Bridging today's Internet Heterogeneity with a Content-oriented Approach

Klaus Wehrle

Protocol Engineering and Distributed Systems**
Wilhelm-Schickard-Institute for Computer Science
Universität Tübingen, Germany
`Klaus.Wehrle@uni-tuebingen.de`

1 Motivation

Despite the success of Internet technology and protocols, their deficiencies become apparent with growing and emerging applications like 3G mobile devices or pervasive environments. In these fields, there is a growing demand for additional support of mobility, multi- or anycast communication, or service composition.

The drawback of existing approaches to these problems like Mobile IP, IP-Multicast, and active networks is twofold: on the one hand, the necessary modifications to existing software systems, in particular operating systems, are substantial and have thus failed to fully disperse through to existing end-systems. On the other hand, central networking infrastructure like routers have to provide support for these techniques. Thus, both their introduction and the increased administration effort cause high costs for Internet service providers.

Another crucial disadvantage is the centralized management and administration of these approaches. They do not adapt well and often only with human intervention to rapid and significant changes in the infrastructure or user demands. As a promising alternative, the last few years have seen the rise of peer-to-peer (*P2P*) technology in which end hosts also form part of the routing and communication infrastructure itself. Its properties of decentralized self-organization, scalability, and fault-resilience make it a promising approach for the problems faced in highly mobile and fluctuating communication environments of personal and pervasive devices. Also in ad-hoc and ubiquitous scenarios, an immediate set-up of services without the need for central instances or provides may be a crucial advantage.

With already existing and emerging device classes and technologies, these environments become increasingly diverse. Thus, a central challenge to maintaining ubiquitous and uniform service accessibility in these environments is their increasing heterogeneity.

On the network layer, the gap between transmission characteristics is widening with high-speed fiber optical links on the one and slow and unreliable wireless technology on the other side of the spectrum which require individual protocol

---

** Parts of this project are funded by DAAD and Landesstiftung Baden-Württemberg.

Dagstuhl Seminar Proceedings 04411
Service Management and Self-Organization in IP-based Networks
http://drops.dagstuhl.de/opus/volltexte/2005/100

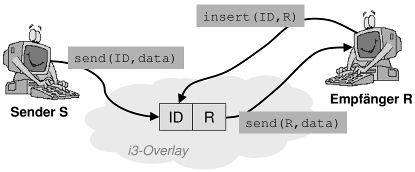

Fig. 1. Communication Paradigm of *i3*: Indirection

support. Furthermore, the widespread use of private networks, address translation, and firewalls subvert the idea of global connectivity IP is based on. This is particularly problematic in seamlessly integrating mobile devices into such fragmented networking environments and providing global access to services for the mobile end systems.

The diversity of applications and data formats make it hard to provide a uniform and consistent usage experience to users. Today's systems do not transcode data transparently to the required format nor do they provide adequate support for multi- or anycast communication becoming increasingly important e.g. in collaborative scenarios.

Heterogeneity on the device and end-system layer comes from the vast number of different hardware platforms and their capabilities. Also the differences in e.g. operating systems, user interfaces, or display facilities necessitate the adaptation of content to the individual properties and needs of devices and users.

Thus, the heterogeneity of today's Internet manifests in many aspects of infrastructure, implementation, and user interaction. In order to make services accessible to users seamlessly and independent of the devices they use, a communication platform is required that provides mobility support, group communication, and service composition.

2 Indirection as Fundamental Principle

The *Internet Indirection Infrastructure* (*i3*) strives to achieve these requirements as a generic distributed platform for implementing arbitrary forms of communication and services based on a decentralized self-organizing infrastructure. Its fundamental principle is to introduce a layer of indirection to the communication between peers. This principle is based on the observation that existing approaches like Mobile IP, IP Multicast, or IP Anycast all use a logical or physical point of indirection to provide their services. By abstracting this observation into a generic mechanism, *i3* supports highly flexible forms of communication and services.

In *i3*, every communication connection runs across at least one point of indirection. These indirection points are identified by globally unique IDs. Thus, *i3* moves away from addressing based on the endpoints of the communication

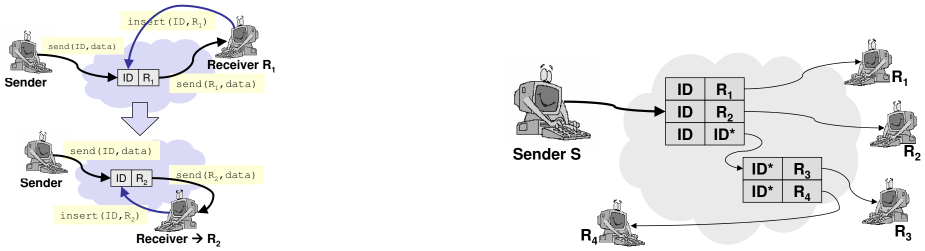

(a) Mobility Support (b) Multicast Support

Fig. 2. Realizing Communication Patterns with *i3*

partners to content-based addressing and data transfer using the IDs of the indirection points.

Senders augment their data with a destination ID to which the packets are routed by *i3*. Receivers insert tuples (*triggers*) of the form (ID, R) where R can be either the IP address and port of an end-host or another *i3*-ID. This information is used by *i3* to forward all data for ID to R, i.e. a receiving end-host or another indirection point.

The nodes participating in the *i3* infrastructure are organized in a distributed hash table (DHT) which provides decentralized and fault resilient routing services with the *i3* IDs acting as keys into the DHT. The IDs and their associated data are not implemented by a fixed node as the structure of the DHT can change dynamically. At the same time, the DHT layer guarantees scalability and reliability. The overhead introduced by the level of indirection can be mitigated for most packets by caching routing information and thus avoiding explicit routing via the DHT.

2.1 Realizing Mobility

Mobility is inherently supported by *i3* as end-hosts are no longer directly addressed. When the network address of a node changes due to a transition to a different physical network, it only needs to update its triggers with its new address. *i3* will then route packets to the end-host using this new address.

2.2 Realizing Multicast

Multicast can be achieved in a similarly straight-forward manner. In *i3*, every ID may be associated with multiple triggers and *i3* forwards data to the destinations in all triggers with the same ID. Thus every end-host interested in data for a specific ID simply inserts a trigger for this ID with its own address as the destination. In order to maintain scalability even for large numbers of end-hosts, custom and flexible multicast trees can be formed using additional triggers that move the cascading of data closer to the receivers.

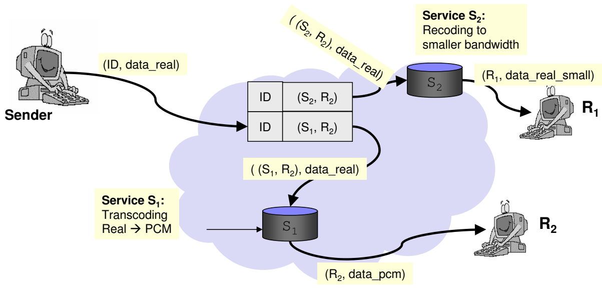

Fig. 3. Flexible Composition of Services with *i3*

2.3 Realizing Anycast

A slight variation of the scheme enables *i3* to support anycast: instead of using exact matching when comparing packet IDs with trigger IDs, partial matching, e.g. longest prefix match, is employed. Thus, a packet is forwarded according to the trigger best matching the packet ID. Different anycast strategies are modeled by choosing trigger IDs appropriately. For example, multiple trigger IDs with the same prefix but a random suffix result in randomized anycast. Of course, choosing trigger IDs based on certain metrics allows for optimizations with regard to e.g. link latency or reliability.

2.4 Realizing Service Composition

With an orthogonal extension to the trigger tuples, service composition can be achieved in *i3*: the destination part may be not only a single address or ID but a stack of them. Similar to IP source routing, this stack is attached to the packets for such a trigger resulting in the packets being routed across all items in the stack. With this approach, flexible chains of services to operate on data can be easily formed. Note how services can reuse other services transparently to the receiver and at the same time receivers can request customized service composition according to their specific requirements.

2.5 Integrating Applications without Modification

By providing a single point of indirection for arbitrary services, *i3* can help to alleviate the problem of restricted network accessibility of end-hosts due to firewalls and NAT routers. An *i3* proxy [10] was developed which can be run inside such a restricted network to give access to services inside and outside the network as a central and single point for authenticating and authorizing network access.

Another important functionality of the *i3* proxy is to integrate legacy applications with the *i3* infrastructure in order to avoid source code modifications.

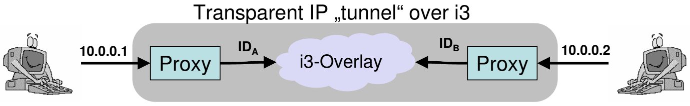

Fig. 4. Transparent Tunneling of IP over *i3*

Both existing client applications such as web browser and server applications such as web servers are thus enabled to transparently make use of the advantages of *i3*.

3 Summary

Shortcomings of the existing IP technology, e.g. lack for mobility and multicast, have in practice not been successfully addressed. With the emergence of small personal networked devices, these shortcomings negatively affect the connectivity and usability of these devices. *i3* tackles these problems as an overlay solution to existing protocols and infrastructure. It provides a homogeneous decentralized service platform in a scalable, reliable, and self-organizing manner. With indirection as its fundamental principle, the diversity of end-systems can be bridged and connectivity across private address space and firewalls can be ensured even in mobile scenarios.

References

1. Steve Deering. Host Extensions for IP Multicasting. RFC 1112, August 1989.
2. Internet Indirect Infrastructure, Project Homepage. http://i3.cs.berkeley.edu/.
3. K. Lakshminarayanan, A. Rao, I. Stoica, and S. Shenker. Flexible and robust large scale multicast using *i3*. Technical Report CS-02-1187, UC Berkeley, 2002.
4. C. Perkins. IP Mobility Support. RFC 2002, October 1996.
5. L. Peterson, T. Anderson, D. Culler, and T. Roscoe. A blueprint for introducing disruptive technology into the internet. In *Proc. of ACM HotNets-I, Princeton, USA*, October 2002.
6. R. Steinmetz and K. Wehrle. Peer-to-Peer-Networking und -Computing. *GI-Spektrum*, 27(1), February 2004.
7. I. Stoica, D. Adkins, S. Zhaung, et al. Internet Indirection Infrastructure. In *Proceedings of ACM SIGCOMM'02*, August 2002. Pittsburgh, PA.
8. I. Stoica, K. Lakshminarayanan, and K. Wehrle. Support for Service Composition in i3. In *Proceedings of ACM Multimedia*, October 2004. New York.
9. Ion Stoica, Robert Morris, David Karger, Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable Peer-To-Peer lookup service for internet applications. In *Proceedings of the 2001 ACM SIGCOMM Conference*, pages 149–160, 2001.
10. Klaus Wehrle, Jayanthkumar Kannan, Ayumu Kubota, Karthik Lakshminarayanan, and Ion Stoica. Supporting legacy applications over i3. Technical Report UCB/CS-04-1342, UC Berkeley, 2004.
11. S. Zhuang, K. Lai, I. Stoica, et al. Host Mobility Using an Internet Indirection Infrastructure. In *Proceedings of ACM MobiSys*, 2003.