Recent Advances in Robust Control Theory

[*new adaptive control methods that improve robustness*]

Michael G. Safonov
University of Southern California

Outline

- Part I: Concepts and Issues
- Part II: Theory
- Part III: Design Studies
- Part IV: Conclusions

Outline

- Part I: Concepts and Issues
- Part II: Theory
- Part III: Design Studies
- Part IV: Conclusions

An Achilles heel of modern system theory has been the habit of 'proof by assumption'

Theorists typically give insufficient attention to the possibility of future observations which may be at odds with assumptions.

Background: Unfalsified Control ONLY 3 elements

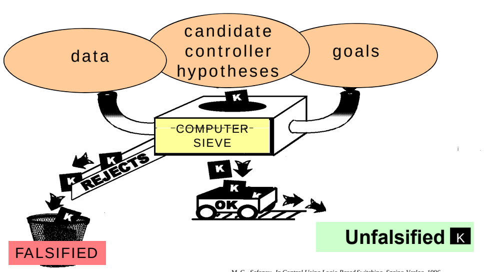

M. G. Safonov. In Control Using Logic-Based Switching, Spring-Verlag, 1996.

Control Design Bottleneck

- We have analytic tools for controlling models

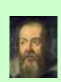

- Need DATA-DRIVEN analytic tools needed to close the design loop at the experimental validation stage

We need DATA-DRIVEN analytic design tools

Two ways to learn & adapt: Observation vs. Introspection

Reality is what we observe. | Reality is an ideal, observable only through noisy sensors.

vs.

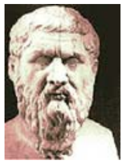

Galileo: open-eyed
"data-driven"
MODELS approximate Observed Data
'Curve-Fitting'

Plato: introspective
"assumption driven"
Data approximates Unobserved TRUTH
'Probabilistic Estimation'

- Traditional control theory ('Platonic'): 
  - contains many assumptions about the plant.
  - some assumptions are unrealistic.
- Unfalsified control theory ('Galilean'): 
  - eliminates hypotheses that are not consistent with evolving experiment data.

Observation vs. Introspection in adaptive control

"Unfalsified: Data Driven"

- goals are modest:
  - no guaranteed predictions of future stability
  - just <u>consistency</u> of goals, decisions and data
- <u>**no**</u> troublesome assumptions, parsimonious formulation:
  - DATA
  - GOALS
  - DECISIONS
- **Fast, reliable, exact**

"Traditional: Assumption Driven"

- goals are ambitious:
  - guaranteed future stability
  - Cost & estimation convergence
- **many <u>troublesome assumptions</u>**
  - "the 'true' plant is in model set",
  - "noise I.I.D."
  - "bounds on parameters, probabilities"
  - …, "linear time-invariant, minimum-phase plant, order < N"
- **Slow, quasi-static, approximate**

Remarkably, some leading "assumption driven" control theorists have held that observed DATA inconsistent with ASSUMPTIONS should be ignored (cf. M. Gevers et al., "Model Validation in Closed-Loop", ACC, San Diego, 1999)

The Behavioral Approach

- Let the data speak...
- **<u>Unfalsify</u>** (validate) models and controllers against <u>hard criteria</u>:
  - Choose criteria expressible directly in terms of observed data (sensor outputs, actuator inputs)
  - Avoid criteria that that rely on noise model assumptions or other prior beliefs

---

UNFALSIFIED HYPOTHESES
:

•The ability of each candidate hypothesis to meet the performance goal is tested directly against evolving real-time measurement data.

* If a hypothesis is a controller, it need not be in

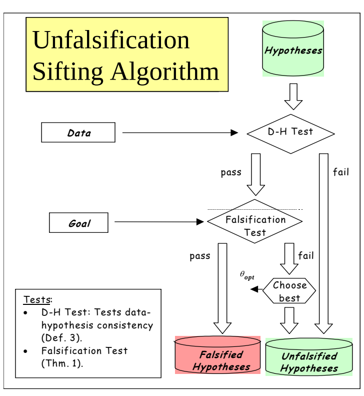

Brugarolas & Safonov, *CCA/CACSD* '99

Impact on SYSID

- Statisticians treat prior probabilities as part of model to be validated/unfalsified
  - validation via one-σ 'confidence intervals'
  - 'Platonic' probability becomes 'Galilean'
- Ljung (e.g., CDC '97) proposed 'confidence interval' reinterpretation of gaussian SYSID
  - model validation/unfalsification

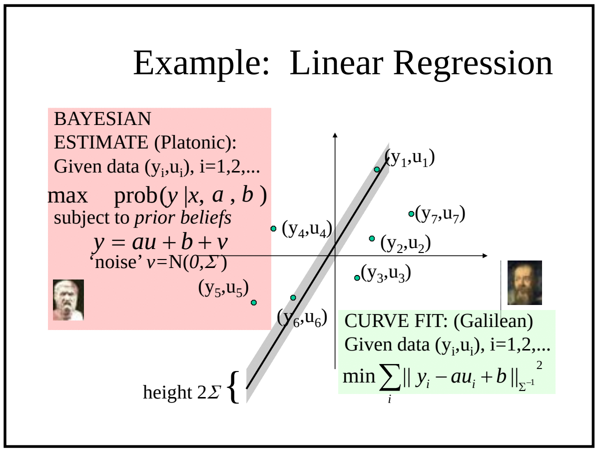

Example: Linear Regression
BAYESIAN ESTIMATE (Platonic):
Given data (y_i,u_i), i=1,2,...
max prob(y | x, a , b)
subject to *prior beliefs*
$y = au + b + v$
'noise' v=N(0,Σ)
(y_1,u_1)
(y_7,u_7)
(y_4,u_4)
(y_2,u_2)
(y_3,u_3)
(y_5,u_5)
(y_6,u_6)
CURVE FIT: (Galilean)
Given data (y_i,u_i), i=1,2,...
$\min \sum_i \| y_i - au_i + b \|^2_{\Sigma^{-1}}$
height 2Σ

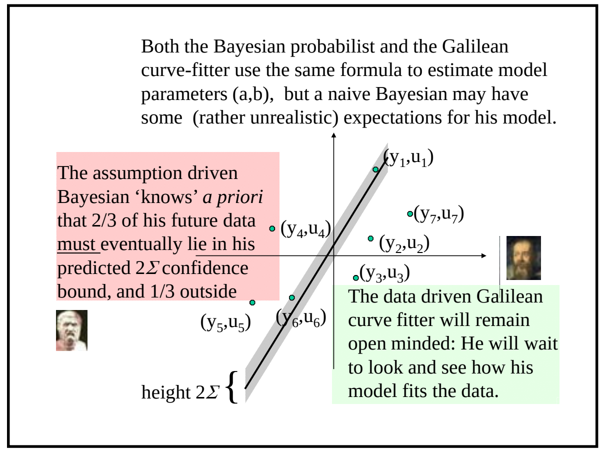

Both the Bayesian probabilist and the Galilean curve-fitter use the same formula to estimate model parameters (a,b), but a naive Bayesian may have some (rather unrealistic) expectations for his model.
The assumption driven Bayesian 'knows' *a priori* that 2/3 of his future data must eventually lie in his predicted 2Σ confidence bound, and 1/3 outside
(y_1,u_1)
(y_7,u_7)
(y_4,u_4)
(y_2,u_2)
(y_3,u_3)
(y_5,u_5)
(y_6,u_6)
The data driven Galilean curve fitter will remain open minded: He will wait to look and see how his model fits the data.
height 2Σ

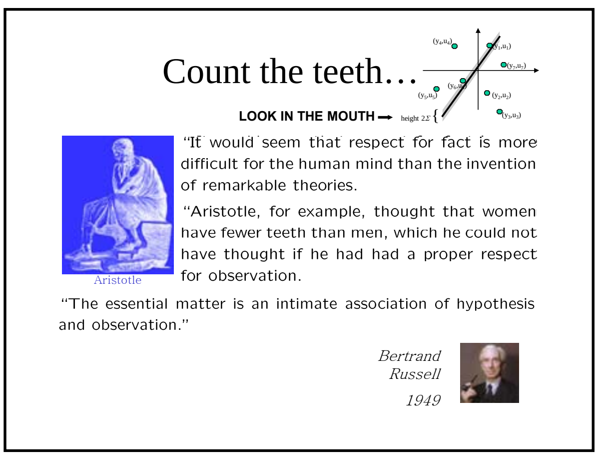

Count the teeth…
LOOK IN THE MOUTH
"It would seem that respect for fact is more difficult for the human mind than the invention of remarkable theories.
"Aristotle, for example, thought that women have fewer teeth than men, which he could not have thought if he had had a proper respect for observation.
Aristotle
"The essential matter is an intimate association of hypothesis and observation."
Bertrand Russell
1949

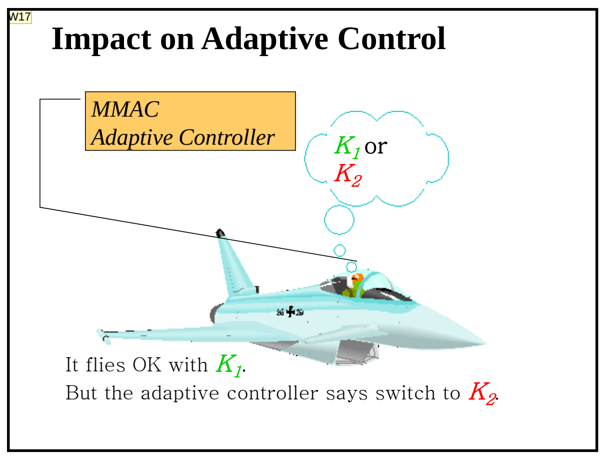

W17
Impact on Adaptive Control
MMAC
Adaptive Controller
K_1 or
K_2
It flies OK with K_1.
But the adaptive controller says switch to K_2.

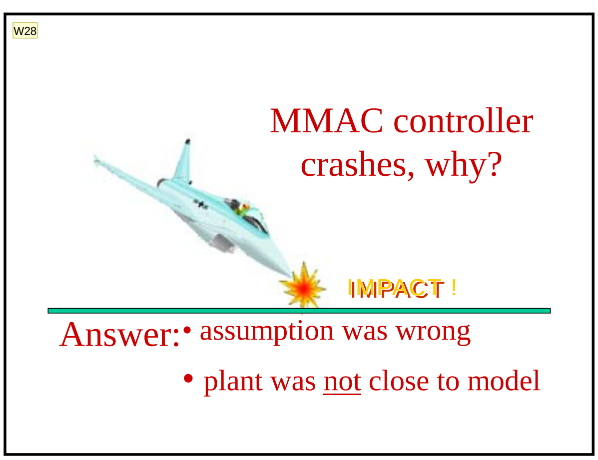

W28
MMAC controller crashes, why?
IMPACT !
Answer:
• assumption was wrong
• plant was not close to model

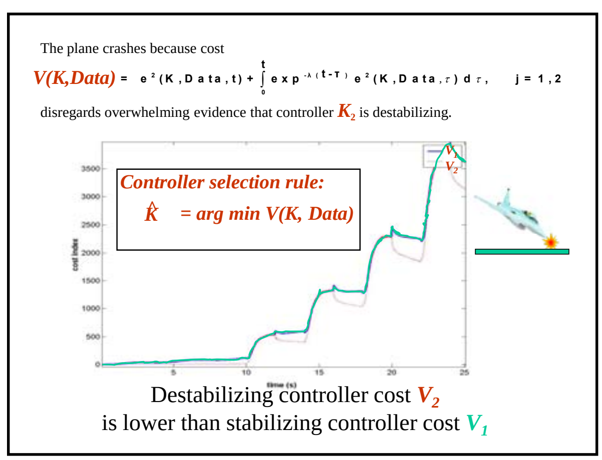

The plane crashes because cost
$V(K,Data) = e^2(K,Data,t) + \int_0^t \exp^{-\lambda(t-\tau)} e^2(K,Data,\tau)\,d\tau, \quad j = 1,2$
disregards overwhelming evidence that controller K_2 is destabilizing.
Controller selection rule:
$\hat{K}$ = arg min V(K, Data)
V_1
V_2
Destabilizing controller cost V_2
is lower than stabilizing controller cost V_1

Slide 16

W17 solution1:find a new monitoring signals(actually it is a new cost funtion,used to solutions by others: change the measure the goodness of controllers(change cost function),hope to be more compatible with the objective of safe switching. [2001', Anderson,multiple model adaptive control with safe switching], but by [2001' Brozenec&Safonov, controller validation], this cost function is not good;

why? and the aircraft crashes finally.

WANG, 5/1/2004

Slide 17

W28 using my anti-example,show that adaptive control with multiple plant model may fail and cause aircraft to crash.

mention that this adpative control method with multiple plant models is used by narendra and morse, and it crashes with the closeness assumption removed.

WANG, 5/13/2004

W1

In life, as in football, you won't go far unless you know where the goal posts are.

— Arnold Glasow —

Goals of adaptive control?

- **Adaptive Control goals**: Ensure stability and desired performance of the closed loop at all times.
- **Problem:** **Model-Mismatch Instability**
  - *If assumptions are wrong, then the adaptive loop can cause instability, even when the initial controller is stabilizing*
  - ***Need** safe adaptive control **laws that recognize** instability, and are not blinded by prior assumptions*

Previous Work: Traditional and recent adaptive methods

- Pre-1980's adaptive: Too many assumptions! (non-realistic): LTI, minimum phase; etc.
- Relaxation of assumptions over the years: alternative methods, deal with poor, unreliable models
- Recently, **switching adaptive schemes**:

❈ **Indirect** (MMAC; Identifier-based supervisory; Localization switching)

❈ **Direct** switching schemes (Fu & Barmish, Mårtensson, Safonov) **reduce *a priori* information** need.

Previous Work without assumptions on plant

- Some pre-routed switching schemes essentially had no assumptions on the plant at all (**Fu & Barmish '86, Mårtensson '85**).
  - –Not very practical - use extensive dense search through parameter space: slow convergence.
  - – Only finitely many ***K***'s allowed
- **Stefanovic-Wang-Paul-Safonov '04** solve safe adaptive control via unfalsified adaptive concept
  - – Fast, practical, but still only finitely many ***K***'s

---

Outline

- Part I: Concepts and Issues
- Part II: Theory
- Part III: Design Studies
- Part IV: Conclusions

Adaptive Control

Every adaptive control law minimizes a cost

Controller selection rule:

$$\hat{K} = \arg\min_{K} V(K, Data)$$

N may be ***infinite***

Continously parameterized ***K***

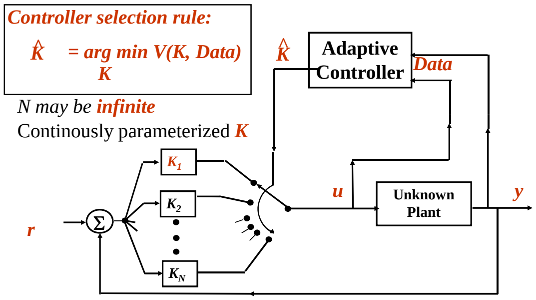

Slide 19

W1 chang "in life'' to 'in adaptive control'
WANG, 5/13/2004

Every adaptive control law minimizes a cost

Controller selection rule:

$$\hat{K} = \arg\min_{K} V(K, Data)$$

Choose cost $V(K, Data)$ so that destabilizing controller K_2 cost is greater than stabilizing K_1

Convergence Theorem

(Wang, Stefanovic, Safonov ACC '04)

Assume

1. $V(K,Data)$ is <u>cost detectable</u>
2. **stabilization is feasible,**

then any adaptive controller of the form

$$\hat{K} = \arg\min_{K} V(K,Data)$$

is stabilizing, <u>irrespective of plant model mismatch</u>.

Proof: We apply the 1992 lemma of Morse, Mayne & Goodwin, but with a new L_{2e}-gain type cost $V(K,Data)$ that <u>*detects* instability without plant assumptions</u>.

Idea for proof Prior Result (continued)

Cost vs. Control Gain K, Time Snapshots

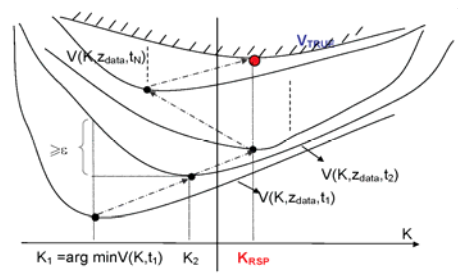

With the L_{2e}-gain type cost function, the finite upper bound condition (a) in the 1992 Morse lemma holds even for big "model mismatch"

This is our 'picture' of the Morse-Mayne-Goodwin proof of their hysteresis lemma. **Q.E.D.**

Assumption 1: Cost Detectability

Definition: If

$$V(K,Data) \to \infty \iff \left(\max \left(\|y\|_\tau^2 + \|u\|_\tau^2 \right) / \|r\|_\tau^2 \right) \to \infty$$

then we the cost function $V(K,Data)$ is **<u>cost detectable</u>**.

Key points:

- This $V(K,Data)$ is an L_{2e}-gain type cost function. Irrespective of plant assumptions, it correctly distinguishes stability vs. instability.
- Prior to Safonov & Tsao (1997), cost-detectable cost functions were not used in adaptive control – making them susceptible to model-mismatch instability.

Assumption 2: Feasibility

Definition: An adaptive control problem is said to be **<u>feasible</u>** if the candidate controller set **K** has at least one stabilizing controller.

- *This is the <u>weakest possible</u> assumption required to guarantee adaptive law $\hat{K}(t)$ converges to a stabilizing controller $K \in$ K*
- *If the initial controller works prior to closing the adaptive loop, then feasibility assumption is satisfied*

Unfalsified Cost Level $V(K,Data)$

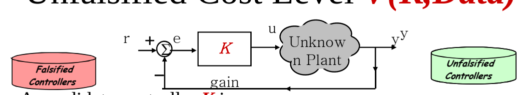

A candidate controller K is **FALSIFIED** with repect if to cost level V_o if

$V(K,Data) = \max \left(\|y\|_\tau^2 + \|u\|_\tau^2 \right) / \|\tilde{r}(K,Data)\|_\tau^2$

when the $Data:=(y,u)$ proves existence of a 'fictitious'

$$\tilde{r}(K, Data) = y + K^{-1}u$$

that would violate performance spec $V(K,Data) < V_o$ if the K were in the loop.

- Controllers may remain **UNFALSIFIED** with repect to cost level V_o until the data proves otherwise.
- Or, fading memory may be used to reinstate previously falsified controllers

Key points:

- A controller need not be in the loop to be falsified.
- Even a single data point can falsify many

Trivial Example

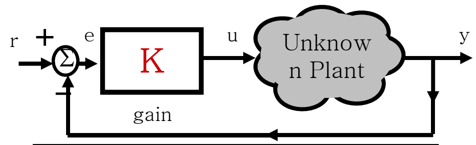

- Plant Data: at time t=0, (u,y)=(1,1)
- Candidate K's: u=Ke real gain
- Goal: |e(t)/r(t)| < V_o = 0.1 for all r(t)

Relations: e=u/K=1/K , r=y+e=y+u/K= 1+1/K

=> K is unfalsified if |1/(1+K)| < 0.1

=> unfalsified K's: K>9 or K<-11

Cost-Detectable *V*(*K*,*Data*) Solves…

Problem of <u>Safe Adaptive Control</u>:
Design adaptive controllers that never destabilize when the adaptive loop is closed (**model-mismatch instability**), with <u>no assumptions</u> on the plant.

Controller selection rule:
$\hat{K}$ = *arg min V(K, Data)*
K

N may be infinite

Continously parameterized *K*

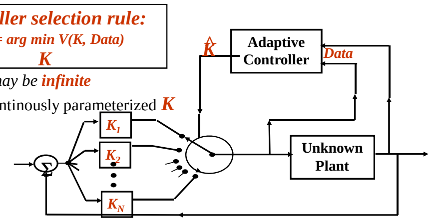

The Behavioral Approach to Adaptive Control

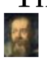

- **<u>Data Driven:</u>** Let the data speak...
  - Don't let modeling beliefs trump observation
- **<u>Unfalsify</u>** (validate) models and/or controllers against <u>hard criteria</u>:
  - Choose criteria expressible directly in terms of observed data (sensor outputs, actuator inputs)
  - Avoid criteria that that rely on "noise model" and other prior beliefs

Unfalsified Adaptive Control

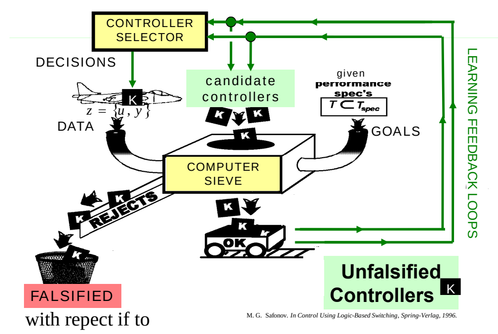

with repect if to cost level V_o

M. G. Safonov. In Control Using Logic-Based Switching, Spring-Verlag, 1996.

Unfalsification

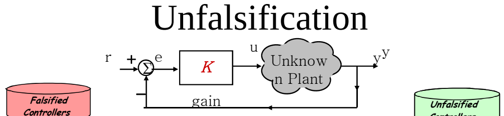

A candidate controller ***K*** is **FALSIFIED** with respect to cost level V_o when the data (y,u) proves existence of a 'fictitious'
$$\tilde{r}(K,y,u) = y + K^{-1}u$$
that would violate performance spec's if the ***K*** were in the loop.

- Controllers may remain **UNFALSIFIED** respect to cost level V_o until the data proves otherwise.
- Or, fading memory may be used to reinstate previously falsified controllers

Key points:
- A controller need not be in the loop to be falsified.
- Even a single data point can falsify many

Outline

- Part I: Concepts and Issues
- Part II: Theory
- Part III: Design Studies
- Part IV: Conclusions

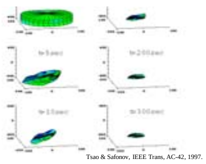

Evolution of unfalsified set

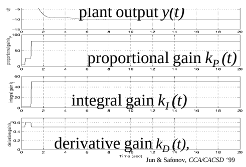

Time responses

Another Example: Missile

Brugarolas, Fromion & Safonov, ACC 1998

- Learns control gains
- Adapts quickly to compensate for damage & failures
- Superior performance

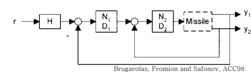

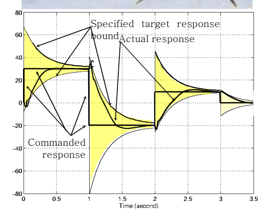

Unfalsified adaptive missile autopilot:

- discovers stabilizing control gains as it flies, nearly instantaneously
- maintains precise sure-footed

Another Example: Satellite

Boeing Satellite Systems 702

- Large flexible appendage of uncertain mode frequency (>25%)
- Relatively fast closed-loop response required
- Reaction Control System (RCS) excites flexible modes

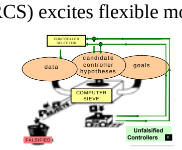

Antenna Mapping Control Problem

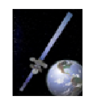

(AIAA GNC Conf., 2003)

PD Control Law: $u = C_p(r-y) - C_d\dot{y}$ (8)

Performace Spec: $\begin{cases} |T_{er}(j\omega)| < |W_1^{-1}(j\omega)| \\ |T_{ur}(j\omega)| < |W_2^{-1}(j\omega)| \end{cases}, \forall\omega$ (9)

In time domain, the performance spec is equivalent to

$$\left\| \begin{bmatrix} \hat{w}_1 * (r-y) \\ \hat{w}_2 * u \end{bmatrix} \right\|_{L_2[0,t)} \Big/ \|r\|_{L_2[0,t)} < 1,\ \forall t,\ \forall r \quad (10)$$

Antenna Mapping Control

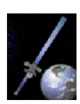

Controller Parameter Vector Switching

When the currently used controller vector is falsified (i.e. not satisfying (12))

==> switch to an element within the unfalsified controller parameter set.

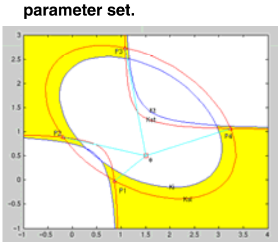

Unfalsified set is in yellow color

An Approach

1. Find the intersections of current set Kt with all previous set Kτ
3. Use a slightly larger set (left plot red-colored) to calculate intersection to avoid switching too frequently
3. Switch to the nearest interstction

Simulation Results

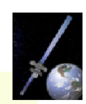

Bode plots of closed-loop S=1/(1+GK)

Dashed blue line: initial transfer function
Red line: W_1^{-1}
Solid blue lines: Ter at each switching, under $W1^{-1}$ at end
Black line: nominal non-adaptiveTer
Green line: final converged unfalsified T_{er}

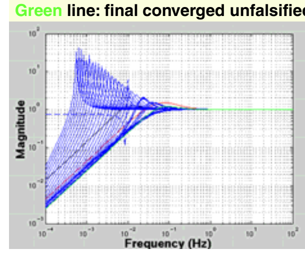

Converged unfalsified Ter lies under non-adaptive Ter and under W_1^{-1}

Simulation Results

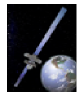

After initial transients, unfalsified controller shows smaller tracking error than the nominal non-adaptive PD controller

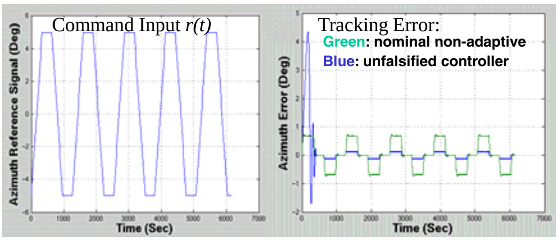

Another Example: Two Link Manipulator

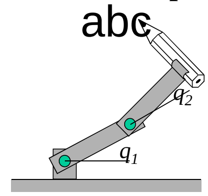

$$u_a = H(q)\ddot{q} + C(q,\dot{q})\dot{q} + g(q)$$
$$= Y(q,\dot{q},\ddot{q})\theta$$

Tsao and Safonov, *Int. J. Adaptive Contr. & Signal Proc.*, 2001

Background: Computed-Torque Robot Control

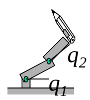

- Control Law

$$u = H(q)\left(\ddot{q}_d + 2\lambda\dot{\tilde{q}} + \lambda^2\tilde{q}\right) + C(q,\dot{q}) + g(q)$$

$\tilde{q} = q - q_d$, tracking error; q_d desired trajectory

- Performance

$$H(q)\left(\ddot{\tilde{q}} + 2\lambda\dot{\tilde{q}} + \lambda^2\tilde{q}\right) = d$$

tracking error converges exponentially fast to a region proportional to the size of disturbance *d*

Simulation

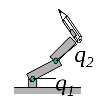

- Tip mass m moves between 2 and 20 every 0.5 sec

at m=2, θ^*=[3.34, 0.97, 1.0392, 0.6]^T
at m=20, θ^*=[30.07, 9.7, 10.392, 6]^T
$\hat{\theta}(0) = [0.5,\ 0.5,\ 0.5,\ 0.5]^T$

$\Theta(0) = \{\theta \mid -200 \le \theta_i \le 200, i = 1,\cdots,4\}$

- q_{d1}=30° (1-cos 2πt), q_{d2}=45°(1-cos 2πt)
- $q(0)=[0, 0.4]^T$, $dq/dt(0)=[0, 0]^T$

Joint 1 Tracking Error(deg)

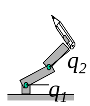

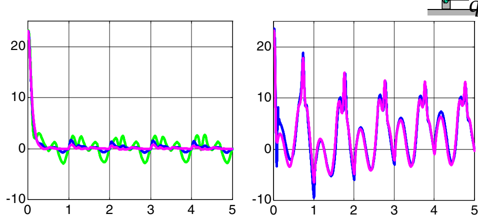

Unfalsified Slotine et al's

τ=0 τ=1/(40π) τ=1/(10π)

Another Example: Unfalsified 'PID Universal Controller'

- Unfalsified adaptive control loop stabilizes in real-time
- Unstable Plant 30 Candidate PID Controllers:

K_I =[2, 50, 100]

K_D = [.5, .6]

K_P = [5, 10, 25, 80, 110]

Example: Adaptive PID

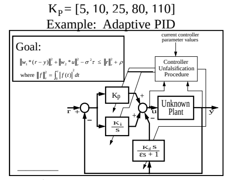

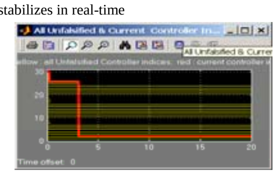

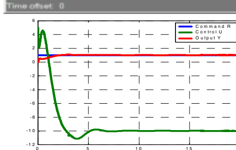

How Unfalsified Works

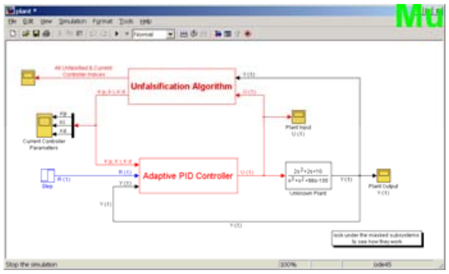

ftp://routh.usc.edu/pub/safonov/PID_Demo_for_MATLAB6.zip

Inside the Unfalsification Algorithm Block

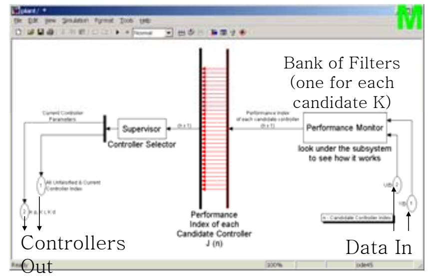

Others peoples' design studies (off-line 'run-to-run'):

- Emmanuel Collins et al. (Weigh Belt Feeder adaptive PID tuning, CDC99)
- Kosut (Semiconductor Mfg. Process run-to-run tuning, CDC98)
- Woodley, How & Kosut (ECP Torsional disk control, adaptive tuning, ACC99
- Razavi & Kurfess, Int. J ACSP, Aug. 2001

USC is the only one doing real-time unfalsified adaptive control.

Outline

- Part I: Concepts and Issues
- Part II: Theory
- Part III: Design Studies
- Part IV: Conclusions

What's Wrong:
Adaptive Control Counter-Examples

Despite "proofs" of stability and convergence, modern adaptive control methods fail to achieve goals when there is large "model mismatch".

Assumptions Fail.

Modern adaptive methods are blinded by assumptions.

53

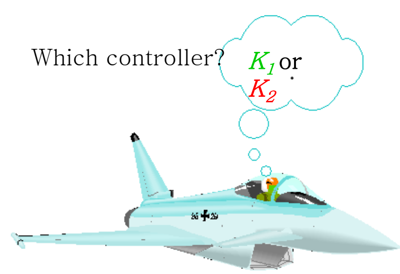

It flies OK with K_1. But the adaptive law says switch *to* K_2.

MRAC chooses c_2 and crashes finally, why?

Wrong assumption: Plant was not close to the model!

W1

In life, as in football, you won't go far unless you know where the goal posts are.

— Arnold Glasow —

THE FIX

Remember your goals:

- Switch (or tune) to a stabilizing controller, based on observation
- Don't be blinded by prior assumptions (expect model mismatch)
- **And, above all, choose a cost function that detects instability when it happens:**

You need *cost detectability*

You need an L_{2e}-gain type cost function ***V(K,Data)***

Background: Unfalsified Control
ONLY 3 elements

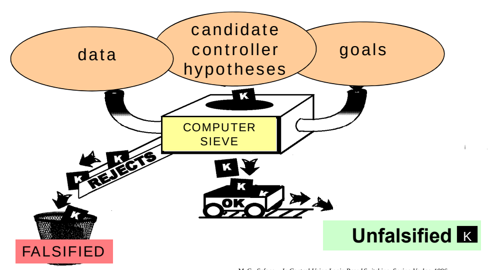

M. G. Safonov. *In Control Using Logic-Based Switching*, Spring-Verlag, 1996.

Unfalsified Adaptive Control

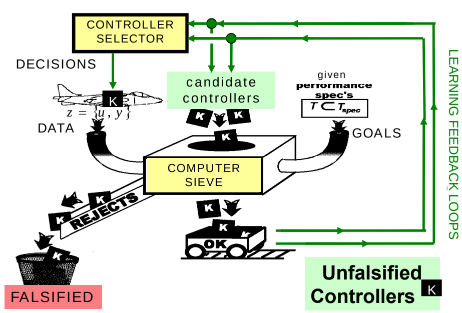

M. G. Safonov. *In Control Using Logic-Based Switching*, Spring-Verlag, 1996.

Conclusions

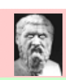

- Introspective models are not adequate for analyzing adaptation

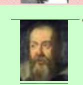

- To explain learning and adaptation, we need the open-eyed, data-driven scientific logic of unfalsification

- Unfalsified adaptive control robustifies
- It remembers where the goal posts are
- No plant assumptions needed

Slide 56

W1 chang "in life" to 'in adaptive control'
WANG, 5/13/2004

Selected References

http://routh.usc.edu

1. M. Stefanovic, A. Paul, and M. G. Safonov. Safe adaptive switching through an infinite controller set: Stability and convergence. *IFAC World Congress*, Prague, July 2005.
2. R. Wang and M. G. Safonov. Stability of unfalsified adaptive control using multiple controllers. In *American Control Conf.*, Portland, OR, June 2005.
3. M. G. Safonov. Robust control, feedback and learning. In Moheimani (ed.), *Perspectives in Robust Control*, pages 283–296. Springer-Verlag, 2001.
4. M. G. Safonov and T. C. Tsao. The unfalsified control concept and learning. *IEEE Trans. Autom. Control*, AC-42, June 1997.

Acknowledgement

Bob Kosut's mid-1980's work on time-domain model validation and identification for control played a key role in laying the foundations of this work, as did later contributions of **Jim Krause, Pramod Khargonekar, Carl Nett, Kamashwar Poolla, Roy Smith** and many others who have advanced the use of validation methods in control-oriented identification. **Tom Mitchell**'s early 1980's "candidate elimination algorithm'' for machine learning is closely related to the unfalsified control methods presented here. My students **Tom Tsao**, **Fabricio Cabral, Myungsoo Jun, Crisy Wang, A. Paul,** and **Margareta Stefanovic** who made it all work. And of course, none of this would have been possible without the superb graduate education that I received at MIT so many years ago under the guidance of first **Jan Willems** and later **Michael Athans**.