Autonomous Vehicle Safety: What Lessons Can we Learn from Aviation?

2019 International Conference on Dependable Systems & Networks
Portland, OR, USA
26 June 2019

http://2019.dsn.org/

Dr. Jay Lala
Sr. Principal Engineering Fellow
Raytheon Company
San Diego, CA 92123

The opinions expressed here are that of the author alone, and do not necessarily represent views of the Raytheon Company.

Bio and Disclaimer

Dr. Jaynarayan Lala is a Sr. Principal Engineering Fellow at the Raytheon Company.

His 43 years of experience includes 4 years as a DARPA Program Manager where he initiated ground-breaking research in intrusion-tolerant and self-healing systems, and a quarter century at Draper Lab in Cambridge, MA, where he architected fly-by-wire fault-tolerant computers for many mission- and safety-critical platforms, including Seawolf Submarine and NASA spacecraft. He made fundamental contributions to the design of safety-critical computers with systems like the Fault-Tolerant Multi-Processor (FTMP) in the seventies and the Fault-Tolerant Parallel Processor (FTPP) in the eighties, while at Draper Lab.

Jay has 4 patents, has published over 50 papers in peer reviewed journals and conferences, and made numerous invited presentations.

He was honored with the Secretary of Defense Medal for Exceptional Public Service for contributions to keeping the US networks secure in 2003.

The Proceedings of the IEEE Paper on FTMP, authored by Hopkins, Smith and Lala, was the recipient of the Jean-Claude Laprie award in 2015.

Jay is a Life Fellow of IEEE and an Associate Fellow of AIAA.

He received Doctor of Science and MS degrees from MIT in Aero & Astro in 1976 and 1973, respectively, and a Bachelor of Tech, with Honors, in Aeronautical Engineering from the Indian Institute of Tech, Bombay in 1971. He was honored in 2018 as a Distinguished Alumnus of IIT, Bombay.

The opinions expressed here are that of the author alone, and do not necessarily represent views of the Raytheon Company.

Topics

- Boeing Max 8 MCAS Case Study
  - How MCAS operates
  - What Boeing is doing to update MCAS
- Autonomous Vehicle Safety Requirements
- Commercial Aircraft Fly-By-Wire Systems
  - A case study in how to design safety-critical systems
- Potential Vehicle Control System Architectures
  - Inspired by autopilots: both digital and analog

Boeing 737 Max 8 Case Study

MCAS: Maneuvering Characteristics Augmentation System

All information for this case study is taken from public sources which are referenced. Author expresses no opinion on the design.

The Boeing 737 MAX MCAS Explained

- *The Maneuvering Characteristics Augmentation System (MCAS) is a flight control law managed by the flight control computer (FCC)*
- *Introduced on the 737 MAX to help it handle like a 737 Next Generation (NG), particularly at slow speeds and high angles of attack (AOA).*

- Sean Broderick, Guy Norris and Graham Warwick,

- Aviation Week & Space Technology, March 20, 2019

Minimize Additional Training for Legacy 737 Pilots

The Boeing 737 MAX MCAS Explained

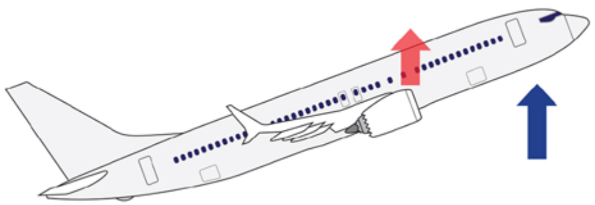

1 | Leap Engines and Pitch-up Moment

The MAX’s larger CFM Leap 1 engines create more lift at high AOA and give the aircraft a greater pitch-up moment than the CFM56-7-equipped NG. The MCAS was added as a certification requirement to minimize the handling difference between the MAX and NG.

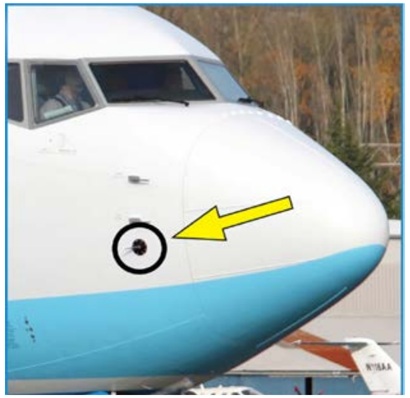

2 | MCAS Activation

The system activates when the aircraft approaches threshold AOA, or stickshaker activation, for the aircraft’s configuration and flight profile. The MAX flight-control law changes from speed trim to the MCAS because the MCAS reacts more quickly to AOA changes.

3 | Angle of Attack Vanes

The MCAS’s primary data sources are the MAX ’s two AOA sensing vanes, one on either side of the nose.

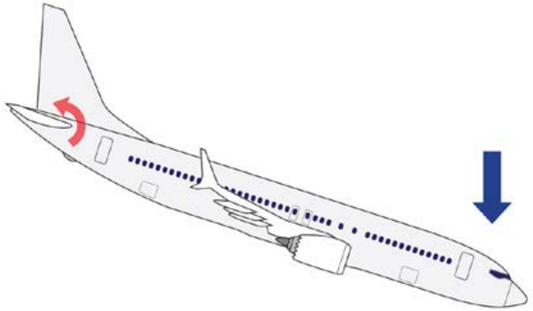

4 | Stabilizer Deflection

When threshold AOA is reached, the MCAS commands 0.27 deg. of aircraft nose-down stabilizer deflection per second for 9.3 sec.—a total of 2.5 units of trim.

MCAS Works Autonomously and with Full Authority:
<u>Is it Safety-Critical</u>?

Stabilizer Trim and Cut-Out Switches

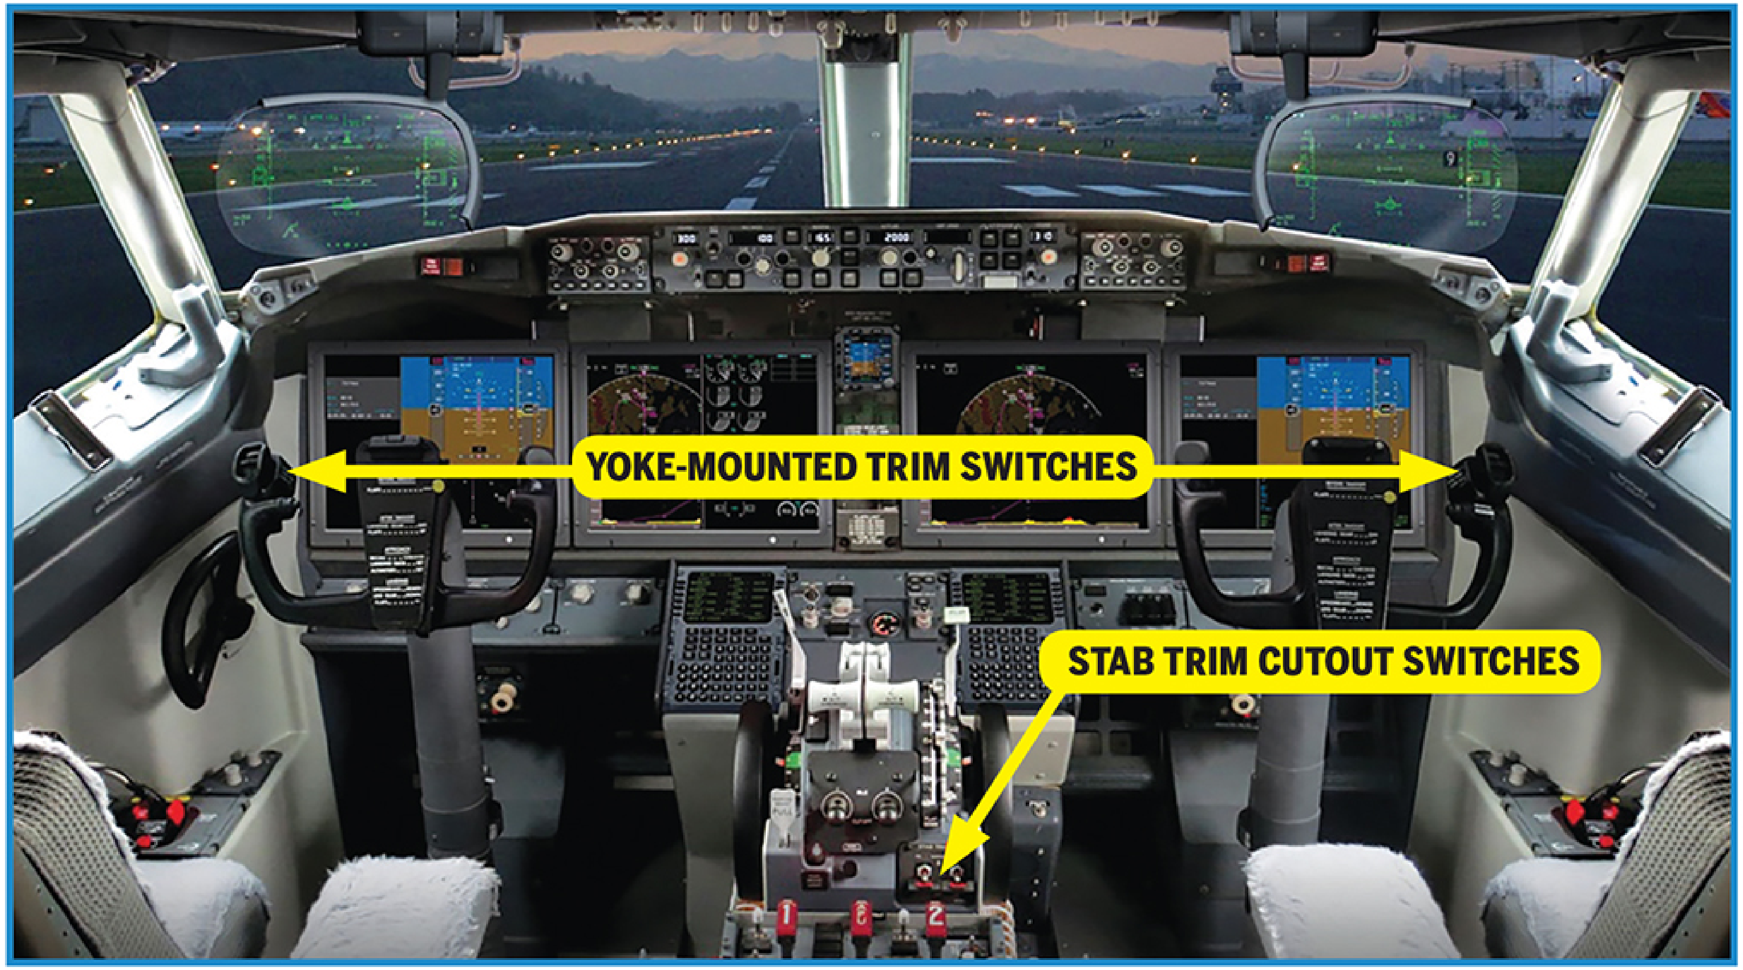

Pilots can Disengage MCAS Manually using Trim Cutout Switches

The Boeing 737 MAX MCAS Explained

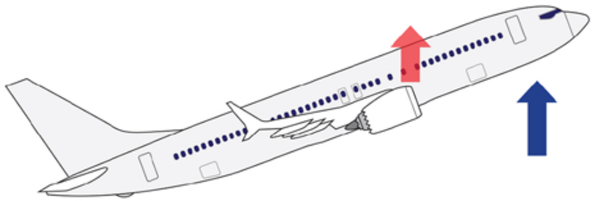

1 | Leap Engines and Pitch-up Moment

The MAX’s larger CFM Leap 1 engines create more lift at high AOA and give the aircraft a greater pitch-up moment than the CFM56-7-equipped NG. The MCAS was added as a certification requirement to minimize the handling difference between the MAX and NG.

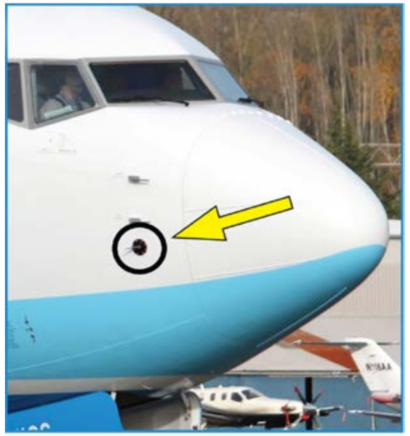

2 | MCAS Activation

The system activates when the aircraft approaches threshold AOA, or stickshaker activation, for the aircraft’s configuration and flight profile. The MAX flight-control law changes from speed trim to the MCAS because the MCAS reacts more quickly to AOA changes.

3 | Angle of Attack Vanes

The MCAS’s primary data sources are the MAX ’s two AOA sensing vanes, one on either side of the nose. Boeing designed the MCAS to receive input from only one of the sensors during each flight. The left and right sensors alternate between flights, feeding AOA data to the FCC and the MCAS. (There was an Optional sensor disagree light in the cockpit).

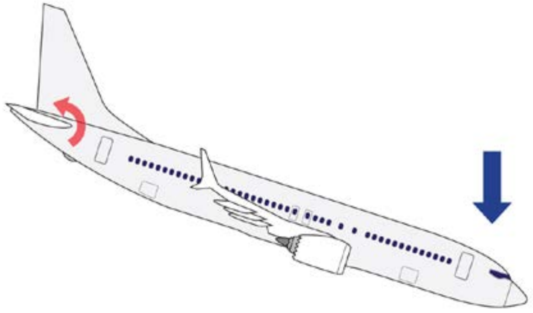

4 | Stabilizer Deflection

When threshold AOA is reached, the MCAS commands 0.27 deg. of aircraft nose-down stabilizer deflection per second for 9.3 sec.—a total of 2.5 units of trim.

Inaccurate AOA data will trigger the MCAS every 5 sec. until the data is corrected or the system is disabled.

AoA is a Single Point of Failure

The Boeing 737 MAX MCAS Explained

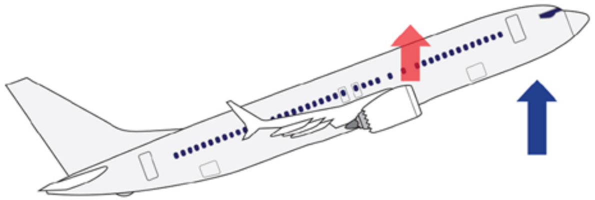

1 | Leap Engines and Pitch-up Moment
The MAX's larger CFM Leap 1 engines create more lift at high AOA and give the aircraft a greater pitch-up moment than the CFM56-7-equipped NG. The MCAS was added as a certification requirement to minimize the handling difference between the MAX and NG.

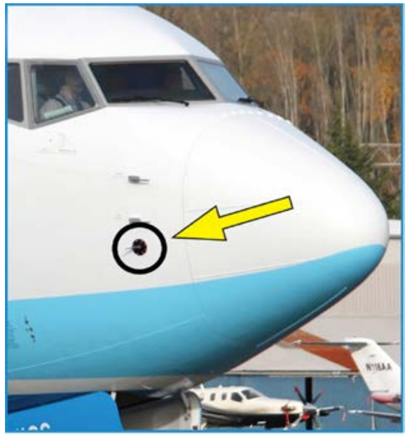

2 | MCAS Activation
The system activates when the aircraft approaches threshold AOA, or stickshaker activation, for the aircraft's configuration and flight profile. The MAX flight-control law changes from speed trim to the MCAS because the MCAS reacts more quickly to AOA changes.

3 | Angle of Attack Vanes
The MCAS's primary data sources are the MAX 's two AOA sensing vanes, one on either side of the nose. Boeing designed the MCAS to receive input from only one of the sensors during each flight. The left and right sensors alternate between flights, feeding AOA data to the FCC and the MCAS. (There was an Optional sensor disagree light in the cockpit).

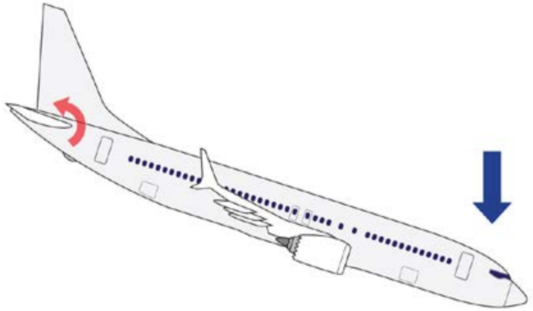

4 | Stabilizer Deflection
When threshold AOA is reached, the MCAS commands 0.27 deg. of aircraft nose-down stabilizer deflection per second for 9.3 sec.—a total of 2.5 units of trim.
Inaccurate AOA data will trigger the MCAS every 5 sec. until the data is corrected or the system is disabled.

MCAS Continues to Push Nose Down Until Manually Disengaged

Root Cause of MCAS Malfunction

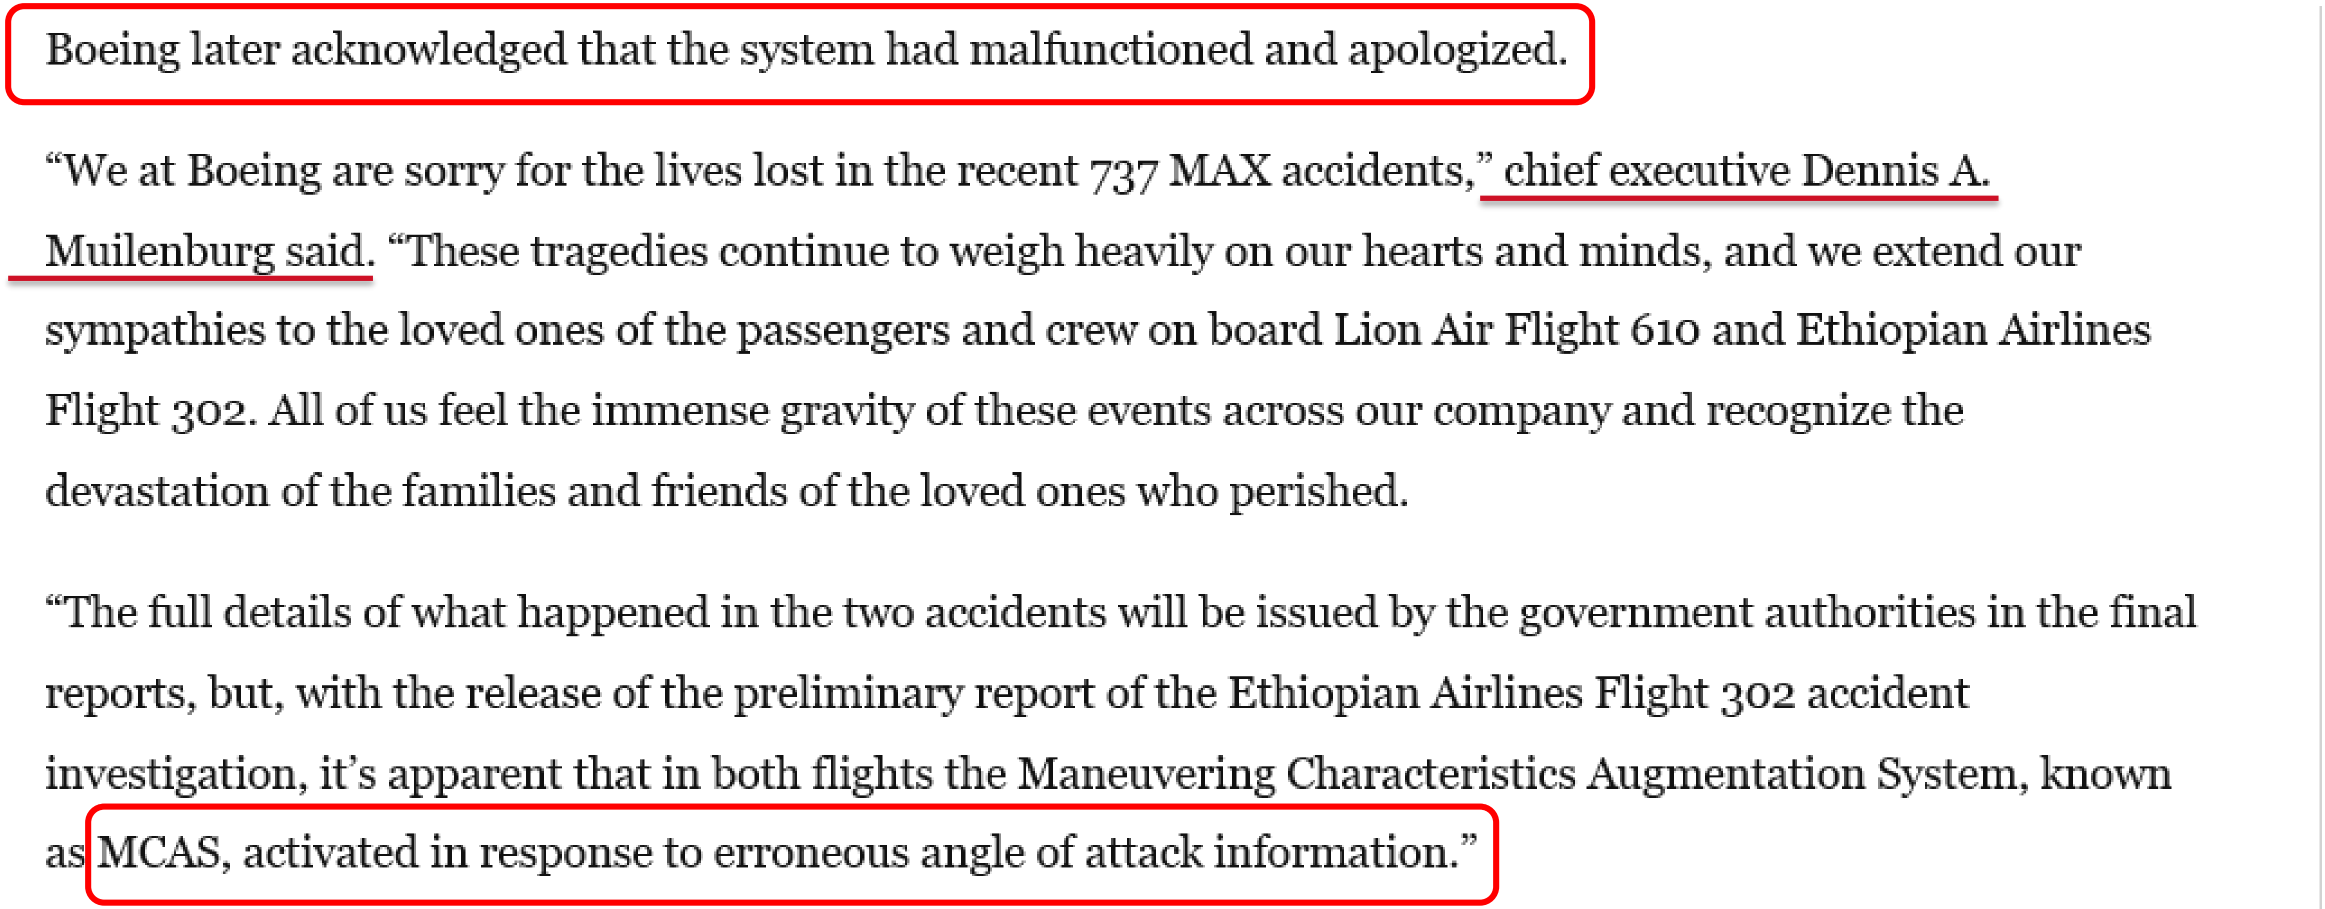

Boeing later acknowledged that the system had malfunctioned and apologized.

“We at Boeing are sorry for the lives lost in the recent 737 MAX accidents,” chief executive Dennis A. Muilenburg said. “These tragedies continue to weigh heavily on our hearts and minds, and we extend our sympathies to the loved ones of the passengers and crew on board Lion Air Flight 610 and Ethiopian Airlines Flight 302. All of us feel the immense gravity of these events across our company and recognize the devastation of the families and friends of the loved ones who perished.

“The full details of what happened in the two accidents will be issued by the government authorities in the final reports, but, with the release of the preliminary report of the Ethiopian Airlines Flight 302 accident investigation, it’s apparent that in both flights the Maneuvering Characteristics Augmentation System, known as MCAS, activated in response to erroneous angle of attack information.”

MCAS Activated in Response to Erroneous AOA Information – Boeing CEO

AOA Sensor Malfunction, A/C Altitude and Pitch

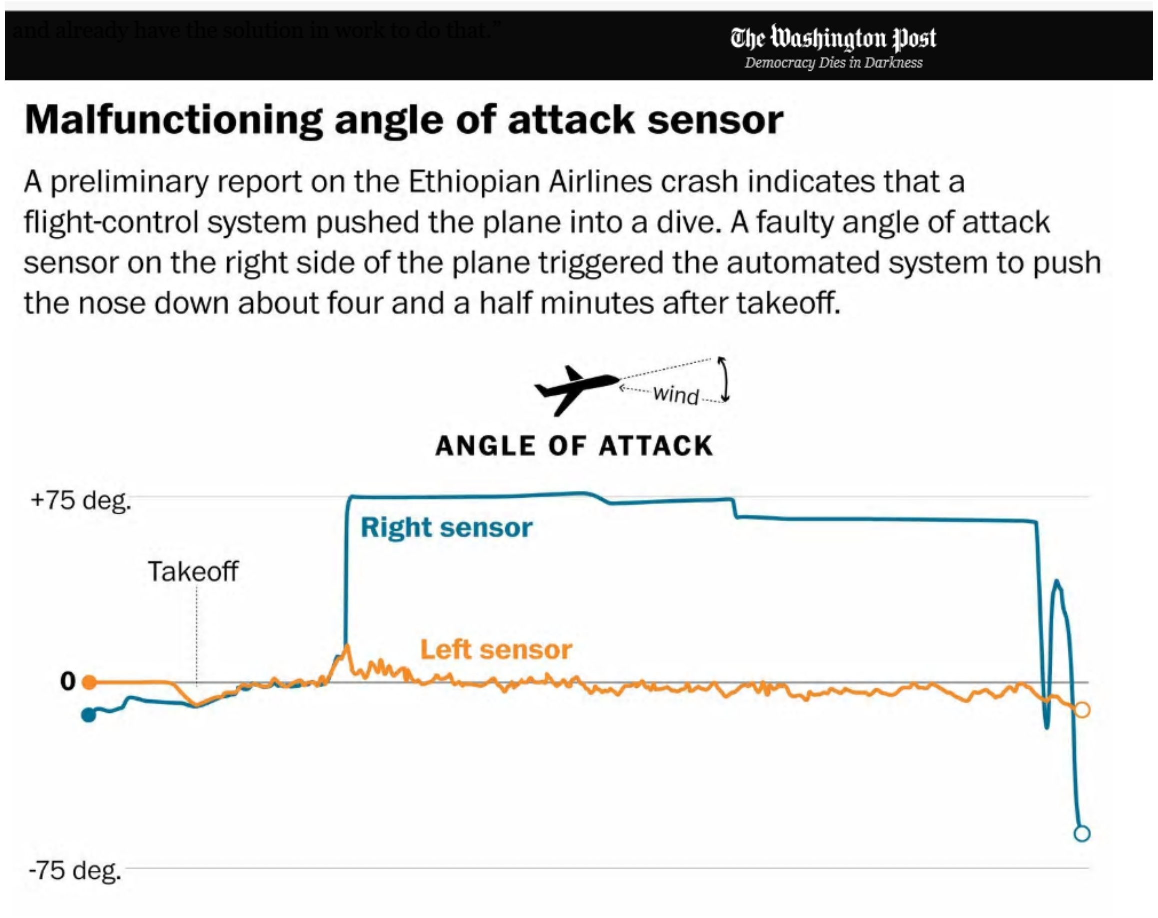

and already have the solution in work to do that.”

The Washington Post
Democracy Dies in Darkness

Malfunctioning angle of attack sensor

A preliminary report on the Ethiopian Airlines crash indicates that a flight-control system pushed the plane into a dive. A faulty angle of attack sensor on the right side of the plane triggered the automated system to push the nose down about four and a half minutes after takeoff.

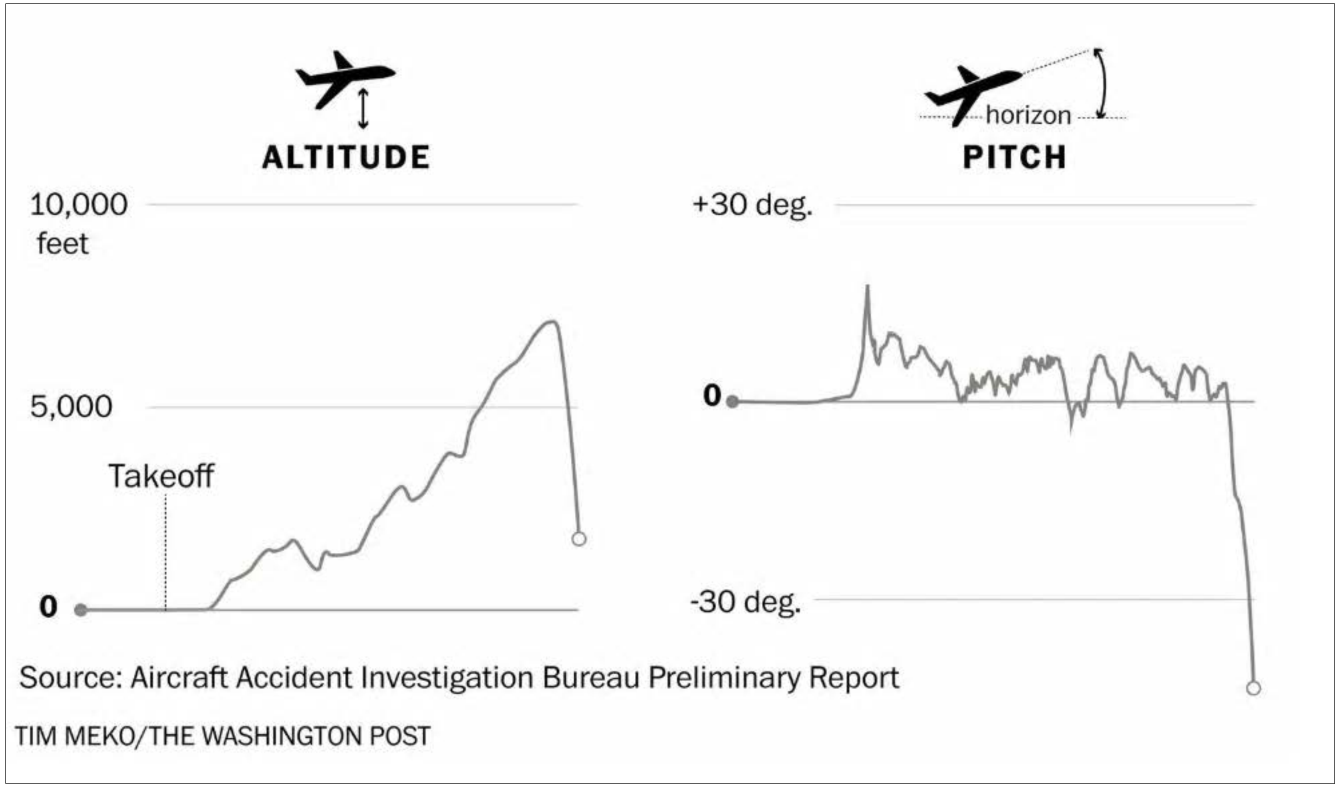

Source: Aircraft Accident Investigation Bureau Preliminary Report

TIM MEKO/THE WASHINGTON POST

Malfunctioning AOA Sensor caused MCAS to Repeatedly Push Airplane Nose Down for 4 ½ Minutes

Cockpit Disagree Warning Light

Boeing delayed fix of defective 737 MAX warning light for three years: U.S. lawmakers

Eric M. Johnson, Reuters

SEATTLE (Reuters) - Boeing Co learned that a cockpit warning light on its 737 MAX jetliner was defective in 2017 but decided to defer fixing it until 2020, U.S. lawmakers said on Friday.

The defective warning light alerts pilots when two sensors that measure the angle between the airflow and the wing disagree.

Boeing spokesman Gordon Johndroe said by email that a company safety review found the absence of the AOA Disagree alert did not adversely impact airplane safety or operation.

“Based on the safety review, the update was scheduled for the MAX 10 entry into service in 2020,” Johndroe said. “We fell short in the implementation of the AoA Disagree alert and are taking steps to address these issues so they do not occur again.”

What is Boeing doing to update MCAS

- MCAS now uses both left and right AOA sensors for redundancy, instead of relying on just one.
- The new software load [P12.1] has triple-redundant filters that prevent one or both angle-of-attack (AOA) systems from sending erroneous data to the FCCs that could falsely trigger the MCAS.
- MCAS cannot trim the stabilizer so that it overpowers elevator pitch control authority.
- If the pilots make electric pitch trim inputs to counter the MCAS, it won’t reset after 5 sec. and repeat subsequent nose-down stab trim commands.
- Cockpit Sensor Disagree Light will be standard equipment.

Cost-Benefit Trade-off of Safety

- Do fault-tolerant and cyber resilient systems cost more? Yes!
  - Additional non-recurring costs: design, development, validation & verification expenses
  - Additional recurring costs: hardware build, integration, and continuing operations & maintenance
- Added costs must be balanced against the adverse consequences of failures
  - Lives lost
  - Compensation to survivors and victims’ families
  - Compensation to system users (Airlines)
  - Govt (SEC) and shareholder lawsuits
  - Missed sales
  - Ruined reputations: Builder and Regulator (FAA)

Dependability is Expensive but lack of it can be Catastrophic

Some Adverse Consequences of MCAS Failures

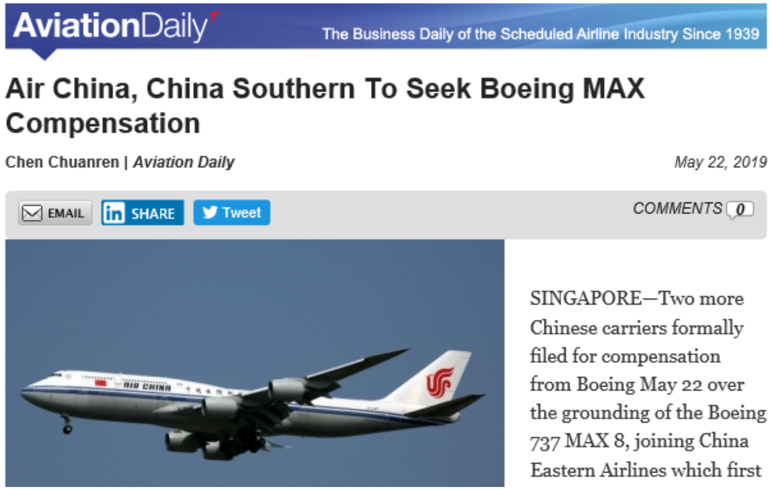

AviationDaily The Business Daily of the Scheduled Airline Industry Since 1939

Air China, China Southern To Seek Boeing MAX Compensation

Chen Chuanren | *Aviation Daily* *May 22, 2019*

SINGAPORE—Two more Chinese carriers formally filed for compensation from Boeing May 22 over the grounding of the Boeing 737 MAX 8, joining China Eastern Airlines which first

> In the most important measure—the 346 lives lost in recent Boeing MAX crashes—the cost is all too well-known. Further, what is evident from the Ethiopian Airlines Flight 302 (ET302) and Lion Air Flight 610 accidents is that those lives should not have been lost—and more could have been done to prevent that.
>
> - Michael Bruno, AW&ST

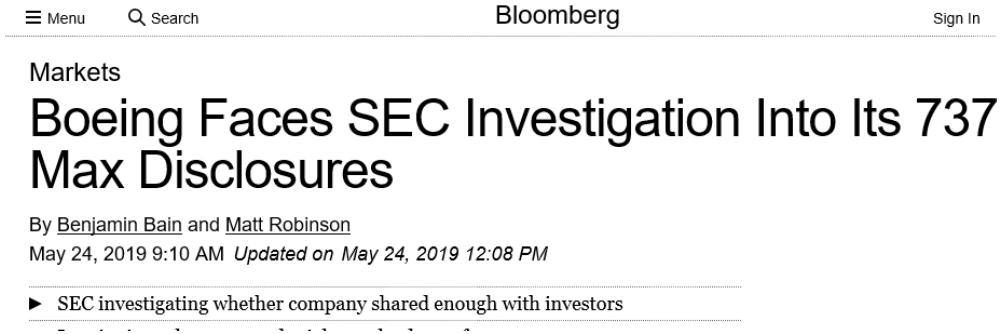

Markets

Boeing Faces SEC Investigation Into Its 737 Max Disclosures

By Benjamin Bain and Matt Robinson
May 24, 2019 9:10 AM *Updated on May 24, 2019 12:08 PM*

- SEC investigating whether company shared enough with investors

BUSINESS | EARNINGS

Boeing Details Financial Hit From 737 MAX Grounding

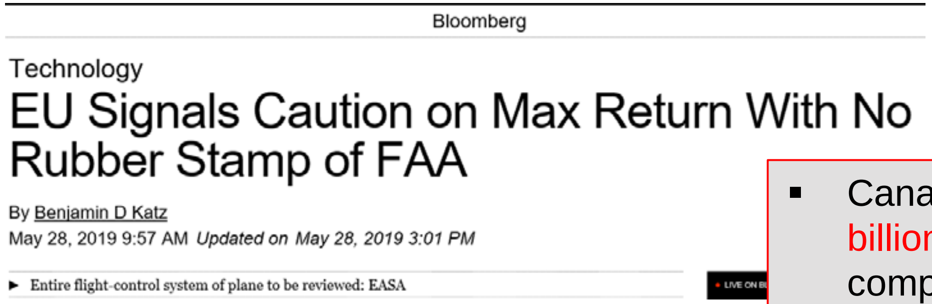

Bloomberg

Technology

EU Signals Caution on Max Return With No Rubber Stamp of FAA

By Benjamin D Katz
May 28, 2019 9:57 AM *Updated on May 28, 2019 3:01 PM*

- Entire flight-control system of plane to be reviewed: EASA

- Canaccord Genuity's Ken Herbert on April 22 estimated there will be around $2.2 billion in one-time costs associated with the groundings and accidents, including compensation to victims' families.
- For every month the groundings continue, it will cost Boeing another $1.2 billion.
- Sheila Kahyaoglu of Jefferies on April 24 estimated that over a quarter, MAX issues could amount to as much as $5 billion.

Autonomous Vehicle Safety Requirements

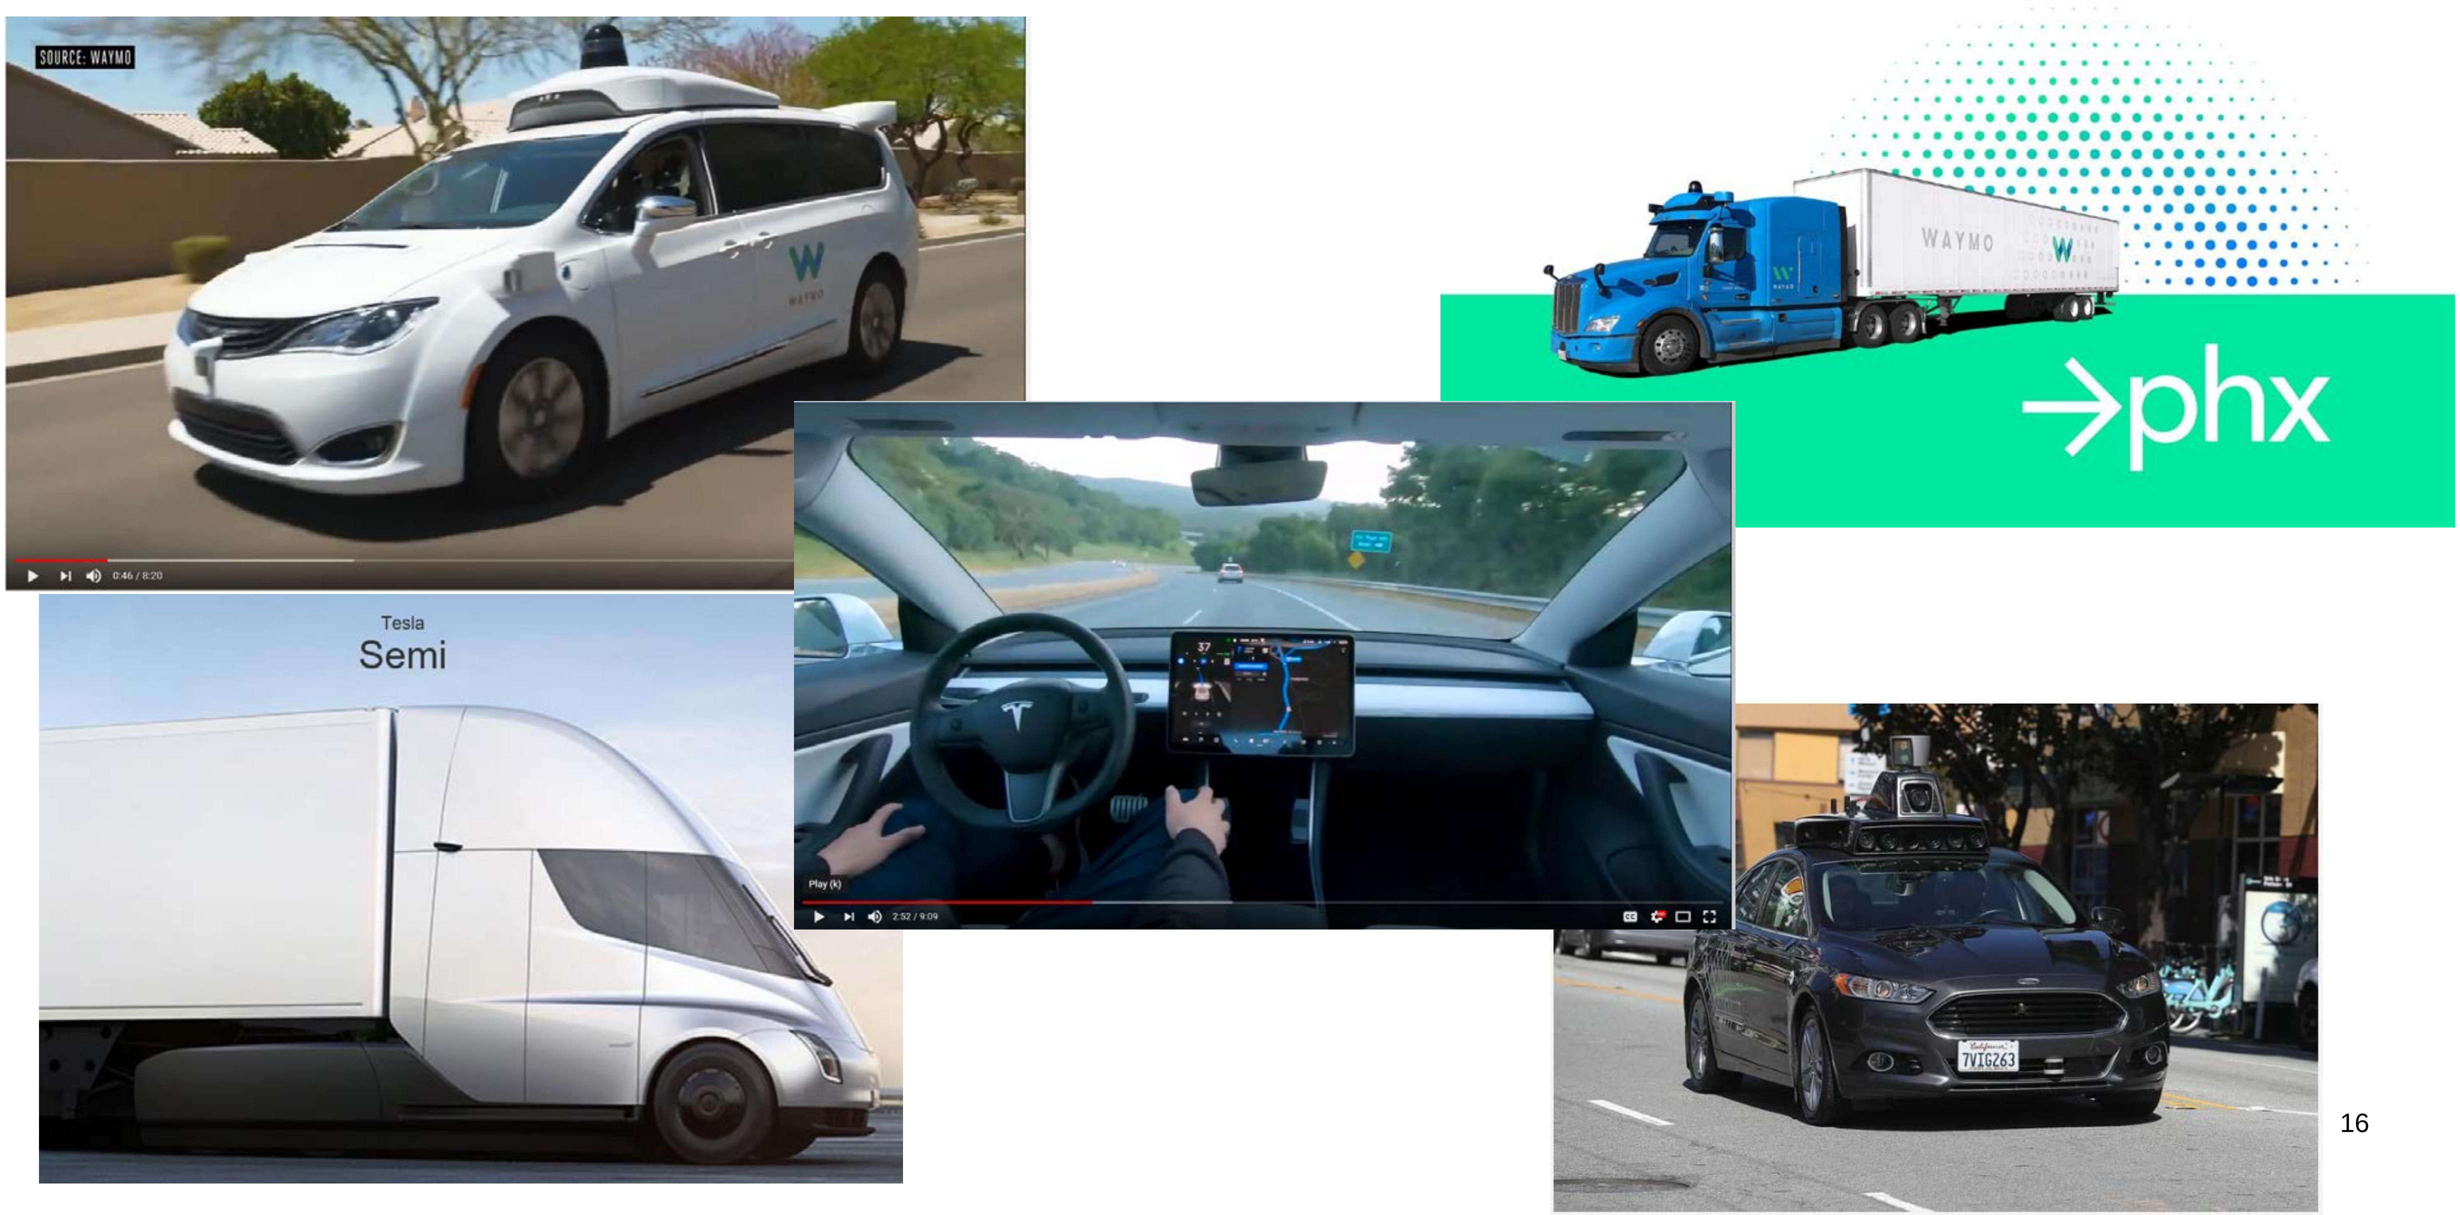

Motor Vehicle Accident Rates* (US)

Year	Deaths	Vehicle Miles Traveled (VMT) (Billions)	Fatalities/ 100 Million VMT	Population	Fatalities/ 100,000 People
2016	37,806	3,174	1.19	323,121,000	11.59

$100X10^6$ Vehicle Miles Traveled = $2.5X10^6$ Vehicle Hrs Traveled @40mph

1.19 Fatalities per $100X10^6$ VMT = $0.5X10^{-6}$ per hr

We use the shorthand Hr to indicate Vehicle Hrs Traveled

US Fatality Rate: $0.5X10^{-6}$ per hr

* US National Highway Traffic Safety Administration (NHTSA)

Autonomous Control System Components

- Sensors: Electro-Optical, Infrared, Radar, GPS, MEMS, Vehicle subsystems (Engine/Brakes/etc) performance, health & status sensors
- Processors: CPUs, GPUs, Software
- Communication: Links to other cars and Traffic Signaling Systems
- Actuators: Commands to Engine, Brakes, Steering
- Algorithms: Catch-all for all the Feedback Control System Functions, incl. sensor processing and correlation, situational awareness, decision making, collision avoidance, etc

Autonomous Control is a Complex System of Systems

Autonomous Control System: Safety Requirements

Several ways to specify requirements:

1. Quantitative Reliability Requirement: Failures/hr
   - Max acceptable prob of control system failure that results in loss of a safety-critical function
2. Ability to disengage and safely stop after one fault: Fail-Safe
3. Ability to continue to provide all safety-critical functions after
   - Any one fault: Fail-Operational
   - Two faults: Fail-Op/Fail-Op or Fail-Op/Fail-Safe
   - …..

Most Safety-Critical Systems Must Meet Both Requirements

Autonomous Intersection

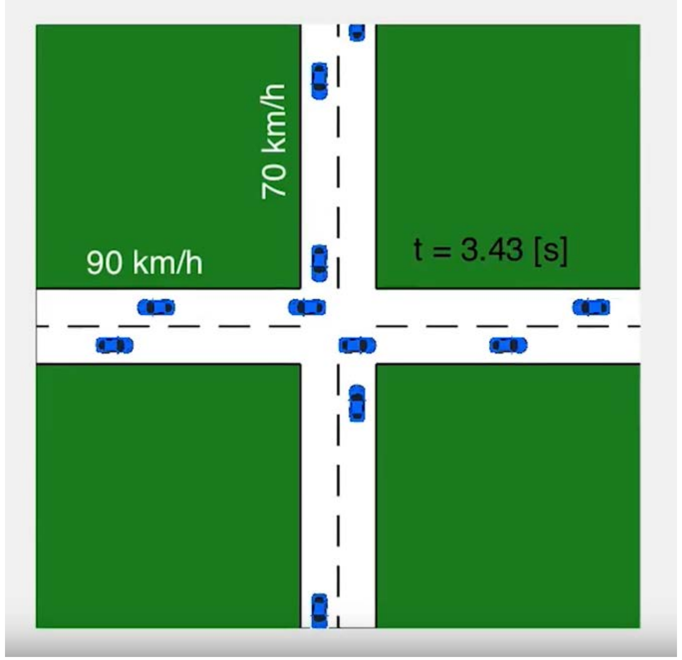

270 msec later

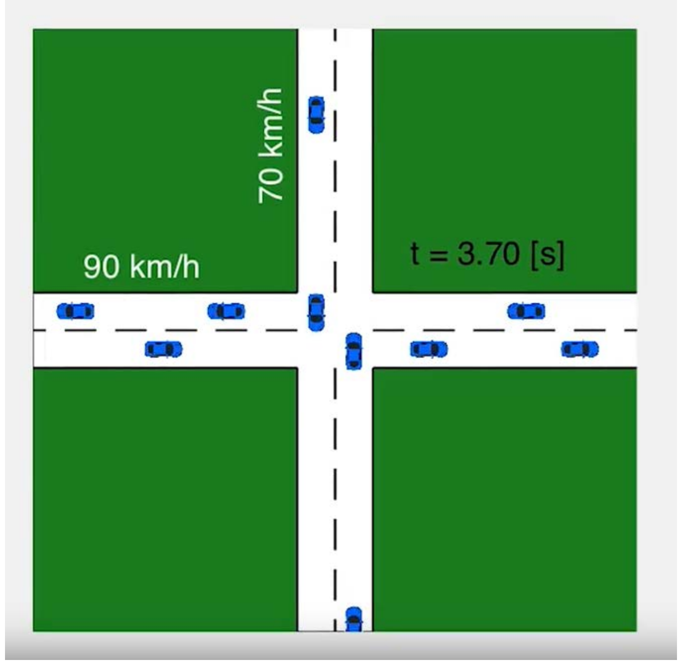

Chalmers University of Technology
Autonomous Intersection: Real & Simulated Traffic
May 17, 2017
https://www.youtube.com/watch?v=fzkv5beS4uk

What could possibly go wrong?

Autonomous Intersection:
6 Lanes in Each Direction

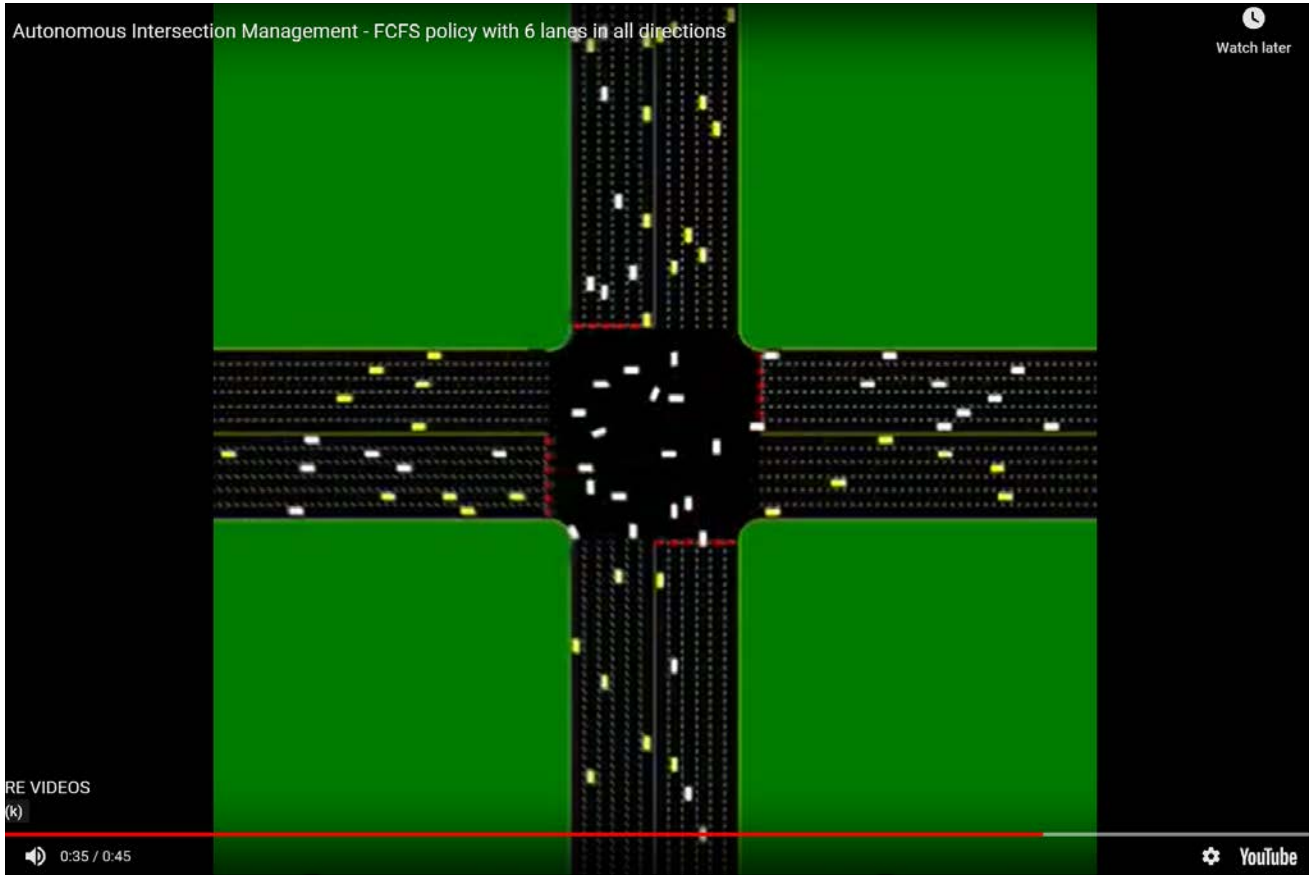
Autonomous Intersection Management - FCFS policy with 6 lanes in all directions
Watch later
RE VIDEOS
(k)
0:35 / 0:45
YouTube

Prof Peter Stone, University of Texas at Austin
Automated Intersection Management (AIM)
March 9, 2017
https://www.nbcnews.com/mach/innovation/self-driving-cars-will-turn-intersections-high-speed-ballet-n731511

What, indeed, could possibly go wrong?

Drivers of Safety Requirements

- Autonomous Vehicle Control System is a hard real-time computer system.
- Under nominal no-fault conditions, must produce correct control commands with low latency.
- In case of faults or errors, system must compensate for these, and still produce correct results in a timely manner.
- Unlike an aircraft fly-by-wire system, vehicle needs to function only for a short time, and in a limited capacity, to configure vehicle into a safe state and move to a safe place.
- A graceful degradation to a limited functionality Fail-Op requirement would seem to be adequate.
- No single point failure.

Graceful Degradation to Limited Capability Fail-Op

Quantitative Reliability Requirements

Case	Safety relative to current manual benchmark	Failure Rate (per hour)	Annual Deaths caused by Control System (US)	Deaths/Day (US)
1	Same as	$0.5X10^{-6}$	37,806	104
2	10X better	$0.5X10^{-7}$	3,780	10
3	100X better	$0.5X10^{-8}$	378	1
4	1,000X better	$0.5X10^{-9}$	38	0.1
5	10,000X better	$0.5X10^{-10}$	4	0.01

- The argument that if fewer people die, and society as a whole is safer, is simplistic.
  - ➢ It is very hard to justify innocent people sacrificing their lives, in the service of others.
  - ➢ Case in point: Reaction to Tempe, AZ; or MCAS

What level of loss of life, *caused by machines*, is acceptable to society?

How reliable are current autonomous vehicle control systems? What’s the evidence?

- There are many claims being made about safety of control systems (not counting CEO tweets) based on simplistic, non-scientific data
- Principally, very limited empirical data on prototype systems
  - Number of vehicle miles drive & Number of accidents
- Relevance of empirical driving record in extrapolating safety predictions
  - How representative are prototypes wrt fully autonomous control systems? (See previous slide on the control system components.)
  - How realistic are the testing conditions? Speed, traffic, weather, visibility, …
  - How good is data collection on control system performance? Unplanned disengagements, minor malfunctions (not resulting in accidents), human taking over control, incorrect decision making (not resulting in accidents), …
  - How many and which corner cases or edge cases were encountered? Outcomes?
- Analytical models: Reliability models, Monte Carlo simulations, Markov state models, Generative Adversarial Networks (GANs),…
- Experimental Data: Fault Injections, Penetration Testing, Zero-day exploits,…

Need Comprehensive Assurance Cases of Projected Safety Claims

Pioneering Work (1960s): Apollo GN&C Computer

- One of the first safety-critical digital computers
- Fault Tolerance
  - Memory parity bit
  - Process recovery
  - 70,000 hrs MTBF (est)
- Specs
  - 40,000 IPS
  - 36,000 Word ROM
  - 2,000 Word R/W Memory
  - 70 lbs; 2 CuFt; 70 W

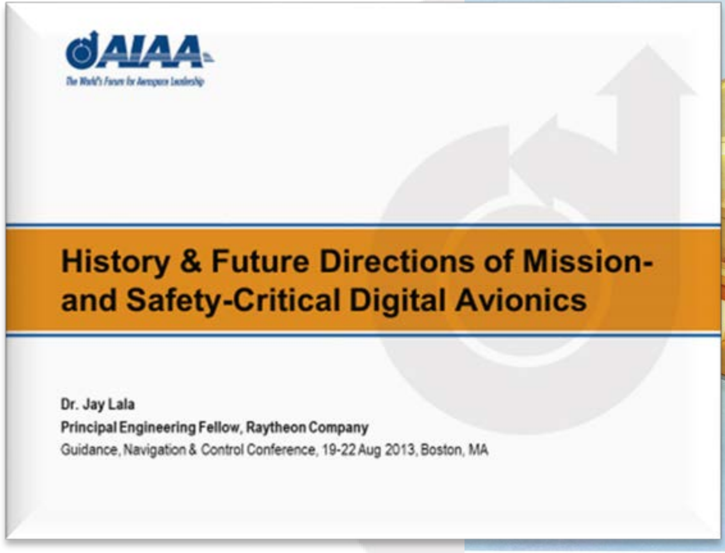

History & Future Directions of Mission- and Safety-Critical Digital Avionics

Dr. Jay Lala
Principal Engineering Fellow, Raytheon Company
Guidance, Navigation & Control Conference, 19-22 Aug 2013, Boston, MA

Designed & Prototyped by: MIT Instrumentation Lab
Manufactured by: Raytheon Company

FAA Memo on Fly-by-Wire Flight Control System – 1974

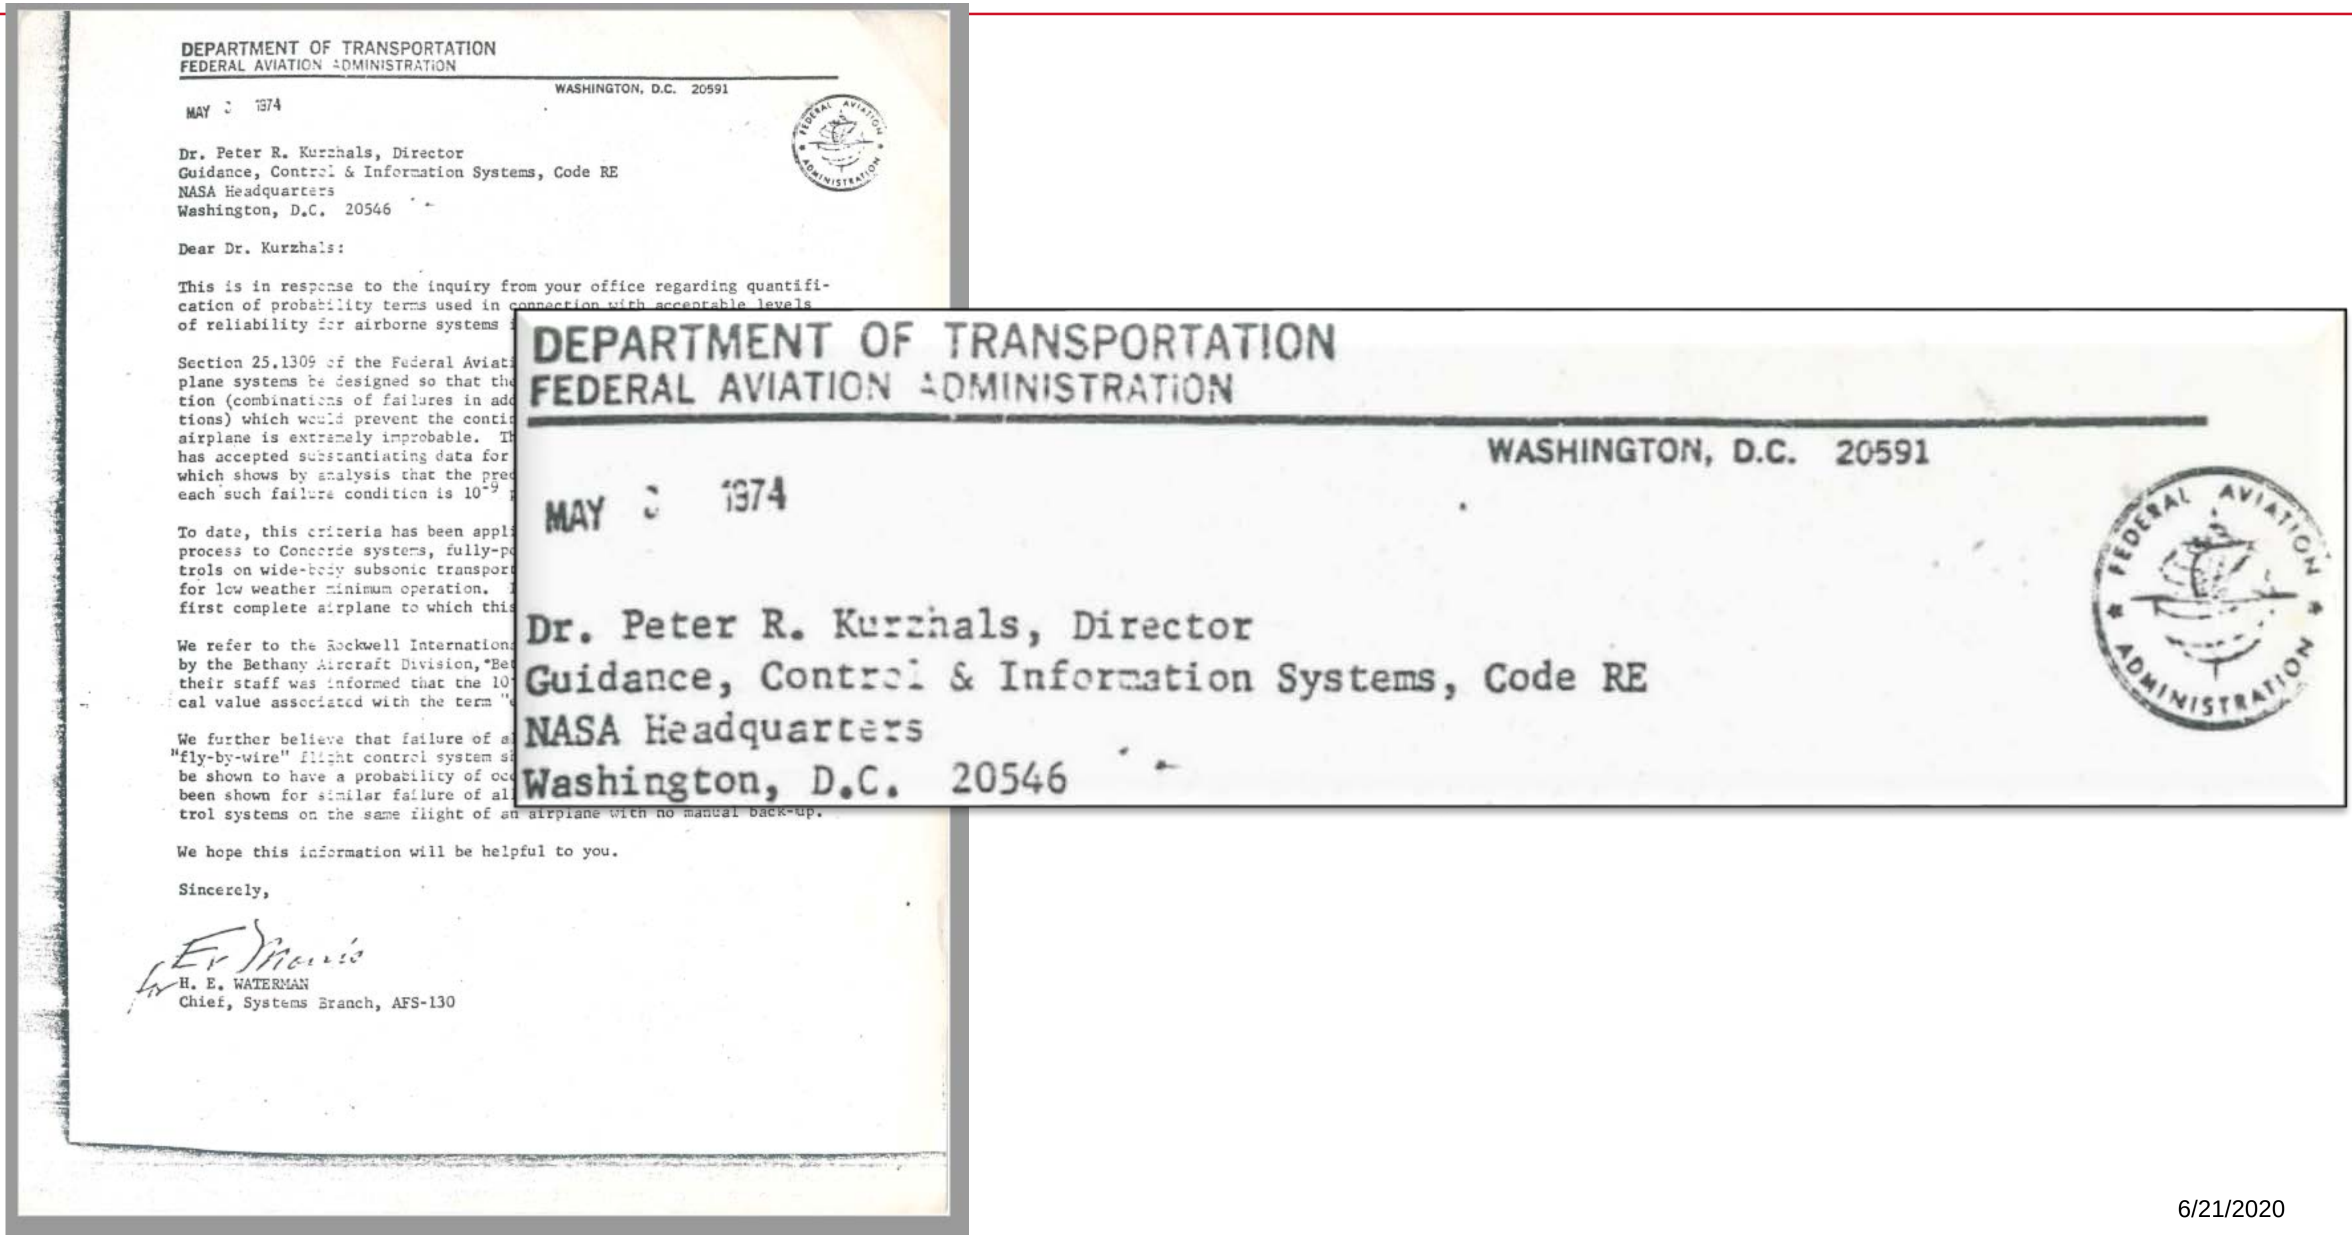

DEPARTMENT OF TRANSPORTATION
FEDERAL AVIATION ADMINISTRATION

WASHINGTON, D.C. 20591

MAY 3 1974

Dr. Peter R. Kurzhals, Director
Guidance, Control & Information Systems, Code RE
NASA Headquarters
Washington, D.C. 20546

Dear Dr. Kurzhals:

This is in response to the inquiry from your office regarding quantification of probability terms used in connection with acceptable levels
of reliability for airborne systems i

Section 25.1309 of the Federal Aviati
plane systems be designed so that the
tion (combinations of failures in add
tions) which would prevent the contin
airplane is extremely improbable. Th
has accepted substantiating data for
which shows by analysis that the pred
each such failure condition is 10^{-9} p

To date, this criteria has been appli
process to Concorde systems, fully-po
trols on wide-body subsonic transport
for low weather minimum operation. I
first complete airplane to which this

We refer to the Rockwell Internationa
by the Bethany Aircraft Division, "Bet
their staff was informed that the 10
cal value associated with the term "e

We further believe that failure of al
"fly-by-wire" flight control system s
be shown to have a probability of occ
been shown for similar failure of all
trol systems on the same flight of an airplane with no manual back-up.

We hope this information will be helpful to you.

Sincerely,

H. E. WATERMAN
Chief, Systems Branch, AFS-130

DEPARTMENT OF TRANSPORTATION
FEDERAL AVIATION ADMINISTRATION

WASHINGTON, D.C. 20591

MAY 3 1974

Dr. Peter R. Kurzhals, Director
Guidance, Control & Information Systems, Code RE
NASA Headquarters
Washington, D.C. 20546

FAA Reliability Requirements for Aircraft Fly-by-Wire Flight Control System - 1974

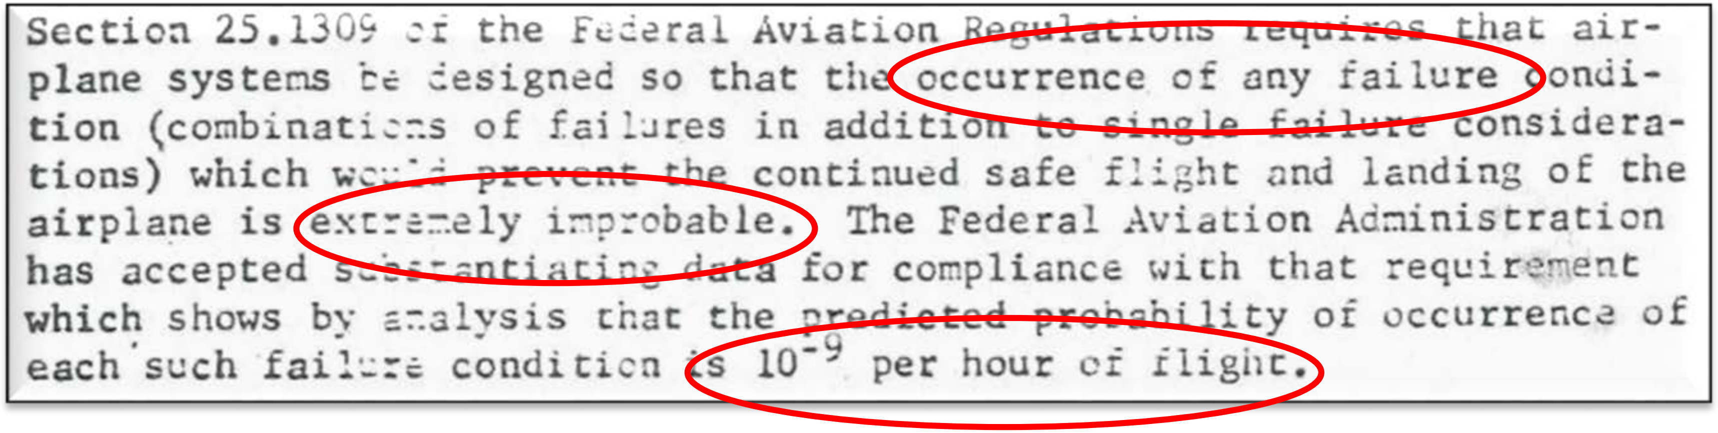

Section 25.1309 of the Federal Aviation Regulations requires that airplane systems be designed so that the occurrence of any failure condition (combinations of failures in addition to single failure considerations) which would prevent the continued safe flight and landing of the airplane is extremely improbable. The Federal Aviation Administration has accepted substantiating data for compliance with that requirement which shows by analysis that the predicted probability of occurrence of each such failure condition is 10^{-9} per hour of flight.

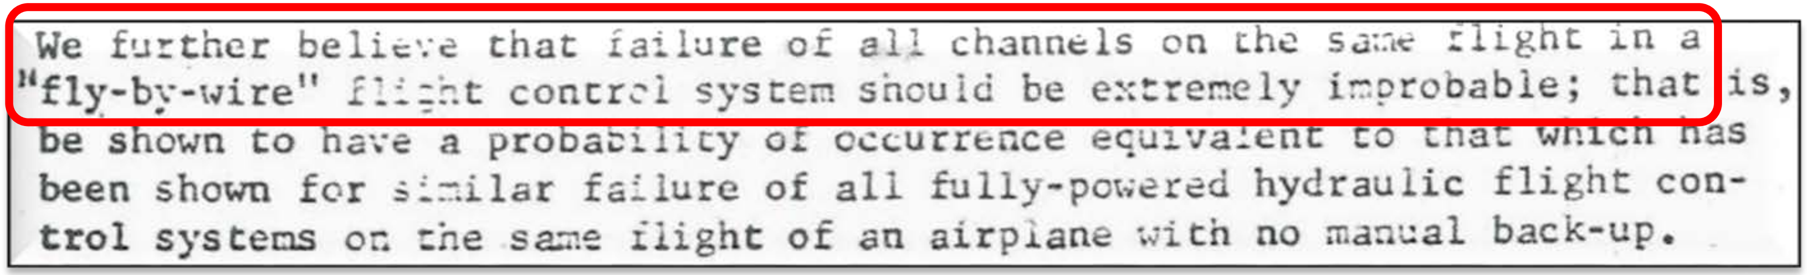

We further believe that failure of all channels on the same flight in a "fly-by-wire" flight control system should be extremely improbable; that is, be shown to have a probability of occurrence equivalent to that which has been shown for similar failure of all fully-powered hydraulic flight control systems on the same flight of an airplane with no manual back-up.

Draper Memo: NASA Interpretation of FAA FBW Requirements

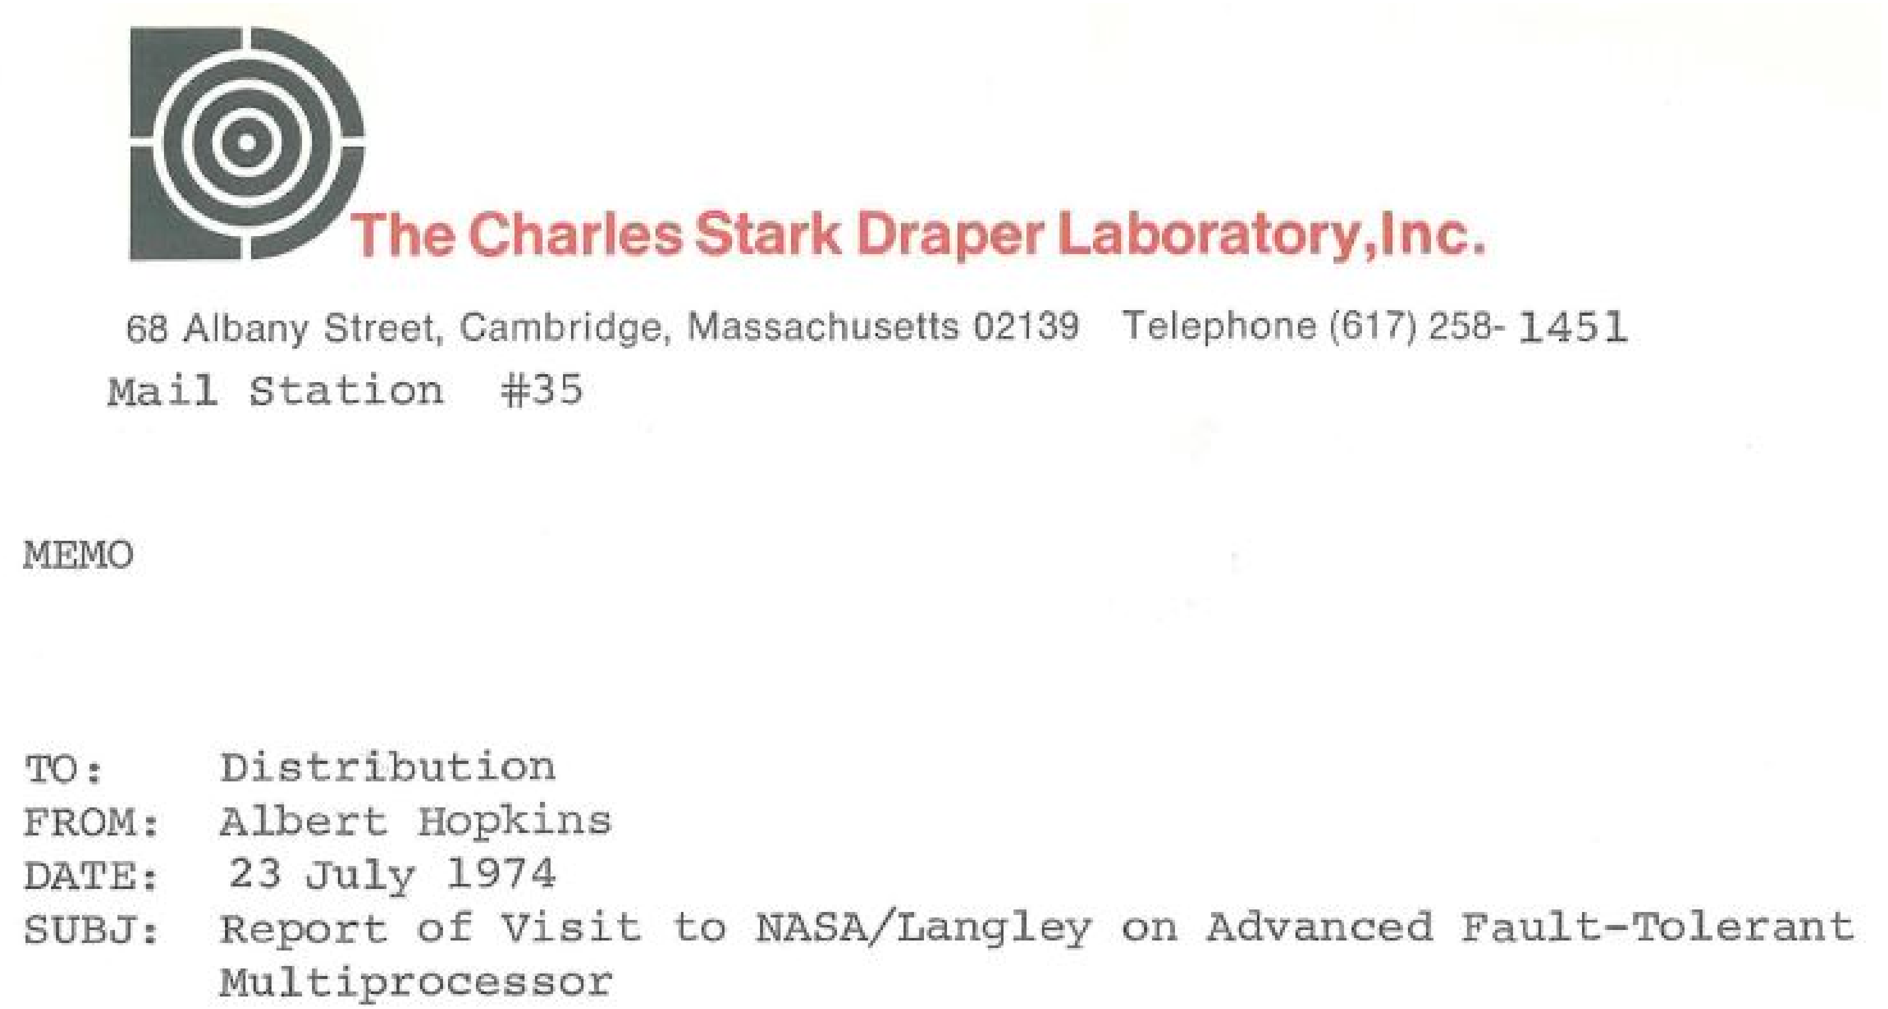

The Charles Stark Draper Laboratory,Inc.

68 Albany Street, Cambridge, Massachusetts 02139 Telephone (617) 258- 1451
Mail Station #35

MEMO

TO: Distribution
FROM: Albert Hopkins
DATE: 23 July 1974
SUBJ: Report of Visit to NASA/Langley on Advanced Fault-Tolerant Multiprocessor

Draper Memo: NASA Interpretation of FAA FBW Requirements

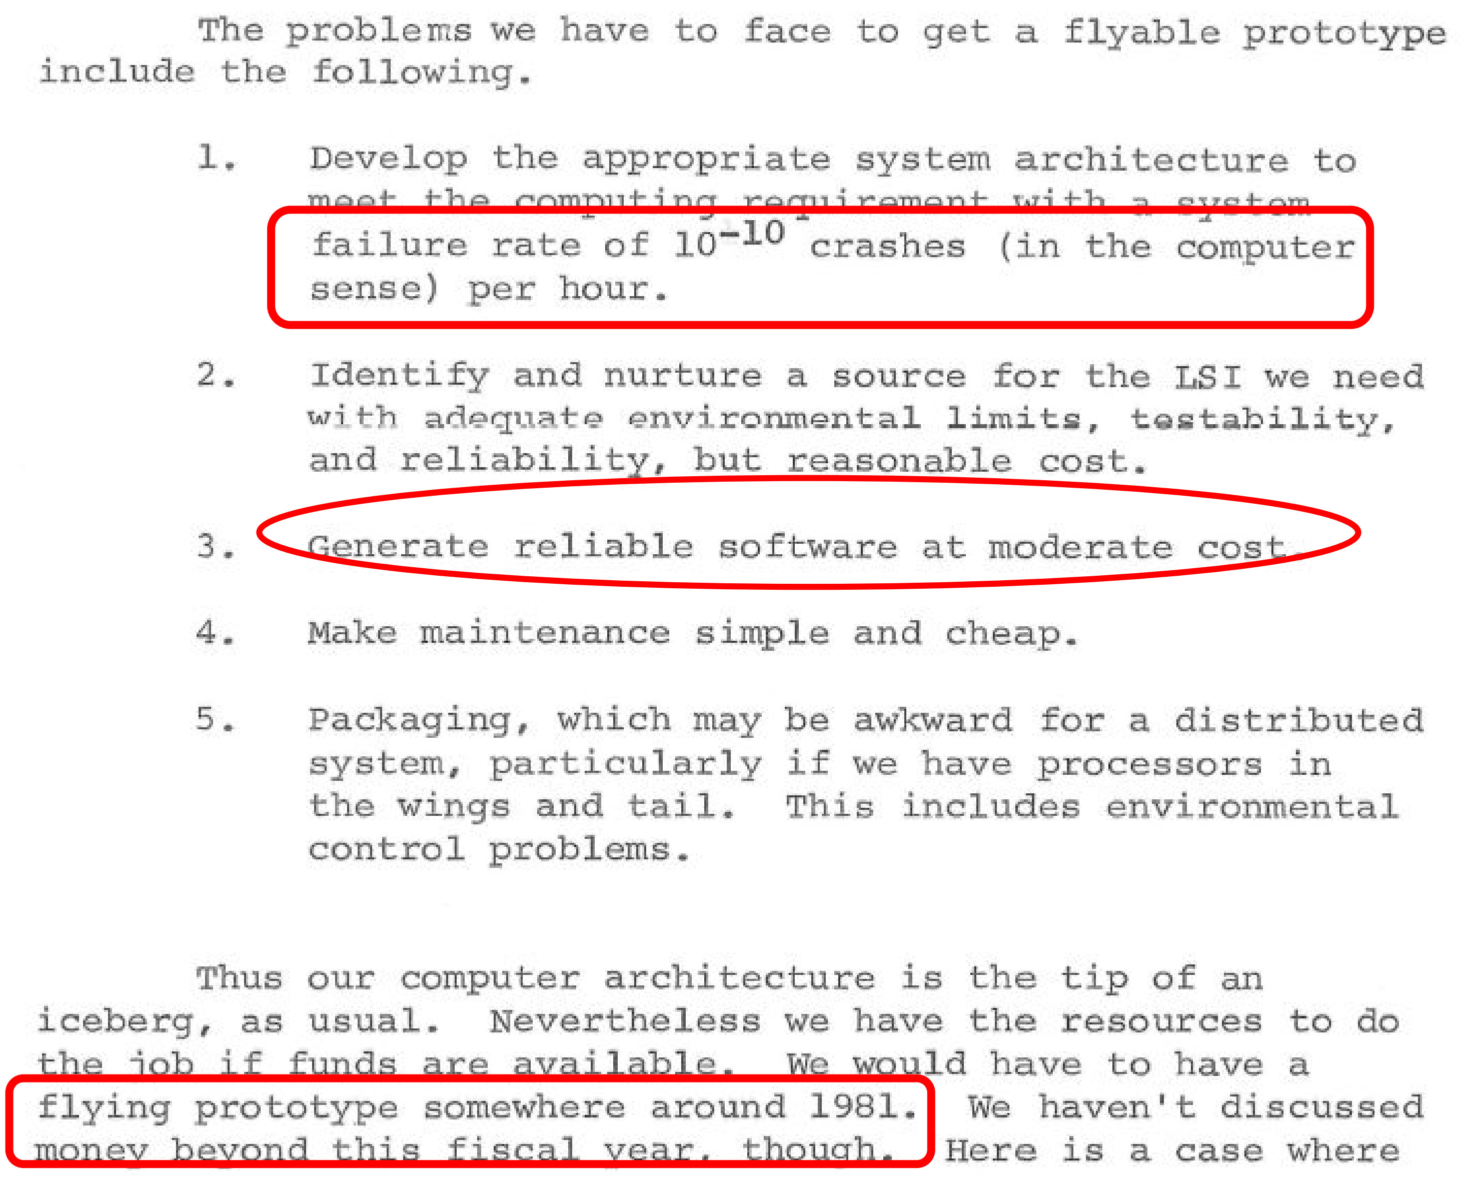

The problems we have to face to get a flyable prototype include the following.

1. Develop the appropriate system architecture to meet the computing requirement with a system failure rate of 10^{-10} crashes (in the computer sense) per hour.

2. Identify and nurture a source for the LSI we need with adequate environmental limits, testability, and reliability, but reasonable cost.

3. Generate reliable software at moderate cost.

4. Make maintenance simple and cheap.

5. Packaging, which may be awkward for a distributed system, particularly if we have processors in the wings and tail. This includes environmental control problems.

Thus our computer architecture is the tip of an iceberg, as usual. Nevertheless we have the resources to do the job if funds are available. We would have to have a flying prototype somewhere around 1981. We haven't discussed money beyond this fiscal year, though. Here is a case where

Draper Fault Tolerant Multi-Processor (FTMP)

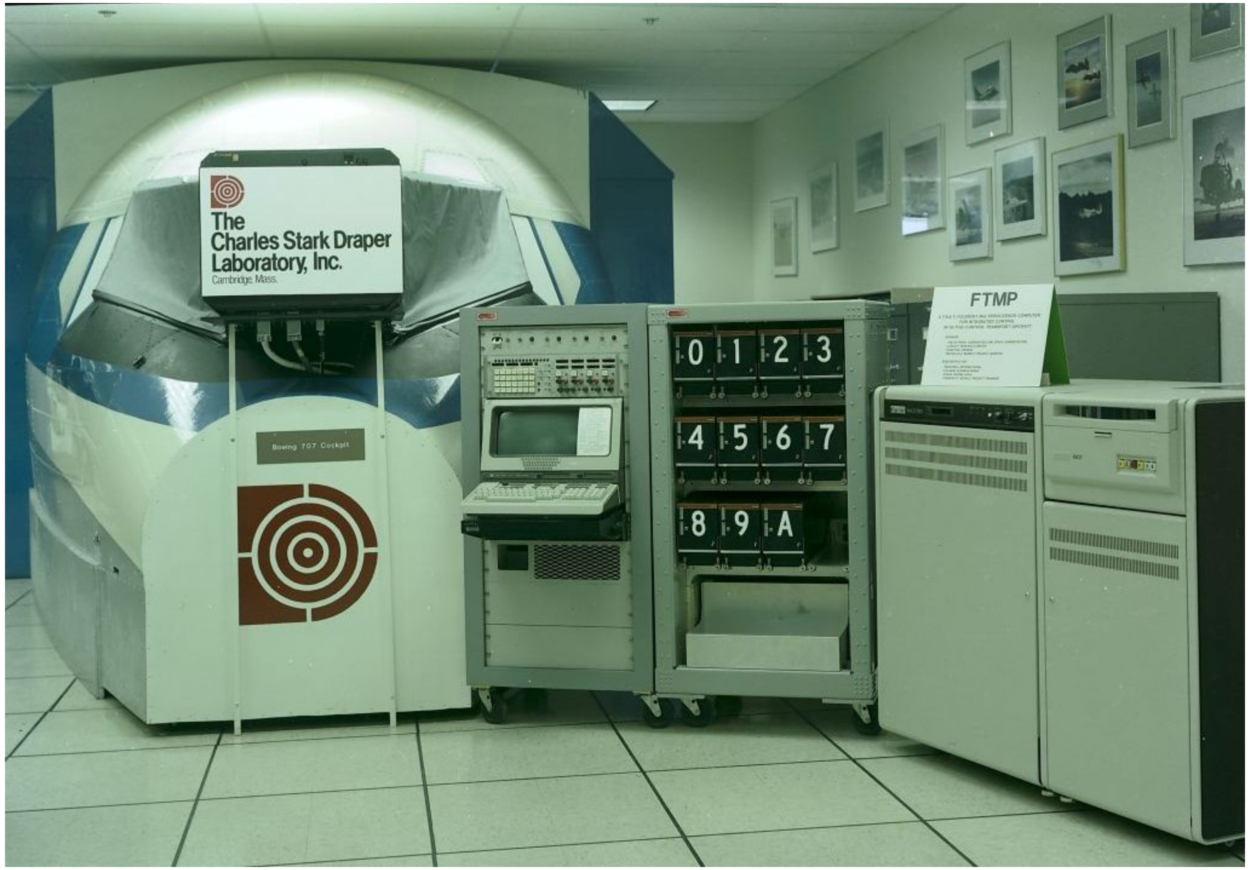

- Highly reconfigurable symmetric multiprocessor architecture
  - Triplex processor and memory triads
  - Hardware voting
  - Automated FDIR
- Validation of $P_f < 10^{-10}$/hr
  - Analytical Markov models
  - Empirical FDIR data collected via HW fault injector

Example Safety-Critical Computers (Draper Lab)

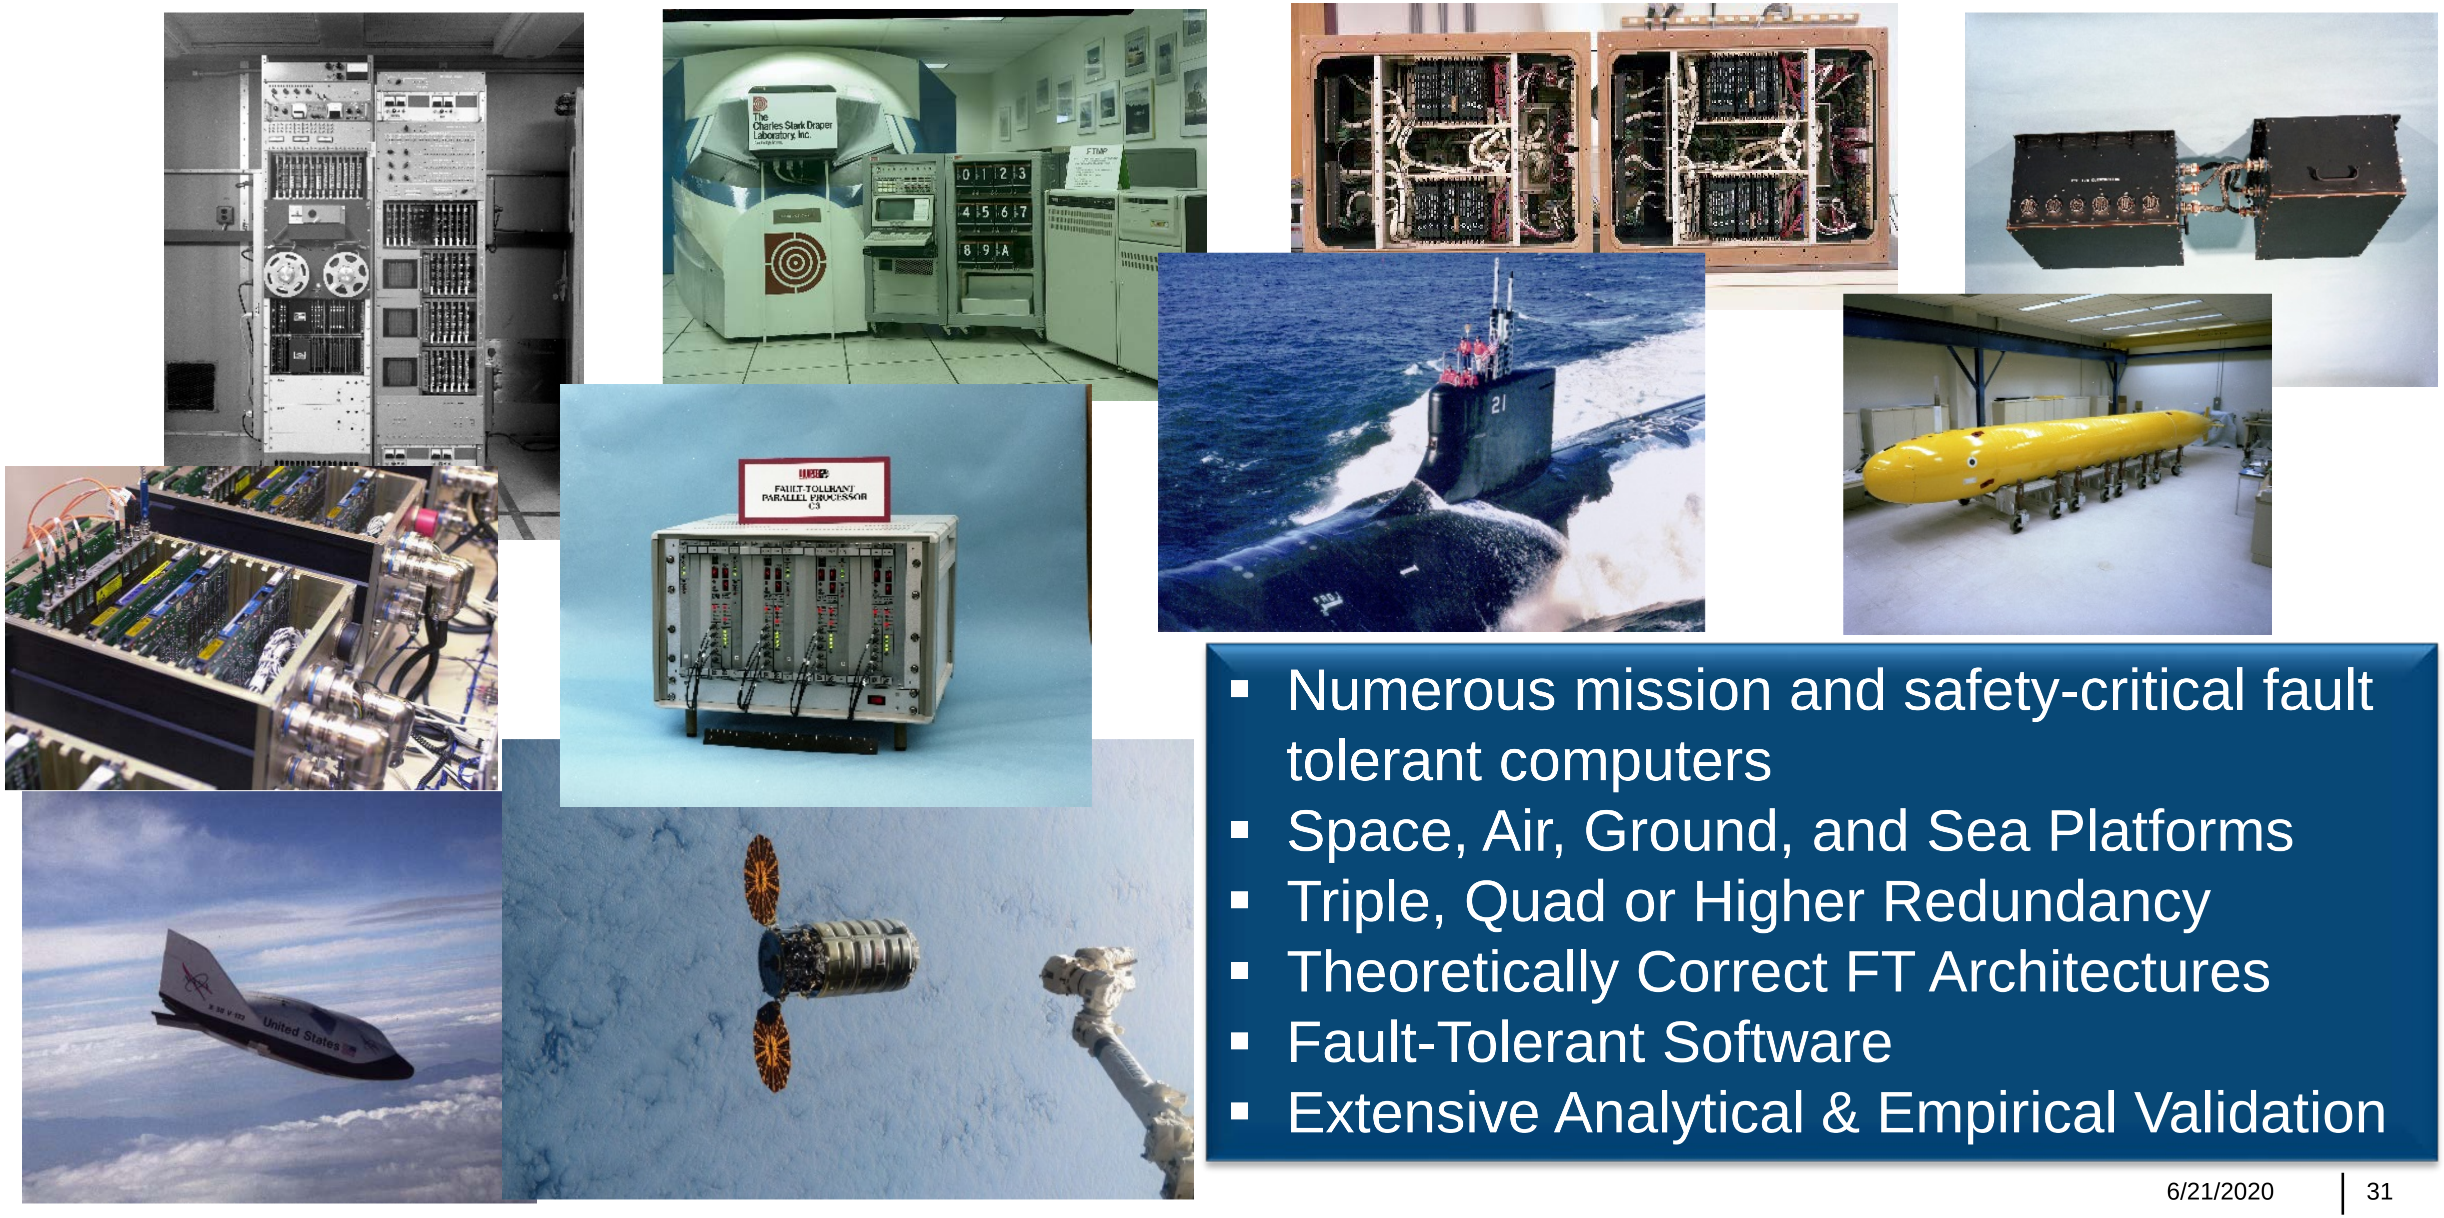
The Charles Stark Draper Laboratory, Inc.
FTMP
0 1 2 3
4 5 6 7
8 9 A
FAULT-TOLERANT PARALLEL PROCESSOR C3
21
United States

- Numerous mission and safety-critical fault tolerant computers
- Space, Air, Ground, and Sea Platforms
- Triple, Quad or Higher Redundancy
- Theoretically Correct FT Architectures
- Fault-Tolerant Software
- Extensive Analytical & Empirical Validation

Is AI/ML Ready for Self-Driving Cars?

Q. Artificial intelligence (AI) already is being widely used on the ground for data mining and trending.
Do you see it starting to move into platforms?

A. In order to get something safety-certified, you have to be able to predetermine what the machine will do in a given scenario, and AI isn't deterministic in that regard.

So I don't think you're going to see AI flying airplanes independently in the near future.

I think it may become a supplemental tool, but there still has to be an overarching system that determines what the airplane does under failed conditions.

- Collins Aerospace CEO Kelly Ortberg Interview w/Aviation Week & Space Technology, June 11, 2019

Collins Aerospace CEO Kelly Ortberg.
Credit: Collins Aerospace

Do we really believe that driving a car in mixed-mode traffic and all kinds of environmental conditions is really simpler than flying an airplane?

Aviation is not yet ready to adopt AI/ML. Why is autonomous vehicle community?

Pioneering Work (early ‘70s): All Weather Autolands

- First generation of jumbo-jets used analog computers to provide “all-weather” autoland capabilities
  - Cat IIIB conditions: zero visibility, zero ceiling
- Architectures were an outgrowth of 60s analog autopilots
- Fault Tolerance
  - DC-10: Duplex channels, each with dual fail-disconnect computers for pitch, roll, and yaw axes
  - B-747: Triple redundant analog computers
  - L-1011: Dual redundant self-checking pair of digital computers
- Ultrahigh reliability had to be sustained for only 2-3 mins

Intrusion Tolerant Systems
Fault Classification & ITS Scope

	NATURE		ORIGIN						PERSISTENCE		Usual Labelling
			Phenomenological Cause		System Boundaries		Phase of Creation				
	Accidental Faults	Intentional Faults	Physical Faults	Human-made Faults	Internal Faults	External Faults	Design Faults	Operational Faults	Permanent Faults	Temporary Faults	
Fault Tolerance	X		X		X			X	X		*Physical Faults*
	X		X			X		X	X		
	X		X			X		X		X	*Transient Faults*
	X		X		X			X		X	*Intermittent Faults*
	X			X	X		X			X	
	X			X	X		X		X		*Design Faults*
	X			X		X		X		X	*Interaction Faults*
ITS		X		X	X		X		X		*Malicious Logic*
		X		X	X		X			X	
		X		X		X		X	X		*Intrusions*
		X		X		X		X		X	

Cyber Resilient Architectures

Prevent Intrusions
(Access Controls, Cryptography, Trusted Computing Base)

But intrusions will occur

Detect Intrusions, Limit Damage
(Firewalls, Intrusion Detection Systems, Virtual Private Networks, PKI)

But some attacks will succeed

Tolerate Attacks
(Redundancy, Diversity, Deception, Wrappers, Proof-Carrying Code, Proactive Secret Sharing)

So the system must reconstitute

Restore System
(Diagnosis, Learning, Reconfiguration, S/W Rejuvenation, Natural Immunity, Reflection)

Trusted Computing Base

Access Control & Physical Security

Cryptography

Multiple Security Levels

1st Generation: Protection

Firewalls

Boundary Controllers

Intrusion Detection Systems

VPNs

PKI

2nd Generation: Detection

Intrusion Tolerance

Big Board View of Attacks Real-Time Situation Awareness & Response

Graceful Degradation

Hardened Operating System

3rd Generation: Tolerance

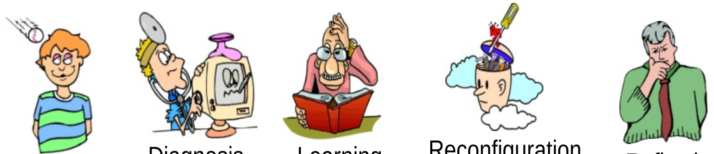

Self-Aware

Diagnosis

Learning

Reconfiguration

Reflection

4th Generation: Regeneration

Summary & Conclusions

- Autonomous vehicle control systems are very complex system of systems.
- They are also hard real-time, safety-critical systems, not unlike commercial airline flight control systems.
- Reliability and safety requirements should be commensurately high.
- Experiences, both good and bad, of the avionics architectures and designs of the past four decades should be leveraged for best solutions.
- Additionally, intrusion tolerance will be an added driver.
- Regulatory oversight and governance will be necessary to create, foster and enforce a culture of safety in automotive sector.

We should aspire to make Autonomous Vehicle as safe as Commercial Aviation