Principles of Software Construction: Objects, Design, and Concurrency

# Design for Robustness: Distributed Systems





#### Administrivia

Reading/quiz for Tuesday.

We are working with selected teams to clean up framework code, will announce/release later today.



#### Where we are

	<i>Small scale:</i> One/few objects	<i>Mid scale:</i> Many objects	<i>Large scale:</i> Subsystems
	Subtype	Domain Analysis 🗸	GUI vs Core 🗸
Design for	Polymorphism 🗸	Inheritance & Del. 🗸	Frameworks and
understanding	Information Hiding,	Responsibility	Libraries 🗸 , APIs 🗸
change/ext.	Contracts 🗸	Assignment,	Module systems,
enange, era	Immutability 🗸	Design Patterns,	microservices
reuse	Types 🗸	Antipattern 🗸	(Testing for)
robustness	Unit Testing 🗸	Promises/	Robustness
	<u> </u>	Reactive P. 🗸	CI ✔, DevOps ✔,
		Integration Testing $\checkmark$	Teams



#### **\ A /I** . . . 4 4

Where	e did we start?				src > T	T\$ ui.ts >	
VVDACCS	<pre>Bailed wee start?  Java-ULjava  Cd Out</pre>		VrPESCRIPT     Scards     Sards     data     ordering     T3 index.ts     figure     () package-lock.json     () package.json     @ README.md     fi tsconfig.json		src> T 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	<pre>IS IX IS uits &gt; import { FlashCard } from './cards/flashcard' import { CardDeck } from './cards/flashcard' import { CardDeck } from './cards/flashcard' import readline from 'readline-sync' interface UI { interface UI { interface UI { function newUI (): UI { function cueAllCards (producer: CardDeck): void { function cueAllCards (producer: CardDeck): void { fort const cardStatus.getCard() const card = cardStatus.getCard() const correcthnswer = cueCard(card) cardStatus.recordResult(correctAnswer) } } function cueCard (card: FlashCard): boolean { function cueCard flashCard): boolean { function cueCard flashCard}: flashCard): function cueCard flashCard}; function cueCard flashCard}</pre>	
	<pre>21</pre>	8 \$	) OUTLINE ) TIMELINE Pimain ⊖ ⊙1∆0 0 Clai	ire 🔐 🔗 I	26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 49 Live Share	<pre>return success return {     /**     /**</pre>	d. RU peScript
🕈 Git 🗄 TODO 🕒 Problems 🔚 Terminal 🧧	📒 CheckStyle 🔨 Build 📚 Dependencies		C Event Log				

. . .

17-214/514



#### We design for robustness even in "small" systems.

- Single-threaded, local systems:
  - Problems are (typically) deterministic
  - Checked vs. unchecked exceptions
- Key ideas:
  - Provide explicit control-flow for normal and abnormal execution
    - Error handling and recovery for the latter
  - Unit testing to increase confidence
    - Cover both typical and atypical/boundary paths



#### We design for robustness even in "small" systems.

- Concurrent, local systems:
  - Non-determinism from thread ordering, asynchronous returns
  - Errors can occur at any shared mutable state
- Key ideas:
  - Reduce mutable state
    - Use atomicity, synchronization everywhere else
  - Organize asynchrony with promises
    - Especially natural in a single-threaded environment



#### Then we forced you to do this...

#### NanoHTTPd





Modern software is dominated by systems composed of [components, APIs, modules], developed by completely different people, communicating over a network!



# For example

- 3rd party Facebook apps
- Android user interface
- Backend uses Facebook data









#### What is a distributed system?

- Multiple system components (computers) communicating via some medium (the network) to achieve some goal
- "Concurrent" (shared-memory multiprocessing) vs. Distributed
  - Agents: Threads vs. Processes
    - Processes typically spread across multiple computers
    - Can put them on one computer for testing
  - Communication: changes to Shared Objects vs. Network Messages



#### **Distributed systems**

- A collection of autonomous systems working together to form a single system
  - Enable scalability, availability, resiliency, performance, etc ...
- Remote procedure calls instead of function calls
   Typically REST API to URL

• Benefits? Drawbacks?



#### **Distributed System Benefits**

Scalability

Very strong encapsulation (only APIs public)

Computation beyond local resources

Independent deployment, operations, and evolution

Also multiple containers on single system

Pay per transaction / storage / use



#### What is a distributed system?

"A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable."

-- Leslie Lamport







17-214/514

A monolithic application puts all its functionality into a single process...



A microservices architecture puts each element of functionality into a separate service...



... and scales by replicating the monolith on multiple servers









... and scales by distributing these services across servers, replicating as needed.



17-214/514

source: http://martinfowler.com/articles/microservices.html





17-214/514

http://christophermeiklejohn.com/filibuster/2021/10/14/filibuster-4.html





1/-214/514

http://christophermeiklejohn.com/filibuster/2021/10/14/filibuster-4.html





17-214/514

19 ISC institute for SOFTWARE RESEARCH

#### Microservices

Building applications as suite of small and easy to replace services

- fine grained, one functionality per service
- (sometimes 3-5 classes)
- composable
- easy to develop, test, and understand
- fast (re)start, fault isolation

Modelled around business domain

Interplay of different systems and languages, no commitment to technology stack

Easily deployable and replicable

Embrace automation, embrace faults

Highly observable



#### 17-214/514

#### **Technical Considerations**

**REST APIs** 

Independent development and deployment

Self-contained services (e.g., each with own database)

• multiple instances behind load-balancer

Streamline deployment







for less-complex systems, the extra baggage required to manage microservices reduces productivity

#### Overhead



but remember the skill of the team will outweigh any monolith/microservice choice

#### 17-214/514

#### Software Architecture vs Design Patterns

Design patterns: Composition and interaction of objects

Architectural pattern: System-level structures, subsystems

Architecture often has focus on system qualities as performance, scalability, robustness, security

Typical architectural patterns/styles: client server, microservice, event-based, pipe and filter

24



# This introduces new challenges when designing for robustness.

- Key ideas:
  - Provide explicit control-flow for normal and abnormal execution
    - Error handling and recovery for the latter
  - Test normal and abnormal execution
- Until now, most of the program was under our control
  - What if something goes wrong and it's not our fault? How can we make a robust system in light of this?
  - How can we test considering all the different components and dependencies?
  - What if the system is too big to test?



# This introduces new challenges when designing for robustness.

- Key ideas:
  - Provide explicit control-flow for normal and abnormal execution
    - Error handling and recovery for the latter
  - Test normal and abnormal execution
- Until now, most of the program was under our control
  - What if something goes wrong and it's not our fault? How can we make a robust system in light of this?
  - How can we test considering all the different components and dependencies?
  - What if the system is too big to test?





## What will you do if

- An API your data plugin uses is temporarily down?
  - Or returns a surprising error code?

- But, you still
   Facebook withdraws its DNS
   Mhat if Facebook mation?
   If recovery via a redirect in case of failure
  - ar data(set), fallback service, caching (e.g., store last Twitter feed)
  - If not, recruit clean-up service Ο
    - Proces, display errors



## **Retry!**

- Exponential Backoff
  - Retry, but wait exponentially longer each time
  - Assumes that failures are exponentially distributed
    - E.g., a 10h outage is extremely rare, a 10s one not so crazy

```
• E.g.:
```

```
const delay = retryCount => new Promise(resolve =>
    setTimeout(resolve, 10 ** retryCount));
```

```
const getResource = async (retryCount = 0, lastError = null) => {
  if (retryCount > 5) throw new Error(lastError);
  try {
    return apiCall();
  } catch (e) {
    await delay(retryCount);
    return getResource(retryCount + 1, e);
  }
    https://www.bayanbennett.com/posts/retrying-and-exponential-backoff-with-promises/
```



## **Retry!**

- Still need an exit-strategy
  - Learn <u>HTTP response codes</u>
    - Don't bother retrying on a 403 (go find out why)
  - Use the API response, if any
    - Errors are often documented -- e.g., GitHub will send a "rate limit exceeded" message



# Proxy Design Pattern

- Local representative for remote object
  - Create expensive obj on-demand
  - Control access to an object
- Hides extra "work" from client
  - Add extra error handling, caching
  - Uses indirection





30

# **Example: Caching**

```
interface FacebookAPI {
   List<Node> getFriends(String name);
}
class FacebookProxy implements FacebookAPI {
   FacebookAPI api;
   HashMap<String,List<Node>> cache = new HashMap...
   FacebookProxy(FacebookAPI api) { this.api=api;}
   List<Node> getFriends(String name) {
      result = cache.get(name);
      if (result == null) {
    }
}
```

```
result = api.getFriends(name);
    cache.put(name, result);
}
return result;
}
```

17-214/514



# **Example: Caching and Failover**

```
interface FacebookAPI {
    List<Node> getFriends(String name);
}
class FacebookProxy implements FacebookAPI {
    FacebookAPI api;
    HashMap<String,List<Node>> cache = new HashMap...
    FacebookProxy(FacebookAPI api) { this.api=api;}
    List<Node> getFriends(String name) {
        try {
             result = api.getFriends(name);
             cache.put(name, result);
             return result;
         } catch (ConnectionException c) {
             return cache.get(name);
```



# **Example: Redirect to Local Service**

```
interface FacebookAPI {
    List<Node> getFriends(String name);
class FacebookProxy implements FacebookAPI {
    FacebookAPI api;
    FacebookAPI fallbackApi;
    FacebookProxy(FacebookAPI api, FacebookAPI f) {
        this.api=api; fallbackApi = f; }
    List<Node> getFriends(String name) {
        try {
             return api.getFriends(name);
         } catch (ConnectionException c) {
             return fallbackApi.getFriends(name);
```



## **Testing Distributed Systems**

- Challenges:
  - Volatility
    - Users are hard to simulate
    - Real-world effects -- things crashing, delays, indicative use/data.
  - Performance
    - Massive databases? Systems with minutes-long start-up times?
    - Very common in ML



# For example:

- 3rd party Facebook apps
  - Android user interface
  - Backend uses Facebook data

How do we test this?







#### Testing in real environments



36 ISC institute for SOFTWARE RESEARCH

#### Eliminating Android dependency

Test driver Code Facebook

```
@Test void testGetFriends() {
   assert getFriends() == ...;
List<Friend> getFriends() {
   Connection c = http.getConnection();
   FacebookAPI api = new FacebookAPI(c);
   List<Node> persons = api.getFriends("john");
   for (Node person1 : persons) {
      . . .
   return result:
```

# That won't quite work

- GUI applications process *thousands* of events
- Solution: automated GUI testing frameworks
  - Allow streams of GUI events to be captured, replayed
- These tools are sometimes called *robots*





# Eliminating Android dependency



#### **Test Doubles**

- Stand in for a real object under test
- Elements on which the unit testing depends (i.e. collaborators), but need to be approximated because they are
  - Unavailable
  - Expensive
  - Opaque
  - Non-deterministic
- Not just for distributed systems!



http://www.kickvick.com/celebrities-stunt-doubles



#### How Test Doubles Help

- 1. Speed: simulate response without going through the API
- 2. Stability: guaranteed deterministic return, reduces flakiness
- 3. Coverage: reliably simulate problems (e.g., return 404)
- 4. Insight: expose internal state
- 5. Development: presume functionality not yet implemented

```
class FakeFacebook implements FacebookInterface {
   void connect() {}
   List<Node> getFriends(String name) {
        if ("john".equals(name)) {
           List<Node> result=new List();
        result.add(...);
        return result;
        }
   }
}
```

## **Types of Test Doubles**

- Most often talk about Mocks and Stubs
  - (technically other categories exist)
- Mocks give you a lot of power
  - Dictate what should be returned when (very broadly construed)
  - Requires framework using reflection
    - E.g., Mockito in Java; Mock functions in Jest\*
- Stubs are way simpler; use when possible





#### Eliminating the Remote Service Dependency



#### Introducing a Double (Stub) Facebook Mock Test driver Code Interface Facebook @Test void testGetFriends() { assert getFriends() == ...; List<Friend> getFriends() { Connection c = http.getConnection(); FacebookInterface api = new FacebookStub(c); List<Node> persor class FacebookStub implements FacebookInterface { for (Node person void connect() {} for (Node per List<Node> getFriends(String name) { ... if ("john".equals(name)) { List<Node> result=new List(); 17-2 result.add(...);

## **REST API Calls and Testing**

Test happy path

Test also error behavior!

- Correct timeout handling? Correct retry when connection down?
- Invalid response detected?
- Graceful degradation?

Need to understand possible error behavior first



#### Fallacies of distributed computing by Peter Deutsch

- 1. The network is reliable.
- 2. Latency is zero.
- 3. Bandwidth is infinite.
- 4. The network is secure.
- 5. Topology doesn't change.
- 6. There is one administrator.
- 7. Transport cost is zero.
- 8. The network is homogeneous.



#### How to test?

- 1. The network is reliable.
- 2. Latency is zero.
- 3. Bandwidth is infinite.
- 4. The network is secure.
- 5. Topology doesn't change.
- 6. There is one administrator.
- 7. Transport cost is zero.
- 8. The network is homogeneous.







#### Handle Errors Locally



17-214/514 http://christophermeiklejohn.com/filibuster/2021/10/14/filibuster-4.html





- Mocks can emulate failures such as timeouts
- Allows you to verify the robustness of system

17-214

```
class FacebookSlowStub implements FacebookInterface {
    void connect() {}
    int counter = 0;
    List<Node> getFriends(String name) {
        Thread.sleep(4000);
        if ("john".equals(name)) {
            List<Node> result=new List();
            result.add(...);
    }
}
```



17-214

```
class FacebookErrorStub implements FacebookInterface {
    void connect() {}
    int counter = 0;
    List<Node> getFriends(String name) {
        counter++;
        if (counter % 3 == 0)
            throw new SocketException("Network is unreachable");
        if ("john".equals(name)) {
            List<Node> result=new List();
            result.add(...);
            return result;
```

## **Chaos Engineering**

Experimenting on a distributed system in order to build confidence in the system's capability to withstand turbulent conditions in production





# **Design: Testability**

- Single responsibility principle
- Dependency Inversion Principle (DIP)
  - High-level modules should not depend on low-level modules; both should depend on abstractions. Abstractions should not depend on details. Details should depend upon abstractions.
- Law of Demeter: Don't acquire dependencies through dependencies.
  - o avoid: this.getA().getB().doSomething()
- Use factory pattern to instantiate new objects, rather than new.
- Use appropriate tools, e.g., dependency injection or mocking frameworks
   17-214/514



## **Designing for Robustness**

- As a *client* of distributed systems (mainly the Internet):
  - No harm trying again (redundancy)
  - Have a backup plan (resiliency)
  - Maintain awareness of what can go wrong (transparency)
    - HTTP status codes, API documentation, keeping tabs on vulnerabilities

#### • Isolation, isolation, isolation

- Use test doubles liberally
- Rely on protocols to contain and manage failures
- Never let one module crash another
  - More pointers coming up



#### Robust Distributed System Design

• Consider reading:

17-214/514

https://www.reactivemanifesto.org

- Yet another meaning for "Reactive"!
- Short guide identifying key principles
  - Goals: robustness, resilience, flexibility
  - Principles: responsiveness, elasticity, message-driven
  - Patterns/Heuristics: isolation, delegation, verbosity, replication, asynchrony



#### **Principle: Modular Protection**

- Errors should be contained and isolated
  - A failing printer should not corrupt a document
  - Handle exceptions locally as much as possible, return useful feedback
  - Don't do this:

ype Exc	eption report
message	
descript	on The server encountered an internal error that prevented it from fulfilling this request.
exceptio	
java.l	ng.NullPointerException
	nl.hu.sp.lesson1.dynamicexample.LogoutServlet.doGet(LogoutServlet.java:39)
	javax.servlet.http.HttpServlet.service(HttpServlet.java:618)
	javax.servlet.http.HttpServlet.service(HttpServlet.java:725)
	org.apache.tomcat.websocket.server.WsFilter.doFilter(WsFilter.java:52)
note Th	full stack trace of the root cause is available in the Apache Tomcat/8.0.5 logs.

17-214/514



#### Principle: Modular Protection

- Online: use HTTP response status codes effectively
  - Don't just hand out 404, 500
    - Unless they really apply
  - Provide and document fall-back options, information
    - Good RESTful design helps



institute fo



# Principle: Delegating Recovery

(Again?)

- Don't make a failing node/module serve a client
  - It needs to clean itself up
  - Forward clients to designated recovery service
    - A bit like the proxy pattern
  - Consider asynchrony
    - Failure is often expensive





#### Principle: Consider Idempotence

- Idempotency: the same call from the same context should have the same result
  - Hitting "Pay" twice should not cost you double!
  - A resource should not suddenly switch from JSON to XML
  - Makes APIs predictable, resilient



#### **Ensuring Idempotence**

- Fairly easy for read-only requests
  - Ensure consistency of read-only data
  - Never attach side-effects to GET requests\*
- Also for updates, deletes
  - Not "safe", because data is mutated
  - Natural idempotency because the target is identified
- How about writing/sending new data?

\*https://twitter.com/rombulow/status/990684463007907840





## **Ensuring Idempotence**

- How about writing/sending new data?
  - Could fail anywhere
    - Including in displaying success message after payment!
  - POST is not idempotent
  - Use <u>Unique Identifiers</u>
  - Server keeps track of requests already handled

```
curl https://api.stripe.com/v1/charges \
```

- -u sk\_test\_BQokikJOvBiI2HlWgH4olfQ2: \
- -H "Idempotency-Key: AGJ6FJMkGQIpHUTX"
- -d amount=2000  $\setminus$
- -d currency=usd  $\$
- -d description="Charge for Brandur" \
- -d customer=cus\_A8Z5MHwQS7jUmZ

https://stripe.com/blog/idempotency



#### **Distributed Systems**

There are entire courses on getting these right; not a goal here But do:

- Understand challenges and solutions to achieving robustness
  - Primarily as a *client* of a distributed system (we all are these days)
  - Test for all scenarios, leveraging test doubles
  - Provide error handling through isolation
- Learn to communicate with, and provide your own, nodes
  - API design
  - Microservices

