

TCSS 558: APPLIED DISTRIBUTED COMPUTING

Distributed Systems Architectures and Middleware Organization

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

OBJECTIVES - 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 - Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
- Class Activity: Architectural Styles
- Chapter 2.2: Middleware Organization
- Chapter 2.3: System Architectures
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.2

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

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

January 19, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.3

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		

January 19, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.4

MATERIAL / PACE

- Please classify your perspective on material covered in today's class (23 respondents):
 - 1-mostly review, 5-equal new/review, 10-mostly new
 - Average - 6.74** (↓ - previous 7.46)
- Please rate the pace of today's class:
 - 1-slow, 5-just right, 10-fast
 - Average - 5.57** (↓ - previous 5.67)

January 19, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.5

FEEDBACK FROM 1/14

- For pervasive, I am not clear about the differences between Ubiquitous computing systems and Sensor networks.***
- Pervasive Computing Systems** consists of systems that have hardware "everywhere"
- We covered three types:
 - Ubiquitous Systems, Mobile Systems, Sensor Networks
- A major distinguishing factor with Ubiquitous Systems is that they consist of many **heterogeneous** devices
 - processors in day-to-day objects → think Home Automation
- Sensor networks
 - Can have heterogeneous devices, but in general many sensor nodes are uniform / similar in nature (**homogeneous**)

January 19, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.6

FEEDBACK - 2

- ***I'm wondering how architectures (Layered and Object-based) are actually built or coded? How are they formed in real life?***
- In general, you would develop the system by following the overall architectural style
- **Layered:**
 - Layers divide functionality to provide a separate of concerns
 - Lower-level layers provide functionality to higher layers
 - Higher-level layers benefit from functionality offered by lower layers
 - Higher-level layers in general do not worry about specific details regarding the implementation of lower layers (i.e. abstraction)
 - This architecture enables extensibility and reusability
- **Object-based:** ...

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L5.7

APPLICATION LAYERING

- Distributed application example: Internet search engine

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L5.8

OBJECT AND COMPONENT BASED ARCHITECTURAL STYLES

Component-based design model emphasizes reuse

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L5.9

OBJECTIVES - 1/19

- Questions from 1/14
- **Assignment 0: Cloud Computing Infrastructure Tutorial**
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 – Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
 - Class Activity: Architectural Styles
 - Chapter 2.2: Middleware Organization
 - Chapter 2.3: System Architectures
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L5.10

ASSIGNMENT 0

- **Preparing for Assignment 0:**
 - Establish AWS Account
 - Standard account – **** request cloud credits from instructor ****
 - Specify "AWS CREDIT REQUEST" as subject of email
 - Include email address of AWS account
 - **AWS Educate Starter account** – some account limitations
 - https://awseducate-starter-account-services.s3.amazonaws.com/AWS_Educate_Starter_Account_Services_Supported.pdf
 - Establish local Linux/Ubuntu environment
- Task 1 – AWS account setup
- Task 2 – Working w/ Docker, creating Dockerfile for Apache Tomcat
- Task 3 – Creating a Dockerfile for haproxy
- Task 4 – Working with Docker-Machine
- Task 5 – For Submission: Testing Alternate Server Configurations

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L5.11


OBJECTIVES - 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 – Architectural Styles
 - **Resource-centered architectures**
 - **Representational state transfer (REST)**
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
 - Class Activity: Architectural Styles
 - Chapter 2.2: Middleware Organization
 - Chapter 2.3: System Architectures
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L5.12

RESOURCE BASED ARCHITECTURES

- Motivation:
 - Increasing number of services available online
 - Each with specific protocol(s), methods of interfacing
 - Connecting services w/ different TCP/IP protocols → integration nightmare
 - Need for specialized client for each service that speaks the application protocol "language"...
- Need standardization of interfaces
 - Make services/components more pluggable
 - Easier to adopt and integrate
 - Common architecture



January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.13

REST SERVICES

- Representational State Transfer (REST)
- Built on HTTP
- Four key characteristics:
 1. Resources identified through single naming scheme
 2. Services offer the same interface
 - Four operations: GET PUT POST DELETE
 3. Messages to/from a service are fully described
 4. After execution server forgets about client
 - Stateless execution

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.14

HYPERTEXT TRANSPORT PROTOCOL (HTTP)

- An ASCII-based request/reply protocol for transferring information on the web
- HTTP request includes:
 - request method (GET, POST, etc.)
 - Uniform Resource Identifier (URI)
 - HTTP protocol version understood by the client
 - headers—extra info regarding transfer request
- HTTP response from server
 - Protocol version & status code →
 - Response headers
 - Response body

HTTP status codes:
 2xx — *all is well*
 3xx — *resource moved*
 4xx — *access problem*
 5xx — *server error*

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.15

REST-FUL OPERATIONS

Operation	Description	
PUT	Create a new resource	(C)reate
GET	Retrieve state of a resource in some format	(R)ead
POST	Modify a resource by transferring a new state	(U)pdate
DELETE	Delete a resource	(D)elete

- Resources often implemented as objects in OO languages
- REST is weak for tracking state
- Generic REST interfaces enable ubiquitous "so many" clients

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.16

EXAMPLE: AMAZON S3

- Amazon S3 offers a REST-based interface
- Requires signing HTTP authorization header or passing authentication parameters in the URL query string
- REST: GET/PUT/POST/DELETE
- SOAP: 16 operations, moving toward deprecation
- Python boto ~50 operations (SDK for Python)
- SDKs for other languages

- AWS SDKs and Explorers
- Set Up the AWS CLI
- Using the AWS SDK for Java
- Using the AWS SDK for .NET
- Using the AWS SDK for PHP and Running PHP Examples
- Using the AWS SDK for Ruby - Version 3
- Using the AWS SDK for Python (Boto)

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.17

REST - 2

- Defacto web services protocol
- Requests made to a URI – uniform resource identifier
- Supersedes SOAP – Simple Object Access Protocol
- Access and manipulate web resources with a predefined set of stateless operations (known as web services)
- Responses most often in JSON, also HTML, ASCII text, XML, no real limits as long as text-based
- curl – generic command-line REST client:
<https://curl.haxx.se/>

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.18

```
// WSDL Service Definition
<?xml version="1.0" encoding="UTF-8"?>
<definitions name="DayOfWeek"
  targetNamespace="http://www.zogswave.com/soapwz/examples/DayOfWeek.wsdl"
  xmlns:tns="http://www.zogswave.com/soapwz/examples/DayOfWeek.wsdl"
  xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:soap11="http://schemas.xmlsoap.org/soap11/">
  <message name="DayOfWeekInput">
    <part name="data" type="xsd:string"/>
  </message>
  <message name="DayOfWeekResponse">
    <part name="dayOfWeek" type="xsd:string"/>
  </message>
  <portType name="DayOfWeekPortType">
    <operation name="GetDayOfWeek">
      <input message="tns:DayOfWeekInput"/>
      <output message="tns:DayOfWeekResponse"/>
    </operation>
  </portType>
  <binding name="DayOfWeekBinding" type="tns:DayOfWeekPortType">
    <soap:binding style="document"
      transport="http://schemas.xmlsoap.org/soap/http"/>
    <operation name="GetDayOfWeek">
      <soap:operation soapAction="getDayOfWeek"/>
      <input
        namespace="http://www.zogswave.com/soapwz/examples"
        encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
      </input>
      <output
        namespace="http://www.zogswave.com/soapwz/examples"
        encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
      </output>
    </operation>
  </binding>
  <service name="DayOfWeekService">
    <documentation>
      Returns the day-of-week name for a given date
    </documentation>
    <port name="DayOfWeekPort" binding="tns:DayOfWeekBinding">
      <soap:address location="http://localhost:8090/dayofweek/DayOfWeek"/>
    </port>
  </service>
</definitions>
```

L5.19

```
// REST/JSON
// Request climate data for Washington

{
  "parameter": [
    {
      "name": "latitude",
      "value": 47.2529
    },
    {
      "name": "longitude",
      "value": -122.4443
    }
  ]
}
```

L5.20

OBJECTIVES - 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 – Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - **Event-based**
 - **Publish and subscribe (Rich Site Summary RSS feeds)**
 - Class Activity: Architectural Styles
 - Chapter 2.2: Middleware Organization
 - Chapter 2.3: System Architectures
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L5.21

PUBLISH-SUBSCRIBE ARCHITECTURES

- Enables separation between processing and coordination
- Types of coordination:


	Temporally coupled (at the same time)	Temporally decoupled (at different times)
Referentially coupled (dependent on name)	Direct Explicit synchronous service call	Mallbox Asynchronous by name (address)
Referentially decoupled (name not required)	Event-based Event notices published to shared bus, w/o addressing	Shared data space Processes write tuples to a shared data space

Publish and subscribe architectures

January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L5.22

PUBLISH-SUBSCRIBE ARCHITECTURES - 2


- **Event-based coordination**
- Processes do not know about each other explicitly
- **Processes:**
 - **Publish:** a notification describing an event
 - **Subscribe:** to receive notification of specific kinds of events
- Assumes subscriber is presently up (*temporally coupled*)
- Subscribers must actively **MONITOR** event bus



January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L5.23

PUBLISH SUBSCRIBE ARCHITECTURES - 3

- **Shared data space**
- Full decoupling (name and time)
- Processes publish “tuples” to shared dataspace (publish)
- Processes provide search pattern to find tuples (subscribe)
- When tuples are added, subscribers are notified of matches
- **Key characteristic:** Processes have no explicit reference to each other



January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L5.24

PUBLISH SUBSCRIBE ARCHITECTURES - 4

- Subscriber describes events interested in
- Complex descriptions are intensive to evaluate and fulfil
- Middleware will:
- Publish matching notification and data to subscribers
 - Common if middleware lacks storage
- Publish only matching notification
 - Common if middleware provides storage facility
 - Client must explicitly fetch data on their own
- Publish and subscribe systems are generally scalable
- What would reduce the scalability of a publish-and-subscribe system?**

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.25

OBJECTIVES - 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 - Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
 - Class Activity: Architectural Styles**
 - Chapter 2.2: Middleware Organization
 - Chapter 2.3: System Architectures
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.26

IN-CLASS ACTIVITY: ARCHITECTURAL STYLES



L5.27

CLASS ACTIVITY 2

- We will form groups of ~2-3 and enter breakout rooms
- Each group will complete a Google Doc worksheet
- Add names to Google Doc as they appear in Canvas
- Once completed, **one person** submits a PDF of the Google Doc to Canvas
- Instructor will score all group members based on the uploaded PDF file
- To get started:
 - Log into your ***** UW Google Account *****
 - Link to shared Google Drive
 - Follow link:
<https://tinyurl.com/y43bflzs>

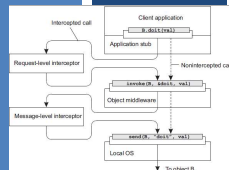
October 7, 2020 TCCS562: Software Engineering for Cloud Computing [Fall 2020]
School of Engineering and Technology, University of Washington - Tacoma L5.28

DISTRIBUTED SYSTEM GOALS TO CONSIDER

- Consider how the architectural change may impact:**
- Availability
- Accessibility
- Responsiveness
- Scalability
- Openness
- Distribution transparency
- Supporting resource sharing
- Other factors...

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.29

CH 2.2: MIDDLEWARE ORGANIZATION



The diagram illustrates the flow of data and control through various middleware layers. It starts with a Client application and Application stub, which interact with Request-level and Message-level interceptors. These interceptors then interact with Object and Message middleware, which in turn interact with Local OS and Subject B. The diagram shows how data and control flow through these layers, with arrows indicating the direction of communication.

January 19, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.30

OBJECTIVES - 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 - Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
 - Class Activity: Architectural Styles
 - Chapter 2.2: Middleware Organization**
 - Chapter 2.3: System Architectures
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma LS.31

MIDDLEWARE ORGANIZATION

- Relies on two important design patterns:
 - Wrappers
 - Interceptors
- Both help achieve the goal of openness

January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma LS.32

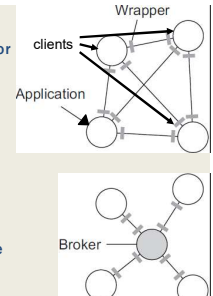
MIDDLEWARE: WRAPPERS

- Wrappers (also called adapters)**
 - WHY?:** Interfaces available from legacy software may not be sufficient for all new applications to use
 - WHAT:** Special "frontend" components that provide interfaces for clients
 - Interface wrappers transform client requests to "implementation" (i.e. legacy software) at the component-level
 - Can then provide modern service interfaces for legacy code/systems
 - Components encapsulate (i.e. abstract) dependencies to meet all preconditions to operate and host legacy code
 - Interfaces parameterize legacy functions, abstract environment configuration (i.e. make into black box)
- Contributes towards system **OPENNESS**
- Example: Amazon S3:** S3 HTTP REST interface
- GET/PUT/DELETE/POST: requests handed off for fulfillment

January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma LS.33

MIDDLEWARE: WRAPPERS - 2

- Inter-application communication
 - Applications may provide unique interface for every client application
- Scalability suffers
 - N applications $\rightarrow O(N^2)$ wrappers
- ALTERNATE: Use a Broker**
 - Provide a common intermediary
 - Broker knows how to communicate with every application
 - Applications only know how to communicate with the broker



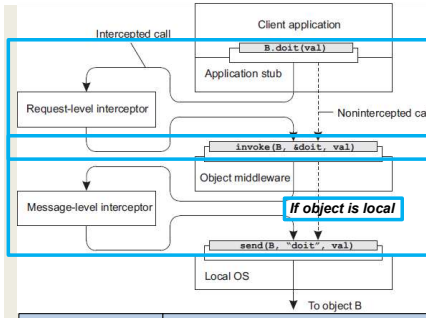
January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma LS.34

MIDDLEWARE: INTERCEPTORS

- Interceptor**
- Software construct, breaks flow of control, allows other application code to be executed
- Interceptors send calls to other servers, or to ALL servers that replicate an object while abstracting the **distribution** and/or **replication**
 - Used to enable remote procedure calls (RPC), remote method invocation (RMI)
- Object A calls method belonging to object B
 - Interceptors route calls to object B regardless of location

January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma LS.35

MIDDLEWARE: INTERCEPTORS - 2



Request-level interceptor transforms:
`B.doit(val)`
 into generic call:
`invoke(B, doit, val)`

Message-level interceptor in middleware sends message through OS (TCP/IP socket) to transfer data:
`send(B, "doit", val)`

Non-intercepted:

January 19, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma LS.36

MIDDLEWARE INTERCEPTION - METHOD

- **MIDDLEWARE:** Provides local interface matching Object B to Object A
- Object A calls Object B's method provided by local interface
- A's call is transformed into a "generic object invocation" by **request-level Interceptor**
- "Generic object invocation" is transformed into a **message** by **message-level Interceptor** and sent over Object A's network to Object B
- Interception automatically routes calls to all object replicas

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.37
------------------	---	-------

MODIFIABLE MIDDLEWARE

- It should be possible to modify middleware without loss of availability
- Software components can be replaced at runtime
- Component-based design
 - Modifiability through composition
 - Systems may have static or dynamic configuration of components
 - Dynamic configuration requires **late binding**
 - Components can be changed at runtime
- Component based software supports modifiability at runtime by enabling components to be swapped out.
- **Does a microservices architecture (e.g. AWS Lambda) support modifiability at runtime ?**

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.38
------------------	---	-------

CH 2.3: SYSTEM ARCHITECTURES

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.39
------------------	---	-------

OBJECTIVES - 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 - Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
 - Class Activity: Architectural Styles
 - Chapter 2.2: Middleware Organization
 - **Chapter 2.3: System Architectures**
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.40
------------------	---	-------

SYSTEM ARCHITECTURES

- Architectural styles (or patterns)
- General, reusable solutions to commonly occurring system design problems
- Expressed as a logical organization of **components** and **connectors**
- Deciding on the system components, their interactions, and placement is a "realization" of an **architectural style**
- System architectures represent designs used in practice

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.41
------------------	---	-------

OBJECTIVES - 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 - Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
 - Class Activity: Architectural Styles
 - Chapter 2.2: Middleware Organization
 - **Chapter 2.3: System Architectures**
 - **Centralized system architectures**
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.42
------------------	---	-------

TYPES OF SYSTEM ARCHITECTURES

- Centralized system architectures
 - Client-server
 - Multitiered
- Decentralized peer-to-peer architectures
 - Structured
 - Unstructured
 - Hierarchically organized
- Hybrid architectures

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.43

CENTRALIZED: SIMPLE CLIENT-SERVER ARCHITECTURE

- **Clients** request services
- **Servers** provide services
- Request-reply behavior

```

sequenceDiagram
    participant Client
    participant Server
    Client->>Server: Request
    Server->>Client: Provide service
    Server-->>Client: Reply
    Client-->>Client: Wait
    
```

- **Connectionless protocols (UDP)**
- Assume stable network communication with no failures
- **Best effort communication:** No guarantee of message arrival without errors, duplication, delays, or in sequence. No acknowledgment of arrival or retransmission
- **Problem:** How to detect whether the client request message is lost, or the server reply transmission has failed
- Clients can resend the request when no reply is received
- **But what is the server doing?**

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.44

CLIENT-SERVER PROTOCOLS

- **Connectionless cont'd**
- Is resending the client request a good idea?
- **Examples:**
Client message: "transfer \$10,000 from my bank account"
Client message: "tell me how much money I have left"
- **Idempotent** - repeating requests is safe
- **Connection-oriented (TCP)**
- Client/server communication over wide-area networks (WANs)
- When communication is inherently reliable
- Leverage "reliable" TCP/IP connections

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.45

CLIENT-SERVER PROTOCOLS - 2

- **Connection-oriented cont'd**
- Set up and tear down of connections is relatively expensive
- Overhead can be amortized with longer lived connections
 - Example: database connections often retained
- Ongoing debate:
 - How do you differentiate between a client and server?
 - Roles are **blurred**
- **Blurred Roles Example:** Distributed databases
- DB nodes both **service** client requests, *and* **submit** new requests to other DB nodes for replication, synchronization, etc.

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.46

TCP/UDP

TCP	UDP
Reliable	Unreliable
Connection-oriented	Connectionless
Segment retransmission and flow control through windowing	No windowing or retransmission
Segment sequencing	No sequencing
Acknowledge segments	No acknowledgement

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.47

CONNECTIONLESS VS CONNECTION ORIENTED

	Connectionless (UDP) <i>stateless</i>	Connection-oriented (TCP) <i>stateful</i>
Advantages		
Disadvantages		

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.48

CONNECTIONLESS VS CONNECTION ORIENTED

	Connectionless (UDP) <i>stateless</i>	Connection-oriented (TCP) <i>stateful</i>
Advantages	<ul style="list-style-type: none"> Fast to communicate (no connection overhead) Broadcast to an audience Network bandwidth savings 	<ul style="list-style-type: none"> Message delivery confirmation Idempotence not required Messages automatically resent - if client (or network) is temporarily unavailable Message sequences guaranteed
Disadvantages	<ul style="list-style-type: none"> Cannot tell difference of request vs. response failure Requires idempotence Clients must be online and ready to receive messages 	<ul style="list-style-type: none"> Connection setup is time-consuming More bandwidth is required (protocol, retries, multinode-communication)

January 19, 2021 TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma LS.49

MULTITIERED ARCHITECTURES

- Where should functionality be distributed?
 - At the client?
 - At the server?

- Why should we consider component composition?

January 19, 2021 TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma LS.50

Bell's Number:

n	k
4	15
5	52
6	203
7	877
8	4,140
9	21,147
n	...

k: number of ways n components can be distributed across containers
 M: Tomcat ApplicationServer
 D: Postgresql DB
 F: nginx file server
 L: Logging server (high O/H)

Resource utilization profile changes from component composition

M-bound RUSLE2 – Soil Erosion Model Webservice

- Box size shows absolute deviation (+/-) from mean
- Shows relative magnitude of performance variance

Two application variants tested

- M-bound: Standard service, M is compute bound
- D-bound: Modified service, D is compute bound

Resource footprint	M-bound	D-bound
Disk sector reads:	21.8%	111.1%
Disk sector writes:	144.9%	145%
Network bytes received:	143.7%	143.9%
Network bytes sent:		

PERFORMANCE IMPLICATIONS OF COMPONENT DEPLOYMENTS

Δ Performance Change: Min to max performance

M-bound:	14%
D-bound:	25.7%

53

MULTITIERED ARCHITECTURES - 2


- M D F L architecture**
- M - is the application server
- M - is also a client to the database (D), fileserver (F), and logging server (L)

Server as a client

January 19, 2021 TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma LS.54

MULTITIERED RESOURCE SCALING

- Vertical distribution
 - The distribution of "M D F L"
 - Application is scaled by placing "tiers" on separate servers
 - M - The application server
 - D - The database server
 - Vertical distribution impacts "network footprint" of application
 - Service isolation: each component is isolated on its own HW
- Horizontal distribution
 - Scaling an individual tier
 - Add multiple machines and distribute load
 - Load balancing



January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.55
------------------	---	-------

MULTITIERED RESOURCE SCALING - 2

- Horizontal distribution cont'd
 - Sharding: portions of a database map to a specific server
 - Distributed hash table
 - Or replica servers

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.56
------------------	---	-------

OBJECTIVES - 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 - Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
 - Class Activity: Architectural Styles
 - Chapter 2.2: Middleware Organization
 - Chapter 2.3: System Architectures
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - Hybrid architectures

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.57
------------------	---	-------

TYPES OF SYSTEM ARCHITECTURES

- Centralized system architectures
 - Client-server
 - Multitiered
- Decentralized peer-to-peer architectures
 - Structured
 - Unstructured
 - Hierarchically organized
- Hybrid architectures

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.58
------------------	---	-------

DECENTRALIZED PEER-TO-PEER ARCHITECTURES

- Client/server:
 - Nodes have specific roles
- Peer-to-peer:
 - Nodes are seen as *all equal...*
- How should nodes be organized for communication?

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.59
------------------	---	-------

STRUCTURED PEER-TO-PEER

- Nodes organized using specific *topology* (e.g. ring, binary-tree, grid, etc.)
 - Organization assists in data lookups
- Data indexed using "semantic-free" indexing
 - Key / value storage systems
 - Key used to look-up data
- Nodes store data associated with a subset of keys

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.60
------------------	---	-------

DISTRIBUTED HASH TABLE (DHT)

- Distributed hash table (DHT) (*ch. 5*)
- Hash function
 - $key(\text{data item}) = hash(\text{data item's value})$
- Hash function "generates" a unique key based on the data
- No two data elements will have the same key (hash)
- System supports data lookup via key
- **Any** node can receive and resolve the request
- Lookup function determines which node stores the key
 - $existing\ node = lookup(key)$
- Node forwards request to node with the data

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.61

FIXED HYPERCUBE EXAMPLE

- Example where topology helps route data lookup request
- Statically sized 4-D hypercube, every node has 4 connectors
- 2 x 3-D cubes, 8 vertices, 12 edges
- Node IDs represented as 4-bit code (0000 to 1111)
- Hash data items to 4-bit key (1 of 16 slots)
- Distance (number of hops) determined by identifying number of varying bits between neighboring nodes and destination

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.62

FIXED HYPERCUBE EXAMPLE - 2

- **Example:** *fixed hypercube*
node 0111 (7) retrieves data from node 1110 (14)
- Node 1110 is not a neighbor to 0111
- **Which connector leads to the shortest path?**

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.63

WHICH CONNECTOR LEADS TO THE SHORTEST PATH?

- **Example:** node 0111 (7) retrieves data from node 1110 (14)
- Node 1110 is not a neighbor to 0111

[0111] Neighbors:
 1111 (1 bit different than 1110) 0011 (3 bits different- bad path)
 0110 (1 bit different than 1110) 0101 (3 bits different- bad path)

- **Does It matter which node is selected for the first hop?**

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.64

DYNAMIC TOPOLOGY

- Fixed hypercube requires static topology
 - Nodes cannot join or leave
- Relies on symmetry of number of nodes
- Can force the DHT to a certain size
- Chord system - DHT (again in ch.5)
 - Dynamic topology
 - Nodes organized in ring
 - Every node has unique ID
 - Each node connected with other nodes (shortcuts)
 - Shortest path between any pair of nodes is ~ order $O(\log N)$
 - N is the total number of nodes

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.65

CHORD SYSTEM

- Data items have m-bit key
- Data item is stored at closest "successor" node with ID \geq key k
- Each node maintains finger table of successor nodes
- Client sends key/value lookup to **any** node
- Node forwards client request to node with m-bit ID closest to, but not greater than key k
- Nodes must continually refresh finger tables by communicating with adjacent nodes to incorporate node joins/departures

January 19, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L5.66

UNSTRUCTURED PEER-TO-PEER

- **No topology:** *How do nodes find out about each other?*
- Each node maintains adhoc list of neighbors
- Facilitates nodes frequently joining, leaving, adhoc systems
- **Neighbor:** node reachable from another via a network path
- Neighbor lists constantly refreshed
 - Nodes query each other, remove unresponsive neighbors
- Forms a "random graph"
- Predetermining network routes not possible
 - How would you calculate the route algorithmically?
- Routes must be discovered

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.67
------------------	---	-------

SEARCHING FOR DATA: UNSTRUCTURED PEER-TO-PEER SYSTEMS

- **Flooding**
- [Node u] sends request for data item to all neighbors
- [Node v]
 - Searches locally, responds to u (or forwarder) if having data
 - Forwards request to ALL neighbors
 - Ignores repeated requests
- Features
 - High network traffic
 - Fast search results by saturating the network with requests
 - Variable # of hops
 - Max number of hops or time-to-live (TTL) often specified
 - Requests can "retry" by gradually increasing TTL/max hops until data is found

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.68
------------------	---	-------

SEARCHING FOR DATA - 2

- **Random walks**
- [Node u] asks a randomly chosen neighbor [node v]
- If [node v] does not have data, forwards request to a random neighbor
- Features
 - Low network traffic
 - Akin to sequential search
 - Longer search time
 - [node u] can start "n" random walks simultaneously to reduce search time
 - As few as n=16,64 random walks sufficient to reduce search time (LV et al. 2002)
 - Timeout required - need to coordinate stopping network-wide walk when data is found...

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.69
------------------	---	-------

SEARCHING FOR DATA - 3

- **Policy-based search methods**
- Incorporate history and knowledge about the adhoc network **at the node-level** to enhance effectiveness of queries
- Nodes maintain lists of preferred neighbors which often succeed at resolving queries
- Favor neighbors having highest number of neighbors
 - Can help minimize hops

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.70
------------------	---	-------

HIERARCHICAL PEER-TO-PEER NETWORKS

- **Problem:**
Adhoc system search performance does not scale well as system grows
- Allow nodes to assume **ROLES** to improve search
- Content delivery networks (CDNs) (*video streaming*)
 - Store (cache) data at nodes local to the requester (client)
 - Broker node - tracks resource usage and node availability
 - Track where data is needed
 - Track which nodes have capacity (disk/CPU resources) to host data
- Node roles
 - **Super peer** - Broker node, routes client requests to storage nodes
 - **Weak peer** - Store data

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.71
------------------	---	-------

HIERARCHICAL PEER-TO-PEER NETWORKS - 2

- Super peers
 - Head node of local centralized network
 - Interconnected via overlay network with other super peers
 - May have replicas for fault tolerance
- Weak peers
 - Rely on super peers to find data
- Leader-election problem:
 - Who can become a super peer?
 - What requirements must be met to become a super peer?

January 19, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.72
------------------	---	-------

OBJECTIVES – 1/19

- Questions from 1/14
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
 - Chapter 2.1 – Architectural Styles
 - Resource-centered architectures
 - Representational state transfer (REST)
 - Event-based
 - Publish and subscribe (Rich Site Summary RSS feeds)
 - Class Activity: Architectural Styles
 - Chapter 2.2: Middleware Organization
 - Chapter 2.3: System Architectures
 - Centralized system architectures
 - Decentralized peer-to-peer architectures
 - **Hybrid architectures**

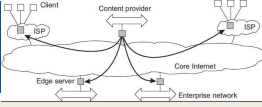
January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.73
------------------	---	-------

TYPES OF SYSTEM ARCHITECTURES

- Centralized system architectures
 - Client-server
 - Multitiered
- Decentralized peer-to-peer architectures
 - Structured
 - Unstructured
 - Hierarchically organized
- **Hybrid architectures**

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.74
------------------	---	-------

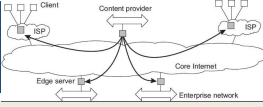
HYBRID ARCHITECTURES



- Combine centralized server concepts with decentralized peer-to-peer models
- **Edge-server systems:**
 - Adhoc peer-to-peer devices connect to the internet through an edge server (origin server)
 - Edge servers (provided by an ISP) can optimize content and application distribution by storing assets near the edge
- **Example:**
 - AWS Lambda@Edge: Enables Node.js Lambda Functions to execute “at the edge” harnessing existing CloudFront Content Delivery Network (CDN) servers
 - <https://www.infoq.com/news/2017/07/aws-lambda-at-edge>

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.75
------------------	---	-------

HYBRID ARCHITECTURES - 2



- **Fog computing:**
 - Extend the scope of managed resources beyond the cloud to leverage compute and storage capacity of end-user devices
 - End-user devices become part of the overall system
 - Middleware extended to incorporate managing edge devices as participants in the distributed system
- Cloud → in the sky
 - compute/resource capacity is huge, but far away...
- Fog → (devices) on the ground
 - compute/resource capacity is constrained and local...


January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.76
------------------	---	-------

COLLABORATIVE DISTRIBUTED SYSTEM EXAMPLE

- **BitTorrent Example:**
 - File sharing system – users must contribute as a file host to be eligible to download file resources
 - Original implementation features hybrid architecture
 - Leverages idle client network capacity in the background
 - User joins the system by interacting with a central server
 - Client accesses global directory from a **tracker** server at well known address to access torrent file
 - Torrent file tracks nodes having chunks of requested file
 - Client begins downloading file chunks and immediately then participates to reserve downloaded content **or network bandwidth is reduced!!**
 - Chunks can be downloaded in parallel from distributed nodes

January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.77
------------------	---	-------

QUESTIONS



January 19, 2021	TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.78
------------------	---	-------