



PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

Socket Emulation over a Publish/Subscribe Network Architecture

Future Internet Symposium 2010

George Xylomenos, Blerim Cici
Mobile Multimedia Laboratory
Athens University of Economics and Business
Greece



Outline

- Motivation
 - PSIRP Implementation Concepts
 - Emulation Options
 - Mapping Addresses to Identifiers
 - Datagram Sockets
 - Stream Sockets
 - Performance
 - Conclusions and Future Work
-



Motivation

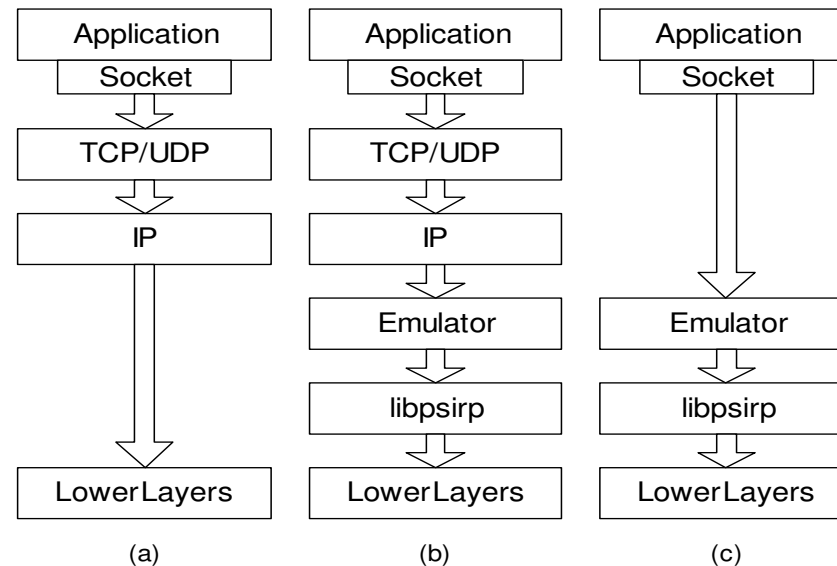
- Most Internet traffic is due to P2P content distribution
 - Users are interested on content, not peers
 - Redesign the Internet to make it information centric
 - Content Centric Networking (CCN)
 - Pub/Sub Internet Routing Paradigm (PSIRP)
 - Co-existence with the current Internet
 - Content distribution applications may be rewritten
 - End-to-end applications need a compatibility mode
 - The PSIRP Socket Emulator: sockets to pub/sub
-

PSIRP Implementation Concepts

- The PSIRP prototype implementation
 - Based on FreeBSD (considering Linux)
 - API provided by the libpsirp library
 - A publication is identified by a SId/RId pair
 - Publication life cycle
 - Psirp_create(): allocate memory for publication
 - Psirp_pub_data(): get a pointer to the memory
 - Psirp_pub(): make a publication version available
 - Psirp_subscribe_sync(): subscribe to publication
 - Psirp_free(): free publication memory
-
-



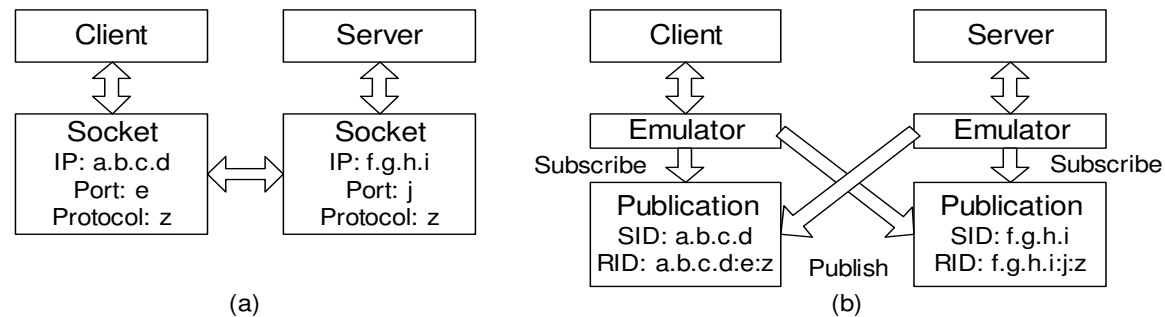
Emulation Options



- Option 1: network layer emulation
 - Translate IP packets into publications
- Option 2: transport layer emulation (preferred)
 - Translate socket calls to libpsirp calls



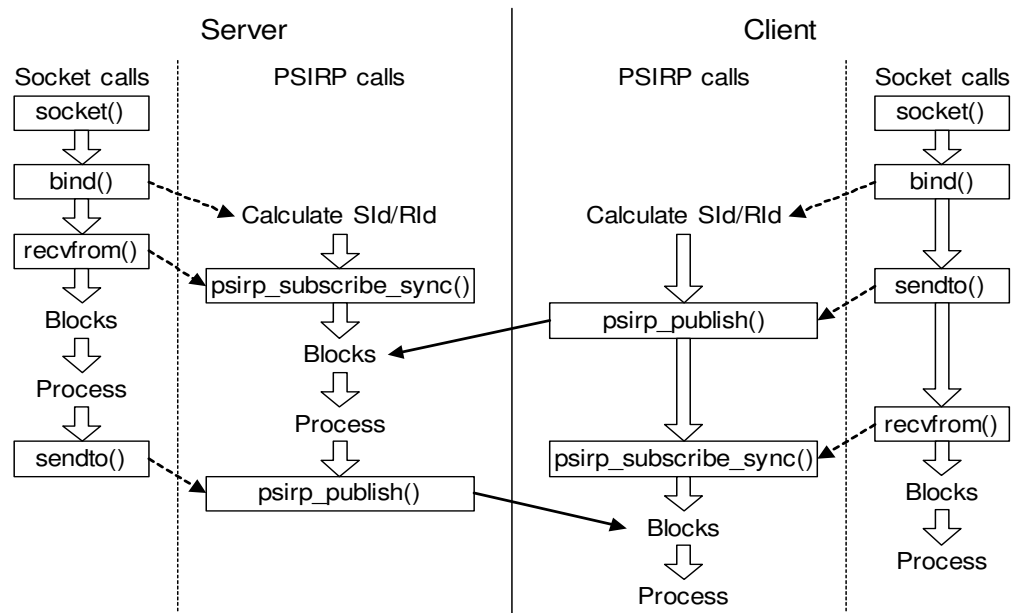
Mapping Addresses to Identifiers



- Sockets use endpoint identifiers
 - Local Port & IP, Protocol, Remote IP & Port
- Map each IP address to a separate SId
 - Publishing means sending data to a machine
- Map each (protocol,port) to a separate RId
 - Publishing means sending data to a (protocol,port)

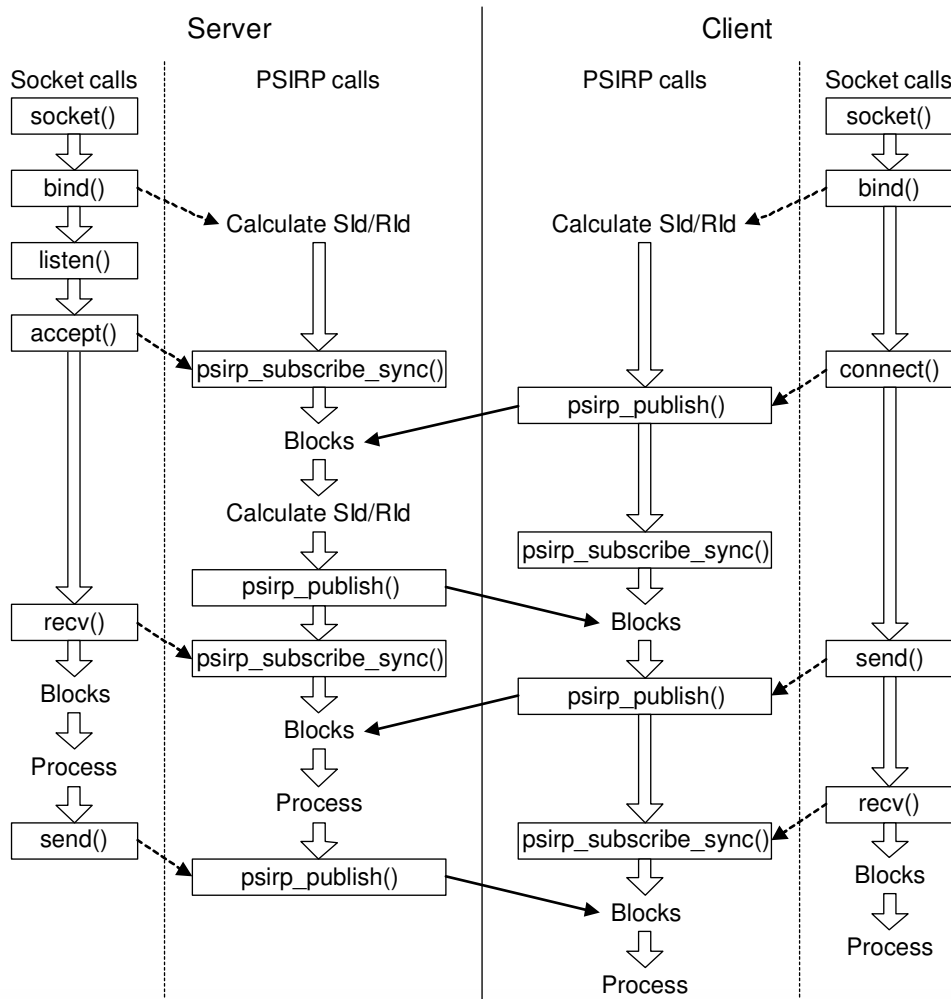


Datagram Sockets



- Each socket endpoint maps to a SId/RId pair
 - Publications to that pair represent UDP messages
 - New messages are mapped to new versions

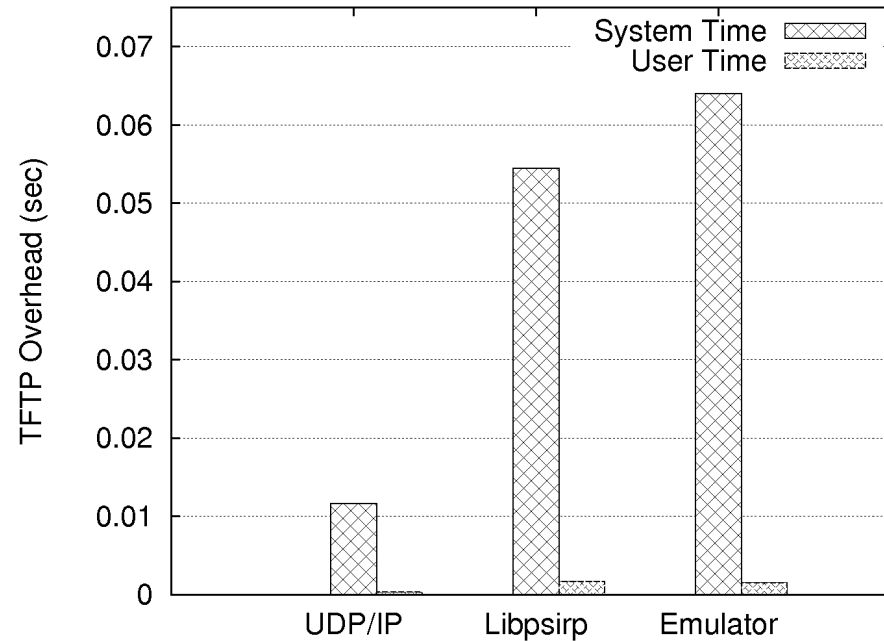
Stream Sockets



- Two SId/RId pairs!
 - Conn. requests
 - Data transfers
- Conn. requests
 - Same as Datagram
 - Half addresses
- Data transfers
 - Full addresses
 - Concurrent sessions



Performance



- Trivial File Transfer Protocol (TFTP)
 - Emulation is much slower than native UDP/IP
 - But not much slower than manual translation (17%)



Conclusions and Future Work

- Information centric Internet architectures
 - More suitable for content distribution
 - Fundamental departure from the current Internet
 - Sockets API emulator for PSIRP
 - Provides compatibility mode for existing applications
 - Similar performance to rewritten applications
 - Future work
 - Exploit reliable PSIRP transport for stream sockets
 - Reconsider the use of SId/RId pairs and versions
-