Time-Critical Planning

ARPI Program Review
June, 1997

Thomas L. Dean
Robert L. Givan
Leslie Pack Kaelbling

Key Idea

Achieve robustness in unpredictable domains by embracing uncertainty

- deterministic models may make fatal approximations
- stochastic models let us hedge our bets

Worst Case World View

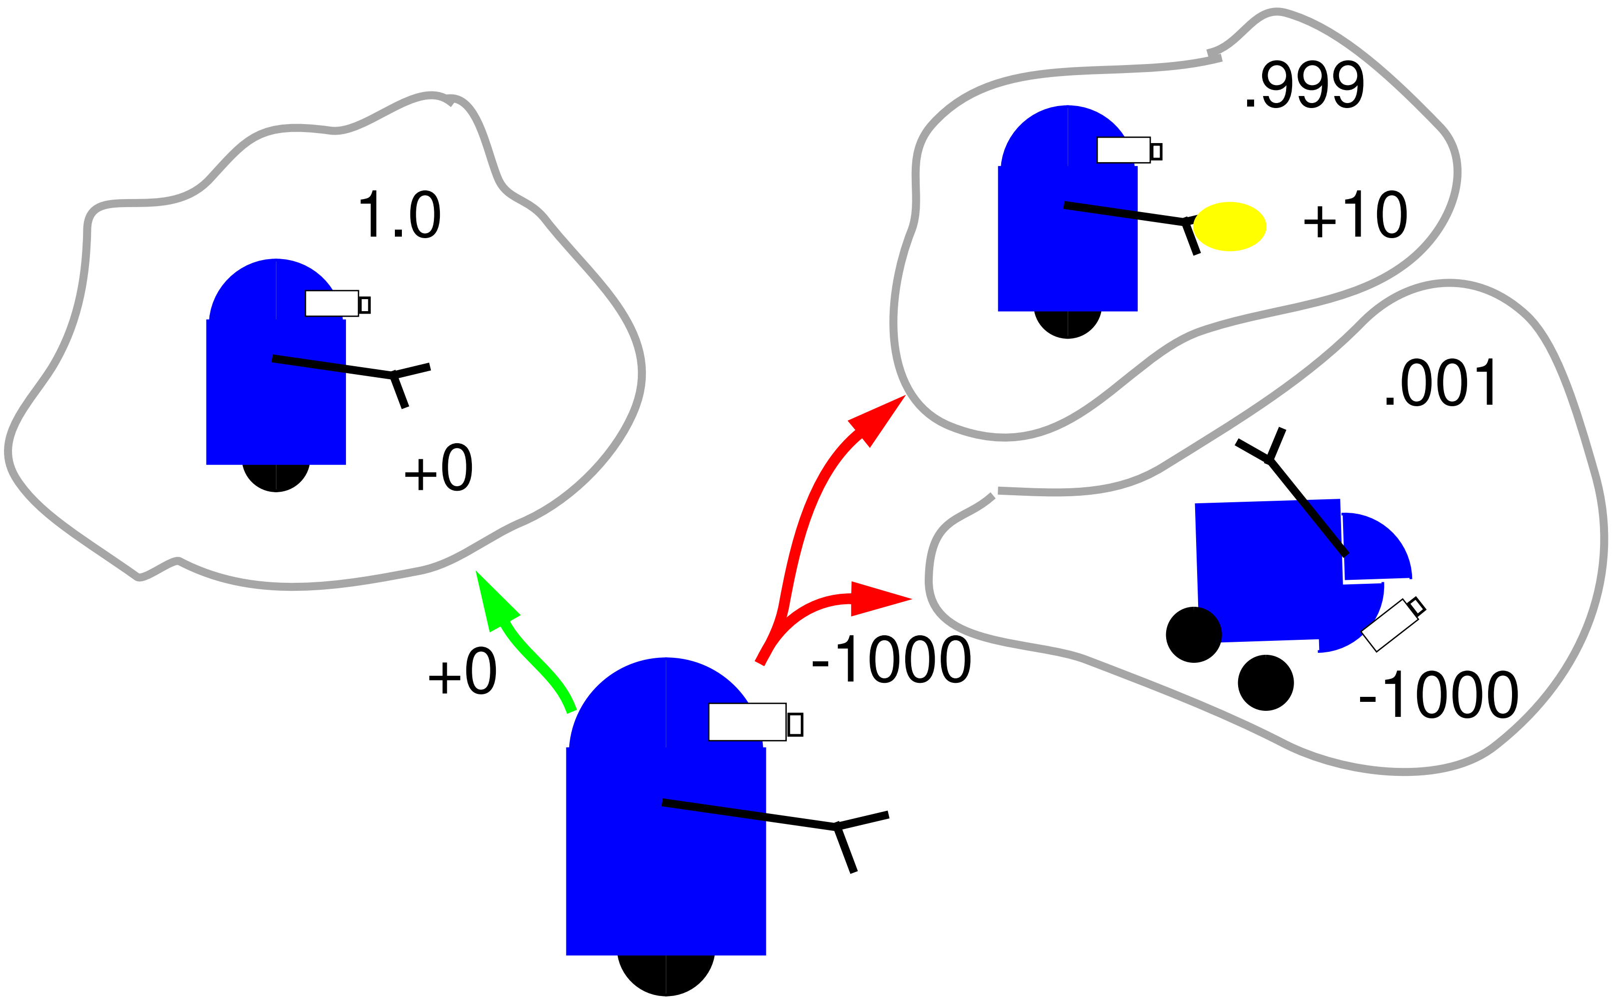

Expected Value World View

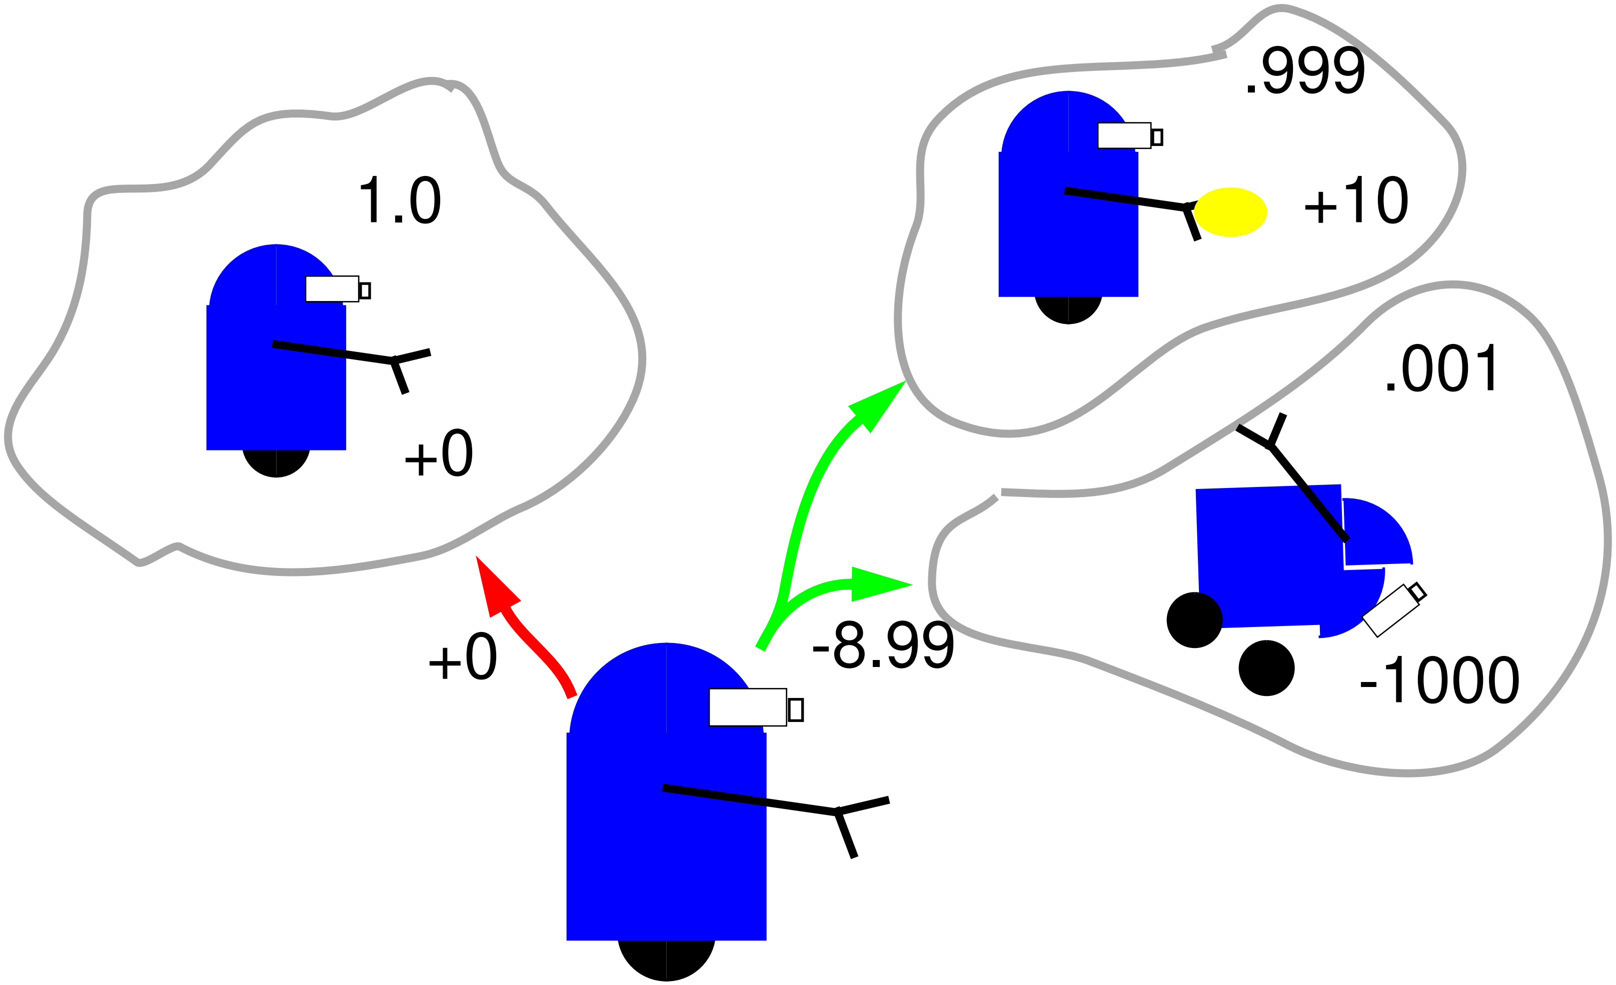

Key Idea

Solve theoretically intractable problems by exploiting structure in the domain

- some apparently large problems may have a much smaller expression and solution
- other large problems can be approximated in a principled way

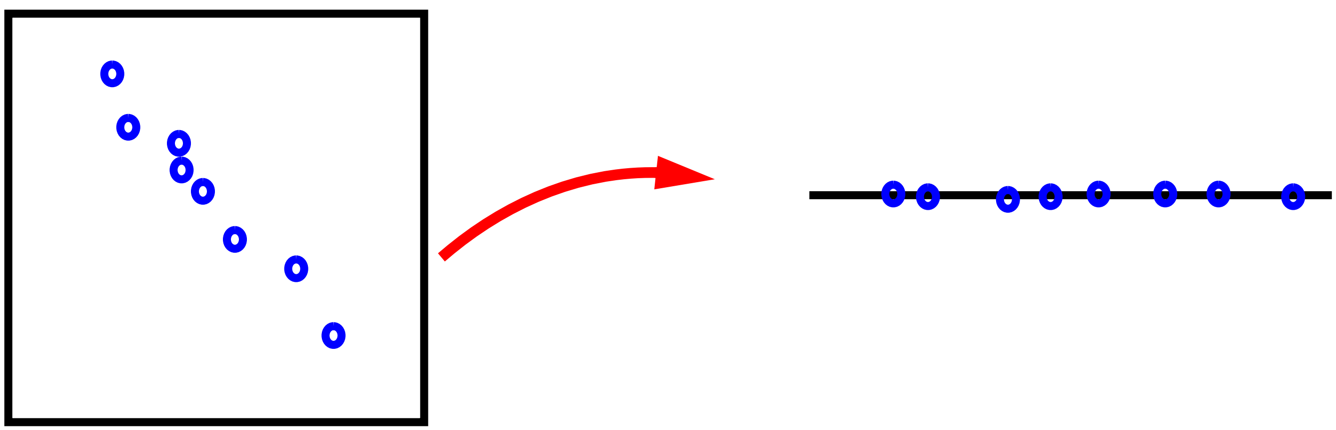

Goals

Solutions for sequential deployment planning problems in very large uncertain domains

- discovery and exploitation of structure in hard problems
- theoretical and empirical analysis

In other words...

Getting Things from Here to There in Multiple Steps

- air-campaign planning
- transportation planning
- troop-deployment planning
- disaster-relief planning
- federal express
- etc...

Air Campaign Scenario

Sudden engagement in theater we've prepared for...

- start with on-the-shelf plan for multi-stage air campaign and current beliefs about situation
- execute first phase of campaign
  - send out aircraft
  - receive intelligence reports
- update situation beliefs
- execute second phase, etc.

Multi-stage uncertainty

It is crucial to consider what might happen on the second stage when planning the first...

If we attack main first, we alert them and they come after us

Secondary is close enough to primary that we can still sneak in

Attack main power station

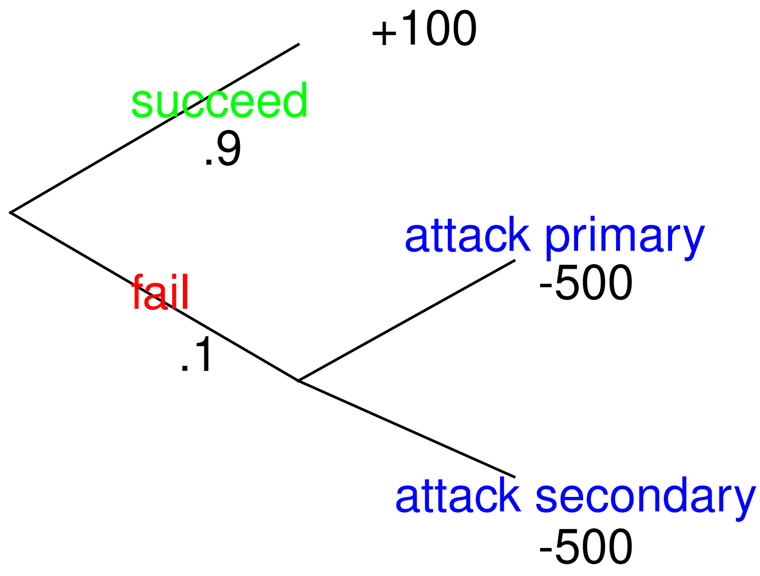

Attack primary distribution site

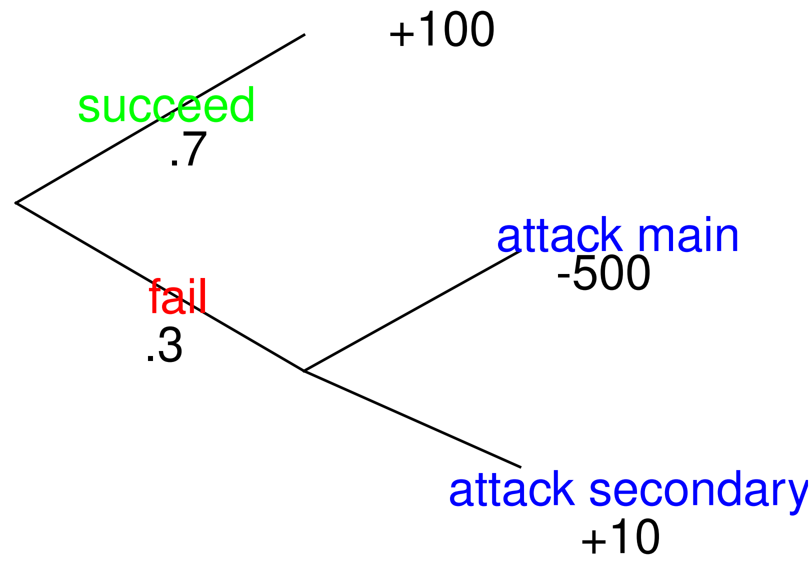

Previous Major Progress

SECT: Plan visualization and evaluation tool for single-stage plans

Plexus: Solve very large MDPs when initial state known

A Tale of Two Planning Models

Classical Model: make a nominal plan; replan on failure

- potentially long planning time
- small plan
- potentially long replanning time

Markov Decision Process Model: plan an action for every possible situation

- long planning time
- huge plan
- instantaneous “replanning”

Plexus

Interpolate between these extremes:

- plan reactions for an envelope of the states
- the more time for planning, the bigger the envelope
- successive approximation
- the bigger the envelope, the more correct and robust the plan

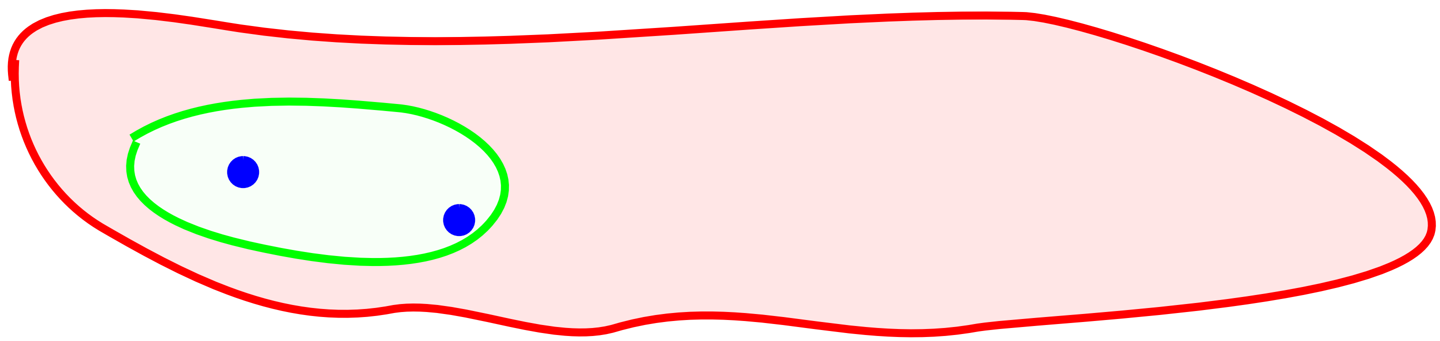

Replanning in Plexus

- most of the time, you're in the envelope: no need to replan
- on exiting the envelope, extend it to include new state and re-solve MDP
- shrink envelope behind and grow in front during execution

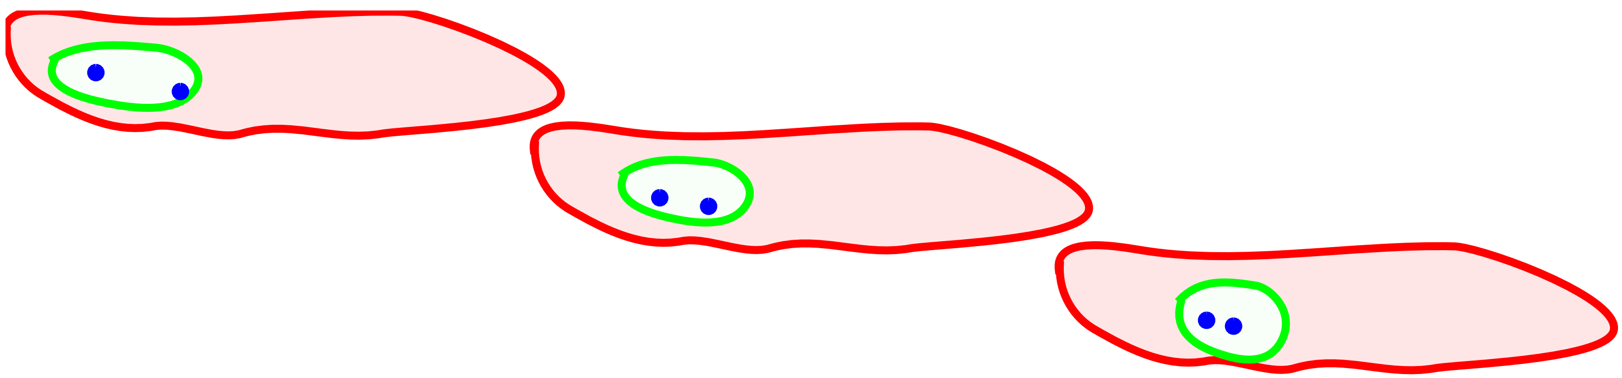

Weaknesses of Current Plexus

- state space is enumerated
- state is completely observable

Partially Observable Markov Decision Processes

- uncertainty in state transitions
  - not sure whether a particular attack *will be* effective
- uncertainty about current state
  - not sure whether a particular attack *was* effective
- choose actions based on the current state of information
  - if we know we took out the powerstation, proceed to next objective
  - otherwise, send out more recon planes

Current Technical Progress

Theoretical contributions

- best exact algorithm
- proof of intractability

Practical contributions

- approximations to optimal solutions
- extremely efficient heuristic solutions
- application to navigation domain

Identify Friend or Foe





- an aircraft is approaching...
- actions: wait, passive recon, active recon, shoot
- knowledge state: probability distribution of types of aircraft

Bridges Out

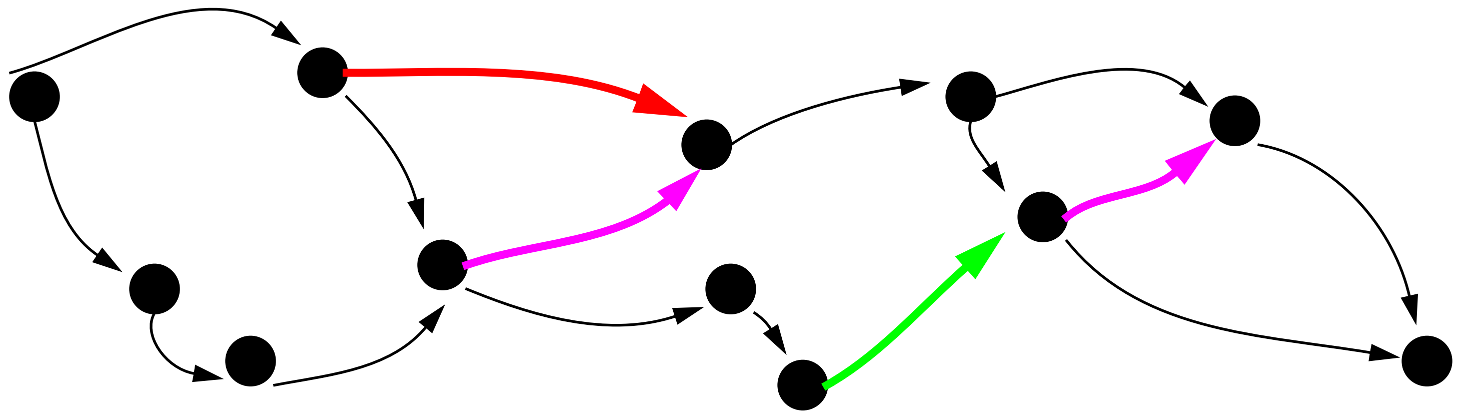

- uncertainty about state of important links
- information obtained when traversal is attempted
- choose actions based on current state of information
- optimize expected path length
- wide applicability

Extremely Large Spaces

- models composed of state variables
- model minimization
- approximate minimization
- value iteration on approximate models
- successive approximation by value-driven partition refinement

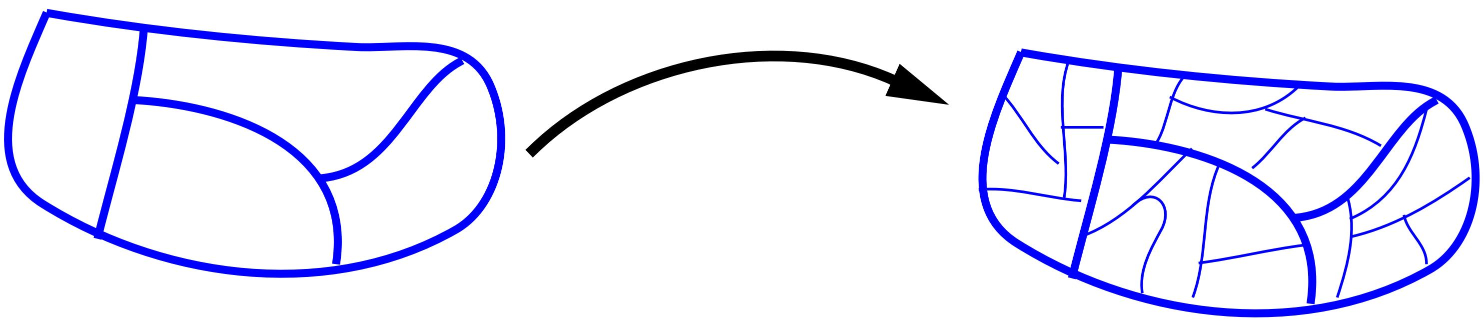

Future Work

- integrate POMDP techniques into Plexus framework
- integrate model minimization and approximation techniques into Plexus framework
- develop framework for handling very large action spaces
- apply to multi-stage aircraft deployment problem

Thrust Areas

We’re formally part of **thrust 4**, working on ***time-sensitive planners***

- reduce replanning time by planning for contingencies in advance
- successive approximation yields anytime planning algorithms

We’re really also in **thrust 2**, working on ***planners that adapt to unpredictable, rapidly changing situations***

- options are kept open through execution
- plan viability is continually reassesed