

DiVinE and DiVinE within

J. Barnat, L. Brim, I. Černá, P. Šimeček, ...



- Introduction
- Programmer's point of view
- User's point of view
- Future plans

Several distributed LTL Model-Checkers

- implemented using various tools
 - spin, maso, diks, ...
- difficult to be used by other users
- incomparable performance
- solve many common problems

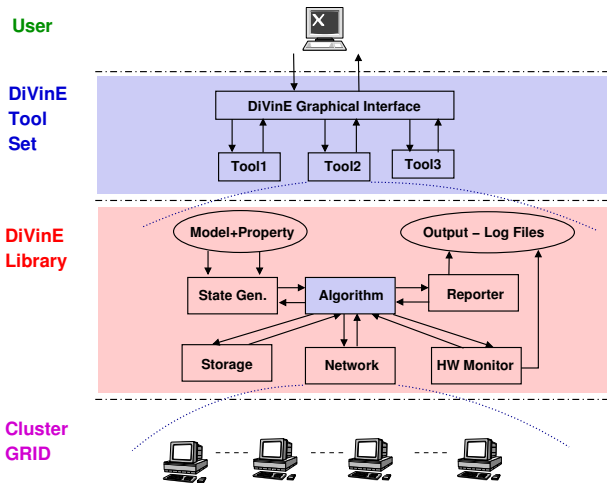
DIVINE

- Distributed Verification Environment

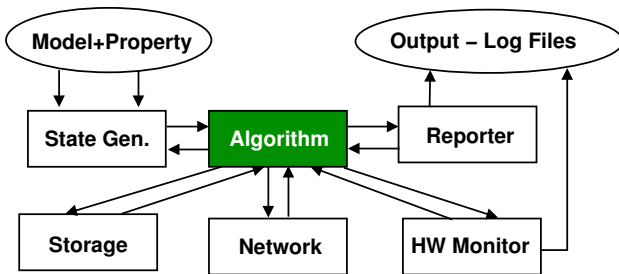
Goals

- Distributed enumerative model-checker
- Development environment
- Platform for experimental evaluation
- Research vehicle

DiVinE Structure



DiVinE from programmer's point of view




Box “algorithm”

- compute the model-checking task
- control computation
- call DiVinE Library functions

Other boxes

- arms and legs of the algorithm
- provide more than 100 useful functions



Graph of synchronous product automaton

- `get_initial_state()`
- `get_succs()`
- `is_accepting()`

Access to inner structure of the model


- partial order reduction
- property automaton decomposition
- static analysis

DiVinE native modeling language

- another modeling language
- nobody wants to learn

Other modeling languages

- separate system class
- methods to test system abilities
 - `can_property_process()`
- Promela (NIPS project)



What is it a State

State

- a piece of memory given by generator
- dynamic size

Appendix

- constant-sized piece of information
- associated with every state

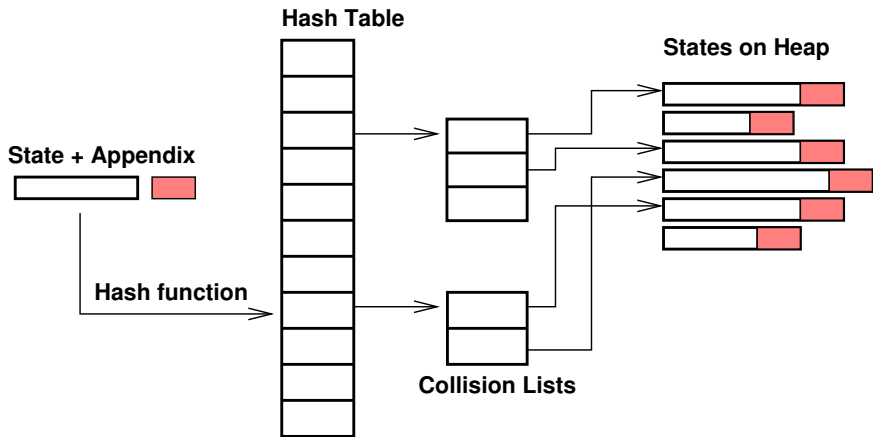
State management

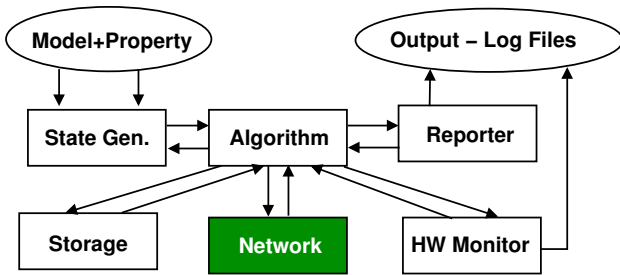
- states organized using hash table
- standard state manipulation functions
- 8 byte state reference

State compression

- no compression
- static Huffman's encoding

Storage structure





Basic network primitives

- send (urgent) message
- barrier synchronization

Receives

- procedure to process user messages
- `process_messages()`

Distributed termination detection

- Safra's algorithm
- busy/idle state
- performed within `process_messages()`
- test for being synchronized
- can exchange data within synchronization
- repeatable

Others

- additional buffers
- partition function
- network statistics
- wrapper for direct network access

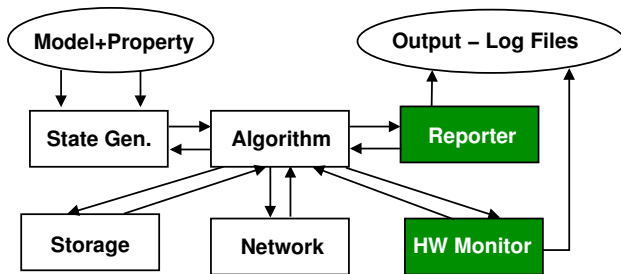
Network Support – Example

```
void process_message(char *buf, int size, int src, ...) {
    state_t state = new_state(buf, size);
    if (!Storage.is_stored(state) {
        Storage.insert(state, state_ref);
        Queue.push(state_ref);
        Distributed.set_busy(); } }
```

```
Distributed.process_user_message = process_message;
```

```
state_t state = System.get_initial_state();
if ( Distributed.partition_function(state) == my_id) {
    Storage.insert(state, state_ref);
    Queue.push(state_ref); }
```

```
while (!Distributed.synchronized()) {
    Distributed.process_messages();
    while (!Queue.empty()) {
        state_ref = Queue.top(); Queue.pop();
        state = Storage.reconstruct(state_ref);
        System.get_succs(state, succs);
        for (size_int_t i=0; i!=succs.size(); ++i) {
            int owner = Distributed.partition_function(succs[i]);
            Distributed.send_message(succs[i].ptr, succs[i].size, ...);
        }
    }
    Distributed.set_idle(); }
```



Type of output

- final result and statistics
- runtime statistics
- error messages

Algorithm dependency

- dependent
 - queue sizes, number of iterations, ...
- independent
 - number of stored states, sent messages, ...

Idea

- stdout belongs to algorithm
- other outputs written to files

Advantages and disadvantages

- persistent
- unified format
- further processing
- slowdown

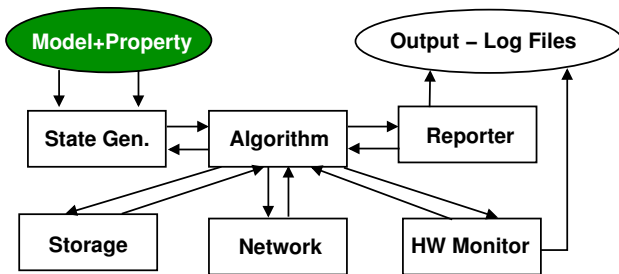
Final report

- single file
- produced before network is finalized

Logfiles

- one file per each workstation
- produced during computation
- POSIX signal + UNIX alarm utility

DiVinE from user's point of view



Processes

- extended FA
- transitions with *guards*, *sync*, *effects*

Interprocess communication


- shared variables
- buffered/unbuffered channels

System

- synchronous, asynchronous
- property Büchi automaton

DiVinE Native Modeling Language – Example

```
...
process cabin
{
  state idle ,mov,open;
  init idle;
  trans
  idle -> mov   {guard v>0;},
  mov -> open   {guard t==p;},
  mov -> mov    {guard t<p; effect p=p-1;},
  mov -> mov    {guard t>p; effect p=p+1;},
  open -> idle  {effect req[p]=0,v=0;};
}
...
system async property LTL_negative_claim;
```



How to make it work

- download from our website
- compile
- run (`divine.*`)

Prerequisites

- Linux cluster
- MPI

divine.owcty

DiVinE Tool Set

OWCTY version 1.0 build 4 (2005/09/21 17:14)

Usage: [mpirun -np N] divine.owcty [options] input_file

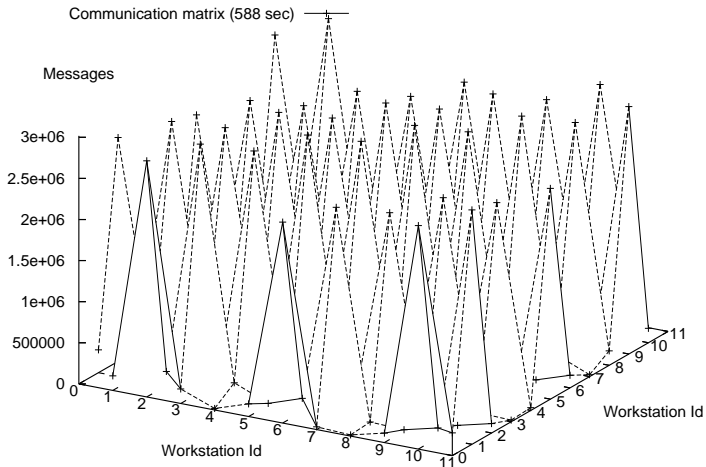
Options:


-V,--version	show version
-h,--help	show this help
-H x,--htsize x	set the size of hash table to (x<33 ? 2^x : x)
-v,--verbose	print some statistics
-q,--quiet	quite mode
-t,--trail	produce trail file
-r,--report	produce report file
-s,--simple	perform simple reachability only
-L,--log	produce logfiles (log period 1 sec)
-X w	sets base name of produced files to w (w.trail ,w.report ,w.00-w.N)

DiVinE ToolSet – More than Algorithms

- Reachability
 - deadlocks, goal states, unreachable code
- Simulator
- `divine.ltl2buchi`
- Utility to draw state-space
- Utility to visualize logfiles
- Initial set of parametrized models

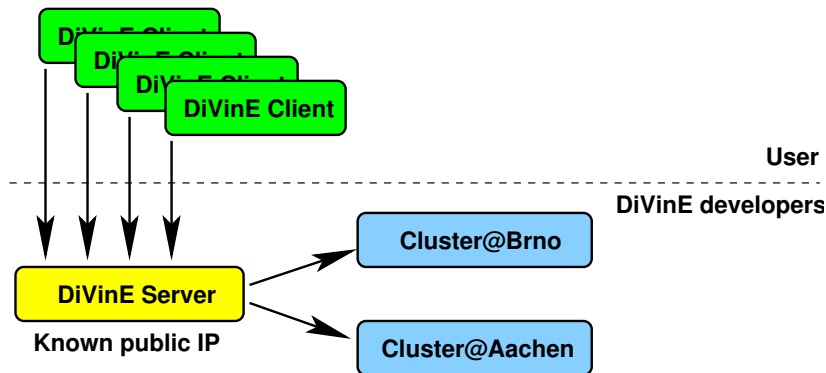
DiVinE ToolSet – Visualized Output Files





Graphical User Interface (GUI)

Server-client application



A very short demo

The screenshot displays the DiVinE application window titled "DiVinE -- user jirik at localhost". The interface is divided into several sections:

- model:** A tree view on the left showing "New property", "New task", and "owcty_reversed".
- Select clusters and computers:** A section with three clusters: "ParaDise laboratory cluster" (1 computer, 21 computers in cluster), "Parsecs cluster" (1 computer, 5 computers in cluster), and "localhost" (1 computer, 2 computers in cluster). Each cluster has a "See details" button.
- Select algorithms:** A list of algorithms with checkboxes: "Token based nested depth first search", "Property driven nested depth first search", "DepS-based driven nested depth first search", "OWCTY reversed" (checked), "Negative cycle detection", and "Back-Level-edge-based cycle detection".
- Select when task is finished:** Radio buttons for "When all algorithms finish" (selected), "When first algorithm finishes", and "When specified algorithms finish" (with a "Choose..." button).

A floating window titled "Cluster ParaDise laboratory cluster" is open, showing "21 of 22 computers currently available." It contains a table with the following data:

Computer Name	Progress
psyche01	0%
psyche02	0%
psyche06	0%
psyche11	0%
psyche14	0%
psyche15	0%
psyche16	0%
psyche17	0%
psyche18	0%
psyche19	0%
psyche20	0%
psyche22	0%
psyche07	0%
psyche03	1%
psyche12	1%
psyche08	2%
psyche04	2%
psyche10	5%
psyche13	5%
psyche05	50%
psyche21	50%
psyche09	unavailable

The Windows taskbar at the bottom shows the system tray with the date and time: "Pá, 11. lis, 13:32".

Future plans

DiVinE as a tool

- stabilize GUI
- extend functionality

DiVinE for programmers

- improve design of library
- optimize implementation
- documentation
- develop and implement new ideas

Dynamic load-balancing

- memory occupation
- work load
- network load

Known techniques

- states partitioning/repartitioning
- queue balancing

Idea

- states to be explored on highly-loaded workstations explored on less-loaded workstations

Problem

- how to access appendix on remote workstations?

`http://anna.fi.muni.cz/divine`

