ICAPS 2011

IPPC Results Presentation

Scott Sanner Sungwook Yoon

NICTA

THE AUSTRALIAN NATIONAL UNIVERSITY

parc
A Xerox Company

Additional domain development by Tom Walsh (ASU)

Main Objective for IPPC 2011

- **More realistically motivated problems**
  - PPDDL cannot represent many probabilistic domains
    - Traffic Control
    - Elevator Control
    - Mars Rovers
  - Needed
    - → concurrency
    - → independent exogenous effects
    - → continuing processes and non-goal rewards
    - → partial observability
    - → distributions that are complex function of state
    - → enumerated, integer, continuous variables **(no competitors)**
  - Required a new language
    - RDDL (new lifted DBN transition semantics)

A Brief History of (ICAPS) Time

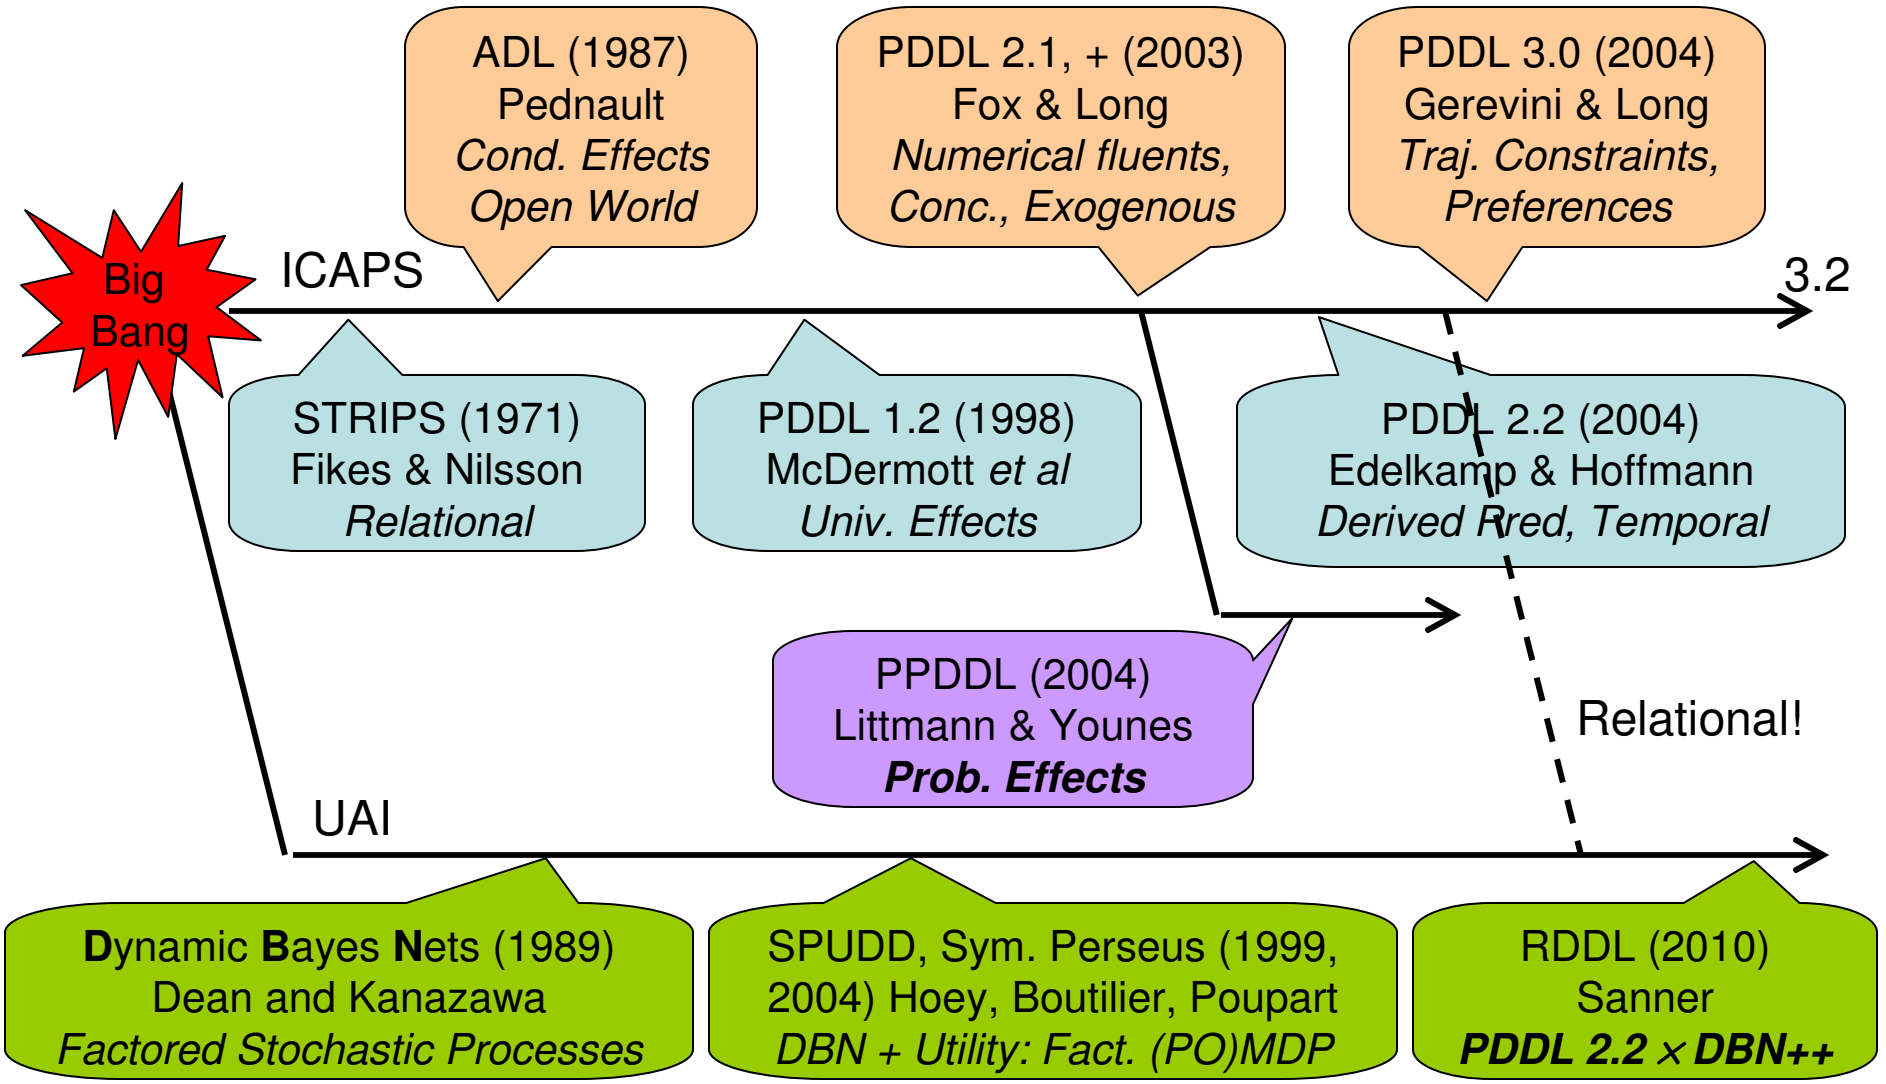

PDDL history from: http://ipc.informatik.uni-freiburg.de/PddlResources

What is RDDL?

- Relational Dynamic Influence Diagram Language
  - Relational [DBN + Influence Diagram]
  - Everything is a fluent!
    - states
    - observations
    - actions
    - derived (stochastic) predicates
  - Uniform expression language

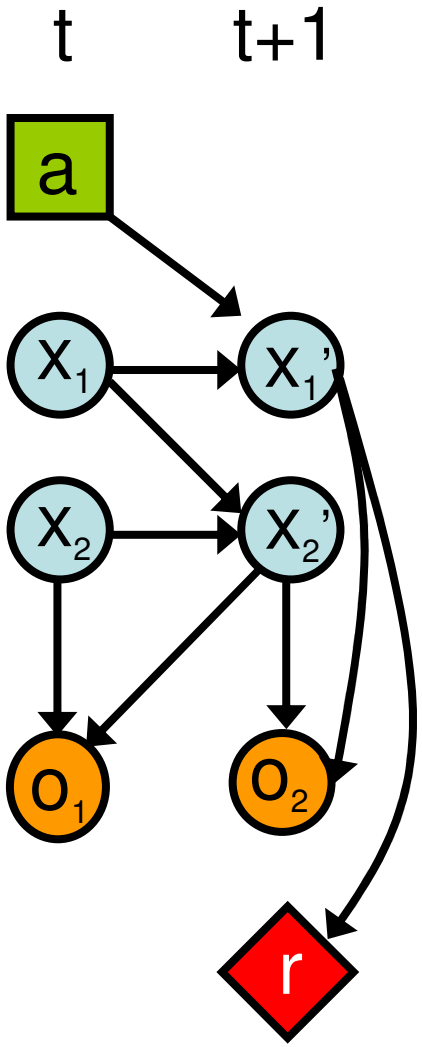

Other Objectives for IPPC 2011

- Translations to draw in different communities
  - Factored MDP / POMDP community
  - ICAPS PPDDL community
  - 11 competitors!
- Single normalized evaluation criteria
  - In previous competitions, a mix of…
    - plan length
    - goal %
    - planner time

    (skipping hard problems could improve domain averages)

RDDLSim Software

Open source & online at
http://code.google.com/p/rddlsim/

RDDL Software Overview

- BNF grammar and parser
- Simulator
- Automatic translations
  - LISP-like format (easier to parse)
  - SPUDD & Symbolic Perseus (boolean subset)
  - Ground PPDDL (boolean subset)
- Client / Server
  - Java and C/C++ sample clients
  - Evaluation scripts for log files
- Visualization
  - DBN Visualization
  - Domain Visualization – *see* how your planner is doing

Domains and Evaluation

- 8 domains
  - **Traffic Control:** highly exogenous, concurrent
  - **Elevator Control:** highly exogenous, concurrent
  - **Game of Life:** highly combinatoric
  - **SysAdmin:** highly exogenous, complex transitions
  - **Navigation:** goal-oriented, determinization killer
  - **Crossing Traffic:** goal-oriented, deterministic if move far left
  - **Skill Teaching:** few exogenous events
  - **Reconnaissance:** few exogenous events

- Conditions
  - 24 hours for all runs
  - 10 instances per domain, 30 runs per instance
  - No discount, finite horizon of 40

- Used average normalized score [0,1]
  - Min: max(random policy, noop policy)
  - Max: best competitor
  - Scores < 0 set to 0

Boolean Traffic

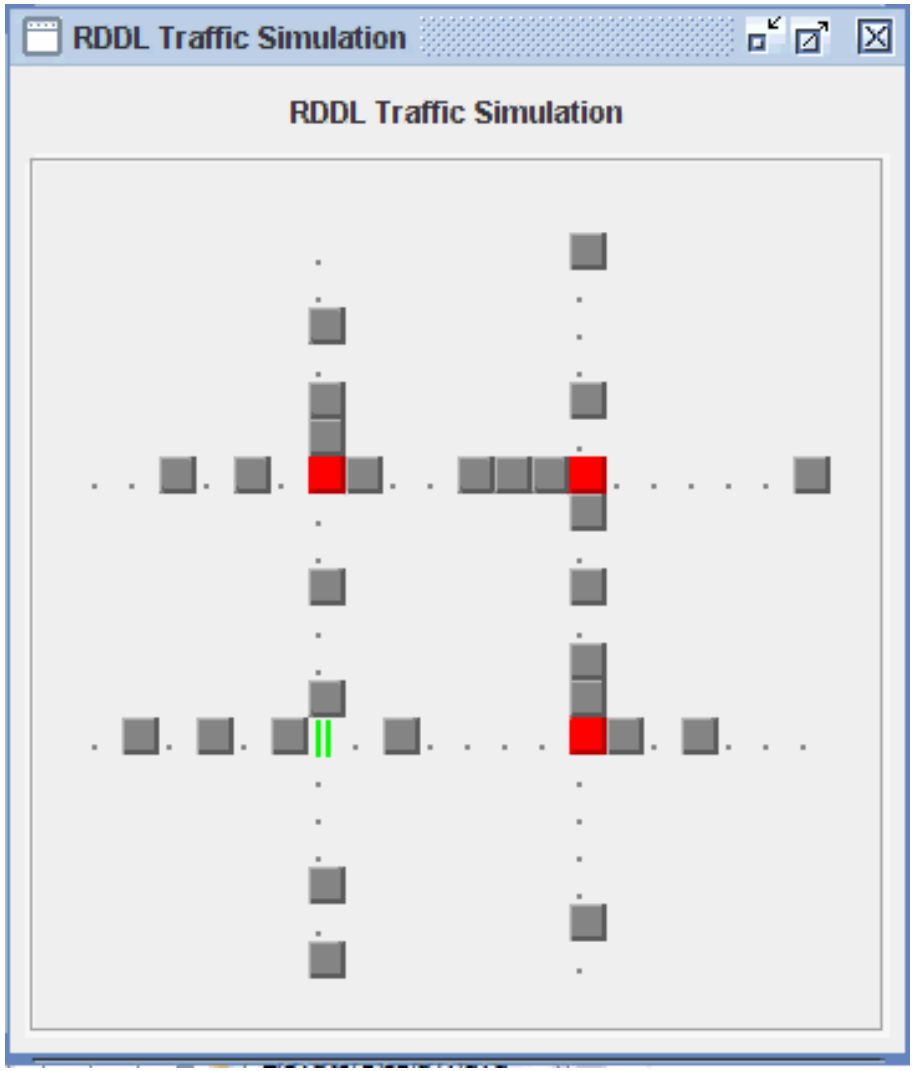

Boolean Elevators

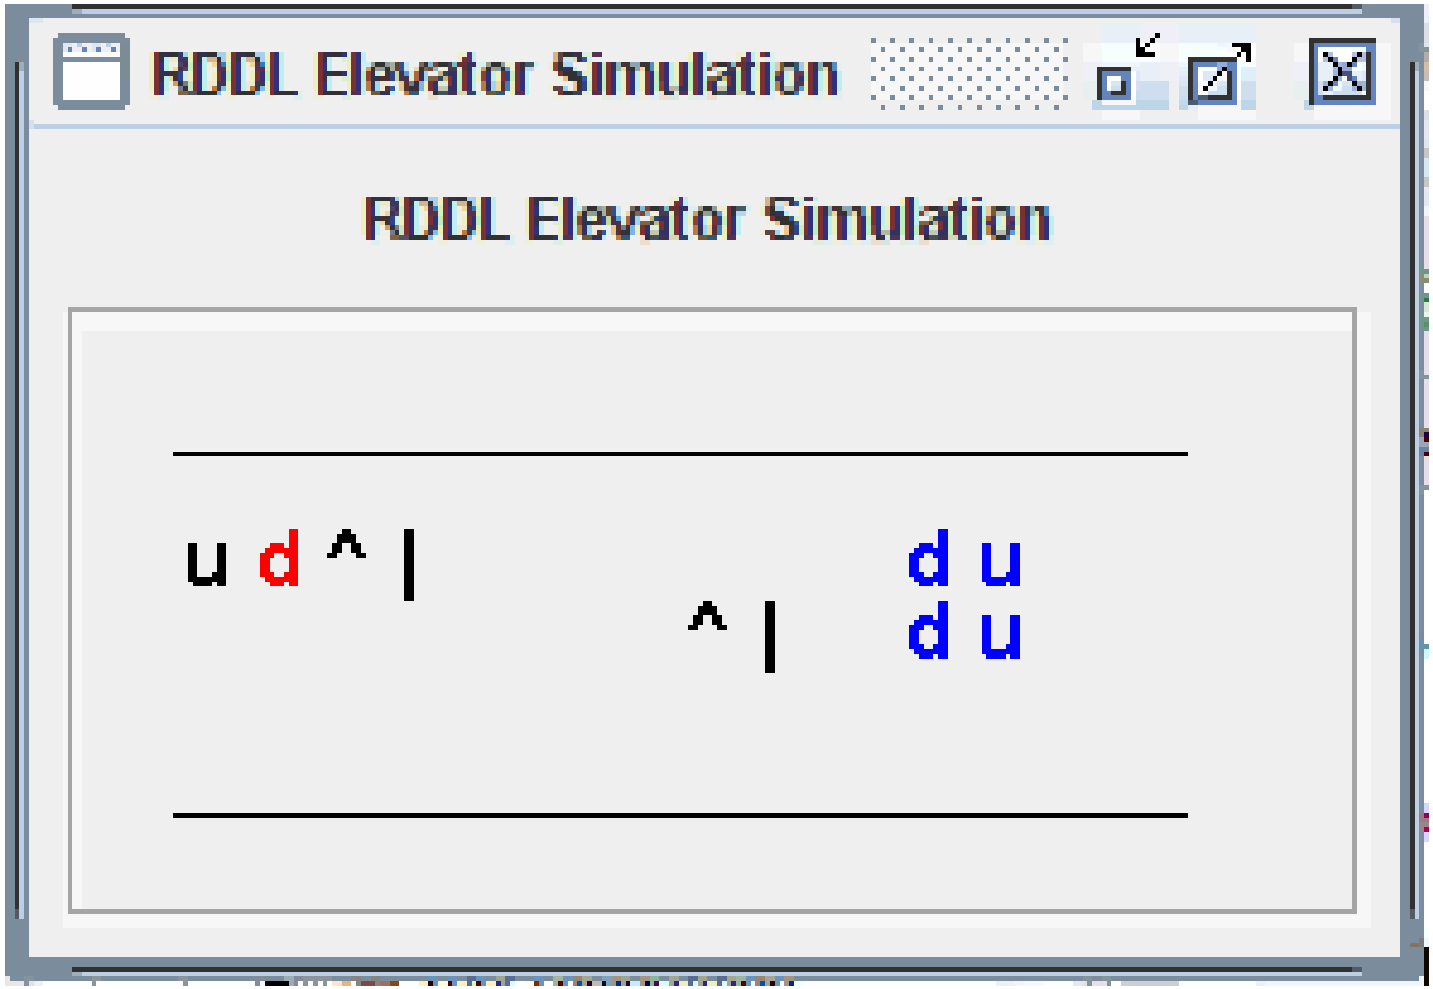

Crossing Traffic (aka Frogger)

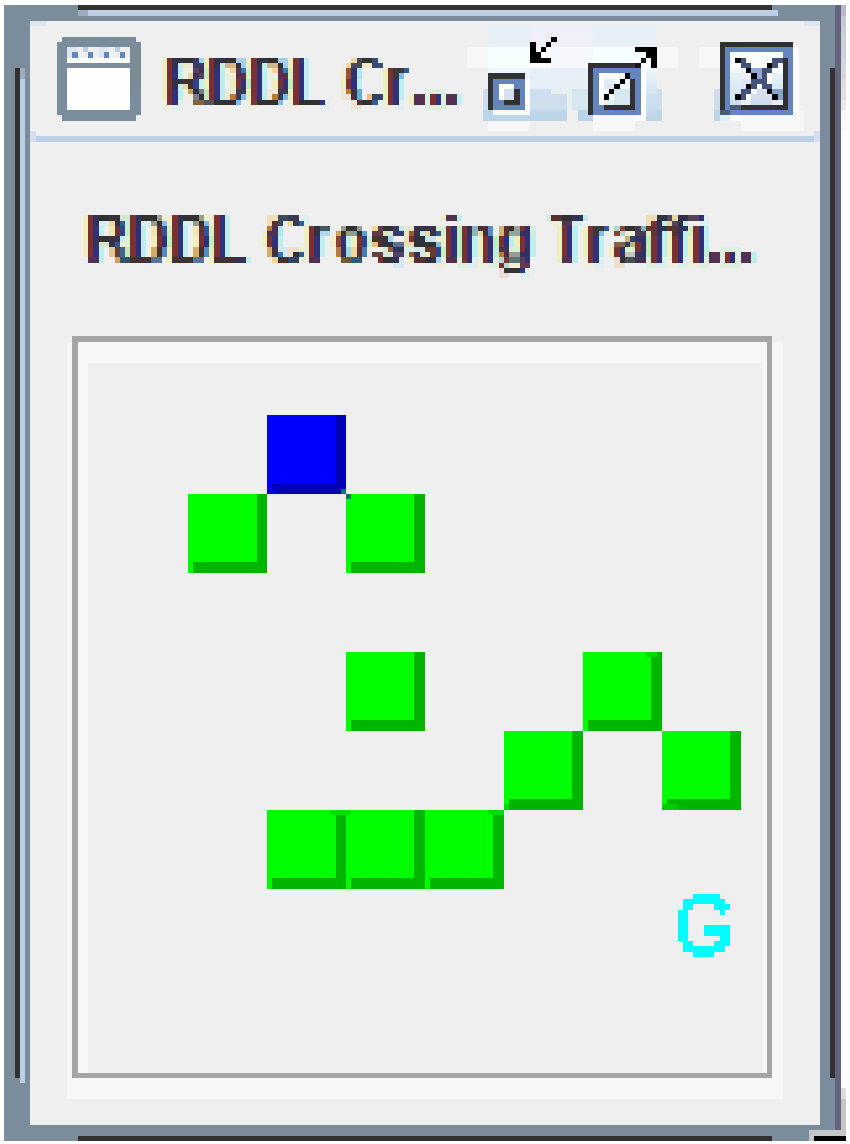

Navigation
(aka deteminization killer)

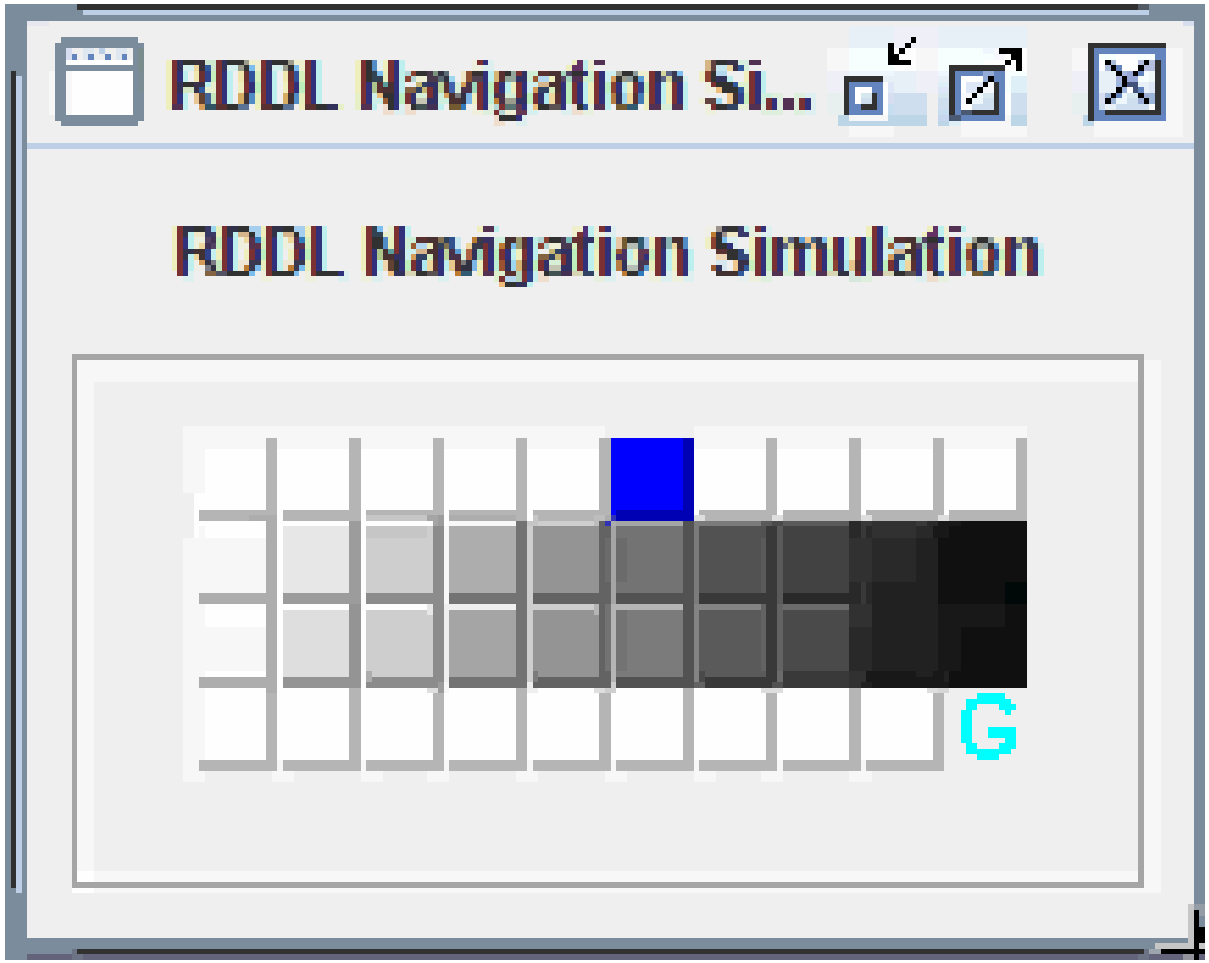

Competition Format

- **Amazon EC2** (Elastic Compute Cloud)
  - 11 instances on demand running for 24 hours
    - Ensures everyone has same computational power
      - Large EC2 instance (7.5Gb RAM, 2 Cores)
    - Everyone has admin access to their machines
  - Just pay for time used
    - received an Amazon EC2 grant of $1000 for competition
      MANY THANKS TO AMAZON FOR THEIR GENEROSITY!!!
    - so running it was free

→ **Highly recommended for future competitions!!!**

Competitors: Boolean MDP Track

Competitors	**Algorithm**
PROST (Eyerich, Keller – Uni. Freiburg)	UCT/Single Outcome Determinization, Caching
Glutton (Kolobov, Dai, Mausam, Weld – UW)	Iterative Deepening RTDP, Caching
MIT-ACL (Ure, Toksoz, Redding, Gemifard – MIT)	RL / Linear Fun. Approx, Feature Discovery
Beaver (Nadamuni, Joshi, Fern, Tadepalli – OSU)	UCT, SPUDD Guidance
SPUDD (Zhu, Grzes, Hoey – Uni. Waterloo)	SPUDD: Value Iteration with ADDS (BASELINE)

Results: Boolean MDP Track

- **1st Place: PROST**
- **2nd Place: Glutton**

PROST (Eyerich, Keller)	**0.874**	± 0.059
Glutton (Kolobov, Dai, Mausam, Weld)	**0.795**	± 0.066
Beaver (Nadamuni, Joshi, Fern, Tadepalli)	**0.245**	± 0.066
MIT-ACL (Ure, Toksoz, Redding, Gemifard)	**0.107**	± 0.055
SPUDD (Zhu, Grzes, Hoey)	**COMMUNICATION BUG**	
SPUDD (Zhu, Grzes, Hoey) (*Post-competition* results after rddlsim server communication bugs fixed, SPUDD unchanged)	**0.297**	± 0.101

Competitors: Boolean POMDP Track

Competitors	**Algorithm**
POMDPX_NUS (Wu, WS Lee, D Hsu – NUS)	SARSOP / UCT (POMCP)
KAIST-AILAB (D Kim, K Lee, K-E Kim – KAIST)	Symbolic HSVI (ADDs), Symmetry Detection
HyPlanClient (Borera, Pyeatt – Texas Tech)	~RTDP-Bel
POND (Bryce, Olsen – USU)	Translation to Conf. Planning, Hindsight Opt
Symbolic Perseus (Poupart, Hoey, Morrison – Uni. Waterloo)	PBVI with ADDs
McGill (Png, Ong, Pineau – McGill)	UCT (POMCP)

Results: Boolean POMDP Track

- **1st Place: POMDPX_NUS**
- **2nd Place: KAIST-AILAB**

POMDPX_NUS (Wu, WS Lee, D Hsu)	**0.590**	± 0.098
KAIST-AILAB (D Kim, K Lee, K-E Kim)	**0.420**	± 0.101
HyPlanClient (Borera, Pyeatt)	**0.168**	± 0.074
POND (Bryce, Olsen)	**0.152**	± 0.058
Symbolic Perseus (Poupart, Hoey, Morrison)	**0.117**	± 0.064
McGill (Png, Ong, Pineau)	**0.034**	± 0.031

Thanks to All Competitors!