HOMEWORK

Taming the 800 pound gorilla - home networking made easier

Joe Sventek, University of Glasgow
joseph.sventek@glasgow.ac.uk

The University of Nottingham · BT · Georgia Tech · University of Glasgow · Imperial College London · Microsoft Research

EPSRC
Engineering and Physical Sciences Research Council

Dramatis Personae

- Nottingham
  - T. Rodden
  - T. Lodge
  - B. Bedwell
  - K. Glover
  - R. Mortier
- Glasgow
  - J. Sventek
  - A. Koliousis
  - O. Sharma
- Imperial
  - N. Dulay
  - D. Pediaditakis
  - M. Sloman
- And others ...
  - P. Tolmie (Nottingham)
  - D. Pezaros(Glasgow)
  - M. Sevegnani (Glasgow)
  - M. Calder (Glasgow)
  - H. Rotsos (Cambridge)
  - E. Lupu (Imperial)

The Problem

- Home networking gear is the most returned consumer electronics item (25%) - Consumers cite technical complexity as the largest barrier to home networking
- Home users are forced to be network administrators using protocols and tools designed in the ‘70’s and ‘80’s for trained administrators
- **Home users lack not only the skills, but also the motivation**
- How would you design network technology based on human considerations, such as understandability, usability, and manageability (in addition to “traditional” considerations such as latency, scalability, etc.)?

Homework Project

- The Approach
  - Create a home network router, connected to the home broadband connection, that passively monitors all traffic in the home network
  - Make this monitored data available in real-time to display, persistence, and reaction applications
  - Provide display and control functionality to home users that is intuitive to navigate and use
  - Iterative design, implementation, and deployment strategy
    - Work with real world users in real domestic settings in the UK and US
    - Prototype a manifestation and deploy in real homes; by understanding how people use the system, inform the next generation of management techniques, network management and modelling approaches
    - Repeat this iterative process three times, growing the number of households as we move through each generation

The Vision

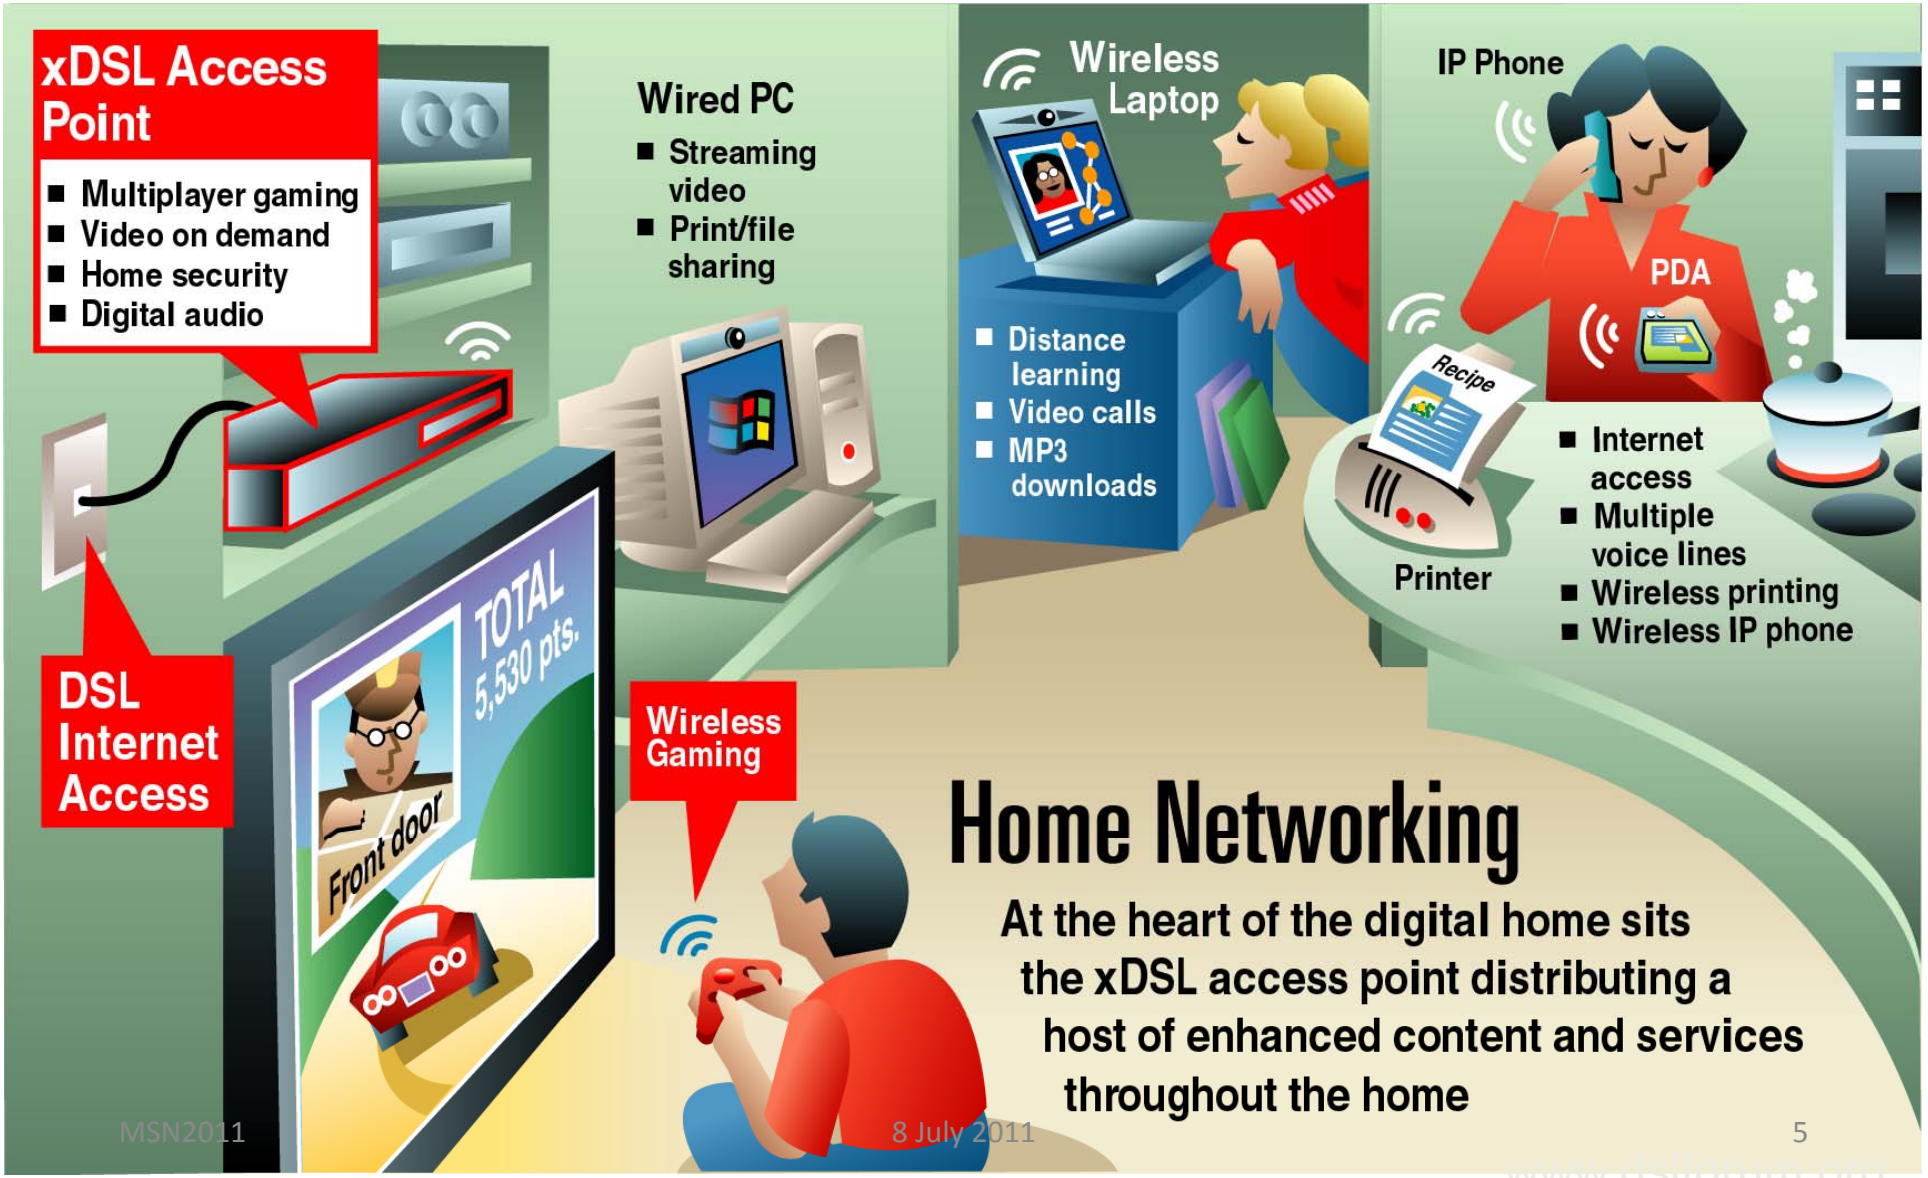

The Complexity

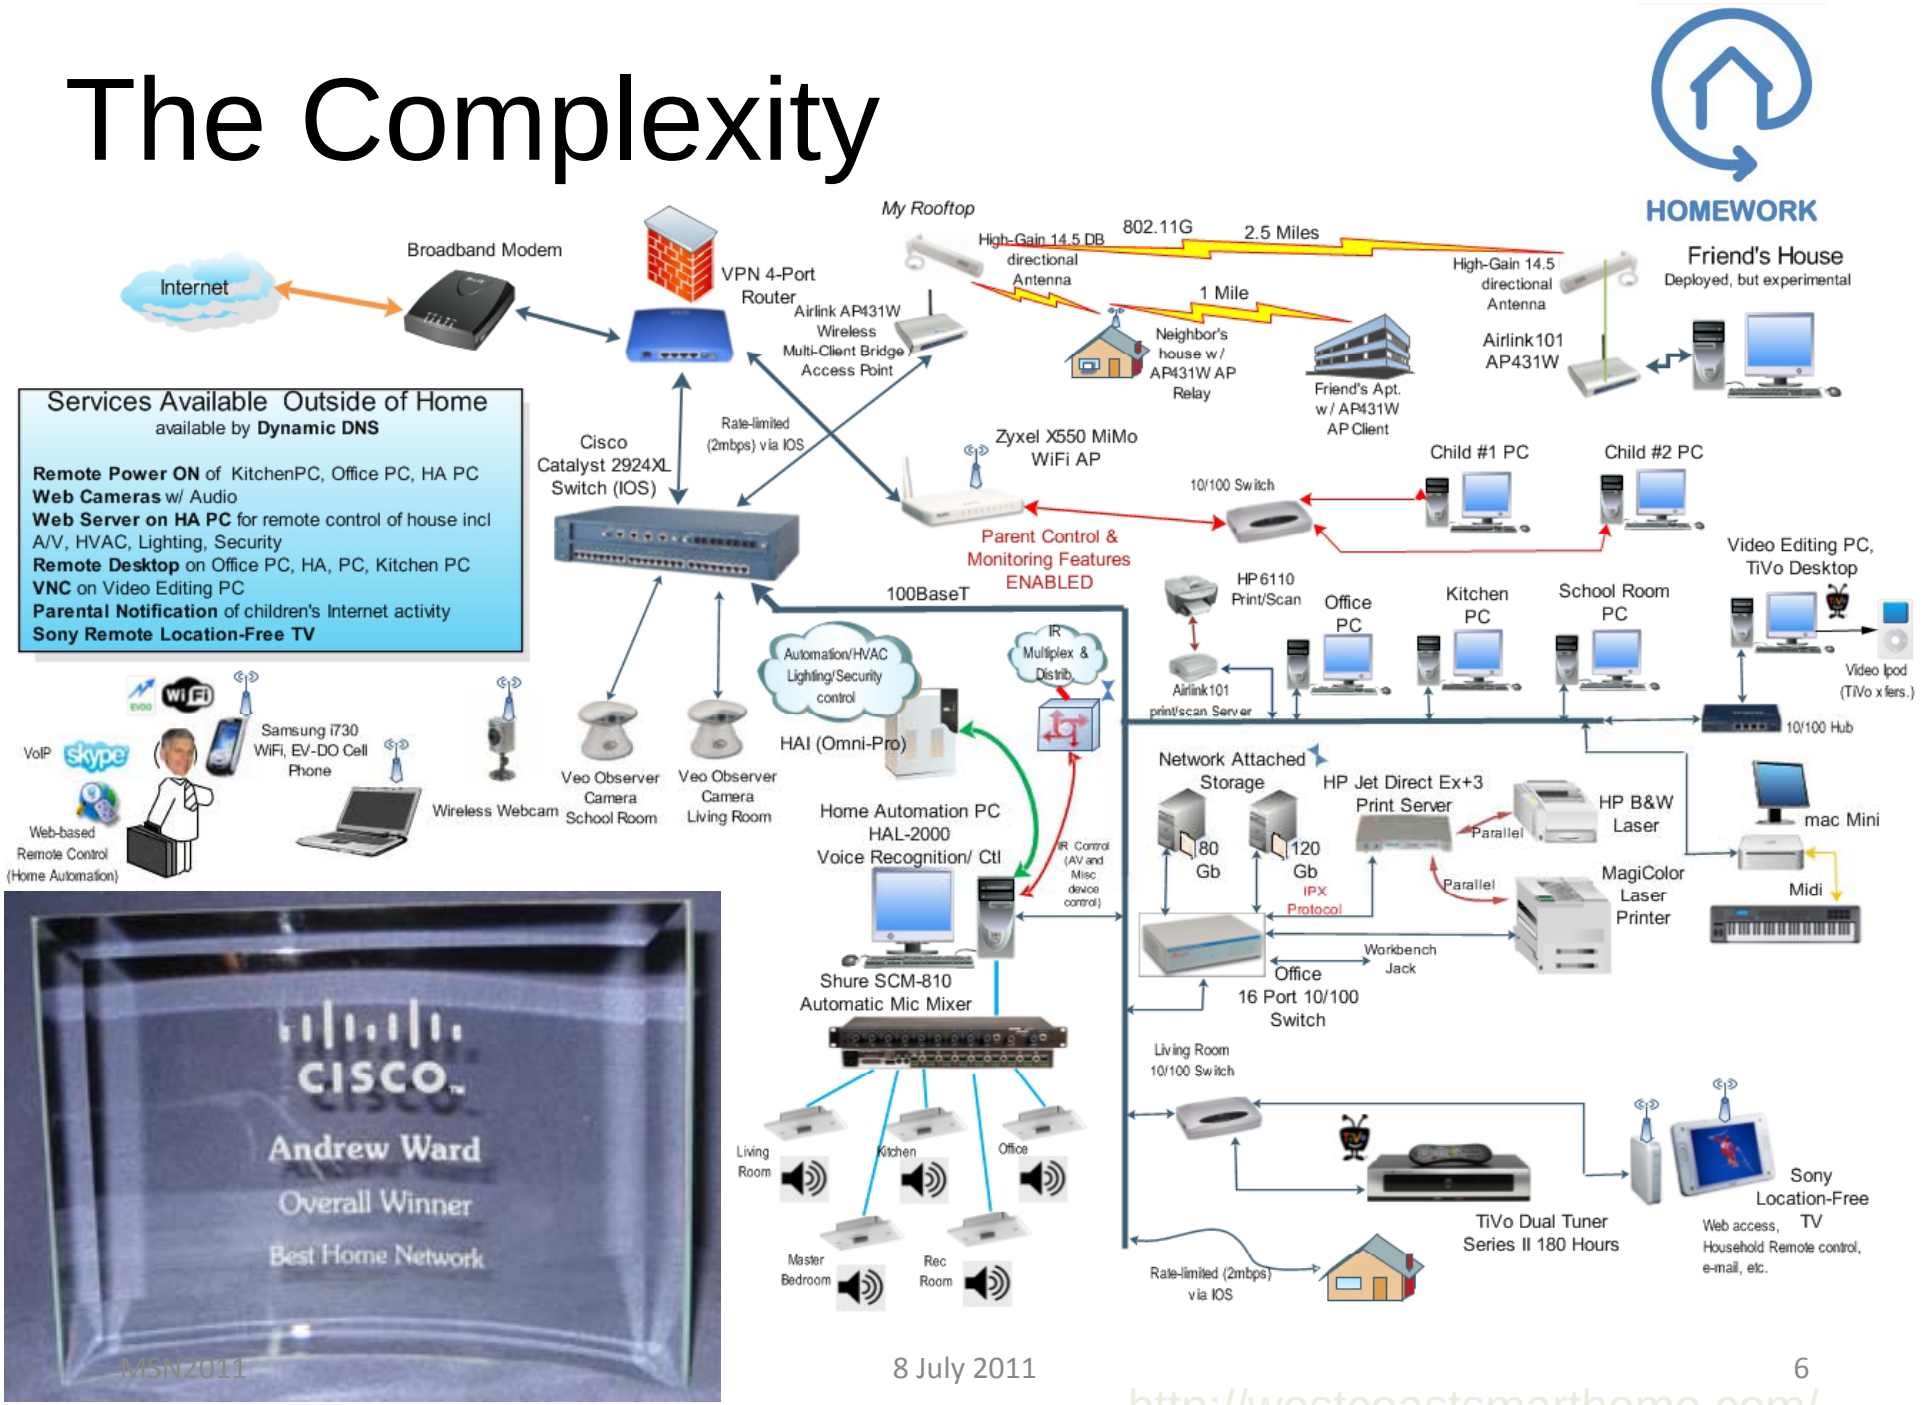

http://westcoastsmarthome.com/

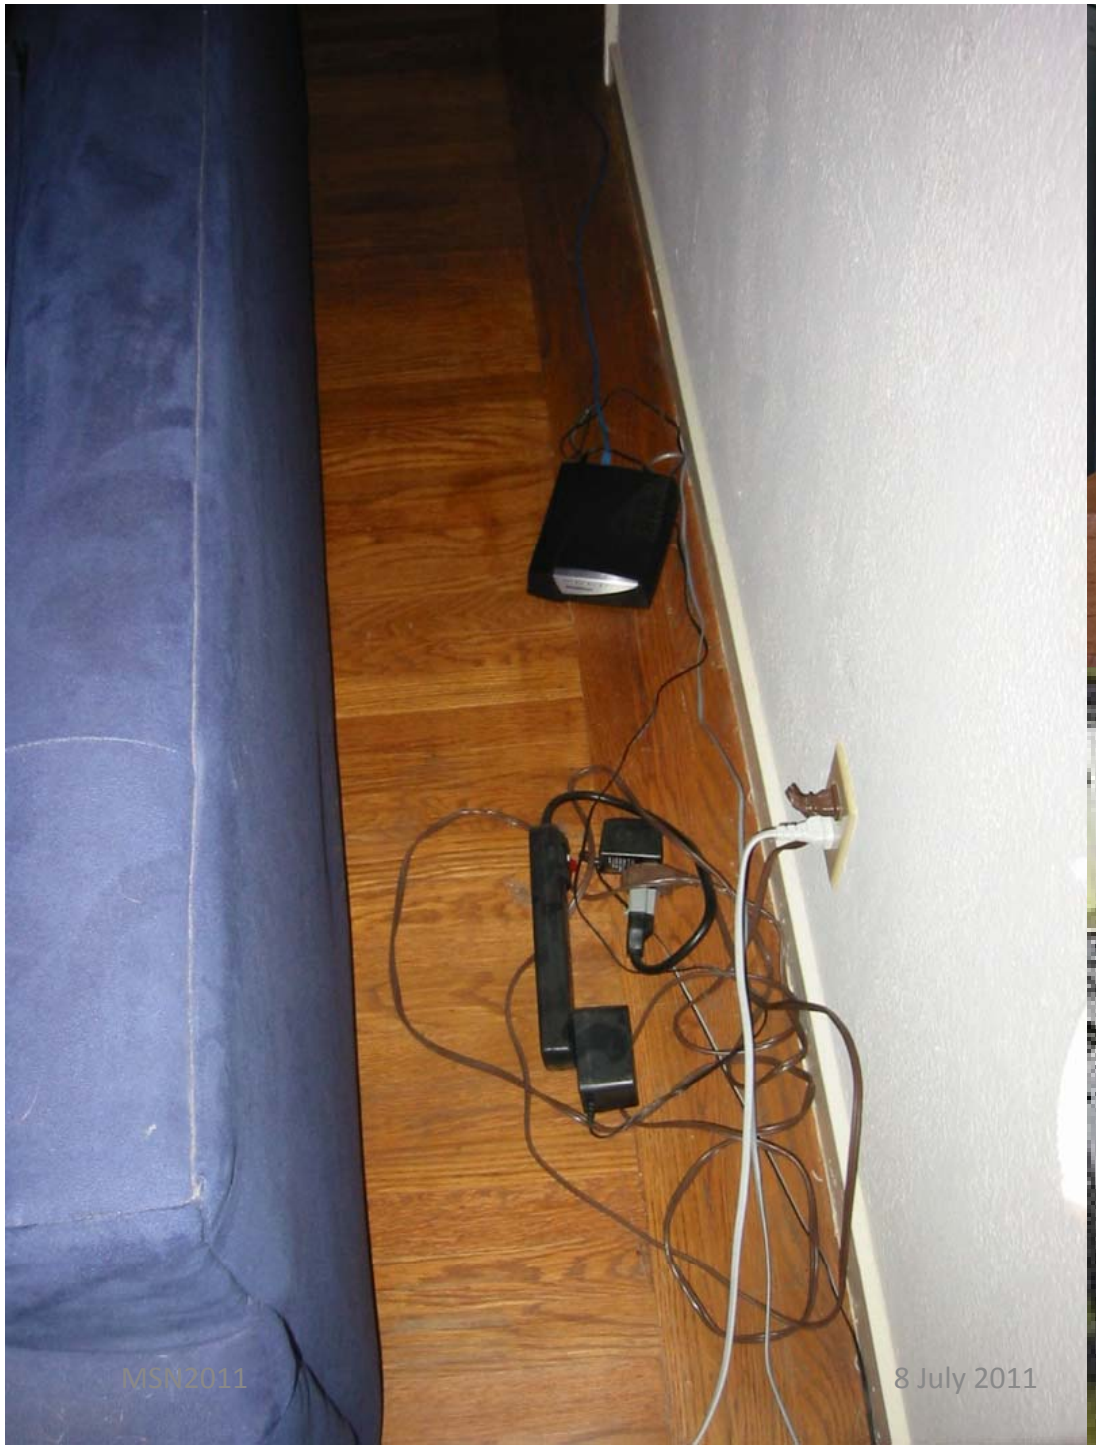


2.4GHz
802.11b
Power
ACT
LINK

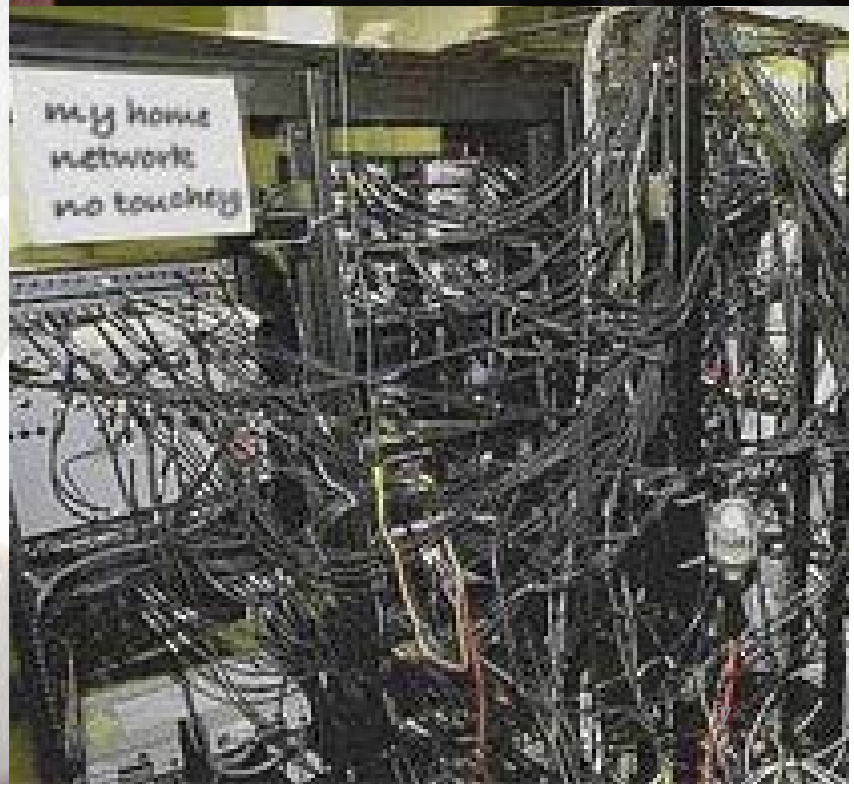
my home
network
no touchy

Technical Home User's View

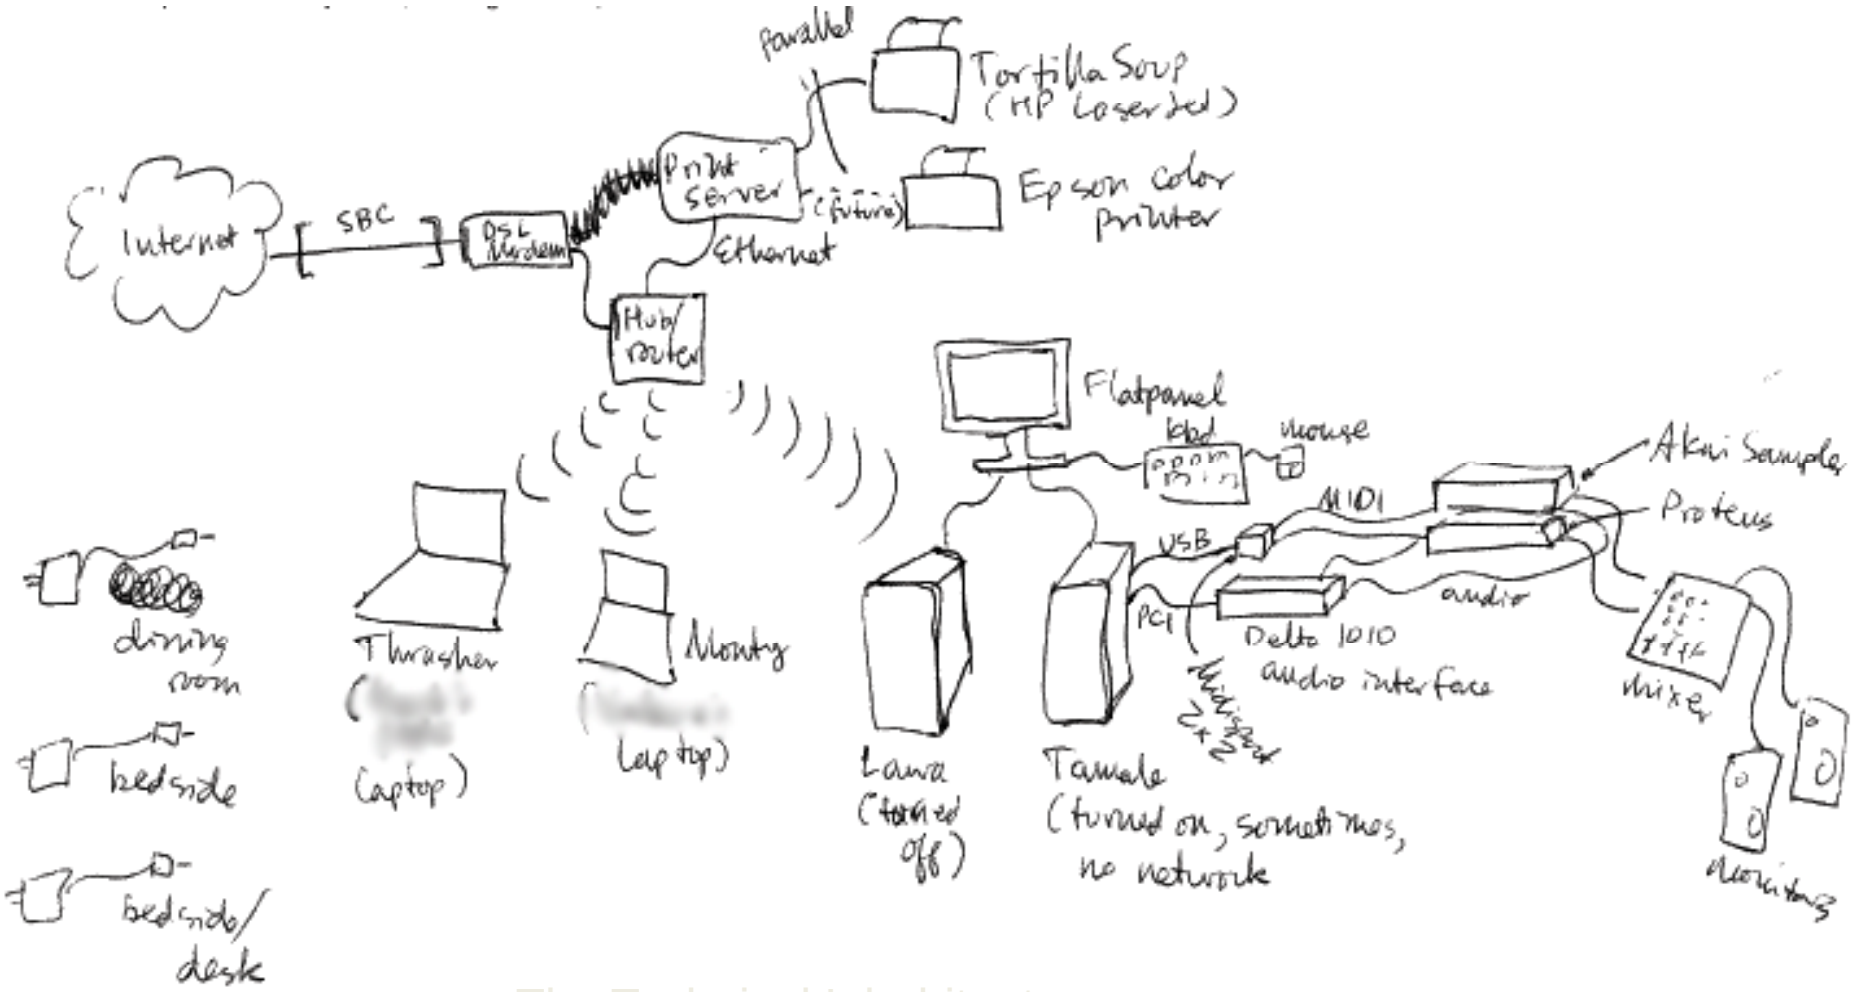

Another User's View

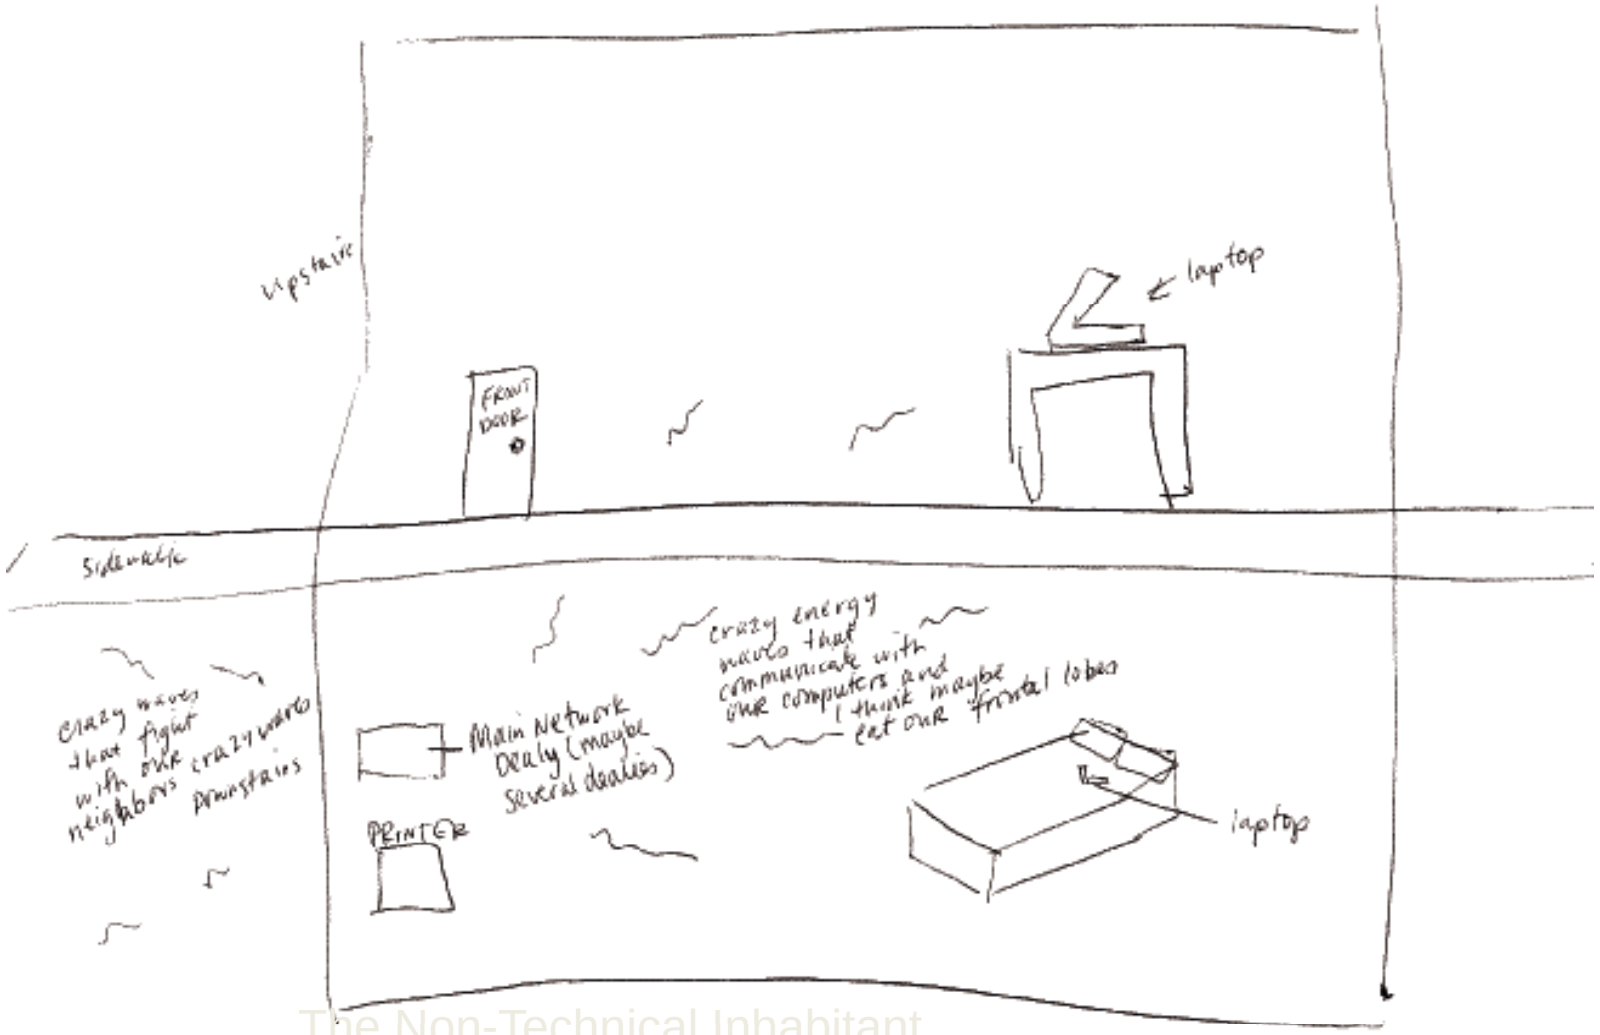

Nature of the Home

- Manageable Size
  - Human activities observable through empirical study
  - Passive local network measurement is possible in the router
  - Scale means we can explore different management models
  - Level of complexity allows us to exploit formal modelling
- Local Edge Network
  - Allows us to consider alternative approaches and architectures
  - Possible to explore localised management policies
- Physical Arrangement
  - Exploit local arrangements & local activities to support management
  - Possible to match human observation and network measures

Four broad desires of household residents

- Understand bandwidth use of the network in order to ***control consumption***
- Understand network activity in order to ***manage performance***
- Respond to demands by ***prioritization*** of network activities and interaction
- Possess systems to ***police the network***

"So if teacher says she's not doing homework and I think she's not doing it because she's spending her life on Facebook, I could block them and say 'I'll let you have them back when your homework's done'. You see that might be handy, if you could do it at the level of saying- You know, it's easy enough for me to say 'show your homework, right that's done', type something, 'right you can have Facebook now'"

Activities to police the network

- Network presentation – determine which devices are currently on the network and what they are doing.
- User notification - trigger requests for residents to intervene when particular activities are taking place that they would consider inappropriate.
- Access control - control access as a matter of principle or policy rather than in response to certain behaviours.
- **N.B. In most domestic settings, user notification followed by interpersonal interaction is the PREFERRED approach. Autonomic responses, specified by policies, are problematic, as such policies are likely to be fluid; ways must be found to enable people to create, manage and amend them without requiring a deep technical understanding**

Information Plane Architecture

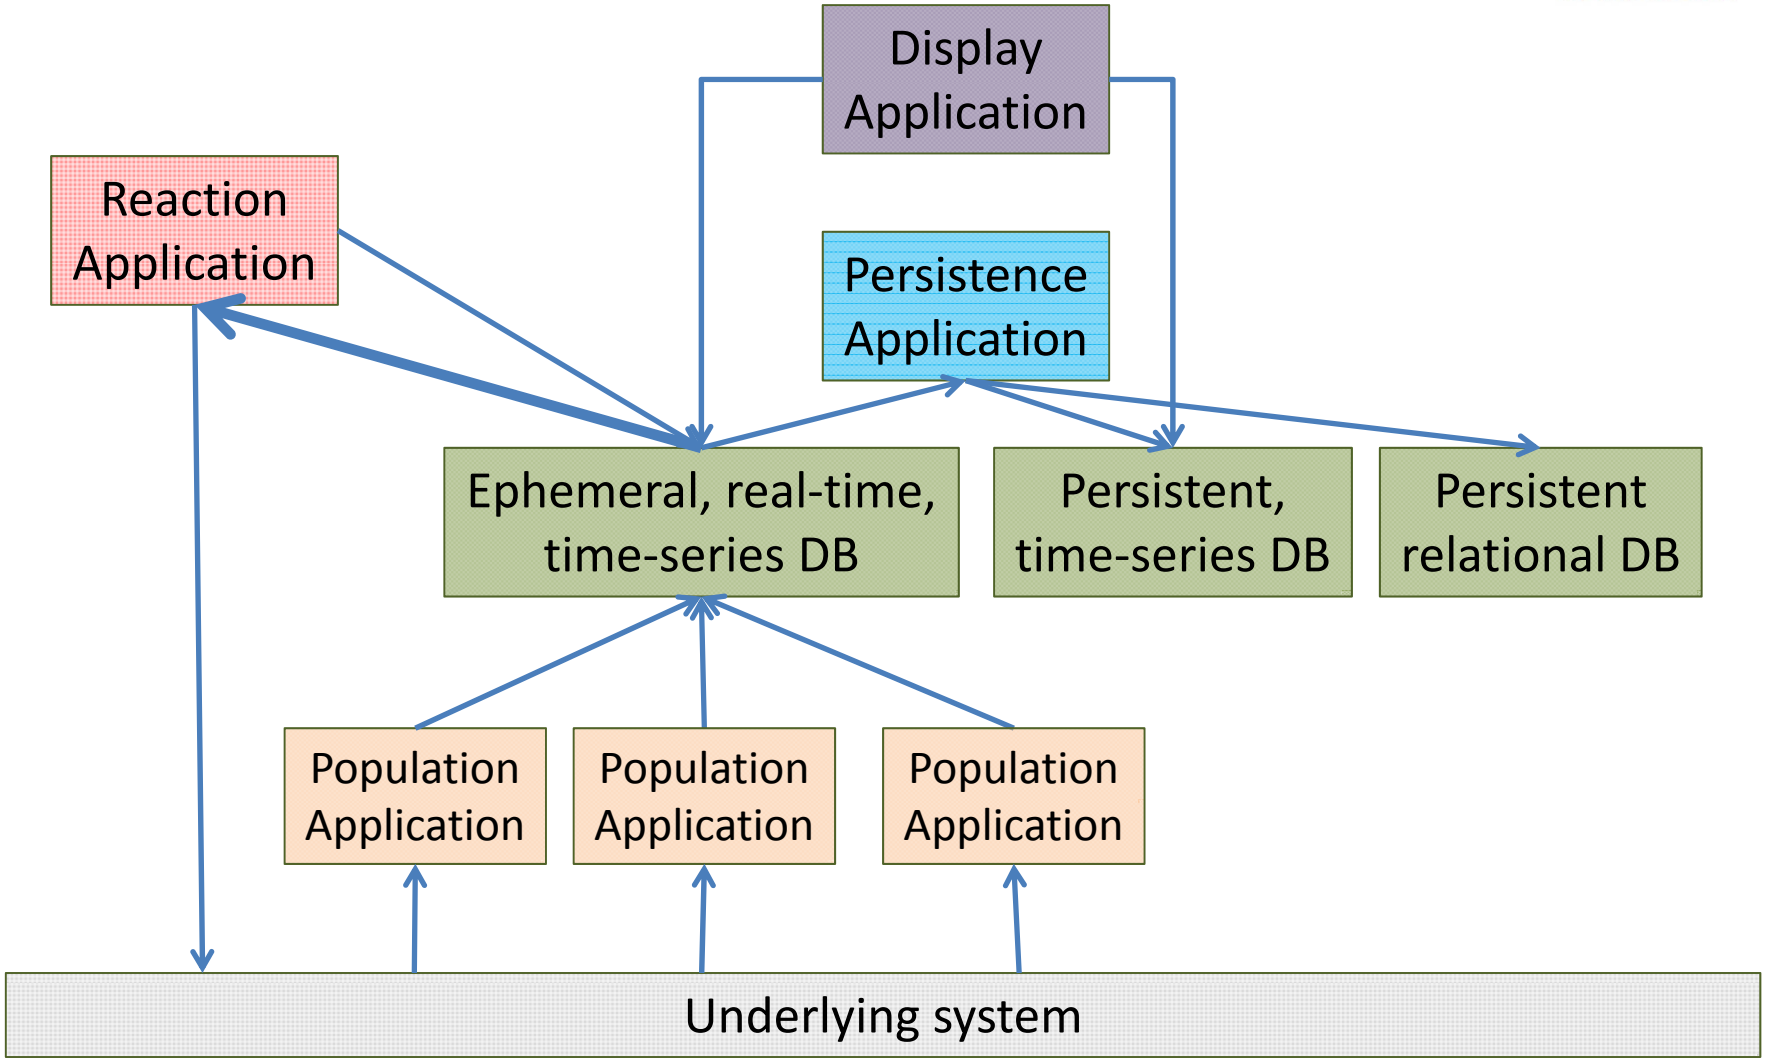

Key component of the information plane

- **Ephemeral, real-time, time-series database**
  - Ephemeral – continuous, large volume of measurements ⇒ cannot possibly make it persistent, so don't even try …
  - Real-time ⇒ must optimize use of resources to keep up with the measurements
  - Time-series ⇒ the primary ordering parameter for the measurement data is the time of occurrence
  - Technology of choice – stream database – enables live querying of recent data
  - Innovative approach – "raw" events are aggregated measures

Standard tables

Table	Attributes	Description
Flows	Protocol, src IP addr, src port, dst IP addr, dst port, # packets, #bytes	A tuple contains the number of packets and the number of bytes associated with a particular flow in the last second.
Links	MAC address, RSSI, #retries, #packets	A tuple contains the average received signal strength, the number of retries, and the number of packets associate with a particular MAC address in the last second
Leases	Action, MAC address, IP address, host name	A tuple denotes either that a DHCP lease has been granted to a particular host (action = “add”) or that a lease has been revoked (action = “del”)

The Homework Router

- Linux-based wireless router for deployment in users' homes
- openvswitch included in Linux kernel to enable interception of packets as they traverse the bridge
- Atom 1.6GHz EeePC 1000H netbook with 2GB of RAM running Ubuntu 10.04
- Ephemeral component runs as a process in the router
- Population, Persistence and Reaction components also run as processes in the router
- Display applications can be run on any device that is connected to the router, either directly over the wireless link, or through the backhaul network if the router's firewall rules enable such interaction.

Homework router (cont)

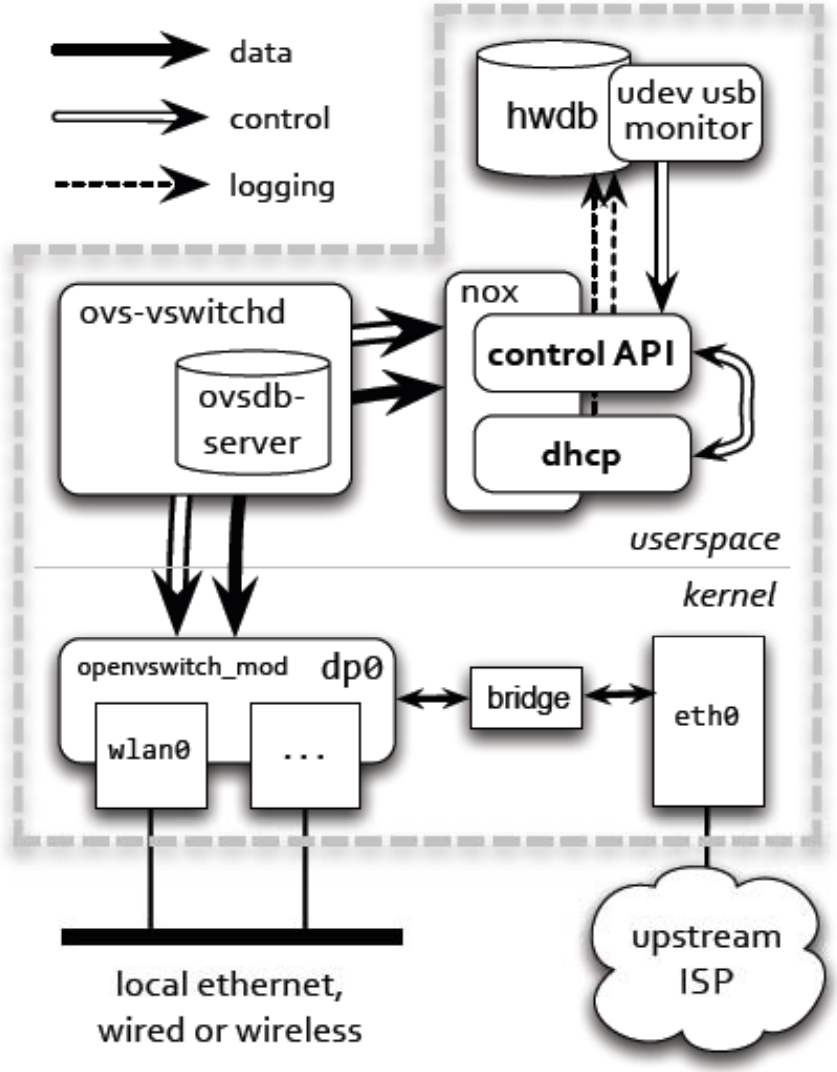

Raw event generation

- Link information obtained using libpcap (RadioTap)
- An additional action in openvswitch passes each packet to a kernel accumulator, which accumulates the following data:
  - Flow records
  - Data about the first N packets in each flow
  - For HTTP packets, the HTTP request header
- A once per second timer interrupt causes the kernel accumulator to write accumulated records to three different devices:
  - /dev/hwdb0 returns flow accumulations (to insert into table Flows)
  - /dev/hwdb1 has statistical information about the first N packets (currently, N = 10) of each flow
  - /dev/hwdb2 has http request headers
- Population applications simply have reads outstanding on these devices; when their reads are satisfied, they format insert commands into relevant tables and then call the Ephemeral component
- Lease information is inserted into hwdb by the DHCP module

Event systems

- Database systems are backward looking (i.e. we can query what has been accumulated to date.)
- Event systems are forward looking (look for particular patterns as events arrive)
- Languages are defined for event systems to enable the construction of automatons that look for particular patterns, and trigger actions when a pattern is matched

Database and Event Unification

- The schema for an HWDB table defines a typed event; the type is represented by the name of the Table
- *ad hoc* select queries can be made on the tuples currently stored in a table
- The table name also represents a stream of these typed events; each such stream can be referenced in an event-processing expression subscribed to the system; as a tuple is added to a table, any such expressions that refer to its associated stream are re-evaluated

CEP Infrastructure

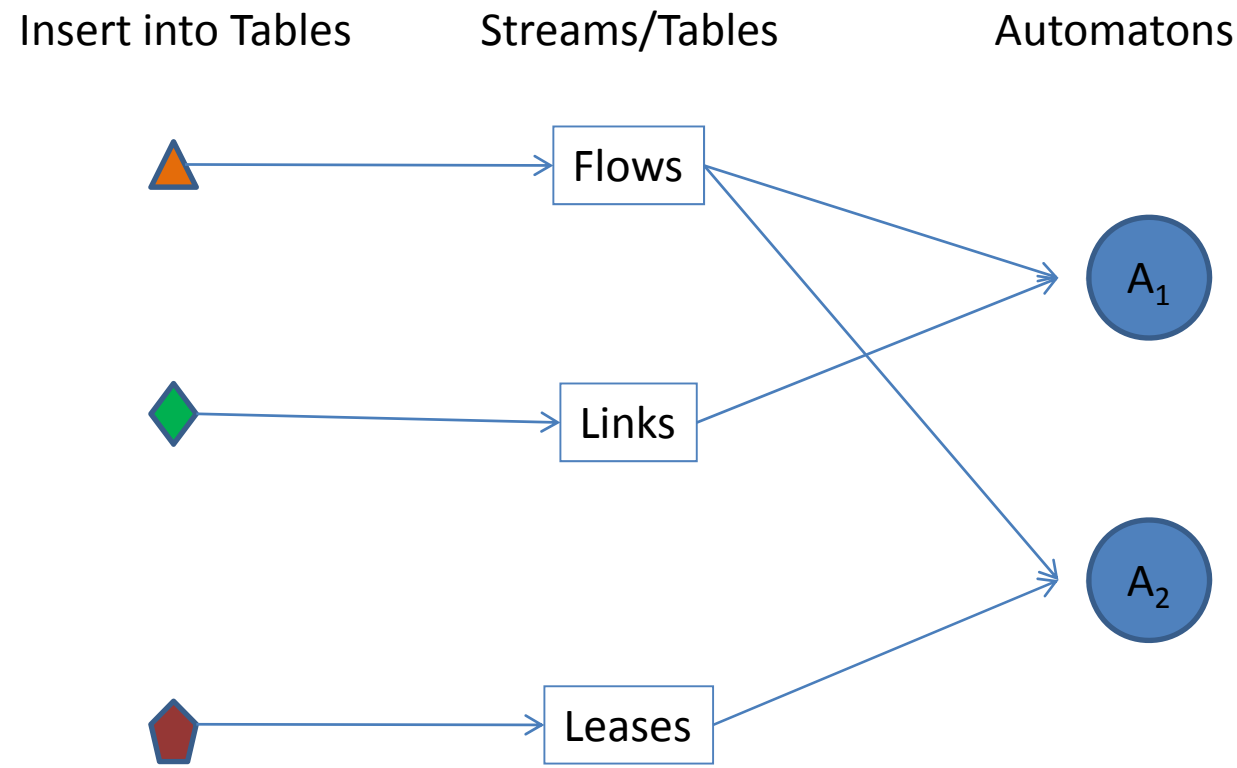

Automaton Example


```
subscribe f to Flows;
map[int,0] bytes, packets, limits;
int nbytes, npackets, bytelimit, hardlimit, limit;
ident key;
initialization {
hardlimit = 100000000; bytelimit = 80000000;
}
behavior {
   key = concat(f.srcadr, f.srcport, f.dstadr, f.dstport, f.proto);
   nbytes = lookup(bytes, key) + f.bytes;
   npackets = lookup(packets, key) + f.packets;
   limit = lookup(limits, key);
   if (limit <= 0) {
      limit = bytelimit; insert(limits, key, limit);
   }
   if (nbytes > limit) {
      send(f.srcadr, f.srcport, f.dstadr, f.dstport, f.proto, bytelimit, hardlimit);
      limit = limit + (hardlimit – limit) / 2;
      insert(limits, key, limit);
   }
   insert(bytes, key, nbytes);
   insert(packets, key, npackets);
}
```

Implementation

- Compiler generates instructions for stack machine
- Each compiled automaton is bound to a separate thread
- When a tuple is inserted into a Table, each automaton thread that has subscribed to that stream is given access to that tuple and awakened
- Upon being awakened, the automaton executes its behavior clause
- Aclient process subscribes an automaton to the database; if the automaton executes a “send” procedure call, this will result in the tuple specified in the send arguments being sent as an RPC to the subscribed process

Live bigraphs

- Bigraphical reactive systems can be used to model both spatial and temporal behaviour of network interactions
- Bigraphs can be used to detect inconsistent network configurations, and configurations that violate system invariants or user-specified access control policies
- Network topologies and constraints on their evolution (according to access control policies) are represented as bigraphs

Live bigraphs (2)

- Bigraphical reactive system models can be generated on-line from the current network topology and activated policies, as recorded in the information plane
- Feedback from real-time checking of these models can be used as events to drive the policy-based management system

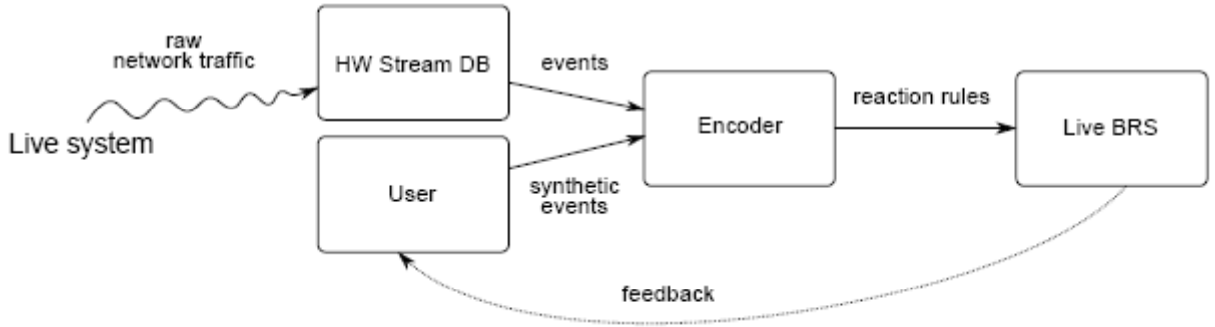

Network presentation

- Situated display – dedicated touch screen running HTML5 application
- Bandwidth contention – Objective C applications running on iPod Touch.

Network presentation: Example situated display

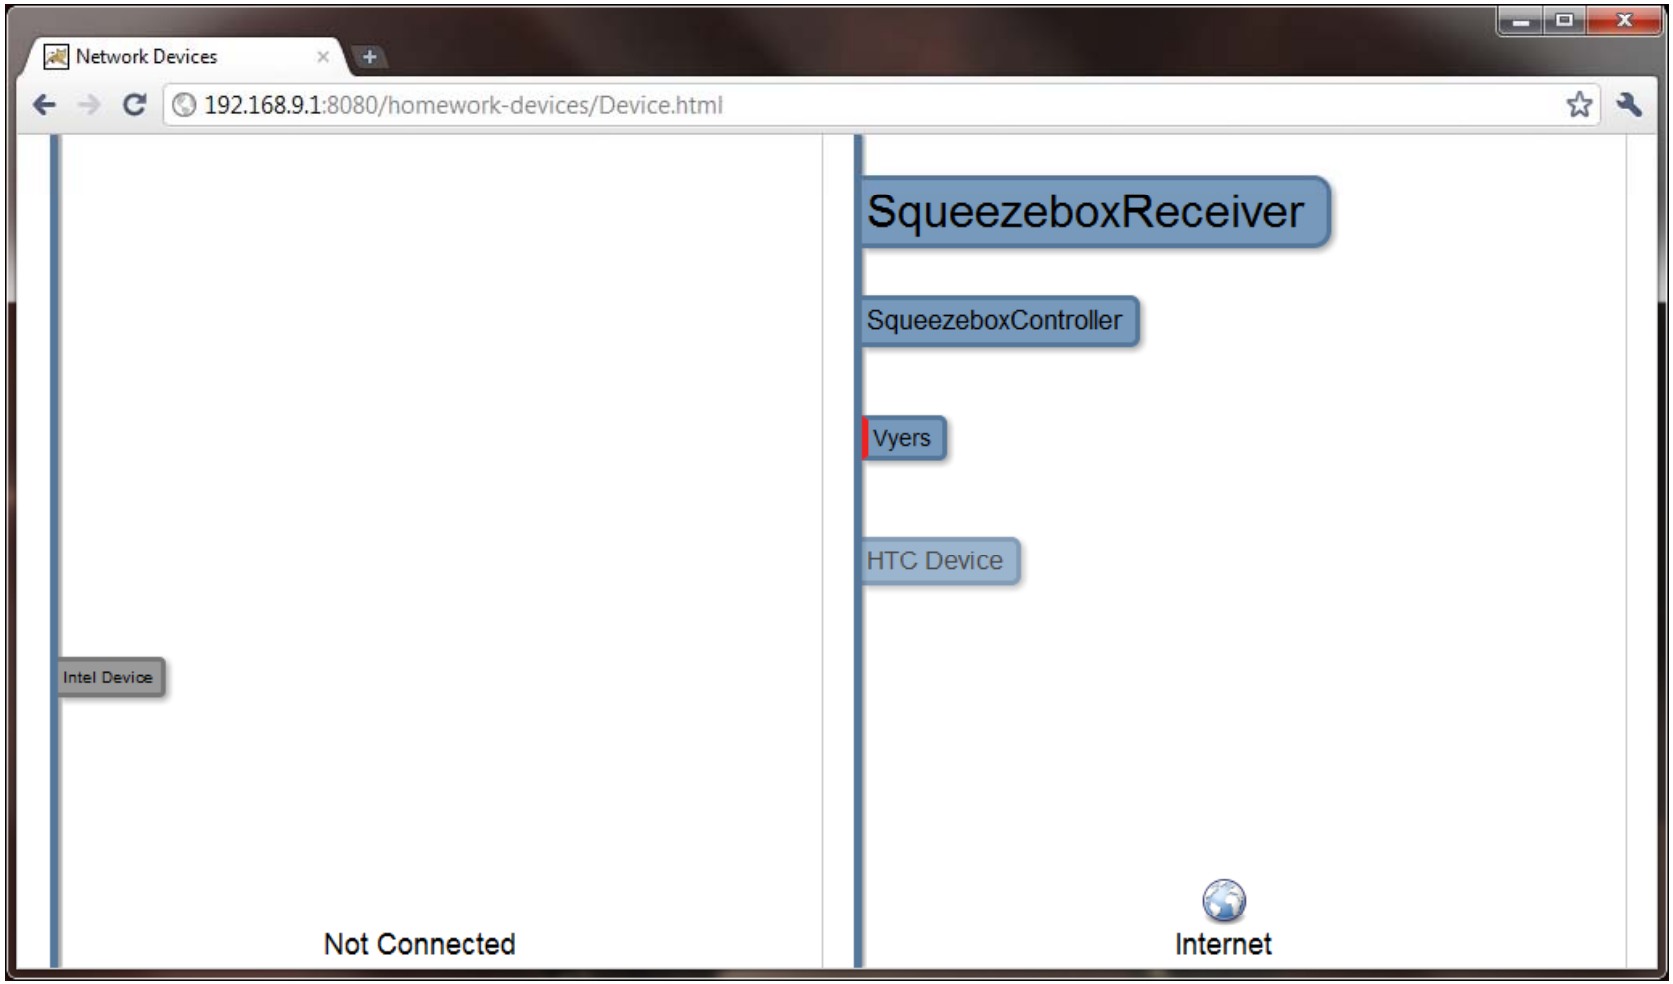

Network presentation:
Details of the example situated display

- Devices granted DHCP leases are shown on the RHS of the display
- Devices that do not have DHCP leases are shown on the LHS
- Display also maps key network characteristics of the devices to particular display features of the labels
  - Wireless signal strength is mapped to the transparency level of the machine label
  - The proportion of bandwidth usage for each machine is mapped to label sizes
  - Packet retries associated with each machine are displayed as red highlights on the label showing those machines that are experiencing traffic loss issues

Network presentation: Machine-specific traffic control

- Users can assert control over the traffic flowing through the home network using a simple drag and drop approach.
- Users can bar a particular machine from the network by dragging its icon from the right hand side of the display to the unconnected area on the left.
- Users can also affect particular equivalence classes of application protocols on a device. The control pane for the machine Vyers allows users to selectively turn off or throttle a number of these equivalence classes.

Network presentation: Bandwidth contention

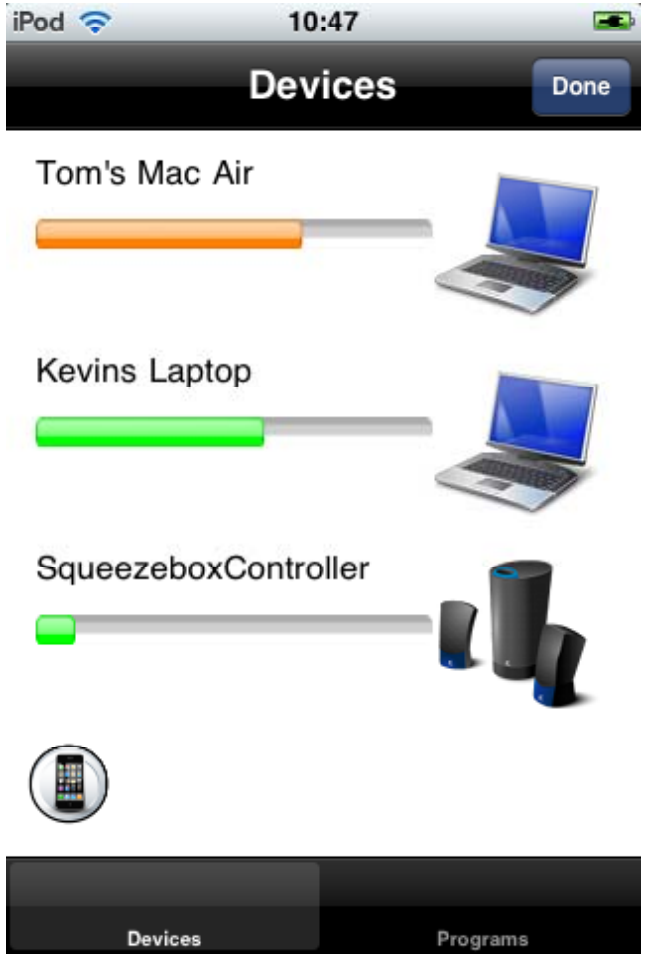

- Who's hogging the bandwidth – a recurring theme in our ethnographic studies
- Simply a display application that periodically retrieves flow tuples, aggregated over src and dst IP addresses, and displays the top N

Network presentation: Drilldown on a particular machine

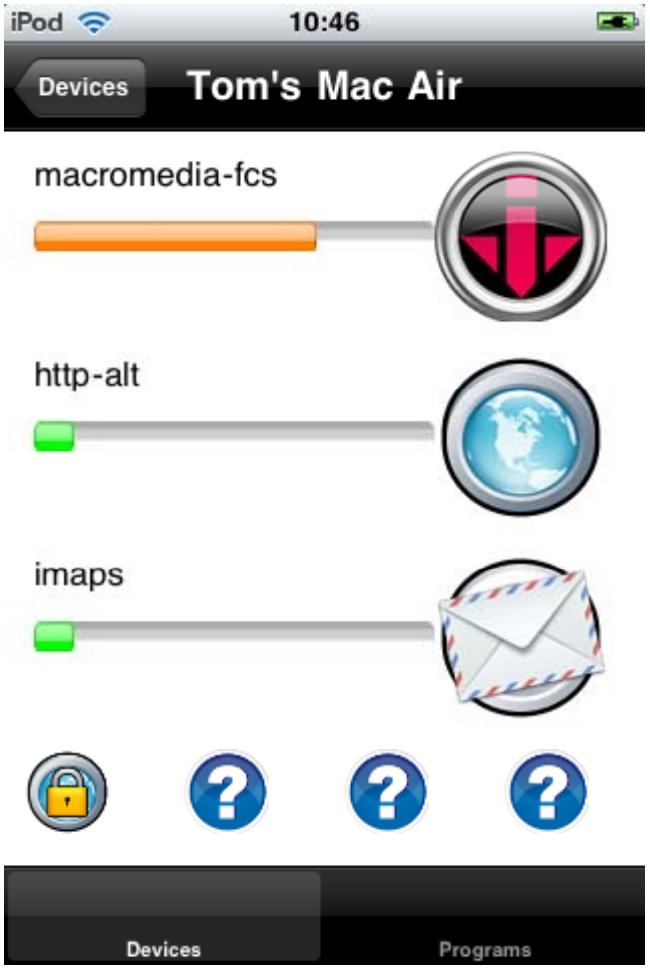

- Classify flows to protocol
- Simply a display application that periodically retrieves flow tuples from or to a particular IP address, and displays the top N
- Can select a particular flow, bringing up a slider display; the user can move the slider and confirm, causing appropriate rate control to be imposed on that flow in the router.

User notification

- User notification for policing purposes can take two general forms:
  - A pattern has been detected that indicates undesirable behaviour
  - Domestic networks require user involvement in what is traditionally an autonomic protocol
- As an example of the latter, consider DHCP – users expressed a desire to be able to ultimately decide if a device should be granted an IP address on the network
- This is an example of inserting people into a networking protocol; to do this, there must be a reasonable default condition if the user does not play his/her role in a timely fashion

User notification: Traditional DHCP

- Upon connection to the network, the client broadcasts a DHCPDISCOVER request to obtain the address of an available DHCP server
- A DHCP server responds with a DHCPOFFER containing, among other things, an IP address, a subnet mask and a gateway address
- The client selects a server and broadcasts a DHCPREQUEST so that all other servers know they have not been selected
- The selected server commits the configuration that it offered the client, and DHCPACKs the client to ensure the client has the correct configuration.

User notification: Extended DHCP

- The homeowner can control which devices are permitted to connect to the home network by interjecting in the protocol exchange on a case-by-case basis.
- We achieve this by manipulating the lease expiry time in our DHCPOFFERs, allocating only a short lease (30s) until the homeowner has permitted the device to connect via suitable user interfaces.
- The short leases ensure that clients will keep retrying until a decision is made; once a device is permitted to connect, we allocate a standard duration lease (1 hour).

User notification: Extended Situated Display

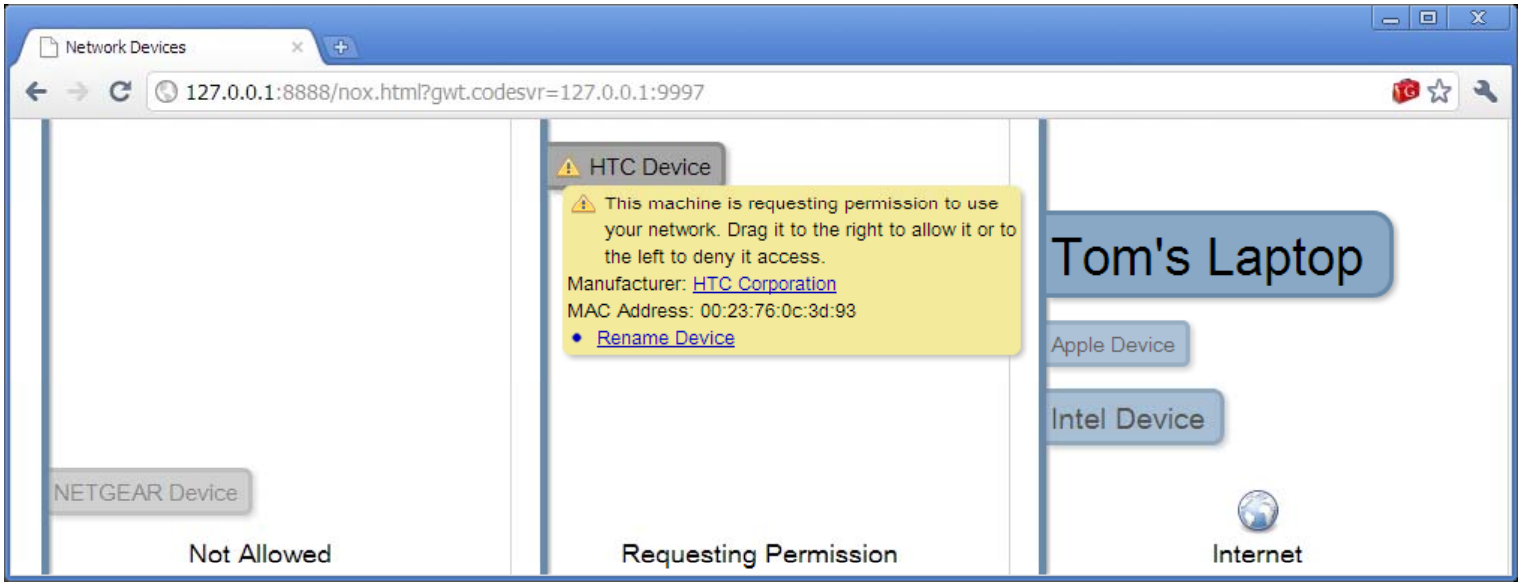

- While awaiting user permission, the device appears in the middle column.
- User grants permission by dragging the device into the right column; denies permission by dragging into the left column

Access control

- Authoring policy specifications requires significant technical understanding.
- Home owners possess neither the technical skill nor the motivation to fight with textually based policy specifications.
- Ways must be found that enable people to create, manage and amend policies without requiring a deep technical understanding.
- We are experimenting with the use of comic strip templates behind “wizard” technology to guide home users in the creation and amendment of policies.


iPad
16:19
43%
when Dad's...
...phone
is being used between these times
let Mum know
TWEET!
by..
the time
I'M WAITING...
the result!

Physically Mediated Access

- Seems a natural metaphor for homeowner/user
- Use USB storage keys with metadata in filesystem
- Can exploit both read- and write-ability of storage
  - Might be interesting to add computation in future
  - More relevant for public access networks
  - E.g., How can you trust the public interface?

In More Detail…

- udev rules fire events into hwdb on un/plug
  - Includes device name, e.g., /dev/sdb
- Handler process
  - Subscribed to hwdb for such events
  - Reacts by un/mounting filesystem on USB stick
- On mount
  - Reads directory names in /
  - Uses API to *permit* those matching eaddr regex
- Can subscribe to other hwdb queries, writing results into per-device directory
  - E.g., Per-device traffic modelling, both current and steady-state

Conclusions

- Domestic network environments are very different from corporate environments
- Home users possess neither the technical skill nor the motivation to fight with non-intuitive technology that does not take their needs into account
- The problem is more than just wrapping a GUI around the technology; the underlying assumptions of the technology are a million miles away from the expectations of home users
- Home users demand visibility of the network activity in the home network in order to address four different requirements: control consumption, manage bandwidth, prioritize network activities and police the network
- The dynamics in domestic environments depends heavily upon interpersonal interactions; autonomic actions, especially of a punitive nature, are always the last resort
- Alternative, visual methods for constructing and amending policy specifications are essential to enable home users to automate policies

Questions?