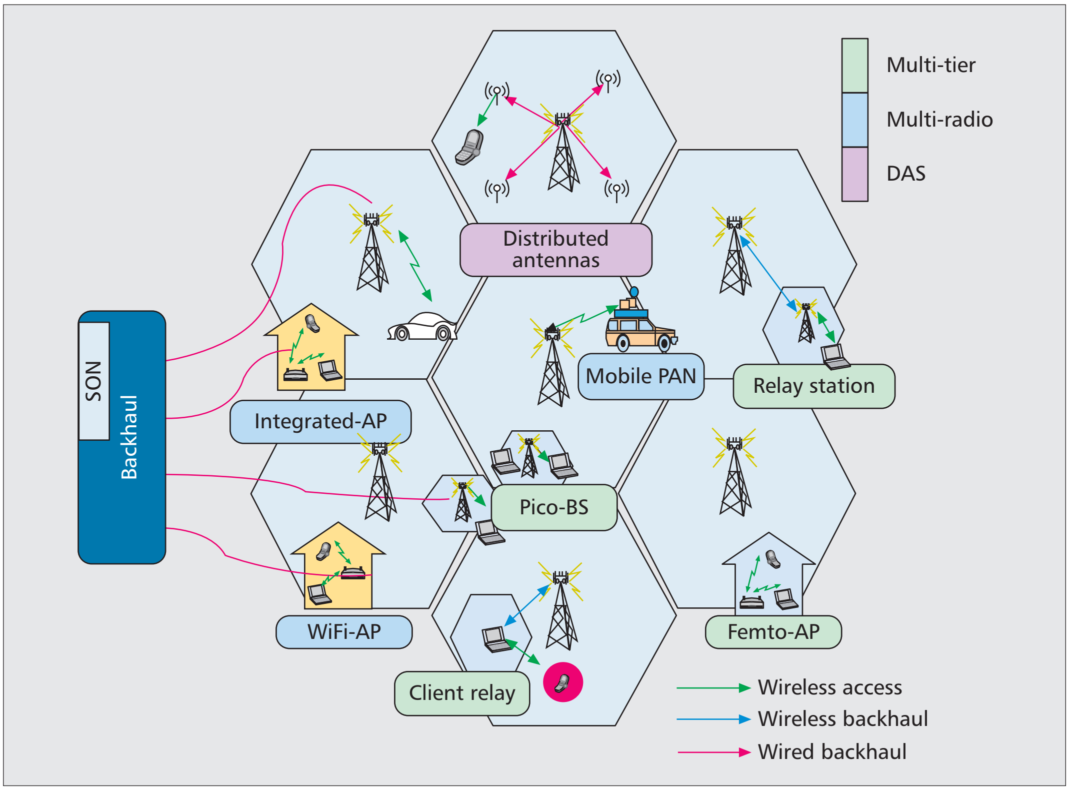

Figure 2. *Hierarchical system architecture [6].*

KEY FEATURES FOR AIR INTERFACE OPTIMIZATIONS

In addition to the architectural innovations discussed earlier, further optimizations of the air interface are needed to efficiently support new M2M traffic characteristics. In this section we discuss features that are unique to M2M, but common to one or more M2M use cases. Table 2 summarizes several key M2M features along with their associated applications and potential impacts on air interface standards [3, 4].

MASS DEVICE TRANSMISSION

This feature deals with the handling of simultaneous or near simultaneous transmission attempts to the access network's base station from an extremely large number of M2M devices. This feature may be required for many use cases such as secured access and surveillance, public safety, healthcare, and metering.

Support for these features may require enhancements to the network entry/re-entry and bandwidth request protocols, link adaptation, (hybrid) automatic repeat request (HARQ/ARQ), and/or the frame structure.

HIGH RELIABILITY

High reliability implies that connectivity and reliable transmission are guaranteed regardless of operating environment (e.g., mobility, channel quality). This feature is required in emergency situations or scenarios where privacy is extremely important (e.g., healthcare, remote payment).

Improved reliability may require changes to the link adaptation protocol or modulation/coding schemes. Other solutions may involve improved interference mitigation, device collaboration, or redundant path establishment.

ENHANCED ACCESS PRIORITY

Priority access is necessary in order to communicate "alarms" in a variety of use cases.

Enhanced access priority may require changes to the bandwidth request, network entry, or ARQ/HARQ protocols. Changes to the frame structure may also be required.

EXTREMELY LOW POWER CONSUMPTION

This feature is required by devices that have no or limited access to power sources, wake only on demand, experience infrequent human or system interaction, or belong to a large network of devices that cumulatively consume a lot of power.

Support for this feature may require updates to control signaling, idle and sleep mode, link adaptation, and uplink (UL) power control. Device collaboration will also reduce power consumption.

SMALL BURST TRANSMISSION

Support for this feature may require changes to burst management, the SMS transmission mechanism, bandwidth request/allocation protocols, channel coding, and/or frame structure. A smaller resource unit may be necessary to transmit an extremely small downlink (DL)/UL burst size.

LOW/NO MOBILITY

Many M2M use cases involve stationary or low mobility devices (e.g., payment, metering, and retail). For these cases, the system should provide simplified or optimized mobility management in order to reduce power consumption and signaling overhead.

Changes to the signaling related to handover preparation and execution may be required to take advantage of low/no mobility. This feature may also impact idle mode.

Monitoring and Security

The nature of M2M deployments makes the system vulnerable to attacks on hardware and software/firmware, compromise of credentials and configuration, and network attacks (e.g., hacking and denial of service).

M2M systems need to be able to detect unusual events (such as changed device location, device damage) and support appropriate levels of authentication for M2M devices and gateways. Enhanced monitoring and security may require changes to the network entry/re-entry procedure.

Addressing Extremely Large Number of Devices

Addressing extremely large numbers of devices may require extending the addressing space or updating the addressing scheme.

Group Control

Group control implies that the system supports group addressing and handling of M2M devices.

Enabling group control of mass devices based on predefined criteria (location, function, etc.) may require changes to group ID allocation, control signaling, paging, sleep-mode initiation, multicast operation, and bandwidth request/allocation. Changes to network entry/re-entry and service flow and connection management protocols may also be necessary.

Time-Controlled Traffic

Time-controlled traffic is transmitted and received at periods of time that are defined well in advance. This type of traffic enables power-saving reductions in the bandwidth request, network entry, idle/sleep mode protocols. The system may release the data connection outside of the access period.

Time-Tolerant Traffic

Time-tolerant traffic can support significant delays in data transmission and reception. This implies that the system can give lower access priority to or defer data transmission of time-tolerant traffic. This feature enables simplifications to the bandwidth request and ARQ/HARQ protocols.

One-Way Data Traffic

When data traffic is "one way," it is only control signaling that is transmitted in the opposite direction. Digital signage and consumer devices are use cases where data may be device-terminated only.

One-way traffic may require changes to the network entry and addressing protocols, and it may enable simplifications to the bandwidth request/allocation protocol. In addition, the receiving procedure of the DL control channel may be simplified for one-way data traffic.

Features	Applications	Standards impacts							
		Sleep and idle mode	Mobility management	Link adaptation	BW request & Allocation	HARQ and ARQ	Frame structure	Network entry	Cooperation
Mass device transmission	Security Metering Tracking		√	√	√	√	√	√	√
High reliability	Health Security	√	√	√	√	√	√		√
Access priority	Health Remote maint & control	√	√		√	√	√	√	
Very low power	Tracking Remote maint & control	√	√	√	√	√	√	√	√
Small data burst	Metering Remote maint & control	√				√			
Low/no mobility	Metering	√	√	√	√			√	√
Monitoring and security	Vehicular Payment	√					√	√	

Table 2. *Summary of air interface optimizations.*

Extremely Low Latency

Extremely low latency requires that both network access latency and data transmission latency be reduced. This feature is required in many emergency situations (e.g., healthcare).

Changes to the bandwidth request and network entry/re-entry protocols may be required to support extremely low latency. This feature may also require changes to the frame structure, ARQ/HARQ, and control signaling.

Infrequent Traffic

Infrequent traffic is common in many M2M use cases. This feature may enable sleep/idle mode improvements that save power and channel resources.

As M2M markets continue to develop, additional features may be added to this list. We expect that the technologies for the embedded Internet will continue to evolve.

Standards Development for M2M

Many standards are moving quickly to support the architectural and air interface changes discussed earlier. Since the requirements and applications are still evolving, most standards bodies are taking a phased approach, where basic M2M features are being standardized and enabled quickly, with additional optimizations expected in later phases as the market grows.

SDO	M2M development
3GPP	**Release 10:** identify requirements and optimize radio and network for features such as low power, congestion and overload control, identifiers, addressing, subscription control and security. **Release 11 and beyond:** network improvements for device to device communication, M2M gateway, enhancements for M2M group and co-located M2M devices, network selection and steering, service requirements and optimizations.
ETSI	**M2M network architecture:** define functional and behavioral requirements of each network element to provide an end-to-end view.
GSMA	**GSM operation for M2M:** define a set of GSM based embedded modules that address operational issues, such as module design, radio interface, remote management, UICC provisioning and authentication, and basic elements costs. Also define use-cases in vertical markets: health, utilities, automotive, and consumer devices.
IEEE	**802.16p (WiMAX):** optimize air interface for low power, mass device transmission, small bursts, and device authentication. Future topics: M2M gateway, co-operative M2M networks, advanced M2M features **802.11 (WiFi):** update air interface to enable use of sub-GHz spectrum **802.15.4 (ZigBee):** air interface optimization for smart grid networks
WiMAX Forum	**Network system architecture specification:** define usages, deployment models with low OPEX, functional requirements based on IEEE 802.16 protocols, and performance guidelines for end-to-end M2M system.
WFA	**Smart grid task group:** promote the adoption of Wi-Fi within the smart grid through marketing initiatives, government and industry engagement, and technical/certification programs **Healthcare task group:** maintain Wi-Fi as the preferred wireless access technology and increase adoption in the Home and Hospital Healthcare market segment.
OMA	**Device manageability:** define requirements for the gateway managed object
TIA	**M2M SW architecture TR50:** develop and maintain access agnostic interface standards for monitoring and bi-directional communication of events and information between smart devices and other devices, applications or networks.
CCSA NITS	**CCSA TC10:** focus on pervasive networks, including general requirements, applications, networking, sensing and related short range RF connectivity. **NITS WGSN:** focus on sensor network interface and data format, ID and security, vertical applications including airport and smart buildings.

Table 3. *Status of global M2M standards development.*

For example, in the first phase, only enhancements that require firmware and software changes (e.g., medium access control [MAC] modifications) are enabled. In later phases, more extensive modifications to the PHY and MAC are expected, which will accommodate advanced requirements such as those for the M2M gateway, which serves as a bridge between multiple protocols.

M2M is dependent on many technologies across multiple industries. Consequently, the required scope of standardization is significantly greater than that of any traditional standards development. As shown in Table 3, the scope of various standards organizations active in M2M ranges from air interface and network architecture (e.g., 3GPP, IEEE, European Telecommunications Standards Institute [ETSI]), to selected vertical applications (e.g., Telecommunications Industry Association [TIA], Wi-Fi Alliance [WFA], GSM Association [GSMA]), to certain critical technical areas (e.g., Open Mobile Alliance [OMA] on device manageability). Collaboration among standards organizations across different industries is therefore essential. Fortunately, the M2M community is starting to recognize this need, and joint efforts and collaborations among standards bodies are increasing.

In addition to developing open interfaces and standard system architectures, M2M ecosystems also need to establish a set of common software and hardware platforms to substantially reduce development costs and improve time to market. Most of the existing proprietary vertical M2M solutions have difficulty scaling. Horizontal developments in the M2M industry are essential for realizing the embedded Internet vision. This is of particular importance to consumer or home M2M applications where the biggest market growth is yet to come.

Conclusion

Mobile Internet is evolving towards embedded Internet. M2M presents both challenges and opportunities to the industry. Although there are significant business and economic motivations for wireless operators and equipment manufacturers to invest in future generations of M2M services, the highly fragmented markets remain a hurdle and risk the forecasted growth of M2M markets. Two things are needed for the embedded Internet vision to materialize: the development of new technologies that scale with the growth of M2M markets, and a broad standardization effort in system interfaces, network architecture, and implementation platforms.

Acknowledgment

The authors would like to thank Geoff Weaver and Boyd Bangerter of Intel Labs for their technology insights and valuable comments.

References

[1] Morgan Stanley, "Internet Trends," March 9, 2010
[2] IEEE C80216-10_0002r7, "Machine to Machine (M2M) Communication Study Report," IEEE802.16 Contribution, May, 2010.
[3] DRAFT-T31-127-R020-v01, "Recommendations and Requirements for WiMAX Machine to Machine (M2M), WiMAX Forum doc., Aug. 2010.
[4] 3GPP TS 22.368 v10.1.0, "Service Requirements for Machine-Type Communications (Stage 1)," Release 10, June 2010.
[5] Cisco Visual Networking Index, Cisco VNI, Oct. 2009.
[6] IEEE C80216-10_0016r1, "Future 802.16 Network: Challenges and Possibilities," Mar. 2010.
[7] S. Yeh, S. Talwar, S. Lee and H. Kim, "WiMAX Femtocells: A Perspective on Network Architecture, Capacity and Coverage," *IEEE Commun. Mag.*, Oct. 2008.
[8] Johansson *et al.*, "A Methodology for Estimating Cost and Performance of Heterogeneous Wireless Access Networks," *PIMRC '07*.

Biographies

Geng Wu (geng.wu@intel.com) is chief architect and director of Wireless Standards in the Wireless Technology Division at Intel Corporation. He has 20 years of experience in the wireless industry. Prior to Intel, he was director of Wireless Architecture and Standards at Nortel Networks, with extensive experience in 3G/4G technology developments. He obtained his B.Sc. in electrical engineering from Tianjin University, China, and his Ph.D. in telecommunications from Université Laval, Canada.

Shilpa Talwar (shilpa.talwar@intel.com) is a principal engineer in the Wireless Communications Laboratory at Intel, where she is conducting research on mobile broadband technologies. She has over 15 years of experience in wireless. Prior to Intel, she held several senior technical positions in the wireless industry. She graduated from Stanford University in 1996 with a Ph.D. in applied mathematics and an M.S. in electrical engineering. She is the author of numerous technical publications and patents.

Kerstin Johnsson (kerstin.johnsson@intel.com) is a senior research scientist in the Wireless Communications Laboratory at Intel, where she conducts research on network, MAC, and PHY optimizations that improve wireless network cost, coverage, and capacity. She graduated from Stanford with a Ph.D. in electrical engineering and has almost 10 years' experience in the wireless industry. She is the author of numerous publications and patents in the field of wireless communications.

Nageen Himayat (nageen.himayat@intel.com) is a senior research scientist with Intel Labs, where she performs research on broadband wireless systems, including heterogeneous networks, cross-layer radio resource management, MIMO-OFDM techniques, and optimizations for M2M communications. She has over 15 years of research and development experience in the telecom industry. She obtained her B.S.E.E from Rice University and her Ph.D. in electrical engineering from the University of Pennsylvania in 1989 and 1994, respectively.

Kevin D. Johnson (kevin.d.johnson@intel.com) is director of Embedded Connected Devices at Intel Corporation. He directs the team that has responsibility for accelerating embedded products and technologies in the M2M industry. He has over 20 years of product, technology, and business experience in the computing and embedded markets. He holds a B.S. in engineering from Oregon State University and a Master's in business from the University of Portland.

Two things are needed for the embedded Internet vision to materialize: the development of new technologies that scale with the growth of M2M markets, and a broad standardization effort in system interfaces, network architecture, and implementation platforms.