

Autonomy Practices: Decision-making Architectures

KISS Workshop on
Engineering Resilient Space Systems

Brian C. Williams, MIT
July 30th, 2012



Publications at: mers.csail.mit.edu

courtesy of JPL

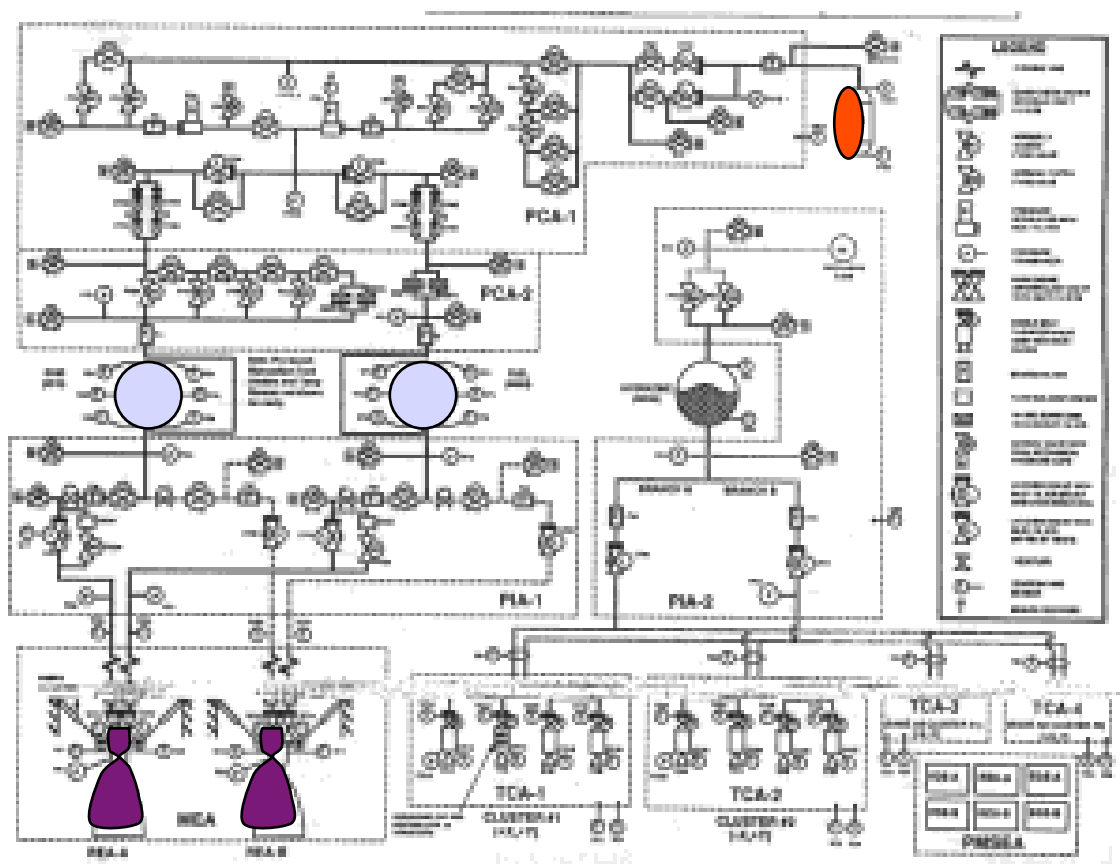
Provide a programming paradigm that enables robotic systems:

- to be commanded simply and intuitively;
- to adapt to uncertainty and failure;
- to communicate at a cognitive level;
- to work fluidly with humans, and
- to manage risk taking effectively.

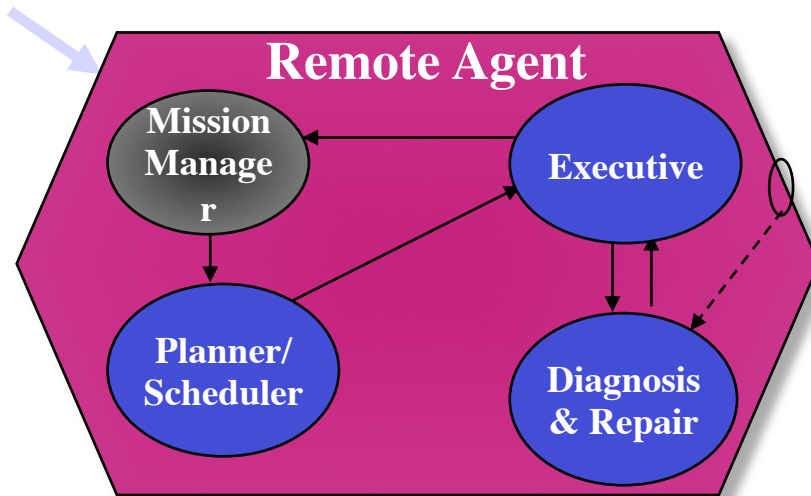
- **Autonomous Operation**
 - Cassini AACS, Remote Agent, MDS, Titan
- **Autonomous Science**
 - Autonomous Sciencecraft Experiment
- **Autonomous Navigation**
 - ClarAty, MERS DIME & Gestalt
- **Mixed Initiate Interaction**
 - MapGen + descendants



- Robust, Goal-directed Execution
- Plan Dispatching
- Diagnosis and Mode Estimation
- Plan Generation



Goals



1. Commanded by giving goals
2. Reasoned from commonsense models
3. Closed loop on goals

[Muscettola et al, AIJ 00;
Williams & Nayak, AAAI 95]

Remote Agent Experiment on Deep Space 1 – May, 1999


May 17-18th experiment: Mission-level Fault Protection

- Generate plan for course correction and thrust
- **Diagnose camera as stuck on**
 - **Power constraints violated, abort current plan & replan**
- Perform optical navigation
- Perform ion propulsion thrust

May 21th experiment: Engineering-level Fault Protection

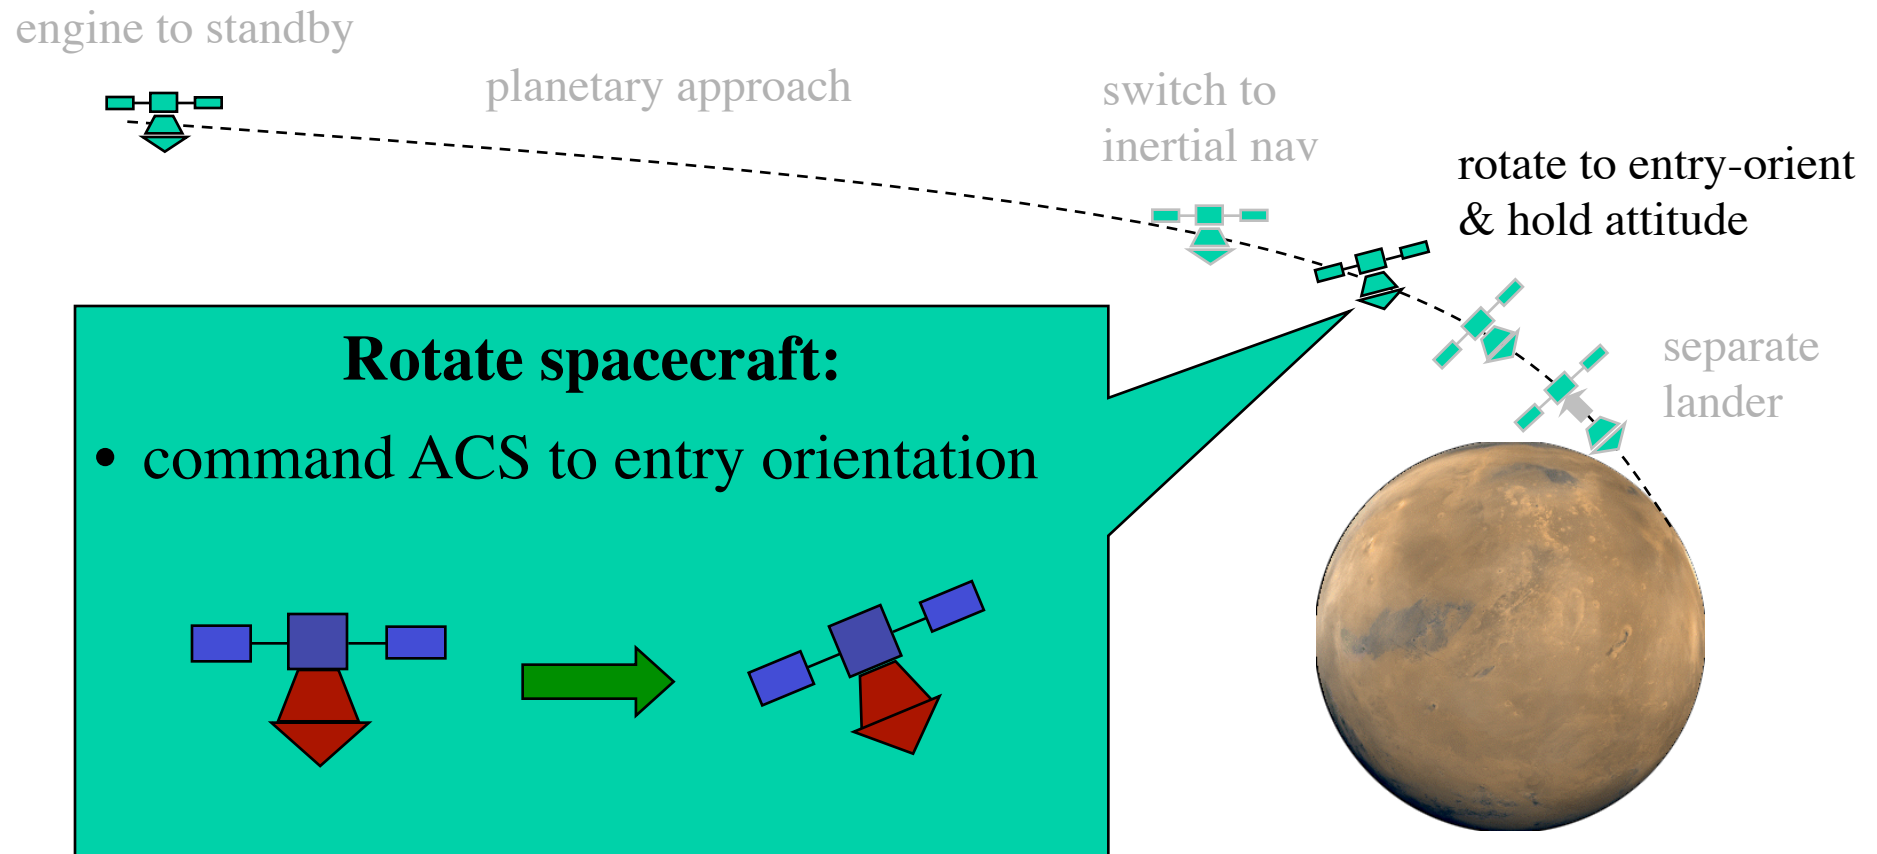
- **Diagnose faulty device and**
 - **Repair by issuing reset.**
- **Diagnose switch sensor failure.**
 - **Determine harmless, and continue plan.**
- **Diagnose thruster stuck closed and**
 - **Repair by switching to alternate method of thrusting.**
- Back to back planning

- An embedded programming **languages** **elevated** to the goal-level through operations on **hidden state** (RMPL).
- A language **executive** that achieves robustness by reasoning over **constraint-based models** (Titan).
- Interfaces that support **natural human interaction** **fluidly** and at the **cognitive level**.



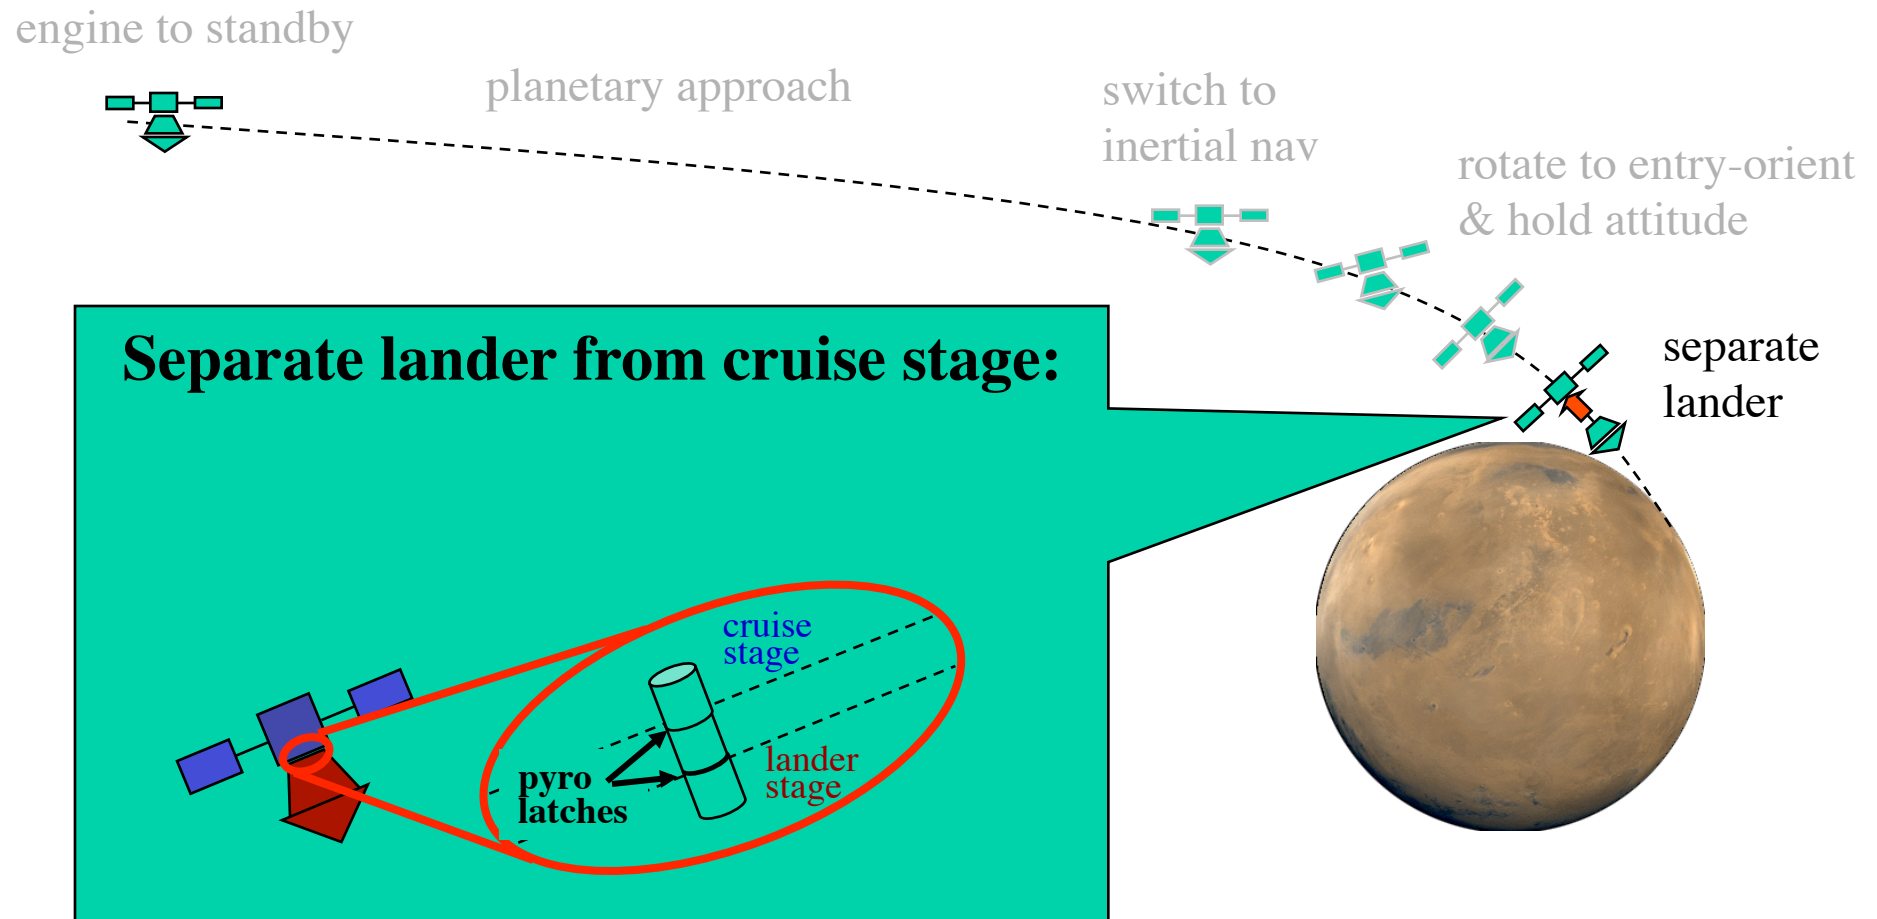
Commanding with Goals: System Engineers Specify Missions in Terms of Evolving States

[Titan Executive,
Williams et al., IEEE Procs 03]



Commanding with Goals: System Engineers Specify Missions in Terms of Evolving States

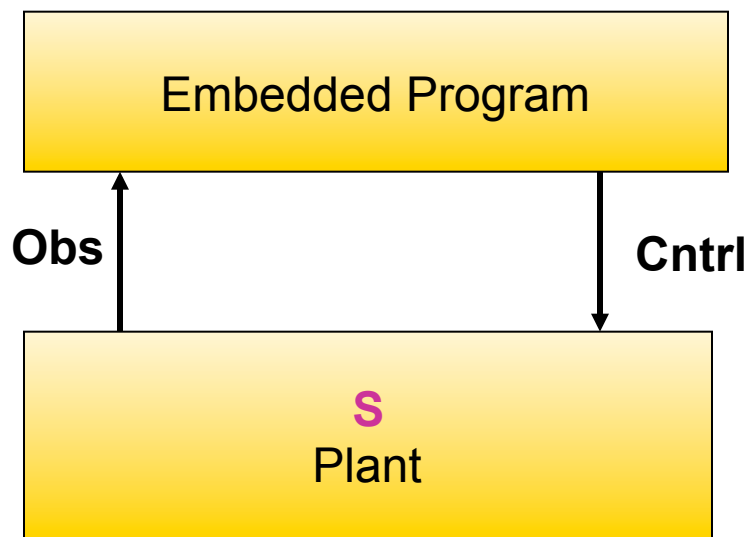
[Titan Executive,
Williams et al., IEEE Procs 03]



Autonomous Systems are Commanded in Terms of Evolving Goal States

Embedded programs evolve actions by interacting with plant sensors and actuators:

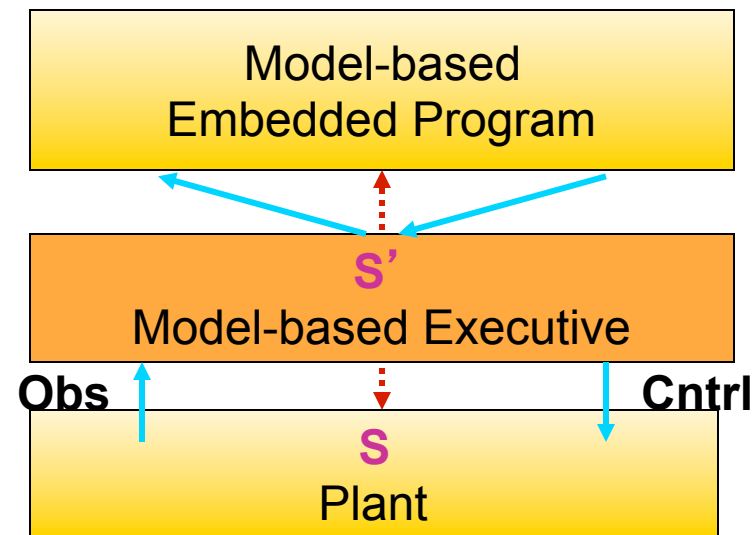
- Read sensors
- Set actuators



Programmer maps between state and sensors/actuators.

Model-based programs evolve abstract states through direct interaction:

- Read abstract state
- Write abstract state




Model-based executive maps between state and sensors/actuators.

Model-based Programs

Specify and Execute Evolving States

Turn camera off and
engine on



OrbitInsert():

do-watching (EngineA = Thrusting OR
EngineB = Thrusting)

parallel {

EngineA = Standby;

EngineB = Standby;

Camera = Off;

do-watching (EngineA = Failed)

{**when-donext** (EngineA = Standby) AND
Camera = Off)

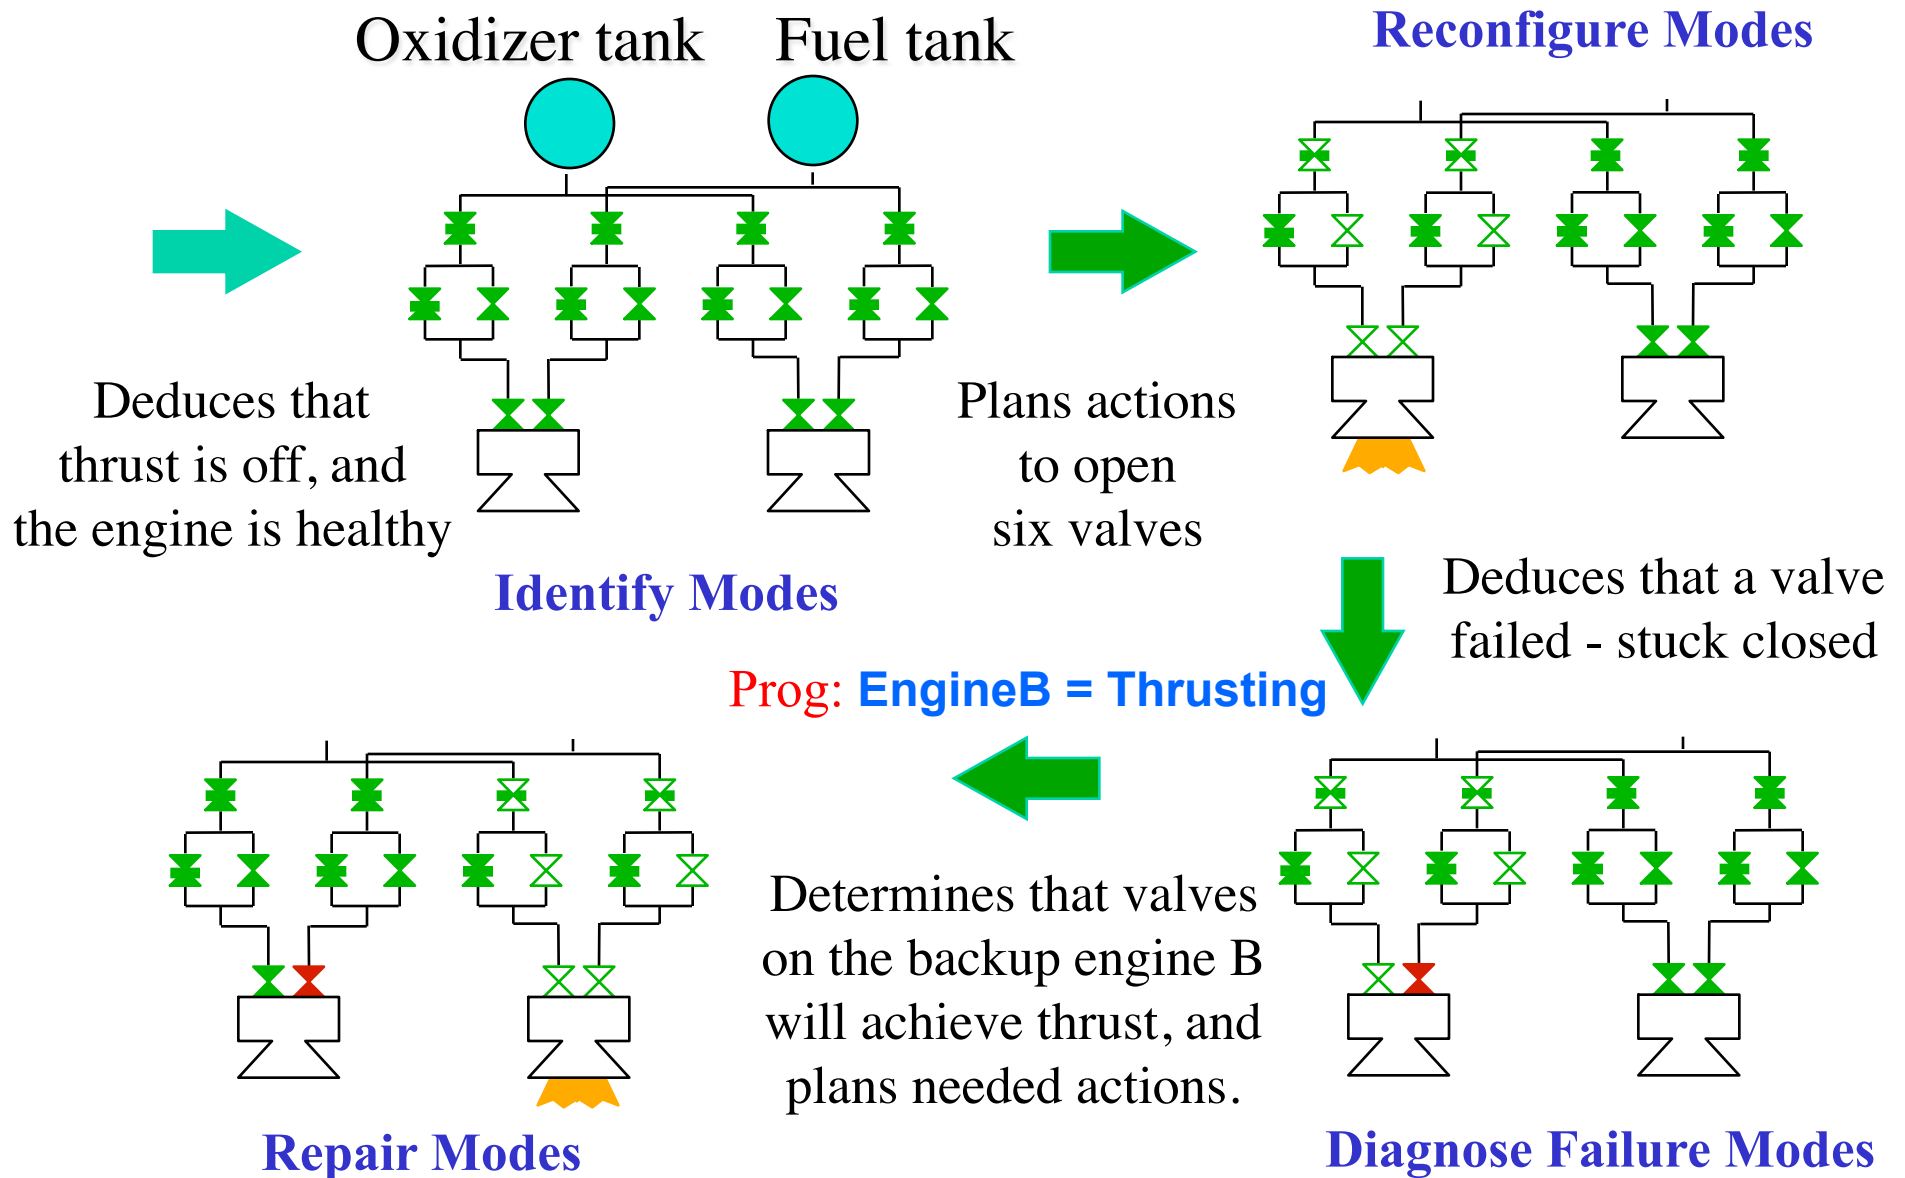
EngineA = Thrusting};

when-donext (EngineA = Failed AND
EngineB = Standby AND
Camera = Off)

EngineB = Thrusting}

[Titan Executive,
Williams et al., IEEE Procs 03]

The program assigns **EngineA = Thrusting**,
and the model-based executive



Behaviors Generated from a Plant Model

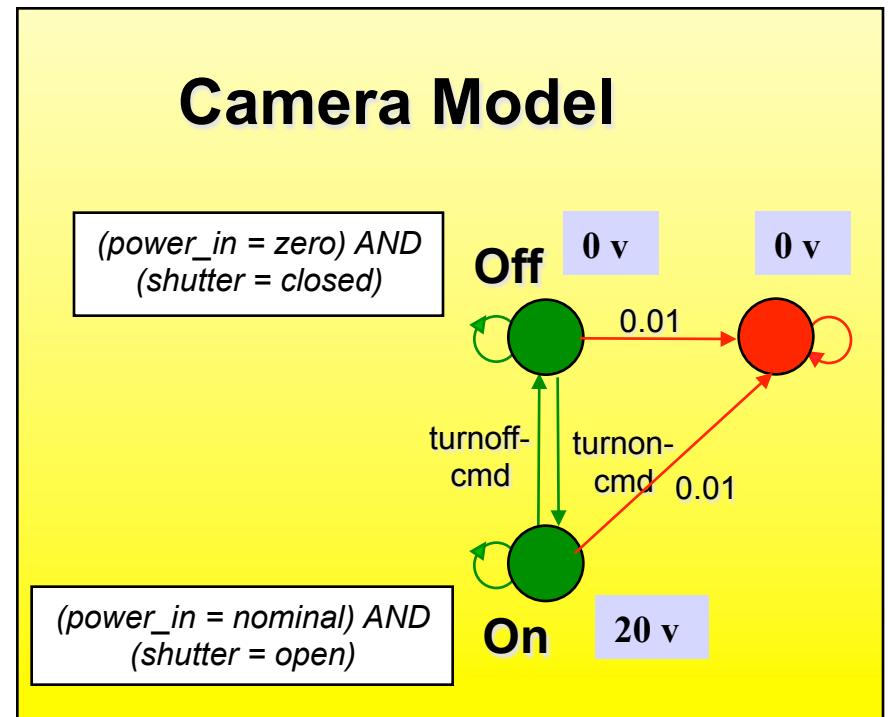
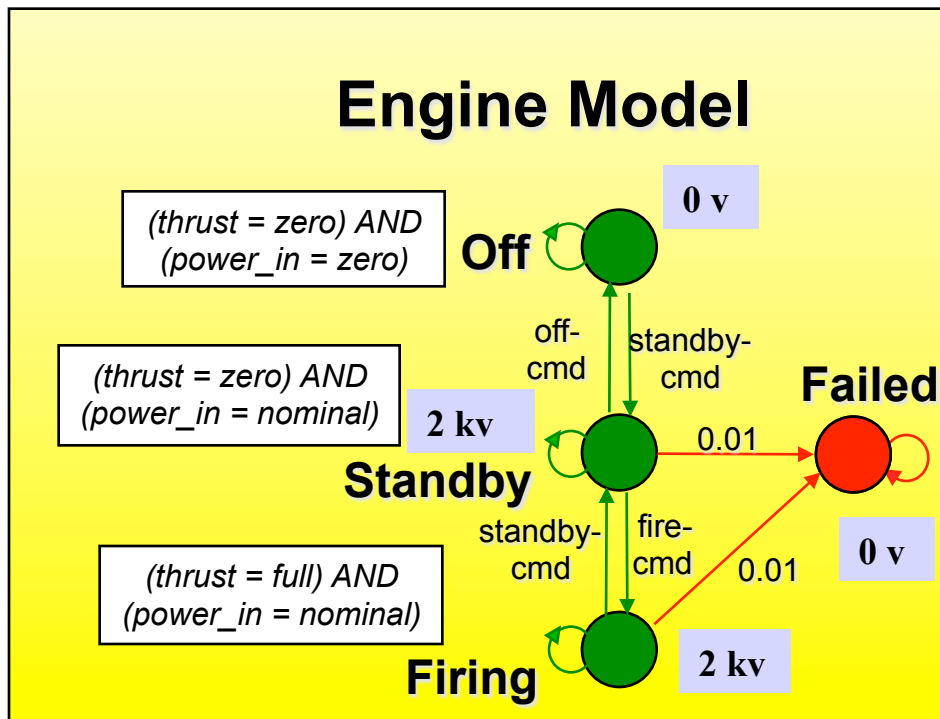
component modes...

[Titan Executive,
Williams et al., IEEE Procs 03]

described by finite domain constraints on variables...

deterministic and probabilistic transitions

cost/reward



one per component ... operating concurrently



Titan Model-based Executive

RM

```
OrbitInsert():  
(do-watching ((EngineA = Firing) OR  
              (EngineB = Firing))  
  (parallel  
    (EngineA = Standby)  
    (EngineB = Standby)  
    (Camera = Off)  
    (do-watching (EngineA = Failed)  
      (when-donext ( (EngineA = Standby) AND  
                    (Camera = Off) )  
        (EngineA = Firing)))  
    (when-donext ( (EngineA = Failed) AND  
                  (EngineB = Standby) AND  
                  (Camera = Off) )  
      (EngineB = Firing))))
```

Generates target goal states
conditioned on state estimates

System Model



Estimates

Goals


Tracks
likely
modes

Tracks least cost
modes achieving
goal states

Observations

Commands

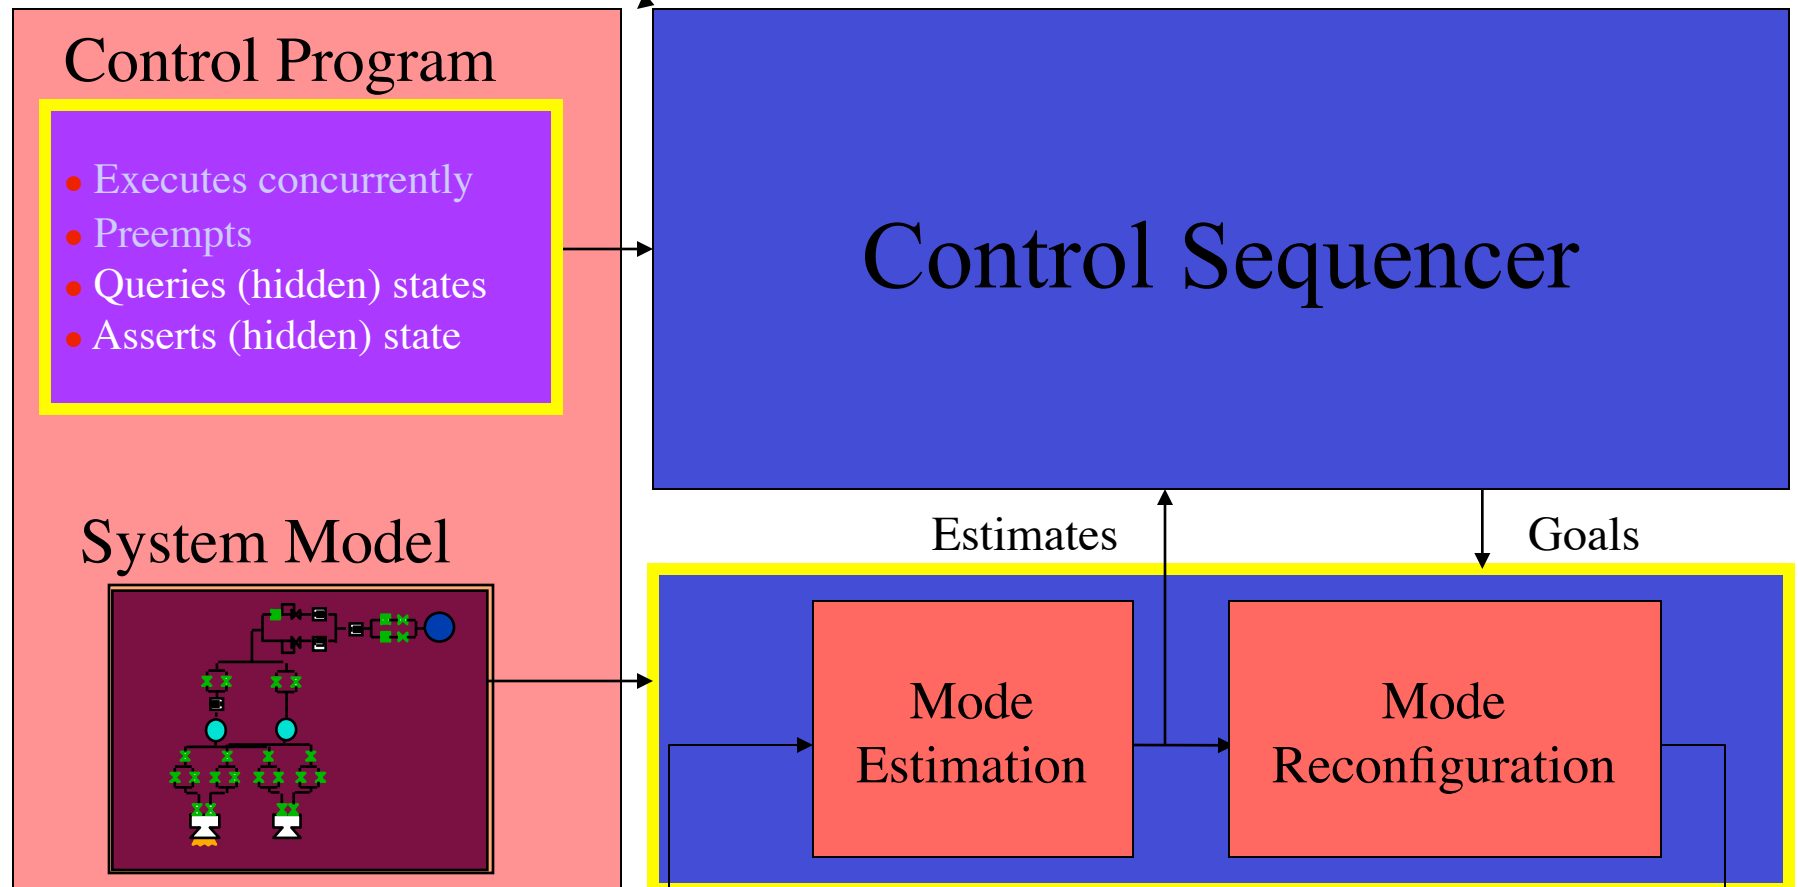
Plant



Control Programs are like State Charts

RMPL Model-based Program

Titan Model-based Executive




A kind of autonomous configuration and health management system.

Architecture similar to

- **Cassini AACCS FP**
- **MDS**

Observations


Plant



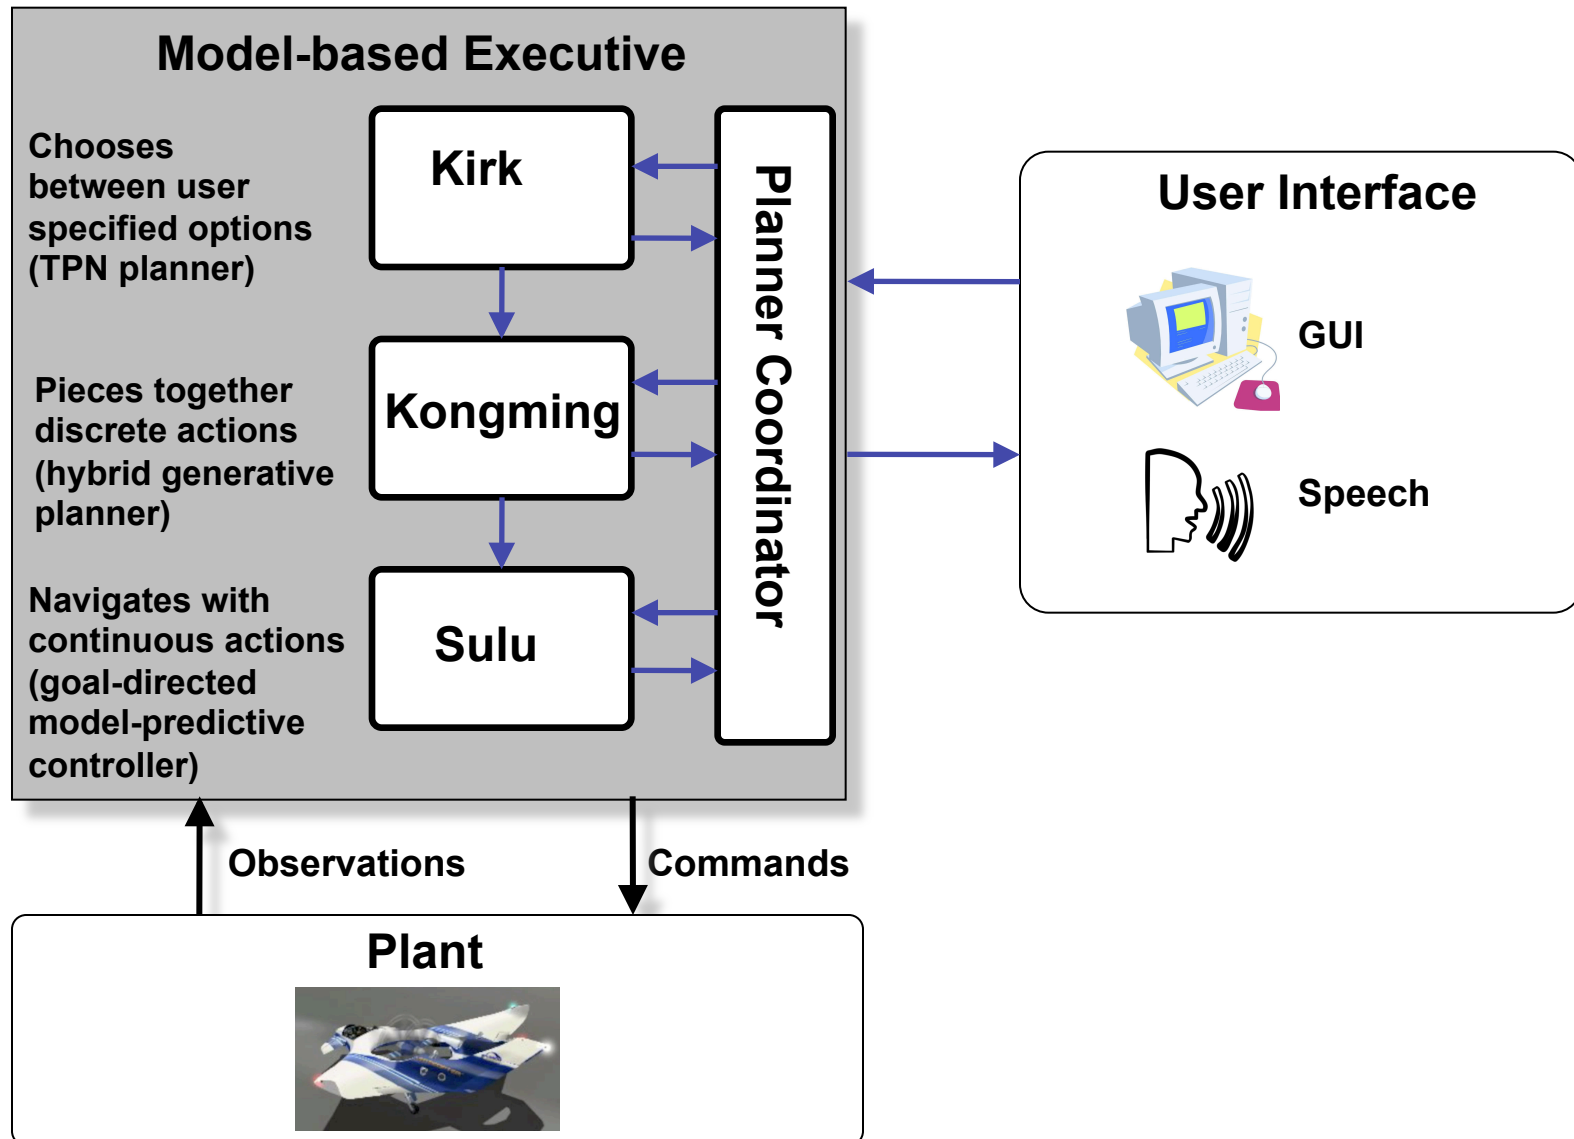
Commands



Navigation, Risk and Mixed Interaction: Personal Transportation System

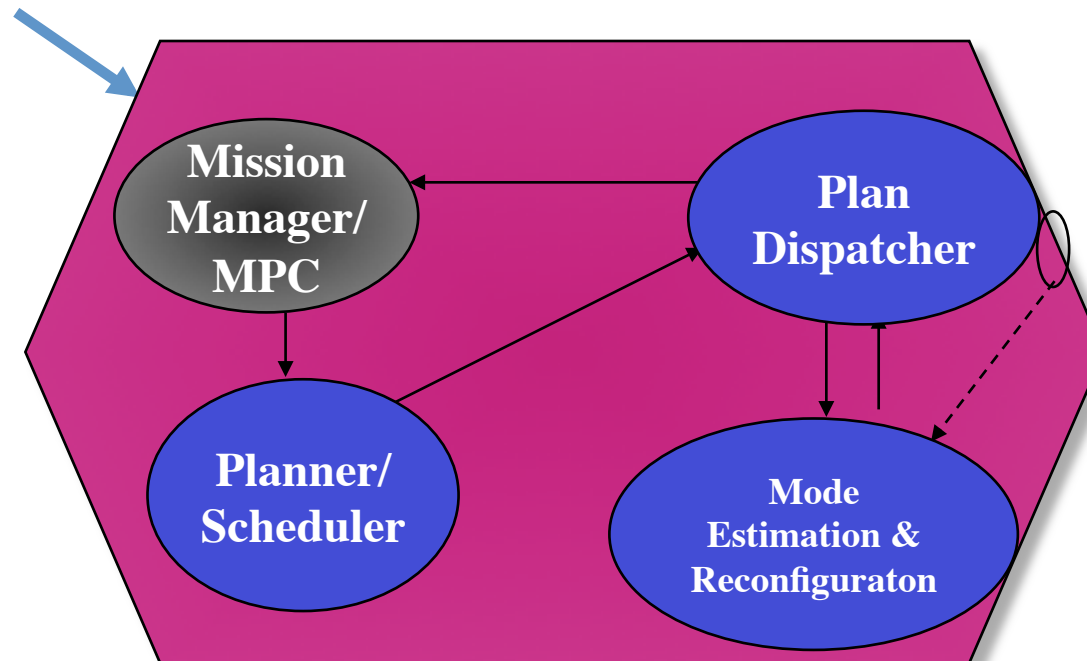


Complements of Branko Sarh, Boeing Research



1. Commanded through time evolved goals.
2. Reasons from commonsense models.
3. Closes loop on goals.
4. Model-based programs specify goals and models.

Goals





- Robust, Goal-directed Execution
- Plan Dispatching
- Diagnosis and Mode Estimation
- Plan Generation

Supports Time Critical Missions



Model-based Embedded & Robotic Systems



Massachusetts Institute of Technology

Supports Robot & Human Coordination



An effective Scrub Nurse:


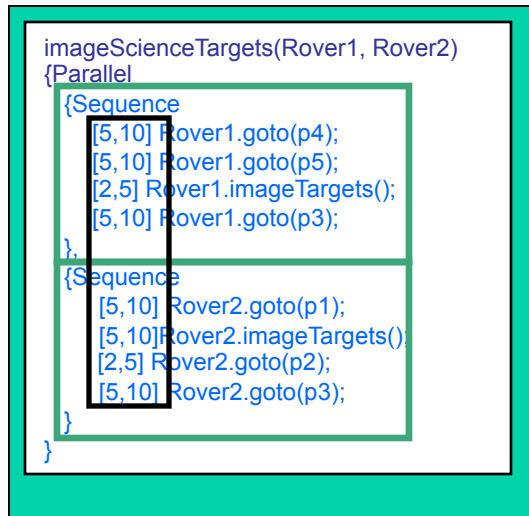
- works hand-to-hand, face-to-face with surgeon,
- assesses and anticipates needs of surgeon,
- provides assistance and tools in order of need,
- responds quickly to changing circumstances,
- responds quickly to surgeon's cues and requests.

[Shah, Conrad and Williams, ICAPS 09]

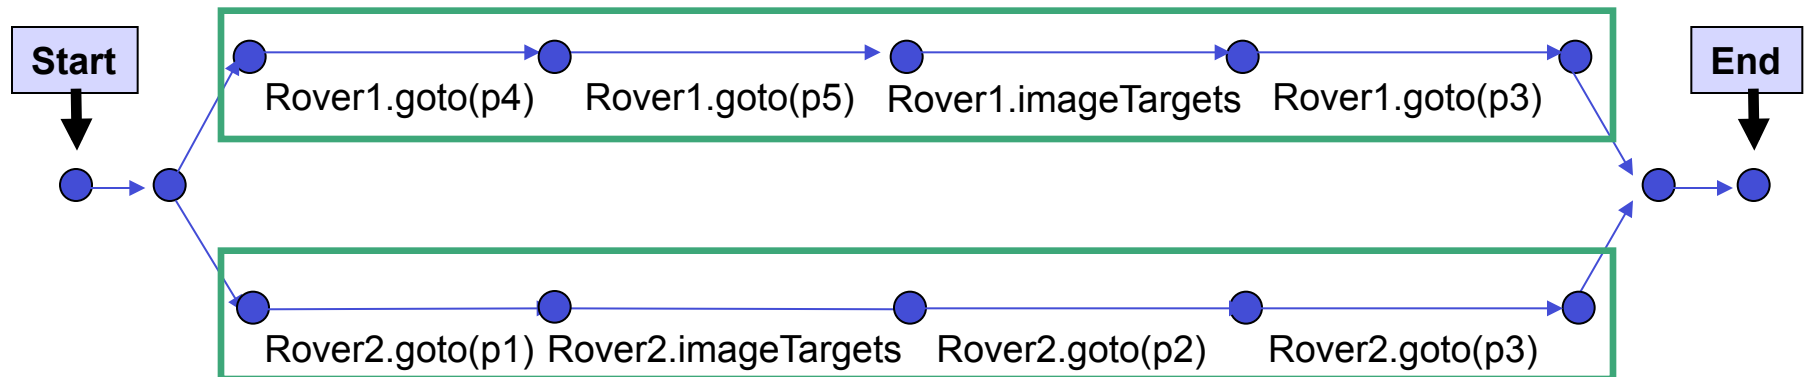
Robust Plan Execution



Model-based Embedded & Robotic Systems



in RMPL [williams et al]



Agents adapt to temporal disturbances in a coordinated manner by scheduling the start of activities on the fly.

[Muscettola, Morris, Tsamardinos, KR 98]



Massachusetts Institute of Technology



Outline: To Execute a Temporal Plan

Schedule Off-line

1. Describe Temporal Plan



2. Test Consistency



3. Schedule Plan



4. Execute Plan

Schedule Online

1. Describe Temporal Plan



2. Test Consistency



3. Reformulate Plan

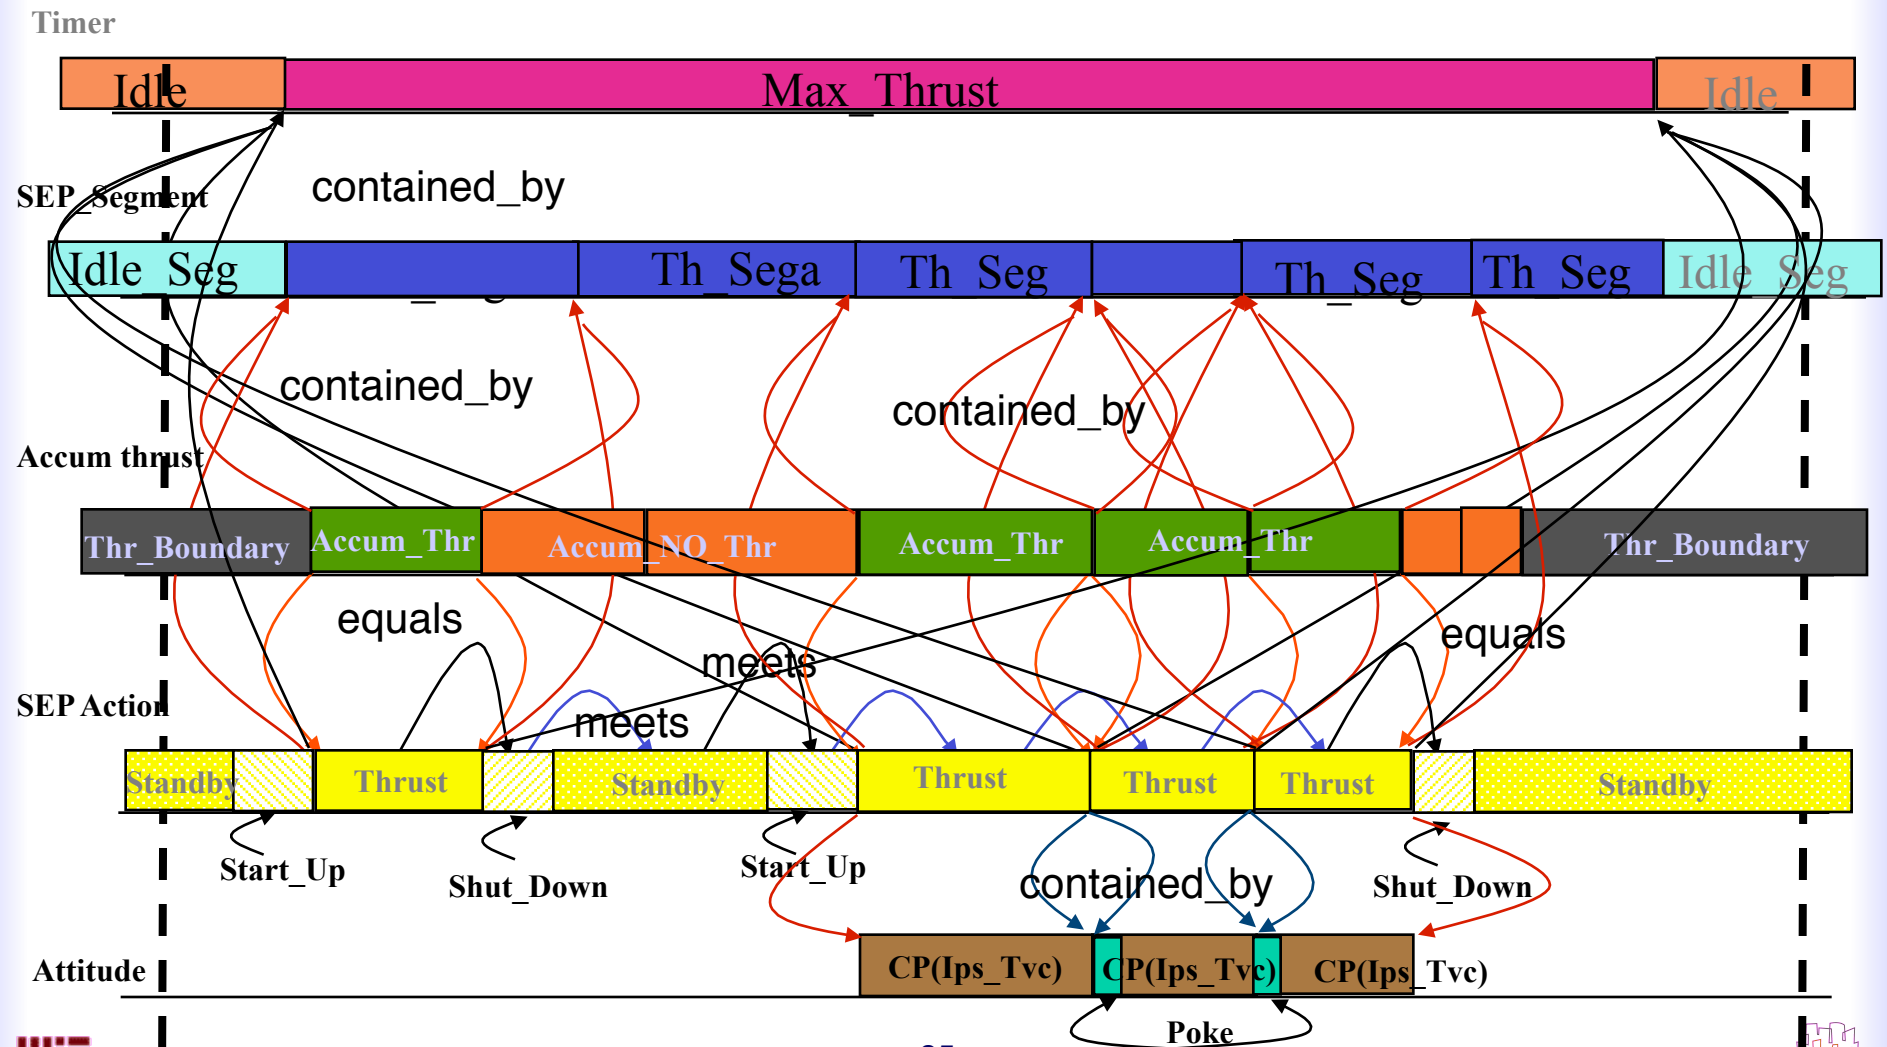


4. Dynamically Schedule Plan

offline
online


Describe Temporal Plan

Example: Deep Space One Remote Agent Experiment



Scheduling a Simple Temporal Network (STN)

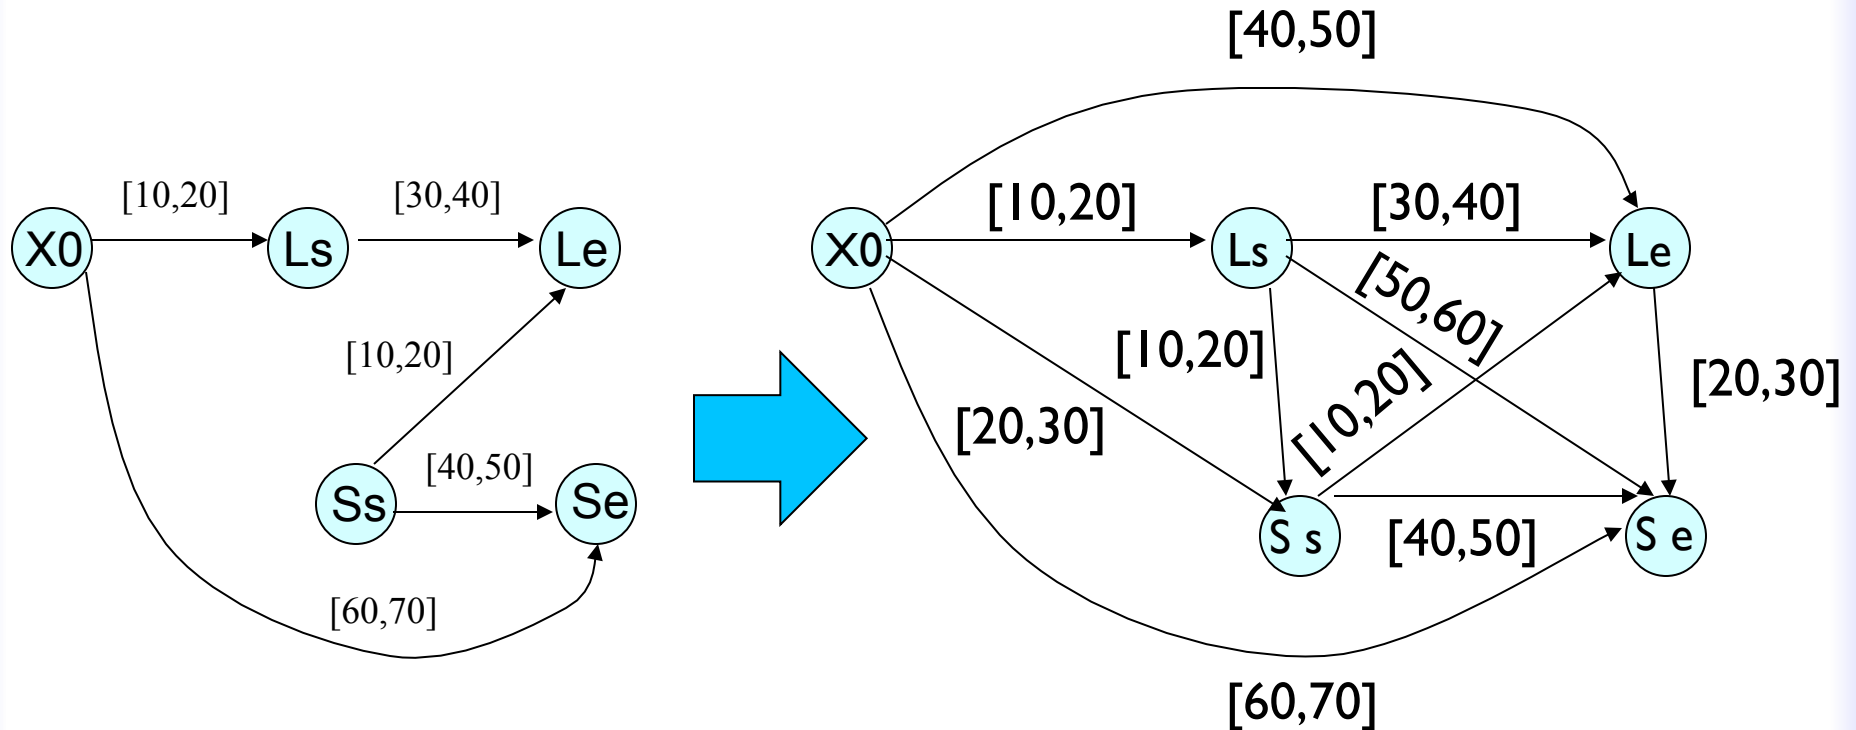
Input: An STN $\langle X, C \rangle$ where $C_j = \langle \langle X_k, X_i \rangle \langle a_j, b_j \rangle \rangle$



Output: An assignment to X satisfying C .

$$A=0, B=2, C=1, D=3$$

Scheduling without Search



Idea: Expose Implicit Constraints in STN

- Input: STN.
- Output: “Decomposable” (Implied) STN \rightarrow Schedule.

Scheduling without Search



Model-based Embedded & Robotic Systems


Input: Decomposable STN (APSP D-Graph)

Output: Schedule (Assignment to X, consistent with STN)

Property: Can assign variables in any order, without backtracking.

Key ideas

- Incrementally tighten feasible intervals, as commitments are made.



Scheduling without Search



Model-based Embedded & Robotic Systems

Input: Decomposable STN (APSP D-Graph)


Output: Schedule (Assignment to X , consistent with STN)

Property: Can assign variables in any order, without backtracking.

Key ideas

- Incrementally tighten feasible intervals, as commitments are made.

- Select value for X_0



Scheduling without Search



Model-based Embedded & Robotic Systems

Input: Decomposable STN (APSP D-Graph)


Output: Schedule (Assignment to X , consistent with STN)

Property: Can assign variables in any order, without backtracking.

Key ideas

- Incrementally tighten feasible intervals, as commitments are made.

- Select value for X_0



Scheduling without Search



Model-based Embedded & Robotic Systems

Input: Decomposable STN (APSP D-Graph)


Output: Schedule (Assignment to X, consistent with STN)

Property: Can assign variables in any order, without backtracking.

Key ideas

- Incrementally tighten feasible intervals, as commitments are made.

- Select value for X0
- Select value for Ls, consistent with X0



Scheduling without Search



Model-based Embedded & Robotic Systems

Input: Decomposable STN (APSP D-Graph)


Output: Schedule (Assignment to X, consistent with STN)

Property: Can assign variables in any order, without backtracking.

Key ideas

- Incrementally tighten feasible intervals, as commitments are made.

- Select value for X0
- Select value for Ls, consistent with X0
- Select value for Le, consistent with X0, Ls
- Select value for Ss, consistent with X0, Ls, Le
- Select value for Se...



Flexible Execution



Model-based Embedded & Robotic Systems

Schedule Off-line

1. Describe Temporal Plan



2. Test Consistency



3. Schedule Plan



4. Execute Plan

offline
online

Problem: delays and fluctuations in task duration can cause plan failure.

Observation: Least commitment temporal plans leave room to adapt.

Flexible Execution adapts through dynamic scheduling [Muscettola et al]
– Assigns time to event when executed.

Flexible Execution



Model-based Embedded & Robotic Systems

Schedule Off-line

1. Describe Temporal Plan



2. Test Consistency



3. Schedule Plan



4. Execute Plan

Schedule Online

1. Describe Temporal Plan



2. Test Consistency



3. Reformulate Plan



4. Dynamically Schedule Plan

offline
online


Dynamic Scheduling by Decomposition?



Model-based Embedded & Robotic Systems

Consider a Simple Example

- Select executable timepoint and assign.
- Propagate assignment to neighbors.



[Muscettola, Morris, Tsamardinos KR98]


Dynamic Scheduling by Decomposition?



Model-based Embedded & Robotic Systems

Consider a Simple Example

- Select executable timepoint and assign.
- Propagate assignment to neighbors.



[Muscettola, Morris, Tsamardinos KR98]


Dynamic Scheduling by Decomposition?



Model-based Embedded & Robotic Systems

Consider a Simple Example

- Select executable timepoint and assign
- Propagate assignment to neighbors



[Muscettola, Morris, Tsamardinos KR98]




Dynamic Scheduling by Decomposition?



Model-based Embedded & Robotic Systems

- Fix by scheduling according to implied orderings.
 - $A \leq C < B < D$



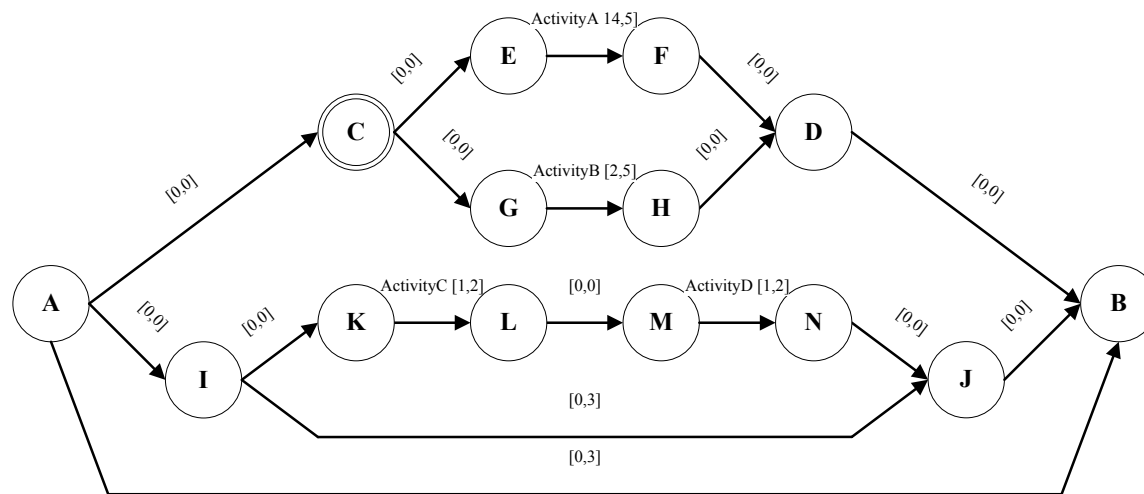
[Muscettola, Morris, Tsamardinos KR98]

Generalizations of Dynamic Execution



Model-based Embedded & Robotic Systems


1. Dynamically choose a) when, b) by whom, c) which method, and d) how.
 2. Achieve consistency wrt models of uncontrollable events.
- Temporal Plan Networks (under Uncertainty)
 - Disjunctive Temporal Networks (under Uncertainty)



Multi-Robot Teamwork



Model-based Embedded & Robotic Systems



Agents choose and schedule activities.

(Someone) Remove one ball from red bin.




Remove one ball from red bin.




OR




$L[32,39] \vee R[42,55]$



Multi-Robot Teamwork



- Off-nominal.
- Partner adapts in response to teammate's failure.




**Kim, Williams, Abramson IJCAI 2001;
Shah and Williams ICAPS 2009;
Conrad, Shah and Williams ICAPS 2010**

Leader & Assistant

Assistant waits to see what Leader will do before acting.

Leader



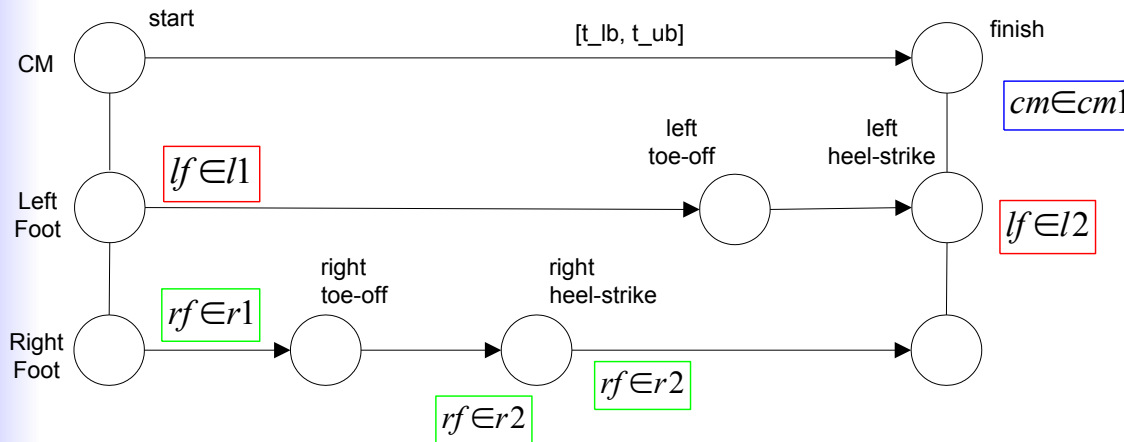

Assistant

Shah and Williams ICAPS 2010

Idea: model leader durations and assignments as uncontrollable (TPNU).



Achieve safety by adaptively executing plans on qualitative poses

Input: Qualitative State Plan

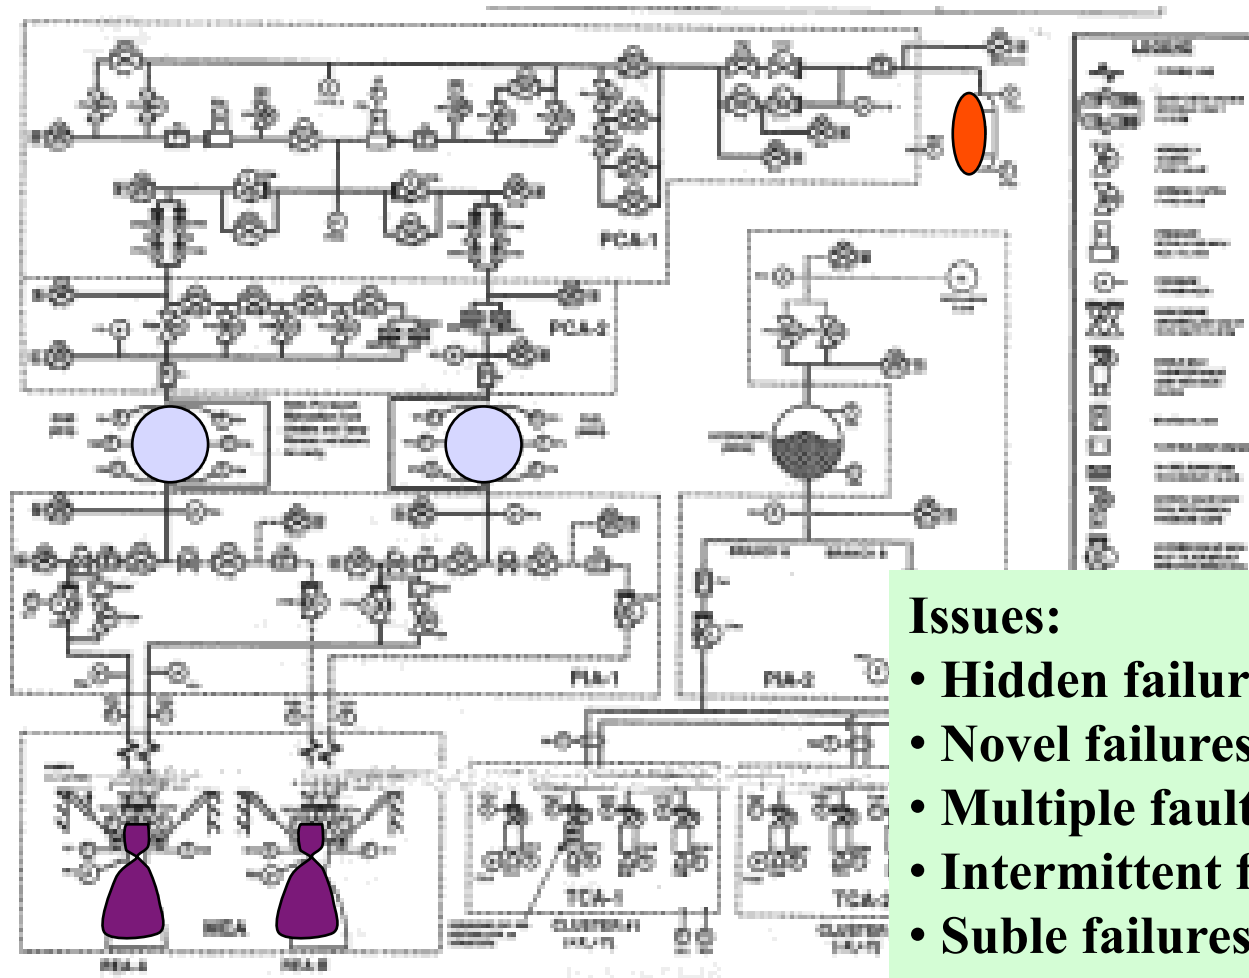


Lateral CM with push disturbance

- Blue - 40 N
- Green - 35 N
- Black - 25 N
- Red - Max allowed displacement



- Robust, Goal-directed Execution
- Plan Dispatching
- Diagnosis and Mode Estimation
- Plan Generation



Issues:


- Hidden failures
- Novel failures
- Multiple faults
- Intermittent failures
- Suble failures.
- HW / SW interactions

- Mode estimates and kernels.
- By divide and conquer (GDE).
- Likely estimates and Conflict-directed A*.
- Mode reconfiguration.
- Estimating probabilistic (hybrid) constraint automata.

Model-based Diagnosis

Input: **Observations** of a system with symptomatic behavior,
and a **model Φ** of the system.

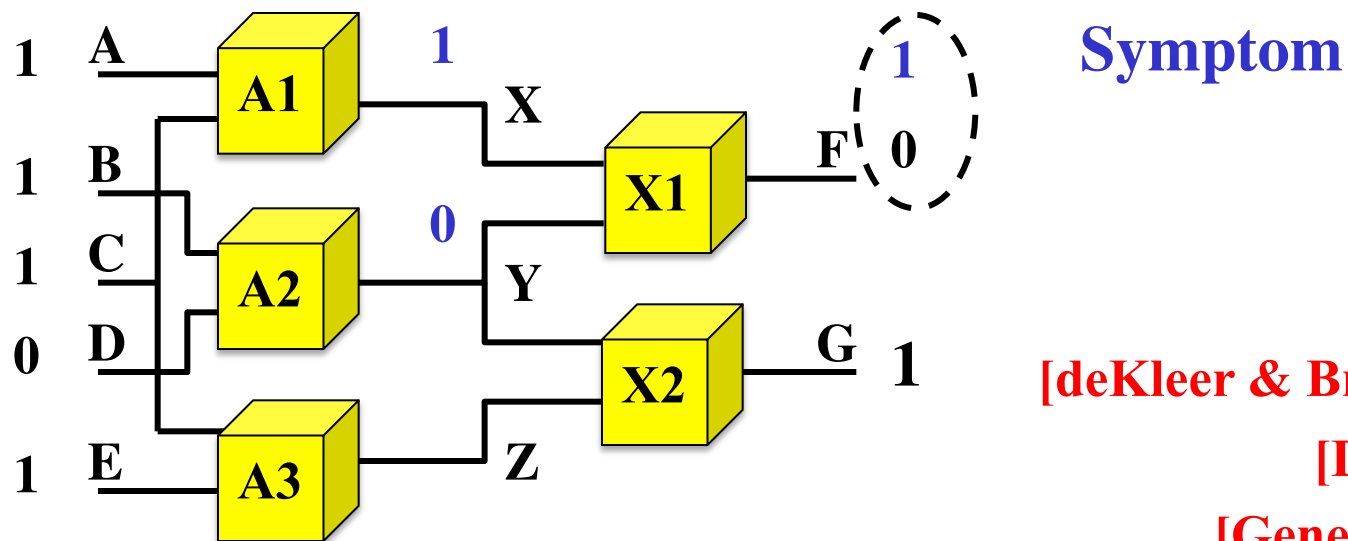
Output: **Diagnoses** that **account** for the **symptoms**.



How Should Diagnoses Account for Novel Failures?

Consistency-based Diagnosis: Given symptoms, find diagnoses that are consistent with symptoms.

Suspending Constraints: For novel faults, make no presumption about faulty component behavior.



[deKleer & Brown, 83]

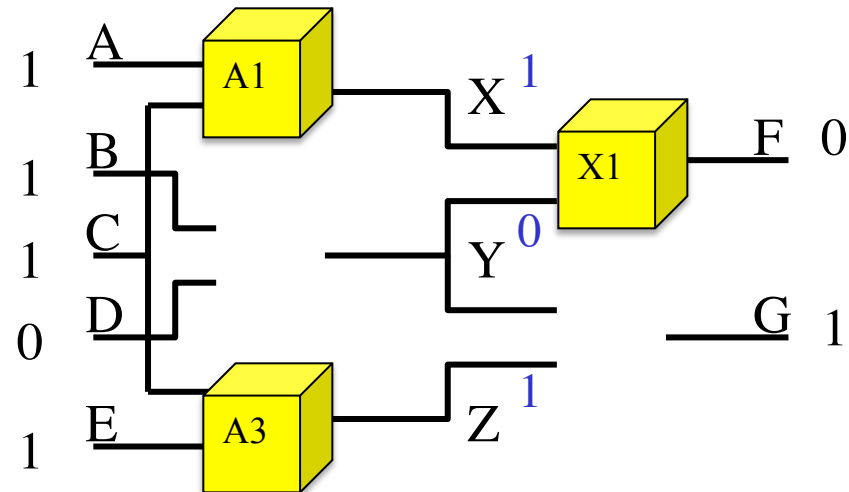
[Davis, 84]

[Geneserth, 84]

Multiple Faults: Identify all Combinations of Consistent “Unknown” Modes

And(i):

- G(i):
Out(i) = In1(i) AND In2(i)
- U(i): *No Constraint*




Diagnosis = {A1=G, A2=U, A3=G, X1=G, X2=U}

- Candidate: Assignment of G or U to each component.
- Diagnosis: Candidate consistent with model and observations.

Incorporating (Failure) Modes: Mode Estimation

Sherlock

[de Kleer & Williams, IJCAI 89]



Idea: Include Nominal, Fault and Unknown Modes

Inverter(i):

- G(i): $\text{Out}(i) = \text{not}(\text{In}(i))$
- S1(i): $\text{Out}(i) = 1$
- S0(i): $\text{Out}(i) = 0$
- U(i):

- Isolates unknown faults.
- Explains known faults.

Compact Encoding: Partial Diagnoses

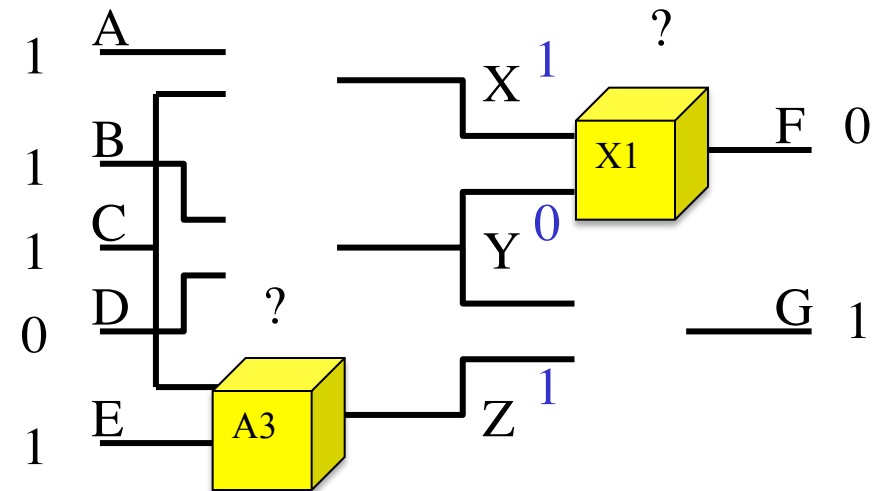
Partial Diagnosis

$\{A1=U, A2=U, X2=U\}$

Partial Diagnosis:

A partial mode assignment M,
that “removes all symptoms.”

- All full extensions of M are diagnoses.



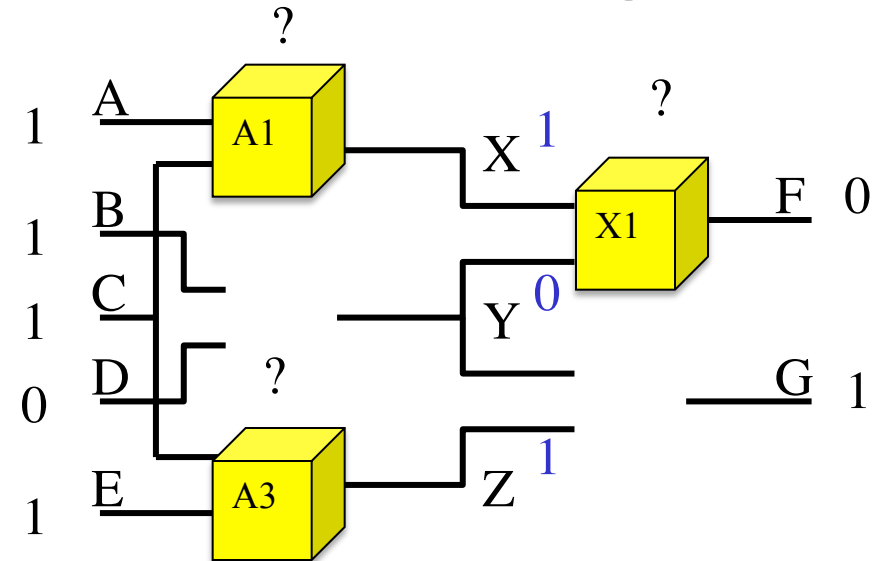
Extensions (Diagnoses):

- $\{A1=U, A2=U, A3=G, X1=G, X2=U\}$
- $\{A1=U, A2=U, A3=G, X1=U, X2=U\}$
- $\{A1=U, A2=U, A3=U, X1=G, X2=U\}$
- $\{A1=U, A2=U, A3=U, X1=U, X2=U\}$

Compact Encoding: Kernel Diagnoses

Kernel Diagnosis

$\{A2=U, X2=U\}$



Partial Diagnosis:

A partial mode assignment M,
that “removes all symptoms”.

- All full extensions of M are diagnoses.

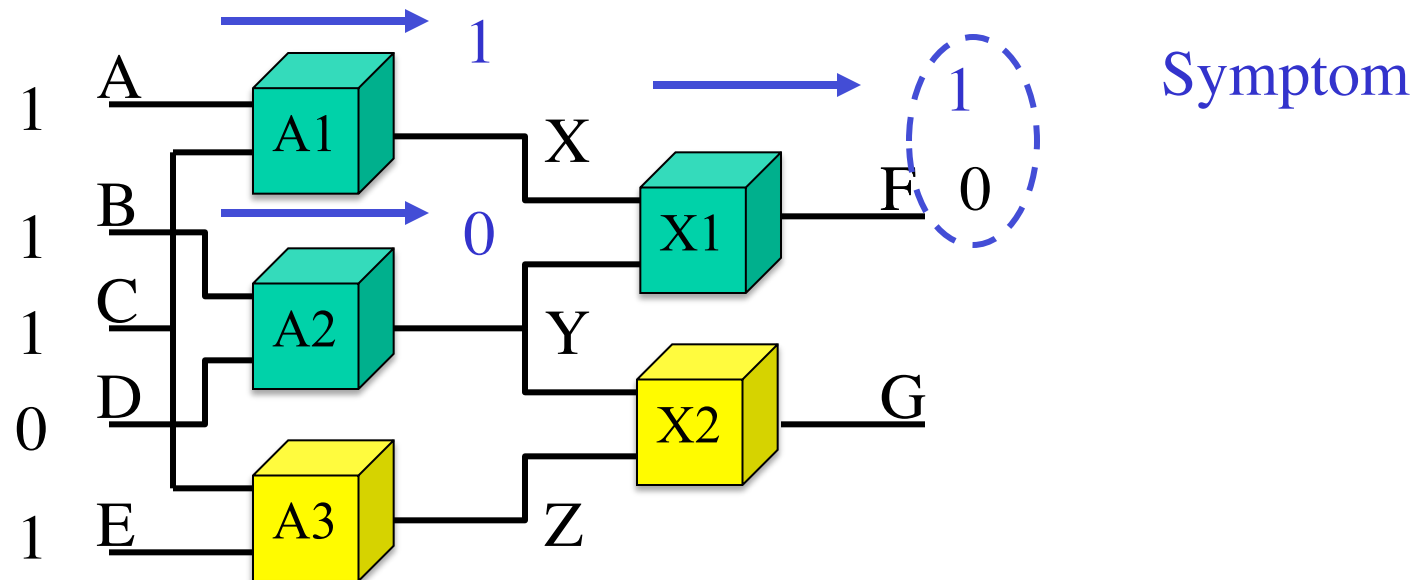
Kernel Diagnosis:

A partial diagnosis K,
no subset of which is a partial diagnosis.

Mode Estimation

- Mode estimates and kernels.
- By divide and conquer (GDE).
- Likely estimates and Conflict-directed A*.
- Mode reconfiguration.
- Estimating probabilistic (hybrid) constraint automata.

Conflicts Explain How to Remove Symptoms



Symptom:

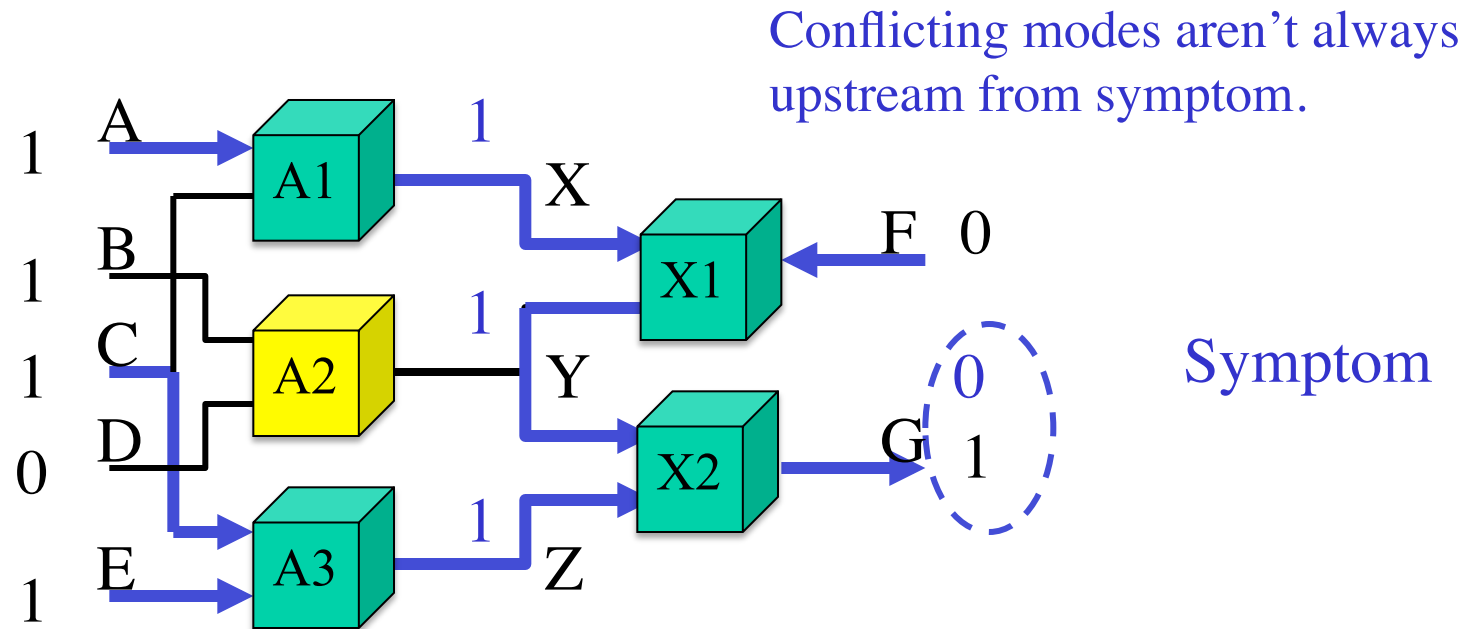
F is observed 0, but predicted to be 1 if A1, A2 and X1 are okay.

Conflict 1: $\{A1=G, A2=G, X1=G\}$ is inconsistent.

→ One of A1, A2 or X1 must be broken.

Conflict: An inconsistent partial assignment to mode variables X.

Second Conflict



Symptom: G is observed 1, but predicted 0.

Conflict 2: $\{A1=G, A3=G, X1=G, X2=G\}$ is inconsistent.

→ One of A1, A3, X1 or X2 must be broken.

Candidate Generation:

Generate Kernels From Conflicts

$\{A1=G, A2=G, X1=G\}$

Conflict 1.

$\{A1=G, A3=G, X1=G, X2=G\}$

Conflict 2.

$\{A1=U, A2=U, X1=U\}$

diagnoses for Conflict 1.

$\{A1=U, A3=U, X1=U, X2=U\}$

diagnoses for Conflict 2.

Kernel Diagnoses =

“Smallest” sets of modes that remove all conflicts.

Candidate Generation: Generate Kernels From Conflicts

$\{A1=G, A2=G, X1=G\}$

Conflict 1.

$\{A1=G, A3=G, X1=G, X2=G\}$

Conflict 2.

$\{A1=U, A2=U, X1=U\}$

diagnoses for Conflict 1.

$\{A1=U, A3=U, X1=U, X2=U\}$

diagnoses for Conflict 2.

Kernel Diagnoses = $\{A1=U\}$



1. Compute cross product.
2. Remove supersets.
 - Old subset New.
 - New subset Old.

“Smallest” sets of modes that remove all conflicts.

Candidate Generation: Generate Kernels From Conflicts

$\{A1=G, A2=G, X1=G\}$

Conflict 1.

$\{A1=G, A3=G, X1=G, X2=G\}$

Conflict 2.

$\{A1=U, A2=U, X1=U\}$

constituents of Conflict 1.

$\{A1=U, A3=U, X1=U, X2=U\}$

constituents of Conflict 2.

Kernel Diagnoses =

$\{X1=U\}$

$\{A2=U, X2=U\}$

$\{A2=U, A3=U\}$

$\{A1=U\}$

1. Compute cross product.
2. Remove supersets.
 - Old subset New.
 - New subset Old.

“Smallest” sets of modes that remove all conflicts.

Mode Estimation

- Estimates and kernels.
- By divide and conquer (GDE).
- Likely estimates and Conflict-directed A*.
- Mode reconfiguration.
- Estimating probabilistic (hybrid) constraint automata.



Mode Estimation as Conflict-directed Best First Search

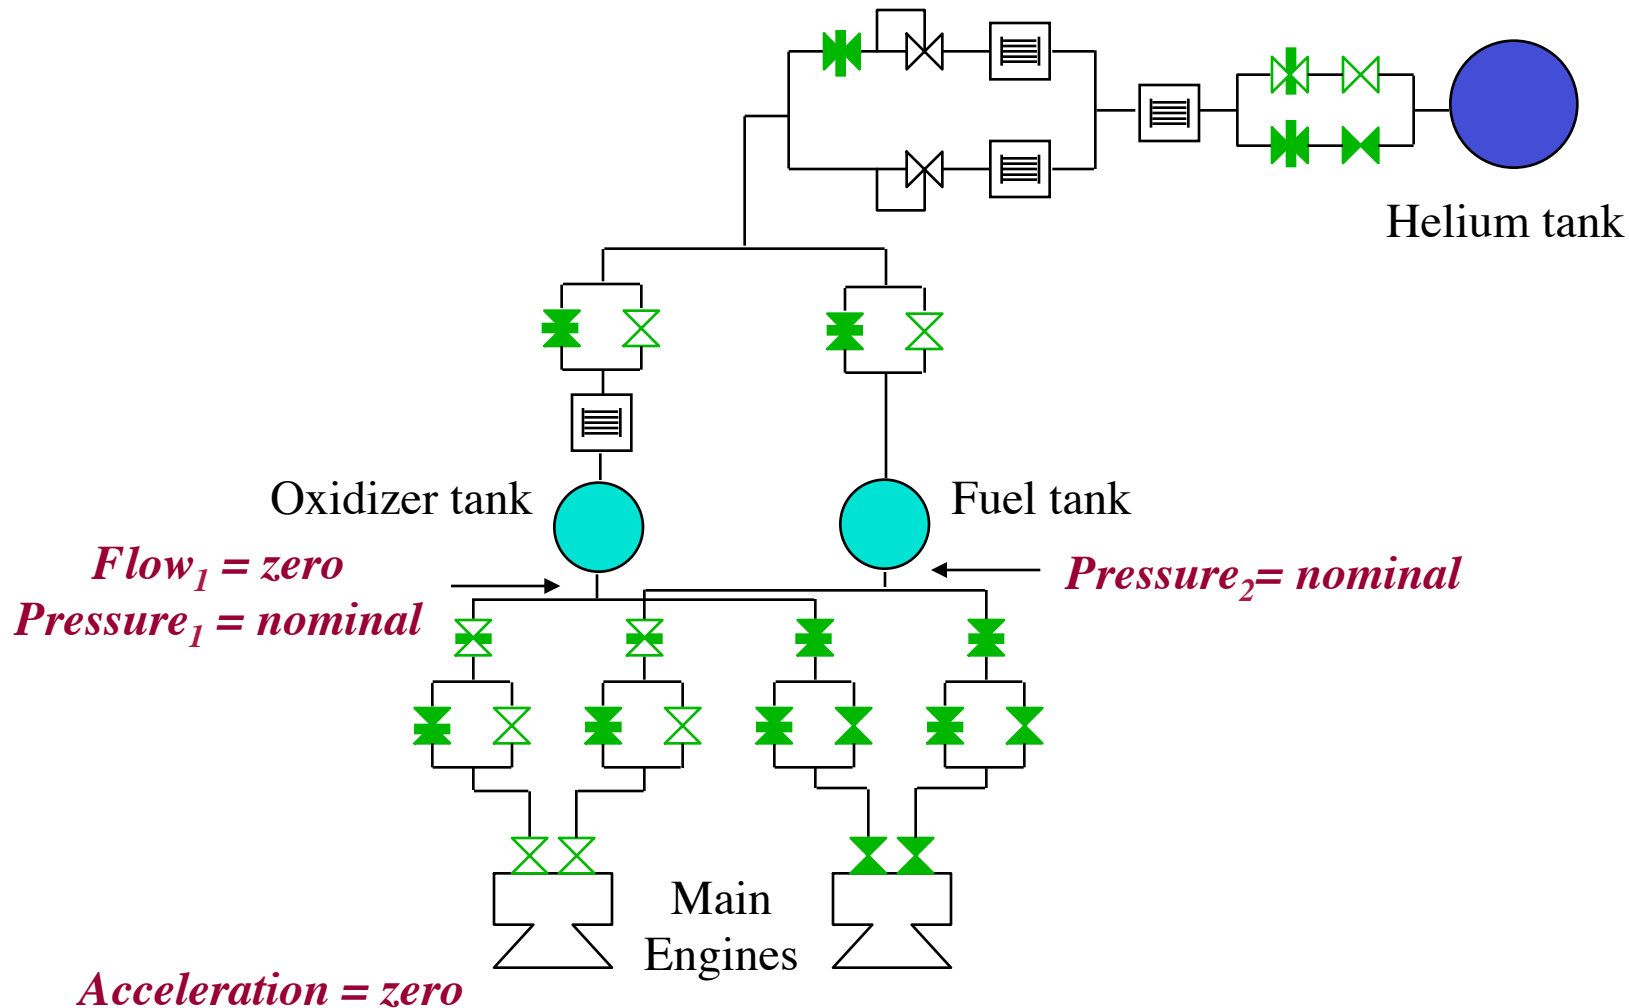


When you have eliminated the impossible,
whatever remains, however improbable,
must be the truth.

- Sherlock Holmes. The Sign of the Four.

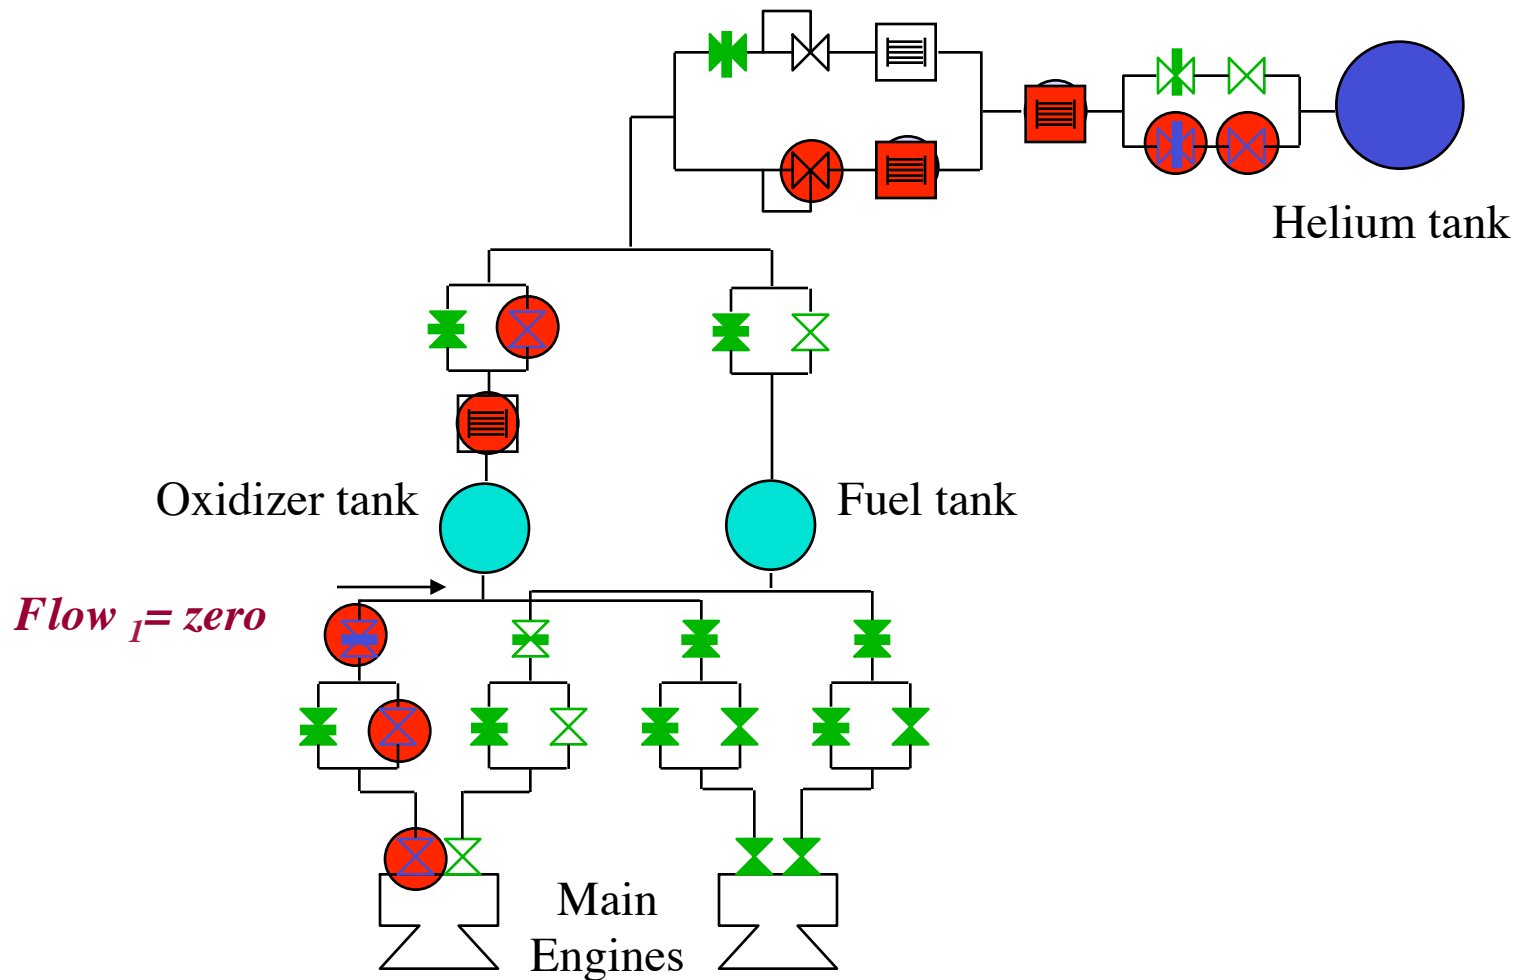
1. Generate most likely hypothesis.
2. Test hypothesis.
3. If inconsistent, learn reason for inconsistency
(a conflict).
4. Use conflicts to leap over similarly infeasible options
to next best hypothesis.

Compare Most Likely Hypothesis to Observations



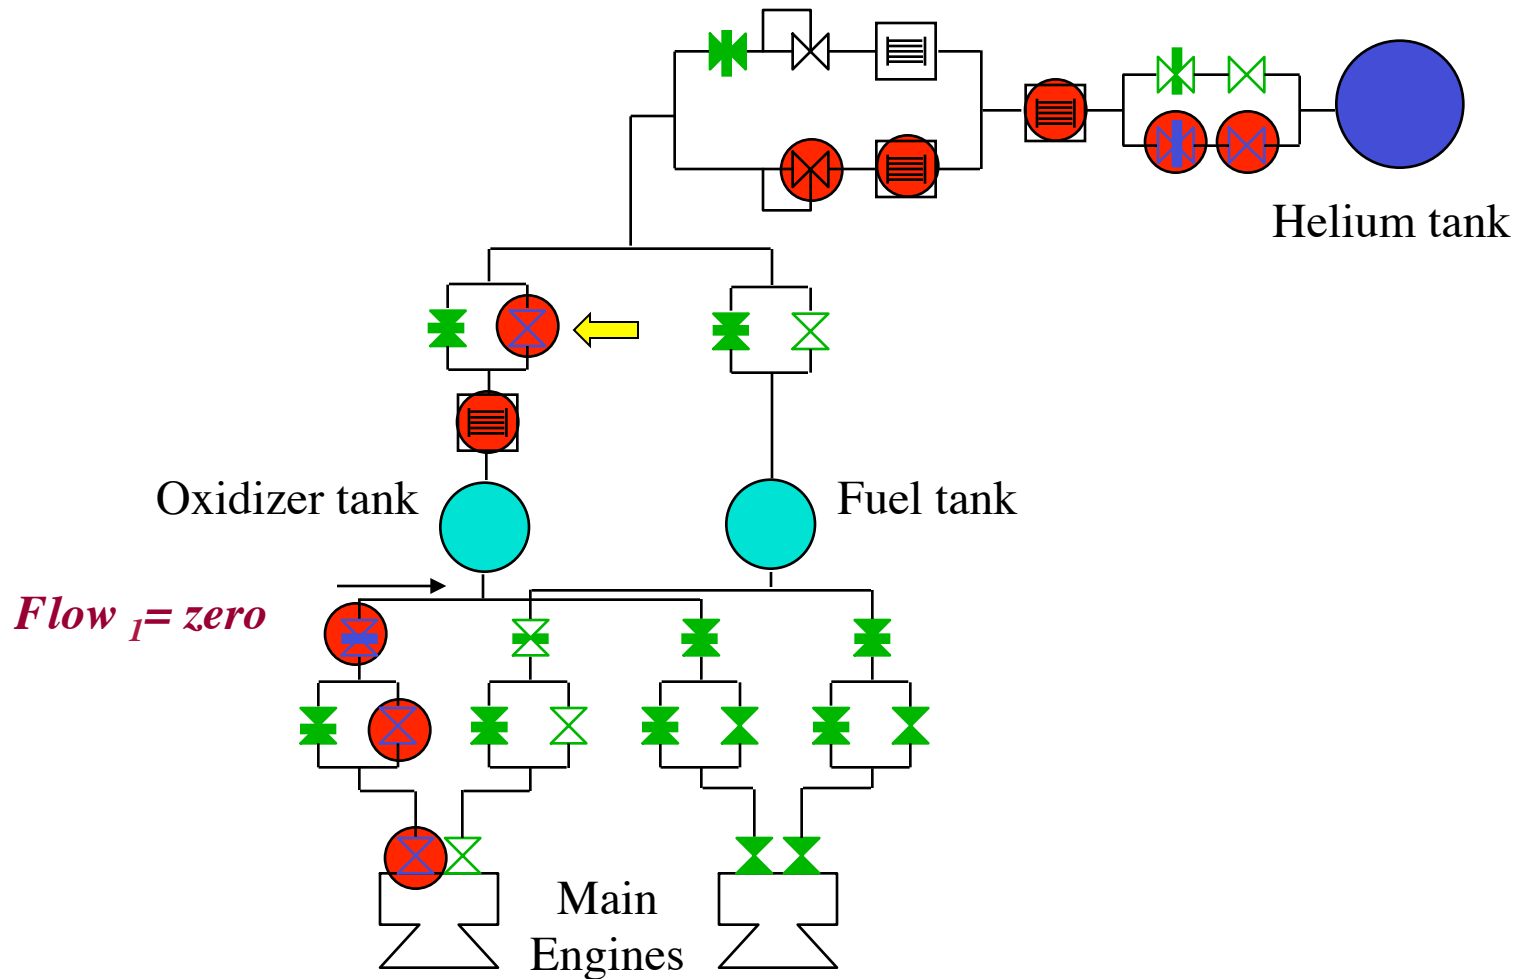
It is most likely that all components are okay.

Isolate Conflicting Information



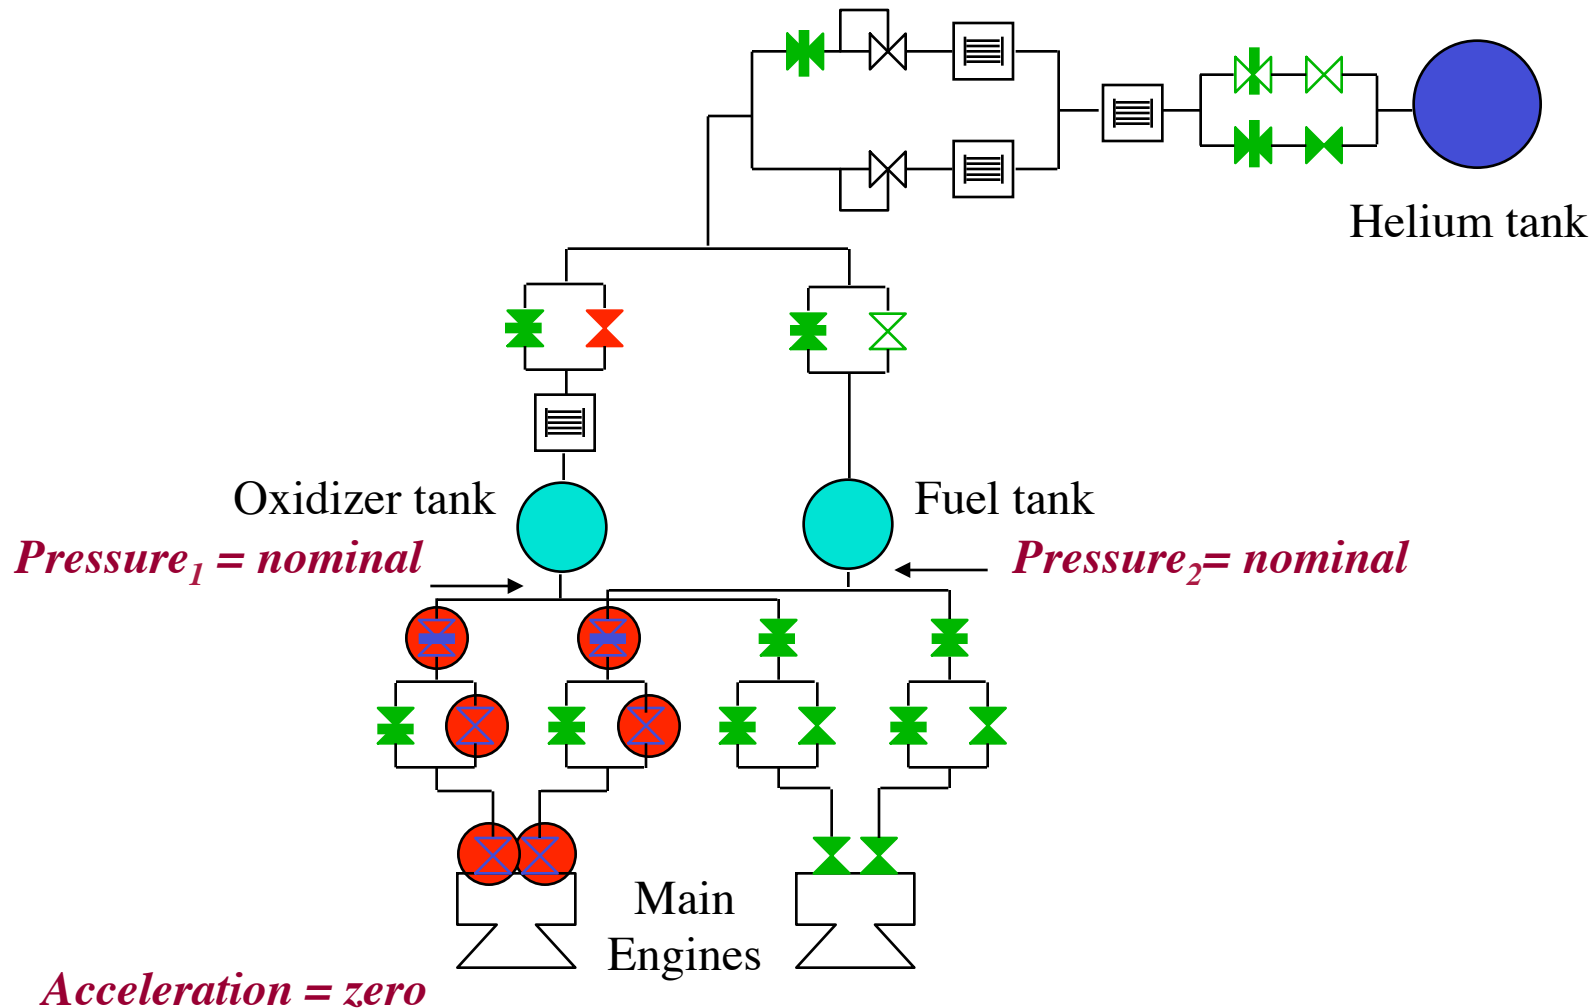
The red component modes *conflict* with the model and observations.

Leap to the Next Most Likely Hypothesis that Resolves the Conflict



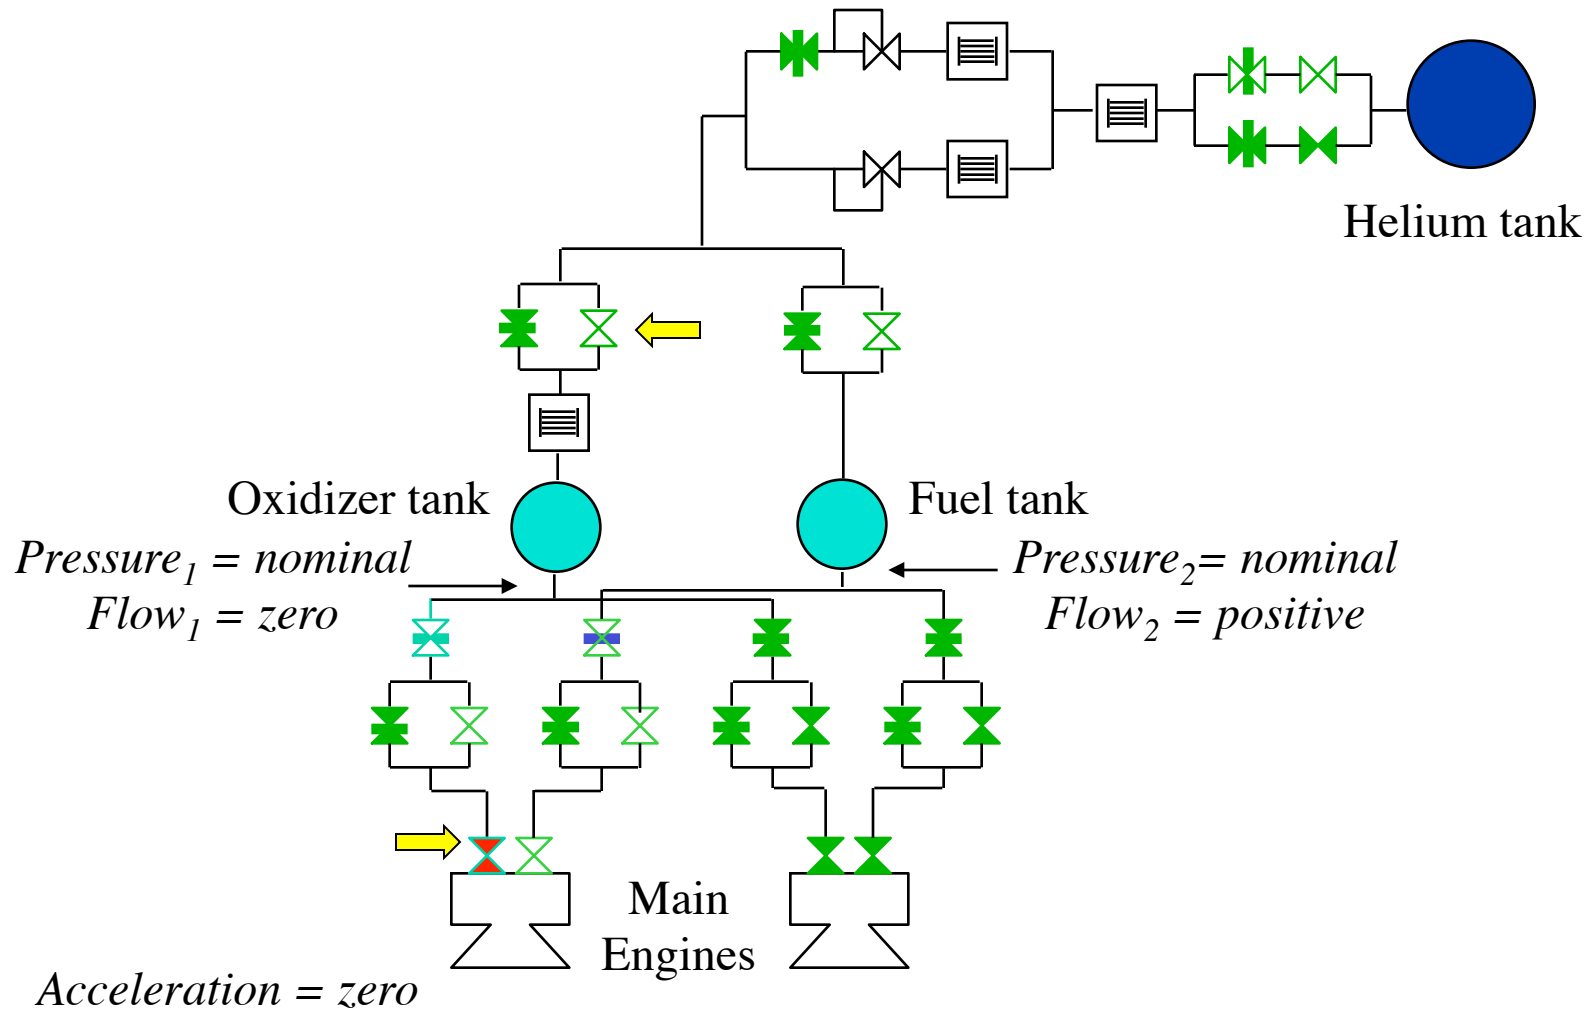
The next hypothesis must remove the conflict.

New Hypothesis Exposes Additional Conflicts



Another conflict, try removing both.

Final Hypothesis Resolves all Conflicts

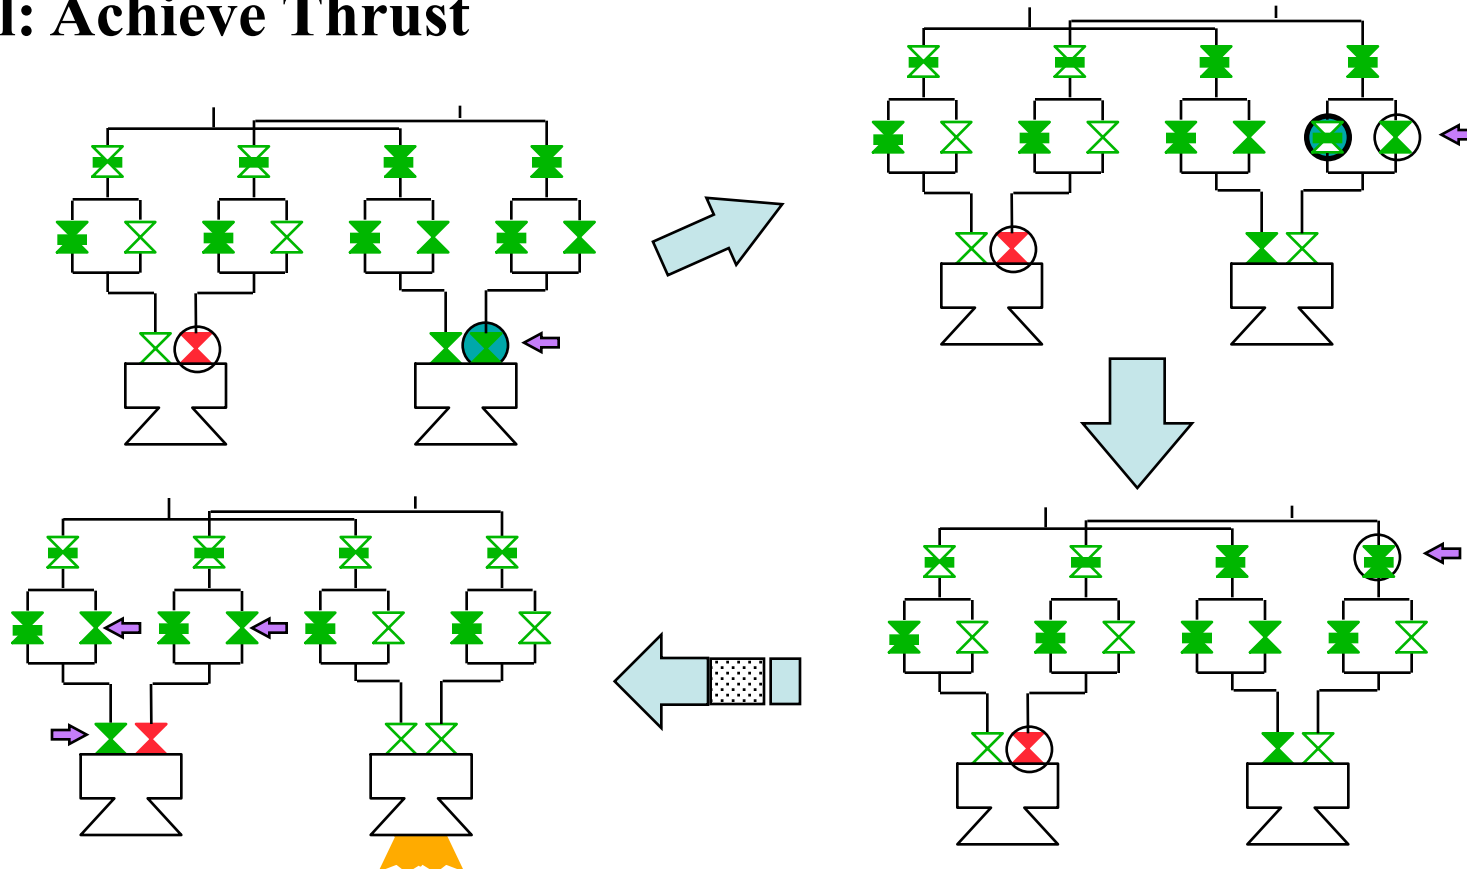


Implementation: Optimal CSPs and Conflict-directed A*.

Reconfiguring Modes using Conflict-directed A*

$\arg \max R_t(Y)$
 s.t. $\Psi(X,Y)$ entails $G(X,Y)$
 s.t. $\Psi(X,Y)$ is consistent
 Y are reachable modes


Goal: Achieve Thrust

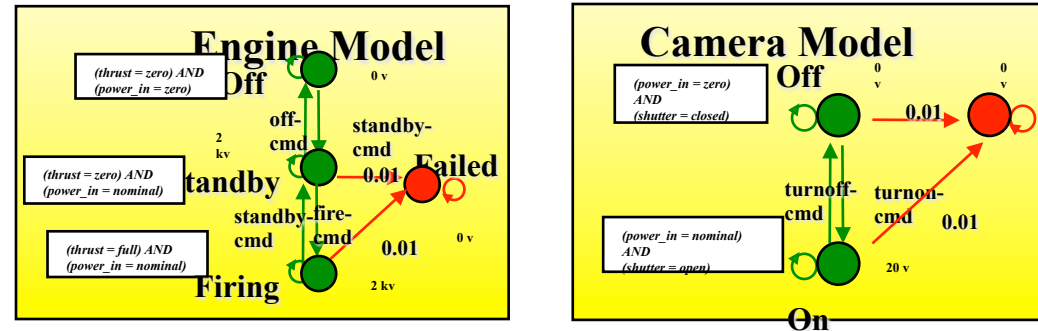


A *conflict* is a partial assignment to mode variables that prevents goal achievement (entails the negation of the goal).

Mode Estimation

- Estimates and kernels.
- By divide and conquer (GDE).
- Likely estimates and Conflict-directed A*.
- Mode reconfiguration.
- Estimating probabilistic (hybrid) constraint automata.






- PCCA encode an HMM compactly through **concurrency** and **constraints**.
- Mode estimation abstracts **state** to **modes**.

Assumes:

- **Transitions** only permitted on **modes**.
- **Transitions** are **conditionally independent**.
- For each time t , all consistent assignments are **equally likely**.

Mode Estimation as Belief State Update for Concurrent PCA




- Assigns a value to each variable (e.g., 3,000 vars).
- Consistent with all state constraints (e.g., 12,000).

- A set of concurrent transitions, one per automata (e.g., 80).
- Previous & Next states consistent with source & target of transitions

1. Infer most likely mode trajectories.
2. Infer distribution on likely mode assignments.

Best-first Trajectory Enumeration (BFTE):



[Williams and Nayak, AAAI-96][Kurien and Nayak, AAAI-00][Williams et al., IEEE '03]



Deep Space One

- Best-first State Enumeration (BFSE):

[Martin, Williams and Ingham, AAAI-05]

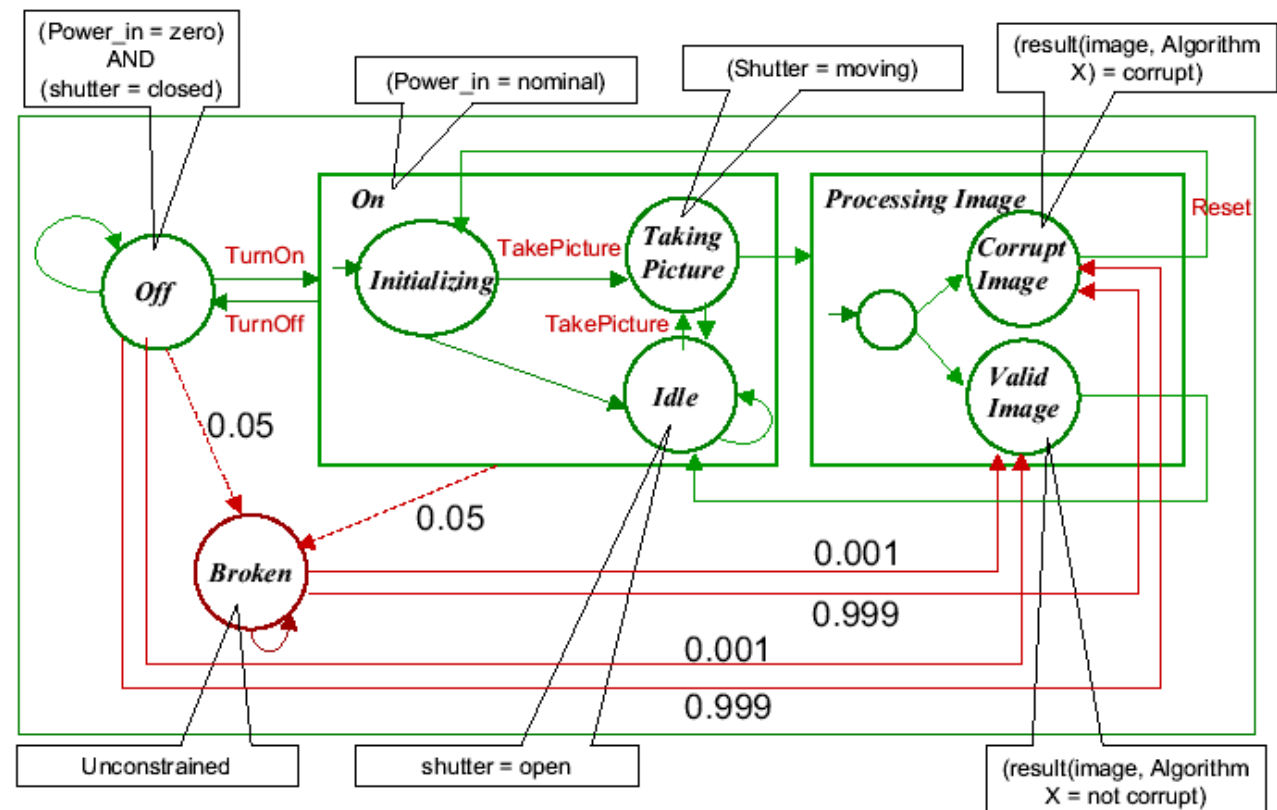


Earth Observing One


- Improves accuracy through compact encoding.
- Accuracy improves runtime!

Monitoring Complex Hardware / Software Systems through Hierarchical Probabilistic Constraint Automata


Example: Rover Image Acquisition



Estimating Hybrid States from Noisy Observations




1. Free
2. Driving
3. Holding
4. Backdriven



Hybrid probabilistic constraint automata

- Stochastic transitions between discrete modes

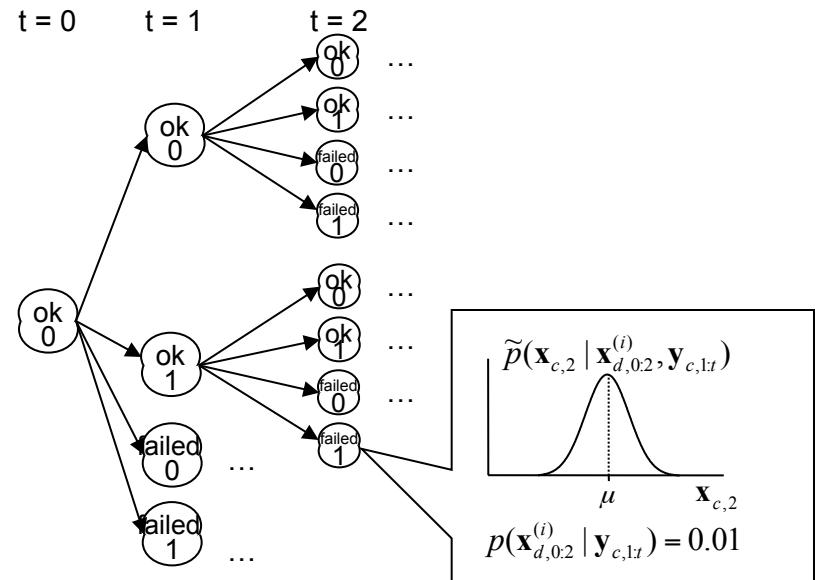


- Different continuous dynamics for each mode

$$\begin{aligned} \mathbf{x}_{t+1} &= f_{\text{nominal}}(\mathbf{x}_t, \mathbf{u}_t, t) + \mathbf{v}_{\text{nominal}}(t) \\ \mathbf{y}_{t+1} &= g_{\text{nominal}}(\mathbf{x}_{t+1}, \mathbf{u}_t) + \omega_{\text{nominal}}(t) \end{aligned}$$

$$\begin{aligned} \mathbf{x}_{t+1} &= f_{\text{failed}}(\mathbf{x}_t, \mathbf{u}_t, t) + \mathbf{v}_{\text{failed}}(t) \\ \mathbf{y}_{t+1} &= g_{\text{failed}}(\mathbf{x}_{t+1}, \mathbf{u}_t) + \omega_{\text{failed}}(t) \end{aligned}$$

Kalman Filters Track Subset of Trajectories



[Blackmore, Funiak, Williams AAAAI 05]

- Robust, Goal-directed Execution
- Plan Dispatching
- Diagnosis and Mode Estimation
- Plan Generation

Planning

Based on slides from David Smith, NASA Ames.

Find:

program of actions that achieves the objective.

Planning

Based on slides from David Smith, NASA Ames.

Find:

program of actions that achieves the objective.



partially-ordered set of actions.

typically unconditional.

no loops.



Goals.

Paradigms

Classical planning,
(STRIPS, operator-based, first-principles)
“generative.”

Hierarchical Task Network planning,
“practical” planning.

MDP & POMDP planning,
planning under uncertainty.

Classical Problem Statement

Propositions:

P_i

Initial Conditions:


P_1

P_2

P_3

P_4

Operators:




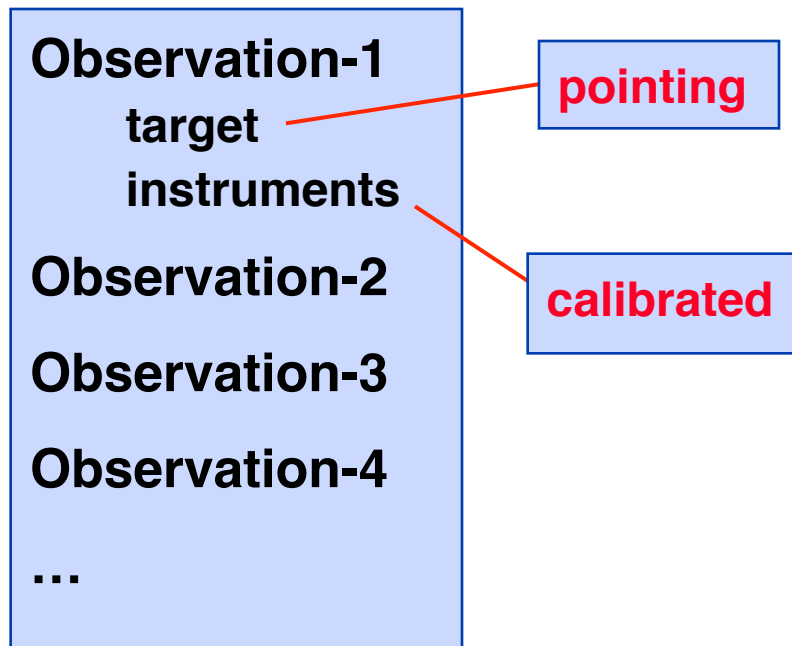
Goals:

$Goal_1$

$Goal_2$

$Goal_3$

Simple Spacecraft Problem




Propositions: Target Pointed To, Camera Calibrated?, Has Image?
Operators: Calibrate, Turn to Y, and Take Image.

Example

Init

p_c

Actions



Goal

I_A

Propositions: Target Pointed To, Camera Calibrated?, Has Image?
Operators: Calibrate, Turn to Y, and Take Image.

Planning Domain Description Language (PDDL)

```
(:action TakeImage :parameters (?target, ?instr)
  :precondition (and (Status ?instr Calibrated)
                    (Pointing ?target))
  :effect       (Image ?target))
```

```
(:action Calibrate :parameters (?instrument)
  :precondition (and (Status ?instr On)
                    (Calibration-Target ?target),
                    (Pointing ?target))
  :effect       (and (not (Status ?instr On))
                    (Status ?instr Calibrated)))
```


```
(:action Turn :parameters (?target)
  :precondition (and (Pointing ?direction)
                    ?direction ≠ ?target)
  :effect       (and (not (Pointing ?direction)
                    (Pointing ?target)))
```

By convention,
parameters start with
“?”, as in ?var.

Planning Paradigms


- From Goals
 - Goal Regression
 - (SNLP, UCPOP, Burton, Europa, Aspen, ...)

Partial order plan




- From Initial State.
 - Heuristic Forward Search
 - (FF, HSP, Colin ...)

Total order plan



- By Solving Constraints.
 - Plan Graphs
 - (SatPlan, Blackbox, Kongming ...)



Proposition Init State Action Time 1 Proposition Time 1 Action Time 2

Continuously replanning as a human helps and hinders. Planner: heuristic forward search.



Based on slides by Dave Smith, NASA Ames

Assumptions of Classic Planning

TakeImage (?target, ?instr):

Pre: **Status**(?instr, Calibrated),

Pointing(?target)

Eff: **Image**(?target)

Calibrate (?instrument):

Pre: **Status**(?instr, On),

Calibration-Target(?target),

Pointing(?target)

Eff: \neg **Status**(?instr, On),

Status(?instr, Calibrated)

Turn (?target):

Pre: **Pointing**(?direction),

?direction \neq ?target

Eff: \neg **Pointing**(?direction),

Pointing(?target)

- Atomic time,
- Agent is omniscient (no sensing necessary),
- Agent is sole cause of change,
- Actions have deterministic effects, and
- No indirect effects.

The Simple Spacecraft Revisited: Complications

Observation-1

priority

time window

target

instruments


duration

Observation-2

Observation-3

Observation-4


...



Objective:

maximize science return.

The Simple Spacecraft Revisited: Complications



More Expressive Planners Include

Time

Resources

Utility

Uncertainty

Hidden State

Indirect Control

Reasoning methods:

STNs or CSPs,

LPs or CSPs,

MDPs or MILPs,

HMMs or BNs,


HMMs or OCSPs,

LPs or RPs.




MAPGEN: Automated Science Planning for MER

NASA Ames



EUROPA
Automated
Planning System




Flight Rules
Engineering
Resource
Constraints
Science
Navigation
DSN/Telcom

Task	Start Time	End Time	Priority	Status
1. OBSERVE	00:00:00	00:05:00	High	Completed
2. MOVE	00:05:00	00:10:00	Medium	In Progress
3. OBSERVE	00:10:00	00:15:00	High	Pending
4. MOVE	00:15:00	00:20:00	Medium	Pending
5. OBSERVE	00:20:00	00:25:00	High	Pending
6. MOVE	00:25:00	00:30:00	Medium	Pending
7. OBSERVE	00:30:00	00:35:00	High	Pending
8. MOVE	00:35:00	00:40:00	Medium	Pending
9. OBSERVE	00:40:00	00:45:00	High	Pending
10. MOVE	00:45:00	00:50:00	Medium	Pending

Sequence
Build



Science Team

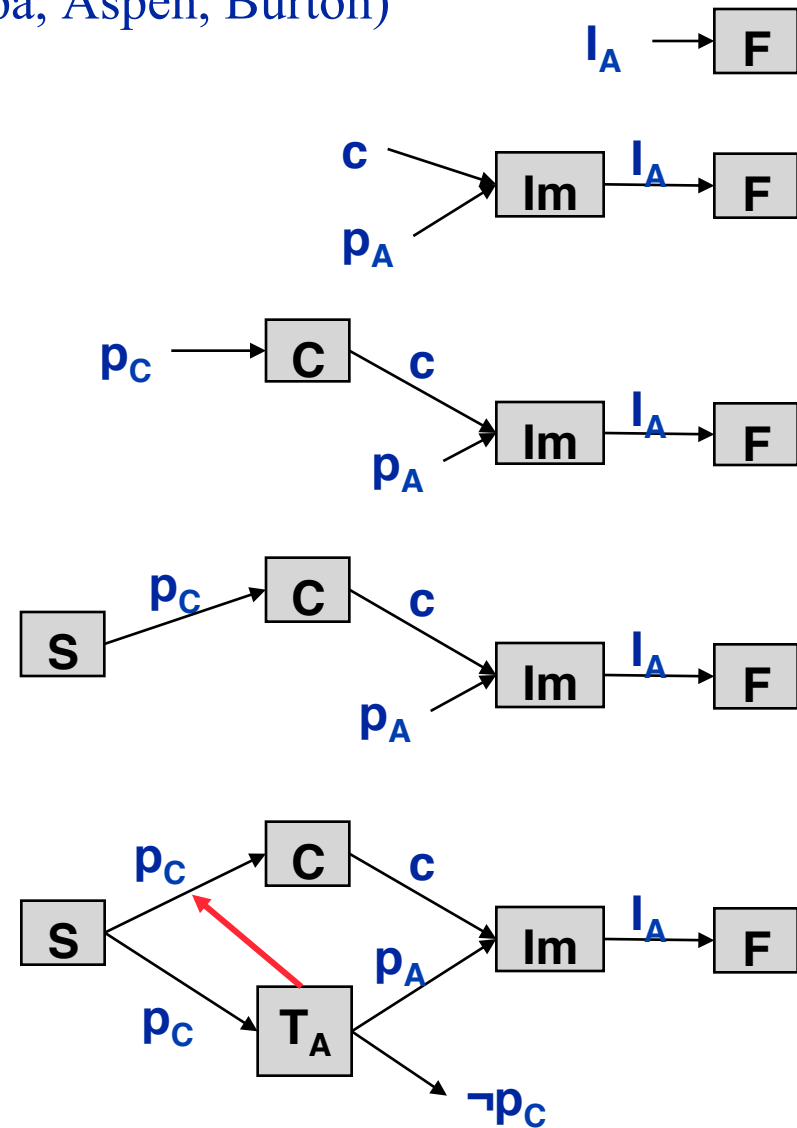



Planning Back from Goals:

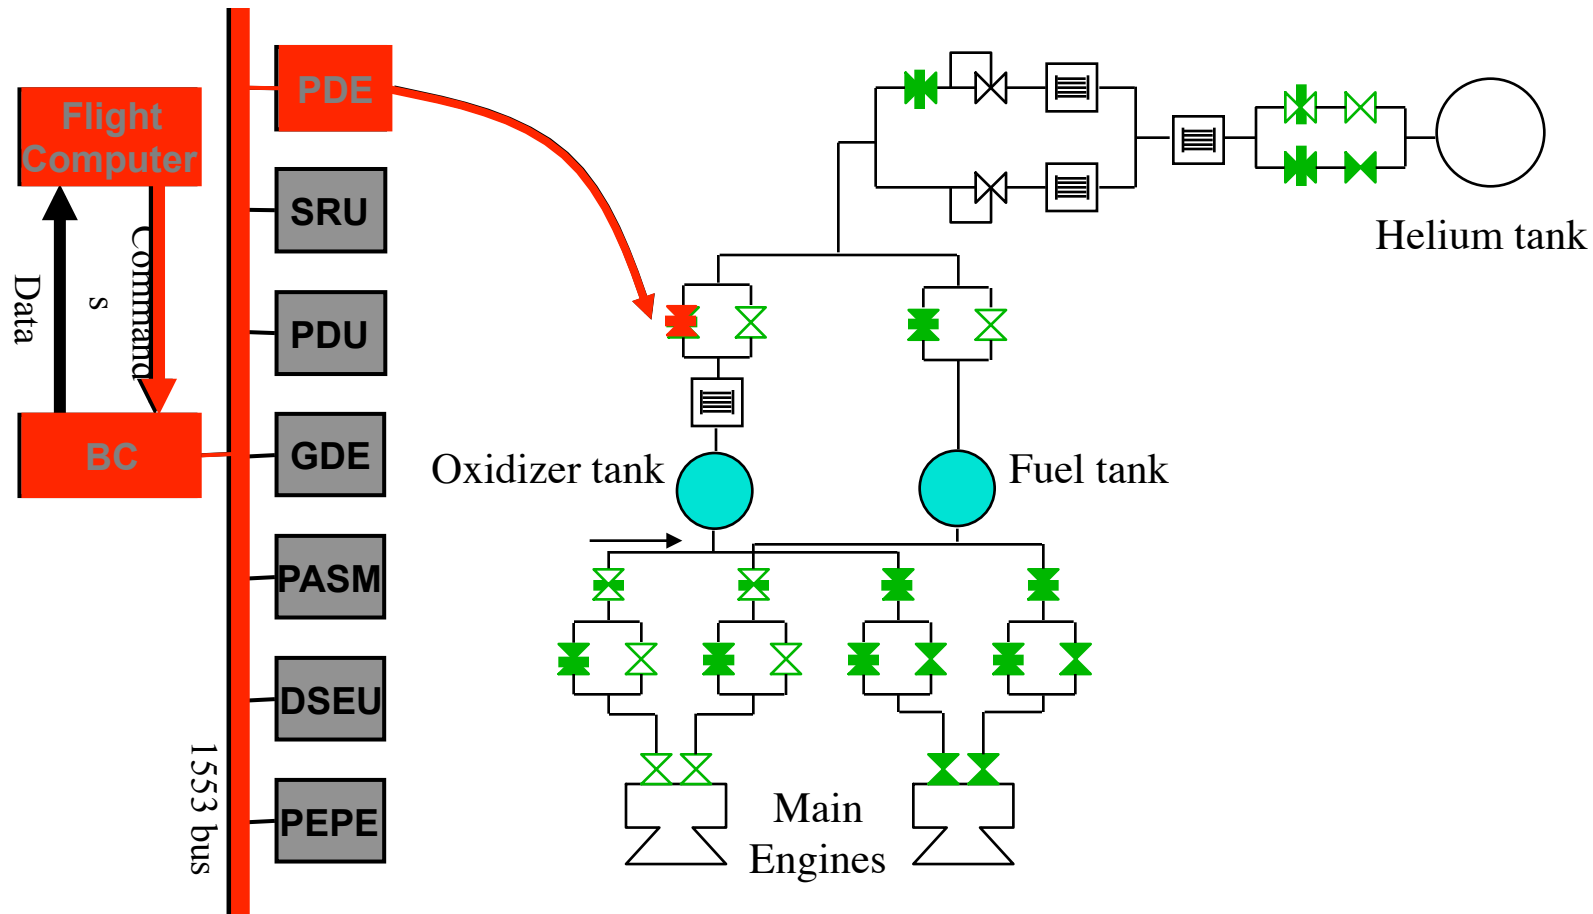
Partial Order Causal Link Planning

(SNLP, UCPOP, Europa, Aspen, Burton)

1. Select an open condition;
2. Choose an op that can achieve it:
Link to an existing instance or
Add a new instance;
3. Resolve threats.




Burton: Reactive Planning with Indirect Effects



**When causal interactions are acyclic,
and actions are reversible,
The first action in the plan can be generated in \sim constant time.**

[Williams & Nayak, IJCAI 97]




00:00 Go to x_1, y_1
00:20 Go to x_2, y_2
00:40 Go to x_3, y_3
...
04:10 Go to x_n, y_n

Command script

Commands

Plant



*“Explore **mapping region** for at least 100m, then explore **bloom region** for at least 50m, then return to **pickup region**. Avoid obstacles at all times.”*


Qualitative State Plan

Model-based Executive

Observations


Commands

Plant



Optimal

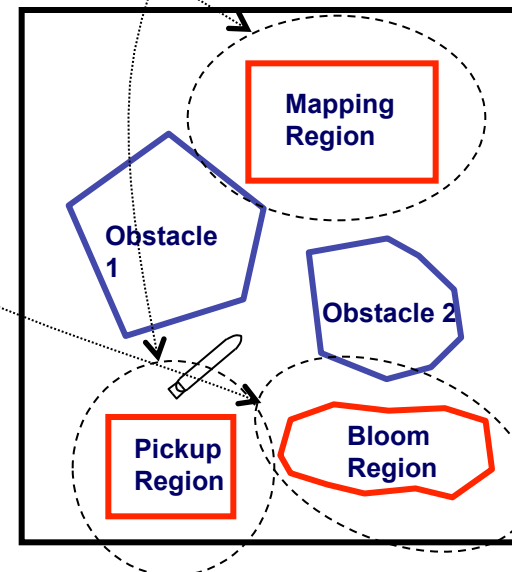
A **qualitative state plan** is a **plan of activities** that specifies desired states rather than executable actions and provides **flexibility** in **state** and **time**.







“Explore bloom region for between 50 and 70 seconds. Afterwards, explore mapping region for between 40m and 50m. End in the pickup region. Avoid obstacles at all times. Complete the mission within 300m.”

Approach: Frame as Model-Predictive Control using Mixed Logic or Integer / Linear Programming.

Leaute & Williams, AAAI 05



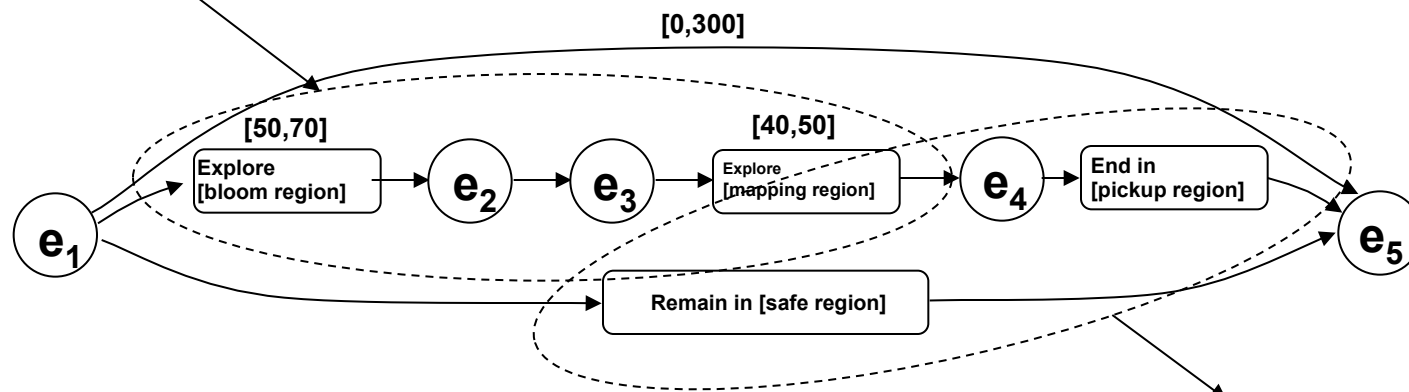
Sulu Depth Navigation for Bathymetric Mapping – Jan. 23rd, 08



94

Problem: Managing Risk within Mission-Guidelines

1. Science Activities



2. Safety Activities


Chance constraints:

1. $p(\text{Remain in [bloom region] fails OR Remain in [mapping region] fails}) < 10\%$.


2. $p(\text{End in [goal region] fails OR Remain in [safe region] fails}) < 1\%$.

Instance of **Chance-constrained Model-based Programming**.

(a) Uniform width safety margin



(b) Uneven width safety margin



(b) results in better path → takes risk when most beneficial

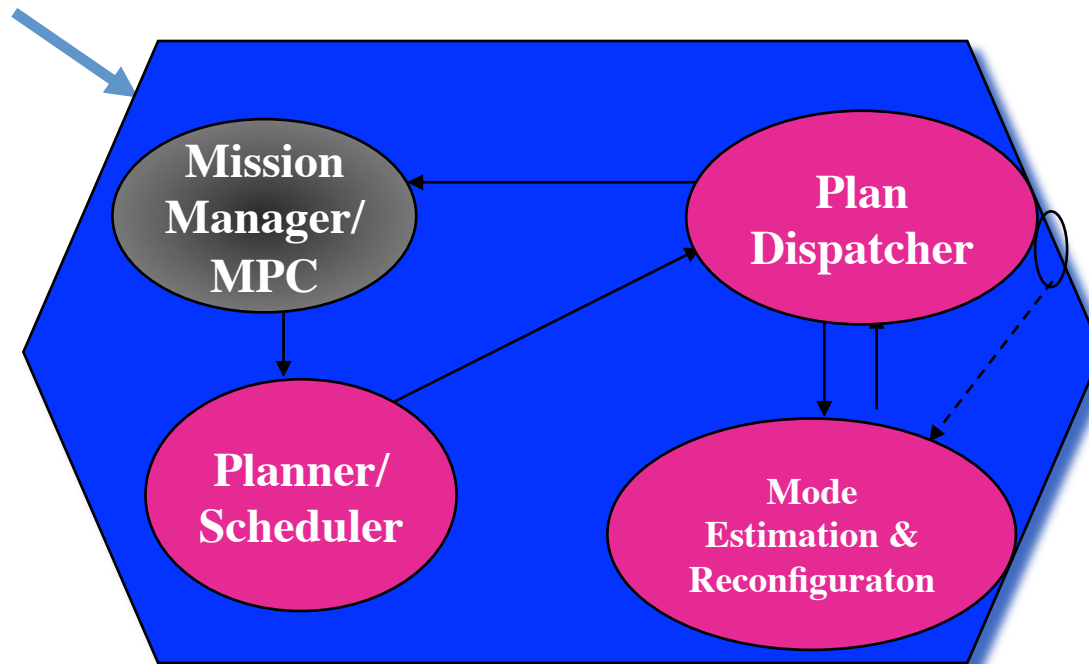
Problem: How do we find the best safety margin?

[Ono & Williams, AAAI 08]

Model-based Executives

1. Commanded through time evolved goals.
2. Reasons from commonsense models.
3. Closes loop on goals.
4. Model-based programs specify goals and models.

Goals



For More: Go to MIT Open Course Ware:

- 16.410 Principles of Autonomy and Decision Making
- 16.412 Cognitive Robotics

QUESTIONS?