Temptations and Difficulties of Protocols for Smart Objects and the Internet

Alexandru Petrescu alexandru.petrescu@cea.fr February 27th, 2011.

Smart Objects and the Internet is an exciting topic of exploration. The excitement has been increasingly stimulated by factors which, in the past, helped the initial spread of the Internet: *miniaturization;* then, the improved and smaller made possible mini and micro-computers – now it logically paves the way to nano-computers[1], with comparable capabilities. Fiction of machines gaining control over humans aside, there is real technology opportunity and business cases to be made out of connecting Smart Objects to the Internet. Casual statements such as '*a body scale brags weight loss*', '*a rabbit warns of weather*', '*a pillow rocks asleep*' recently led to technology developments, business cases, visions, all conditioned by one single fundamental *if*: if these objects were connected to the Internet. Sure, one could have a rabbit plush equipped with a pressure sensor making it warn – sometimes – of storm; however, make same rabbit `listen ()` to the freely available supercomputed weather models and 10-minute satellite radars (on the Internet) and suddenly that plush is much more dependable.

In a context of products already marketed as Internet objects, what is there still to be excited about for the protocol designer? One should consider objects further along the miniaturization scale. The plush-sized Objects (10cmx10cm) can be considered as barely a beginning towards much smaller devices. But when looking closer one sees it as a limit as well: it's harder and harder to put IP into something smaller than 5cmx5cm integrating plush's Wireless, battery, CPU and memory for an IP stack. The typical trade-offs when satisfying physical board size led to eliminating the Operating System features and the protocol Stack together with it[2]. Or, without an IP stack, there is no Internet. Some of the smallest available independently running devices are GPS receivers (just a bit bigger than a large fingernail) – they run two wireless interfaces and maybe a wired one, yet no IP. Should the GPS receiver be connected to the Internet? Should every sensor device out there be connected to the Internet, act as an End node?

Designing protocols to connect small devices to the Internet is a respectable endeavor. Some of the representative developments witnessed at the IETF showed consensual code achievements; yet most are subject to a common pitfall – that of ignoring the infinitely large while focusing too much on the smaller matters. It is like the smallness parts were a clean canvas on which one *could* draw IP without ICMP. Some of these small parts are so small that they couldn't ever run TCP – why connecting them to the Internet then?

Another fallacy is that of ignoring the overarching importance of the link layer. Some of the devices on which IP is supposedly running simply could not run a MAC layer of some IEEE links, for the same reasons of memory size. Some link layers do support link-layer multicast natively some others don't. The link layers should be used according to their specifications, and not trespassed (e.g. resist the temptation of using IP routing to extend PHY range).

It is difficult to use 16byte long addresses on a device which may have 256bytes of memory.

In a system where every byte counts, it is strange to use 256bits of addressing and 96bits of authentication to just send a bit-valued (1 bit!) of presence or alarm.

Acknolwledging some of these factors may lead to designs different than that of Internet: maybe a different addressing scheme together with a different routing mechanism, which could still link together a large number of Smart Objects and the Internet.

It may very well be the case that not all Objects need to be IP-addressable. In today's Internet, there's an IP address for the NIC and a completely different memory address for byte 6 at location 0x02, although they're on the same computer and NAT evil is not being used. If every CPU needed to access every disk sector location then IPv6 space would be largely unfit. It may be that the only entities which need to be IP-addressable are those who have enough compute power to compute a CRC in time enough.

---

[1] The term nano-computers should not be taken literally to mean nano-meter scale – that is way beyond the silicium transistor based computer technology of today. The same with nano-communications.

[2] It is hard to relate code size to board size, it is hard to trade one for another, because a bit is like a point – it has no dimensions. Wherease there can be no bit-centimeter constant, there could probably exist a bit-second-Joule constant somewhere.