

Lecture 04 · SENG 41283 · Distributed and Cloud Computing

Interprocess Communication

Synchronous/Asynchronous · Sockets · UDP · TCP · Marshalling · Multicast

1. Introduction

UDP provides a message-passing abstraction: independent packets called datagrams. TCP provides a two-way stream with no message boundaries, the building block for producer-consumer communication.

Motivation	Type	Description
Tailored for application needs	Distributed hash tables	One of the most prominent classes of overlay network, offering a service that manages a mapping from keys to values across a potentially large number of nodes in a completely decentralized manner (similar to a standard hash table but in a networked environment).
	Peer-to-peer file sharing	Overlay structures that focus on constructing tailored addressing and routing mechanisms to support the cooperative discovery and use (for example, download) of files.
	Content distribution networks	Overlays that subsume a range of replication, caching and placement strategies to provide improved performance in terms of content delivery to web users; used for web acceleration and to offer the required real-time performance for video streaming [www.kontiki.com].
Tailored for network style	Wireless ad hoc networks	Network overlays that provide customized routing protocols for wireless ad hoc networks, including proactive schemes that effectively construct a routing topology on top of the underlying nodes and reactive schemes that establish routes on demand typically supported by flooding.
	Disruption-tolerant networks	Overlays designed to operate in hostile environments that suffer significant node or link failure and potentially high delays.
Offering additional features	Multicast	One of the earliest uses of overlay networks in the Internet, providing access to multicast services where multicast routers are not available; builds on the work by Van Jacobsen, Deering and Casner with their implementation of the Mbone (or Multicast Backbone) [mbone].
	Resilience	Overlay networks that seek an order of magnitude improvement in robustness and availability of Internet paths [nms.csail.mit.edu].
	Security	Overlay networks that offer enhanced security over the underlying IP network, including virtual private networks, for example, as discussed in Section 3.4.8.

Figure 1.1 Overview of interprocess communication: UDP datagrams, TCP streams, and socket API.

2. Characteristics of IPC

Synchronous and Asynchronous Communication

Synchronous: both send and receive are blocking. The sending process blocks until the corresponding receive is issued; the receiving process blocks until a message arrives.

Asynchronous: the send is non-blocking (process proceeds as soon as the message is copied to a local buffer); receive can be blocking or non-blocking.

Message Destinations

Messages are sent to (Internet address, local port) pairs. A local port has exactly one receiver but can have many senders. Servers publicise their port numbers for use by clients.

Reliability and Ordering

- **Reliable:** guaranteed delivery despite a reasonable number of dropped packets; messages arrive uncorrupted and without duplication.
- **Ordered:** some applications require messages delivered in sender order.

3. Sockets

Both UDP and TCP use the socket abstraction as an endpoint for communication. A process must bind its socket to a local port and Internet address before receiving. Each computer has $2^{16} = 65,536$ possible port numbers.

4. UDP Datagram Communication

Transmitted without acknowledgement or retries; if a failure occurs the message may not arrive. Non-blocking sends; blocking receives (unless a timeout is set).

- **Message size:** most environments impose an 8 KB limit; larger messages must be fragmented.
- **Timeouts:** set on sockets to prevent indefinite blocking.
- **Failure model:** omission failures (dropped packets) and ordering failures.
- **Use cases:** DNS lookups · Voice over IP (VoIP).

5. TCP Stream Communication

A two-way stream between pairs of processes. A connection is established before communication. TCP handles: lost messages (ACK + retransmit), flow control (blocking the writer when the reader is too slow), duplication and ordering (sequence numbers), and provides integrity and validity guarantees.

- **Use cases:** HTTP (80/443) · FTP (22) · Telnet (23) · SMTP (25/587).

6. External Data Representation and Marshalling

Data structures must be flattened to byte sequences before transmission. Different computers may store integers differently (big-endian vs little-endian). **Marshalling:** assembling data items into a form suitable for transmission. **Unmarshalling:** disassembling received bytes at the destination.

- **CORBA CDR:** external representation for structured and primitive types for CORBA RMI.
- **Java Object Serialisation:** flattening of any object or tree of objects; Java-only.
- **XML:** textual format for structured data; used in web services messages.

7. Multicast Communication

A single operation sends a message from one process to each member of a group. IP multicast is available only via UDP. Class D addresses (224.0.0.0–239.255.255.255). Uses: fault tolerance via replicated services; service discovery; replicated data; event notifications.

Exercise

Compare and contrast CORBA CDR, Java Object Serialisation, and XML. Write a descriptive blog article on External Data Representation and Marshalling.