Internet of things

What's up and why?

Patrik Fältström – Office of the CTO

paf@cisco.com

This is not something new

The Trojan Room Coffee Pot

A (non-technical) biography

by Quentin Stafford- Fraser

Several people have asked about the origins of the Trojan Room coffee pot. It started back in the dark days of 1991, when the World Wide Web was little more than a glint in CERN's eye. I was working on ATM networks in a part of the Computer Lab known as the Trojan Room, (a name which, perhaps, causes some amusement to American readers). There were about fifteen of us involved in related research and, being poor, impoverished academics, we only had one coffee filter machine between us, which lived in the corridor just outside the Trojan Room. However, being highly dedicated and hard-working academics, we got through a lot of coffee, and when a fresh pot was brewed, it often didn't last long.

Some members of the 'coffee club' lived in other parts of the building and had to navigate several flights of stairs to get to the coffee pot; a trip which often proved fruitless if the all-night hackers of the Trojan Room had got there first. This disruption to the progress of Computer Science research obviously caused us some distress, and so *XCoffee* was born.

In the Trojan Room there were several racks of simple computers used in the testing of our networks. One of these had a video frame-grabber attached and was not being used at the time. We fixed a camera to a retort stand, pointed it at the coffee machine in the corridor, and ran the wires under the floor to the frame-grabber in the Trojan Room. Paul Jardetzky (now working in California) then wrote a 'server' program, which ran on that machine and captured images of the pot every few seconds at various resolutions, and I wrote a 'client' program which everybody could run, which connected to the server and displayed an icon-sized image of the pot in the corner of the screen. The image was only updated about three times a minute, but that was fine because the pot filled rather slowly, and it was only greyscale, which was also fine, because so was the coffee.

Just more things, and more fun!

Healthcare

Defense

Energy Saving (I2E)

Predictive maintenance

Improve Productivity

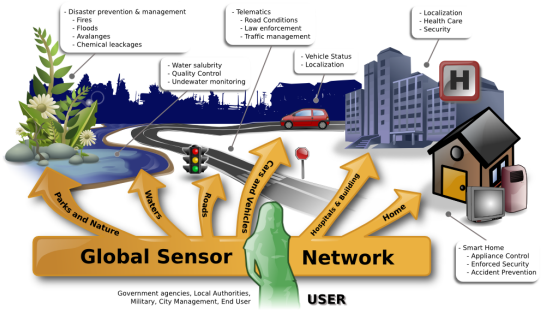

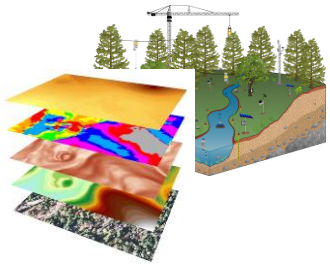

New Knowledge

Intelligent Building

Agricultural

Smart Cities

Industrial Automation

High-Confidence Transport and assets tracking

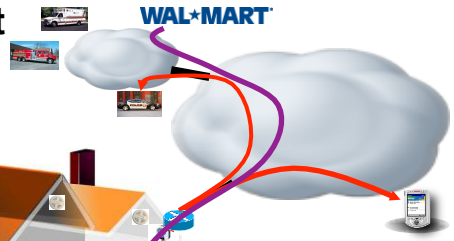

Smart Grid

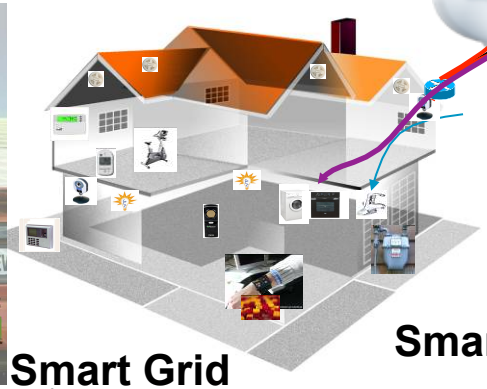

Smart Home

We build this with the Internet Protocol

- For any kind of reach, noone is building it with anything else than the Internet Protocol, IP
- Reach is important, because we are not talking about a remote control to a TV, or a cordless phone
- We know how IP works, we have IP networks all over the place
- We will have local networks, with specialized protocols, but they have to at some time change to IP
- The more IP we do, the more we will reuse what is already done, because it is tested and it works

One solution: IP

- IP is independent of the physical layer.
- Does IP work on highly constrained devices (meaning small and with limited memory) such as a 16-bit microcontroller with tens of Kbytes of RAM and possibly battery operated?
- Absolutely: this has been very successfully demonstrated and there ARE several deployments!
- The suggested IP stack only needs about 4K RAM
- Will there be new protocols related to IP?
  - Many of the existing protocols can be used with no additional cost
  - New protocols will be implemented only when and where needed

One physical layer will not fit all

- Different requirements

  Range, power, bandwidth, frequency band, media, security, …

- Lead to different physical layers

  Radio, ethernet, powerline, …

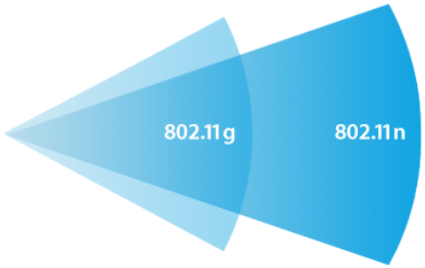

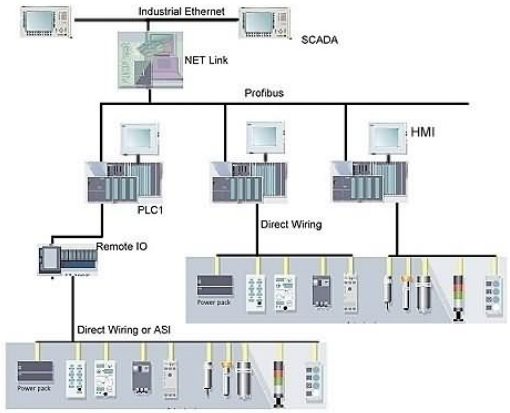

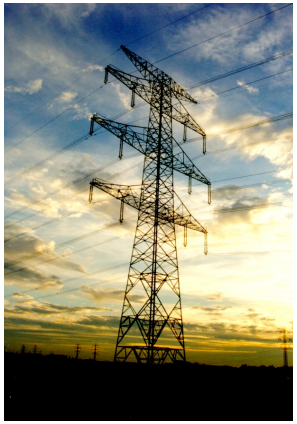

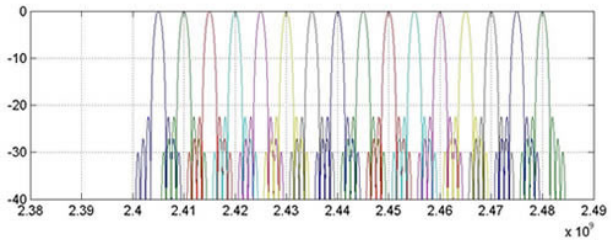

IP is **the** common interface

New applications pretty much every day … but …

- The number of ***proprietary*** solutions has literally ***exploded***: Zigbee, Z-Wave, Xmesh, SmartMesh/TSMP, … at many layers (physical, MAC, L3) and most chip vendor claim to be compatible with their own ***standard***
- Many non-interoperable "solutions" addressing specific problems ("*My application is specific" syndrome*)
  - Different **Architectures**,
  - Different **Protocols**

=> Deployments are limited in **scope** *and* **scale,**

Sensor Networks - Usually a constrained environment requiring adaptation

- Network and sensors that play together
- Energy consumption is a major issue
- Limited processing power (CPU, memory),
- Prone to failures => very dynamic topologies,
- When mobile => increase the dynamic nature of topology,
- Data processing may be required on the node itself,
- Sometimes deployed in harsh environments,
- Potentially deployed at very large scale,
- Must be self-managed (auto-discovery, self-organizing networks)

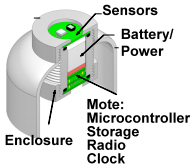

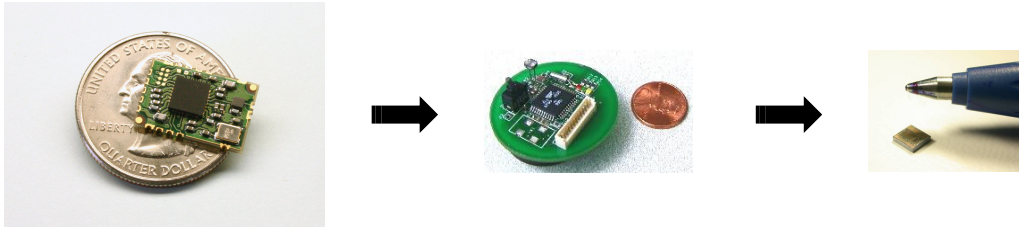

Where the key work is done

- IETF:
  - IPv6 over low power radio
  - Routing Over Low power and Lossy networks
- ISA:
  - Standards for machine to machine communication
- IPSO:
  - IP for smart objects alliance
  - http://www.ipsoalliance.org/

An evolution of the Internet

- "Internet of things", "The Future Internet" and such, is just terminology for evolution of the existing Internet we have **today.**
- Internet exists, Internet grows. We all connect new things to the Internet every day. What we talk about is nothing new. When something is connected to the Internet, it ends up being part of the Internet.
- At the same time, we must learn from history.
- Architectures should be open, and that include standards development. This to foster innovation and competition.

Policy and regulation, RFID

- An RFID tag passes a reader, and generate an event. The event can be recorded, but **where**, and **why**?
- Many solutions exist for finding repository. One of them is a DNS based resolution, called ONS. **Other solutions** under development, including p2p variants.
- There might be problems to access that data. **Cost?** Is it cheaper to add a new tag than use an existing tag?
- Privacy issues; what mechanisms exist to define **acceptable** tracking of RFID tags?
- **What is acceptable, and what is not, should be defined by, and evaluated with, existing regulation!**

What is important for a success?

- Ensure that the architecture does not lead to specific business models
- The Internet Protocol is ubiquitous and offers a cost-effective means to create innovative architectures to support multiple industry sectors and business models
- Finding the location for the information associated with a “tag” is a different issue from being able to access that data
- Existing policies can be applied to new problems, instead of inventing technology specific regulation
- The Internet of Things is already here

We call it The Internet!

Questions

Patrik Fältström – Office of the CTO
paf@cisco.com