A Data Architecture for Consumer RFID Applications

Gaetano Borriello

Department of Computer Science & Engineering
University of Washington

(some material courtesy of C. Diorio,
UW/CSE Faculty and Founder of Impinj, Inc.)

Outline

- RFID Basics
- The UW RFID Ecosystem
- Our Initial Applications
- Data Architecture
- Future Work

Identifying Objects

- Use innate characteristics of the object
  - Faces, fingerprints, etc.
- Add markings to it
  - Engraving, bar codes
- Use non-visual modalities
  - RF, magnetic media

Can a specific instance be identified or just the class?

Basics

- Radio-frequency identification
  - Using radio frequency (RF) signals to identify (ID) an object
  - Does not require line-of-sight
- Tags are attached to an object
  - ID number in tag uniquely identify an object, not just its class
  - Current tags use 64 to 128 bits
  - Can include other information besides ID
    - Current state
    - Location
    - History

RFID Basics

- RFID systems comprise tags and readers
  - Tags are placed on objects
  - Readers interrogate tags
- Tags can be active or passive
  - Active tags: battery, expanded capability, longer range
  - Passive tag: receives power from RF field, limited capability

Active Tag

Passive Tag

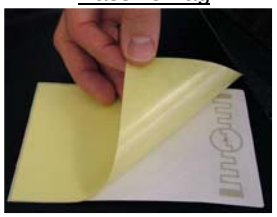

Communication between tag and reader

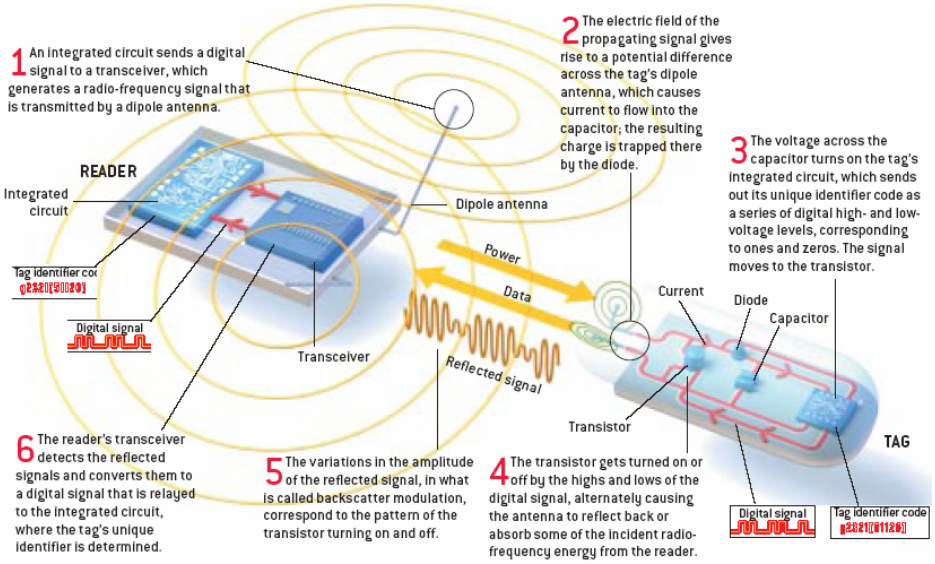

From Scientific American (1/2004)

Different types of tags for different uses

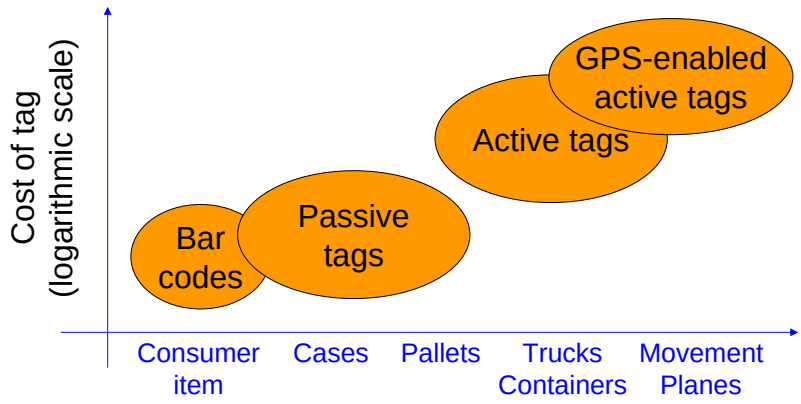

RFID versus Bar Codes

- RFID can replace bar codes
  - Don't need line-of-sight
  - Harder to spoof
  - Don't smudge
  - Tags can be rewriteable
  - Unique serial numbers
- EPC can replace UPC
  - 96 bit code gives 12,200,000,000,000,000,000 unique identifiers for every human being alive today

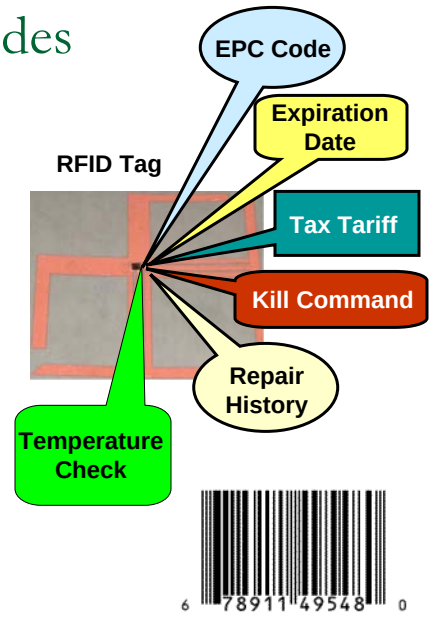

The elements of an RFID system

- Tags
  - Carries the ID number and very limited processing capability
- Readers
  - Tag communicates ID to “reader”
  - Readers emit RF and are regulated – differently around the world
- Networking infrastructure
  - Reader is connected to a network and communicates tag IDs to interested parties
- Databases
  - Collect the “read events” and log them with time/place
- Applications and their user interfaces
  - Browse the database looking for correlations and patterns

Applications in Supply-Chain

- Automation
  - RFID will boost throughput
    - Barcode scanning is slow
    - Barcode is applied manually
  - RFID = agile supply chain
    - Reduces out-of-stock
    - Reduces shrinkage
    - Increases inventory control
- Package tracking
- Airline tickets, luggage
- Pharmaceuticals
- Anti-counterfeiting
- Asset tagging, archiving

Market-Driven Technology

- 6 of top 7 retailers worldwide support RFID
  - > $1 trillion revenue
- 100s of manufacturers and retailers

10 Trillion Tags?

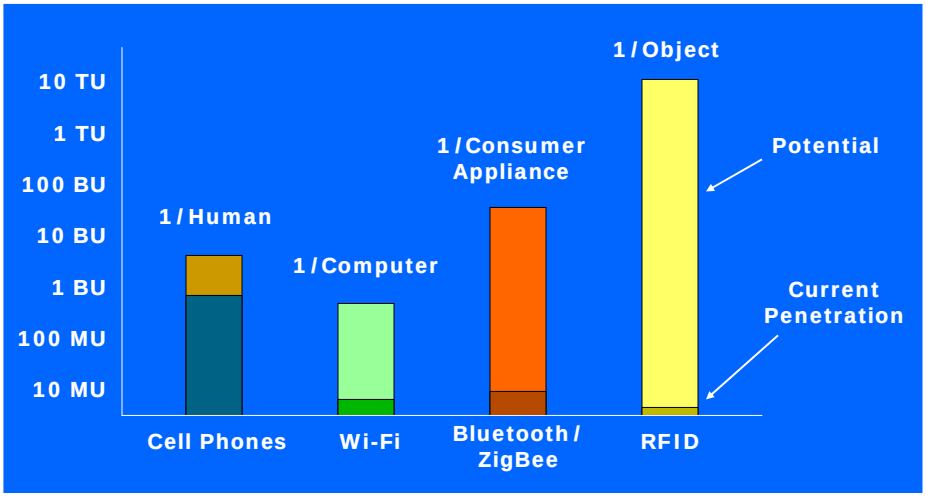

A wide variety of tags

Example: Impinj ZumaRFID™ Chip

- Field-rewritable tag
- 8m read range; 6m write range
- >500 tags/sec read rate; >15 tags/sec write rate
- Designed for dense-tag environments

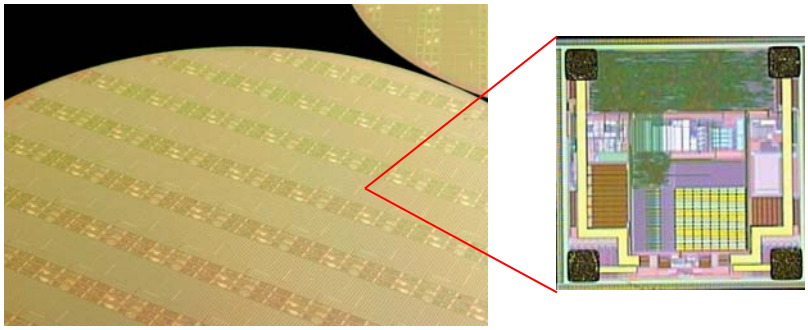

Example – Impinj ZumaRFID™ Tags

UW CSE RFID Ecosystem

- Create a microcosm of a world saturated with identifiable objects
- Explore applications, systems, and social implications
- Do it while there is still time to learn and adapt
- Use our building to do it

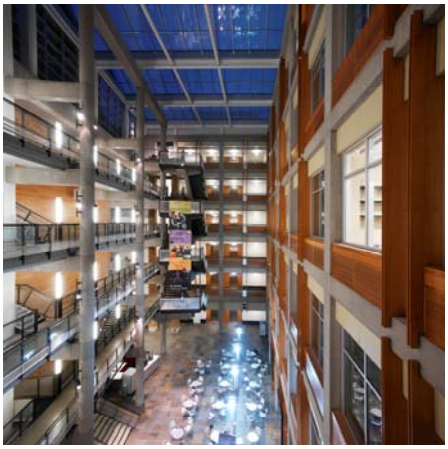

The opportunity

- Focus on supply chain has left other areas open
  - Will RFID tags have broader impact than simply replacing bar codes on more expensive items?
  - Will we be able to manage the data generated by all the readers that will be installed?
  - Can we devise valuable and/or interesting consumer applications through tagging of objects?
  - Can we architect systems to make it easier to safeguard privacy and address security concerns?

The research question

- What are the implications
  – for technology, business and society
  – of having a “number on everything”?
  - Merge physical and virtual worlds
  - Every object is an index into a world-wide database
  - Every object has its own history
  - Track object over its entire lifetime
  - Analyze trends in user habits
- The ecosystem is a way to get a view into this future

An RFID ecosystem

- CSE will equip the Allen Center building with RFID readers and its occupants with tags
- Sensing: Static readers, mobile tags (100s of readers)
  - Tags attached to objects of interest (computers, books, lab equipment, etc.)
  - Readers deployed at regular locations (hallway ends, etc.)
  - Mainly used for sensing and detection
- Actuation: Mobile readers, static tags (1 reader/person)
  - Tags attached to building locations of interest
    - offices, posters, etc.
  - Users equipped with phone-based NFC readers
  - Mainly used for actuation
- Goal: provide a living laboratory for interesting RFID systems research

Steps in rolling out the ecosystem

- 1. Testing of likely readers and tags

Testing issues

- Mounting of antenna
- Distance between antennas and tags
- Orientation of tag to antenna
- Material that tag is attached to
- Density of tags and material around them
- Rate of transit of tag past antenna

Steps in rolling out the ecosystem

- 2. Acquisition and installation (we are now in this phase)

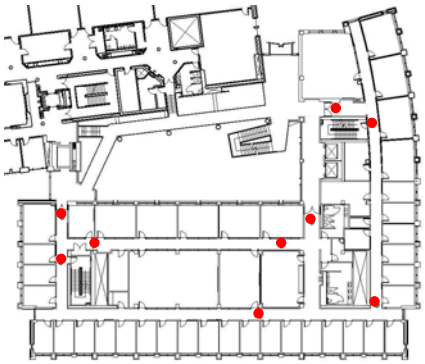

Installation

- FCC requires standoff to people (> 9in.)
  - Ceiling mounting, behind wall or glass, etc.
- Orientation of antennas to cover hallways
  - Multiple antennas at angles to each other
- Position in hallways
  - Try to identify a small enough region where tag is likely to be

Steps in rolling out the ecosystem

- 3. Develop data architecture (prototype in place)

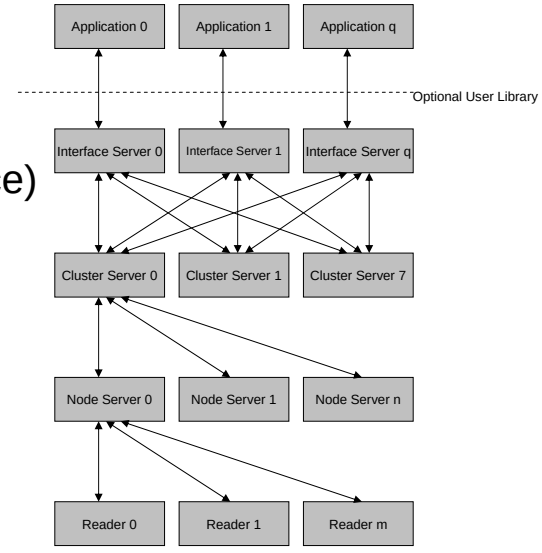

Data architecture issues

- Privacy, privacy, privacy
  - Lots of tags already exist that aren't ours
  - Revocability of data
  - Data mining
- Scaling to more readers and more apps
  - Don't tax basic infrastructure
  - Incremental addition of resources

Steps in rolling out the ecosystem

- 4. Deploy applications for study and use (we've done a couple so far)

Deployment issues

- Develop mainstay applications
  - Likely to always be needed
- Develop API for new experimental apps
  - Allow anyone to develop their own easily
- Privacy controls
  - Limiting what can be discovered and mined

Data architecture

- Three main concerns
  - Privacy including revocability
  - Scalability to more infrastructure and more apps
  - API for app development

What are we afraid of?

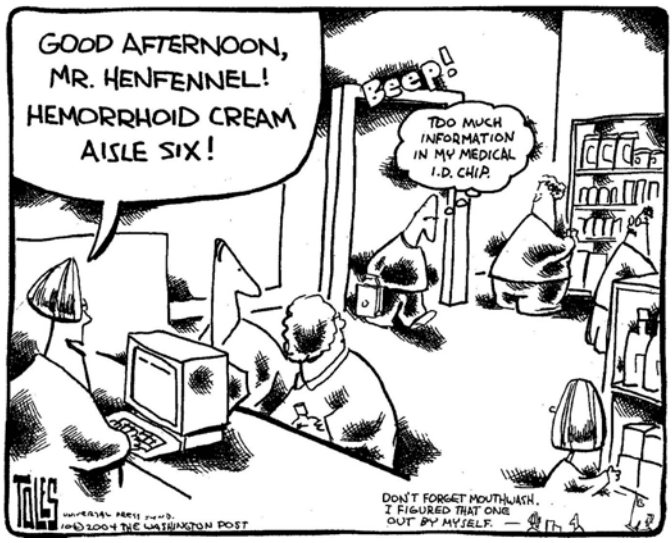

It is even more serious . . .

- Ability to track objects
- Ability to track people through their objects
- Ability to mine associations
  - People to objects
  - People to people
- Fears
  - Targetted advertising
  - "Big Brother" government
  - Personal security

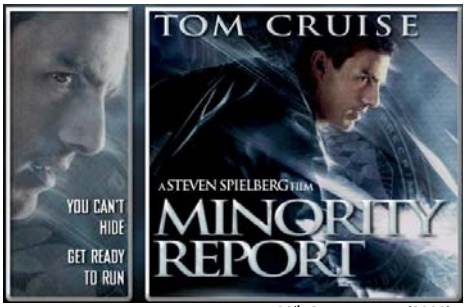

20th Century Fox (2002)

Data mining

- What are the implications of tagging everything?
  - Can use other data to link tag data
    - E.g., unknown tag passes by a reader,
      credit card transaction nearby,
      same tag seen again later,
      same credit card used nearby that location,
      tag may belong to that person,
      next time I see tag it implies that person is present
- Motivation for password protected tags and scramble tags (that change their ID)

Application: Object Tracker (finding lost objects, inventory control)

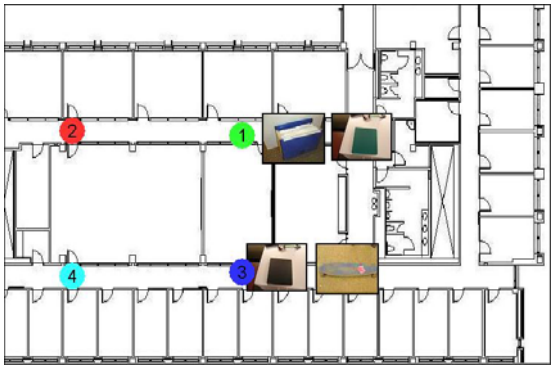

Application: Human Social Dynamics

- Study of evolution of human interactions
  - Paths walked, objects transferred, etc.

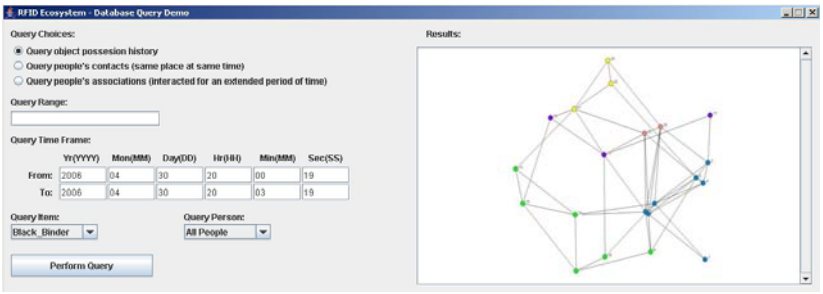

Application: Reminder System

tags on objects and server on wrist

user passes by a reader antenna on wall

alarm if object that should be present is not

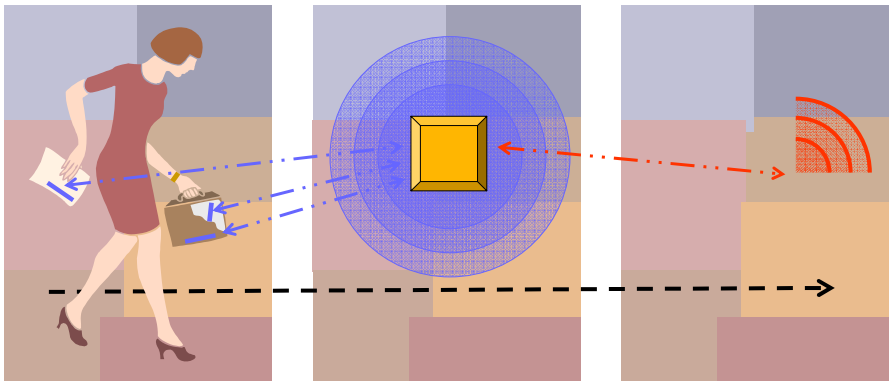

Application: Woodland Park Zoo

- Track visitors to zoo
- Pattern through zoo
- How groups break up and merge
- Attractor exhibits (where people go first)
- Personalized web pages for visit and other items of possible interest – extend visit

Data architecture

- Four layers
  - Node servers
  - Cluster servers
  - Interface servers
  - Application servers

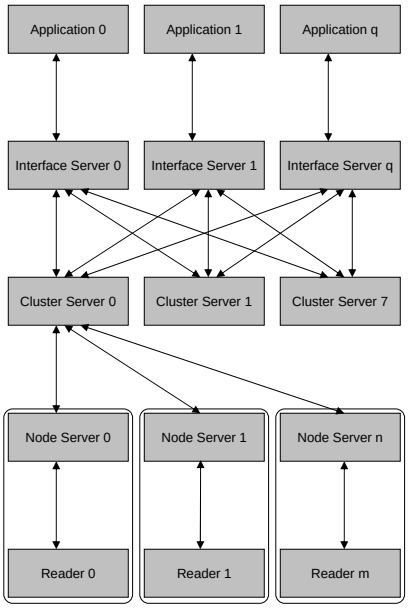

Node servers

- Reside inside the reader hardware
  - Our readers run Linux and have code easily added
  - Can broadcast tag reads in the vicinity of the reader
  - Interacts with permissions server to determine if tags are registered
    - Caches tagID permissions for efficiency
- Tag's privacy registration determines whether:
  - Tag can be routed to the cluster server for permanent storage
  - Tag can be broadcast locally
- Streams "storable" tag reads to Cluster Server (each node server is paired with one cluster server), and buffers if Cluster Server unavailable

Privacy

- Users must explicitly authorize each application to access data generated from one of their tags
- Centralized model
  - Tag ID changes at each Ecosystem layer
    - System uses multiple encrypted/scrambled IDs to create a blind system so that the information from various parts of the system can not be combined
  - Node servers encrypt all storable tags using ecosystem key
  - Cluster servers encrypt tag events using application's public key
- Broadcast model
  - When broadcasting a TagID, a random number is appended to the TagID and encrypted with the TagID itself for broadcast, only a user that knows the TagID can make sense of the packet

Cluster Server

- Routes tag reads from node servers to interface servers for enabled applications
- Stores all incoming tag reads to local DB for archiving
- Periodically checks its node servers to make sure they are still running

Interface Server

- Each app gets its own ‘client’ of the RFID ecosystem
  - Provisioned by whoever deploys app
- Events are generated on this server based on the application’s setting of event-generation parameters
- Compute higher-level events based on raw tag read data from cluster servers, thereby removing computation load off the centralized cluster servers
- Maintain connections with applications
- Stream events to applications
- Respond to event queries by applications (possibly on archived data)

Application Servers

- Actually run application logic and user interfaces
- Could be same physical server as their own interface server
- Provisioned by whoever deploys app

Reasons for the layered design

- Compartmentalize reader features (node server)
- Separate privacy concerns into a permission server
- Make the system scalable by not trying to do to much in a standard way (use interface servers to generate specialized events for apps)
- Retain control of data stores in cluster servers
- Make it easy to incorporate a data stream processing system e.g., Borealis (at the cluster server)

Interface Server: Event Hierarchy

- Interface servers provide a hierarchical event structure
  - Uses low-level events to generate higher-level events
  - Starting from tag reads to whatever is defined by applications
  - Some will be fairly common and are pre-defined
- Applications are given additional tag meta data
  - Tag owner
  - Item of attachment

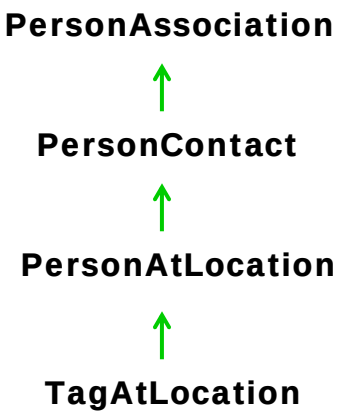

Events

- Each event type generated by an interface server can be adjusted by the application via interface server parameters
  - Events can be turned on/off
  - Even parameters can be changed, e.g., different timeout for defining the presence of a tag
  - New events can be created, e.g., tag followed a path through building

Administrative Tools

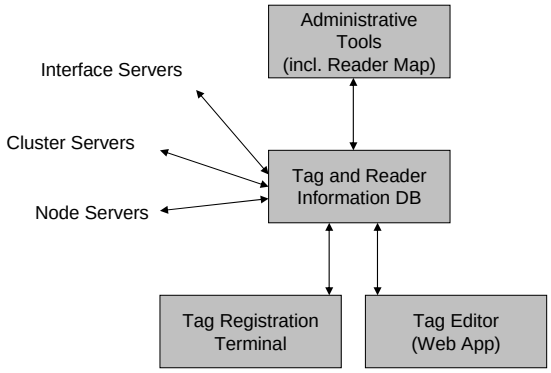

Registration Terminal

- All ecosystem RFID tags must be registered (with owner and item they are attached to)
  - All unregistered tags are ignored by the system

Tag Information Editor

- All users can edit the privacy information associated with their tag
- Storable/broadcastable
- List of trusted applications that can use the tag – others never see it
- Standard templates

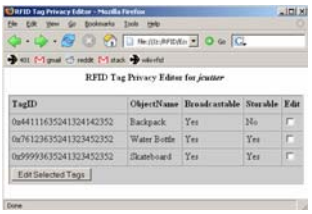

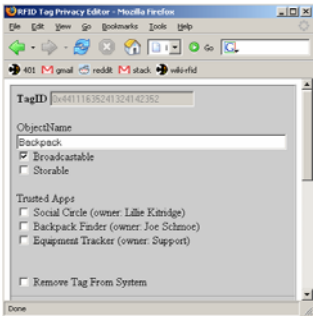

Reader Administration Tool

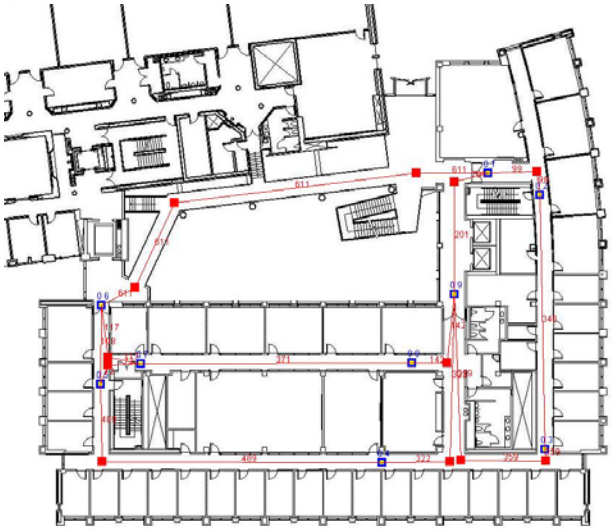

Robustness

- All servers in the system actively work to re-establish lost connections.
- All data received from each input server is buffered until output server connection is re-established.
- All input servers are determined at runtime from a database. Additional servers can be added on the fly.

Audit Trails

- When tag data is deleted – some events will also need to be deleted
  - Any event that was determined from the data being deleted should also be removed
- Event hierarchy helps maintain relationships between events
- E.g., instead of generating a four-way contact we generate 6 two-way contacts
  - Makes it easy to delete data from the system (even for one tag)
  - Only have to delete dependent events and not generate replacement events

Future work

- Tags are getting more interesting all the time
  - Tags with sensors – motion, temperature, etc.
- Tag have increasing memory
  - Database can be distributed into tags
- Tag data security
  - Passwords, registrations – who keeps it all straight
- Study of real application deployments
  - Can users manage their tags? What are the right default behaviors?

Application: Elder Care (w/ Intel Seattle)

- What objects people use is a good indicator of what they are doing
- Track objects using WISP tags
- Use to infer activities of daily living (ADLs)
- Trend analysis – mix/duration of activities

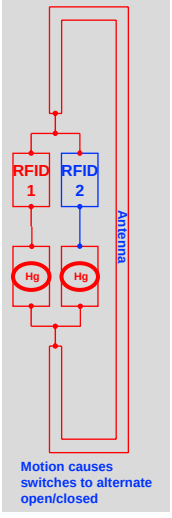

Thank you for this opportunity!