

# On Name-oriented Publish/Subscribe in the Future Internet\*

Thomas C. Schmidt  
HAW-Hamburg  
Dept. Informatik  
Email: t.schmidt@ieee.org

Matthias Wählisch  
Freie Universität Berlin & link-lab  
Institut für Informatik  
Email: waehlich@ieee.org

Dominik Charousset, Sebastian Meiling  
HAW Hamburg  
Dept. Informatik  
Email: {first.last}@haw-hamburg.de

## I. INTRODUCTION

The search for future directions in Internet development has drawn significant attention to name-based communication concepts that are built upon the publish/subscribe paradigm. Inspired by the Web use case and widely deployed content distribution networks, but also by real-time applications like IPTV or MMORGs, proposals for name-oriented content networking abandon the current Internet model of pairwise inter-connecting end nodes.

Publish/subscribe approaches in its various forms [1] separate the act of content publishing from its consumption and thereby decouple senders from receivers in one or several of the dimensions (a) communication state, or (b) content location, or (c) time of access, while names act as the tie between communicating parties. Distributed pub/sub systems face the major challenges of

- 1) scalability
- 2) transparent rendezvous
- 3) compliance to application needs.

In this work, we analyze and discuss the facets of name-oriented multicast in a general publish/subscribe Internet.

## II. PUBLISH/SUBSCRIBE INTERNET MODELS

### A. Information Centric Networking

In recent discussions, the concept of *Information Centric Networking* (ICN) shaped a vision of a future publish/subscribe Internet. In ICN, consumers shall retrieve content by name directly from a network that provides storage, caching, content-based rendezvous, and searching at times. Several proposals have been presented, among them TRIAD [2], DONA [3], NDN [4], [5], PSIRP/LIPSIN [6], and NetInf [7]. The schemes differ in naming/addressing, routing/rendezvous, security/authentication, forwarding/caching, and minor design choices. They jointly consider a decoupling in space and time, on the price of communication states that are coupled to the forwarding plane by data-driven events.

\*This work is part of the HVMcast project (<http://hamcast.realmv6.org>), which is funded by the German Federal Ministry of Education and Research. HVMcast is part of G-Lab (<http://www.german-lab.de>).

### B. Multicast

Multicast is the traditional approach to publish/subscribe on the Internet layer [8], but remained technology-bound and lacks a uniform naming. In contrast to ICN, multicast enforces a highly efficient, scalable data forwarding for simultaneously operating senders and receivers. ICN also facilitates the consumption of a single (cached) data copy by multiple receivers, but requires content subscribers to act nearly synchronized in order to profit from tree distribution. Occasionally, traditional multicast is proposed to improve the ICN efficiency, e.g., in MultiCache [9]. Multicast can be implemented without data-driven events.

### C. HVMcast

HVMcast, our approach to hybrid adaptive mobile multicast [10], introduces a uniform naming scheme along with a common API [11], and thereby grants simultaneous access to varying distribution technologies. The application programmer and the end user operate on transparent, content-centric identifiers, which will be mapped to technology-specific addresses or names. HVMcast focuses on the integration of pluralism in network service deployment. It dynamically selects distribution technologies provided at the current environment and hides technology-specific treatment at the socket level. This ‘Virtualized Socket’ architecture is visualized in Figure 1. The HVMcast system architecture aims for ease of deployment, availability, and efficiency. Our approach is evolutionary and rather extends the network architecture than making changes to existing protocols.

## III. DISCUSSION AND FUTURE DIRECTIONS

### A. Distribution Paradigms

An important goal of future communication in a publish/subscribe Internet must be seen in a flexible decoupling of parties. While multicast separates senders from receivers in space and state, ICN approaches re-introduce states to (partially) decouple content distribution in time. ICN’s opt for network states must be considered a strong design decision. We investigated performance aspects of the different content-centric approaches in detail, dedicating special focus to the impact of state management from mapping, routing and forwarding. Our comparative analysis revealed that data-driven state management in ICN severely threatens the efficiency,

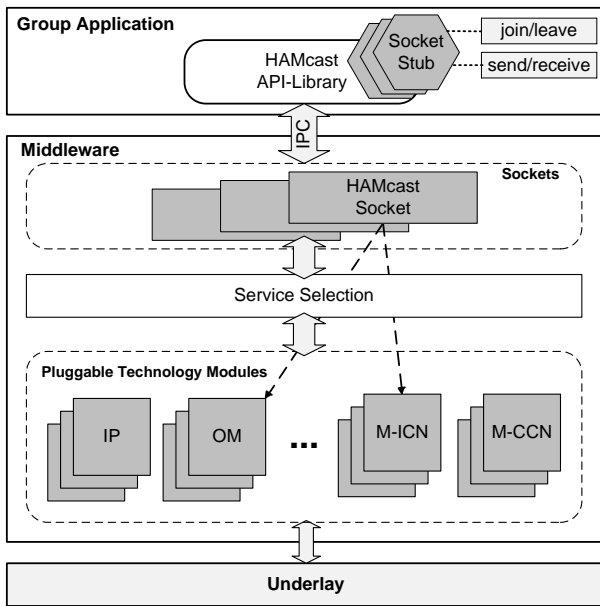


Fig. 1. System-centric Multi-service Architecture of HAMcast

stability and security of the corresponding systems [14], [15], while a pre-cached naming states in HAMcast do not show visible impact.

From a conceptional perspective, HAMcast may seamlessly extend into the information-centric networking realm [12]. Cache-based content replication may be located at gateways within the system middlewares. However, future directions in name-oriented content distribution need to specify detailed mechanisms with a focus on flexible access to technologies that can provide the scalability of decoupled states, a transparent replication of content, and a controllable separation in time.

### B. Programming Paradigms

Functions of the base network layer will only show practical relevance, if a general, easy and widespread programming interface is present at the community of application developers. Such API needs embedding in an appropriate programming paradigm and support of common languages.

The service abstraction provided by HAMcast along with the common API presents a first step towards a programming abstraction. However, application access at a higher level than the common API is needed to comply to developers needs. We report on early work to raise transparent group communication to the distributed message passing paradigm of Actors. Our approach extends the Actor model by a publish/subscribe-based, technology-independent and extensible group communication layer [13]. We developed `libcppa`<sup>1</sup>, a C++ library that seamlessly extends the language by loosely coupled message passing primitives for Actors.

This not only widens the realm of Actors to Internet-size scalability, but may fill a common gap in today's mainstream

languages that do not provide a common interface for the publish/subscribe paradigm. Future work will elaborate on enhanced concepts of universal applicability.

### REFERENCES

- [1] P. T. Eugster, P. A. Felber, R. Guerraoui, and A.-M. Kermarrec, "The Many Faces of Publish/Subscribe," *ACM Comput. Surv.*, vol. 35, no. 2, pp. 114–131, June 2003.
- [2] M. Gritter and D. R. Cheriton, "An Architecture for Content Routing Support in the Internet," in *Proc. USITS'01*. Berkeley, CA, USA: USENIX Association, 2001, pp. 4–4.
- [3] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A Data-Oriented (and beyond) Network Architecture," *SIGCOMM Computer Communications Review*, vol. 37, no. 4, pp. 181–192, 2007.
- [4] L. Zhang, D. Estrin, J. Burke, V. Jacobson, and J. D. Thornton, "Named Data Networking (NDN) Project," PARC, Tech.report ndn-0001, 2010.
- [5] V. Jacobson, D. K. Smetters, J. D. Thornton, and M. F. Plass, "Networking Named Content," in *Proc. of the 5th Int. Conf. on emerging Networking EXperiments and Technologies (ACM CoNEXT'09)*. New York, NY, USA: ACM, Dec. 2009, pp. 1–12.
- [6] P. Jokela, A. Zahemszky, C. E. Rothenberg, S. Arianfar, and P. Nikander, "LIPSIN: Line Speed Publish/Subscribe Inter-networking," in *Proc. of the ACM SIGCOMM 2009*. New York, NY, USA: ACM, 2009, pp. 195–206.
- [7] B. Ahlgren *et al.*, "Second NetInf Architecture Description," 4Ward EU FP7 Project, Tech.report D-6.2 v2.0, 2010.
- [8] S. E. Deering and D. R. Cheriton, "Multicast Routing in Datagram Internetworks and Extended LANs," *ACM Trans. Comput. Syst.*, vol. 8, no. 2, pp. 85–110, 1990.
- [9] K. Katsaros, G. Xylomenos, and G. C. Polyzos, "MultiCache: An Overlay Architecture for Information-centric Networking," *Comput. Netw.*, vol. 55, no. 4, pp. 936–947, Mar. 2011.
- [10] S. Meiling, D. Charousset, T. C. Schmidt, and M. Wählisch, "System-assisted Service Evolution for a Future Internet – The HAMcast Approach to Pervasive Multicast," in *Proc. of IEEE GLOBECOM 2010, Workshop MCS 2010*. Piscataway, NJ, USA: IEEE Press, Dec. 2010, pp. 913–917.
- [11] M. Wählisch, T. C. Schmidt, and S. Venaas, "A Common API for Transparent Hybrid Multicast," IRTF, IRTF Internet Draft – work in progress 04, January 2012. [Online]. Available: <http://tools.ietf.org/html/draft-irtf-samrg-common-api>
- [12] T. C. Schmidt and M. Wählisch, "Why We Shouldn't Forget Multicast in Name-oriented Publish/Subscribe," Open Archive: arXiv.org, Technical Report arXiv:1201.0349v1, 2012. [Online]. Available: <http://arxiv.org/abs/1201.0349v1>
- [13] D. Charousset, S. Meiling, T. C. Schmidt, and M. Wählisch, "A Middleware for Transparent Group Communication of Globally Distributed Actors," in *Middleware Posters 2011*. New York, USA: ACM, DL, Dec. 2011.
- [14] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Bulk of Interest: Performance Measurement of Content-Centric Routing," in *Proc. of ACM SIGCOMM*. New York: ACM, August 2012, SIGCOMM Poster – accepted for publication.
- [15] —, "Backscatter from the Data Plane — Threats to Stability and Security in Information-Centric Networking," Open Archive: arXiv.org, Technical Report arXiv:1205.4778v1, 2012. [Online]. Available: <http://arxiv.org/abs/1205.4778v1>

<sup>1</sup>See <http://www.realmv6.org/libcppa.html> for the latest release and documentation of the Actor library.