

# An Information-centric Publish-Subscribe Internet

Addressing the question “Why are we really doing Information-centric Networking?”

Dirk Trossen, Computer Laboratory

# Outline

- WHY are we actually doing ICN?
- WHAT could be the real potential?
- HOW could we do it (VERY high-level!)
- Demo

# We All Know About Video: Staggering Numbers

- 2 billion videos watched on YouTube every day
  - 35 hrs uploaded on YouTube every minute
- 20 mio videos uploaded to FB every month
- 68 mins average iPlayer usage per week
  - 3.1 mio daily iPlayer viewers in Sep 2010
- ...in all this, mobile usage just started to take off!
  - Mobile video traffic will exceed 50 percent of total mobile data for the first time 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

# SO WHAT: The Internet Has Always Been About Information – And It Copes Well With It!

**That is correct... (to a point to be discussed)**


**BUT:** Economics have changed the possible starting points for a design

- Computing and storage resources are NOT scarce anymore; this led to an almost ubiquitous availability of processing and memory
- Information availability has changed attitude of users
  - WHAT is primary, WHO and WHERE mostly secondary! Information is often not locked anymore behind portals

⇒ **Location loses its meaning!**

⇒ **There is desire to fully optimize resource usage (wherever they are)**

# Distributed Systems: A Multi-Dimensional Optimization Space



IP has held us captive in the corner!

Cloud Computing tries to break out (but costly to loose the IP restriction)

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

# 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 and computation (and with it storage) into a single design approach for any resulting distributed system**

**NOTE: the abstraction of an information graph is a choice we made to align application concepts with net concepts**

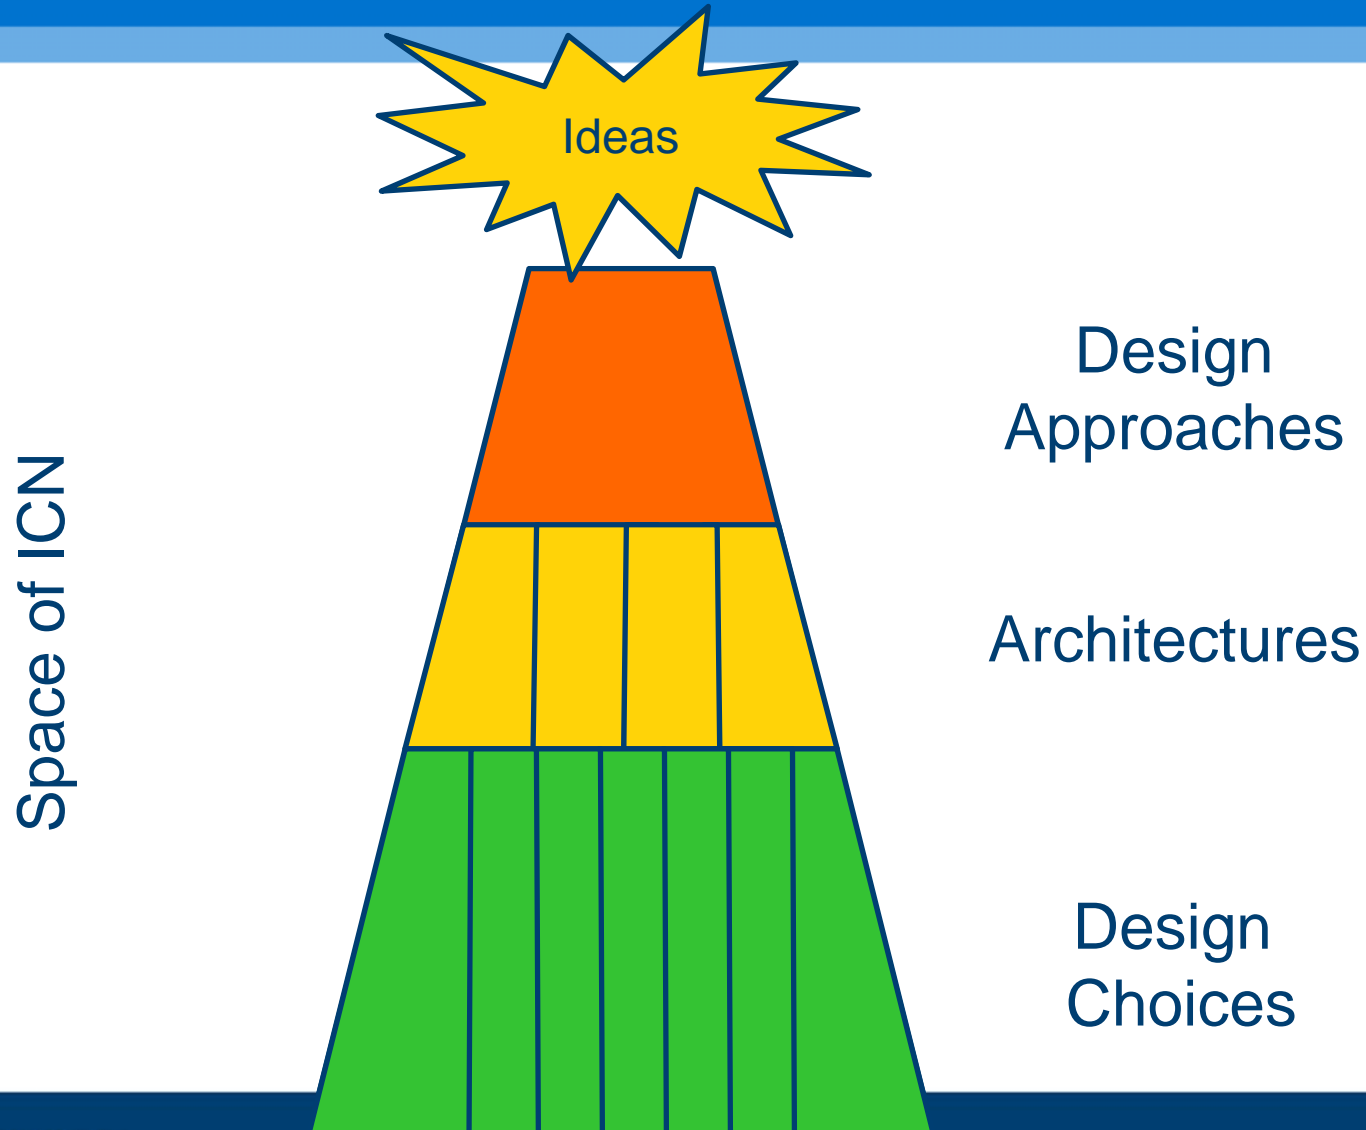
# What Are The Promises of this?

- More Resilient and robust
- More Flexible
- More Efficient
- Greener
- Better aligning interests (e.g., economic, security, social)
  - What about more private (if wanted)?

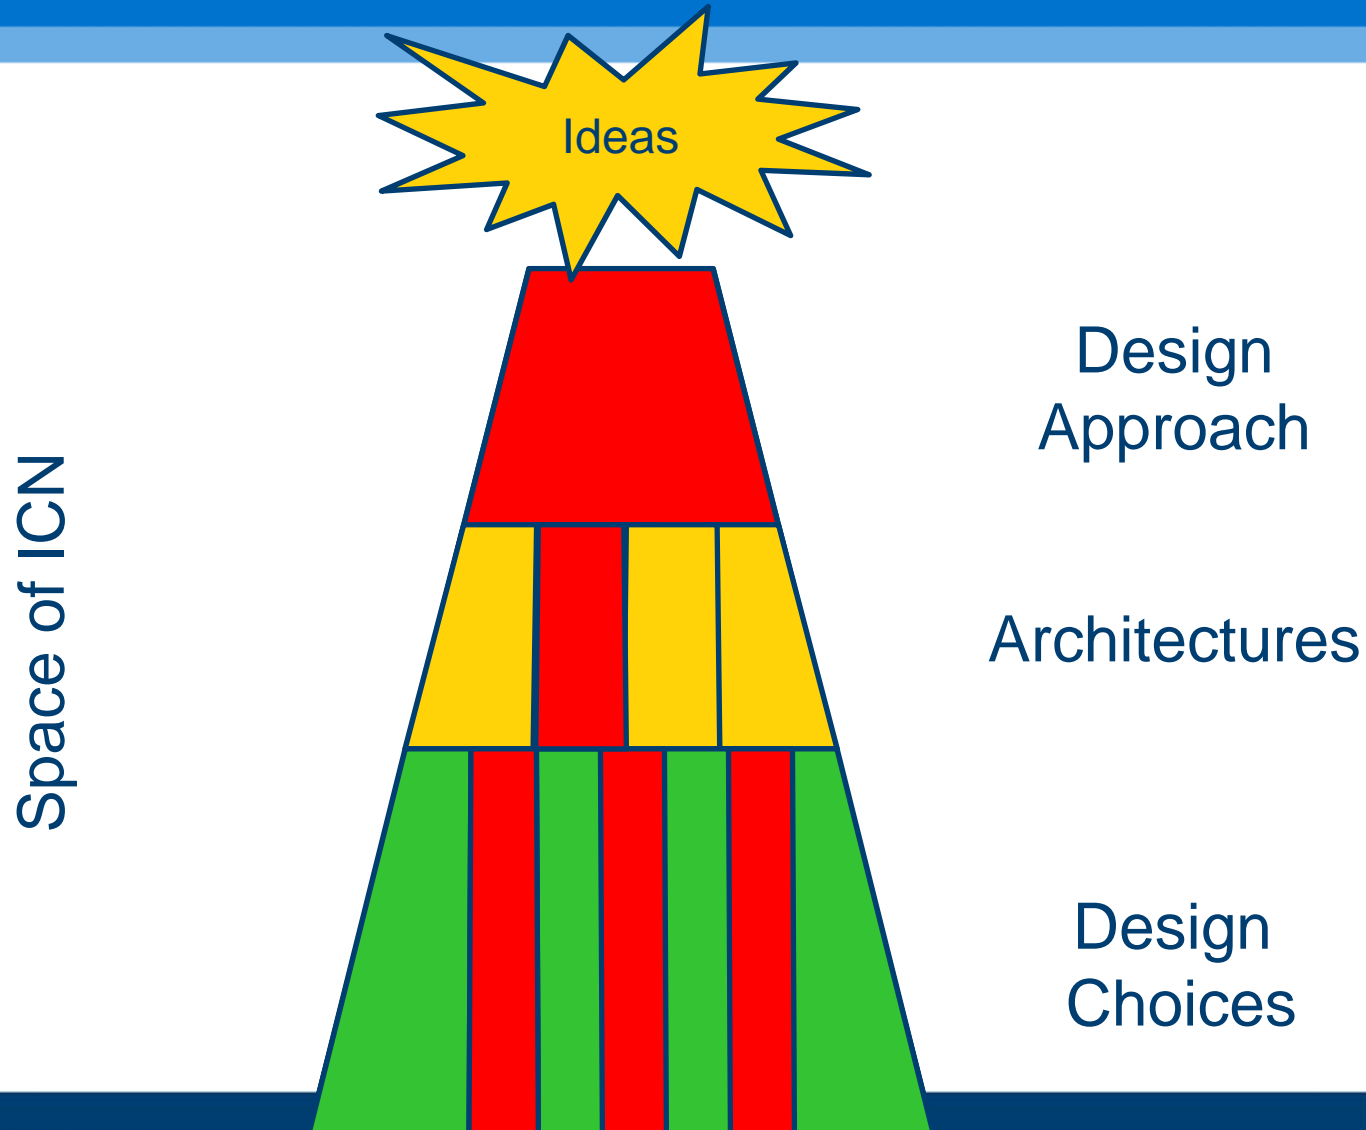
**BEWARE:** We have NO conclusive evidence on any of this...but early indications!



# The Breadth of the Challenge



# The Coverage of PURSUIT



# Starting Point: Solving Problems in Distributed Systems

- One wants to solve a problem, each of which might require solving another problem
  - **Examples:**
    - Send data from A to B(s), involving fragmentation along the link(s)
    - Disseminate a video over a local network
  - Problems involve “*a collection of information that*” an implementation “*can use to decide what to do*”, which is to implement a problem solution (\*)
- > Computation in distributed systems is all about *information dissemination* (pertaining to a task at hand)

\*REF: S. J. Russell, P. Norvig, “Artificial Intelligence: A Modern Approach”, 2nd Edition, Pearson Educ., 1998

# Desired System Properties...

- **Manipulation of (structured) information flows for computational purposes**
  - Expose service model and provide late binding (*WHAT->WHO*)
- **Modularity within a single computational problem**
  - Provide modular core functions (*enable optimization*)
- **Modularity across computational problems**
  - Provide rigorous but flexible layering (*deconstrain constraints*)

REF: CHIANG, M., LOW, S. H., CALDERBANK, A. R., AND DOYLE, J. C. Layering as Optimization Decomposition: A Mathematical Theory of Network Architectures. Proceedings of the IEEE (2007)

## ...Translated into Design Tenets...

- Provide means for identifying individual information (items)
  - Can be done via labeling or naming
- Provide means for scoping information
  - Allows for forming DAGs (directed acyclic graphs)
- 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

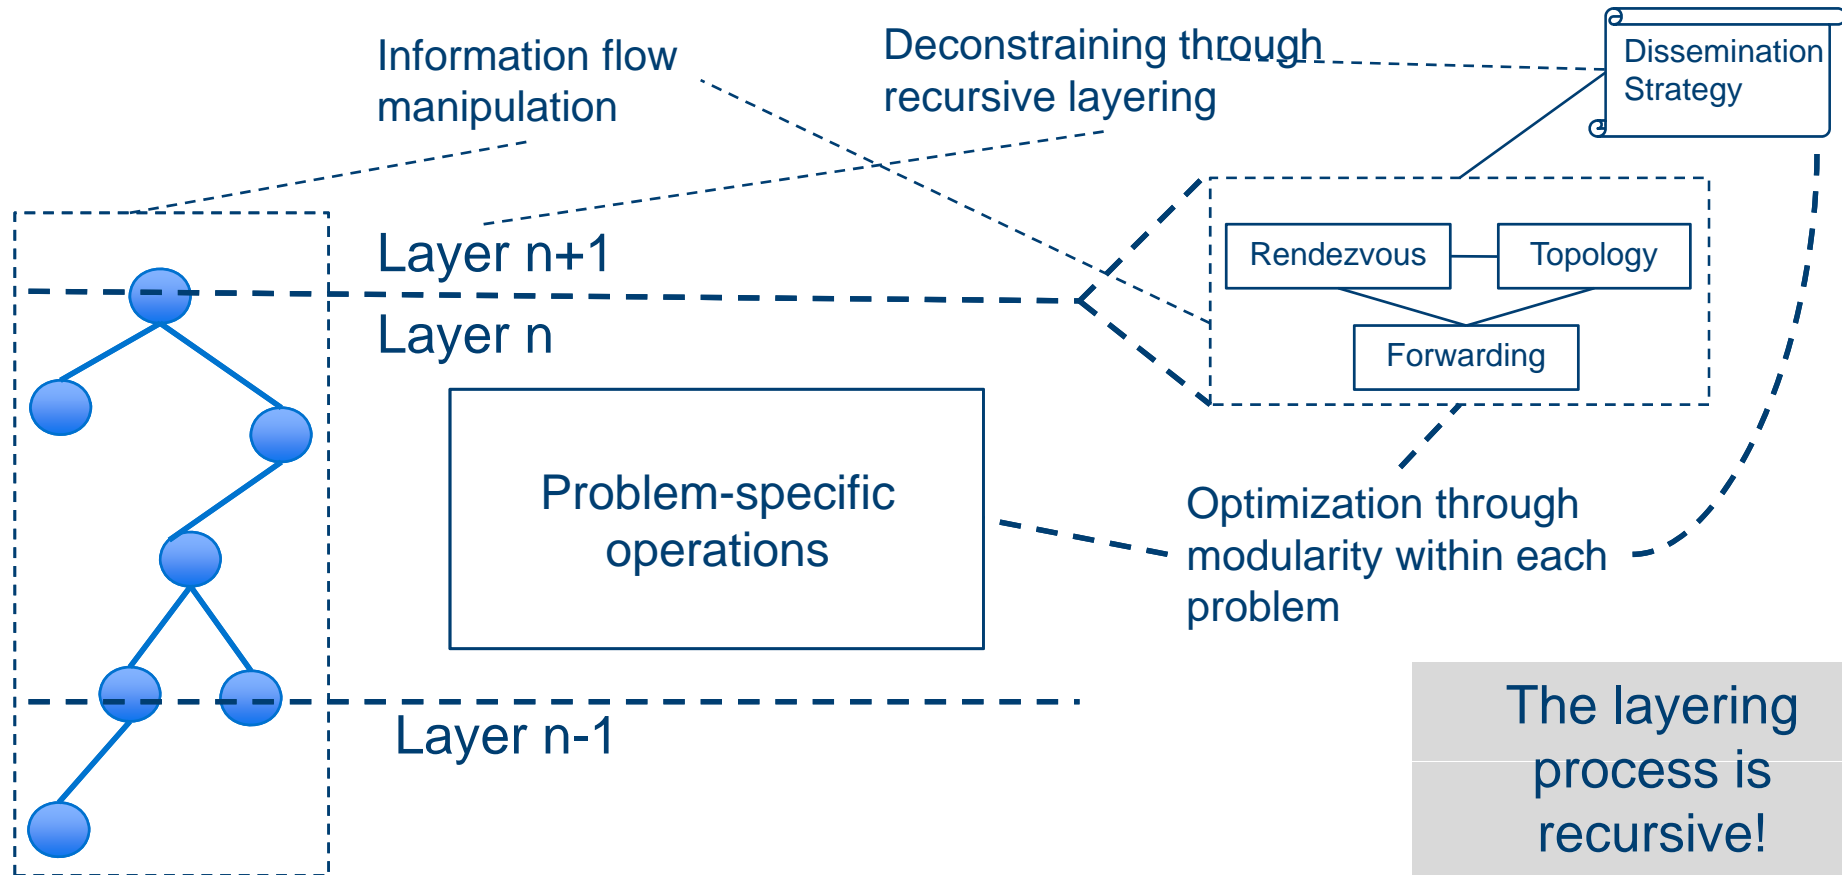
## ...With An E2E Principle...

*The problem in question can be implemented through an assembly of sub-problem solutions, whose individual dissemination strategies are not in conflict with the ones set out by the problem in question.*

- Hence, problems are assembled to larger solutions by recursively applying the scoping invariant of the functional model!
- Conflicts are avoided through design and re-design, e.g., via standards procedures!
- Can extend this to runtime reconciliation!

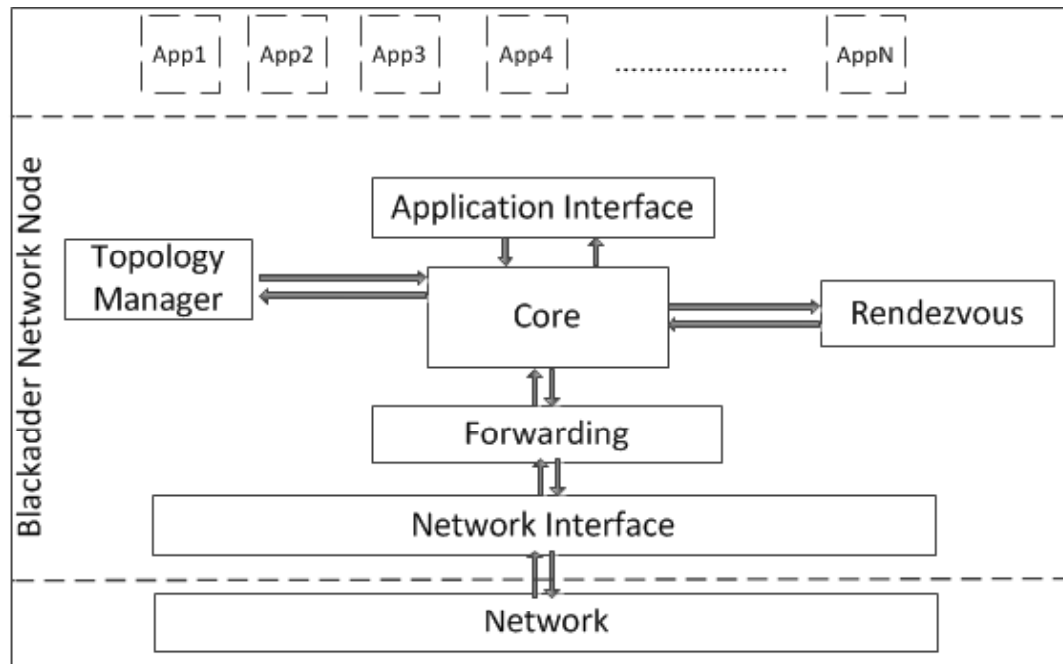
*NOTE: I leave it as a thought exercise to relate this to the IP E2E principle!*

# ...And Placed into a Layered Model



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

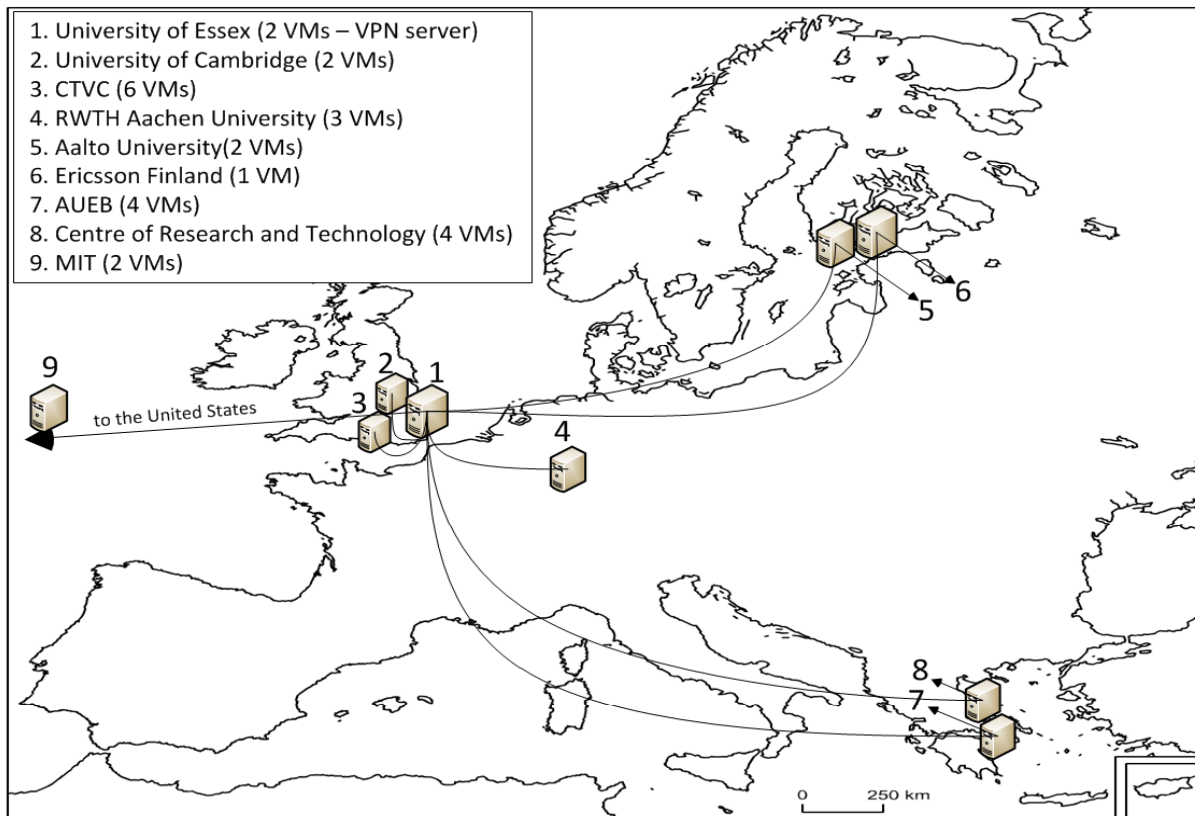
# Prototype: Blackadder



- Implemented as **Click** elements
- User and kernel space
- Implements functional model (i.e., can replace strategies easily)
- Open source at <https://github.com/georgeparisis/blackadder>
- Domain-local throughput reaches 1GB/s



# Test Beds



- 9 sites
- Tunneled via openVPN
- Topologies configurable

Also:

- Dedicated 1GB/s test bed with 15 nodes
- Planetlab

# What is the Take-Away Here?

- Information-centric networking is **NOT** about disseminating information *because the Internet is no good at it!*
- Information-centric networking is about utilizing the entire design space provided by information (& storage) as well as computation
  - Aided by technological developments that made computation (and storage) ubiquitously available
- To get there, we need to re-think how we design/build systems
  - We have first results, working prototypes, and a growing test bed
- **More Info:** <http://www.fp7-pursuit.eu>