Autonomy in the Real World

Air Transportation Systems Architecting

March 18th, 2004

Nicholas Roy

AERO | ASTRO

What is Autonomy?

- **Autonomous**
  - From the Greek: “auto” + “matos”
    “self-willed”
- **Robot**
  - From the Czech/Polish word: “robota” and “robotnik”
    “labour” and “workman”
  - First used in Capek’s play “R.U.R.”

What is Autonomy?

- Autonomy software performs the sophisticated reasoning and decision making needed to accomplish user goals with limited human intervention.
- **Autonomous** is much more than **Automated**
- **Automated**: low-level, mechanical decisions
  - (if-then, control law)
  - designed for a limited class of situations.
- **Autonomous**: sophisticated system-level decisions.
  - can deal with many situations, including the unexpected.
  - can deal with situations that automated systems cannot.

Unstructured Worlds

Groundhog

Movie courtesy of S. Thrun

Pearl

Why Autonomy?

- Inhospitable environments
- Remote environments
- High-precision tasks
- High-fatigue tasks
- Disagreeable tasks

Fitts' List

Attribute	Machine	Human
Speed	Superior	Comparatively slow
Power Output	Superior in level in consistency	Comparatively weak
Consistency	Ideal for consistent, repetitive action	Unreliable, learning & fatigue a factor
Information Capacity	Multi-channel	Primarily single channel
Memory	Ideal for literal reproduction, access restricted and formal	Better for principles & strategies, access versatile & innovative
Reasoning Computation	Deductive, tedious to program, fast & accurate, poor error correction	Inductive, easier to program, slow, accurate, good error correction
Sensing	Good at quantitative assessment, poor at pattern recognition	Wide ranges, multi-function, judgment
Perceiving	Copes with variation poorly, susceptible to noise	Copes with variation better, susceptible to noise

- **inductive** and **deductive**. Induction is usually described as moving from the specific to the general, while deduction begins with the general and ends with the specific; arguments based on experience or observation are best expressed inductively, while arguments based on laws, rules, or other widely accepted principles are best expressed deductively.

Slide courtesy of M. Cummings

Hollnagel, 2000

Different Kinds of Autonomy

- Model-based
  - Williams, Latombe
- Control theory
  - So many....

- Machine learning, probabilistic models
  - Thrun, Leonard

- Reactive
  - Brooks
- Behaviour-based
  - Arkin, Mataric

“Sensor-based” Autonomy

Agent acting in
the real world

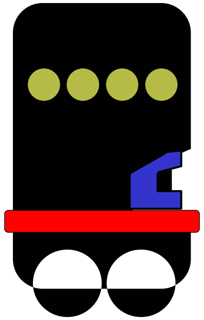

“Sensor-based” Autonomy

Agent taking actions

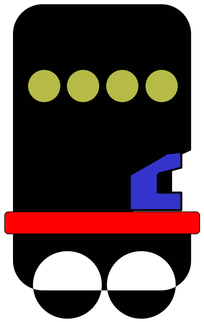

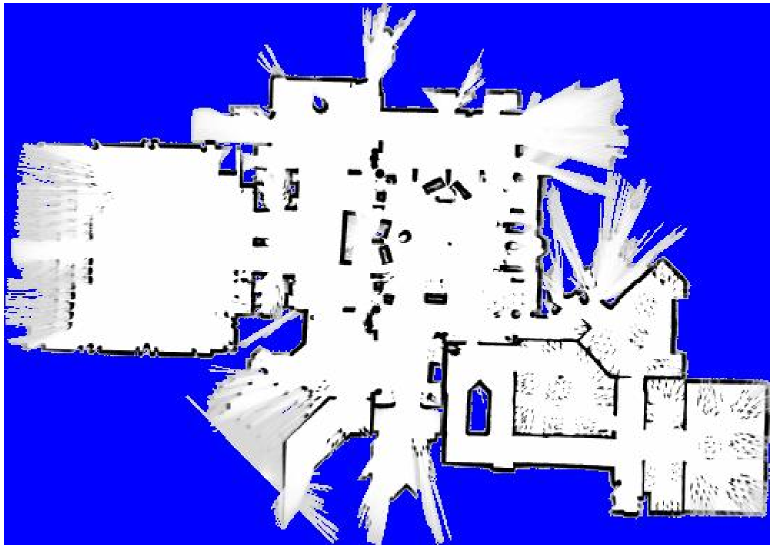

Environmental Model

"Sensor-based" Autonomy

Agent trying to achieve some goal

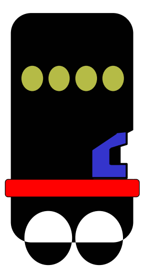

Probabilistic State Estimation

Probabilistic Decision Making

Machine Learning

Environmental Model

Minerva

- Don't oversell Minerva – control problem/hci was really simple, interactivity was very minimal.

- Many of you may recognize Minerva, a robot that I had the privilege of working on a few years ago. Minerva was deployed in the Smithsonian in Washington, giving tours of various exhibits to visitors. The robot took requests from people for different tours, navigating autonomously from exhibit to exhibit. As you can see, these were some fairly demanding conditions for a mobile robot, crowded and full of children.

- Nevertheless, Minerva was able to handle these conditions robustly for a number of reasons. and chief among them…

Another view of Autonomy

	State estimation	Decision making
System operations	Reconfiguration	Fault Diagnosis
Mobile robotics	Global position estimation, mapping	Motion Planning
Classical control	Local state estimation	Classical control

Slide courtesy of Brian Williams

Remote Agent Architecture

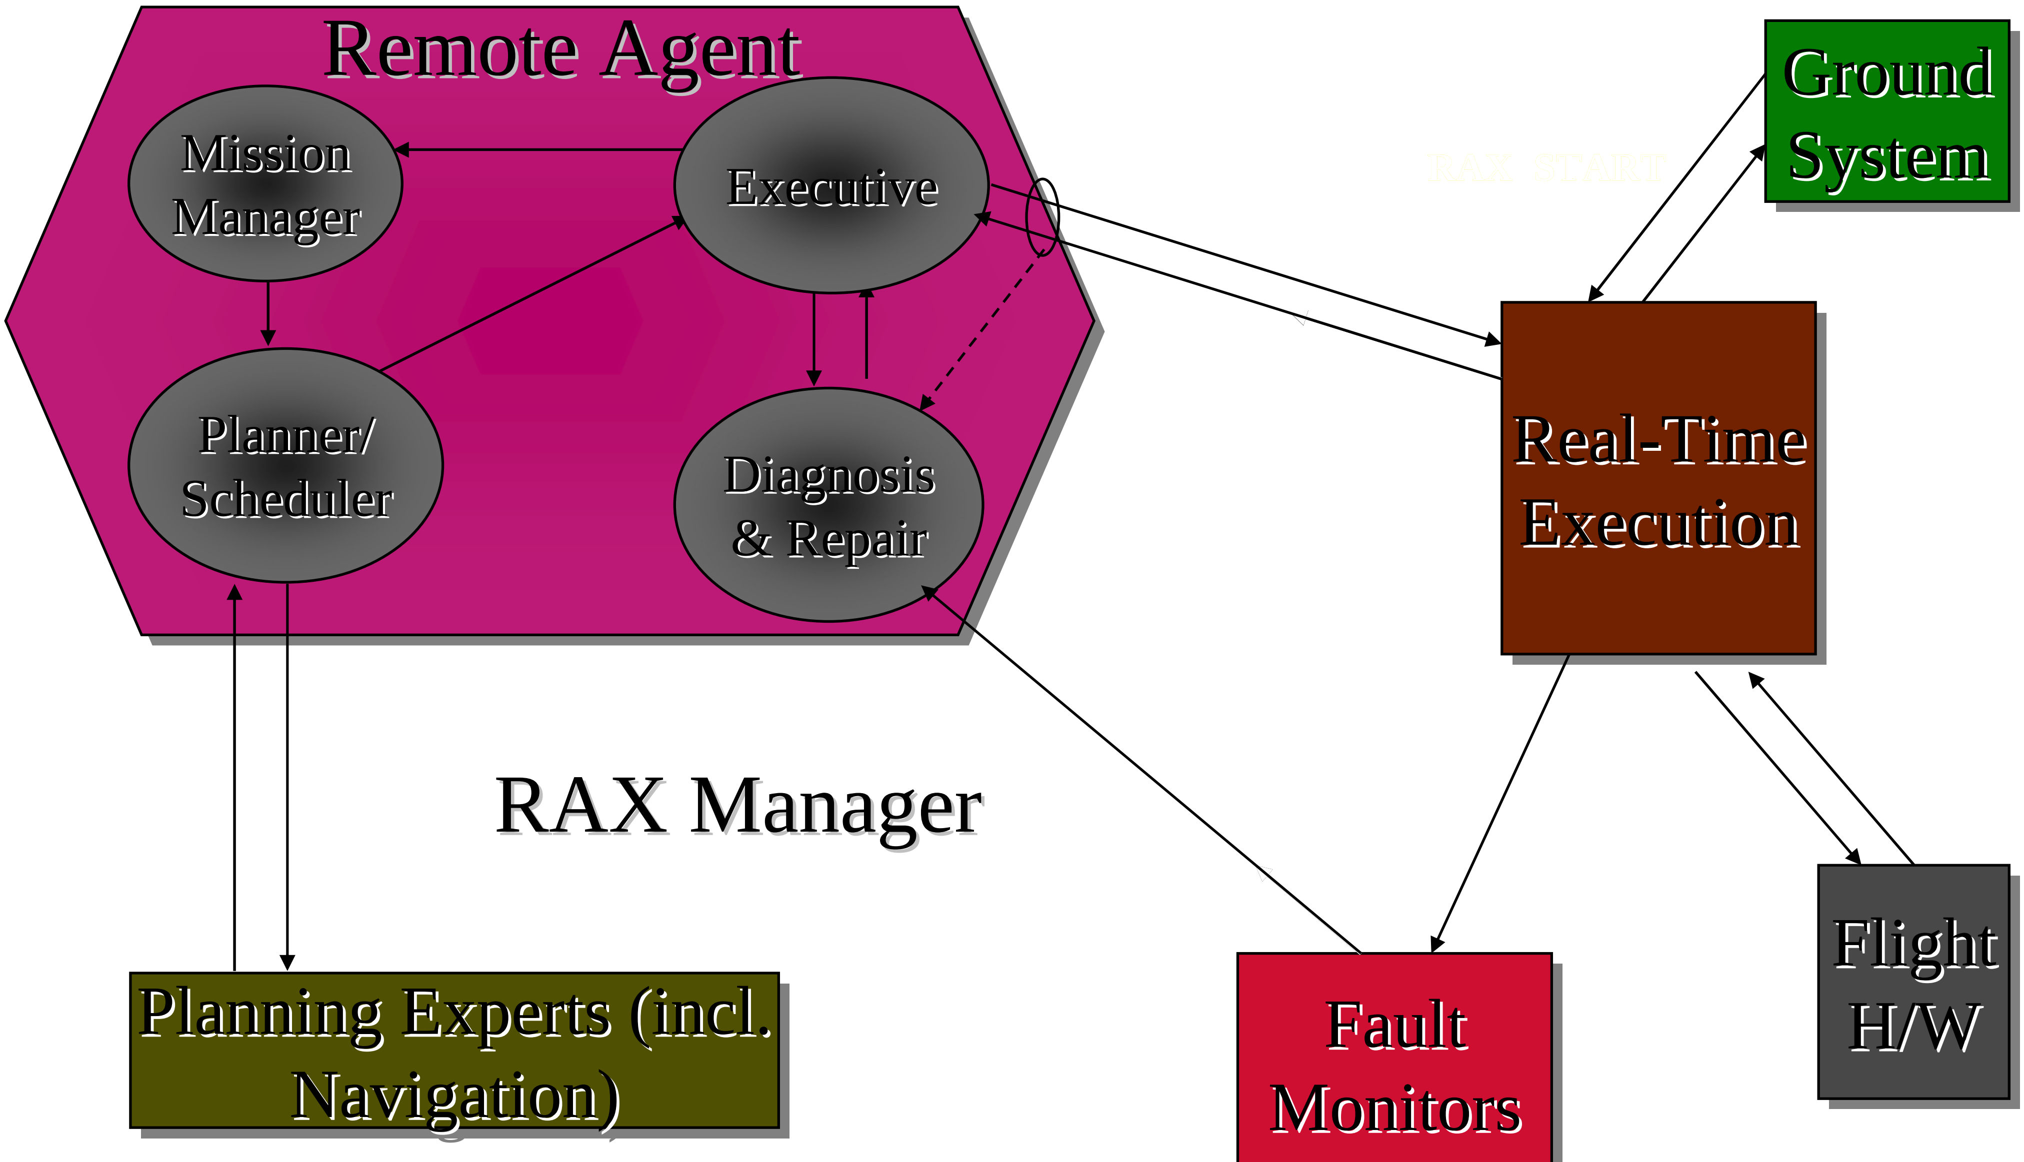

Slide courtesy of Brian Williams

Executive Requests Plan

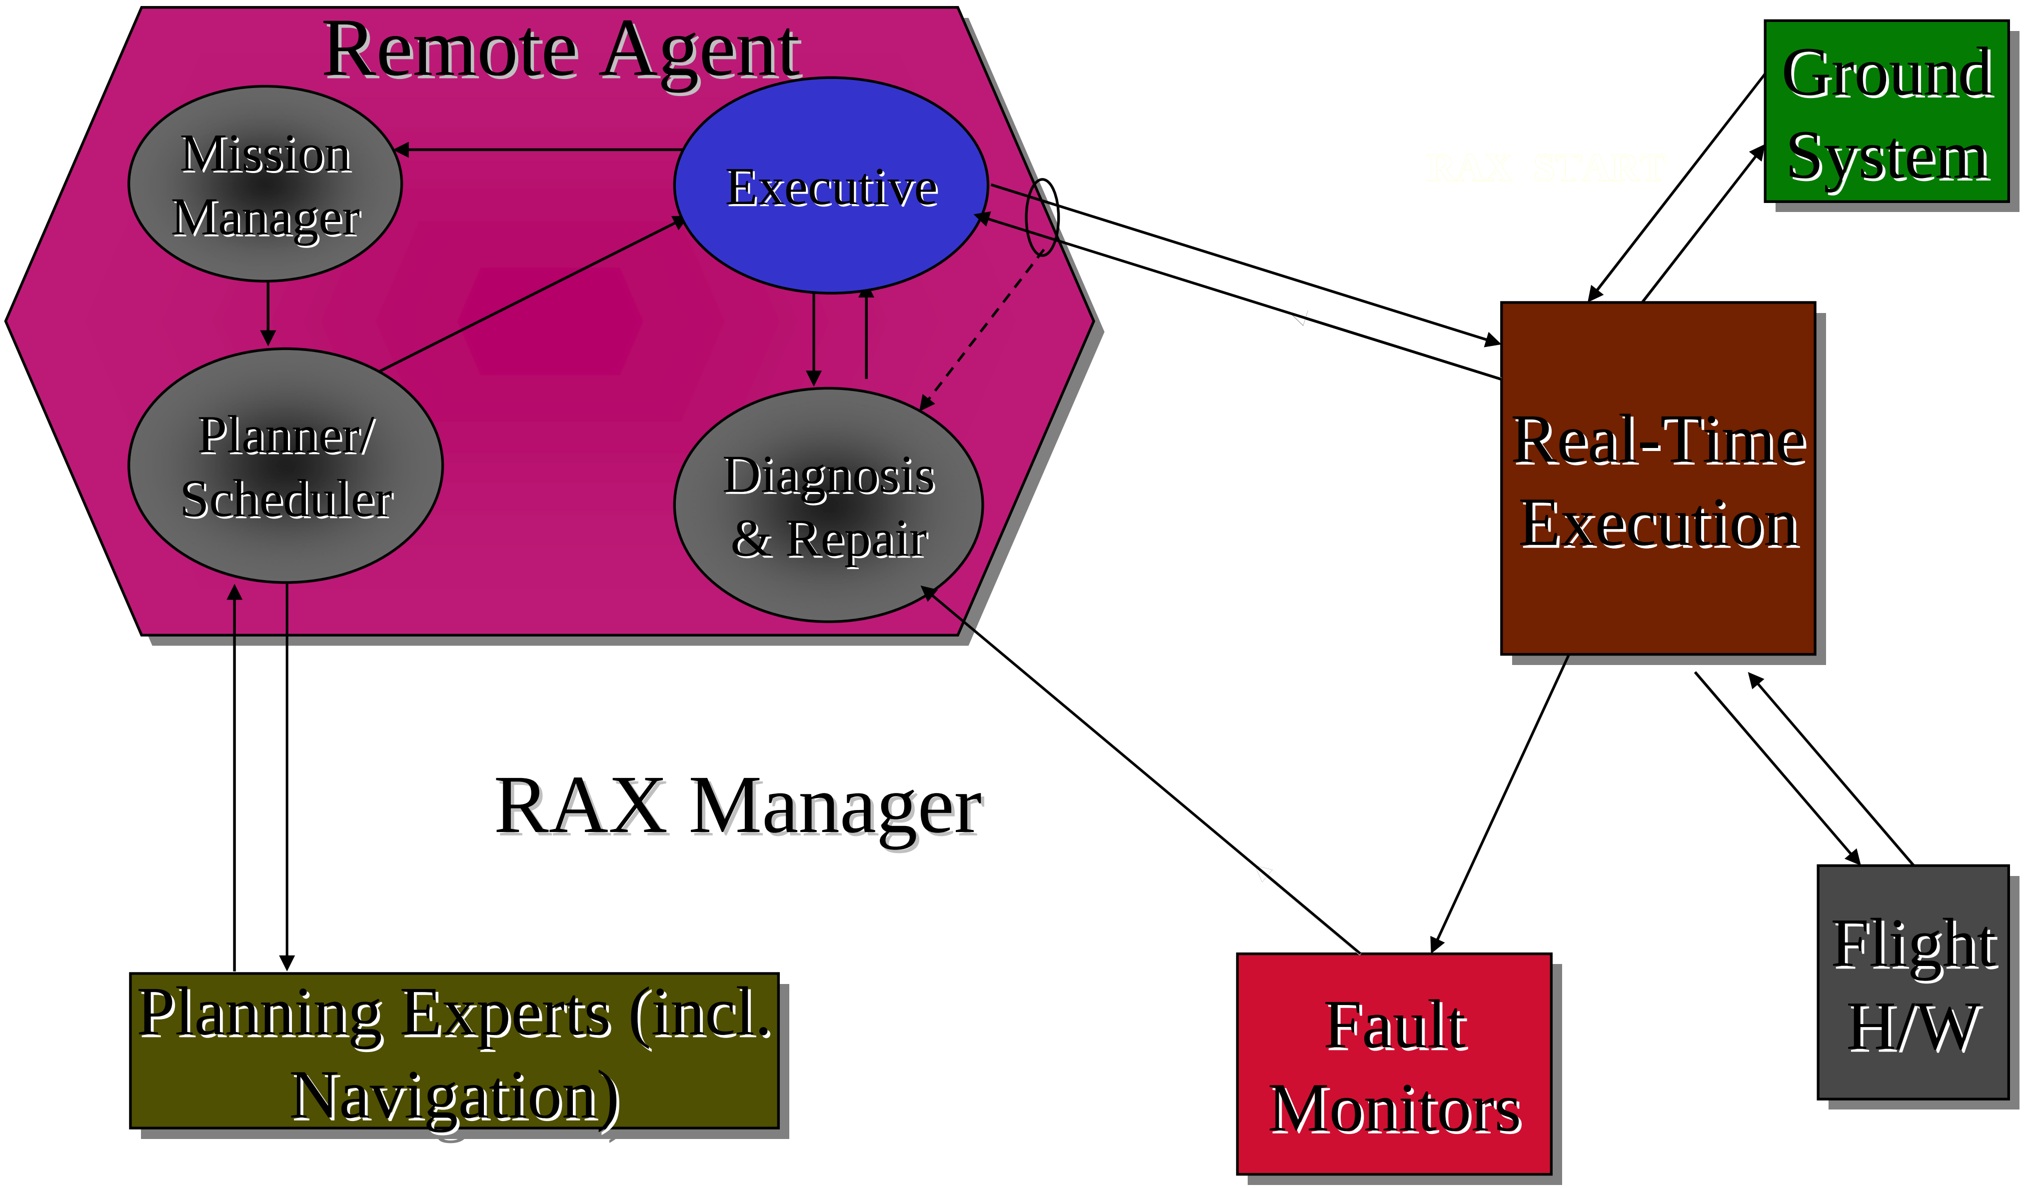

Slide courtesy of Brian Williams

Mission manager establishes goals, planner generates plan

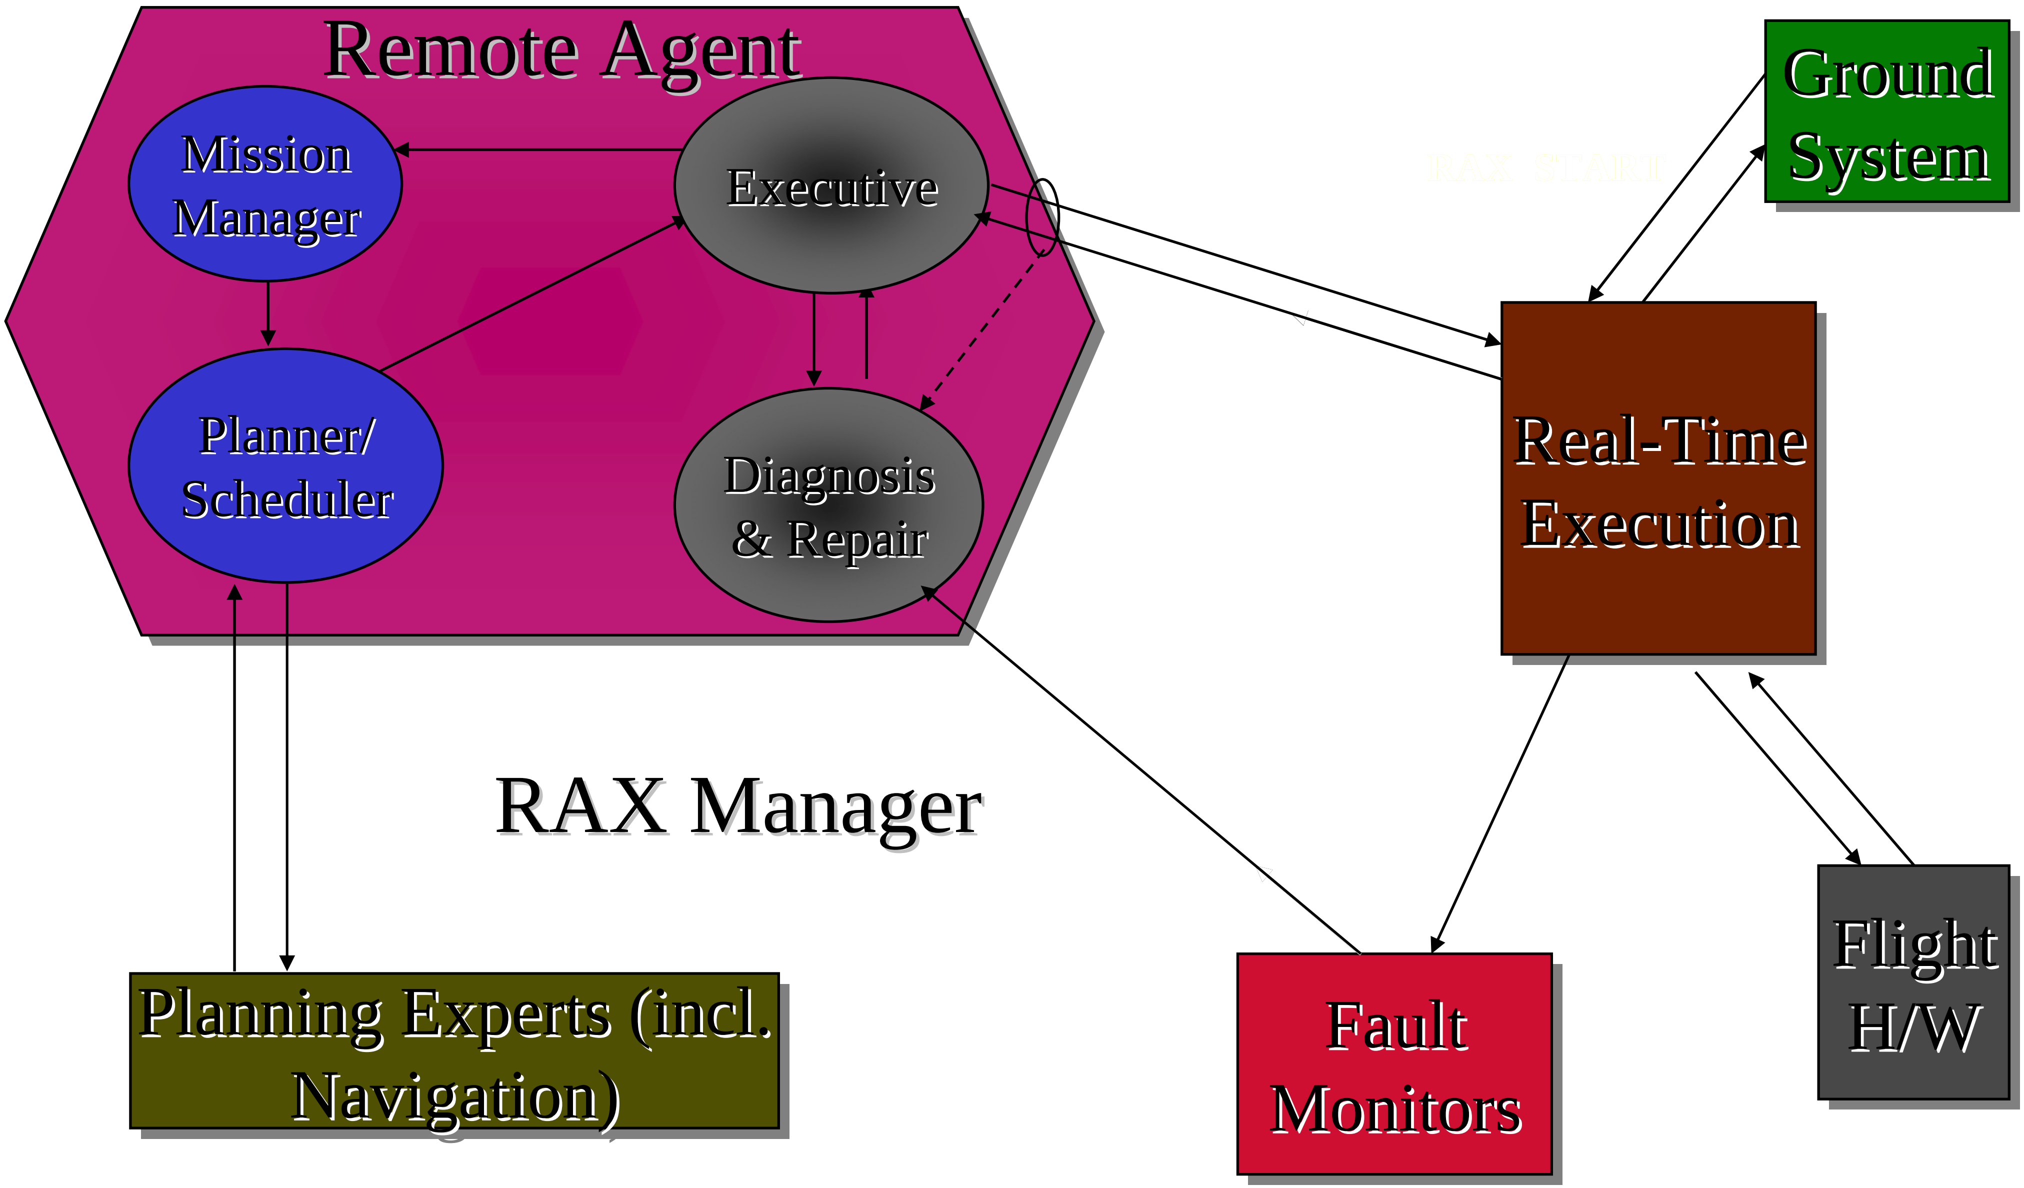

Slide courtesy of Brian Williams

Executive executes plan

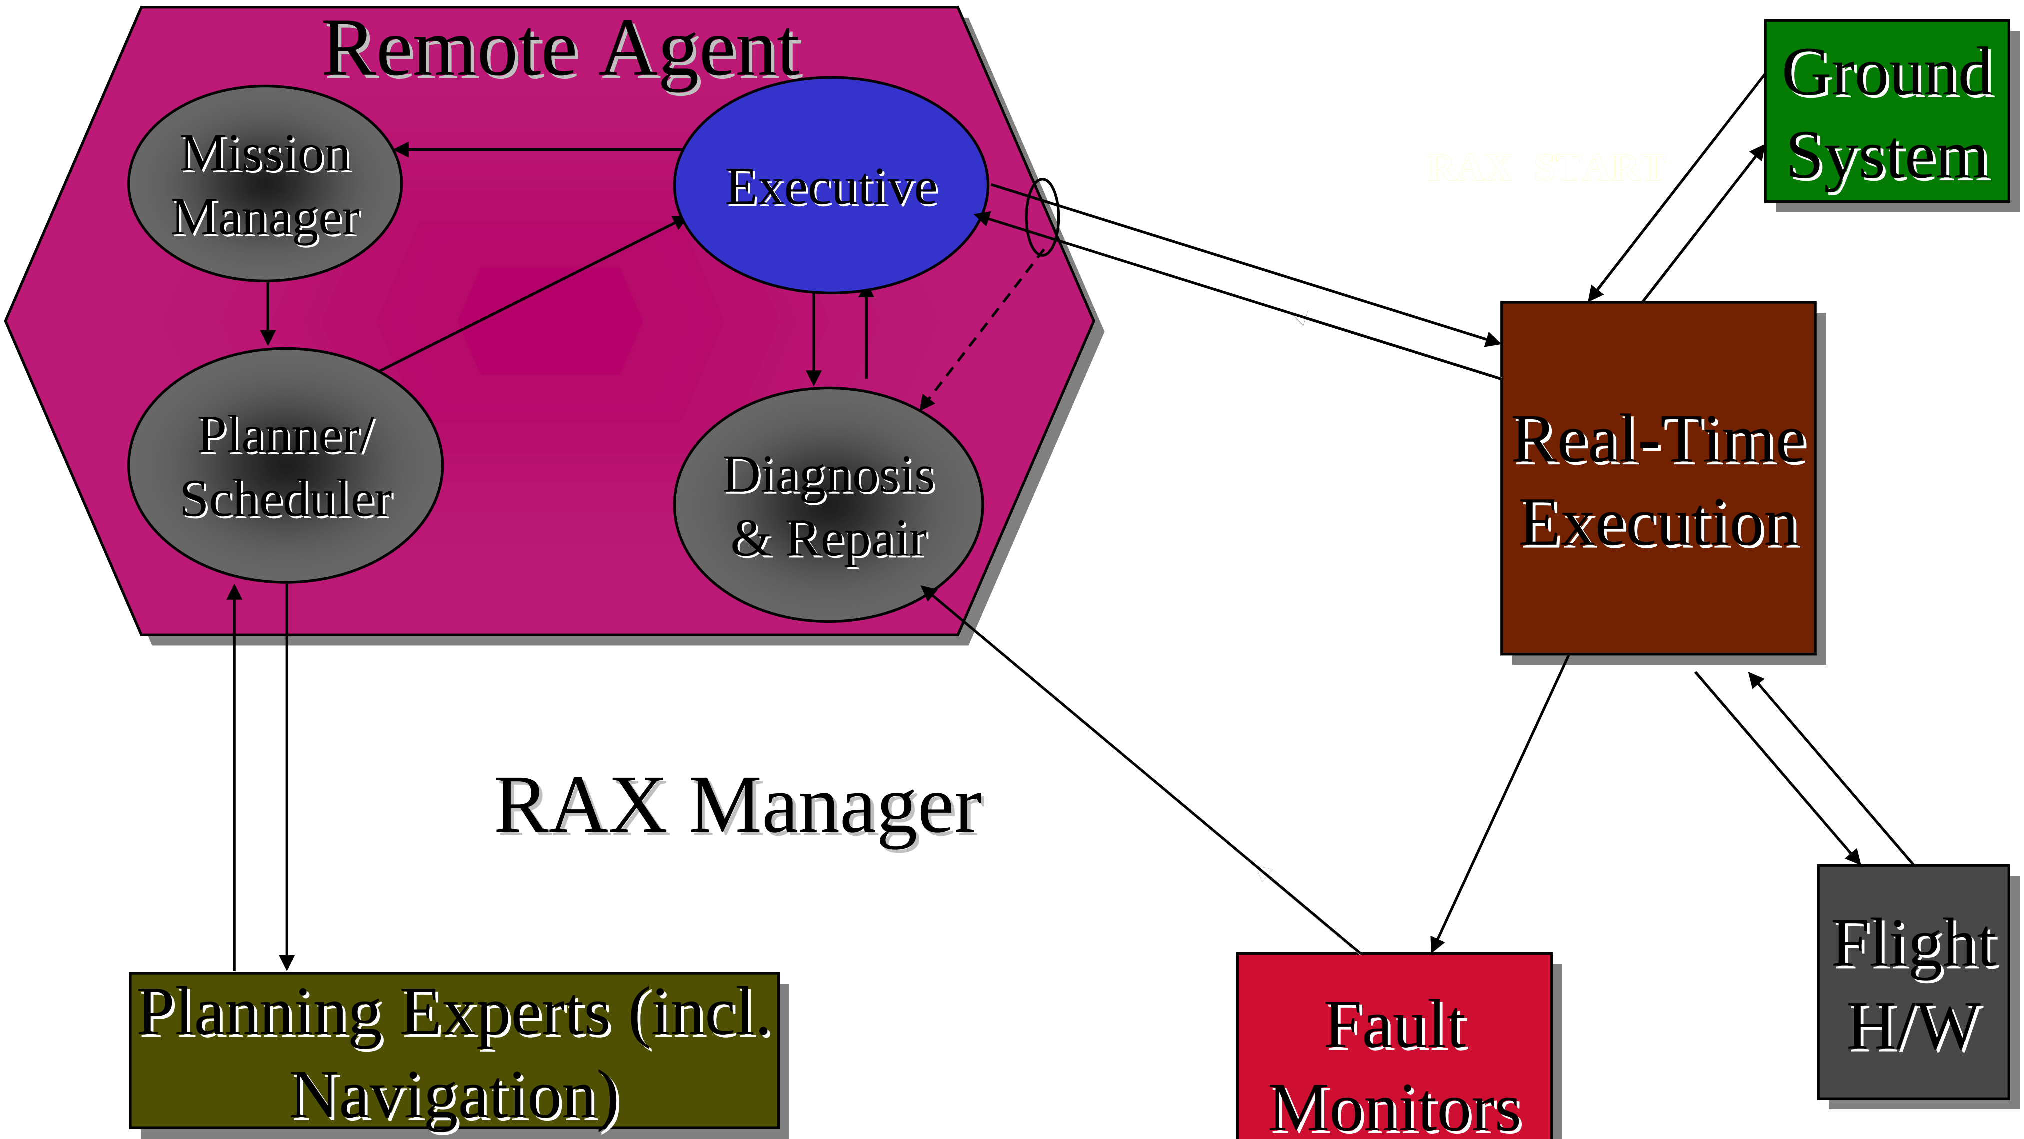

Slide courtesy of Brian Williams

Diagnosis System monitors and repairs

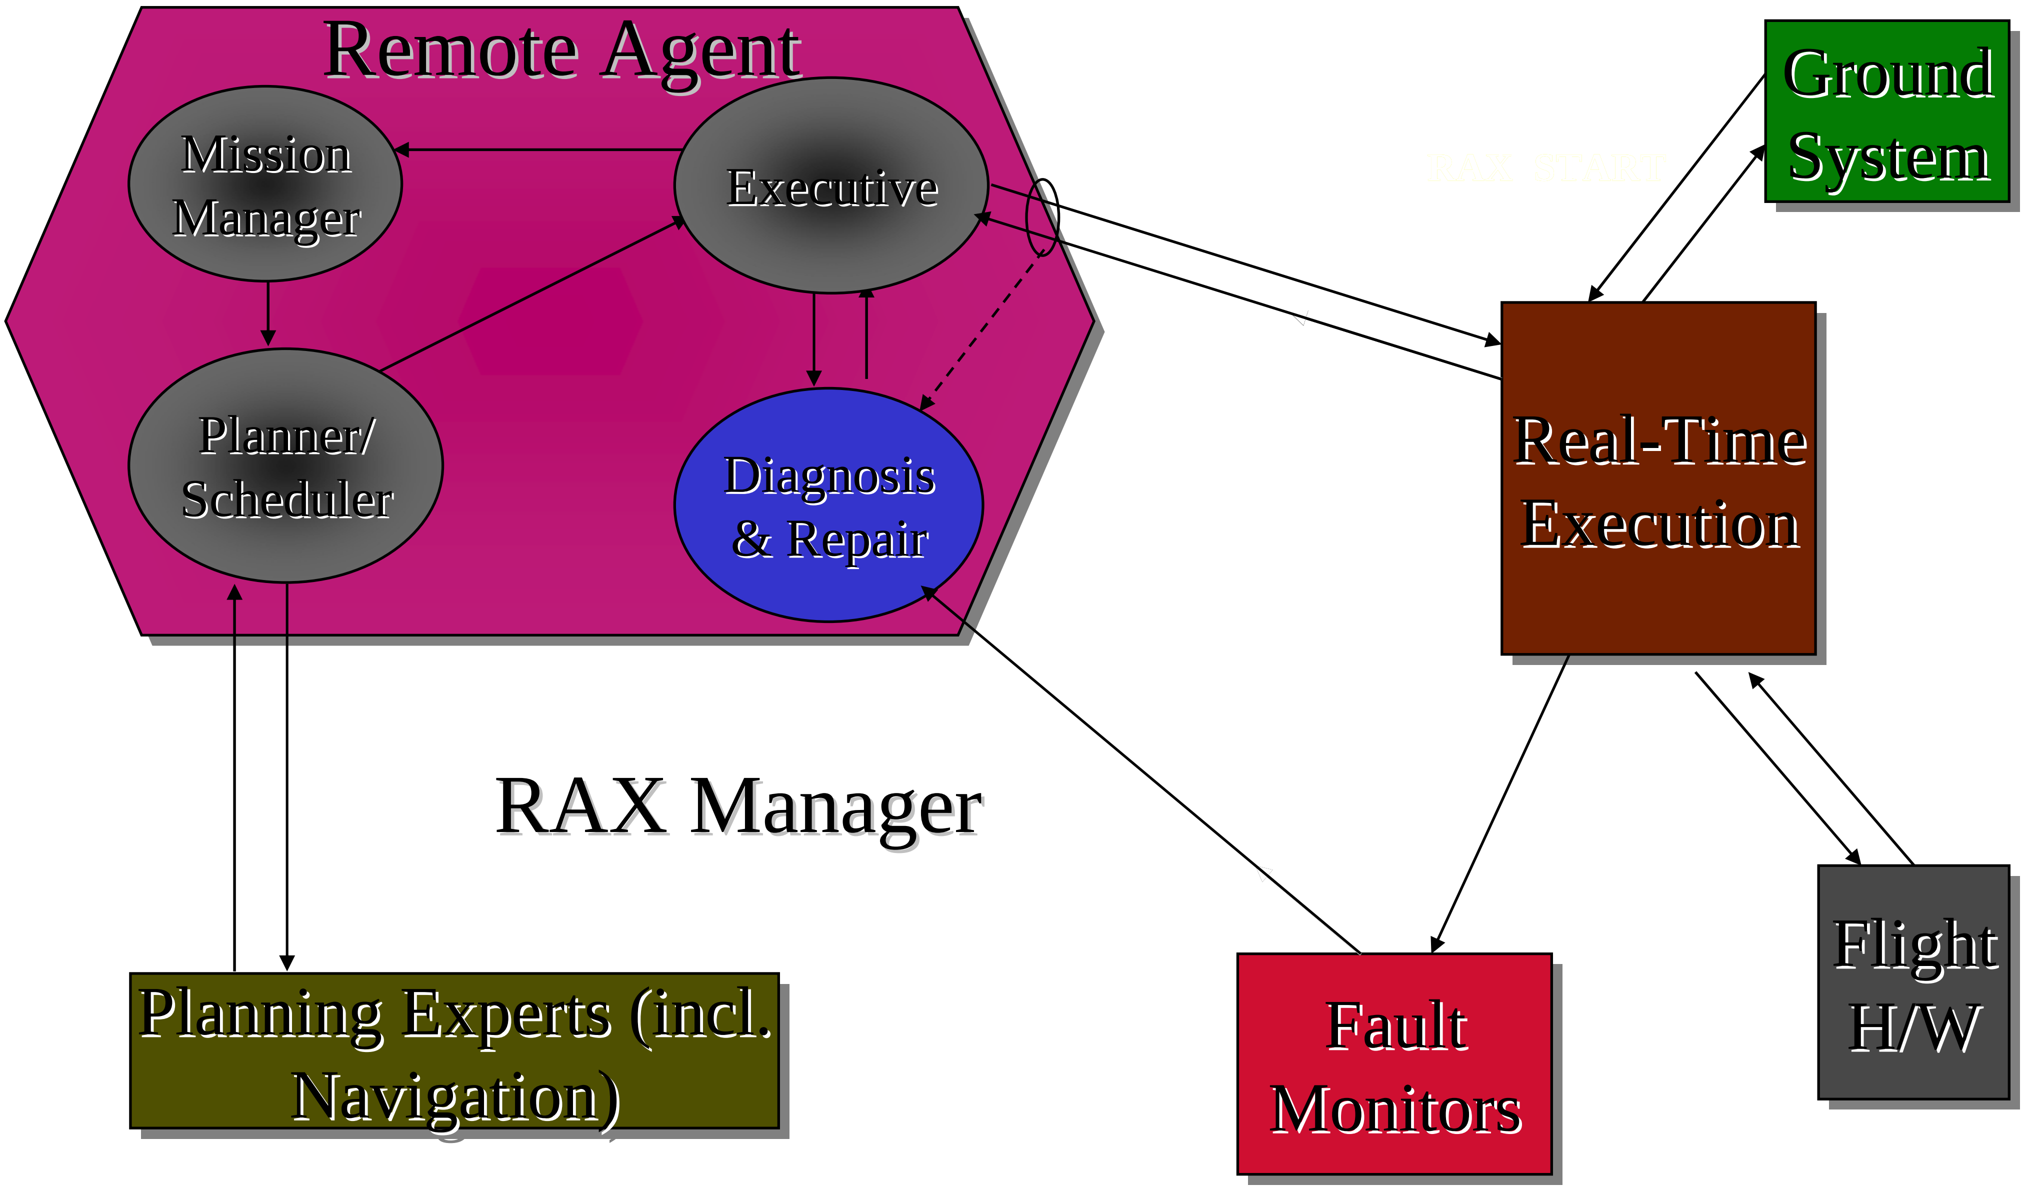

Slide courtesy of Brian Williams

Challenges of Autonomy in the Real World

Wide range of sensors

Noisy sensors

World dynamics

Adaptability

Incomplete information

Robustness under uncertainty

Possibly add more points under reliability

Mars Polar Lander Failure

Leading Diagnosis:

- Legs deployed during descent.
- Noise spike on leg sensors latched by software monitors.
- Laser altimeter registers 40m.
- Begins polling leg monitors to determine touch down.
- Latched noise spike read as touchdown.
- Engine shutdown at ~40m.

Programmers often make commonsense mistakes when reasoning about hidden state.

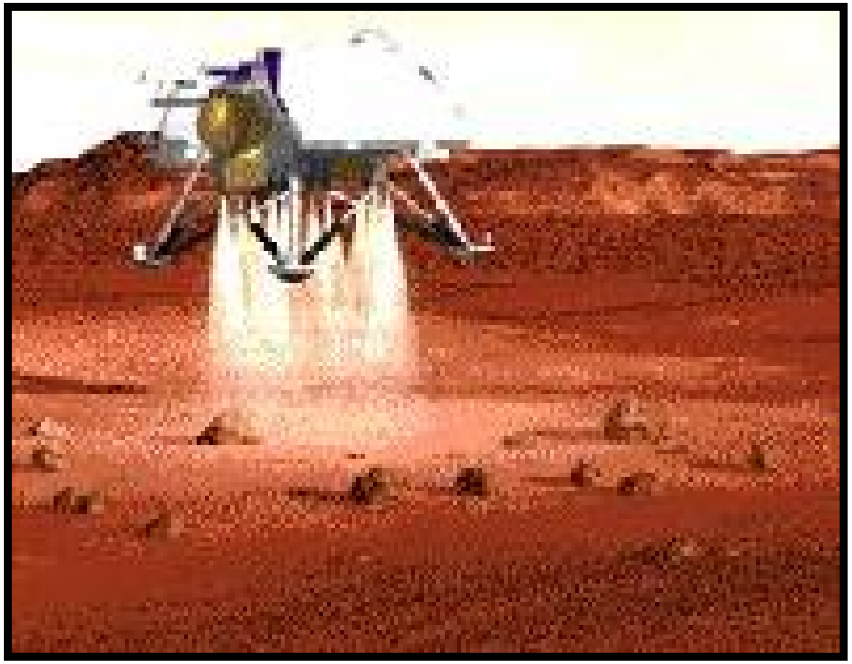

Reactive Model-based Programming Language (RMPL)

Support programmers with embedded languages that avoid these mistakes, by reasoning about hidden state automatically.

Slide courtesy of Brian Williams

Modelling Complex Behaviours through Probabilistic Constraint Automata

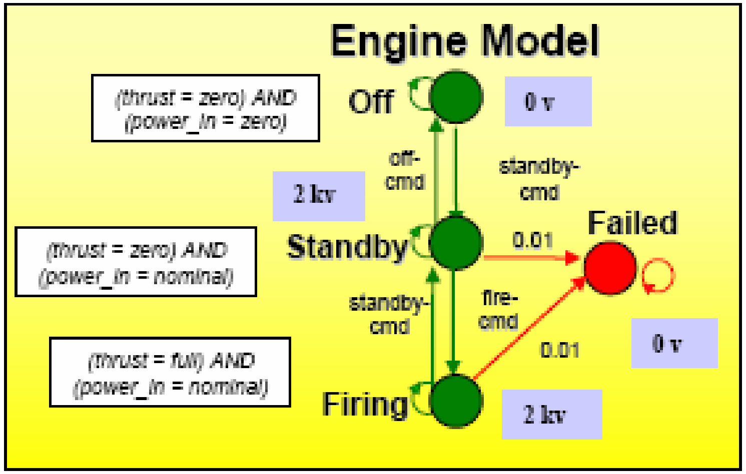

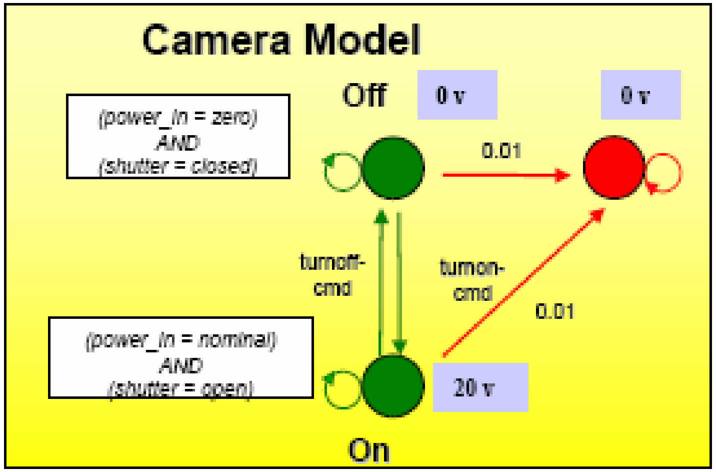

- Complex, discrete behaviours
  - modelled through concurrency, hierarchy and timed transactions
- Anomalies and uncertainty
  - modelled by probabilistic transitions
- Physical interactions
  - modelled by discrete and continuous constraints

Slide courtesy of Brian Williams

The Curse of Dimensionality

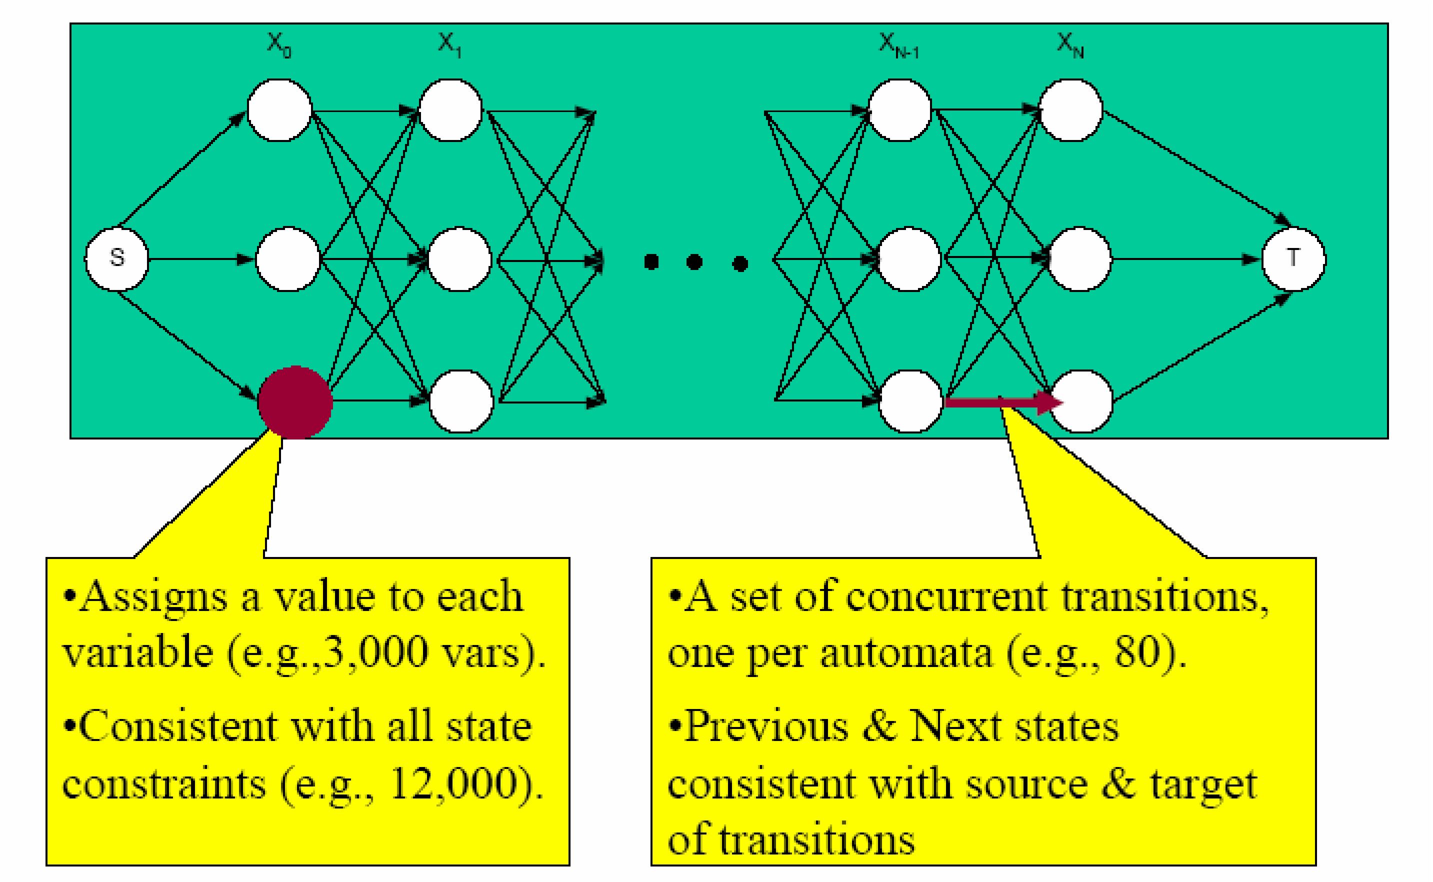

Slide courtesy of Brian Williams

Many problems aren't so hard

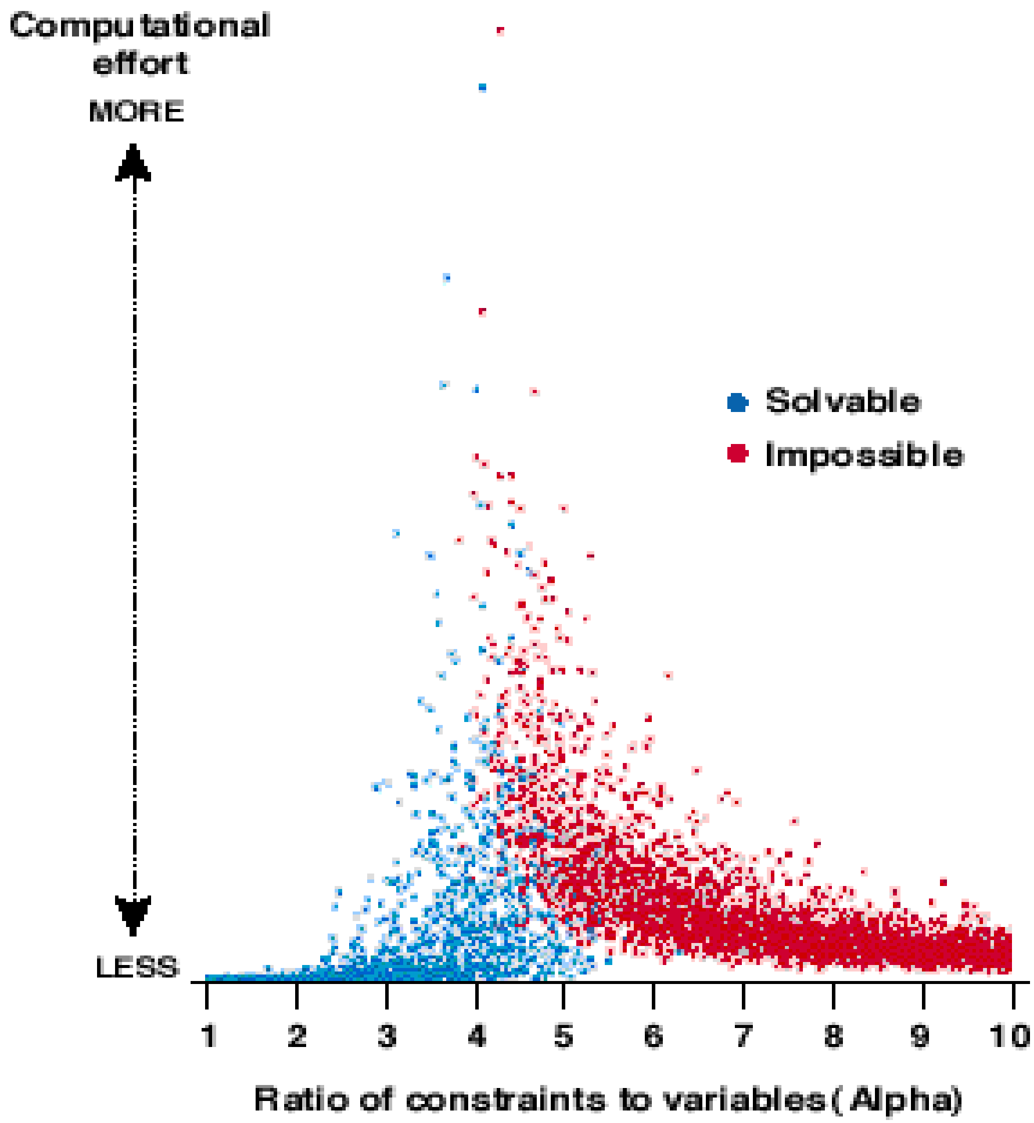

Slide courtesy of Brian Williams

Probabilistic State Estimation

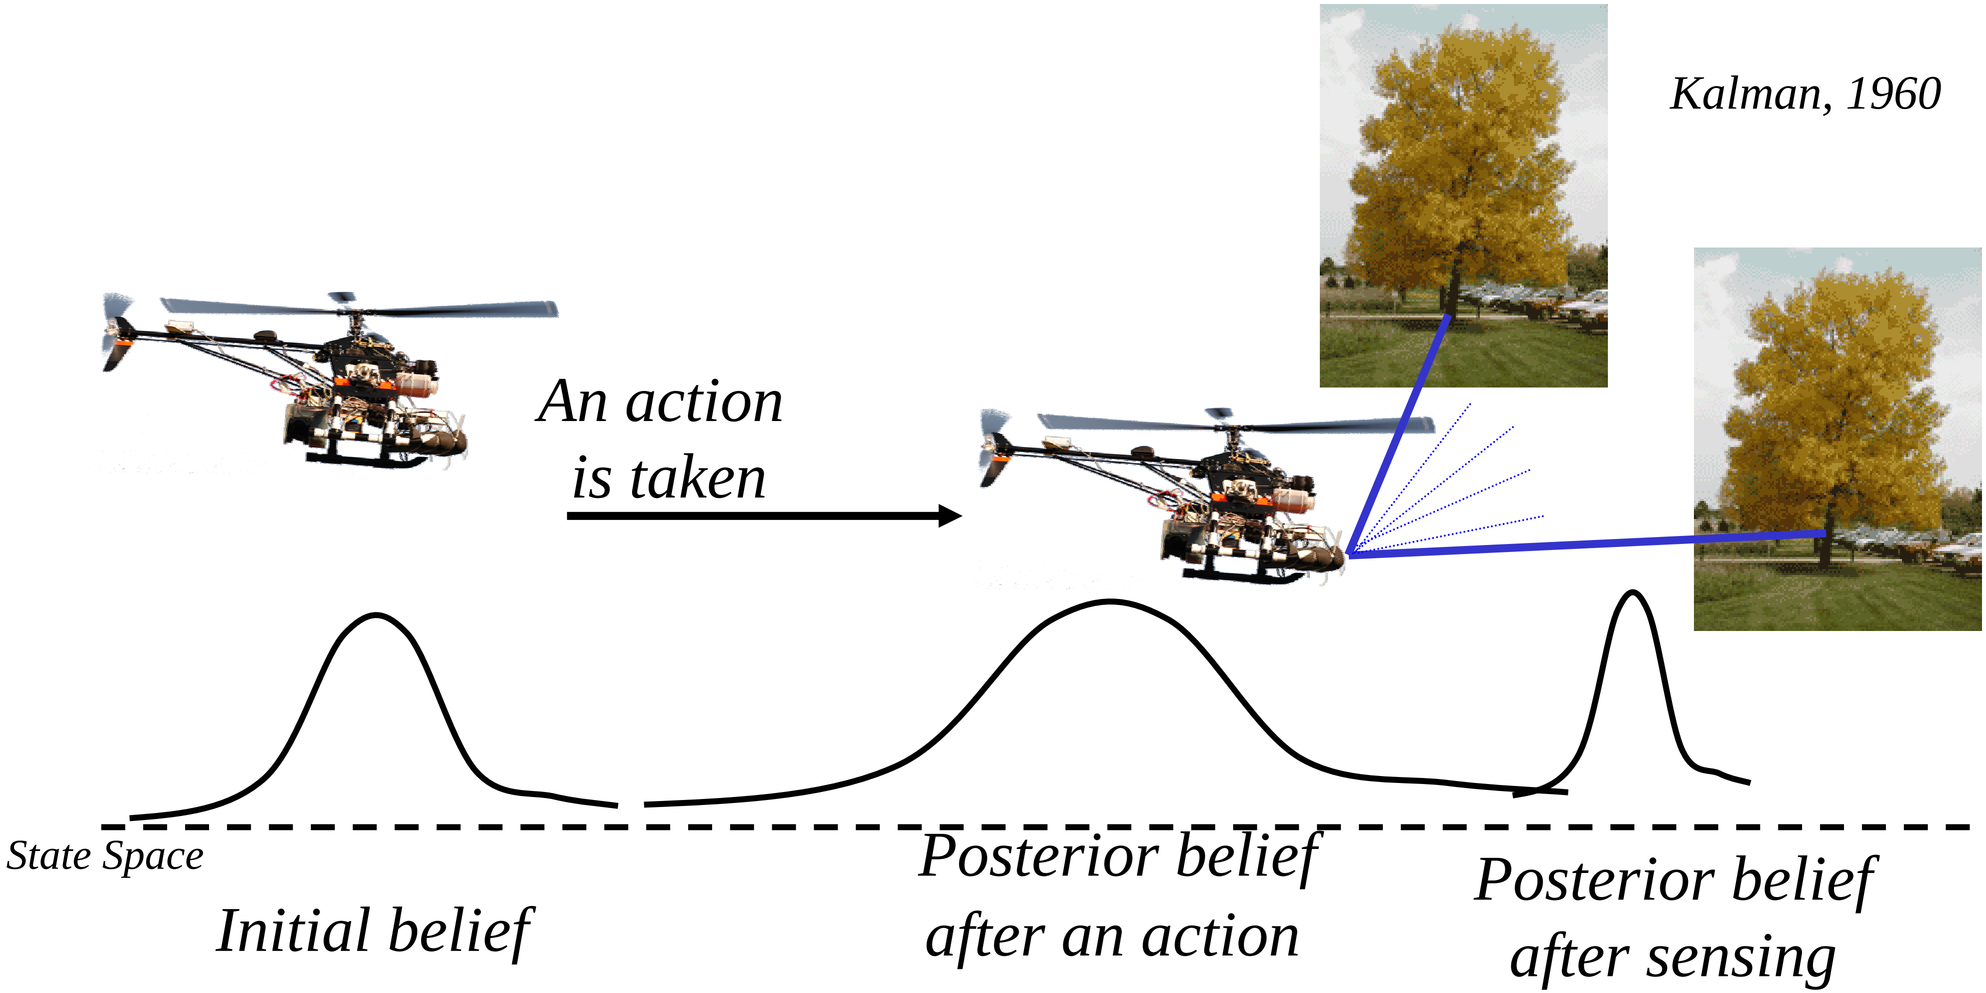

Monte Carlo Localization

Fox, Dellaert,Burgard &Thrun 2000

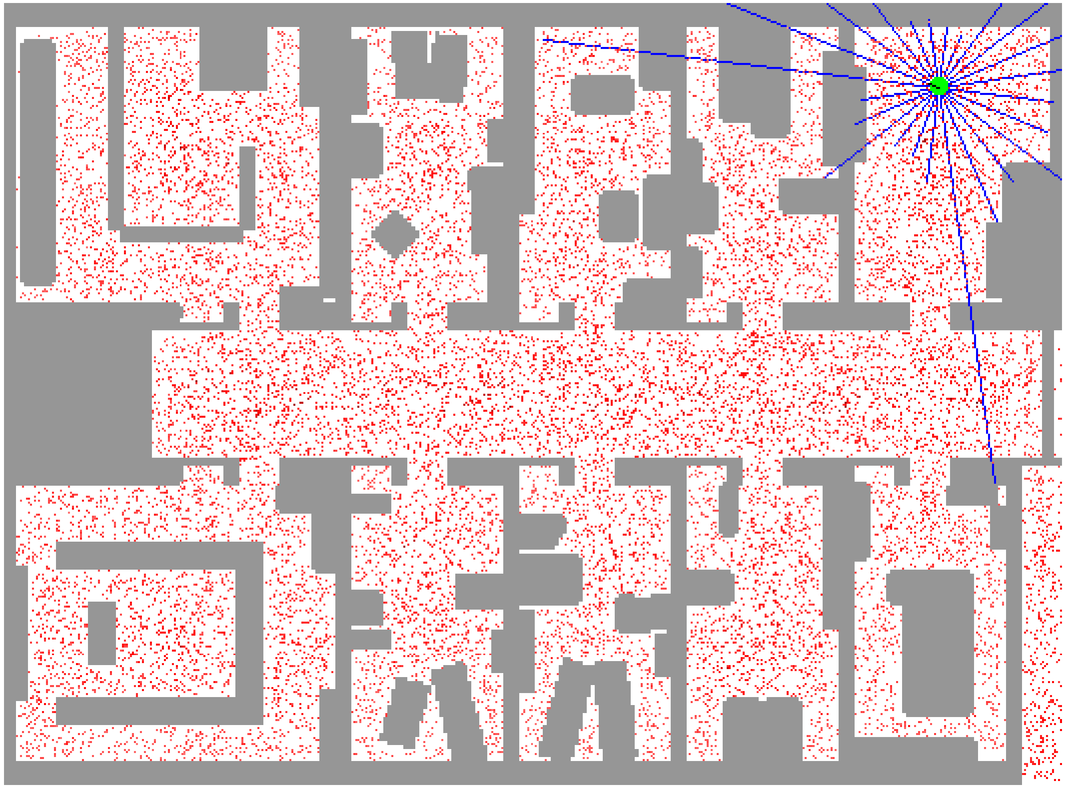

Image courtesy of Dieter Fox

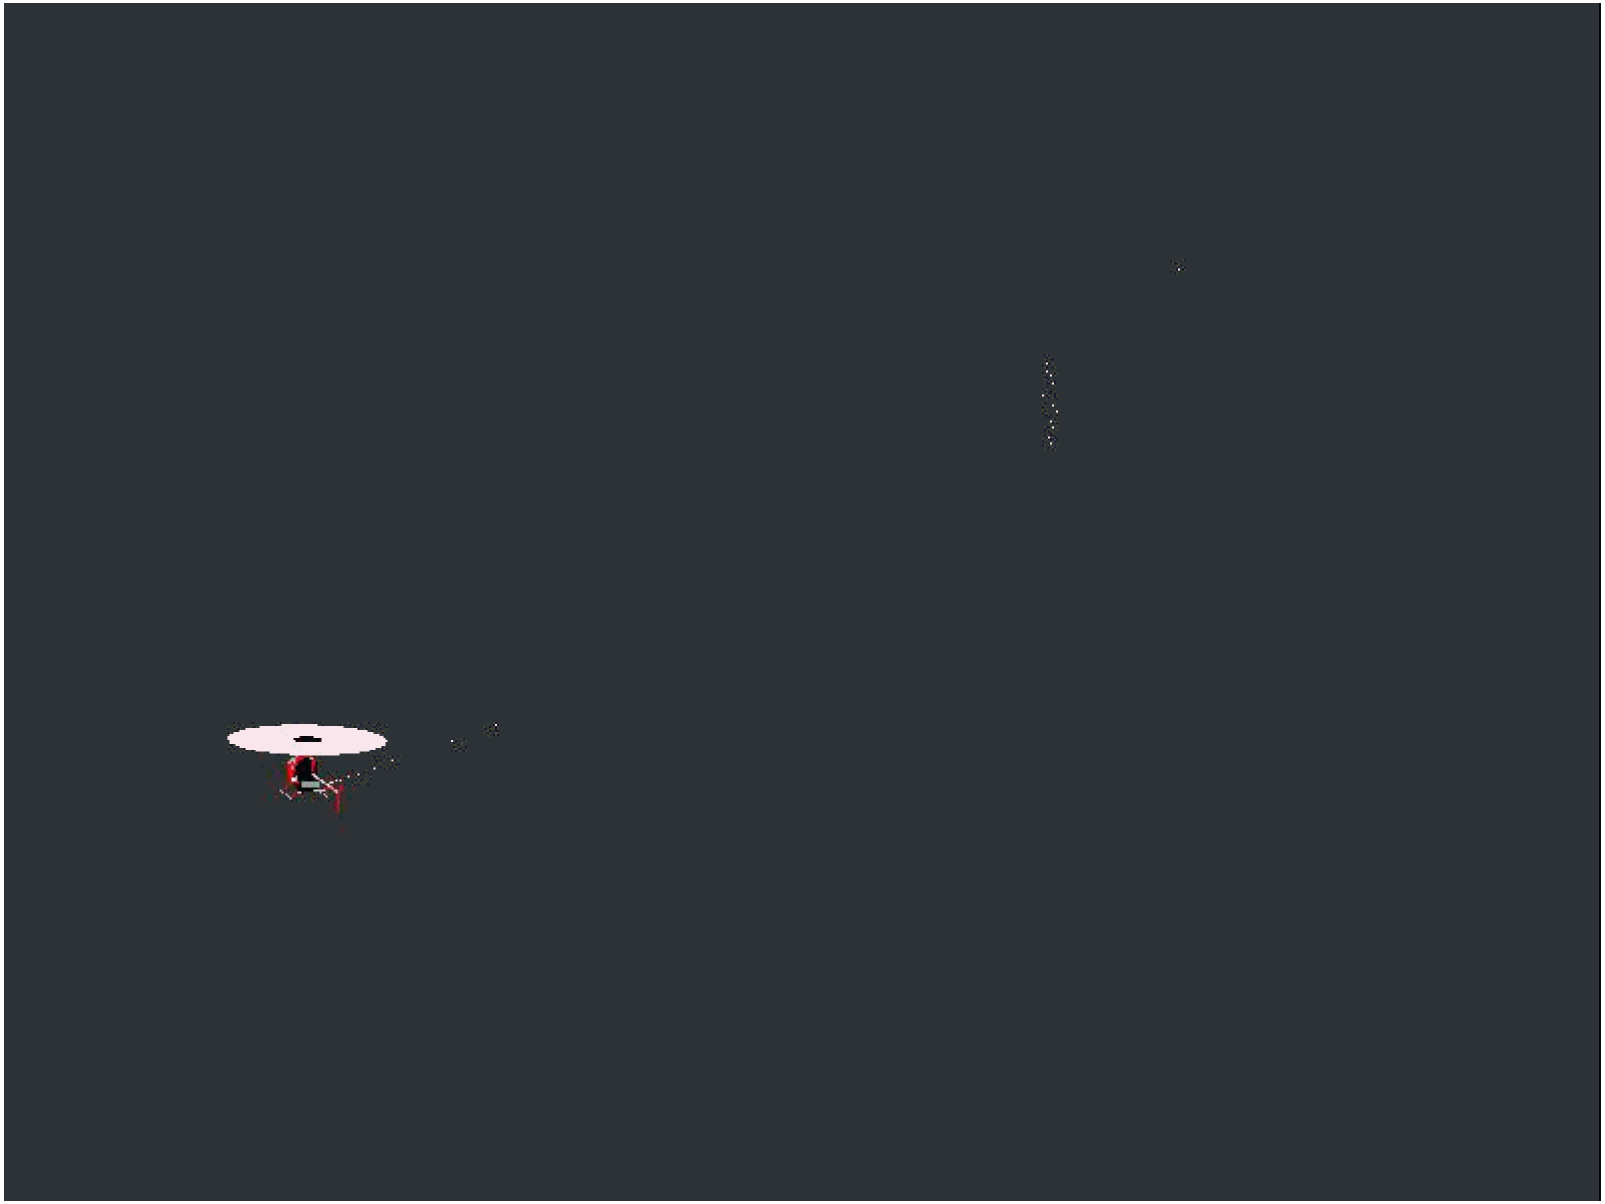

Movie courtesy of S. Thrun

Perception and Control

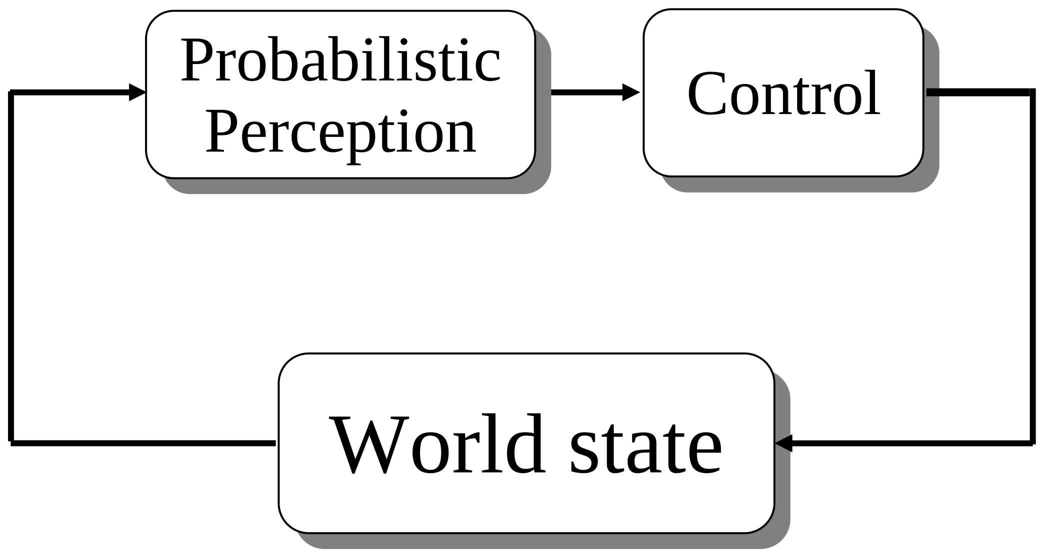

- But, these mobile robots haven't used the same kinds of probabilistic techniques for their control.
- Engineering diagram
- Prob perception -> argmax -> control
- SLAP & CPF not everybody's doing it.
- For the purposes of the talk, just going to concentrate on the control aspect.
- What usually happens is that the most likely state is assumed to be the correct state. In some cases, the system thresholds the uncertainty of the observation as a way to reject outliers, but it's really hard to actually integrate probabilities into the control loop.
- The question is, why does this matter?

Reliable Navigation

- Conventional trajectories may not be robust to localization error

Estimated robot position ●
True robot position ●
Goal position ●

Well, sometimes, good state estimation just isn't good enough.

Perception and Control

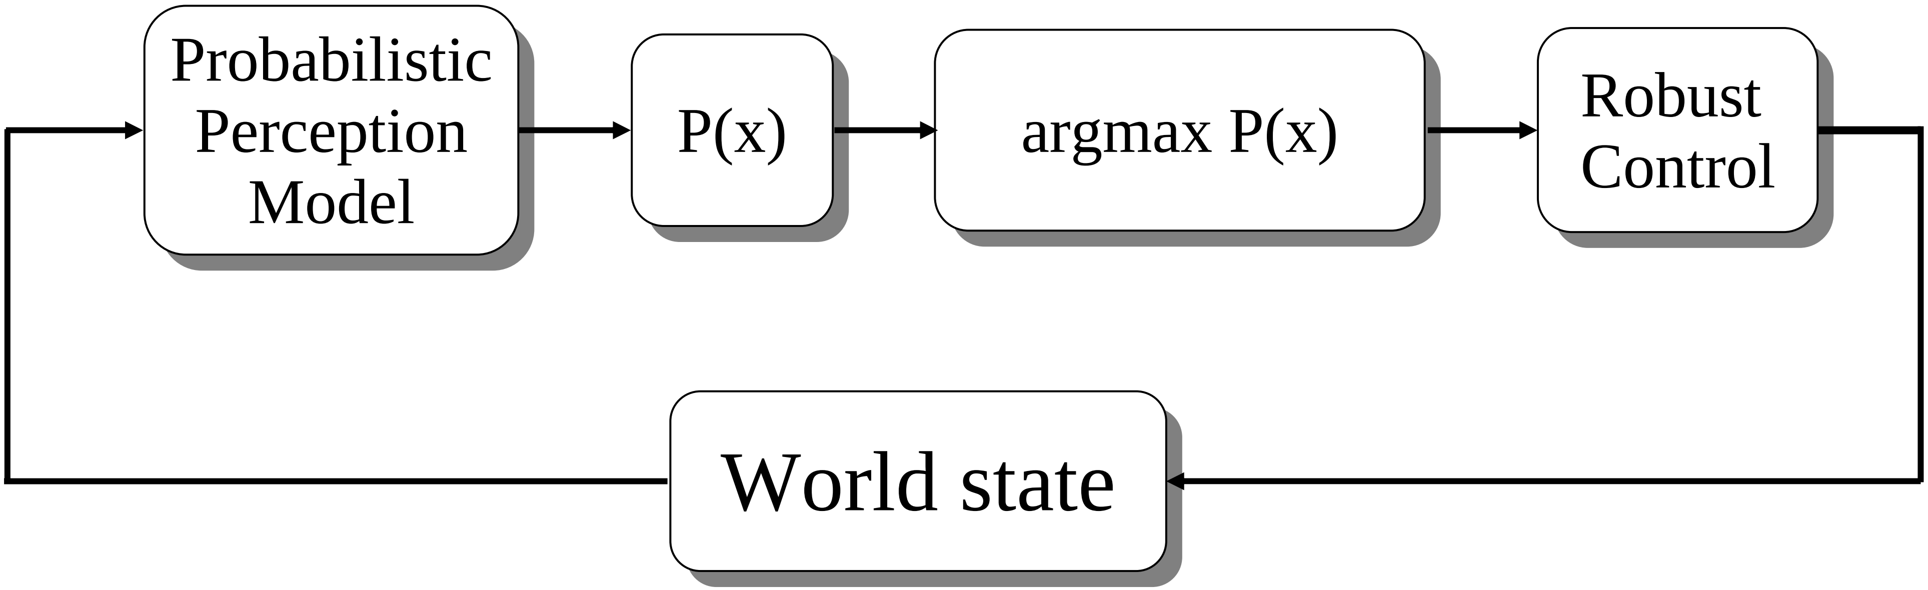

- But, these mobile robots haven't used the same kinds of probabilistic techniques for their control.
- Engineering diagram
- Prob perception -> argmax -> control
- SLAP & CPF not everybody's doing it.
- For the purposes of the talk, just going to concentrate on the control aspect.
- What usually happens is that the most likely state is assumed to be the correct state. In some cases, the system thresholds the uncertainty of the observation as a way to reject outliers, but it's really hard to actually integrate probabilities into the control loop.
- The question is, why does this matter?

Coastal Navigation

POMDPs

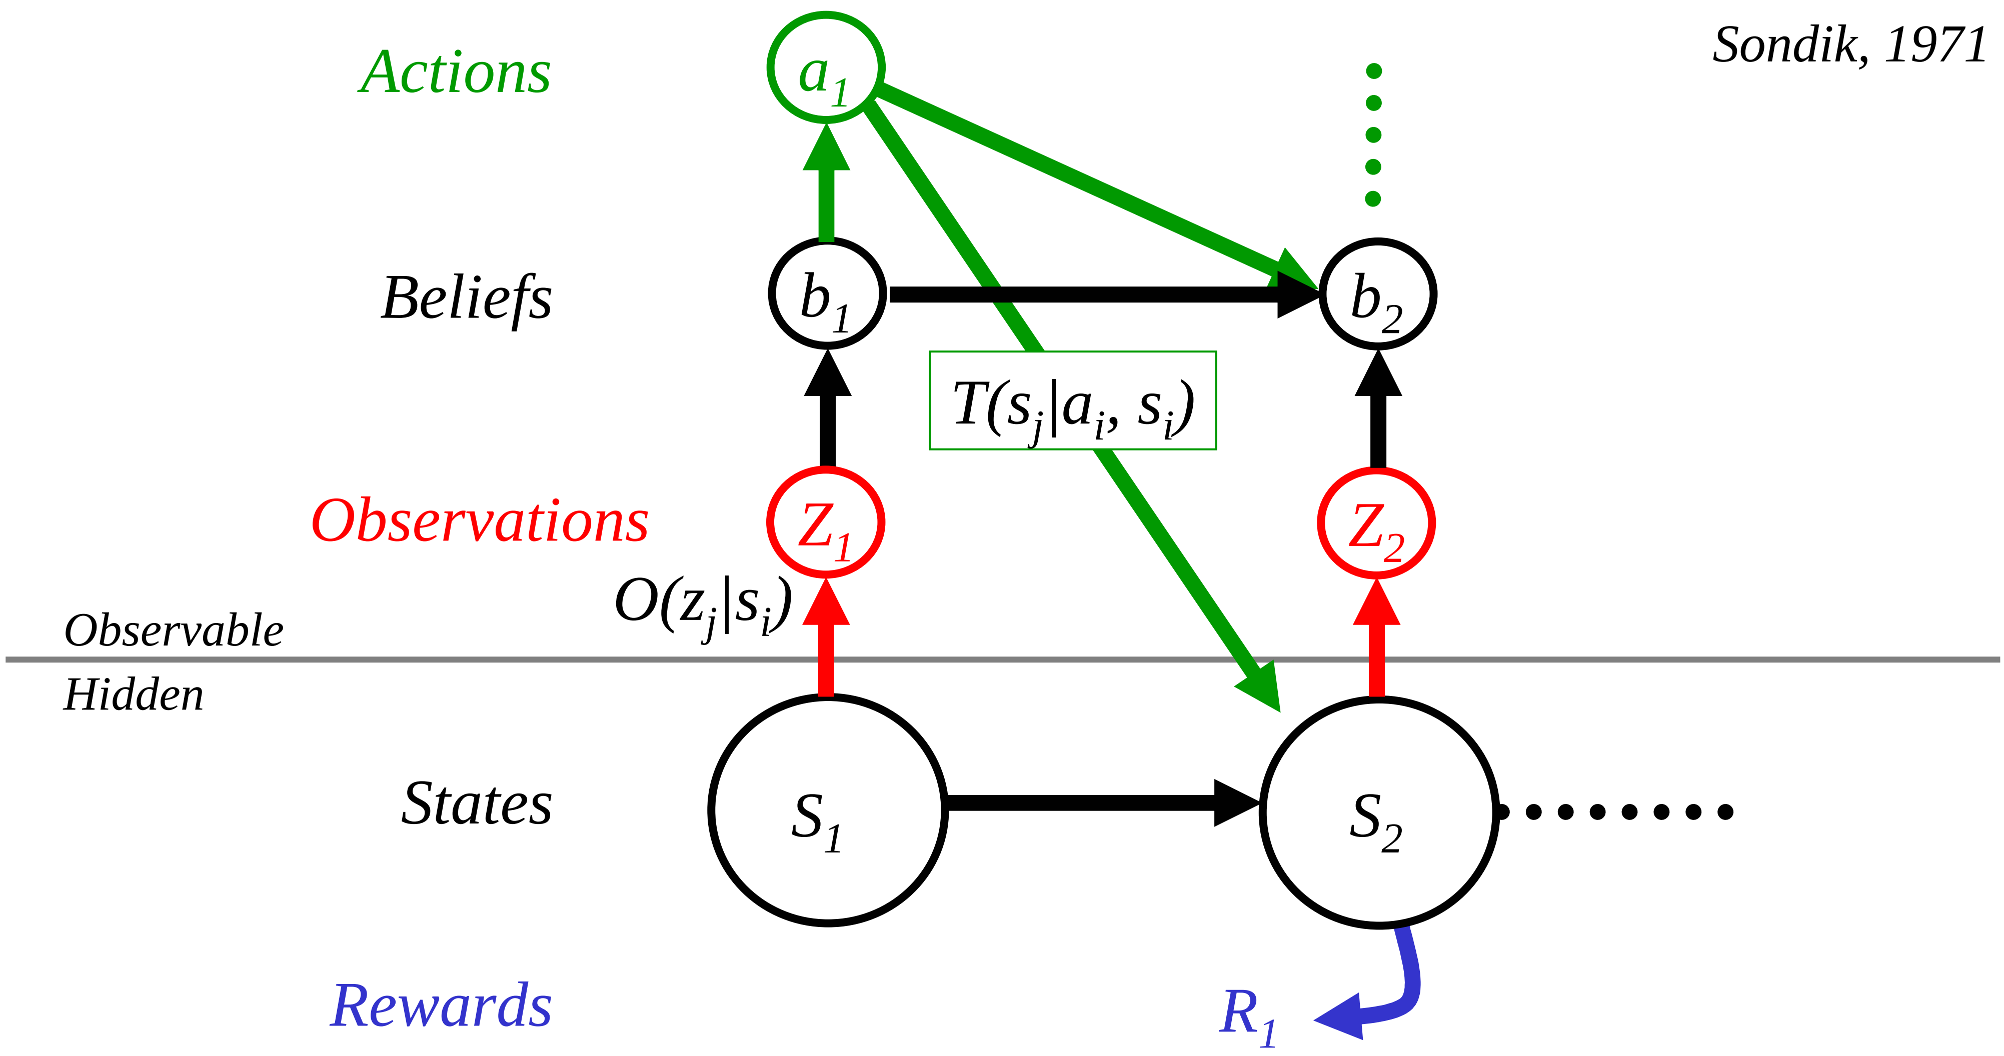

Sondik, 1971
Actions
a_1
Beliefs
b_1
b_2
$T(s_j|a_i, s_i)$
Observations
Z_1
Z_2
$O(z_j|s_i)$
Observable
Hidden
States
S_1
S_2
Rewards
R_1

Navigation as a POMDP

Controller chooses actions based on probability distributions

Kalman, 1960

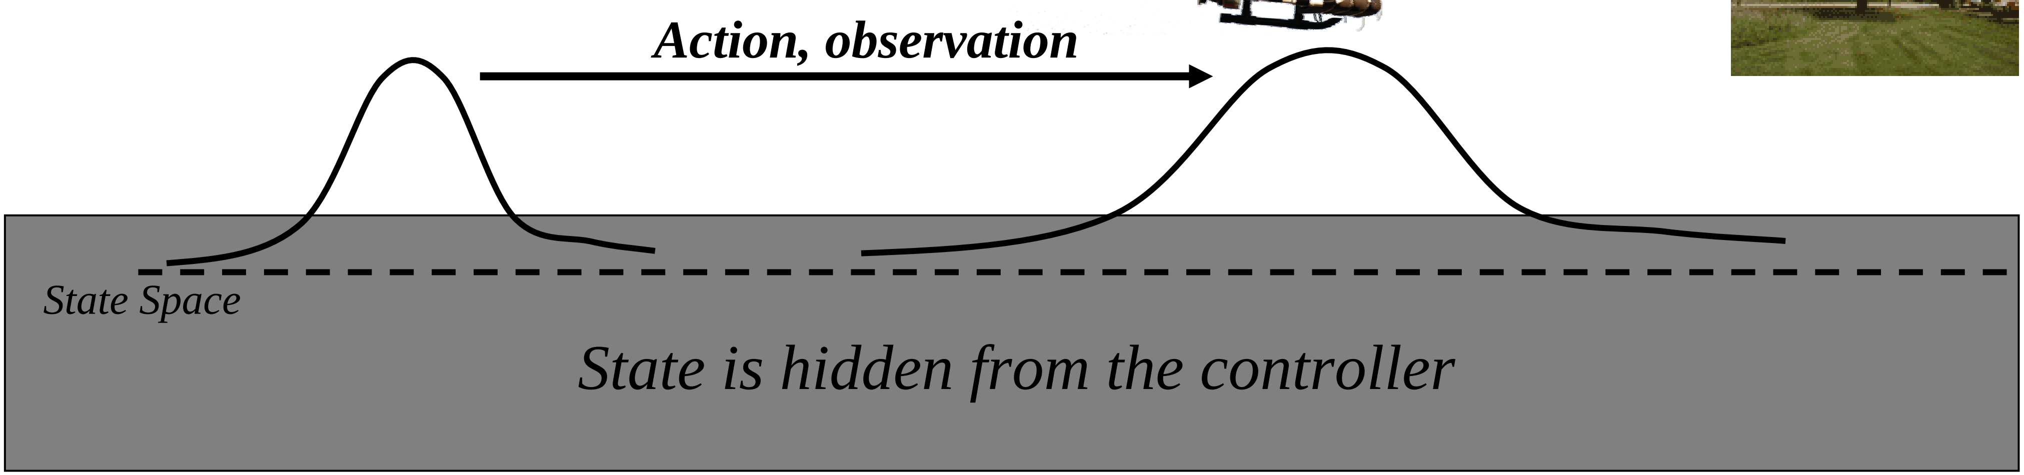

POMDP Advantages

- Models information gathering
- Computes trade-off between:
    - Getting reward
    - Being uncertain

Nursebot Pearl

Assisting Nursing Home Residents

Longwood, Oakdale, May 2001
CMU/Pitt/Mich Nursebot Project

Predicted Health Care Needs

- By 2008, need 450,000 additional nurses:
  - Monitoring and walking assistance
    30 % of adults 65 years and older have fallen this year

Cost of preventable falls:
$32 Billion US/year

Alexander 2001

- Intelligent reminding

Cost of medication non-compliance:
$1 Billion US/year

Dunbar-Jacobs 2000

Canadian numbers assumed to be similar.

Dialogue Management using POMDPs

- Unobserved state space is user's desired task
- Observations are utterances reported from speech recognition
- Actions are: robot motion, speech acts
- Reward: maximised for satisfying user task

Human-Robot Interaction

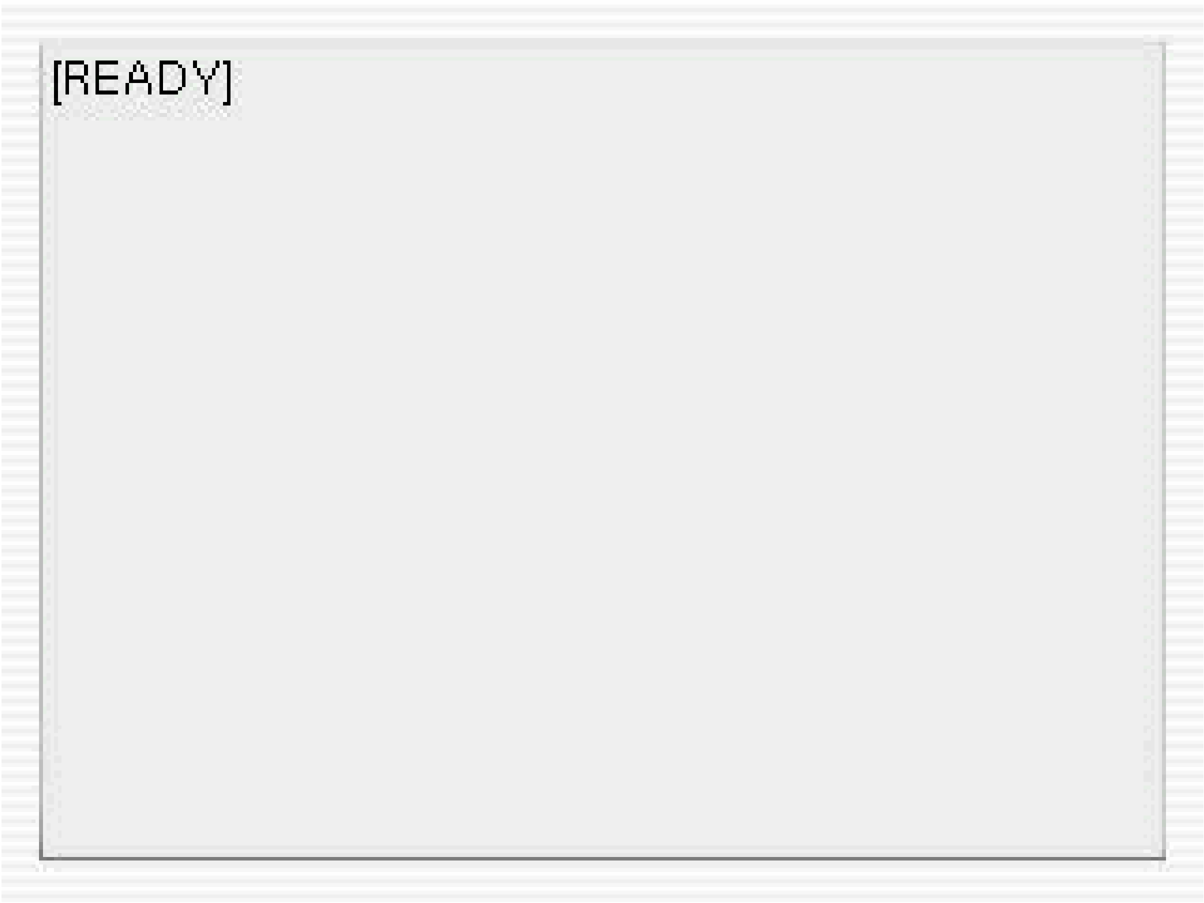

- This movie is a transcript of a speaker-independent voice recognition system, CMU Sphinx. You can see (and hear) that many utterances are just noise. But, with an awareness of the noise level of some utterances, a good dialogue manager might be able to handle this kind of input without too many errors.

- So the question is, how can a robot integrate probabilistic reasoning into both its perception and control?

POMDP Dialogue Manager

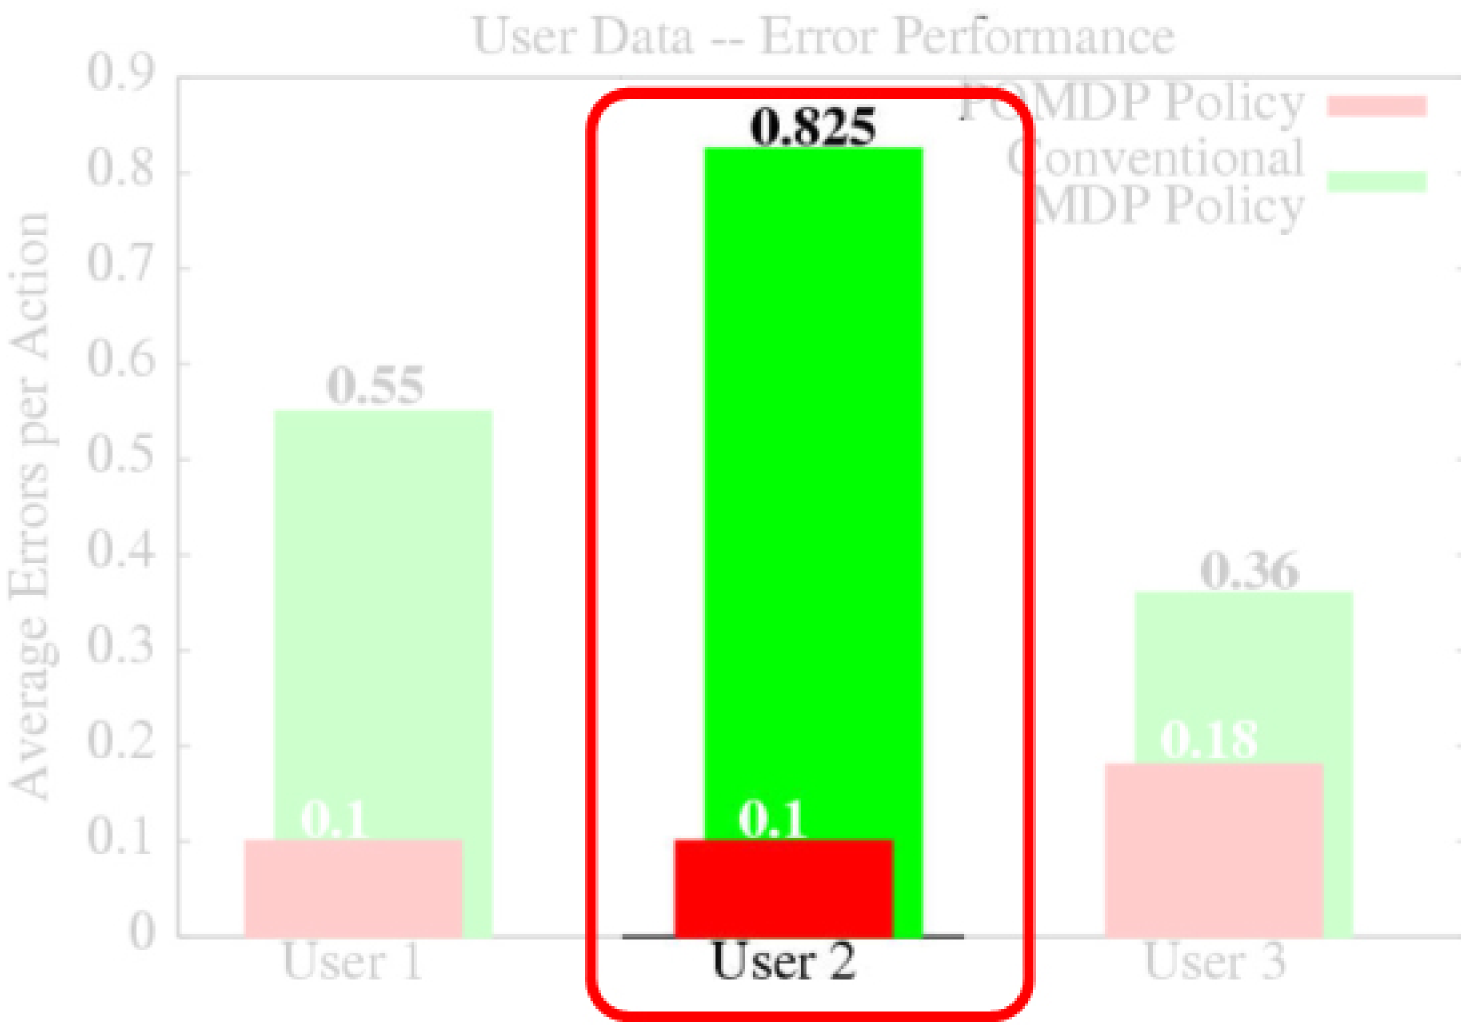

This is a 600 state POMDP. The only reason this POMDP is solvable is by using our techniques. Conventional techniques

Mixed-Initiative Planning

- Brings to table:
  - mechanisms for human involvement in plan generation
  - language for explaining choices to human
  - look-ahead search of options and consequences
- Lacks
  - execution of plans
- Citations
  - Ferguson, et al 1996
  - Burstein and McDermott, 1996
  - Pollack and Horty, 1999
  - Myers, 1996

Slide courtesy of Dave Kortenkamp

Adjustable Autonomy

- Brings to table:
  - execution of some plans
  - automatic hand-off to humans
- Lacks
  - full spectrum of control
  - verification
  - understanding
- Citations
  - Barber, et al 2000
  - Bonasso, et al 1997
  - Dorais, et al 1998
  - Kortenkamp, et al 2000
  - Musliner and Krebsbach, 1999
  - Thurman, et al 1997

Slide courtesy of Dave Kortenkamp

What you should know

- Why use autonomy
- What it can do
- When it's likely to fail
- Where autonomy stops and human begins
- System Architectures
- Trade-offs: sensors vs. computation
  robustness vs. computation
  complexity vs. capability
  adaptability vs. determinism

What are the big problems?

- Large systems
- Multi-agent systems
- Large-scale models
- Long-term models
- Cost

- 3-D models
- Dynamic models
- Interacting with people
- Changing the environment