

Introduction to Distributed Systems

Introduction



Benoît Garbinato

distributed object programming lab

Tuesday, March 14, 2006

Distributed systems

networks distributed
"As long as there were no ~~machines~~, programming was no problem
networks distributed
at all; when we had a few weak ~~computers~~, programming became a
networks
mild problem and now that we have gigantic ~~computers~~,
distributed
programming has become an equally gigantic problem. In this
sense the electronic industry has not solved a single problem, it has
only created them - it has created the problem of using its products."

Edgster Dijkstra, The Hummel Programmer.
Communication of the ACM, vol. 15, no. 10.
October 1972. Turing Award Lecture.

Tuesday, March 14, 2006

Historical background

- Hardware became continuously cheaper
- Cheap and fast networks emerged
- The example of Unix:
 - 1969 K. Thompson & D. Ritchie develop Unix as a multi-users system on PDP-7
 - 1979 B. Joy enhances Unix with interprocess communication facilities (BSD Unix)
 - 1980's Sun Microsystems used BSD Unix as operating systems for its workstations

Approach of this course (1)

- This course teaches distributed systems from both a practical and a theoretical perspective
 - “In theory, there is not difference between theory & practice. In practice, there is.”
- The practitioner needs the theoretical perspective to understand the implicit assumptions hidden in the technologies, and their consequences
- The theoretician needs the practical perspective to validate that theoretical models, problems & solutions work in accordance to existing technologies

Approach of this course (2)

To achieve this, we will approach distributed systems through four complementary views:

- The model view
- The interaction view
- The architecture view
- The algorithm view

The model view

- What distributed entities?
E.g., processes, objects, threads, etc.
- What time assumptions?
E.g., synchronous, asynchronous, etc.
- What failure assumption?
E.g., crash-stop, malicious, etc.

The interaction view

- What interaction paradigm?
E.g., message passing, shared memory, etc.
- What reliability guarantees?
E.g., best-effort, reliable, secure, etc.

The architecture view

- What level of decentralization?
E.g., client/server, multi-tier, etc.
- What level of separation of concerns?
E.g., library-based, container-based, etc.

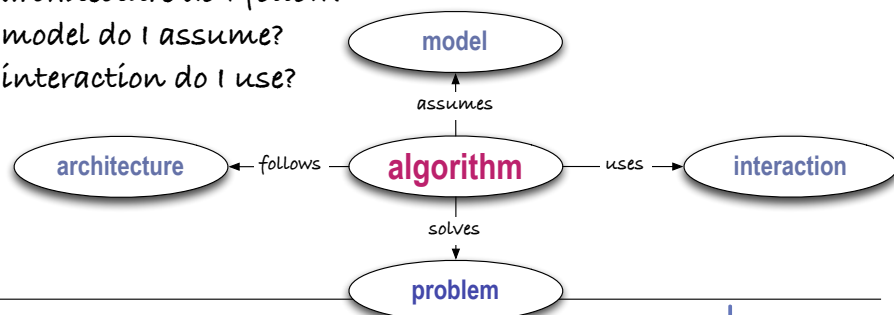
The algorithm view

- What problem?
E.g., internet payment, consensus, etc.
- What algorithm?
E.g., two phase commit, sliding window, etc.
- What complexity and what performance?
E.g., NP-complete, polynomial, etc.

The big picture

When implementing a distributed program, you will always end up writing some algorithm. In doing so, you will have to answer the following questions:

- What problem am I trying to solve?
- What architecture do I follow?
- What model do I assume?
- What interaction do I use?



Layered abstractions (1)

- Sometimes, the system you are building is (yet) another abstraction level to ease the programming distributed applications.
E.g., middleware, transactional monitor, etc.
- In this case, your problem is expressed in terms of the interaction you want to provide at your level.
- To avoid confusion, you thus have to clearly identify your origin & your target interactions, models and architectures, respectively.

Layered abstractions (2)

Assume you want to devise an algorithm implementing remote procedure calls

- Target interaction: remote procedure call
- Target model: partially synchronous crash-stop
- Target architecture: client/server (middleware-level)

- Origin interaction: unreliable message passing (e.g., UDP)
- Origin model ↔ Target model
- Origin architecture: peer-to-peer (os-level)

Technologies in this course

- The Java Programming platform
- Internet protocols (TCP, UDP)
- Unix (Linux, Mac OS X, etc.) or Windows

Introduction © Benoît Garbinato



Tuesday, March 14, 2006

Content & calendar

	10:00 - 12:00	12:00 - 13:00
March 14	Introduction Java in a nutshell	Get familiar with Java & lab tools
March 21	Concurrent Programming	Concurrent Chat
March 28	Remote Method Invocation	
April 4	Asynchronous Messaging	RMI Chat
April 11		
April 25	Network Programming	JMS Chat
May 2	Distributed Algorithms	Broadcast Chat
May 9		
May 16	Tales from the academic world	
May 23	Tales from the real world	
May 30	Q & A	
June 6		
June 13		
June 20		
Legend:	Course	
	Exercise	

Introduction © Benoît Garbinato



Tuesday, March 14, 2006

Course form

- Each Tuesday:
 - from 10 to 12 : principles
 - from 12 to 13 : exercises
- Evaluation :
 - Written final exam
 - Bonus based on exercises


Introduction © Benoit Garbinato




Tuesday, March 14, 2006

Exercises


You will start from a concurrent application and you will distribute it using various programming abstractions, e.g., remote method invocations, sockets, message-oriented middleware, etc.



client-side business logic



programming abstractions



server-side business logic

Introduction © Benoit Garbinato



Tuesday, March 14, 2006

For further information

- <http://lpdwww.epfl.ch/teaching/ids.html>
- sebastien.baehni@epfl.ch
- bastian.pochon@epfl.ch
- [newsgroup: epfl.ssc.ids](mailto:epfl.ssc.ids)
- benoit.garbinato@unil.ch
- <http://www.hec.unil.ch/dop>

Introduction © Benoit Garbinato



Tuesday, March 14, 2006

Questions?

Tuesday, March 14, 2006