Certification for Adaptive Controls

John Rushby

Computer Science Laboratory
SRI International
Menlo Park CA USA

Classical Control

- We have a plant that we wish to control
- The desired state is given by the input i
- The actual state is observed as the output o
- The controller looks at the difference (or error) between these, and their history, and computes a control input c that will bring the error to 0

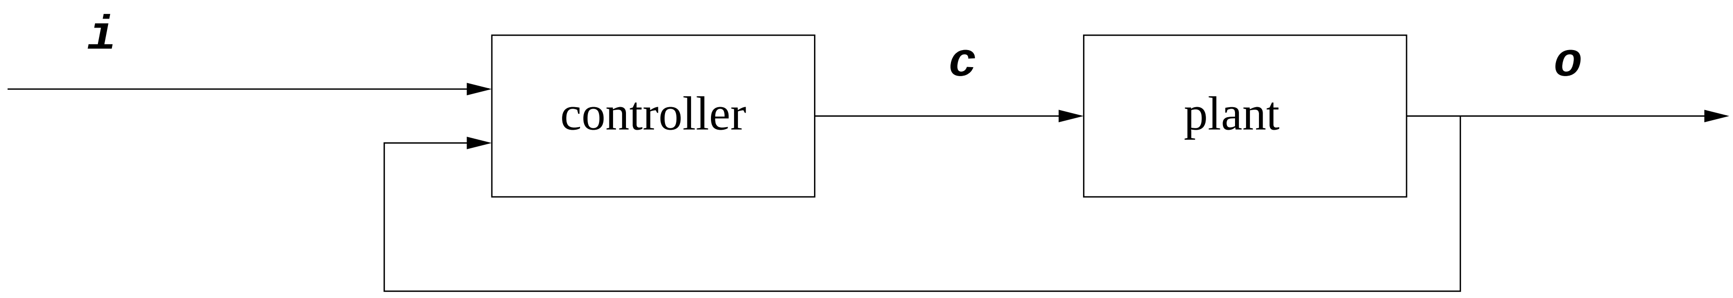

Certification for Classical Control (1)

- The controller should have nice properties
  - Always smoothly bring the error to 0
  - With no overshoot, or thumping etc.
- Classical treatment: stability
- CS treatment: Lyapunov functions
- The controller is designed wrt. some model of the plant
- The properties are verified wrt. this model
- Model might not be completely accurate for this airplane
  - Actuator performance
  - Rivets, dents, paint, dirt on the surfaces
  - Weight, and weight distribution etc.
- So you show the controller is fairly robust wrt. these
- Phase and gain margins are used for this

Certification for Classical Control (2)

- The controller is implemented as software
- DO-178B provides guidelines for this
- Basically, code must implement exactly what is specified
- Should be deterministic, traceable to requirements etc.
- The control algorithm has to be safe
- Its implementation must be correct
- All validated by flight test

Adaptive Control

- The controller is designed wrt. some model of the plant
- If the model is inaccurate, or the plant changes, we could try to adapt the controller by adjusting its internal parameters
- The adaptation mechanism typically performs some kind of machine learning
- Problem is, we now have two components sharing the control task and they could get in each other's way

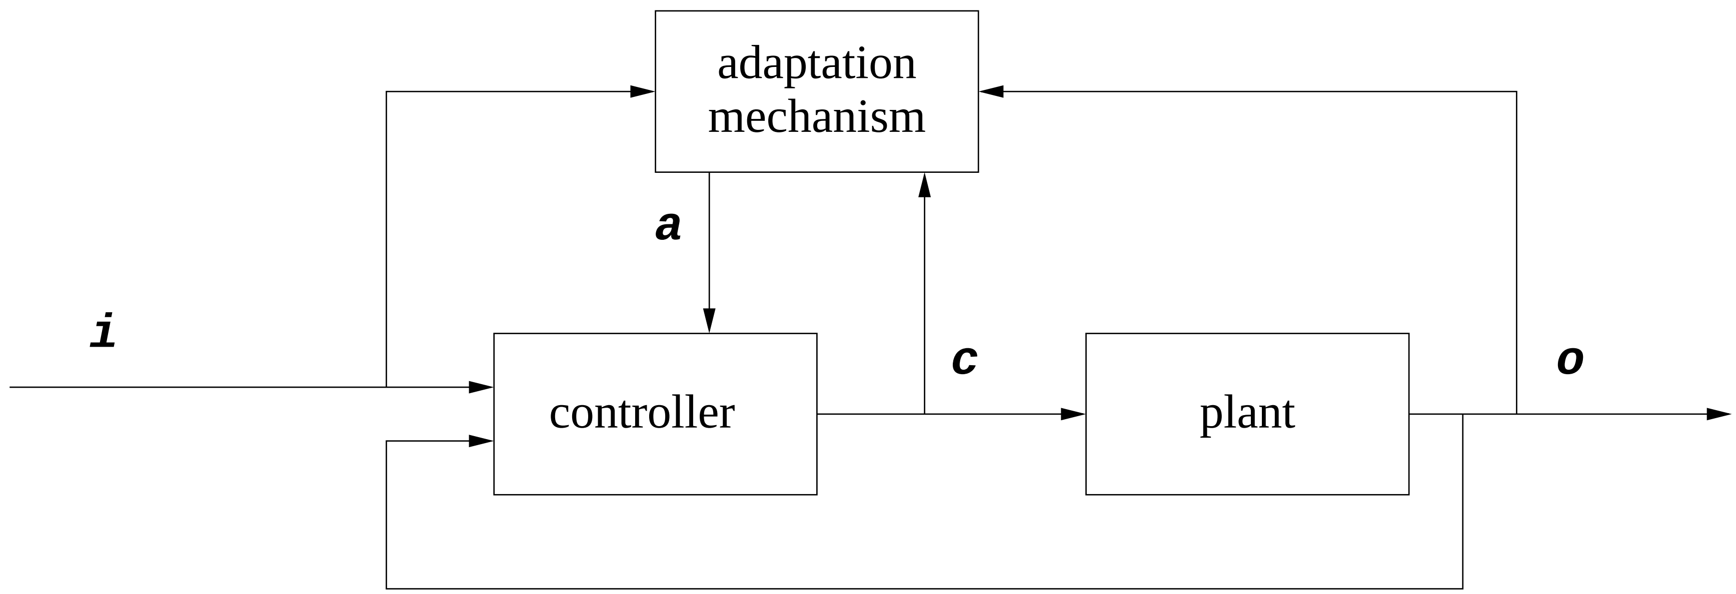

Direct Model Reference Adaptive Control (MRAC)

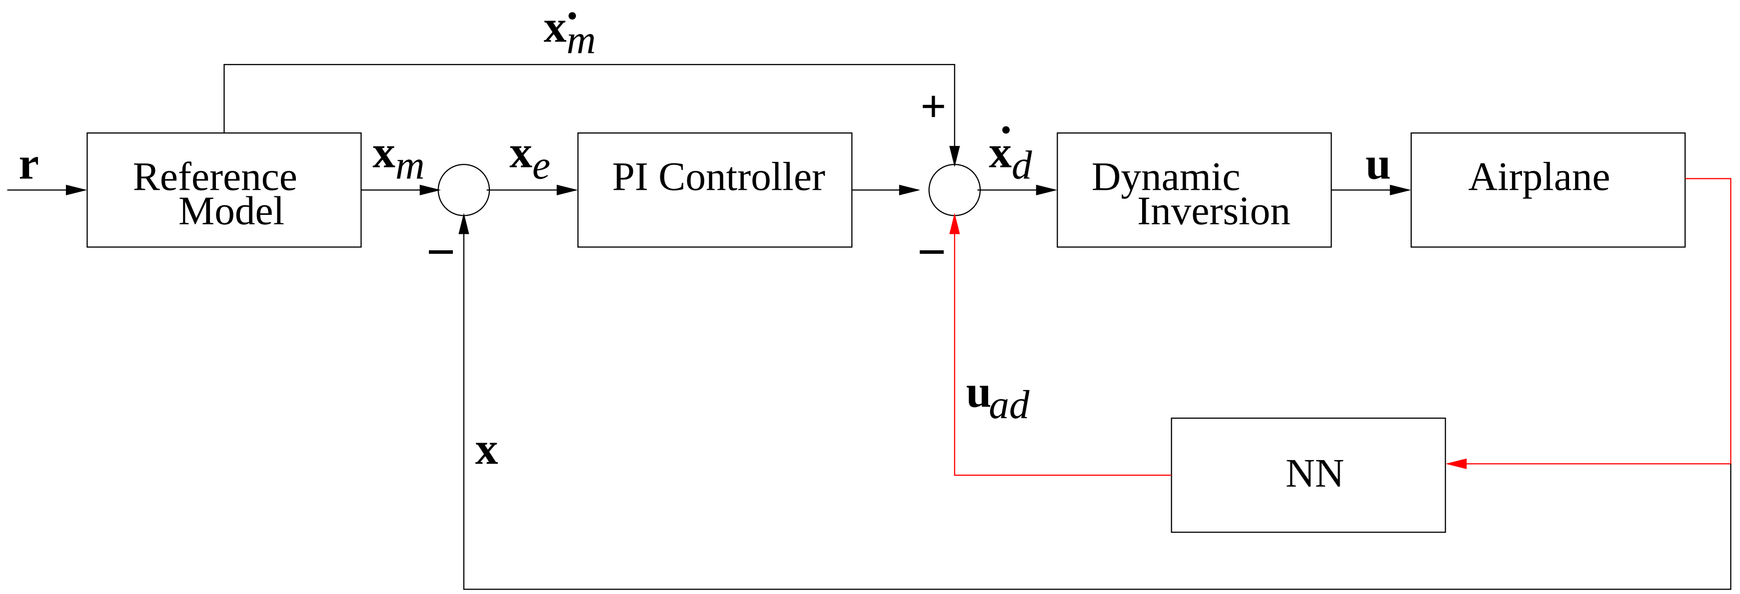

NN is Neural Net

Indirect Model Reference Adaptive Control (MRAC)

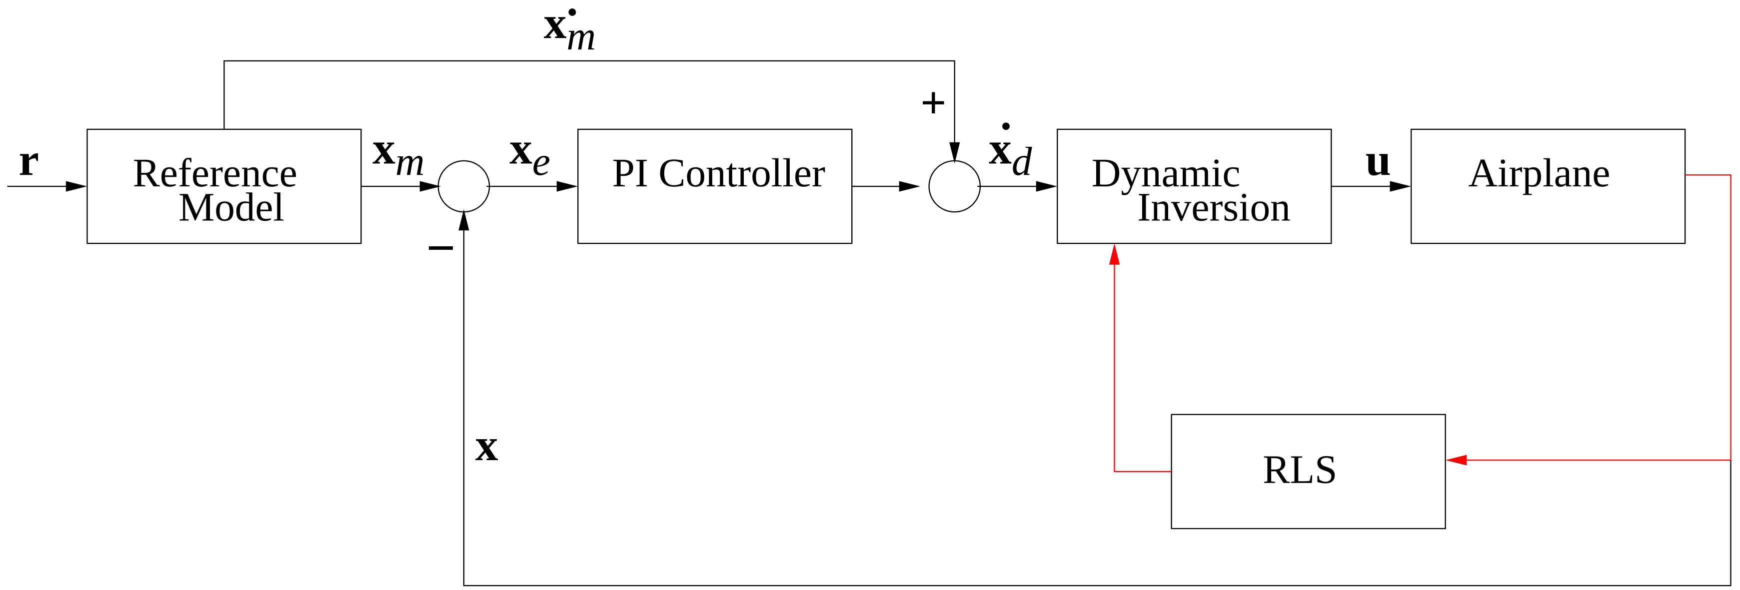

RLS is Recursive Least Squares

Motivation For Adaptive Control

- The plane suffers damage or extreme failures
- The plane is in an unexpected attitude (e.g., inverted)
- Improve efficiency by optimizing trim for this plane
- Reduce gain scheduling
  - Different conditions require different controllers
    low, slow, heavy vs. high, fast, light
  - Usually same controller, different parameters (gains)
  - Often as many as 30 different gain schedules
  - Each as to be certified, must move/blend between them
- To provide lifetime employment for control engineers

Certification Difficulties for Adaptive Control

- Bad experience: X15 crash and death of its pilot due to adaptive control
- Intellectual complexity: we have two components sharing the control task and they could get in each other's way
  - Could be overcome with advanced control theory
- Departure from certification guidelines: we cannot verify stability etc. wrt. a model (the model is learned at runtime)
  - Could be overcome with advanced control theory
- Departure from certification guidelines: it's not a deterministic implementation of a fixed algorithm
- So what can we do?

Certification of Adaptive Controls For Damaged Aircraft (1)

- No matter how the control system works, there must be some assumptions about the nature/extent of damage underlying its operation and hence its certification
- Within the assumptions it is conceptually a standard certification problem
- Outside the assumptions we provide weak assurance (simulations) that the adaptation does OK
- It is almost impossible to state useful damage assumptions
  - Any part of any one flight surface (did it come off cleanly or is it flapping?)
  - Any one actuator (would do better to build in more fault tolerance)
- So assumption may as well be that the airplane is undamaged

Certification for Damaged Aircraft (2)

- Two plausible architectures
  - Classical control for the undamaged case
  - Adaptive control for the damaged case
  - Automatic/manual switchover
- *versus*
  - Adaptive controller for both cases
  - It's a single controller but we only certify its behavior for the undamaged case
- Automated switchover is impossible to certify in my view, and pilots would never use a manual one
- Full time adaptive control runs into the certification difficulties mentioned before
- But there's a way out

Certification for Damaged Aircraft (3)

- Lui Sha's Simplex Architecture
- A certified controller provides a protection envelope
- An untrusted controller operates inside this envelope
- Monitor a Lyapunov function (works like a guardrail)
- When the system bumps against the guardrail,
  the certified controller takes over
- It's (sort of) known how to certify and analyze the reliability of monitored systems like this
- In the damaged case, we remove the guardrail
  (but then the same switchover problem as before)

Certification for Damaged Aircraft (4)

- Seems we really do need to verify an adaptive controller
- Ashish Tiwari has mechanically verified properties about indirect MRAC using Lyapunov functions
- One approach: assume/guarantee
    - Assuming the adaptation is small, the classical part of the controller guarantees stability
    - And assuming classical part operates nicely, the adaptation is guaranteed to be small
- Could consider a variant where a monitor constrains the adaptation to be small, remove the monitor for "Hail Mary"
- We still have the problem that the implementation is not deterministic and does not comply with DO-178B

Certification of Adaptive Control
To Reduce Gain Scheduling and Improve Trim

- Here the Simplex Architecture could work well
- Use crude but safe classical controllers to provide the protection envelope
    - Could have many fewer gain schedules, since the controllers merely need to be safe, not good
- An adaptive controller then operates in the protected envelope of the classical controllers
- This is quite attractive: the crude classical controllers should be less expensive to develop and certify than traditional ones, yet we get the benefit of adaptive control

Discussion

- Proponents of adaptive control often cite the Sioux City DC-10 (controlled by differential engine thrust following loss of hydraulics), and Pittsburgh 737 (rudder hardover) crashes
  - In both these cases, a better airplane is the preferred solution
- They also cite loss of control accidents resulting from upsets and unusual attitudes
  - Not clear you need to tinker with primary controls here
  - Want an outer loop that knows acrobatic maneuvers
- So I don't buy these motivations for adaptive control
- Adaptive control within the protection envelope of a conventional controler (i.e., simplex architecture) is attractive for improving trim and reducing gain scheduling
- Could switch off the protection for "Hail Mary" situations