### ICAPS 2012 Tutorial

### Discrete and Continuous Planning Domain Modeling in RDDL

#### Scott Sanner



### Observation

- Planning languages direct 5+ years of research
  - PDDL and variants
  - PPDDL
- Why?
  - Domain design is time-consuming
    - So everyone uses the existing benchmarks
  - Need for comparison
    - Relatively little planner code is released
    - Only means of comparison is on competition benchmarks

#### • Implication:

- We should choose our languages & problems well...

#### **Current Stochastic Domain Language**

- PPDDL
  - more expressive than PSTRIPS
  - for example, probabilistic universal and conditional effects:

(:action put-all-blue-blocks-on-table :parameters () :precondition () :effect (probabilistic 0.9 (forall (?b) (when (Blue ?b) (not (OnTable ?b)))))

- But wait, not just BlocksWorld...
  - Colored BlocksWorld
  - Exploding BlocksWorld
  - Moving-stacks BlocksWorld
- Difficult problems *but* where to apply solutions???





#### More Realistic: Logistics

• Compact relational PPDDL Description:



(:action load-box-on-truck-in-city :parameters (?b - box ?t - truck ?c - city) :precondition (and (BIn ?b ?c) (TIn ?t ?c)) :effect (and (On ?b ?t) (not (BIn ?b ?c))))

- Can instantiate problems for any domain objects
  - 3 trucks: 🖡 🖡 🖡 2 planes: 🦗 🐜 3 boxes: 🖱 🖱 🖱
- But wait... only one truck can move at a time???
  - No concurrency, no time: will FedEx care?

# What stochastic problems should we care about?

### Mars Rovers



Continuous

– Time, robot position / pose, sun angle, ...

- Partially observable
  - Even worse: high-dimensional partially observable

### **Elevator Control**

- Concurrent Actions
  - Elevator: up/down/stay
  - 6 elevators: 3^6 actions
- Exogenous / Non-boolean:
  - Random integer arrivals (e.g., Poisson)
- Complex Objective:
  - Minimize sum of wait times
  - Could even be nonlinear function (squared wait times)
- Policy Constraints:
  - People might get annoyed if elevator reverses direction





### Traffic Control



- Concurrent
  - Multiple lights
- Indep. Exogenous Events Partially observable
  - Multiple vehicles

- **Continuous Variables** 
  - Nonlinear dynamics
- - Only observe stoplines

Can PPDDL model these problems?

No? What happened?

#### A Brief History of (ICAPS) Time



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

# What would it take to model more realistic problems?

Let's take a deeper look at traffic control...

#### Birth of RDDL: Solving Traffic Control



## What's missing in PPDDL,

PRIMARY REASON FOR DEPARTURE FROM PPDDL

- Need Unrestricted Concurrency:
  - In PPDDL, would have to enumerate joint actions
  - In PDDL 2.1: restricted concurrency
    - conflicting actions not executable
    - when effects probabilistic, some chance most effects conflict

       really need unrestricted concurrency in probabilistic setting
- Multiple Independent Exogenous Events:
  - PPDDL only allows 1 independent event to affect fluent
    - E.g, what if cars in a queue change lanes, brake randomly?

#### Need a way to resolve conflicting effects... solution will be a Relational DBN

# What's missing in PPDDL, F

Could be added to PPDDL as well

- Expressive transition distributions:
  - (Nonlinear) stochastic difference equations
    - Gaussian noise
- Partial observability:
  - In practice, only observe stopline



### What's missing in PPDDL, P

Could be added to PPDDL as well

- Distinguish fluents from nonfluents:
  - E.g., topology of traffic network
  - Lifted planners must know this to be efficient!
- Expressive rewards & probabilities:
  - E.g., state and action dependent sums / products over domain objects (+1 for each computer running)
- Global state-action constraints:
  - Concurrent domains need global action preconditions
    - E.g., two traffic lights cannot go into a given state
  - In logistics, vehicles cannot be in two different locations
    - Regression planners need state constraints!

### Is there any hope?

Yes, but we need to borrow from factored MDP / POMDP community...

#### A Brief History of (ICAPS) Time



# What is RDDL?

- Relational Dynamic Influence Diagram Language
  - Relational
     [DBN + Influence Diagram]
- Think of it as Relational SPUDD / Symbolic Perseus
  - But lifted



Key task: how to specify lifted distributions & reward?

#### **RDDL** Grammar

Let's examine BNF grammar in infinite tedium!

OK, maybe not. (Grammar <u>online</u> if you want it.)

#### **RDDL** Examples

Easiest to understand RDDL in use...

#### How to Represent Factored MDP?

#### Current State and Actions

Next State and Reward



### **RDDL** Equivalent

// Define the state and action variables (not parameterized here) pvariables { p : { state-fluent, bool, default = false }; q : { state-fluent, bool, default = false }; r : { state-fluent, bool, default = false }; Can think of a : { action-fluent, bool, default = false }; transition }; distributions as "sampling // Define the conditional probability function for each // state variable in terms of previous state and actio instructions" cpfs { p' = if (p ^ r) then Bernoulli(.9) else Bernoulli(.3); q' = if (q ^ r) then Bernoulli(.9) else if (a) then Bernoulli(.3) else Bernoulli(.8); r' = if (~q) then KronDelta(r) else KronDelta(r <=> q); }; // Define the reward function; note that boolean functions are // treated as 0/1 integers in arithmetic expressions

reward = p + q - r;

#### A Discrete-Continuous POMDP?



#### A Discrete-Continuous POMDP, Part I

```
// User-defined types
types {
    enum_level : {@low, @medium, @high}; // An enumerated type
};
pvariables {
   p : { state-fluent, bool, default = false };
    q : { state-fluent, bool, default = false };
   r : { state-fluent, bool, default = false };
    i1 : { interm-fluent, int, level = 1 };
    i2 : { interm-fluent, enum_level, level = 2 };
    o1 : { observ-fluent, bool };
    o2 : { observ-fluent, real };
    a : { action-fluent, bool, default = false };
};
cpfs {
    // Some standard Bernoulli conditional probability tables
   p' = if (p ^ r) then Bernoulli(.9) else Bernoulli(.3);
    q' = if (q ^ r) then Bernoulli(.9)
                    else if (a) then Bernoulli(.3) else Bernoulli(.8);
    // KronDelta is a delta function for a discrete argument
    r' = if (~q) then KronDelta(r) else KronDelta(r <=> q);
```

#### A Discrete-Continuous POMDP, Part II



### RDDL so far...

- Mainly SPUDD / Symbolic Perseus with a different syntax <sup>(2)</sup>
  - A few enhancements
    - concurrency
    - constraints
    - integer / continuous variables
- Real problems (e.g., traffic) need lifting
  - An intersection model
  - A vehicle model
    - Specify each intersection / vehicle model once!

# Lifting: Conway's Game of Life

(simpler than traffic)

- Cells born, live, die based on neighbors
  - < 2 or > 3 neighbors: cell dies
  - 2 or 3 neighbors: cell lives
  - 3 neighbors  $\rightarrow$  cell birth!
  - Make into MDP
    - Probabilities
    - Actions to turn on cells
    - Maximize number of cells on







http://en.wikipedia.org/wiki/Conway's\_Game\_of\_Life

• Compact RDDL specification for *any* grid size? Lifting.



#### A Lifted MDP



#### Nonfluent and Instance Defintion



### Power of Lifting

Simple domains

can generate

complex DBNs!

#### non-fluents game3x3 { non-fluents game2x2 { Current State and Actions Intermediate @ Level 1 domain = game of life; domain = game of life; set(x3, y1) set(x3, y2) objects { objects { x\_pos : {x1,x2,x3} set(x3, y3) Next State and Reward x pos : {x1,x2}; y\_pos : {y1,y2,y3}; y\_pos : {y1,y2}; alive'(x3, y1) }: }; alive'(x3, y2) non-fluents { non-fluents { alive(x3, y3) NEIGHBOR(x1,y1,x1,y2); alive'(x3, y3) PROB REGENERATE = 0.9; NEIGHBOR(x1,y1,x2,y1); NEIGHBOR(x1,y1,x2,y2); NEIGHBOR(x1,y1,x1,y2); alive(x3, y1) NEIGHBOR(x1,y2,x1,y1); NEIGHBOR(x1,y1,x2,y1); NEIGHBOR(x1.v2.x2.v1): NEIGHBOR(x1,y1,x2,y2); count-neighbors(x3, y2) NEIGHBOR(x1,y2,x2,y2); NEIGHBOR(x1,y2,x2,y3); alive(x3, y2) count-neighbors(x3, y1) NEIGHBOR(x1,y2,x1,y1); NEIGHBOR(x1,y2,x1,y3); NEIGHBOR(x1,y2,x2,y1); NEIGHBOR(x1,y3,x1,y2); NEIGHBOR(x1,y2,x2,y2); NEIGHBOR(x1,y3,x2,y2); count-neighbors(x3, y3) NEIGHBOR(x1,y3,x2,y3); **Reward Function** NEIGHBOR(x2,y1,x1,y1); NEIGHBOR(x2,y1,x1,y1); alive(x2, y1) NEIGHBOR(x2,y1,x1,y2); count-neighbors(x2, y2) NEIGHBOR(x2,y1,x1,y2); NEIGHBOR(x2,y1,x2,y2); NEIGHBOR(x2,y1,x2,y2); NEIGHBOR(x2,y1,x3,y2); alive(x2, y3) alive'(x2, y2) NEIGHBOR(x2,y2,x1,y1); NEIGHBOR(x2,y1,x3,y1); NEIGHBOR(x2,y2,x1,y2); set(x2, y2) NEIGHBOR(x2,y2,x1,y1); NEIGHBOR(x2,y2,x2,y1); alive'(x2, y1) NEIGHBOR(x2,y2,x1,y2); set(x2, y1) }: NEIGHBOR(x2,y2,x1,y3); NEIGHBOR(x2.v2.x2.v1): alive(x2, y2) NEIGHBOR(x2,y2,x2,y3); Current State and Actions count-neighbors(x2, y1) Intermediate @ Level 1 NEIGHBOR(x2,y2,x3,y1); set(x1, y1) Next State and Reward alive'(x2, y3) count-neighbors(x2, y3) NEIGHBOR(x2,y2,x3,y2); NEIGHBOR(x2,y2,x3,y3); set(x1, y2)alive'(x1, y1) NEIGHBOR(x2,y3,x1,y3); count-neighbors(x1, y1) NEIGHBOR(x2,y3,x1,y2); alive'(x1, y2) alive(x1, y2) count-neighbors(x1, y3) alive(x1, y2) NEIGHBOR(x2,y3,x2,y2); NEIGHBOR(x2,y3,x3,y2); NEIGHBOR(x2,y3,x3,y3); alive(x1, y1) count-neighbors(x2, y1) count-neighbors(x1, y2) NEIGHBOR(x3,y1,x2,y1); NEIGHBOR(x3,y1,x2,y2); count-neighbors(x1, y1) set(x2, y3) NEIGHBOR(x3,y1,x3,y2); alive'(x1, y1) alive(x1, y1) count-neighbors(x2, y2) NEIGHBOR(x3,y2,x3,y1); NEIGHBOR(x3,y2,x2,y1); alive'(x1, y2) count-neighbors(x1, y2)alive(x2, y2) NEIGHBOR(x3.v2.x2.v2): set(x1, y1) **Reward Function** NEIGHBOR(x3,y2,x2,y3); alive(x1, y3) alive(x2, y1) NEIGHBOR(x3,y2,x3,y3); alive'(x1, y3) alive'(x2, y1) NEIGHBOR(x3,y3,x2,y3); set(x1, y2) NEIGHBOR(x3,y3,x2,y2); NEIGHBOR(x3,y3,x3,y2); set(x1, y3) set(x2, y1) alive'(x2, y2)

}

set(x2, y2)

}

#### Complex Lifted Transitions: SysAdmin SysAdmin (Guestrin et al, 2001)

- Have n computers  $C = \{c_1, ..., c_n\}$  in a network
- State: each computer c<sub>i</sub> is either "up" or "down"



- **Transition:** computer is "up" proportional to its state and # upstream connections that are "up"
- Action: manually reboot one computer
- Reward: +1 for every "up" computer

#### Complex Lifted Transitions SysAdmin (Guestrin et al, 2001)

pvariables {

```
REBOOT-PROB : { non-fluent, real, default = 0.1 };
    REBOOT-PENALTY : { non-fluent, real, default = 0.75 };
    CONNECTED(computer, computer) : { non-fluent, bool, default = false };
    running(computer) : { state-fluent, bool, default = false };
    reboot(computer) : { action-fluent, bool, default = false };
};
                                      Probability of a
cpfs {
                                     computer running
                                     depends on ratio of
  running'(?x) = if (reboot(?x))
                                        connected
     then KronDelta(true) // if
                                                        then must be running
                                     computers running!
     else if (running(?x)) // else
                                                       network properties
        then Bernoulli(
         .5 + .5*[1 + sum_{?y} : computer\} (CONNECTED(?y,?x) ^ running(?y))]
                 / [1 + sum_{?y : computer} CONNECTED(?y,?x)])
        else Bernoulli(REBOOT-PROB);
};
reward = sum_{?c : computer} [running(?c) - (REBOOT-PENALTY * reboot(?c))];
```

# Lifted Continuous MDP in RDDL: Simple Mars Rover

![](_page_33_Figure_1.jpeg)

### Simple Mars Rover: Part I

types { picture-point : object; };

#### pvariables {

![](_page_34_Figure_3.jpeg)

to make multirover?

# Simple Mars Rover: Part II

#### cpfs {

// Noisy movement update **xPos'** = xPos + xMove + Normal(0.0, MOVE\_VARIANCE\_MULT\*xMove); **yPos'** = yPos + yMove + Normal(0.0, MOVE\_VARIANCE\_MULT\*yMove); White noise, variance // Time update proportional to distance moved **time'** = if (snapPicture) then DiracDelta(time + 0.25) Fixed time for picture else DiracDelta(time + [if (xMove > 0) then xMove else -xMove] + \if (yMove > 0) then yMove else -yMove]); Time proportional to distance moved };

# Simple Mars Rover: Part III

// We get a reward for any picture taken within picture box error bounds
// and the time limit.

```
reward = if (snapPicture ^ (time <= MAX_TIME))
          then sum_{?p : picture-point} [
             if ((xPos >= PICT_XPOS(?p) - PICT_ERROR_ALLOW(?p))
                 ^ (xPos <= PICT_XPOS(?p) + PICT_ERROR_ALLOW(?p))</pre>
                 ^ (yPos >= PICT_YPOS(?p) - PICT_ERROR_ALLOW(?p))
                 ^ (yPos <= PICT_YPOS(?p) + PICT_ERROR_ALLOW(?p)))</pre>
             then PICT VALUE(?p)
             else 0.0]
                               Reward for all pictures taken
           else 0.0;
                                   within bounding box!
state-action-constraints {
                                                    Cannot move and take
                                                    picture at same time.
        // Cannot snap a picture and move at the same w
        snapPicture => ((xMove == 0.0) \land (yMove == 0.0));
};
```

#### How to Think About Distributions

- Transition distribution is stochastic program
  - Similar to BLOG (Milch, Russell, et al), IBAL (Pfeffer)
  - Leaves of programs are distributions
    - Think of SPUDD / Sym. Perseus decision diagrams as having Bernoulli leaves
- *Procedural* specification of sampling process
  - Use intermediate DBN variables for storage
  - E.g., drawing a distance measurement in robotics

true-distance

10

()

- **boolean** *Noise* := sample from **Bernoulli** (.1)
- real Measurement := If (Noise == true)
  - Then sample from Uniform(0, 10)
  - Else sample from Normal(true-distance,  $\sigma^2$ )

Convenient way to write complex mixture models and conditional distributions that occur in practice!

## RDDL Recap I

- Everything is a fluent (parameterized variable)
  - State fluents
  - Observation fluents
    - for partially observed domains
  - Action fluents
    - supports factored concurrency
  - Intermediate fluents
    - derived predicates, correlated effects, ...
  - Constant nonfluents (general constants, topology relations, ...)
- Flexible fluent types
  - Binary (predicate) fluents
  - Multi-valued (enumerated) fluents
  - Integer and continuous fluents (from PDDL 2.1)

# RDDL Recap II

- Semantics is ground DBN / Influence Diagram
  - Unambiguous specification of transition semantics
    - Supports unrestricted concurrency
  - Naturally supports independent exogenous events
- General expressions in transition / reward
  - Logical expressions  $(\land, \lor, \Rightarrow, \Leftrightarrow, \forall, \exists) <$ Logical expr. {0,1} so can use in
  - Arithmetic expressions  $(+,-,*,/, \Sigma_x, \Pi_x)$  arithmetic expr.
  - In/dis/equality comparison expressions (=,  $\neq$ , <,>,  $\leq$ ,  $\geq$ )
  - Conditional expressions (if-then-else, switch)
  - Basic probability distributions
    - Bernoulli, Discrete, Normal, Poisson

 $\sum_{x}$ ,  $\prod_{x}$  aggregators over domain objects extremely powerful

# RDDL Recap III

- Goal + General (PO)MDP objectives
  - Arbitrary reward
    - goals, numerical preferences (c.f., PDDL 3.0)
  - Finite horizon
  - Discounted or undiscounted
- State/action constraints
  - Encode legal actions
    - (concurrent) action preconditions
  - Assert state invariants
    - e.g., a package cannot be in two locations

### **RDDL** Software

Open source & online at <a href="http://code.google.com/p/rddlsim/">http://code.google.com/p/rddlsim/</a>

## Java 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
  - Evaluation scripts for log files
- Visualization
  - DBN Visualization
  - Domain Visualization see how your planner is doing

#### Visualization of Boolean Traffic

![](_page_43_Figure_1.jpeg)

#### Visualization of Boolean Elevators

![](_page_44_Figure_1.jpeg)

### **RDDL** Domains

- Boolean track
  - 8 domains (including traffic & elevators)
  - 10 instances per domain from IPPC
  - Generators for any size instance!
- General track (bool, integer, continuous)
   Range of problems (Mars Rover, concurrent)
  - Where I hope future IPPC focuses...

### Ideas for other RDDL Domains

- UAVs with partial observability
- (Hybrid) Control
  - Linear-quadratic control (Kalman filtering with control)
  - Discrete and continuous actions avoided by planning
  - Nonlinear control
- Dynamical Systems from other fields
  - Population dynamics
  - Chemical / biological systems
  - Physical systems
    - Pinball!
  - Environmental / climate systems
- Bayesian Modeling
  - Continuous Fluents can represent parameters
    - Beta / Bernoulli / Dirichlet / Multinomial / Gaussian
  - Then progression is a Bayesian update!
    - Bayesian reinforcement learning

### Submit your own Domains in RDDL!

Field only makes true progress working on realistic problems

#### Future RDDL Extensions?

- Elementary functions
  - sin, cos, log, exp, sqrt
- Effects-based specification?
  - Easier to write than current fluent-centered approach
  - But how to resolve conflicting effects in unrestricted concurrency
- Binomial / Multinomial
  - Need a vector fluent type when sampling vectors of counts

#### • Object fluents

- Much harder than PDDL 3
- Distribrutions over indefinite number of objects
  - Perhaps can borrow ideas from BLOG (Milch et al)
- Timed processes?
  - Continuous time stochastic differential equations
  - Asynchronous concurrency + time quite difficult

Enjoy RDDL! (no lack of difficult problems to solve!)

**Questions**?