

Implementation of a Common API for Transparent Hybrid Multicast - Design & Performance Aspects -

Matthias Wählisch, Thomas C. Schmidt

Stig Venaas

{waehlich, t.schmidt}@ieee.org,
stig@cisco.com

Agenda

1. Background: Draft common multicast API
2. Design Aspects
3. Performance Evaluation
4. Open issues
5. Conclusion

Current Multicast Diversity

Group communication services exist in a variety of:

- o Different flavors
 - Any Source Multicast vs. Source-specific Multicast
- o Different technologies
 - IPv4 vs. IPv6, multicast tunnels, etc.
- o Different layers
 - Native multicast vs. overlay distribution

Implications

- o Programmers decide on technology at coding time
 - How do they know about the multicast deployment state at this time?
 - o Applications provide their own solutions to overcome inter-domain deployment problem
 - Increases complexity & introduce redundancy
 - o Lack of efficiency
 - Reasonable to assume no global IP-layer multicast
- Difficult to write multicast applications that run everywhere & use most efficient group service

Recall: Common Multicast API Draft

Idea: Move complexity from application to the system level

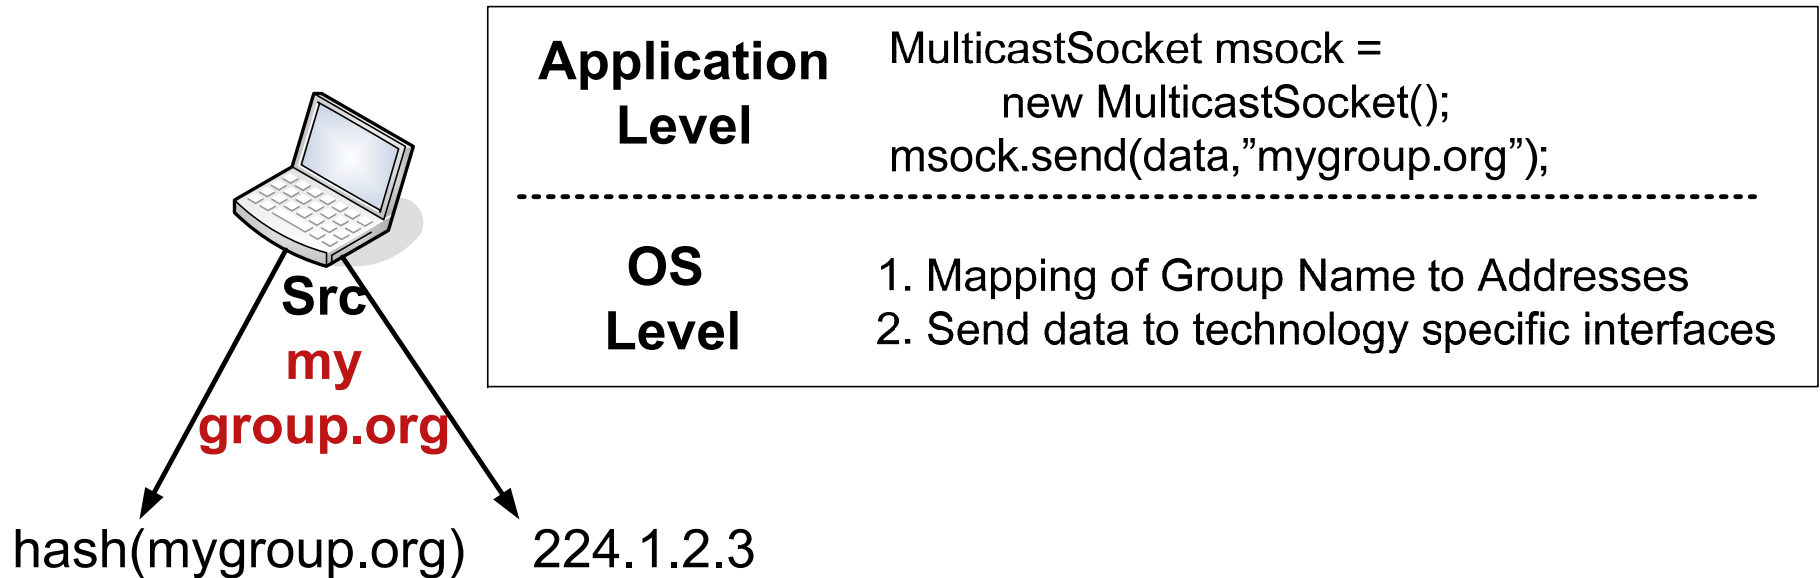
The current draft provides ...

- o a common multicast API on app. layer that abstracts group communication from distribution technologies
- o abstract naming and addressing by multicast URIs
- o mapping between naming and addressing
- o definition of protocol interaction to bridge multicast data between overlay and underlay

Status

- o Version 00/01 presented at IETF 76, Hiroshima
- o Update version 02 presented at IETF 77, Anaheim
 - Interesting work, but needs extended motivation
- o Current version: 04
 - Version 03 submitted before 78th IETF
- o Version 05: Working on an update, which also includes insights of the NBS BoF
 - Will be available some weeks after this IETF meeting

Common Multicast API & Middleware



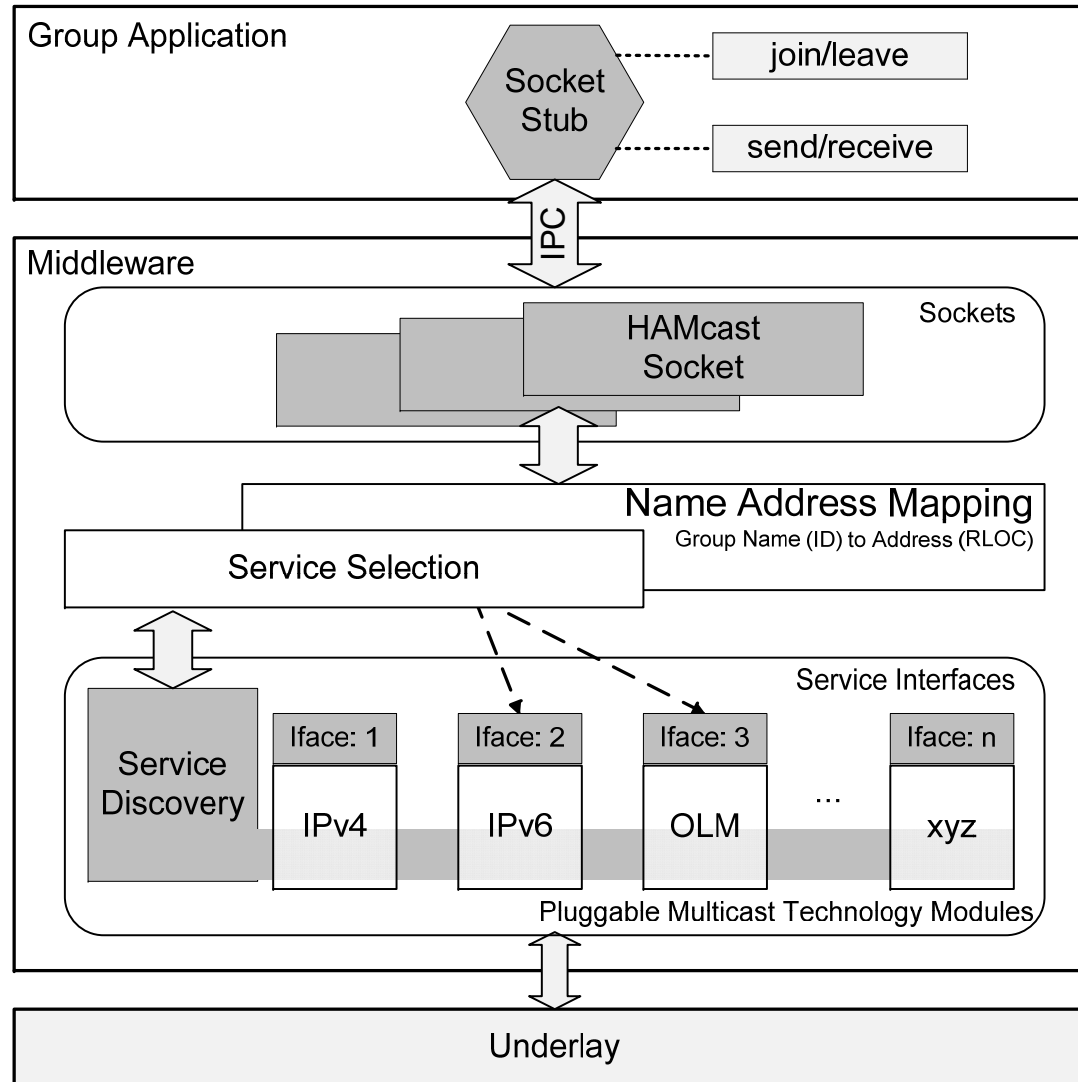
- o Mapping of names to technologies at run-time
- o Late binding

Implementation Report: The HAMcast stack

Design Aspects:

- o Flexibility for future extensions and adaption to new network technologies
 - No recompilation or changes to existing application using the HAMcast stack
- o Easy integration of new programming languages
 - Only mapping of API calls & 'signaling' to OS layer
- o Separate implementation of general multicast logic from technology-specific multicast instantiation

Overview HAMcast Stack



Middleware: Rough Overview

- o Unique daemon process instantiated once per host
- o Implemented in C++
- o Service functions realized by single modules
 - Can be loaded by the system if required
- o Technology-specific service interfaces
 - Implement a specific multicast understanding, e.g., IPv4/6, OLM, ...
 - Includes service discovery
- o Send/receive functions implemented as asynchronous calls

Module Example: Service Discovery

- o Implemented by each technology module
- o IP-layer multicast:
 - Discovery based on libcap sniffer library
 - Observation of general IGMP/MLD queries & PIM Hello
- o DHT-based multicast:
 - Unicast overlay is part of HAMcast stack
 - Inquires group and neighbor states
- o Gateways (e.g., AMT):
 - Future work

Prototype Status

- o First working version presented at EuroView'10
- o Hybrid distribution of streaming video using IPv4 + Scribe
- o Monitoring tool to discover and visualize hybrid multicast tree structure
 - Usage of API service calls

Performance Evaluation

Aim:

- o Basic evaluation of stack implementation
 - We do not focus on distribution technology
- o Comparison with native unicast

Setup:


- o Single source, single receiver scenario; Linux OS
- o Constant bit stream; 100 Mb/s Uplink
- o Different packet sizes
- o Average over several runs until convergence

Caveat


- o **Ongoing** evaluation & optimization
- o Performance results are a **first** view:
Optimization still to come

CPU Performance Send/Receive

Linux HAMcast middleware CPU time
Linux HAMcast application CPU time




Linux app+middleware CPU time
Linux native application CPU time




Send

Receive

Data Rate



Packet Rate



HAMcast stack at a Glance

- o Middleware implemented in C++
- o Programming libraries for C++ (and Java)
- o IPv4 and IPv6 module for native IP-layer multicast
- o Scribe implementation for overlay multicast
 - Based on Chimera P2P network stack
- o Service discovery for IPv4/v6 multicast & OLM
- o Prior to optimization: Moderate performance results

Open Issues

- o More detailed analysis of stack performance
- o Optimize implementation of HAMcast stack
- o Development of libraries for further programming languages
- o Service selection strategy:
 - Prefer most efficient technology
 - Do not distribute in parallel
- o SSM support

Thank you ...

- o Please, read the upcoming version 05
- o Interest in the implementation project?
- o **More feedback** is needed by RG members!