Byzantine Filtering

- From WG10.4 in Siena
  (and our 2003 SAFECOMP and 2004 IEEE DASC papers)
  - Byzantine fault propagation "physics" and example
  - Combating Byzantine Generals' fault propagation
    - Masking (blocks Byzantine signals via dominant logic)
      - Two-of-Three voter example
      - Can be done only with ***completely*** independent sources (***completely*** independent sources are very rare)
    - Filtering (converts a Byzantine signal to non-Byzantine)
      - Buried within all real Byzantine tolerance mechanisms
      - Needs to be tested to determine coverage
  - Byzantine filter testing idea
  - But, can this be done with a practicable number of tests?
    - How can proof-of-coverage testing be reduced?
      - An answer that reduces amplitude test range 

- Braided Ring: A network to exploit Byzantine filtering 

Digital Circuitry Behavior

There is no such thing as digital circuitry, …
there is just analog circuitry driven to extremes.

This allows the possibility of a digital logic signal being “1/2”.

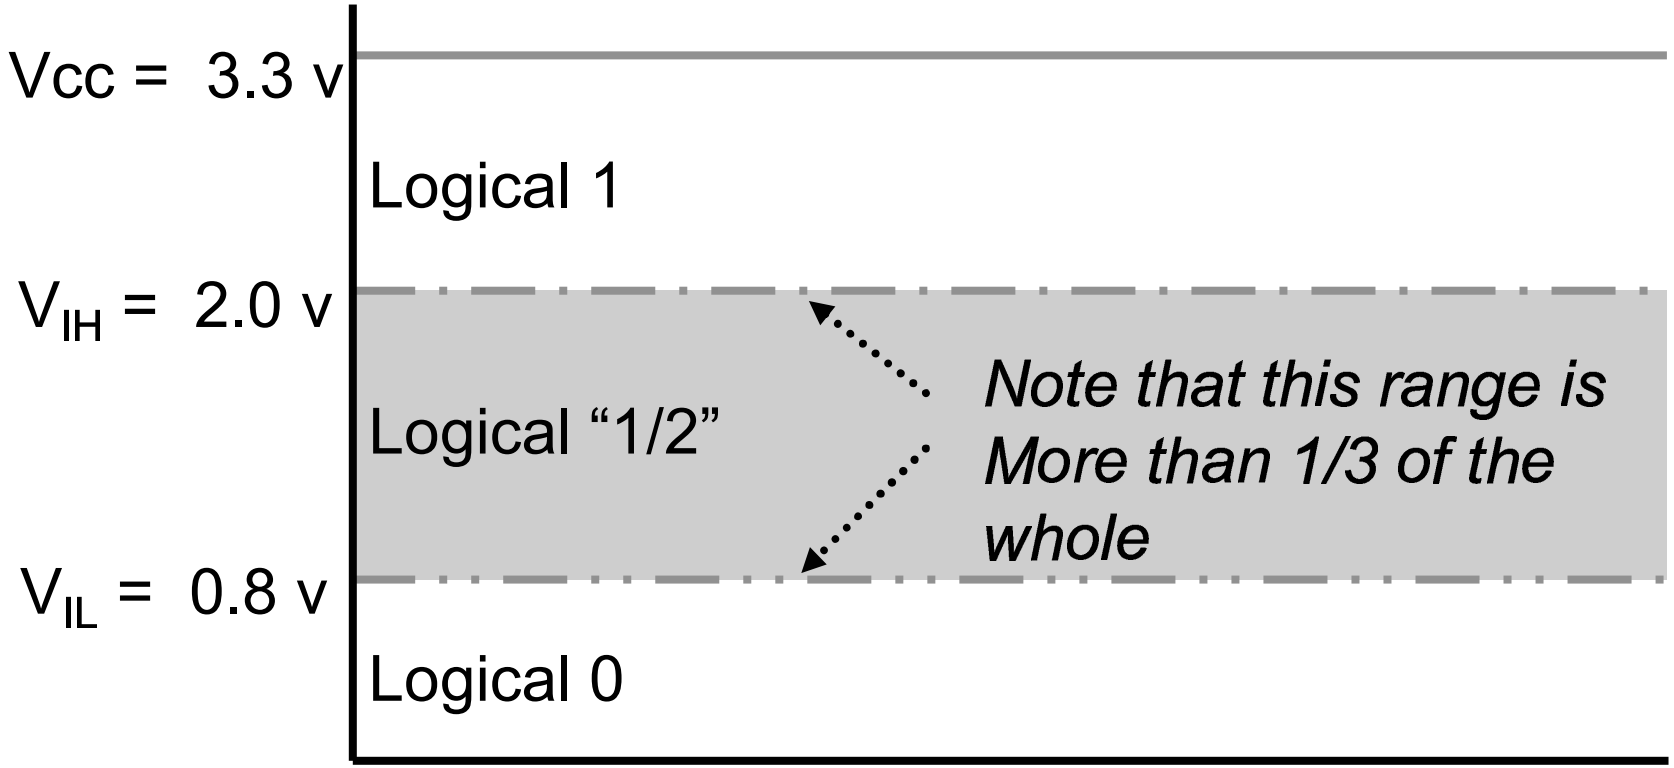

Logic Gate Transfer Function with “1/2” Noise

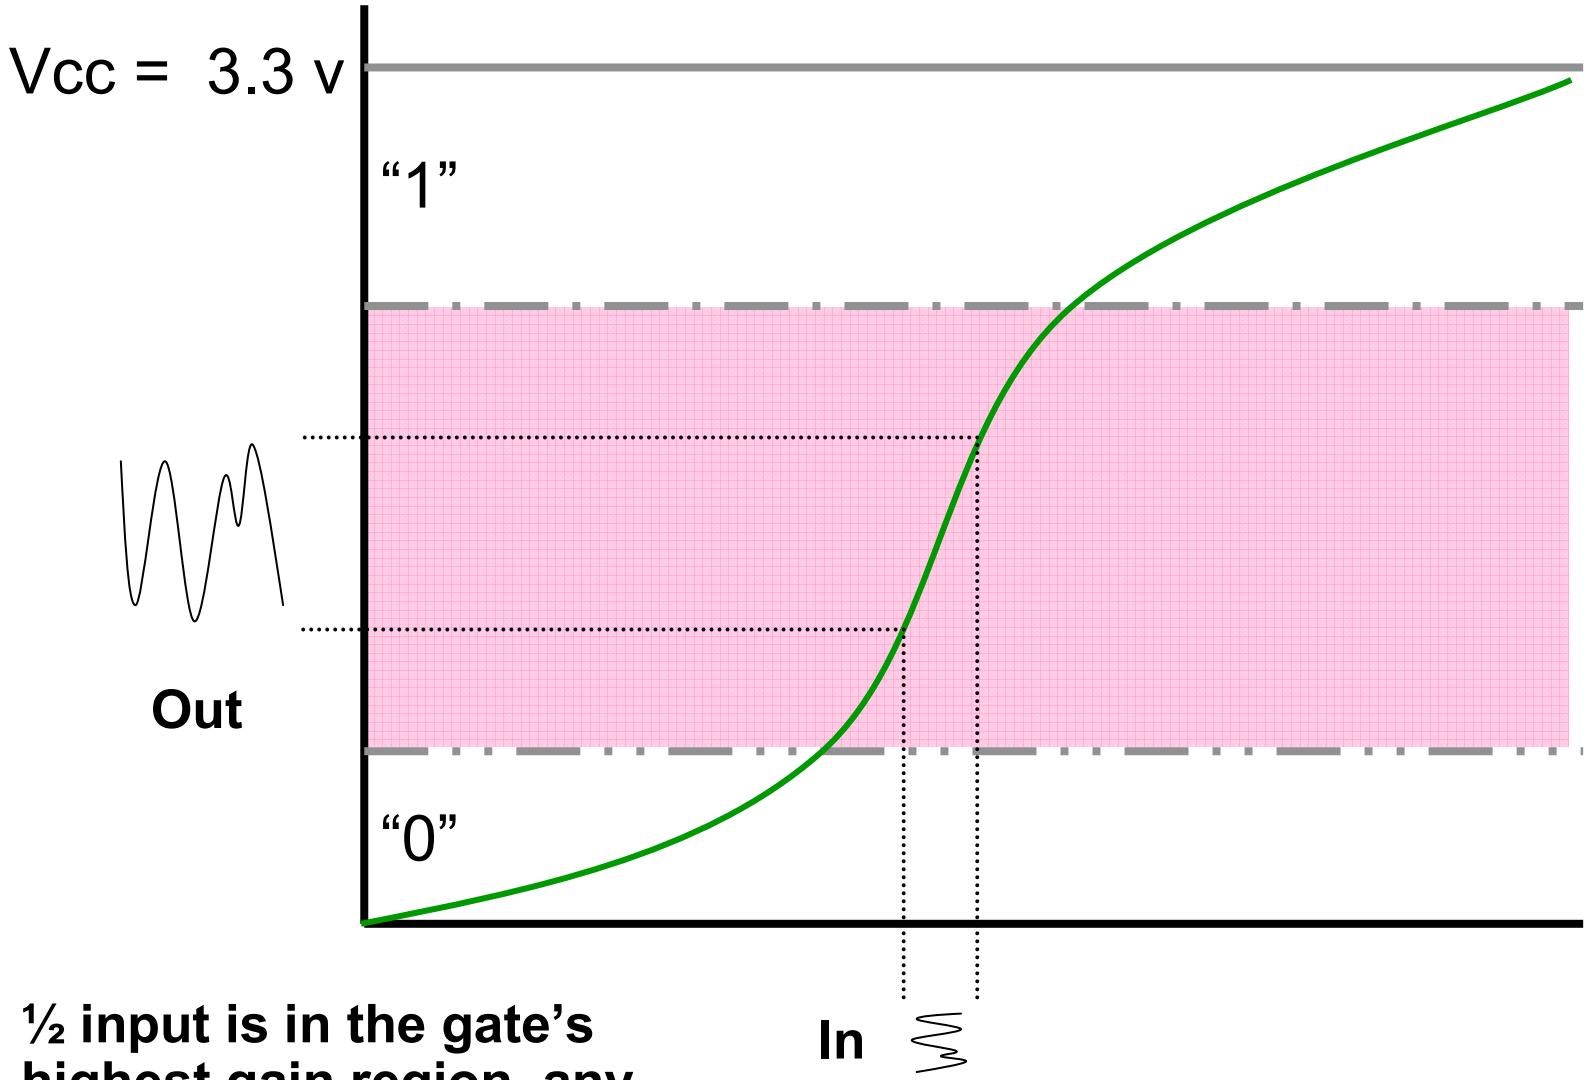

½ input is in the gate’s highest gain region, any noise is greatly amplified

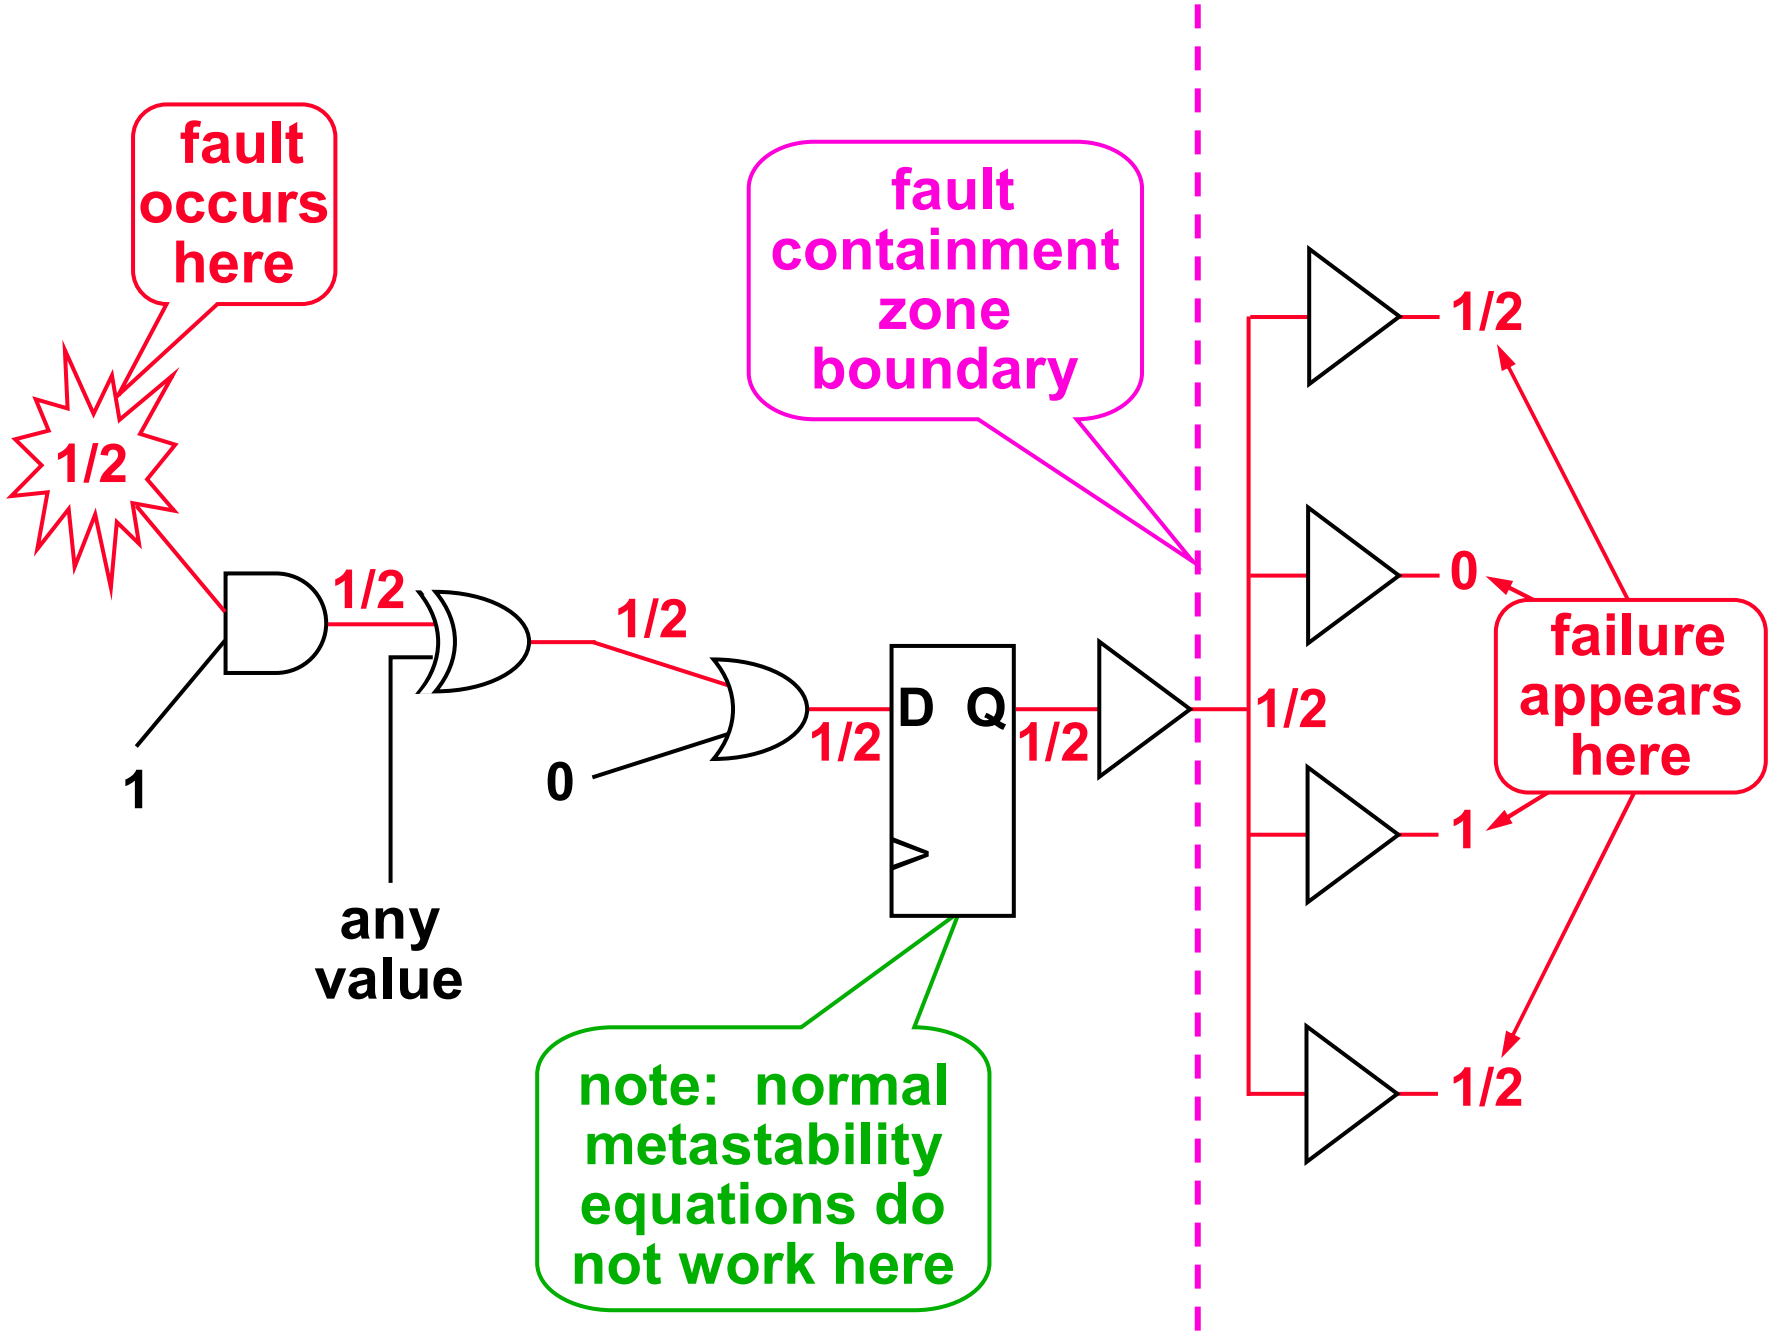
fault occurs here
1/2
1
1/2
any value
1/2
0
1/2
D
Q
1/2
fault containment zone boundary
1/2
0
1
1/2
failure appears here
note: normal metastability equations do not work here

Byzantine Filtering

- Bit-level (waveform) filtering
  - Implemented with any combination of:
    - Schmitt triggers
    - Synchronizers
      (same as used to mitigate metastability)
    - Glitch filters
    - . . . (almost any technique to reduce noise)
  - Perfect coverage impossible
  - Need to determine coverage of implementation
    - Typical pessimistic system Byzantine failure probability is 10^{-5} (10 nodes, 10 critical components in each node with 10^{-7} probability of failure)
    - Typical system requires < 10^{-10} probability of failure
    - Typical coverage needs to be 0.99999

“Byzantine” Fault Injection

Concept:

Create a suitably representative set of faulty waveforms

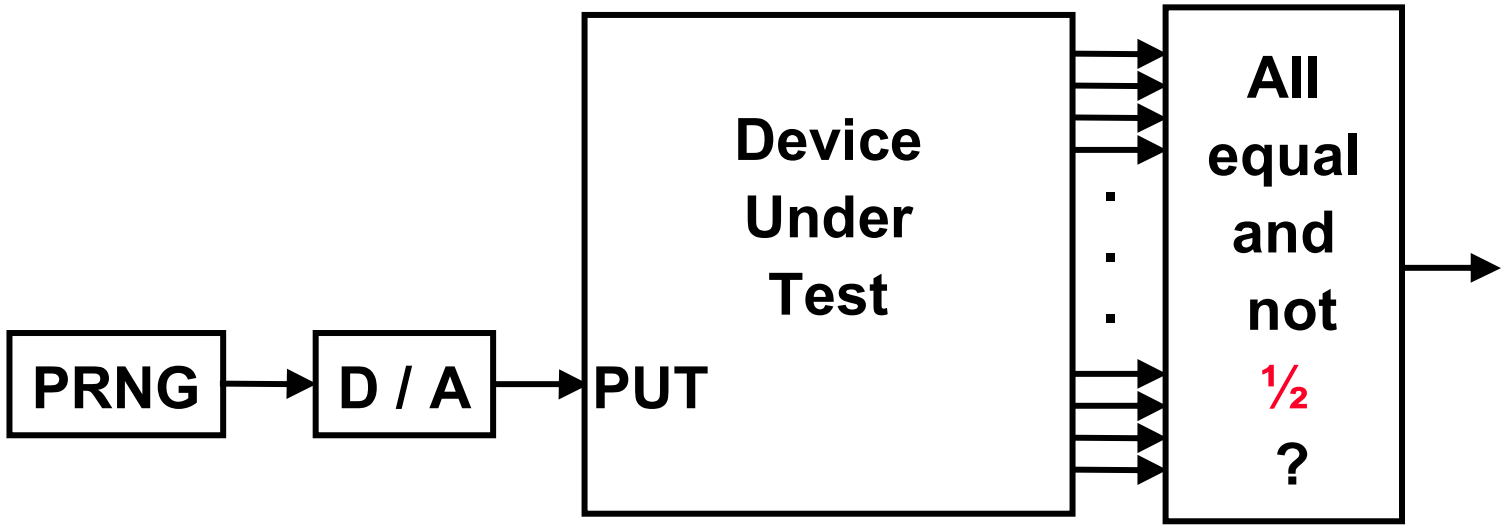

Acronyms:

PRNG = Pseudo Random Number Generator
D / A = Digital to Analog Converter
PUT = Port Under Test

Bit Cell Decision Threshold

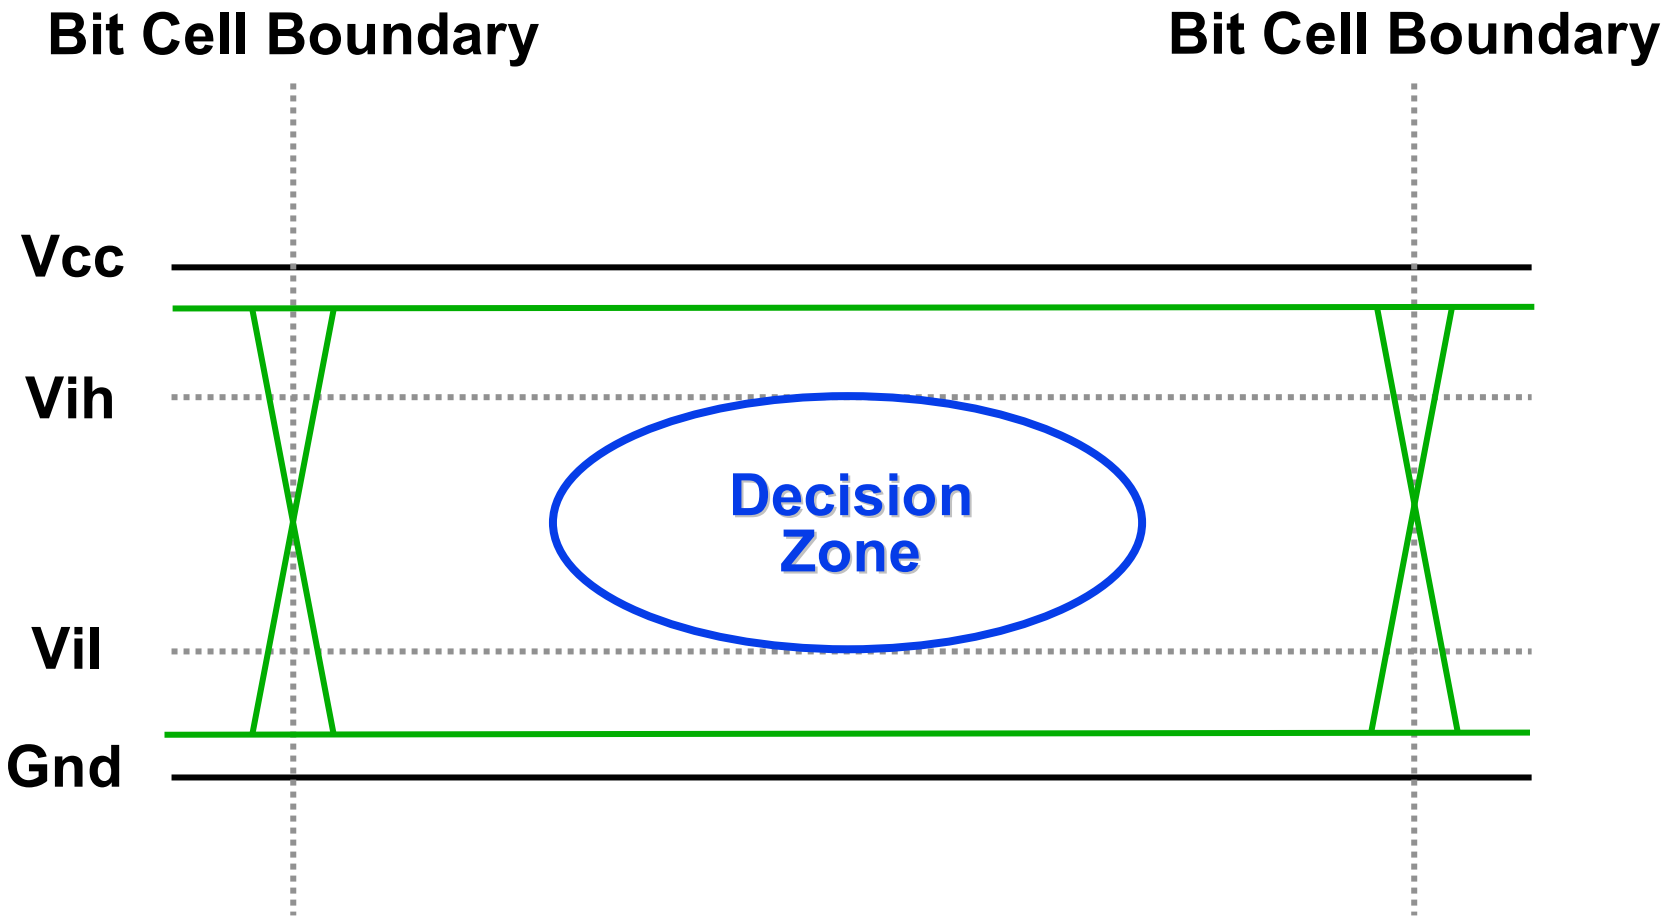

Remove Test Sample Points Past Hold Time

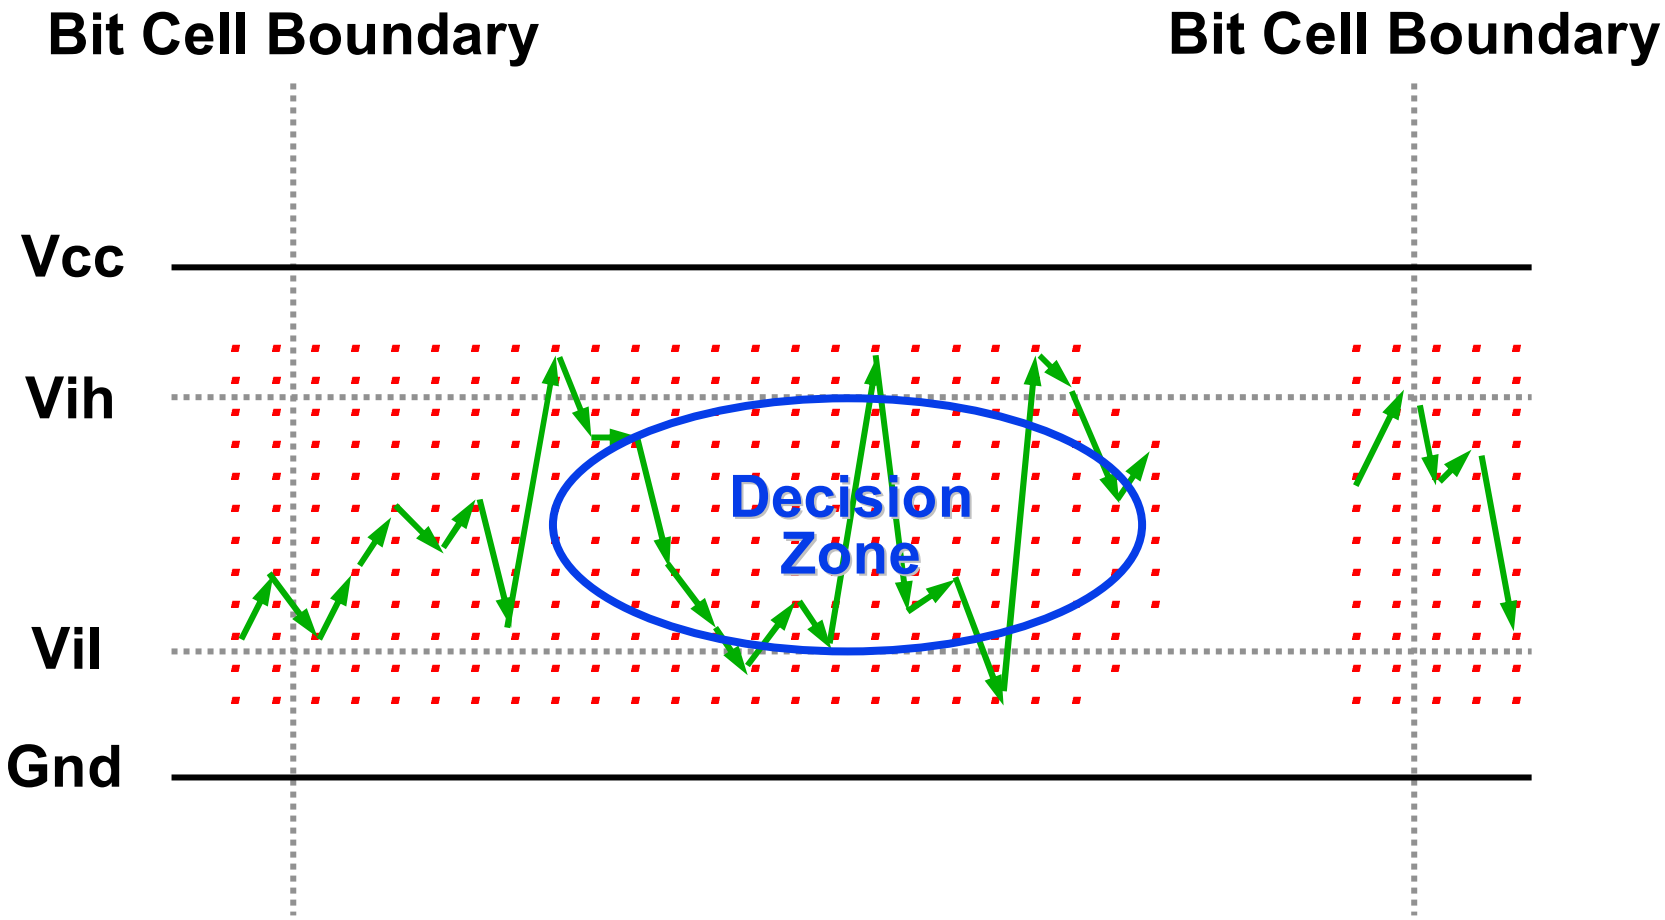

Test Amplitude Reduction with Known Threshold

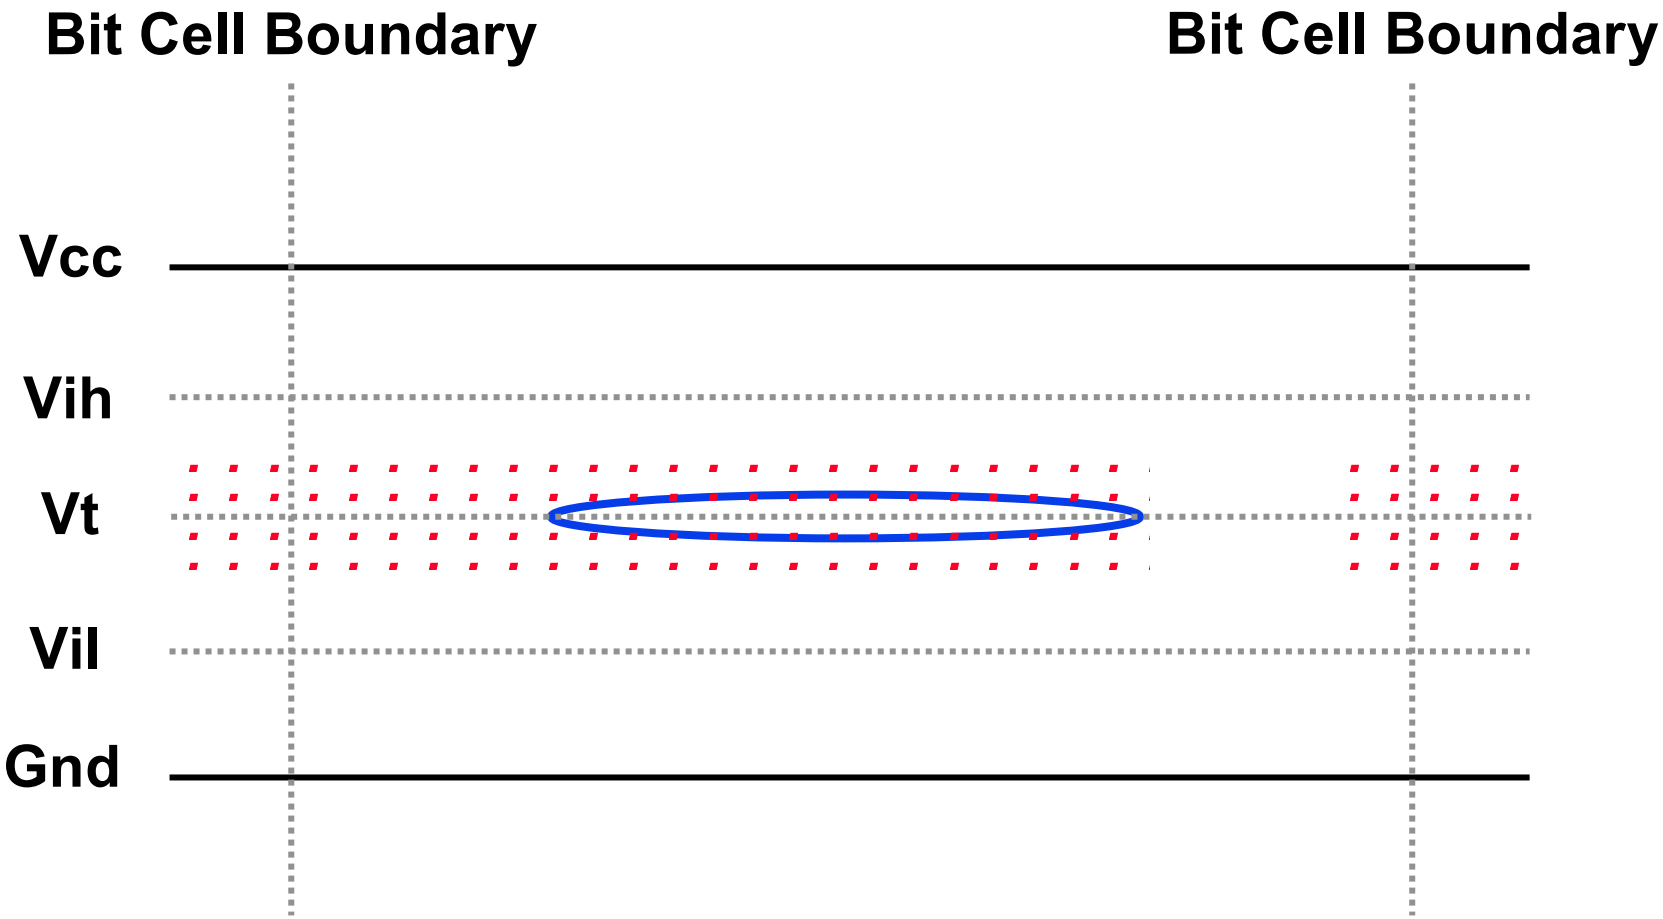

“1/2” DUT Output Signal Rejection

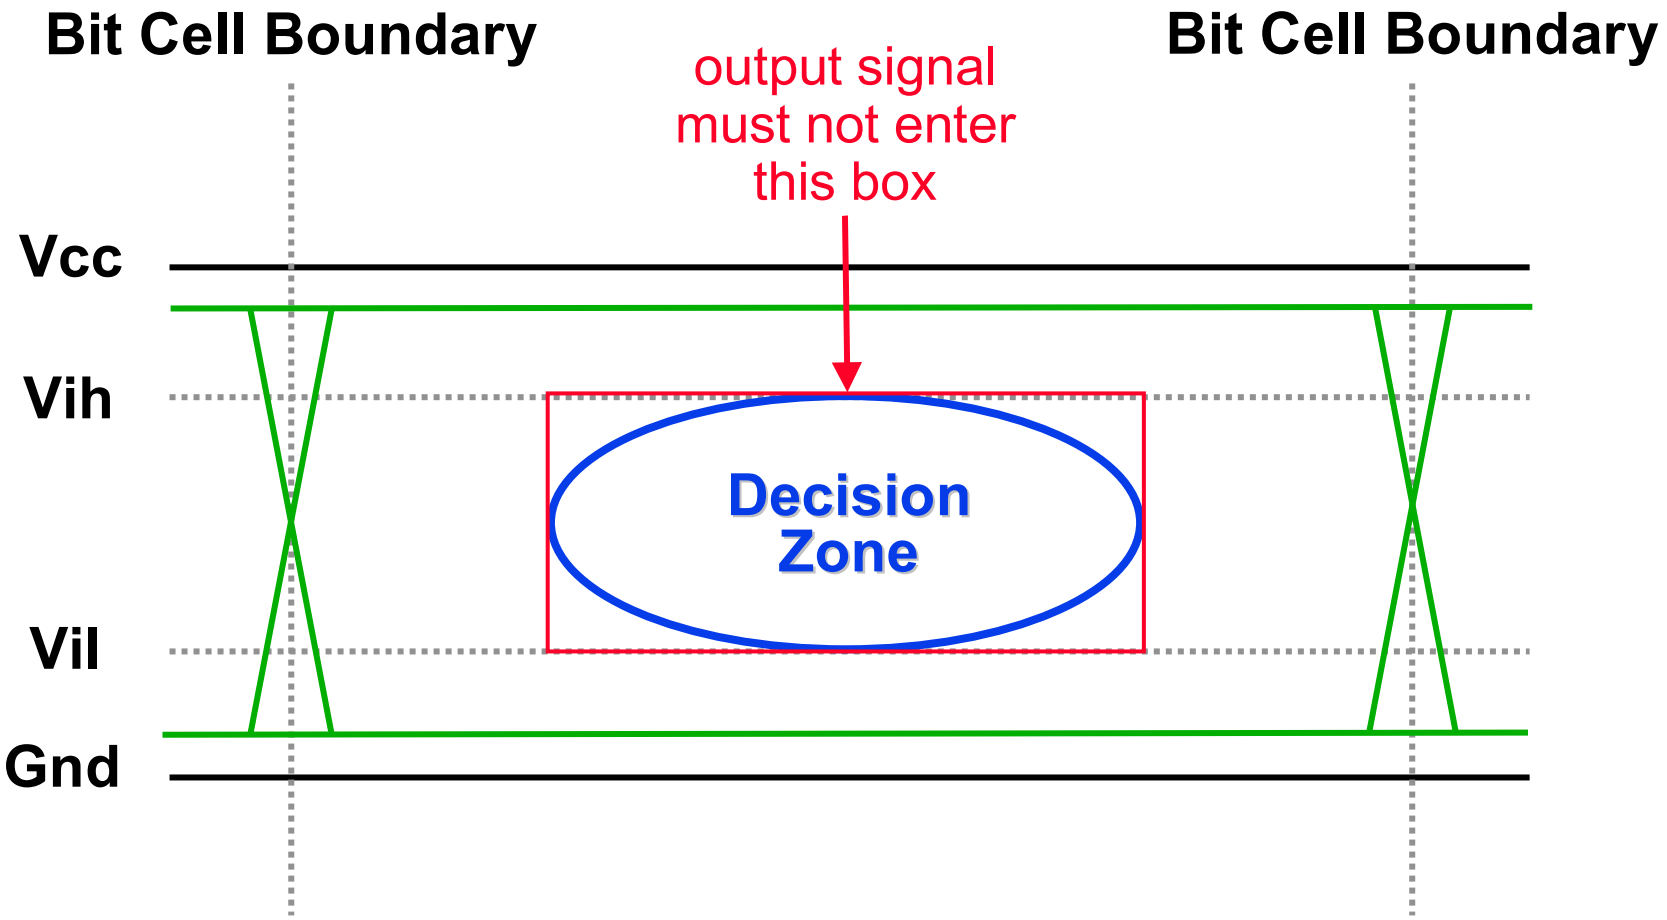

How to Find a Device's Input Threshold

- Connect a device's output back to its input such that the loop has an odd number if inversions in it.
  - This creates an oscillator.
- Add an integrator with a very large time constant.
  - This filters out the oscillations.
- Integrator's output settles on a value which is the input's threshold voltage.

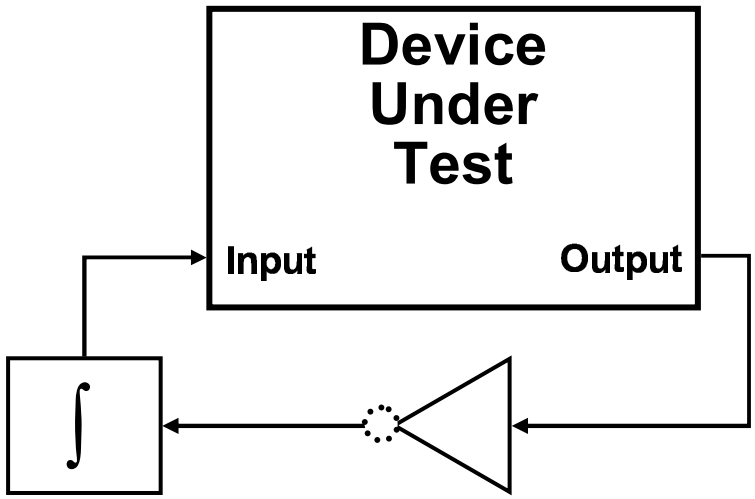

Completion of the Tester

- Add the pseudo-exhaustive bit pattern trajectories to the input feedback (with a reduced number of amplitude test points).
- Either latch the integrator's output before applying the test patterns or make sure the test pattern's are "DC balanced" over the time constant of the integrator.

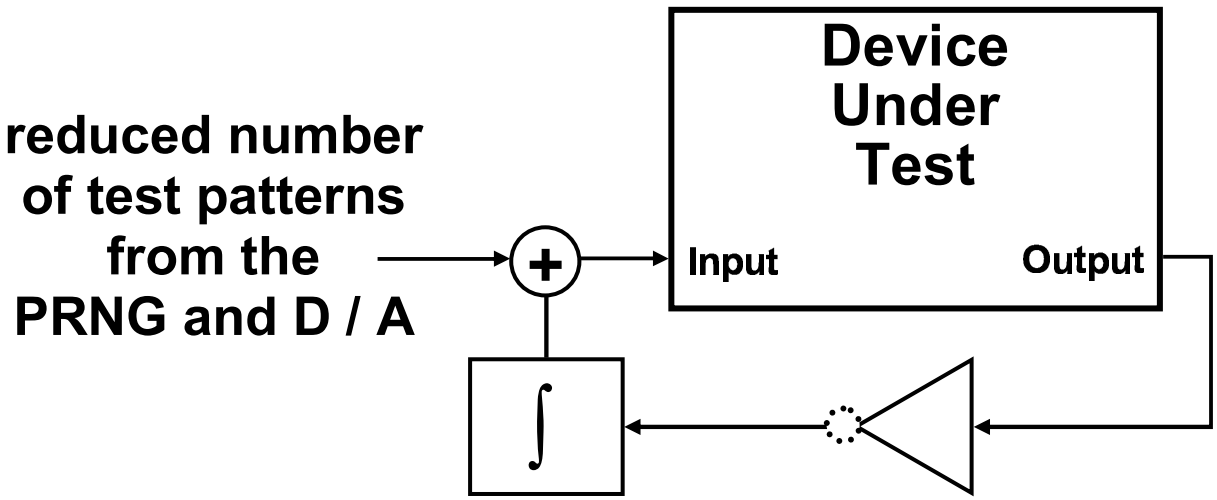

The Path of Low-Cost Dependable Systems

- Redundant Bus
  - Not sufficient due to spatial proximity faults (unavoidable via routing)
  - Serious issues with babbling and masquerade failures
    - Local Guardians not truly independent

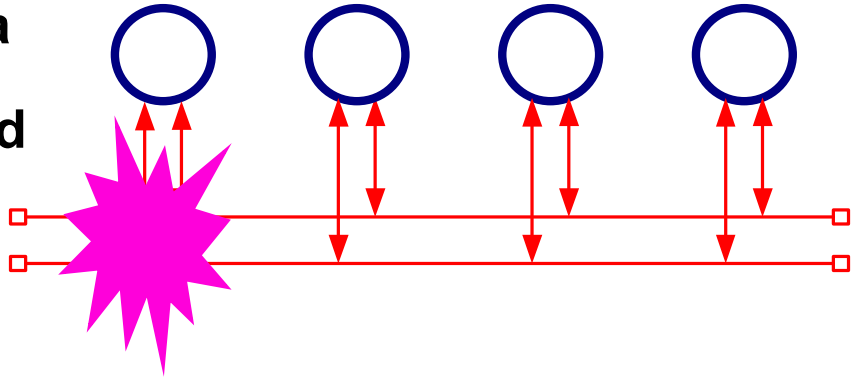

- Redundant Star
  - Independent guardians (G)
  - Reshaping in Guardians performs Byzantine filtering
  - Dual architecture does not allow arbitrarily faulty components
  - What dependability level can be reached?
    - Weakest link principle says 10^{-6}
  - 10^{-6} is not good enough
    - Try argue bizarre fault mode has lower probability
    - Or use triplex (not low cost)

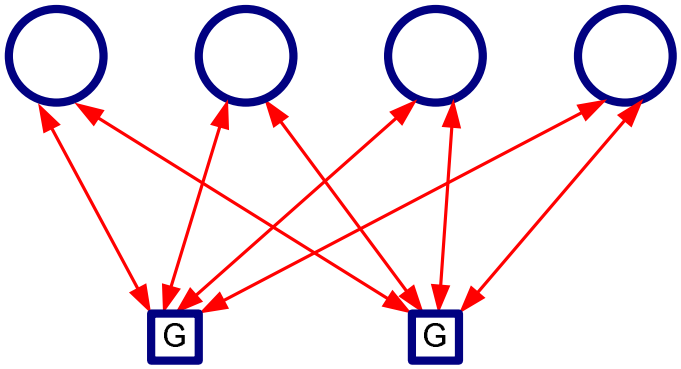

The Path of Low-Cost Dependable Systems

- Redundant Bus
  - Not sufficient due to spatial proximity faults (unavoidable via routing)
  - Serious issues with babbling and masquerade failures
    - Local Guardians not truly independent

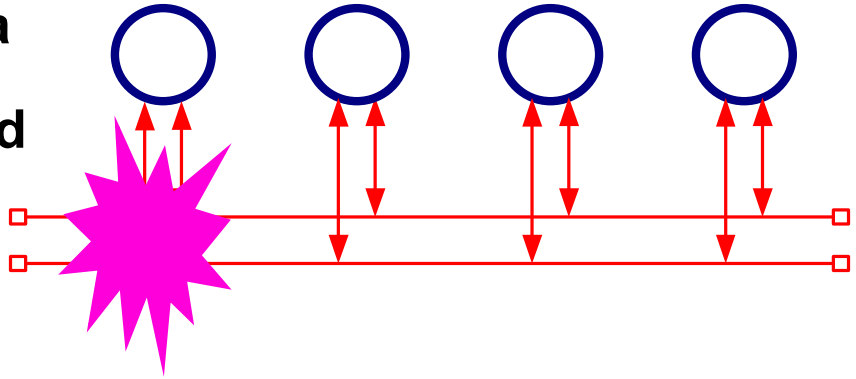

- Redundant Star
  - Independent guardians (G)
  - Reshaping in Guardians performs Byzantine filtering
  - Dual architecture does not allow arbitrarily faulty components
  - What dependability level can be reached?
    - Weakest link principle says 10^{-6}
  - 10^{-6} is not good enough
    - Try argue bizarre fault mode has lower probability
    - Or use triplex (not low cost)

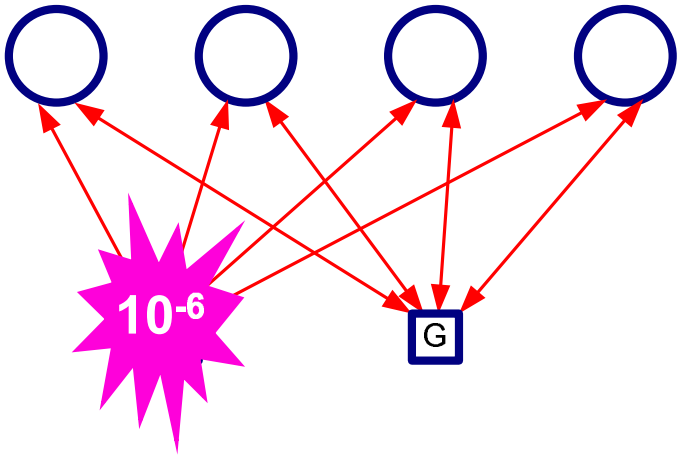

Braided Ring

- Begins with traditional braided ring
  - Each node has links to two nearest neighbors and links to two next nearest neighbors
  - Each link is 2x unidirectional
  - Four link paths from each source to each destination (used for availability only)
- Adds these new ideas
  - Eliminate half of transmitters and ~1/4 of wire length
  - Uses Byzantine filtering during bit regeneration in each node
  - Does a bit-for-bit compare of each node's output vs input
    - Miscompares set a failed flag in the tail of bad messages
    - Nearly 100% coverage of regeneration errors
  - …

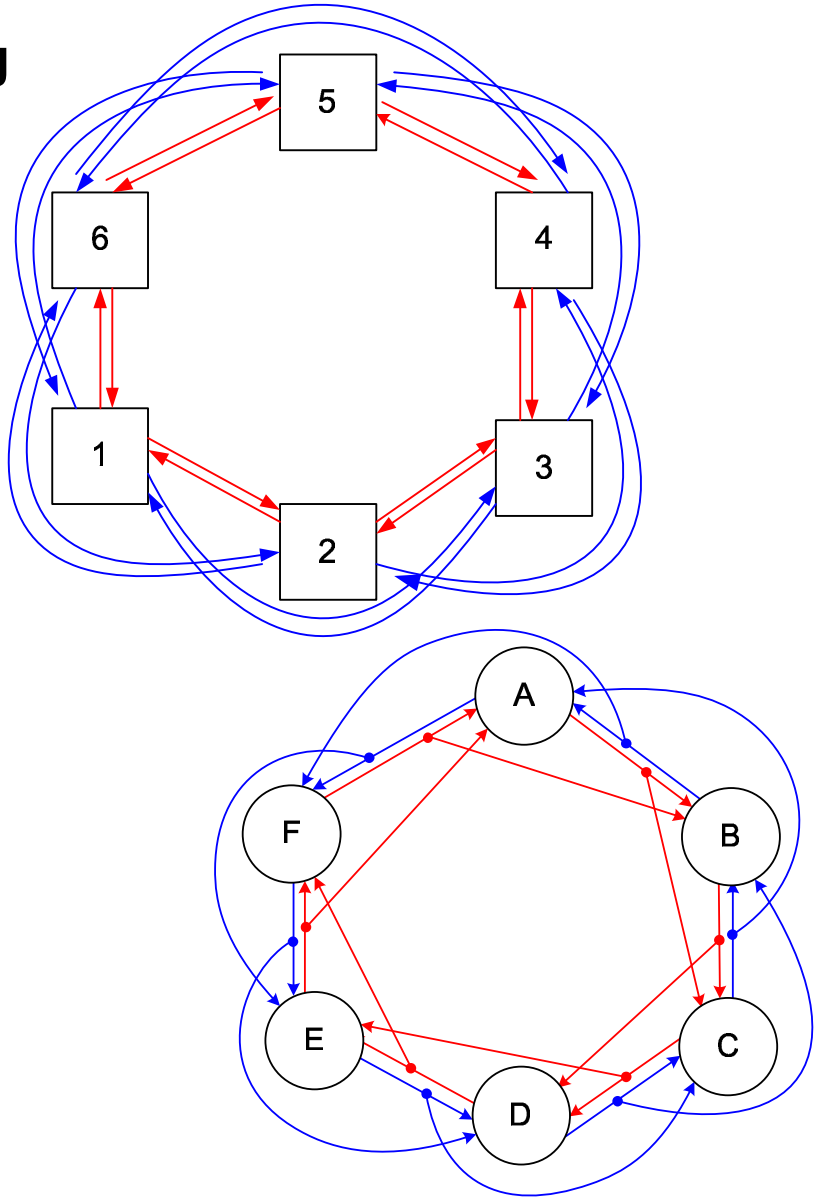

Availability versus Integrity

- **The Availability "OR"**
  - **If miscompare, arbitrarily select one**
  - **Goal is readiness for correct service**

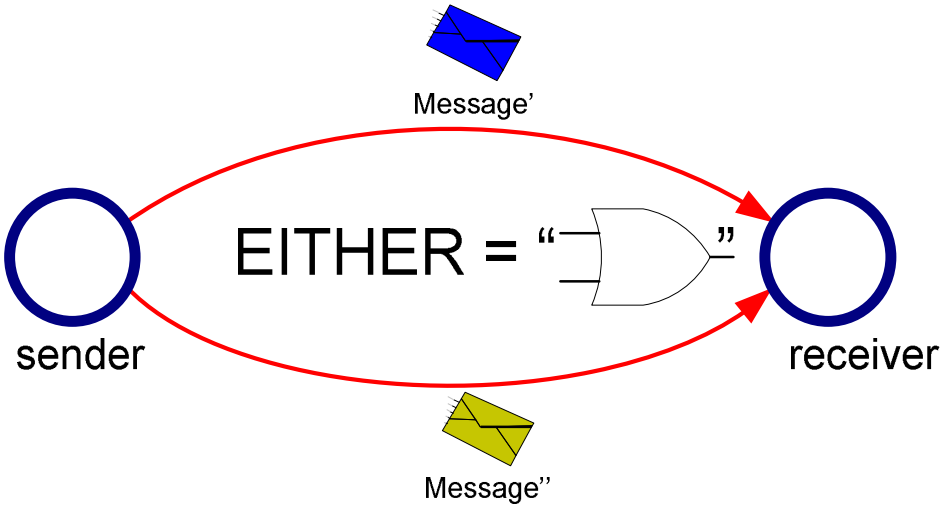

- **The Integrity "AND"**
  - **If miscompare, reject both**
  - **Goal is absence of improper alterations**

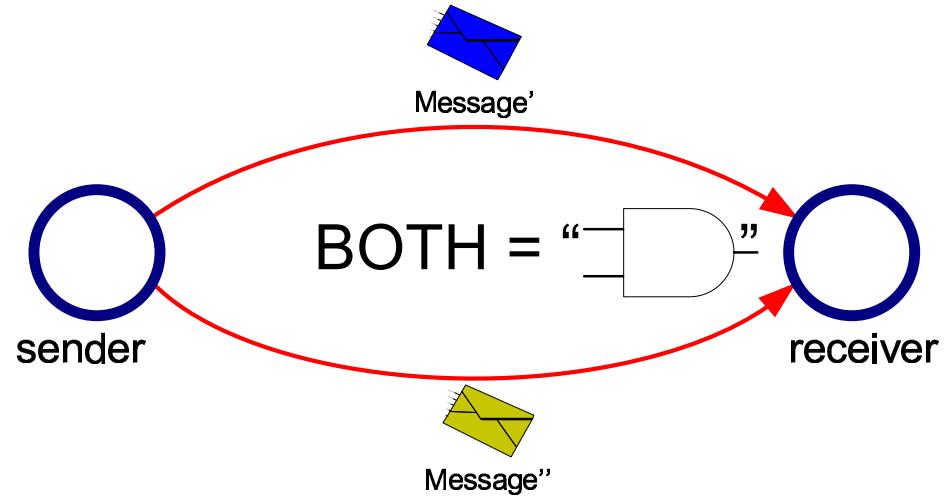

Simple 2x replication gives you availability or integrity but not both!

Checking for Node Failures

- **Errors of components are detected by doing a bit-for-bit compare of node's input versus output**

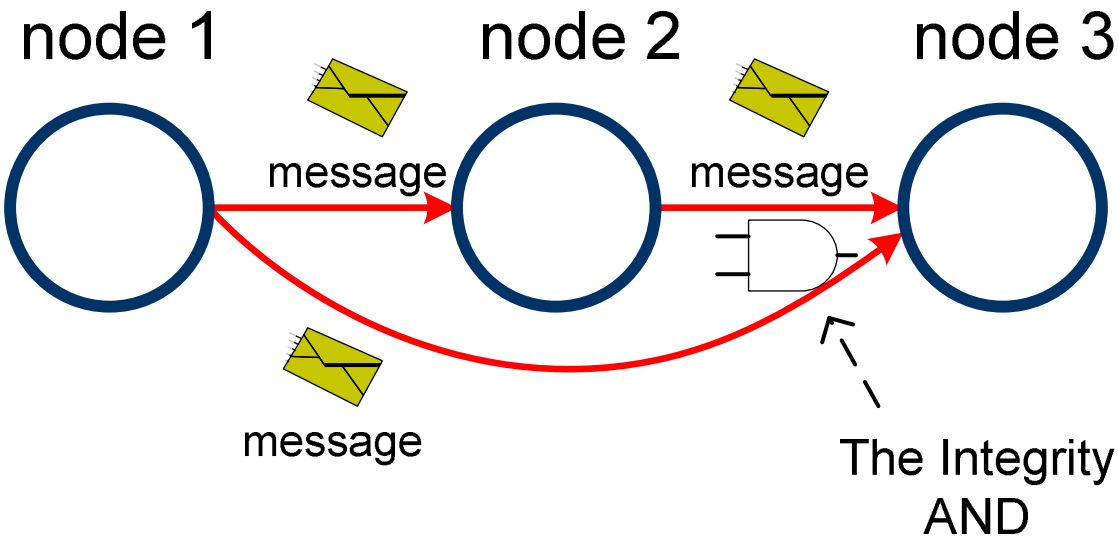

Checking for Node Failures

- **Errors of components are detected by doing a bit-for-bit compare of node's input versus output**
- **Comparison of protocol behavior (timing)**

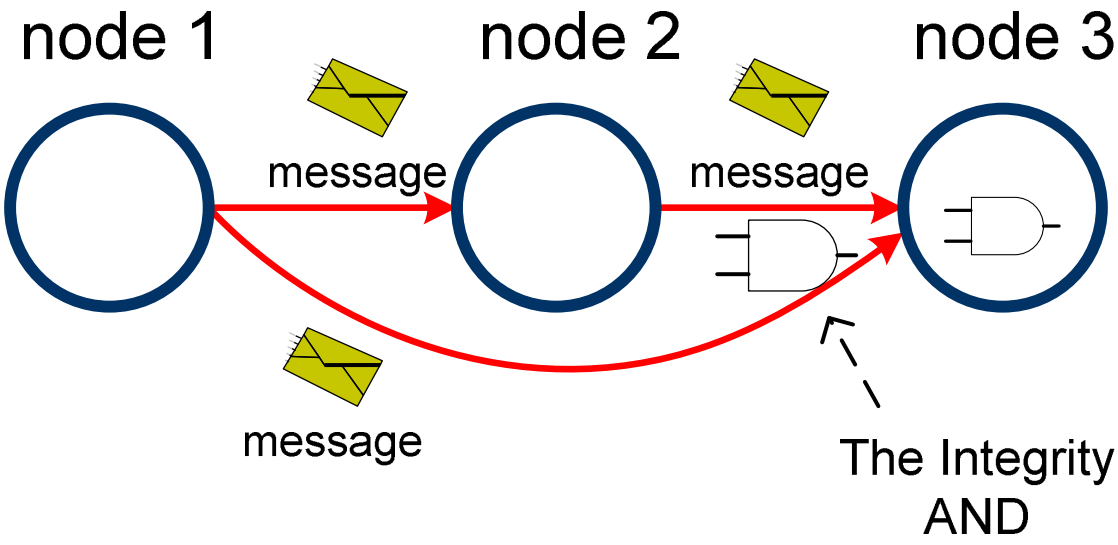

Braided Ring is a Full Coverage Architecture

- **Full coverage data propagation**
  - **"true" 10^{-9}**
- **Neighboring nodes perform guardian function**
- **No need for separate silicon for guardian**
  - **Saves silicon and thus cost**

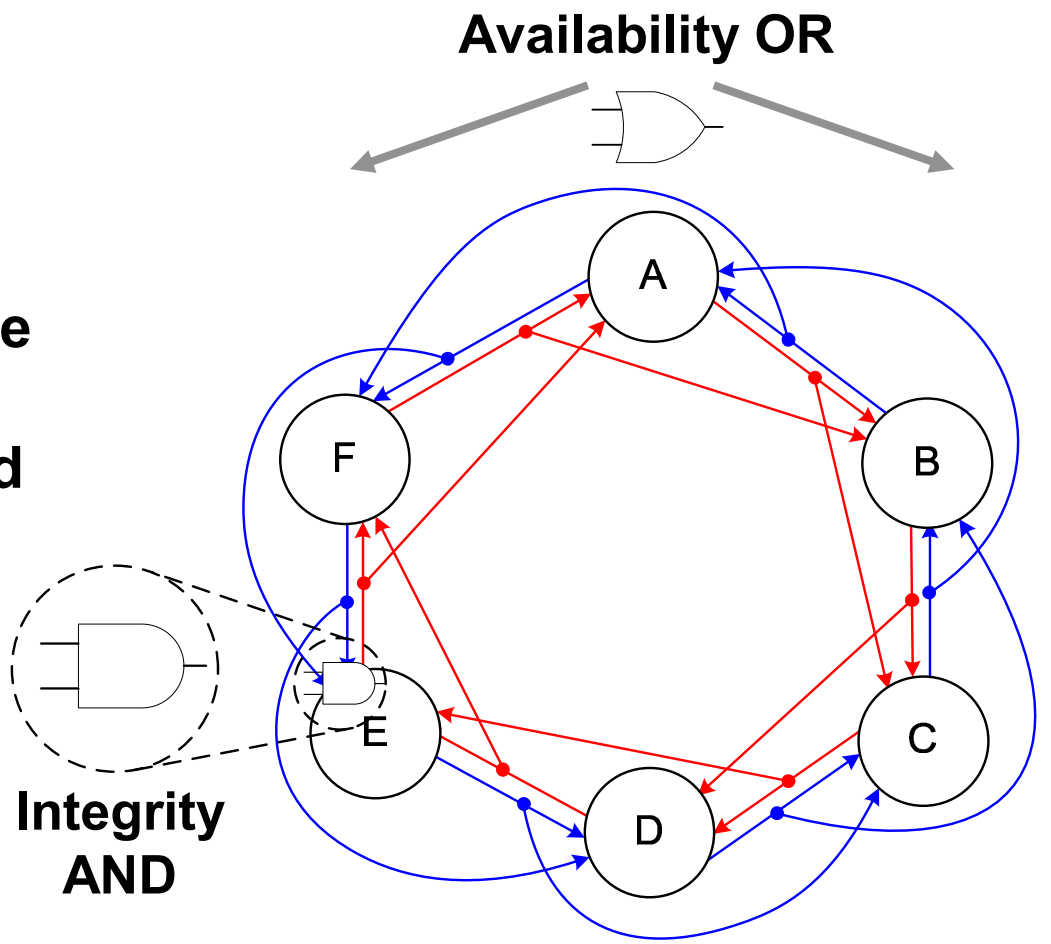

One of Many Ways to Cable the Braided Ring

- "Skip" links to next nearest neighbors can be routed via nearest neighboring nodes
- Useful when bundling cable costs are a significant part of wiring costs

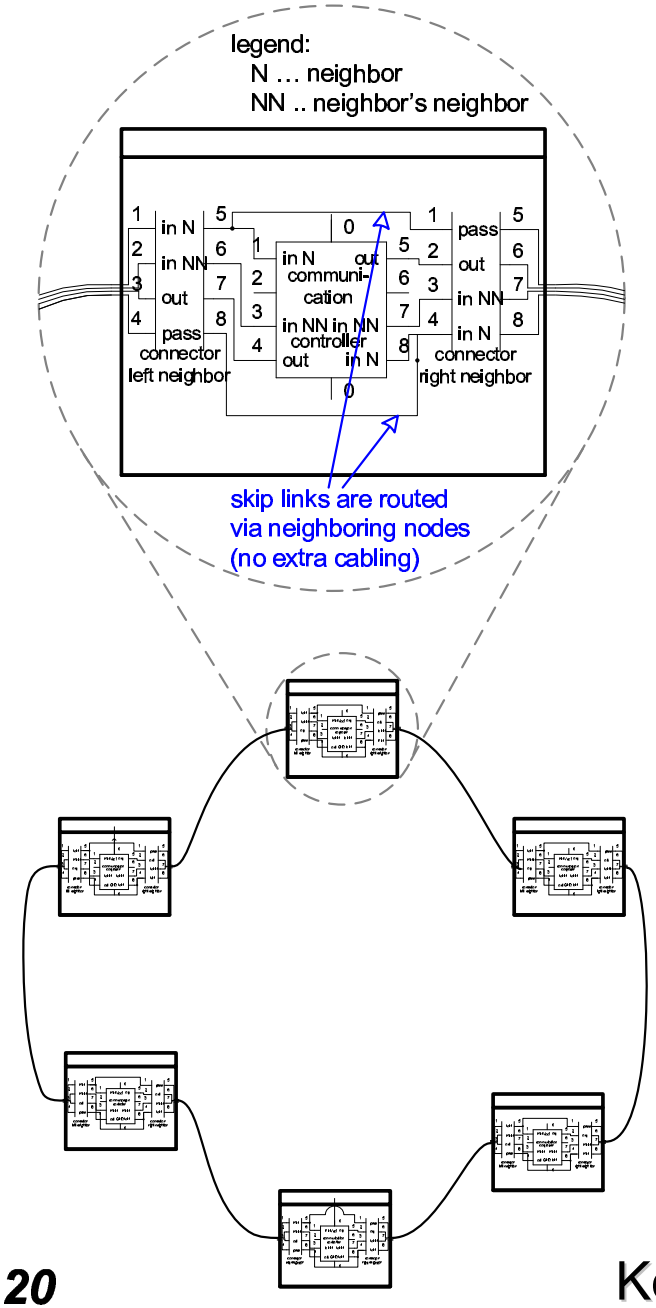

Detection Containment Bypasses

- Need to detect internal and external bypasses of devices entrusted to do fault (error) containment.

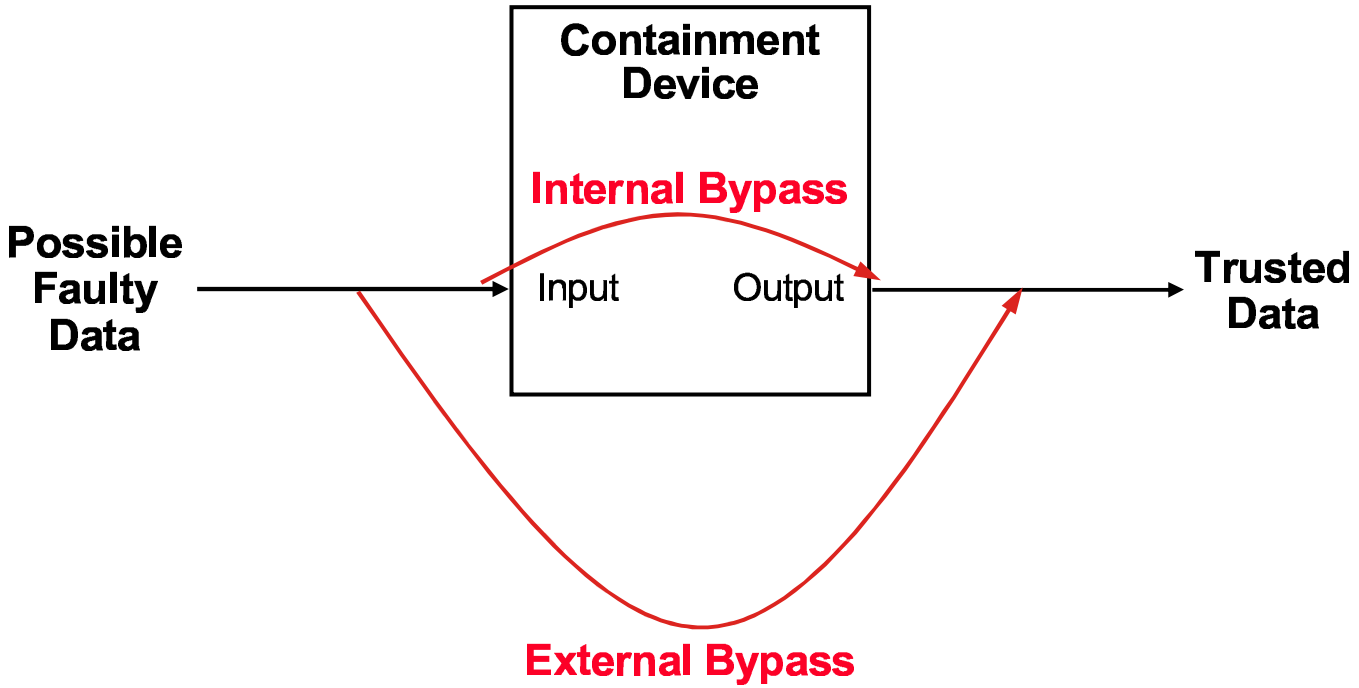

Detection Containment Bypasses

- Need to detect internal and external bypasses of devices entrusted to do fault (error) containment.

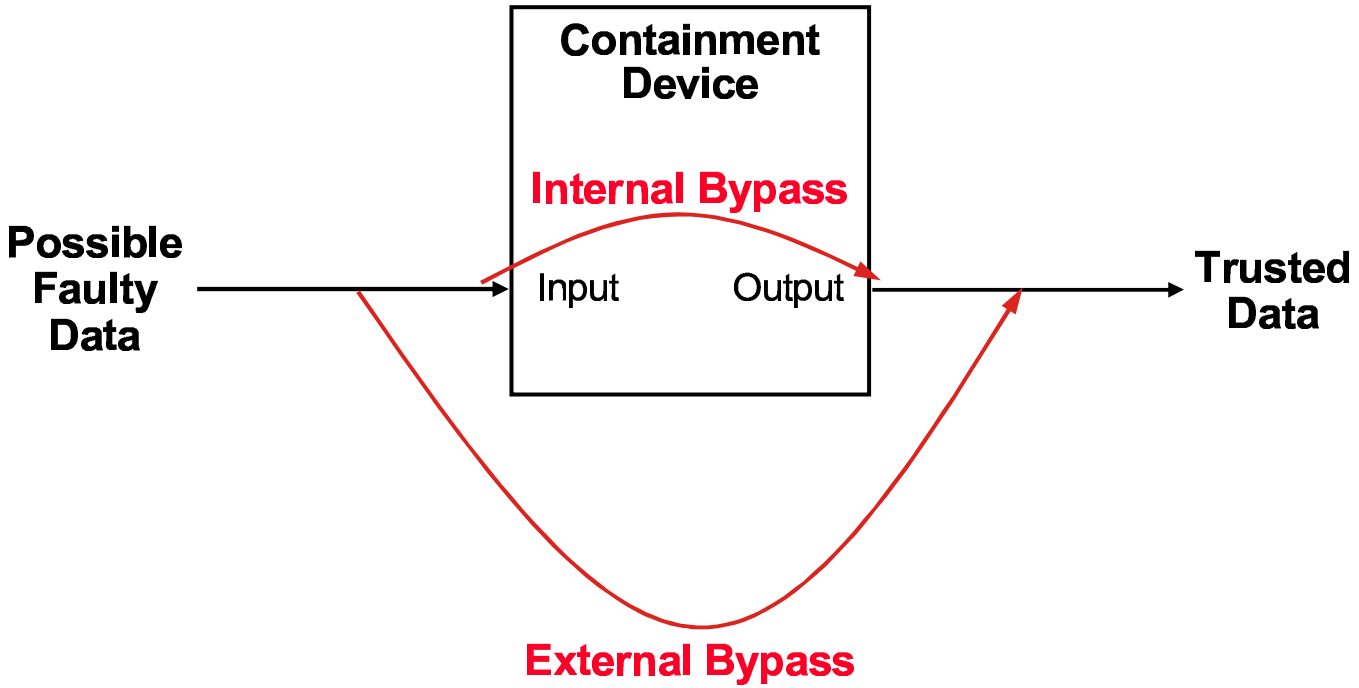

- Solution: “Encrypt” each link differently

Ring Example Using the Minimum 6 Keys

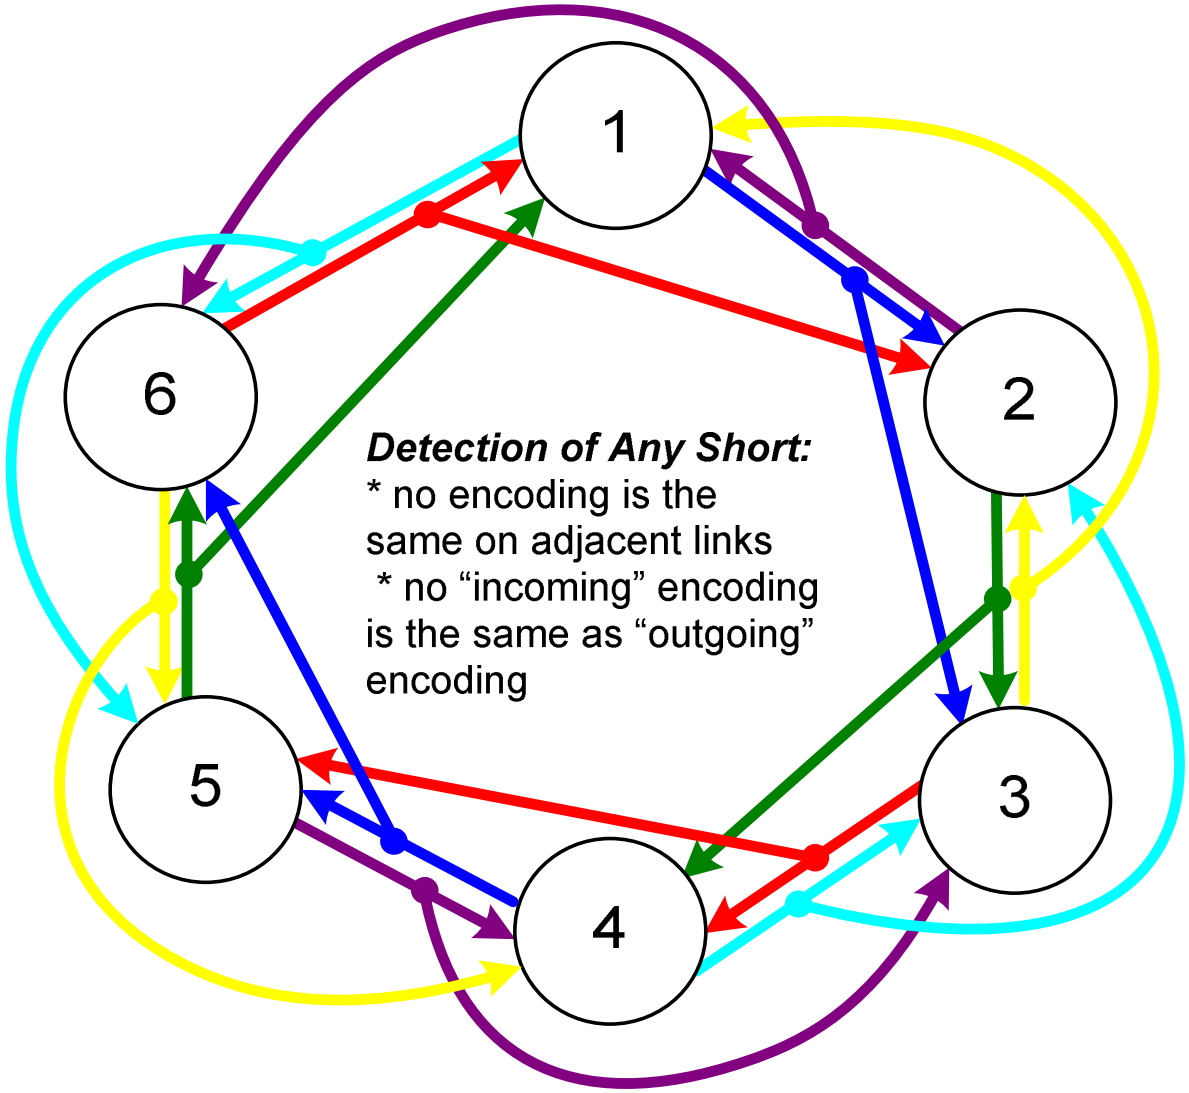

legend: different color represents different encoding

Benefits of a Braided Ring

- Compared to a bus topology system
  - Survives a proximately fault
  - Babble and masquerade faults stopped by neighbors
    - No problem with untrustable local guardians
    - No need to add another integrated circuit for guardianship
    - No electrical fault isolation needed for network interface
- Compared to a star topology system
  - No need for additional (triplex) central components
    - Less cost
    - Less unreliability
  - Less costly wiring
    - Cable has to go only to nearest neighbor, not all the way to a central star
- Optimally cheap Byzantine solution?