TTP Forum, Munich 15 November 2002

Based on SC200, 12 Nov 2002,

Dagstuhl seminar on Dependability of Component Based Systems, 3–8 November 2002

NASA Langley 12, 13 September 2001

Modular Safety Analysis And Modular Certification

John Rushby

Computer Science Laboratory
SRI International
Menlo Park, California, USA

Background

- We're part of a NASA/Honeywell project looking at new verification and certification methods for modular avionics
- The specific context is Honeywell Tucson's MAC (Modular Aerospace Controls) architecture
  - Based on TTTech's TTA (Time Triggered Architecture)

  That is being used for engine controls
  - Aeromachi, due to fly soon
  - F110 for F16
- We're also involved in formal verification of TTA algorithms and properties

System Safety Cases

- Dependability, safety, and related attributes are system-level (i.e., emergent) properties
- Hence, safety cases and certification, consider the full system
  - E.g., in civil air transport, the smallest unit of certification is a complete airplane or engine
- However, there is informal reuse of component-level arguments
- This has been aided by the federated architectures employed
  - Each function (autopilot, autothrottle, FMS etc.) has its own fault-tolerant platform (cf. brakes, steering)
  - Little interaction between them, limited fault propagation

  Safety analysis and certification can focus on component-system interactions, rather than component-component

The Trend Towards Integration

- But safety-critical systems are now becoming integrated
  - Resources shared between functions
  - Stronger interactions among them

  More functionality at less cost
- Avionics
  - Integrated Modular Avionics (IMA)
  - Modular Aerospace Controls (MAC)
- Similar trend in automobiles
  - Integrated steering, brakes, suspension by wire
- And everything else (chemical plants, assembly lines. . .)
- New hazards from fault propagation, and unintended emergent behavior

The Trend Towards Components

- Existence of standardized IMA platforms (e.g., TTA)
  - And interfaces (e.g., APEX, OSEK)

  Encourages development of semi-standardized components (e.g., autopilot or brake-by-wire modules)
- Integrators and customers want these components to be packaged with a certification argument
- Or to come already "pre-certified"
- System certification then largely a composition of component-level pre-certification arguments
- Need a notion of modular safety analysis and certification to support these constructs
- RTCA SC200 and Eurocae WG60 are holding joint meetings to develop a basis for modular certification

Component-Based Systems

- For less critical systems, and other kinds of properties, engineering design and analysis is generally based on components
- Components are characterized by specified properties at their interfaces
- System design and analysis then considers the composite behavior of many components operating within some structure (the design, or architecture)
  - Components are treated as "black boxes"
  - That is, we reason about their properties and interfaces
  - And do not need to open up their internal designs
- Can we (or why can't we) do this for avionics and cars?

Modular Safety Analysis: First Interpretation

- Safety argument for system Y with component/subsystem X makes use of separately certified claims about X's properties at its interfaces
- As opposed to opening up the design of X
- Pictorially:

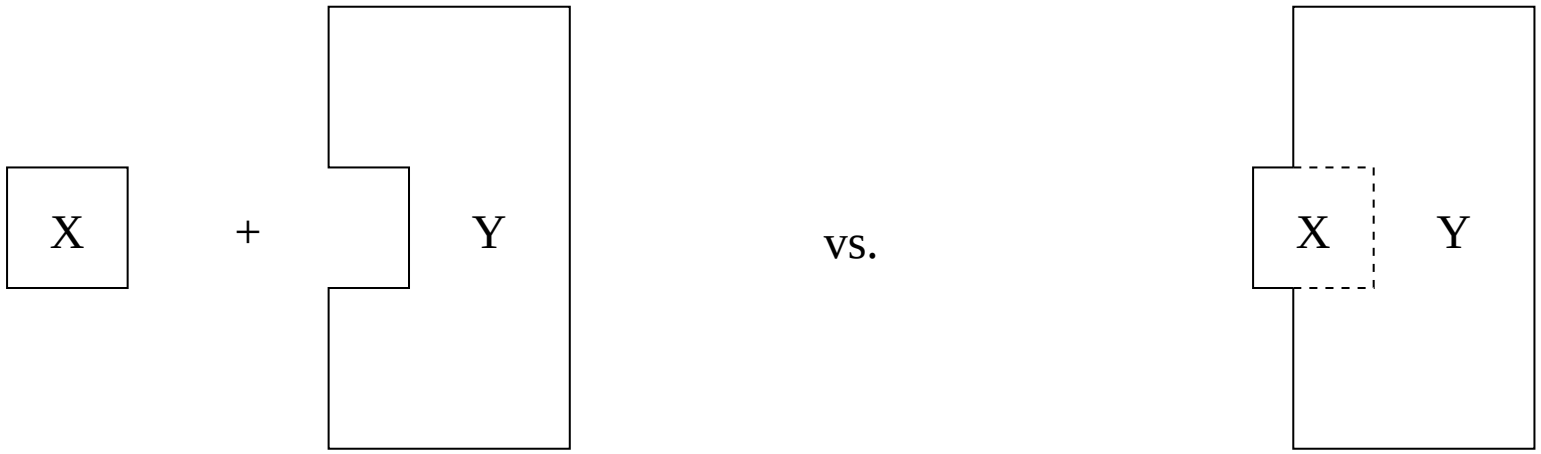

- Think of X as some safety function of automobile Y

Modular Safety Analysis: Second Interpretation

- May be infeasible to analyze safety of X separately from Y
- But suppose that X_1 and X_2 have both separately been analyzed with Y
- Can we establish if it is safe to use X_1 and X_2 together in Y?

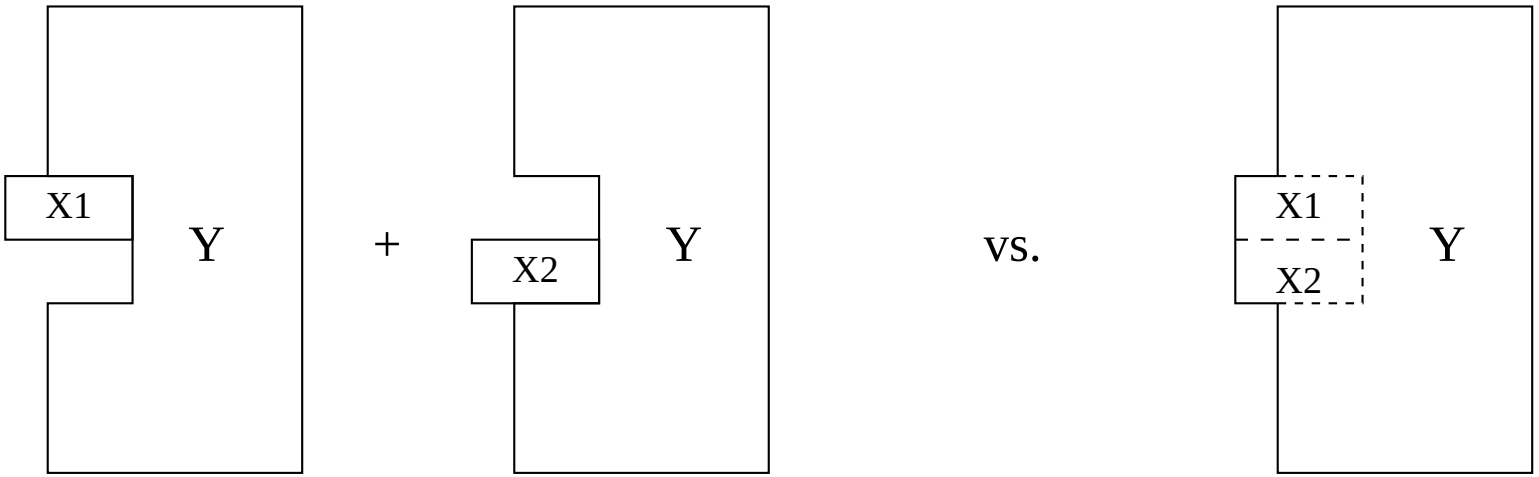

- Could still be difficult if X_1 and X_2 have strong interactions (e.g., jet engine controller and thrust reverser)

Component-Based Certification: The Difficulty

- Most of the effort in safety analysis is not in showing that things work right when all is going well
  - But in showing that the system remains safe in the face of hazards
- In the case of X+Y, we have to consider all the hazards that X can pose to Y and vice-versa
- And similarly with X_1 and X_2
- Hazards include malfunction and unintended function as well as loss of function
- Difficult to anticipate all the hazards X might pose to Y, or X_1 to X_2 without knowing quite a lot about Y or X_2, and vice-versa

Examining The Difficulty

- Compositional design and analysis requires that
  - All interactions are via the interfaces
  - The specifications capture all relevant behavior at those interfaces
- But. . .
- Hazards may bypass the traditional notion of "interface"
  - Think of Concorde's tires
- Hazards may include unspecified behavior over known interfaces
  - Think of babbling on a shared bus
- Compounded by the fact that we are dealing with embedded control systems

Generic Embedded Control System

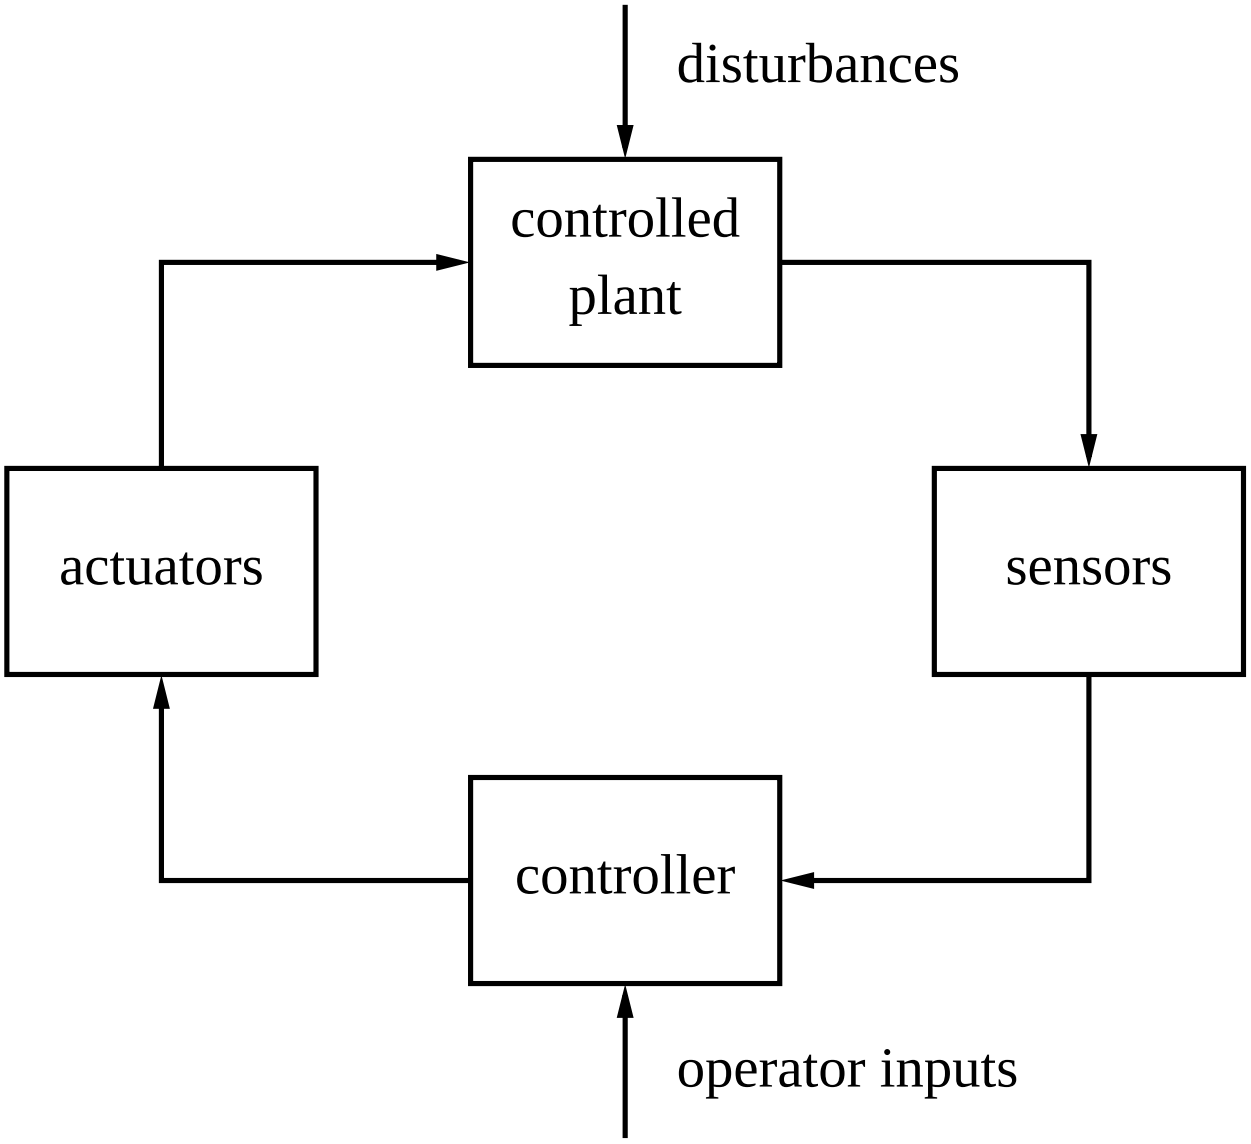

Now Suppose We have Two Of Them

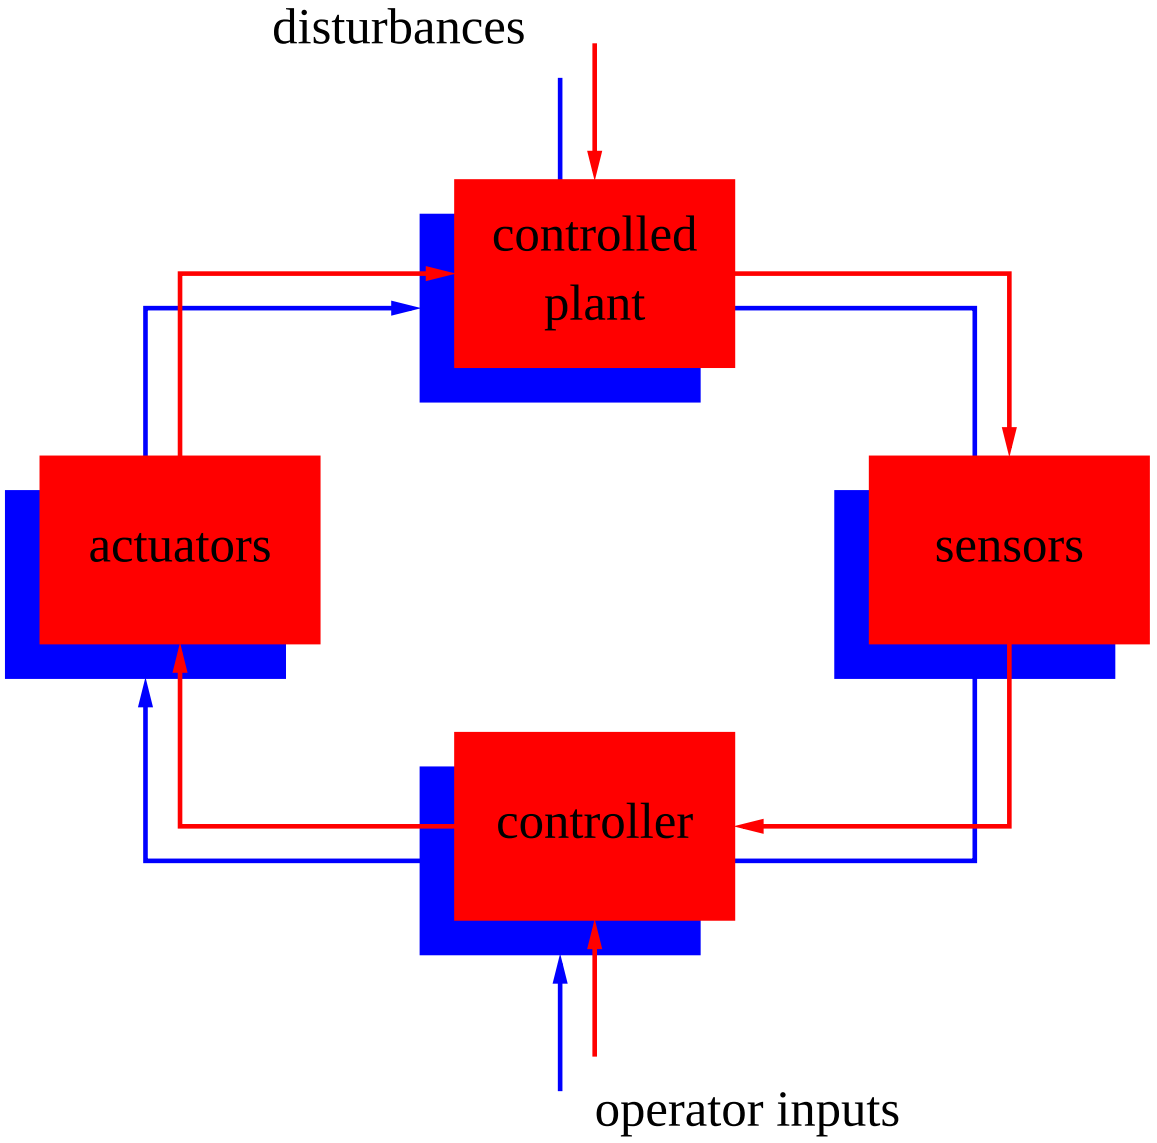

Must Consider Their Interaction (Including the Effects of Faults)

potential interaction
through physical coupling

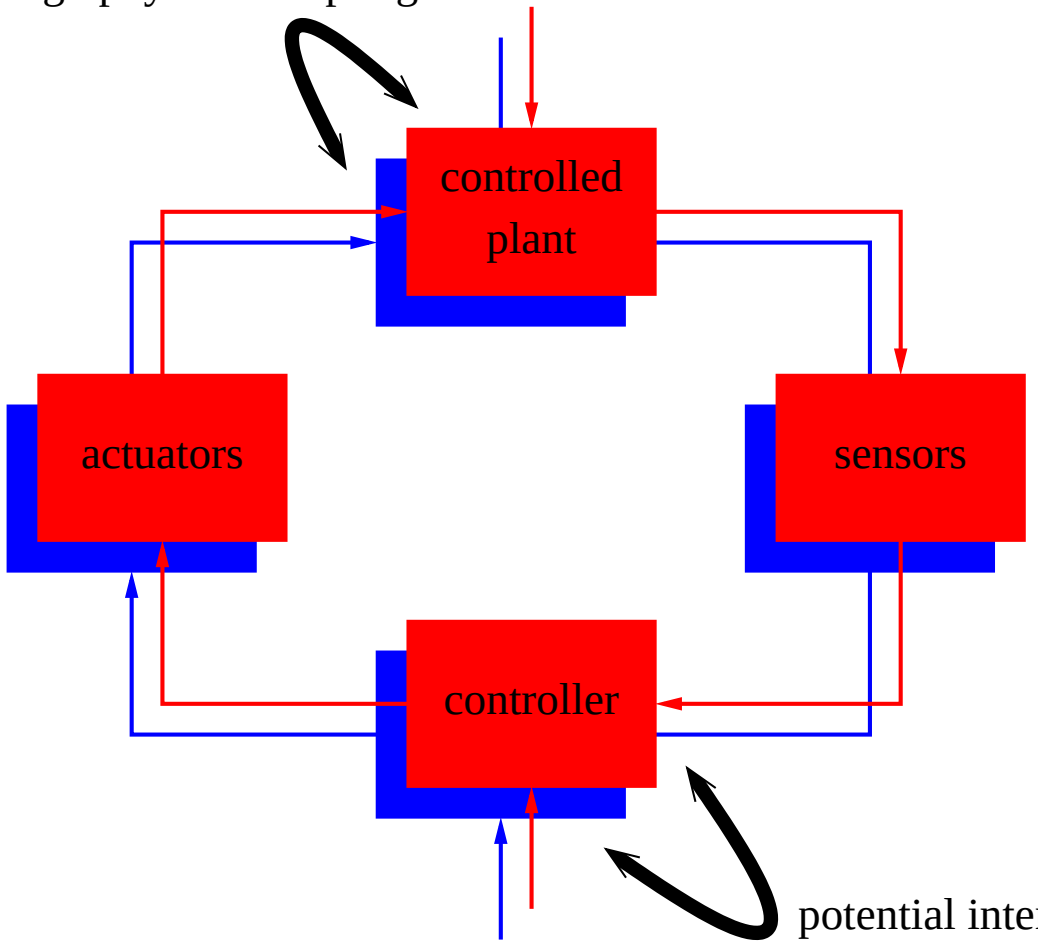

potential interaction
through shared resources

Interaction Through Physical Coupling Of The Plants

- Must be examined by the techniques of hazard analysis
  - FHA (Functional hazard analysis), HAZOP, etc.
  - See ARP 4754, 4761
- E.g., engine controller and thrust reverser
  - No reverse thrust when in flight (system-component)
  - No movement of the reverser doors when thrust above flight idle (component-component)
- Similar for brakes and steering
- No avoiding the need for holistic analysis here
  - Though some component-system and component-component analyses may be routine

Interaction Through Coupling of the Control Systems

- Interaction
  - One controller might provide data to another
  - One controller might assume something about the other

  Could get more complicated if controllers coordinate
  - E.g., integrated brakes, suspension, steering in cars

  Or one compensates for another
  - E.g., extreme toe-in by steering to compensate for brake failure
- Shared resources
  - One controller might monopolize computational or communications resources
  - Or change those assumed private by another
  - Or look at private data (hidden interface)

In The Worst Case, Could Lose All Independence

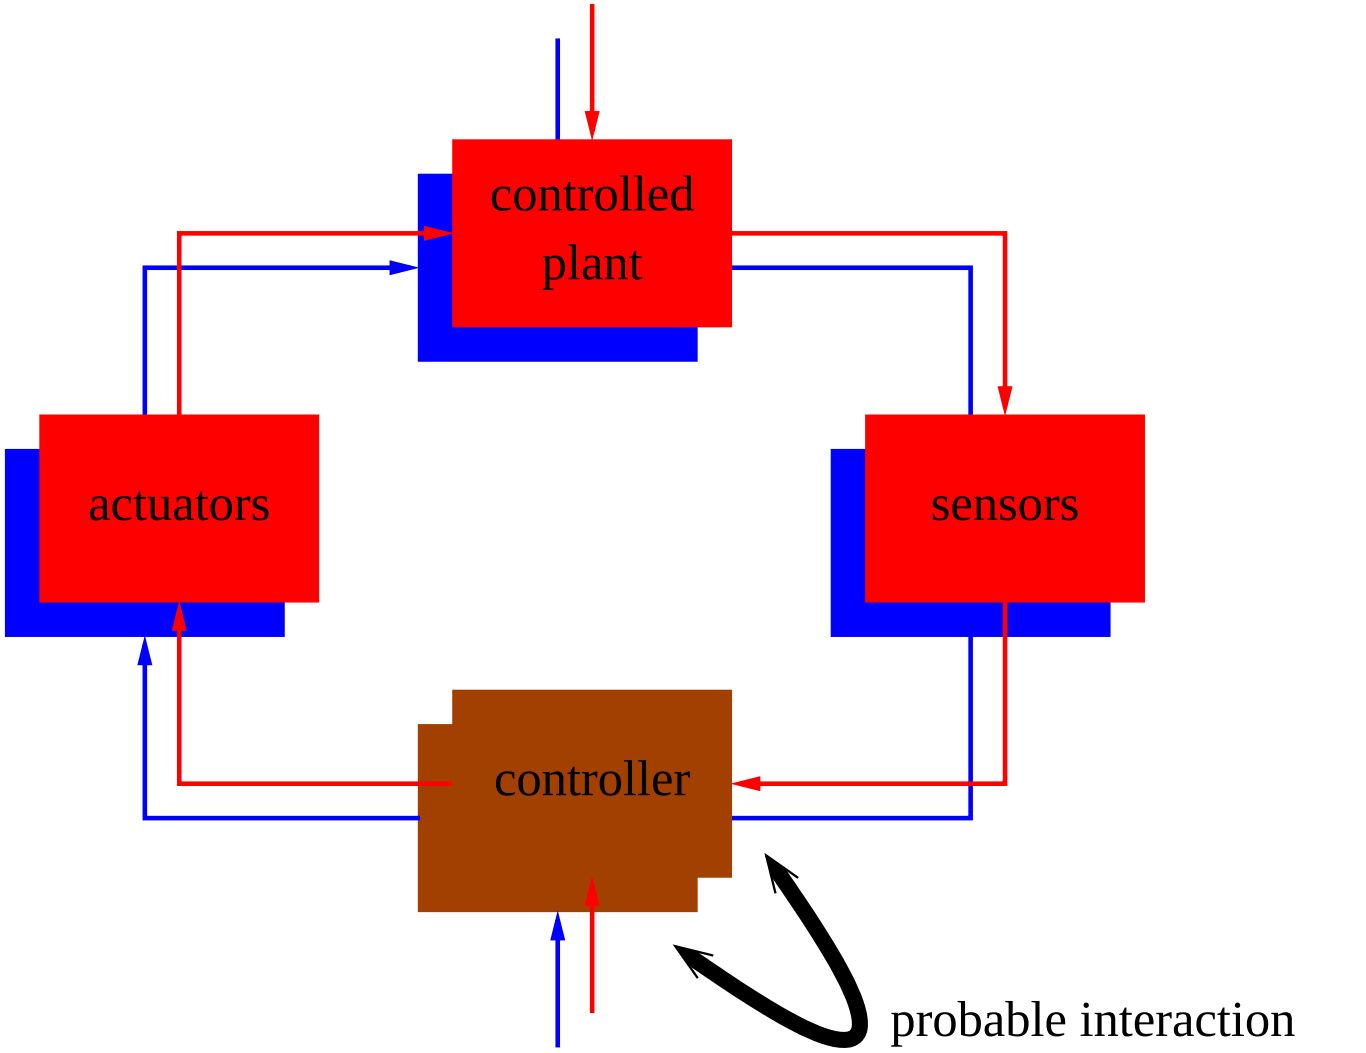

Modular Certification for Combinations of Software

- Need to examine the interactions possible due to integration and resource sharing among separate control software
- New hazards from fault propagation, and unintended emergent behavior
- If we can control fault propagation due to resource sharing
  - Traditionally accomplished by federated architecture
  - Whereas IMA relies on partitioning (see later)
- Then we can focus on the interaction among components
  - Design of each component relies on the interface specifications of the others
  - Each guarantees its specification on basis of assumptions about those of the others

Modular Certification for Combinations of Software

- Suppose X_1 and X_2 are software for separate functions in Y
  - E.g., X_1 is the software for the engine controller and X_2 is the software for the thrust reverser
- Each is certified for its part in Y on the basis of assumptions about what the other does

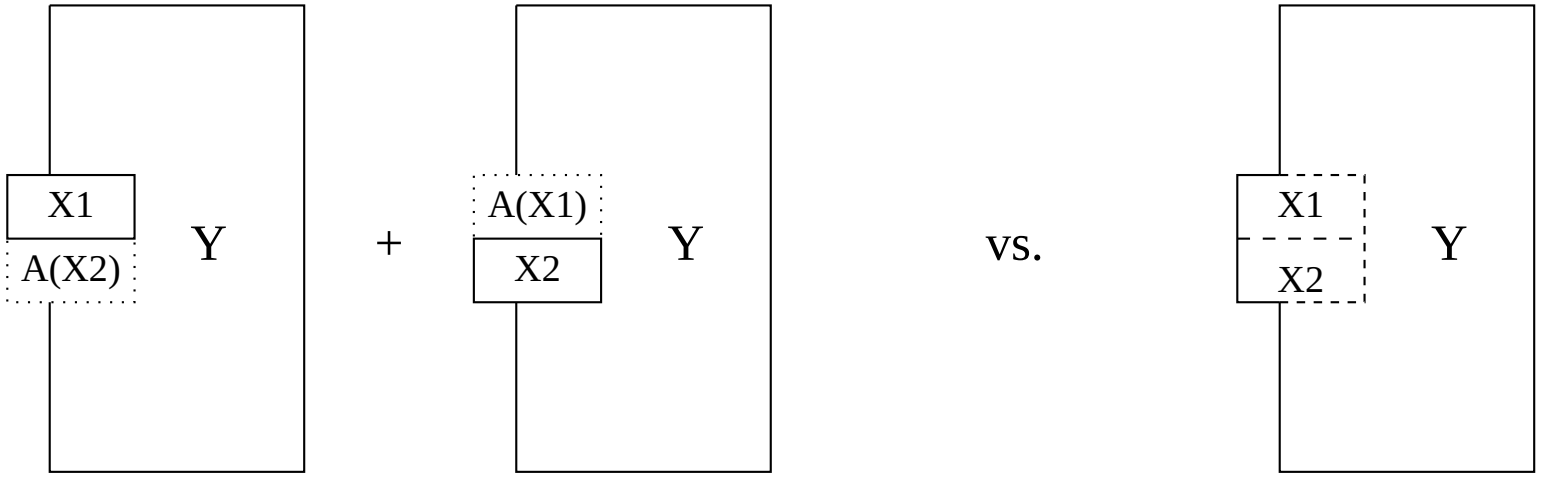

- Want to conclude that their actual combination is safe

Modular Safety Analysis
By Assume-Guarantee Reasoning

- The idea of verifying properties of one component X_1 on the basis of assumptions about another X_2 (and, in general, X_3, $X_4 \ldots$), and vice-versa is called assume-guarantee reasoning and is fairly well developed in computer science
- But there are two challenges:
  - The approach looks circular, so how do we get a sound method? (not pursued here, see final slides)
  - Assume-guarantee reasoning is used for verification, but we are interested in safety analysis and certification

Modular Safety Analysis for Combinations of Software

- Certification differs from verification in that we have to take failures and hazards into account
- Maybe we can split the problem into normal/abnormal cases
  - Can we certify the normal operation of X_1 and X_2 by assume/guarantee methods?
    - ⋆ E.g., the thrust reverser does its thing assuming the engine controller supplies some sensor readings
    - ⋆ Yes, this is similar to other assume/guarantee applications
  - Can we do something similar for the hazards?
    - ⋆ That is, provide guarantees on the behavior of X_1 when X_2 does not fulfill its assumptions?
    - ⋆ This is the hard one!

Assume/Guarantee Reasoning Over Hazards

How can X_2's failure to fulfill its guarantees disturb X_1?

Interface failure: X_2 affects X_1 through channels other than their defined interface

- E.g., X_2 corrupts or monopolizes resources used by X_1 (memory, communications bandwidth)

Behavioral failure: X_1 depends on data or control values from X_2 in such a way that its computation becomes hazardous when X_2 fails to satisfy its guarantees

- E.g., X_1 relies on sensor data from X_2 with no independent backup

Function failure: X_1 cannot guarantee the safety of its function if X_2 allows the system it controls to malfunction

These are the only hazards X_2 can pose to X_1

Controlling The Hazards

- We must control these three classes of hazards
- Some require that X_1 and X_2 have certain properties
- And some require architectural properties of the environment in which X_1 and X_2 operate

The Environment Must Enforce Partitioning

- Interface failure cannot be allowed; X_1, X_2,... must operate within an environment that enforces partitioning
  - Such as the traditional federated architecture
  - Or an IMA or MAC architecture such as SAFEbus or TTA
- Must ensure that no failure of X_2 can affect the basic operation or timing of X_1, nor its communications with nonfaulty X_3, X_4,...
- Top-level requirement specification for partitioning:
  - Behavior perceived by nonfaulty components must be consistent with some behavior of faulty components interacting with it through specified interfaces
- A partitioning architecture must be guaranteed to satisfy this requirement (one of the things TTA does for you)

Normal and Abnormal Assumptions and Guarantees

- In most concurrent programs one component cannot work without the other
  - E.g., in a communications protocol, what can the sender do without a receiver?
- But the software of different aircraft and automobile functions must not be so interdependent
  - In fact, X_1 must not depend on X_2
  - In worst case, X_1 must provide safe operation of its function in the absence of any guarantees from X_2
  - Though may need to assume some properties of the function controlled by X_2 (e.g., thrust reverser may not depend on the software in the engine controller, but may depend on engine remaining under control)

Normal and Abnormal Assumptions and Guarantees (ctd)

- X_1 should provide a graduated series of guarantees, contingent on a similar series of assumptions about X_2
  - These can be considered the normal behavior of X_1 and one or more abnormal behaviors
- A component may be subjected to external failures of one or more of the components with which it interacts
  - Recorded in its abnormal assumptions on those components
- A component may also suffer internal failures
  - Documented as its internal fault hypothesis
- Hypotheses must encompass all possible faults
  - Including arbitrary or Byzantine faults
  - Unless these can be shown infeasible (another thing that TTA does for you)

Components Must Avoid Behavioral and Function Failure

True guarantees: under all combinations of failures consistent with its internal fault hypothesis and abnormal assumptions, the component must be shown to satisfy one or more of its normal or abnormal guarantees.

Safe function: under all combinations of faults consistent with its internal fault hypothesis and abnormal assumptions, the component must be shown to perform its function safely (e.g., if it is an engine controller, it must control the engine safely)

- Where "safely" means behavior consistent with the safety case assumption about the function concerned

Avoiding Domino Failures

- If X_1 suffers a failure that causes its behavior to revert from guarantee $G(X_1)$ to $G'(X_1)$
- May expect that X_2's behavior will revert from $G(X_2)$ to $G'(X_2)$
- Do not want the lowering of X_2's guarantee to cause a further regression of X_1 from $G'(X_1)$ to $G''(X_1)$ and so on

  Controlled failure: there should be no domino effect.

  Arrange assumptions and guarantees in a hierarchy from 0 (no failure) to j (rock bottom). If all internal faults and all external guarantees are at level i or better, component should deliver its guarantees at level i or better

  This subsumes true guarantees

Summary

- We have seen that modular safety analyses and certification depends on three properties
  - Partitioning
  - Safe function
  - Controlled failure (which subsumes true guarantees)
- TTA provides partitioning, hence creates the environment for modular analysis
- Note that when the functions are not coupled, we get preservation of prior properties
- Sheds some light on Perrow's concern over interactive complexity and strong coupling

Future Development

- This is work in progress
- Need to try it out in practice
  - Simplified engine controller/thrust reverser
  - MAC applications (subfunctions rather than functions)
  - Membership and synchronization in TTA
- Should try to derive specific partitioning requirements from the top-level specification

Technical Aside

- What is a suitable formulation of assume-guarantee reasoning?
- $\langle p \rangle X \langle q \rangle$ asserts that if X is part of a system that satisfies p, then the system must also satisfy q

$$\frac{\langle P_2 \rangle X_1 \langle P_1 \rangle \quad \langle P_1 \rangle X_2 \langle P_2 \rangle}{\langle true \rangle X_1 || X_2 \langle P_1 \wedge P_2 \rangle}$$

- Naïve rules like this are unsound

Sound Assume Guarantee Rules

- Consider a specific example: clock synchronization and group membership in TTA
  - Each depends on the other
  - But the circularity is broken by time
  - Membership at time t depends on synchronization at $t-1$
- Similar time offset in dependency of X_1 in this frame on assumptions about X_2 in previous frame
- This suggests that we could modify our previous circular rule to read as follows, where P_j^t indicates that P_j holds up to time t

$$\frac{\langle P_2^t \rangle X_1 \langle P_1^{t+1} \rangle \quad \langle P_1^t \rangle X_2 \langle P_2^{t+1} \rangle}{\langle true \rangle X_1 || X_2 \langle P_1 \wedge P_2 \rangle}$$

A Practical Sound Circular Rule

- The idea is good, but we want a rule that supports systematic reasoning from one time point to the next
- A formulation that has this character has been introduced by McMillan; here H is a "helper" property, $\Box$ is the "always" modality of Linear Temporal Logic (LTL), and $p \rhd q$ means that if p is always true up to time t, then q holds at time $t+1$ (i.e., p fails before q)

$$\frac{\langle H \rangle X_1 \langle P_2 \rhd P_1 \rangle \quad \langle H \rangle X_2 \langle P_1 \rhd P_2 \rangle}{\langle H \rangle X_1 || X_2 \langle \Box(P_1 \wedge P_2) \rangle}$$

Soundness and Utility

- Have formally verified this rule in PVS
  (for the 2-process case)
- Notice that $p \rhd q$ can be written as the LTL formula $\neg(p \,\mathsf{U}\, \neg q)$, where U is the LTL "until" operator, so the antecedents can be discharged by model checking in the case of finite systems

Completeness

- McMillan's rule is sound, but Namjoshi and Trefler show that it is incomplete
- They show that the following extended version of the rule is both sound and complete

$$\frac{\langle H\rangle X_1\langle (A_2 \wedge P_2) \rhd ((A_2 \supset P_1) \wedge A_1)\rangle \quad \langle H\rangle X_2\langle (A_1 \wedge P_1) \rhd ((A_1 \supset P_2) \wedge A_2)\rangle}{\langle H\rangle X_1 || X_2\langle \Box(P_1 \wedge P_2)\rangle}$$

  Here A_1 and A_2 are auxiliary properties; by taking both to be *true*, this more complicated rule reduces to the previous one

- Probably not an issue in practice