Developing Information Networking Further: From PSIRP to PURSUIT

Nikos Fotiou[1], Pekka Nikander[2], Dirk Trossen[3], and George C. Polyzos[1]

[1] Athens University of Economics and Business, Athens, Greece
{fotiou,polyzos}@aueb.gr
[2] NomadicLab, Ericsson Research, Finland
pekka.nikander@ericsson.com
[3] Computer Laboratory, University of Cambridge
dirk.trossen@cl.cam.ac.uk

Abstract. PSIRP (Publish-Subscribe Internet Routing Paradigm) is an EU FP7 funded project that has developed a clean-slate architecture for the future Internet, based on the publish-subscribe primitives (rather than the send-receive ones), all the way down to the core networking functions. The PSIRP vision is a pure information-centric Internet architecture, possibly providing remedies to many of the current Internet problems. In PSIRP, all is information and everything is about information. Content-based identities, recursive application of ideas, cryptographic techniques, and the Trust-to-Trust principle are all extensively used to achieve the design goals. Furthermore, incentive compatibility and socio-economic considerations are guiding the design from the outset, to ground the project in reality and to provide credible and viable potential deployment paths. The project has developed, implemented, and preliminarily evaluated solutions for rendezvous, topology formation and routing, and information forwarding, with ongoing work currently focusing in experimenting.

A new (also EU FP7 funded) follow-on project, PURSUIT (Publish-Subscribe Internet Technologies), will refine and further explore and expand PSIRP's vision. We believe that this will eventually lead to a more complete architecture and protocol suite, thereby providing for more extensive performance evaluation and investigations on scalability. This paper provides an overview of the PSIRP concepts and the developed architecture, along with some key results, and outlines the research directions of the PURSUIT project, focusing on the project goals and its expected outcomes.

Keywords: PSIRP, PURSUIT.

1 Introduction

The current Internet architecture cannot effectively and efficiently handle various challenges, including security, mobility, scalability, quality of service, and economics, due to the focus and shortcomings in its original design [7]. The Internet, even though based on packet switching, a critical departure from circuit

switching used in the telephone network, still resembles significantly to the telephone network in many design choices — and since its inception this design has been hardly changed. A key goal of the Internet was to efficiently interconnect mainframes and minicomputers, and to provide efficient remote access to them. This end-to-end approach and especially its specific practical implementation, however, have been identified as a root cause of many limitations of the current Internet architecture. Various add-ons, such as NATs, Mobile IP, CDNs, p2p overlays, etc., all violate, in various ways, several aspects of the original Internet architecture in order to provide answers to features that were not part of the original design (or the original requirements). Moreover, the original Internet architecture and protocols were developed assuming a benign and cooperative environment, which is far from today's reality, where competition is widespread and lack of trust and security threats, such as malware, denial of service attacks, and phising, have become more and more prevalent.

More recently, it has been observed that information has become at the heart of almost all communication and, in particular, of Internet use [10]. This information-centric usage of the Internet raises various architectural challenges, many of which are not effectively handled by the current Internet architecture [29]. The challenges include medium-independent information access, tussle mediation through information governance, privacy and accountability through controlled information dissemination, and information scarcity.

It seems that it is time for a shift from the current Internet, which interconnects machines, towards a new Internet that interconnects information. PSIRP (Publish-Subscribe Internet Routing Paradigm), an EU FP7 funded research effort, has created, implemented, and initially evaluated a clean-slate, information oriented future Internet architecture, aiming at overcoming most limitations of the current Internet. Its goal is to bring out the role of information as the main building block of the (future) Internet. This new architecture is based on a paradigm completely different from the current one. PSIRP is based on pure, through-the-stack application of the Publish-Subscribe paradigm.

The Publish-Subscribe paradigm is an alternative to the commonly used Send-Receive paradigm. The communication architectures that are build based on this paradigm are composed of three basic elements: publishers, subscribers, and a network of brokers [6]. Publishers are information owners or disseminers that advertise information availability by issuing publication messages. Subscribers are information consumers, expressing their interest for specific information items by issuing subscriptions. Between the publishers and subscribers there exist a network of brokers that is responsible for routing publication and subscription messages, for matching publications with subscriptions, as well as for initiating the information forwarding process from publishers towards subscribers. The broker at which the publication-subscription matching takes place (for a particular publication) is known as the rendezvous point (RP) (of that publication). The existence of such rendezvous points allows for publisher-subscriber decoupling both in terms of location and time, location-identity split, multihoming and mobility, and anonymity.

The PSIRP project envisions a publish-subscribe-based, information-centric Future Internet architecture, in which the current Internet limitations will be removed through the design of a set of new core network mechanisms, rather than by add-on fixes. It envisions a robust, reliable, and extensible Internet. PSIRP has produced a specific internetworking architecture, with specific proposals in many key areas, such as rendezvous, topology formation and routing, and information forwarding. It includes integrated support for anycast and multicast, caching, multihoming and mobility, and security and privacy. A PSIRP prototype has been implemented and is available under open source license terms [24]. Moreover, a PSIRP testbed across Europe is already being established for testing native PSIRP applications.

The PURSUIT project can be considered as a continuation of the PSIRP work. Starting in fall 2010, PURSUIT is expected to revisit the PSIRP architecture, to develop missing or only partially developed components, and to produce and evaluate alternative designs, as well as to expand dissemination and exploitation. A significant challenge is to scale the practical experimental network from the current few nodes to hundreds or more. In this paper, we provide an overview of the PSIRP architecture and present its core elements and functions. The rationale and solutions that PSIRP has developed, in order to address the requirements of a Future Internet architecture, are being presented, along with our vision for their planned culmination in PURSUIT.

The remainder of this paper is organized as follows. Section II introduces PSIRP and its architecture, as well as its key contributions. In section III, the main goals of the PURSUIT project are highlighted, and, in Section IV, related work in other projects, worldwide, is presented. Finally, in Section V, our conclusions as well as thoughts for additional work are presented.

2   PSIRP

The PSIRP project investigates a clean-slate internetworking architecture focusing on the information itself, rather than on the location of the information. The system architecture is built around information from the viewpoints of 'semantics' (i.e. meaning) and 'scope' (i.e. breadth of coverage). Multicast, caching in the network, mobility of end-points and information are considered mainstream and not exceptional cases that should be handled with add-on mechanisms. Moreover, security is diffused throughout all the levels of the architecture.

The PSIRP architecture abides to the following principles [27]:

- PSIRP specific principles. Here, two concepts are fundamental: the recursive nature of *information semantics*, and the notion of *information scope* as a concept of reachability.
  - Information is organised in a *multi-hierarchical* manner.
  - *Scopes* act as a means of information organisation; there are mechanisms that allow for scoping information on different levels of semantics (e.g. via means of rendezvous, discovery, search, and others). This *limits the*

> *reachability of information* to the parties having access to the particular mechanism (e.g. rendezvous) that implements the scoping.
> - The architecture is *neutral* with regard to the semantics and structure of data.
> - The architecture is *receiver-driven*; no entity shall be delivered data unless it has agreed to receive it beforehand through appropriate signalling methods (limitations to this principle can apply, for instance, at the physical level).

- General principles. These principles are shared with other architectural approaches, but are still fundamental to the PSIRP architecture.
  - Trust-to-trust principle: Functions are implemented at points of the network that can be considered trustworthy from the user's (of the function) perspective.
  - Market creation: The architecture supports potentially multi-dimensional metrics of compensation that will enable effective market creation.
  - Multi-dimensionality: The effectiveness of any architectural solution is evaluated under a (multi-dimensional) metric along the dimensions of financial, social, environmental, and economical benefits.
  - Evolvability: Any solution should enable proper evolution beyond itself.
  - Minimality: No entities are introduced beyond the necessity to implement the design principles of the architecture.

2.1 Components and Functions

Information is the core component of the PSIRP architecture; in PSIRP everything is information [28], ranging from small ('meaningless') chunks of data to complex information items such as documents, images and videos. Information items are organized within *scopes*. Scopes may denote physical structures, e.g., a corporate network, or logical structures, e.g., my friends in *facebook*. This allows for information location and control over information dissemination. Every information item is uniquely identified by a pair of identifiers; a rendezvous identifier (RId) and a scope identifier (SId). The RId is information item specific and it has to be unique within a scope, whereas the SId denotes the scope in which an information item belongs. Scopes can be included within each other, creating a flexible structure. Both RIds and SIds are flat and endpoint independent. Flat and endpoint independent labels seem to be a natural choice for information oriented architectures as they clearly separate location from identity, allowing for properly incorporating mobility, multicasting, and multihoming into the architecture, as well as a more comprehensive notion of identity [2].

In order for a publisher to issue a publication in PISRP, it has to know the SId of a scope within which it wants the publication to be published, as well as to create a RId for the publication. The publication's RId is then forwarded to the rendezvous node of the SId rendezvous network, which manages the publication's RId. Figure 1 illustrates a publication example. A publisher creates a publication that she wants to publish in her blog scope. A publication message is created, carrying the publication's RId as well as the blog's SId. This publication message
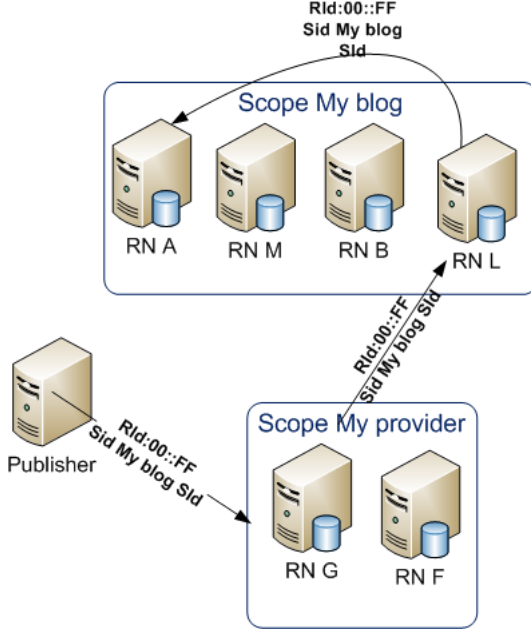
Fig. 1. Publication in a PSIRP network

is forwarded to the rendezvous node RN L, which is responsible for managing the blog's scope. In this case, the blog's scope consists of four rendezvous nodes, one of which (RN A) is the rendezvous point for this specific RId; therefore the publication message ends up at this node.

Subscriptions in PSIRP follow a similar process. A subscriber learns the RId and the SId of a desired piece of information and issues a subscription message towards the appropriate rendezvous point. When his subscription message reaches the rendezvous point, a forwarding path is created from the publisher towards the subscriber and the desired piece of information is sent through this path. The forwarding path creation process, as well as the forwarding procedure, follow an MPLS-like, label-based approach [28]. Each active publication is assigned a forwarding identifier (FId) that denotes the forwarding path that this item has to follow.

PSIRP's operation is based on three basic functions that are recursively executed in all layers of the architecture. These functions are the *rendezvous*, the *topology and routing*, and the *forwarding* functions [25]. The rendezvous function is responsible for matching subscribers' interests with publications. The topology function monitors the network topology and detects changes using various techniques, depending on the layer of the architecture within which it is implemented. Within the same module, the routing function is responsible for creating information delivery paths. Finally, the forwarding function implements information forwarding throughout the delivery paths.

2.2   Key Outcomes

PSIRP has studied the potential of having an Internet scale architecture where every information item is identified by flat, endpoint-independent identifiers. The notion of the organization of information items within scopes, introduced in this project, is expected to offer the desired scalability. Moreover, PSIRP has introduced the usage of algorithmically related information identifiers, allowing for information item grouping [28].

PSIRP has achieved line-speed label-based forwarding with the usage of an in-packet Bloom-filter (iBF) based mechanism [11]. An iBF is a data structure that encodes the data delivery path in a compact manner. Although constant in size, iBFs allow for Internet scale data delivery paths, enabling at the same time loop prevention, fast path recovery, and small-scale multicast. iBFs may be constructed in an information-flow-specific manner, with every delivery path being encoded in a flow-identifier dependent way. This makes it almost impossible for an attacker to create crafted iBFs that would lead to DoS attacks or to information leakage. As a backwards compatible option, iBFs formation is also possible through the usage of an MPLS-like approach, named MPSS [30]. MPSS is an approach designed to compute a delivery path or tree on the graph of the network, to compute an iBF for it, and to optionally allocate resources. Moreover, the NetFPGA implementation of the basic iBFs demonstrated the applicability of this technique over high-speed (1Gbps) links [17].

PSIRP has also studied the security requirements for a publish-subscribe based Internet architecture [21], as well as the possibilities for applying existing work on cryptographic protocol analysis in a pure publish-subscribe architecture [23]. Moreover, new security mechanisms have been developed. As an example, PLA [19] is a novel mechanism for protecting the networking architecture, based on the assumption that per-packet public key cryptographic operations are possible at wire speed in high speed networks, due to new cryptographic algorithms and advances in semiconductor technology. PLA can also be used to provide accountability, billing and Internet-wide roaming while maintaining privacy and reducing the burden on the access networks providers [20]. At the higher layers of the architecture, information-oriented solutions have also been investigated. For example, Information ranking [9] is a mechanism developed for ranking information items based on positive votes. This mechanism achieves better results, compared to the commonly used user's ranking-based mechanism, when it comes to polluted information pieces isolation, as well as to spam prevention in publish-subscribe architectures [8].

In order to evaluate PSIRP's efficacy as well as the applicability of PSIRP principles, an overlay architecture on top of access network routers was designed, supporting multicast, caching, and mobility. The goal of this architecture is to allow *proxy* participation of end-hosts into an inherently information-centric overlay architecture. In this way, operators are expected to gain control of the traffic in their network and benefit from the multicast data delivery mode, while improving the resulting data delivery structures by eliminating data forwarding from end-hosts. The information-centric character of the proposed architecture

was further enhanced by a distributed caching scheme designed to operate in conjunction with the multicast delivery mechanism so that content is provisioned either via multicast or via unicast from a nearby cache location [16]. This overlay variant was evaluated through the simulation of large scale deployments and the results verified expectations. As far as the multicast tree creation is concerned, this overlay variant shows a significant improvement in the resulting tree properties, such as path stretch and link stress, as well as a reduction in the overlay network establishment and maintenance overhead, at the cost of increased forwarding overhead for the access routers [12]. Moreover, the caching mechanism achieves a high cache hit ratio, reaching values up to 98,5% when it comes to large cache sizes. Finally, it was found that when it comes to mobile nodes, our architecture achieves 50% less packet loss than Mobile IPv6 as well as almost zero resume time [14,15].

Taking a step further, since a synchronized Internet-wide deployment of any network architecture seems infeasible, particular attention was paid to the evaluation of this overlay architecture as it gradually gets deployed across the Internet. In this direction, the specifics of an *incremental deployment* process, in which the ISPs progressively add the overlay publish-subscribe functionality on top of their existing infrastructure, were investigated. This investigation included both the dimensions of *inter-ISP* and *intra-ISP* incremental deployment. The simulation results demonstrated that network operators have strong incentives to adopt the proposed architecture, such as significant gains in terms of transmission load in their network and end user satisfaction. Even after initial deployment, network operators supporting the proposed architecture have a consistent advantage over the remaining ones. The simulation results also show that sparse intra-domain deployments are sufficient, indicating that an operator can reap the benefits of the proposed architecture by deploying it at only 25% of its routers, regardless of what other operators do [13].

In order to perform experiments with native PSIRP applications, a multi-site PSIRP testbed is being created. Figure 2 shows the PSIRP testbed setup. The testbed spans across Europe – Athens, Sofia, Aachen, Essex, Cambridge, and Helsinki – and there are plans for expanding it to U.S., with M.I.T. participation. It includes VPN links, as well as two fibre links. Every node in the PSIRP testbed runs a PSIRP network prototype and applications, such as file transfer, real-time voice and video streaming, which will be tested and evaluated.

3   PURSUIT

PURSUIT aims at revisiting the PSIRP architecture, filling in missing components and improving the existing ones, by taking advantage of all the lessons learned and of all the experience that has been accumulated during PSIRP's span. All layers of the architecture, including the link and the physical layers, of both wireless and wireline technologies, will be considered with an *information-only* and *address-less* approach.
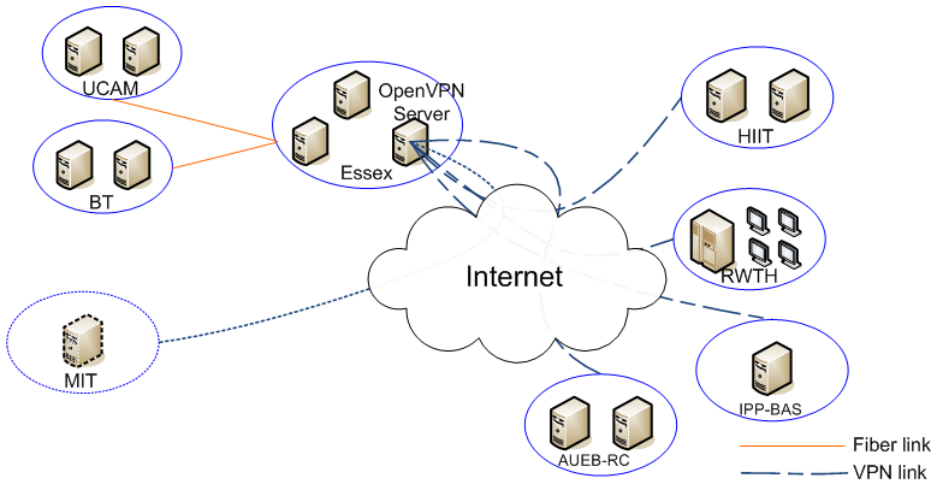
Fig. 2. PSIRP Testbed

The PURSUIT project will abide to the Design for Tussle approach [4], enabling flexibility when it comes to existing as well as to new tussles. The information-centric nature of this new architecture is expected to allow better resource utilization in both wireless and wireline environments. More specifically, in wireless environments, better spectrum allocation and effective mobility handling are anticipated, whereas in wireline environments, PURSUIT is expected to harvest the full potential of the future transparent optical networks. In particular, in-network caching will be further investigated and enhanced, leading to better and more effective information distribution, including completely new transport protocols. Throughout the project security will play crucial role, with trust, privacy, and information accountability being in the research spotlight.

3.1   Main Goals

The flat-label-based information identification introduced in PSIRP will be maintained in the PURSUIT architecture. Every information item in PURSUIT will be associated with – at least – one scope, and the information organization will follow the same principles as in the PSIRP architecture. One of the primary goals of PURSUIT is the development of solutions and mechanisms that will enable providers to take advantage of these new information structures and provide innovative services that will offer them competitive advantages. Moreover, these solutions and mechanisms should be able to resolve the tussles that will occur in this information oriented environment.

In PSIRP, the notion of the algorithmic identifiers was introduced. They provide a compact and convenient way to identify pieces of information collectively, but also to identify individual sub-pieces of information. In PURSUIT, it is expected that this notion will be further studied and ultimately it will be provided

as a core service of the architecture. The possibility of correlating information items in an algorithmic manner opens the floor for new fragmentation techniques that will lead to the deployment of effective caching schemes. Information items will be intelligently distributed and cached. PURSUIT aims at providing caching as an in-network mechanism, in addition to the caching services that will be offered by various providers, achieving this way significant resource utilization improvement.

Caching is only one way to achieve the goal of better resource utilization, with multicast being another. In PURSUIT, we will also focus on the lower layers and on mechanisms that will enable better resource allocation and utilization at the link and physical layers, by taking advantage of the information structures that at the higher layers. We will try to show how wireless and wireline network access methods can benefit from a pure information oriented architecture.

The declared goal of PSIRP to consider mobility as a normal condition and to handle it with in-network mechanisms rather than with add-ons, remains in PURSUIT. Multicast and caching are expected to facilitate mobility support and enhance performance. Moreover, the recursive nature of the rendezvous, topology, routing, and forwarding functions is expected to allow for handling mobility as close to the mobile node as possible, meanwhile hiding it from the core network.

In PSIRP, PLA and information-centric security mechanisms were introduced. PURSUIT's goal is to enhance the notion of trust, to provide accountability mechanisms even at the physical layer, while maintaining privacy. In addition, to develop a framework that will enable the flexible definition of policies that will dictate topology formation, forwarding decisions and information dissemination, at all levels of the architecture.

In order for a publish-subscribe solution to be deployed, proper evaluation, understanding of the socio-economics forces as well as a realistic migration strategy are needed. PURSUIT anticipates achieving all of them. The PSIRP testbed, the large scale simulation tools whose development started during PSIRP, as well as the socio-economic work and the overlay variant of PSIRP are expected to be the driving forces that will lead to the achievement of this goal. Moreover, the PURSUIT team will seek collaborations with other projects developing similar architectures and technologies in order to join forces towards a realistic evaluation framework.

3.2 Expected Outcomes

One of the main outcomes of the PURSUIT project is expected to be a suite of information-centric protocols, solutions, and mechanisms that will handle all aspects of inter-networking, including mobility, caching, transport, flow and error control. PURSUIT's anticipation is that all these methods will have been properly tested and evaluated under realistic conditions, and, therefore, they will be relatively ready for deployment in the future Internet.

PURSUIT is expected to have a significant impact on the techniques used in order to build wireless and wireline networks. Applying an information-only

approach to the physical as well as to the link layer will lead to a new form of network access. This new form of network access will cooperate with the higher layers of the architecture, which will also be organized around information (but for which this seems less of a dramatic departure from the status quo).

New models and tools regarding security and privacy will be created during the PURSUIT project. This new security architecture is going to consider new forms of threats and risks and their impact has to be modelled and evaluated. Furthermore, the novel security mechanisms that are expected to be created throughout the duration of this project will have to be tested and evaluated using new analytical techniques as well as new simulation frameworks.

A project such as PURSUIT is expected to have a significant socio-economic impact. An important outcome of this project will be estimation and evaluation of this impact. This estimation is planned to be achieved by identifying the tussles that will occur in this new form of inter-networking. PURSUIT will come with new socio-economic models as well as with solutions that aim towards enabling new market opportunities and a healthy socio-economic playground.

Finally, the PURSUIT project is expected to release various prototypes which, with the appropriate APIs, will allow researchers to develop new solutions that will take advantage of the PURSUIT architecture. The PURSUIT testbed, which will be deployed across all partners' sites, is hoped to play a significant role in achieving this objective.

4 Related Work

Significant, fundamental research has been undertaken in the last few years, investigating core issues of inter-networking and many of the known shortcomings of the current Internet technologies and practices.

The Internet Indirection Infrastructure (i3) [26] and the Host Identity Protocol (HIP) [22] introduced indirection as a solution to the problems that point-to-point communication poses to mobility, multicast, and multihoming. i3 implements an IP overlay network that replaces the point-to-point communication model with a rendezvous-based paradigm, where senders send packets to a specific rendezvous-point, while receivers issue triggers on specific packet identifiers. HIP introduces a new layer in the internetwork stack between the IP layer and the transport layer. This new layer decouples host identity from location identity. PSIRP uses similar concepts through the rendezvous and topology formation processes. The problem of routing based on flat information identifiers rather than on hierarchical location-based identifiers has also been studied in the Data-Oriented Network Architecture (DONA) [18] and in the Routing on Flat Labels (ROFL) [2] projects. PSIRP borrows the information identification concept of DONA, but chooses a separate inter-domain architecture with slow and fast paths. Moreover, PSIRP extends ROLF towards flat identifiers within hierarchical *scopes*, which are expected to offer faster information scoping, better scalability, and piecewise dissemination.

Various research projects are currently aiming at redesigning the Internet with an information-centric or content-centric perspective. CCNx [3] is a research

effort that proposes routing based on hierarchical naming. In CCNx consumers ask for content by broadcasting 'Interest' packets that contain the name of the content requested. Any 'Data' packet whose content name is a suffix of the name in the 'Interest' packet is conspired that it satisfies this interest. PSIRP, on the other hand, introduces flat label identifiers organized into scopes, allowing for a variety of naming approaches to be layered on top of the internetworking architecture. Moreover, although overlaying PSIRP over the current Internet is possible, it is the declared goal of the PSIRP project to investigate a native solution that could replace the current inter-networking technology. This leads to a focus on inter-domain functions, which is not found in CCNx.

4WARD [1], another EU FP7 funded on-going research project, also advocates an information-centric Internet which will enable network diversity, allowing various types of networks to co-exist and cooperate in a smooth and cost-efficient manner. It envisions an Internet where networks will be self-manageable and network paths will be an active networking component that will be able to affect transport services. 4WARD borrows concepts from DONA in terms of labelling, and intends to shed light on business aspects, similarly to the socio-economic work in PSIRP. Finally, COMET [5] is a more recent (EU FP7 funded) research effort (launched in January 2010), aiming at creating a unified approach to content location, access, and distribution, irrespectively of the intermediary used. COMET, similarly to PURSUIT, plans on using some PSIRP, or similar, concepts. Nevertheless, COMET is more focused at the higher layers of the architecture.

5 Conclusion

The Publish-Subscribe Internet Routing Paradigm (PSIRP) project, concluding in the Fall of 2010, has developed a clean-slate information-centric architecture for the future Internet, based on publish-subscribe (rather than send-receive) primitives, all the way down to the core networking functions. Cryptographic techniques, layerless recursive design, and the Trust-to-Trust principle are extensively used to achieve the design goals. Incentive compatibility and socio-economic considerations are guiding the design, providing potential for a credible and viable deployment path. Solutions for rendezvous, topology formation and routing, information forwarding, including support for multicast, caching, mobility, and security, have been developed, preliminarily evaluated, implemented in prototypes, and are currently being experimented with within the PSIRP project. A new, follow-on project, PURSUIT (Publish-Subscribe Internet Technologies), will further refine and expand PSIRP's technologies, eventually leading to a more complete architecture and protocol suite and more extensive performance evaluation and investigation of scalability. In this paper, we have provided an overview of PSIRP concepts, the developed architecture, and briefly discussed key results. We have also outlined the research directions of PURSUIT, focusing on its goals and expected outcomes.

References

1. 4WARD. Web site (2010), `http://www.4ward-project.eu`
2. Caesar, M., Condie, T., Kannan, J., Lakshminarayanan, K., Stoica, I.: ROFL: routing on flat labels. ACM SIGCOMM Computer Communication Review 36(4), 374 (2006)
3. CCNx. Web site (2010), `http://www.ccnx.org`
4. Clark, D., Wroclawski, J., Sollins, K., Braden, R.: Tussle in cyberspace: defining tomorrow's internet. IEEE/ACM Transactions on Networking 13(3), 462–475 (2005)
5. COMET. Web site (2010), `http://www.comet-project.org`
6. Eugster, P.T., Felber, P.A., Guerraoui, R., Kermarrec, A.M.: The many faces of publish/subscribe. ACM Comput. Surv. 35(2), 114–131 (2003)
7. Feldmann, A.: Internet clean-slate design: what and why? SIGCOMM Computer Communication Review 37(3), 59–64 (2007)
8. Fotiou, N., Marias, G., Polyzos, G.C.: Fighting Spam in Publish/Subscribe Networks Using Information Ranking (2010)
9. Fotiou, N., Marias, G., Polyzos, G.C.: Information Ranking in Content-Centric Networks (2010)
10. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L.: Networking named content. In: CoNEXT 2009: Proc. of the 5th International Conference on Emerging Networking Experiments and Technologies, pp. 1–12 (2009)
11. Jokela, P., Zahemszky, A., Esteve Rothenberg, C., Arianfar, S., Nikander, P.: LIPSIN: line speed publish/subscribe inter-networking. In: SIGCOMM 2009: Proc. of the ACM SIGCOMM 2009 Conference on Data Communication, pp. 195–206 (2009)
12. Katsaros, K., Bartsotas, N., Xylomenos, G.: Router assisted overlay multicast (2009)
13. Katsaros, K., Stais, C., Xylomenos, G., Polyzos, G.C.: On the incremental deployment of overlay information centric networks (2010)
14. Katsaros, K., Fotiou, N., Polyzos, G.C., Xylomenos, G.: Overlay multicast assisted mobility for future publish/subscribe networks (2009)
15. Katsaros, K., Fotiou, N., Polyzos, G.C., Xylomenos, G.: Supporting mobile streaming services in future publish/subscribe networks (2009)
16. Katsaros, K., Polyzos, G.C., Xylomenos, G.: A hybrid overlay multicast and caching scheme for information-centric networking (2010)
17. Keinänen, J., Jokela, P., Slavov, K.: Implementing zFilter based forwarding node on a NetFPGA. In: Proc. of NetFPGA Developers Workshop (2009)
18. Koponen, T., Chawla, M., Chun, B.G., Ermolinskiy, A., Kim, K.H., Shenker, S., Stoica, I.: A data-oriented (and beyond) network architecture. ACM SIGCOMM Computer Communication Review 37(4), 192 (2007)
19. Lagutin, D.: Redesigning internet-the packet level authentication architecture. Licentiate's Thesis in Computer Science, Helsinki University of Technology, Espoo, Finland (2008)

20. Lagutin, D., Tarkoma, S.: Cryptographic Signatures on the Network Layer - an Alternative to the ISP Data Retention. In: Proc. of IEEE Symposium on Computers and Communications, ISCC, Riccione, Italy (2010)
21. Lagutin, D., Visala, K., Zahemszky, A., Burbridge, T., Marias, G.: Roles and Security in a Publish/Subscribe Network Architecture (2010)
22. Moskowitz, R., Nikander, P.: Host identity protocol (HIP) architecture. Tech. rep., RFC 4423 (May 2006)
23. Nikander, P., Marias, G.: Towards Understanding Pure Publish/Subscribe Cryptographic Protocols. In: Sixteenth International Workshop on Security Protocols, Cambridge, England (2008)
24. PSIRP. Web site (2010), `http://www.psirp.org`
25. Särelä, M., Rinta-aho, T., Tarkoma, S.: RTFM: Publish/subscribe internetworking architecture. In: Proc. of the ICT Mobile Summit (2008)
26. Stoica, I., Adkins, D., Ratnasamy, S., Shenker, S., Surana, S., Zhuang, S.: Internet Indirection Infrastructure. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 191–202. Springer, Heidelberg (2002)
27. Tarkoma, S. (ed.): PSIRP deliverable 2.2, conceptual architecture of PSIRP including subcomponent descriptions, D2.2 (2008), `http://www.psirp.org/`
28. Tarkoma, S. (ed.): PSIRP deliverable 2.3, architecture definition, component descriptions, and requirements, D2.3 (2008), `http://www.psirp.org/`
29. Trossen, D., Sarela, M., Sollins, K.: Arguments for an information-centric internetworking architecture. SIGCOMM Comput. Commun. Rev. 40(2), 26–33 (2010)
30. Zahemszky, A., Jokela, P., Särelä, M., Ruponen, S., Kempf, J., Nikander, P.: MPSS: Multiprotocol Stateless Switching (2010)