UNIVERSITY OF CAMBRIDGE

From Panopticon to Fresnel Dispelling A False Sense of Security

Jon Crowcroft,

http://www.cl.cam.ac.uk/~jac22

With Ian Brown, OII

Introduction

- Marconi Prof@U o Cambridge Computer Laboratory
- Spent 22 years at University College London (CS)
  - Home of Jeremy Bentham (creator of panopticon)
  - Presented my first paper here@UKC in 1982 (t - 30 years:)
- Topic - ubiquity of sensors & privacy
- Sub topics -
  - differential privacy - limitations
  - Privacy of graph data - limitations
  - privacy by design
- Thanks to Siani Pearson of HP for Monday's Keynote
  - provides perfect background, tutorial & definitions!

UNIVERSITY OF CAMBRIDGE

What are you doing?

UNIVERSITY OF CAMBRIDGE

Where are you now

UNIVERSITY OF CAMBRIDGE

Bringing it all back home

Total Situational Awareness

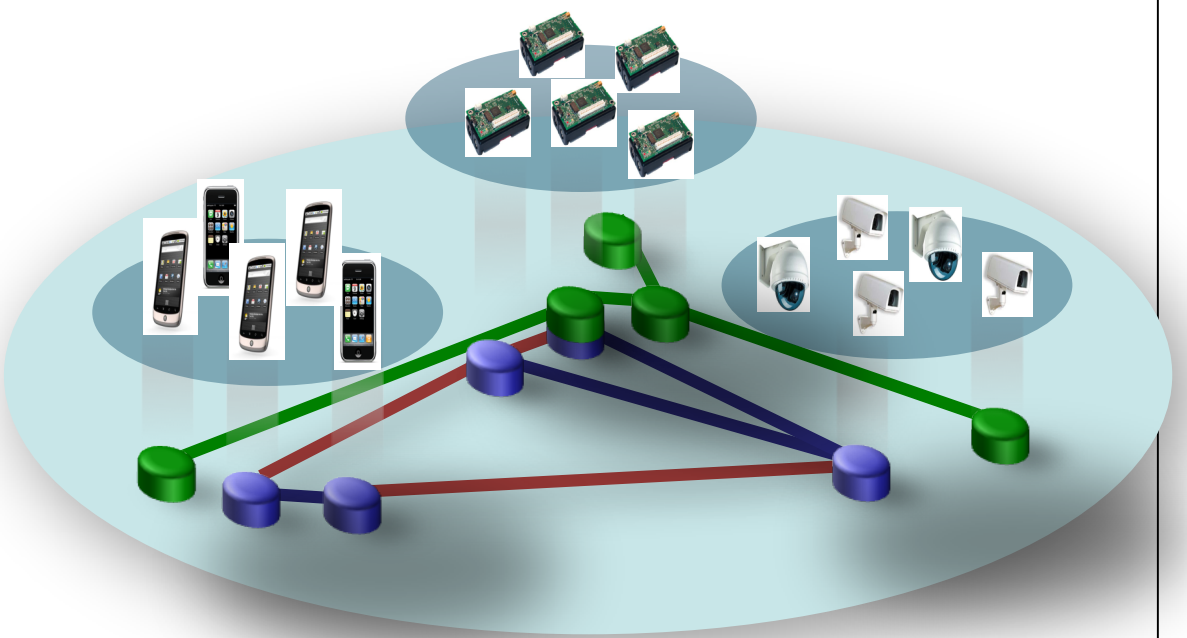

Federating Sensor Nets

UNIVERSITY OF CAMBRIDGE

- Fresnel is an EPSRC funded project
- To Federate Sensor Nets
- But provide isolation and privacy
- And tools (intellectual and technical) for
- ***Privacy by design***

Fresnel

UNIVERSITY OF CAMBRIDGE

Isolation

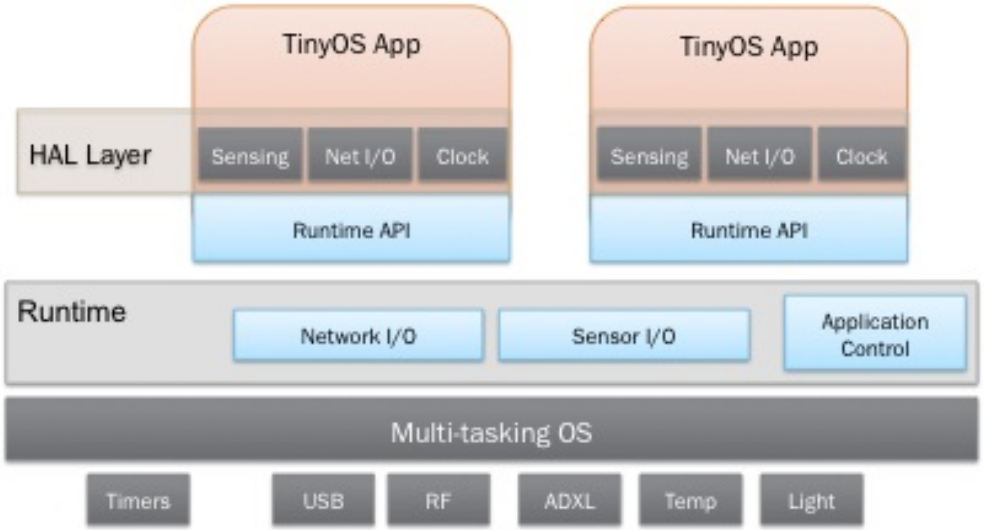

Fresnel Net

UNIVERSITY OF CAMBRIDGE

Virtualisation

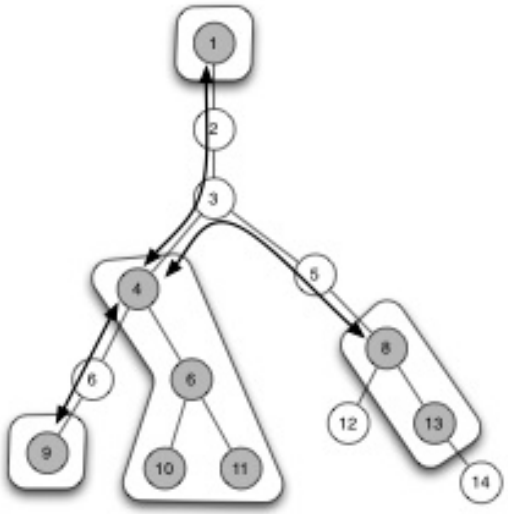

(a)

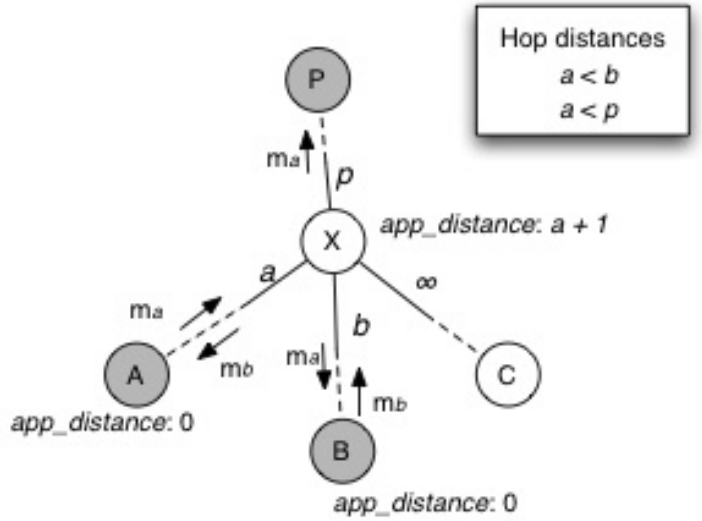

(b)

Fresnel App

UNIVERSITY OF CAMBRIDGE

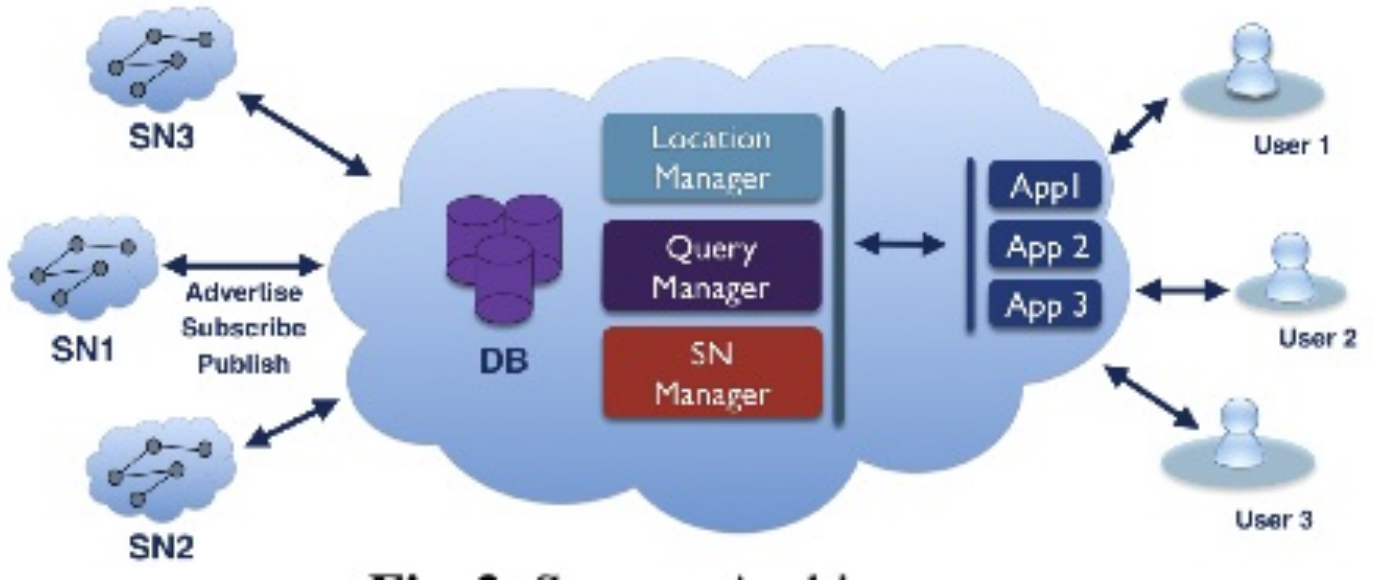

Deployment

UNIVERSITY OF CAMBRIDGE

- 34 Nodes (29 permanent)
- 10 Rooms
- iMote2 Sensors attached to desks
- Two applications
  - Desk occupancy
  - Environmental monitoring

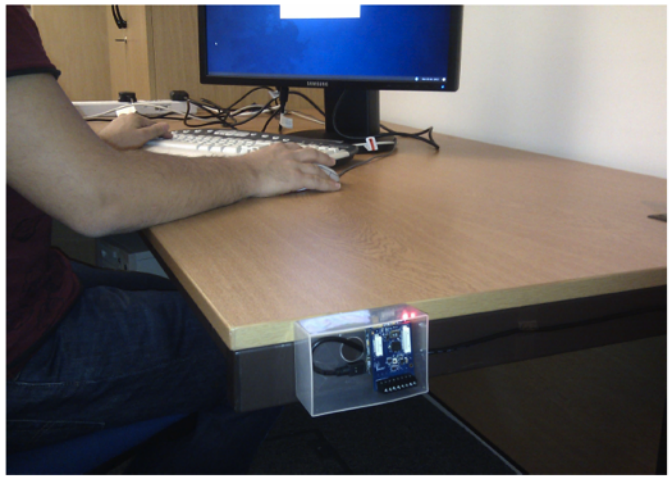

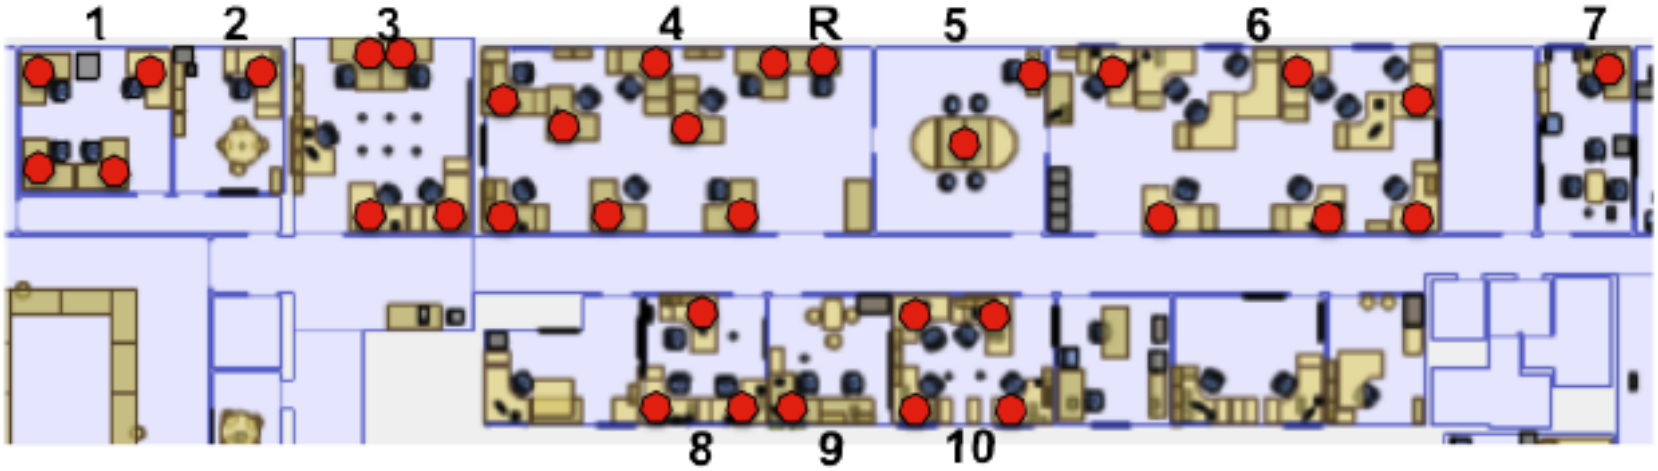

UNIVERSITY OF CAMBRIDGE

Deployment Results

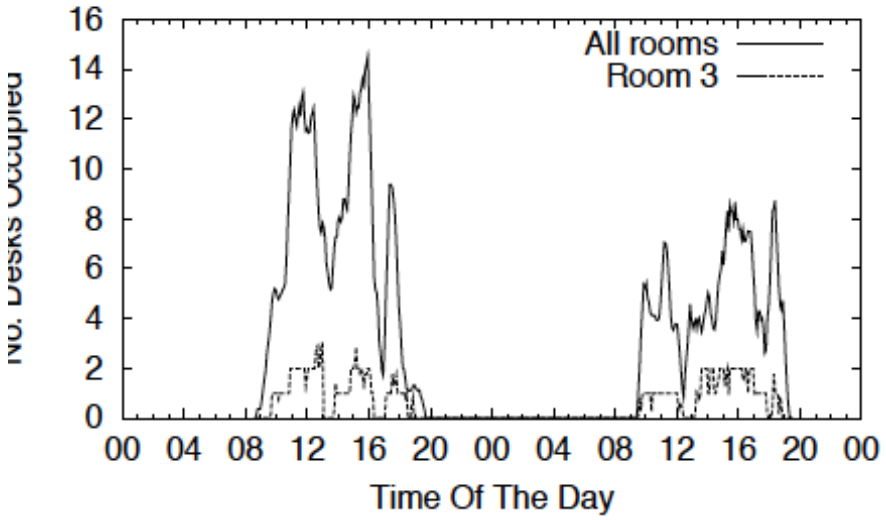

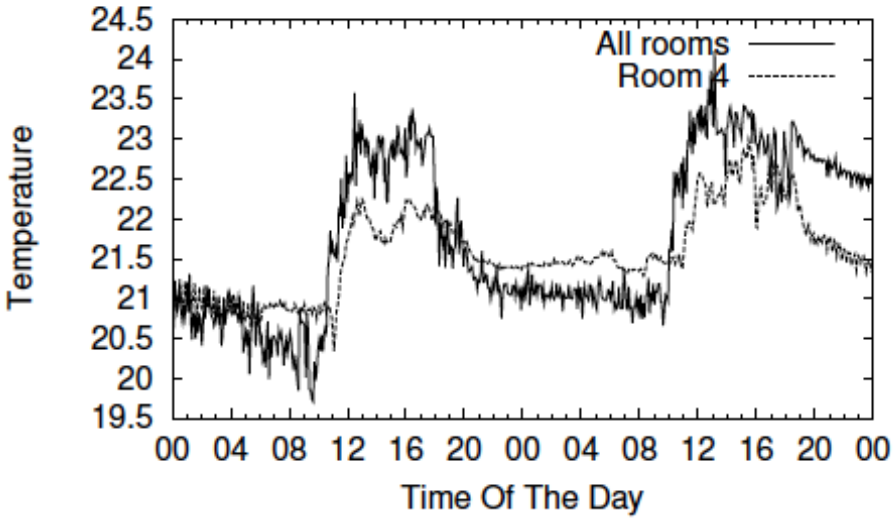

Fresnel Use

UNIVERSITY OF CAMBRIDGE

We can see you

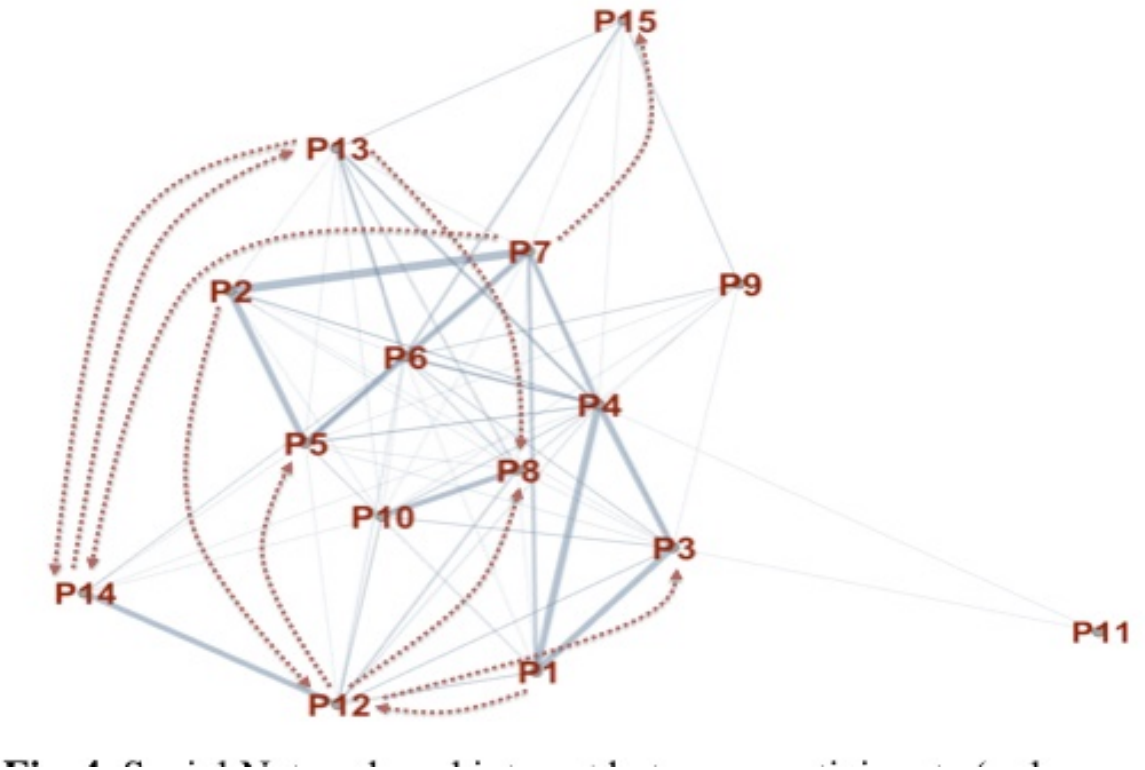

Fresnel Worry

UNIVERSITY OF CAMBRIDGE

Pan-panopticon

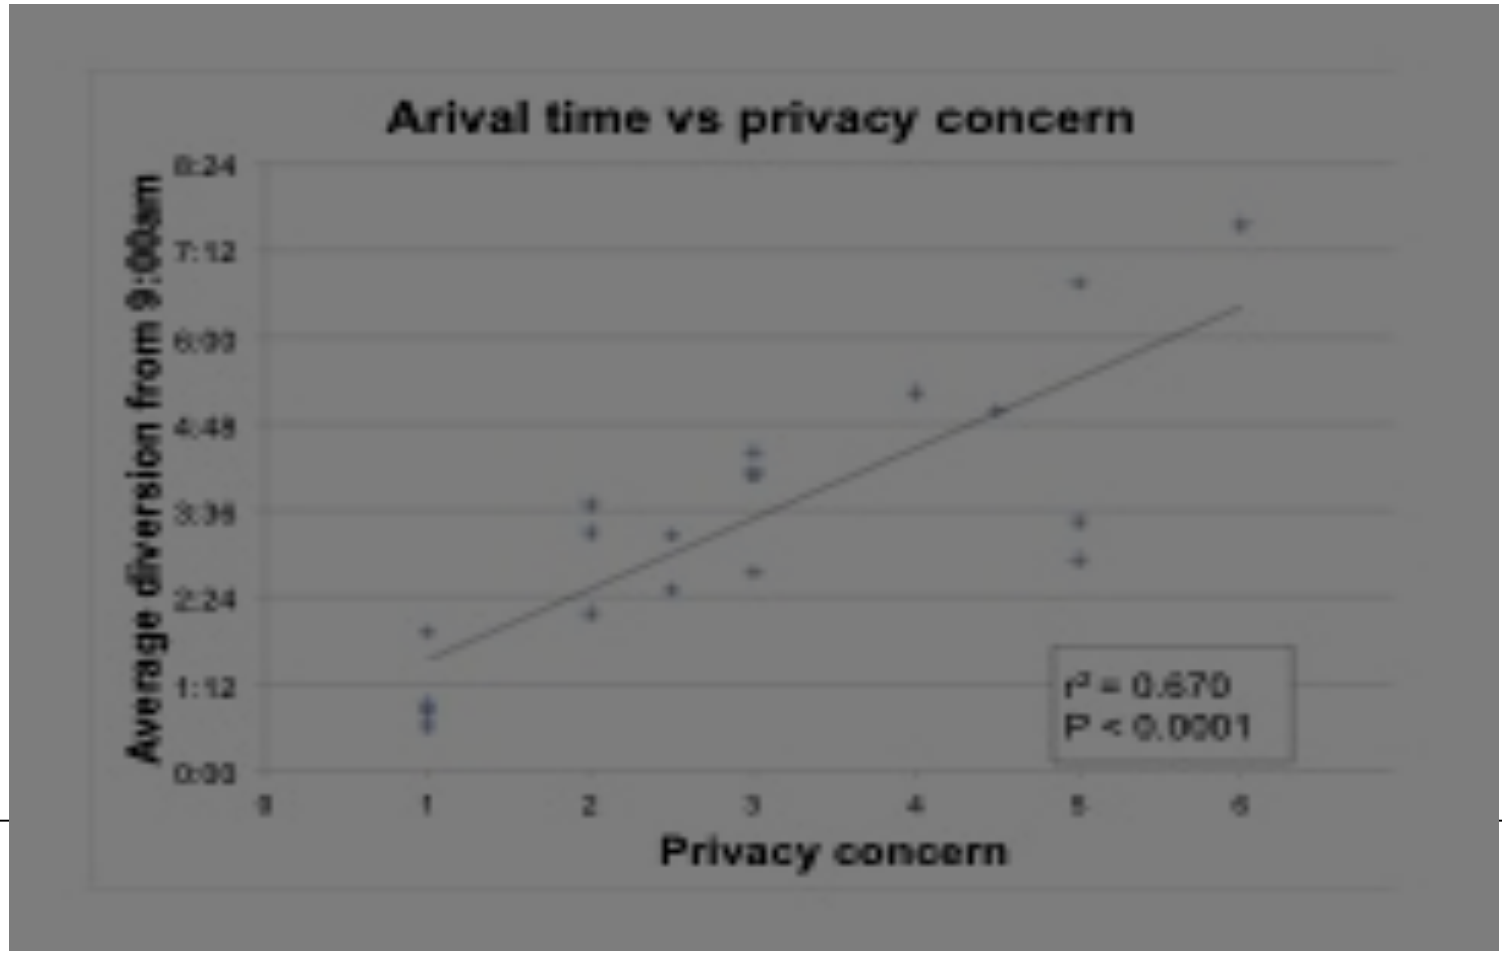

Take Homes

UNIVERSITY OF CAMBRIDGE

- Why am I telling you about this?
- Be afraid…you are’nt just under the fresnel lens
- It’s a Pan-panopticon…be very afraid :-(

Aggregation of Sources. problem

UNIVERSITY OF CAMBRIDGE

- But what about the data, once its bought back to the cloud?
- Well, then we need different tools - which is what we talk about next…

Through a Graph, Darkly

UNIVERSITY OF CAMBRIDGE

- A Manifesto for Personal Privacy
- But still allowing data driven research
- And evidence based policy
- And all that targetted advertising:-)
- N.B. Of particular concern in UK right now, as

1. Open Data Initiatives (including NHS)
2. ICO is pushing for naïve anonymisation!!
3. Communications Capability Dev. Programme

    Aka Communications Data Bill is a v. bad idea … … …

Sharing Data

UNIVERSITY OF CAMBRIDGE

- We’re increasingly asked to share data
  - Government Open Data
  - Nature June article
  - EPSRC (&NIH&etc) funding
- We’re increasingly asked to be ethical
  - See Neuhaus&Webmoor “agile ethics for massified research and visualization”
  - Informed consent?
  - Anonymization/Data Privacy?
- Social Media very tempting study
  - Graph data & Personal Identifying Information(PII)

Complexity & Value

UNIVERSITY OF CAMBRIDGE

- Do people really understand?
  - Recent court cases on Terms & Conditions
  - Judgement was no (spread betting case)
  - 40 pages of legalise is not comprehensible
  - And therefore not valid
- Do people know the value of their PII?
- Well, there's some disagreement:-
  - Preibusch&Jentzsch/Harasser: Monetizing Privacy
  - Yes, but its pretty cheap
  - Brown et al
  - No, but then when they lose it it's very expensive

But graph data is *so* interesting...

UNIVERSITY OF CAMBRIDGE

- Social Network is incredibly tempting
  - Twitter, Fb, WWW, Co-authorship
  - Social Science, Medicine, Commercial Motives….
- Characteristics like degree distribution, assortativity, betweenness very useful:
  - Info flow (percolation, gossip, epidemic)
  - Viral marketting (opinion dynamics)
  - Attack/defense/immunise/quarantine
- Can we anonymize
  - Bonneau "8 friends are enough" say naïve is no use
  - Backstrom "Wherefore Art Thou R3579X" say even quite subtle, no
  - Sala "Sharing Graphs using Differentially Private Graph Models", IMC 2011, say, finally, yes, but ***take great care*** *only _model_*

UNIVERSITY OF CAMBRIDGE

Nodes: people, Links: relations

- Looking at an abstract graph hides reality
- Node data is PII
- Its personal
- But collection of edge/link data can be used to identify nodes
- Even if PII is protected

Anonymizing node data

UNIVERSITY OF CAMBRIDGE

- If data is separate from graph, then anonymization is feasible.
- Risk of re-identification of records if not careful statistically
- Differential Privacy…

UNIVERSITY OF CAMBRIDGE

Differential *Piracy* example

- Imagine we have a database of pirates.
- If we query for a very tall pirate with a long beard, we are asking to identify a unique record ("Long John Silver"
- If we ask "How many pirates in Penzance?" we are safe, as there are lots
- Or if we ask for the number of 1 legged pirates who also have parrots?
- But don't ask for the pirate with the prosthetic hand, coz that even tells you his name…

UNIVERSITY OF CAMBRIDGE

Piracy Preserving DBase

name	port	Parrot	Wooden leg	Height
x	penzance	y	y	1.75
y	penzance	y	y	1.74
z	penzance	y	y	1.76
Dread pirate roberts	?	n	n	1.80
Hook	neverland			1.65
shakespeare	airport			1.60
sparrow	hollywood			1.50
Long john silver	Treasure island	y	y	2.00

UNIVERSITY OF CAMBRIDGE

Piracy Preserving DBase

#name	port	Parrot	Wooden leg	Height
xxx	penzance	y	y	1.75
yyy	penzance	y	y	1.74
zzz	penzance	y	y	1.76
Dread pirate roberts (*)	?	n	n	1.80
foo	neverland			1.65
bar	airport			1.60
baz	hollywood			1.50
fie	Treasure island	y	y	2.00

UNIVERSITY OF CAMBRIDGE

Adding graph edges messes this

- Link data represents a lot of attacks on hash of name:

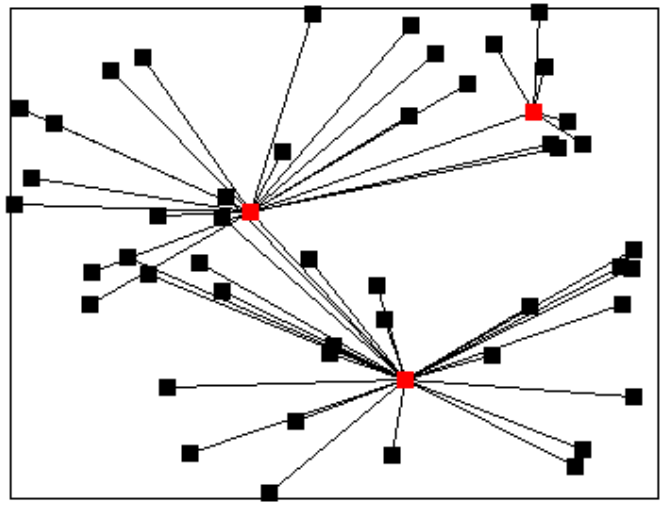

UNIVERSITY OF CAMBRIDGE

K-clique analysis reveals…

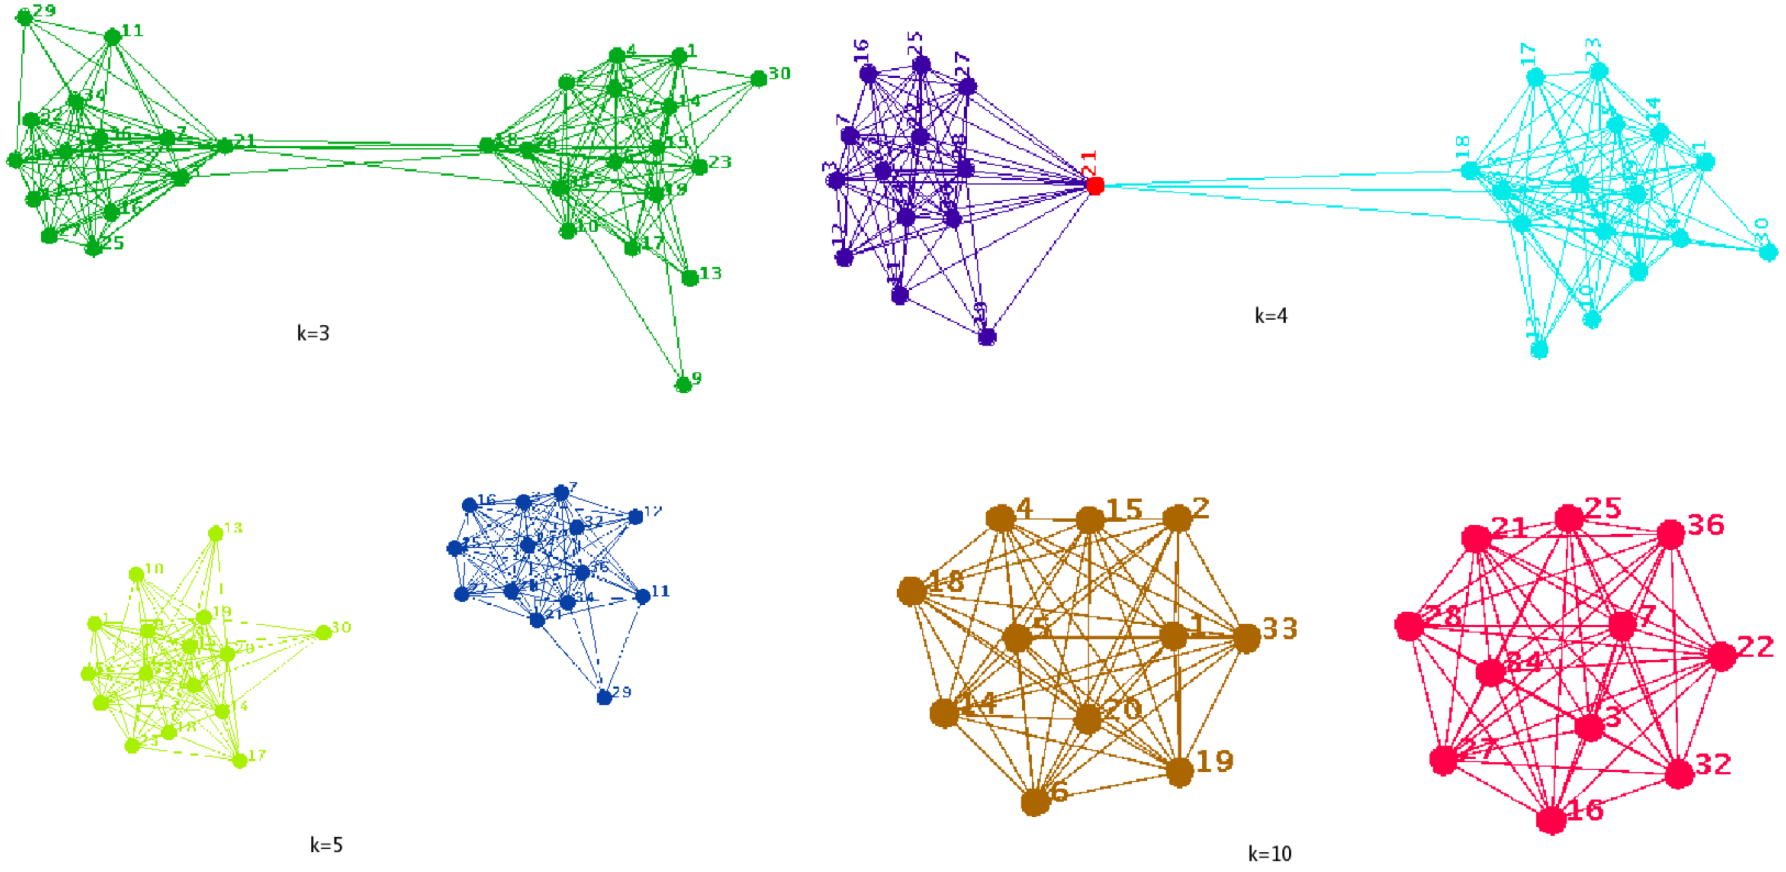

Minor External Knowledge.

UNIVERSITY OF CAMBRIDGE

- If I know that graph is of 2 years of undergrads,
- And the names of the class reps for y1&y2…
- I can infer those nodes identity trivially
- In the intersection of two cliques
- (the class reps meet together once per week)
- If this was student health record,
- I would have re-identified those two students… …

UNIVERSITY OF CAMBRIDGE

Lots more graph properties...

- Degree of nodes
- All the centrality types (including spectral etc)
- If links have properties too (strength, as in recommendation or reputation, or age, or other)
- Worse than ever!

Worse to come…

UNIVERSITY OF CAMBRIDGE

- Dunbar's # - 150
  - So if friend id is 32 bits, your friend list is 4800 bits on average
  - So the attack surface for identifying you is **huge**
- Worse Still - you have lots of "edges"

Hypergraphs...

UNIVERSITY OF CAMBRIDGE

- You have an edge for each type of relationship
  - kin, friend, colleague
  - Co-author of work
  - Co-located (e.g. paid congestion charge same time, used oyster card on same journey, checked in on foursquare same place)
  - Pay tax together, live at same postcode,
  - Sent SMS, IM, Email, Phone call, cell phone call from location
  - Same smart meter address

Re-identification is trivial

UNIVERSITY OF CAMBRIDGE

- Anyone in possession of 8 (see Anderson et al) I-Ds a graph of one set of edge type, with access to "anonymized" any other graph edge types, can re-identify the whole thing
- E.g. Tesco's clubcard can re-identify your whole health net…..

Dynamics

UNIVERSITY OF CAMBRIDGE

- Forgetting might help reduce attack surface
- Remove edges from old (&therefore less trusted)

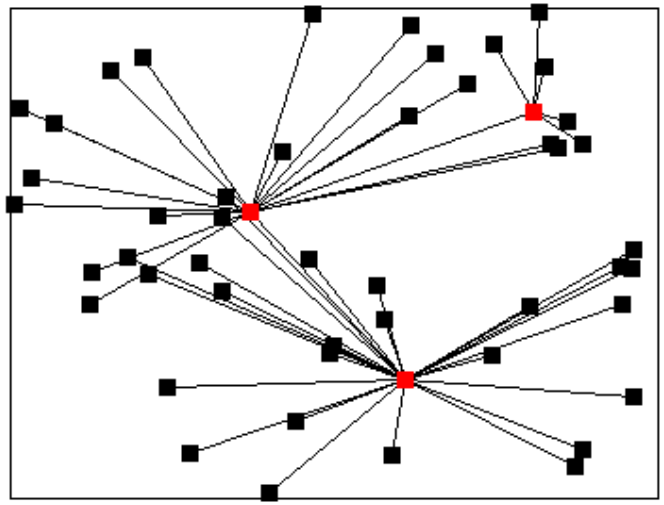

Manifesto

UNIVERSITY OF CAMBRIDGE

- Separate storage of node PII and link data
- Always crypt PII
- Decentralize nodes *and* links
- Partition PII by role
  - Kin, friend, worl, school
  - Health, finance, gov, social
- Use differrential privacy on graph data as well as node data
- Make it easy to understand
  - Maybe add forgetting

Take homes

UNIVERSITY OF CAMBRIDGE

- Doesn’t have to be all central
  - Cannot guarantee safe way to share graphs (sorry:-(
  - Can use Differential Privacy for node data records (without graph)
  - Can do diff priv on graphs but need to take care on
  - multiple priv pres queries can still snowball
- Epidemiologists don’t need our bank data, government don’t need our social data, we don’t need your health statistics

Questions?

UNIVERSITY OF CAMBRIDGE

- Remembering I'll be recording the Q&A
- And who you are:-)

System Architecture

UNIVERSITY OF CAMBRIDGE

- The Fresnel Team…


MIGHTY IS THE MOWEEN COOK

Realisations

UNIVERSITY OF CAMBRIDGE

- Typical sensor networks are **fit-for-purpose**
  - **Single-app:** Inflexible to updates and addition of services
  - **Single-user:** High cost of deployment and maintenance

Structural Health Monitoring

Smart Buildings

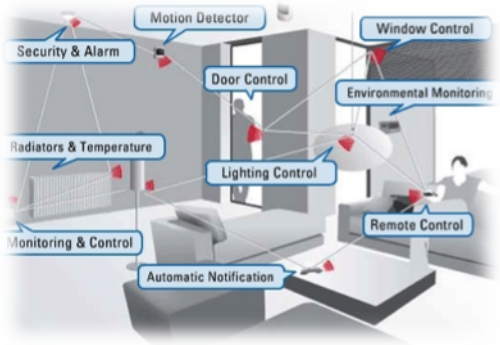

Mobile Urban Sensing

Shared Sensing Infrastructure

UNIVERSITY OF CAMBRIDGE

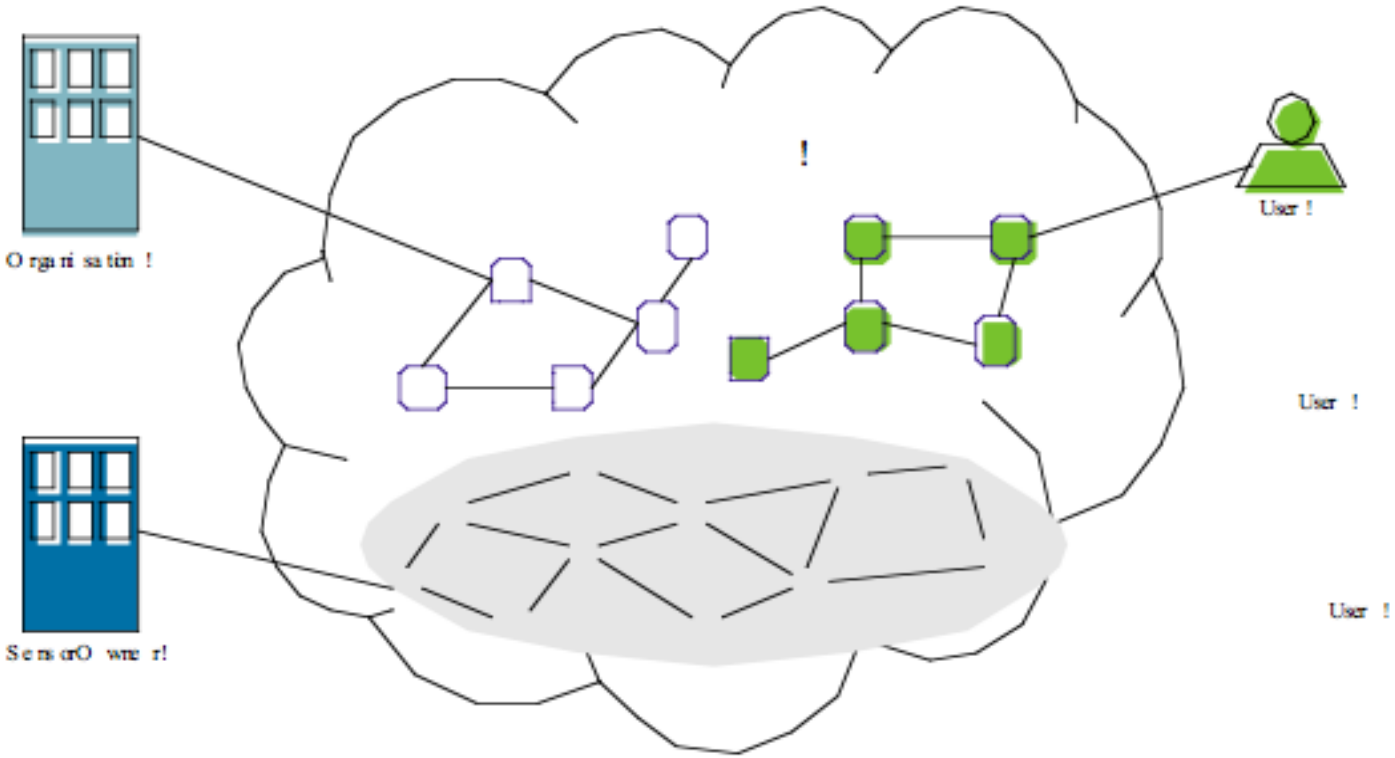

Management

UNIVERSITY OF CAMBRIDGE

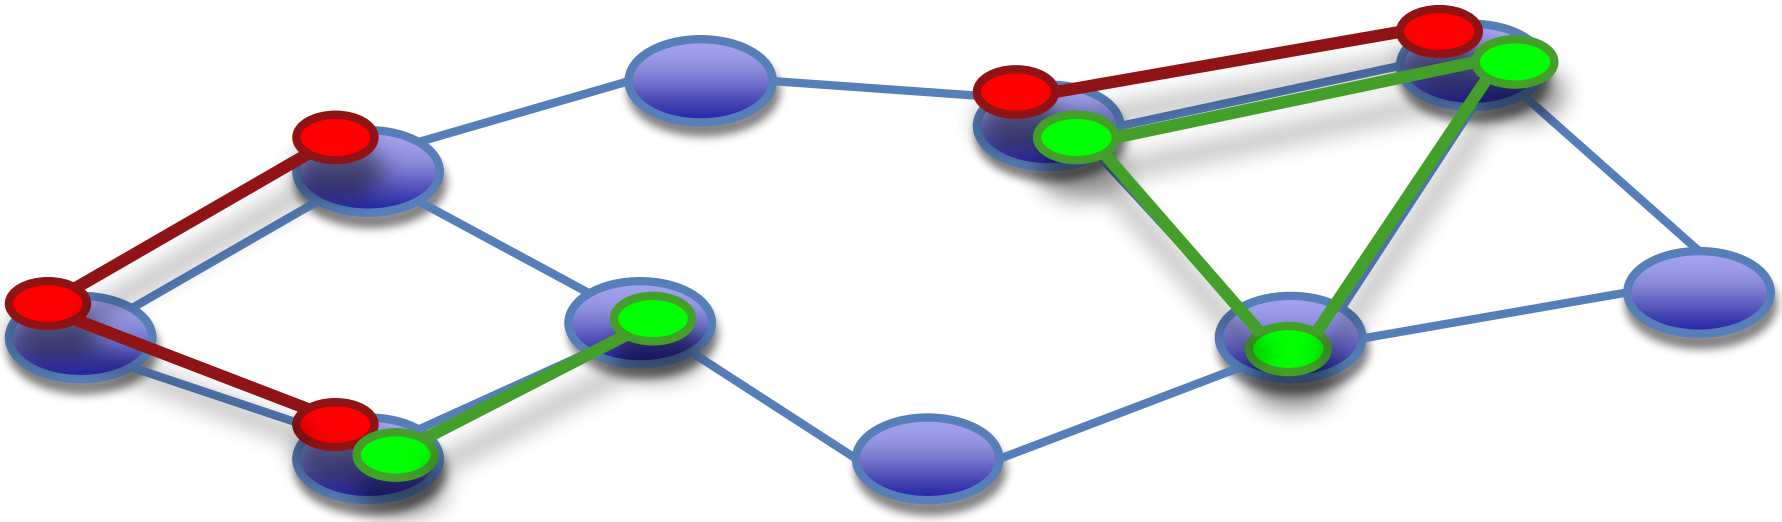

- Over-the-air installation

UNIVERSITY OF CAMBRIDGE

In-node Virtualization

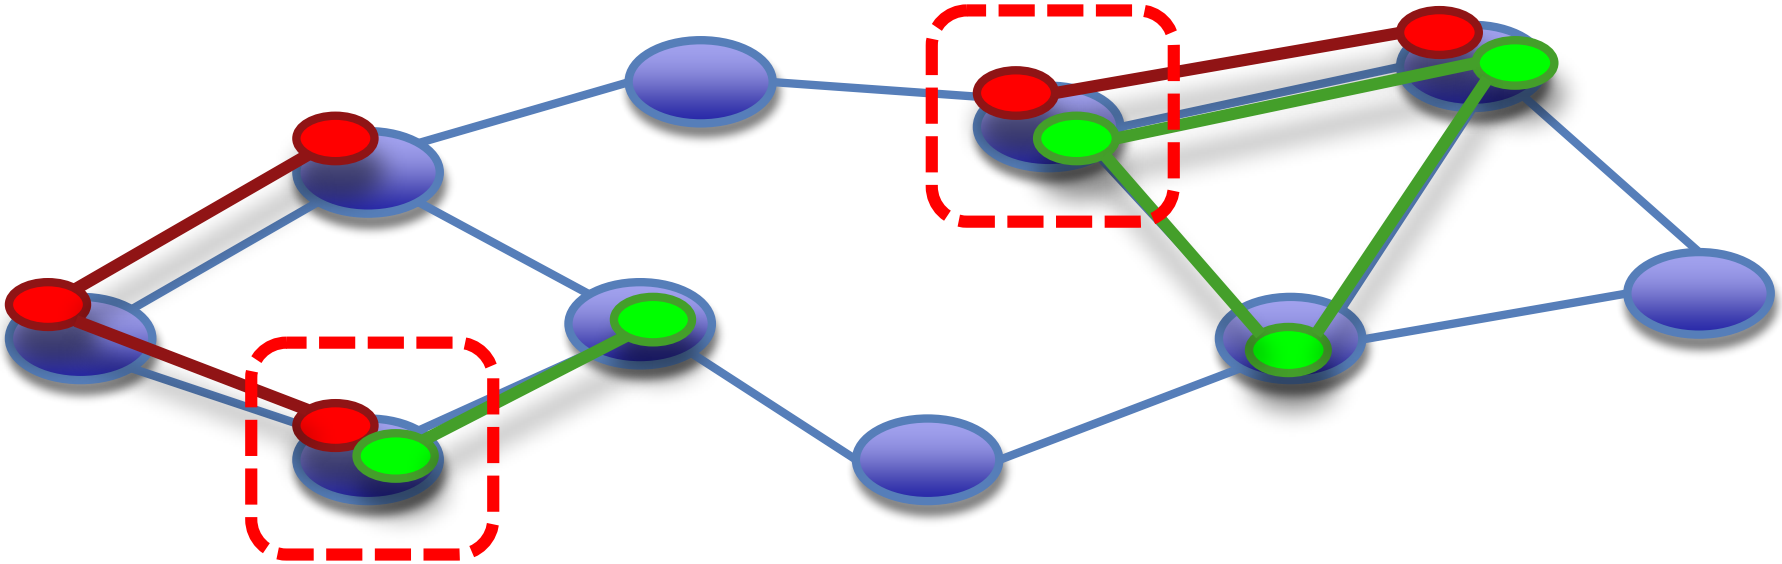

- Multiple applications co-live on the same node

Network support

UNIVERSITY OF CAMBRIDGE

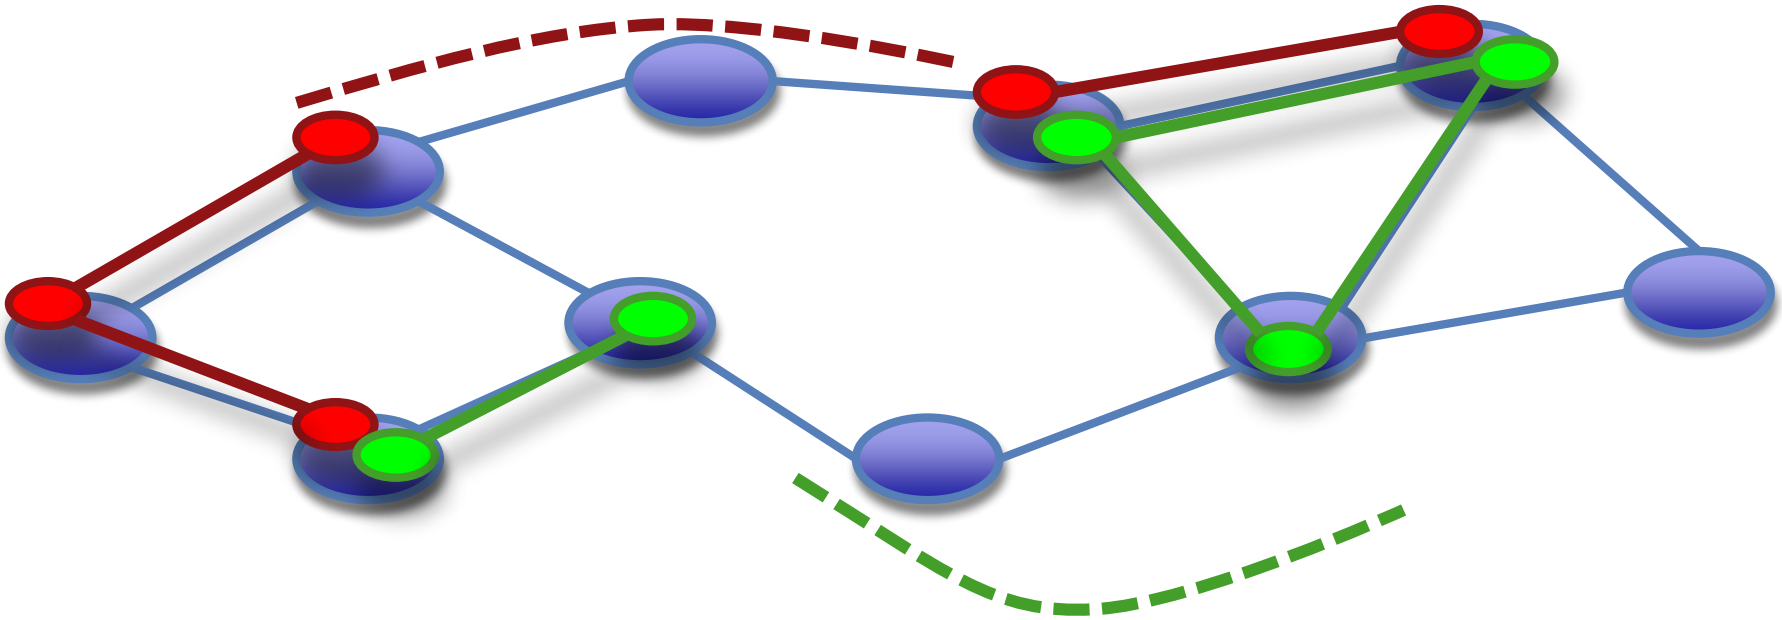

- **Isolation:** Applications should operate as if they are running over a dedicated sensor node/network.
- **Overlay virtual sensor network** for each application

UNIVERSITY OF CAMBRIDGE

The SenShare Platform

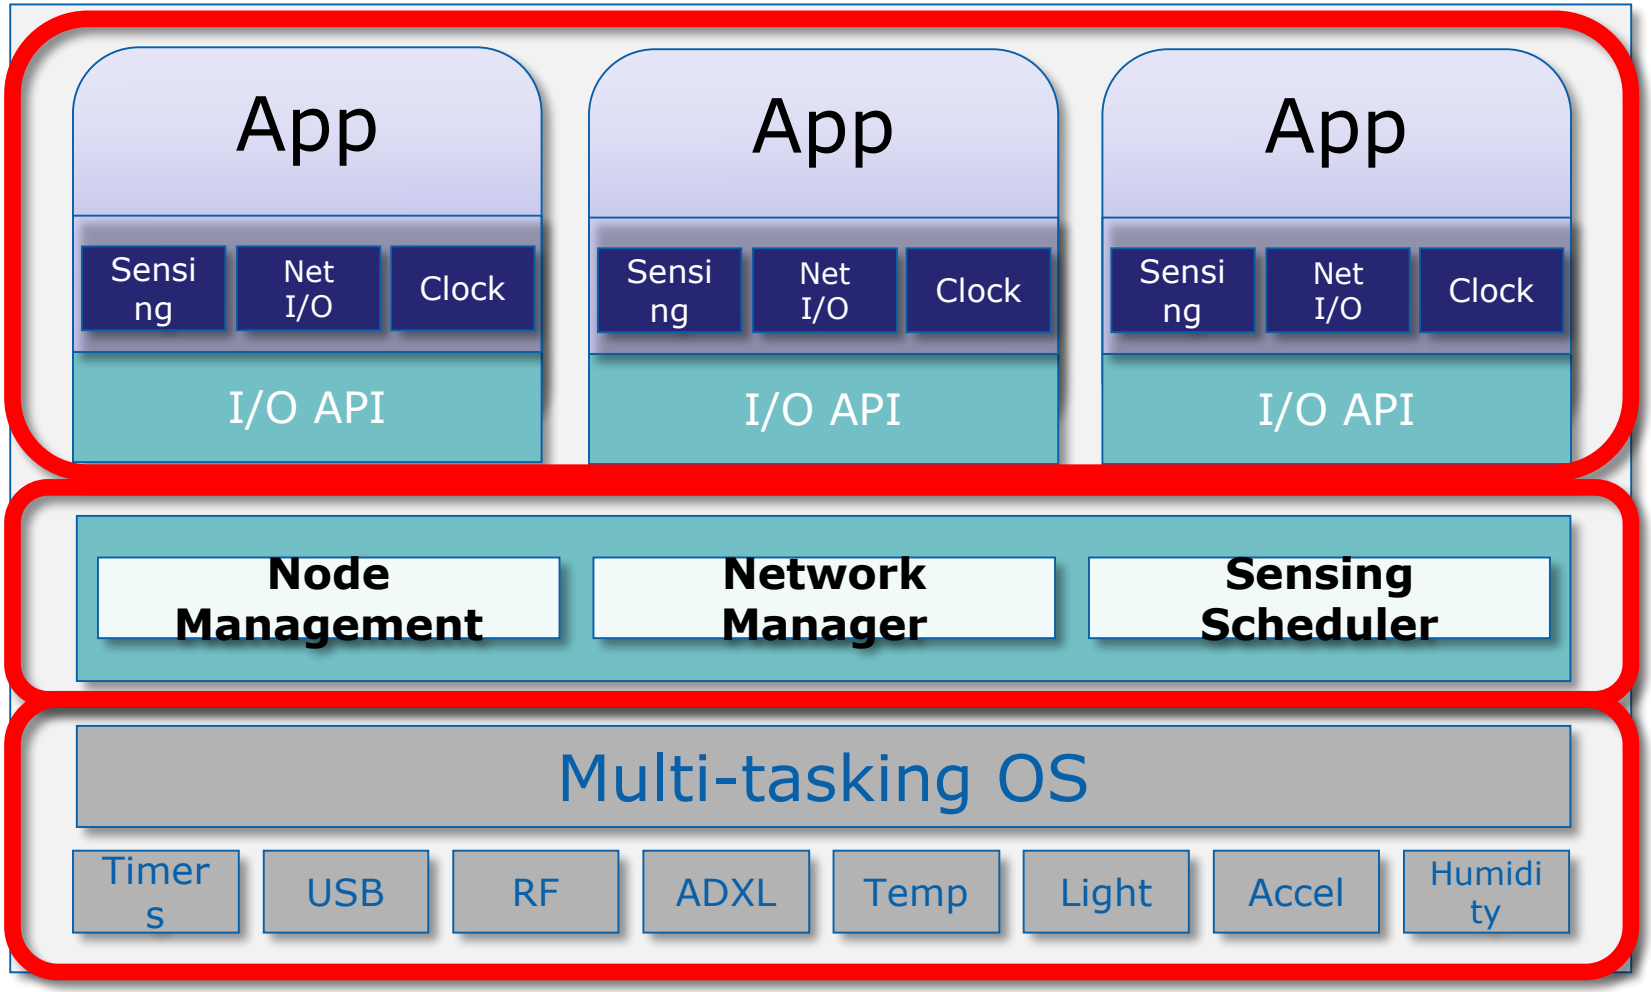
App
Sensing
Net I/O
Clock
I/O API
App
Sensing
Net I/O
Clock
I/O API
App
Sensing
Net I/O
Clock
I/O API
Node Management
Network Manager
Sensing Scheduler
Multi-tasking OS
Timers
USB
RF
ADXL
Temp
Light
Accel
Humidity

UNIVERSITY OF CAMBRIDGE

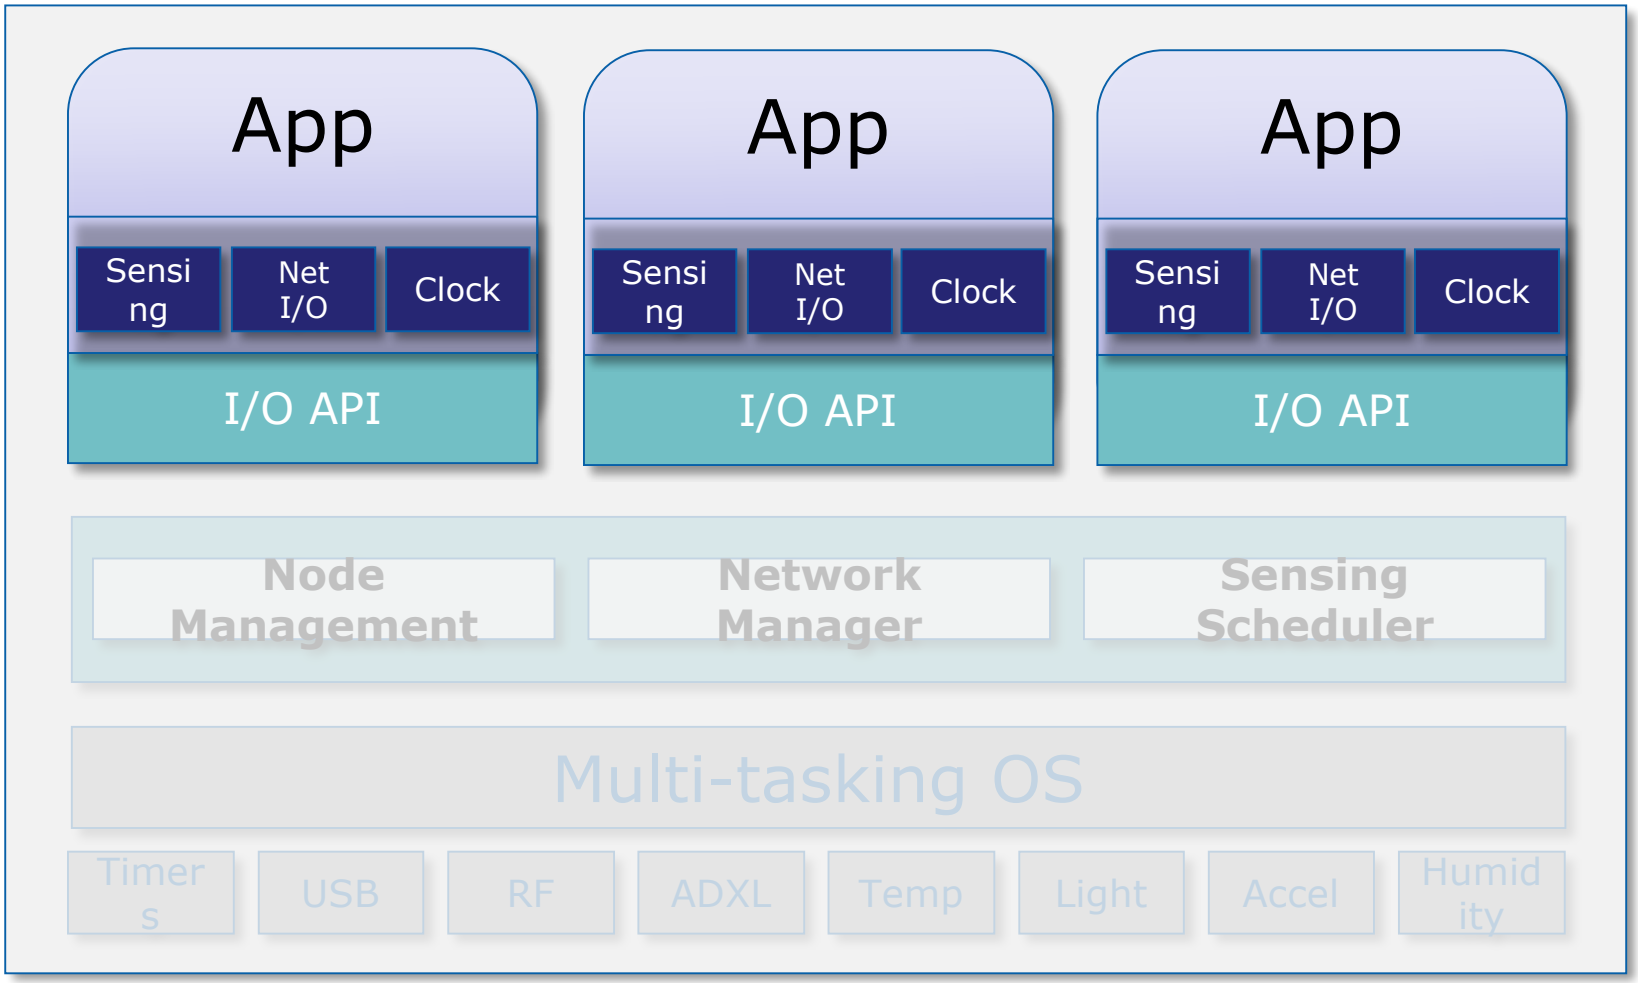

UNIVERSITY OF CAMBRIDGE

Application Support

- Support:
  - TinyOS
  - C/C++
  - ~~Future: Contiki~~

- Only the I/O is virtualized (not the binary)

TinyOS support

UNIVERSITY OF CAMBRIDGE

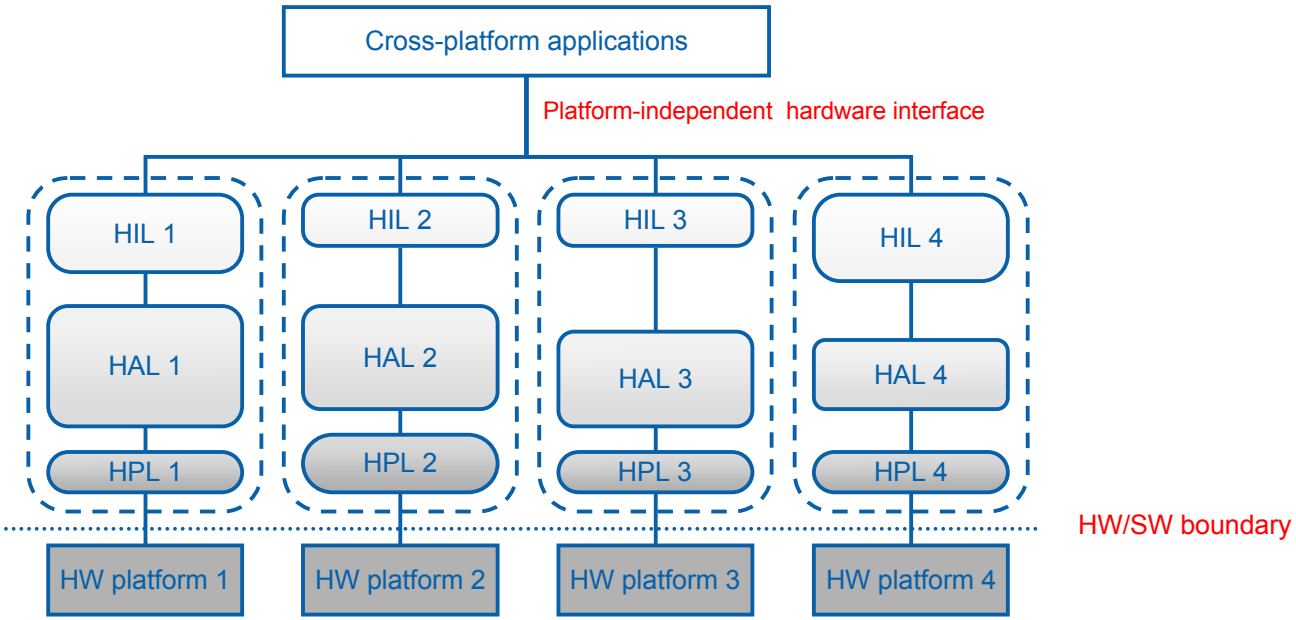

- The SenShare API is integrated with the TOS platform
- Application's source code does not require modifications if HIL is used

TinyOS support

UNIVERSITY OF CAMBRIDGE

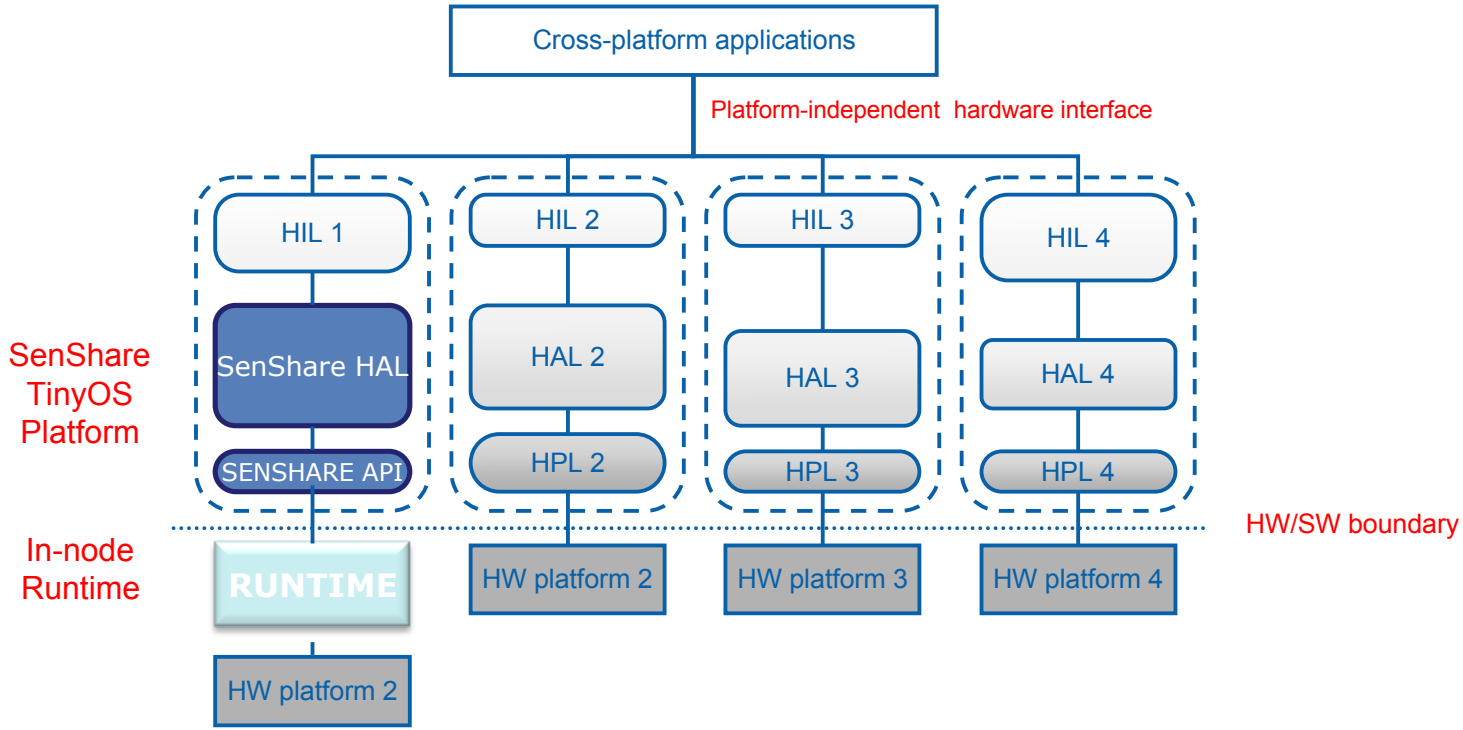

- The SenShare API is integrated with the TOS platform
- Application's source code does not require modifications if HIL is used

UNIVERSITY OF CAMBRIDGE

C/C++ Application API

- Applications are linked with our library
- Hardware Independent API to:
  - Read/write to sensors
  - Send/receive network message
  - Interact with the runtime

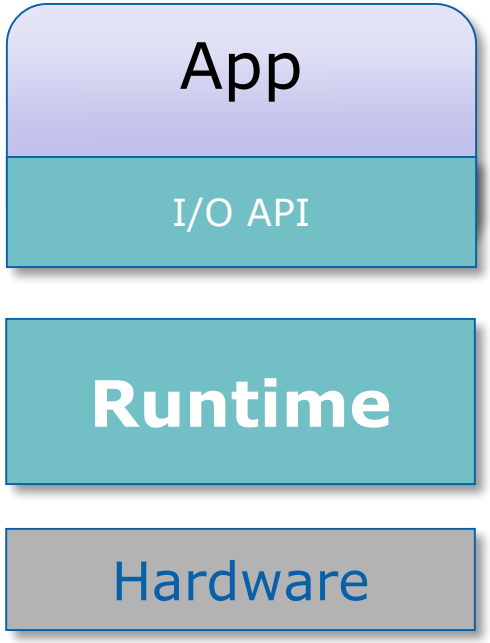

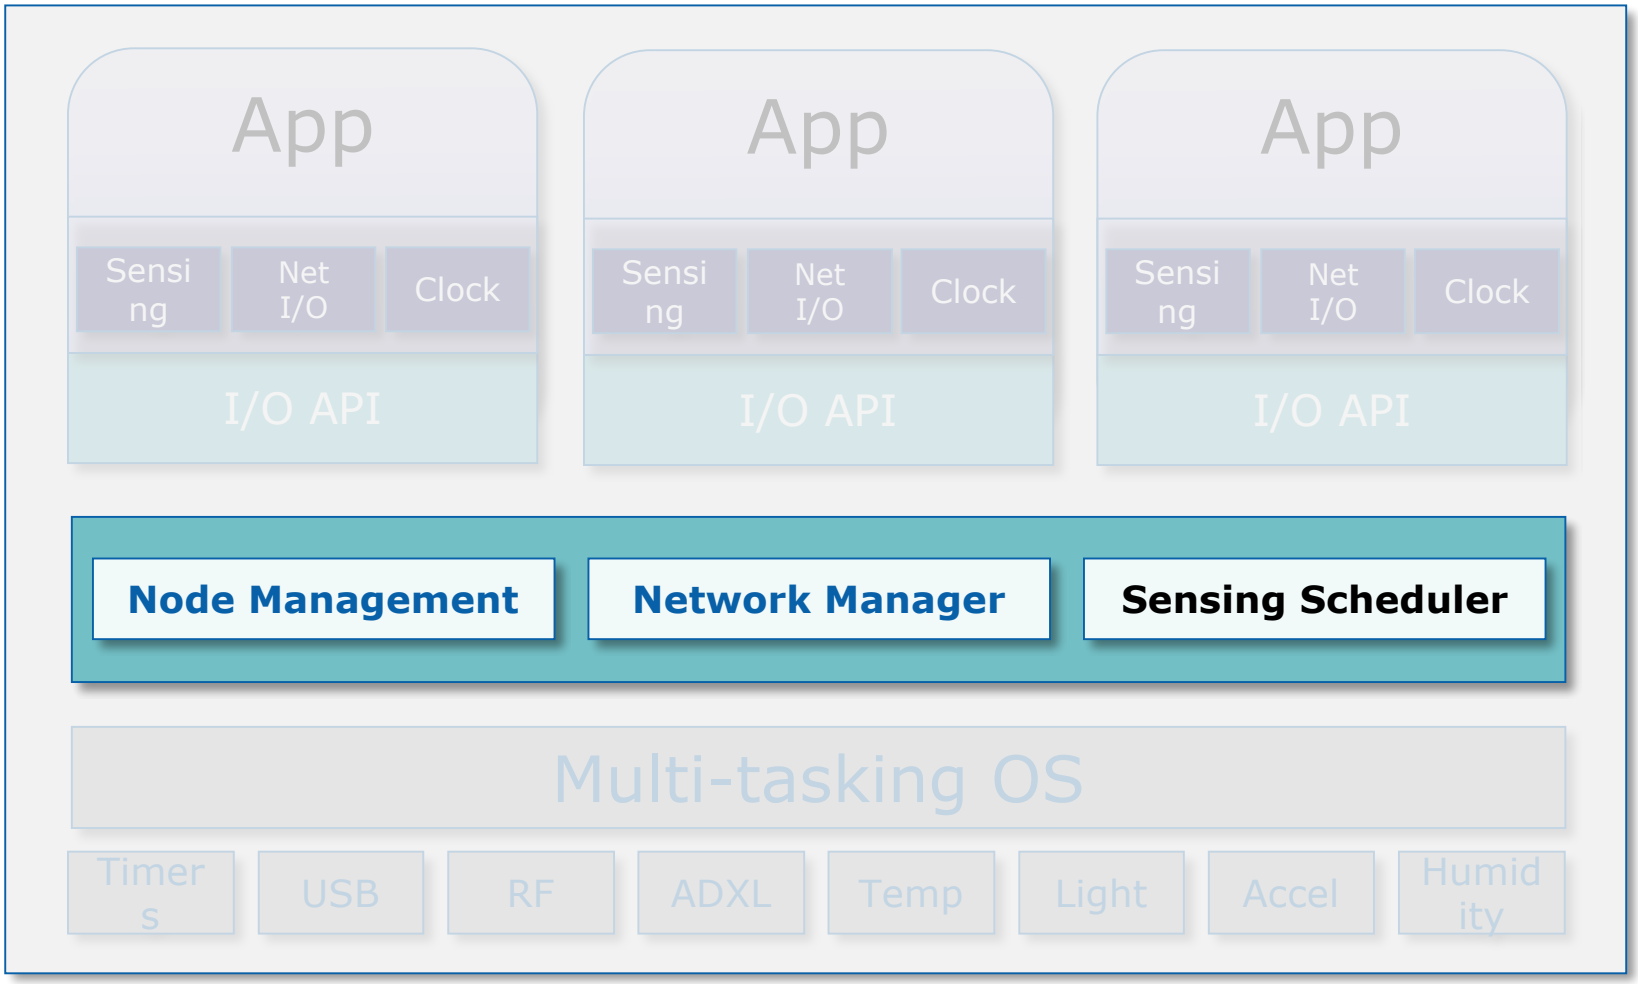
App
Sensing
Net I/O
Clock
I/O API
App
Sensing
Net I/O
Clock
I/O API
App
Sensing
Net I/O
Clock
I/O API
Node Management
Network Manager
Sensing Scheduler
Multi-tasking OS
Timers
USB
RF
ADXL
Temp
Light
Accel
Humidity

UNIVERSITY OF CAMBRIDGE

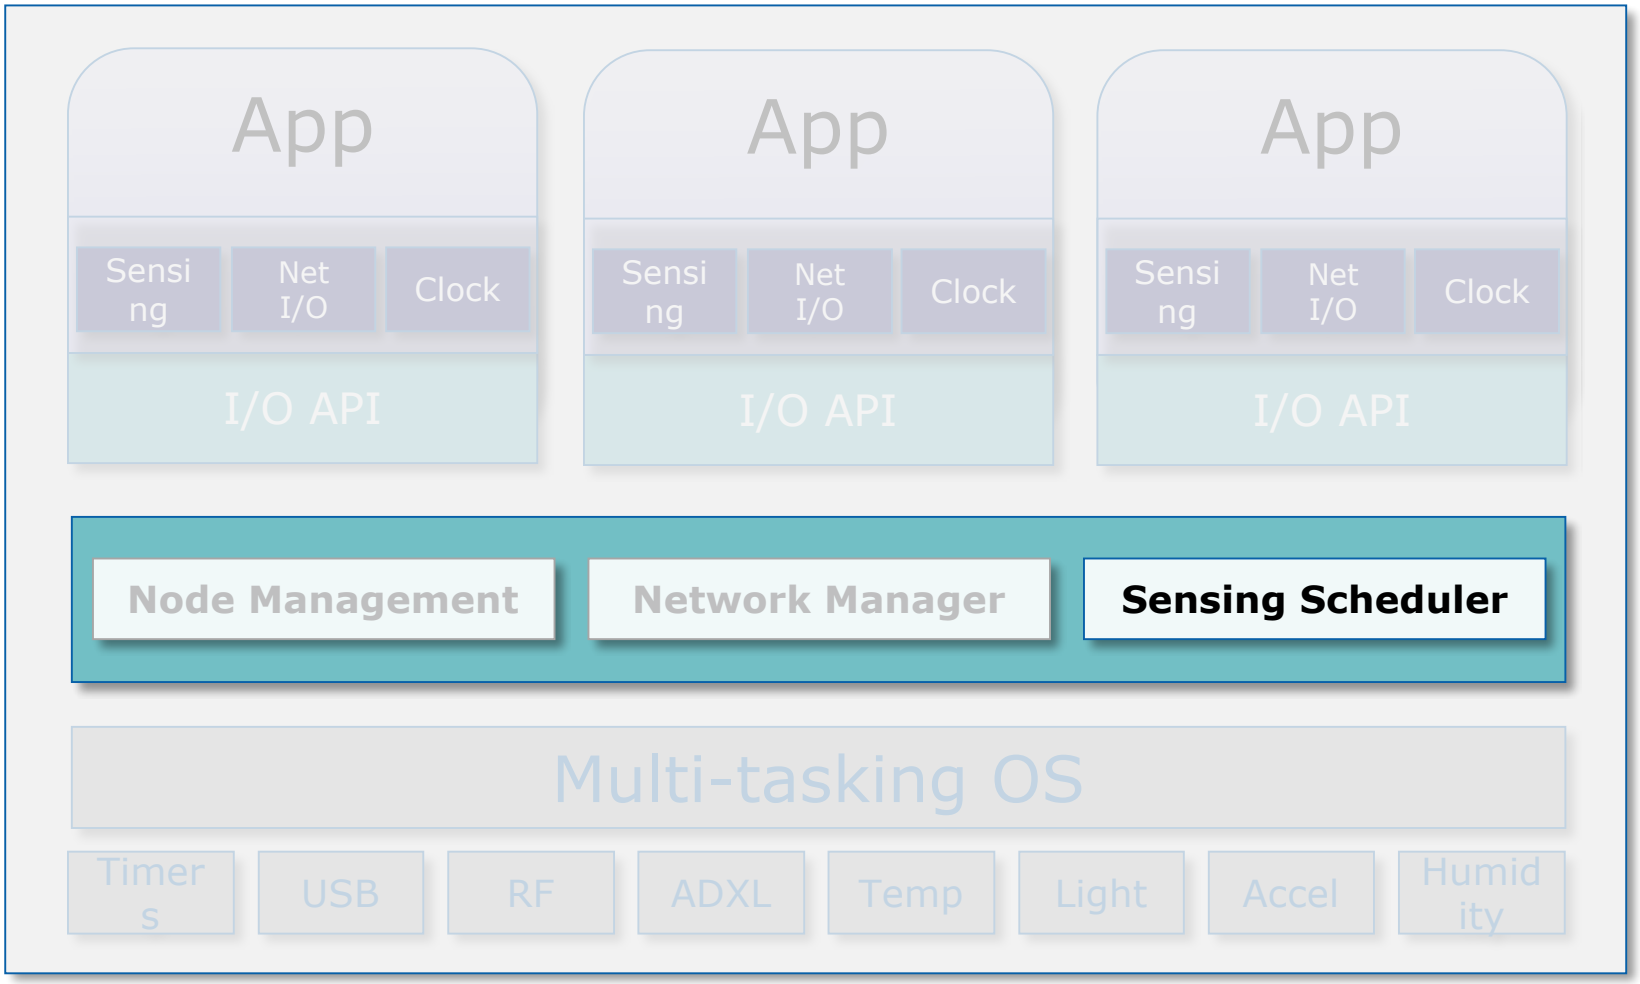

UNIVERSITY OF CAMBRIDGE

Sensing Scheduler

- Virtualize I/O
- Advantages:
  - Common way to support all platforms
  - Access policies (privacy)
  - Throttle applications
  - Isolation
- Multiplexing
  - Asynchronous requests
  - Bursting queue
  - Serve all requests simultaneously

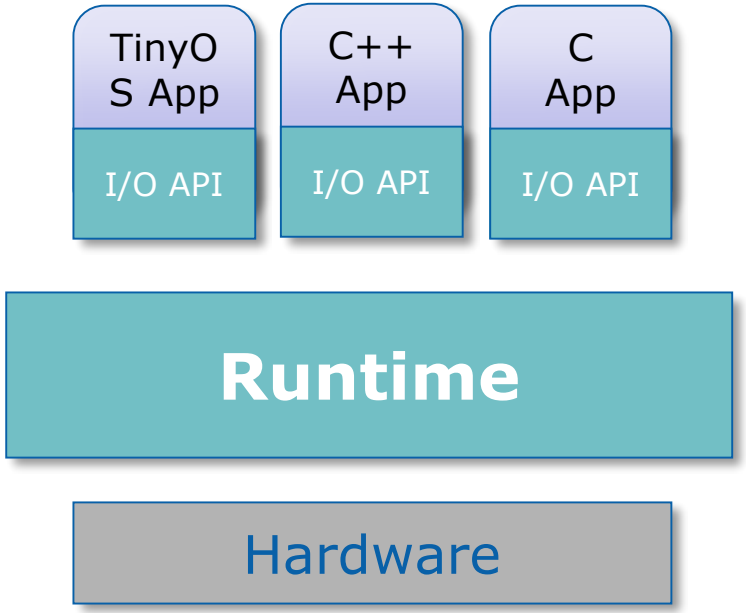

UNIVERSITY OF CAMBRIDGE

Virtualization Penalty

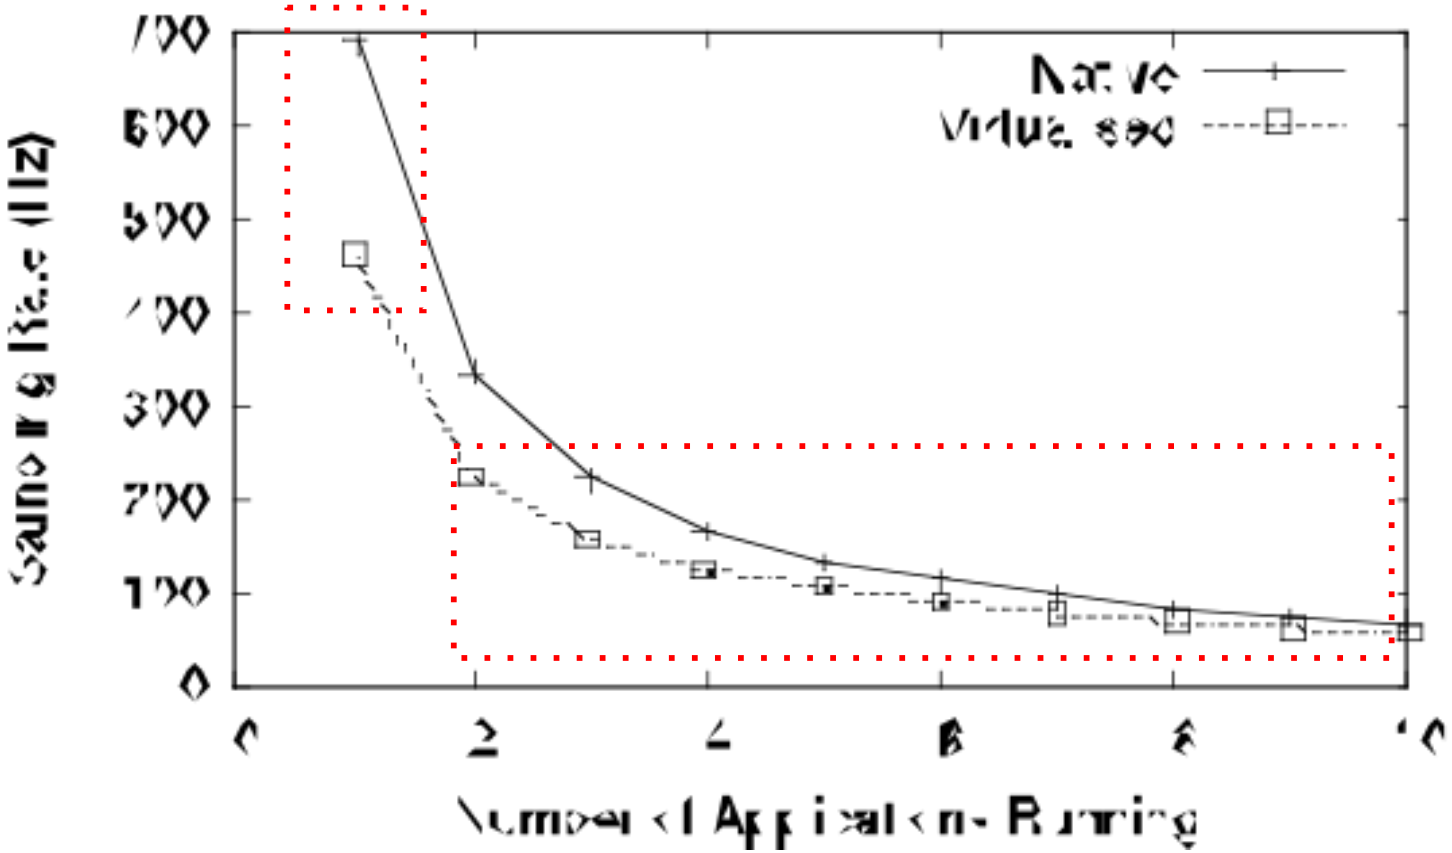

UNIVERSITY OF CAMBRIDGE

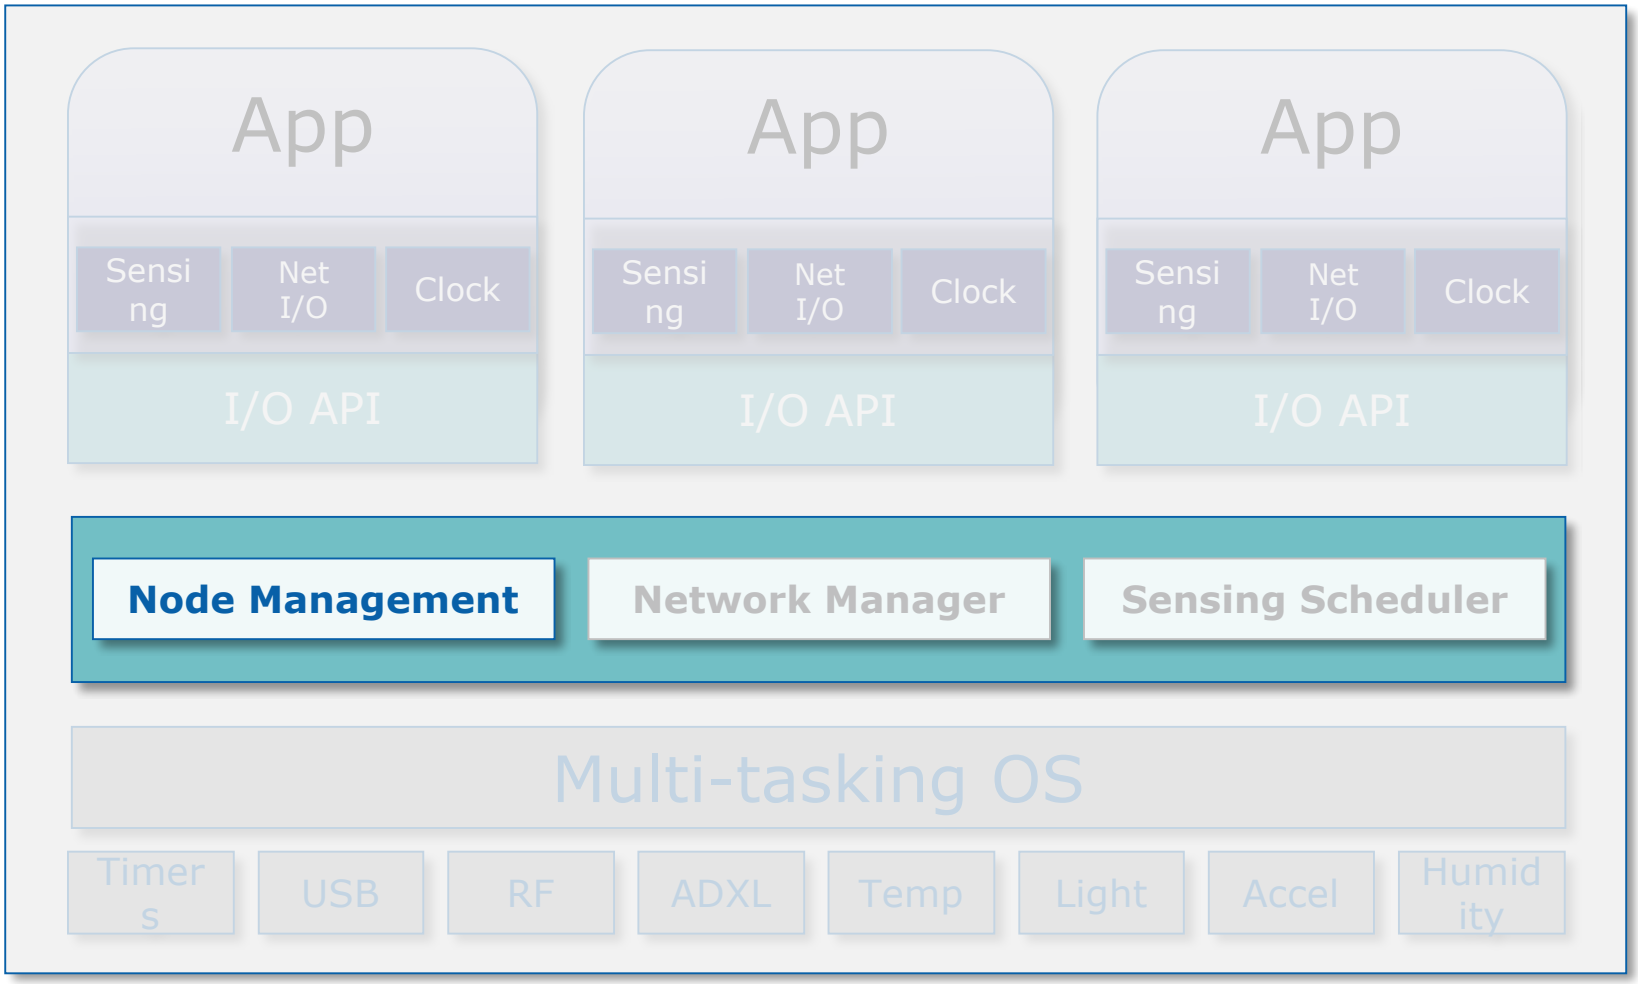

UNIVERSITY OF CAMBRIDGE

Node Management

- Allow the network owner and the users to
  - Configure Nodes
  - Change access policies
  - See what is running
  - Debug
  - Deploy/Install applications
  - Start/Stop applications

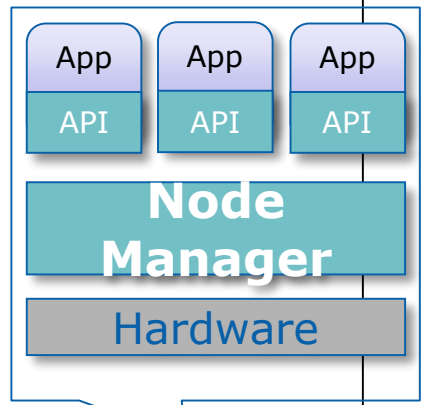

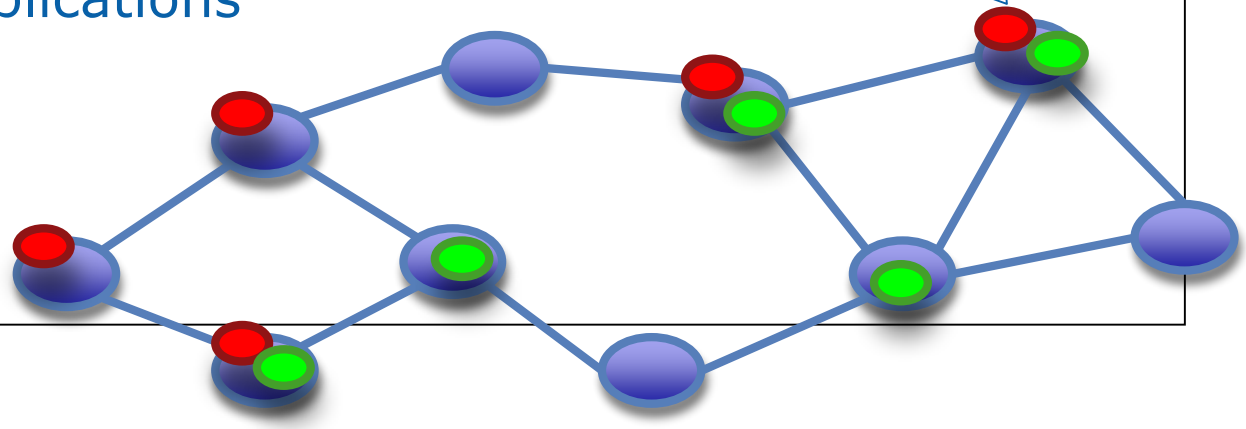

Network Support

UNIVERSITY OF CAMBRIDGE

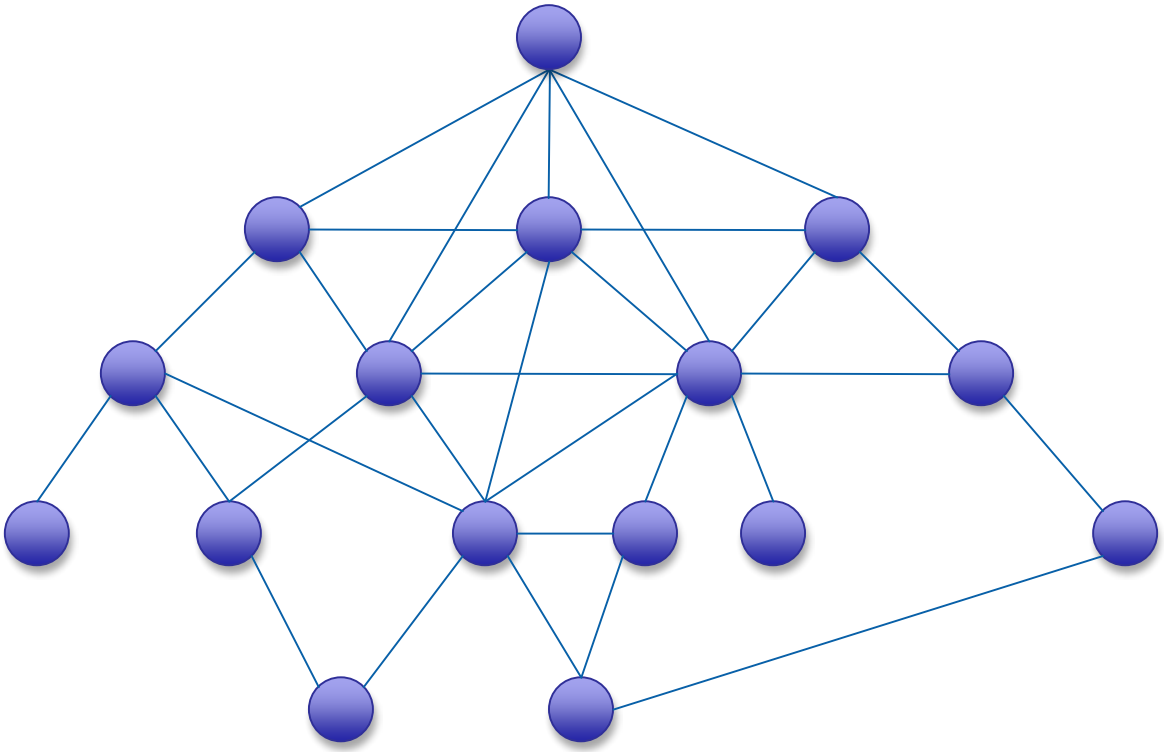

UNIVERSITY OF CAMBRIDGE

Selective Control

ROOM_ID	Select a room
ROOM_TYPE	Select a room type (corridor, office, etc)
NODE_ID	Select individual nodes
NODE_TYPE	Select specific hardware (e.g., imote2)
SENSOR	Select specific sensor (e.g. temperature)
POWER_TYPE	Select on power type (battery, permanent)
AVAIL_POWER	Select on remaining power
NETWORK_LOAD	Select on average traffic
CPU_LOAD	Select on average CPU load
AVAIL_MEMORY	Select on available memory
AVAIL_STORAGE	Select on available storage

- Examples:
  - <SELECT **ALL;** SENSOR=temp **AND** ROOM_TYPE=office **AND NOT** NODE_ID=5>
  - <SELECT **ONE**; SENSOR=temp **AND (**ROOM_ID=1 **OR** ROOM_ID=2)>

Selective Control

UNIVERSITY OF CAMBRIDGE

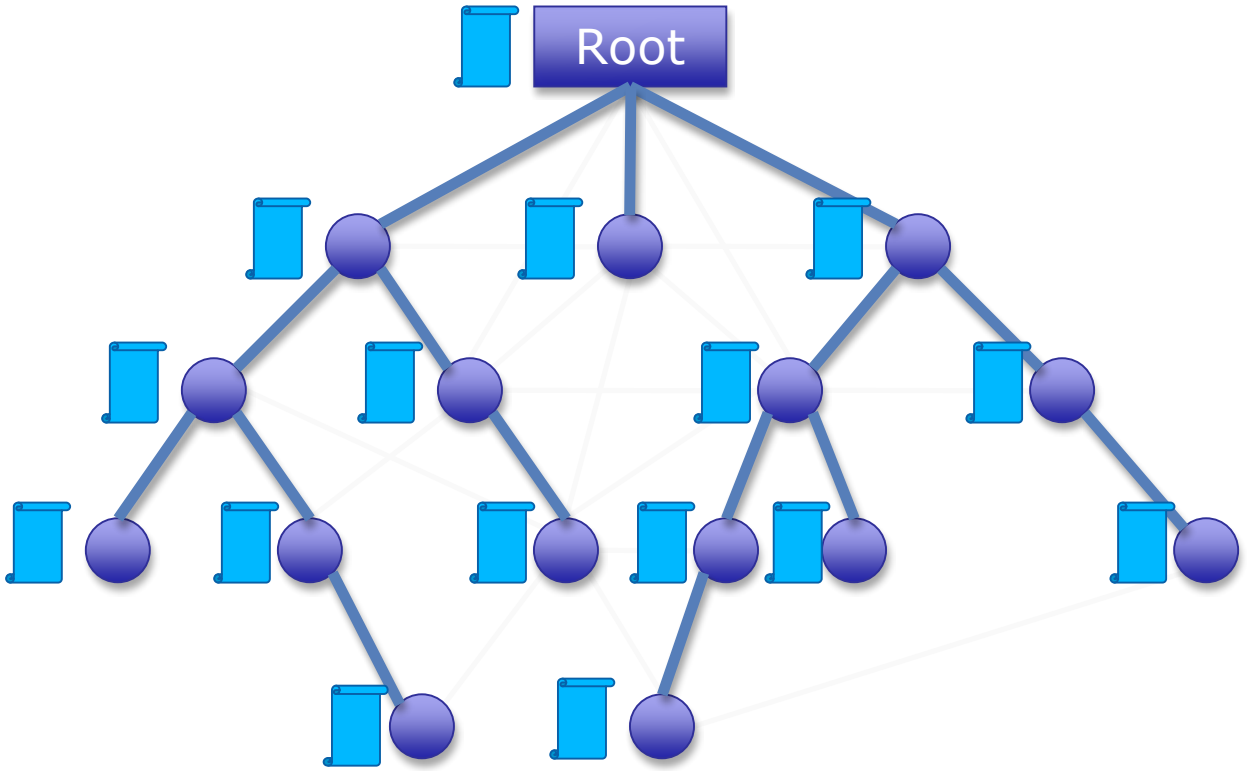

Selective Deployment and Installation

UNIVERSITY OF CAMBRIDGE

- Install new applications ***over-the-air***
- Modified Deluge protocol*
- The CTP tree is used again

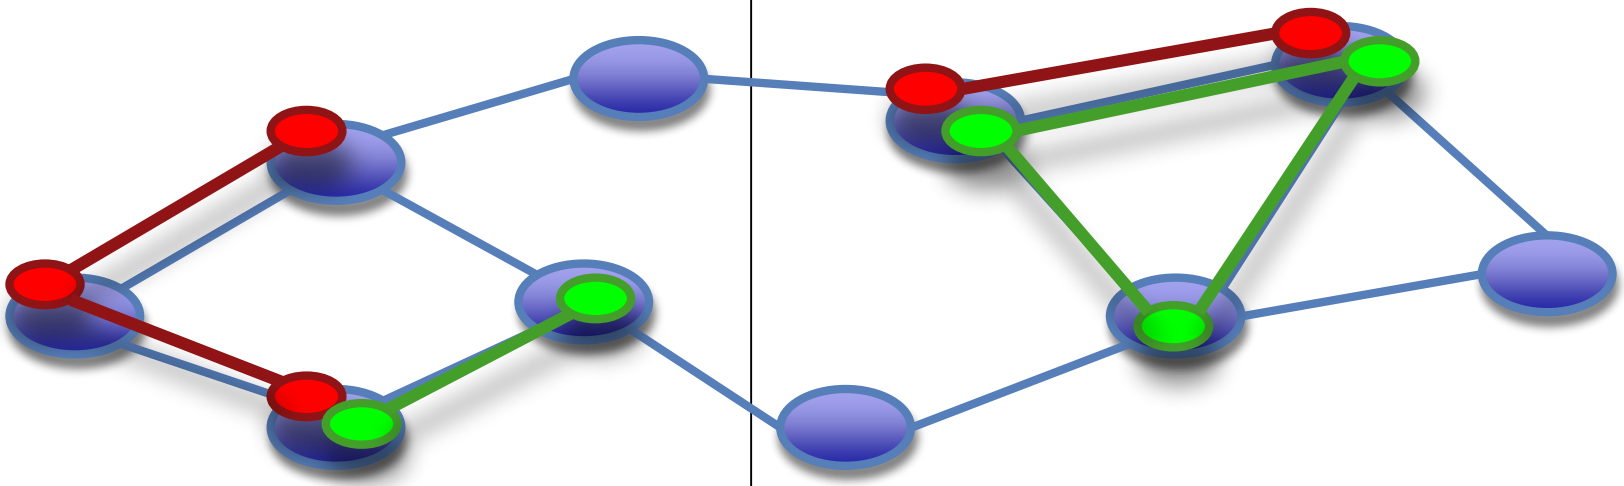

* J.W.HulandD.Culler.Thedynamicbehaviorofadataadisseminationprotocolfor network programming at scale. In Proceedings of the 2nd international conference on Embedded networked sensor systems, SenSys '04, pages 81–94, New York, NY, USA, 2004. ACM.

UNIVERSITY OF CAMBRIDGE

Participating nodes

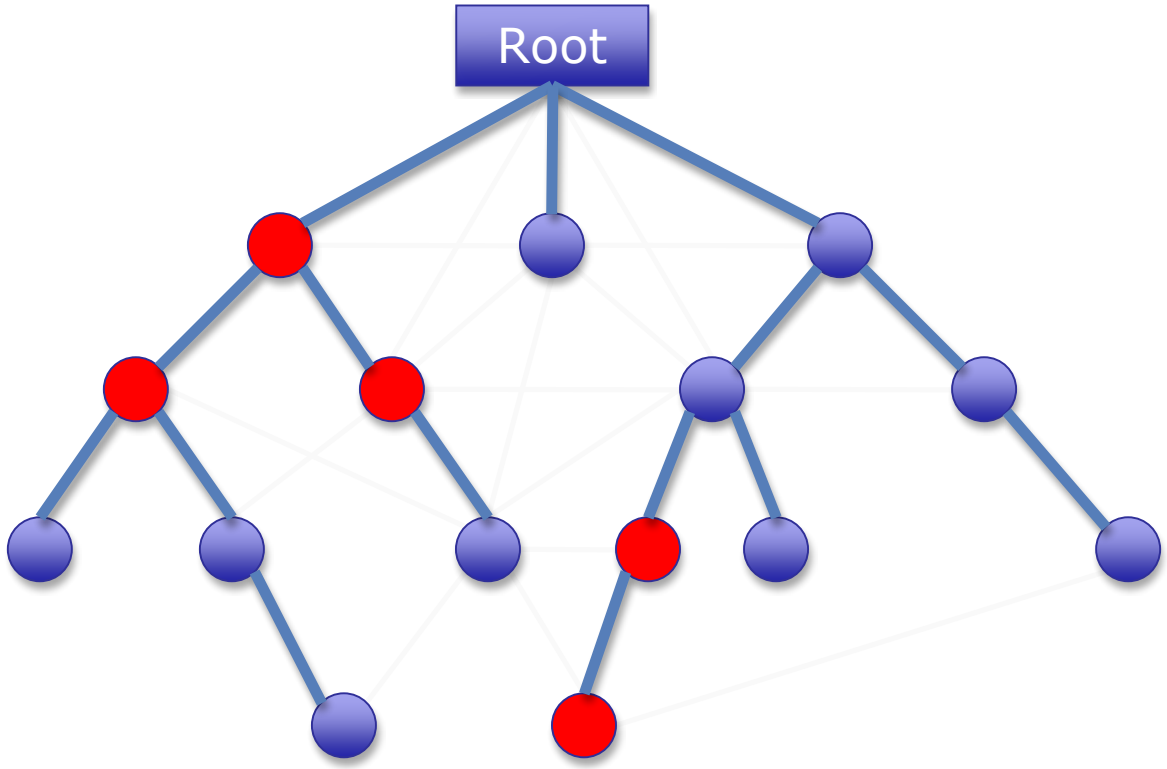

UNIVERSITY OF CAMBRIDGE

Data Broadcast

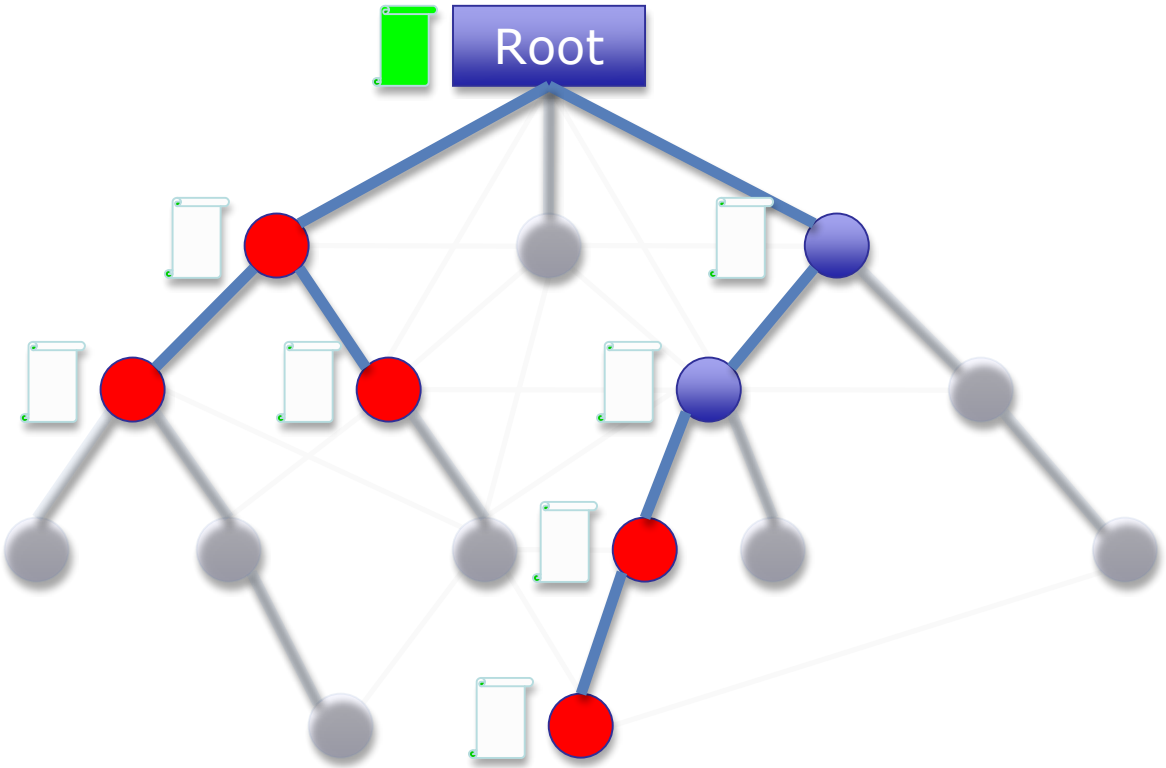

UNIVERSITY OF CAMBRIDGE

Data Broadcast

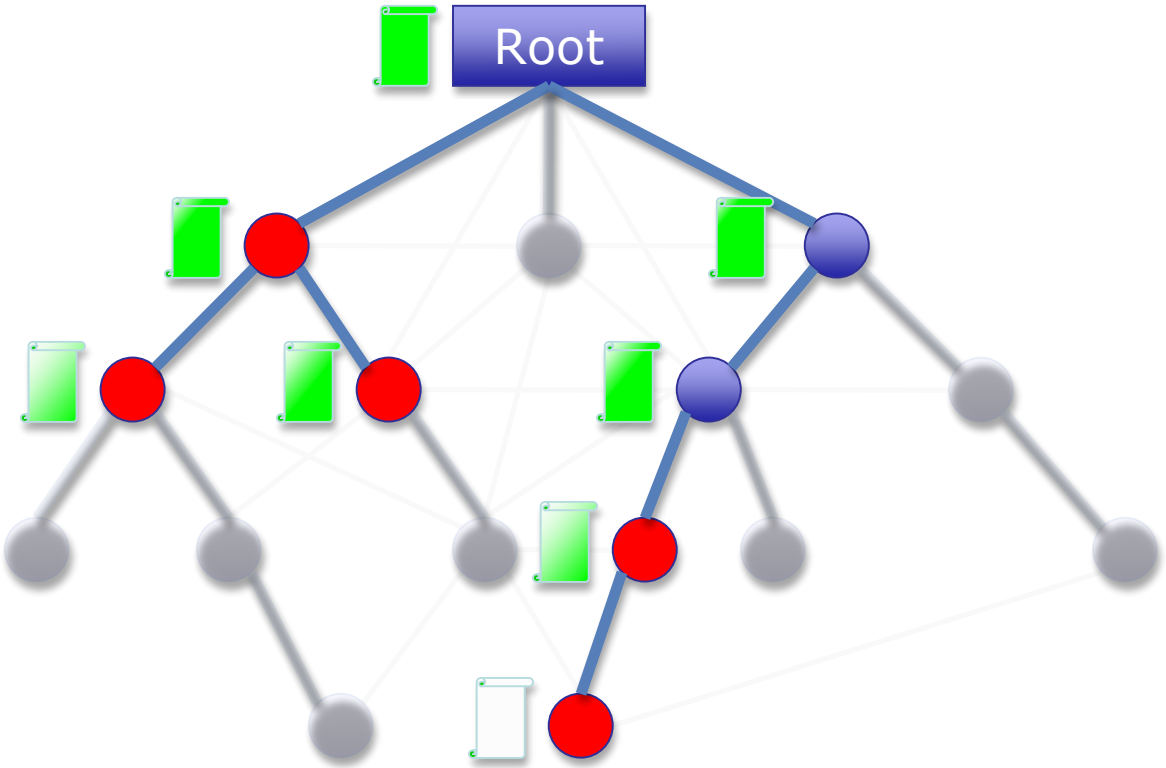

Data Broadcast

UNIVERSITY OF CAMBRIDGE

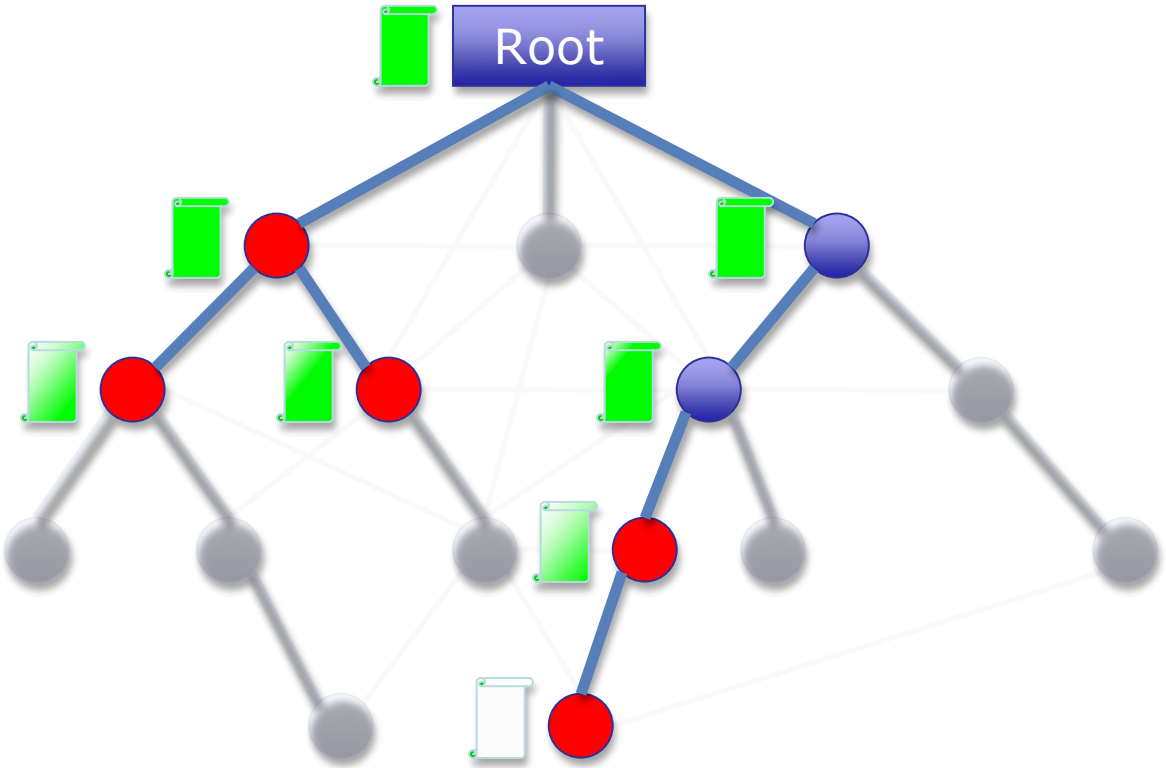

UNIVERSITY OF CAMBRIDGE

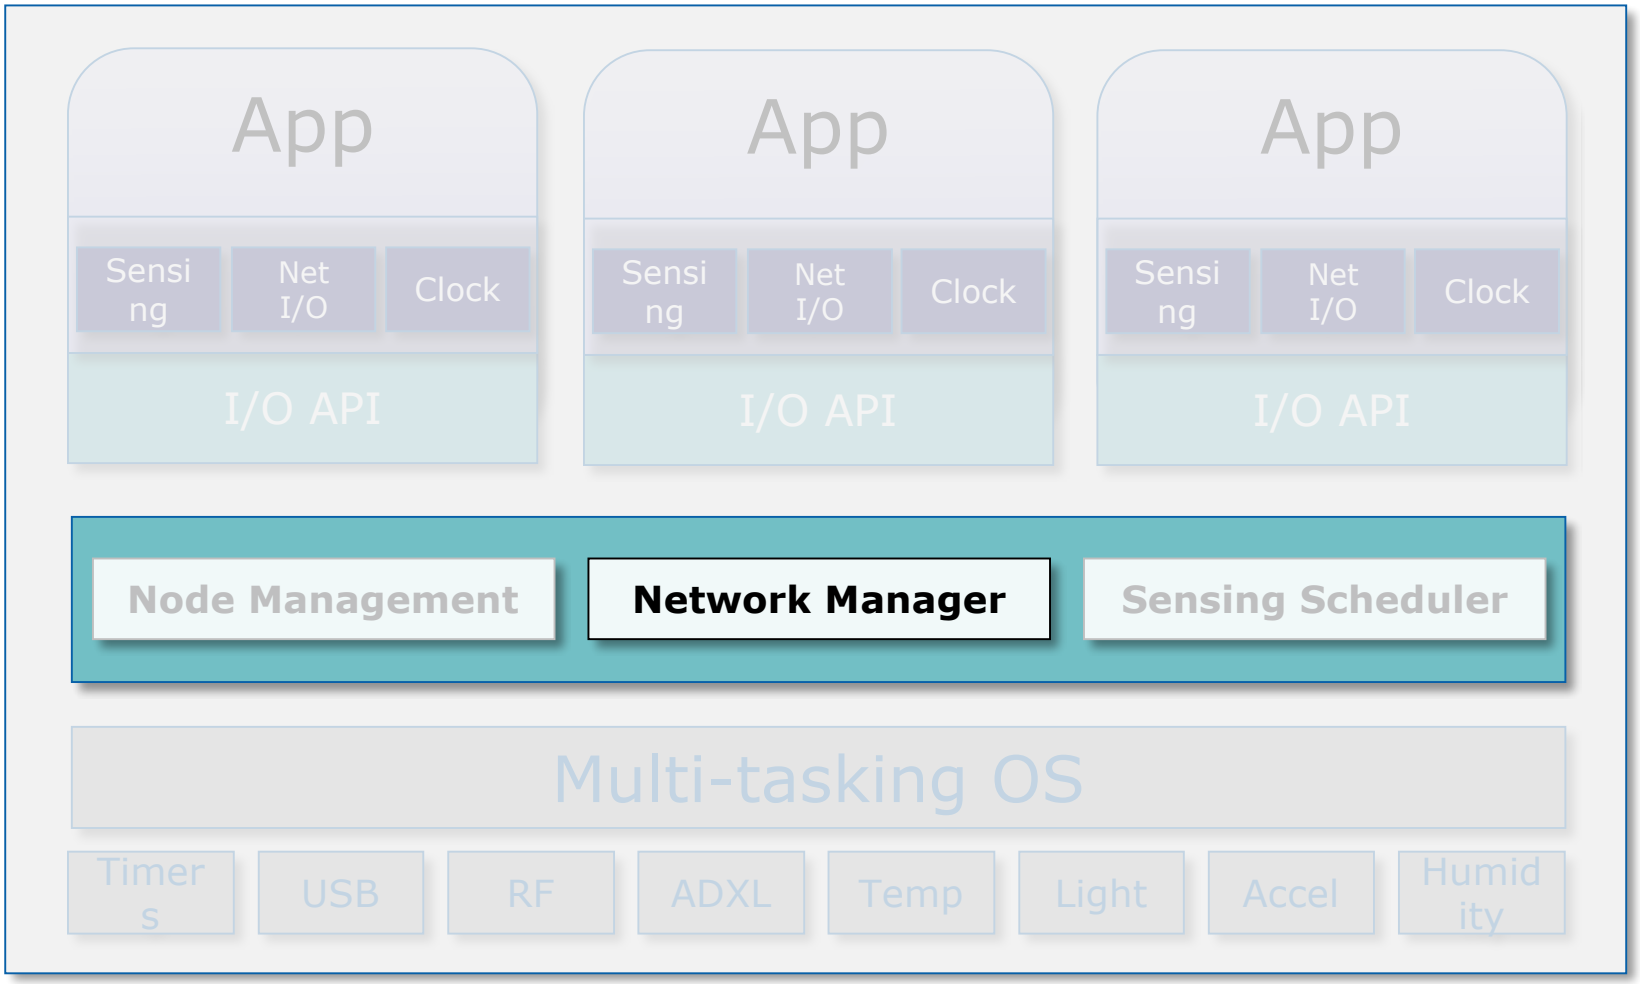

Traffic Isolation

UNIVERSITY OF CAMBRIDGE

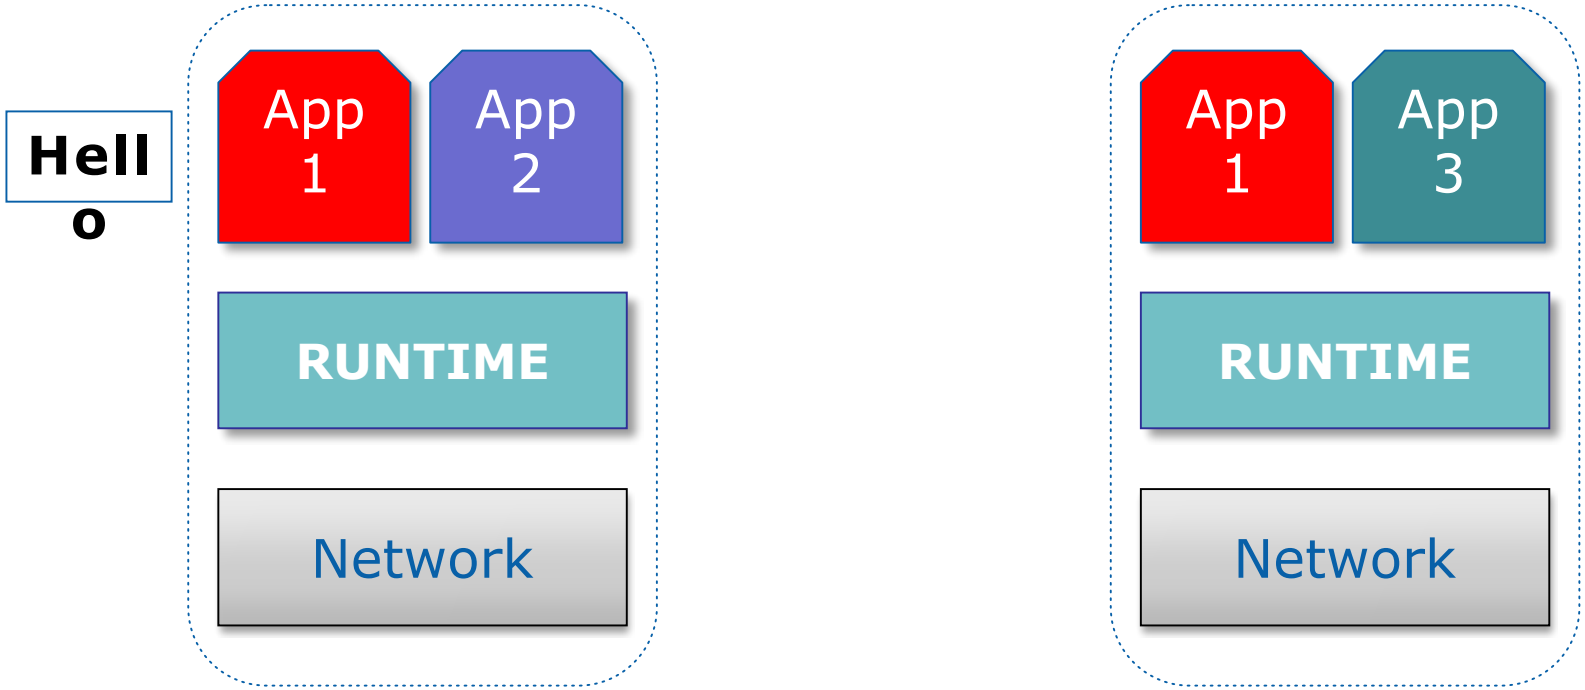

Traffic Isolation

UNIVERSITY OF CAMBRIDGE

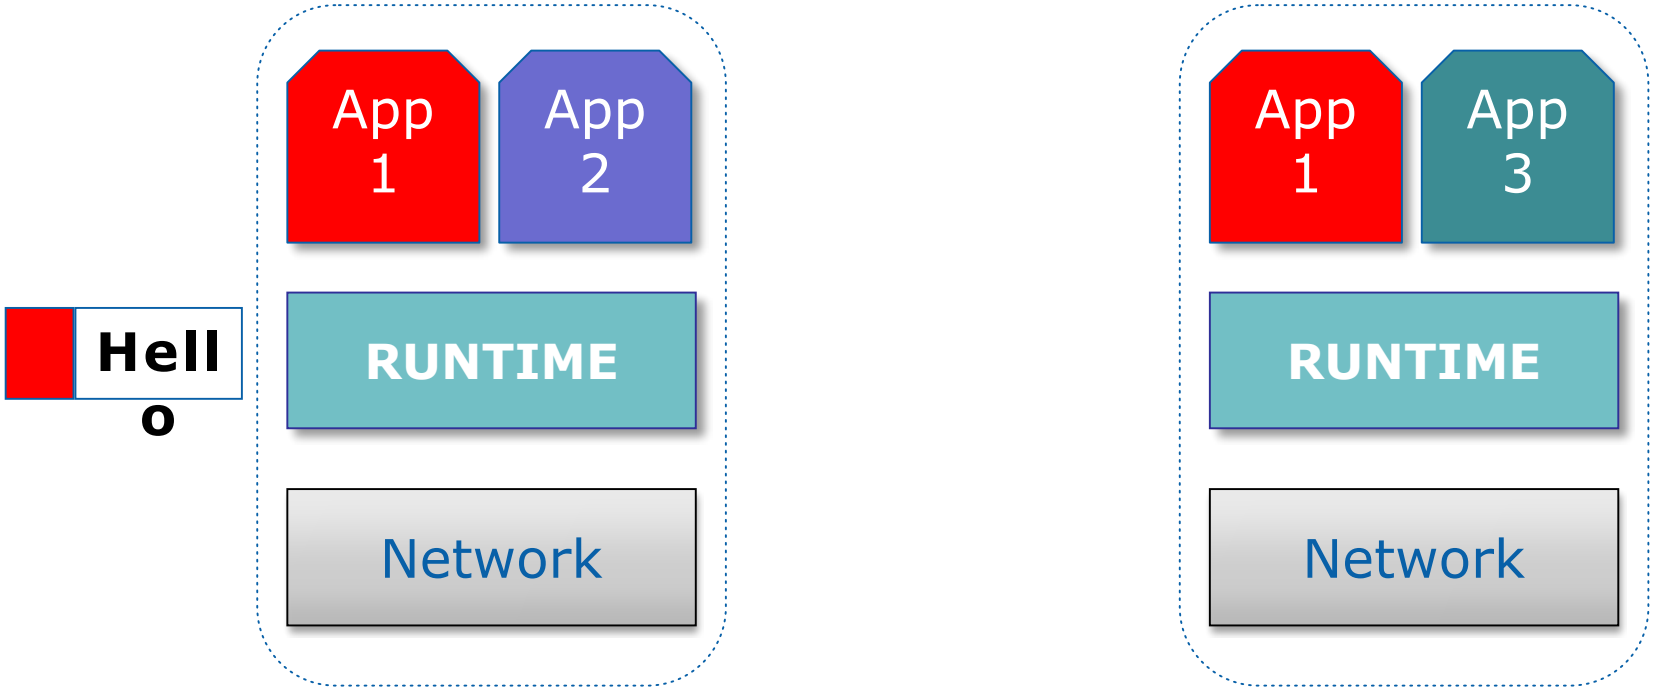

Application header: `{app id, seq no, origin, destination}`

UNIVERSITY OF CAMBRIDGE

Overlay sensor networks

- Connect disconnected blobs

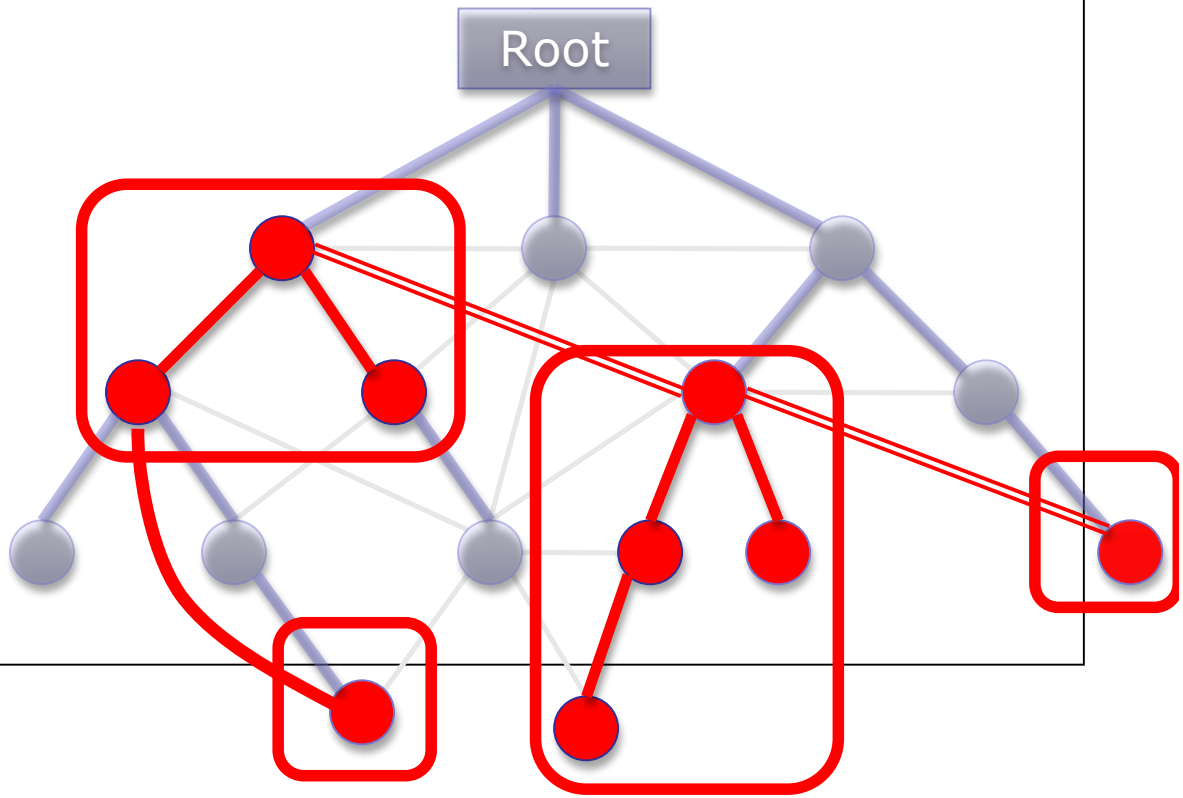

UNIVERSITY OF CAMBRIDGE

Overlay sensor networks

- Connect disconnected blobs
- Root over the CTP tree

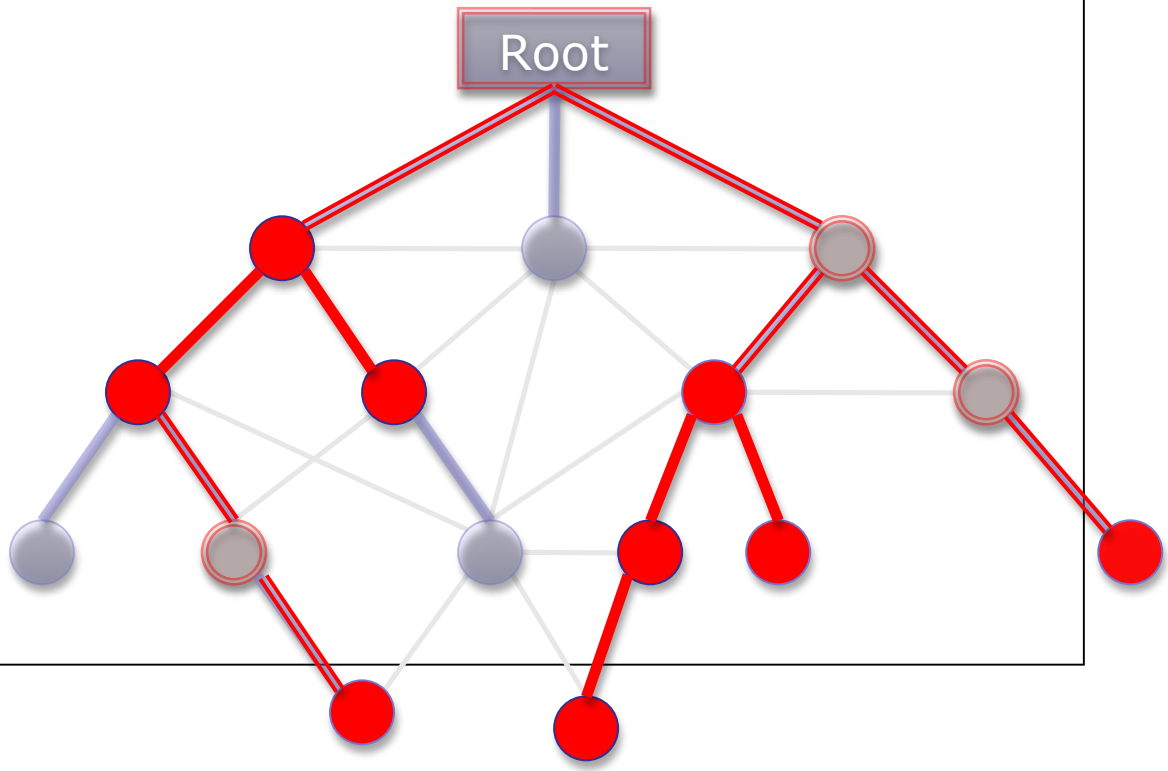

UNIVERSITY OF CAMBRIDGE

Overlay sensor networks

- After the installation of a new application
  - Each node measures their distance to the closest cluster edge lower down the routing tree.

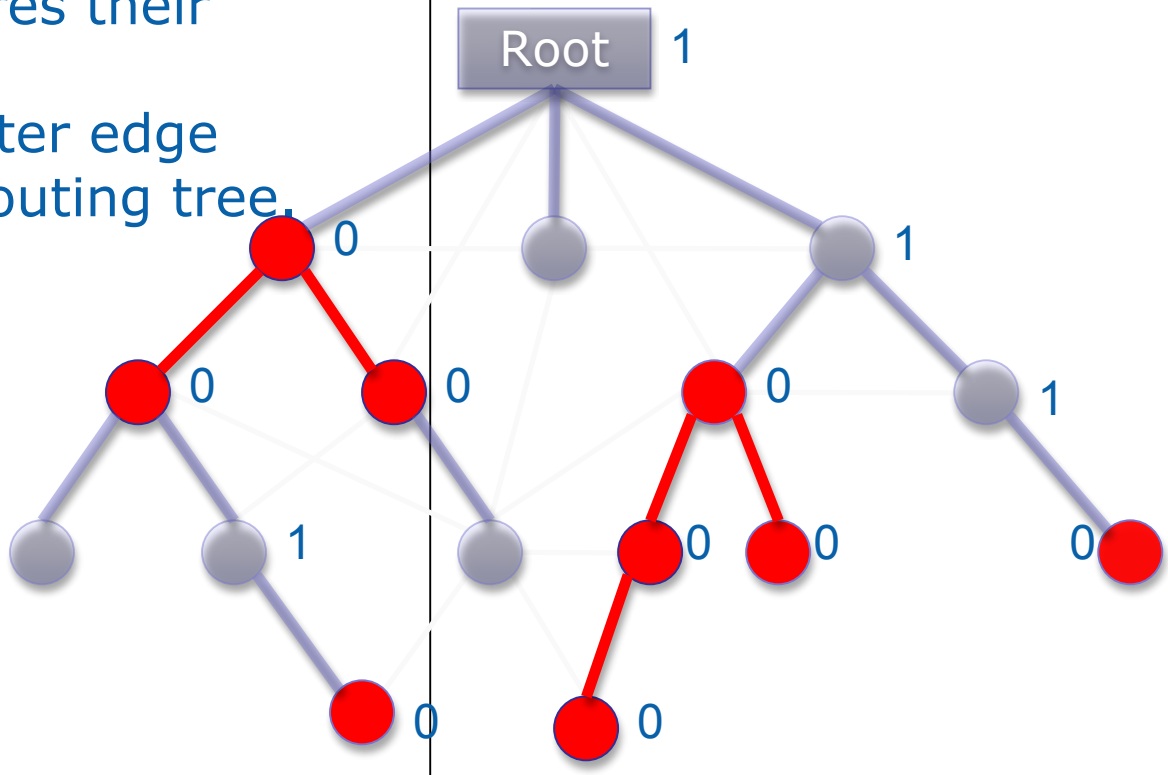

UNIVERSITY OF CAMBRIDGE

Overlay sensor networks

Message_received_from_node (A)

- if (local_distance == 0)
  - accept_message
- else if (distance(A) > (local_distance – 1)
  - send_to_closest_neighbour
- else
  - send_to_all_links(x) with distance(x)>distance(A)
  - send_to_parent

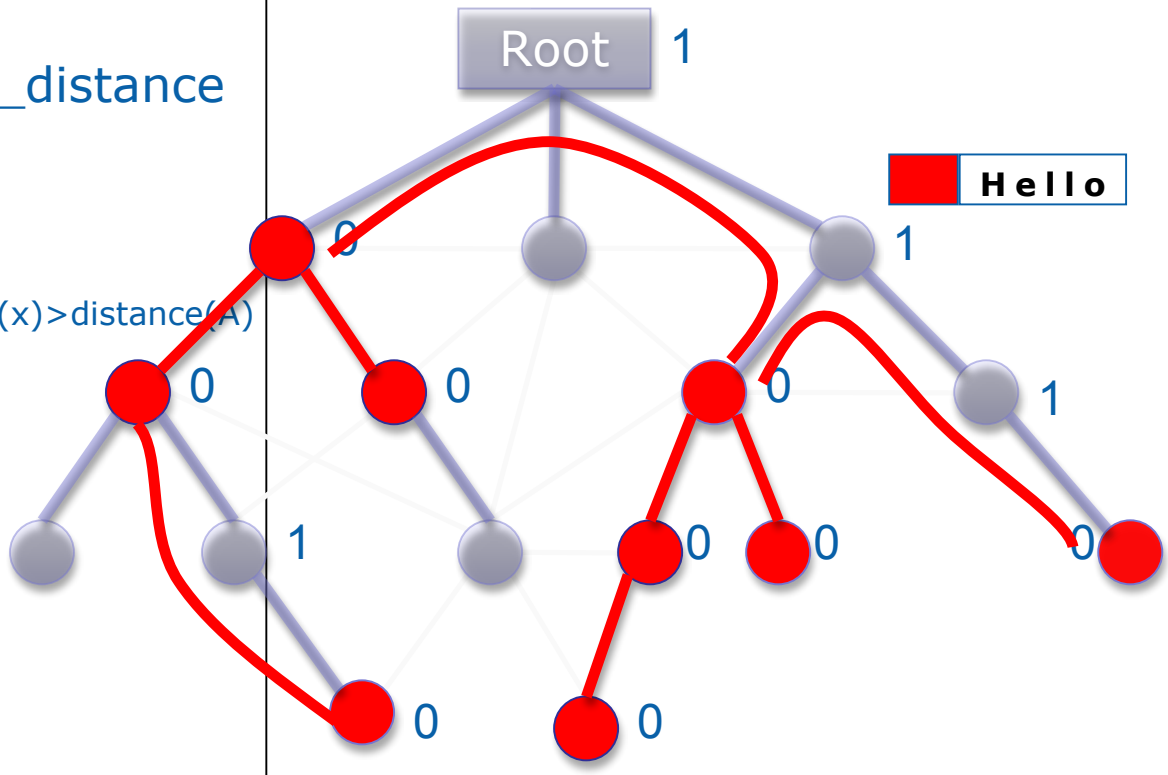

Other applications

UNIVERSITY OF CAMBRIDGE

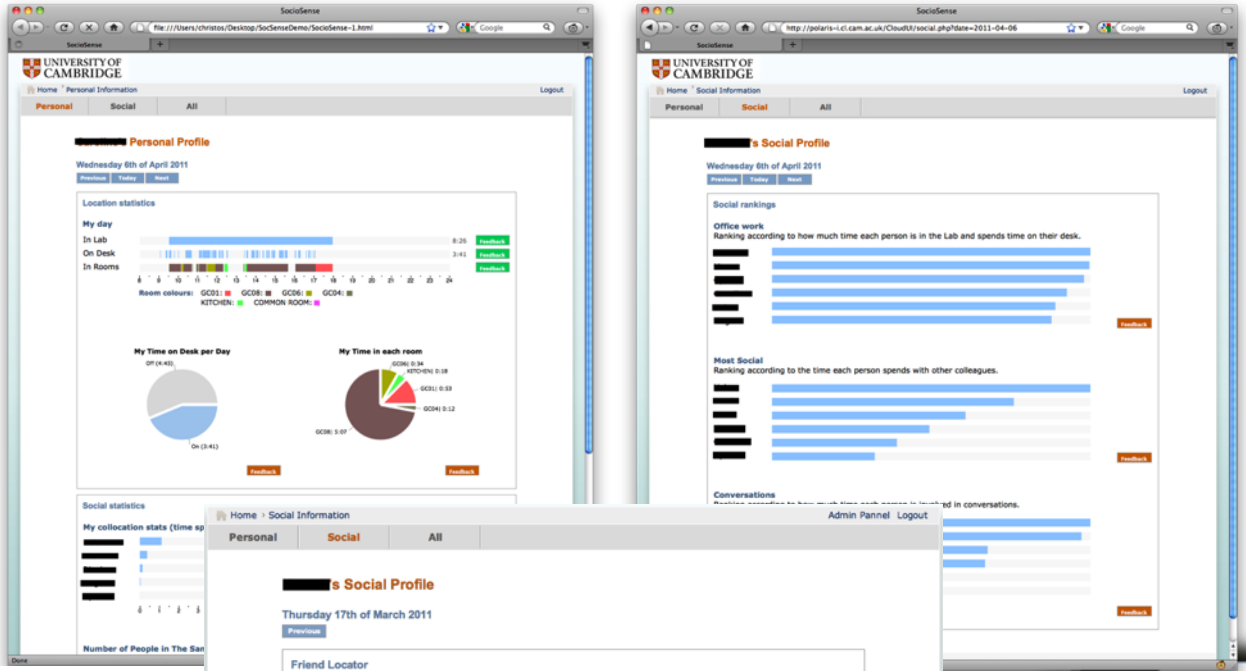

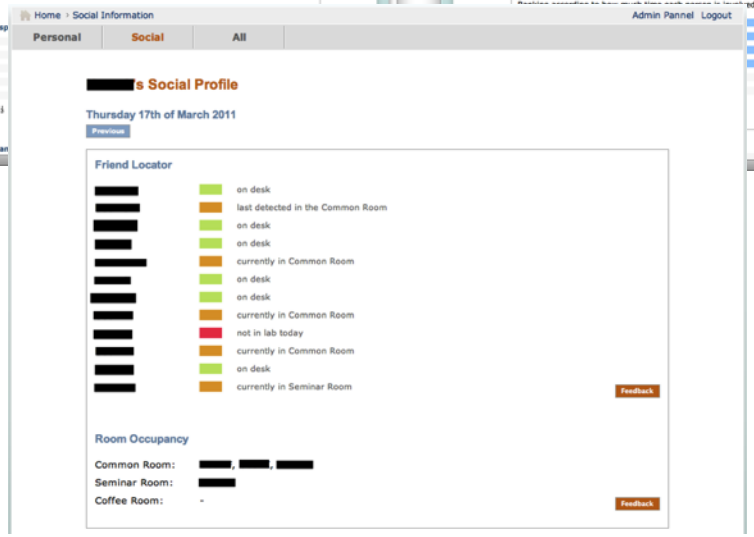

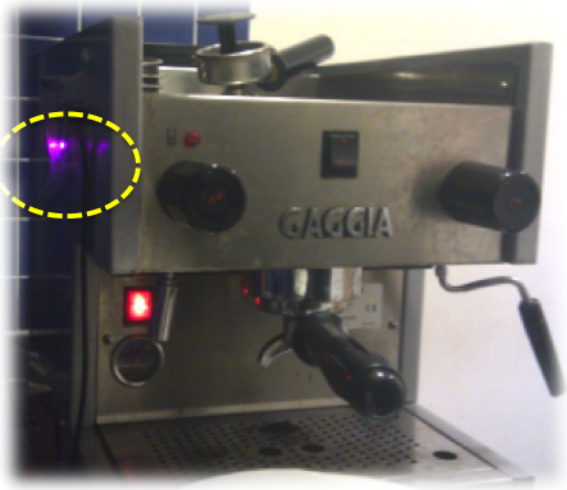

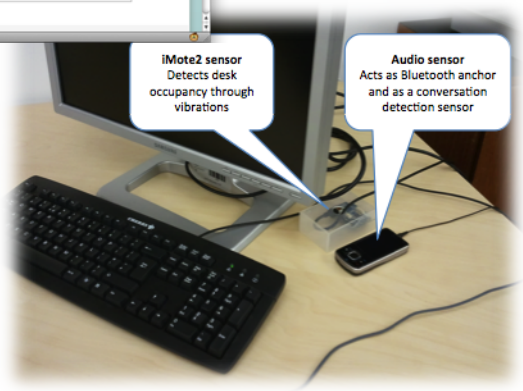

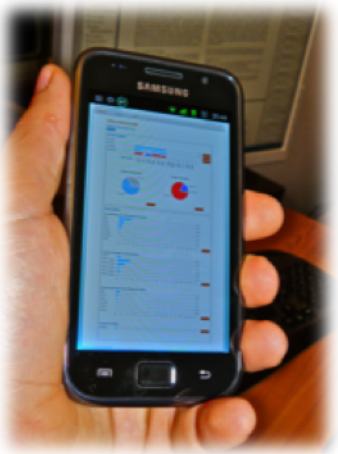

Take away points

UNIVERSITY OF CAMBRIDGE

- We need to break the current, fit-for-just-one-purpose, model of SN.
- Allow users to install applications on shared infrastructure
- Provide the APIs, platform and tools that are required to support this.
- Do privacy right.

Future work

UNIVERSITY OF CAMBRIDGE

- Future work
  - Mobile phones
  - Federation
  - Privacy/security issues
  - Economic model

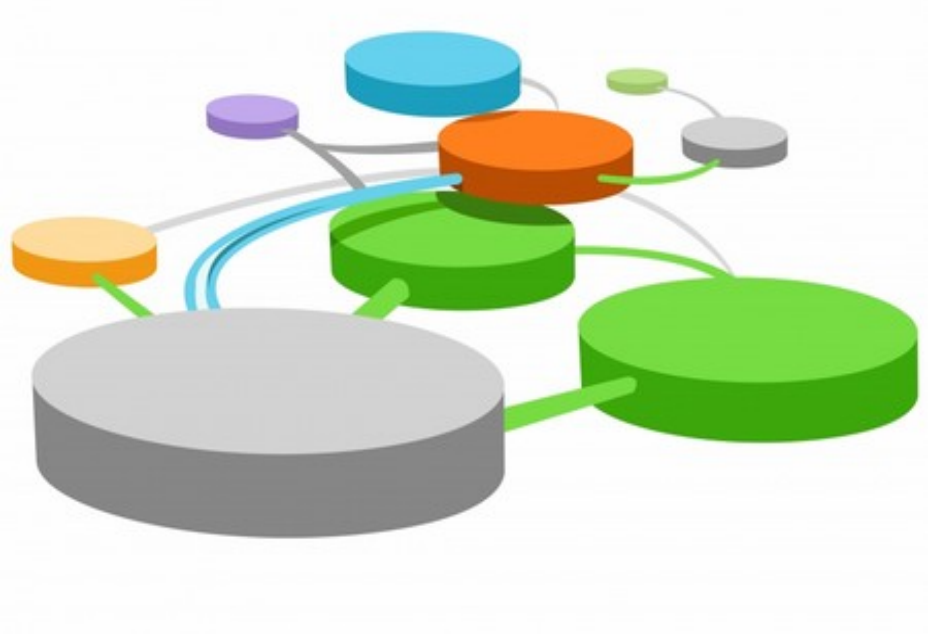

UNIVERSITY OF CAMBRIDGE

Thanks!
Questions?

Download Link:
http://www.cl.cam.ac.uk/research/srg/netos/fresnel/