

The ROme OpTimistic Simulator: A Tutorial



SAPIENZA
UNIVERSITÀ DI ROMA

Alessandro Pellegrini
Francesco Quaglia

High Performance and Dependable
Computing Systems Group
Sapienza, University of Rome

PADABS 2013

Discrete Event Simulation (DES)

- A simulation model is described in terms of:
 - *Simulation state*, describing the current state of the system
 - *Events*, associated to particular actions/changes in the system
 - *State transitions*, which modify the state depending on the executed events
- Based on event-driven programming
 - Events are dispatched to the associated event handlers which implement the model's logic
- A *discrete event* occurs at an instant in time, producing a change in the system
- DES represents the operation of a system as a chronological sequence of events
- An event cannot be scheduled in the past

DES Building Blocks

- **Clock**

- Keep track of the current simulation time (independently of the measuring unit)
- Being *discrete*, time hops to the next event's time

- **Events List**

- At least the *pending event set* must be maintained by the simulation architecture
- Events can arrive at a higher rate than they can be processed

- **Random-Number Generators**

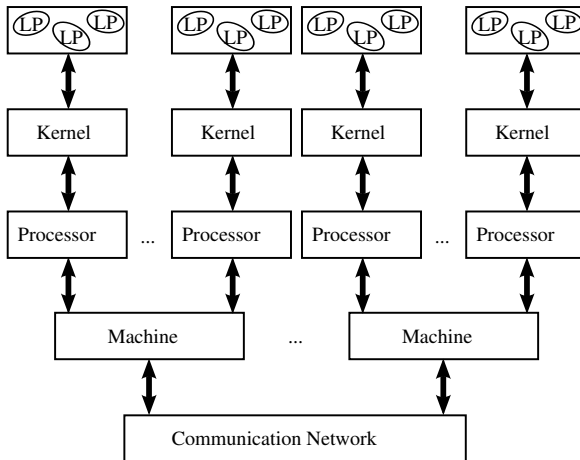
- Simulation often rely on distributions, in order to model real world's aspects

- **Statistics Collection**

- **Ending Condition**

- Real systems can often run forever, so the designer of the model must decide when the simulation will halt

Going *parallel*: PDES Logical Architecture



Going *parallel*: PDES Object Model

- Simulation model is partitioned into *simulation objects*
- Simulation objects model a portion of the space and/or agents
- Disjoint States: Message Passing to represent interactions

$$S = \bigcup_{i=1}^{numLP} S_i \quad \forall i, j \ i \neq j : S_i \cap S_j = \emptyset$$

- Logical Processes (LP) implement event handlers, used to dispatch events to simulation objects
- A simulation kernel schedules the execution of a particular LP for processing an event at a certain simulation object

Going *parallel*: Optimistic PDES


- Events are executed speculatively: processed events can be committed or uncommitted
- The *commitment horizon* is associated with GVT value
- Relatively independent of lookahead
- Resource utilization approaches 100%
- It is faster than the critical path

The Synchronization Problem

- The greatest opportunity arises from processing events from different LPs concurrently on different processors
- Is *correctness* always ensured?

The Synchronization Problem

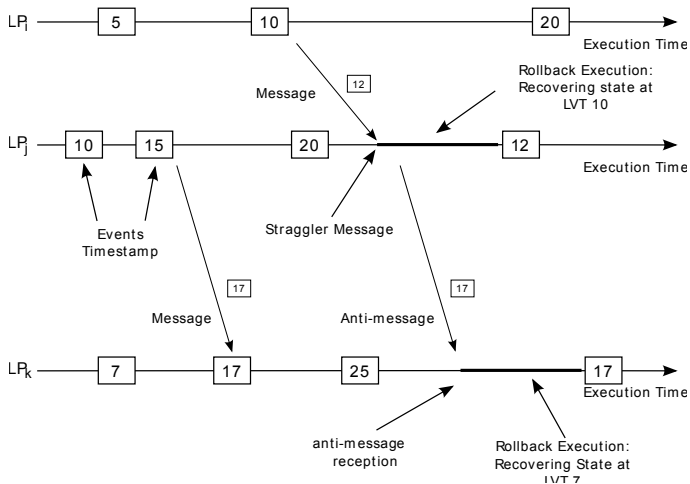
- The greatest opportunity arises from processing events from different LPs concurrently on different processors
- Is *correctness* always ensured?



Optimistic Synchronization: Time Warp

- There are no state variables that are shared between LPs
- Communications are assumed to be reliable
- Messages might not be received in timestamp order
- **Local Control Mechanism**
 - Events not yet processed are stored in an *input queue*
 - Events already processed are not discarded
- **Global Control Mechanism**
 - A-posteriori detection of causality violation
 - Event processing can be **undone** (*rollback*)
 - Reverse computation
 - Simulation State Checkpoint/Restore

Rollback



The ROME OpTimistic Simulator (ROOT-Sim)



- Simulation Platform built according to the Time Warp Synchronization Protocol
- Supports ANSI-C programming
- Targets both simulation efficiency and model development transparency
- Comes bundled as a static library

<http://www.dis.uniroma1.it/~hpdcs/ROOT-Sim/>

ROOT-Sim API: Objects/Model Description

- The ROOT-Sim API is based on a reduced set of call/callback functions:
 - `ProcessEvent()` (callback) – Used to give control to the simulation model (to a LP). It passes an event to be dispatched to some simulation object
 - `ScheduleNewEvent()` (call) – Allows to inject a new event in the system
 - `OnGVT()` (callback) – Used to perform analysis on a committed state, and for termination detection
- The simulation model is written in ANSI-C
- The LPs' simulation state is stored into dynamically-allocated memory
- The simulation is started via a special INIT event

An Example Simulation Model: Data Definition

```
1 #include <ROOT-Sim.h> // The ROOT-Sim header file
2
3 #define PACKET 1 // Event definition
4 #define DELAY 120
5 #define PACKETS 1000000 // Termination condition
6
7 typedef struct _event_content_t {
8     time_type sent_at;
9 } event_content_t;
10
11 typedef struct _lp_state_t{
12     int packet_count;
13 } lp_state_t;
```

An Example Simulation Model: Event Processing

```
1 void ProcessEvent(unsigned int me, time_type now, unsigned int
   event, event_t *content, unsigned int size, lp_state_t *state){
2
3     event_t new_event;
4     time_type timestamp;
5
6     switch(event) {
7
8         case INIT: // must be ALWAYS implemented
9             state = (lp_state_t *)malloc(sizeof(lp_state_t));
10            state->packet_count = 0;
11            timestamp = (time_type)(20 * Random());
12            ScheduleNewEvent(me, timestamp, PACKET, NULL, 0);
13            break;
14
15
```

An Example Simulation Model: Event Processing (2)

```
16     case PACKET: {
17         state->packet_count++;
18         new_event_content.sent_at = now;
19         int recv = FindReceiver(TOPOLOGY_MESH);
20         timestamp = now + Expent(DELAY);
21         ScheduleNewEvent(recv, timestamp, PACKET, &new_event,
22                         sizeof(new_event));
23     }
24 }
```

An Example Simulation Model: Termination Detection

```
1 bool OnGVT(unsigned int me, lp_state_t *snapshot) {  
2     if (snapshot->packet_count < PACKETS)  
3         return false;  
4     return true;  
5 }
```


ROOT-Sim Facilities

Already Released:

- Simulation state can be scattered over dynamically allocated memory
- Supports Full, Incremental and Autonomic Logging
- Consistent and Committed Global State management
 - Termination detection
 - Consistent statistics collection
- Transparently-rollbackable statistical library
- Topology library

ROOT-Sim Facilities

Already Released:

- Simulation state can be scattered over dynamically allocated memory
- Supports Full, Incremental and Autonomic Logging
- Consistent and Committed Global State management
 - Termination detection
 - Consistent statistics collection
- Transparently-rollbackable statistical library
- Topology library

Currently under development:

- Shared state management (via global variables)
- Consistent output streams management
- Load Balancing
- Load Sharing (via multithreading)

It's now time...

...for a live demo!



Thanks for your attention

Questions?

pellegrini@dis.uniroma1.it

<http://www.dis.uniroma1.it/~pellegrini>

<http://www.dis.uniroma1.it/~ROOT-Sim>