Model Predictive Control and the Unreasonable Effectiveness of Replanning

Sanjiban Choudhury

Cornell University Founded A.D. 1865 | Cornell Bowers CIS Computer Science

Landscape of Planning / Control Algorithms

Low-level control

High-level path planning

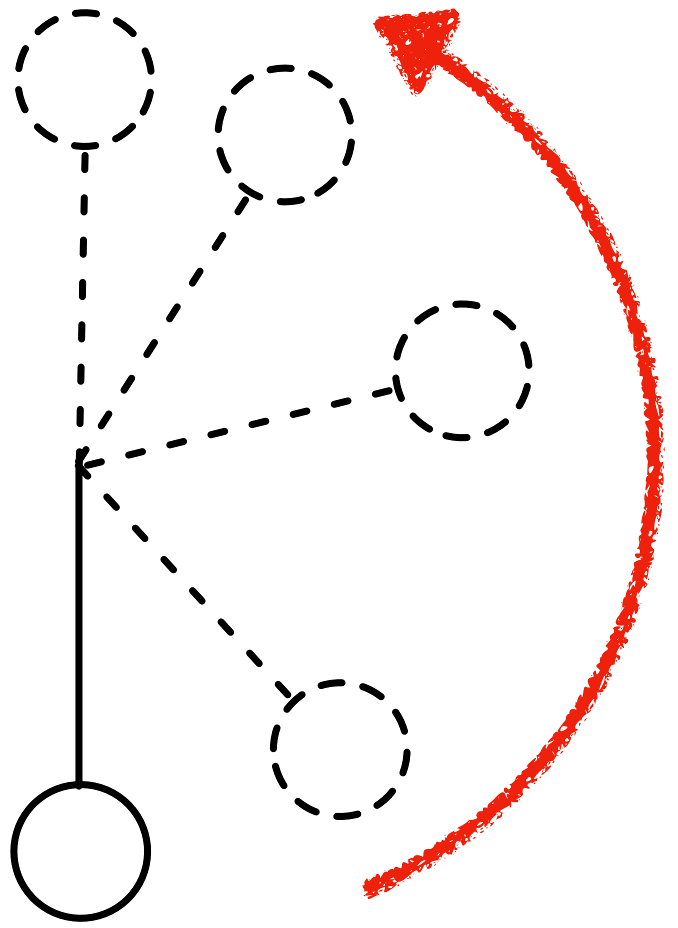

LQR

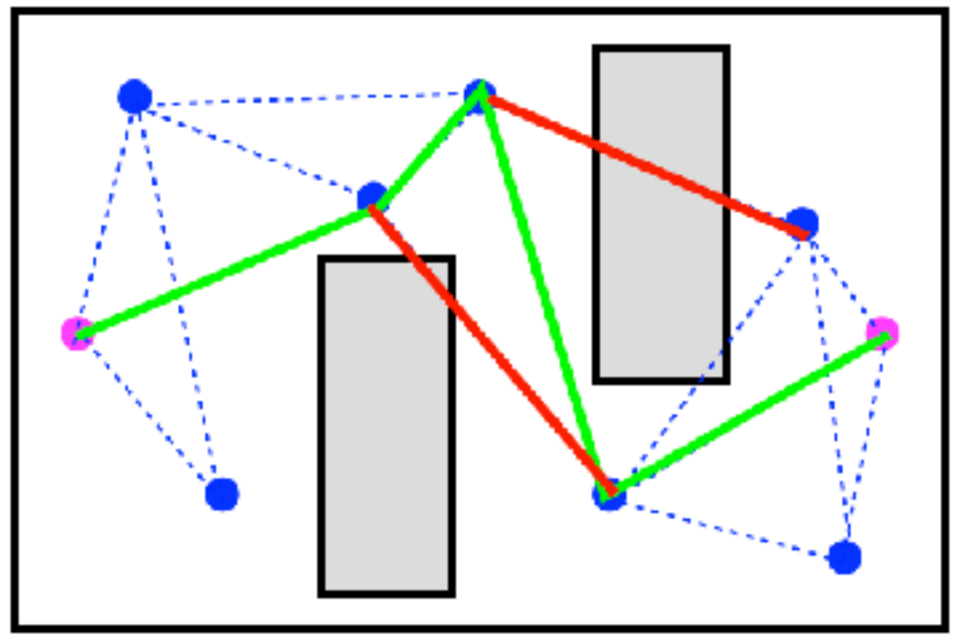

LazySP

Goal: Plan for a real-world helicopter

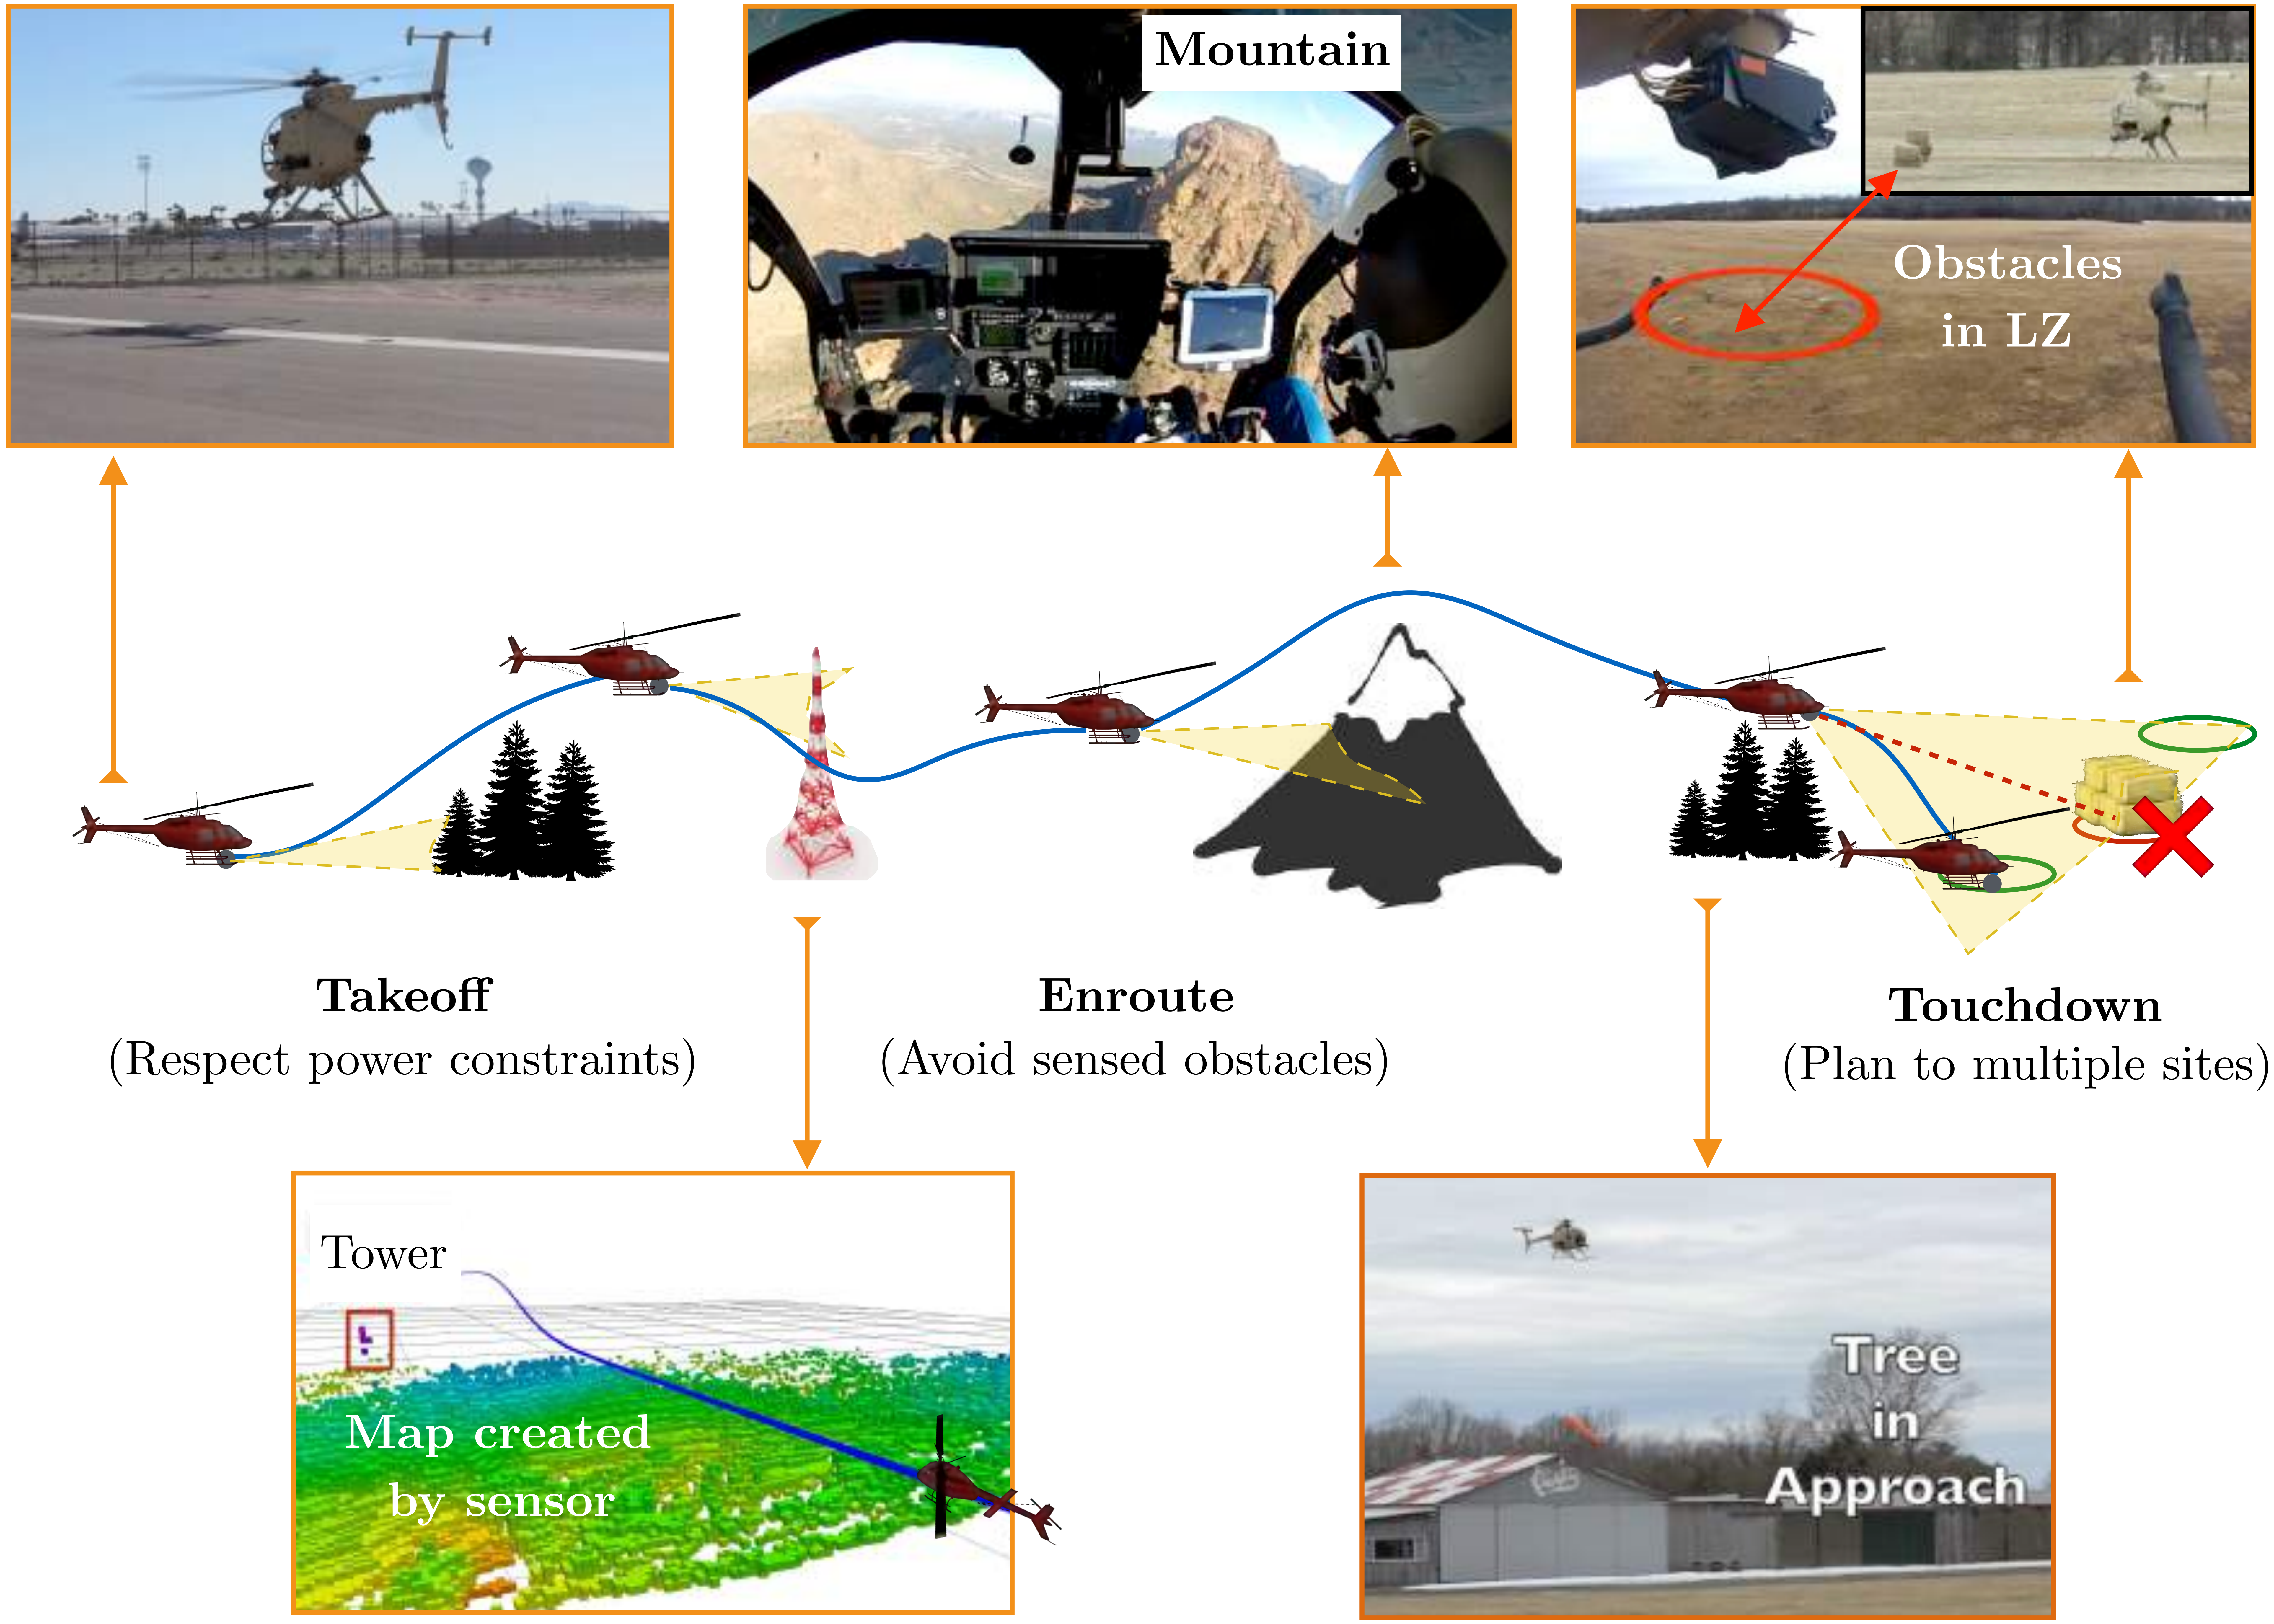
Mountain
Obstacles
in LZ
Takeoff
(Respect power constraints)
Enroute
(Avoid sensed obstacles)
Touchdown
(Plan to multiple sites)
Tower
Map created
by sensor
Tree
in
Approach

Recap: Solving a MDP

$$\min_{a_0,\ldots,a_{T-1}} \sum_{t=0}^{T-1} c(s_t, a_t) \qquad s_{t+1} = \mathcal{T}(s_t, a_t)$$

(Solve for a sequence of actions)

(Sum over all costs)

(Transition function)

Brainstorm: Challenges in solving MDP for helicopter

$$\min_{a_0,\ldots,a_{T-1}} \sum_{t=0}^{T-1} c(s_t, a_t) \qquad s_{t+1} = \mathcal{T}(s_t, a_t)$$

(Solve for a sequence of actions)

(Sum over all costs)

(Transition function)

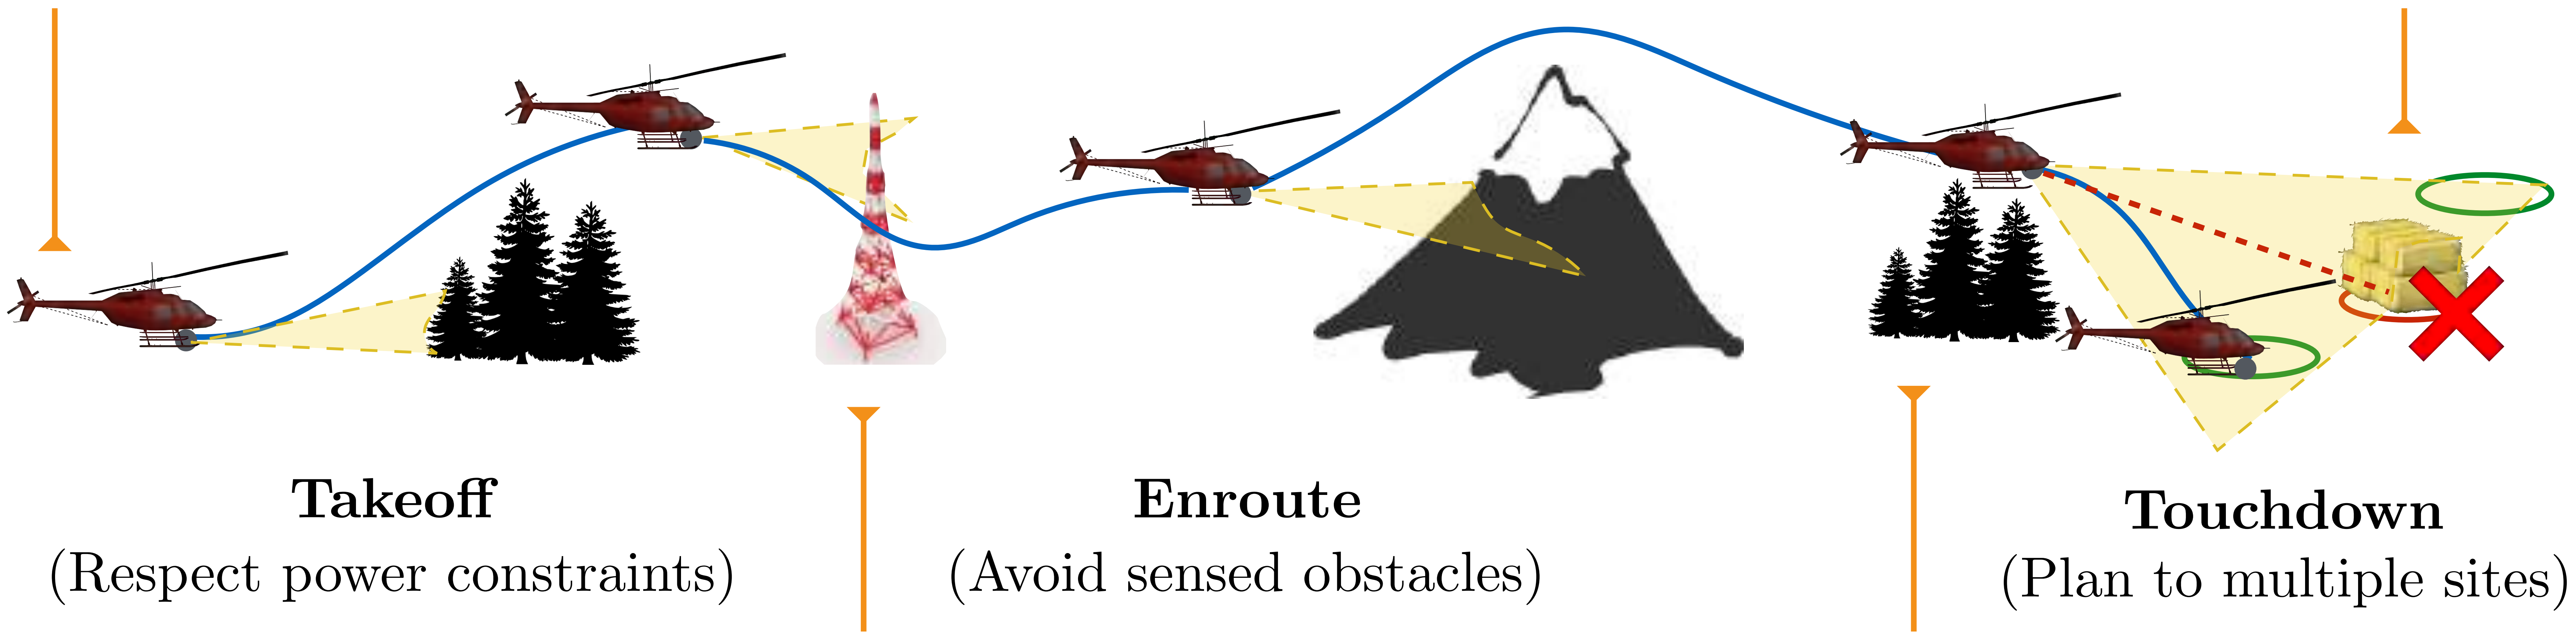

The Big Challenges

Problem 1: Don't know the terrain ahead of time!

Problem 2: Don't have a perfect dynamics model!

Problem 3: Not enough time to plan all the way to the goal!

The Big Challenges

Problem 1: Don't know the terrain ahead of time!

Problem 2: Don't have a perfect dynamics model!

Problem 3: Not enough time to plan all the way to the goal!

Activity!

Brainstorm!

Find a sequence of actions to go from start to goal.

The helicopter can only sense upto 1km.

How should it deal with unknown terrain? What assumptions can it make?

What is the problem mathematically?

$$\min_{a_0, \ldots, a_{T-1}} \sum_{t=0}^{T-1} c(s_t, a_t) \qquad s_{t+1} = \mathcal{T}(s_t, a_t)$$

(Solve for a sequence of actions)

(Sum over all costs)

(Transition function)

Is the transition function fully known?

If not, then how can we solve the optimization problem?

Idea: Plan with an optimistic model

$$\min_{a_0,\ldots,a_{T-1}} \sum_{t=0}^{T-1} c(s_t, a_t)$$

(Solve for a sequence of actions)

(Sum over all costs)

$$s_{t+1} = \hat{\mathcal{T}}(s_t, a_t)$$

(Optimistic Model)

Assume that any unknown space is fully traversable.

Update model as you get information from real world. Replan!

Plan optimistically and replan as you learn more about the world

Be Optimistic and Replan!

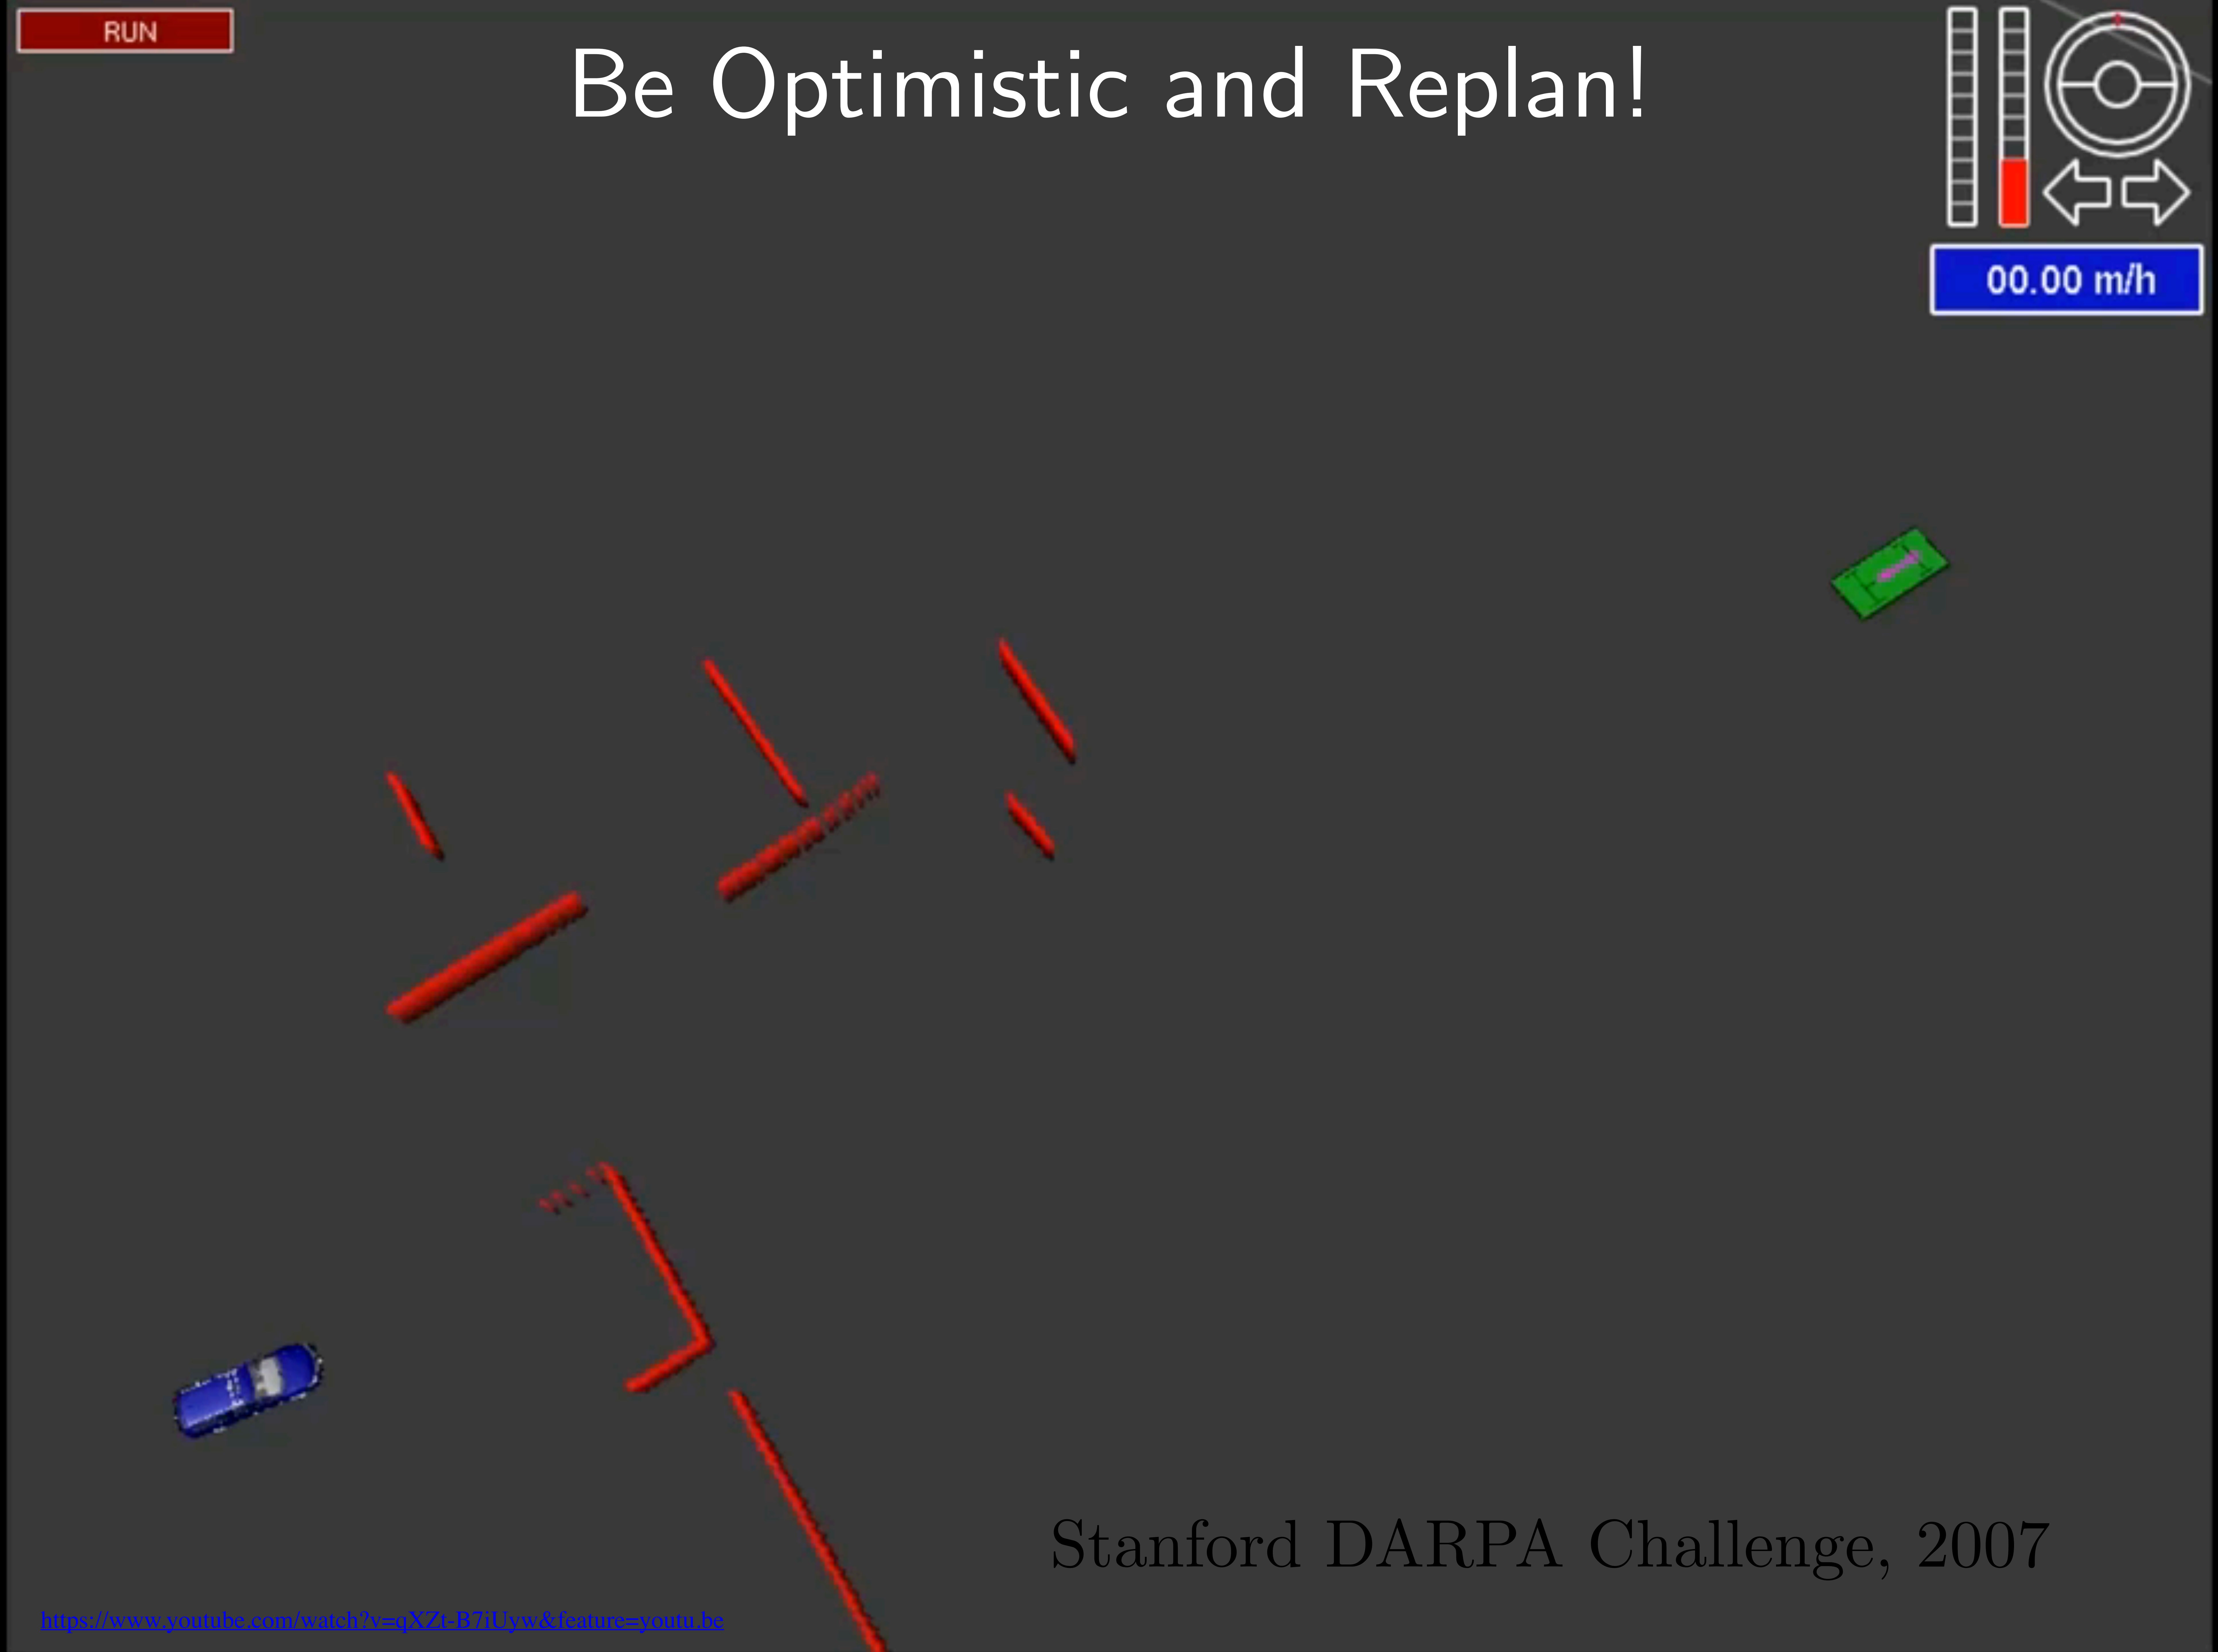

Stanford DARPA Challenge, 2007

https://www.youtube.com/watch?v=qXZt-B7iUyw&feature=youtu.be

Model Predictive Control (MPC)

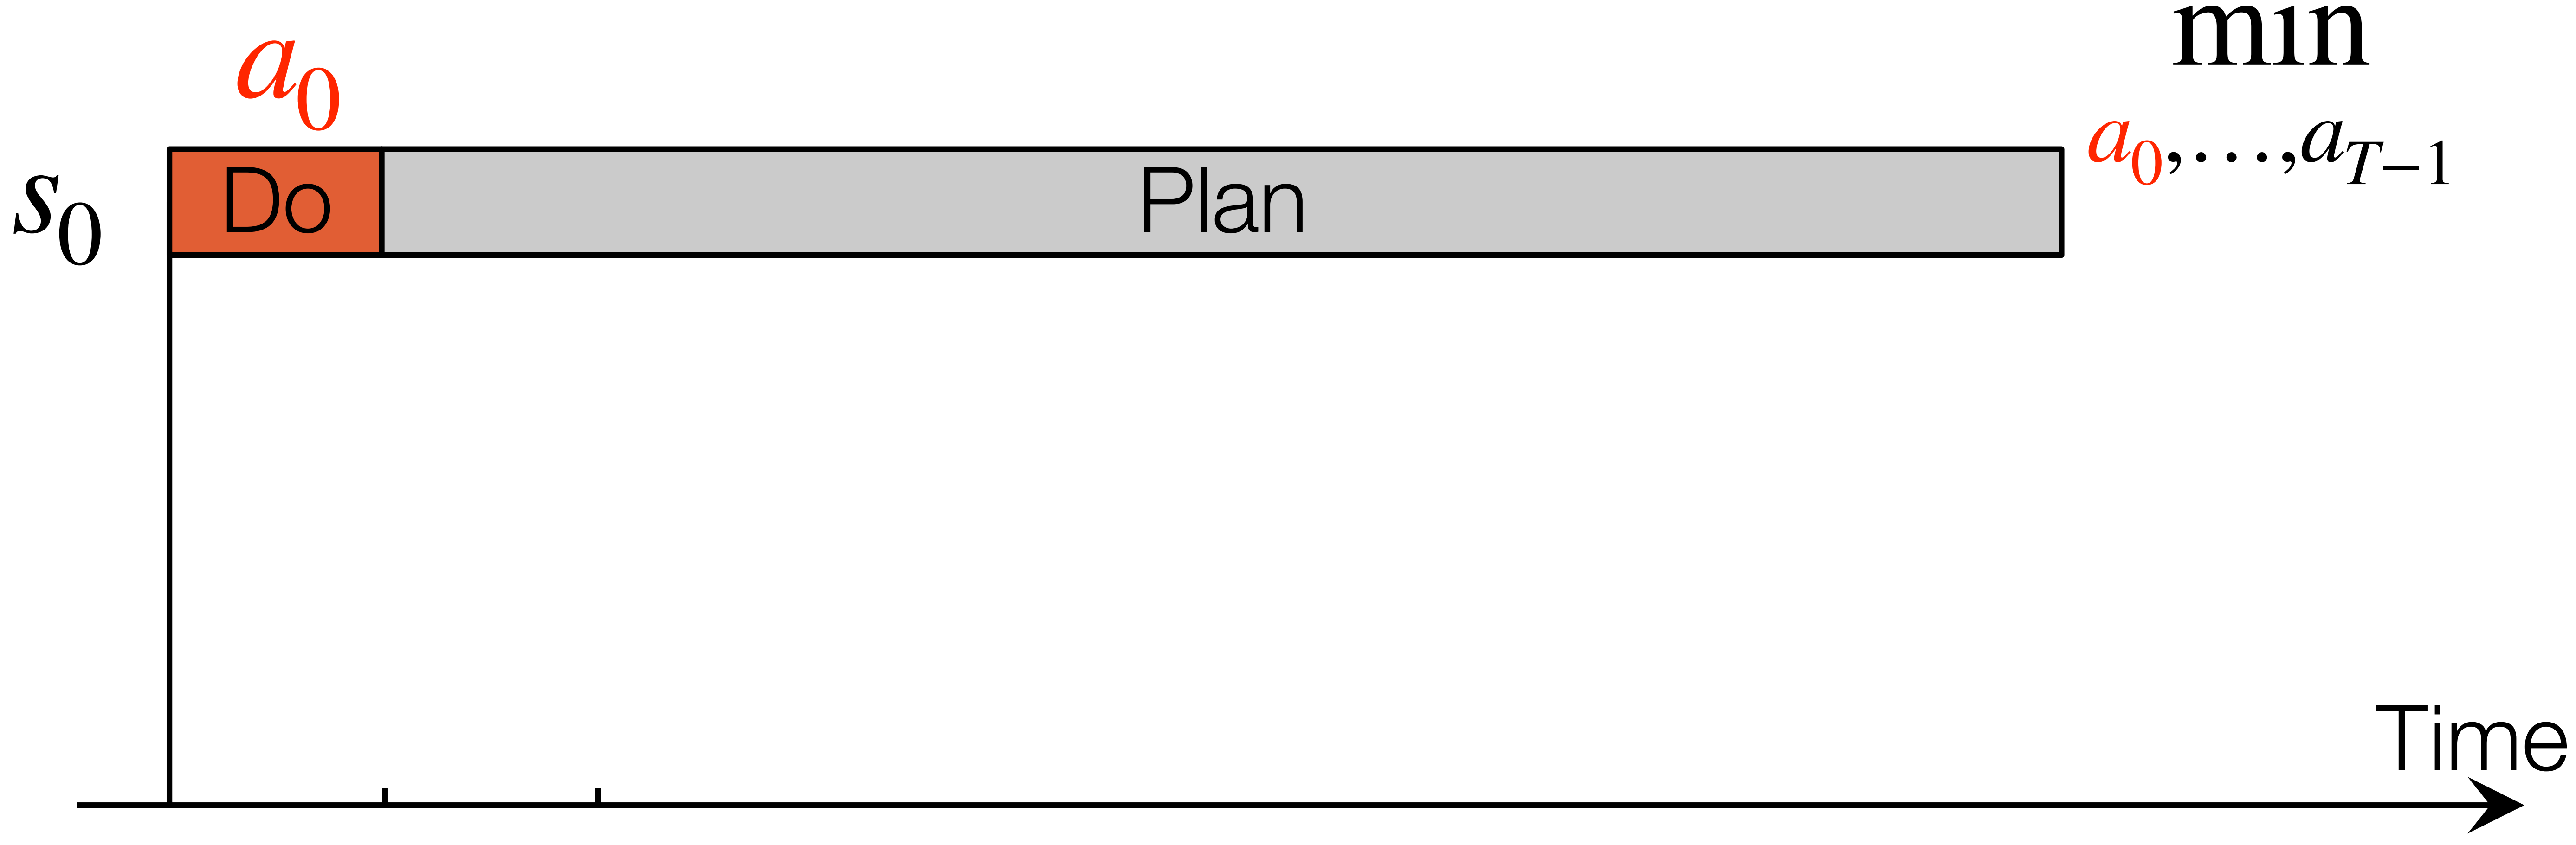

$$\min_{a_0,\ldots,a_{T-1}} \sum_{t=0}^{T-1} c(s_t, a_t)$$

Step 1: Solve current MDP (plan) to find a sequence of actions

Step 2: Execute the first action in the real world and update MDP

Step 3: Repeat!

Model Predictive Control (MPC)

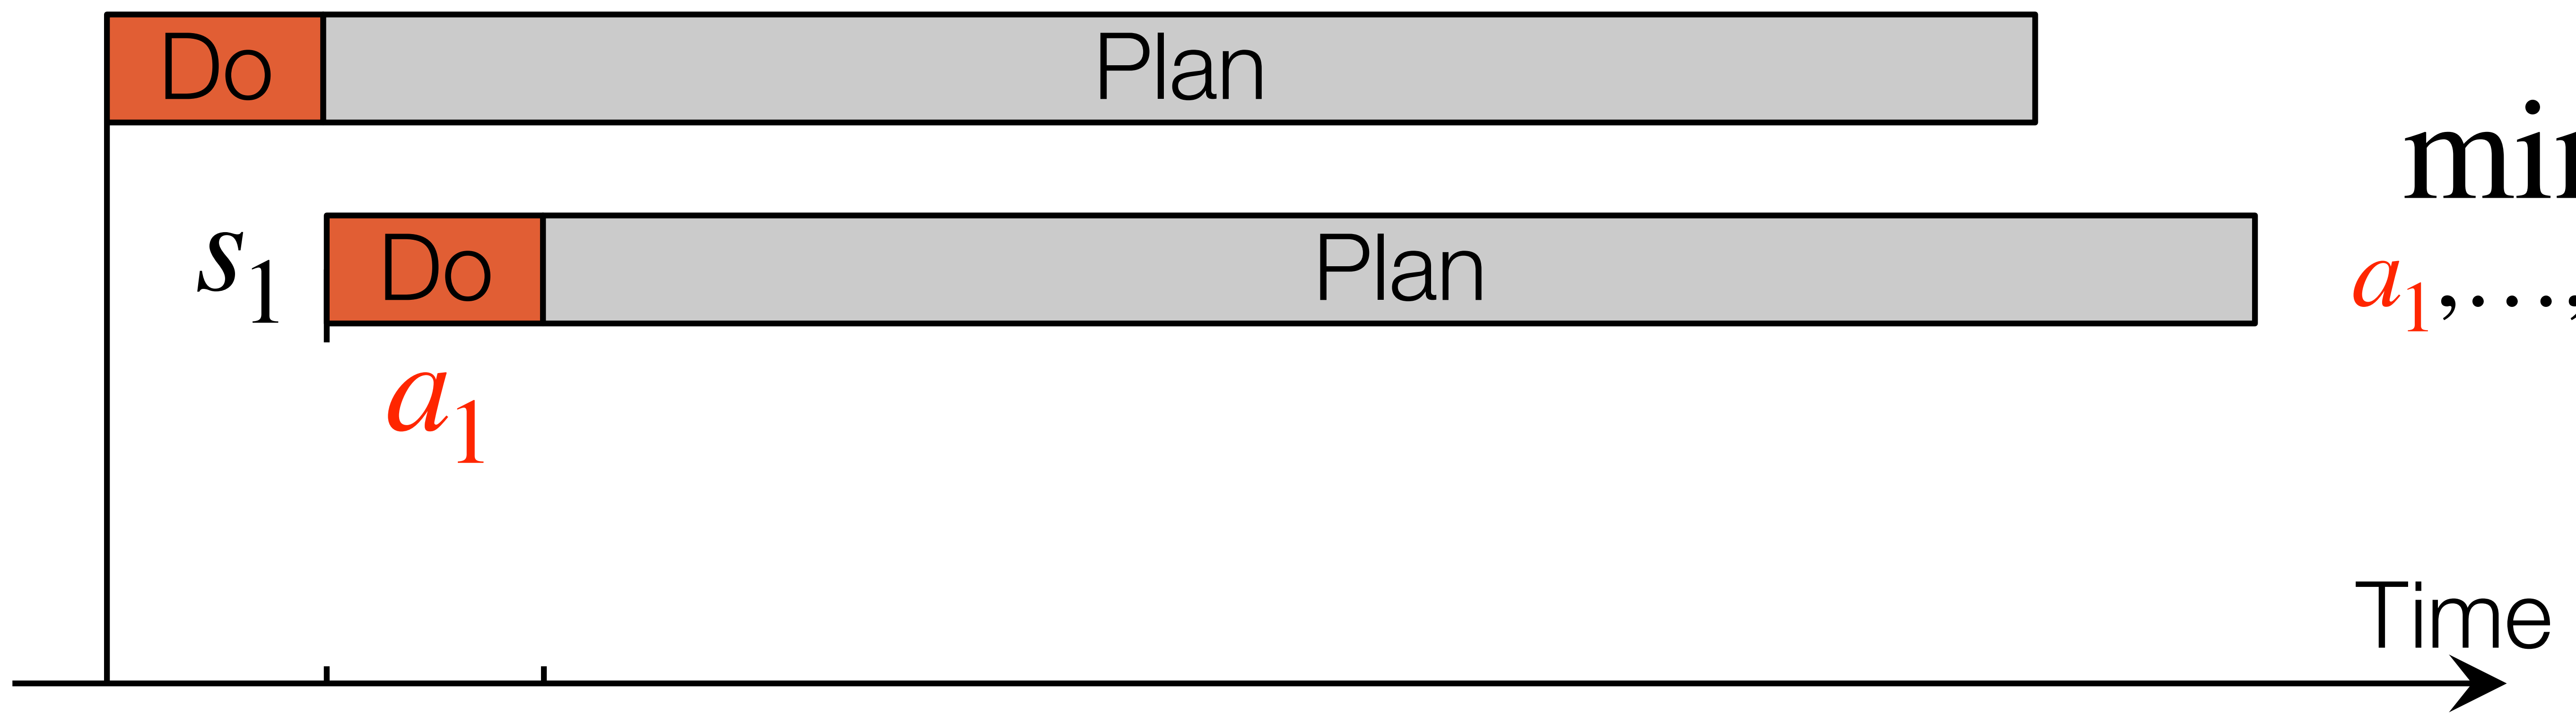

$$\min_{a_1,\ldots,a_T} \sum_{t=1}^{T} c(s_t, a_t)$$

Step 1: Solve current MDP (plan) to find a sequence of actions

Step 2: Execute the first action in the real world and update state

Step 3: Repeat!

Model Predictive Control (MPC)

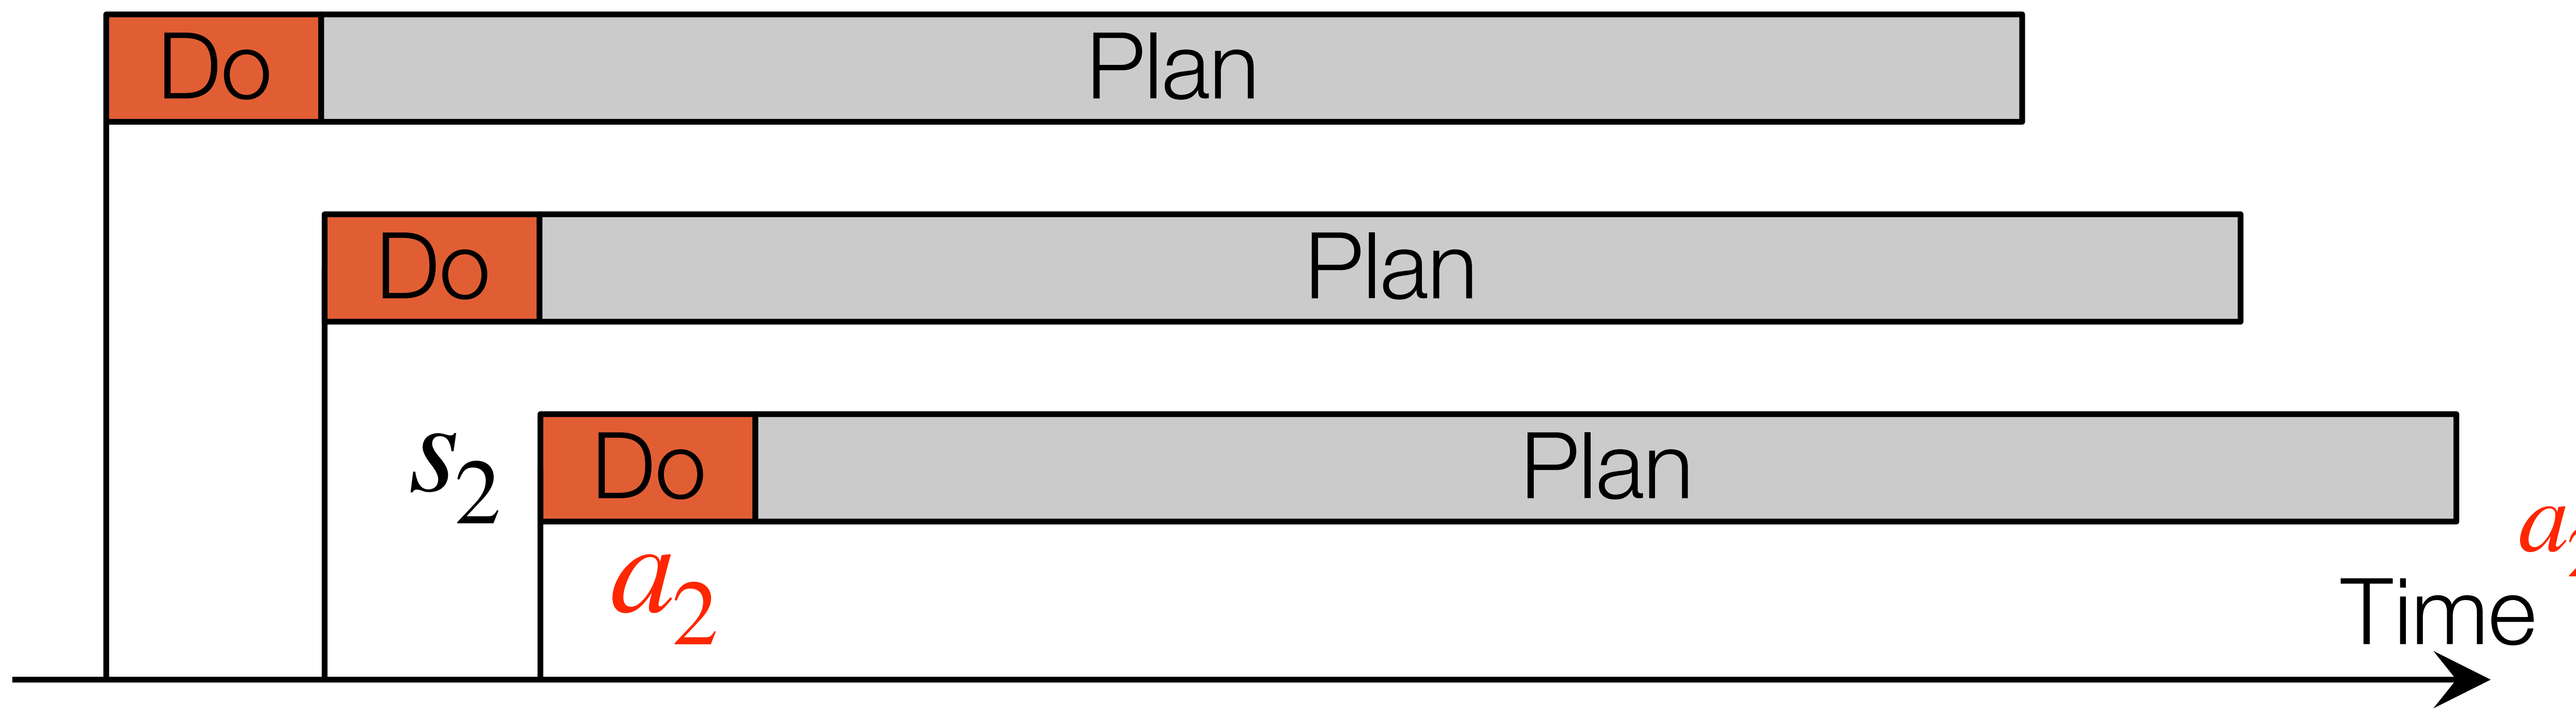

$$\min_{a_2,\ldots,a_{T+1}} \sum_{t=2}^{T+1} c(s_t, a_t)$$

Step 1: Solve current MDP (plan) to find a sequence of actions

Step 2: Execute the first action in the real world and update state

Step 3: Repeat!

The Big Challenges

Problem 1: Don't know the terrain ahead of time!

Problem 2: Don't have a perfect dynamics model!

Problem 3: Not enough time to plan all the way to the goal!

Problem 2: Don't have a perfect dynamics model!

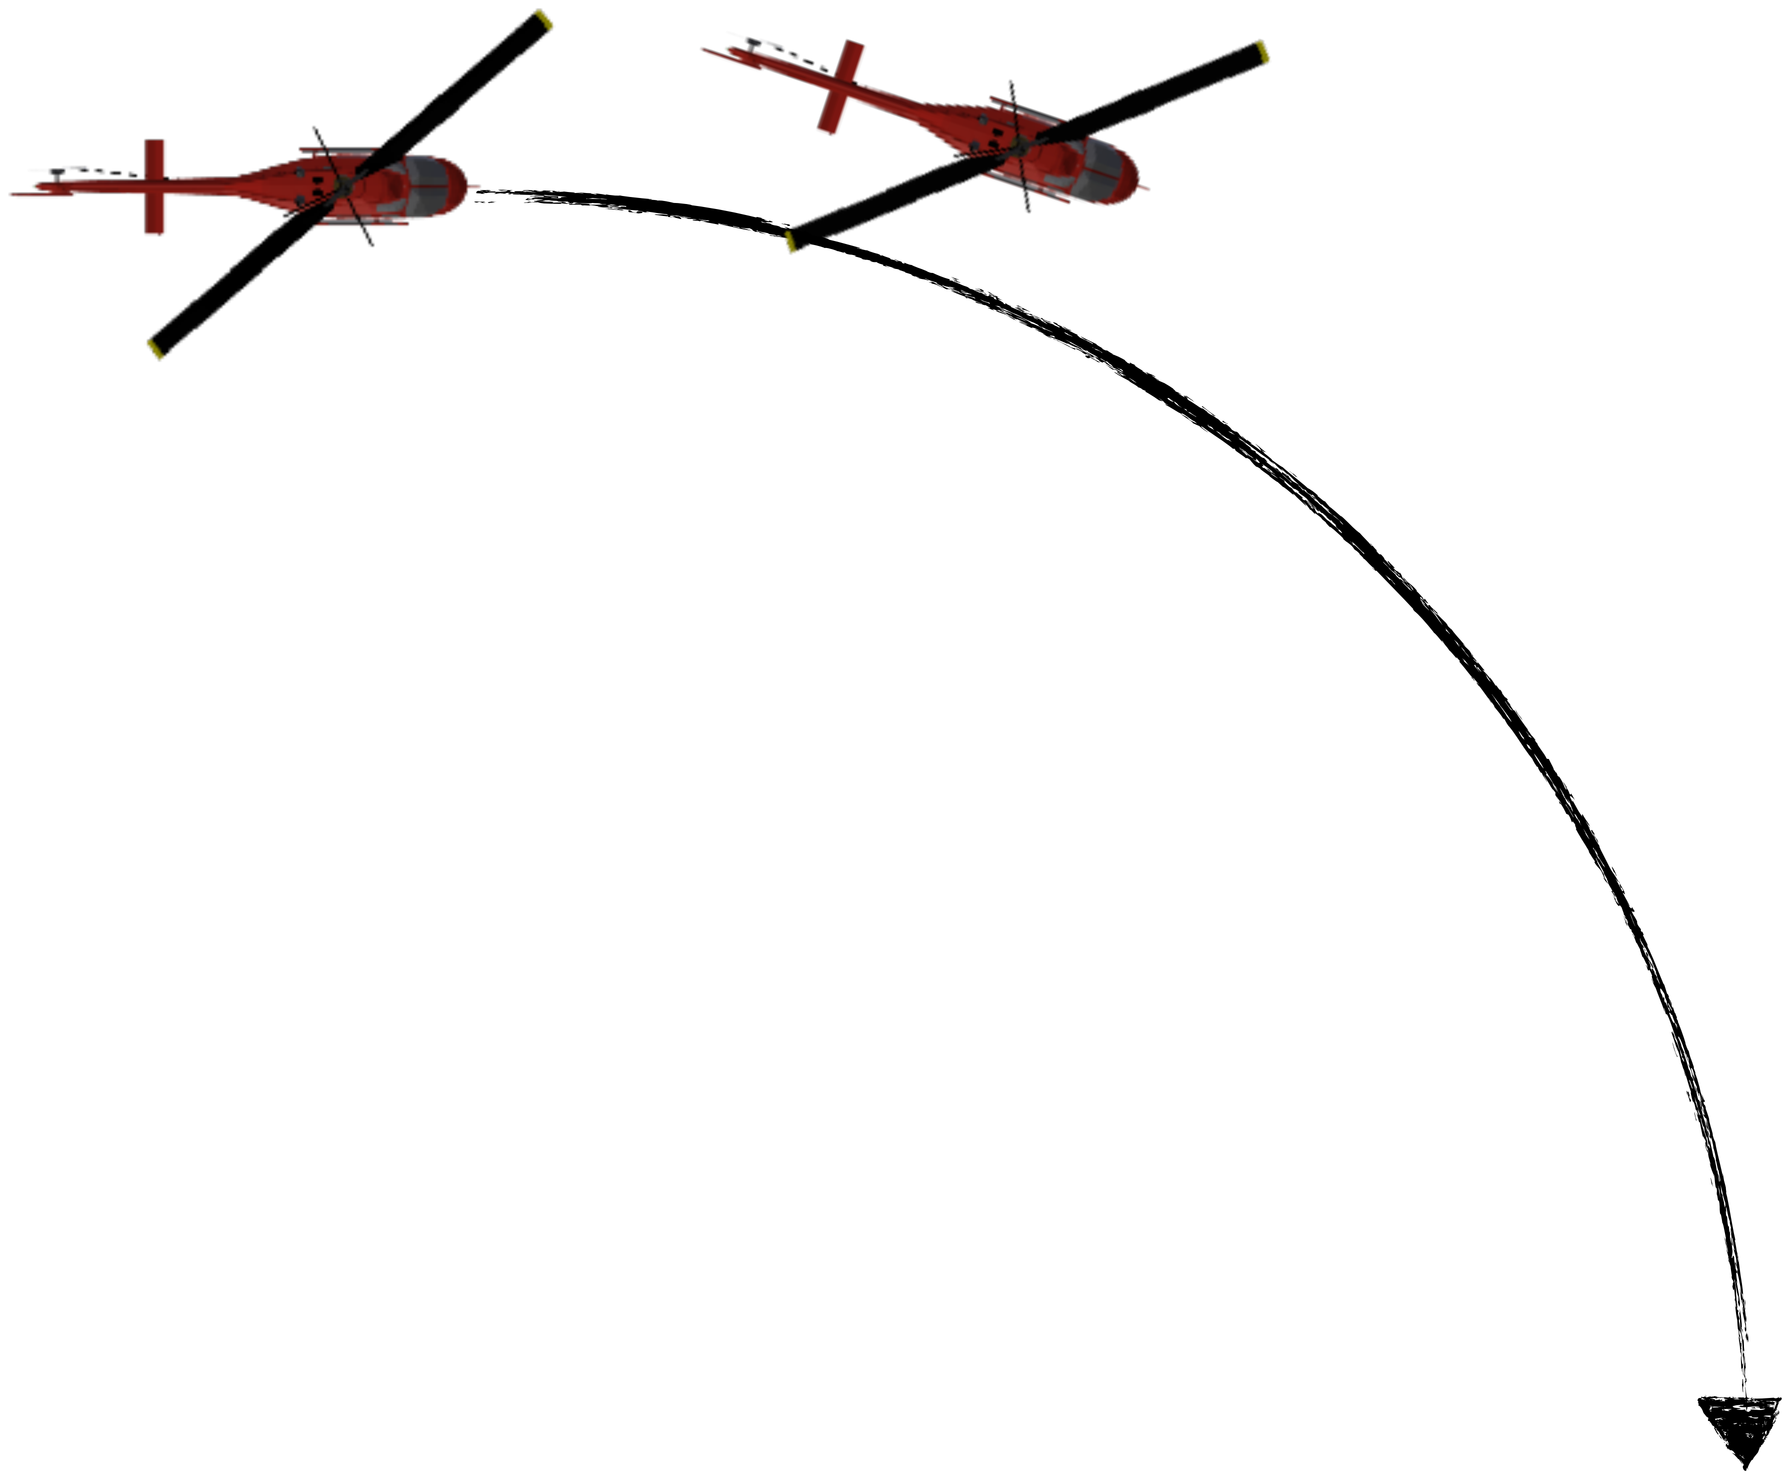

Let's say there is an unknown gust of wind pushing you off the path

What is the problem mathematically?

$$\min_{a_0,\ldots,a_{T-1}} \sum_{t=0}^{T-1} c(s_t, a_t) \qquad s_{t+1} = \mathcal{T}(s_t, a_t)$$

(Solve for a sequence of actions)

(Sum over all costs)

(Transition function)

Is the transition function fully known?

Problem 2: Don't have a perfect dynamics model!

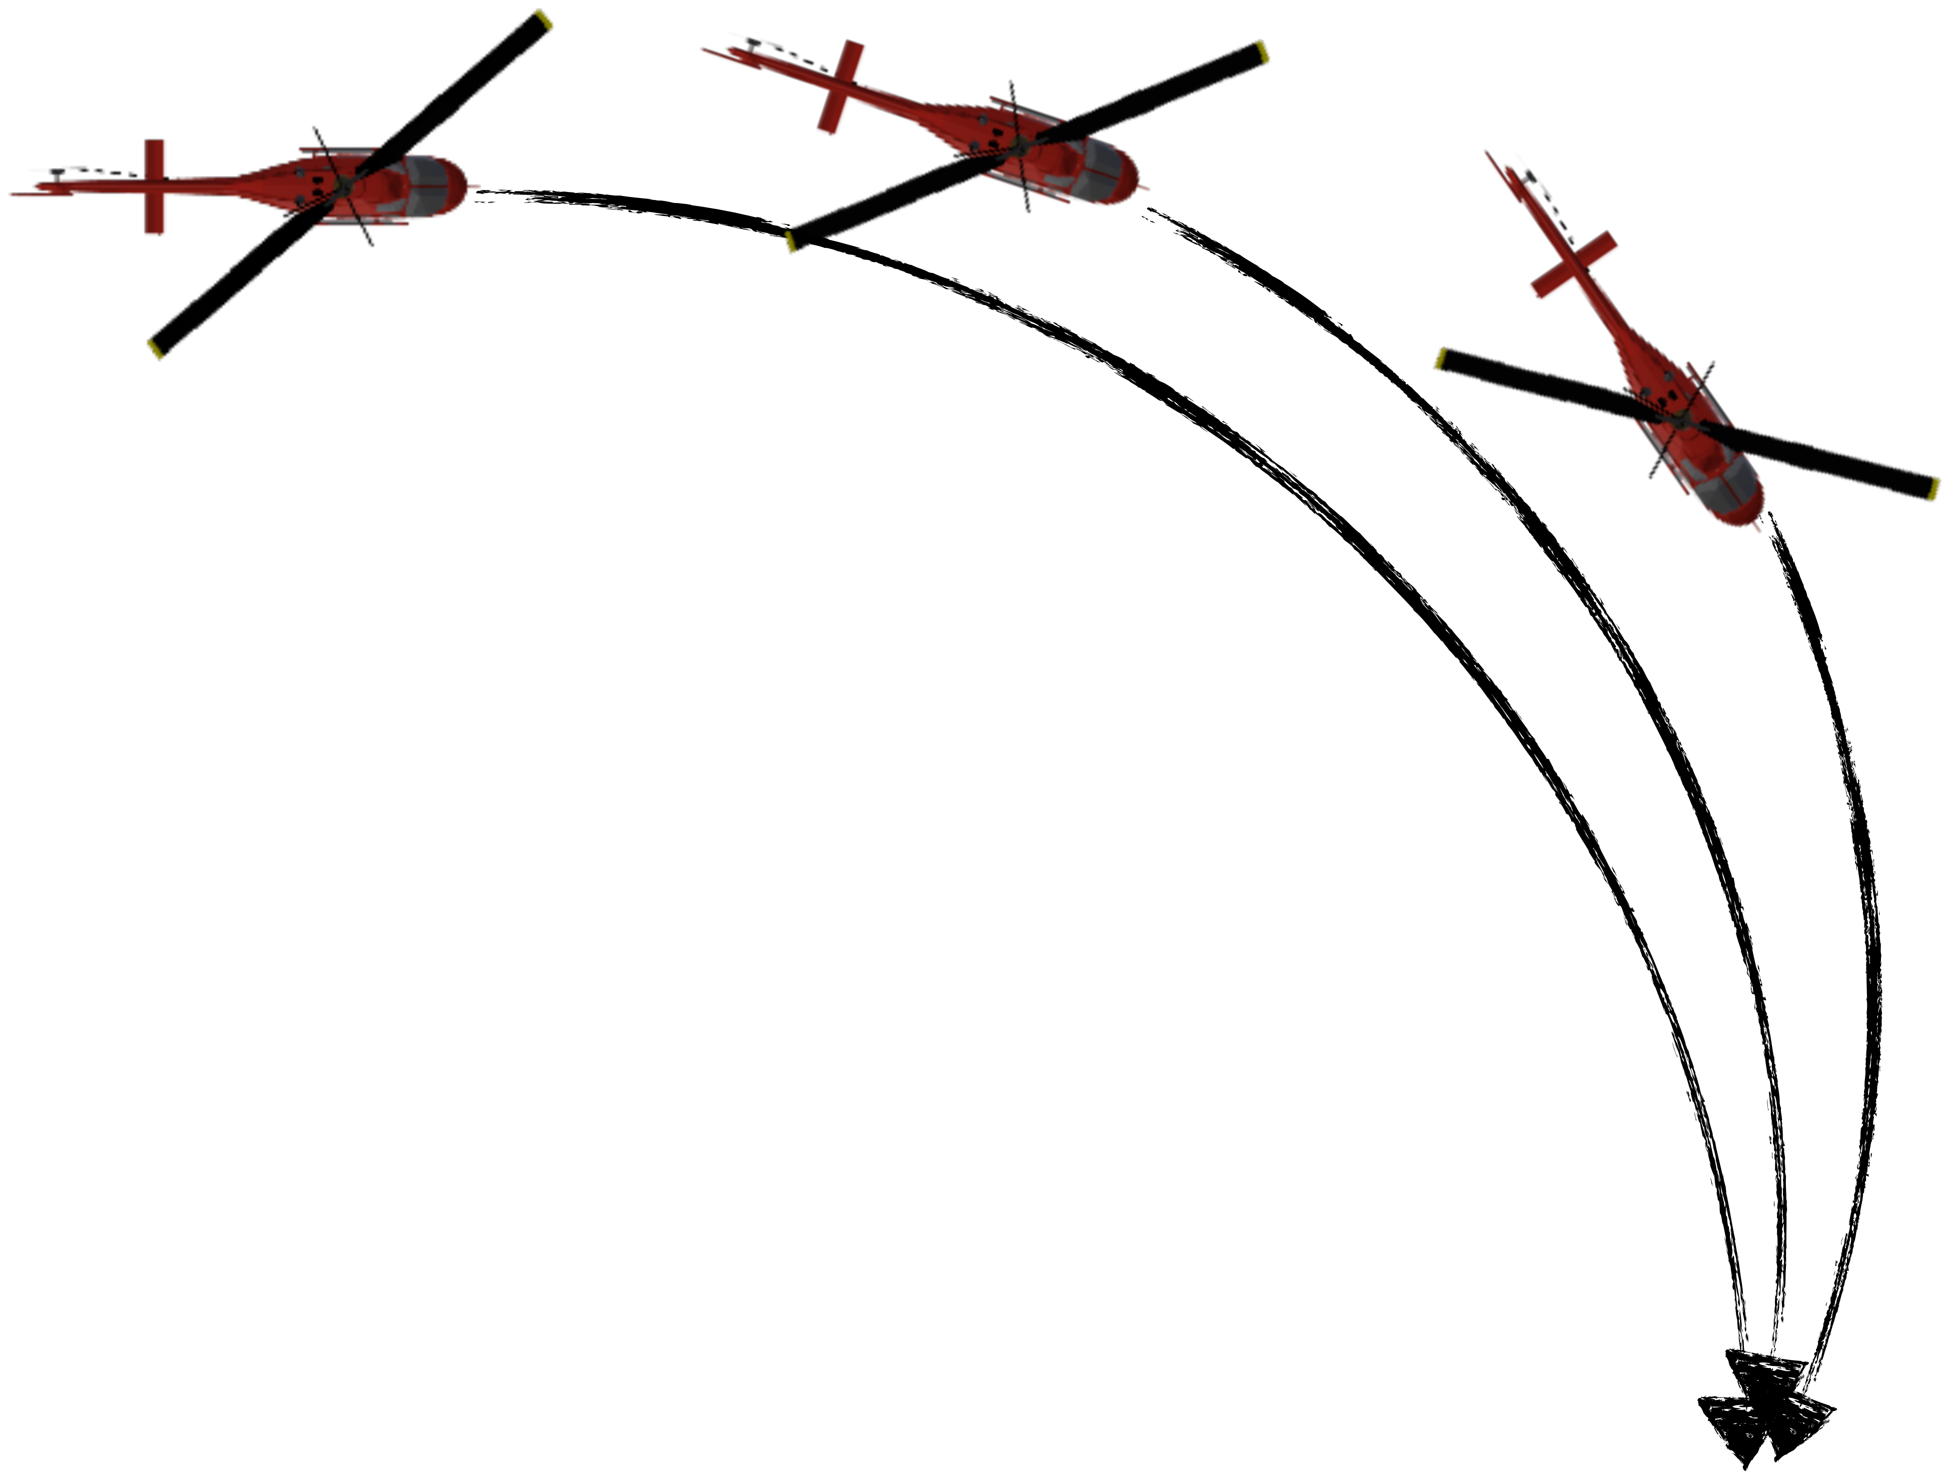

Plan with incorrect transition model and replan!

Theorem:
An optimal policy in an incorrect model has bounded suboptimality in the real model

The Big Challenges

Problem 1: Don't know the terrain ahead of time!

Problem 2: Don't have a perfect dynamics model!

Problem 3: Not enough time to plan all the way to the goal!

Problem 3: Not enough time to plan all the way to goal!

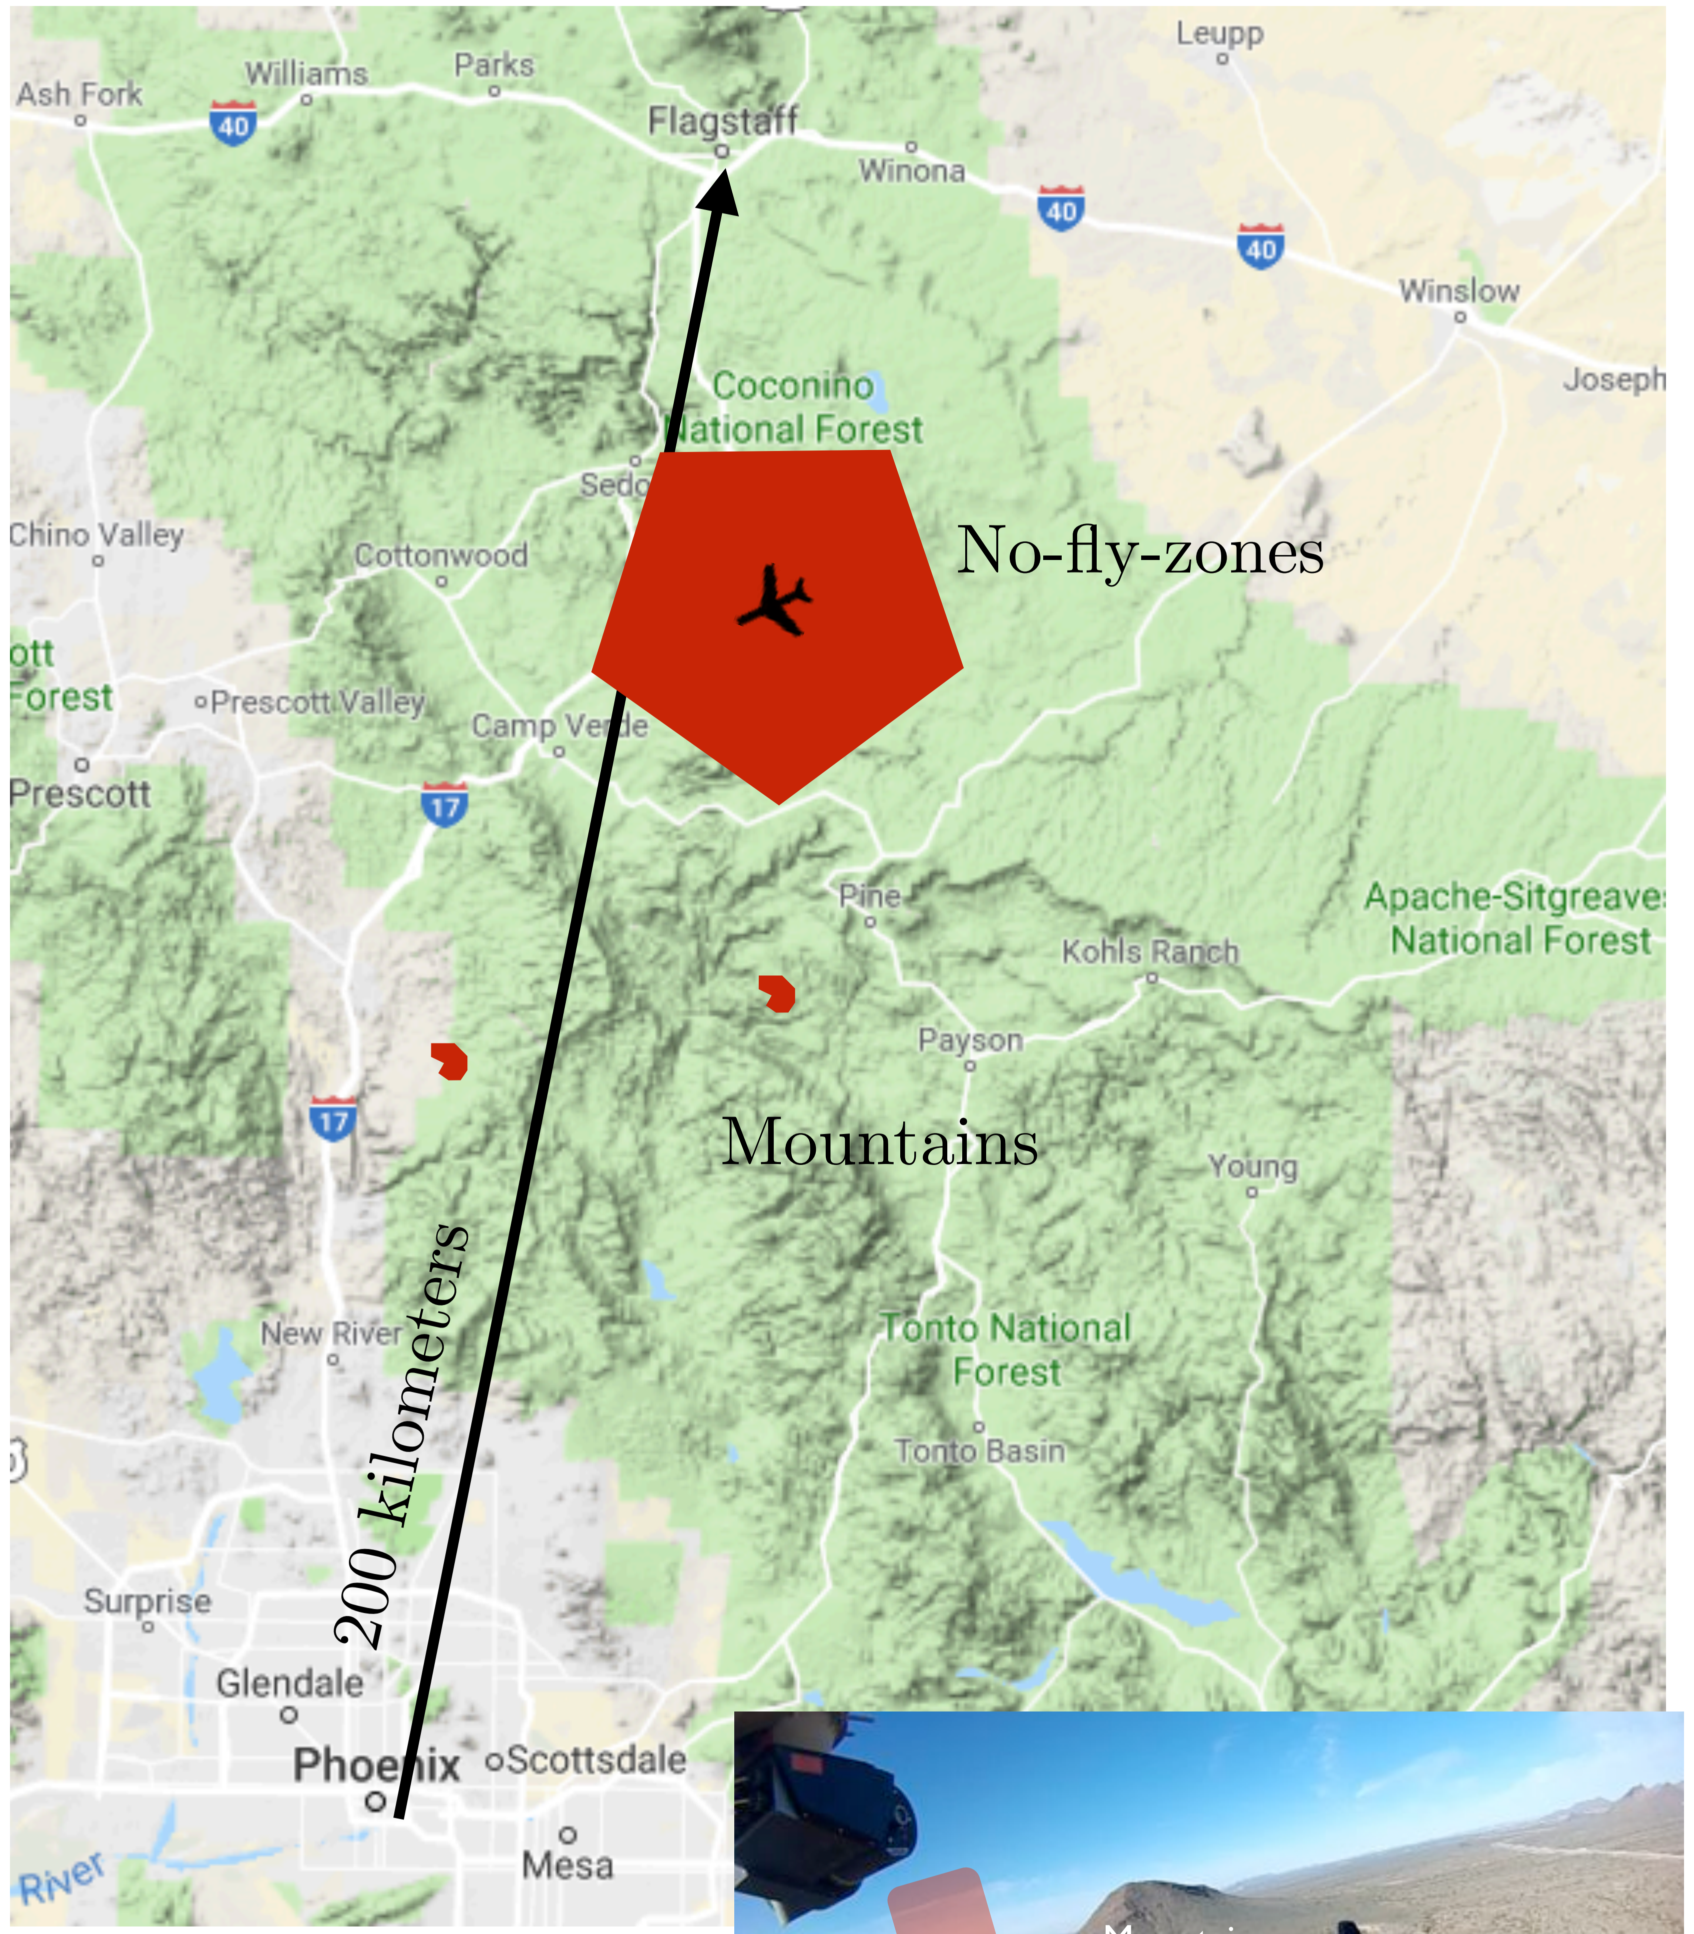

Example mission:

Fly from Phoenix to Flagstaff as fast as possible (200 km)

Problem:

Take forever to plan at high resolution ALL the way to goal

What is the problem mathematically?

$$\min_{a_0,\ldots,a_{T-1}} \sum_{t=0}^{T-1} c(s_t, a_t)$$

(Solve for a sequence of actions)

(Sum over all costs)

How large can T be?

What if we planned till a shorter time horizon T'?

$$\min_{a_0,\dots,a_{T'-1}} \sum_{t=0}^{T'-1} c(s_t, a_t)$$

(Solve for a sequence of actions)

(Sum over all costs)

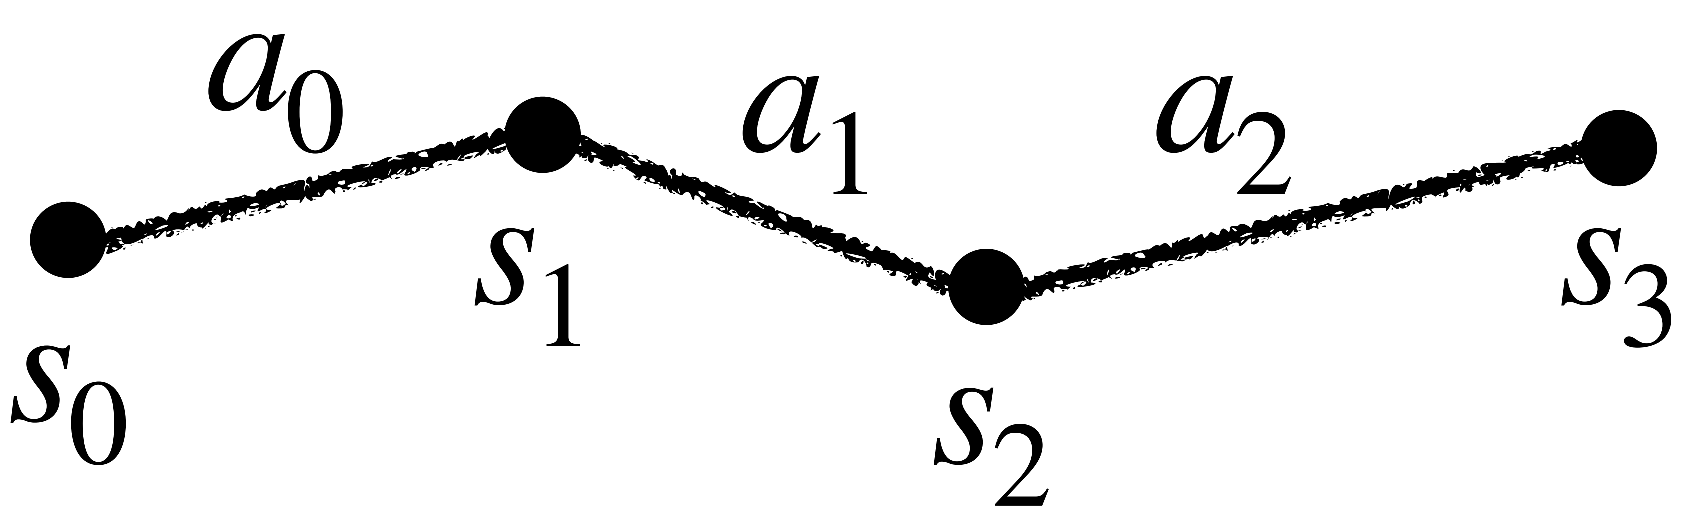

Is this even allowed???

Would we get the same solution for a_0?

We have to add in a terminal value for the final state

$$\min_{a_0,\ldots,a_{T'-1}} \sum_{t=0}^{T'-1} c(s_t, a_t) \quad + V^\star(s'_T)$$

(Solve for a sequence of actions)

(Sum over all costs)

(Optimal value of state $s_{T'}$)

Can we compute the optimal value V*?

If not, how can we approximate it

Idea: Use a global planner to approximate $\hat{V}^{\star}$

$$\min_{a_0,\ldots,a_{T'-1}} \quad \sum_{t=0}^{T'-1} c(s_t, a_t) \quad + \hat{V}^{\star}(s'_T)$$

(Solve for a sequence of actions)

(Sum over all costs)

(Approximate value of state $s_{T'}$)

For example: Run a 2D planner from s_T to the goal

Use the cost of that plan to compute approximate value