
AMICT 2011
Petrozavodsk State University
28.4.2011

PURSUIT
Publish-Subscribe Internet Technology

Professor Arto Karila
Helsinki Institute for Information Technology (HIIT)
Finland
arto.karila@hiit.fi

Fundamentals of the Internet

- **Collaboration**
 - Reflected in forwarding and routing
 - **Cooperation**
 - Reflected in trust among participants
 - **Endpoint-centric services** (mail, FTP, even web)
 - Reflected in E2E principle
- ⇒ **IP with full end-to-end reachability**

VS.

Reality in the Internet Today

- Trust erosion through phishing, spam, viruses
 - Current technology economically favors senders
 - Receivers are forced to carry the cost of unwanted traffic
 - Do endpoints really matter?
 - Information more important
 - Endpoint-centric services move towards information retrieval through, e.g., CDNs
- ⇒ **Ossification of IP-based architecture**

Hypothesis: Importance of Information Requires Information-centric Networks

Application developers care about information concepts

- Creation of information topologies of various kinds

-> Endpoint-centric networking structures are inadequate

- Topological network changes too slow in timescale
- Topological network boundaries often not aligned with information topologies (in particular in cross-organisation scenarios)
- Overlaying possible but restricted in (developer) scalability

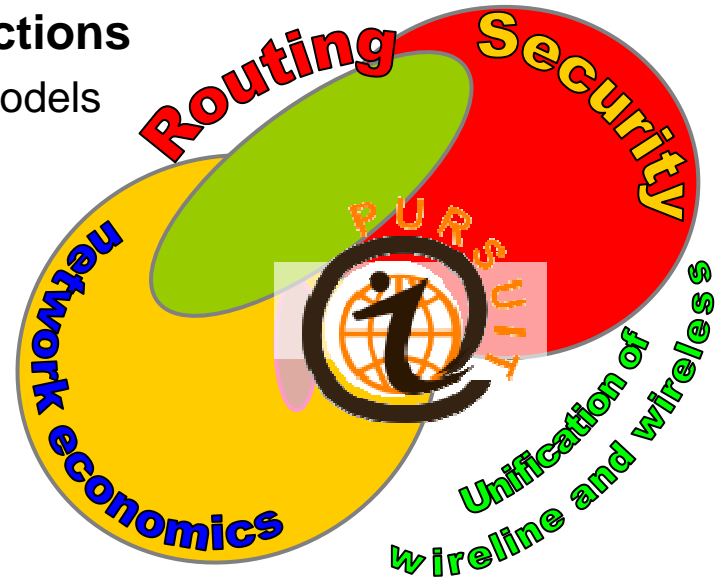
**⇒ If it is all about information,
why not route on information?**



Vision

Envision a system that dynamically adapts to evolving concerns and needs of their participating users

- **Provides an improved impedance match between net & svc/apps**
 - Better aligned with today's application concepts
- **Provides tussle delineation of crucial functions**
 - Better suited for future (unknown) business models
- **Enables optimized sub-architectures**
 - Better suited for various access technologies
- **Provides high performance**
- **Scales to the needs of the Future Internet**



Potential Impacts on End User

- Relevant information at your fingertips
 - Wherever, from whoever, through whatever access, on whatever device
- More natural form of communication
 - Emulates sensing, processing, actuation
- Ability to avoid information overload
 - Tackle attention scarcity problem
- Increased security & privacy
 - Only relevant information gathered & provided to user

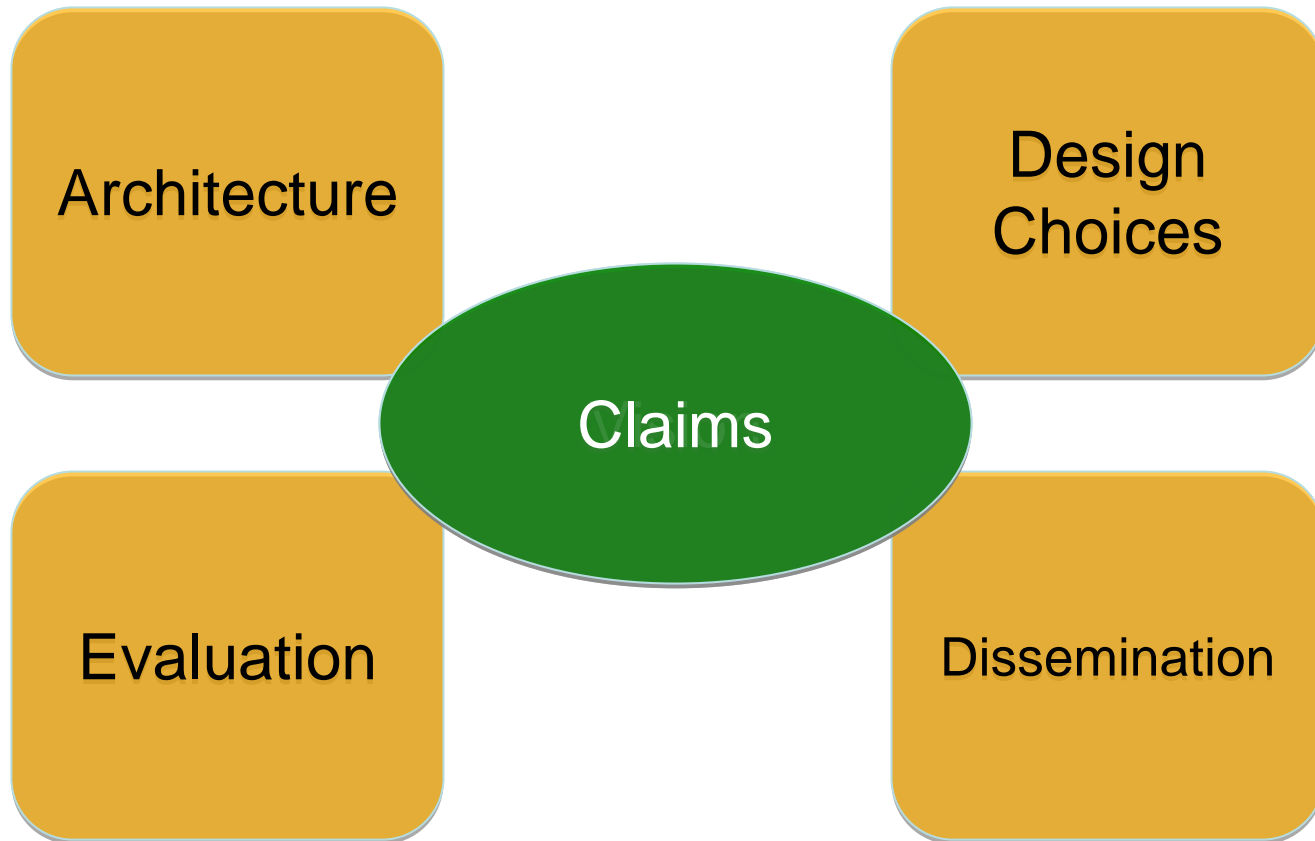


Potential Impacts on Industry

- Increased caching
 - Could lead to price decline for transit traffic (death of Tier-1)
 - Could lead to decline of managed memory (death of CDN)
 - Opportunity to operate networks more efficiently (locally)
- Increased policy compliance
 - Visibility of 'items' on routing level
 - Opportunity of flexible policy enforcement on routing level
- Increased low-level search capability
 - Move from crawling approach to information routing (advance today's search engines)
 - Opportunity to eliminate broken links (increase relevance)
- More flexible services
 - Individual information items allow for faster mash-ups across traditional value chains, e.g., retail, content, health, government
 - Opportunity of real-time collaboration



Our Main Challenges



Our Claims: As Formulated So Far

Design, develop and evaluate a novel information-centric pub/sub-based internetworking architecture that:

- Provides an improved impedance match towards application-level concepts
- Provides tussle delineation of crucial functions
- Enables optimization of sub-architectures
- Provides high performance
- Scales to the needs of the Future Internet

Provide a sound architectural framework for information-centric networking

Main thrusts:

- Invariants and their specific or general viability
- Translate invariants into coherent set of concepts
- Provide a set of coherent architectural arguments for their viability
 - In particular the proper (socio-)economic arguments

Our Main Challenges: Design Choices

Develop a set of design choices to support our architectural claims

Main thrusts:

- Rendezvous throughout all (recursive) levels of the architecture
- Inter-domain topology formation
- Topology management (focus on optical and wireless)
- Forwarding
- Caching & Transport
- Information-centric middlewares

Our Main Challenges: Evaluation

Provide the required proofs for our architectural claims

Main thrusts:

- Implementation (prove that it runs – and performs)
- Simulation (prove that it scales – and performs)
- Socio-economics (prove that its design is viable)
- Economics (prove that it is economically sensible)

Our Main Challenges: Dissemination

Provide the required tools for disseminating our results

Main thrusts:

- Implementation (a tool to create a community)
- Test bed (a place to meet and try out)
- Website (a place to exchange)
- Course material (a tool to educate the new generation)
- Exploitation strategies (a tool to convince the stakeholders)

Publications and presentations are means to an end for all the above

Main Design Principles...

- **Everything is Information**
 - Higher-level information semantics are constructed as graphs of information
- **Information is scoped**
 - Provide a simple mechanism for structuring data and limiting the reachability of information to the parties having access to the particular mechanism that implements the scoping
- **Functionality is scoped**
 - Functions to disseminate information implement a scoped strategy!
- **Scoped information neutrality**
 - Within each scope of information, data is only forwarded based on the given (scoped) identifier
- **Ensure balance of power**
 - No entity is provided with data unless it has agreed to receive those beforehand

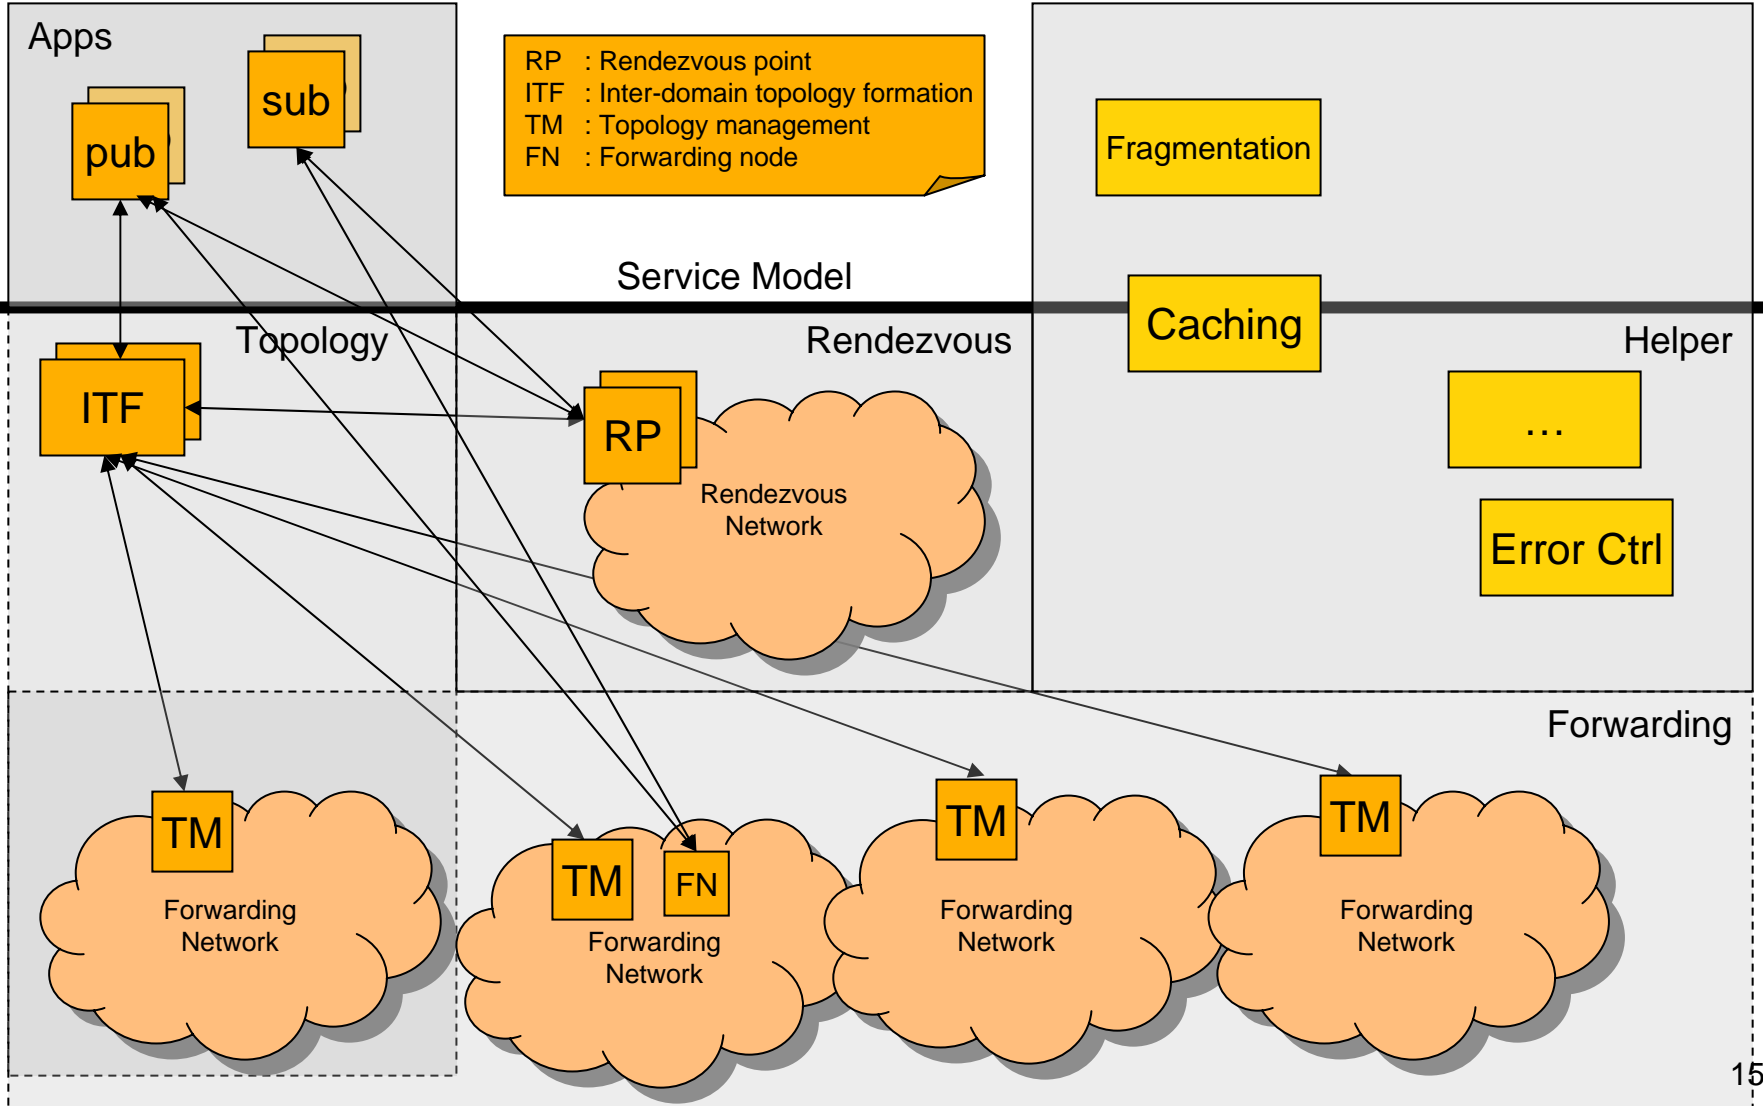
...Translating into Architecture Invariants

- **Flat-label referencing:** identify anything as information
- **Scoping:** group information and functions (including scopes themselves)
- **Pub/sub service model:** anything is delivered by pub/sub
- **Separation of functions:** each scope provides functions for finding (**rendezvous**), constructing (**topology**) and delivering (**forwarding**)
 - Can be implemented jointly for optimization reasons
- **Dissemination strategy per scope:** the implementation of the functions is described by a dissemination per scope
 - Inherited by each sub-scope as default reconciliation

... Leading to A High-Level Architecture

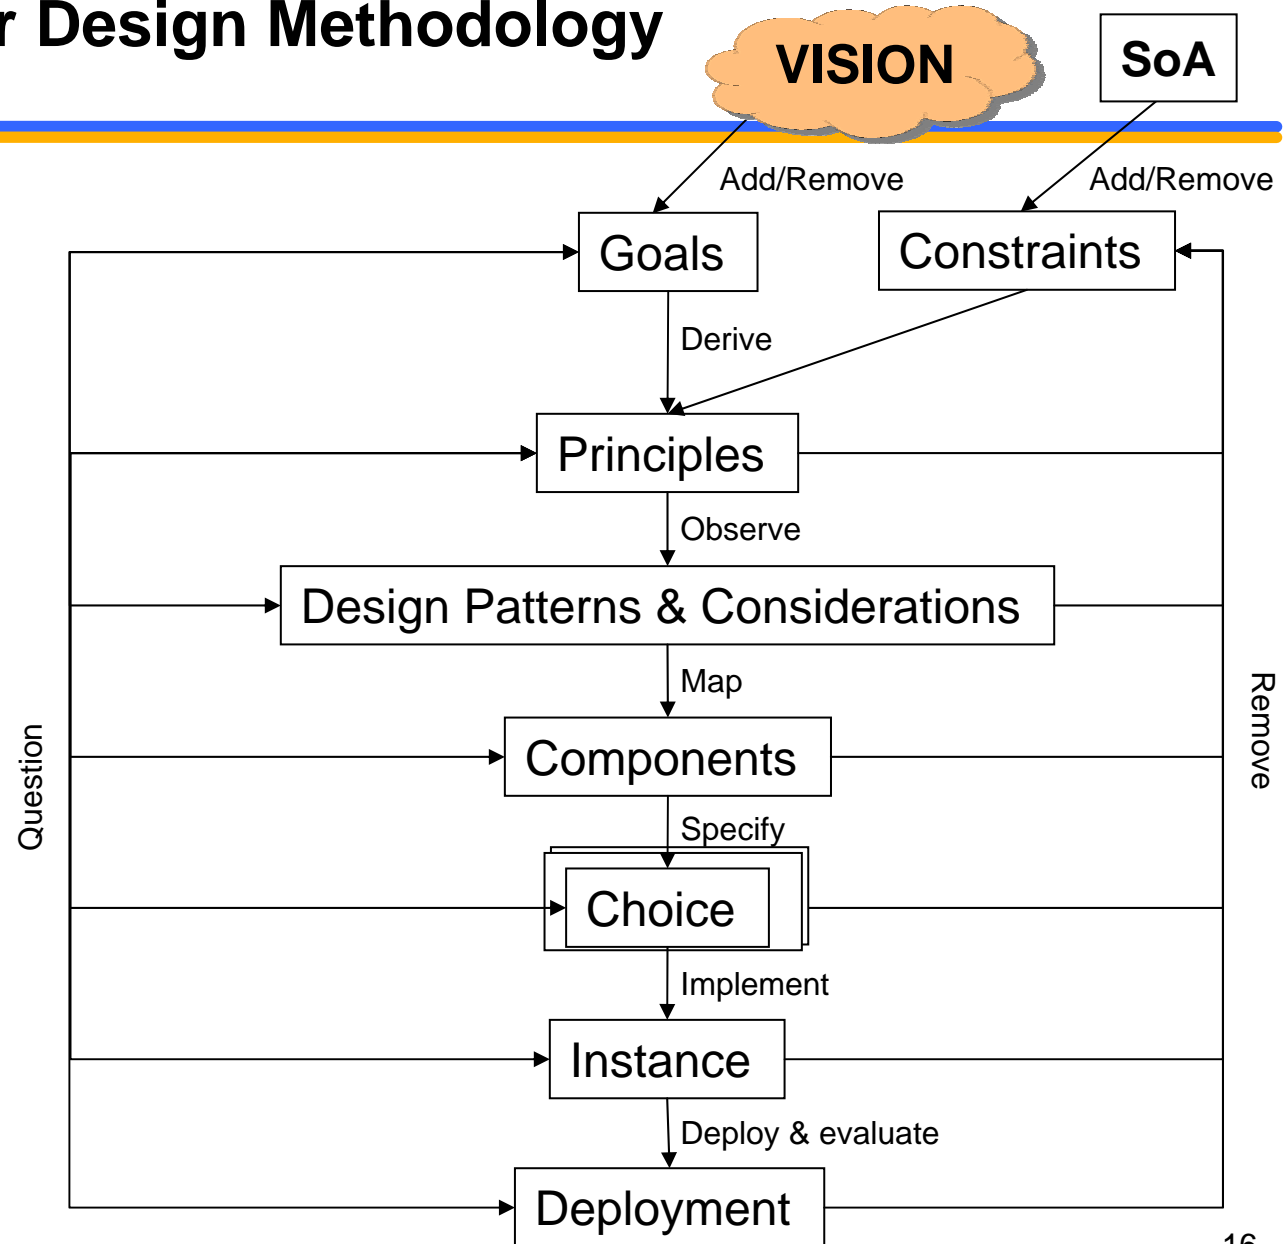
Node Architecture

Network Architecture



Our Design Methodology

- Combination of top-down and bottom-up (rationalization and development)
- Several rounds of consolidation
- Getting into (early) deployment already!





Project Objectives

- Specify, implement and test an internetworked pub/sub architecture
 - **clean-slate design** approach with deployment and migration realism
- Build on successful work being done
 - Utilize PSIRP project results from FP7 call2
 - Build on architectural concepts, implementation & test bed
- Perform qualitative and quantitative evaluation
 - Security and socio-economics important!
 - Migration and incentive scenarios important (e.g., overlay)!
- The results will be widely published
 - Open source code for the Future Internet
- Engage with FI community
 - Engage openly through public blogs, wikis and twitter

Project Overview

Project Coordinator

Arto Karila
Aalto University
Tel: +358 50 384 1549 Fax: +358 9 694 9768
Email: arto.karila@hiit.fi

Technical Manager

Dirk Trossen
Cambridge University
Tel: +44 7918 711695 Email: dirk.trossen@cl.cam.ac.uk

Partners:

- Aalto University (FI)
- Cambridge University (GB)
- RWTH Aachen University (DE)
- Oy L M Ericsson Ab (FI)
- Athens University of Economics & Business (GR)
- Essex University (GB)
- CTVC Ltd. (GB)
- Centre for Research and Technology Hellas (GR)

Duration: 09/2010 – 02/2013 **Contract No:** INFSo-ICT-257217
Total Cost: €4.9m **EC Contribution:** €3.7m

WP1 Management (Aalto)

WP2 Architecture Design
(UCAM)

WP3 Implementation,
Prototyping and Testing (LMF)

WP4 Validation & Tools
(RWTH)

WP5 Dissemination &
Exploitation (AUEB)

Project website: www.fp7-pursuit.eu
Twitter: @fp7pursuit