PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# Academic Dissemination and Exploitation of a Clean-slate Internetworking Architecture: The Publish-Subscribe Internet Routing Paradigm

# TR10-0003

Document Properties:

| | |
|---|---|
| Title of Contract | Publish-Subscribe Internet Routing Paradigm |
| Acronym | PSIRP |
| Contract Number | FP7-INFSO-IST 216173 |
| Start date of the project | 1.1.2008 |
| Duration | 33 months, until 30.09.2010 |
| Document Title: | Academic Dissemination and Exploitation of a Clean-slate Internetworking Architecture: The Publish-Subscribe Internet Routing Paradigm |
| Date of preparation | 22.10.2010 |
| Author(s) | Mark Ain (AALTO-HIIT) |
| Responsible of the deliverable | Mark Ain |
| Target Dissemination Level: | PU |
| Status of the Document: | Completed |
| Version | 1.00 |
| Document location | http://www.psirp.org/ |
| Project web site | http://www.psirp.org/ |

SEVENTH FRAMEWORK PROGRAMME

European Commission
Information Society and Media

**AALTO UNIVERSITY**
School of Science and Technology
Faculty of Information and Natural Sciences
Degree Programme of Computer Science and Engineering

Mark Ain

# ACADEMIC DISSEMINATION AND EXPLOITATION OF A CLEAN-SLATE INTERNETWORKING ARCHITECTURE:

## THE PUBLISH-SUBSCRIBE INTERNET ROUTING PARADIGM

Master's Thesis
Espoo, Finland June 25, 2010

Supervisors: Professor Antti Ylä-Jääski, Aalto University
Professor Øivind Kure, Norwegian University of Science and Technology

Instructor: D.Sc. Arto Karila, Helsinki Institute for Information Technology

**AALTO UNIVERSITY** MASTER´S THESIS
School of Science and Technology ABSTRACT
Faculty of Information and Natural Sciences
Degree Programme of Computer Science and Engineering

| | | | |
|---|---|---|---|
| **Author:** | Mark Ain | | |
| **Title of thesis:** | Academic Dissemination and Exploitation of a Clean-slate Internetworking Architecture: The Publish-Subscribe Internet Routing Paradigm | | |
| **Date:** | June 25, 2010 | **Pages:** | 125 (12 + 94 + 19) |
| **Professorship:** | Data Communications Software | **Code:** | T-110 |
| **Supervisors:** | Professor Antti Ylä-Jääski, Aalto University<br>Professor Øivind Kure, Norwegian University of Science and Technology | | |
| **Instructor:** | D.Sc. Arto Karila, Helsinki Institute for Information Technology | | |


**Abstract:**

This thesis is about disseminating and exploiting a clean-slate internetworking architecture through advanced academic courses. We design, execute, and report on the results of an academic course-based engagement and dissemination pilot-trial for a revolutionary clean-slate internetworking architecture: the Publish-Subscribe Internet Routing Paradigm (PSIRP).

ICT capabilities have evolved considerably since the inception of the Internet's predecessor, ARPANet, spawning a variety of new usage demands and operating conditions and incentives for which the original Internet architecture was never intended. Although the Internet is highly successful, its architecture has arguably become ossified and it is ridden with operational problems stemming from the obsolescence of its endpoint-centric send-receive underpinnings. In the face of increasing demands, size, and complexity, a revolutionary architectural alteration is warranted to adapt the Internet to current usage trends and holistically address its growing range of operational problems. One potential revolutionary solution that serves as the centerpoint of this thesis is the clean-slate architecture proposed by the EU FP7 PSIRP project.

It is imperative, however, that future Internet research is supplemented by efficient dissemination and exploitation activities so that technological enhancements are dutifully enacted. Unfortunately, the dissemination and exploitation of innovations from PSIRP and other revolutionary architecture proposals has been notoriously problematic.

Our academic course-based engagement approach performed very well in practice and achieved a commendable participant completion rate. Participants exhibited a good grasp of PSIRP material and provided positive qualitative and quantitative feedback regarding the quality of the course. We also gained positive feedback on the overall success of the course from a reputable group of external evaluating experts. Moreover, our observations provide useful information for future undertakings of this kind. We surmise that these results are a meaningful indicator that our methods offer a promising means by which to disseminate the innovations of a clean-slate internetworking architecture and educate the public with the end goal of promoting controlled change and progression of the Internet.



**Keywords:** academic, clean-slate, dissemination, exploitation, Internet, publish-subscribe, PSIRP;
**Language:** English

# Acknowledgements

First and foremost, I would like to thank my instructor and supervisor D.Sc. Arto Karila from the Helsinki Institute for Information Technology (HIIT). Arto extended me the opportunity to become involved with HIIT in May 2008 and his friendship and mentorship have been invaluable to me. Arto's experience and guidance were instrumental to my completion of this thesis.

I would also like to thank my thesis supervisors D.Sc. Antti Ylä-Jääski (AU) and D.Sc. Øivind Kure (NTNU) for their help coordinating and guiding my thesis, D.Sc. Dirk Trossen (UCAM), D.Sc. Sasu Tarkoma (UH), and D.Sc. Jussi Kangasharju (UH) for their helpful direction throughout the duration of my work, and M.Sc. Walter Wong for his brilliant work lecturing our two PSIRP courses.

The work documented in this thesis was conducted as part of the Dissemination and Exploitation work package (WP5) within the EU FP7 PSIRP project. I would like to acknowledge all of the technical and administrative members of the PSIRP consortium for their hard work and unprecedented innovation throughout the duration of the project. Working alongside you has been the learning experience of a lifetime and I really cherish my time spent in such a great working environment.

Last but certainly not least, I want to extend a special thank you to Mona Nordaune, Anna Stina Sinisalo, Eija Kujanpää, and the remainder of the administrative staff of the NordSecMob Programme. Your help has been invaluable to my colleagues and I throughout the past two years and it's largely through your hard work and dedication that our experience in the program has been so rewarding. I feel confident recommending the NordSecMob Programme knowing that future participants will be in your capable hands.

Sincerely,

*Mark Ain*

*Espoo, Finland* *June 25, 2010*

# Abbreviations and acronyms

The following acronyms appear frequently and/or are undefined within the body of this document.

| | |
|---|---|
| **4WARD** | (pseudo-acronym) |
| **ACK** | Acknowledge |
| **AId** | Application Identifier |
| **API** | Application Programming Interface |
| **ARP** | Address Resolution Protocol |
| **ARPANet** | Advanced Research Projects Agency Network |
| **AS** | Autonomous System |
| **BGP** | Border Gateway Protocol |
| **BSD** | Berkeley Software Distribution |
| **CCN** | Content-Centric Networking |
| **CDN** | Content Distribution Network |
| **CIDR** | Classless Inter-Domain Routing |
| **CN** | Correspondent Node |
| **cr** | Credit (ECTS) |
| **DDoS** | Distributed Denial-of-Service |
| **DHCP** | Dynamic Host Configuration Protocol |
| **DHT** | Distributed Hash Table |
| **DiffServ** | Differentiated Services |
| **DONA** | Data-Oriented Network Architecture |
| **DoS** | Denial-of-Service |
| **Dx.x** | Deliverable x.x (PSIRP) |
| **EC** | European Commission |
| **ECTS** | European Credit Transfer System |

| | |
|---|---|
| **EGP** | Exterior Gateway Protocol |
| **EIFFEL** | (pseudo-acronym) |
| **FA** | Foreign Agent |
| **FId** | Forwarding Identifier |
| **FP7** | Framework Program 7 |
| **GNU** | GNU's Not Unix |
| **HA** | Home Agent |
| **HI** | Host Identity |
| **HIP** | Host Identity Protocol |
| **HIT** | Host Identity Tag |
| **$i^3$** | Internet Indirection Infrastructure |
| **ICT** | Information and Communication Technology |
| **IETF** | Internet Engineering Task Force |
| **IGP** | Interior Gateway Protocol |
| **IntServ** | Integrated Services |
| **IP** | Internet Protocol (IPv4) |
| **IPSec** | Internet Protocol Security |
| **IPTV** | Internet Protocol Television |
| **IPv4** | Internet Protocol version 4 |
| **IPv6** | Internet Protocol version 6 |
| **LIPSIN** | Line Speed Publish-Subscribe Internetworking |
| **MCP** | Microsoft Certified Professional |
| **MCSE** | Microsoft Certified Systems Engineer |
| **MIP** | Mobile Internet Protocol |
| **MN** | Mobile Node |
| **NIRA** | New Internet Routing Architecture |
| **PKI** | Public Key Infrastructure |
| **PSIRP** | Publish-Subscribe Internet Routing Paradigm |

| | |
|---|---|
| **Pub-Sub** | Publish-Subscribe |
| **QoS** | Quality of Service |
| **RId** | Rendezvous Identifier |
| **ROFL** | Routing on Flat Labels |
| **SEATTLE** | Scalable Ethernet Architecture for Large Enterprises |
| **SHOK** | (pseudo-acronym) |
| **SId** | Scoping Identifier |
| **SPAM** | Stupid Pointless Annoying Message |
| **SSL** | Secure Socket Layer |
| **SYN** | Synchronize |
| **TCP** | Transmission Control Protocol |
| **TLS** | Transport Layer Security |
| **TRIAD** | Translating Relaying Internet Architecture integrating Active Directories |
| **UDP** | User Datagram Protocol |
| **URI** | Uniform Resource Identifier |
| **URL** | Uniform Resource Locator |
| **VPN** | Virtual Private Network |
| **WP** | Work Package |

# Contents

# List of tables

# List of figures

# 1 Introduction

This section provides an overview of our work and introduces the underlying factors that form the foundation for our efforts. This includes the motivation and basis for our course-based dissemination pilot-trial, our research hypothesis and execution plan, expected results, strategic and operational objectives, the scope and limitations of the work, and the structure of the remainder of this document.

## 1.1 Overview

Although the Internet is highly successful, its architecture has arguably become ossified and it is ridden with operational problems stemming from the obsolescence of its endpoint-centric send-receive underpinnings (e.g. SPAM, denial-of-service attacks, unicast and source-retrieval overhead, IP semantic overload etc.). In the face of increasing demands, size, and complexity, the Internet community has a growing view that these problems are in fact due to an increasing incompatibility between the Internet's original underlying architecture and modern usage demands. In this respect, we can no longer afford to implement evolutionary "patchwork" solutions; a revolutionary architectural update is necessary to modernize the Internet.

A notable revolutionary internetworking approach around which the work of this thesis is based is the European Union (EU) Framework Program 7 (FP7) Publish-Subscribe Internet Routing Paradigm (PSIRP) project (Section 3). PSIRP is referred to as a "clean-slate" architecture because its design and functionality start from scratch, taking nothing, not even IP, for granted.

The PSIRP project has been very successful and produced a technically-sound architectural redesign of the Internet that performs well and addresses the Internet's current problems both on paper and in practice. However, there have been considerable difficulties disseminating and exploiting the work. Many people, even experienced technical professionals, have a hard time understanding project concepts and thus fail to gain an adequate appreciation of the corresponding results. Moreover, there is a large discrepancy between the duration of the PSIRP project and the timeframe of its ambitions, making it difficult to spur ongoing dissemination, exploitation, and ubiquitous deployment.

Through project experiences, research of related literature, and networking amongst experienced academic and industrial partners, we have come to the conclusion that inspiring an understanding of the Internet's current problems and the need for revolutionary solutions is paramount in addressing the aforementioned exploitation barriers. In this, academic dissemination and application development appear to be especially useful. Future generations of engineers and researchers need to be educated about clean-slate architectures and gain practical usage experience so that they will be able to understand them and develop them further.

This thesis is about the dissemination and exploitation of clean-slate Internet architectures through academic courses. We justify the importance of dissemination and exploitation processes to future Internet research and examine the suitability of academics in this capacity by designing, executing, and reporting on the results of an academic course-based engagement and dissemination pilot-trial for the EU FP7 PSIRP project. This trial consists of two courses, focusing on information dissemination and application development, respectively, which are designed to educate advanced graduate and postgraduate students on the PSIRP project and relevant background. While PSIRP was used as the test case, we believe that the results will be applicable to any fundamentally new internetworking paradigm.

We will undertake the following exercises:

1) Discuss the Internet's inception and past developmental milestones and reveal how this history and pertinent technical and socio-economic factors have led to the Internet's general state of ossification (Sections 2.1 and 2.2).

2) Provide an overview of the current Internet's most prominent operational problems (Section 2.3).

3) Present notable attempted evolutionary solutions and revolutionary solution proposals designed to address the Internet's shortcomings, most importantly the FP7 PSIRP project (Sections 2.4, 2.5, and 3).

4) Discuss the importance of dissemination and exploitation processes within the ICT division of the European Union's 7$^{th}$ Framework Program and future Internet research, deduce how these processes are critical to overcome the ossification factors discussed in Section 2.2, and ascertain that academics are a promising venue for the dissemination of future Internet research (Section 4).

5) Design, execute, and present the results of a comprehensive academic engagement and dissemination pilot-trial within FP7 PSIRP (Sections 5, 6, 7, and 8).

In this, we adopt a chronological structure which effectively guides the reader through each of the aforementioned points, using the argumentation of a given point to introduce and justify the next. Exercises 1, 2, and 3 (Sections 2 and 3) are included as background and may be omitted by readers who are merely interested in the PSIRP academic pilot-trial.

## 1.2 Motivation

The current Internet's endpoint-centric send-receive communication paradigm stemmed from the need for a network that could support simplistic point-to-point communication and resource-sharing between a limited number of trusted endpoints. As this environment and the capabilities of its constituent systems grew, so did the demands that its users placed upon them. Despite the fact that the current Internet is successful and highly functional, its underlying mechanisms are almost wholly based on the designs of its predecessor, ARPANet, and have remained largely unchanged for over a quarter century. This fact is exemplified by the appearance of a variety of previously unforeseen problems resulting from the interplay of host demands, offered services, and the underlying network architecture (Section 2.3).

It now seems clear that we are reaching a critical point of operation similar to those observed in the past where a significant alteration to the Internet's core architecture (e.g. DNS, TCP/IP, CIDR etc.) was necessary to resolve the inadequacies that surfaced. However, we can no longer afford to implement **evolutionary** solutions that merely serve as "patchwork" to a static underlying architecture which is ultimately responsible for the Internet's functional problems. We believe that a **revolutionary** modification based on a wholly new underlying communication paradigm is warranted to holistically address the Internet's shortcomings and bring the Internet in line with modern usage demands.

The Internet's developmental history and a variety of technical and socio-economic factors have ultimately led to the ossification of its architectural foundations (Sections 2.1 and 2.2). It is consequently increasingly difficult, if not impossible, to devise and successfully deploy technically sound solutions on the Internet. In effect, any new attempts to modify the underlying Internet architecture have been nearly impossible to manifest due to a lack of:

1) **Understanding** (i.e. little or no appreciation of existing problems and their potential solutions)
2) **Perceived urgency** (i.e. no recognized impending operational catastrophe)
3) **Motivation** (i.e. no perceived benefit or gap to fill)

These deficiencies are even more apparent when dealing with revolutionary Internet architecture proposals due to the fact that they are almost exclusively non-incremental in nature and there is a large discrepancy between the duration of these projects and the timeframes of their ambitions. **The end result has been a worsening state of problematic operation and overall stagnation in the Internet itself.**

The first step to elicit change in the face of a scientific innovation is to address public **understanding.** That is, one must educate external communities and promote an awareness of the problem(s) that the innovation intends to solve, with the goal of fostering knowledge of the innovation and creating an understanding of why and how the solution in question is well-suited to address the given problem(s). In the realm of Internet development, understanding is a cornerstone without which perceived urgency and motivation are impossible, and it then becomes apparent that **dissemination and exploitation processes** are crucial to future Internet and ICT research because they provide the chief means to educate the general public and create this understanding of existing problems and potential solutions. From this, perceived urgency and motivation must follow if a solution is to be ubiquitously deployed.

Unfortunately, PSIRP and other revolutionary Internet architectures (and even evolutionary solutions) all suffer from a common detriment: **the dissemination and exploitation of advanced internetworking research results has been notoriously problematic.** Consequently, it has been nearly impossible to promote understanding and inspire a sense of perceived urgency and motivation to apply these results for the benefit of the Internet. As such, commercial deployments and eventual global acceptance have been virtually non-existent. These problems have been clearly evident within the PSIRP project as its dissemination and exploitation work package has struggled to effectively reach third parties.

We surmise that a root-cause of these failures is **ineffective and/or insufficient dissemination and exploitation practices**, and in light of this, we believe that dissemination and exploitation, chiefly to educate and create understanding, must be reinforced as cornerstone activities in the fields of future Internet and ICT research. Fostering an understanding of current and foreseeable problems and their potential solutions is the key starting point through which perceived urgency, motivation, and finally controlled progression, can be stimulated.

## 1.3 Problem statement

**Given an Internet which has historically struggled to implement evolutionary solutions, how can one effectively disseminate and exploit the technological innovations offered by a revolutionary clean-slate internetworking architecture so as to effectively inspire an atmosphere of understanding, perceived urgency, motivation, and finally, controlled progression?**

## 1.4 Research hypothesis

Based on a variety of tenets presented throughout Sections 1 through 5 and summarized in Section 5.2.1, **we believe that academic courses offer a promising means to engage external audiences and disseminate and exploit the innovations of clean-slate internetworking architectures with the long-term goal of promoting controlled progression on the Internet**. We believe that a course-based dissemination pilot-trial within the FP7 PSIRP project will yield the following key results, as evaluated by the course staff and an external panel of evaluating experts:

1) Positive participant performance
2) Positive participant feedback
3) Positive final evaluations with regard to the overall success of the course

Our academic pilot-trial within PSIRP is a base case; we believe that our results will also apply to other clean-slate internetworking architectures. Details pertaining to the dissemination course (course code T-110.6120), the external expert panel, the nature of these results and their interpretation etc. are introduced in the remainder of this section and discussed in detail in Sections 5 and 6.

## 1.5 Execution

It is our intention to institute two successive special-topic courses within the Faculty of Information and Natural Sciences of the School of Science and Technology at Aalto University in Espoo, Finland, whose aims will be to promote

1) **information dissemination** (course code T-110.6120) and
2) **application development** (course code T-110.6100),

respectively, for the EU FP7 PSIRP project. These courses will take place during the spring 2010 term (January – May) and will be targeted towards advanced graduate and postgraduate students who possess a thorough background in ICT. **Due to time and resource constraints, this thesis will chiefly focus on the information dissemination aspect of the project and only touch briefly upon the application development course in Section 7.2.**

An external **expert panel** consisting primarily of doctoral-level researchers who possess extensive experience within FP7 PSIRP and related fields will be convened to oversee the design, operation, and termination of these courses. Through the application of documented systematic forecasting and consensus techniques (e.g. the Delphi Method), it will be the responsibility of this panel to:

1) Design and validate the courses' objectives, structures, contents, operating methods, and assessment measures.
2) Oversee the instruction of the courses and monitor their progression.
3) Document and analyze the performance of the courses based on participant performance, participant feedback, and comments from overseeing staff.
4) Correlate said performance as an indicator of the suitability of this approach towards disseminating and exploiting a clean-slate internetworking architecture.

The validity of this approach has been verified by the Center of Excellence at Aalto University's Faculty of Information and Natural Sciences as nominated by the National Academy of Finland. Full details are provided in Sections 5.2 and 5.3.

## 1.6 Expected results

We believe that our academic course-based dissemination approach (discussed in detail in Sections 5.2 and 5.3) will yield the following key results:

1) Positive participant performance based on submissions evaluated by the course staff and expert panel (and, within the context of the PSIRP application development course T-110.6100, automated code analysis tools); we also highly anticipate the presentation of original results within T-110.6100
2) Positive participant feedback as evaluated by the course staff and expert panel
3) Positive final evaluations by the course staff and expert panel

In this, we hope to demonstrate the validity of our unique course-based engagement approach and show that academic courses are a promising means by which to disseminate and exploit a clean-slate internetworking architecture.

## 1.7 Objectives

The objectives of our academic engagement initiative within the FP7 PSIRP project fall into two categories: **strategic** and **operational**.

### 1.7.1 Strategic objectives

Strategic objectives convey the overall goals of the academic engagement initiative explored in this thesis with respect to long-term project interests. These objectives were formulated by ascertaining long-term project goals and adopting a "black box" approach with respect to the courses, allowing no feedback from the courses to affect the nature of our strategic objectives.

In conducting our research and preparing this document, we have sought to achieve the following three strategic objectives:

**OBJECTIVE #1: Compile background information on the inception and pertinent developmental history of the Internet and compose a brief literary review of prominent evolutionary and revolutionary solution proposals which serves as a useful reference for the reader.**

Our means for achieving this objective:

1) Provide a brief history of the Internet that highlights how the events surrounding its inception and the demands of users at the time contributed to its foundational endpoint-centric send-receive design (Section 2.1).
2) Highlight milestone modifications during the past 40 years of Internet development and characterize their evolutionary nature in response to impending operational limitations (Section 2.2).
3) Demonstrate that the core Internet architecture has essentially become ossified as a result of various technical and socio-economic conditions (Section 2.2).
4) Identify notable problems plaguing the current Internet as a result of modern usage demands, introduce notable evolutionary and revolutionary solution proposals, and through this demonstrate the plausible need for a revolutionary clean-slate redesign (Sections 2.3, 2.4, and 2.5).
5) Provide an overview of the FP7 PSIRP project (Section 3).

We intend to complete these ancillary tasks through the preliminary sections in this thesis in such a manner that the PSIRP engagement work fits as a subsequent logical step. Note that these points have also been used to a large degree in designing the operational objectives for the dissemination course T-110.6120 (Section 5.3.1).

**OBJECTIVE #2: Justify the importance of dissemination and exploitation processes to future Internet research, and design, execute, and report on the results of an academic dissemination and exploitation pilot-trial for the FP7 PSIRP project in order to present evidence supporting (or refuting) the conclusion that this approach constitutes a promising route by which to disseminate and exploit a clean-slate internetworking architecture.**

Our means for achieving this objective:

1) Present reputable evidence that substantiates the importance of dissemination and exploitation processes to the ICT division of the European Union's 7th Framework Program and future Internet research (Section 4).

2) Justify the importance of academia as a potential means of dissemination and exploitation for advanced future Internet innovations (Section 4.2.1).

3) With regard to the **PSIRP information dissemination course T-110.6120** offered at Aalto University during period III, spring 2010, demonstrate:

    a. Positive participant performance based on submissions evaluated by the course staff and expert panel

    b. Positive participant feedback as evaluated by the course staff and expert panel

    c. Positive final evaluations by the course staff and expert panel

The brunt of the application development course T-110.6100 will not be addressed in this thesis.

**OBJECTIVE #3: Provide an account of notable events which took place during the course T-110.6120, document our lessons learned, and provide recommendations for similar projects in the future.**

Our means for achieving this objective is based on the experiences of the course staff and expert panel and the results of objective #2.

### 1.7.2 Operational objectives

Operational objectives (Sections 5.3.1 and 7.2.1) convey the objectives of the academic courses which were designed with the intention of achieving our strategic objectives and serving the interests of the course staff and participants. These goals were evaluated by the course staff and expert panel so as to lay the framework for our envisioned courses.

## 1.8 Scope and limitations

We believe that dedicated academic courses are a relatively underemphasized tool within the realm of dissemination and exploitation for clean-slate internetworking architectures. To the best of our knowledge, this is the first instance where research of this kind has been conducted. We are unaware of any other case involving the adoption of a dual-course approach in an attempt to determine the suitability of academic dissemination and development activities towards exploiting a clean-slate internetworking architecture.

Unfortunately, the qualitative nature of this type of research poses limitations on the objectivity of our methods and results. Our application development course T-110.6100 does allow for some objective assessments of participant submissions, but the concreteness of information dissemination and academic course success are marginal at best and it is difficult in practice to conclusively prove if exploitation is achieved and to what extent. This makes it difficult to gauge the overall effectiveness of the dissemination course T-110.6120 and the combined impact of both courses.

The chief cause of these uncertainties is the presence of extrinsic factors that ultimately stem from humanistic involvements. Examples include but are not limited to:

- **Course characteristics:** Content selection, operating methods, assessment measures etc.
- **Human characteristics:** Instructors, participants, instructor-participant interaction etc.

etc.

The biases introduced by these factors are unquantifiable and it would be impossible (and arguably pointless) to attempt to empirically account for all of their effects. One must also remember that any similar undertakings to which our results may apply will also face similar external influences, and we can therefore argue that they are essentially inalienable and must be accepted as part of the experimentation environment. Thus, as with most any qualitative research, the correct course of action is to properly characterize and mitigate the effects of immovable factors so that they do not significantly impact the credibility of our work and its results. The following sections explain the steps we have taken to achieve this with respect to key facets of our work.

### 1.8.1 Definition of the strategic objectives

Our strategic objectives were designed so as to be definite while placing minimal restrictions on their interpretability. Bearing this in mind, our target objective is **not** to decisively prove that academic course-based dissemination and development is an optimal (or otherwise) exploitation method for a clean-slate internetworking architecture. The definitive structure of this argument is too rigid to conclusively verify in practice. Instead, we will demonstrate that our course T-110.6120 receives positive (or otherwise negative) results based on justifiable design decisions and operating methods, and argue that these results give an **indication of the suitability** of our approach for the exploitation of a clean-slate internetworking architecture, taking into account academia's long successful history in similar capacities.

This goal is entirely flexible because it places no objective restrictions on the conclusions that might be drawn from our subjectively-attained results. In effect, our experimentation simply serves to illustrate that a particular result is realistically achievable under a given set of conditions. The nature of that result, positive or negative, and its validity, will ultimately be determined and reasonably justified by our expert panel (Section 5.2.3).

### 1.8.2 Definition of the working methodology

As with any qualitative research, it is important to recognize that humanistic involvement is not only a potential source of subjectivity but also one of the only ways to mitigate its effects. One manner to achieve this mitigation is to obtain the support of a quorum of qualified, reputable experts whose consensus lends credibility to one's work. In our case, we have convened a panel of trained professionals (chiefly at the doctoral-level) with extensive experience in the FP7 PSIRP project and related fields. By applying documented methods of systematic forecasting and consensus, it is the responsibility of this panel to methodically complete the tasks discussed in Section 1.5. The impartial consensus of this panel is understood to provide credible and reasonably correct results by means of the Delphi Principle.

## 1.9 Thesis structure

The remainder of this document is structured as follows:

Section 2 provides background which covers the pertinent operational aspects of the Internet from its beginnings until the present. We first provide a brief overview of the inception of the Internet and discuss how the interplay between the systems and usage requirements of the time lay the foundation for the endpoint-centric send-receive communication paradigm that governs virtually all internetworking today. We also cover the milestone changes over the past 40 years of Internet development, characterizing their evolutionary nature in response to critical operational problems, and reveal how the Internet architecture has essentially become ossified due to a variety of technical and socio-economic conditions. Lastly, we present some of the most important problems plaguing the Internet today along with notable attempted evolutionary solutions and prominent revolutionary architecture proposals.

Section 3 provides a comprehensive overview of the FP7 PSIRP project, including its motivation, working structure, participants, administrative aspects, and chiefly, its technical components.

*The background and literary review in Sections 2 and 3 are intended to serve as a useful but non-obligatory reference for the reader. Those readers who are familiar with this information and merely interested in the PSIRP-related dissemination and exploitation portion of our work can omit these parts and begin with Section 4.*

Section 4 briefly discusses the nature and importance of dissemination and exploitation activities to EU FP7 ICT and future Internet research. We also ascertain how the involvement of academia is potentially crucial to these processes within the latter context.

Section 5 summarizes PSIRP's Dissemination and Exploitation work package (WP5) and outlines our academic exploitation plan as an extension to its current policies. This includes the basis for our work (a summary of the key background in preceding sections), our execution plan, the composition of our expert panel, and details pertaining to the operational objectives, structure, content, and operating methods of our course T-110.6120.

Section 6 presents the final performance results obtained from the course T-110.6120 and ascertains the degree of success of our academic dissemination approach. We discuss participant performance as evaluated by our expert panel and deduce the correlation of this performance with respect to our strategic and operational goals. Supplementary participant and staff feedback is also included and we conclude by discussing our main observations and lessons learned throughout the duration of the course.

Section 7 discusses related research and potential future work, including preliminary results from our PSIRP application development course T-110.6100. We offer concluding remarks in Section 8.

Course-related documents and other raw results are included as appendices and referenced as necessary.

# 2 Background

This section is intended to serve as a condensed literary review that covers the historical events and operational issues that have shaped the formation of the current Internet since its inception. We show why the Internet is built the way it is and how its worsening architectural rigidity has led to the prominent operational problems and proposed solutions observed today.

We first summarize the occurrences which prompted the beginning of computer networking as we know it today, leading to the now-ubiquitous endpoint-centric send-receive communication paradigm which was adopted by the Internet's predecessor, ARPANet. Subsequently, we discuss the notable (successful) developmental milestones of the Internet architecture over the past 40 years and reveal that these change-points are subject to three important commonalities:

- The need for definite modifications was only accepted once a critical breaking point in the Internet architecture had been reached.
- The modifications were devised and agreed-upon at the last minute.
- The modifications were arguably **evolutionary** in nature, not **revolutionary.**

*"… in many ways, the Internet only just works. The number of ways in which it only just works seems to be increasing with time, as non-critical problems build. The main question is whether it will take failures to cause these problems to be addressed, or whether they can start to be addressed before they need to be fixed in an ill co-ordinated last-minute rush."*

- Courtesy of [Han2006]

With this, the state of the modern Internet beyond its most recent successful evolutions serves to show that **its architecture has essentially become ossified** as a result of historical architectural designs, evolutionary rigidity, misaligned operating incentives, unforeseen evolving usage demands etc. The notions of understanding, perceived urgency, and motivation, discussed in Sections 1.2 and 4, are central to addressing these problems and the three commonalities listed above.

We also provide an account of the most notable problems of the Internet today and present several attempted evolutionary solutions (most of which have arguably failed) that were designed to fix these problems by "patching" the Internet's original shortcomings in response to modern usage demands.

Lastly, we provide a brief overview of some of the most prominent revolutionary Internet architecture proposals.

## 2.1 The inception of the Internet

The original basis for modern computer communication networks can be traced back to the publication of the first papers and books on packet switching theory in the early 1960's [Lei2010]. In 1965, Professor Leonard Kleinrock of the University of Central Los Angeles (UCLA), then a graduate student at the Massachusetts Institute of Technology (MIT), collaborated with Thomas Merrill to interconnect two computers in Massachusetts and California over a simple circuit-switched dial-in line, effectively creating the world's first wide area network (WAN) for data [Lei2010]. The experiment proved that time-shared computers could in fact exchange resources effectively over long distances. However, the experiment also showed that the circuit-switched network paradigm was wholly unprepared to handle the resulting workload, in spite of modern-day arguments to the contrary [Fer2003].

It was this realization along with the need for point-to-point computer communication that in 1967 led to the publication of the first plan for the world's first packet-switched network: the Advanced Research Projects Agency Network (ARPANet) [Lei2010].

Over the next decade, ARPANet grew to include multiple major universities, government agencies, and scientific institutions. ARPANet's original implementation offered reliable message delivery by means of the 1822 protocol for host-to-host communication and the Network Control Program (NCP) for combined addressing and transport. By 1982, ARPANet contained over 200 nodes and as expansion continued it was becoming clear that the architecture and its protocols were lacking [Int2010].

## 2.2 Characterizing evolutions and ossifications

*ACKNOWLEDGEMENTS: We would like to especially acknowledge the contributions of M. Handley [Han2006] to our preparation of this section.*

**Name-address translation** in the ARPANet was originally accomplished through the `hosts.txt` file which contained mappings that allowed each host to convert human-friendly hostnames to machine-friendly addresses representing logical network destinations. The fatal flaw in this approach was its lack of scalability: any change to the network topology required a new entry in the `hosts.txt` file of every host, which effectively amounted to the inefficient mass-redistribution of identical data. As a result, the Domain Name System (DNS) was conceived in 1983 and deployed in the mid-1980's (the exact date is arguable based on the implementation) to alleviate the distribution of common identifier-locator information and effectively distribute its administration. This is one of the first and most notable examples of a last-minute evolutionary fix to an internetworking problem which had reached critical proportions: ARPANet engineers were forced to implement a solution to bypass an underlying architectural scalability limitation, and yet the core issue, the dichotomy between human-friendly hostnames and machine-friendly network addresses, was ignored and is still largely unaddressed today.

Even though the dichotomy between network locators and human-friendly names had been seemingly rectified, ARPANet's continued growth soon began to shed light on other scalability limitations. Sporting **combined addressing and transport**, the NCP relied exclusively on the underlying ARPANet to provide reliable message delivery and it could not address networks or machines beyond a destination Interface Message Processor (IMP, i.e. a router). Thus, the loss of packets or a forwarding node would completely destroy any application-level activity. This was further exacerbated by the fact that open-architecture networking was widely unknown at the time and the only method available for federating networks was to interconnect machines at the circuit level and synchronously pass bits point-to-point. As the number of network nodes and links increased, so did the amount of traffic and potential for failures. ARPANet had reached a critical point where an overhaul would be necessary in order to accommodate a growing amount of diverse clients.

*"The ARPANet was very successful, but it was also clear that flexibility should be of prime importance in the design of a general-purpose successor… and as a result reliability was separated from addressing and packet transfer in the design of the Internet protocol suite, with IP being separated from TCP."*

- Courtesy of [Han2006]

The development of the TCP/IP stack was a direct result of the aforementioned scalability and flexibility considerations, and on January 1st, 1983, Flag Day, the entire ARPANet was switched to TCP/IP. ARPANet contained approximately 400 nodes at the time and it was arguably the last time that a non-iterative alteration to the core Internet was possible. Following Flag Day, ARPANet was split into two separate networks: the Military Network (MILNET) and the remaining ARPANet (research).

By 1987 the number of ARPANet hosts exceeded 10,000 and traffic and congestion began to reach critical proportions. The network was prone to occasionally reach a state known as **congestion collapse**, where the infrastructure was routing and switching traffic at full capacity without completing any useful work. The root cause was deemed to be TCP's error-control and retransmission strategy. In retrospect, traffic loads and congestion are protocol-independent and affect the global Internet beyond singular traffic flows. Nevertheless, it was too late in ARPANet's development to implement a core architectural change to reflect this, such as adding a new independent layer to the network stack. The end-to-end principle was effectively preserved by relegating functional changes to the network endpoints, and as such, the obvious easy incremental fix was to work with TCP itself, leading to the beginnings of TCP congestion control. TCP's congestion control mechanisms are arguably one of the most successful incremental changes that sustained the Internet into and beyond the 1990's. In spite of this, the change was purely last-minute "patchwork"; a flexible protocol-independent architectural scheme to handle traffic overloads was simply impossible to conceive and deploy at the time.

In the early 1990's the greater business community began to recognize the Internet's value to commerce and the general industry. The diversity of users (e.g. residential, commercial, academic, governmental etc.), their interests, and the resulting outpour of services necessitated the introduction of policy enforcement mechanisms for traffic handling. The large tier-1 telecommunications providers who controlled the Internet 's chief interconnects were also intent on serving their interests by instituting fine-grained controls on the passage of different types of traffic through their networks. This need for **policy routing** effectively led to the development of the Border Gateway Protocol (BGP) which has been used ever since to circulate routing information amongst the autonomous systems (AS) of the Internet. Routing scrutiny further increased through the development of dedicated intra and inter-AS routing protocols, interior gateway protocols (IGPs) and exterior gateway protocols (EGPs), respectively.

During these periods of intense growth throughout the late 1980's and early 1990's, people began to notice that **address availability** in the once-thought inexhaustible IPv4 space was becoming frighteningly limited, primarily due to inefficient allocation practices. The original unicast allocation scheme used network number ranges to create three major network classes, A, B, and C, which allocated 8, 16, and 24 bits to network addressing, respectively, leaving the remaining bits for hosts addresses. The problem with this approach was its coarse granularity.

*ASIDE: Class D and E network ranges were reserved for multicast and research, respectively.*

Class A subnets allowed for ~16.7 million hosts, a number far too large and expensive for most small and medium enterprises (SMEs). At the other end of the scale, class C subnets allowed for only 254 hosts, which was inadequate for all but the smallest organizations. Class B networks were quickly becoming exhausted as a result. The proposed solution was to discard these network classes altogether and institute an allocation and routing policy based on variable-length "classless" subnets. Explicit subnetting information was to be included with every routed IP address and effectively dictate the size of the network and host portions of the address. This Classless Inter-Domain Routing (CIDR) mechanism required changes to all end-hosts and routers in the Internet, a seemingly impossible leap for successful deployment and global acceptance. The key reason for CIDR's success is its **backwards compatibility** with classful routing which provided for a relatively unhindered migration path. CIDR is another example of a last-minute evolutionary fix to an impending limitation on the Internet; it is arguably the last wide-scale architectural change deployed on the Internet.

The shortage of IPv4 addressing space was recognized as early as the late 1980's and a variety of intentionally and unintentionally-mitigating technologies have since emerged, most notably revised allocation policies, the aforementioned CIDR, the Dynamic Host Configuration Protocol (DHCP), and Network Address Translation (NAT). The nearly universal permeation of NAT "middleboxes" is an example of a forced post-CIDR evolutionary coping mechanism. NAT is one of the most ubiquitously deployed and successful IP-aggregating technologies to date despite the fact that it effectively violates the Internet's fundamental end-to-end principle described in [Blu2001] and breaks end-to-end connectively. Many NAT-traversal techniques have been devised but differences amongst NAT implementations have prevented any single approach from achieving universal success and global acceptance.

Table 2.1 summarizes our discussion of the Internet's key historical problems and their evolutionary solutions and lists the associated ossification factors.

**Tab. 2.1 - Key evolutions and underlying ossification factors in the Internet**

| Limitation(s) | Solution(s) | Key underlying ossification(s) |
|---|---|---|
| Name-address translation | DNS | ▪ Network vs. human naming dichotomy |
| Scalability, routing inflexibility, combined addressing and transport | TCP/IP | ▪ Endpoint-centrism<br>▪ Rigid core protocol stack |
| Congestion | TCP congestion control | ▪ Lack of built-in protocol-independent QoS<br>▪ Rigid core protocol stack |
| Traffic control | BGP, IGPs + EGPs | ▪ Endpoint-centrism<br>▪ Send-receive communication paradigm |
| Address space exhaustion | CIDR, NAT, DHCP etc. | ▪ IPv4 |
| Security | Various | ▪ Endpoint-centrism<br>▪ Send-receive communication paradigm<br>▪ Rigid core protocol stack |

IP aggregation as provided by NAT is in high demand due to a shortage of IP addresses and NAT stands out as a particularly malignant evolution because of its ubiquity and availability. Internet users will continue to deploy NAT middleboxes in response to IP allocation policies and shortages, and the resulting breaks in end-to-end connectivity make it nearly impossible to successfully deploy new applications and network solutions. This can lead to only one thing: stagnation.

Commercial interests have also played a very insidious role in the later developmental stages previously discussed. As commerce and big business began to dominate the Internet, they also started to dominate the allocation of resources which shaped the Internet's development. Business and political motivators, primarily profit, took priority over architectural soundness. This not only fueled unchecked "good enough" patchwork (e.g. NAT) but also led to one of the most detrimental stagnation cycles in the Internet's history, Figure 2.1, which is ever rampant today.

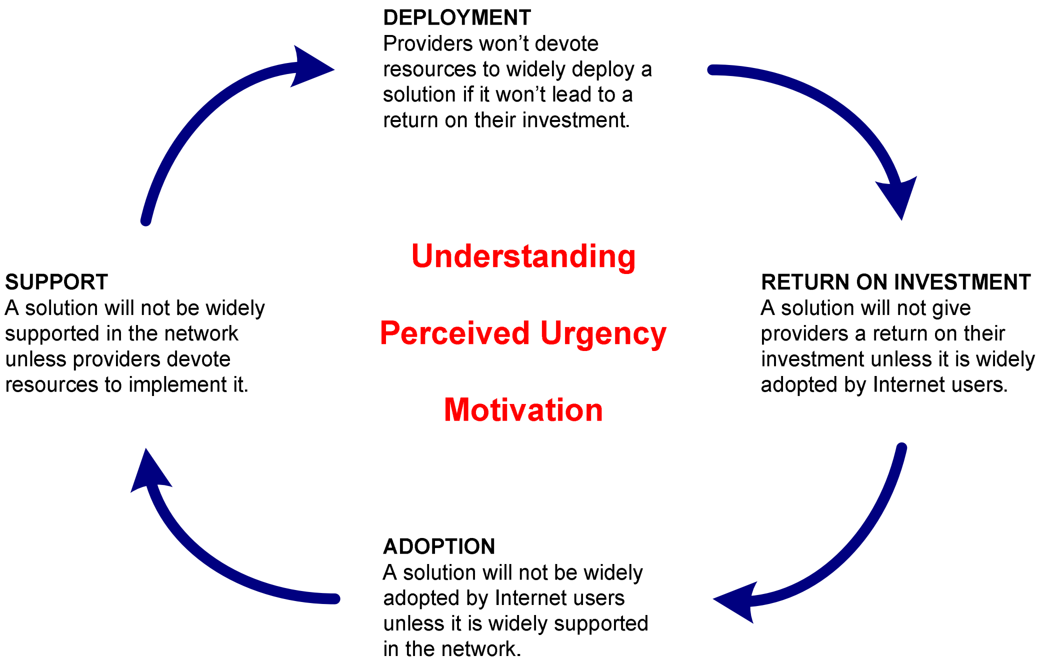


Fig. 2.1 - Internet stagnancy feedback loop

Figure 2.1 is chiefly the result of stakeholder tussle and the situation worsens as the Internet increases in size and more corporate and political involvement occurs. Understanding, perceived urgency, and motivation are key inciting factors that are necessary to break this cycle of events.

In effect, **the Internet has become ossified**, chiefly due to historical architectural designs, evolutionary rigidity, misaligned operating incentives, and unforeseen evolving usage demands. A plethora of evolutionary patchwork (e.g. DNS, TCP/IP, policy routing, CIDR, NAT, overlays etc.) has eased the Internet's growing pains but ultimately only masked the limitations of the underlying architecture in the face of modern usage trends. The end result is an ever-increasing amount of operational problems in response to increasing demands, growth, and complexity.

## 2.3 Problems in the current Internet

*ACKNOWLEDGEMENTS: We would like to especially acknowledge the contributions of M. Handley [Han2006] to our preparation of this section.*

Although the Internet has grown to provide remarkable functionality that spans extensively across both physical and societal bounds, the core architecture has in essence remained unchanged since its inception nearly 40 years ago. In spite of the successful evolutionary fixes shown in Table 2.1, a wealth of technical, social, and political problems still amass the world's largest network due to the fact that its foundation has not evolved to match the usage requirements placed upon it by its users. This fact is exemplified by the appearance of a variety of previously unforeseen technical issues resulting from the interplay of hosts, offered services, and the network architecture itself.

Here we discuss (in no particular order) some of the most notable technical problems that plague the Internet today.

### 2.3.1 Sender empowerment

The packet switching experiments carried out by Leonard Kleinrock and Thomas Merril in the 1960's and the subsequent introduction of the ARPANet were organized in response to the need for efficient resource-sharing amongst the computers of the time. These computers were physically large and possessed extremely limited resources by modern standards. Data was stored, input, and output by physical means such as punch cards or punch tape; computers served to perform some meaningful operation on data and not as a means of data storage and retrieval. Therefore, the original Internet was devised on the basis of a send-receive communication paradigm because it was simple, arguably obvious, and well-suited for the non-content-centric needs of the time. However, computers have evolved significantly since this time period, with content storage and retrieval being at the forefront of the Internet revolution.

The current dominant internetworking solution, the IP suite, works reasonably well for most existing demands but suffers from a number of limitations. One of the most notable Internet design aspects that has turned detrimental is the imbalance of powers in favor of the sender of information, who is overly trusted. The network accepts anything that the sender wants to send and will make a best effort to deliver it to the receiver, regardless of the nature of the traffic and whether or not the receiver actually wants to receive the packets. This has led to increasing problems with unsolicited traffic (e.g. SPAM, DoS etc.) and general security, forcing companies and users to conceal their e-mail addresses and place their systems behind firewalls and other types of middleboxes.

#### *2.3.1.1 SPAM*

SPAM remains one of the most common nuisances affecting many Internet services, most prominently e-mail, but also instant messaging, newsgroups, forums, search engines etc., and not without good reason: the entry barrier to mass SPAM-ing is low, the only significant overhead being that related to setting up and maintaining sending resources. This has led to a multi-billion dollar SPAM business that provides yet another venue for identity theft, malware, exploits, phishing, DoS etc. to permeate global networks.

Despite appearing to be little more than a bothersome hindrance in many situations, the full repercussions of SPAM on the Internet are intimidating: some experts estimate SPAM-induced global losses of upwards of $130 billion USD in 2009 (~€103 billion at the peak USD→EUR exchange rate in 2009, unadjusted for inflation) [Fer2009]; a 2009 report by McAfee Inc. [McA2009] estimates current SPAM volumes at ~62 trillion messages per year, with hypothetical server-side filtering saving ~135TWh of energy per year, the rough equivalent of 17 million metric tons of $CO_2$ emissions.

A variety of solutions exist to address SPAM, including but not limited to digital signatures, server and client-side filtering mechanisms, algorithmic and dynamic addresses, message propagation policies etc. The aforementioned technologies ultimately fail to wholly address the problem because they are implemented as rough afterthoughts to solve a growing complication whose foothold lies at the root of the Internet's send-receive communication paradigm.

#### 2.3.1.2 DoS

Denial-of-service (DoS) attacks are one of the most apparent and arguably destructive repercussions of the Internet's send-receive communications paradigm. The network delivers traffic from source to sink without any regard of whether or not the sink actually wants the packets in question. This level of sender empowerment allows a malicious user to completely exhaust a target's resources, typically by flooding the target with traffic until even the most rudimentary network connectivity is unavailable.

In a distributed denial-of-service (DDoS) attack, a malicious entity mounts a DoS attack using a large number of hijacked hosts, known as "bots." DDoS is particularly threatening as a single malicious user can potentially simulate a much larger number and craft an enormous amount of attack traffic; there are little to no countermeasures available since the enabling factor for the attack is the Internet's core communicational paradigm itself. Moreover, hijacked hosts may not even be aware that they are participating in the attack. Legitimate services (e.g. DNS, ICMP etc.) can also be used as a platform by which to mount DoS attacks and the potential for source address spoofing further exacerbates the situation.

As we'll see in Section 3, the PSIRP architecture was specifically designed to make all forms of DoS a thing of the past.

### 2.3.2 Infrastructure trustworthiness

Largely as a result of sender empowerment, general security over the Internet and its hosts is marginal at best, leading to questionable trustworthiness throughout the architecture:

- IP is by default best-effort and insecure. Encryption and authentication add-ons such as IPSec, VPNs, public key infrastructures (PKIs), application-level solutions etc. mitigate some security problems but are ultimately only patchwork fixes. Moreover, these solutions attempt to secure the endpoint-based channels though which data is exchanged and not the data itself.
- As seen in DDoS attacks, sender empowerment further enables malicious users to cause harm by hijacking the machines of well-intentioned users. In this respect, user intention is not well respected on the current Internet.
- Increasingly malicious and virulent malware, viruses, Trojans etc. are being developed on a daily basis to undermine even the latest security countermeasures. Sender empowerment enables such malware to further exert its negative effects through unknowingly-compromised systems.

etc.

### 2.3.3 Application deployment

The end-to-end principle should in theory make it easy for developers to deploy applications across a multitude of end hosts without having to worry about interactional problems across the network. Unfortunately, evolutionary developments and patchwork solutions like NAT have broken end-to-end functionality on many levels. This leads to a development stagnation cycle similar to that illustrated in Figure 2.1.

*"There is a vicious circle – application developers will not use a new protocol (even if it is technically superior) if it will not work end-to-end; OS vendors will not implement a new protocol if application developers do not express a need for it; NAT and firewall vendors will not add support if the protocol is not in common operating systems; the new protocol will not work end-to-end because of lack of support in NATs and firewalls."*

- Courtesy of [Han2006]

In effect, the best applications to address a given usage demand may never see widespread adoption. Problems and usage demands are addressed as they manifest themselves, and at the last minute, using the solution which is most readily available.

### 2.3.4 Congestion control

Traffic and congestion are protocol-independent and affect the global Internet beyond singular traffic flows. Nevertheless, congestion control was implemented at the transport layer via TCP because it was too late in the Internet's development to change its core protocol stack.

Although TCP congestion control is a largely successful incremental evolution that has sustained the growth of the Internet in form and function, a variety of insufficiencies have surfaced as the Internet has continued to expand over the past two decades:

- TCP only **reacts** to congestion; it does not necessarily proactively prevent it. TCP's convergence times have proven insufficient as bandwidth and flow capacities continue to increase and link characteristics are increasingly dynamic.
- Application and per-flow requirements have changed considerably in the past 20 years and exposed a variety of security, performance, and compatibility limitations in TCP's congestion control algorithms.
- TCP generally performs poorly over links with high "bandwidth*delay"-products because it is too slow to converge to the maximum transmission rate and backs-off too aggressively when it detects congestion. This results in significant added overheads, even under the most optimistic uncorrected bit error rate conditions in modern equipment.
- TCP congestion control was expressly designed for wired environments where bit error rates are low and congestion is the chief cause of perceived packet loss. TCP congestion control is highly unsuited for wireless operation because it reacts to bit-error-induced packet loss as though congestion had occurred, inappropriately reducing transmission rates and hampering overall network performance. This has proven to be a significant detriment as wireless networking has proliferated since the 1990's.

etc.

Developers have attempted to address these problems by releasing alternative TCP versions with updated congestion control algorithms. These offerings have seen limited success, chiefly due to TCP's tight integration in operating system kernels and the difficulties associated with incremental deployments that require core network compatibility.

### 2.3.5 Inter-domain routing

The original policy-routing mechanisms of BGP were a reaction to the abundance of users and potential for commercial competition over the Internet, and BGP operation is centered about the fact that Internet ASs are separate and equal entities and route-path information is commercially-sensitive. BGP attempts to avoid unnecessarily releasing this information and is therefore often subject to a certain degree of misconfiguration, security vulnerabilities, slow convergence, etc.

### 2.3.6 Mobility

*"Mobility raises five fundamental problems:*

1) ***Locating the mobile host or service:*** *Before any communication can be initiated, the desired end-point must be located and mapped to an addressable destination.*
2) ***Preserving communication:*** *Once a session has been established between end points (typically applications), communication should be robust across changes in the network location of the end points.*
3) ***Disconnection gracefully:*** *Communicating applications should be able to rapidly discern when a disconnection at either end, or a network partition, causes communication to be disrupted.*
4) ***Hibernating efficiently:*** *If a communicating host is unavailable for a significant period of time, the system should suspend communications, and appropriately reallocate resources.*
5) ***Reconnecting quickly:*** *Communicating peers should detect the resumption of network connectivity in a timely manner. The system should support the resumption of all previously established communication sessions without much extra effort on the part of the applications.*

*…We argue that a complete – and useful – solution must address all of these issues."*

- Courtesy of [Sno2001]

The issue of mobility is almost wholly unaddressed in the **original** Internet architecture and is largely fueled by IP semantic overload, which is essentially the root cause of problems #1 and #2 above. The need for constant maintenance of an addressable locator is inherent of the separation between human-friendly identifiers and network locators. The two most prominent attempts to address this problem are MIP and HIP, which have both arguably failed due to a lack of deployment incentives. [Sno2001] also argues that universal naming should be avoided as mobile applications will tend to resort to the naming schemes which best suit their unique operating needs.

*"The last three – disconnection, hibernation, and reconnection – have received little attention outside of the file system context..."*

- Courtesy of [Sno2001]

There are little to mechanisms embedded in the network to freeze, disconnect, reconnect, and resume sessions amongst logically distinct hosts. The difficulty of these tasks is further exacerbated by the fact that a variety of dangerous security deficiencies are common in most transport protocols during session freezing and re-establishment [Aur2004]. [Kha2008] is an example of a session freezing and re-establishment method using the existing TCP/IP protocol stack, although this is still a patchwork solution which is likely subject to functional and security-related obstructions.

Considering Section 2.2, it is almost certain that we will never see the proliferation and widespread adoption of a solution that addresses the five aforementioned aspects of mobility whilst operating over an unchanged Internet protocol stack. It seems more logical at this stage to address mobility as a key functional requirement in a revolutionary architecture that simultaneously resolves all of the Internet's most pressing problems.

### 2.3.7 Multi-homing

The importance of multi-homing has become more apparent as the Internet has grown in size and function; reliability, transparent-failover, and load-sharing often necessitate multi-homed connections. However, the mere presence of multiple IP prefix announcements on a wide scale removes the benefits of hierarchical IP aggregation and fail-over mechanisms are ill-suited to preserve higher-level network functions in the face of multiple underlying links.

### 2.3.8 Address space

It has been nearly 30 years since the inception and widespread deployment of the TCP/IP protocol suite over the Internet. The facilitators of Flag Day likely never imagined that a 32 bit addressing space would be in danger of exhaustion so soon after wide spread deployment. Increased numbers of wireless devices, always-on connections, higher Internet adoption rates, and inefficient address allocation policies are chiefly to blame, despite the implementation of various coping mechanisms (e.g. NAT, DHCP, improved allocation policies, address reclamation projects etc.).

The development and planned deployment of IPv6, Figure 2.2, was thought to be a long-term solution to this problem [Han2006], but unfortunately the actual deployment trend, Figure 2.3, has been wholly insufficient.

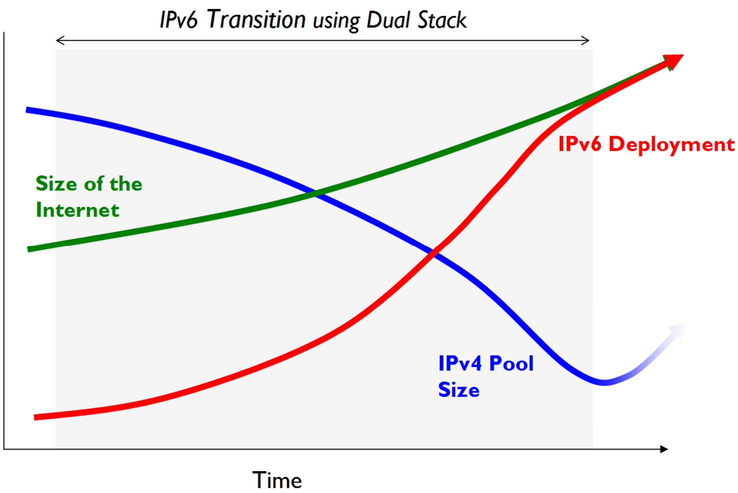


Fig. 2.2 - Planned IPv6 transition, non-linear scale (courtesy of [Hus2008])

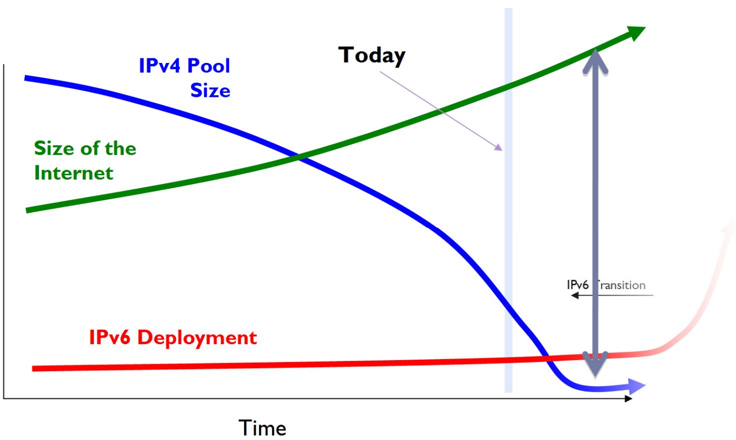


Fig. 2.3 - Actual IPv6 transition, non-linear scale (courtesy of [Hus2008])

### 2.3.9 Identifier-locator unification

The original framework for the Internet was designed to support simple insecure point-to-point communication between two stationary end hosts in a relatively small trusted network environment. Very little attention was paid to the notion of end-host mobility or the consequences of separating human-friendly names from network addresses. The situation is quite different today, with mobile nodes numbering in the millions and diversifying at increasing rates. And yet, a node's topological locator (i.e. its IP address) also serves as its unique name to identify that node within the network. Therefore, when a node changes its location, its name also changes due to the hierarchical nature of IP, and thus end-to-end reachability is broken.

The workings of DNS should in theory fix this problem as the local DNS server in a given network acts as a rendezvous point by which a statically named node's globally routable dynamic IP can be discovered. Unfortunately, the DNS server tree is ill-suited as a mobility mechanism because it caches hostname-to-IP bindings to improve lookup efficiency and reduce network overhead. It can take days for a record update to propagate globally.

Other solutions that address the identifier-locator split include Mobile IP (MIP), the Host Identity Protocol (HIP), and distributed hash tree (DHT)-based systems. However, none has been globally implemented to date (with the possible exception of DHT-based file sharing mechanisms) and communications involving mobility have been severely limited as a result.

### 2.3.10 QoS

Maintaining proper quality of service (QoS) over the Internet has been an increasingly important aspect of service delivery as the Internet has grown to encompass more users and advanced services. Certain applications necessitate varying degrees of control over traffic flow characteristics such as bandwidth, latency, jitter, packet loss etc. so that their services can be delivered in an optimal manner, Figure 2.4.

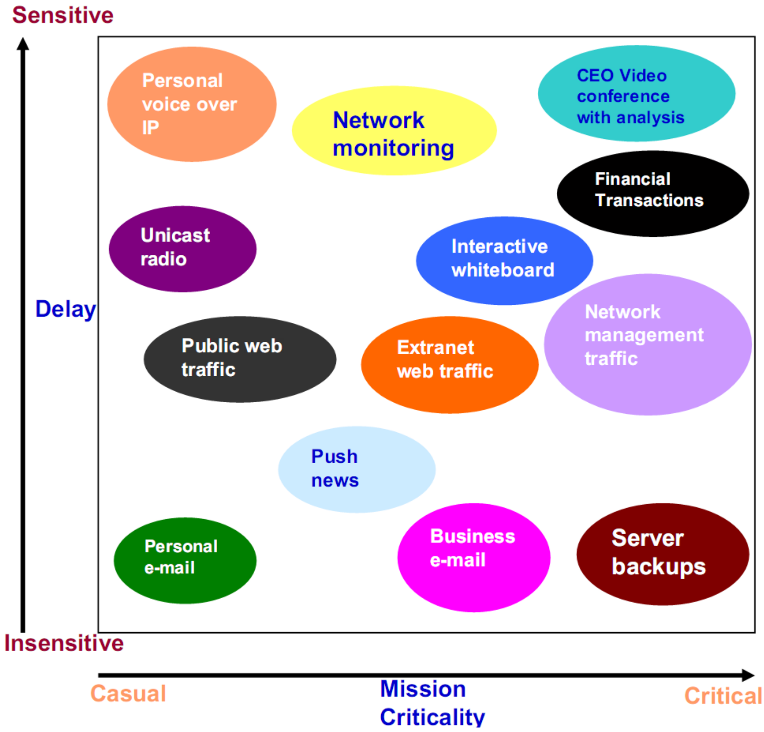


Fig. 2.4 - Network applications: latency vs. loss-sensitivity (courtesy of [Sie2009])

The Internet currently has no globally-implemented reliable end-to-end QoS infrastructure. The two QoS schemes which were standardized to address the aforementioned requirements, DiffServ and IntServ (Section 2.4.7), each suffer from unique technological limitations and architectural struggles. DiffServ allows operators to control the treatment of packets but does not guarantee any particular level of service or policy adherence across network boundaries. IntServ reserves network resources and attempts to guarantee the conditions of a network flow end-to-end, although the process is complex, resource intensive, and requires supportive cooperating routers across autonomous networks.

QoS policies are often in conflict amongst multiple operators due to governance and charging policies. In conjunction with the incentives offered by peering, these problems only serve to further complicate the state of QoS on the Internet and effectively negate the possibility of a network-wide solution which satisfies the needs of all of the parties involved. Moreover, most application-level QoS solutions that attempt to circumvent network involvement are typically proprietary and operate over an underlying network that is still best-effort at its core.

Most importantly, abstracting the concatenation of multiple distinct links between intermediary devices into a single end-to-end link severely limits consistent QoS and robustness end-to-end, thus rendering the aforementioned link characteristics of limited value in constructing predictable QoS solutions.

### 2.3.11 Unicast and source-retrieval overhead

Source-retrieved (i.e. non-cached) unicast streams are the simplest instance of endpoint-centric send-receive communications and have accounted for a majority of the Internet's traffic. This wasn't a problem in the Internet's early years, when the network contained few hosts, mass content delivery was virtually unknown, and the need to share identical data (e.g. request redundancy) was also very limited. Circumstances have changed considerably today, with topologically-distant hosts numbering in the billions and requests for identical data being increasingly common (e.g. streaming audio/video, breaking news, software updates etc.). In this growingly dynamic environment, it is highly inefficient to serve redundant requests with their own unique traffic flows (i.e. unicast) that are each rooted at the data source (i.e. source-retrieval) regardless of the location of the sink. In tandem, concurrent source-retrieved unicast streams severely deplete network resources and limit the scalability of the Internet in both form and function. As the amount of hosts and content increases, network and service providers will increasingly rely on multicasting (Section 2.4.2) and caching to keep network overheads in check. Although there are overlay solutions (such as DHTs and the Akamai CDN [Aka2007]) that provide variants of these functionalities within the network, it is becoming more apparent that further benefits could be realized by making multicast and caching native operating modes within the underlying architecture. This is another aspiration of the PSIRP project.

## 2.4 Attempted evolutionary solutions

In light of the Internet's growing range of problems, the following prominent evolutionary solutions were designed to address some of the functional holes, described in Section 2.3, which appeared when the Internet's usage demands outgrew its intended architectural capabilities. These solutions have largely **failed** due to lackluster global deployment and worldwide acceptance. These failures can be attributed to one or more of the following:

- A lack of understanding and/or perceived urgency and/or motivation
- Internet stakeholder tussle (i.e. why spend resources to implement support if others won't do the same for you?)
- Technological shortcomings

### 2.4.1 ECN

The Explicit Congestion Notification (ECN) [Ram2001] extension to IP was officially finalized by the IETF in 2001 and enables end-to-end congestion notifications between collaborating endpoints. In the event of high traffic loads, ECN-enabled routers can set the ECN bit of received packets to indicate the beginnings of network congestion to the destination host. Consequently, the destination informs the source node that it has received an ECN and the source must henceforth behave as though packets had been dropped in the network. This typically entails a reduction in the sender's congestion window as per TCP congestion control.

The brilliance of ECN lies in its attempt to make TCP address congestion **before** its effects (i.e. traffic loss) are manifested, as opposed to traditional TCP congestion control which merely reacts **after** congestion causes lost or delayed packets. A recently proposed variant of ECN, ECN+ [Kuz2009], enables congestion control of TCP SYN/ACK packets and supposedly dramatically improves the performance of short-lived connections.

Although ECN has been shown to improve virtually all facets of network queuing performance [Kuz2005], it has seen relatively low deployment and acceptance amongst interoperating parties, particularly because it requires both end-host and network compatibility. Another problem is the fact that outdated or poorly designed network equipment will sometimes simply drop packets that have the ECN bit enabled, instead of disregarding the ECN bit and forwarding the packets anyway. Furthermore, routers must have some form of active queue management technology enabled in order to properly detect congestion and set ECN bits [Kuz2005], and as per Section 2.3.10, QoS over the global Internet is still problematic. Source nodes may also choose to behave maliciously and ignore ECN-notifications sent by the sink in an attempt to obtain enhanced throughput.

### 2.4.2 Multicast

Multicasting is a powerful tool to reduce network traffic overhead and resource usage in cases where multiple recipients are being served identical unreliable real-time streams. Multicasting assigns a class D group address to a set of recipients and sends a singular traffic stream through the network until an optimal local splitting point (i.e. a rendezvous point) is reached, after which the stream is demultiplexed and the clients are each served with identical unicast streams. The splitting point is typically a high-capacity local aggregator (e.g. a multicast router). By definition, multicast traffic is connectionless and is thus transported over UDP in virtually all circumstances.

The Internet Group Management Protocol (IGMP) [Fen1997] is the most commonly implemented multicasting-group-management protocol, operating between end hosts and the local rendezvous point for the purpose of establishing and managing multicast group membership.

Protocol Independent Multicast (PIM) is the most widely implemented family of protocols for constructing multicast distribution trees. Unlike IGMP, PIM operates over the core network between multicast-enabled routers and builds optimized forwarding trees. Four different variations are available:

- **PIM-SM:** PIM-Sparse-Mode [Fen2006] creates a unidirectional distribution tree rooted at a local rendezvous-point. Participation is based on membership requests (i.e. end hosts must elect to join the tree).
- **PIM-DM:** PIM-Dense-Mode [Ada2005] creates a unidirectional distribution tree rooted at a local rendezvous point. Participation is based on multicast flooding and pruning (i.e. end hosts must elect to withdraw from the tree).
- **BIDIR-PIM:** Bidirectional-PIM [Han2007] builds bidirectional distribution trees.
- **PIM-SSM:** PIM-Source-Specific-Multicast [Bha2003] builds distribution trees rooted at a single source.

Multicast is most commonly used in LANs to distribute real-time audio and video (e.g. IPTV) or for concurrent file distribution. Large-scale use has been very limited due to setup overhead and performance, deployment, and security concerns. The demand amongst end-consumers is also relatively limited because the problems of inefficient unicast resource usage are not pressing to solve from the perspective of end-users who have little to no investment in the network.

### 2.4.3 IPv6

The Internet Protocol version 6 (IPv6) [Dee1998] was first proposed in 1991, entered into the Internet Standards Track in 1995, and officially standardized by the IETF in 1998. IPv6 offers several improvements over IPv4:

- **Larger address space:** IPv6 addresses are 128 bits in length ($\sim 3.4 * 10^{38}$ available addresses), whereas IPv4 addresses are only 32 bits ($\sim 4.3 * 10^{9}$ available addresses).
- **Stateless address auto-configuration:** IPv6 hosts can automatically assign themselves an address without a stateful configuration method such as DHCP.

- **Standardized multicast:** Unlike IPv4, IPv6 includes provisions for multicast as part of the base protocol specification.
- **Mandatory security:** The use of IPSec is mandatory in IPv6.
- **Reduced processing overhead:**
  - Simpler packet headers; most fields have been relegated to the extension header as options.
  - Routers are no longer required to perform fragmentation.
  - No checksum (assumed provided by other layers).
  - Path routers no longer need to compute the time a packet has been queued and modify its time-to-live (TTL) field accordingly.
- **Improved mobility:** Support for MIPv6 (avoids MIPv4 detriments) and Network Mobility (NEMO).
- **Extensibility:** The extension header of IPv6 allows for future expansion and variable length options.
- **Jumbo packets:** IPv6 supports packets as large as $2^{32} - 1$ bytes.

etc.

Despite boasting an advanced feature set which was standardized over a decade ago, IPv6 packet loads represent less than one hundredth of one percent of all Internet traffic and growth remains slow in western countries, Figures 2.2 and 2.3 [IP2010]. In short:

- **There is a lack of understanding:** The average Internet user is not aware of the ramifications of the IPv4 address shortage.
- **There is a lack of perceived urgency:** Even the educated Internet user has yet to significantly notice the ramifications of IPv4 address exhaustion; if the effects are present (e.g. increases in Internet subscription prices), they are often difficult to properly attribute to a shortage of IPv4 address space.
- **There is a lack of motivation:** The world's network and service providers are reluctant to invest in IPv6 infrastructures because they don't foresee any meaningful immediate return on their investment.

### 2.4.4 MIP

Mobile IP (MIP) was standardized in 2002 (MIPv4) [Per2002] and 2004 (MIPv6) [Joh2004] with the intention of overcoming the problems induced by the identifier-locator split in IP host mobility. MIP allows any third parties, known as correspondent nodes (CNs), to maintain communications with a mobile node (MN) via a static destination IP address, even if the MN periodically migrates from one network to another. The MIP specification calls for the use of "agents" within logical networks to keep track of MNs and their movements:

- **Home Agent (HA):** Keeps track of mobile nodes whose home address is within the HA's network.
- **Foreign Agent (FA):** Keeps track of all roaming nodes within its network and assigns them care-of-addresses (CoAs).

Note that the terms "home" and "foreign" denote the point of view of a single MN. That is, a given MN has only a single HA located in its home network and only foreign networks contain FAs. However, the HA of a given MN may serve as the FA for a visiting MN. The HA and FA are logical devices and may or may not be implemented in the same physical device within a given network.

The HA is always updated as to the location of the mobile nodes for which it is responsible. When a CN wishes to communicate with an MN, it addresses traffic to the MN's home address. These packets are intercepted by the HA and routed to the MN's local FA which forwards the packets to the MN. The MN can send return traffic through the HA to the CN using its permanent home address as the source address (reverse tunneling), or it can send the packets directly to the CN (triangle routing, more efficient). The latter is possible only if the CN's network is equipped with routers that do not perform ingress filtering.

MIP suffers from several problems which have contributed to its meager adoption. Technological shortcomings, security vulnerabilities, and suboptimal routing are still major concerns. Moreover, HA-FA and FA-FA handoffs associated with host mobility can be subject to high latency and tunneling packets through the HA and FA is resource intensive. The most significant roadblock to deployment occurs due to stakeholder tussle:

- Why implement MIP support in MNs if no autonomous networks are willing to spend resources to implement a FA?
- Why spend resources to implement a FA if visiting MNs don't support MIP?

### 2.4.5 HIP

The principles behind the Host Identity Protocol (HIP) were first discussed by the IETF in 1999. A working group and a research group were established in 2004 and by 2009 the base HIP specification [Mos2008] was essentially complete, although more ongoing development is still needed.

As its name implies, HIP adds a “waist” layer to the TCP/IP protocol stack between the network and transport layers. This layer is responsible for establishing a binding between hosts and public keys for the purpose of alleviating IP semantic overload and facilitating mobility, multi-homing, end-to-end security etc. HIP uses hosts’ public keys to implement a new name space of host identifiers (HIs) which are used as a binding point for transport layer connections. HIs are presented as 128 bit host identity tags (HITs) output by a hash function (a size of 128 bits was chosen to enable future compatibility with IPv6 address fields). End hosts address connection endpoints using HIs instead of IP addresses; HIs are translated to IP addresses by the host kernel, effectively providing an identifier-locator split and integrating cryptographic security into protocol negotiations.

DNS was originally the intended means of distributing and locating HIs amongst potential communicators, although this approach inherits all of the problems which plague the aging domain name system (e.g. lackluster security, excessive record propagation times etc.). Ericsson Research introduced a more elegant solution termed $HI^3$ which essentially combines HIP and the Internet Indirection Infrastructure’s ($I^3$, section 2.5.1.1) DHT-based approach, using the latter as a control plane to bootstrap connectivity via HIs.

Despite boasting a highly innovative feature set and a number of updated IETF Internet Drafts, HIP has only a few known implementations in the industry at the time of this writing and ongoing adoption shows little signs of increasing. Global deployment is difficult (e.g. PKI overhead, architectural ossification, complexities inherent of mass adoption etc.) and the issues of understanding, perceived urgency, and motivation amongst daily Internet users are still largely unaddressed.

### 2.4.6 IPSec

The Internet Protocol Security (IPSec) [Atk1995] suite was originally defined in 1995 for the purpose of providing authentication, confidentiality, and key management at the network layer. IPSec can be implemented in one of two modes:

- **Transport Mode:** The default mode most commonly implemented between end stations; IPSec headers are directly added to secure IP packets.
- **Tunnel Mode:** Encapsulates entire IP packets within IP packets that are protected by IPSec headers. Essentially functions as a VPN. Most commonly used between gateways.

The Internet Key Exchange (IKE) protocol [Har1998] is responsible for setting up a shared security association (SA) between IPSec endpoints. Cryptographic keys are derived through a shared session secret obtained via Diffie-Hellman (DH) key exchange. IPSec adds one or less commonly both of the following additional headers to protect IP packets:

- **Authentication Header (AH):** Provides integrity, data origin authentication, and potentially replay protection and non-repudiation.
- **Encapsulating Security Payload (ESP):** Provides integrity, data origin authentication, replay protection, and confidentiality.

The extent of the protection gained through the AH and/or ESP depends on a variety of factors, most importantly the version of IP in use and whether transport or tunnel mode is implemented. The selected hashing functions and encryption ciphers are also important in determining the level of protection provided.

While somewhat popular amongst enterprises seeking to leverage secure Internet communications, public use of IPSec is low due to its complexity and setup overhead, and other more practical solutions (e.g. SSL, TLS etc.) tend to dominate everyday Internet use.

### 2.4.7 DiffServ and IntServ

Differentiated Services (DiffServ) [Nic1998] and Integrated Services (IntServ) [Bra1994] are two disparate attempts at creating a standardized QoS system over the Internet infrastructure. As network usage and service complexity continually increase, QoS has become an integral part of modern networking as the base Internet's best-effort delivery guarantees are often insufficient for differing types of service traffic and operator-consumer business models. QoS methods were devised to provide more predictable propagation characteristics (e.g. bandwidth, latency, jitter, loss etc.) with the goal of enabling improved services, pricing models, and performance.

DiffServ addresses network QoS by relegating complex processing from the Internet core to its endpoints. Packets are classified by edge devices according to inbound and/or outbound policies which tag packets with a numerical identifier that is observed by core routers when determining input and/or output queuing priorities. This allows different types of traffic to be efficiently prioritized within the core without requiring complex processing amongst core devices. IP packets are tagged via the Differentiated Services CodePoint (DSCP) field, originally known as the Type-of-Service (ToS) field (obsolete), and provisions exist that allow cross-marking and translation amongst many equivalent marking fields in different protocols, including but not limited to:

- Ethernet Class-of-Service (CoS)
- Multi-Protocol Label Switching (MPLS) Experimental (EXP)

etc.

DiffServ is lightweight and does not implement signaling or resource reservation. Packet markings may change throughout network travel and thus end-to-end QoS is not provided across autonomous networks.

IntServ adopts a more complex strategy involving end-to-end resource reservation and guaranteed traffic propagation characteristics from source to sink. The Resource Reservation Protocol (RSVP) [Bra1997] is responsible for dictating the source's requested flow parameters to the network and signaling the network's acceptance or rejection of this request to the source. If the request is accepted, the requested network resources (e.g. links, bandwidth, queue priority, time etc.) are reserved from end-to-end and provide what is essentially a dedicated path across the network with predictable propagation characteristics. The network ensures that the source does not violate the limitations imposed on its traffic flow. Once communication is complete, the resources are freed and the path is torn down.

DiffServ is relatively simple for providers to implement although adoption is hindered by the fact that different autonomous systems and service providers are able to modify packet tags within their own networks. On the other hand IntServ has seen very limited deployment because it is resource-intensive and lacks scalability (e.g. per-flow reservation and tear-down, core router processing overhead and state demands etc.).

### 2.4.8 Distributed hash tables

A distributed hash table (DHT) [Res2006] provides exactly what its name implies: a data storage and retrieval system based on an indexed hash table with decentralized properties. DHTs are made-up of two key components:

- **Nodes:** Individual member systems comprising the DHT.
- **Keyspace:** An abstract address space (e.g. the set of all 128 bit strings) that delineates the DHT's storage structure.

DHT storage and retrieval functions as follows:

1) A keyspace partitioning scheme is applied to assign addresses from within the keyspace to the participating nodes. A node's address uniquely determines its responsibility for a section of the keyspace and a self-healing algorithm ensures that the entire keyspace is "owned" despite membership changes in the DHT.

2) An overlay network is created in which nodes form routes amongst each other according to a pre-determined algorithm; typically, source nodes form links to other destination nodes whose addresses within the keyspace conform to certain properties with respect to the source node's address. An overlay is by definition built on top of an underlying network such as IP, which implies that network reachability amongst the nodes is required before the overlay can be formed.

3) A node wishing to store a given piece of data *d* within the DHT begins by computing the data's index *h(d, m)*, where *h()* is a globally-known hashing function and *m* is meta-data or null. The node then issues a *put* operation to store *d* within the node that is responsible for the keyspace containing the index *h(d, m)*. *d* is greedily routed by every node towards the neighbor whose address is closest to but not beyond that of the ultimate destination node. This process continues until the destination is reached. Due to the full-coverage of the keyspace partitioning mechanism and the self-healing properties of the DHT, the destination will always eventually be found.

4) A node wishing to retrieve a given piece of data *d* from the DHT performs the same index-resolution action as described in step #3 and issues a *get* request towards the node that is responsible for the index of *d*. The request is greedily routed to the destination node which in turn responds to the request by sending the data *d* towards the requesting node. In some cases the data may not be found even if the appropriate source node is discovered.

DHTs are an innovative and unique concept which will surely be remembered in history as one of the most important ICT discoveries of the late 1990's. Their most important characteristics are decentralized storage and routing, scalability, robustness, and provisions for anonymity.

Unfortunately, DHTs suffer from several drawbacks that reduce their suitability as mainline replacements to traditional hierarchical IP networking. Most importantly, DHTs are overlays and are completely dependent on the topology management, routing, and forwarding functions of subordinate network layers. This overlay functionality also imposes a completely arbitrary binding between the DHT keyspace and underlying network addressing scheme, meaning that DHT operations between two nodes located within the same subnet may be subject to multiple DHT hops. This in turn leads to situations where traffic may unnecessarily propagate across many geographically-distant networks when the source and sink may in fact be within the same local area network (LAN). Attempts to introduce hierarchical address aggregation or fine-tuned control of data storage often compromises the beneficial characteristics of the DHT.

Modern examples of DHT implementations include Kamdelia [May2002], Chord [Sto2001], Tapestry [Zha2004], Pastry [Row2001], and Canon [Gan2004]. The popular BitTorrent [Bit2010] CDN implements DHT functionality based on Kamdelia [May2002].

## 2.5 Proposed revolutionary solutions

We are now faced with the increasing realization that **we are using an Internet whose core architecture was never designed to cope with the demands that users are placing on it today.** We can no longer afford to implement **evolutionary** solutions that only serve as patchwork to a static underlying architecture which is ultimately responsible for the Internet's functional problems. We believe that a **revolutionary** modification based on a wholly new underlying communication paradigm is warranted.

The idea of a revolutionized Internet is certainly not new. The need was slowly recognized as early as the mid 1990's as the World-Wide-Web (WWW) prevailed and modern usages such as content storage and retrieval, mobility, multi-homing etc. began to bring the Internet's endpoint-centric send-receive inadequacies to light. This led to considerable research into innovative architecture designs over the past decade. These proposals suggest radical changes to the core functions that support the current Internet, including routing, forwarding, addressing, naming etc. Moreover, a variety of new functional solutions were proposed based on state-of-the-art developments in areas such as mobility, context-aware computing, distributed data storage and processing, ubiquitous computing etc.

The following are brief overviews of some of the most prominent revolutionary Internet architecture proposals of the past decade.

### 2.5.1 DHT-based approaches

The following architecture proposals are primarily based on DHTs as described in Section 2.4.8.

#### *2.5.1.1 i³*

As its name implies, the Internet Indirection Infrastructure (i³) [Sto2004] adds send-receive indirection within the Internet so that sending operations are effectively decoupled from reception. A rendezvous-based data exchange communication paradigm is implemented whereby sink nodes associate themselves with an identifier and insert and remove "triggers" to these identifiers at will within the network to indicate their willingness to receive data or the lack thereof, respectively. Source nodes wishing to communicate with a particular sink send data to the sink's chosen identifier(s). This process is illustrated in the figure below.

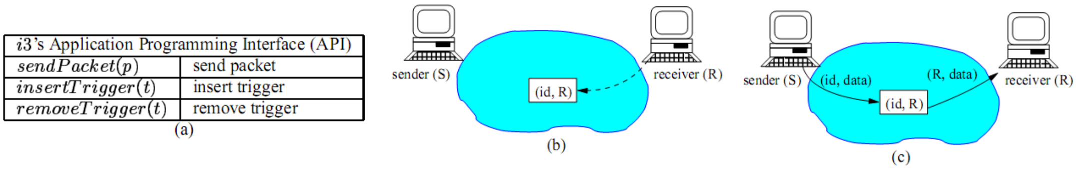


**Fig. 2.5 - i3 API (a), trigger (b), and send-receive (c) (courtesy of [Sto2004])**

Rendezvous-based indirection provides a natural means to support a wide variety of desirable features, including host mobility, anycast, multicast etc. The communication paradigm is essentially publish-subscribe (see Section 3.2 for an explanation of pub-sub communication within the context of PSIRP) and a variety of prominent services and applications (e.g. CDNs) become relatively trivial to implement.

i3 is based on the Chord DHT optimized with caching and location-aware trigger enhancements which greatly reduce suboptimal routing overheads associated with arbitrary bindings between DHT nodes and the underlying network topology.

#### *2.5.1.2 ROFL*

Routing on Flat Labels (ROFL) [Cae2006] assesses the suitability of a Chord-based routing architecture composed of an entirely flat keyspace. In this respect, the notion of location is not only diverged from identity but effectively discarded altogether. For added security and authentication, the ROFL identifier space is intermingled with public keys. A node's identity is given by the hash of its public key (as in HIP).

This approach claims the following key advantages:

1) There is no need for any additional infrastructure since identifiers are routable and need not be resolved to anything.

2) Packet delivery is not dependent on anything beyond the DHT-dictated delivery path.

3) Hierarchy, topological locality etc. are not involved in identifier assignments.

4) Network access controls can be implemented directly based on identities instead of hierarchical network addresses.

Unfortunately, scalability and efficiency are relatively poor and [Cae2006] effectively serves to show that an architecture based on flat-labeling has only a limited extent of feasibility.

#### 2.5.1.3 SEATTLE

Most people overlook the fact that the main purpose of IP is to alleviate the scalability problems of simple Ethernet by segregating Ethernet networks with routers and introducing hierarchical routable addresses. However, IP routing also negates Ethernet's most prominent advantages: easy manageability and configuration-free operation. Unlike IP, Ethernet uses flat MAC addresses that are not tied to host location. The Scalable Ethernet Architecture for Large Enterprises (SEATTLE) [Kim2008] investigates the possibility of extending flat Ethernet bridging to large networks by introducing a one-hop DHT based on keyed MAC addresses.

The SEATTLE DHT is composed of switches; end hosts do not participate. Once a host connects to the network, its local switch applies a hashing function separately to the host's MAC and IP addresses and stores them separately in the DHT along with the switch's identity and the host MAC, respectively. End-to-end connectivity is bootstrapped by instructing the local switch to locate the MAC of the destination node within the DHT using the destination's IP as a key. Subsequently, the destination's MAC is used to discover the destination's local switch in the DHT and forward an initial packet to the destination. Lastly, control messages are forwarded from the destination to the source in order to setup a shortest path for future forwarding. This procedure is illustrated in Figure 2.6.

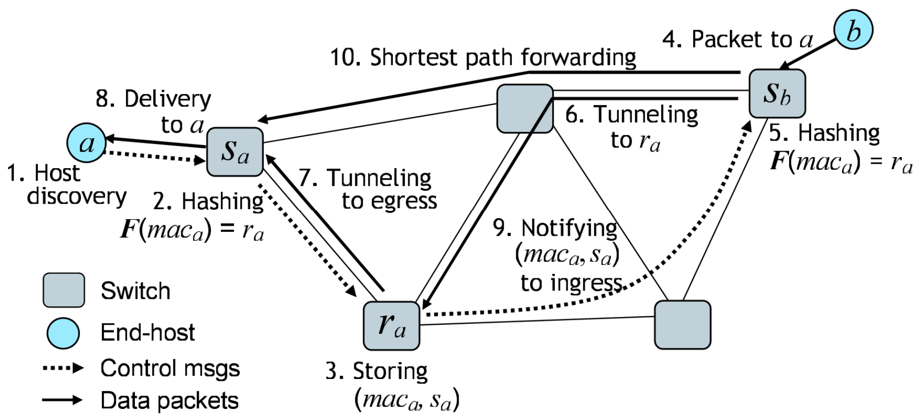


Fig. 2.6 - SEATTLE end-to-end communication (courtesy of [Kim2008])

SEATTLE's approach is advantageous because it is backwards compatible with existing applications and infrastructure and per-packet processing overheads are reportedly low. Most importantly, SEATTLE leverages the scalability of DHTs to address the instances where Ethernet and IP fail to scale. This includes ARP broadcasting, DHCP configuration, IP mobility etc. Nevertheless, the limitations of DHTs are still present and SEATTLE does not appear to circumvent the send-receive communication paradigm.

### 2.5.2 TRIAD

The Translating Relaying Internet Architecture integrating Active Directories (TRIAD) [Che2000] proposal is unique because it actually seeks to remove the need to transition to IPv6, instead relying chiefly on NAT to relieve address exhaustion. TRIAD exclusively uses names and URLs for end-to-end identification and leverages this “content layer” to integrate directory, routing, and connection-setup in order to route based on content. TRIAD also implements “path-based routing” based on a shim protocol which is layered on top of IPv4.

In principle, TRIAD provides end-to-end communication and preserved semantics across NAT without sacrificing the benefits of NAT. Moreover, TRIAD’s shim protocol extends IP and allows the Internet to scale potentially indefinitely without mandating a shift to IPv6. The directory service efficiently supports human-friendly name identification and authentication and communications are supposedly subject to a low space-time overhead.

To the best of our knowledge, the latest available TRIAD publications are over a decade old and the architecture is likely deprecated.

### 2.5.3 NIRA

The New Internet Routing Architecture (NIRA) [Yan2003] attempts to resolve the scalability and fault-isolation limitations of the current Internet’s routing system. A variety of these problems occur due to commercial aspects of the modern Internet and a lack of choice in wide-area providers from the user perspective. NIRA designers stipulate that offering user choices in domain-level routing will stimulate competition and foster innovation, and NIRA attempts to provide the technical foundation to achieve these objectives.

Within NIRA, users are treated as abstract agents under a strict hierarchical addressing scheme rooted at the provider. The 128 bit IPv6 syntax is applied and name lookup is achieved by a Name-to-Route Lookup Service (NRLS). The Topology Information Propagation Protocol (TIPP) supports scoped route propagation and provides a consistent link-state-like network view with fast convergence. Failure handling is implemented through a combination of router feedbacks, timeouts, route switching, and TIPP notifications.

Although NIRA represents an innovative approach to redefining traffic routing over the Internet, potential incentive and participation problems may occur amongst users and ISPs. Moreover, comprehensive QoS is a possible uncertainty and associated configurations may be complex.

### 2.5.4 DONA

The Data-Oriented Network Architecture (DONA) [Kop2007] proposes an information-centric network architecture based on static flat self-certifying names of the form *P:L* composed by concatenating a hash *P* of the principal's public key with a label *L* chosen by the principal to provide uniqueness. Two primitives, *FIND(P:L)* and *REGISTER(P:L)*, are instantiated to indicate an interest in and the intention to serve the object denoted by *P:L*, respectively. These functions are supported by network devices known as "data handlers" which are responsible for name resolution and data caching.

This approach is similar to PSIRP although DONA is still based on an underlying IP infrastructure, which in principle means that it may be susceptible to the existing limitations of IP networking.

### 2.5.5 4WARD

The 4WARD project [4WA2008] is funded under the EU's 7$^{th}$ Framework Program and began at approximately the same time as the PSIRP project. The aim of the 4WARD project is to develop a new generation of dependable and interoperable wireless and wireline network technologies. The technologies are slated to offer excellent adaptability, cost-effectiveness, and developmental suitability.

The 4WARD architecture employs a naming scheme similar to that proposed by DONA and leverages publish-subscribe communication instead of send-receive [Ach2009]. Although not exclusively based on a DHT, 4WARD uses a DHT implementation to resolve network names and resource locations.

### 2.5.6 CCN

Content-centric Networking (CCN) [Jac2009] addresses the limitations of the current Internet's host-to-host communication paradigm by proposing a content-centric model whereby data is routed and publicly authenticable by use of public key cryptography.

CCN introduces two types of packets:

1) **interests**, which indicate an intent to receive certain data, and

2) **data**, which satisfy corresponding interest packets.

CCN organizes and names content in a hierarchical fashion similar to that observed with current Internet URLs. Backwards compatibility is maintained through the ingenious use of the Open-Shortest Path First (OSPF) and Intermediate-System Intermediate-System (IS-IS) routing protocols to transfer content amongst routers whose engines are capable of participating in the CCN. OSPF and IS-IS employ standard Abstract Syntax Notation 1 (ASN.1) type-length-value (TLV) encodings to distribute routing information, and these TLVs are suitable to communicate CCN Uniform Resource Identifiers (URIs). The brilliance of this approach lies in the fact that non-participating routers will ignore the seemingly invalid CCN TLVs and their presence should in no way compromise the existing network infrastructure beyond necessitating the occasional dropped packet.

The CCN forwarding model is constructed such that there is a one-to-one mapping between **interest** and **data** packets. An interest is nullified or "consumed" by matching data. Coupled with publicly-authenticable data, this is a powerful countermeasure against DoS and content substitution attacks. Moreover, application experiments involving VoIP demonstrate excellent performance, security, efficiency, and disruption tolerance.

# 3 The Publish-Subscribe Internet Routing Paradigm

*The work discussed in this section has been funded by the EU FP7 PSIRP project, contract INFSO-ICT-216173. The administrative overview is derived from [PSI2009a]. The technical overview is based on the work of the PSIRP architecture and implementation teams and is largely excerpted from [Tar2009]. The author of this thesis is a primary contributor to both of these sources.*

## 3.1 Administrative overview

The PSIRP project [PSI2008a] is an EU FP7 [EC2010a] project funded by the European Commission (EC) with a lifetime of 33 months (30 months base + 3 month extension) beginning in January 2008. Its ambition is to investigate major changes to the network layer of the current Internet, leading to the replacement of this and other low-level layers for the purpose of adopting a new form of internetworking: **information-centric publish-subscribe**.

### 3.1.1 Motivation

Experts all over the world are beginning to agree that a fundamental reform of the Internet's paradigms and core technologies is needed to cope with the challenges presented in the new millennium. The endpoint-centric send-receive communication paradigm that underlies the current Internet, its predecessors, and even the telephony network, places rather arbitrary topological constraints on the delivery of information. With the observed increase of information-centric services such as the World-Wide-Web (WWW) and contemporary applications such as sensor networks, the inflexibility of endpoint-oriented topologies increasingly places a burden on solution developers that need circumvention by virtue of an increasing number of overlays. This leads to a lack of flexibility and increasing architectural rigidity. Moreover, these modernistic applications bring to light a number of inadequacies manifested in areas such as mobility, naming, security, routing, scalability etc. which can currently only be addressed by rough "patchwork" solutions on top of an ossified core protocol stack.

**The worst consequence is that the full range of possibilities offered by the Internet is not being exploited and trust in its proper operation has been lost.**

The PSIRP project ascertains that the main reason for the shortcomings of current IP-based internetworking is deeply embedded in its underlying communication paradigm rather than in its operational characteristics. Thus, PSIRP aims to resolve these issues by investigating a new information-centric publish-subscribe communication paradigm that could potentially serve as the flexible foundation for a new Internet. For this, we consider any communication scenario as being constituted by the production, retrieval, and consumption of information, all of which are surrounded by the concerns of the different parties involved in a particular situation. The notion of intention becomes crucial in the establishment of communications between two parties, where the match of intentions, and not the reachability of endpoints, is the key to successful communication. With this, the PSIRP approach moves away from the currently sender-driven IP model and towards a receiver-controlled operating philosophy.

In such receiver-controlled pub-sub networking, senders “publish” what they want to send and receivers “subscribe” to the publications that they want to receive. In principle, no one receives any material to which they have not explicitly expressed an interest by way of subscription. The result is a powerful yet flexible infrastructure with a high degree of resiliency.

One can observe that a large share of the Internet’s usage is already pub-sub in nature:

- Update dissemination (e.g. software, web feeds etc.)
- News delivery
- General media broadcasting (e.g. audio/video feeds and IPTV)
- Periodic and aperiodic messaging services
- E-mail

etc.

In addition, contemporary areas of research such as sensor networks and context awareness also rely on pub-sub communications to provide services to end users.

It seems promising to derive a new Internet architecture based on an information-centric pub-sub paradigm, leading to a redesign of all Internet communication layers. In such a new Internet, **multicast** and **caching** will be the norm, as opposed to unicast and source-retrieval, and **security** and **mobility** will be designed into the architecture rather than added as afterthoughts.

Placing information at the center of design considerations has many impacts on the PSIRP architecture. For instance, the support for multicast and anycast primitives for data delivery becomes a native mode of operation in an information-centric world. In addition, many traditional difficulties, such as mobility and multi-homing, become easier to solve, some (e.g. distributed file storage and retrieval) even trivial. With clients indicating their willingness to receive data by subscribing to it, it is the aim of the network to act as a substrate for the ensuing data delivery process from potentially distributed data sources. However, the scalability of information and interest-oriented networking is still a major challenge, and PSIRP intends to lead the Internet to a scalable information-centric communication foundation.

### 3.1.2 Structure

Project work is divided into five work packages (WPs) which each span across several project participants:

- **WP1 – Management:** General project management, technical management, and management of intellectual property rights (IPR).
- **WP2 – Architecture Design:** Technical design of PSIRP's information-centric pub-sub architecture.
- **WP3 – Implementation, Prototyping, and Testing:** Development and testing of a prototype implementation of the proposed architecture.
- **WP4 – Validation and Tools:** Development and application of tools for quantitative and qualitative validation of the proposed architecture.
- **WP5 – Dissemination and Exploitation:** Investigation of industry requirements and constraints, development of a migration path, engagement with the wider research community, and general dissemination of project results.

The foundation of this segregation is an iterative "life-cycle" engineering approach composed of systematic development stages that are synchronized amongst all work packages. The ultimate goal is to avoid isolated development "silos" which can in turn lead to problematic integration at the concluding stages of the project. Thus far, this methodology has shown great flexibility and proven to be extremely effective throughout the duration of the project and its terminating phases [PSI2008b].

### 3.1.3 Participants

The FP7 PSIRP project's membership has included two leading European telecom vendors, one of the largest telephone companies, and several highly-ranked academic institutions:

- Athens University of Economics and Business (AUEB)
- BT (formerly British Telecom)

- Ericsson Hungary Limited (ETH) **(until 31.12.2009)**
- Helsinki University of Technology **(now Aalto University)** – Helsinki Institute for Information Technology (TKK-HIIT)
- Institute for Parallel Processing, Bulgarian Academic of Science (IPP-BAS)
- Nokia-Siemens Networks Finland (NSNF)
- Oy LM Ericsson Ab (LMF)
- RWTH Aachen University (RWTH)
- University of Cambridge (UC) **(as of 1.1.2010)**

### 3.1.4 Miscellaneous facts

- **Duration:** 30 months base (January 2008 – June 2010) + 3 month extension
- **Total Cost:** €4.1m
- **EC Contribution:** €2.5m
- **Coordinator:** Helsinki University of Technology (TKK) – Helsinki Institute for Information Technology (HIIT)

## 3.2 Technical overview

The PSIRP project will redesign the entire Internet architecture from the pub-sub point of view, taking nothing, not even IP, for granted. PSIRP's work will focus on the intersection of security, scalability, trust, usability, and network economics, in order to design and develop efficient and effective solutions. **Multicast** and **caching** will replace unicast and source-retrieval, while **security** and **mobility** will be embodied directly into the foundation of the architecture rather than being added as afterthoughts [PSI2008d] [PSI2009a].

PSIRP aspires to change the routing and forwarding fabric of the global Internet so as to operate entirely based on the notion of information (associated with a notion of **identifiers** to support fabric operation) and its surrounding concerns, explicitly defining the **scope** of the information and directly addressing information (via **rendezvous identifiers**) as opposed to addressing physical network endpoints, Figure 3.1.

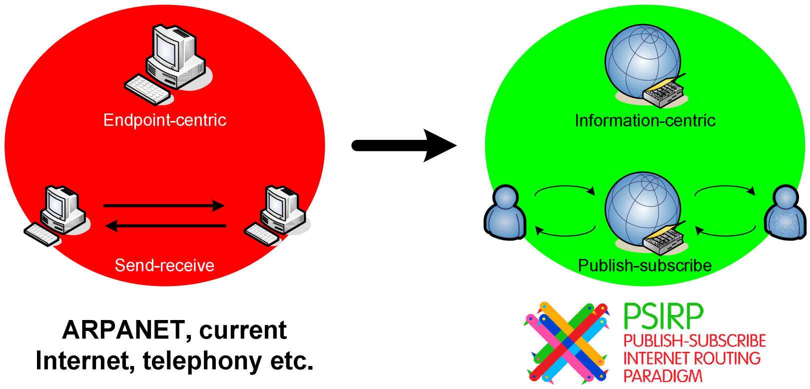


Fig. 3.1 - Endpoint-centric send-receive vs. information-centric publish-subscribe

The PSIRP model emphasizes information-centric operation: data pieces are explicitly addressed through identifiers serving as high-level designations/resolvers to lower-level schemas, and scoping mechanisms that can define information inter-networks and relationships within a global information taxonomy, Figure 3.2. Information is embedded immediately into the network and it is the only effective element in need of direct user-manipulation; the physicality of the network (i.e. endpoint locations) need not be known directly.

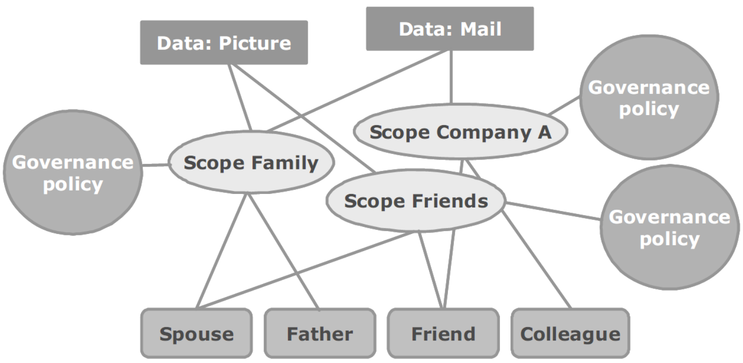


Fig. 3.2 - Scoping information

Another important aspect of the PSIRP architecture is that it is receiver-driven. We adopt the approach that the receiver should always have control over what they receive and we cascade this view throughout the core of the PSIRP component wheel and its multiple operational elements, Figure 3.3. A receiver must elect to join (i.e. subscribe) to an identifier before it can receive any information. Sending (i.e. publishing) as well as receiving operations are thus decoupled between the senders and receivers in both time and space. PSIRP not only intends to move the functionality of many existing publish-subscribe systems onto the internetworking layer but also base all communications throughout the architecture on this paradigm.

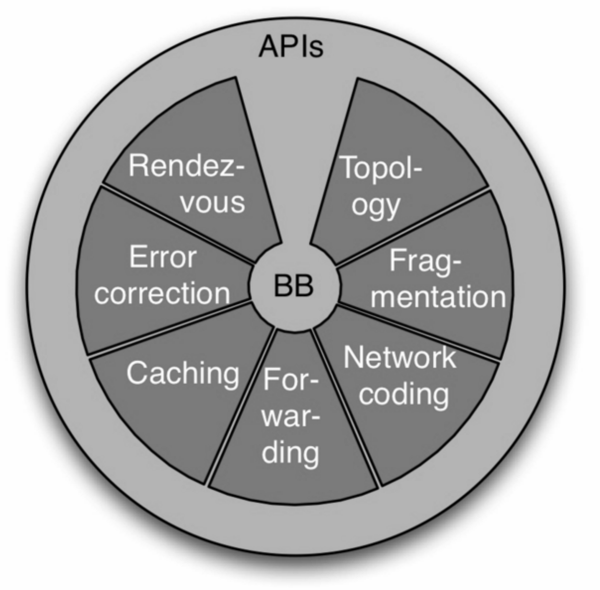


**Fig. 3.3 - The PSIRP component wheel**

The PSIRP conceptual architecture is based on a modular and extensible core, called the PSIRP component wheel. The architecture does not have the traditional stack or layering of common telecommunications systems, but rather components that may be decoupled in space, time, and context, surrounding a "blackboard" (BB) for pub-sub operations. Above the component wheel, we have APIs that facilitate accessibility to and implementation of different networking features that are available in the system. Figure 3.3 illustrates the typical components needed in the wheel for inter-domain operation: rendezvous, error correction, caching, forwarding, network coding, fragmentation, and topology.

The idea of such a layer-less network stack has been proposed before (e.g. in the Haggle architecture [Hag2007]). The novelty of the PSIRP proposal is to use publish-subscribe style interaction throughout the conceptual architecture and thus support a layer-less and modular protocol organization. This organization is primarily achieved through the efficient structuring of information identifiers and their interactions amongst network elements, Figure 3.4, offering ample flexibility for future expansion.

We can view the global network of information as an acyclic graph of related pieces of data, each identified and scoped by some identifiers. Identifiers define the relationships between the pieces of information in the different operating levels of the PSIRP architecture, such as the application or networking level. With this in mind, we propose the following classes of identifiers:

- **Application identifiers (AId):** Used directly by publishers and subscribers to manipulate applications that interact with the network. These identifiers are in a human-readable format and serve to simplify network operations from the user and application points of view.
- **Rendezvous identifiers (RId):** Used to bridge higher level identifiers with lower layer identifiers. A rendezvous identifier is implicitly associated with a well-defined (but not necessarily fixed) data set consisting of one or more publications. The data sets may also have associated metadata, which may include scoping information and other useful information, either for the ultimate receivers or for network elements.
- **Scope identifiers (SId):** A type of rendezvous identifier used to delimit the reachability of information through contextual definition (e.g. person, location, time, genre etc.) and associated governance policies (e.g. access rights, usage policies etc.). Scoping information is associated with a publication, determining the elements of the rendezvous system that act on published data and therefore defining the information network that the publication belongs to. A publication may be associated with one or more scopes.
- **Forwarding identifiers (FId):** Used to define network transit paths that are used to transport publications across the network. The breadth of reference of FIds is variable, potentially limited to single hops or dynamically expandable to encompass full multicast trees. This relatively open structuring scheme allows concurrent use of FIds to support flexible routing mechanisms based on source routing, anycast, multicast etc.

The functional relationships between PSIRP identifiers and network operations is shown in Figure 3.4.

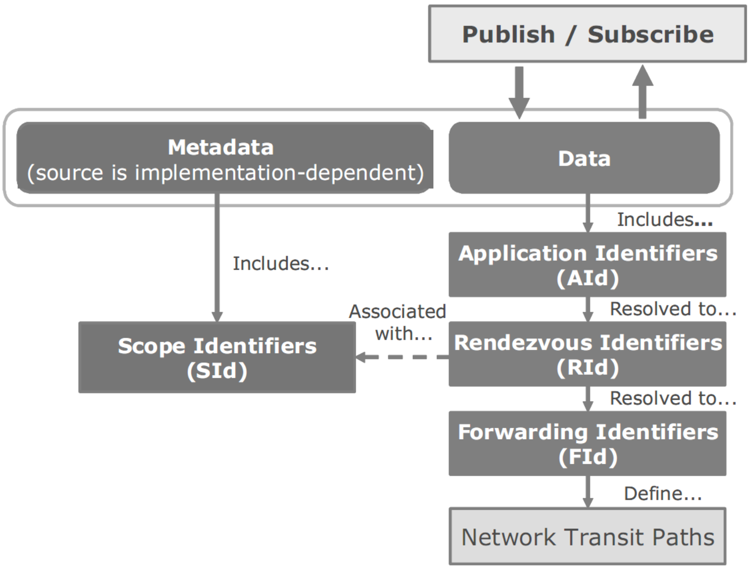


Fig. 3.4 - PSIRP identifiers and network functionality

## 3.3 Prototype

Largely through the work of Oy LM Ericsson Ab (LMF), the PSIRP consortium has developed a prototype to demonstrate the capabilities of its newly devised pub-sub architecture. The prototype, known as Blackhawk, is based on a 64 bit FreeBSD operating system and features a kernel-integrated Bloom filter forwarding engine. Full details are available in [Jok2009].

Bloom filters [Bro2003] were first devised by Burton Howard Bloom in 1970 as a space-efficient probabilistic data structure for data-set aggregation. A Bloom filter is defined as an array of size *m* bits; if the Bloom filter is empty, all bit positions are set to zero. To insert an item into the Bloom filter, *k* different hashes with an output size of *m* bits are taken and the bit positions of the Bloom filter corresponding to each of the *k* numerical hash result are set to "1." Multiple items can be added in this fashion and overlapping bit positions are set to "1." Verifying an item's membership in the Bloom filter is achieved by checking the values of the bit positions corresponding to its *k* hash outputs. In this fashion, Bloom filters offer a very efficient storage mechanism for a large number of data items without the possibility of false negatives. False positives are however possible since multiple item insertions can overlap all of the bit positions of a non-member item.

The Blackhawk forwarding implementation applies the Bloom filter concept in a slightly different manner, with the resulting structures, known as zFilters, being used as forwarding identifiers (FIds), Figure 3.5.

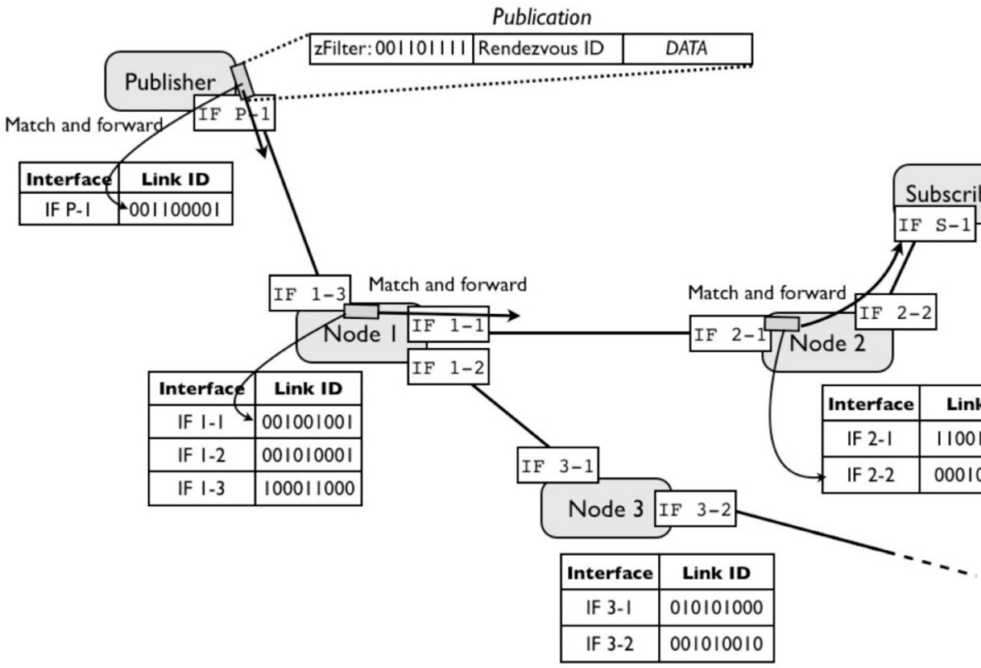


**Fig. 3.5 - zFilter forwarding (courtesy of [Jok2009])**

Rather than identifying interfaces, Blackhawk FIds identify unidirectional links. Thus, every bidirectional network link is indentified by two link identifiers (LIDs), one for each direction of data travel. The LIDs can be either locally generated (e.g. at random, based on unique node information etc.) or assigned by a central authority such as the network administration (e.g. directly, algorithmically etc.). The LIDs are 256 bits in size and 5 unique bits are set to "1." A forwarding path amongst a network of nodes and bidirectional links is created by computing the logical OR product of the LIDs of the involved links and performing strict source routing with this result. Upon receiving an incoming packet, a node checks the packet's destination zFilter and determines through which outgoing interfaces to forward the packet by separately computing the logical AND operation of the destination zFilter and the outgoing LID for each of its interfaces. If the output matches the LID, then it is included in the zFilter (or it is a false positive) and the packet is forwarded through that interface. If the output does not match the LID, the LID is not part of the zFilter and the packet is not forwarded through the corresponding interface. The end result is a dynamically expandable multicast forwarding tree which can potentially span the entire network.

Figure 3.5 illustrates the forwarding of a publication containing data, an RId, and a zFilter (i.e. a FId). The first receiving node (Node 1) performs an individual logical AND operation with the packet's zFilter and the LID of each of its interfaces. If the result indicates that a particular LID is included in the packet's zFilter, the packet is forwarded through the corresponding interface.

The presence of false-positives and "broadcasting" through full-zFilters have yet to be fully addressed, although a variety of mechanisms (e.g. link identity tags, or LITs) provide a reasonably thorough means of achieving statistically unique zFilters [Jok2009].

The zFilter forwarding mechanism is uniquely efficient in that it is almost wholly based on simple Boolean operations and can be performed on-the-fly, thus enabling cut-through forwarding (i.e. an incoming packet can be forwarded to its outgoing interface(s) before it has been fully received by the ingress interface). This makes the zFilter forwarding implementation notably fast and relatively frugal with network resources.

In accordance with EU FP7 dissemination policies, the first code release of the PSIRP framework implementation was made under the GNU Public License version 2 (GPLv2) and Berkeley Systems Distribution (BSD) licenses [Kjä2010]. In time, separately developed functional modules (e.g. rendezvous, topology, host-internal blackboard-based publication management and packet forwarding etc.) will be merged with Blackhawk to provide fully-functional architectural demonstrator.

# 4 The importance of dissemination and exploitation

Dissemination and exploitation routes for research innovations typically fall under two prevailing categories:

- **Passive** methods openly release information to the general public but are usually not accompanied by active engagements or ongoing support on the part of the source party, with the possible exception of periodic updates to ensure that public materials are up-to-date; passive dissemination is predominantly achieved through publications, open documentation, advertisements, conference proceedings etc.
- **Active** methods require the source party to actively support their dissemination efforts and offer ongoing direct interaction with external consumers; examples include educational courses, seminars, structured and monitored funding schemes (e.g. venture capitalism), supported technology demonstrators, recruitment events etc.

Research projects will typically employ both passive and active methods to disseminate and exploit their work. Whereas passive forms are easy to institute and virtually universally present – most all responsible research is properly documented – **active dissemination is both far more demanding and rewarding for the source party because the general public can be actively engaged and guided**, typically on more than one occasion, directly by knowledgeable and enthusiastic personnel from the source party. Section 5.1 discusses notable examples of both passive and active dissemination practices within the context of FP7 PSIRP.

Whether passive or active, the dissemination and exploitation of observations, techniques, and results is an invaluable aspect of any research project. Unless the public at large can openly access and make use of this information, it will be virtually impossible to educate external professionals and the innovations in question will have little chance of being exploited. This limits the development potential for both existing innovations and innovations which have yet to be discovered but are achievable through continuations of current research.

This section depicts the importance of dissemination and exploitation processes within the two research sectors that are most relevant to PSIRP and this thesis: the ICT division of EU FP7 and general future Internet research. We also revisit the deficiencies that are typically associated with dissemination and exploitation of revolutionary Internet architectures and reveal the **linkage between these deficiencies and the involvement of academics within future Internet research**. With this, based on the developments of the PSIRP consortium and reviews of related literature, we propose that **dedicated academic engagements offer a unique opportunity to combat these change-barriers through efficient dissemination and exploitation venues**.

## 4.1 FP7 ICT

The ICT research unit within the EU's 7th Framework Program (FP7) [EC2010a] is one of the European Commission's principal instruments to strengthen Europe's information society and media policies. A variety of project proposals are reviewed by the European Commission on a regular basis in response to strategically published calls for research and development in key ICT fields. Under strict scrutiny, the proposing consortiums with the most thorough and promising proposals are granted funding to proceed with their project plans. European and non-European institutions are encouraged to advance their research in a cooperative environment where directed progress is carefully monitored. With this, the European FP7 has successfully promoted research and related education with the end goal of making Europe one of the world's leading centers for progress and scientific innovation [Ehl2007].

The "knowledge triangle" – research, **education (i.e. creating understanding)**, and innovation – is at the forefront of the FP7 effort [EC2010b].

*"Knowledge lies at the heart of the European Union's Lisbon Strategy to become the 'most dynamic competitive knowledge-based economy in the world'. The 'knowledge triangle' – research, education and innovation – is a core factor in European efforts to meet the ambitious Lisbon goals. Numerous programmes, initiatives and support measures are carried out at EU level in support of knowledge.*

*The Seventh Framework Programme (FP7) bundles all research-related EU initiatives together under a common roof playing a crucial role in reaching the goals of growth, competitiveness and employment;"*

- Courtesy of [EC2010b]

The three components of the knowledge triangle – research, education, and innovation – clearly model the core chronological sequence that is required to successfully realize global change from preliminary future Internet research (Section 1.2) and arguably most any general ICT research. That is: **research** yields **innovations** and innovations must be disseminated; the first step to elicit change in the face of a scientific innovation is to **educate** the public in order to promote an awareness of the problem(s) that the innovation intends to solve, with the goal of fostering knowledge of the innovation and creating an **understanding** of why and how the solution in question is well-suited to address the given problem(s). In turn, this understanding fuels **perceived urgency** and **motivation**, and these three aspects eventually lead to **controlled progression**.

In this, research and development of the caliber conducted within FP7 are clearly not suitable in isolation; coordinated dissemination and exploitation efforts are necessary to share ideas, thinking, and results, and reduce fragmented development and overlapping research [Ehl2007]. Dissemination, chiefly to spread knowledge, is a cornerstone activity in the field of ICT as it serves a starting point for further research, industrial collaboration, and worldwide growth.

The European commission is very well aware of this and has a strict stance on dissemination and exploitation within FP7:

*"Participants in projects funded under the Seventh Framework Programme (FP7) are required to use and disseminate the results generated by the project ("foreground"). Dissemination is meant to promote the results as swiflty* [sic] *and effectively as possible to benefit the whole community and avoid duplication of R&D efforts.*

*Dissemination is also important to the interests of the participants. In fact, an adequate description of 'the potential impact through the development, dissemination and use of project results' at the application stage may positively affect the evaluation of the project proposal, thus contributing to its consideration for funding. Moreover, prompt and effective dissemination of the project results may help participants in the subsequent market exploitation and establishment of fruitful collaborations."*

- Courtesy of [IPR2010]

It then becomes apparent that **dissemination and exploitation processes are crucial to FP7 because they provide the chief means to fulfill the educational requirements set forth by the European Commission and promote the permeation of research results amongst external parties.** The PSIRP consortium is strictly responsible for ensuring that these requirements are met within the context of the PSIRP project.

## 4.2 Future Internet research

In Sections 2.1 and 2.2, we revealed how the Internet's developmental history and a variety of technical and socio-economic factors have ultimately led to the ossification of its architectural foundations. Consequently, it is increasingly difficult, if not impossible, to devise and successfully deploy technically sound solutions on the Internet. In effect, any new attempts to modify the underlying Internet architecture have been nearly impossible to manifest due to a lack of:

1) **Understanding** (i.e. little or no appreciation of existing problems and their potential solutions)
2) **Perceived urgency** (i.e. no recognized impending operational catastrophe)
3) **Motivation** (i.e. no perceived benefit or gap to fill)

These deficiencies are clearly apparent in the context of future Internet research because it is almost exclusively constituted of revolutionary and non-incremental solution proposals.

We've also discussed how dissemination and exploitation are critical to invoke an understanding of problems and their potential solutions amongst the general public. Nevertheless, PSIRP and other revolutionary Internet architectures all suffer from a common detriment: the dissemination and exploitation of advanced internetworking research results has been notoriously problematic. Consequently, without understanding, it has been impossible to inspire a sense of perceived urgency and motivation to apply these results for the benefit of the Internet, and commercial deployment and eventual global acceptance have been virtually non-existent.

We surmise that a root-cause of these failures is ineffective and/or insufficient dissemination and exploitation practices. In light of this, we believe that dissemination and exploitation, chiefly to educate and spread knowledge, must be reinforced as cornerstone activities within future Internet research and development.

### 4.2.1 The involvement of academics

Future Internet research is a notably challenging field due to its inherent technical complexity and the extent to which it impacts the world. Clean-slate internetworking proposals push the limits of creativity and technological advancement and these approaches are notoriously difficult to develop and understand because they require extensive experience and a robust academic background.

As a bare minimum, one must possess a thorough technical background and understand the fundamental sciences with which form the basis of computer networking and packet-switching theory (e.g. mathematics, electronics, algorithmics etc.). This type of education serves as a foundation to enable effective learning in the face of new computing concepts. Furthermore, one must also be familiar with current internetworking technologies and possess a historical understanding of the Internet, its underlying design, technologies and mechanisms, developmental milestones etc., chiefly to understand the mistakes of the past and appreciate how the Internet must grow to meet the demands of the new millennia.

In addition, a variety of fundamental fields and social sciences are elementarily involved in pioneering technological innovation. The ability to construct, comprehend, and apply abstract concepts from these disciplines is paramount to guide the direction(s) of technical development in the hands of human ingenuity. Future Internet research is no exception, especially in light of the fact that the Internet is arguably one of the most globally far-reaching constructs in terms of social and economic consequence. Initial clean-slate internetworking developments are largely intangible and innovation is almost always promoted through implementations that are based on designs and abstractions in fields such as physics, mathematics, systems theory, machine learning, human-computer interaction, psychology etc.

Largely as a result of the aforementioned complications, typical dissemination and exploitation methods appear relatively ill-suited when it comes to educating the general public on the innovations produced by advanced future Internet projects. **Where research is complex and heavy academic backgrounds are requisite, strong academic involvement is typically needed to disseminate the corresponding results.** This has been clearly evident within the PSIRP project as its dissemination and exploitation work package has struggled to effectively reach third parties.

The main problem disseminating revolutionary Internet solutions such as PSIRP is two-fold:

- **Complexity:** Even experienced professionals are apprehensive over embracing clean-slate technologies because their complex disparity in comparison with traditional internetworking practices makes it difficult for unacquainted users to gain a thorough understanding of their workings. Thus, the intended audience will not appreciate the problems at hand and have little to no understanding of their potential solutions. *ASIDE: Most researchers and engineers in the 1960's (and, in some cases, even into the 1990's!) had a great deal of trouble grasping the notion of packet-switching in contrast with the traditional circuit-switching communication paradigm made popular by the early switched phone system.*

*Today we are in fact faced with the same problem when it comes to endpoint-centrism vs. information-centrism and send-receive vs. publish-subscribe!*

- **Timespan:** The timeframes of these projects are not aligned with those of their real-world ambitions and implementations. It is thus very difficult to create a sustainable line of exploitation that will eventually lead to successful deployment and controlled progression

Academic environments are unique because they not only foster an initial understanding of low-level concepts and complex innovations (and have a long successful history in this regard) but also provide a learning vehicle that can continue fostering knowledge indefinitely beyond the conclusion of a technical project. Academic institutions have long been the primary point of inception and ongoing development for innovative ideas and structures, and courses in academia are specifically formulated to train individuals in unfamiliar new concepts and methods. Academia is effectively designed for state-of-the-art dissemination, both through professional networking and traditional coursework, and has historically proven itself to be well-suited as a starting point in this respect.

Moreover, academia provides an atmosphere that is traditionally further removed from the motivators which typically serve commercial and industrial sectors (e.g. profit, politics etc.). This is especially important in the context of Internet progression as it provides a means to combat misaligned incentives and resource allocation that continue to worsen the Internet's state of ossification and problematic operation (Section 2.2). Furthermore, universities, research institutions, and their members and alumni are and will continue to be majority stakeholders in future Internet research [Rob2006] (a majority of the PSIRP consortium's members are universities or university-based) and their involvement in its ongoing development cannot be limited exclusively to technological endeavors. Universities are abundant sources of intellectual capital whose main value is realized through efficient dissemination to present and future audiences [AAU2009]. Without academic exploitation, this capital is wasted [AAU2009].

Dissemination is a paramount consideration in the context of innovation and concerning the best practices of ICT education [Ell2002]. The recent introduction of enhanced digital mediums and the increased permeation of information-dissemination mechanisms in ICT throughout the world have shed light on academia's increased responsibility to take a more active role exploiting the knowledge that it produces [AAU2009]. The modern industry is heavily reliant on academia, not only for preparatory instruction, but also as a source of applicable technological innovation and ongoing training of the workforce [Fas2000]. The importance of academic-industrial liaison should not be underemphasized, as is the need for a developmental feedback loop between academia's innovations and the practical applications of the industry [Fas2000].

As discussed in Section 4.1, the European Commission stringently requires dissemination and exploitation from its FP7-funded projects. Similarly, it is becoming increasingly well recognized that obtaining optimal results from funded technology ventures requires dedicated dissemination and exploitation processes that are tightly interwoven with academic environments. A 2009 report by the Association of American Universities, the Association of Research Libraries, the Coalition for Networked Information, and the National Association of State Universities and Land Grant Colleges [AAU2009] states:

*"In a networked environment* [such as an FP7 project consortium] *one maximizes technology investments by integrating dissemination functions directly into existing university technology environments.*

*…*

*Campuses should initiate discussions involving administration and faculty about modifying current practices… to ensure that broad dissemination of the research and scholarly work produced by its faculty occurs."*

- Courtesy of [AAU2009]

Taking this information and the points expressed earlier in Section 4 into account, we believe that academic future Internet project members and research institutions have a responsibility to explore the most effective means of fostering an understanding of current and foreseeable Internet problems and their potential solutions. This is the key starting point through which perceived urgency, motivation, and finally controlled progression, can be stimulated. **Academic means of dissemination, chiefly to educate and create understanding, appear to be a promising route to exploit the innovations of future Internet research.**

# 5 FP7 PSIRP exploitation

The PSIRP project has set out to re-examine some of the crucial fundamentals of the current Internet, leading to the outline, specification, and early implementation of a possible future Internet architecture – a future that focuses on the intersection of security, scalability, trust, usability, network economics, and a balance of power(s) in communication. In this, PSIRP WP5: Dissemination and Exploitation represents a major engagement effort as mandated by the European Commission.

This section provides a short overview of WP5 and presents the details of our dissemination course T-110.6120. WP5's policies and engagements reveal a very limited range of dedicated educationally-oriented dissemination activities amongst PSIRP's academic partners. With this, we reiterate our reasoning for extending PSIRP's dissemination and exploitation processes through academics, laying the framework for our envisioned courses as a promising extension of WP5's existing agenda. Lastly, we give a thorough overview of T-110.6120, discussing its execution plan, operational objectives, structure, content, and operating methods.

## 5.1 WP5: Dissemination and Exploitation

*ACKNOWLEDGEMENTS: We would like to especially acknowledge the contributions of the PSIRP WP5 team and their deliverables [PSI2008e] [PSI2008f] [PSI2008g] [PSI2009b] to our preparation of this section.*

WP5 is overseen by the PSIRP engagement team whose responsibility is to work towards a methodology within the project that defines dissemination and exploitation to the wider community beyond the PSIRP consortium. This primarily includes

1) thoughts around the openness of PSIRP **documentation** (i.e. passive dissemination), and

2) organized **engagements** (i.e. active dissemination) including workshops, conferences, development and directional sessions etc. with particular communities (e.g. academic, industrial, governmental etc.),

with the end goals of disseminating PSIRP results and eventually circumventing the barriers to global change on the Internet.

### 5.1.1 Documentation

All developments in PSIRP are documented through technical deliverables to the European Commission as well as through research papers and proceedings from various engagements (Section 5.1.2). As agreed in PSIRP's FP7 contracts, most of PSIRP's deliverables are openly available. PSIRP makes use of its online resources, including a public homepage [PSI2008a] and semi-public Wiki [PSI2008b], to establish a partition that holds documents to be released to the greater community. This mechanism has proven to be a valuable tool throughout the duration of the project for the purpose of engaging external partners in meaningful discussions and providing records of our work and results. However, comprehensive documentation alone is not sufficient to effectively disseminate project results; dedicated active engagements are necessary to motivate and direct outside communities.

### 5.1.2 Engagements

Since January 2008 the PSIRP consortium has participated in over 20 external dissemination events and published over 35 scientific papers through prestigious venues which include ACM's SIGCOMM and ReArch. The following sections outline some of the most prominent of these dissemination markets.

#### *5.1.2.1 European projects*

The **ICT SHOK** program represents a global research effort involving Europe, the United States, China, and others, to concentrate key research resources in order to develop future internetworking technologies and inspire new global business ecosystems based on advanced ICT foundations. The work of FP7 PSIRP has largely contributed to the success of ICT SHOK, mainly through its Information Networking division focusing on the storage, dissemination, and access of information.

The Athens University of Economics and Business (AUEB), a partner in the PSIRP consortium, maintains membership in the FP7 **"Euro-NF: Anticipating the Network of the Future – From Theory to Design"** project supported under the Networks-of-Excellence (NoE) funding scheme designed to help integrate European research fragments. AUEB staff have presented PSIRP developments to their Euro-NF peers on several occasions, contributed to the Euro-NF vision document, and actively disseminated PSIRP material through several Euro-NF events and workshops. The extent of Euro-NF involvement represents a major route for ongoing marketing and research conditioning throughout the EU.

PSIRP staff have also contributed to the FP7 **EIFFEL** support action which has engaged in reciprocal ongoing development with PSIRP.

PSIRP has directly contributed requirements, results, and processes for the development of technical platform solutions in the **OneLab2** EU-funded research project. In addition, PSIRP and OneLab2 have compiled a joint technical report which outlines requirements and conceptual solutions for a framework to evaluate inter-domain networking solutions such as the rendezvous solution of PSIRP. PSIRP has also held development sessions with OneLab2 in an effort to devise experiments centered around the Bloom filter forwarding mechanism of the PSIRP Blackhawk prototype.

#### 5.1.2.2 International initiatives and projects

The **University of Essex** has recently made the PSIRP Blackhawk prototype available via its campus-wide wireless network. Over 2500 students have access to the test bed from within their dormitories and administrators hope to generate substantial testing in the future.

The **University of Campinas (UniCamp)** has been actively experimenting with the PSIRP prototype, documenting bugs in the code and actively disseminating PSIRP material via workshops and guest lectures in postgraduate courses. UniCamp has also implemented a Firefox plug-in for Blackhawk allowing users to directly handle SId/RId subscriptions from within a familiar browser environment. This is a major achievement as it represents a first step to create a user-friendly interface that will better enable the public to interact with and understand the PSIRP prototype. It is our intention to actively employ the Firefox plug-in within the PSIRP application deployment course T-110.6100 (Section 7.2).

PSIRP has also been collaborating with the Communications Futures Program (CFP) at the **Massachusetts Institute of Technology (MIT)** and produced a joint paper discussing argumentation for a new internetworking architecture such as that provided by PSIRP, along with a whitepaper on identity in information-centric networking.

#### 5.1.2.3 Academic partners

Five of the current members of the FP7 PSIRP consortium are universities or university-based and exert a strong influence within their academic circles. This section covers the current extent of academic dissemination and exploitation activities in the project.

The **Athens University of Economics and Business (AUEB)** has three faculty members and a total of nine graduate and postgraduate students involved in the PSIRP project. AUEB has published over half-a-dozen workshop and conference papers related to the PSIRP effort and incorporated PSIRP concepts in existing graduate and undergraduate curricula in communications, multimedia technology, and distributed systems. There are also plans for several PSIRP "spin-off" projects.

**RWTH Aachen University** and its Department of Wireless Networks have been actively involved in PSIRP-related research, notably in the fields of network coding and publish-subscribe topology management. Offered courses in ad-hoc networking and mobile computing have also been redesigned to incorporate content-centric networking concepts and a number of graduate and postgraduate students have been recruited into the project. RWTH is also a leading proponent of the merits of PSIRP's publish-subscribe technologies in other computing domains such as cognitive wireless networking and wireless sensing. Recently, RWTH has also begun exploiting PSIRP results through several industry collaboration projects and usage studies.

The **Institute for Parallel Processing at the Bulgarian Academic of Science (IPP-BAS)** has devoted a seminar to PSIRP and created a dedicated Internet forum space for the project in the Bulgarian Research and Education network. A single graduate student has been recruited and components of the PSIRP architecture have been discussed in a "Global Networks" graduate course.

### 5.1.3 Limitations

Despite its time-tested documentation practices and an extensive and reasonably successful engagement agenda, WP5 has still experienced some difficulties reaching third parties on different levels. Our project experience has shown that these difficulties arise primarily because our audiences are either apprehensive or more commonly have difficulty gaining an adequate understanding of PSIRP concepts.

PSIRP's academic partners have involved graduate and postgraduate students in the project, and, to a limited extent, integrated PSIRP material within **existing academic coursework**. However, a lack of **dedicated academic coursework** among WP5 efforts is apparent. PSIRP concepts may be integrated and conveyed somewhat through existing course offerings but they will typically not achieve a standing within the course subject matter that is central enough to effectively exploit the project. That is, the overall goals and content of existing academic courses overshadow our intended PSIRP-specific dissemination targets.

On the other hand, **dedicated academics** are a relatively unexplored dissemination venue within PSIRP and uniquely enable the use of passive and active dissemination tactics for the focused exploitation of project results over a potentially indefinite time period. That is, **dedicated academic courses allow project staff to bring PSIRP to the forefront of attention, using project documentation as a principal learning tool**

**and actively engaging a motivated audience through lectures and other ongoing events**. In these respects, the need for dedicated academic coursework stands out as a strikingly obvious area for potential improvements.

*"Lectures are probably the best teaching method in many circumstances and for many students; especially for communicating conceptual knowledge, and where there is a significant knowledge gap* [as in PSIRP] *between lecturer and audience."*

- Courtesy of [Cha2006]

While PSIRP serves as the centerpoint of these analyses, we strongly believe that our results will be applicable to any fundamentally new internetworking paradigm.

## 5.2 Academic exploitation plan

This section presents the academic exploitation plan explored in this thesis. We summarize pertinent information from past sections which acts as the basis for our intended work and cover the details of our execution plan, expert panel, operational objectives, structure and content selection, and general operating methodology.

### 5.2.1 Basis

Despite significant activities in European projects, international initiatives, and its academic partners (Section 5.1), PSIRP has yet to use dedicated academic courses as a means to educate the general public and disseminate its results. Recent developments, largely summarized in Sections 4 and 5.1.3, have reinforced the notion that academics may provide a key dissemination outlet which is less affected by the main detriments observed through traditional dissemination and exploitation venues. These tenets, summarized below, form the basis that substantiates the course-based academic dissemination approach explored in this thesis:

Section 1.2 – With regard to eliciting change in the face of a scientific innovation:

- The first step to elicit change in the face of a scientific innovation is to promote an understanding of the problems it intends to solve, with the goal of fostering knowledge of the innovation and creating an understanding of why and how the solution in question is well-suited to address the given problem(s).
- Dissemination and exploitation are requisite to create the aforementioned understanding.

Section 2.2 – With regard to the progression of the Internet:

- Ineffective and/or insufficient dissemination and exploitation practices are a root cause of the many technical and socio-economic factors that have obstructed the progression of the Internet.

Section 4 – With regard to the nature of dissemination and exploitation processes:

- Dissemination and exploitation activities fall into two predominant categories: passive and active; the latter are typically more demanding and rewarding than the former.

Section 4.1 – With regard to the importance of dissemination and exploitation to the ICT division of the European Commission's 7th Framework Program:

- The "knowledge triangle" – research, education, and innovation – is at the forefront of the EU's FP7 effort and clearly models the core chronological sequence that is required to successfully realize global change from preliminary future Internet research and arguably most any general ICT research.
- Research and development of the caliber conducted within FP7 are not suitable in isolation; coordinated dissemination and exploitation efforts are necessary to share ideas, thinking, and results, and reduce fragmented development and overlapping research.
- The European Commission has strict guidelines enforcing the dissemination of results and innovations from FP7 projects.

Section 4.2 – With regard to the importance of dissemination and exploitation to the future Internet:

- It has been increasingly difficult, if not impossible, to devise and successfully deploy technically sound solutions on the Internet.
- Revolutionary Internet architectures have experienced problematic dissemination and exploitation and lackluster global acceptance.
- Understanding is the key starting point through which perceived urgency, motivation, and finally controlled progression, can be stimulated.

Section 4.2.1 – With regard to the involvement of academia in future Internet dissemination and exploitation:

- Future Internet research is a notably challenging field; it is inherently complex and requires an extensive experience and a solid academic background to grasp.

- Where research is complex and heavy academic backgrounds are requisite, strong academic involvement is typically needed to disseminate the corresponding results. In this, there is a distinct linkage between the deficiencies observed in disseminating revolutionary Internet solutions and the lack of appropriate academic involvement.
- Typical dissemination, exploitation, and marketing methods may not be well-suited to disseminate the type of results produced by advanced future Internet projects to the general public; inherent complexity and misaligned timeframes are of key importance.
- Academic environments foster an initial understanding of low-level concepts and complex innovations (having a long successful history in this regard) and provide a learning vehicle which can continue fostering knowledge indefinitely beyond the conclusion of a technical project.
- Academic institutions have long been the primary point of inception and ongoing development for innovative ideas and structures.
- Academia is especially designed to educate the general public on previously unknown and/or advanced subjects and it has historically proven itself to be well-suited as a starting point to spur public understanding in this respect.
- Academic environments are unique as compared to the general industry because they provide a learning atmosphere that is further removed from the motivators which typically serve commercial and industrial sectors (e.g. profit, politics etc.); this is especially important in the context of Internet progression as it provides a means to combat misaligned incentives and resource allocation.
- Universities, research institutions, and their members and alumni represent majority stakeholders in future Internet research.
- Universities are abundant resources of intellectual capital whose main value is realized through efficient dissemination to present and future audiences.
- The modern industry is heavily reliant on academia, not only for preparatory instruction, but also as a source of applicable technological innovation and ongoing training of the workforce.
- The importance of academic-industrial liaison cannot be overemphasized, as is the need for a feedback loop between academia's innovations and the industry's practical implementations in cutting edge research.

- Technology investments such as those instituted by FP7 are maximized by directly integrating exploitation and dissemination functions into university technology environments.

Section 5 – With regard to the past and current state of dissemination and exploitation efforts within the FP7 PSIRP project:

- The PSIRP project has an extensive engagement work package which has attempted to disseminate and exploit PSIRP results through open documentation and engagement events.
- Five of the current members of the FP7 PSIRP consortium are universities or university-based and exert a strong influence within their academic circles.
- FP7 PSIRP and most other revolutionary clean-slate architecture projects have hardly touched upon dedicated academic courses as a means of disseminating and exploiting their results.
- PSIRP results are not necessarily well integrated and conveyed through existing academic courses as their content overshadows PSIRP-specific dissemination and exploitation targets.
- Dedicated academic courses allow project staff to bring PSIRP to the forefront of attention, using project documentation as a principal learning tool and actively engaging a motivated audience through lectures and other ongoing events.

**With this, there is substantial evidence pointing towards the potential benefit of dedicated academic dissemination methods for the clean-slate internetworking architecture proposed by the PSIRP project. We also believe that these analyses and results will hold true for other future Internet architectures.**

### 5.2.2 Execution

Given the tenets outlined in Section 5.2.1, we have elected to expand PSIRP's exploration of academic means of satisfying the dissemination and exploitation requirements set forth by the European Commission by instituting two successive special-topic courses within the Faculty of Information and Natural Sciences of the School of Science and Technology at Aalto University in Espoo, Finland, whose aims will be to promote:

1) **information dissemination** (course code T-110.6120) and
2) **application development** (course code T-110.6100),

respectively, for the EU FP7 PSIRP project. These courses will take place during the spring 2010 term (January – May) and will be targeted towards advanced graduate and postgraduate students who possess a thorough background in ICT. **Due to time and resource constraints, this thesis will chiefly focus on the information dissemination aspect of the project and only touch briefly upon the application development course in Section 7.2.**

### 5.2.3 Expert panel

An external expert panel, Table 5.1, consisting primarily of doctoral-level researchers who possess extensive experience within FP7 PSIRP and related fields will be convened to oversee the design, operation, and conclusion of these courses.

**Tab. 5.1 - Expert panel membership**

| Member | Title(s) and Affiliation(s) | | |
|---|---|---|---|
| [1] D.Sc. A. Karila | Principal Scientist | - | Helsinki Institute for Information Technology |
| | Coordinator | - | FP7 PSIRP |
| D.Sc. S. Tarkoma | Professor | - | University of Helsinki |
| | Principal Scientist | - | FP7 PSIRP |
| D. Sc. J. Kangasharju | Professor | - | University of Helsinki |
| | Principal Scientist | - | FP7 PSIRP |
| D.Sc. D. Trossen | Senior Researcher | - | University of Cambridge |
| | Technical Manager | - | FP7 PSIRP |
| M.Sc. M. L. Markkula | Research Manager | - | Helsinki Institute for information Technology |
| | Project Manager | - | FP7 PSIRP |
| [1] M.Sc. W. Wong | Visiting Researcher | - | Oy LM Ericsson Ab |
| | Visiting Researcher | - | FP7 PSIRP |
| [1,2] M.Sc. M. Ain | Researcher | - | Helsinki Institute for information Technology |
| | Researcher | - | FP7 PSIRP |

[1] Indicates a staff member of T-110.6120

[2] "M.Sc." title granted from an accredited program prior to the completion of this thesis

Through the application of documented systematic forecasting and consensus techniques (e.g. the Delphi Method), it will be the responsibility of this panel to:

1) Design and validate the courses' objectives, structures, contents, operating methods, and assessment measures.
2) Oversee the instruction of the courses and monitor their progression.
3) Document and analyze the performance of the courses based on participant performance, participant feedback, and comments from overseeing staff.

4) Correlate said performance as an indicator of the suitability of this approach towards disseminating and exploiting a clean-slate internetworking architecture.

The validity of this approach has been verified by the Center of Excellence at Aalto University's Faculty of Information and Natural Sciences as nominated by the National Academy of Finland.

The determinations of the expert panel are conveyed directly within the analyses and results in subsequent sections. Additional notes from the expert panel are included following the main material where appropriate.

## 5.3 T-110.6120: Special course in pub-sub internetworking

The course T-110.6120 serves to address the need for effective dissemination of FP7 PSIRP material through an academic approach that embodies the structured learning environment and professor-student relationship of a university-level lecture course. In this, we hope to use academic courseworks to create an understanding of the project and its underpinnings.

In this section we outline T-110.6120's operational objectives and provide an overview of the design decisions which contributed to the objectives, structure, content selection, and operating measures of the course. Refer to Appendices A through E for the course syllabus and additional material.

### 5.3.1 Operational objectives

The need for PSIRP and the functionality it provides is deeply rooted in the idiosyncrasies of the current Internet. The Internet's developmental history, current problems, attempted solutions, operating conditions, usage demands etc. all serve as a basis that guides PSIRP's development. As such, it was agreed early in the planning stages of the course that we would need to give the participants a suitable amount of background in these areas in order to create the proper foundation to introduce the PSIRP project and facilitate the dissemination of its components. Our research into the aforementioned background led to the material of Section 2, which when followed by an overview of PSIRP, formed the operational objectives of the course:

1) Provide a brief history of the Internet that highlights how the events surrounding its inception and the demands of users at the time contributed to its foundational endpoint-centric send-receive design.

2) Highlight milestone modifications during the past 40 years of Internet development and characterize their evolutionary nature in response to impending operational limitations.
3) Demonstrate that the core Internet architecture has essentially become ossified as a result of various technical and socio-economic conditions.
4) Identify notable problems plaguing the current Internet as a result of modern usage demands, introduce notable evolutionary and revolutionary solution proposals, and through this demonstrate the plausible need for a revolutionary clean-slate redesign.
5) Provide a comprehensive overview of the FP7 PSIRP project which includes the foundations of its information-centric pub-sub communication paradigm, design tenets, architectural components, prototype implementations, future outlooks etc., and conclude with practical demonstrations and a panel discussion.

These objectives were formulated based on the academic experience of project staff and the expert panel for the primary purpose of creating a beneficial learning environment and achieving the strategic objectives of the dissemination and exploitation effort (Section 1.7.1).

### 5.3.2 Structure and content selection

Our operational objectives identify four core sequential areas of background which precede the PSIRP dissemination component of the course. The following sections briefly discuss the selected structure and content (i.e. lectures slides and reading) of these five course topics and indicate their weightings within the course.

#### *5.3.2.1 Internet inception*

**Weighting: ~5%**

The origins of the Internet and its underlying endpoint-centric send-receive paradigm ultimately prompted its developmental design stages. This long developmental history is directly attributed to the starting circumstances of the Internet, and it is therefore imperative that course participants have a solid grasp of this background before the Internet's developmental milestones, modern problems, and potential solutions are introduced. The content of this section includes primarily the material discussed in Section 2.1. [Han2006] is assigned as requisite reading due to its reputable in-depth coverage of relevant subject matter.

#### 5.3.2.2 Internet development

**Weighting: ~5%**

Subsequent to gaining an understanding of the Internet's inception, it is necessary to understand the major developmental stages of the Internet that took place in response to the interplay between the Internet's original design, evolutionary growth, and evolving usage demands. These architectural changes effectively sustained the architecture and enabled the ubiquitous global network that is observed today. This section of the course chiefly includes the major developmental milestones discussed in Section 2.2 and references to their IETF RFCs (Request for Comments).

#### 5.3.2.3 Internet ossification

**Weighting: ~5%**

The inception of the Internet's endpoint-centric send-receive foundations and its subsequent developmental changes led to its current state of ossification in response to evolving usage demands and various technical and socioeconomic trends. This portion of the course illustrates the nature of these developments (i.e. evolutionary), the circumstances that prompted their invention (i.e. critical operational limitations), and their implementation (i.e. last-minute), culminating in the rigidity of the Internet's core protocol stack. This information serves as a foundation by which the Internet's key problems, attempted evolutionary solutions, and proposed revolutionary solutions, can be introduced. [Han2006] is a reliable and reputable overview of this material.

#### 5.3.2.4 Internet problems and solutions

**Weighting: ~5%**

This portion of the course serves to exemplify the end results of the Internet's ossification and the apparent need for revolutionary progression. This constitutes the bulk of Sections 2.3, 2.4, and 2.5. The Internet's prominent problems and evolutionary solutions are relatively easy to characterize to the intended audience, all of whom have a reasonable awareness of these issues due to their background in accordance with the course prerequisites (Appendix A). On the other hand, we've seen that it is particularly difficult to educate the general public on the innovations of revolutionary Internet architectures. As such, due to time and resource limitations, we opted to choose a single modern architecture proposal to serve as an example of a revolutionary solution so as to avoid overwhelming participants in preparation for the dissemination of PSIRP material. [Jac2009] was selected as the sole reading assignment since the expert panel considers it to be one of the most creative and best-written information-centric architecture proposals. Furthermore, [Jac2009] is a relatively recent publication and it was produced by the Palo Alto Research Center (PARC) [PARC2010], a well-known and established source.

#### 5.3.2.5 PSIRP dissemination

**Weighting: ~80%**

We devoted a majority of the course lectures to the PSIRP dissemination effort, covering the state-of-the-art [PSI2008c], general aspects of the PSIRP architecture (e.g. information centrism, publish-subscribe communication etc.) [PSI2009a] [Tar2009], PSIRP architectural components [PSI2009a], prototype implementations [Jok2009] [Kjä2010], and other key focal areas such as security and mobility [PSI2009c]. **These subjects were selected by course staff and the expert panel according to their immediate relevance and the success of their corresponding weightings within PSIRP's mandatory dissemination deliverables to the European Commission.**

The following reading materials were prescribed in the indicated order:

1) [Tar2009]
2) [PSI2008c]
3) [PSI2009a]
4) [Jok2009]
5) [PSI2009c]

The final two lectures included a hands-on demonstration of the PSIRP Blackhawk prototype and a panel discussion involving key project staff. The purpose of these sessions was to provide a tangible demonstration of PSIRP technologies followed by a structured open discussion, allowing participants and experienced staff to interact, exchange ideas, discuss future work etc.

### 5.3.3 Participant requirements

Participants had to meet the following requirements in order to pass the course:

- **Mandatory lecture attendance:** The concept of mandatory attendance is alien to most students at the target level of the course. However, the course was designed to be lecture-intensive and stress active discussion and participation as part of PSIRP's course-driven dissemination ideology. We recognized that without a mandatory attendance policy some students might have opted to disregard lecture sessions and tend to the lecture materials on their own time (this type of behavior is common among university students). We sternly sought to avoid this possibility as students would likely not benefit from the course without being guided through the material by an experienced lecturer. [Cha2006] exemplifies the merits of this lecture-driven approach to learning:

*"Lectures are probably the best teaching method in many circumstances and for many students; especially for communicating conceptual knowledge, and where there is a significant knowledge gap* [as in PSIRP] *between lecturer and audience.*

*…*

*The formal structure of a lecture therefore artificially focuses more attention and generates authority for the lecturer to make their communications more memorable. Furthermore, to allow the potential for repeated interactions to allow trust to develop between lecturer and class, it is much more educationally-effective for lectures to be given as a course rather than as one-off interactions."*

- Courtesy of [Cha2006]

Extenuating circumstances and individual absence requests were handled on a case-by-case basis.

- **Submission of a weekly assignment:** Each participant was required to submit a weekly "learning diary" (Appendices B and C) which was graded "pass" or "fail" by the course staff based on the following factors:
  - Appropriateness and correctness of submitted material
  - Reasonable observance of length guidelines
  - Timeliness of the submission
  - Overall impressions on the quality of the submission and effort of the participant

  etc.

  The course staff and expert panel concluded that this type of weekly assignment would keep participants thinking about the week's material and enable a reasonable assessment of their comprehension and opinions. Participants were required to receive a grade of "pass" for all submissions in order to receive course credit.
- **Completion of a weekly survey:** Course participants were required to rate each lecture and its requisite reading material on a weekly basis (Appendix C) using a

Likert scale of 1 to 5 [Tro2005]. Students were explicitly informed that the nature of their responses would have absolutely no bearing on their successful completion of the course. The tools used to administer the surveys [TKK2010] prevented participants from observing the responses of their peers, although participant responses were not anonymous when viewed by course staff. The survey results were never discussed with participants in any capacity, aside from displaying an anonymous summary chart (Table 6.2) for informational purposes at the conclusion of the course. We do not expect any significant perturbation of our results on these grounds.

- **Completion of weekly reading assignments:** The course staff and expert panel compiled a selection of project deliverables and scientific papers as mandatory reading prior to each week's lectures (Appendix A). These papers were chosen based on the experience of the expert panel and the success of the readings within their corresponding distribution venues. For example, the PSIRP project deliverables have consistently received excellent reviews from the European Commission which justify their inclusion in the course material.
- **Completion of a final assignment:** A final "learning-portfolio" was designed by the course staff and expert panel in order to provide an assessment of how well the course participants assimilated the course material. The assignment was evaluated on a pass/fail basis. Details are provided in Appendix D.

### 5.3.4 Participant evaluation

As with most any new and complex innovation, clean-slate internetworking approaches such as PSIRP have been at somewhat of a disadvantage when it comes to educating the public and gaining worldwide acceptance. These downfalls are especially important in academia because they can seriously hinder course participants' performances. **The manner by which participant performance is assessed is one of the chief determinants of participant behavior, which inexorably influences the effectiveness of the PSIRP dissemination effort.**

It is well known that evaluation and assessment measures are a primary motivating factor for students. Student motivation is directly correlated to learning and dissemination effectiveness, and in this respect, designing a fair and motivating assessment scheme is a primary consideration in PSIRP's academic exploitation efforts.

Since evaluations are largely influential to their futures, students are often most tempted to focus on earning higher grades instead of learning material in a manner that suits them best. This illustrates the dichotomy which is characteristic of learning environments where performance is often highly subjective: a given student will almost always choose to approach a course in the manner which yields the best assessment results because their corresponding abilities will almost always be judged based on those results in the future, even if their method of approaching the course is usually not the one which gives the student an optimal learning experience.

We recognized these considerations when designing PSIRP's academic courses and sought to create an evaluation system which would ensure a positive learning experience for the participants and facilitate a high degree of dissemination. We devised the following key points as an agenda for student evaluations:

- Eliminate perceptions of inter-student competition to the greatest extent possible.
- Promote a low-stress environment that does not emphasize arbitrary performance measures.
- Enable individual students to focus on learning material in the manner which suits them best.
- Institute evaluation measures which guarantee that the course participants will be motivated to complete the tasks assigned by overseeing staff.

Some of these goals are in conflict to a certain extent. In the absence of enforced assessments and/or negative reinforcement, all but the most self-motivated students will choose to pursue a study route which requires the least work to achieve the most positively-measurable results (i.e. influential assessments and credits). On the other hand, strenuous evaluations are time consuming for staff and add to the stresses and biases which lead students to simply strive towards the best evaluations.

All students are differently-abled and motivated, leading to different perceptions of what constitutes a reasonable balance of time spent towards effective learning vs. time spent in an attempt to achieve the best evaluation. Note that these goals do overlap to a certain extent. The expert panel and course staff considered this observation at length in conjunction with the aforementioned agenda points and eventually came to the conclusion that overlapping and combining these student perceptions into a single "acceptable" margin would yield a grading scheme which would afford the most flexibility in appealing to all students. **Pass/fail** grading appears to do this by condensing students' performance categories into two extremes. Students are effectively free to go about completing course tasks as they please and learn in the manner which suits them best, so long as they achieve the minimal baseline performance set by the course staff in accordance with the expected capabilities of the entire student body. This greatly diminishes perceptions of direct inter-student competition. It seems reasonable to assume that this baseline “passing” threshold confers an acceptable amount of flexibility to students such that their chosen approaches to satisfying the course demands will be closer to their optimal learning methods.

Our assessment of a pass/fail marking scheme and its usefulness is supported by several studies, industry examples, and usages amongst notable academic institutions.

Pass/fail grading is commonly adopted in ICT learning through industry certification programs when achieving certification is the single most important performance measure. The details of underlying assessment measures and metrics are typically unimportant to the participant and future employers. Likewise in the PSIRP dissemination effort, the dissemination target itself is most important, and not the course participants’ unique learning styles and the assessments thereof. Moreover, a traditional numeric evaluation (or any sort of translation derived thereof) may also imply a degree of comparison between exams, participants, conceptual difficulty etc. which is completely subjective or even wholly inappropriate.

One example of a pass/fail grading strategy in ICT learning is the Microsoft Corporation’s abandonment of traditional scoring in its industry certification exams (e.g. MCP, MCSE etc.) nearly a decade ago in favor of a simple pass/fail scheme [Eck2002].

Many reputable universities (including Aalto University) have assessed and implemented pass/fail grading in study programs which largely share the complexity and innovativeness of future Internet research. It is hard to find a better example than medical programs, whose curricula are certainly one of the most demanding in higher education. A recent study [Blo2009] published in Academic Medicine has shown that switching to a pass/fail grading system during the first two years of medical school confers numerous benefits to students, including improved psychological well-being and satisfaction, with no significant reduction in performance in courses, clerkships, test scores, residency placement, or attendance. The following is a listing of several prominent North American medical schools and their corresponding student requirements where pass/fail marking is used exclusively [AAMC2010]:

- **McGill University, Faculty of Medicine:** Required basic sciences, basic science electives.
- **Johns Hopkins University School of Medicine:** Required basic sciences, basic science electives.
- **Mayo Medical School:** Required basic sciences, basic science electives.
- **Harvard Medical School:** Required basic sciences, basic science electives.

- **Stanford University School of Medicine:** Required basic sciences, basic science electives, required clinical clerkships, elective clinical clerkships.
- **University of Toronto, Faculty of Medicine:** Required basic sciences, basic science electives, required clinical clerkships, elective clinical clerkships.

Moreover, a 2006 study of first-year medical students at the Mayo Medical School in Rochester, Minnesota, comparing students graded using traditional and pass/fail systems, revealed that the latter students had less perceived stress and greater group cohesion than their counterparts who were graded on a typical 5-point scale [Roh2006].

The research of the expert panel coupled with these indicators provides strong evidence that pass/fail marking is reasonable for our purposes. However, there are some historical results which seem to indicate to the contrary:

*"Experience with the pass/fail system at the University of Alberta and a review of the literature has shown that (1) pass/fail does not seem to motivate the student to learn; (2) students do not use it as a vehicle to explore outside their major; and (3) students do fewer of the assigned readings and attend fewer classes in courses elected under pass-fail than they do with courses elected under the conventional grading system.."*

- Courtesy of [Ott1972]

*"College students voluntarily took all their courses or one course on a pass-fail basis. The mean grade point average (GPA) before conversion to pass-fail for freshman taking all their courses on a pass-fail basis was 1.67 (C-), which is significantly lower than the 2.26 (C+) for controls who wanted but were denied pass-fail grading. Even after returning to conventional grading the former pass-fail students continued to get significantly lower grades than controls. Juniors taking one course on a pass-fail basis received significantly lower grades before conversion, in their pass-fail course (mean 2.07) than did controls who wanted but were denied pass-fail grading (mean 2.4). There was no compensatory improvement in the grades received in non-pass-fail courses."*

- Courtesy of [Gol1971]

We can observe several mitigating factors in these results:

- Nearly 40 years have passed since these studies were undertaken and many contributing factors may have changed in that time. Universities, the industry, student values, social norms, the state of ICT, information availability etc. have clearly evolved to the point where these results may no longer be valid.

- These studies were not directed within future Internet research or ICT in general.
- Lecture attendance, as discussed in [Ott1972], is irrelevant due to our mandatory attendance policy (however, issues pertaining to prescribed readings are discussed in Section 6.4).
- The results in [Gol1971] merely conclude that students subjected to traditional grading schemes achieve higher marks in these respective grading schemes. The questionable validity of these marks, the corresponding assessment measures, and their significance with respect to student learning and stress remains unaddressed. Conformity to assessment measures is a likely culprit.
- The results in [Gol1971] are based on translating a traditional grading scheme to a pass/fail system by selecting an (arbitrary?) “pass” threshold (note the use of the word “conversion” in the quoted passage). The subjectivity of this approach is undeniable. There is no guarantee that the threshold is appropriately selected or even appropriate to the material being taught. The pass/fail students are still effectively subjected to traditional grading means and may even be aware of the threshold, enabling them to purposefully dedicate only the minimum required effort to achieve a passing grade (our experience with this problem is discussed in Section 6.4).
- The subjects in these studies were arguably less advanced than the graduate and postgraduate participants of our dissemination course, and thus it is reasonable to assume that their performance may be poorer.

etc.

There is obviously no conclusively perfect marking scheme that guarantees optimal dissemination and exploitation of advanced ICT research. Nevertheless, we believe that pass/fail grading constitutes a reasonable approach for our purposes. Our observations and experiences in this respect are detailed in Section 6.4.

# 6 Results

General participant statistics for T-110.6120 are shown in Table 6.1. The course achieved a student pass-rate just above 90% and lecture attendance was acceptable.

**Tab. 6.1 - T-110.6120 general statistics**

| **Total participants** | **21** |
|---|---|
| Passing | 19 (90.5%) |
| Failing | 2 (9.5%) |
| Final assignment failure | 1 |
| No final assignment submission | 1 |

| **Absences** | **8** |
|---|---|
| Single lecture | 7 |
| Multiple lectures | * 1 |

* One participant was excused from four lectures due to severe illness. The student recovered and submitted a supplementary assignment based on the contents of the missed lectures to compensate for their absence.

## 6.1 Participant performance

The overall performance of the course participants was assessed by evaluating the quality of their contributions to lecture discussions, weekly learning diary submissions, and final assignments, in this order of ascending weighted precedence. These assessments were performed primarily by the course staff and reviewed by the expert panel.

The weekly lecture discussions were unfortunately less lively than anticipated. The students were attentive to the lecturer but showed little motivation to contribute their own thoughts when asked to discuss a given topic. However, this is typically the case in most academic environments; all but the most interested students are usually apprehensive over engaging in lecture discussions because they either have little motivation to do so or they fear that revealing their opinions may result in negative reprisals. The experiences of the course staff and expert panel in academia support this conclusion, although this remains a definite point for improvement in future offerings of the course.

The learning diary submissions were very good overall. Students were generally quite adept at picking up the main points of each lecture and summarizing what they felt they had learned. We observed minor problems when students were asked to convey their opinions on the course material and whether there were any particular topics that they strongly agreed or disagreed with. Most students abstained or offered few insightful comments, indicating a lack of motivation to work beyond minimal assignment requirements. The students may not have taken the opportunity to think for themselves seriously. Individual analyses and commentary are of course highly encouraged and we are adamant about improving this aspect of the course in the future.

The participants also showed a good degree of resiliency and insightfulness in their final assignment submissions. They correctly identified the Internet's notable problems and evolutionary and revolutionary solutions, and exhibited an excellent understanding of the intricacies of endpoint-centric send-receive communications as well as PSIRP's information-centric pub-sub approach. We were also very pleased with the students' ability to explain the PSIRP architecture and its components. Participant submissions contained a good degree of depth and appropriate technical explanations covering the state-of-the-art and areas such as PSIRP identifiers, rendezvous and scoping, the Blackhawk prototype, zFilter forwarding etc.

**The expert panel has concluded that the participants' final assignments reflect a good degree of learning that is comparable to that which is observed in traditional successful courses. Based on the judgments of the course staff and expert panel, we can conclude to a reasonable degree of certainty that students profited from the course and gained a level of understanding that is at or above what is expected in order to gain ECTS credits from a traditional successful academic course in ICT.**

## 6.2 Participant feedback

The final tabulated results of the participant feedback surveys are shown in Table 6.2. Raw participant feedback is provided in Appendix E.

There appears to be no consistent linear trend in the lecture ratings, although the reading ratings appear to drop throughout the progression of the course. We attribute this to three aspects:

- The complexity of the readings.
- The lengthiness of the readings.
- Our unfortunate lack of clarity when communicating the sections of each reading which were especially pertinent to a given lecture (see Section 6.4).

It is worthwhile to note that the two selected non-PSIRP readings were most highly ranked by the course participants. While this suggests that their inclusion in the course is warranted, it also indicates that the PSIRP deliverables may need to be altered in order to be more suitable for this type of academic environment, despite their successful reception by the European Commission. This possibility is discussed in Section 6.4 in accordance with the three aspects listed above.

**Tab. 6.2 - T-110.6120 overall participant feedback**

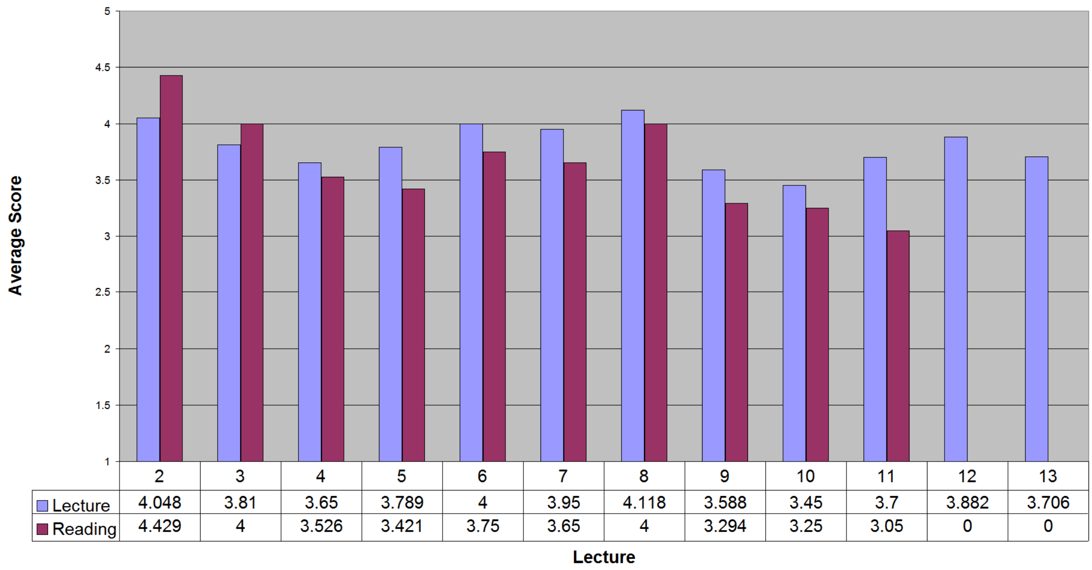


| | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Lecture | 4.048 | 3.81 | 3.65 | 3.789 | 4 | 3.95 | 4.118 | 3.588 | 3.45 | 3.7 | 3.882 | 3.706 |
| Reading | 4.429 | 4 | 3.526 | 3.421 | 3.75 | 3.65 | 4 | 3.294 | 3.25 | 3.05 | 0 | 0 |

### Average Lecture Score

3.81 / 5.00 = 76.2%

### Best-rated Lectures

1) 4.12/5.00 Lecture 8: zFilter-based forwarding: LIPSIN
2) 4.05/5.00 Lecture 2: Evolutionary and revolutionary approaches
3) 4.00/5.00 Lecture 6: Architectural components 1

### Average Reading Score

3.64 / 5.00 = 72.8%

**Best-rated Readings**

1) 4.43/5.00 Lecture 2: [Han2006]
3) 4.00/5.00 Lecture 3: [Jac2009] TIE
3) 4.00/5.00 Lecture 8: [PSI2009a] TIE

**Based on these figures, the course staff and expert panel have concluded to a reasonable degree of certainty that participant sentiments towards the course are subjectively positive overall. Based on their experiences in academia, the course staff and expert panel have also expressed the view that these results are comparable or superior to those commonly observed from student bodies of traditional successful ICT courses.** This reflects positively on the lecture materials, reading materials, lecturer performance, and the overall management of the course, and is especially encouraging in light of the fact that this was a pilot-trial. Moreover, we have gained a broad array of constructive criticisms and observations from the course (Section 6.4) which provide ample means to improve our approach in the future.

**The most important observation is that this data proves that course-based academic dissemination for a clean-slate internetworking architecture can feasibly achieve above-average participant feedback when compared to traditional successful ICT courses.**

## 6.3 Staff feedback

The course staff were asked to provide a brief account of their experiences with the course and an overall assessment of its success. Their comments follow:

*"I was quite satisfied with the number and quality of students that we could attract at such a short notice. The atmosphere in the classroom was attentive and open. While the discussion could have been more lively, the attendance was good and the students mentally present. I am convinced that most of the students got a lot out of this course and that it will help them orient to new clean-slate approaches in general – not just pub/sub and PSIRP."*

– D.Sc. A. Karila

Lecturer, T-110.6120

*"In my opinion the participants performed very well. There were many discussions during the lectures, showing that the students were interested in the topic and motivated to absorb the new ideas. They were really amazed with the new concepts, especially about the routing mechanism based on zFilters, and in the Blackhawk prototype implementation. Based on their questions during the lectures and final assignment, the students showed that they had a good understanding about the work done in PSIRP. Many of them also asked information about the PSIRP codecamp course [T-110.6100, Section 7.2], showing that they are interested in testing the new concepts in the prototype.*

*I believe the dissemination and exploitation plan through the university was a complete success. Many students understood the new concepts developed in the PSIRP project, especially the problems the current Internet faces and how these new concepts in PSIRP solve them. Another argument supporting the success of the dissemination plan is the interest of most students to continue with the second part of the course with the codecamp."*

- M.Sc. W. Wong

Lecturer, T-110.6120

**In general, the course staff's feedback is predominantly positive and they believe that the course has achieved its primary goal of disseminating PSIRP material.**

## 6.4 Observations and lessons learned

Overall, the course participants, staff, and expert panel were generally contented by the end results of the course. The participants demonstrated a good grasp of the most important material, subjective feedback was strong, and reactions amongst FP7 PSIRP personnel were generally positive.

In conjunction with comments from the course participants and staff, we have identified several key issues which are in need of further investigation and/or improvement in future undertakings of the course.

**Pass/fail marking** proved to be somewhat problematic, despite the fact that the premises discussed in Section 5.3.4 appear relatively sound. Although we believe that our pass/fail grading scheme served as a useful tool to ensure that participant motivations were properly directed (i.e. the students should want to learn, as opposed to gain the best possible mark), we noted some cases where students appeared to learn the expected passing threshold by observing the varying quality of their passing submissions over time. It was consequently easier for these students to attempt to minimize their time spent by devoting only the minimum amount of effort required to achieve a passing mark. Moreover, participants were also likely aware that the staff would be more reluctant to fail students toward the end of the course since:

- Students had already successfully completed a significant quantity of work.
- The students' prior submissions were at the passing level, making a single late-failure an outlier which could likely be negotiated or even discarded.
- Students were likely aware of the course's significance to FP7 PSIRP and the staff's goal of achieving a low failure rate.

etc.

We observed these issues to some extent during T-110.6120 and now believe that the application of an exclusively pass/fail marking system may not be the best option in this type of environment. To address this problem, future undertakings should potentially consider the simultaneous application of traditional and pass/fail grading schemes in order to promote their benefits (i.e. reduced participant stress, reduced perceptions of inter-participant competition, improved focus on individualized learning, motivation to complete assigned work etc.) and negate their disadvantages. One possible approach to accomplish this is to use a traditional grading scheme with pass-fail aggregation and minimum passing marks. That is, students are motivated to achieve a minimum acceptable grade for each of their submissions and strive to achieve a minimum overall passing mark in the course, which will then simply be listed as "pass" on their final transcript. Inter-student competition may be beneficial in this context since it allows the course staff to normalize individual student assessments according to the relative performance of the course body as a whole. In such an environment, assessment thresholds are less static and more dependent on the performance of an individual student in comparison to their peers.

We also noted significant difficulties **motivating students to participate in lecture discussions**, which were a central component of the course. One can only assume that this problem stems from a lack of meaningful participation incentives beyond the encouragements of the lecturer. The use of a pass/fail grading scheme adds to the problem as students were well aware of the subjectivity involved in gauging discussion participation and the corresponding low passing threshold which would be necessary to account for this within the marking scheme. Because of this, discussions were short and less frequent than desired. Future undertakings should consider applying techniques which better motivate participant discussions using positive and/or negative reinforcement methods. One possibility is to request that each participant prepare and shortly explain their position on a relevant course topic, possibly under direct staff evaluation or with a student-selected opponent. With this, we anticipate lengthier and more involved deliberations.

The concept of **mandatory attendance** is naturally alien to graduate and postgraduate students and we noted some minor enforcement problems. The course staff had no preset policy for dealing with student absences, having instead chosen to handle absences on a case-by-case basis. We expected that this approach would be relatively successful based on the advanced level of the participants: there were a total of seven excused absences throughout the duration of the course (i.e. seven different students who missed a single lecture), one absence of four consecutive lectures due to serious illness (the student was asked the complete a compensatory assignment), and no unexcused absences.

Due to the innovative nature of the course and its material, the stressed importance of group discussion, and the implementation of pass/fail grading, a mandatory attendance policy was chosen and deliberately announced prior to the beginning of the course in order to ensure that participants would follow the instruction of the course staff and not attempt to complete the required material exclusively on their own time.

Given the results in [Cha2006] and the experiences of the course staff and expert panel, we suggest continuing a mandatory attendance policy in future undertakings, although the enforcement policy should perhaps be more clearly defined.

We consider the course's reading materials as a central component of the project's dissemination effort. As such, we were somewhat distraught when we discovered that it was virtually impossible to **verify whether or not the students had properly completed their out-of-class reading tasks**. In theory, students could potentially take advantage of the fact that

1) their knowledge of the reading was only evaluated through the learning diaries and the final assignment, and
2) the lecture slides provided ample material to counteract the necessity of completing all of the requisite reading tasks.

It was unfortunately impossible to selectively change the reading policy to account for this problem once the course was underway. We suggest that future undertakings implement short informal assessments at the beginning of each lecture to ensure that students have read the assigned materials. Once again, these assessments should be constructed so as to ensure that students are not stressed to gain the highest possible evaluation. Instead, they need to be designed to account for the wide breadth of abilities and potential focuses exhibited by a broad array of students when reading a technical document, and be easily passable by any participant who genuinely takes an interest in the tested material. One interesting possibility would be to integrate this intention with the aforementioned approach to improve course discussions, thus simultaneously addressing the need for lecture participation and completion of requisite reading.

In retrospect, the **credit allotment** for the course (4cr) may have been too generous, especially if one takes into account the meager lecture discussions and our inability to verify completion of the assigned readings. Aalto University abides by the European Credit Transfer System (ECTS) whereby 1cr is equivalent to ~27 working hours. The course's per-credit breakdown follows:

- **Lectures (1cr):** 13 total * ~2 hours each = 26 hours
- **Reading (1cr):** 10 total * ~3 hours each = 30 hours
- **Learning Diaries (1cr):** 6 total * ~4 hours each = 24 hours
- **Contingency (1cr)** (e.g. surveys, studies, final assignment, optional reading etc.)

There is a margin of leniency in this assessment as this was a pilot-trial and it was difficult to anticipate the effective workload that would be demanded from the participants. We expect that our analyses and the student feedbacks on the lectures and readings will enable us to better gauge requirements for future undertakings, leading to a more thorough basis by which credits can be assigned.

As is the case with most academic pilot-trials, we could have benefitted from **additional planning time and resource availability** prior to and during the course. Issues such as online-learning environment preparation and maintenance, administrative access restrictions, official grade submissions, lecture material preparation and coordination, guest speaking arrangements etc. were more time consuming that expected. This resulted in sometimes problematic coordination among students, staff, and guest speakers, leading to occasional repetitions among lectures and somewhat broad scopes in the reading materials. We anticipate that these issues will improve considerably in the future as we gain further experience offering this type of course and the administration of Aalto University acclimates itself to the unique operating methods inherent of research-dissemination courses. Most of these administrative overheads will likely significantly decrease if the course becomes a standard offering within Aalto University's Faculty of Information and Natural Sciences, which is our long-term goal.

Another side effect resulting from reduced planning time and resource availability is the **broadness of the reading** selections assigned to students. Certain materials were upwards of 50 pages in length (project deliverables) and it was often difficult to direct students to focus on certain specific parts of the material. Given more time, course staff should opt to specifically delineate the portions of lengthy PSIRP deliverables that students should read prior to each lecture. We suspect that this would lead to an improvement amongst student ratings of these readings.

Lastly, we are of the opinion that **participant submissions should have more concise and enforced length guidelines**. Some student submissions were significantly longer or shorter than expected and/or not written in the manner requested by the staff. This wasn't a serious problem and no disciplinary actions were taken. Nevertheless, we recommend investigating the correlation between the expected writing style, length recommendations, and pertinent material of the learning diaries so more effective requirements can be instituted along with a more strenuous enforcement policy.

# 7 Related and future work

This section discusses considerations for future undertakings of T-110.6120 and provides an initial overview of the T-110.6100 PSIRP application development codecamp and its preliminary results. Recommendations for general related work follow.

## 7.1 T-110.6120: Special course in pub-sub internetworking

The course staff and expert panel have identified several components related to PSIRP's academic dissemination efforts which specifically warrant further investigation.

Most importantly, our approach could benefit from more conclusive research by specialists in fields such as marketing, information dissemination, human learning, psychology etc. so as to provide more thorough requirements and recommendations to improve upon our pilot-trial. It would also be helpful to better determine the effects of subjective variables (e.g. inclusion of selected content, effectiveness of the instructors, suitability of the learning environment, staff and participant feedback tactics etc.) on the dissemination component of the course.

In the context of future Internet research, it would be beneficial to carry out and compare similar experiments with other revolutionary architecture projects and related ICT research initiatives throughout Europe and other parts of the world. The European Commission's central role coordinating the European Framework Programs is an ideal starting point in this regard.

Resources should also be devoted to investigate the possibility of developing standardized academic dissemination and exploitation tools based on our research for future EU Framework Program projects. A dedicated engagement and dissemination framework would help offload these demanding processes from project consortiums and free resources for additional development work. However, we must note that a venture of this sort would likely require an extensive research and approval process under the European Commission. As we have seen in [AAU2009], technology investments such as those instituted by FP7 are maximized by directly integrating exploitation and dissemination functions into university technology environments. The major concern is that standardized dissemination and exploitation processes would lead to an operational gap between these activities and dedicated project development efforts. Steps must be taken to ensure that these processes are tightly integrated so that only work overheads, and not cooperation, are reduced.

## 7.2 T-110.6100: Special course in pub-sub application development

T-110.6100 is an attempt at disseminating PSIRP material through active hands-on application development using PSIRP's Blackhawk prototype. The course is run as a codecamp, emphasizing an intense learning environment and dedicated development over a period of 2 – 3 weeks. This type of course will allow us to determine the user-friendliness of the PSIRP paradigm and gauge its potential for success in the open development community. Another primary goal is to gain preliminary performance measures pertaining to the paradigm's ability to handle existing communication demands and developer methods.

In this section we briefly cover the operational objectives of the course and the preliminary results we've obtained thus far. For the sake of brevity, full course details are relegated to the syllabus in Appendix F.

### 7.2.1 Operational objectives

The course staff and expert panel have devised the following operational objectives:

1) Provide one or more short introductory sessions outlining the nature of PSIRP's information-centric publish-subscribe networking approach and the functionality of the Blackhawk prototype and available APIs.

2) Assign participants a series of development projects designed to give a comprehensive view of the capabilities of information-centric pub-sub internetworking.

3) Arrange a creative open environment which gives course participants the freedom to experiment with the Blackhawk prototype and employ the API to develop unique applications and services according to their own ideas.

4) Provide supportive assessment meetings in which participants freely demonstrate, analyze, and constructively evaluate their solutions and those of their peers.

As with T-110.6120, we intend to focus on creating a low-stress open development environment that encourages participants to take an active interest in the course and avoid working towards a static staff-defined performance baseline. The observations of

Section 6.4 will be particularly important in this capacity. Most importantly, participants who present ambitious application and service designs will be allowed to continue development beyond the end of the codecamp and receive additional ECTS credits for their efforts. With this, we hope to gauge the suitability of the prototype towards existing internetworking needs and promote its progression through user creativity and innovation.

### 7.2.2 Preliminary results

As the codecamp is still in progress, we have yet to obtain comprehensive results and feedback. This section includes an overview of the first of two development projects assigned to students and the results of automated vulnerability analyses performed on the associated student code submissions.

The first development project (A1) consists of a web service "mashup" whereby students must implement a simple client-server communication model over pub-sub using the Blackhawk prototype's Python API. A publisher periodically fetches a news feed from a selected web page and publishes the content for subscribers. New versions denote updated content. Two subscriber implementations are required. One subscriber is only capable of subscribing to the publication once the publication has been made (i.e. synchronous polling), and the other must be able to subscribe prior to the act of publication and receive the data immediately when the publication takes place (i.e. asynchronous interrupt). The mashup concept is illustrated in Figure 7.1.

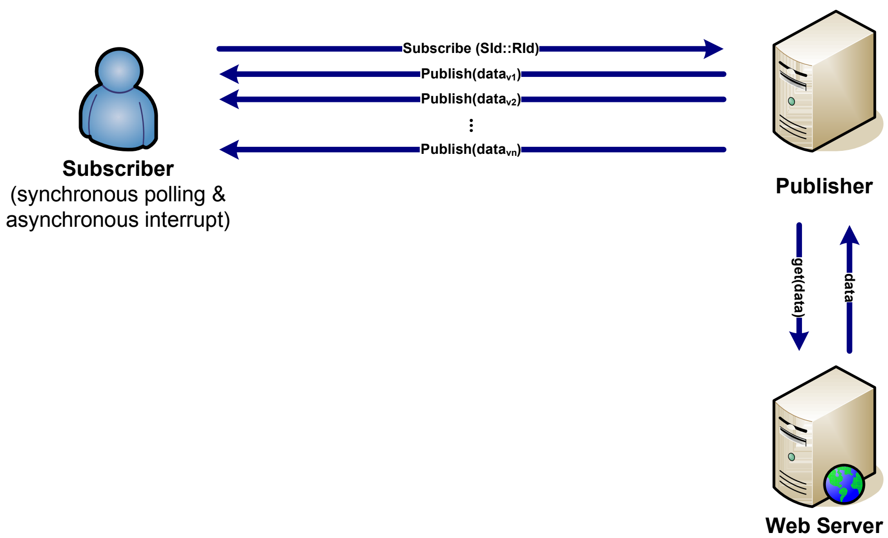


Fig. 7.1 - Pub-sub web service mashup

The participants' A1 submissions were evaluated using the Rough Auditing Tool for Security (RATS) version 2.3 [RATS2010]. The standard vulnerability databases were used and the warning level for threat detection was set to 3 (the most strict, on a scale of 1 to 3). Function calls accepting external input were also included in the analysis. These options correspond to the following execution command:

```
C:\>rats <filename> –w 3 –i
```

Results follow:

| **Participant** | **Implementation** | **Vulnerabilities** |
|---|---|---|
| 1 | Publisher | Possible input overflow (one occurrence) |
| | Subscriber | None found |
| 2 | Publisher | None found |
| | Subscriber | None found |

Although automated tools can typically only provide a very rudimentary and limited code analysis, we are generally pleased with these basic results. The student implementations show virtually no vulnerabilities and the potential security problems that have been discovered are attributable to coding style, not the prototype or communication paradigm, and can thus be easily fixed.

Initial participant comments are also positive and students have been generally well-able to understanding information-centric pub-sub communications and adopt PSIRP principles in application development.

With these results, we have good reason to believe that Blackhawk and its APIs will show promise in the hands of developers and sustain PSIRP's dissemination and exploitation efforts. Efforts are underway to institute the PSIRP-specific offering of T-110.6100, in conjunction with T-110.6120, as a standardized course within Aalto University's Faculty of Information and Natural Sciences.

## 7.3 The future of the Internet

The ossification factors and stagnancy feedback loop illustrated in Figure 2.1 are at the core of the Internet's problems must be addressed directly. We need to stress the criticality of impending operational limitations in the Internet so that people are motivated to act now, and not later. The key questions are:

- How can one change any of the states of the cycle in such a manner that the loop is broken?
- What externalities are necessary to elicit the aforementioned change(s)?

It is clear that dissemination and exploitation are of key importance here through the education of key stakeholders, increased perceptions of urgency and motivation by end users etc. In this, we need to better understand the interplay between network openness and the need for regulation to reinforce standardization and ubiquity. Ongoing research into dissemination and exploitation processes, their motivators, optimizations, and relationship to addressing the commonalities listed in the introduction of Section 2, is of the utmost importance, as is accounting for the mismatch between the length of future Internet projects and the timeframes of their ambitions.

# 8 Conclusions

With respect to the three strategic objectives listed in Section 1.7.1, we have reached the following conclusions:

| **OBJECTIVE #1: Compile background information on the inception and pertinent developmental history of the Internet and compose a brief literary review of prominent evolutionary and revolutionary solution proposals which serves as a useful reference for the reader.** |
|---|
| We have assembled a brief but informative overview of this information which serves as the requisite background for this thesis and the PSIRP course T-110.6120. The information included in Section 2 was arguably instrumental as a foundation for our dissemination effort and without a doubt contributed immensely to the success of this thesis and T-110.6120. |

| **OBJECTIVE #2: Justify the importance of dissemination and exploitation processes to future Internet research, and design, execute, and report on the results of an academic dissemination and exploitation pilot-trial for the FP7 PSIRP project in order to present evidence supporting (or refuting) the conclusion that this approach constitutes a promising route by which to disseminate and exploit a clean-slate internetworking architecture.** |
|---|
| Through extensive networking with project partners and research of related literature, we have compiled a variety of information which conveys the importance of dissemination and exploitation processes to future Internet research, the FP7 PSIRP project in particular, and the importance of academics in this capacity.<br><br>With regard to the PSIRP information dissemination course T-110.6120 offered at Aalto University during period III, spring 2010, we have produced the following key results and associated conclusions:<br><br>1) **Positive participant performance based on submissions evaluated by the course staff and expert panel.** The members of the course staff and expert panel are of the opinion that the course participants' overall performance based on their final assignments, weekly learning diary submissions, and lecture discussions, is comparable to that of successful productive participants in traditional academic courses in ICT. The panel and the Aalto University |

administration have thus concluded that the participants who have successfully met the course requirements have gained a sufficient knowledge of the course material to receive official ECTS credit.

2) **Positive participant feedback as evaluated by the course staff and expert panel.** As per the results in Section 6.2, the expert panel has concluded that the feedback collection methods and results are acceptably unbiased and comparatively above-average in the context of traditional academic courses in ICT. Furthermore, the fact that these results were obtained from a pilot-trial gives a good indication that future improvements may potentiate even better results.

3) **Positive final evaluations by the course staff and expert panel.** The course staff have offered optimistic statements regarding their experiences with the course. The expert panel has also expressed their optimism with regard to these results.

In light of these results, the expert panel, the PSIRP consortium, and the author have chosen to report the successful results of the course to the European Commission and will continue to explore the potential for similar undertakings in the future. **Most importantly, these results prove that an academic dissemination course for a clean-slate internetworking architecture can feasibly produce participant performance and feedback figures that are comparable or superior to those of a traditional successful academic course in ICT.**

**OBJECTIVE #3: Provide an account of notable events which took place during the course T-110.6120, document our lessons learned, and provide recommendations for similar projects in the future.**

This information (Section 6.4) is central this thesis and contains a broad array of observations and constructive criticisms that will serve as a useful guide when designing future offerings of T-110.6120 and T-110.6100.

We are thus confident that our academic dissemination effort has been successful. We also feel confident in recommending academic course-based means for disseminating and exploiting clean-slate internetworking architectures. Academic courses provide an important "first-level" to educate the general public about the problems of the Internet and their potential solutions, and our results serve to show that academic courses are a feasible venue by which to disseminate a clean-slate internetworking architecture.

# Appendix A: T-110.6120 syllabus

## COURSE DESCRIPTION

The Internet has evolved to be dominated by content distribution and retrieval. People increasingly want to access information – not hosts. However, the Internet is still based on naming hosts and addressing their network interfaces. We want to be able to specify what we wish to receive instead of where it shall be retrieved from. This leads into information-centric networking where content is named, routed, and cached. Another major problem with the current Internet is that it is working on the terms of the sender, which leads to unsolicited traffic, including SPAM and denial-of-service (DoS) attacks. We need to restore the balance empowering the consumers of information.

The publish-subscribe (pub-sub) paradigm is a proposed solution to the needs described above. Publications are named and you only receive the publications that you have subscribed to. Since January 2008, TKK-HIIT is coordinating an EU FP7 project PSIRP (www.psirp.org), where a new Internet architecture has been designed, implemented, and validated based entirely on the publish-subscribe paradigm. Rendezvous IDs identify publications within scopes specified by Scope IDs. There is a working prototype of the system, called Blackhawk, with an API for Python. Experiences gained in the project demonstrate that an Internet architecture can be built on the publish-subscribe paradigm and that certain applications, such as BitTorrent, become almost trivial to implement on it.

Many leading Internet researchers believe that the shift towards information-centric networking and publish-subscribe-type approach is inevitable. This new paradigm will require new skills from developers of applications and services. The purpose of this course is to give the students an introduction to pub-sub and the foundations of PSIRP.

The course is lectured twice a week during the first half of the spring semester (teaching period III) and it is primarily targeted to graduate and postgraduate students.

The course is lectured by docent Arto Karila and other researchers of the PSIRP project and assisted by Mark Ain (M.Sc.).

## PREREQUISITES

There are no mandatory prerequisites for the course. However, the course is targeted to senior and graduate students, and its successful passing requires an understanding of internetworking and its concepts. T-110.4100 Computer Networks or something similar is recommended.

**TENTATIVE SCHEDULE**

The course is lectured in lecture hall T2 of the T building, Konemiehentie 2, Otaniemi. The regular lecture times are on Monday 14 – 16 and Wednesday 12 – 14. Please note that the first lecture is on Thursday 21.1. 14 – 16.

| # | Date | Lecturer | Topic | Requisite reading |
|---|---|---|---|---|
| 1. | Thu 21.1. | AK | Practical arrangements<br>Why the Internet only just works – problems with the current Internet | 1 |
| 2. | Mon 25.1. | AK | The evolutionary approach e.g. CCN<br>The clean-slate approach; new paradigms of internetworking | 2 |
| 3. | Wed 27.1. | AK | State-of-the-Art | 3 |
| 4. | Mon 1.2. | AK | Introduction to the publish-subscribe paradigm and its central concepts | 4 |
| 5. | Wed 3.2. | AK | Overview of the PSIRP architecture | 5 |
| 6. | Mon 8.2. | WW | Architectural components: Identifiers, Algorithmic IDs, Node-internal Architecture, Helper Functions, Rendezvous | 5 |
| 7. | Wed 10.2. | WW | Architectural components continued: Topology Management and Formation, Forwarding, Network Attachment | 5 |
| 8. | Mon 15.2. | PJ, WW | zFilter-based Forwarding<br>LIPSIN | 6 |
| 9. | Wed 17.2. | DL, KV, WW | Security architectures | 7 |
| 10. | Mon 22.2. | ST, WW | Inter-domain Topology Formation | 7 |
| 11. | Wed 24.2. | WW | Mobility | 7 |
| 12. | Mon 1.3. | JK, PJ, WW | The Blackhawk prototype.<br>Demonstration | - |
| 13. | Wed 3.3. | AK, PN, WW | Conclusions, future directions, panel discussion | - |

AK - Arto Karila

DL - Dmitrij Lagutin

JK - Jimmy Kjällman

KV - Kari Visala

PJ - Petri Jokela

PN - Pekka Nikander

WW - Walter Wong

ST - Sasu Tarkoma

## READING MATERIAL

All reading material and lecture slides are openly available through the course homepage on Noppa.

1) M. Handley: *Why the Internet only just works*
2) V. Jacobson et al.: *Networking Named Content*
3) PSIRP D2.1: *State-of-the-Art Report and Technical Requirements*
4) PSIRP D2.2: *Conceptual Architecture of PSIRP Including Sub-component Descriptions*
5) PSIRP D2.3: *Architecture Definition, Component Descriptions, and Requirements*
6) P. Jokela et al.: *LIPSIN: Line Speed Publish/Subscribe Inter-Networking*
7) PSIRP D2.4: *Update on the Architecture and Report on Security Analysis*

## REQUIREMENTS

The course will graded pass/fail.

Passing the course requires:

- Active participation in the lectures (mandatory attendance)
- Completion of weekly learning diaries
- Completion of weekly surveys
- Completion of weekly readings
- Filling in a questionnaire at the beginning and at the end of the course

The grade is determined by the number and quality of the learning diaries and questionnaires submitted.

**We reserve the right to modify these requirements within reason in response to unforeseen circumstances throughout the progression of the course.**

# Appendix B: T-110.6120 learning diary instructions

**INSTRUCTIONS**

| | |
|---|---|
| ANSWERS: | Answer the questions in the provided template based on this week's lecture and reading material. Provide your answers in **bullet-point** form and include as much summarizing information along with each point as you feel is necessary to justify your answer. |
| LENGTH: | The approximate recommended length in words of each of your responses is shown at the end of the question e.g. (300+ words). Bear in mind that this is only a guideline. If you can express your ideas concisely using fewer words, or if you genuinely feel that you can't provide the required amount of material, you are welcome to state your reasoning and move on, but **be prepared to defend your position**. You should also try to avoid overly lengthy responses. On average, your submission should be ~1000 words. |
| FILENAME: | Save your submission as a PDF and name the file *lastname_week#.pdf* in lowercase letters. Write the week number using 2 digits so we can sort the filenames properly. For example, if I were submitting my diary for week 2, I would name it *ain_02*. NOTE: the week number refers to the week of the course, NOT the week of the year! If you are confused, follow the numbering used by the parent week folder in Optima. |
| DEADLINE: | The submissions are due by **11:59pm on Friday of the following week** (i.e. the deadline for your learning diary for week 2 is on Friday of week 3). |
| SUBMISSION: | Upload your submission to the "Learning Diaries" folder within the proper week folder in Optima (https://optima.tkk.fi/). |
| MISC: | **This is your chance to express your thoughts about what we did this week. Tell us what you think! Remember, there are no right or wrong answers; your submission is graded PASS or FAIL based on how much effort you put in.** If for some reason you agreed ahead of time that you could not attend a lecture this week, make a note of it where necessary. |

# Appendix C: T-110.6120 learning diary and weekly survey templates

### LEARNING DIARY QUESTIONS

Answer the following questions based on the instructions provided.

1) In **your** opinion, what were the main points of this week's lectures? (100+ words)
2) What did **you** learn? Elaborate on your answers from question 1 and feel free to introduce additional material. (300+ words)
3) What did **you** find most interesting? (300+ words)
4) Are there any arguments from the lectures or reading material that **you** particularly agree or disagree with? Explain. (Variable)
5) Rate the quality of the lecture slides from 1 (very poor) to 5 (excellent) and justify your response; do this individually for both lectures. What did **you** like? What did **you** dislike? What would **you** change? (100+ words)
6) Rate the quality of the reading material from 1 (very poor) to 5 (excellent) and justify your response; do this individually for both lectures. What did **you** like? What did **you** dislike? What would **you** change? (100+ words)
7) OPTIONAL: Do you have any other comments or questions? (Variable)

### SURVEY FORMAT

The course participants were asked to complete a short survey following each week's lectures to rate the quality of the lectures (i.e. topics discussed, slides, audience engagement etc.) and requisite reading material. Participants were instructed that their responses would have absolutely no bearing on their grades in the course.

**1) Lecture *lecture_#* (*lecture_date*) – rate the quality of the lecture:**

1 (Poor) 2 (Passable) 3 (Acceptable) 4 (Good) 5 (Excellent)

**2) Lecture *lecture_#* (*lecture_date*) – rate the quality of the reading material which was due for this lecture:**

1 (Poor) 2 (Passable) 3 (Acceptable) 4 (Good) 5 (Excellent)

# Appendix D: T-110.6120 final assignment

Course participants were required to attend a scheduled final assignment session approximately one week following the final lecture. The students were provided with the instructions detailed below and forbidden access to any aids. The allotted time was approximately 2 hours, although an extra 15 – 30 minutes of writing time was granted on an as-needed basis.

NOTE: Students were instructed that the percentage figures following each question were NOT grade weightings but rather an indicator of the average amount of time and material that should be devoted to each question.

**INSTRUCTIONS**

---

- **Write an essay (you MAY use bullet-point form) which covers ALL of the following topics:**
    1) **Is there anything wrong with the current Internet? What (if anything) are its current problems? Why do they occur? Give a quick overview of some solutions. (20%)**
    2) **What is information centric-networking? What is publish-subscribe networking? Provide a *short* example of both. (10%)**
    3) **Give a quick overview of LIPSIN. (20%)**
    4) **What is PSIRP (motivation, architecture, prototype etc.)? (50%)**
- Max ~3 pages. **You may NOT use any aids!**
- Save your submission to PDF as "*lastname_final*" and upload it to Optima **before** you leave. You may leave before 2pm if you finish early.
- Your submission is graded PASS/FAIL.
- **REMEMBER: We are NOT concerned with minute details or overwhelming technical accuracy.** We are trying to gauge how well you have absorbed the course material and how well you can support your arguments.

# Appendix E: T-110.6120 raw participant feedback

The tables show the overall results of the course survey listed in Appendix C.

| | Lecture Ratings | | | | | | |
|---|---|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 | Avg. | Responses |
| Lecture 1 | N/A | N/A | N/A | N/A | N/A | N/A | N/A |
| Lecture 2 | 0 | 0 | 4 | 12 | 5 | 4.05 | 21 |
| Lecture 3 | 0 | 1 | 4 | 14 | 2 | 3.81 | 21 |
| Lecture 4 | 0 | 2 | 4 | 13 | 1 | 3.65 | 20 |
| Lecture 5 | 0 | 1 | 4 | 12 | 2 | 3.79 | 19 |
| Lecture 6 | 0 | 0 | 4 | 12 | 4 | 4.00 | 20 |
| Lecture 7 | 0 | 0 | 4 | 13 | 3 | 3.95 | 20 |
| Lecture 8 | 0 | 0 | 3 | 9 | 5 | 4.12 | 17 |
| Lecture 9 | 0 | 0 | 8 | 8 | 1 | 3.59 | 17 |
| Lecture 10 | 0 | 3 | 6 | 10 | 1 | 3.45 | 20 |
| Lecture 11 | 0 | 1 | 6 | 11 | 2 | 3.70 | 20 |
| Lecture 12 | 0 | 1 | 4 | 8 | 4 | 3.88 | 17 |
| Lecture 13 | 0 | 2 | 4 | 8 | 3 | 3.71 | 17 |
| **TOTAL** | **0** | **11** | **55** | **130** | **33** | **3.81** | **229** |

| | Reading Ratings | | | | | | |
|---|---|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 | Avg. | Responses |
| Lecture 1 | N/A | N/A | N/A | N/A | N/A | N/A | N/A |
| Lecture 2 | 0 | 0 | 0 | 12 | 9 | 4.43 | 21 |
| Lecture 3 | 1 | 1 | 1 | 12 | 6 | 4.00 | 21 |
| Lecture 4 | 0 | 1 | 7 | 11 | 0 | 3.53 | 19 |
| Lecture 5 | 0 | 1 | 10 | 7 | 1 | 3.42 | 19 |
| Lecture 6 | 0 | 0 | 6 | 13 | 1 | 3.75 | 20 |
| Lecture 7 | 0 | 1 | 6 | 12 | 1 | 3.65 | 20 |
| Lecture 8 | 0 | 0 | 4 | 9 | 4 | 4.00 | 17 |
| Lecture 9 | 0 | 2 | 9 | 5 | 1 | 3.29 | 17 |
| Lecture 10 | 1 | 3 | 6 | 10 | 0 | 3.25 | 20 |
| Lecture 11 | 2 | 3 | 7 | 8 | 0 | 3.05 | 20 |
| Lecture 12 | N/A | N/A | N/A | N/A | N/A | N/A | 0 |
| Lecture 13 | N/A | N/A | N/A | N/A | N/A | N/A | 0 |
| **TOTAL** | **4** | **12** | **56** | **99** | **23** | **3.64** | **194** |

# Appendix F: T-110.6100 syllabus

## COURSE DESCRIPTION

The Internet has evolved to be dominated by content distribution and retrieval. People increasingly want to access information, not hosts, and the nature of the content is important, not its physical location. However, the Internet is still based on naming hosts and addressing their network interfaces. We want to be able to specify what we wish to receive instead of where it shall be retrieved from. This leads to **information-centric** networking where **content** is named, routed, and cached. Another major problem with the current internet is that it empowers the sender of information, which leads to unsolicited traffic, including spam and denial-of-service attacks. We need to restore the balance empowering the consumers of information.

The **publish-subscribe (pub-sub) paradigm** is a proposed solution to the needs described above. Publications are named and you only receive the publications that you have subscribed to. Since January 2008, TKK-HIIT is coordinating the EU FP7 Publish-subscribe Internet Routing Paradigm (PSIRP) project (www.psirp.org) where a new internet architecture has been designed, implemented, and validated based entirely on the publish-subscribe paradigm. A series of identifiers are used to work with applications, identify publications, match publisher and subscriber interests, scope information, form unicast/anycast/multicast distribution trees etc.

There is a working prototype of the system, called **Blackhawk**, with an API for **C** and **Python**. Experiences gained in the project demonstrate that an Internet architecture can be built on the publish-subscribe paradigm and that certain applications, such as BitTorrent, become almost trivial to implement on it.

Many leading Internet researchers believe that the shift towards information-centric networking and publish-subscribe-type approach is inevitable. This new paradigm will require new skills from developers of applications and services. During teaching period III (18 January – 5 March, 2010) there was a lecture course on publish-subscribe internetworking, to which this hands-on codecamp forms a natural continuation. The purpose of this codecamp is to give students an introduction to the information-centric pub-sub communication paradigm by giving them an opportunity to develop innovative applications, services, and interfaces using the PSIRP Blackhawk prototype. The codecamp includes a few guidance lectures in the first week, forming an introduction to information-centric pub-sub networking, PSIRP, the Blackhawk prototype, and the Python API.

The bulk of the codecamp consists of two programming assignments made in teams of two students. It is also possible to work individually. The teams will develop a few relatively simple programs and interfaces, primarily in the Python programming language, designed to demonstrate some fundamental pub-sub concepts on top of the API of the Blackhawk prototype. Students are expected to provide a short (less than 2 pages) report on the progress/functionality of their work, problems encountered, proposed improvements to the API etc.

The last assignment is a small application and user-interfacing project on a topic chosen by the team and approved by the codecamp staff. The grade of this part is determined based on the originality of the selected topic, the programming difficulty, and an analysis of how well the pub-sub paradigm and its current implementation support this kind of software project and how the API could be modified to better suit the intended purpose. **An ambitious and highly original final project which is submitted in an unstable or even inoperative state will earn a higher grade than a stable but simplistic application with little evidence of effort.**

**NOTE: With staff approval, students who have chosen an ambitious design for A2 may continue working beyond the stated deadline. The top final assignment submission(s) may be rewarded with additional credits.**

We will have demonstration and assistance sessions beginning after the first two weeks of the codecamp. You may contact the responsible professor by e-mail at any time for technical assistance. At the end of the codecamp, there is a seminar where the teams present their individual projects and discuss their main results.

## STAFF INFORMATION

The course is lectured by Walter Wong (M.Sc.) wong@hiit.fi and possibly other researchers in the PSIRP project and assisted by Mark Ain (M.Sc.) mark.ain@hiit.fi.

## PREREQUISITES

There are no mandatory prerequisites for the course. The course is targeted to senior graduate and postgraduate students and its successful passing requires an understanding of internetworking and its concepts as well as good programming skills. Participation in the preceding lecture course **T-110.6120: Special Course in Data Communications Software – Publish-Subscribe Internetworking** is highly recommended.

## PROGRAMMING ENVIRONMENT

The programming assignments are carried out on the Blackhawk prototype, which currently runs on the FreeBSD operating system (64-bit version). This is the supported programming environment for the codecamp.

Each team will have to install their own virtual machine using Sun Microsystems' VirtualBox Emulator. A disk image of the entire FreeBSD system with Blackhawk incorporated in the kernel is available via www.psirp.org.

## ASSIGNMENTS

To pass the codecamp, students will have to successfully demonstrate the following two projects:

- A1: Implementation of client/server web service mashup over pub/sub
- A2: Student-proposed technology demonstrator (subject to staff approval; more details to come)

## REQUIREMENTS & GRADING

The codecamp implements a traditional grading system (scale 0 – 5) with **pass/fail aggregation** and **minimum required grades**. You must meet the following requirements in order to pass the codecamp:

- You must complete and pass ALL assignments with a **grade of 2 or higher**; this includes submitting a (brief) analysis report with each assignment
- You must actively participate in the seminar at the end of the course and receive a participation **grade of 2 or higher**
- You must fill in a short a feedback survey at the conclusion of the course
- **Your OVERALL grade average must be 3 or higher**

If you accomplish the above, you will receive a PASS mark (it will be entered as "pass" in Oodi, not as a numerical grade)

**NOTE: We reserve the right to add or remove requirements, within reason, throughout the duration of the course should unexpected circumstances arise.**

## LOCATION

Codecamp sessions will be held in Maarintalo in Otaniemi.

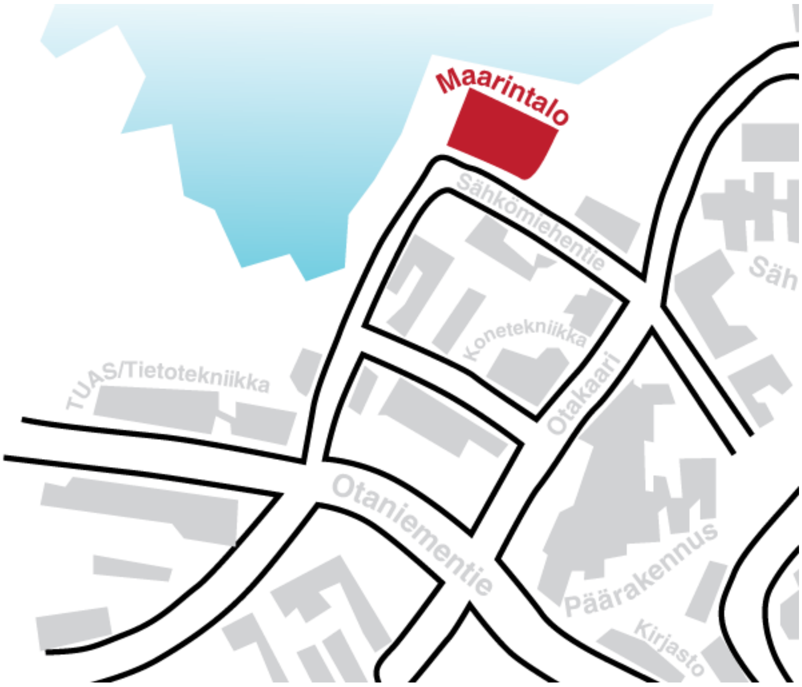


## SCHEDULE

| # | Date/Time | Room | Topics | Due |
|---|---|---|---|---|
| 1. | Wed 28.4<br>12-16 | Maari-M | • PSIRP project presentation<br>• Crash course in Python (if required)<br>• Hands-on Blackhawk installation and configuration<br>• Blackhawk API presentation and examples<br>• Hands-on Blackhawk programming<br>• A1 – Implementation of client/server paradigm over publish/subscribe<br>• Discussion and implementation of A1 | |
| 2. | Mon 3.5<br>12-16 | Maari-M | • A2 – Topics discussion<br>• A1 – Help session | |
| 3. | Wed 5.5 | Maari-M | • A1 – Demo session<br>• A2 – Proposal presentation | A1 |

| | | | | |
|---|---|---|---|---|
| | 12-16 | | | |
| 4. | Fri 7.5<br>12-16 | Maari-E | • A2 – Help session (if required) | |
| 5. | Mon 10.5<br>12-16 | Maari-M | • A2 – Help session (if required) | |
| 6. | Wed 12.5<br>12-16 | Maari-M | • A2 – Help session (if required) | |
| 7. | Fri 14.5<br>12-16 | Maari-M | • * A2 – Demo session and seminar | * A2 |

* With staff approval, students who have chosen an ambitious design for A2 may continue working beyond the deadline.

# Bibliography


[4WA2008] The 4WARD Consortium, "The FP7 4WARD Project," 2008; available at: http://www.4ward-project.eu [accessed on 4 June 2010].

[AAMC2010] Association of American Medical Colleges, "Curriculum Directory," 2010; available at: http://services.aamc.org/currdir/section1/grading1.cfm [accessed on 4 June 2010].

[AAU2009] Association of American Universities (AAU), Association of Research Libraries (ARL), The Coalition for Networked Information (CNI) et al., "The University's Role in the Dissemination of Research and Scholarship – A Call to Action," Feb. 2009; available at: http://www.arl.org/bm~doc/disseminating-research-feb09.pdf [accessed on 4 June 2010].

[Ach2009] M. Achemlal, P. Aranda, Z. Boudjemil et al., "Draft Architectural Framework," technical report, FP7 4WARD Project, Apr. 2009; available at http://www.4ward-project.eu/index.php?s=file_download&id=35 [accessed on 4 June 2010].

[Ada2005] A. Adams, J. Nicholas, and W. Siadak, "Protocol Independent Multicast – Dense Mode (PIM-DM): Protocol Specification (Revised)," IETF RFC 3973, Jan. 2005.

[Aka2007] Akamai Technologies, "Akamai: Powering a Better Internet", 2007; available at: http://www.akamai.com [accessed on 4 June 2010].

[Atk1995] R. Atkinson, "Security Architecture for the Internet Protocol," IETF RFC 1825, Aug. 1995.

[Aur2004] T. Aura, P. Nikander, and G. Camarillo, "Effects of Mobility and Multihoming on Transport-Protocol Security," in proc. of IEEE Symposium on Security and Privacy, Berkeley/Oakland, California, USA, pp. 12-26, May 2004.

[Bha2003] S. Bhattacharyya, "An Overview of Source-Specific Multicast (SSM)," IETF RFC 3569, July 2003.

[Bit2010] BitTorrent Inc., "BitTorrent," Jan. 2010; available at: http://www.bittorent.com [accessed on 4 June 2010].

[Blo2009] R. A. Bloodgood, J. G. Short, J. M. Jackson et al., "A Change to Pass-Fail Grading in the First Two years at One Medical School Results in

Improved Psychological Well-Being," *Academic Medicine*, vol. 84, no. 5, pp. 655-662, May 2009.

[Blu2001] M. S. Blumenthal and D. D. Clark, "Rethinking the design of the Internet: The end to end arguments vs. the brave new world," *ACM Transactions on Internet Technology (TOIT)*, vol. 1, no. 1, pp. 70-109, Aug. 2001.

[Bra1994] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," IETF RFC 1633, June 1994.

[Bra1997] R. Braden (Ed.), L. Zhang, S. Berson et al., "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," IETF RFC 2205, Sept. 1997.

[Bro2003] A. Broder and M. Mitzenmacher, "Network Applications of Bloom Filters: A Survey," *Internet Mathematics*, vol. 1, no. 4, pp. 485-509, 2003.

[Cae2006] M. Caesar, T. Condie, J. Kannan et al., "ROFL: Routing on Flat Labels," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4, pp. 363-374, Sept. 2006.

[Cha2006] B. G. Charlton, "Lectures are an effective teaching method because they exploit human evolved 'human nature' to improve learning," *Medical Hypotheses*, vol. 67, no. 6, pp. 1261-1265, 2006.

[Che2000] D. R. Cheriton and M. Gritter, "TRIAD: A Scalable Deployable NAT-based Internet Architecture," Mar. 2000; available at: http://www-dsg.stanford.edu/papers/triad/triad.html [accessed on 4 June 2010].

[Dee1998] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," IETF RFC 2460, Dec. 1998.

[EC2010a] European Commission, "Seventh Framework Programme (FP7)," 2010; available at: http://cordis.europa.eu/fp7/home_en.html [accessed on 4 June 2010].

[EC2010b] European Commission, "The main objectives of FP7: Specific programmes," 2010; available at: http://cordis.europa.eu/fp7/understand_en.html [accessed on 4 June 2010].

[Eck2002] E. Eckel, "Microsoft exams move to pass/fail grading," Feb. 2002; available at: http://articles.techrepublic.com.com/5100-10878_11-1044179.html [accessed on 4 June 2010].

[Ehl2007] U. Ehlers, “REPORT: ICT research in European Union,” technical report, Winds Latin America, Mar. 2007.

[Ell2002] A. Ellis, S. Markham, J. Munro et al., “Disseminating Innovation and Best Practice in ICT Education: the ICT-ED Database,” technical report, Monash University, Victoria, Australia, Nov. 2002.

[Fas2000] Y. Fassin, “The Strategic Role of University – Industry Liaison Offices,” *Journal of Research Administration*, Summer-Fall 2000; available at http://www.entrepreneur.com/tradejournals/article/80350501.html [accessed on 4 June 2010].

[Fen1997] W. Fenner, “Internet Group Management Protocol, Version 2,” IETF RFC 2236, Nov. 1997.

[Fen2006] B. Fenner, M. Handley, H. Holbrook et al., “Protocol Independent Multicast – Sparse Mode (PIM-SM): Protocol Specification (Revised),” IETF RFC 4601, Aug. 2006.

[Fer2003] P. M. Fernández and N. McKeown, “The performance of circuit switching in the Internet,” *OSA Journal of optical networking*, vol. 2, no. 4, pp. 83-96, Mar. 2003.

[Fer2009] Ferris Research, “Cost of Spam is Flattening – Our 2009 Predictions”, 2009; available at: http://www.ferris.com/2009/01/28/cost-of-spam-is-flattening-our-2009-predictions/ [Accessed on 4 June 2010].

[Gan2004] P. Ganesan, K. Gummadi, and H. Garcia-Molina, “Canon in G Major: Designing DHTs with Hierarchical Structure,” *ICDCS*, Mar. 2004.

[Gol1971] R. M. Gold, A. Reilly, R. Silberman et al., “Academic Achievement Declines Under Pass-Fail Grading,” *The Journal of Experimental Education*, vol. 39, no. 3, pp. 17-21, 1971.

[Hag2007] Haggle partners, “FP6 Haggle Deliverable D1.2: Specification of the CHILD-Haggle,” Aug. 2007; available at: http://www.haggleproject.org/deliverables/D1.2_final.pdf [Accessed on 4 June 2010].

[Han2006] M. Handley, “Why the Internet only just works,” *BT Technology Journal*, vol. 24, no. 3, July 2006.

[Han2007] M. Handley, I. Kouvelas, T. Speakman et al., “Bidirectional Protocol Independent Multicast (BIDIR-PIM),” IETF RFC 5015, Oct. 2007.

[Har1998] D. Harkins and D. Carrel, "The Internet Key Exchange (IKE)," IETF RFC 2409, Nov. 1998.

[Hus2008] G. Huston, "Beyond the IPv4 Internet," May 2009; available at: http://www.potaroo.net/presentations/2009-05-11-Madrid.pdf [accessed on 4 June 2010].

[Int2010] Internet Society, "Histories of the Internet," 2010; available at http://www.isoc.org/internet/history [accessed on 4 June 2010].

[IP2010] The IPv6.net team, "IPv6.net," 2010; available at http://ipv6.net [accessed on 4 June 2010].

[IPR2010] IPR-Helpdesk Project, "Dissemination of Foreground under FP7," 2010; available at: http://www.fp7.org.tr/tubitak_content_files/267/IPR/Dissemination_of_Foreground_under_FP7.pdf [accessed on 4 June 2010].

[Jac2009] V. Jacobson, D. K. Smetters, J. D. Thornton et al., "Networking Named Content," in proc. of CoNEXT 2009, Rome, Italy, pp. 1-12, Dec. 2009.

[Joh2004] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6," IETF RFC 3775, June 2004.

[Jok2009] P. Jokela, A. Zahemsky, C. E. Rothenberg et al., "LIPSIN: Line Speed Publish/Subscribe Inter-Networking", in proc. of ACM SIGCOMM 2009 Conference on Data Communication, Barcelona, Spain, pp. 195-206, Aug. 2009.

[Kha2008] J. Khan, S. Davu, and R. Zaghal, "High performance mobility without agent infrastructure for connection oriented service," *Pervasive and Mobile Computing*, vol. 4, no. 4, pp. 526-545, August 2008.

[Kim2008] C. Kim, M. Caesar, and J. Rexford, "Floodless in SEATTLE: A Scalable Ethernet Architecture for Large Enterprises," in proc. of ACM SIGCOMM 2008, Seattle, WA, USA, pp. 3-14, Aug. 2008.

[Kjä2010] J. Kjällman, "Blackhawk Publish/Subscribe Prototype for FreeBSD," Jan. 2010; available at: http://wiki.hiit.fi/display/psirpcode/Blackhawk [accessed on 4 June 2010].

[Kop2007] T. Koponen, M. Chawla, B. G. Chun et al., "A Data-Oriented (and Beyond) Network Architecture," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 181-192, Oct. 2007.

[Kuz2005] A. Kuzmanovic, "The Power of Explicit Congestion Notification," in proc. of the 2005 Conference on Applications, Technologies, Architectures, and

Protocols for Computer Communications, Philadelphia, PA, USA, pp. 61-72, Aug. 2005.

[Kuz2009] A. Kuzmanovic, A. Mondal, S. Floyd et al., "Adding Explicit Congestion Notification (ECN) Capability to TCP's SYN/ACK Packets," IETF RFC 5562, June 2009.

[Lei2010] B. M Leiner, V. G. Cerf, D. D. Clark et al., "A Brief History of the Internet," 2010; available at http://www.isoc.org/internet/history/brief.shtml [accessed on 4 June 2010].

[May2002] P. Maymounkov and D. Mazieres, "Kademlia: A Peer-to-peer Information System Based on the XOR Metric," Lecture Notes in Computer Science (LNCS), vol. 2429, Jan. 2002.

[McA2009] McAfee Inc., "The Carbon Footprint of Email Spam Report," 2009; available at: http://resources.mcafee.com/content/NACarbonFootprintSpam [accessed on 4 June 2010].

[Mos2008] R. Moskowitz, P. Nikander, P. Jokela et al., "Host Identity Protocol," IETF RFC 5201, Apr. 2008.

[Nic1998] K. Nichols, S. Blake, F. Baker et al., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," IETF RFC 2474, Dec. 1998.

[Ott1972] D. J. Otto, "A Study of the Pass/Fail Grading System," 1972; available at: http://www.eric.ed.gov/ERICDocs/data/ericdocs2sql/content_storage_01/0000019b/80/39/43/06.pdf [accessed on 4 June 2010].

[PARC2010] Palo Alto Research Center Inc., "Palo Alto Research Center," 2010; available at: http://www.parc.com/ [accessed on 4 June 2010].

[Per2002] C. Perkins, "IP Mobility Support for IPv4," IETF RFC 3344, Aug. 2002.

[PSI2008a] The PSIRP Consortium, "The PSIRP Project Homepage", 2008; available at: http://www.psirp.org/ [accessed on 4 June 2010].

[PSI2008b] The PSIRP Consortium, "The PSIRP Project Wiki," Jan. 2008; available (access restricted) at: http://wiki.psirp.org/ [accessed on 4 June 2010].

[PSI2008c] A. Karila, M. Ain, D. Lagutin et al., "FP7 PSIRP Deliverable D2.1: State-of-the-Art Report and Technical Requirements," June 2008; available at: http://www.psirp.org/publications [accessed on 4 June 2010].

[PSI2008d] M. Ain, S. Tarkoma, D. Trossen et al., "FP7 PSIRP Deliverable D2.2: Conceptual Architecture of PSIRP Including Subcomponent

Descriptions," Aug. 2008; available at: http://www.psirp.org/publications [accessed on 4 June 2010].

[PSI2008e] H. Flinck, D. Trossen, and A. Karila, "FP7 PSIRP Deliverable D5.1: Project presentation," June 2008; available at: http://www.psirp.org/publications [accessed on 4 June 2010].

[PSI2008f] D. Trossen, P. Nikander, P. Jokela et al., "FP7 PSIRP Deliverable D5.2: Dissemination and Exploitation Plan," June 2008; available at: http://www.psirp.org/publications [accessed on 4 June 2010].

[PSI2008g] H. Flinck, Pekka Nikander, D. Trossen et al., "FP7 PSIRP Deliverable D5.3: Dissemination Report," Sept. 2008; available at: http://www.psirp.org/publications [accessed on 4 June 2010].

[PSI2009a] M. Ain, S. Tarkoma, D. Trossen et al., "FP7 PSIRP Deliverable D2.3: Architecture Definition, Component Descriptions, and Requirements," Feb. 2009; available at: http://www.psirp.org/publications [accessed on 4 June 2010].

[PSI2009b] H. Flinck, A. Karila, S. Tarkoma et al., "FP7 PSIRP Deliverable D5.4: Dissemination and Exploitation Report," Aug. 2009; available at: http://www.psirp.org/publications [accessed on 4 June 2010].

[PSI2009c] M. Ain, D. Trossen, K. Visala et al., "FP7 Deliverable D2.4: Update on the Architecture and Report on Security Analysis," Sept. 2009; available at: http://www.psirp.org/publications [accessed on 4 June 2010].

[Ram2001] K. Ramakrishnan, D. Floyd, and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP," IETF RFC 2001, Sept. 2001.

[RATS2010] Fortify Software Inc., "Welcome to RATS – Rough Auditing Tool for Security," 2010; available at: http://www.fortify.com/security-resources/rats.jsp [accessed on 4 June 2010].

[Res2006] E. Rescorla, "Introduction to Distributed Hash Tables," 2006; available at: https://www2.ietf.org/proceedings/65/slides/plenaryt-2.pdf [accessed on 4 June 2010].

[Rob2006] M. M. Roberts, "Lessons for the Future Internet: Learning from the Past," *EDUCAUSE Review*, vol. 41, no. 4, pp. 16-25, July/Aug. 2006; available at: http://www.educause.edu/EDUCAUSE+Review/EDUCAUSEReviewMagazineVolume41/LessonsfortheFutureInternetLea/158065 [accessed on 4 June 2010].

[Roh2006] D. E. Rohe, P. A. Barrier, M. M. Clark et al., "The Benefits of Pass-Fail Grading on Stress, Mood, and Group Cohesion in Medical Students," *Mayo Clinic Proceedings*, vol. 81, no. 11, pp. 1443-1448, Nov. 2006.

[Row2001] A. Rowstron and P. Druschel, "Pastry: Scalable, Decentralized Object Location and Routing for Large-Scale Peer-to-Peer Systems," in proc. of Middleware 2001, Heidelberg, Germany, pp. 329-250, Nov. 2001.

[Sie2009] M. Siekkinen, "Quality of Service," October 2009; available at: https://noppa.tkk.fi/noppa/kurssi/t-110.5116/luennot/T-110_5116__4.pdf.pdf [accessed on 4 June 2010].

[Sno2001] A. C. Snoeren, H. Balakrishnan, and M. F. Kaashoek, "Reconsidering Internet Mobility," in proc. of 8th Workshop on Hot Topics in Operating Systems (HotOS), Elmau/Oberbayern, Germany, pp. 41-46, May 2001.

[Sto2001] I. Stoica, R. Morris, D. Karger et al., "Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications," in proc. of ACM SIGCOMM 2001, San Diego, CA, USA, pp. 149-160, Aug. 2001.

[Sto2004] I. Stoica, D. Adkins, S. Zhuang et al., "Internet Indirection Infrastructure," *IEEE/ACM Transactions on Networking,* vol. 12, no. 2, pp. 205-218, Apr. 2004.

[Tar2009] S. Tarkoma, M. Ain, and K. Visala, "The Publish/Subscribe Internet Routing Paradigm (PSIRP): Designing the Future Internet Architecture," *Towards the Future Internet*, IOS Press, pp. 102-111, 2009.

[TKK2010] The Helsinki University of Technology, "Information systems," Jan. 2010; available at http://www.tkk.fi/en/studies/information_systems/ [accessed on 4 June 2010].

[Tro2005] W. M. K. Trochim and J. P. Donnelly, "Likert Scaling," *The Research Methods Knowledge Base 3rd Ed.*, Atomic Dog Publishing, 2005.

[Yan2003] X. Yang, "NIRA: a new Internet routing architecture," in proc. of ACM SIGCOMM Workshop on Future Directions in Network Architectures, Karlsruhe, Germany, pp. 301-312, Aug. 2003.

[Zha2004] B. Y. Zhao, L. Huang, J. Stribling et al., "Tapestry: A Resilient Global-Scale Overlay for Service Deployment," *IEEE Journal on Selected Areas in Communications*, vol. 22, no. 1, pp. 41-53, Jan 2004.