Capturing Safety Intent Through Assurance Cases

Arnab Ray, PhD
Senior Research Scientist
Fraunhofer USA Center for Experimental Software Engineering

Rance Cleaveland, PhD
Professor,
Department of Computer Science
University of Maryland

Fraunhofer USA
Center for Experimental Software Engineering

Safety Intent

- Assertion: My system is safe
- Argument: I have safety requirements, I followed a development standard, I did some testing (here are my tests)

What's wrong with this picture?

The Problem

- Severe gaps in the safety arguments
  - What is safety in the context of your system? [Validation]
  - How did you come up with the safety requirements? [Validation]
  - How much can we trust you? [Validation]
  - How does your testing connect to the safety requirements? [Verification]

A Proposed Solution

- Assurance Cases for Safety (aka safety cases)
  - *Assurance goal*
  - *Context*
  - *Evidence that the goal has been satisfied*:
  - *Strategy*
    - Links the evidence to the goal in a logically consistent and coherent manner

This Talk

- Lays out a safety case framework that argues for “safety” in a comprehensive way
- Illustrates our framework with a medical device example

Definition of Safety (Medical Domain)

- Safety Intent: Does not harm the patient (i.e. it cannot do something bad)
  - e.g. introduce an air bubble into bloodstream
  - definition of safety provided by regulation

Example Used

- The Generic Infusion Pump (GIP) Project
- Goal: Create an exemplar set of hazards, requirements, models for GIPs
- Example: GPCA (Generic Patient Controlled Analgesic Pump)

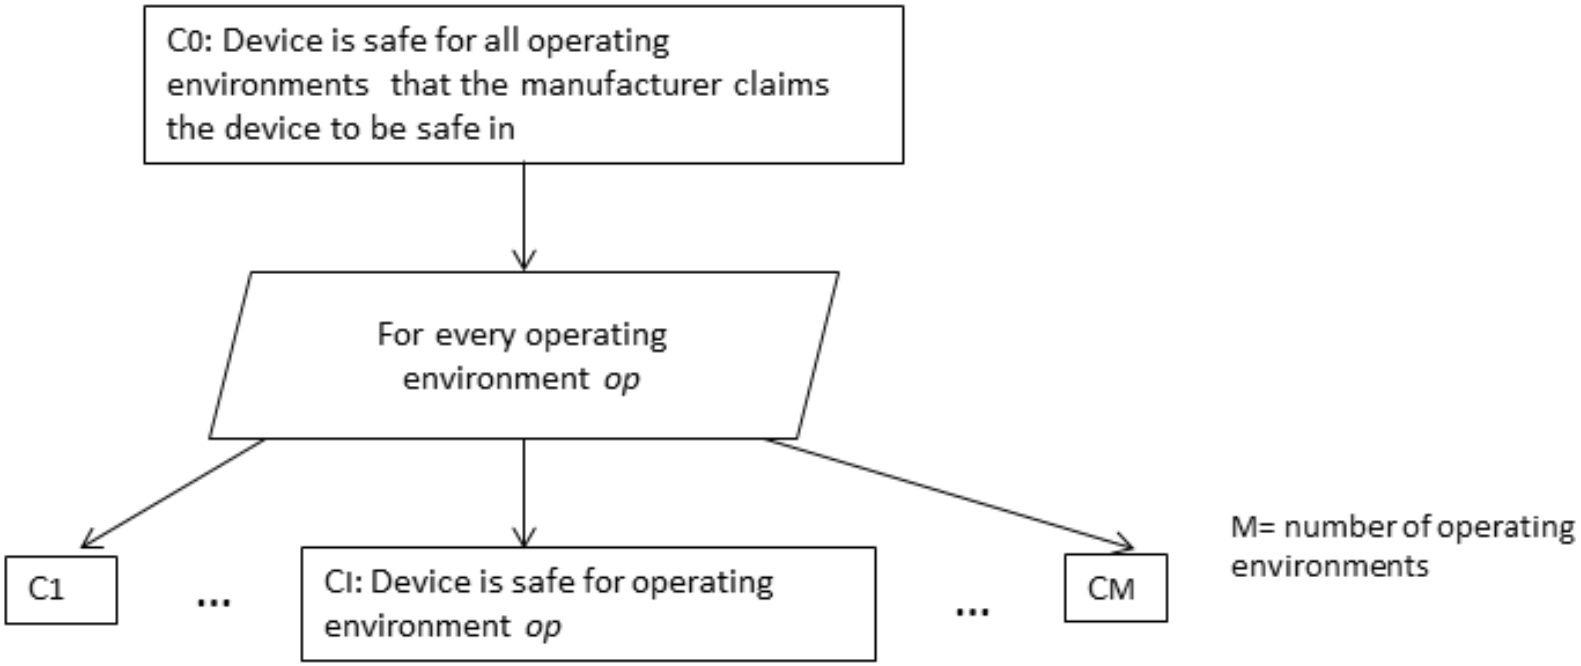

Note: Safety arguments vary by operating environment

A PCA pump safe for home may not be safe in a moving van !

What is Safe?

- In order to claim a device does nothing “bad”
  - Comprehensively define “bad” (bad=anything that causes injury or death to human beings i.e. hazards)

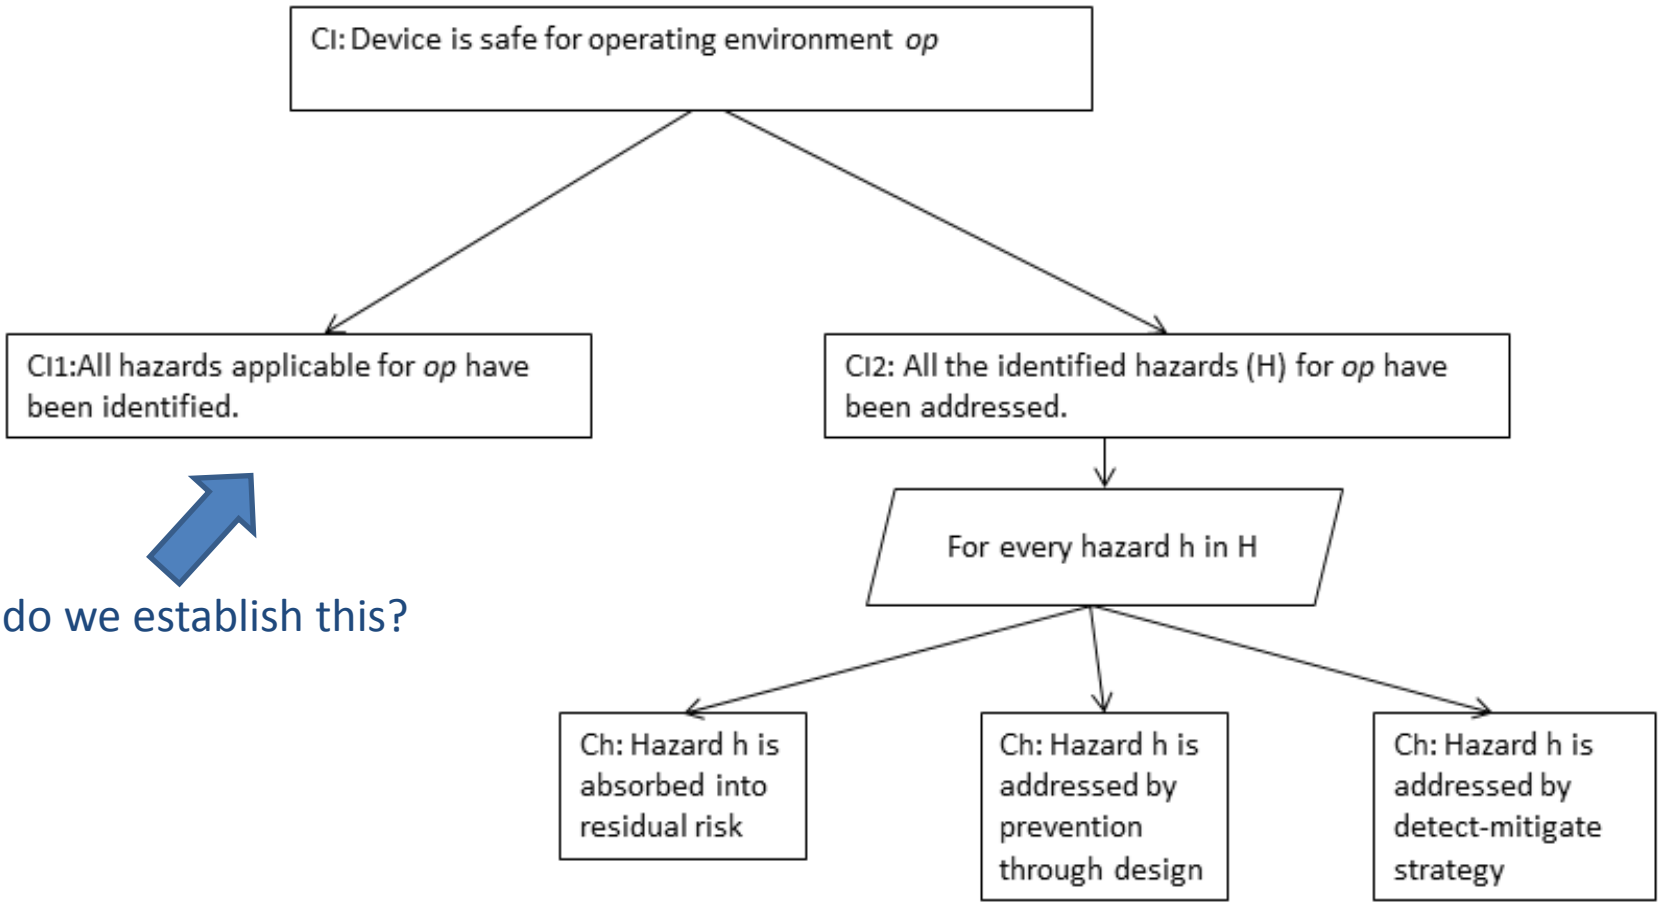

How do we establish this?

All hazards?

- Theoretically impossible to claim all hazards have been identified
- Strategies for arguments
  - Reference to standards
  - Past adverse events (“We handle all of them”)
  - Predicate device (“We handle same set of hazards as this product on market”)

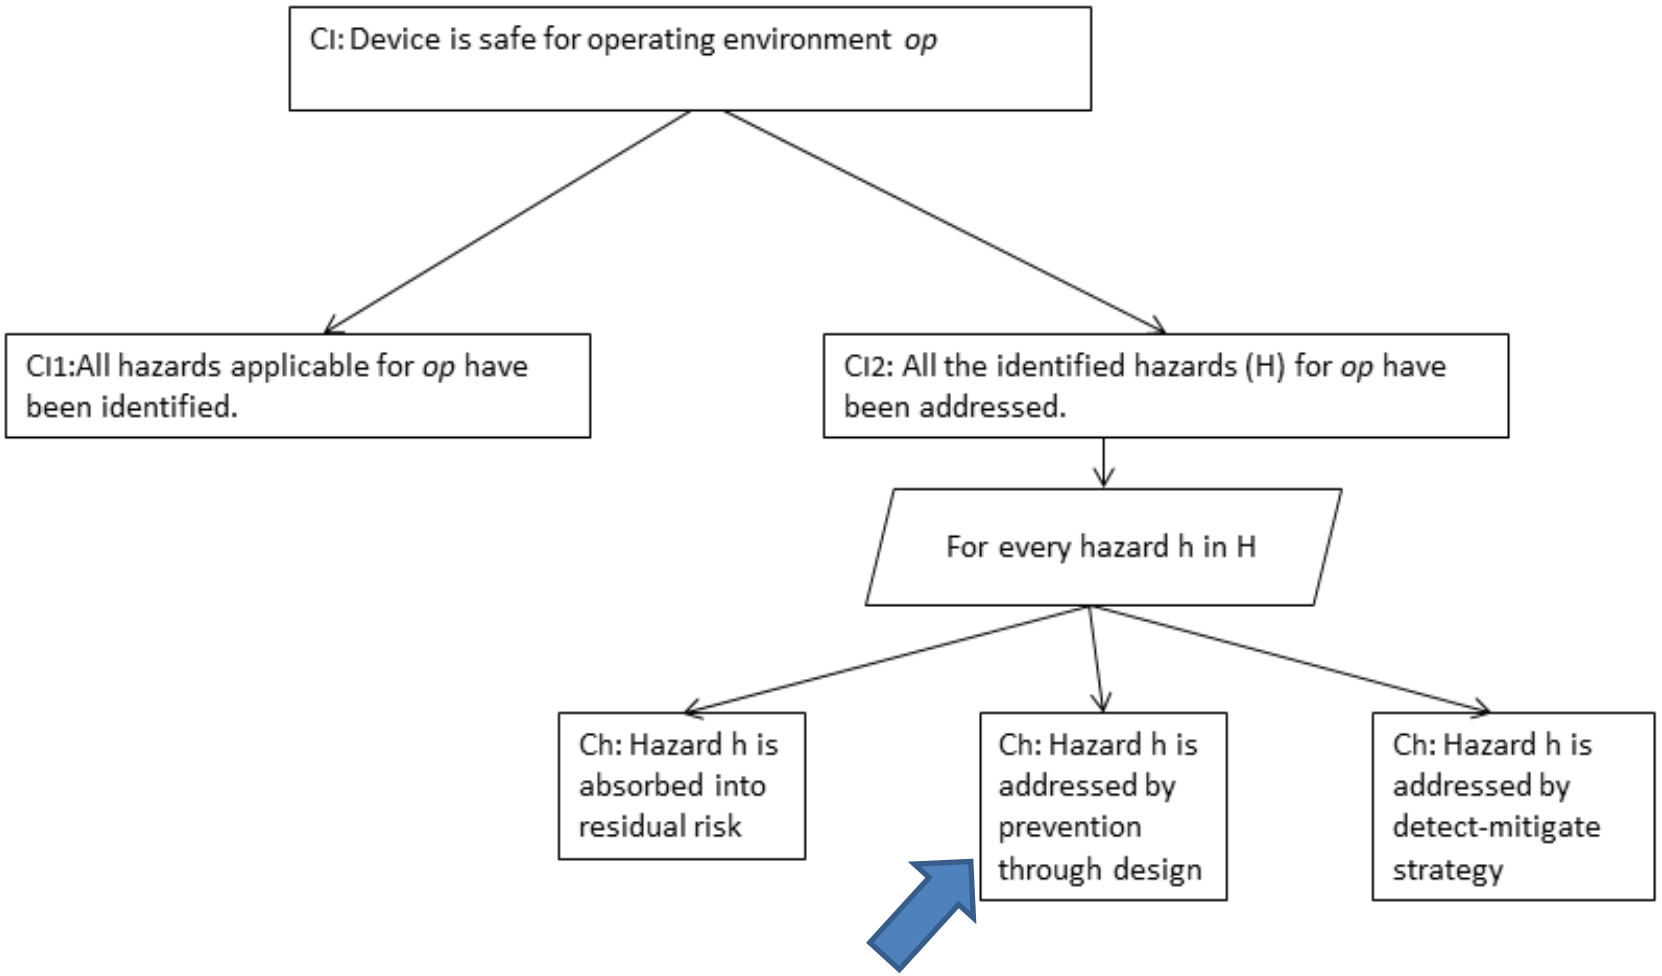
CI: Device is safe for operating environment *op*
CI1:All hazards applicable for *op* have been identified.
CI2: All the identified hazards (H) for *op* have been addressed.
For every hazard h in H
Ch: Hazard h is absorbed into residual risk
Ch: Hazard h is addressed by prevention through design
Ch: Hazard h is addressed by detect-mitigate strategy

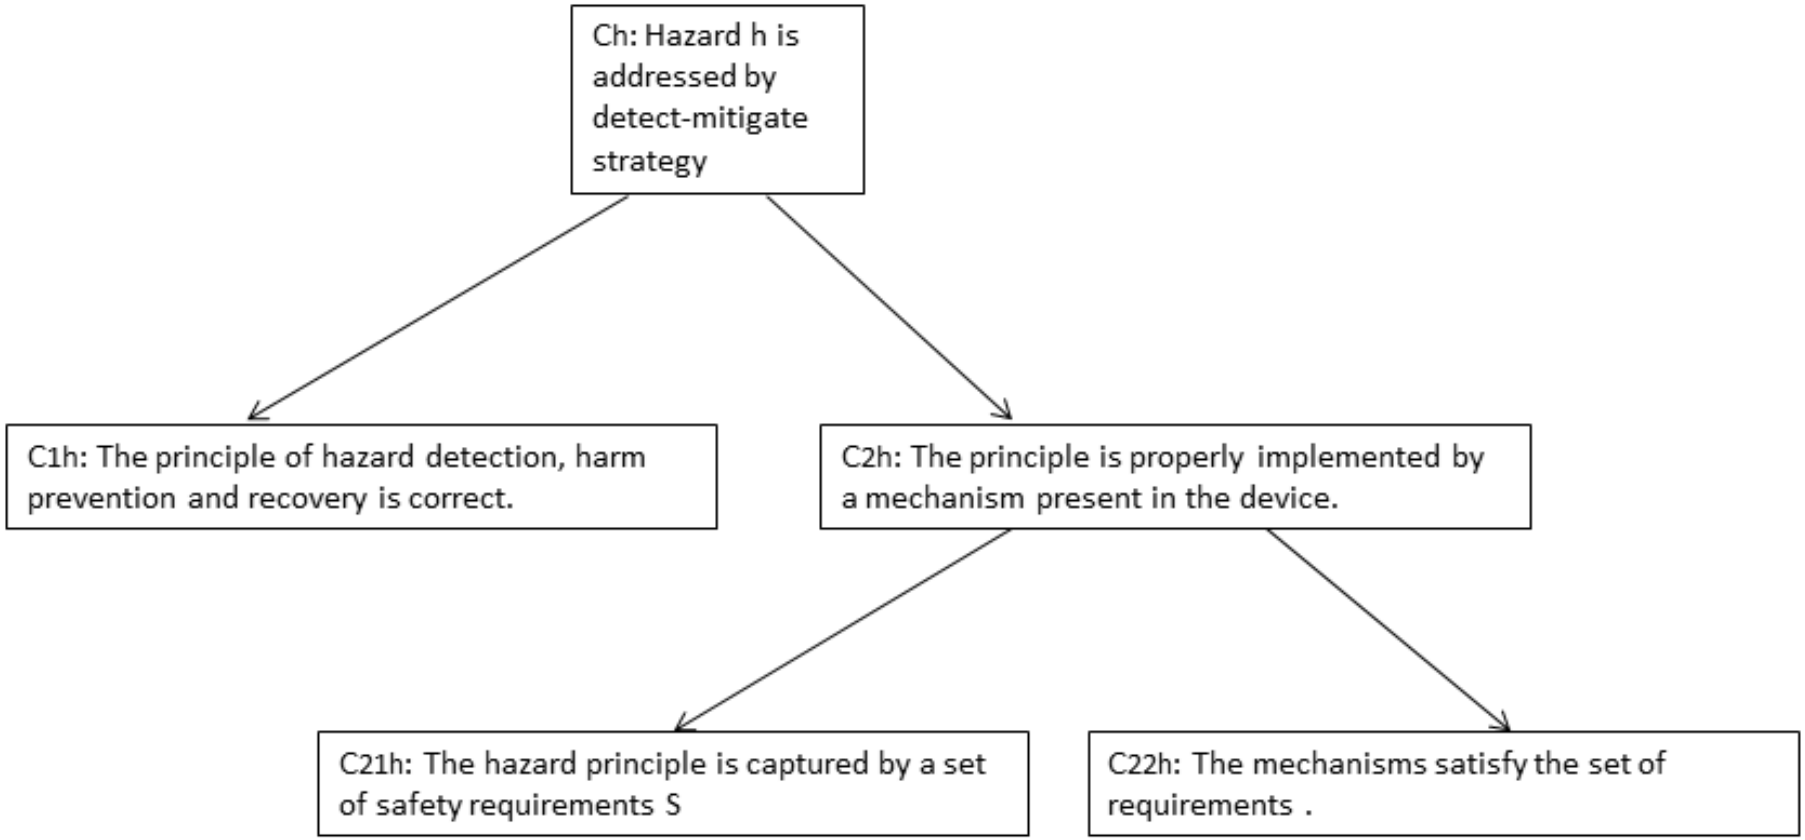
Ch: Hazard h is addressed by detect-mitigate strategy
C1h: The principle of hazard detection, harm prevention and recovery is correct.
C2h: The principle is properly implemented by a mechanism present in the device.
C21h: The hazard principle is captured by a set of safety requirements S
C22h: The mechanisms satisfy the set of requirements .

Example

- The principle: “If bubble size is greater than X microns, then hazard air-in-line has occurred. The patient is not impacted if infusion is stopped before bubble reaches bloodstream and he is notified ”
  - Need to establish that this principle is correct

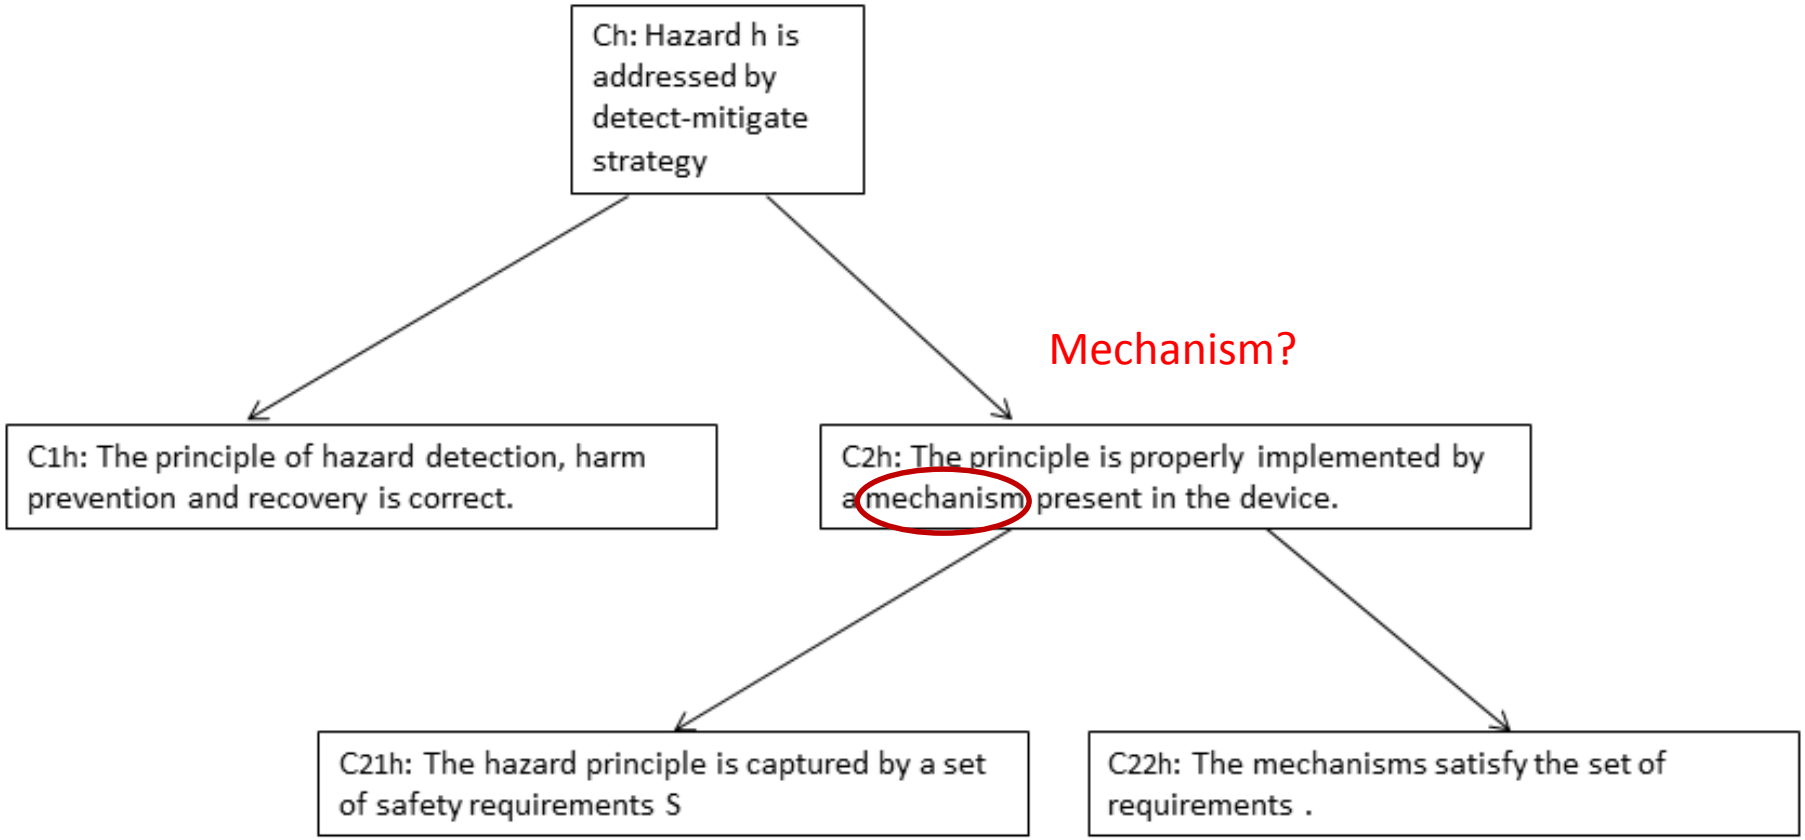
Ch: Hazard h is addressed by detect-mitigate strategy
Mechanism?
C1h: The principle of hazard detection, harm prevention and recovery is correct.
C2h: The principle is properly implemented by a mechanism present in the device.
C21h: The hazard principle is captured by a set of safety requirements S
C22h: The mechanisms satisfy the set of requirements .

Mechanism

- Exclusively mechanical or electrical
- Exclusively software (e.g. a range check for drug safe limits)
- Combination of all of them (mechanical + electrical+ software)

Example

- Sensor is mechanism that detects bubble size
- Once safe limit is crossed, signal goes to software controller
- Controller
  - sends message to alarm module
  - stops mechanical pump

Proof Obligations?

- Entire mechanism is able to detect bubble size appropriately
- (Time from bubble introduction to detection) + (Time from detection to stoppage of infusion)< Safe limit such that bubble does not reach bloodstream

Safety Requirements

- A number of mechanism-specific constraints on implementations

- R1: *An air-bubble must be detected by sensor within "t" time units of its introduction*.
- R2: *The controller software can transition from an infusion mode to an alarming mode within "s" time units of hazard detection by sensor.*
- R3: *No infusion should be possible in the alarming mode*.
- R4: *An alarm should be sufficiently loud to be heard*.
- R5: *The time between the detection of an air-bubble and its entry into the patient's bloodstream is more than s+t time units.*

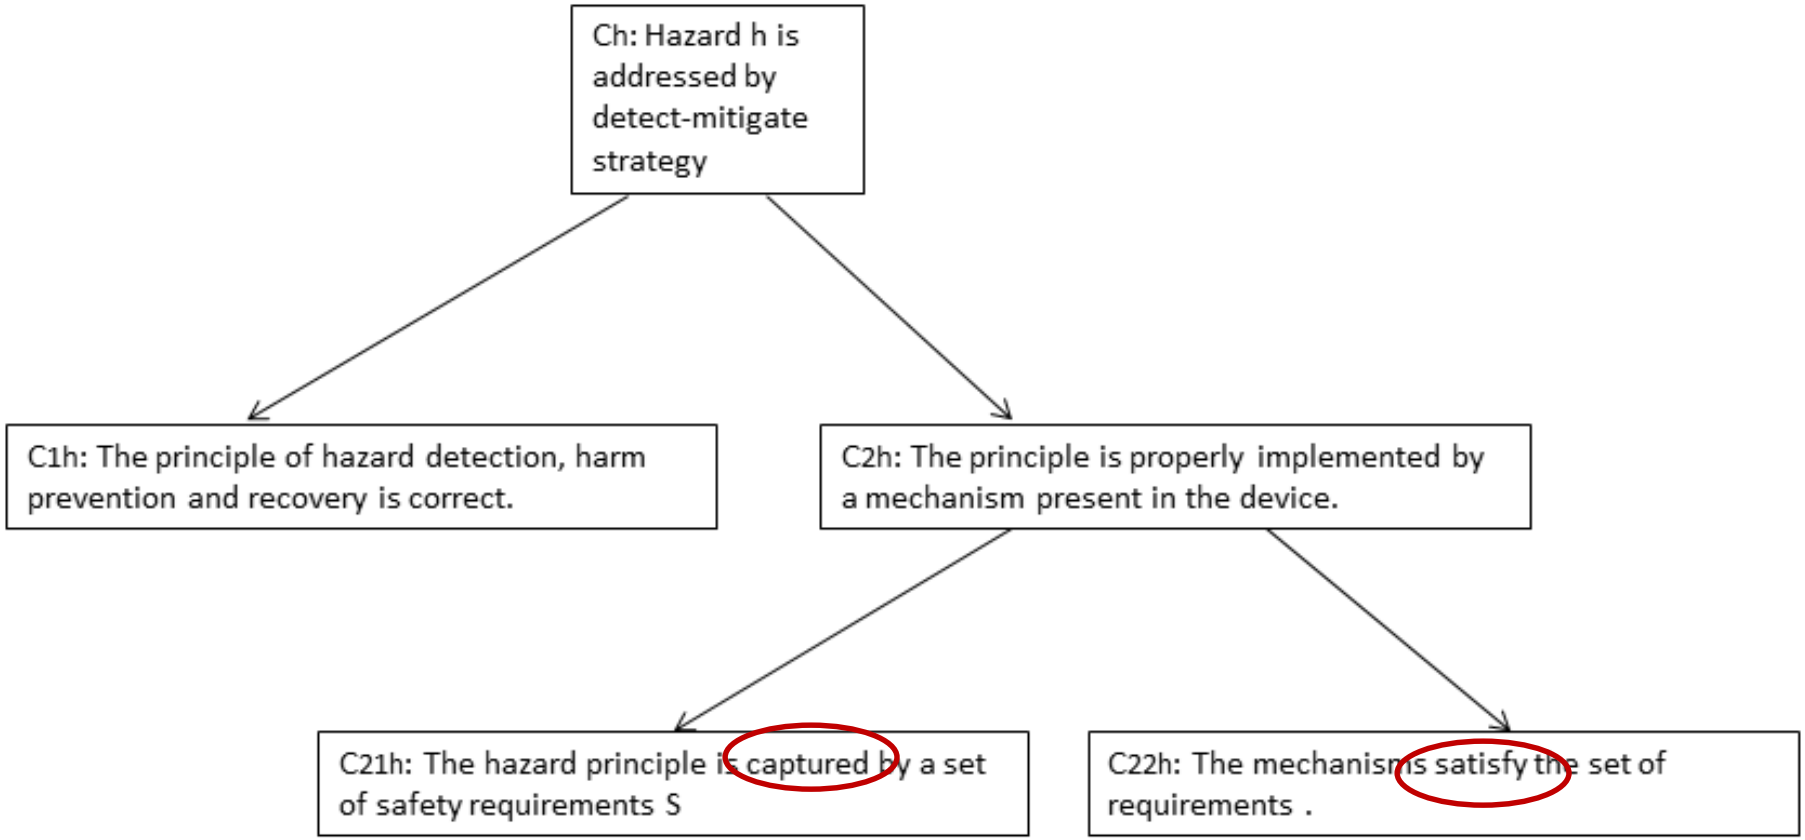
Ch: Hazard h is addressed by detect-mitigate strategy
C1h: The principle of hazard detection, harm prevention and recovery is correct.
C2h: The principle is properly implemented by a mechanism present in the device.
C21h: The hazard principle is captured by a set of safety requirements S
C22h: The mechanisms satisfy the set of requirements .

Safety Requirements

- Set of safety requirements
  - is relevant (no safety requirement not linked to a hazard)
  - is exhaustive (all aspects of the principle of hazard detection, harm prevention and recovery have been translated to requirements)
  - is trustworthy (the safety requirements are internally consistent i.e. do not contradict each other)

Mechanisms Satisfy Requirements

- Depends on the mechanism as to how its behavior is captured
  - Behavior of fully mechanical & electrical systems can be captured by specifications (motor speed, voltage rating etc)
- Software systems are more problematic

More Sub-claims

- "The software system satisfies the set of safety requirements" may broken down into sub-claims with a development standard (e.g. IEC 62304) as reference
  - One sub-claim for every step of the process
  - Overall compliance with standard

Conclusions

- You can’t start at safety requirements
- You need to document every step of the reasoning chain
- You need to arrange it in a safety case

Supplementary Slides Follow

The Regulatory Process

- **510(k)**: device to be marketed is as **safe** and **effective**, that is, substantially equivalent (SE), to a legally marketed device that is not subject to premarket approval (PMA)
- **PMA**: Approval based on a determination by FDA that the PMA contains sufficient valid scientific evidence that provides reasonable assurance that the device is **safe** and **effective** for its intended use or uses

The Food And Drug Administration

- Federal body charged with the responsibility of “protecting the public health by assuring the safety, efficacy and security of human and veterinary drugs, biological products, **medical devices**, our nation’s food supply, cosmetics, and products that emit radiation”

External Infusion Pumps

- An infusion pump infuses fluids, medication or nutrients into a patient's circulatory system
- Problematic class of devices responsible for a number of adverse events every year
- Includes insulin pumps, patient-controlled analgesic pumps

Manufacturers & Assurance Cases

- “More regulatory overhead”
- “Do I have to redo everything I have in terms of pictures?”
- “Where should I start?”
- “What would be acceptable evidence for the FDA?”
- “How deep should we argue?”

Our Thesis

- In any “approval worthy” device submission, the safety assurance case already exists, albeit in an implicit and undocumented form
- Safety assurance case: Formally and explicitly codifies the logical trail of reasoning for a device’s safety

The Paper

- Outlines an approach for safety assurance case argumentation
  - Goal: Serves as the logical glue for different parts of the submission