

Lecture 6 P2P with TomP2P <http://tomp2p.net/doc> Introduction into P2P

1. Introduction

What is TomP2P
History and project information

Introduction

TomP2P
A P2P-based high performance key-value pair storage library

- **TomP2P history**
 - ▶ TomP2P v1: Created in 2004 and used for a distributed DNS project
 - ▶ This version used blocking IO operations (1 thread / socket)
 - ▶ TomP2P v2: Apache MINA (java.nio framework) / 6K LoC
 - ▶ Not well designed for non-blocking operations (event-driven)
 - ▶ TomP2P v3: Redesigned for non-blocking operations
 - ▶ Switched to Netty / 14K LoC, 6K LoC JUnits
 - ▶ TomP2P v4: API refinements, new features
 - ▶ Current release (preview) 4.1
 - ▶ Latest feature (work in progress) MapReduce
 - ▶ 22K LoC, 8K LoC JUnits

0. Lecture Overview

1. Introduction
 1. What is TomP2P
 2. History and project information
2. Example
3. Fundamental concepts
 1. XOR-based iterative routing
 2. Futures
 3. API Overview
4. Components with examples
 1. DHT with examples
 2. Tracker / PEX with examples
5. Advanced Topics
 1. NAT (UPNP/NAT-PMP)
 2. Security
 3. Replication
 4. SimGrid
 5. Direct data / persistent connections
 6. Android
6. References

Introduction

TomP2P
A P2P-based high performance key-value pair storage library

- **TomP2P is an *extended* DHT**
 - ▶ Distributed hash table concept → `put(key, value) / get(key)`
 - ▶ Extended DHT operations → `put(key1, key2, value) / add(key, value)`
- **TomP2P features (v.4.1)**
 - ▶ Java6 DHT implementation with non-blocking IO
 - ▶ Replication (direct / indirect)
 - ▶ Mesh-based distributed tracker
 - ▶ Stores multiple values for one key (examples follow)
 - ▶ Storage is memory-based or disk-based

Introduction

TomP2P
A P2P-based high performance key-value pair storage library

- **Academic background (CSG - UZH):**
 - ▶ Used in EU projects: EC-GIN, EMANICS, SmoothIT
 - ▶ Used in research projects: FastSS, LiveShift, PSH, B-Tracker, DRFS
- **<http://tomp2p.net>**
 - ▶ For questions: mailinglist (<http://lists.tomp2p.net/cgi-bin/mailman/listinfo>)
 - ▶ Specific questions: bocek-at-ifi.uzh.ch or tom-at-tomp2p.net
 - ▶ Documentation: <http://tomp2p.net/doc/> (TomP2P v4.0)
Overview: <http://en.wikipedia.org/wiki/TomP2P>
 - If something is missing, ask!
 - ▶ Development: <https://github.com/tomp2p>
 - Feature request possible if good reasons provided
- **Demo: how to setup TomP2P with Eclipse/git/maven**

2. Example

Example and Demo

Example

- Demo: a simple put / get example
- Package net.tomp2p.examples.Examples

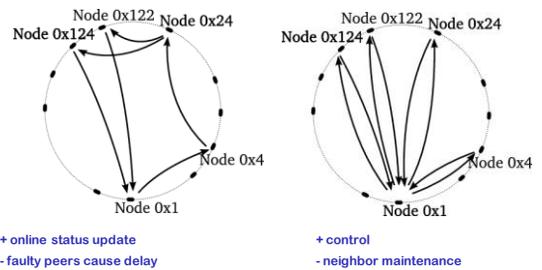
```
public static void examplePutGet(Peer[] peers) throws IOException, ClassNotFoundException
{
    Number160 nr = new Number160(rnd);
    FutureDHT futureDHT = peers[30].put(nr, new Data("hallo"));
    futureDHT.awaitUninterruptibly();
    System.out.println("peer 30 stored [key: "+nr+", value: \"hallo!\"]");
    futureDHT = peers[77].get(nr);
    futureDHT.awaitUninterruptibly();
    System.out.println("peer 77 got: \"**\" + futureDHT.getData().getObject() + \"\ for the key "+nr);
    // the output should look like this:
    // peer 30 stored [key: 0x8992a603029824e810fd7416d729ef2eb9ad3cfc, value: "hallo"]
    // peer 77 got: "hallo" for the key 0x8992a603029824e810fd7416d729ef2eb9ad3cfc
}
```

3. Fundamental Concepts

XOR-based iterative routing Futures API Overview

Fundamental Concepts

- Recursive routing vs. iterative routing

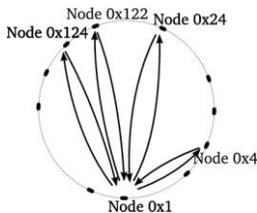


Fundamental Concepts

- TomP2P: iterative XOR-based routing
 - ▶ Node and data item unique 160bit identifier
 - ▶ Keys are located on the nodes whose node ID is closest to the key
 - ▶ Search for a key:
 - Lookup in neighbor table for closest peer (e.g. peers with ID: 0x1, 0x2, 0x3, 0x4)

My ID	Neighbor ID	Distance (XOR)
1	2	3
1	3	2
1	4	5

- Difference to Pastry: one metric, no leaf set / routing table



Fundamental Concepts

- TomP2P iterative XOR-based routing
 - ▶ Neighbors stored in 159 "bags", bag has capacity c (Kademlia, c=20)
 - ▶ Routing takes $O(\log n) \rightarrow M03$, slides 12
 - ▶ By default UDP, message header 56 bytes
 - ▶ Configuration options (RoutingConfiguration.java)
 - ▶ directHits – used for get() operations. (routing sends digest)
 - ▶ forceTCP – use TCP instead of UDP
 - ▶ maxSuccess, maxFailure – stop conditions
 - ▶ parallel – number of parallel connections
 - ▶ maxNoNewInfoDiff – stop condition. Stops if no new information was reported. Difference to minimumResults (e.g. for get(key))
 - ▶ For the CT - don't worry, default settings are fine ☺

Fundamental Concepts

- All distributed operations use futures
- Future objects
 - ▶ Keeps track of future events, while the "normal" program flow continues → `addListener()` or `await()`
 - ▶ `await()`: Operations are executed in same thread
 - ▶ `addListener()`: Operations are executed in same or other thread
- Demo: blocking operation

(`net.tomp2p.examples.Examples`)

```
public static void exampleGetBlocking(Peer[] nodes, Number160 nr)
{
    FutureDHT futureDHT = nodes[77].get(nr);
    //blocking operation
    futureDHT.awaitUninterruptibly();
    System.out.println("result: "+futureDHT.getObject());
    System.out.println("this may *not* happen before printing the result!");
}
```

Fundamental Concepts

- Demo: non - blocking operation
- ### (`net.tomp2p.examples.Examples`)
- ▶ New utilities necessary (loops as recursions)
 - ▶ Advise: use `addListener(...)` as much as possible!
 - ▶ `operationComplete(...)` must be **always** called

```
public static void exampleGetNonBlocking(Peer[] nodes, Number160 nr)
{
    FutureDHT futureDHT = nodes[77].get(nr);
    //non-blocking operation
    futureDHT.addListener(new BaseFutureAdapter<FutureDHT>() {
        @Override
        public void operationComplete(FutureDHT future) throws Exception {
            System.out.println("result: "+future.getObject());
        }
    });
    System.out.println("this may happen before printing the result!");
}
```

Fundamental Concepts

- Future utilities
 - ▶ `FutureForkJoin(int nr, boolean cancel, K... Forks)`
 - Joins already "forked" futures. Waits until all or `nr` future finished. If `nr` reached, futures may be cancelled (e.g. abort download)
 - ▶ `FutureLateJoin(int nrMaxFutures, int minSuccess)`
`FutureLateJoin()`
 - No need to add the futures in the constructor, can be added later
 - ▶ `FutureWrapper()`
 - A placeholder for futures that are created later
- ForkJoin in Java7
 - ▶ Fork and join framework – future utilities in TomP2P focus on join, forking is done "manually"

Fundamental Concepts

- Fun with futures: loops

```
Future loop() {
    Future future = new Future();
    recLoop(future);
    return future;
}

void recLoop(Future future) {
    int active = 0;
    for (int i = 0; i < parallel; i++) {
        //if future finished, it will be set to null
        if (futureResponses[i] == null) {
            active++;
            futureResponses[i] = doSomething();
        }
        else if (futureResponses[i] != null) active++;
    }
    if (active == 0) future.setDone();
    FutureForkJoin<FutureResponse> fp = new FutureForkJoin<FutureResponse>(1, futureResponses);
    fp.addListener(new BaseFutureAdapter<FutureForkJoin<FutureResponse>>() {
        @Override
        public void operationComplete(FutureForkJoin<FutureResponse> future)
            throws Exception {
            boolean finished = evaluate(future);
            if (finished) future.setDone();
            else recLoop(future);
        }
    });
}
```

Fundamental Concepts

- API Overview: `Peer.java`
 - ▶ Basic methods for DHTs
 - `put(key, value), get(key)`
 - ▶ Additional methods in TomP2P:
 - For initial connection: `bootstrapBroadcast()` / `bootstrap(IPaddress, port) / discover(IPaddress, port, port)`
 - Requires to specify `set*DataReply(...): send(peeraddress, value) / send(peerconnection, value) / send(key, value)`
 - Data manipulation: `add(key, value) / putIfAbsent(key, value) / digest(key) / remove(key)`
 - Tracker operations: `getFromTracker(key) / addToTracker(key, value)`
 - Used mostly internally `parallelRequests(...)`



Fundamental Concepts

- Configurations used in the API
 - ▶ TomP2P can store multiple values for a key
 - `put(location_key, content_key, value) → content_key` specified in Configurations
 - `get(location_key) → returns a map with [content_key, value]`
 - `add(location_key, value) → is translated to put(location_key, hash(value), value)`
 - ▶ TomP2P support domains
 - Avoid collision for same keys
 - Domains are used for protection (more details later)
 - Domains specified in Configurations
 - `put(key, domain, value) → get(key, domain)`

Fundamental Concepts

Configurations Example

```

Number160 nr = new Number160(rnd);
ConfigurationStore cs = Configurations.defaultStoreConfiguration();
cs.setDomain(Number160.createHash("my_domain"));
cs.setContentKey(new Number160(11));
FutureDHT futureDHT = peers[30].put(nr, new Data("hallo"), cs);

public static ConfigurationStore defaultStoreConfiguration()
{
    ConfigurationStore config = new ConfigurationStore();
    config.setRequestP2PConfiguration(new RequestP2PConfiguration(3, 5, 3));
    config.setRoutingConfiguration(new RoutingConfiguration(5, 10, 2));
    config.setDomain(DEFAULT_DOMAIN);
    config.setContentKey(Number160.ZERO);
    config.setStoreIfAbsent(false);
    config.setProtectDomain(false);
    config.setSignMessage(false);
    config.setRefreshSeconds(0);
    config.setAutomaticCleanup(true);
    return config;
}
    
```

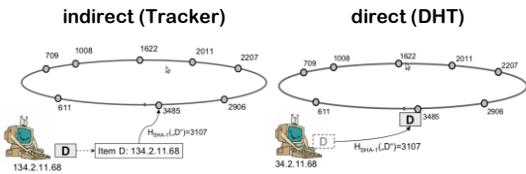
5. Components with Examples

DHT
Tracker

Components with Examples

DHT vs. Tracker

- ▶ M03, slide 23: DHT "stored by value" – direct storage
- ▶ M03, slide 24: Tracker "stored by reference" – indirect storage



Components with Examples

B-Tracker

- ▶ Centralized tracker – one peer gets traffic
- ▶ DHT: store reference on 20 peers – 20 peers gets traffic
- ▶ PEX: exchange information every minute (push)
- ▶ B-Tracker, every downloading peer becomes a tracker → forms mesh
 - Better balance of load
 - To avoid duplicates send compressed list of known peers
- ▶ B-Tracker in TomP2P enabled by default
- ▶ Currently tests with B-Tracker in Vuze

Components with Examples

Demo: Tracker with exchange of popular items (net.tomp2p.examples.ExampleTracker)

- ▶ Create 100 peers, 3 peers have initially each a song
- ▶ M03 slide 26: peer joining / bootstrap

```

11 public class ExampleTracker
12 {
13     public static void main(String[] args) throws Exception
14     {
15         peers = new ArrayList<>();
16         peers.add(new Peer(1, "134.2.11.68", 10000));
17         peers.add(new Peer(2, "134.2.11.68", 10000));
18         peers.add(new Peer(3, "134.2.11.68", 10000));
19         peers.add(new Peer(4, "134.2.11.68", 10000));
20         peers.add(new Peer(5, "134.2.11.68", 10000));
21         peers.add(new Peer(6, "134.2.11.68", 10000));
22         peers.add(new Peer(7, "134.2.11.68", 10000));
23         peers.add(new Peer(8, "134.2.11.68", 10000));
24         peers.add(new Peer(9, "134.2.11.68", 10000));
25         peers.add(new Peer(10, "134.2.11.68", 10000));
26         peers.add(new Peer(11, "134.2.11.68", 10000));
27         peers.add(new Peer(12, "134.2.11.68", 10000));
28         peers.add(new Peer(13, "134.2.11.68", 10000));
29         peers.add(new Peer(14, "134.2.11.68", 10000));
30         peers.add(new Peer(15, "134.2.11.68", 10000));
31         peers.add(new Peer(16, "134.2.11.68", 10000));
32         peers.add(new Peer(17, "134.2.11.68", 10000));
33         peers.add(new Peer(18, "134.2.11.68", 10000));
34         peers.add(new Peer(19, "134.2.11.68", 10000));
35         peers.add(new Peer(20, "134.2.11.68", 10000));
36         peers.add(new Peer(21, "134.2.11.68", 10000));
37         peers.add(new Peer(22, "134.2.11.68", 10000));
38         peers.add(new Peer(23, "134.2.11.68", 10000));
39         peers.add(new Peer(24, "134.2.11.68", 10000));
40         peers.add(new Peer(25, "134.2.11.68", 10000));
41         peers.add(new Peer(26, "134.2.11.68", 10000));
42         peers.add(new Peer(27, "134.2.11.68", 10000));
43         peers.add(new Peer(28, "134.2.11.68", 10000));
44         peers.add(new Peer(29, "134.2.11.68", 10000));
45         peers.add(new Peer(30, "134.2.11.68", 10000));
46         peers.add(new Peer(31, "134.2.11.68", 10000));
47         peers.add(new Peer(32, "134.2.11.68", 10000));
48         peers.add(new Peer(33, "134.2.11.68", 10000));
49         peers.add(new Peer(34, "134.2.11.68", 10000));
50         peers.add(new Peer(35, "134.2.11.68", 10000));
51         peers.add(new Peer(36, "134.2.11.68", 10000));
52         peers.add(new Peer(37, "134.2.11.68", 10000));
53         peers.add(new Peer(38, "134.2.11.68", 10000));
54         peers.add(new Peer(39, "134.2.11.68", 10000));
55         peers.add(new Peer(40, "134.2.11.68", 10000));
56         peers.add(new Peer(41, "134.2.11.68", 10000));
57         peers.add(new Peer(42, "134.2.11.68", 10000));
58         peers.add(new Peer(43, "134.2.11.68", 10000));
59         peers.add(new Peer(44, "134.2.11.68", 10000));
60         peers.add(new Peer(45, "134.2.11.68", 10000));
61         peers.add(new Peer(46, "134.2.11.68", 10000));
62         peers.add(new Peer(47, "134.2.11.68", 10000));
63         peers.add(new Peer(48, "134.2.11.68", 10000));
64         peers.add(new Peer(49, "134.2.11.68", 10000));
65         peers.add(new Peer(50, "134.2.11.68", 10000));
66         peers.add(new Peer(51, "134.2.11.68", 10000));
67         peers.add(new Peer(52, "134.2.11.68", 10000));
68         peers.add(new Peer(53, "134.2.11.68", 10000));
69         peers.add(new Peer(54, "134.2.11.68", 10000));
70         peers.add(new Peer(55, "134.2.11.68", 10000));
71         peers.add(new Peer(56, "134.2.11.68", 10000));
72         peers.add(new Peer(57, "134.2.11.68", 10000));
73         peers.add(new Peer(58, "134.2.11.68", 10000));
74         peers.add(new Peer(59, "134.2.11.68", 10000));
75         peers.add(new Peer(60, "134.2.11.68", 10000));
76         peers.add(new Peer(61, "134.2.11.68", 10000));
77         peers.add(new Peer(62, "134.2.11.68", 10000));
78         peers.add(new Peer(63, "134.2.11.68", 10000));
79         peers.add(new Peer(64, "134.2.11.68", 10000));
80         peers.add(new Peer(65, "134.2.11.68", 10000));
81         peers.add(new Peer(66, "134.2.11.68", 10000));
82         peers.add(new Peer(67, "134.2.11.68", 10000));
83         peers.add(new Peer(68, "134.2.11.68", 10000));
84         peers.add(new Peer(69, "134.2.11.68", 10000));
85         peers.add(new Peer(70, "134.2.11.68", 10000));
86         peers.add(new Peer(71, "134.2.11.68", 10000));
87         peers.add(new Peer(72, "134.2.11.68", 10000));
88         peers.add(new Peer(73, "134.2.11.68", 10000));
89         peers.add(new Peer(74, "134.2.11.68", 10000));
90         peers.add(new Peer(75, "134.2.11.68", 10000));
91         peers.add(new Peer(76, "134.2.11.68", 10000));
92         peers.add(new Peer(77, "134.2.11.68", 10000));
93         peers.add(new Peer(78, "134.2.11.68", 10000));
94         peers.add(new Peer(79, "134.2.11.68", 10000));
95         peers.add(new Peer(80, "134.2.11.68", 10000));
96         peers.add(new Peer(81, "134.2.11.68", 10000));
97         peers.add(new Peer(82, "134.2.11.68", 10000));
98         peers.add(new Peer(83, "134.2.11.68", 10000));
99         peers.add(new Peer(84, "134.2.11.68", 10000));
100        peers.add(new Peer(85, "134.2.11.68", 10000));
101        peers.add(new Peer(86, "134.2.11.68", 10000));
102        peers.add(new Peer(87, "134.2.11.68", 10000));
103        peers.add(new Peer(88, "134.2.11.68", 10000));
104        peers.add(new Peer(89, "134.2.11.68", 10000));
105        peers.add(new Peer(90, "134.2.11.68", 10000));
106        peers.add(new Peer(91, "134.2.11.68", 10000));
107        peers.add(new Peer(92, "134.2.11.68", 10000));
108        peers.add(new Peer(93, "134.2.11.68", 10000));
109        peers.add(new Peer(94, "134.2.11.68", 10000));
110        peers.add(new Peer(95, "134.2.11.68", 10000));
111        peers.add(new Peer(96, "134.2.11.68", 10000));
112        peers.add(new Peer(97, "134.2.11.68", 10000));
113        peers.add(new Peer(98, "134.2.11.68", 10000));
114        peers.add(new Peer(99, "134.2.11.68", 10000));
115    }
116 }
    
```

Components with Examples

Demo: Tracker with exchange of popular items

- ▶ Although demo uses `await()`, try not to use it

Demo: Store popular items in DHT (net.tomp2p.examples.ExampleDHT)

- ▶ Tracker vs. DHT what is better for the CT? You decide!

Further interesting aspects for the challenge task:

- ▶ Automate downloads
- ▶ Suggestions evaluated by the user
- ▶ How to do this more anonymous: music list from a peer is known
- ▶ Incentives
- ▶ Spamming the system with bogus suggestions

6. Advanced Topics

NAT (UPNP/NAT-PMP)
 Security
 Replication
 SimGrid integration
 Direct data connection / persistent connection
 Android

Advanced Topics

• NAT

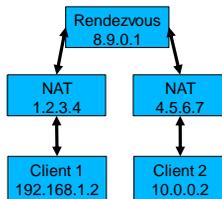
- ▶ Network Address Translation – breaks end-to-end
- ▶ “If nothing else, [NAT] can serve to provide temporarily relief while other, more complex and far-reaching solutions are worked out” (RFC 1631 - The IP Network Address Translator (NAT))
- ▶ Easy solutions: UPNP / NAT-PMP
 - Both configure port forwarding, but UPNP is more
 - UPNP: discover devices - uses broadcasting to find router (Simple Service Discovery Protocol)
 - UPNP: configure devices - uses HTTP and XML to configure portforwarding (Internet Gateway Device Protocol)
 - NAT-PMP: protocol made for configuring port-forwarding, but no discover (how to find router?)

Advanced Topics

- NAT: Difficult solution: rendezvous / relay peer which does “hole punching”, in worst case relay traffic.

• Hole punching

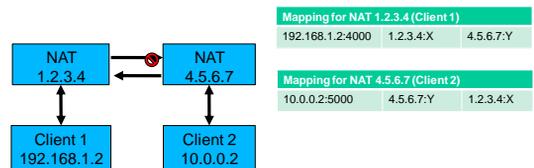
- ▶ Client 1 wants to connect to Client 2 (both clients maintain connection to Rendezvous)
- ▶ Client 1 sends connection request to Rendezvous → Rendezvous send connection request to Client 2 and the outgoing port X that Client 1 will use and send to Client 1 what outgoing port Y will be used by Client 2 (guess!)



Advanced Topics

• Hole punching

- ▶ Client 1 sends request to NAT 4.5.6.7 that will fail – no mapping, however, Client 1 creates a mapping with that request
- ▶ Client 2 send a request to Client 1 (1.2.3.4:X) – success!



Advanced Topics

• NAT example in TomP2P, the easy solution

- ▶ TomP2P supports NAT-PMP and UPNP, no holepunching or relaying
- ▶ Before bootstrap: `peer.discover(PeerAddress);`
- ▶ How it works: (1) send request how others peers sees our IP
 - If other peers sees the same IP as we see, we are fine
 - If not, we are most likely behind a NAT
- ▶ (2) do UPNP, if it fails, do NAT-PMP, if it fails, no connection
- ▶ (3) If it works test connection, send request to other peer to contact us using the port we just set up.
- ▶ (4) If we get contacted by this peer within 5 sec, port-forwarding works.
- ▶ Manual setup possible using `Bindings.java`

• No Demo, did not bring my NAT device

Advanced Topics

• Security in TomP2P

- ▶ Signature-based, no data encryption
- ▶ Messages are signed using SHA1 with DSA
- ▶ Sybil attacks!
 - This attack creates large number of identities, may collude

• How to prevent Data from being overwritten

- ▶ Domain and entry protection, requires cooperation
- ▶ `StorageGeneric.setProtection(...)`

For domains and entries		
protectionEnabled	ALL	NONE
protectionMode	NO_MASTER	MASTER_PUBLIC_KEY

Advanced Topics

• Domain protection

- ▶ Set public key `new PeerMaker(PublicKey)`
 - `Enable=ALL, Mode=NO_MASTER` → every peer can protect domains, first come first served
 - `Enable=NONE, Mode=NO_MASTER` → no peer can protect domains
 - `Enable=ALL, Mode=MASTER_PUBLIC_KEY` → every peer can protect domains, the owner can claim domain
 - `Enable=NONE, Mode=MASTER_PUBLIC_KEY` → no peer can protect domains except the owner
- ▶ Owner of domain `0x1234` is peer where `0x1234 == hash(public_key)`
- ▶ Same concept for entries
- ▶ Tracker should have no domain protection and entry protection set to `Enable=NONE, Mode=MASTER_PUBLIC_KEY` → WiP

• Demo



Advanced Topics

▶ Demo 1 (`net.tomp2p.examples.ExampleSecurity`):

- ▶ 3 peers, all with public keys
- ▶ Setup for domains: `Enable=ALL, Mode=MASTER_PUBLIC_KEY`
- ▶ (1) peer1 stores data in domain2 → success
- ▶ (2) peer3 wants to store data in domain2 → fail
- ▶ (3) peer2 wants to store data in domain2 → success

▶ Demo 2 (`net.tomp2p.examples.ExampleSecurity`):

- ▶ 3 peers, all with public keys
- ▶ Setup for domains: `Enable=NONE, Mode=MASTER_PUBLIC_KEY`
- ▶ (1) peer1 stores data in domain2 → success
- ▶ (2) peer3 wants to store data in domain2 → success
- ▶ (3) peer2 wants to store data in domain2 → success
- ▶ (4) peer3 wants to store data in domain2 → fail

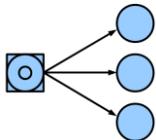


Advanced Topics

• Replication

- ▶ Enough replicas
- ▶ Direct replication
 - Originator peer is responsible
 - Periodically refresh replicas
 - Example: tracker that announces its data

- Responsible for X
- Originator of X
- Close peers to X



• Problem

- ▶ Originator offline → replicas disappear. Content has TTL, e.g. `data.setTTLSeconds(15)`

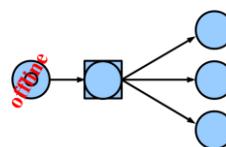


Advanced Topics

• Indirect Replication

- ▶ The closest peer is responsible, originator may go offline
 - Periodically checks if enough replicas exist
 - Detects if responsibility changes

- Responsible for X
- Originator of X
- Close peers to X



• Problem

- ▶ Requires cooperation between responsible peer and originator
- ▶ Multiple peers may think they are responsible for different versions → eventually solved



Advanced Topics

• Replication Demo (`net.tomp2p.examples.ExampleDirectReplication`)

- ▶ Direct replication – for `put()` and `add()`
 - `ConfigurationStore.setRefreshSeconds(2);`
 - Stop replication if in progress: `futureDHT.shutdown();`
- ▶ Direct replication for `remove()`
 - `ConfigurationRemove.setRefreshSeconds(2);`
 - `ConfigurationRemove.setRepetitions(2);`
 - Stop replication after 4 seconds, 2 repetitions
- ▶ Indirect replication (`net.tomp2p.examples.ExampleIndirectReplication`)
 - Set when creating peers
 - `PeerMaker.setEnableIndirectReplication(true);`
 - Two types of events: (1) peer learns about closer peer (2) peer checks frequently for enough replicas



Advanced Topics

• SimGrid integration

- ▶ Scalable simulation of distributed systems
- ▶ Publish over 100 papers that include SimGrid
- ▶ SimGrid vs. real network
- ▶ For TomP2P: simulates network with many peers
 - Defined in XML files: `platform.xml` and `deployment.xml`
- ▶ Logging in console
- ▶ Current issue in jMSG: threads, threads, threads!



• Demo: how to use it with TomP2P

- ▶ Get the Eclipse workspace: <http://tomp2p.net/dev/simgrid/> (Linux x64)
- ▶ 10'000 peers are OK, to simulate more, kernel parameter tuning



Advanced Topics

- **Direct data and persistent connections**
 - ▶ All connections in TomP2P are RPC and very short-lived
 - Open connection – request – reply – close connection
 - ▶ Direct data as seen in the tracker example → keep alive
 - ▶ Direct send (PeerAddress, ...) or with routing send(key, ...);
 - ▶ Always use setObjectDataReply() or setRawDataReply()
 - Object serializes object to byte[] (easy)
 - Raw exposes (Netty) buffer to the user for your own protocol (more work)
 - ▶ Persistent connections set by the user
 - Only for direct send send(PeerAddress, ...)
- **Demo with persistent connections**
(net.tomp2p.examples.ExamplePersistentConnection)

Advanced Topics

- **TomP2P with Android (early research)**
 - ▶ CSG: early adopter



Advanced Topics

- **TomP2P with Android ICS 4.0.3**
 - ▶ Latest Android is ~Java6 (source code) compatible
 - ▶ Extra work (permissions, IPv4)
 - ▶ TomP2P with multiple emulators
 - Redirect ports!
 - telnet to all emulators:
redir add udp:x:y
redir add tcp:x:y
 - Connect to 10.0.2.2!
- **TomP2P on Android:**
Demo with local peers



7. References

4. References

- **TomP2P homepage**
 - ▶ <http://tom2p.org>
- **Kademlia**
 - ▶ S. C. Rhea and J. Kubiatowicz: Probabilistic Location and Routing: IEEE INFOCOM 2002, New York, NY, USA, pp. 1248-1257, June 2002.
- **Sybil attack**
 - ▶ Douceur, John R. (2002). "The Sybil Attack". International workshop on Peer-To-Peer Systems. Retrieved 31 March 2011.
- **Sybil attack counter measurements**
 - ▶ Raul Landa , David Griffin , Richard G. Clegg , Eleri Mykoraist , Miguel Rio. "A Sybilproof Indirect Reciprocity Mechanism for Peer-to-Peer Networks"
- **B-Tracker**
 - ▶ Hecht, F.V.; Bock, T.; Stiller, B.; B-Tracker: Improving load balancing and efficiency in distributed P2P trackers, Peer-to-Peer Computing (P2P), 2011 IEEE International Conference on , vol., no., pp.310-313, Aug. 31 2011-Sept. 2 2011
- ▶ **ForkJoin in Java**
 - ▶ <http://docs.oracle.com/javase/7/tutorial/parallel/incremental/forjoin.html>
- ▶ **Hole punching**
 - ▶ <http://www.baycoms.ru.com/bu/holepunch/>