
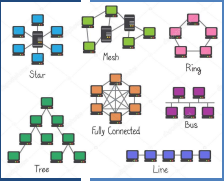


TCSS 558:
APPLIED DISTRIBUTED COMPUTING

Chapter 4 - Processes

Wes J. Lloyd
School of Engineering
and Technology
University of Washington - Tacoma



1

OBJECTIVES

- Midterm Review
- Assignment 1 – questions
- Feedback from 2/13
- Chapter 4: Communication
 - Chapter 4.1: Foundations
 - Chapter 4.2: Remote Procedure Call
 - Chapter 4.3: Message Oriented Communication

February 20, 2020

TCSS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.2

2

MATERIAL / PACE

- Please classify your perspective on material covered in today's class (7 respondents):
 - 1-mostly review, 5-equal new/review, 10-mostly new
 - Average – 5.14
- Please rate the pace of today's class:
 - 1-slow, 5-just right, 10-fast
 - Average – 5.14

February 20, 2020

TCSS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.5

5

FEEDBACK FROM 2/13

February 20, 2020

TCSS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.6

6



CH. 4 COMMUNICATION

L13.7

7

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 20, 2020

TCSS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.8

8



CH. 4.1: FOUNDATIONS

L13.9

9

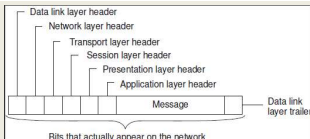
LAYERED PROTOCOLS

- Distributed systems lack shared memory
- All distributed system communication is based on sending and receiving low-level messages
 - $P \rightarrow Q$
- **Open Systems Interconnection 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 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.10

10

LAYERED PROTOCOLS: OSI MODEL



- 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

February 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.11

11

MIDDLEWARE PROTOCOLS

- Middleware is reused by many applications
- Provide needed functions applications are built and depend upon
 - For example: communication frameworks/libraries
- Middleware offer many general-purpose protocols
- Functionality is reusable by **MANY** applications
- Middleware protocol examples:
 - **Authentication protocols:** supports granting users and processes access to authorized resources
 - Doesn't fit as an "application specific" protocol
 - Considered a "Middleware protocol"

February 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.12

12

MIDDLEWARE PROTOCOLS - 2

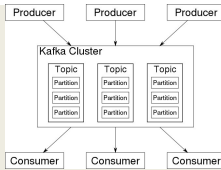
- **Distributed commit protocols**
 - Coordinate a group of processes (nodes)
 - Facilitate all nodes carrying out a particular operation
 - Or abort transaction
 - Provides distributed atomicity (all-or-nothing) operations
- **Distributed locking protocols**
 - Protect a resource from simultaneous access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

February 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.13

13

MIDDLEWARE PROTOCOLS - 3

- **Message queuing 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



February 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.14

14

MIDDLEWARE PROTOCOLS - 3

- **Message queueing services**
 - Support synchronization of data

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

Middleware functions are general, application-independent in nature

Middleware protocol functions are so commonly needed they are offered in reusable frameworks / libraries

February 20, 2020
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.15

15

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?**

February 20, 2020
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.16

16

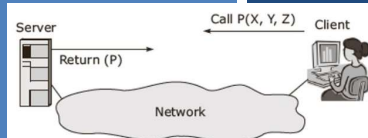
TYPES OF COMMUNICATION - 2

- **Asynchronous communication**
 - Client does not block, continues doing other work
- **Synchronous communication**
 - Client blocks and waits
- **Three types of blocking**
 1. Until middleware notifies it will take over delivering *request*
 2. Sender may block until *request* has been delivered
 3. Sender waits until *request* is processed and result is returned
- **Persistence + synchronization (blocking)**
 - Common scheme for message-queueing systems
 - Block until message delivered to queue
- **Consider each type of blocking (1, 2, 3). Are these modes connectionless (UDP)? connection-oriented (TCP)?**

February 20, 2020
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.17

17

CH. 4.2: RPC



February 20, 2020
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.18

18

RPC – REMOTE PROCEDURE CALL

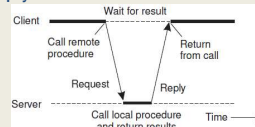
- In a nutshell,
- Allow programs to call procedures on other machines
- Process on **machine A** calls procedure on **machine B**
- Calling process on **machine A** is suspended
- Execution of the called procedure takes place on **machine B**
- Data transported from caller (**A**) to provider (**B**) 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 20, 2020
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.19

19

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 20, 2020
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.20

20

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.21

21

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.22

22

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.23

23

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

BIG-ENDIAN							
...	00	01	02	03	04	05	06
	a	a+1	a+2	a+3	a+4	a+5	a+6
							a+7
LITTLE-ENDIAN							
...	07	06	05	04	03	02	01
	a	a+1	a+2	a+3	a+4	a+5	a+6
							a+7

February 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.24

24

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.25

25

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

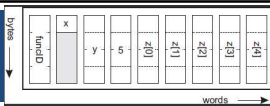
February 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.26

26

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 20, 2020
TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.27

27

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 20, 2020
TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.28

28

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

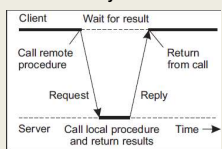
February 20, 2020
TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.29

29

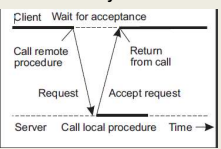
RPC VARIATIONS

- 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/server asynchronous RPC



February 20, 2020
TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.30

30

RPC VARIATIONS - 2

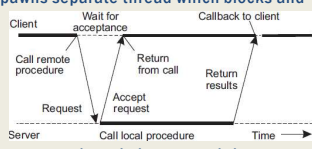
- What are tradeoffs for synchronous vs. asynchronous procedure calls?
 - For a local program
 - For a distributed program (system)
- Use cases for asynchronous procedure calls
 - Long running jobs allow client to perform alternate work in background (in parallel)
 - Client may need to make multiple service calls to multiple server backends at the same time...

February 20, 2020
TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.31

31

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 20, 2020
TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma
L13.32

32

L13.33

L13.34

L13.35

L13.36

L13.37

L13.6

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.39

39

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.40

40

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.41

41

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

February 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.42

42

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.43

43

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 20, 2020

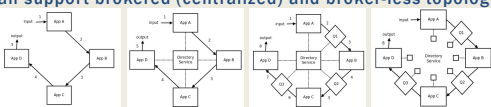
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.44

44

ZEROMQ – SOCKET LIBRARY

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



February 20, 2020


TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.45

45

ZEROMQ – 2

- ZeroMQ is **TCP-connection-oriented communication**
- 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



February 20, 2020

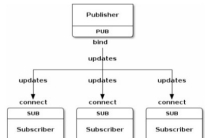
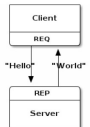
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.46

46

ZEROMQ - PATTERNS

- **Request-reply pattern**
 - Traditional client-server communication (e.g. RPC)
 - Client: request socket (**REQ**)
 - Server: reply socket (**REP**)
- **Publish-subscribe pattern**
 - Clients **subscribe** to messages **published** by servers
 - As in event-based coordination (Ch. 1)
 - Supports multicasting messages from server to multiple
 - Client: subscribe socket (**SUB**)
 - Server: publish socket (**PUB**)



February 20, 2020

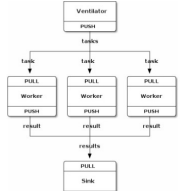
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.47

47

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 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.48

48

QUEUEING ALTERNATIVES

- Cloud services
 - Amazon Simple Queueing Service (SQS)
 - Azure service bus
- Open source frameworks
 - Nanomsg
 - ZeroMQ

February 20, 2020

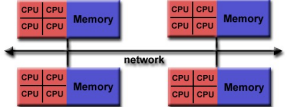
TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.49

49

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



February 20, 2020

TCCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.50

50

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 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.57

57

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 Interface**
 - 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 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.58

58

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 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.59

59

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 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.60

60

MESSAGE BROKER ORGANIZATION

February 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.61

61

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 20, 2020 TCS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma L13.62

62

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 20, 2020

TCSS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.63

63

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 20, 2020

TCSS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.64

64

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 20, 2020

TCSS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.65

65

QUESTIONS



February 20, 2020

TCSS558: Applied Distributed Computing [Winter 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.66

66