THE UNIVERSITY OF BRITISH COLUMBIA — DEPARTMENT OF COMPUTER SCIENCE — LABORATORY FOR COMPUTATIONAL INTELLIGENCE

A Scalable Framework for Representation and Reasoning in Large Scale, Spatial-Temporal Planning Problems

http://www.cs.ubc.ca/~crowley   by: Mark Crowley, PhD Candidate   crowley@cs.ubc.ca

LST Problems (Forestry Planning)

Large Scale Spatial Temporal (LST) Planning Problems

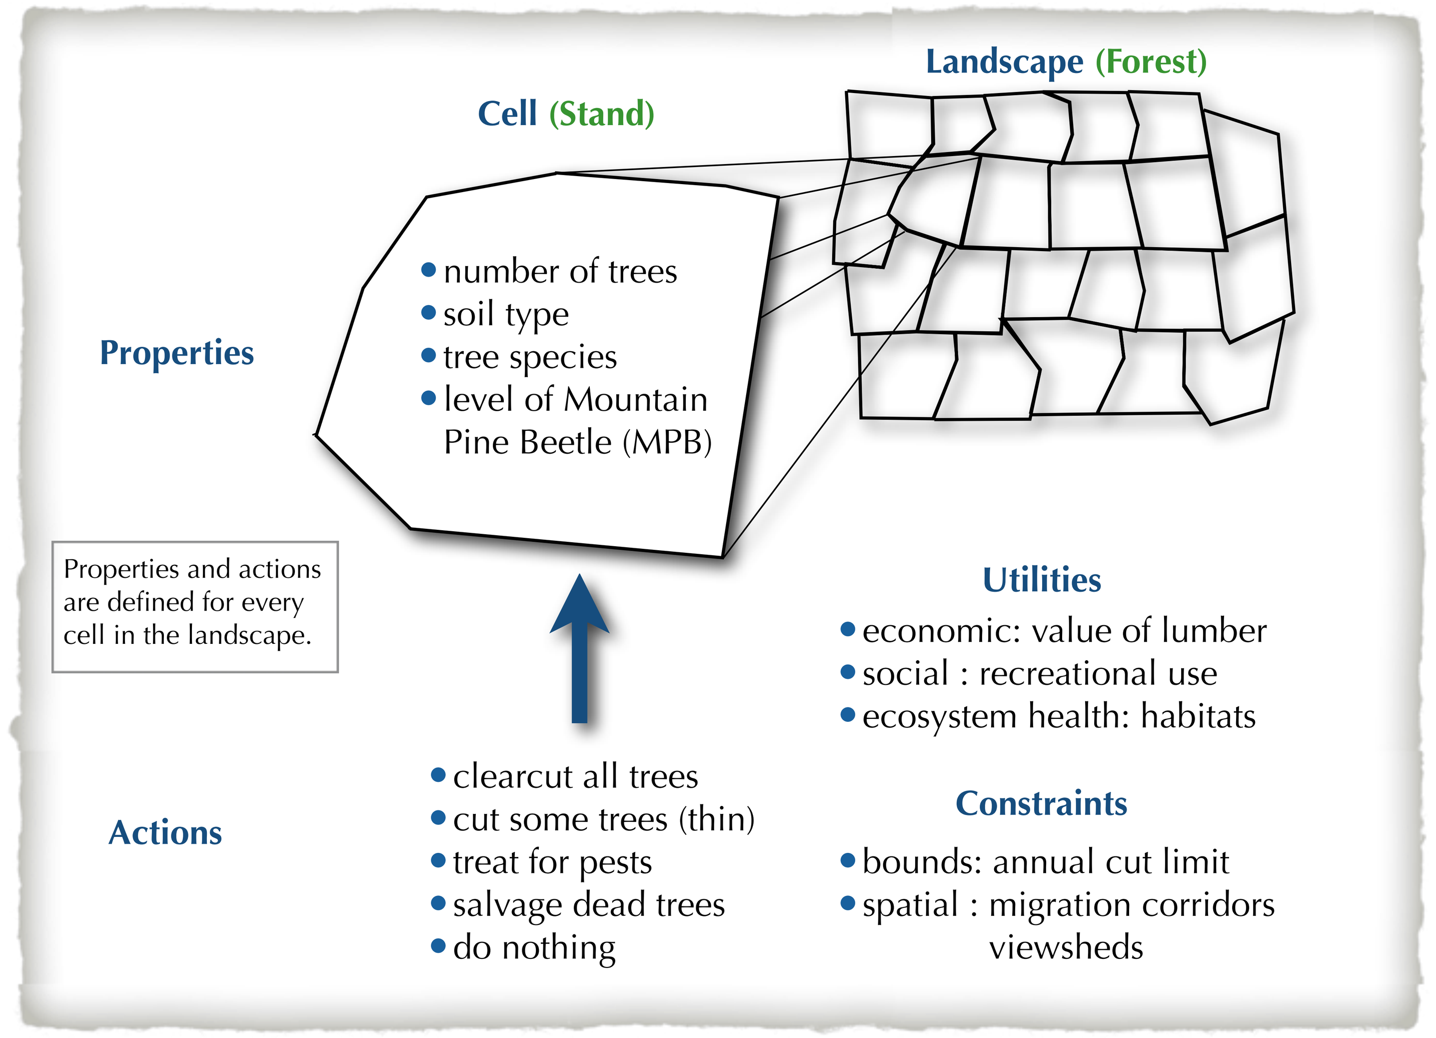

Uncertainty

- state of landscape, soil, temperature known with high certainty
- state of tree properties are **aggregated estimates** from aerial, satellite and ground data
- state of **previous year's** MPB infestations estimated based on aerial surveys of "red tree" counts.
  - 1 to 2 year time lag
  - high uncertainty about precise level of infestation for individual stands

Dynamics

- **Tree growth :** model with yield curves, fairly deterministic
- **MPB growth :** spread determined by infestation level, wind direction, temperature. Beetles spread spatially across the forest landscape from stand to stand or fly over 20km to other stands
  - Models exist that estimate the risk of spread based on these factors for a stand given data about nearby stands

Mountain Pine Beetle

(Dendrochtonus ponderosae)

- **Life cycle :** 1 year
- **Dispersal distance :** up to 20km
- **Origin :** naturally endemic
- **Trees at risk :** lodgepole pine and other pine in BC and increasingly in Alberta
- **Method of attack :** burrows under bark, cuts off water, leaves fungus and eggs
- **Current infestation :** over 11 million ha
- **Weaknesses:** thin forests, young trees, fire, extreme cold (-40C)

MPB makes planning problem more complex
- MPB spread affects utility by killing trees
- Actions affect MPB spread by altering the trees they can spread to

Existing Approaches

In forestry several methods are currently used to perform harvest planning:
- Linear models of independent cells - solved with Linear Programming
- Nonlinear models solved with Stochastic Local Search (Meta-heuristics)
  - Simulated Annealing
  - Genetic Algorithms

Neither of these approaches deal well with the issues of **uncertainty** or **spatial interrelation.** But **MPB infestations** makes the impact of these issues much greater. There is a desire[3] in forestry for a more realistic treatment of uncertainty and spatial interrelation in their models.

Our Goals

To provide a general framework for
- **Representation** - how to represent LST problems compactly in a way that supports decision making
- **Reasoning** - how to perform efficient reasoning to find good policies in LST problems

We want to take account of
- uncertainty in the model
- spatial interrelations in the data

We also want to enable use of existing simulations from the problem domain so that simulation/modeling is separated from general planning.

Increasing Abstraction Over Time

In long term planning, future states are highly uncertain and plans will be recomputed long before they are used. Thus, an abstraction schedule could be followed as future time steps are considered. The schedule could be determined:
- **as needed :** to balance performance and accuracy
- **ahead of time :** to indicate level of interest at points in the future

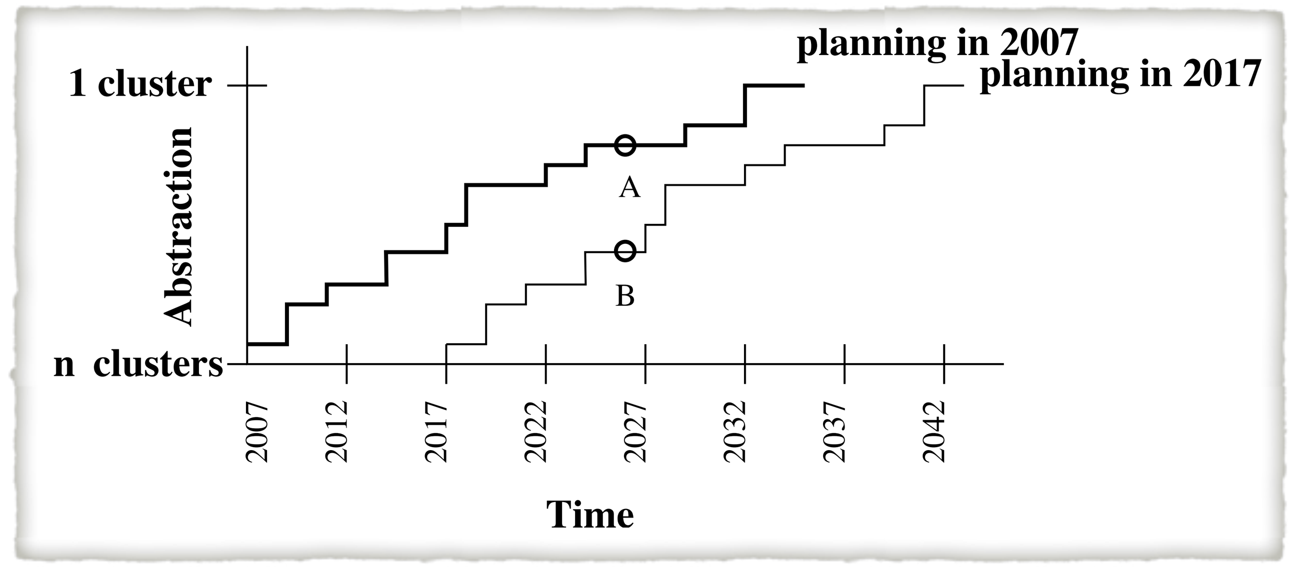

Representation as a DBN

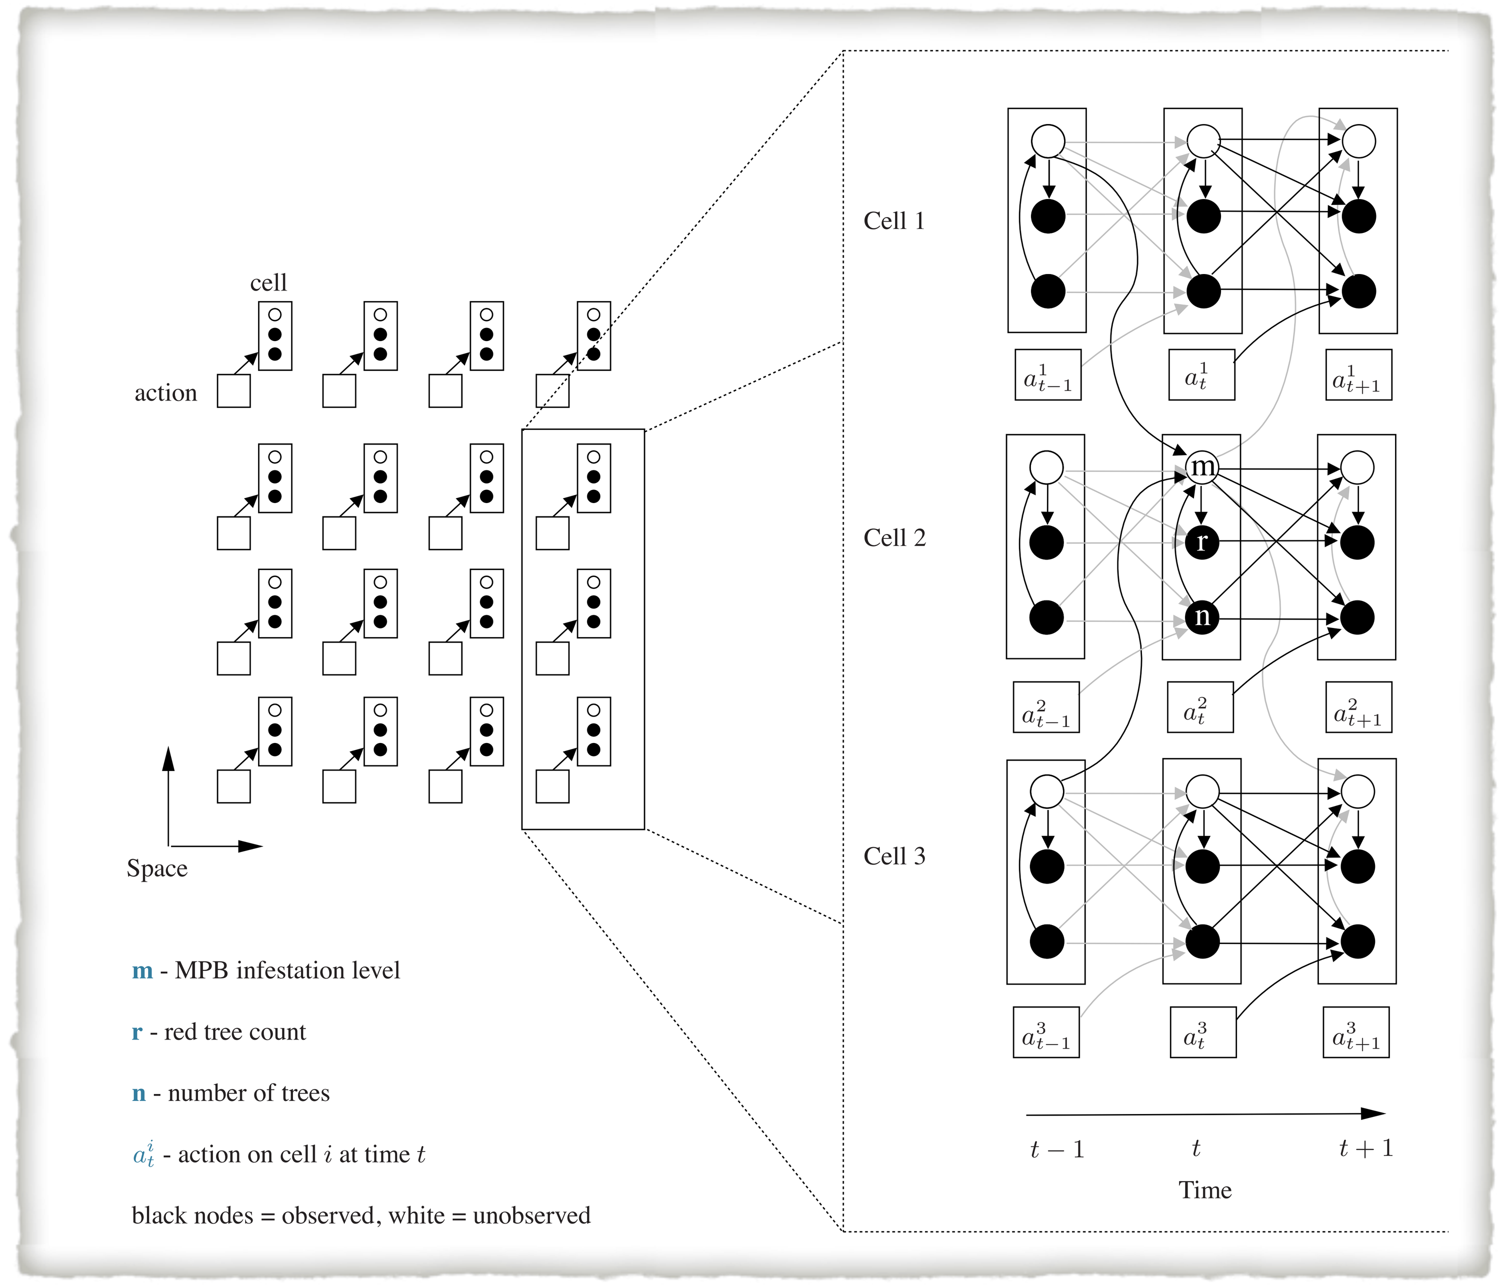

State Abstraction

We will reduce the size of the state space by reducing the number of cells that need to be considered:
- A **cluster of cells** with similar properties are mapped to a typical cell or class called a **prototype**
- Clustering will be performed on cell features including:
  - local cell properties - number of trees, level of MPB infestation, tree spacing
  - spatial relations between cells - average level of MPB in neighbourhood, proximity to roads/water

The goal is to find a clustering that minimizes the prediction error if cells are all replaced with their prototypes. We will then use these cell prototypes in planning.

Neighbourhod Sampling

We will use the cell prototype mapping to learn useful spatial structure for efficient planning by observing that :
- most spatial relationships have an upper bound on their distance of influence (e.g. MPB spread no more than 20 km)

Thus, we can use prototypes to create **sample neighbourhoods** that are much smaller than the landscape but still contain useful spatial relationships.

A	A	B	A	C
B	C	B	C	C
B	A	A	B	B
A	C	B	C	B
B	A	B	B	A

C	B	C
A	A	B
B	C	B

C	B	C
B	A	B
B	B	A

B	C	C
A	B	B
B	A	B

Hierarchical Abstraction Model

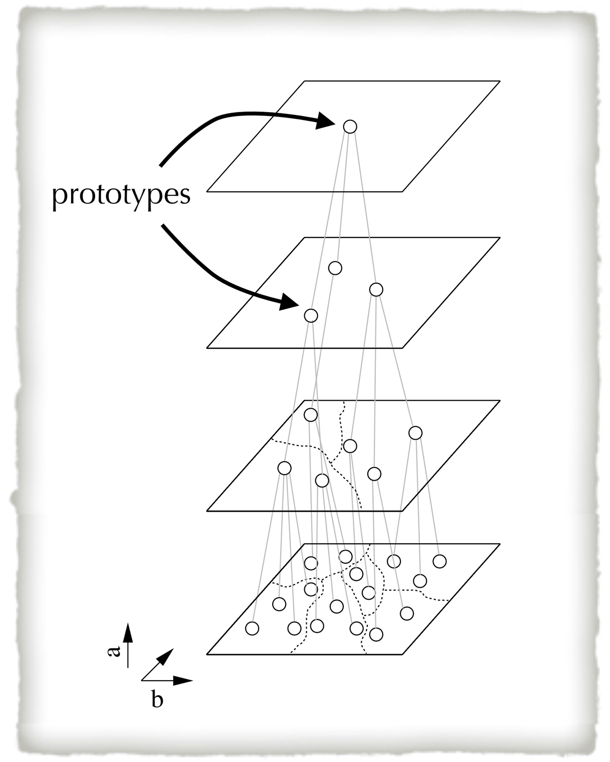

We can maintain a representation at multiple levels of abstraction by adjusting the number of prototypes used:
- The **highest level** of abstraction is represented by a single prototype, one cluster, averaging out all proper ties across the landscape
- The **lowest level** of abstraction is represented by a prototype for every cell in the landscape

Advantages:
- size of state is **adjustable**
- planning in the real world is **carried out hierarchically** anyways
  - strategic planners deal with forest-wide plans
  - tactical planners take these plans and look more closely at local conditions

Related Approaches

- Hierarchical Reinforcement Learning [1] produces polices at different scales of time and uses them to inform each other
- Model based solutions
  - Belief state POMDP solution methods, especially point based methods such as **Perseus**[9]
  - **Symbolic Perseus**[8] can handle millions of states. We have evaluated this method and found it did not scale well with a simple LST problem
- Direct Policy Search [5][10] (such as **PEGASUS**[7])
  - search directly to improve a parameterized policy
  - deals well with simulated dynamics rather than having full transistion model
- Multiagent/concurrent decision making[2][4]
  - many actions taken in parallel seen as seperate agents sharing knowledge
  - effective for compactly representing the value function or the policy itself
  - connected to representations of value functions and policies as linear combinations of basis functions
- Relational MDPs [6] represent relations between classes of individuals and form general plans based on those classes. Related to neighbourhood sampling idea.

References

[1] T. G. Dietterich. An overview of MAXQ hierarchical reinforcement learning. *Lecture Notes in Computer Science*, 1864/2000, 2000.

[2] C. Guestrin, D. Koller, and R. Parr. Multiagent planning with factored MDPs. In *NIPS-14*, 2001.

[3] J.P.Kimmins et al. *Possible Forest Futures.* Working Paper 2005-11. Pacific Forestry Centre, Canadian Forest Service, National Resources Canada. Victoria, Canada., 2005.

[4] B. Marthi, S. Russell, D. Latham, and C. Guestrin. Concurrent hierarchical reinforcement learning. In *The International Joint Conference on Artificial Intelligence (IJCAI-05)*, Edinburgh, Scotland, 2005.

[5] R. Martinez-Cantin, N. de Freitas, A. Doucet, and J. Castellanos. Active policy learning for robot planning and exploration under uncertainty. In *Robotics: Science and Systems (RSS)*, 2007.

[6] Mausam and D. S. Weld. Solving relational MDPs with first-order machine learning. In *AAAI-04*, San Jose, CA, 2003.

[7] A. Y. Ng and M. Jordan. PEGASUS: A policy search method for large MDPs and POMDPs. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence.*, 2000.

[8] P. Poupart. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes.* PhD thesis, University of Toronto, 2005.

[9] M. T. J. Spaan and N. Vlassis. Perseus: Randomized point-based value iteration for POMDPs. *Journal of Artificial Intelligence Research*, 24:195–220, 2005.

[10] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(2):229–256, 1992.