

15-440 Distributed Systems Recitation 3

Tamim Jabban

جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Project 1

- Involves creating a *Distributed File System (DFS)*: *FileStack*
- Stores data that does not fit on a single machine
- Enables clients to perform operations on files stored on **remote servers (RMI)**

جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar


Entities

- Three main entities in FileStack:
 - **Client:**
 - Creates, reads, writes files using RMI
 - **Storage Servers:**
 - Physically hosts the files in its local file system
 - **Naming Server:**
 - Runs at a predefined address
 - Maps file names to Storage Servers
 - Therefore, it has *metadata*

جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Architecture


- FileStack will boast a Client-Server architecture:



جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Communication


- Registration phase



جامعة Carnegie Mellon University Qatar

Communication


- Post registration, the Naming Server responds with a list of *duplicates* (if any).



جامعة Carnegie Mellon University Qatar

Communication


- System is now ready, the Client can invoke requests.



جامعة Carnegie Mellon University Qatar

Communication


- Client requests a file (to read, write etc...) from the Naming Server.



جامعة Carnegie Mellon University Qatar

Communication


- Depending on the operation, the **Naming Server** could either perform it, or, respond back to the **Client** with the **Storage Server** that hosts the file.



جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Communication

- After the **Client** receives which **Storage Server** hosts the file, it contacts that **Server** to perform the file operation.



جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Communication

- When a **Client** invokes a method, it basically invokes a **remote method** (and hence, **Remote Method Invocation**)
 - This is because the logic of the method resides on the server
- To perform this remote invocation, we need a library: **Java RMI**
- RMI allows the following:**
 - When the **client** invokes a request, it is **not aware of where it resides** (local or remote). It only knows the **method's name**.
 - When a **server** executes a method, it is **oblivious to the fact that the method was initiated by a remote client**.

جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

RMI

- The RMI library is based on two important objects:
 - Stubs:**
 - When a client needs to **perform an operation**, it invokes the method via an object called the "**stub**"
 - If the operation is **local**, the stub just calls the *helper function that implements this operation's logic*
 - If the operation is **remote**, the stub does the following:
 - Sends (marshals) the method name and arguments** to the appropriate server (or skeleton),
 - Receives the results (and unmarshals),**
 - Reports them back to the client.**

جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

RMI

- The RMI library is based on two important objects:
 - Skeletons:**
 - These are **counterparts** of stubs and reside reversely at the **servers**
 - Therefore, each **stub** communicates with a corresponding **skeleton**
 - It's responsible for:**
 - Listening** to multiple clients
 - Unmarshalling** requests (**method name & method arguments**)
 - Processing** the requests
 - Marshalling & sending results** to the corresponding stub


جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Interfaces

- Servers declare all their methods in **interfaces**
- Such interfaces contain a subset of the methods the server can perform


جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Naming Server Interfaces



جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Storage Server Interfaces




جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Creating Stubs & Skeletons

- For a client to create a **Stub**, it needs:
 - An **interface** of the corresponding **Skeleton**
 - Network address** of the corresponding **Skeleton**
- For a server to create a **Skeleton**, it needs:
 - An **interface**
 - A **class that implements the logic of the methods** defined in the given interface
 - Network address of the server


جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Naming Server Skeletons & Stubs




جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Storage Server Skeletons & Stubs




جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Storage Server Skeletons & Stubs



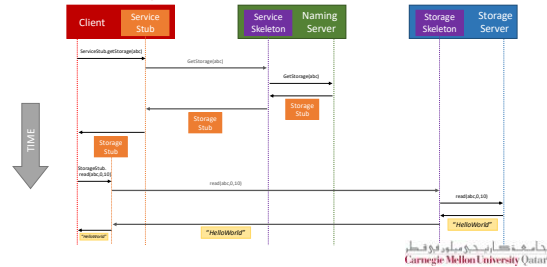
جامعة كارنيجي ميلون قطر
Carnegie Mellon University Qatar

Simple Stub-Skeleton Communication



جامعة كارنيجي ميلون
Carnegie Mellon University Qatar

Full Example: Client Read



جامعة كارنيجي ميلون
Carnegie Mellon University Qatar

Creating a Stub

- In Java, a stub is implemented as a **dynamic proxy**
- A proxy has an associated **invocation handler**
- **Example:** `getStorage` in Figure 2:
 - When `getStorage` is invoked on the **Service Stub**, the **proxy** encodes the method name (`getStorage`) and the argument(s) (file `'abc'`)
 - The proxy sends the encoded data to the **invocation handler**
 - The **invocation handler** determines if it is a **local** or **remote** procedure, and acts accordingly (as how it was shown earlier)
- Go over `java.lang.reflect.Proxy` via the JavaDocs!

جامعة كارنيجي ميلون
Carnegie Mellon University Qatar