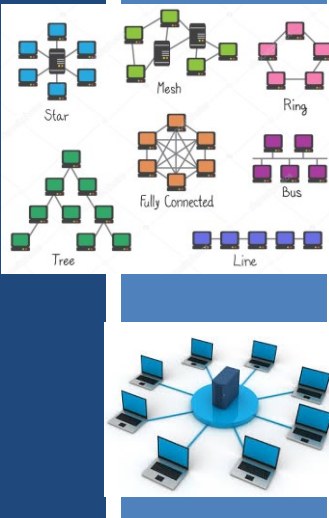# TCSS 558:
APPLIED DISTRIBUTED COMPUTING

Ch. 4 - Communication

Wes J. Lloyd
School of Engineering
& Technology (SET)
University of Washington - Tacoma

1

OBJECTIVES – 2/16

- **Questions from 2/14**
- Midterm Review
- Assignment 1: Key Value Store
- Chapter 4: Communication
  - Chapter 4.1: Foundations
  - Chapter 4.2: RPC (light-review)
  - Chapter 4.3: Message Oriented Communication

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.2

2

ONLINE DAILY FEEDBACK SURVEY

- Daily Feedback Quiz in Canvas – Available After Each Class
- Extra credit available for completing surveys *ON TIME*
- Tuesday surveys: due by ~ Wed @ 10p
- Thursday surveys: due ~ Mon @ 10p

TCSS 558 A › Assignments

Winter 2021

Home

Announcements

Assignments

Zoom

Chat

Search for Assignment

▼ Upcoming Assignments

🚀 TCSS 558 - Online Daily Feedback Survey - 1/5
Not available until Jan 5 at 1:30pm | Due Jan 6 at 10pm | -/1 pts

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.3

3

TCSS 558 - Online Daily Feedback Survey - 1/5

Due Jan 6 at 10pm   Points 1   Questions 4
Available Jan 5 at 1:30pm - Jan 6 at 11:59pm 1 day   Time Limit None

Question 1	0.5 pts

On a scale of 1 to 10, please classify your perspective on material covered in today's class:

1	2	3	4	5	6	7	8	9	10

Mostly
Review To Me

Equal
New and Review

Mostly
New to Me

Question 2	0.5 pts

Please rate the pace of today's class:

1	2	3	4	5	6	7	8	9	10

Slow

Just Right

Fast

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.4

4

MATERIAL / PACE

- Please classify your perspective on material covered in today's class (25 respondents):
- 1-mostly review, 5-equal new/review, 10-mostly new
- **Average – 6.70** (↑ - *previous 6.14*)

- Please rate the pace of today's class:
- 1-slow, 5-just right, 10-fast
- **Average – 5.61** (↓ - *previous 5.84*)

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.5

5

TENURE TRACK FACULTY CANDIDATE RESEARCH SEMINARS – EXTRA CREDIT

- Tuesday February 21 – 12:30pm – Cherry Parkes room TBA
- Thursday February 23 – 12:30pm – Cherry Parkes room TBA
- Wednesday March 1 – 12:30pm – Cherry Parkes room TBA
- Friday March 3 – 12:30pm – Cherry Parkes room TBA

- Earn up to 2.5% extra credit added to the overall course grade
- Scored out of 5 total points
- First seminar – earn 2 points
- 2nd, 3rd, 4th seminar – earn 1 point each

Add to final course grade:
(Total_seminar_points / 5 points) * 2.5%

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.6

6

FEEDBACK FROM 2/14

- *About DNS example on slide page 21. I want to confirm if my understanding is correct:*
- *My understanding is that it means it takes 22.458ms to ping www.google.com when I connect to UW wirelessly, and takes 1.278ms to ping google in EC2 instances (equivalent to a virtual machine), right? Is this because the virtual machine (us-east-1) is located close to the google server ...?*
- YES, this is correct
- The ec2 instance is in us-east-1 Virginia region
- The ec2 instance accesses a local google server which is fetched using a DNS server lookup
- The ec2 instance was seen to have ~17.5x less network latency than the WA laptop using a wireless network

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.7

7

FEEDBACK - 2

- *What does application agnostic vs application specific approach mean?*
- *Application agnostic* refers to designs or solutions to a problem that can work for any/every application
- They are general/generic solutions that are broadly applicable
- These solutions can be delivered via middleware because no application specific attributes are leveraged for the solution
- They may not be as performant, or they may feature overhead compared to application specific solutions
- *Application specific* refers to designs or solutions to a problem that only works in the context of a specific application
- The solution leverages specific aspects of the application to solve the problem
- Other applications may not be able to use the solution

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.8

8

OBJECTIVES – 2/16

- Questions from 2/14
- **Midterm Review**
- Assignment 1: Key Value Store
- Chapter 4: Communication
  - Chapter 4.1: Foundations
  - Chapter 4.2: RPC (light-review)
  - Chapter 4.3: Message Oriented Communication

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.9 |

9

WE WILL RETURN AT 2:58 PM

13

OBJECTIVES – 2/16

- Questions from 2/14
- Midterm Review
- **Assignment 1: Key Value Store**
- Chapter 4: Communication
  - Chapter 4.1: Foundations
  - Chapter 4.2: RPC (light-review)
  - Chapter 4.3: Message Oriented Communication

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.14 |

14

ASSIGNMENT 1
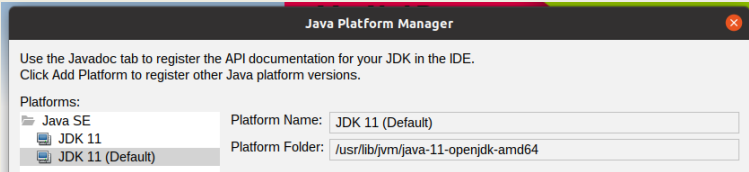
- Team signup posted on Canvas under 'People'

- GenericNode.tar.gz includes Dockerfile examples
- GenericNode.tar.gz assumes Java 11

- TCP/UDP/RMI Key Value Store

- Implement a "GenericNode" project which assumes the role of a client or server for a Key/Value Store

- Recommended in Java 11 LTS

- Client node program interacts with server node to put, get, delete, or list items in a key/value store

| February 14, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L12.15 |

15

USING JAVA 11 IN NETBEANS

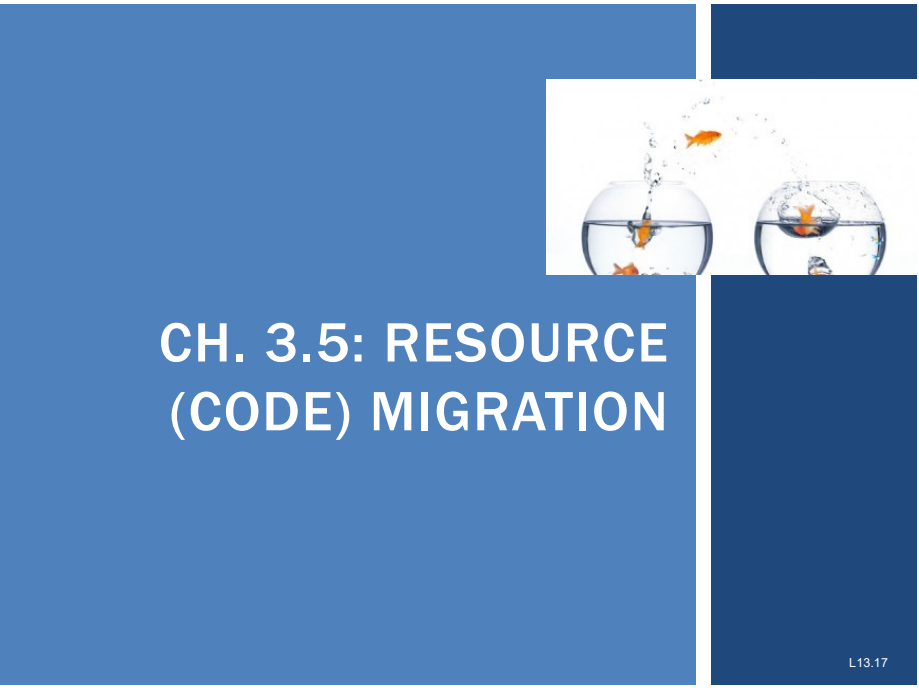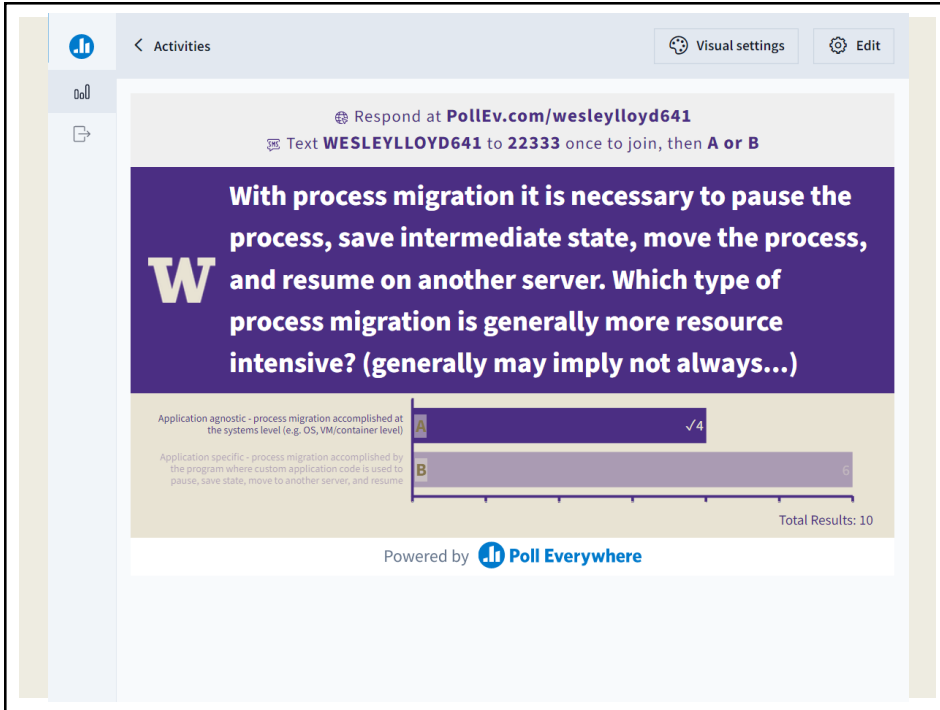- In Netbeans IDE, under Tools menu, 'Java Platforms', be sure to install and select JDK 11

Java Platform Manager

Use the Javadoc tab to register the API documentation for your JDK in the IDE.
Click Add Platform to register other Java platform versions.

Platforms:

Java SE
  JDK 11
  JDK 11 (Default)

Platform Name: JDK 11 (Default)

Platform Folder: /usr/lib/jvm/java-11-openjdk-amd64

- On left-hand Project menu, right-click on 'GenericNode' project
- Select Properties
- Under Build | Compile, be sure Java Platform is JDK 11
- Under Sources, be sure Source/Binary Format is 11

February 14, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L12.16

16

CH. 3.5: RESOURCE (CODE) MIGRATION

L13.17

17

18

OBJECTIVES – 2/16

- Questions from 2/14
- Midterm Review
- Assignment 1: Key Value Store
- Chapter 4: Communication
  - Chapter 4.1: Foundations
  - Chapter 4.2: RPC (light-review)
  - Chapter 4.3: Message Oriented Communication

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.19 |

19

CH. 4 COMMUNICATION

L13.20

20

CHAPTER 4

- **4.1 Foundations**
  - **Protocols**
  - **Types of communication**
- **4.2 Remote procedure call**
- **4.3 Message-oriented communication**
  - **Socket communication**
  - **Messaging libraries**
  - **Message-Passing Interface (MPI)**
  - **Message-queueing systems**
  - **Examples**
- **4.4 Multicast communication**
  - **Flooding-based multicasting**
  - **Gossip-based data dissemination**
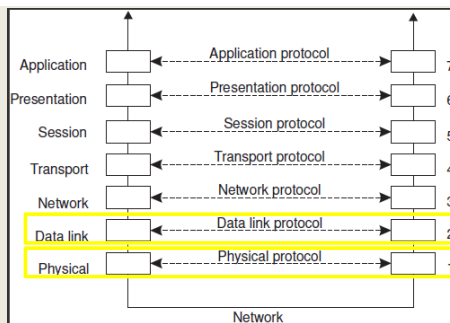
Reviews and builds on content from Ch. 2/3

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.21

21

CH. 4.1: FOUNDATIONS

L13.22

22

LAYERED PROTOCOLS

- Distributed systems lack shared memory
- All distributed system communication
  is based on sending and receiving low-level messages
  - P → Q

- **O**pen **S**ystems **I**nterconnection Reference Model
  (OSI Model)
  - Open systems communicate with any other open system
  - Standards govern format, contents, meaning of messages
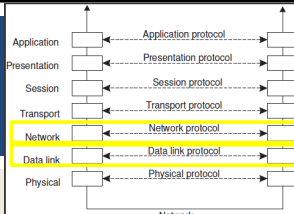 - Formalization of rules forms a **communication protocol**

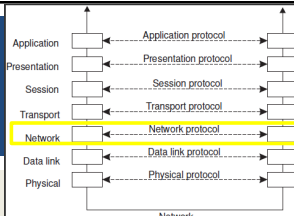| February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.23

23

LAYERED PROTOCOLS - 2

- Protocols provide a **communication service**

- **Two service types:**

  - **Connection-oriented**: sender/receiver establish connection, negotiate parameters of the protocol, close connection when done
  - Physical example: telephone

  - **Connectionless**: No setup. Sender sends. Receiver receives.
  - Physical example: Mailing a letter

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.24

24

OSI MODEL REVISITED

- Physical layer: just sends bits → ... 0 0 0 1 0 1 1 0 1 1 ...
- Data link layer: Groups bits into frames
  - Provides error correction via *checksum*
  - Special bit pattern at start/end of frame

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.25

25

OSI MODEL - 2

- Data link layer:
  - **Checksum**: computed by adding all bytes in frame in particular way
  - Added to message
  - Receiver removes checksum, recomputes checksum, and compares
  - If receiver and sender agree, frame is considered correct
  - Receiver can request failed frames to be resent
  - Frames assigned sequence numbers *in the header*
- Network layer:
  - Sometimes referred to as the *Internet layer*
  - On WANs sending msgs between client/server requires routing
  - Provides addressing using IPV4 (32-bit), IPV6 (64-bit)

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.26 |

26

OSI MODEL - 3

- Network layer:
  - Helps with routing network traffic
  - Shortest route (# of hops) may not be the best route
  - Minimizing delay (latency) is paramount
  - Routing algorithms: use long-term average network conditions, or try to adapt to changing conditions
  - ICMP Protocol: Internet Control Message Protocol
  - Not typically used for sending data, but for diagnostic/control purposes
  - ICMP Examples: (*ping*, *traceroute*)

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.27 |

27

OSI MODEL - 4

- **Internet Control Message Protocol (ICMP)**
  - **8 bytes header: 4 fixed, 4 variable**

ICMP Header Format

Offsets	Octet	0								1								2								3						
Octet	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Type								Code								Checksum															
4	32	Rest of Header																															

- Example message types:
- 0- echo reply (**PING**), 3- destination unreachable, 4- source quench (congestion control), 5- redirect message, 8- echo request (**PING**), 9- router advertisement
- Others: 10 (router solicitation), 11 (time exceeded), 12 (parameter problem), 13 (timestamp), 15 (info request), 16 (info reply), 17 (address mask request), 18 (address mask reply), 30-39 (**traceroute**), 40 (security failures), 42 (ext echo request)…255

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.28

28

OSI MODEL - 5

- **Transport layer:**
  - **Provides reliable connections**
  - **Reorganizes packets arriving out of sequence**
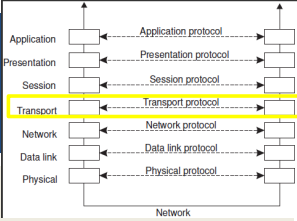 - **Requests delivery of missing packets**

1. **Breaks application layer protocol messages into pieces to transmit**
2. **Assigns messages sequence numbers**
3. **Sends all messages**

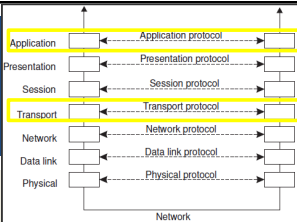| February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.29

29

OSI MODEL - 6

- Transport layer provides an infallible "message pipe"
  - Put messages in
  - Always come out undamaged, in correct order

- Transport layer protocols:
  - TCP: Transmission Control Protocol (connection-oriented)
  - UDP: Universal Datagram Protocol (connectionless)
    - Point to point messaging
      - Send a message to a point
    - No session/connection

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.30

30

OSI MODEL - 7

- Other transport protocols
  - Real-time transport protocol (RTP): real-time data, no data delivery guarantee
  - Streaming Control Transmission Protocol (SCTP): alternative to TCP

- Higher-level protocols:
- **Session layer**: mechanisms for opening, closing, managing session between communicating processes
- **Presentation layer**: deals with syntactical meaning of messages
  - Presentation services convert data among formats, for example:
  - from extended binary coded decimal interchange code (EBCDIC) to ASCII
- **Application layer**: protocols that don't fit into other layers
  - Many protocols: FTP, SFTP, HTTP, etc. etc.

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.31

31

OSI MODEL - 8

- **Each OSI layer contributes overhead bits to the message**

- **Layers append data to front (and maybe end) of the message**

- **Receiver strips off headers as the message goes up the OSI model stack:**

  physical → data-link → network → transport → application

32

PROTOCOL STACK

- **Collection of layers used for communication from OSI model**

33

MIDDLEWARE PROTOCOLS

- Middleware is reused by many applications
- Protocols provided more broadly by middleware
- Provide needed functions applications are built and depend upon
  - For example: communication frameworks/libraries
- Middleware offer many general-purpose protocols

- Middleware protocol examples:
  - <u>Authentication protocols</u>: supports granting users and processes access to authorized resources
  - Doesn't fit as an "application specific" protocol
  - Considered a "Middleware protocol"

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.34

34

MIDDLEWARE PROTOCOLS - 2

- <u>Distributed commit protocols</u>
  - Coordinate a group of processes (nodes)
  - Facilitate all nodes carrying out a particular operation
  - Or abort transaction
  - Provides distributed atomicity (all-or-nothing) operations
- <u>Distributed locking protocols</u>
  - Protect a resource from simultaneous access from multiple nodes
- <u>Remote procedure call</u>
  - One of the oldest middleware protocols

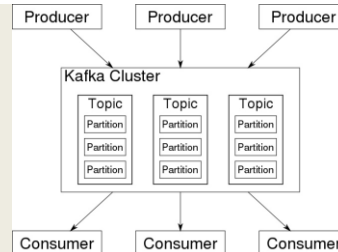| February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.35

35

MIDDLEWARE PROTOCOLS - 3

- **Message queueing services**
  - **Support synchronization of data streams**
  - **Transfer real-time data**
  - **Distributed and scalable implementation**

- **Multicast services**
  - **Scale communication to thousands of receivers spread across the Internet**

36

MIDDLEWARE PROTOCOLS - 3

- **Message queueing services**

KEY: middleware protocols offer functionality to satisfy the software requirements of *many* applications

Middleware functions are general, application-independent in nature

Functions are so commonly needed they are offered in reusable frameworks / libraries
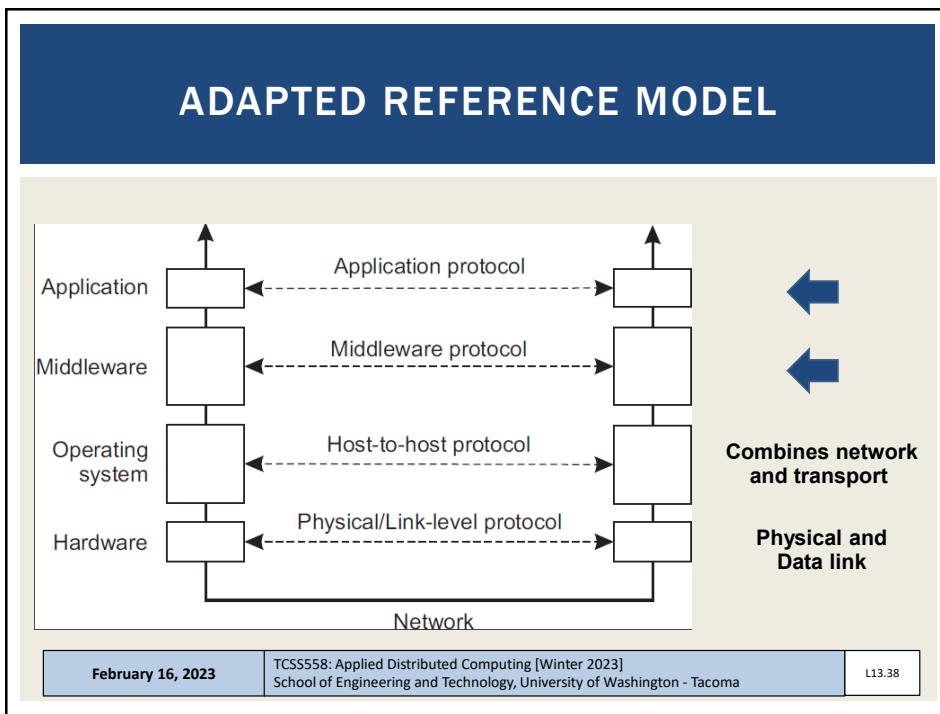
37

ADAPTED REFERENCE MODEL

Application	← Application protocol →	
Middleware	← Middleware protocol →	
Operating system	← Host-to-host protocol →	**Combines network and transport**
Hardware	← Physical/Link-level protocol →	**Physical and Data link**

Network

38

TYPES OF COMMUNICATION

- Persistent communication
  - Message submitted for transmission is stored by communication middleware as long as it takes to deliver it
  - Example: email system (SMTP)
  - Receiver can be offline when message sent
  - Temporal decoupling (delayed message delivery)

- Transient communication
  - Message stored by middleware only as long as sender/receiver applications are running
  - If recipient is not active, message is dropped
  - Transport level protocols typically are transient (*no msg storage*)

- **What OSI protocol level is the SMTP Protocol?**

39

TYPES OF COMMUNICATION - 2

- Asynchronous communication
  - Client does not block, continues doing other work
- Synchronous communication
  - Client blocks and waits
- Three types of blocking (*synchronous*)
  1. SHORTEST:
     Until middleware notifies it will take over delivering ***request***
  2. MEDIUM:
     Sender may block until ***request*** has been delivered
  3. LONGEST:
     Sender waits until ***request*** is processed and result is returned

- Persistence + synchronization (blocking)
  - Common scheme for message-queueing systems
  - ***Publish message to queue***: block until message delivered to queue

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.40

40

41

OBJECTIVES – 2/16

- Questions from 2/14
- Midterm Review
- Assignment 1: Key Value Store
- Chapter 4: Communication
  - Chapter 4.1: Foundations
  - Chapter 4.2: RPC (light-review)
  - Chapter 4.3: Message Oriented Communication

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.42

42

CH. 4.2: RPC (LIGHT-REVIEW)

L13.43

43

RPC – REMOTE PROCEDURE CALL

- In a nutshell,
- Allow programs to call procedures on other machines
- Process on __machine A__ calls procedure on __machine B__
- Calling process on __machine A__ is suspended
- Execution of the called procedure takes place on __machine B__
- Data transported from caller __(A)__ to provider __(B)__ and back __(A)__.
- No message passing is visible to the programmer
- __Distribution transparency__: make remote procedure call look like a local one
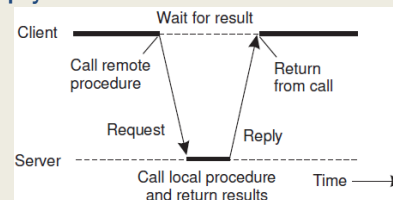- `newlist = append(data, dbList)`

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.44 |

44

RPC - 2

- Transparency enabled with client and server "stubs"
- Client has "stub" implementation of the server-side function
- Interface exactly same as server side
- But client __DOES NOT HAVE THE IMPLEMENTATION__
- __Client stub__: packs parameters into message, sends *request* to server. Call blocks and waits for reply
- __Server stub__: transforms incoming *request* into local procedure call
- Blocks to wait for *reply*
- Server stub unpacks *request*, calls server procedure
- *It's as if the routine were called locally*

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.45 |

45

RPC - 3

- Server packs procedure *results* and sends back to client.

- Client "*request*" call unblocks and data is unpacked

- Client can't tell method was called remotely over the network... *except for network latency...*

- Call abstraction enables clients to invoke functions in alternate languages, on different machines

- Differences are handled by the RPC "framework"

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.46 |

46

RPC STEPS

1. Client procedure calls client stub
2. Client stub builds message and calls OS
3. Client's OS send message to remote OS
4. Server OS gives message to server stub
5. Server stub unpacks parameters, calls server
6. Server performs work, returns results to server-side stub
7. Server stub packs results in messages, calls server OS
8. Server OS sends message to client's OS
9. Client's OS delivers message to client stub
10. Client stub unpacks result, returns to client

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.47 |

47

PARAMETER PASSING

- **STUBS**: take parameters, pack into a message, send across network

- Parameter marshaling:
- `newlist = append(data, dbList)`
- Two parameters must be sent over network and correctly interpreted

- Message is transferred as a series of bytes
- Data is serialized into a "stream" of bytes
- Must understand how to unmarshal (unserialize) data

- Processor architectures vary with how bytes are numbered: Intel (right→left), older ARM (left→right)

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.48

48

RPC: BYTE ORDERING

- Big-Endian: write bytes left to right (ARM)

- Little-endian: write bytes right to left (Intel)

- Networks: typically transfer data in Big-Endian form

- Solution: transform data to machine/network independent format

- Marshaling/unmarshaling: transform data to neutral format

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.49

49

RPC: PASS-BY-REFERENCE

- Passing by value is straightforward
- Passing by reference is challenging
- Pointers only make sense on local machine owning the data
- Memory space of client and server are different

- Solutions to **RPC pass-by-reference**:
1. Forbid pointers altogether
2. Replace pass-by-reference with pass-by-value
   - Requires transferring entire object/array data over network
   - **Read-only optimization**: don't return data if unchanged on server
3. Passing global references
   - Example: file handle to file accessible by client and server via shared file system

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.50 |

50

RPC: DEVELOPMENT SUPPORT

- Let developer specify which routines will be called remotely
   - Automate client/server side stub generation for these routines

- Embed remote procedure call mechanism into the programming language
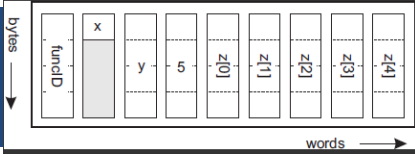   - E.g. Java RMI

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.51 |

51

STUB GENERATION

- `void func(char x; float y; int z[5])`
- 1-byte character transmits with 3-padded bytes
- Float sent as whole word (4-bytes)
  - Array as group of words, proceed by word describing length
  - Client stub must package data in specific format
  - Server stub must receive and unpackage in specific format

- Client and server must agree on representation of simple data structures: int, char, floats w/ little endian
- RPC clients/servers: must agree on protocol
  - TCP? UDP?

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.52 |

52

STUB GENERATION - 2

- Interfaces are specified using an Interface Definition Language (IDL)

- Interface specifications in IDL are used to generate language specific stubs

- IDL is compiled into client and server-side stubs

- Much of the plumbing for RPC involves maintaining boilerplate-code

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.53 |

53

LANGUAGE BASED SUPPORT

- Leads to simpler application development

- Helps with providing access transparency
  - Differences in data representation, and how object is accessed
  - Inter-language parameter passing issues resolved:
    → *just 1 language*

- Well known example: *Java Remote Method Invocation* RPC equivalent embedded in Java

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.54 |

54

RPC VARIATIONS

- RPC: client typically blocks until reply is returned
- Strict blocking *unnecessary* when there is no result

- Asynchronous RPCs
  - When no result, server can immediately send reply

Client/server synchronous RPC

Client   Wait for result

Call remote   Return
procedure   from call

Request   Reply

Server   Call local procedure   Time →
          and return results

Client/server asynchronous RPC

Client   Wait for acceptance

Call remote   Return
procedure   from call

Request   Accept request

Server   Call local procedure   Time →

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.55 |

55

RPC VARIATIONS – 2

- What are tradeoffs for synchronous vs. asynchronous procedure calls?
  - For a local program
  - For a distributed program (system)

- Use cases for asynchronous procedure calls
  - Long running jobs allow client to perform alternate work in background (in parallel)
  - Client may need to make multiple service calls to multiple server backends at the same time...
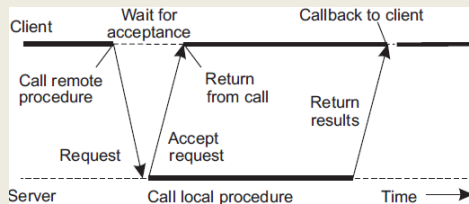
| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.56 |

56

TYPES OF ASYNCHRONOUS RPC

- **Deferred synchronous RPC**
  - Server performs **_CALLBACK_** to client
  - Client, upon making call, spawns separate thread which blocks and waits for call

- **One-way RPCs**
  - Client **does not wait** for _any_ server acknowledgement – it just goes...

- **Client polling**
  - Client (_using separate thread_) continually polls server for result
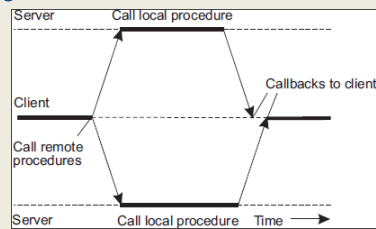
| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.57 |

57

MULTICAST RPC

- Send RPC request *simultaneously* to group of servers
- Hide that multiple servers are involved
- Consideration:
  Does the client need all results or just one?
- Use cases:
  - Fault tolerance – wait for just one
  - Replicate execution – verify results, *use first result*
  - Divide and conquer - multiple RPC calls work in parallel on different parts of dataset, client aggregates results

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.58 |

58

RPC EXAMPLE:
DISTRIBUTED COMPUTING ENVIRONMENT (DCE)

- **DCE**: basis for Microsoft's distributed computing object model (DCOM)
- Used in Samba, ***cross-platform*** file and print sharing via RPC
- Middleware system – provides layer of abstraction between OS and distributed applications
- Designed for Unix, ported to ***all*** major operating systems
- Install DCE middleware on set of heterogeneous machines – distributed applications can then access shared resources to:
  - Mount a windows file system on Linux
  - Share a printer connected to a Windows server
- Uses client/server model
- All communication via RPC
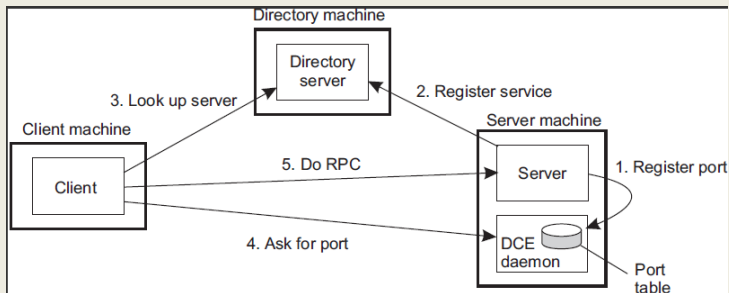- DCE daemon tracks participating machines, ports

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.59 |

59

DCE CLIENT-TO-SERVER BINDING

- **Server name comes from directory server**
- **Server port comes from DCE daemon**
  - **DCE daemon has a well known port # client already knows**

60

EXTRA: DCE – CLIENT/SERVER DEVELOPMENT

1. **Create Interface definition language (IDL) files**
   - **IDL files contain Globally unique identifier (GUID)**
   - **GUIDs must match: client and server compare GUIDs to verify proper versions of the distributed object**
   - **128-bit binary number**

2. **Next, add names of remote procs and params to IDL**

3. **Then compile the IDL files**
   Compiler generates:
   - **Header file (interface.h in C)**
   - **Client stub**
   - **Server stub**

61

EXTRA: DCE – BINDING CLIENT TO SERVER

- For a client to call a server, server must be registered
  - *Java: uses RMI registry*
- Client process to search for RMI server:
  1. Locate the server's host machine
  2. Locate the server (i.e. process) on the host
- Client must discover the server's RPC port

- **DCE daemon:** maintains table of (server,port) pairs

- When servers boot:
1. Server asks OS for a port, registers port with DCE daemon
2. Also, server registers with directory server, separate server that tracks DCE servers

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.62 |

62

OBJECTIVES – 2/16

- Questions from 2/14
- Midterm Review
- Assignment 1: Key Value Store
- Chapter 4: Communication
  - Chapter 4.1: Foundations
  - Chapter 4.2: RPC (light-review)
  - **Chapter 4.3: Message Oriented Communication**

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.63 |

63

CH. 4.3: MESSAGE-ORIENTED COMMUNICATION

Apache ActiveMQ

L13.64

64

MESSAGE ORIENTED COMMUNICATION

- RPC assumes that the _client_ and _server_ are running **at the same time…** (temporally coupled)
- RPC communication is typically **synchronous**

- When client and server are not running at the same time
- Or when communications should not be **blocked…**

- **This is a use case for message-oriented communication**
  - Synchronous vs. asynchronous
  - Messaging systems
  - Message-queueing systems

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.65

65

SOCKETS

- Communication end point
- Applications can read / write data to
- Analogous to file streams for I/O, but *network streams*

Operation	Description
socket	Create a new communication end point
bind	Attach local address to socket (IP / port)
listen	Tell OS what max # of pending connection requests should be
accept	Block caller until a connection request arrives
connect	Actively attempt to establish a connection
send	Send some data over the connection
receive	Receive some data over the connection
close	Release the connection

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.66

66

SOCKETS - 2

- Servers execute 1st - 4 operations (socket, bind, listen, accept)
- Methods refer to C API functions
- Mappings across different libraries will vary (*e.g. Java*)

Operation	Description
socket	Create a new communication end point
bind	Attach local address to socket (IP / port)
listen	Tell OS what max # of pending connection requests should be
accept	Block caller until a connection request arrives
connect	Actively attempt to establish a connection
send	Send some data over the connection
receive	Receive some data over the connection
close	Release the connection

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.67

67

SERVER SOCKET OPERATIONS

- **Socket**: creates new communication end point

- **Bind**: associated IP and port with end point

- **Listen**: for connection-oriented communication, non-blocking call reserves buffers for specified number of pending connection requests server is willing to accept

- **Accept**: blocks until connection request arrives
  - Upon arrival, new socket is created matching original
  - Server spawns thread, or forks process to service incoming request
  - Server continues to wait for new connections on original socket
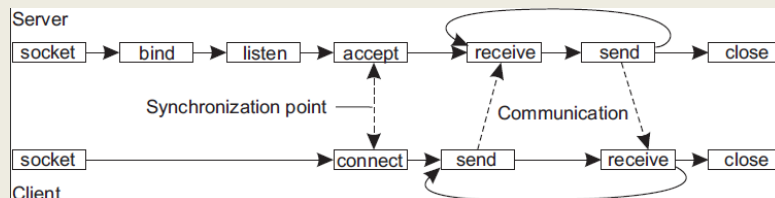
| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.68 |

68

CLIENT SOCKET OPERATIONS

- **Socket**: Creates socket client uses for communication
- **Connect**: Server transport-level address provided, client blocks until connection established
- **Send**: Supports sending data (to: server/client)
- **Receive**: Supports receiving data (from: server/client)
- **Close**: Closes communication channel
  - Analogous to closing a file stream

Server
socket → bind → listen → accept → receive → send → close

Synchronization point   Communication

socket → connect → send → receive → close
Client

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.69 |

69

SOCKET COMMUNICATION

- Sockets provide primitives for implementing your own TCP/UDP communication protocols

- Directly using sockets for transient (non-persisted) messaging is very basic, can be brittle
  - Easy to make mistakes...

- Any extra communication facilities must be implemented by the application developer

- More advanced approaches are desirable
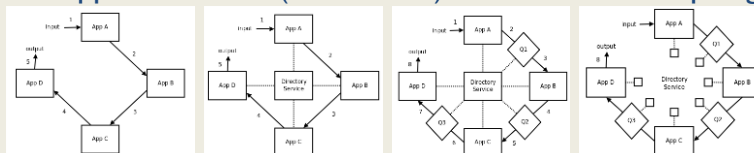 - E.g. frameworks with support common desirable functionality

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.70

70

ZEROMQ – SOCKET LIBRARY

- (0MQ) High performance intelligent **socket library**
- **zero broker, zero latency, zero admin, zero cost, zero waste**
- Provides a message queue
- **Builds upon** functionality of traditional sockets   ØMQ
- Implementation in C++
  - 30+ language bindings provided
- Enables support for various messaging patterns
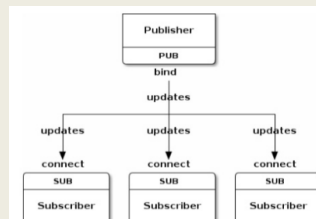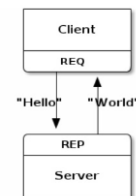- Can support brokered (centralized) and broker-less topologies

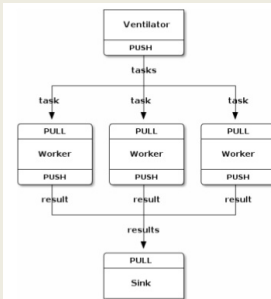| February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.71

71

ZEROMQ – 2

- **ZeroMQ is <u>TCP-connection-oriented communication</u>**

- **Provides socket-like primitives with more functionality**
  - **Basic socket operations *abstracted* away**
  - **Supports many-to-one, one-to-one, and one-to-many connections**
  - ***Multicast* connections (one-to-many – single server socket simultaneously "connects" to multiple clients)**

- **Asynchronous messaging**

- **Supports pairing sockets to support communication patterns**

72

ZEROMQ - PATTERNS

- **<u>Request-reply pattern</u>**
  - **Traditional client-server communication (e.g. RPC)**
  - **Client: request socket (<u>REQ</u>)**
  - **Server: reply socket (<u>REP</u>)**

- **<u>Publish-subscribe pattern</u>**
  - **Clients subscribe to messages published by servers**
  - **As in event-based coordination (Ch. 1)**
  - **Supports multicasting messages from server to multiple**
  - **Client: subscribe socket (<u>SUB</u>)**
  - **Server: publish socket (<u>PUB</u>)**

73

ZEROMQ – PATTERNS - 2

- **Pipeline pattern (FIFO-queue)**
  - Analogous to a producer/consumer bounded buffer
  - Producing processes generate results, push to pipe
  - Consuming processes consume results, pull from pipe
  - Producers: push socket (**PUSH** socket)
  - Consumers: pull socket (**PULL** socket)
  - Push- distributes messages to all pull clients evenly
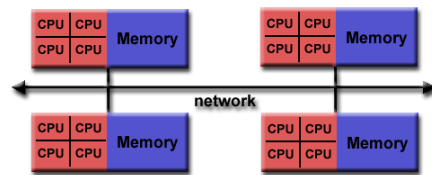 - Consumers pull results from pipe and push results downstream

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.74 |

74

QUEUEING ALTERNATIVES

- **Cloud services**
  - **Amazon Simple Queueing Service (SQS)**
  - **Azure service bus**

- **Open source frameworks**
  - **Nanomsg**
  - **ZeroMQ**

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.75 |

75

MESSAGE PASSING INTERFACE (MPI)

- MPI introduced – version 1.0 March 1994
- Message passing API for parallel programming: *supercomputers*

- Communication protocol for parallel programming for:
  Supercomputers, High Performance Computing (HPC) clusters

- Point-to-point and collective communication

- Goals: high performance, scalability, portability

- Most implementations
  in C, C++, Fortran

76

MOTIVATIONS FOR MPI

- Motivation: sockets insufficient for interprocess
  communication on large scale HPC compute clusters and
  super computers

  - Sockets at the wrong level of abstraction
  - Sockets designed to communicate over the network using
    general purpose TCP/IP stacks
  - Not designed for proprietary protocols
  - Not designed for high-speed interconnection
    networks used by supercomputers,
    HPC-clusters, etc.
  - Better buffering and synchronization needed

77

MOTIVATIONS FOR MPI - 2

- Supercomputers had proprietary communication libraries
  - Offer a wealth of efficient communication operations

- All libraries mutually incompatible

- Led to significant portability problems developing parallel code that could migrate across supercomputers

- Led to development of MPI
  - To support transient (non-persistent) communication for parallel programming

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.78 |

78

MPI FUNCTIONS / DATATYPES

- Very large library, v1.0 (1994) 128 functions

- Version 3 (2015) 440+

- MPI data types:

- Provide common mappings

MPI datatype	C datatype
MPI_CHAR	signed char
MPI_SHORT	signed short int
MPI_INT	signed int
MPI_LONG	signed long int
MPI_UNSIGNED_CHAR	unsigned char
MPI_UNSIGNED_SHORT	unsigned short int
MPI_UNSIGNED	unsigned int
MPI_UNSIGNED_LONG	unsigned long int
MPI_FLOAT	float
MPI_DOUBLE	double
MPI_LONG_DOUBLE	long double
MPI_BYTE	
MPI_PACKED	

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.79 |

79

COMMON MPI FUNCTIONS

- MPI - no recovery for process crashes, network partitions
- Communication among grouped processes: (groupID, processID)
- IDs used to route messages in place of IP addresses

Operation	Description
MPI_bsend	Append outgoing message to a local send buffer
MPI_send	Send message, wait until copied to local/remote buffer
MPI_ssend	Send message, wat until transmission starts
MPI_sendrecv	Send message, wait for reply
MPI_isend	Pass reference to outgoing message and continue
MPI_issend	Pass reference to outgoing messages, wait until receipt start
MPI_recv	Receive a message, block if there is none
MPI_irecv	Check for incoming message, **do not block!**

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.80

80

MESSAGE-ORIENTED-MIDDLEWARE

- **Message-queueing systems**
  - Provide extensive support for _**persistent**_ asynchronous communication
  - In contrast to transient systems
  - Temporally decoupled: messages are eventually delivered to recipient queues

- Message transfers may take minutes vs. sec or ms

- Each application has its own private queue to which other applications can send messages

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.81

81

MESSAGE QUEUEING SYSTEMS: USE CASES

- Enables communication between applications, or sets of processes
  - User applications
  - App-to-database
  - To support distributed real-time computations

- Use cases
  - Batch processing, Email, workflow, groupware, routing subqueries

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.82 |

82

MESSAGE QUEUEING SYSTEMS

- Scenarios:
  (a) Sender/receiver both running

  (b) Sender running, receiver offline

  (c) Sender offline, receiver running

  (d) Sender/receiver both offline

- Queue persists msgs, and attempts to send them but no one may be available to receive them...

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma | L13.83 |

83

MESSAGE QUEUEING SYSTEMS - 2

- **Key:** Truly persistent messaging
- Message queueing systems can persist messages for awhile and senders and receivers can be offline

- **Messages**
- Contain *any* data, may have size limit
- Are properly addressed, to a destination queue

- **Basic Inteface**
- PUT: called by sender to append msg to specified queue
- GET: blocking call to remove oldest msg from specified queue
  - Blocked if queue is empty
- POLL: Non-blocking, gets msg from specified queue

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.84

84

MESSAGE QUEUEING SYSTEMS ARCHITECTURE

- **Basic interface cont'd**
- NOTIFY: install a callback function, for when msg is placed into a queue. Notifies receivers

- **Queue managers**: manage individual message queues as a separate process/library
- Applications get/put messages only from *local* queues
- Queue manager and apps share local network
- **ISSUES:**
- How should we reference the destination queue?
- How should names be resolved (looked-up)?
  - Contact address (host, port) pairs
  - Local look-up tables can be stored at each queue manager

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.85

85

MESSAGE QUEUEING SYSTEMS ARCHITECTURE - 2

- **ISSUES**:
- How do we route traffic between queue managers?
  - How are name-to-address mappings efficiently kept?
  - Each queue manager should be known to all others

- **Message brokers**
- Handle message conversion among different users/formats
- Addresses cases when senders and receivers don't speak the same protocol (language)
- Need arises for message protocol converters
  - "Reformatter" of messages
- Act as application-level gateway

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.86

86

MESSAGE BROKER ORGANIZATION

Plugins to convert messages between APPs
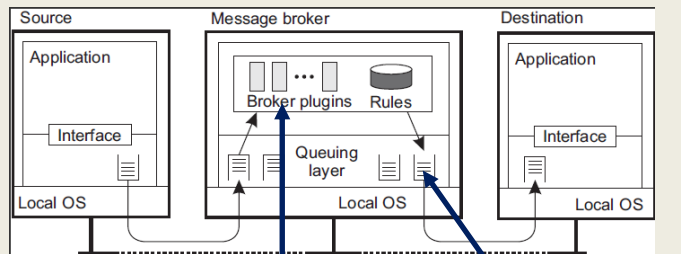
Application-level Queues

February 16, 2023	TCSS558: Applied Distributed Computing [Winter 2023] School of Engineering and Technology, University of Washington - Tacoma	L13.87

87

AMQP PROTOCOL

- Message-queueing systems initially developed to enable legacy applications to interoperate
- Decouple inter-application communication to "open" messaging-middleware
- Many are proprietary solutions, *so not very open*
- e.g. Microsoft Message Queueing service, Windows NT 1997

- **Advanced message queueing protocol (AMQP)**, 2006
- Address openness/interoperability of proprietary solutions
- Open wire protocol for messaging with powerful routing capabilities
- Help *abstract* messaging and application interoperability by means of a generic open protocol
- Suffer from incompatibility among protocol versions
- pre-1.0, 1.0+

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.88 |

88

AMQP - 2

- Consists of: Applications, Queue managers, Queues

- **Connections:** set up to a queue manager, TCP, with potentially many channels, stable, reused by many channels, long-lived

- **Channels:** support short-lived one-way communication

- **Sessions:** bi-directional communication across two channels

- **Link:** provide fine-grained flow-control of message transfer/status between applications and queue manager

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.89 |

89

AMQP MESSAGING

- AMQP nodes: producer, consumer, queue
- Producer/consumer: represent regular applications
- Queues: store/forward messages

- Persistent messaging:
- **Messages** can be marked *durable*
- These messages can only be delivered by nodes able to recover in case of failure
- Non-failure resistant nodes must reject durable messages
- **Source/target** nodes can be marked *durable*
- Track what is durable (node state, node+msgs)

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.90 |

90

MESSAGE-ORIENTED-MIDDLEWARE EXAMPLES:

- **Some examples:**
- RabbitMQ, Apache QPid
  - Implement Advanced Message Queueing Protocol (AMQP)

- Apache Kafka
  - **Dumb broker** (message store), similar to a distributed log file
  - **Smart consumers** – intelligence pushed off to the clients
  - Stores stream of records in categories called topics
  - Supports voluminous data, many consumers, with minimal O/H
  - Kafka **does not track** which messages were read by each consumer
  - Messages are removed after timeout
  - Clients must track their own consumption (*Kafka doesn't help*)
  - Messages have key, value, timestamp
  - Supports high volume pub/sub messaging and streams

| February 16, 2023 | TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma | L13.91 |

91

QUESTIONS

February 16, 2023

TCSS558: Applied Distributed Computing [Winter 2023]
School of Engineering and Technology, University of Washington - Tacoma

L13.92

92