The Internet of Things is not a scaled down version of the Internet

A Thinnect Whitepaper
Jürgo-Sören Preden, PhD

Change of the Computing Landscape

The shift to IoT that is happening now can be compared to the shift from mainframes to personal computers in the 1980s. With the advent of personal computers the architectures of hardware and software were changed as the multitude of users formerly running their programs on a single processor (in the mainframe) now transitioned to single users commanding a single processor. Within a decade it became possible to run several applications on the same processor in a pseudo-parallel fashion. As the next paradigm shift the Internet enabled connecting the individual processors (personal computers) across the world together via a network, providing quick and simple access to information. While the internet changed the way information was transported to and from the personal computer the fashion in which individual computers were used had not change much since the days of the mainframe – the user was still entering the commands via a user interface and waiting for a response in return.

With adoption of IoT the number of 'connected processors' (to the internet) will increase many fold but also the way the computers are used will change fundamentally. No longer will we be interacting directly with the computers, instead there will be thousands of computers embedded in the environment around us, running in the background, expected to provide a persistent service. While there may be benefits in some applications from the use of centralized computing, this is not the case with IoT. In applications that are closely interacting with the physical world (which the IoT applications are) there are far more benefits to be gained from localizing computing, some of which were witnessed during the transition from mainframes to personal computers e.g., latency, cost savings, availability, etc. Today's technology offers a unique opportunity to implement IoT in a radically new way distributing the computation and bringing intelligence to the very edge of the network, enabling the devices to operate autonomously, providing persistent services needed for realizing IoT applications.

Real World Aspects of IoT

The devices that make up Internet of Things are embedded computing devices, which communicate via wireless interfaces. The fact that the number of these devices is going to be high and that they are going to be in close proximity means that protocols such as or

similar to IEEE 802.15.4 and Bluetooth are also going to prevail as the communication medium for IoT devices. The reasons for this are derived from basic physics – the signal strength used for communication must be low to limit interference between devices, and the packet size is going to be small to limit the probability of packet loss. Another factor that will have a great influence on the design of IoT devices is the market's desire to have many of these devices battery powered. Battery powering IoT devices will substantially reduce the costs and complexity involved in the deployment of these devices – removing the cabling otherwise required extends the options available for placing these devices.

Based on the projected very high numbers of IoT devices there will also be considerable market pressure to achieve very low pricing, resulting in the use of more cost effective microcontrollers, limiting processing power and memory size. All of the factors listed above result in significant design constraints for IoT system designs, addressing these constraints will to a great extent determine how successful the Internet of Things will become. To extract from the above, the main properties of IoT are: limited communication bandwidth, unreliable communication channels, unpredictable communication paths (and limited ability of the devices to store communication packets), limited memory sizes and processing power.

Another crucial consideration is that the purpose of IoT devices fundamentally differs from that of most of the conventional devices connected to the Internet today. Instead of sharing data for services which typically can tolerate quite high delays, IoT will be monitoring and controlling the real world. Aside from increasing the need for responsiveness (or low latency communication) this will also create high expectations for reliability, beyond those being offered by the conventional internet, today. Finally, since much of the data being harvested/processed will be highly proprietary, this will also place higher demands on security.

While it is no doubt desirable to view IoT as a mere extension of the conventional internet, and therefore implementable using existing internet methodology, after reviewing the above it should be apparent that whilst attempting to adapt existing technologies and protocols offers lesser design demands, settling for this approach will yield a solution, which will not provide the features and properties we are expecting.

IP Connectivity in IoT

One of the more questionable trends (worth exploring) in IoT lies in providing IP-level connectivity to the very edge of the network. This means that it will be possible to communicate IP packets from the Internet to every sensor and actuator by the router at the edge of the network. While in theory this would offer clear benefits, such as the ability to use known and internet-proven protocols and removing the need for application-level knowledge at the network gateway (referred to as 'border routers' in IoT networks), there are also many drawbacks to this approach. The technical challenges involved with providing IP connectivity to the end devices start with security and end with the constraints of the resource-limited networks (such as IEEE 802.15.4 or Bluetooth LE mesh) in processing and communicating IP packets. Security is a critical topic which deserves its own, stand-alone

analysis, beyond this paper. The vulnerabilities that can be exploited using various types of attacks on resource constrained networks that are IP-enabled are numerous. Even the simplest type of attack – the Denial of Service (DoS) attack would be extremely straight forward to mount on a resource-constrained network; which it is challenged even by handling regular application traffic.

The ability of resource constrained networks to process and communicate IP packets represents a very serious challenge. Although the protocol for supporting IP connectivity at the edge has been standardized with 6LoWPAN and the RPL routing protocol, the scalability of the methods it describes is highly questionable. 6LoWPAN with RPL fails to address 2 distinct issues, one of which is the inefficiency of node-to-node communication in the network (as all communication gets routed through a border router at the edge of the network) and the other issue comes from the bandwidth required by IP communication compared with the available bandwidth in resource-constrained networks.

The limitations of bandwidth availability in resource-constrained networks will be clear when we compare the packet size of IP packets with the available packet size on a resource-constrained network, such as IEEE 802.15.4. The maximum packet size of an 802.15.4 packet is 127 bytes, but one must also consider the overheads associated with asynchronous, wireless communication. The frame overhead of the 802.15.4 packet can be 25 bytes, the size of the compressed IP header as defined by the 6LoWPAN protocol can be 12 or 20 bytes, leaving just 90 or 82 bytes for the payload. If link layer security is also used an additional overhead of 21 bytes is added, leaving just 69 or 61 bytes for the payload. The packet sizes are visualized on the figure below.

Frame overhead	IP header	Link layer security	Payload
25 B	12 B	21 B	69 B

Frame overhead	IP header	Link layer security	Payload
25 B	20 B	21 B	61 B

Figure 1. 6LoWPAN packet sizes – header vs payload

As IPv6 requires the size of the maximum transmission unit to be at least 1280 bytes, an IP packet clearly does not fit in a single 802.15.4 packet. The IP packet must therefore be fragmented, meaning that it will take up to 19 packets on the 802.15.4 network to communicate a single IP packet. While this might not be an issue in ideal conditions, one must also account for the relatively high error rates typically experienced on wireless networks in potentially noisy environments (industry sources indicate that an error rate of about 25% is typical) resulting in further overhead being introduced by retransmissions. In a practical example consider the case of a mesh network which consists of 10 devices in a 10 hop configuration: the theoretical maximum packet rate under typical conditions is about 4 IP packets per second. If bidirectional communication occurs (which is the norm for IoT networks), then the throughput will be substantially lower (can be reduced to half the

throughput, resulting in a throughput of 2 interactions with a sensor device per second). It must be noted that this is the maximum throughput for the entire network that is connected to the internet via a single gateway (a border router in the 6LoWPAN terminology). As the throughput of the network is reduced the communication delay for every single packet is increased.

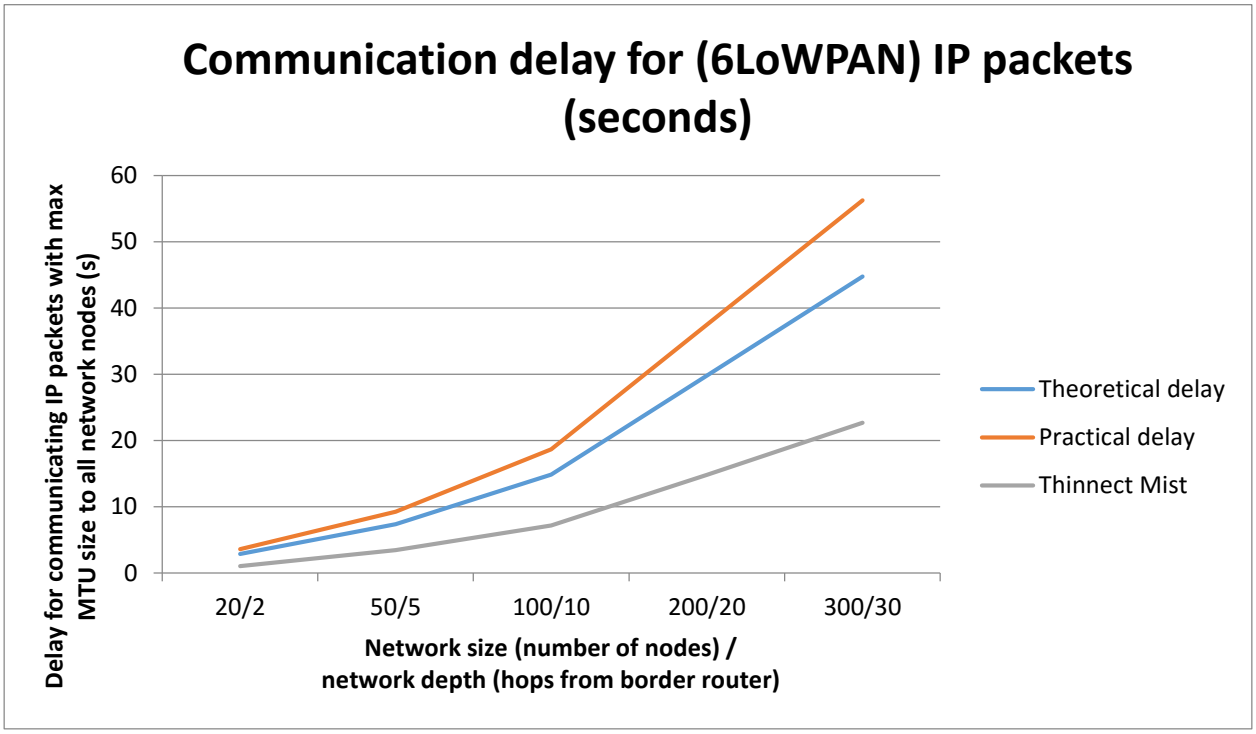

Figure 2. Communication delay in a 6LoWPAN network

As indicated above the average packet loss rate in a 802.15.4 network is 25%, which is caused by interferences, collisions and other factors. Implementing a protocol with high overhead on such an inefficient transport aggravates the inefficiency of the network, potentially rendering is unusable.

Although these performance figures are already low, one must also consider the fact that the buffer sizes of resource constrained devices are also limited, which means that the rate at which packets can be communicated to the network is limited by the throughput of the network (there are no efficient flow control methods available for these networks). If there are several IP clients to a single mesh network the total number of packets directed to the mesh network cannot exceed the theoretical maximum throughput of the network. So even if the network is able to survive with one dedicated client, it may collapse when another client is added to the network. As the clients are not coordinated there is no way to guarantee the operation of the network in a multiple client configuration as the communication breaks down when the practical maximum throughput of the network is reached. This means that it will be extremely difficult to guarantee the operation of a resource constrained mesh network that supports IP communication.

Another crucial shortcoming of 6LoWPAN with RPL is the latency introduced by the routing protocol. In IoT networks that are part of the physical world the latency that can be tolerated in actuation loops is quite low – the delay between the human flipping a switch and the lights turning must be minimal for the human to perceive these events as simultaneous, clearly a delay of several seconds is unacceptable. However, in a 6LoWPAN/RPL network the delay in a non-trivial network (even in ideal conditions) can amount to seconds when non-storing routing mode is used in nodes (which is the typical case). Below, on Figure 3 the network delay between adjacent nodes is depicted. A real world example of adjacent nodes is a light switch and a light in a room, as in case of 6LoWPAN/RPL all communication is directed via the gateway a substantial delay is introduced, which increases as the network size is increased.

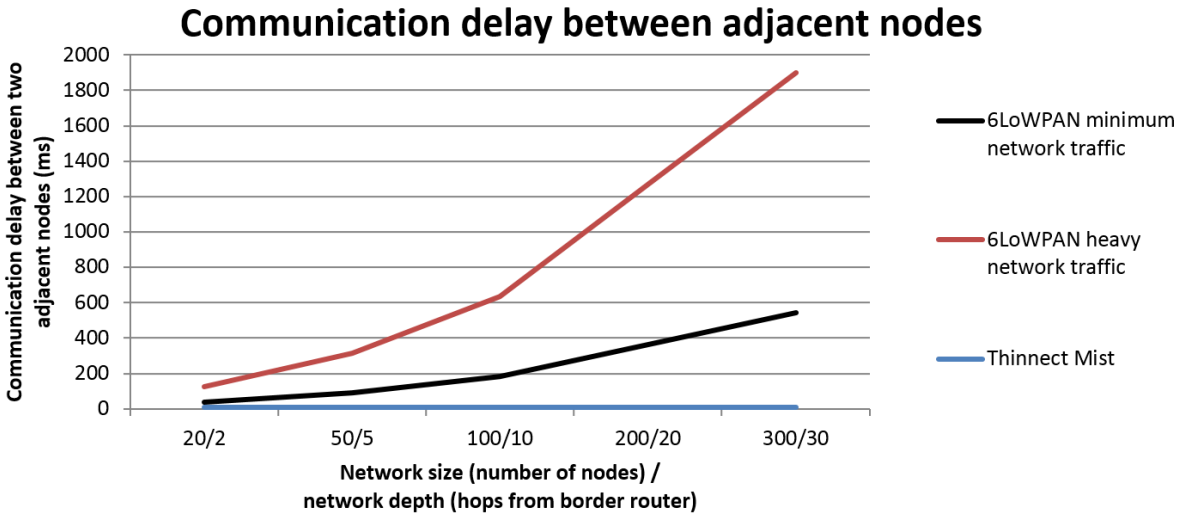

Figure 3. Communication delay between adjacent nodes in 6LoWPAN/RPL and Thinnect Mist networks

One must also keep in mind that the graphs on Figure 3 represent the ideal conditions with optimal timings and minimum interferences in the network, in unfavorable conditions the delay is going to be greater.

As the local delays in the 6LoWPAN/RPL network are affected by the network size an important effect comes to play in these networks. Even though a network mat meet the performance requirements as it is deployed, as additional devices (e.g., sensors for detecting presence or measuring environmental parameters) are added to the network, its performance may degrade to a level where the network is not able to provide a service at an acceptable level.

Alternative to IP in IoT

A feasible alternative to using Internet protocols in IoT is to use the Mist computing concept with dedicated binary protocols in the mesh network (i.e., the network of sensors and actuators, such as movement sensors, light switches and lights). Mist computing with binary protocols introduces several orders of magnitude lower bandwidth requirements with

interaction patterns suited for low bandwidth networks. In Mist computing the devices act using high level rules provided by the system manager, minimizing communication overhead and latency. The Mist computing approach differs fundamentally from the request-response type of interaction (as enabled for example by the COAP protocol) in IP-enabled mesh networks (inspired by interactions that occur in the Interne).

The Language for IoT

Just as important as the communication protocols used is the ontology or "language" that is used to designate the resources available in an IoT network and the data and commands exchanged within the network. Using a standard language in the entire application across all the networks and devices (e.g., mobile devices, sensors, Cloud servers) that the application spans ensures that all application components can be integrated with minimal effort, that the application can be extended easily and that individual IoT devices can be used across multiple applications. Initiatives, such as the one being driven by the Open Interconnect Consortium (OIC) strive to do exactly this – standardize a language for IoT applications. The languages for data representation were not important with Internet as it is mostly people who are the consumers of information and they are good at interpreting the spoken languages, which are used in the Internet to convey information. There was an effort with the Semantic Web to formalize the data and information on the Internet but as there was no real requirement for this, the initiative did not get too far.

The language defined by OIC was originally intended to be used on IP networks, however the same language can be used just as well with Mist computing, yielding the benefits of reduced bandwidth usage and optimized application behavior while maintaining compatibility across devices and networks. Just as with 6LoWPAN the communication to a mesh IoT network is handled by dedicated gateways, which encapsulate the payload of IP packets in a dedicated binary protocol and vice versa – packets in binary format originating from end devices can be translated to IP packets at the gateway and seamlessly communicated to their destinations in the internet.

When every IoT device makes its resources available using the Mist computing approach and is able to communicate in a standard language (such as the one specified by the OIC), future proof applications can be created. Using a standard language across all applications ensures that individual devices (such as movement sensors) can be reused in multiple applications, saving on the purchase and deployment costs and thereby lowering the total cost of applications for the user. This approach will significantly improve the 'cost/benefit equation' for IoT enhancing its attractiveness and adoption by the market in general.

This document summarizes some of the aspects pertinent to applying Internet technologies in the IoT domain. If you would like to receive further information, please contact

Jurgo Preden,
CEO
jurgo@thinnect.com