

ICN – WHAT NEXT?

Q-ICN workshop 20.08.2014

FROM FORECAST AND PROMISES...

We All Know About Video: Staggering Numbers

- Over 4 billion hrs of videos watched on YouTube every month
 - 72 hrs uploaded on YouTube every minute
 - 70% of traffic from outside US
- The 2012 Olympics broke all records
 - BBC delivered 2.8 petabytes on its busiest day, 700Gb/s during the B. Wiggins' gold
- 74 mins average BBC iPlayer TV usage per week
 - 1.6 mio daily iPlayer viewers in July 2011
- ...in all this, mobile usage just started to take off!
 - YouTube mobile traffic tripled in 2011

...With Staggering Forecasts (Cisco)

- Annual global IP traffic will reach the zettabyte threshold by 2015
- The average smartphone will generate 1.3 GB of traffic per month in 2015 (26x)
- In 2015, there will be 6 million Internet households worldwide generating over a terabyte per month in traffic
- By 2012 Internet video will account for over 50 percent of consumer Internet traffic

Hypothesis

*A systems approach that operates on **graphs of information** with a **late** (as late as possible) binding to a location at which the **computation** over this graph is going to happen, enables the full potential for **optimization!***

This systems approach requires to marry information & computation (and with it storage) into a single design approach for any resulting distributed system

What Are The Promises of this?

- More Resilient and robust
- More Flexible & more Efficient, possibly greener
- Better aligning interests (e.g., economic, security, social)
 - What about more private (if wanted)?
- New services, such as in lifestyle management in pervasive healthcare

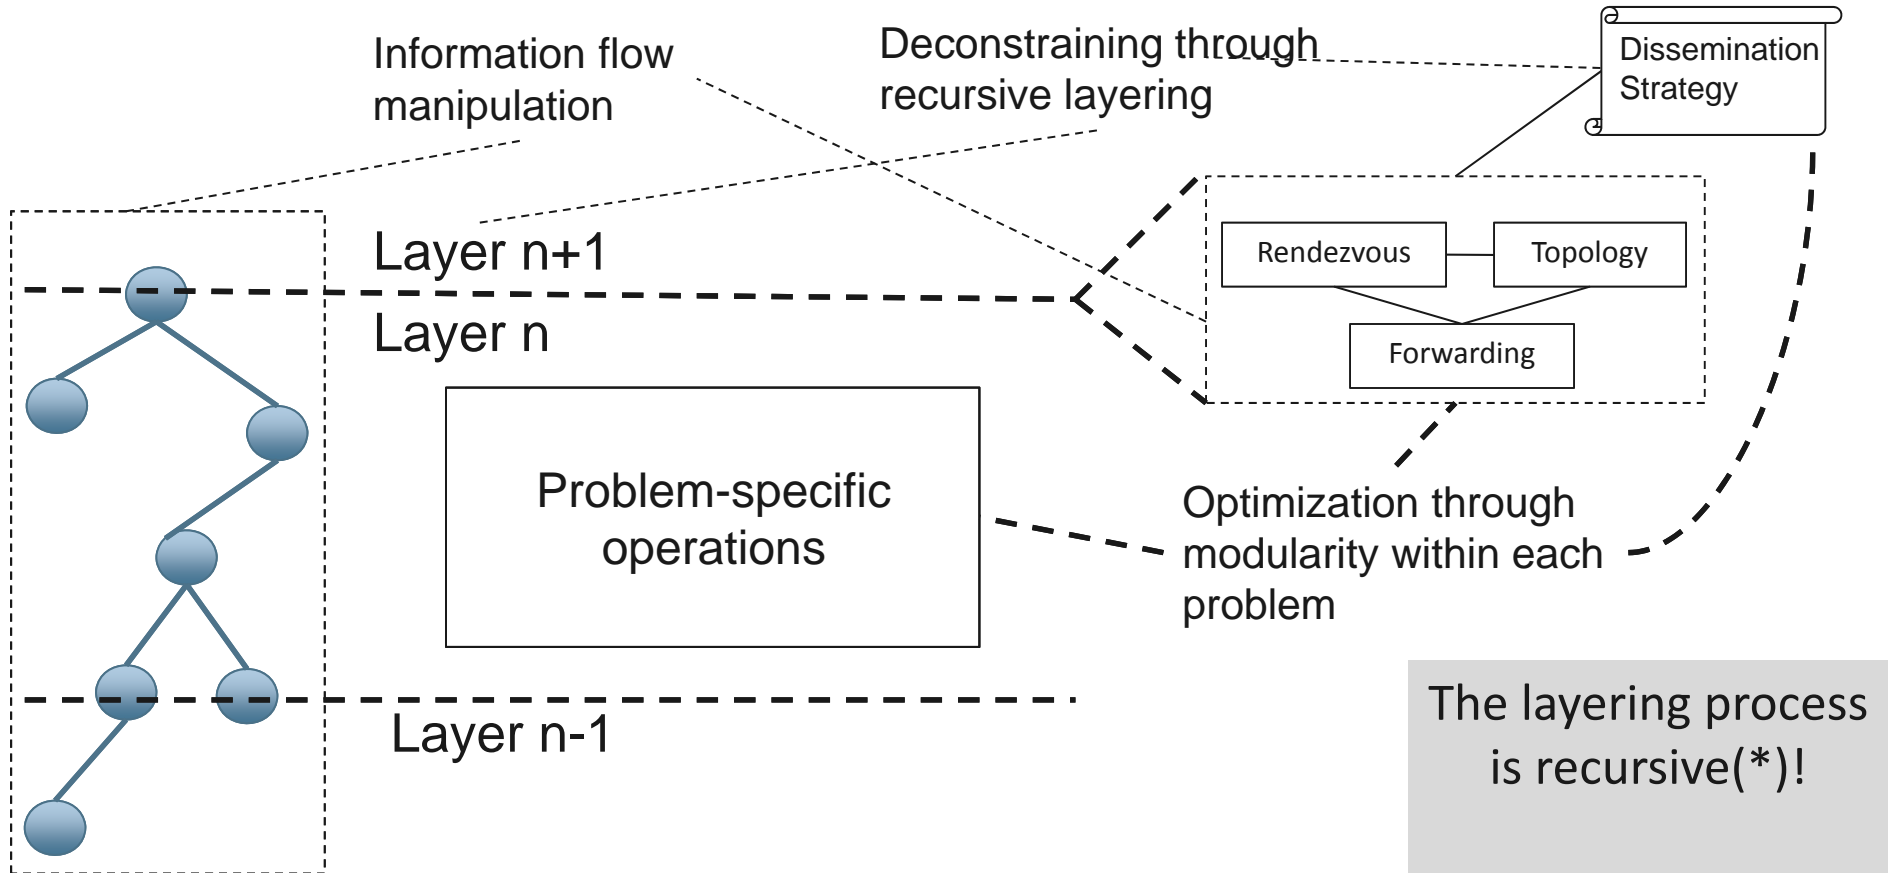
TO BE CLEAR: We have NO conclusive evidence on any of this...but early indications!

**...OVER STARTING POINTS AND
COMMON CONCEPTS**

Possible Design Tenets...

- Provide means for identifying individual information (items)
 - Can be done via labeling or naming (latter requires mapping mechanism)
- Provide means for scoping information
 - Allows for forming DAGs (directed acyclic graphs) of information
- Expose service model
 - Can be pub/sub
- Expose core functions
 - Rendezvous, topology management, and forwarding
- Common dissemination strategy per sub-structure of information
 - Define particulars of functional implementation and information governance

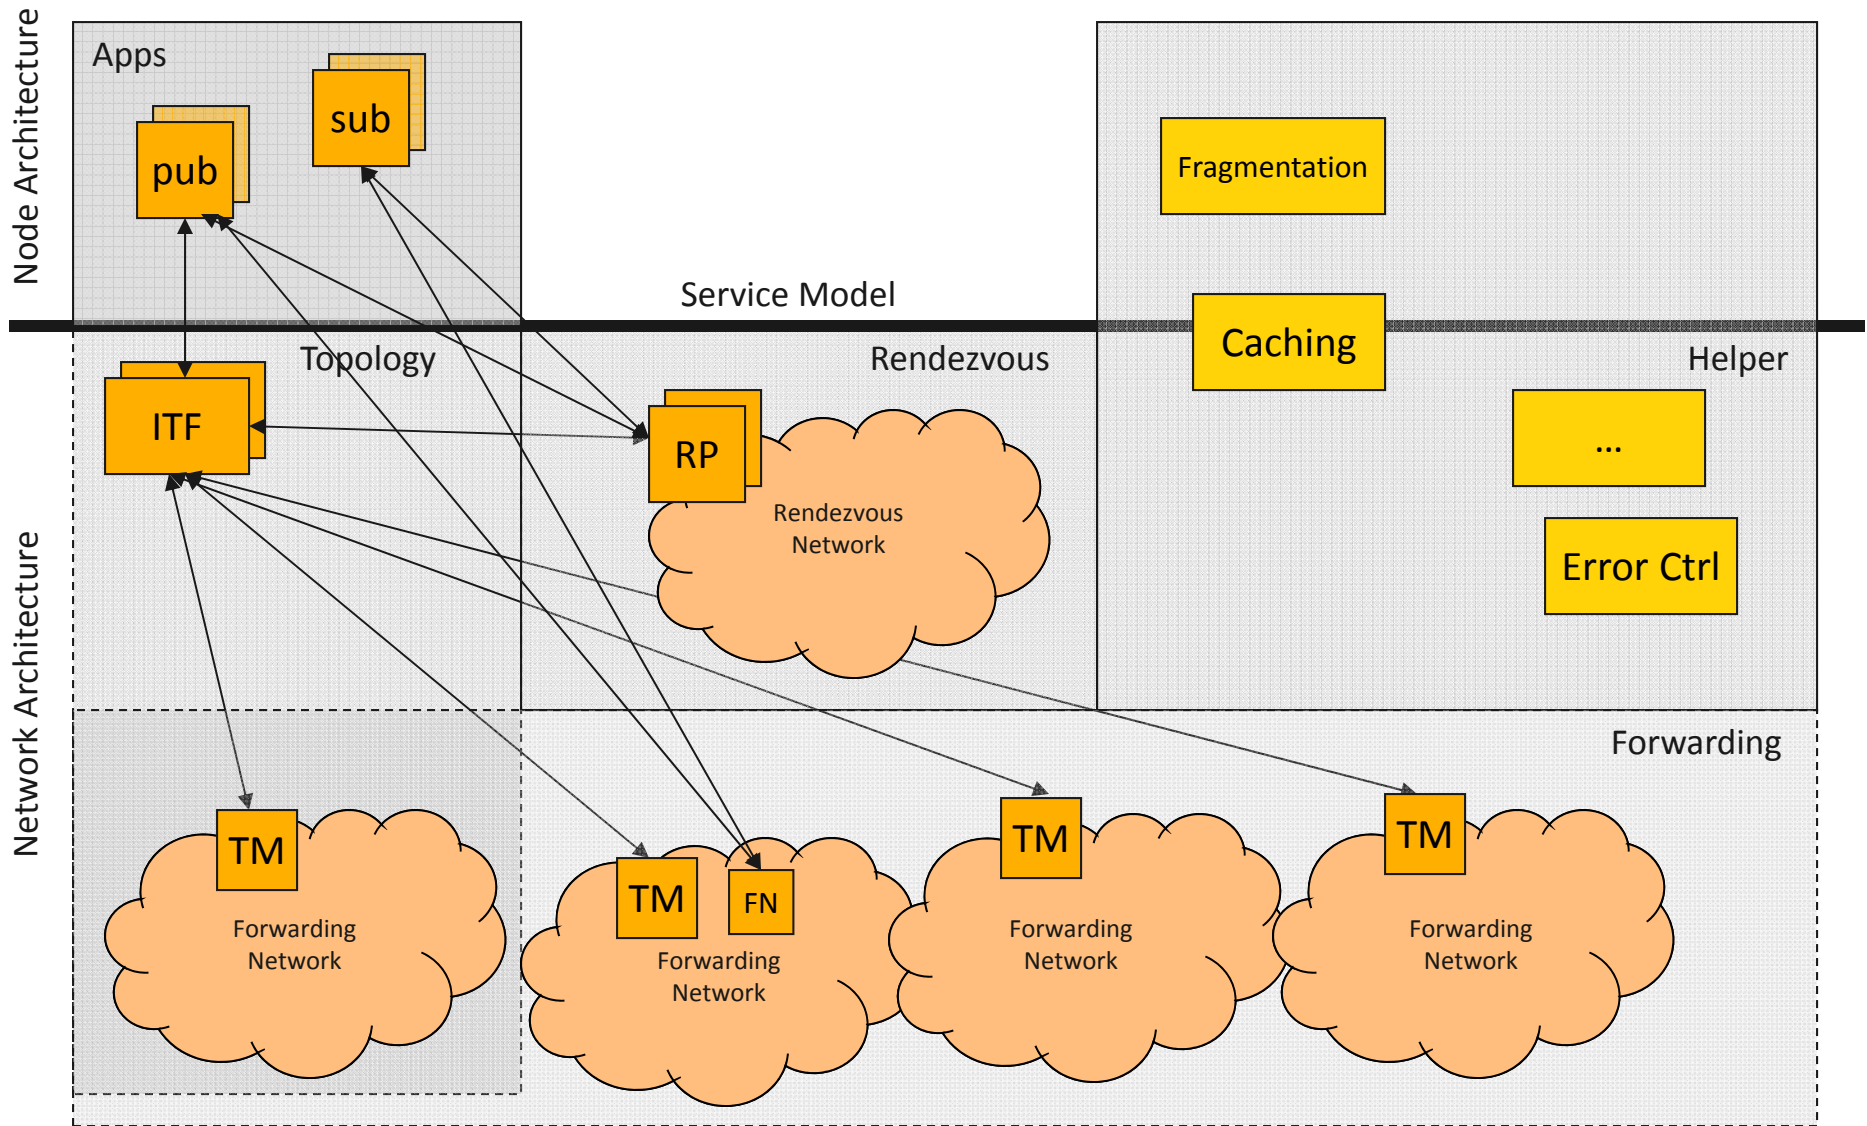
Possible Layered Model



(*) REF: DAY, J. Patterns in Network Architecture - A Return to Fundamentals. Prentice Hall, 2008

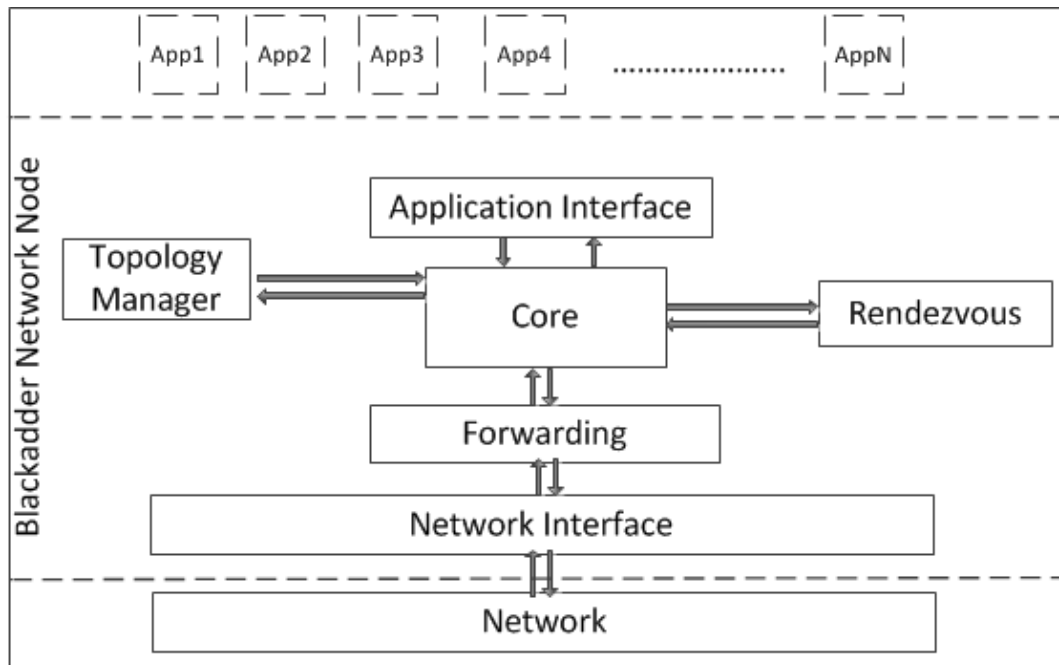
Possible Global Architecture(s)

RP : Rendezvous point
 ITF : Inter-domain topology formation
 TM : Topology management
 FN : Forwarding node



...TO PROTOTYPES AND TESTBEDS

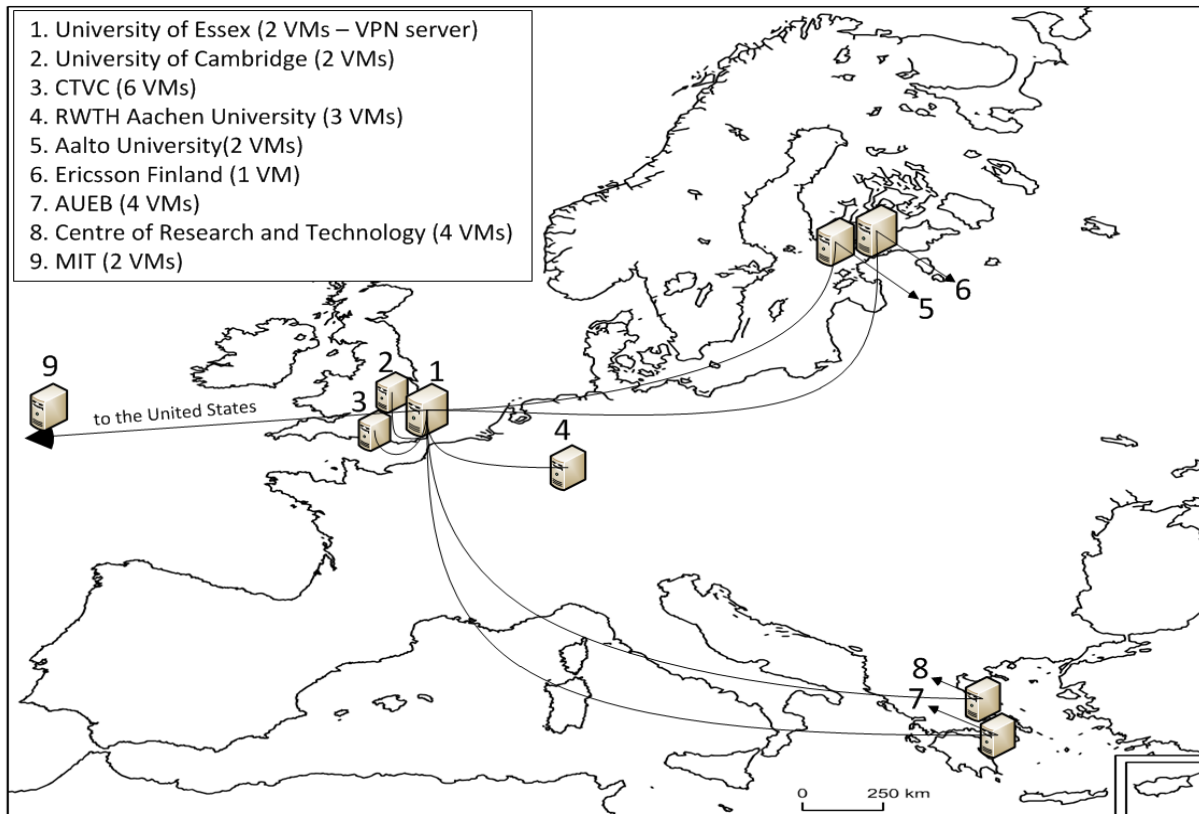
Possible Prototypes: Blackadder@PURSUIT



- Implements design tenets
- Based on **Click** router platform (*)
 - Easy user/kernel space support
 - Easy porting onto other OSes
 - Easy plugging into ns-3
- Available at <https://github.com/fp7-pursuit/blackadder>
- Domain-local throughput beyond 1GB/s

(*) REF: E. Kohler, R. Morris, B. Chen, J. Jannotti, F. Kaashoek. The click modular router. ACM Trans. Comput. Syst. 18, 3 (August 2000), 263-297.

Possible Test Beds



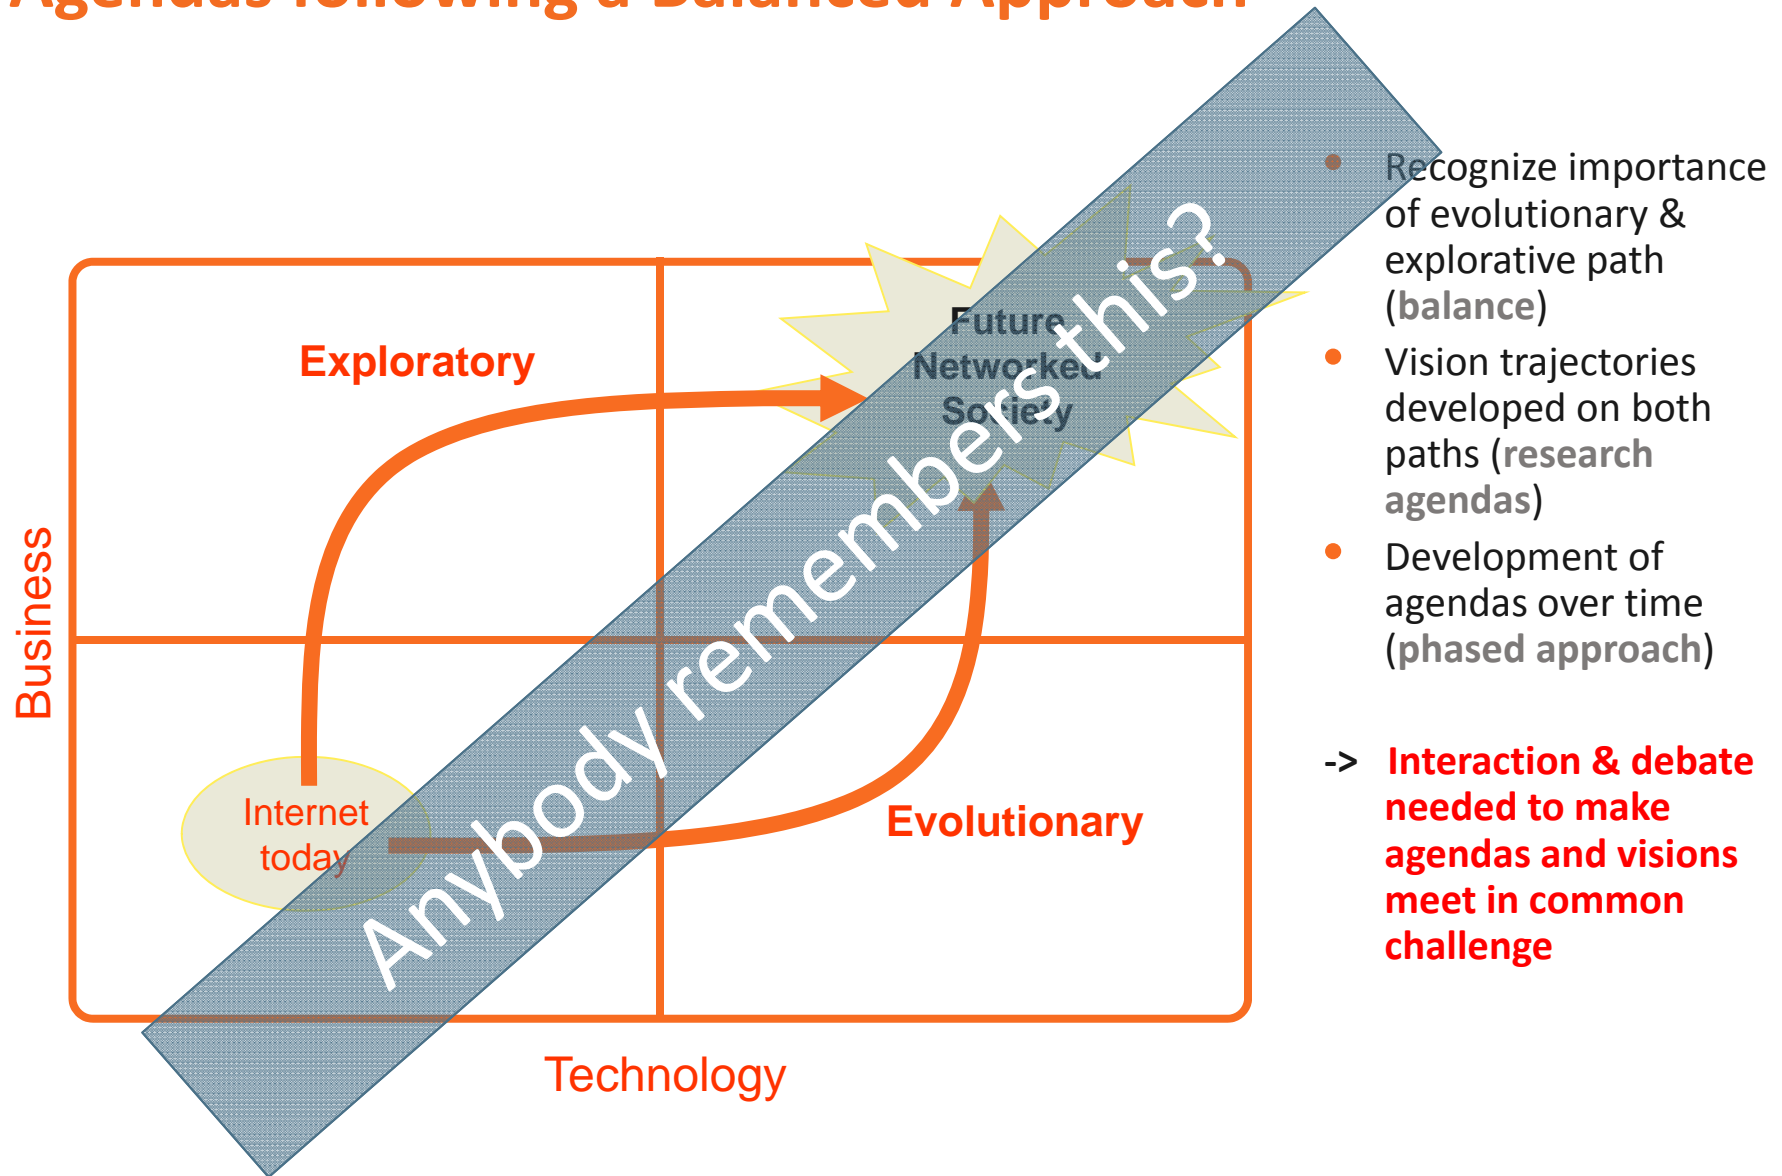
- 10 international sites
 - tunneled via openVPN with configurable topologies

Other test beds

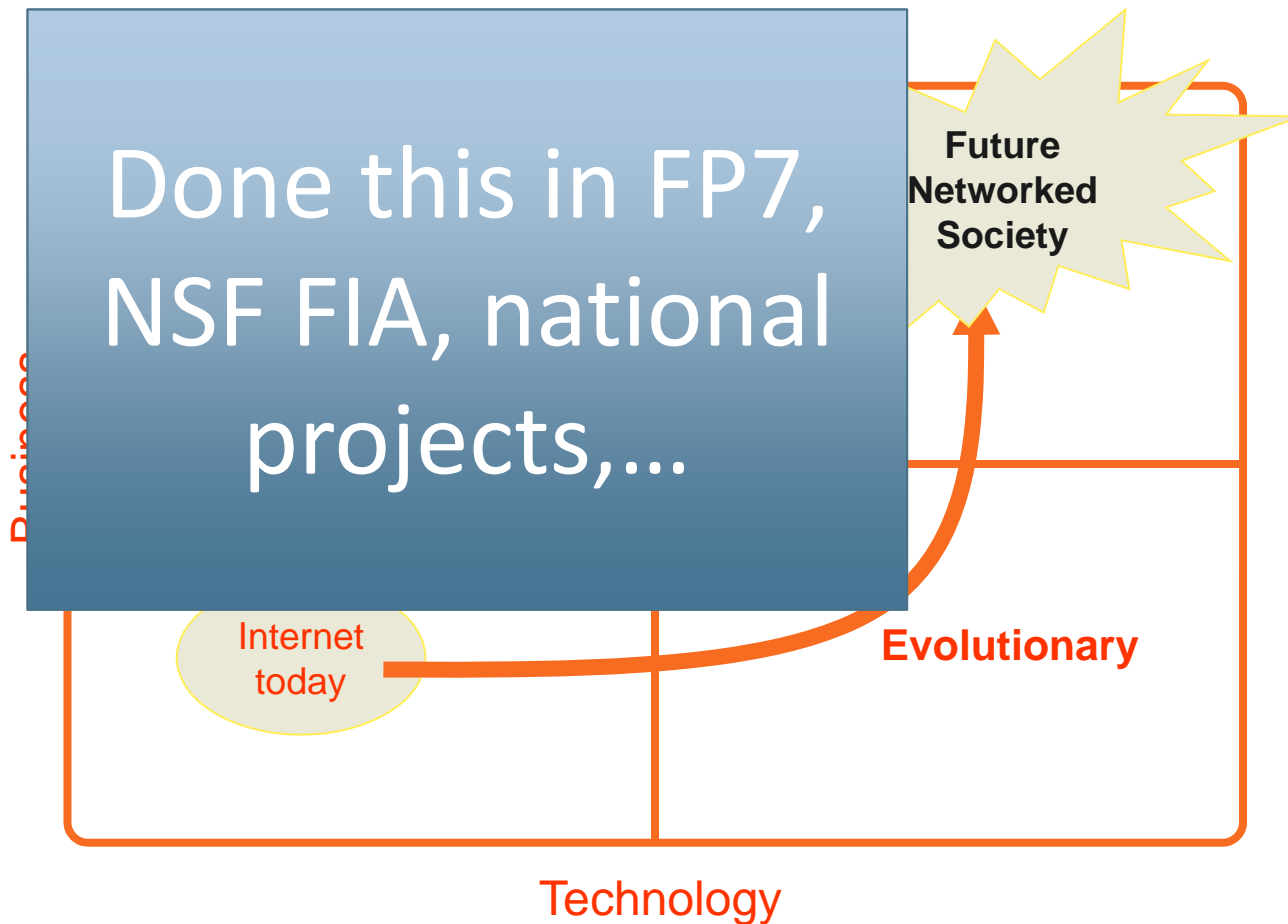
- Dedicated 1GB/s test bed with 15 nodes
- Planetlab (>100 nodes)
- Emulated topologies via ns-3

SO: WHAT IS NEXT?

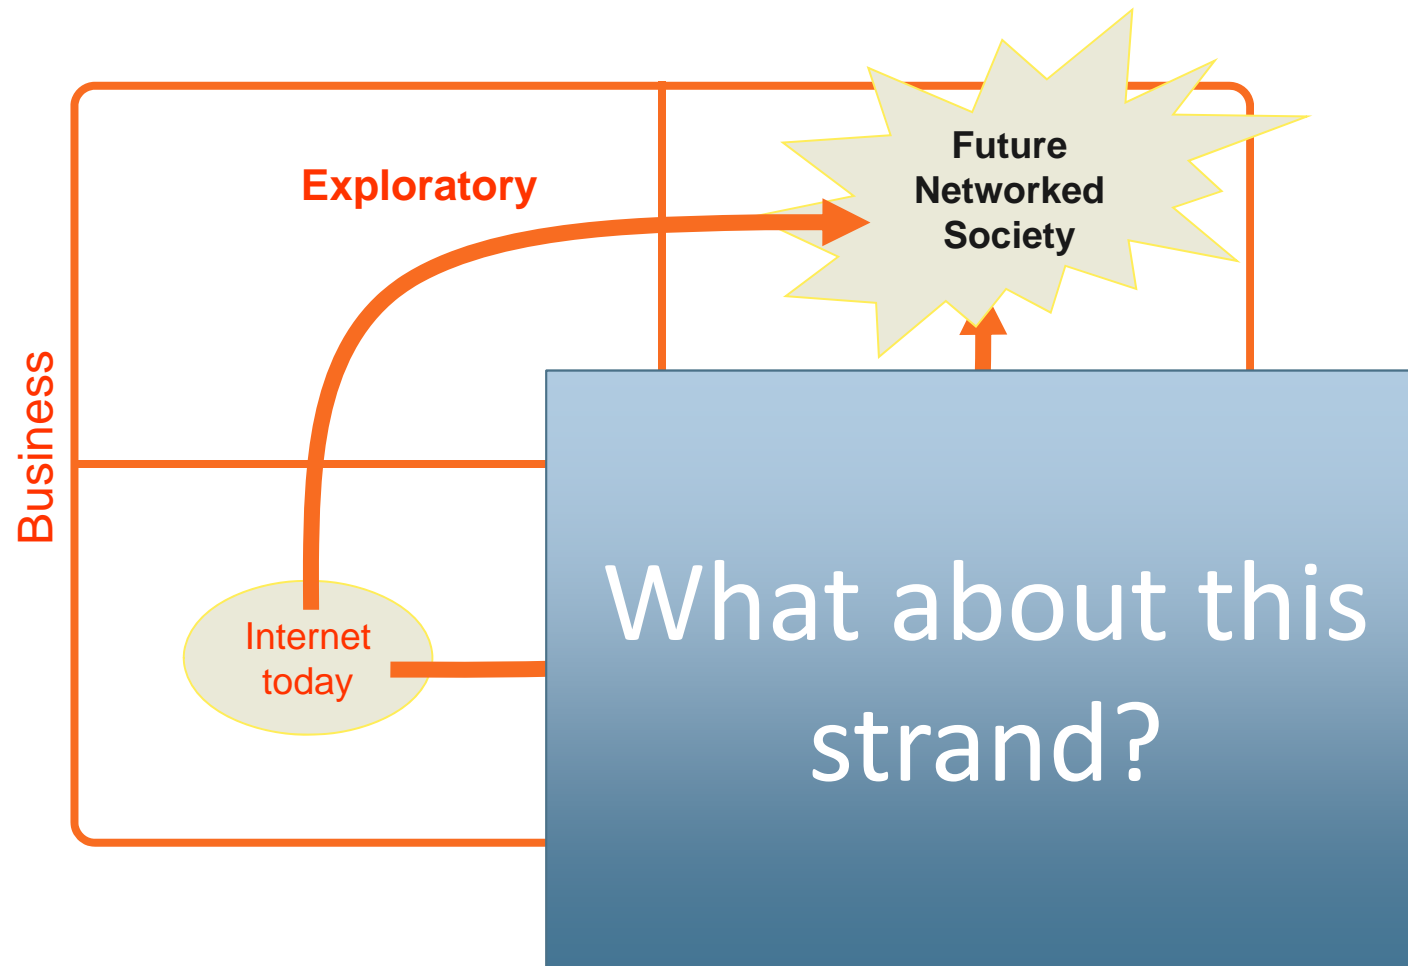
Needed: A Phased Approach for Developing Research Agendas following a Balanced Approach



Needed: A Phased Approach for Developing Research Agendas following a Balanced Approach

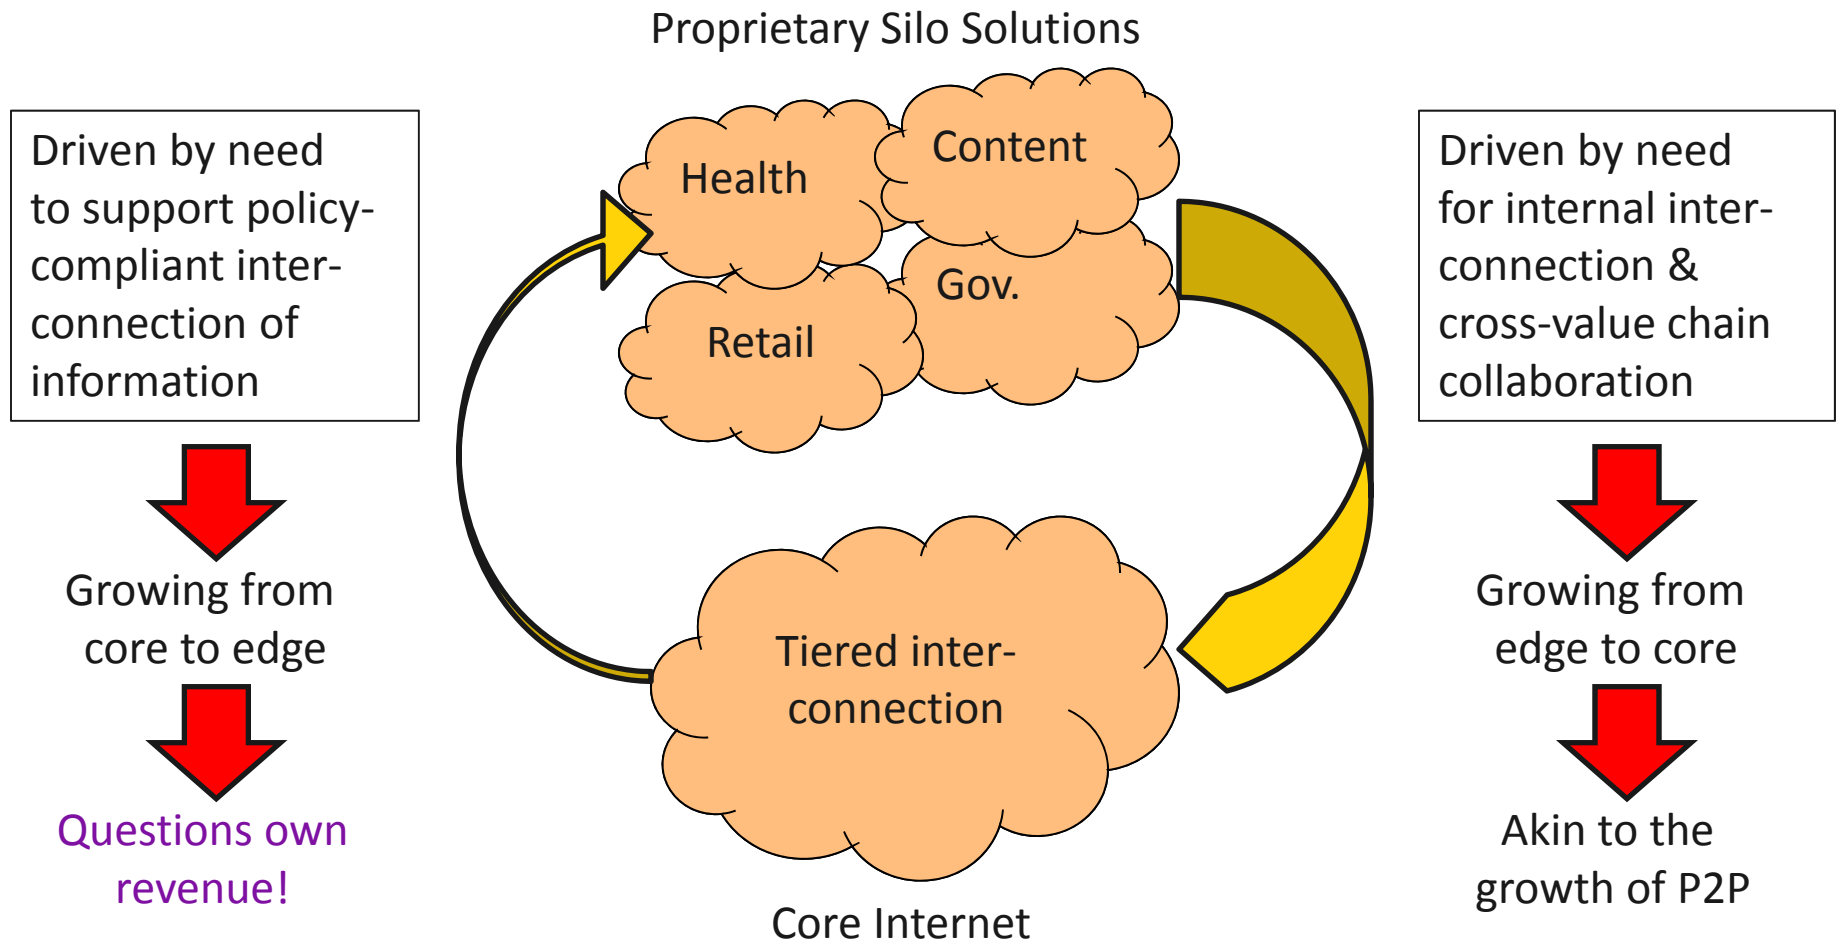


Needed: A Phased Approach for Developing Research Agendas following a Balanced Approach



Core->Edge vs Edge->Core: How likely?

PSIRP study in 2009



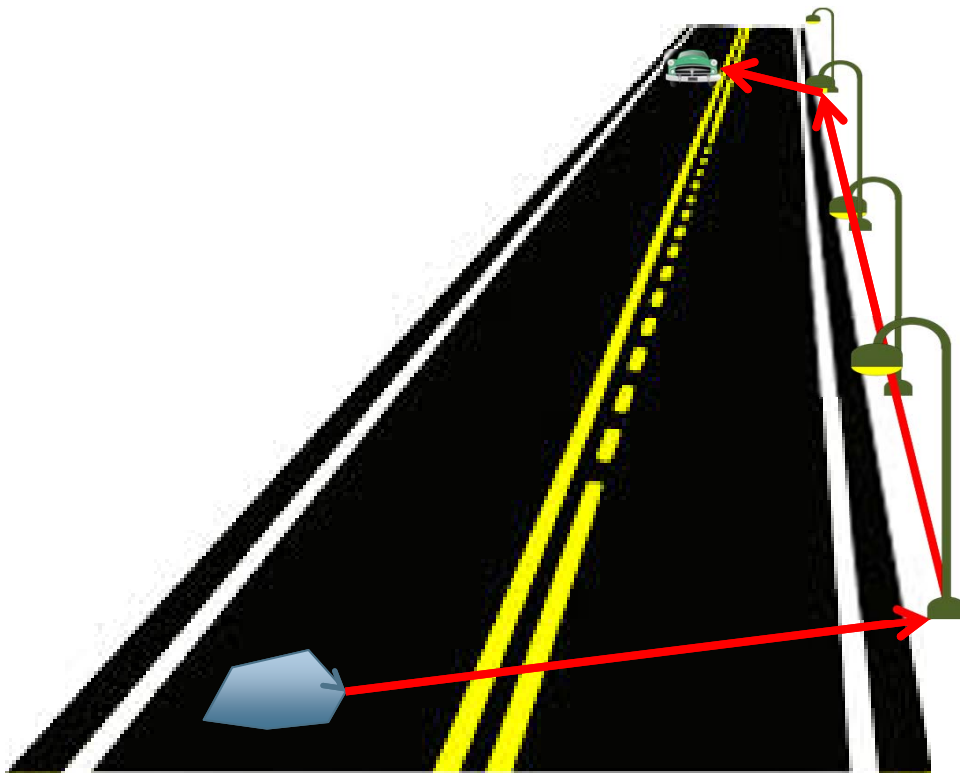
LET'S RE-TELL THE STORY

...AND FORGET ABOUT CONTENT!

SmartAnything: (Some) Drivers

- Disintegration of (product) assets
 - Move from individual product offerings towards aggregation of information markets
 - Requires to establish an (audit) trail of the value that is extracted from the usage of the information
- Efficiency is paramount due to
 - Energy constraints of involved devices
 - Energy budgets imposed by regulation & policy
- Virtualization of ownership
 - **Today:** ownership -> control
 - **Tomorrow:** virtual ownership -> control


Use Case: Ice Warning System



- Ice warning system with road sensors
- Lamppost gateways for local area sensors
- Application to warn n lampposts ahead for safety maneuvers!
- What am I interested here?
 - E2e connectivity?
- When am I interested?
 - Immediate delivery (**latency!**)
 - Moderate validity
 - **NOTE:** information might arrive after I passed lamppost!
 - Low-level accountability desirable

Use Case: See What You Don't See

- Augmented reality system for drivers and passengers
- Lamppost gateways for local area sensors
- Application to provide audiovisual information from n lampposts ahead for safety but also entertainment purposes!
 - Dangers but also interesting surroundings!



THE OPPORTUNITY OF LATE BINDING

That's How Today's Internet Works

- **Network asks:** "Who do you want to talk to?"
- **You say:** "I want to talk to X!"
- Network will do everything it can (best effort) to ensure that you can talk to (any) X you like!
- **How do I get to information Y?** Either I know that X has Y or I might need to interrogate several candidates!
 - Ineffective since I need a good knowledge of X as a candidate (search engine?) or requires looking for it myself (again, search engine?)

Let's Assume the Following Network

- **Network asks:** "What do you want?"
- **You say:** "I want to retrieve Y!"
- Network will do everything it can (best effort) to provide you with the right Y
 - Getting to candidate(s) X becomes a **late binding** task!
- **How do I get to information Y?** I simply ask the question and assume that the candidates have made the information availability known
 - Does not require any knowledge on my side who has and does not require interrogation of candidates (at least for the immediate answer)

Opportunities of Late Binding

- Does not require E2E address space (just needs temporal relationship)
 - lends it self to localised mobility management
- Allows the network to provide alternative locations to the answer
 - Does not require interrogation by those who want
 - Lends itself to differentiated delivery based on network-level optimization criteria (i.e., the retrieval of Y can be optimized by those who own the resources so that inefficiency is accounted for by those who create it), e.g.,
 - Localized caching of highly sought after information
 - Locally optimal delivery mechanisms as a differentiator
 - Can accommodate situations where originator of the information is unavailable
 - **Archiving** becomes a native service of the network
- Multi-point/path becomes natural mode of communication, not exception,
 - Interesting for cases where question is asked at similar times and answered provided by several parties
- Largely removes the need for (bot-based) search
 - translation of human concepts onto scalable net-level concepts still required

The Usual Question

**What can we do that we can't do today?
Nothing, really!**

- **Assertion:** both models are isomorphic from an end user perspective!
- *HOWEVER: what drives the validity of this assertion is (i) latency and (ii) efficiency!*
- In other words: In a world where (i) and (ii) might matter, the assertion is no longer valid since (i) can make the isomorphism impossible and (ii) undesirable!

Return to Ice Warning Scenario

- Current Internet requires full E2E contacting of lamppost *current+n* further down the road (or some central repository where all values are stored)
 - Implement mini-CDN where lamppost readings are properly 'seeded' (placed in lamppost CDN cache) -> is anybody really proposing this???
 - Implement topological addressing of some kind to shorten time to completion for E2E contacting (do we need some agreement for this?)
- Late binding has lamppost *current+n* publish their ice readings periodically with lamppost *current* subscribing to it
 - Use locally efficient delivery mechanism, independent from network-level information abstraction
 - Use agreed topological information model, independent from network-level delivery mechanism
 - Lamppost-local delivery is fast (time to completion), while delivery across lamppost is less time-critical
- **Opportunities:** latency could make it just possible while being efficient

What are the Main Tenets of a Late Binding System?

- Common information abstraction
 - Agree on information model
 - Agree on operations for retrieving information
- Common core functions
 - Where is information X?
 - What (comms) resources do I need to get X?
 - How do I get X from where it is right now?
- Freedom to optimize the realization of the core functions
 - Building block approach
 - Layering approach (in order to scale)

HOW TO PRACTICALLY MOVE FORWARD

Redo the Internet Architecture as We Know It...

- Boiling the ocean
 - Practically impossible: the core->edge model
 - Where has it gotten us throughout FP7?
- Never done in any other case either
 - IP has been an overlay for many years – still is in many places
- Re-thinking the architecture is closer to what we need
 - A thought experiment, best funded by an agency as basic research?!

Lucky for us: this has been happening for some years now

...or Driving ICN Concepts into the Right Places?

- Societal Challenges
 - M2M, IoT, ...
 - > create the right edge pull to make ICN happen
- Technological renewals at large(r) scale*
 - 5G, programmable networks
 - > create the right islands of deployments to make ICN happen

* BTW: this need for renewal is essentially also a societal challenge, namely that of increasing energy demands for ICT systems!!

5G: AN OPPORTUNITY

What is 5G?

I Imagine a Discussion in the 1930s



“...our next generation of calculators will provide even more digits, faster wheels, and temporary mechanical memory...”



“...just imagine an infinite tape on which you store symbols, read by a recording head...”

Not Alone in the View

Statement of the ETP Whitepaper Draft

...we see the role of 5G as that of providing a **universal communication environment** that enables to **address the wider societal challenges** such as in transport, automotive, societal safety, employment, health, environment, energy, manufacturing and food production. This will be achieved through **flexibly aligning stakeholder incentives** by virtue of being truly **programmable**, secure, dependable, privacy preserving, ubiquitous, and flexible, while minimizing the costs per bit by efficiently harnessing all available communication capabilities and reducing the system power consumption, e.g. by harvesting accessible energy from the environment and other means.

Or in other words: finally and truly turn the Internet into a gigantic, universal computer that executes anything anywhere at anytime*

*"Turing, the Internet and a Theory for Architecture: A (Fictional?) Tale in Three Parts"
at <http://www.sigcomm.org/sites/default/files/ccr/papers/2012/July/2317307-2317317.pdf>

A Fundamental Shift: Truly Universal

3GPP traffic classes

- Conversational
- Streaming
- Interactive
- Background (or anything else)

Optimized environment to provide rather limited services

5G traffic scenarios

- Wide range of traffic requirements
- Often unknown at the time of design
- Utilizing storage, comms and computation resources, driven by Moore's Law

Anything else becomes the standard class!

The wireless industry is at the point of crossing the rubicon of becoming a general, distributed computing platform for everything rather than staying a specialized execution platform for the few

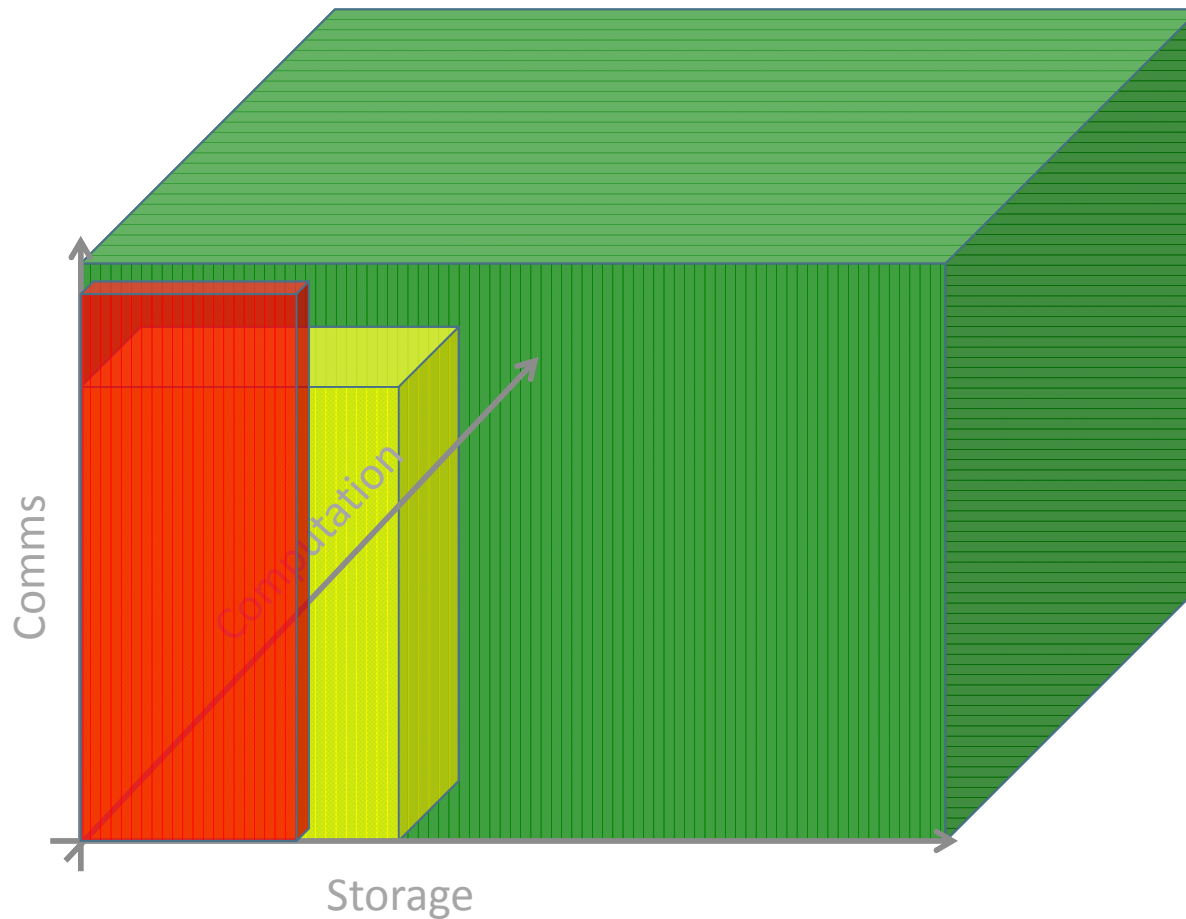
Universal and for Wider Societal Challenges

- Wider set of stakeholders
 - Individual and communities of people.
 - SMEs, Corporations, not-for-profit and social organizations.
 - Digital asset owners, such as public transport and utilities authorities and organizations.
 - **Vertical sectors** like energy, health, manufacturing, robotics, environment, broadcast, content and creative industries, transport, smart cities.
 - **Municipalities** and public administrations.
 - **Public safety** organizations and defense bodies.

-> leading to wider set of requirements (possibly at the same time across the same or parallel 5G systems)

NOTE: could create the edge pull for ICN concepts to appear

One Core Concept: Utilize ANY Resource ANYWHERE at ANYTIME



IP has held us captive in the corner!

Cloud Computing tries to break out of the barriers

This is where we want to get to – total freedom to optimize at runtime!

Orchestration of policies in a **programmable** network substrate enables true **flexibility**

Conclusions

- We've come a long way in ICN
 - From promises over concepts to working prototypes at remarkable speed
- Basic ICN research needs to continue
 - There are unsolved problems
 - ...and there's need for fundamental research in networking!

BUT: ICN needs to address the evolutionary line, thinking revolutionary in order to embed key ICN concepts into various opportunities

- 5G is one such opportunity (but not the only one)!