

Fault Tolerance in Distributed Systems: An Introduction

Distributed Systems
Sistemi Distribuiti

Andrea Omicini
andrea.omicini@unibo.it

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
ALMA MATER STUDIORUM – Università di Bologna a Cesena

Academic Year 2015/2016

These Slides Contain Material from [TvS07]

Slides were made kindly available by the authors of the book

- Such slides shortly introduced the topics developed in the book [TvS07] adopted here as the main book of the course
- Some of the material from those slides has been re-used in the following, and integrated with new material according to the personal view of the teacher of this course
- Every problem or mistake contained in these slides, however, should be attributed to the sole responsibility of the teacher of this course

Outline

1 Introduction

2 Basic Concepts



Failure in Distributed Systems

Partial failure

- A typical feature of distributed systems is the notion of *partial failure*
- One component may fail, while the rest of the systems keeps running
- While the functionality guaranteed by the failed component is compromised, this does not necessarily holds for the other components, as well as for the overall system

Engineering distributed systems with failure

- When engineering a distributed systems, a twofold goal is possible
 - reducing the impact of failure of a single component on the others, and on the overall system performance
 - exploiting partial failure to recover from failure

Outline

1 Introduction

2 Basic Concepts

Dependable Systems

Main features of dependable systems

- Availability
- Reliability
- Safety
- Maintainability

Dependability is closely related to fault tolerance

Availability

Definition

Availability refers to the property that a system is ready for immediate use

This means...

- ... that availability refers to the probability that a system is operating correctly at any given moment, ready to provide users with its functions
- So, a highly-available system is a system that is most likely to be ready and working at any given instant of time

Reliability

Definition

Reliability refers to the property that a system can run continuously without failure

This means...

- ... that reliability is defined in terms of a time interval, rather than of a instant – as in the case of availability
- So, a highly-reliable system is a system that is most likely to keep on running for a long period of time

Safety

Definition

Safety refers to the situation that when a system temporarily fails to operate correctly, nothing catastrophic happens

This is...

- ... a very difficult property to be defined, and to be ensured as well

Maintainability

Definition

Maintainability refers to how easily a failed systems can be repaired

This means...

- ... that maintainability is closely related to availability
- So, a highly-maintainable system may also show a high degree of availability

Faults I

Failure

- A system is said to *fail* when does not behave as promised
- An *error* is a part of a system state that might have caused a failure
- The cause of an error is a *fault*

Faults II

Fault tolerance and dependable systems

- Building a dependable system closely relates to controlling faults
- One may distinguish between
 - preventing faults
 - removing faults
 - forecasting faults
- In distributed system, the most important issue is *fault tolerance*
- as the property of a system to provide its function even in the presence of faults

Faults III

Sorts of faults

- *Transient faults* occur once then disappear
- *Intermittent faults* occur, vanishes of its own accord, then reappears, and so on
- *Permanent faults* keep on existing until the faulty component is replaced /fixed

Failure Models

Type of failure	Description
Crash failure	A server halts, but is working correctly until it halts
Omission failure <i>Receive omission</i> <i>Send omission</i>	A server fails to respond to incoming requests A server fails to receive incoming messages A server fails to send messages
Timing failure	A server's response lies outside the specified time interval
Response failure <i>Value failure</i> <i>State transition failure</i>	A server's response is incorrect The value of the response is wrong The server deviates from the correct flow of control
Arbitrary failure	A server may produce arbitrary responses at arbitrary times

Different types of failures
[TvS07]

Failure Masking by Redundancy

Idea

- Hiding failures from other processes
- The key technique for masking faults is *redundancy*

Three kinds of redundancy

- Information redundancy
 - e.g., extra bits
- Time redundancy
 - e.g., redos after transaction aborts
- Physical redundancy
 - typical in biological systems

References I



Andrew S. Tanenbaum and Marteen van Steen.
Distributed Systems. Principles and Paradigms.
Pearson Prentice Hall, Upper Saddle River, NJ, USA, 2nd edition,
2007.

Fault Tolerance in Distributed Systems: An Introduction

Distributed Systems
Sistemi Distribuiti

Andrea Omicini
`andrea.omicini@unibo.it`

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
ALMA MATER STUDIORUM – Università di Bologna a Cesena

Academic Year 2015/2016

