

Towards a Simple, Versatile, Distributed Low-Power Wireless M2M Infrastructure

Jeffrey Dungen, Traian Antonescu, Pier-Olivier Genest
reelyActive

Montréal, Québec, Canada

Email: jeff@reelyactive.com, traian@reelyactive.com, pier-olivier@reelyactive.com

Abstract—Existing wireless M2M infrastructure based on cellular and WiFi networks is often unsuitable for the growing number of simple, inexpensive, low-power connected devices. Low-power wireless technologies are examined in order to identify the considerations for a suitable low-power wireless M2M area network infrastructure. A novel design with the versatility to support multiple technologies and to easily extend coverage is presented and its implementation detailed. Results, including real-time location capability, are promising and development continues to support emerging applications and wireless technologies.

I. INTRODUCTION

Machine-to-Machine (M2M) technology enables the communication of information from one device to another. A wired or wireless network serves as the medium for information transfer. Wireless M2M commonly employs cellular or WiFi networks which represent a global infrastructure thanks to their extensive deployment. However, for wireless devices on a constrained power or cost budget, these networks are often unsuitable and, moreover, no suitable global infrastructure exists [1], [2].

With predictions of tens of billions of devices joining the Internet of Things (IoT) over the coming years, of which a significant proportion are likely to be of the simple, low-power variety, dedicated wireless networks that connect these devices with M2M gateways become essential. These are classified as M2M area networks as per the ETSI M2M specification [3] and reside in the device/gateway domain. ETSI M2M, one of many standardization efforts, is particularly relevant due to its focus on the service middleware layer [4].

This paper addresses the known need [5], [6] for heterogeneous M2M area networks which support the growing number and variety of resource-constrained wireless devices. We present our ongoing work and design process towards a simple, versatile low-power wireless M2M area network infrastructure. M2M software and interfaces are outside of the scope of this paper but are briefly treated in the discussion nonetheless.

II. LOW-POWER WIRELESS TECHNOLOGIES

It is reasonable to assume that a significant proportion of the IoT devices predicted to come online in the coming years will be simple, compact and mobile. Low-power wireless technologies provide an energy-efficient means for such devices to exchange information. For wireless devices that

TABLE I
COMPARISON OF SELECT LOW-POWER RADIO TECHNOLOGIES

Protocol	Max Packet Size	Transfer Rate	Band
BLE	27 bytes	1Mbps	2.4GHz
IEEE 802.15.4	128 bytes	250kbps (max)	Several
DASH7	256 bytes	200kbps	433MHz

depend on a limited energy source such as a coin cell battery, communication efficiency can spell the difference between a longevity measured in days or years. The list of low-power wireless technologies includes Bluetooth Low Energy (BLE), ZigBee, Z-Wave, ANT+, Nike+, DASH7 and many proprietary protocols. These technologies are often optimized for an application such as Personal Area Networks (PAN), home automation, fitness or remote controls, among others. With so many competing options, today there is no single standard for low-power wireless devices.

Our interest is in technologies that enable M2M or Internet-connectivity. In the case of low-power wireless devices, there is an opportunity to facilitate their convergence and interoperability through contextual awareness, flexible routing and optimized protocols.

Table I compares key characteristics of a selection of such technologies, including BLE, IEEE 802.15.4 (the physical layer and media access control which serves as a basis for ZigBee and 6LoWPAN, among others) and DASH7 [7]. In all cases, the transmission range of a given technology is highly variable based on the environment, modulation frequency, antenna types, etc. In our experience, it is reasonable to consider the range of the given technologies to be on the order of tens of metres in real-world environments.

From Table I it is clear that technologies vary greatly in the use of frequency bands, however their packet sizes are all on the order of tens of bytes and transfer rates are on the order of hundreds of kbps. This generalization can also be extended to many of the proprietary protocols which are governed by the same radio regulations, and, obviously, the same laws of physics.

Figure 1 illustrates the how the described low-power wireless technologies compare in terms of range and data transfer rate with technologies such as WiFi and cellular. The latter support high data transfer rates at the expense of increased device complexity and power consumption. Ultra Narrow




Fig. 1. Comparison of Wireless M2M Technologies




Fig. 2. Multiple nodes for multiple technologies

Band (UNB) is an exceptional low-power wireless technology which can achieve much higher range at the expense of reduced transfer rates on the order of 1kbps, for example [8]. This limits UNB technology to applications requiring minimal, sparse data transfer.

III. CONSIDERATIONS FOR A STANDARD LOW-POWER WIRELESS M2M INFRASTRUCTURE

While the idea of designing a single infrastructure for all low-power wireless devices may seem an ambitious challenge given the landscape of competing technologies, there is promise in the fact that these nonetheless share several common characteristics, as identified in the previous section. Based on these findings we are able to enumerate the considerations for such an infrastructure, which will consist of repeating elements. For the purposes of this discussion, we will use the term ‘node’ to represent a single network-connected element which implements a given wireless technology.

A. Support for multiple technologies

In the absence of a single, standard technology, a low-power wireless infrastructure must support at least the most popular technologies for a given context. Figure 2 illustrates the infrastructure required to support multiple technologies at a given location. Each of the three nodes in the example (labeled as N) implements a distinct technology (indicated by shading and lettered subscript). Devices (labeled as D) in range of a compatible node would be able to connect to the node of the corresponding technology. The infrastructure must serve as a technology-agnostic communication funnel.

B. Range and coverage

In order to provide seamless wireless coverage of an area, nodes must be distributed such that their ranges overlap




Fig. 3. Multiple nodes for extended coverage

without dead spots. Since the technologies of interest share a range on the order of tens of metres, they too will require a spacing distance of the same order of magnitude. Figure 3 illustrates the infrastructure required to extend the wireless coverage of a given technology (in this case, technology ‘A’).

C. Power

Each node must operate its radio continuously in order to detect and connect with devices in range. This requires an uninterrupted source of power for perpetual operation.

D. Network connectivity and throughput

Each node will require network connectivity to transport packets to the gateway. For resource-constrained wireless devices, communication favours an uplink bias and typically occurs in bursts to minimize energy consumption. The infrastructure must provision sufficient throughput to simultaneously support a potentially large quantity and variety of wireless devices.

E. Mobility

The architecture must support device mobility, specifically the case where the device may be in communication range of multiple nodes simultaneously.

IV. REELYACTIVE INFRASTRUCTURE DESIGN

Based on the considerations identified in the previous section, we present a novel design for a low-power wireless M2M area network infrastructure premised upon two generalizations. First, it is not uncommon for many nodes to exist in proximity of one another, both to implement different technologies and to extend their coverage. And, second, each of these nodes will require a connection to both power and network. Our solution is the combination of a connection topology we call a ‘reel’ and a family of nodes we call ‘reelceivers’.

A. Reel

A reel is a linear daisy chain topology for the interconnection of reelceivers. Power and network connectivity are provided at the start of the chain, and propagate to each subsequent reelceiver over a single cable, as shown in Figure 4. In this topology, a reelceiver is connected to no more than two of its peers. Reelceivers are able to receive data from one neighbour and transmit to the other such that messages propagate by repetition toward the network connection point. Conversely, messages from the network reach each reelceiver in parallel, simultaneously.




Fig. 4. Reel Functional Elements




Fig. 5. Reelceiver Functional Blocks




Fig. 6. Reel Packet for Radio Decoding by Multiple Reelceivers

An advantage of the reel is in the radio-technology-agnostic connectivity that it provides. All reelceiver types can coexist on a reel and can be added in series without requiring additional configuration.

B. Reelceiver

A reelceiver is a radio transceiver of a given low-power wireless technology. Reelceivers connect to one another via a single cable in a reel topology. A reelceiver converts decoded radio packets into reel packets for transmission to the network and vice-versa. Figure 5 illustrates the functional blocks of a reelceiver. The power block consumes power from the reel. The reel communication block decodes reel packets from the network and decodes and propagates reel packets from its neighbour, if present, en route to the network. The radio communication block decodes and encodes radio packets.

Reelceivers are able to detect the received signal strength (RSSI) of each radio decoding, appending this information, as well as their identity, to the resulting reel packet communicated with the network. Moreover, because each reelceiver must forward packets from its neighbour toward the network, it is possible for reelceivers to append this information to reel packets in transit. In the case where multiple reelceivers of a given technology decode the same radio transmission, a single reel packet is communicated with the network. This minimizes overhead and can simplify processing at the packet destination. Figure 6 illustrates the structure of a reel packet where the ID & RSSI component is repeated once for each reelceiver to decode the given radio packet.

V. REELYACTIVE INFRASTRUCTURE IMPLEMENTATION

The infrastructure design presented in the previous section has been implemented in hardware at reelyActive by




Fig. 7. 915MHz Reelceiver PCB and Enclosure

the authors. Here we describe the technical details of the implementation.

A. Reel

The reel is implemented using easily-sourced Cat5e cables with RJ-45 connectors. These cables consist of four twisted pairs which are assigned as follows: one pair each for communication to and from the network, and one pair each for power and ground. The pairs are assigned as per the IEEE 802.3af (Mode B) specification for Power over Ethernet (PoE) [9].

Respect for the PoE pair assignment assures safe failure should the reel be accidentally connected to a PoE (Mode B) power sourcing equipment (PSE). Potential differences from 5VDC to 45VDC are supported across the power and ground pairs, with up to 60VDC safely tolerated. However, neither PoE nor Ethernet are implemented as these require relatively complex components. Instead, differential, unidirectional serial communication using the ANSI-standard RS-422 protocol represents a simple, cost-effective alternative.

The selection of serial communication signaling rate and input voltage results in a tradeoff between reel length and throughput as we will show in the simulation results.

B. Reelceiver

Reelceivers have been implemented using the Texas Instruments CC1110 and CC2541 System-on-a-Chip (SoC). These implement proprietary sub-1GHz and BLE/proprietary 2.4GHz technologies respectively. Reel voltages are handled and converted using a high-efficiency switching regulator. An internal ceramic antenna is used. All reelceivers share a common, compact form factor with dimensions of 76mm by 25mm. Figure 7 shows a 915MHz proprietary reelceiver PCB and enclosure.

VI. SIMULATION AND TEST RESULTS

Initial prototypes of the hardware implementation described in the previous section were completed in April 2012. Systems have been in continuous operation since the following month

TABLE II
REEL THROUGHPUT FOR GIVEN WIRELESS TECHNOLOGIES

Wireless Technology	Wireless Throughput	Reel Signaling Rate
reelyActive	64kbps	160kbps
BLE	375kbps	781kbps

TABLE III
RS-422 SIGNALING RATE VS. MAXIMUM CABLE LENGTH

Signaling Rate	Maximum Cable Length
230.4kbps	434m
460.8kbps	217m
921.6kbps	109m

and have been deployed in real-world environments for a variety of applications. In this section we present the pertinent simulation and test results.

A. Throughput

The data throughput of the reel must meet or exceed that of the wireless technologies it implements or else packets may be dropped. Table II presents the maximum throughput of two currently supported wireless technologies and the corresponding minimum signaling rate required on the reel. In this exercise, packets are assumed to contain only the advertiser address as relevant payload. The given throughput is an upper bound representing the unrealistic case where a continuous stream of radio packets immediately follow one another in time without collisions.

The minimum reel signaling rate is calculated by accounting for the additional four bytes (overhead, ID and RSSI) of the reel communication protocol and the two overhead bits (start and stop) per asynchronous serial byte transmission. In order for a single reelceiver to support the maximum theoretical BLE throughput, the nearest standard signaling rate of 921.6kbps would be recommended. Initial implementations have been limited to the standard rate of 230.4kbps for compatibility with legacy off-the-shelf serial device servers. Pending third-party hardware support, or a device of our own design, 921.6kbps will be implemented.

B. Maximum Reel Connection Length

The maximum reel connection length between two reelceivers is limited by the characteristics of the RS-422 protocol. An industry rule of thumb states that the data signaling rate (in bps) multiplied by the cable length (in metres) should not exceed 10^8 [10]. Table III lists the maximum cable lengths for the implemented and recommended signaling rates based on this rule. Reel connections of at least 100m are supported, which could be envisaged for outdoor line of sight deployments where wireless range is on this order of magnitude, as we show later in this section.

C. Maximum Overall Reel Length

The maximum overall length of a reel is limited by the input voltage and losses due to cables and connections. Additionally,

the maximum reel connection length distances described in Table III must be respected. Assuming the ideal case where connection losses are negligible and the maximum input voltage of 45VDC is used, the maximum length of a reel can be estimated based on DC transmission line losses. A reelceiver requires a minimum of 5VDC for correct operation, and, at this voltage, it is experimentally determined to draw up to 40mA of current. The maximum resistance of the DC transmission line is therefore given by Equation 1.

$$R_{max} = \frac{V_{loss}}{I} = \frac{45V - 5V}{40mA} = 1k\Omega \quad (1)$$

And the maximum reel length can be calculated as per Equation 2 based on the characteristic DC loop resistance of Cat5e cable, which is typically no more than $200\Omega/km$.

$$Distance_{max} = \frac{R_{max}}{R_{DCLoop}} = \frac{1k\Omega}{200\Omega/km} = 5km \quad (2)$$

This theoretical limit of 5km is two orders of magnitude greater than the reel lengths typical of deployments to date. Reels of fifteen reelceivers over nearly 100m have been successfully deployed with 24VDC input voltage. In practice, input voltages beyond 24VDC are rarely used due to the lack of economical, off-the-shelf power supplies. Resistive losses at connections are observed to be minor but non-negligible.

D. Wireless Range

The maximum range between transmitters and reelceivers is influenced by many factors. It is nonetheless possible to estimate the maximum range under ideal conditions. Consider 915MHz reelyActive active RFID tags and reelceivers. At 1m range, with line of sight, typical RSSI is -60dBm. Given that the maximum sensitivity of the reelceiver is approximately -100dBm, this leaves a power budget of 40dBm (or 10^{-4} times the power at 1m). From the Friis transmission equation, where power is proportional to $\frac{1}{R^2}$, the maximum wireless range can therefore be estimated at 100m.

In practice, outdoors with line of sight, decodings have indeed been observed at over 100m. Indoor range without line of sight is typically observed at 10m to 20m depending on the environment.

VII. REEL INFRASTRUCTURE IN THE M2M SPACE

In the previous sections we have described the architecture, characteristics and experimental results of the reel infrastructure, based on which we may now position the design among M2M technologies and standards.

A. Area Network Scope

The wireless range of the reel is similar to that of WiFi, the most prevalent wireless Local Area Network (LAN) infrastructure. Reel infrastructure is best classified as a LAN for low-power wireless devices. As a short range communication system for heterogeneous wireless technologies, the infrastructure is positioned between the smartphone-based PAN and the WiFi-based LAN. In fact, reels have the potential to offload

lower-throughput devices from the Radio Access Network (RAN) of both. Connectivity can be brought closer to the wireless devices thanks to the relatively low cost per bit of the reelceiver.

B. Interfaces

The reel infrastructure is typically connected to a Wide Area Network (WAN) via a gateway. As per ETSI M2M standards, operation among Service Capability Layers (SCL) is supported by means of interfaces following a RESTful approach such as dIa for device/gateway service capabilities, and mId and mIa for network service capabilities [11]. Although the reel does not itself implement REST, the gateway connected to its serial stream can implement a REST interface on its behalf. Similar to what a dIa reference point does for a gateway SCL, this software interface enables integration within the M2M area network or with the M2M core through the WAN.

An alternative approach is to process the reel packets remotely. A variety of commercial serial device servers can encapsulate the RS-422 serial stream in UDP or TCP/IP packets for transmission to a remote server over Ethernet, WiFi or Cellular networks.

In either approach, software may act as an avatar for the wireless devices, as well as the reel and its reelceivers. This enables local and remote monitoring and, assuming that each device can be uniquely identified, the software may implement M2M or IP protocols on its behalf. For this reason, all reely-Active devices have a unique EUI-64 identifier in anticipation of a future software mapping to IPv6. The development of this software is an ongoing work and is outside of the scope of this paper.

VIII. DISCUSSION AND ONGOING WORK

A. Initial Applications

Initial deployments have focused on active RFID applications using tags of our own design. The reel architecture provides a simple, accessible means for identifying and locating devices in space. In a typical application, the reel is connected to an off-the-shelf serial server which encapsulates reel packets in UDP packets for transport to the reelyActive cloud service. The cloud service receives and interprets the packets, and stores the wireless payload (including the unique device identifier), RSSI and timestamp. The collected information is available for consumption by higher level applications via a RESTful API.

B. Emerging Applications

An emerging focus is the provision of ambient connectivity for the rapidly increasing number of BLE devices on the market [12]. BLE-enabled smartphones and tablets represent a mobile, distributed PAN infrastructure for devices such as smart tags, sensors and beacons. BLE reelceivers can offload the mobile RAN by providing equivalent connectivity to these BLE devices. Moreover, reel infrastructure enables real-time point-of-interest location for any BLE-advertiser, which extends to the smartphones and tablets themselves.

As heterogeneous low-power wireless M2M area networks gain in prevalence, we can expect resource-constrained devices to migrate from technologies such as WiFi and even cellular to more efficient low-power technologies [7]. This is especially true for Internet of Things devices which are often intended to reside in human spaces traditionally served by WiFi.

The development of reelceivers which support additional wireless technologies such as IEEE 802.15.4 will likely be driven by market adoption and demand. It is interesting to note that even long-range technologies such as UNB could fit the reelceiver model. In this case, only the node density required to achieve seamless coverage would differ.

C. Real-Time Location Capability

The reel architecture was inherently designed for device mobility. Not only does it provide ambient connectivity and seamless extension of coverage, but it also includes real-time location capability. The location of a transmitting device is estimated as the area surrounding the reelceiver with the highest RSSI. By strategically placing reelceivers at human-recognizable points of interest, semantic location is possible, for example: "the item is in the supply closet". A given location granularity can be achieved by adjusting the spacing and placement of reelceivers.

Since, as we have shown, all of the RSSI values are consolidated into a single reel packet, location can easily be determined at the gateway level. This enables spatial contextual-awareness at the gateway which may be used for local management of peer to peer connectivity.

D. Distribution

Our team are excited about the possibility for reels to be deployed as distributed, crowdsourced infrastructure for ambient connectivity. Given the positive results of experimentation to date, the configuration-free set-up-and-forget simplicity of the system and the relatively low BOM costs, it is not unimaginable to envisage a distributed network of public, Internet-connected reel picocells.

IX. CONCLUSION

In this paper we have presented the interest for a low-power wireless M2M infrastructure. We have proposed a novel, versatile design capable of supporting multiple technologies and providing extensive coverage, and have detailed its implementation. Initial results are promising and development continues not only on hardware and firmware, but also on the software which empowers the simplest of devices with M2M/IP connectivity.

ACKNOWLEDGMENT

The authors would like to thank their family and friends for their continued support even in the absence of an understanding for the ambition to "connect things"!

REFERENCES

- [1] Yan Zhang *et al*, "Home M2M networks: Architectures, standards, and QoS improvement", *IEEE Commun. Mag.*, vol. 49, no. 4, pp. 44-52, Apr. 2011.
- [2] OECD (2012), "Machine-to-Machine Communications: Connecting Billions of Devices", *OECD Digital Economy Papers*, no. 192, OECD Publishing, Jan. 2012.
- [3] *Machine-to-Machine communications (M2M); Functional architecture*, ETSI TS 102 690 v1.2.1, 2013.
- [4] Kim Chang *et al*, "Global Wireless Machine-to-Machine Standardization", *IEEE Commun. Mag.*, vol. 15, no. 2, pp. 67, Mar. 2011.
- [5] Geng Wu *et al*, "M2M: From Mobile to Embedded Internet", *IEEE Internet Comput.*, vol. 49, no. 4, pp. 37, Apr. 2011.
- [6] Michael Starsinic, "System Architecture Challenges in the Home M2M Network", *Long Island Systems Applications and Technology Conference*, Farmingdale, NY, 2010, pp. 6.
- [7] Phil Smith, "Comparing Low-Power Wireless Technologies", CSR PLC, Cambridge, UK, Tech. Rep. CS-213199-AN, May 24, 2011.
- [8] "SIGFOX Flyer", Sigfox, Labege, France. [Online]. Available: <http://www.sigfox.com/fr/download/fichiers/118/Flyer-SIGFOX-v1.2.pdf>
- [9] *Ethernet*, IEEE 802.3af-2003, Jun. 2003.
- [10] Manny Soltero, Jing Zhang and Chris Cockril, (2002, Jun.), "RS-422 and RS-485 Standards Overview and System Configurations", Texas Instruments, TX. [Online]. Available: <http://www.ti.com/lit/an/slla070d/slla070d.pdf>
- [11] *Machine-to-Machine communications (M2M); m1a, d1a and m1d interfaces*, ETSI TS 102 921 v1.2.1, 2013.
- [12] *2012 Annual Report*, Bluetooth SIG, 2012, pp. 5.