

Distributed Systems Debugging

Jiří Patera

University of West Bohemia in Pilsen
Czech Republic



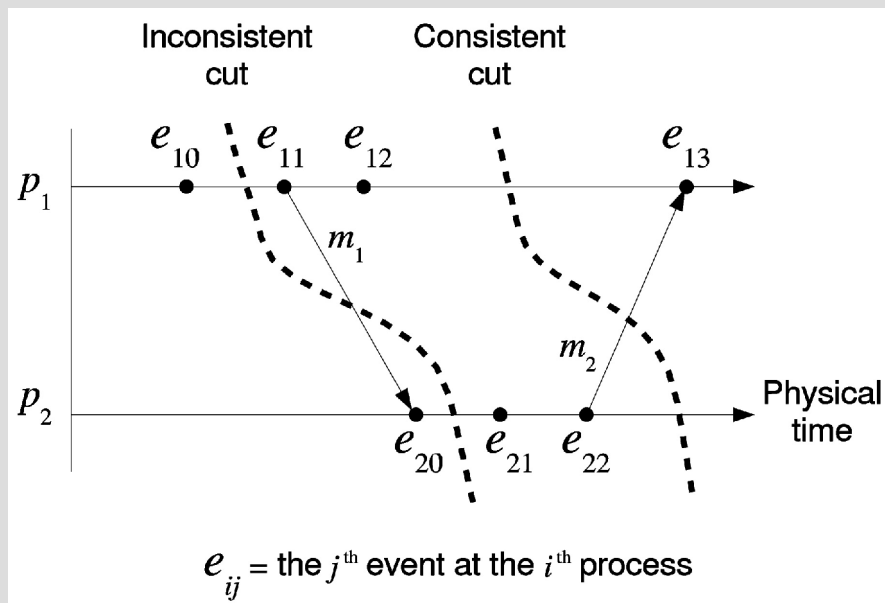
Overview

- Introduction
- Model of a distributed computation
- Predicates and their connection to CTL*
- Global state lattice
- Predicate detection algorithms
- On-line vs. off-line debugging
- Summary
- Further Work

Introduction

- What is a *distributed system*?
 - A system with components that must co-operate on some task and which are distributed over a private or public network.
- Drawbacks of such systems are:
 - Absence of *global time*
 - Absence of *global state*

Consistency



- Inconsistent cut
- Orphan message
- Consistent cut
- Consistent state

Predicates

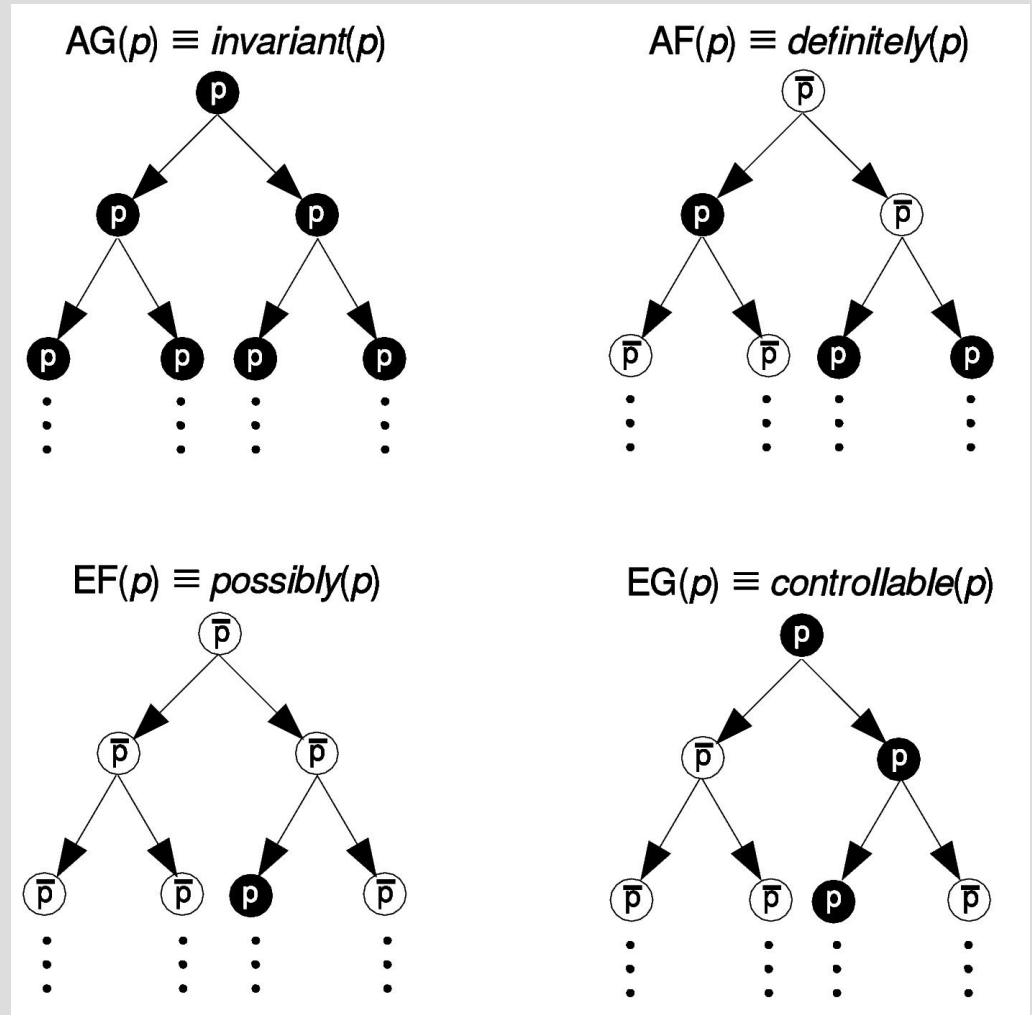
- Detecting a boolean condition in a particular node of the system equals to evaluating a *local state predicate*.
- Detecting a condition in a distributed system equals to evaluating a *global state predicate*.
- There are four classes of predicates:
 - Stable, Observer-Independent, Linear, and Semi-Linear

CTL* Temporal Logic

- It provides operators for the description of how the truthfulness of predicates changes with time.
- CTL = Computation Tree Logic
- CTL* is an extended variation of CTL
- Two path quantifiers are: **A** (all paths), **E** (there exists a path)
- Five temporal operators are: **F** (future), **G** (global), **U** (until), **R** (release), **X** (next).

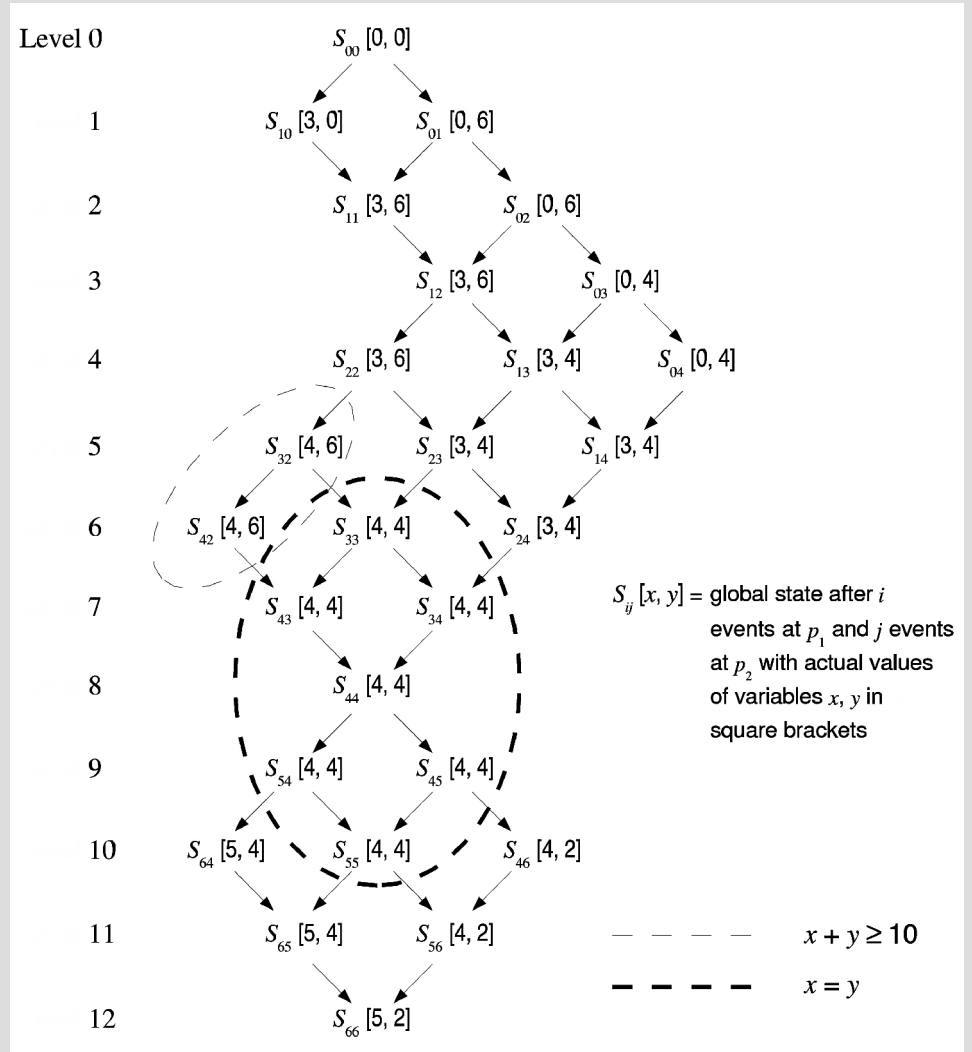
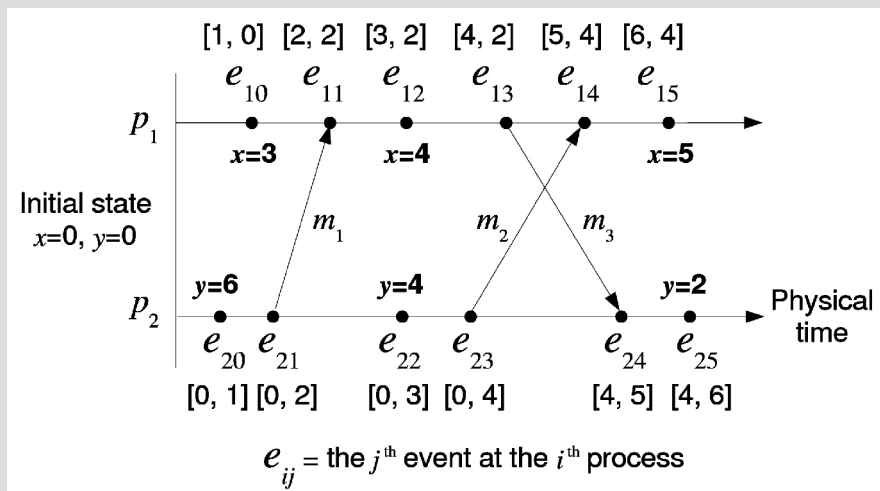
Predicates and CTL*

- Four operators are used to describe the predicate detection problem.
- Connection of predicates and CTL* operators.



Global State Lattice

- It is composed of consistent global states only.



Predicate Detection Algorithms

- A review of references to existing algorithms is presented in the paper.
- It points out that the efficient algorithms for the detection of **regular** and **linear** predicates under the **AF** temporal operator have not been published yet.

Model of Observation

- The detection of a global predicate is divided among a ***checker*** and ***non-checker processes***.
 - Non-checker processes send messages to the checker. They are situated at computation nodes and monitor local variables and messages.
 - The checker process searches for a consistent cut that satisfies a predicate.

Model of Control

- Every process has associated with it a *supervisory process*.
 - This process observes the underlying process and controls it by delaying, disabling, or by changing the order of incoming and outgoing messages.
- Useful for testing the system under some desired conditions.

On-line vs. Off-line

- On-line debugging
 - We have no information on the future of the computation, we know only the past.
 - Thus, we cannot predict (e.g. deadlocks).
- Off-line debugging
 - We have all information not only on the past of the computation but also on its future.
 - So we can predict future (i.e. The things that have happened in the recorded computation).

Summary

- Global state lattice checking approach
 - Advantage is that it checks all possible linearizations.
 - Drawback is that its complexity grows exponentially with the number of processes.
- The record/replay approach
 - Advantage is that it is not as complex as the previous one.
 - Drawback is that it checks only one particular linearization.

Further Work

- Detailed investigation and analysis of approaches to the debugging of distributed Java programs.
- Proposal of a new method for the debugging of distributed programs, a method which is based on a modification of the Java bytecode.
- Experimental implementation and testing of the proposed method.

Thank You

Thank you for your attention...

Now it is time for your questions.