Independent events

- Definition: Events A and B defined on an experiment are said to be (stochastically) mutually independent if

$$P(A \cap B) = P(A)P(B)$$

- Sometimes people say “A is independent of B” instead, but independence is mutual: A is independent of B if and only if B is independent of A

What’s with the “stochastically”?

- $P(A \cap B) = P(A)P(B)$ for stochastically independent events
- If we believe that events A and B are physically independent, then we insist that the probability measure must assign probabilities to events in such a way that this equality holds
- But, equality can hold even for events that are provably physically dependent… such events are stochastically independent

Physical stochastic

- Physical independence is, in essence, a property of the events themselves
- We believe that events A and B are physically independent and express this independence via $P(AB) = P(A)P(B)$
- Stochastic independence is a property of the probability measure
- Stochastic independence does not necessarily mean that the events are physically independent

Consequences of independence

- If A and B are mutually independent events, then $P(AB) = P(A)P(B)$
- This is equivalent to each of the following
  - $P(AB^c) = P(A)P(B^c)$
  - $P(A^cB) = P(A^c)P(B)$
  - $P(A^cB^c) = P(A^c)P(B^c)$
- In other words, A and B^c are mutually independent, as are A^c and B, and as for A^c and B^c, why, they are independent too!

Mutually is as mutually does…

- If A and B are mutually independent events, then $P(AB) = P(A)P(B)$
- If A and B are mutually exclusive events, then $P(AB) = 0$
- Mutually exclusive events cannot be mutually independent
- Mutually independent events cannot be mutually exclusive
- Except in trivial cases: $P(A)$ or $P(B)$ is 0

Conditional = unconditional

- If A and B are mutually independent events, then $P(AB) = P(A)P(B)$
- If $P(A) > 0$, we get that

$$P(B|A) = P(AB)/P(A) = P(B)$$

- The conditional probability of B given A is the same as the unconditional probability!
- Knowing that A occurred does not cause any “updating” of the chances of B
- Similarly, $P(A|B) = P(AB)/P(B) = P(A)$

Independence of three events

- Events A, B, and C are said to be mutually independent if all four of the following conditions hold:
  - $P(AB) = P(A)P(B)$
  - $P(AC) = P(A)P(C)$
  - $P(BC) = P(B)P(C)$
  - $P(ABC) = P(A)P(B)P(C)$
- The first three conditions do not imply the fourth, nor does the fourth imply the first three

Independence of multiple events I

- $\{A_1, A_2, \ldots, A_n\}$ is said to be a collection of mutually independent events if
  - $P(A_1A_2\ldots A_n) = P(A_1)P(A_2)\ldots P(A_n)$
  - every subcollection containing two or more of the A_i's is also a collection of independent events
- This is a recursive definition: the product rule applies to the "big" intersection and also to all smaller ones as well

$$P(A_iA_j\ldots A_m) = P(A_i)P(A_j)\ldots P(A_m)$$

Independence of multiple events II

- $\{A_1, A_2, \ldots, A_n\}$ is said to be a collection of mutually independent events if all 2^n of the following equations hold:

  $P\{(A_1)^*(A_2)^* \ldots (A_n)^*\}$

  $= P\{(A_1)^*\}P\{(A_2)^*\} \ldots P\{(A_n)^*\}$
- Here, $(A_i)^*$ represents either A_i or $(A_i)^c$ (the same on both sides of each equation)
- Two choice for each $(A_i)^*$ 2^n equations

Independence of multiple events III

- $P\{(A_1)^*(A_2)^* \ldots (A_n)^*\}$

  $= P\{(A_1)^*\}P\{(A_2)^*\} \ldots P\{(A_n)^*\}$

  implies that the product rule also applies to every subcollection containing two or more of the $(A_i)^*$s
- $P\{(A_i)^*(A_j)^* \ldots (A_m)^*\} =$

  $P\{(A_i)^*\}P\{(A_j)^*\} \ldots P\{(A_m)^*\}$

Boolean functions

- Any Boolean function of independent events $\{A_1, A_2, \ldots, A_n\}$ is independent of any other Boolean function as long as they do not include any events in common
- Example: If A, B, C are independent events, then A B is independent of C
- Example: If {A,B,C,D,E,F,G,H} are independent events, then A C, B H, D, and E FG are independent events

Modeling in reliability of systems

- Components of engineering systems are prone to failure in various ways
- Careful modeling of failure modes is extremely important
- Careful modeling of how a component failure affects system operation is also extremely important
- Most "errors" in reliability analyses arise from incorrect modeling of failures

Analysis and design of reliable systems

- How reliable is a system that has been constructed from unreliable components?
- How can we design reliable systems from unreliable components?
- How can we design systems that continue to operate even when some of their components have failed?
- How to mask component failures?

Reliability of digital systems I

- Logic circuits can fail in various ways
- Failures are called faults
- Stuck-at-zero fault and stuck-at-one fault
- A stuck-at-one gate produces the correct output for some values of the inputs!
- If the gate that is stuck is internal to the circuit, then at times the gate failure can be masked if the same cells in Karnaugh map are covered by other implicants

Reliability of digital systems II

- Consider an AND-OR circuit with inputs x, y, and z, and output xy yz xz
- If the output OR gate is stuck-at-one, the output is incorrect only when two or more of the inputs is zero
- If the AND gate computing xy is stuck-at-0, the output is incorrect only when x = y = 1 and z = 0

Reliability of digital systems III

- Logic circuit design to mask possible internal faults is an important research area
- Fault-tolerant circuit design
- Fault-tolerant computing
- Fault recovery

Triple Modular Redundancy (TMR)

- Triple modular redundancy (TMR) systems are often used when very high reliability is desired (e.g. military systems, financial systems)
- Three identical circuits are used to compute the specified Boolean function
- All three circuits have identical inputs
- Because of gate failures, the outputs may not be identical

The majority gate

- Three identical circuits are used to compute the specified Boolean function
- Due to gate failures, the three outputs may not be identical
- A majority gate is used to decide which is the correct output
- Majority gate output is 1 if two or more of its inputs are 1; output is 0 if two or more of the inputs are 0

More about the majority gate

- Majority gate output is 1 if two or more of its inputs are 1; output is 0 if two or more of the inputs are 0
- The majority gate is a simple AND-OR circuit — if the inputs are denoted by x, y, and z, then the output is xy yz xz
- A majority gate is just the carry output of a full adder
- Majority gate is a voting machine!

How TMR works

- If the Boolean function has value 1, then all three identical circuits produce output 1 assuming they have not failed
- If one circuit fails and produces an output of 0, while the other two are working correctly and produce an output of 1, the majority gate output is a 1
- Similarly for the case when the Boolean function has value 0

Summary of TMR system

- The majority gate output is guaranteed to be correct as long as no more than one of the three identical circuits has failed
- Could the majority gate itself fail?
- Yes, the majority gate can also fail
- The Boolean functions computed in triplicate by TMR systems are chosen to be fairly complicated ones
- Majority gate is a simple circuit

Analysis of TMR system I

- Assume that the majority gate cannot fail
- Assume that each of the three circuits fails with probability p and that the failures are independent
- The independence of failures is a critical assumption
- If the circuits share a common power supply or a common circuit board, failures may not be independent

Analysis of TMR system II

- The majority gate output is incorrect when two or more of the circuits has failed
- The probability that two or more of the circuits fails is $3p^2(1–p) + p^3 = 3p^2 – 2p^3$
- P(majority gate output is incorrect) $= 3p^2 – 2p^3$
- For small values of p, $3p^2 – 2p^3 << p$
- The majority gate output is much more likely to be correct than any circuit output

Analysis of TMR system III

- P(majority gate output is incorrect) $= 3p^2 – 2p^3 << p$
- Example: If $p = 10^{–4}$, then the majority gate output is incorrect with probability $≈ 3×10^{–8}$
- The majority gate output is much more likely to be correct than any circuit output
- TANSTAAFL! The costs of a TMR system are at least three times as much

What about majority gate failure?

- Assume that the majority gate can fail with probability $q << p$
- $q << p$ because the majority gate is a simple circuit as compared to the others
- Majority gate failure means the output is not the majority value of its inputs
- If majority gate fails, and two or more of the inputs are incorrect because of circuit failures, the output is correct!

Use theorem of total probability

- P{majority gate output is incorrect}
  = P{incorrect | majority gate OK}P{OK}
  + P{correct | majority gate NOK}P{NOK}
  $= (3p^2 - 2p^3)(1-q) + (1 - (3p^2 - 2p^3))q$
  $\approx q$ in most cases of interest
- Example: If $p = 10^{-4}$ and $q = 10^{-6}$, then $(3p^2 - 2p^3) \approx 3\times10^{-8}$ and the probability of error is $\approx q$

Moral

- Failures in the majority gate dominate the error probability
- Possible solutions: make the majority gate as reliable as possible
- Law of diminishing returns
- The solution actually used: triplicate the majority gate as well!
- 2 or 3 outputs are wrong only when at least 2 circuits (or majority gates) fail

Block diagram of a TMR system

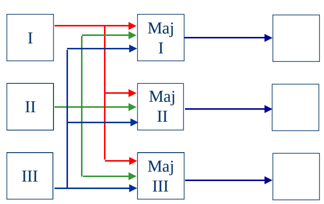

Communication Networks

- Communication networks consist of links over which messages can pass between nodes (or terminals or hosts)
- Messages to distant nodes have to pass over multiple links and multiple nodes
- The intermediate links can fail
- The intermediate nodes can fail
- What is the probability that two nodes can communicate?

Links only; we don't do nodes

- A node that fails makes all the links connected to it inoperable
- All paths between other nodes that pass through the failed node are not available
- In the simple analyses presented in this course, we assume that nodes do not fail
- The only failures that we are concerned with are the link failures

Communication over parallel links

- The diagram illustrates two parallel links connecting a transmitter and a receiver
- The transmitter can send messages as long as at least one link is viable
- Both links must fail in order for communication to be impossible

Some notation

- Let V_i denote the event that link #i is viable that is, link #i has not failed
- The complementary event is F_i
- V ($F = V^c$) denotes the event that there is (is not) a communication path from T to R

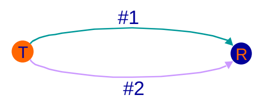

As easy as rolling off a log…

- $V = V_1 \quad V_2$
- $F = F_1F_2$
- $P(F) = P(F_1F_2) = P(F_1)P(F_2)$

  if the events F_1 and F_2 can be assumed to be independent
- Similarly, $P(V) = P(V_1) + P(V_2) - P(V_1V_2)$

  $= P(V_1)+P(V_2) - P(V_1)P(V_2)$

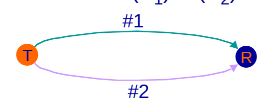

Did matters improve?

- $P(F) = P(F_1F_2) = P(F_1)P(F_2)$ for independent failures
- $P(F) << \min\{P(F_1), P(F_2)\}$
- Example: If $P(F_1) = P(F_2) = 10^{-3}$, then $P(F) = 10^{-6}$
- Improvement in P(V) is not as "dramatic"

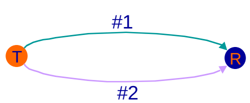

Path over serial links

- The diagram illustrates two serial links connecting a transmitter and a receiver
- The transmitter can send messages as long as both links are viable
- Communication is impossible if either link fails

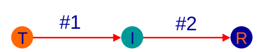

Even easier than rolling off a log…

- $V = V_1V_2$
- $F = F_1 \quad F_2$
- $P(V) = P(V_1V_2) = P(V_1)P(V_2)$

  if the events V_1 and V_2 can be assumed to be independent
- Similarly, $P(F) = P(F_1) + P(F_2) - P(F_1F_2)$

  $= P(F_1)+P(F_2) - P(F_1)P(F_2)$

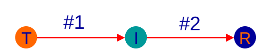

Matters did not improve

- $P(V) = P(V_1V_2) = P(V_1)P(V_2)$ for independent failures
- $P(V) < \min\{P(V_1), P(V_2)\}$
- $P(F) = P(F_1) + P(F_2) - P(F_1)P(F_2)$
- If $P(F_1) = P(F_2) = 10^{-3}$, then $P(F) = 10^{-3} + 10^{-3} - 10^{-6} \approx 2\times10^{-3}$

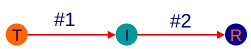

Summary

- We discussed the reliability of systems
- We discussed triple modular redundancy (TMR) systems at some length
- We discussed communication networks and how the link failure probabilities can be used to compute the system failure probabilities

