Redundancy and Dependability Evaluation

EECE 513: Design of Fault-tolerant Systems

Learning Objectives

- At the end of this lecture, you will be able to
  - Define combinatorial models of reliability
  - Evaluate the reliability of series, parallel systems
  - Evaluate reliability of non-series, parallel systems
  - Evaluate standby redundancy schemes
  - Model failures using the exponential distribution
  - Evaluate the reliability of TMR and TMR Simplex
  - Understand the pitfalls of single voter in TMR

Why learn this stuff (again) ?

- Probability & statistics is at the heart of what we do in reliability
  - No such thing as a perfectly reliable system
  - Engineering design requires quantification of trade-offs and design choices
  - Failures are rare events (hopefully), so you need the tools of probability to reason about them
  - Modeling errors using statistical distributions is the de-jure model for evaluating systems

Combinatorial Modeling

- System is divided into non-overlapping modules
- Each module is assigned either a probability of working, P_i, or a probability as function of time, $R_i(t)$
- The goal is to derive the probability, P_{sys}, or function $R_{sys}(t)$ of correct system operation
- Assumptions:
  - module failures are independent
  - once a module has failed, it is always assumed to yield incorrect results
  - system is considered failed if it does not satisfy minimal set of functioning modules

Learning Objectives

- At the end of this lecture, you will be able to
  - Define combinatorial models of reliability
  - Evaluate the reliability of series, parallel systems
  - Evaluate reliability of non-series, parallel systems
  - Evaluate standby redundancy schemes
  - Model failures using the exponential distribution
  - Evaluate the reliability of TMR and TMR Simplex
  - Understand the pitfalls of single voter in TMR

Reliability of series systems

- Reliability R_I = Prob that component i works
- Reliability R_s = Prob that system works

Assume that components fail independently.
Component i fails with probability $p_{i.}$

$R_{I\,=}$ Prob that comp. i works = $1 - p_i$

R_s = Prob that system works = $R_1.R_2.R_3....R_n$ = ¶$\mathbf{R_i}$

Reliability of parallel systems

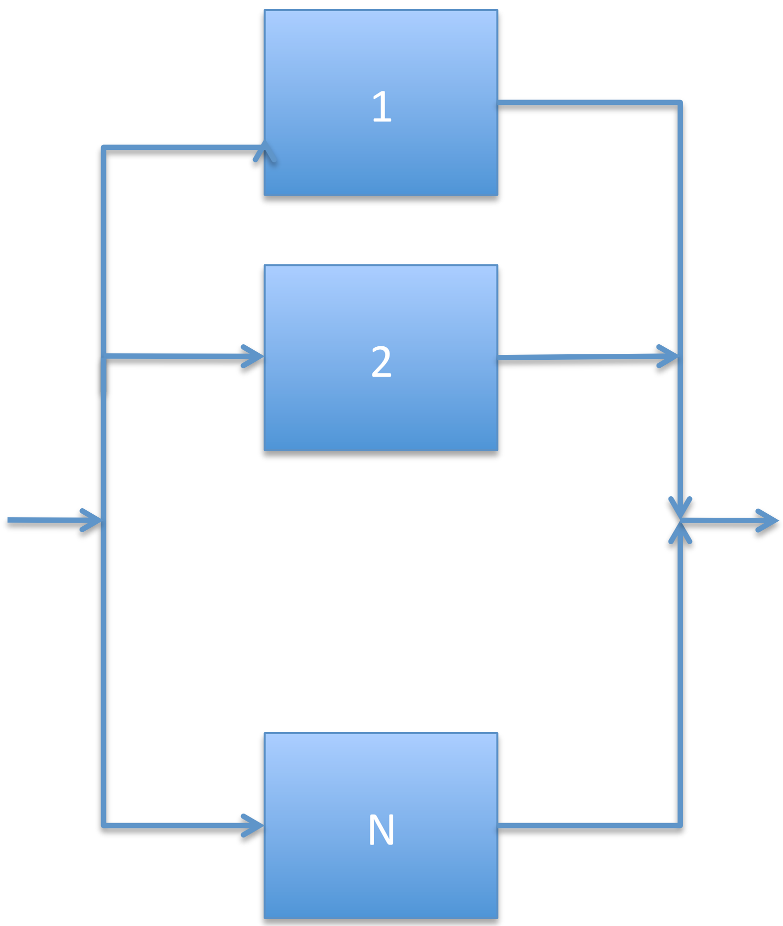

- Assume that components fail independently
- Probability that system works =

R_p = 1 - Probability of all components failing

$= 1 – P_1P_2P_3...P_n$

$= 1 – \Pi_0^n(1 – R_i)$

Series-Parallel Effects

- Series system Reliability
- Parallel system Reliability

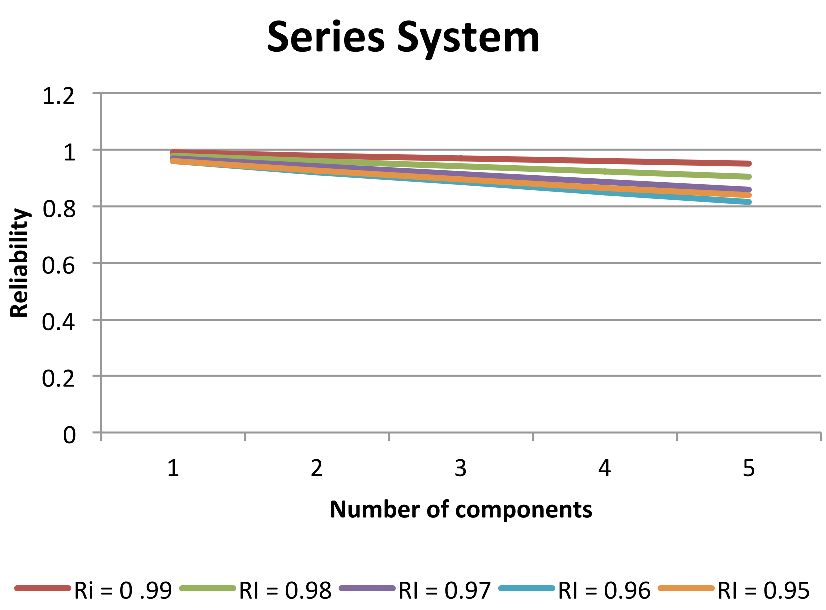

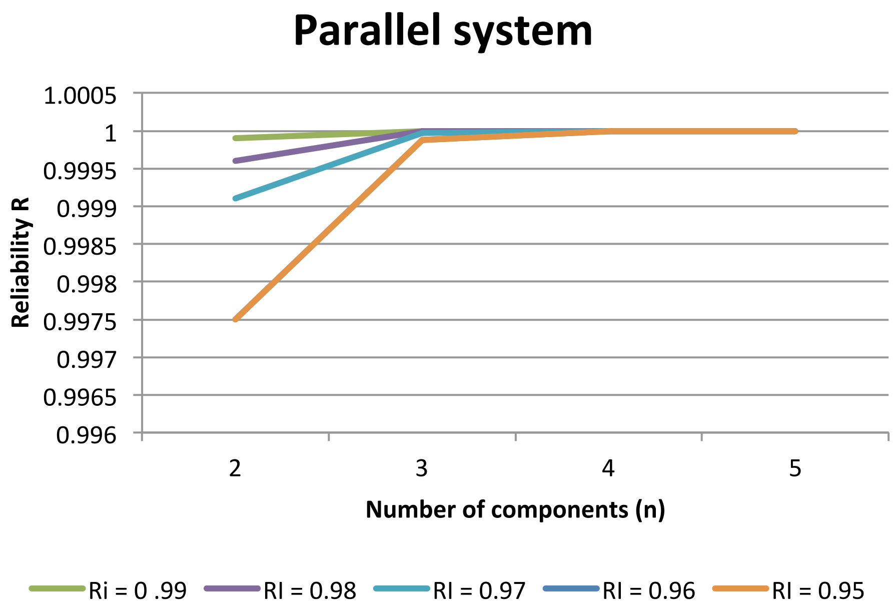

Reliability decreases with N steadily

Reliability increases sharply with N

Exercise

- Calculate the reliability of the following system

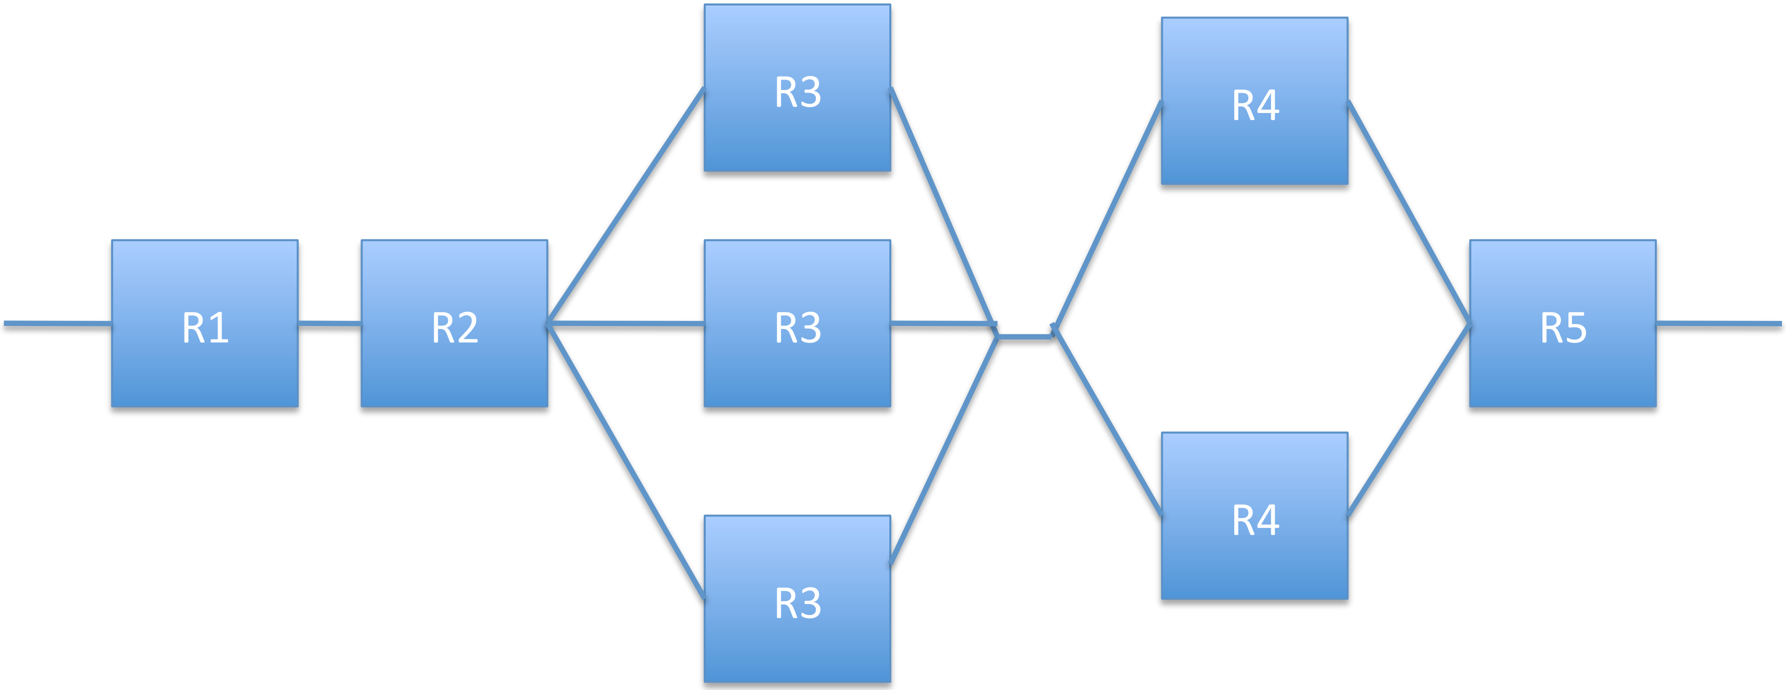

Learning Objectives

- At the end of this lecture, you will be able to
  - Define combinatorial models of reliability
  - Evaluate the reliability of series, parallel systems
  - Evaluate reliability of non-series, parallel systems
  - Evaluate standby redundancy schemes
  - Model failures using the exponential distribution
  - Evaluate the reliability of TMR and TMR Simplex
  - Understand the pitfalls of single voter in TMR

Non-Series Parallel System

- Consider the following system: It is neither a series nor parallel system. So how do we evaluate its reliability ?

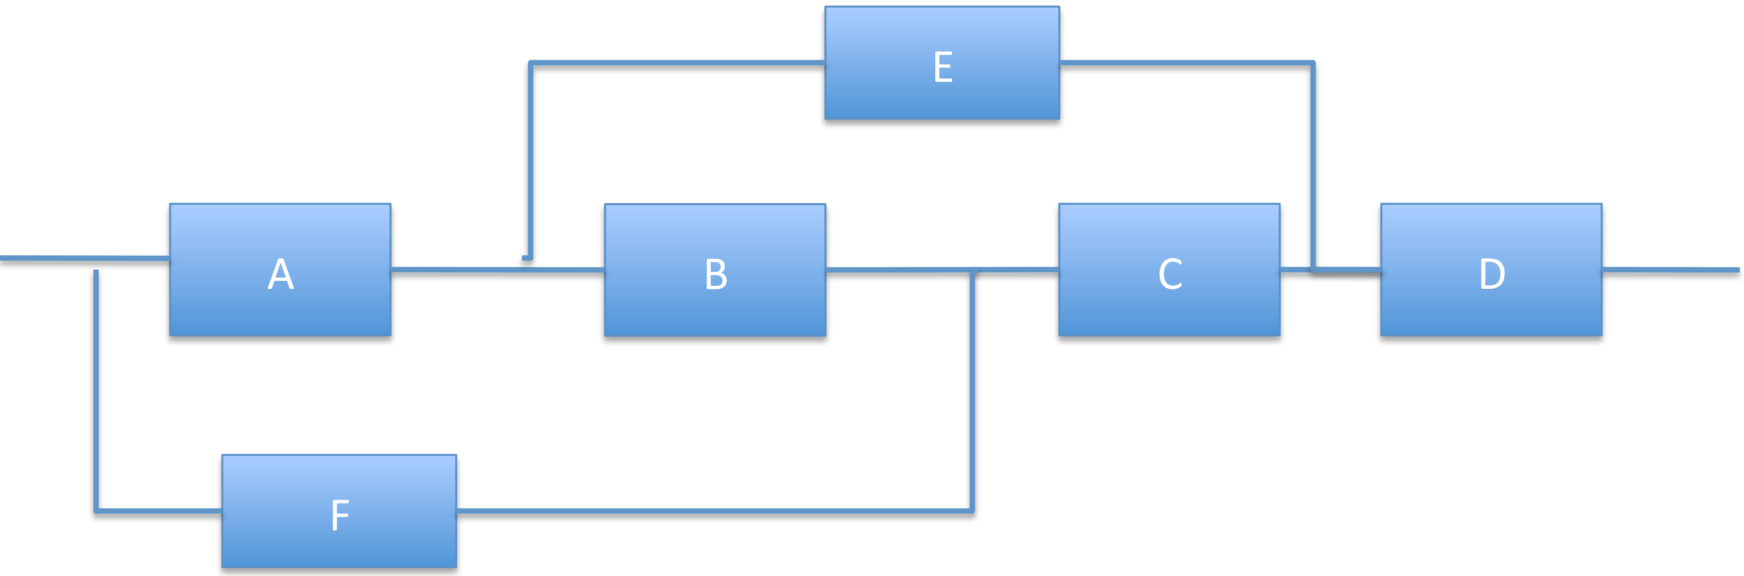

Non-Series Parallel System - 1

Start by picking a module 'm' in the system

P(sys works) = P(sys works | m works) P(m works) + P(sys works | m fails) P (m fails)

R_{sys} = P(sys works | m)R_m + P(sys works | m')(1 – R_m)

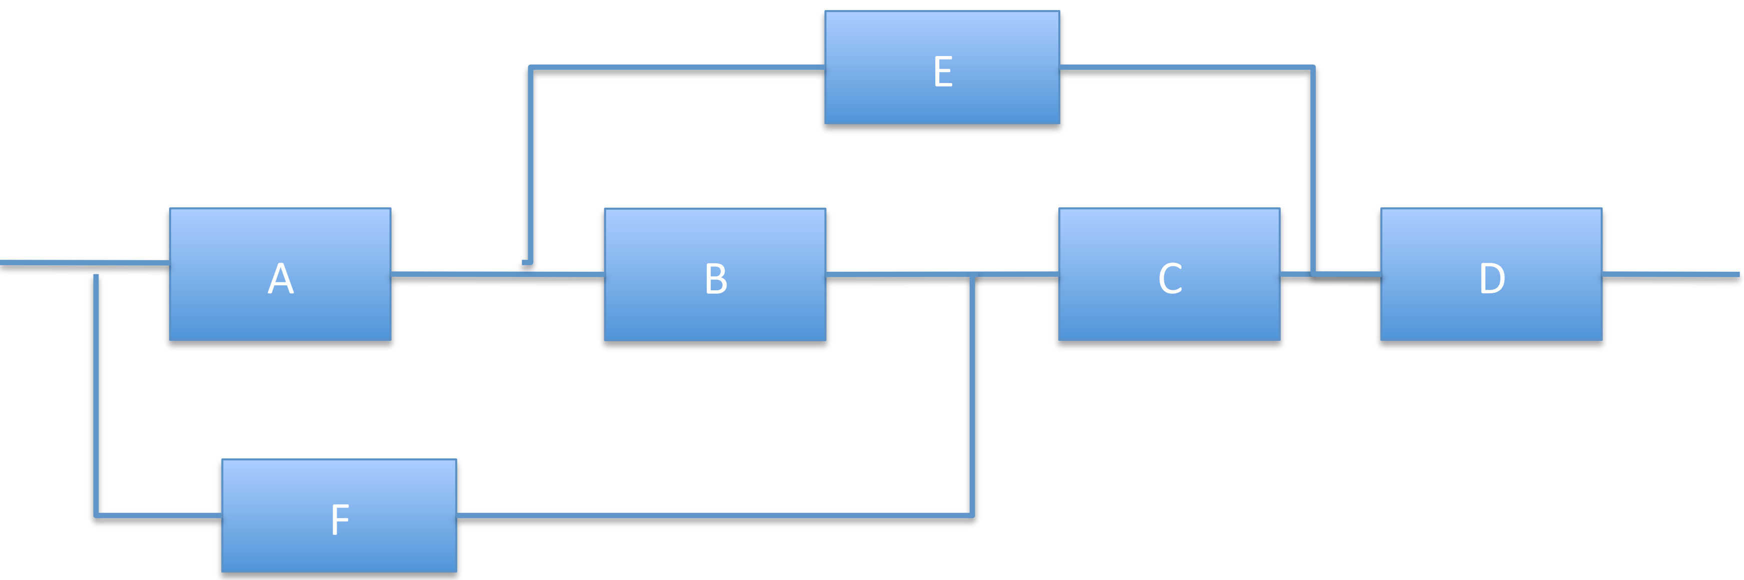

Non-Series Parallel System - 2

Pick 'm' to be module E. Consider the case where E fails. Then,

$$P(\text{sys works}|E') = [\,1 - (1 - R_A R_B)(1 - R_F)]R_C R_D$$

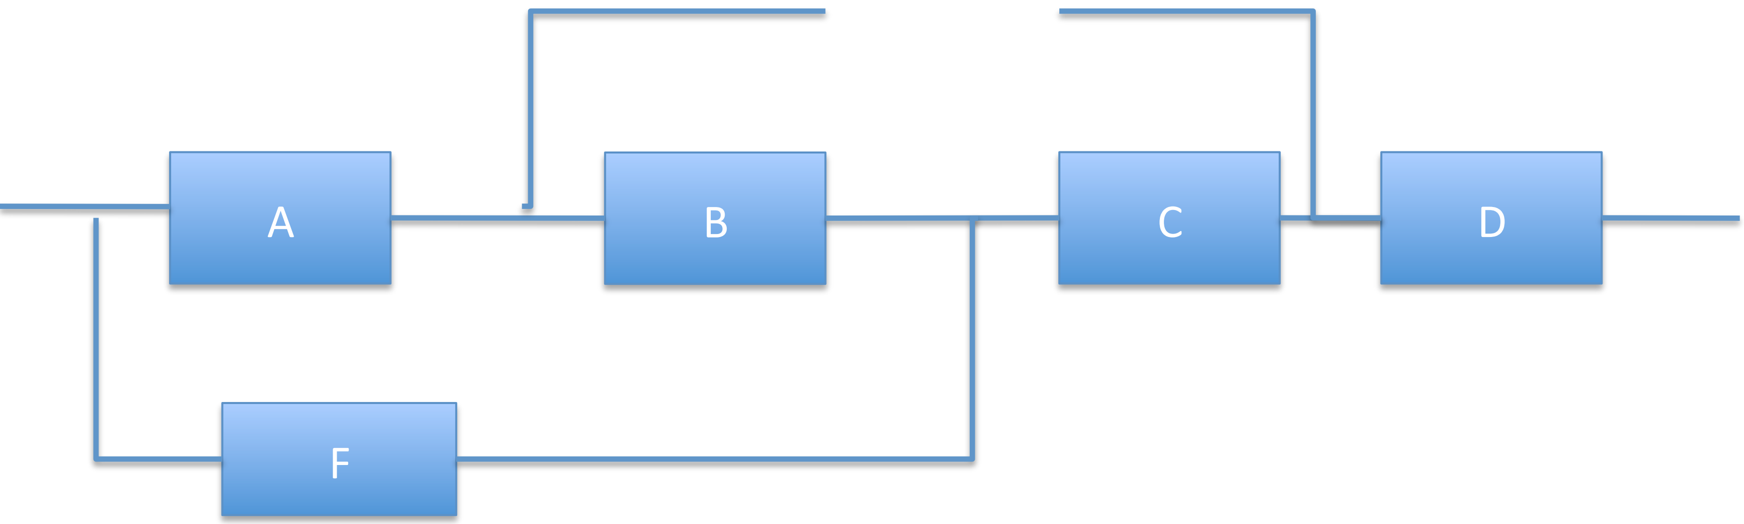

Non-Series Parallel System - 3

Now, consider the case where E works. This is equivalent to shorting the line. However, this does not make the problem any easier. So do the same thing again with a different module.

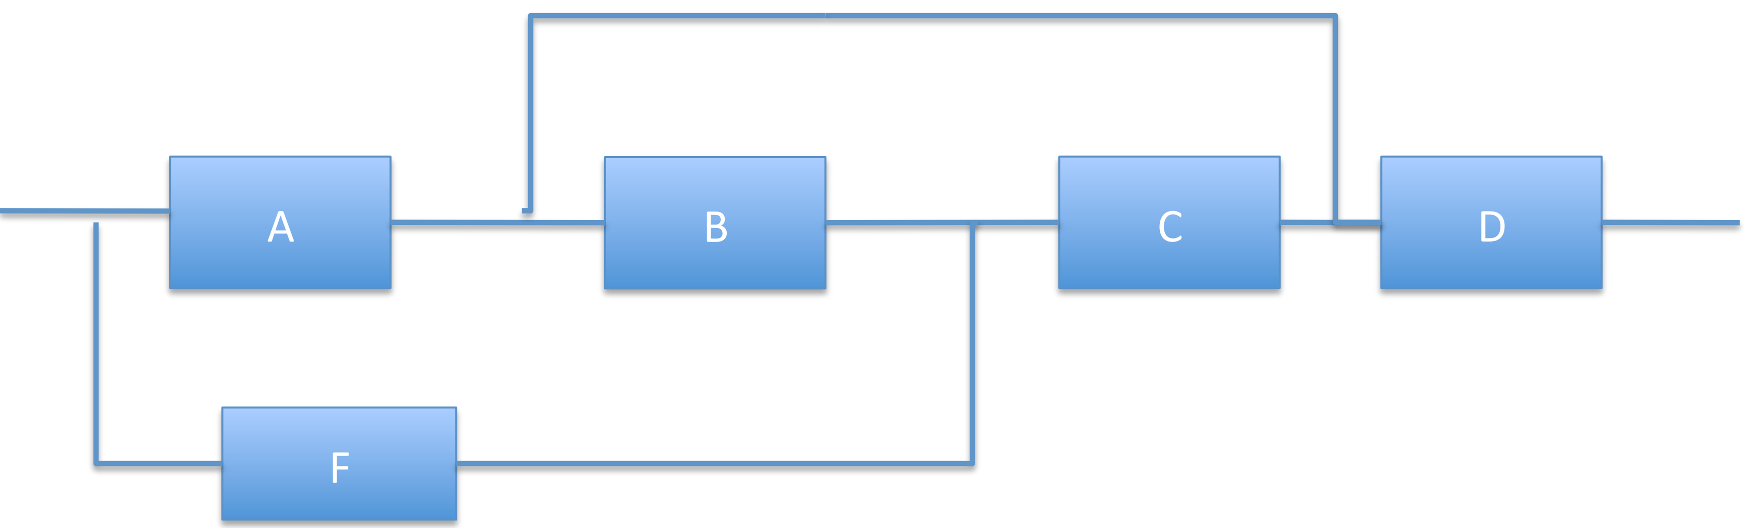

Non-Series Parallel System - 4

Let's pick module B. When B fails,

$$P(\text{sys works} | E \text{ and } B') = [\,1 - (1 - R_A)(1 - R_F R_C)]R_D$$

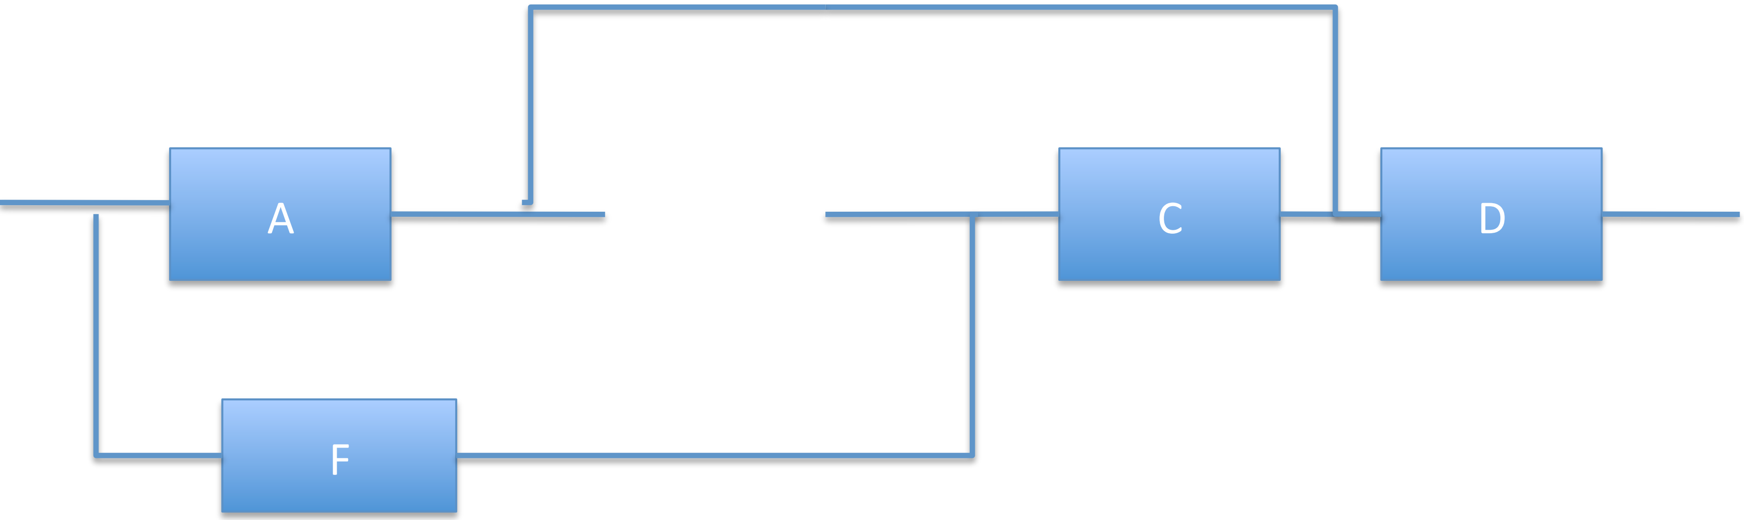

Non-Series Parallel System - 5

When B works, it is still the same as before:

$P(\text{sys works} | \text{E and B}) = [\, 1 - (1 - R_A)(1 - R_F R_C)]R_D$

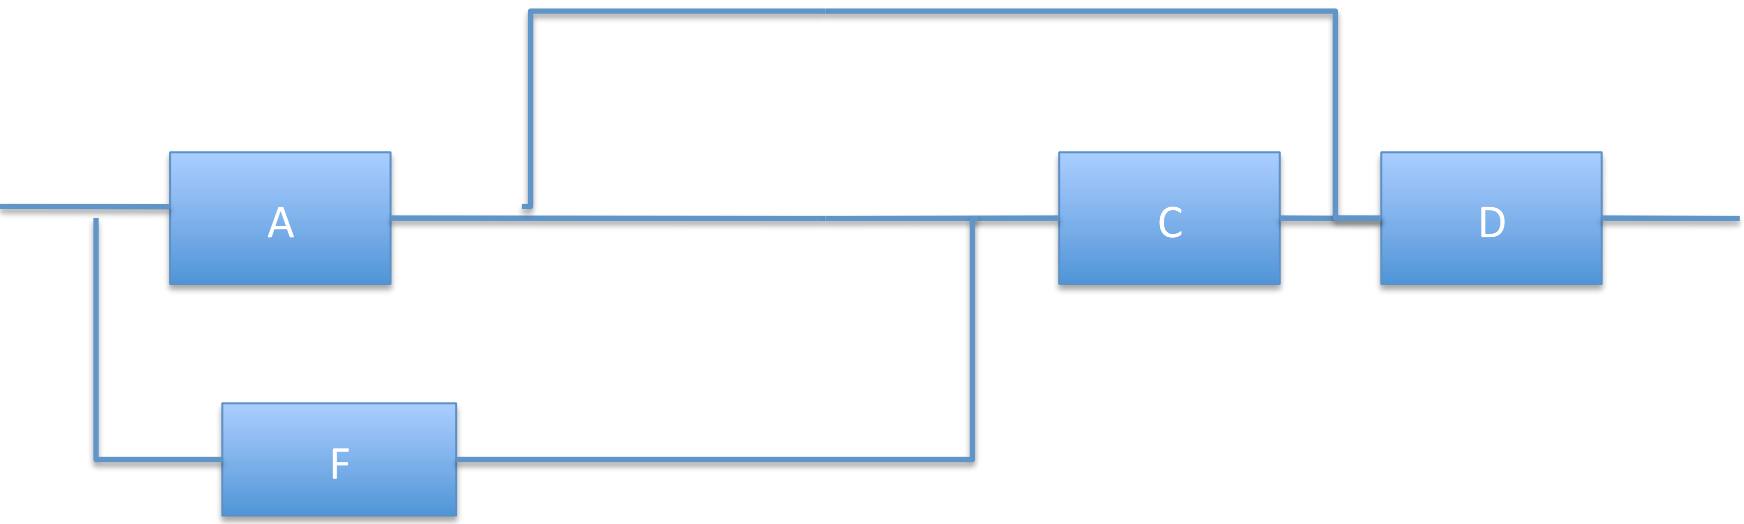

Putting it all together - 1

$P(\text{sys works} \mid E') = [\,1 - (1 - R_A R_B)(1 - R_F)]R_C R_D$

$P(\text{sys works} \mid E \text{ and } B') = [\,1 - (1 - R_A)(1 - R_F R_C)]R_D$

$P(\text{sys works} \mid E \text{ and } B) = [\,1 - (1 - R_A)(1 - R_F R_C)]R_D$

$P(\text{sys}|E) = [\,1 - (1 - R_A)(1 - R_F R_C)]R_D[R_B + R_B'] =$

$= [\,1 - (1 - R_A)(1 - R_F R_C)]R_D$

$P(\text{sys works}) = P(\text{ sys works} \mid E)\, R_E$

$+ P(\text{ sys works} \mid E')(1 - R_E)$

Putting it all together - 2

R_{sys} = P(sys works) =

$[\ 1 - (1 – R_A R_B)(1 – R_F)]R_C R_D\ (1 – R_E)$

$+ [\ 1 - (1 – R_A)(1 – R_F R_C)]R_D(1 – R_E)$

$= [\ R_E R_D + R_C R_D – R_E R_C R_D][R_A + R_F R_C – R_A R_F R_C]$

- Let $R_A = R_B = R_C = R_D = R_E = R_F = R$, then

$R_{sys} = R^6 – 3R^5 + R^4 + 2R^3$

Some tips for non-series-parallel

- The above problem would have been much simpler if I'd picked module A initially
  - Try it yourselves, result should be the same
- Choosing the initial module is crucial.
- Heuristics:
  - Pick modules that are on as many paths as possible
  - Think about whether a single module prevents the system from becoming a serial/parallel system

Non-series-parallel systems

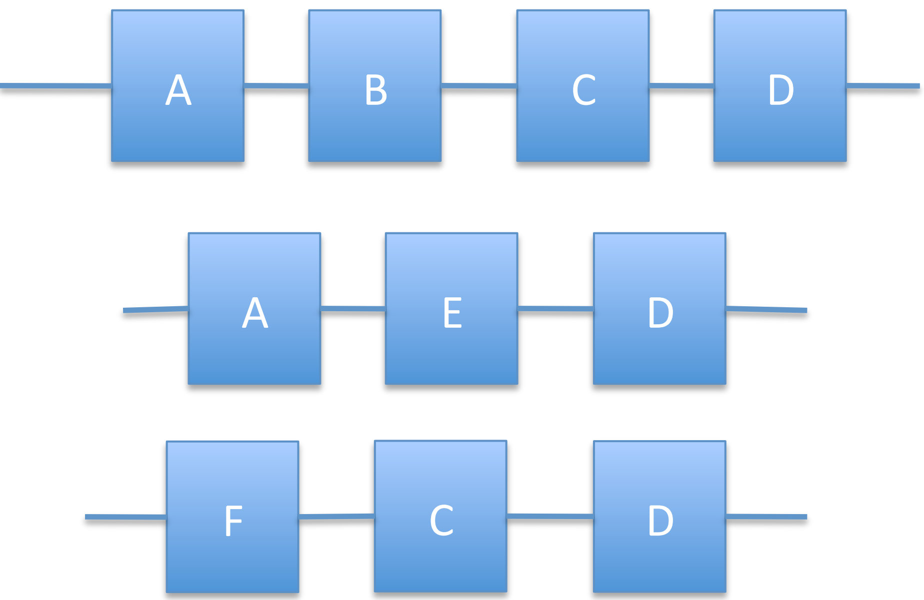

$$Rsys <= 1 – (1 – R^4)(1 – R^3)(1 – R^3)]$$
$$<= 2R^3 + R4 – R^6 -2R^7 + R^{10}$$

- Sometimes all you want is an upper-bound on reliability
- Consider each path separately and treat it as a parallel system (of the paths)
- Why is it an upper bound ?

Learning Objectives

- At the end of this lecture, you will be able to
  - Define combinatorial models of reliability
  - Evaluate the reliability of series, parallel systems
  - Evaluate reliability of non-series, parallel systems
  - Evaluate standby redundancy schemes
  - Model failures using the exponential distribution
  - Evaluate the reliability of TMR and TMR Simplex
  - Understand the pitfalls of single voter in TMR

Standby Redundancy - 1

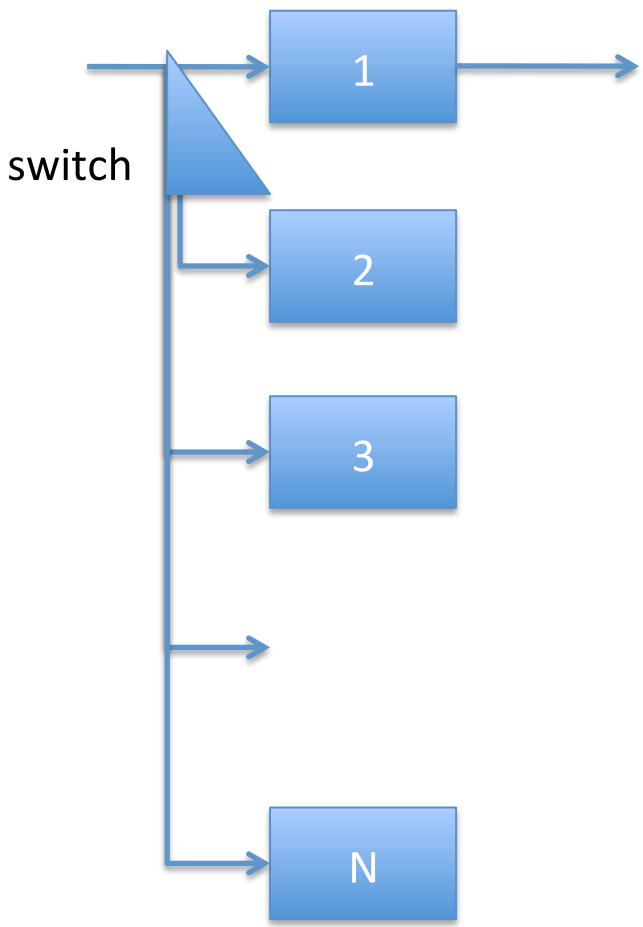

- Consider a parallel system in which another component is activated if and only if the current one fails. A switch detects failure of the system and reconfigures the system around it.

Standby Redundancy - 2

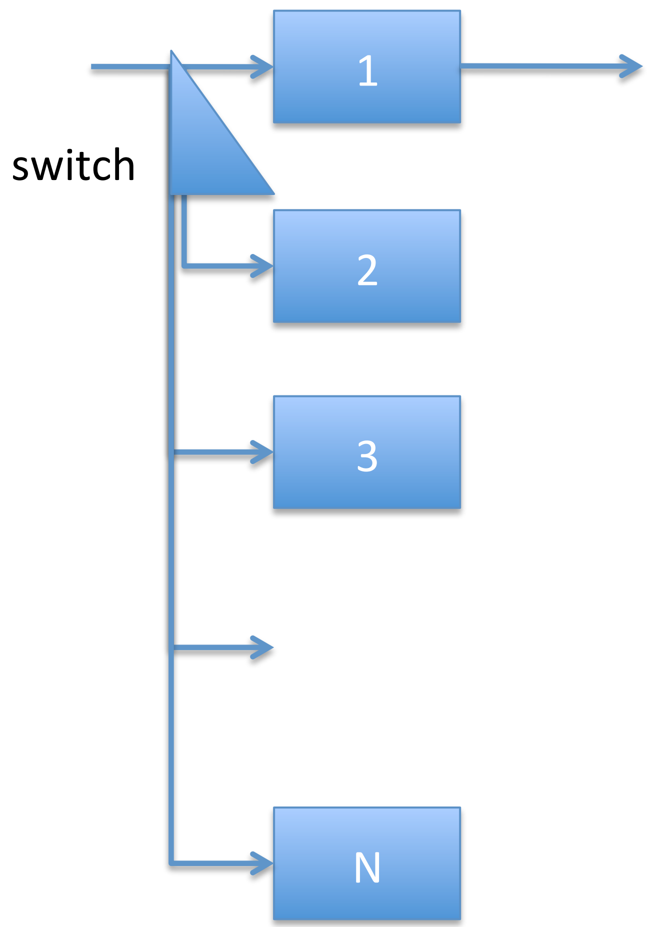

Let each module have a reliability of R. Assume that failures are independent, and that the detection coverage of the switch is c.

$$R_{sys} = R_m + R_m \sum_{i=1}^{n-1} c^i (1 - R_m)^i$$

Standby Redundancy - 3

- Consider a system with standby redundancy where each component has reliability 0.9. Assume that the coverage of the detection mechanism is 0.99. How many modules will you need to achieve a reliability of 0.999 ?

Standby Redundancy : Coverage

- Reliability decreases sharply as coverage drops and saturates !

	c=0.99, R = 0.90	c = 0.99, R = 0.70	c = 0.80, R = 0.90	c = 0.80, R = 0.90
n = 2	0.989	0.908	0.972	0.868
n = 4	0.999	0.988	0.978	0.918
n = inf	0.999	0.996	0.978	0.921

Learning Objectives

- At the end of this lecture, you will be able to
  - Define combinatorial models of reliability
  - Evaluate the reliability of series, parallel systems
  - Evaluate reliability of non-series, parallel systems
  - Evaluate standby redundancy schemes
  - Model failures using the exponential distribution
  - Evaluate the reliability of TMR and TMR Simplex
  - Understand the pitfalls of single voter in TMR

Reliability in terms of CDF

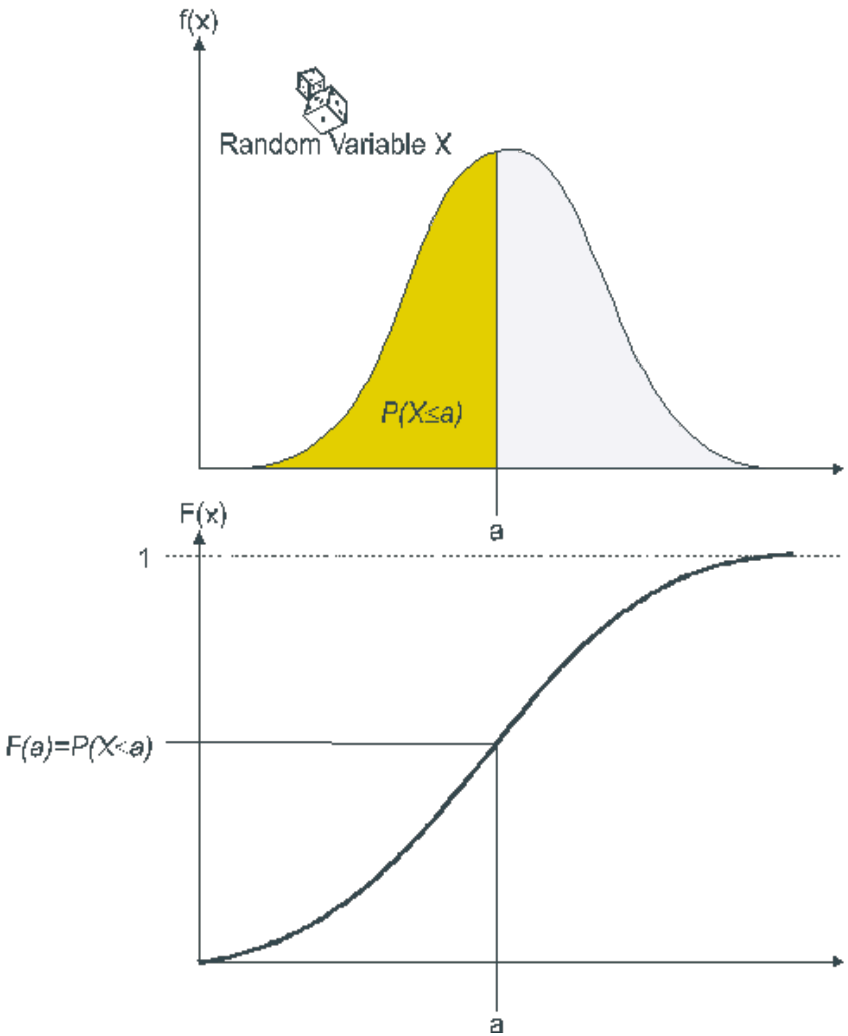

- Let the random variable X denote the lifetime or time to failure of a component.
- **Reliability R(t)** = Prob that component survives up to time t

  $= P(X > t) = 1 - F(t)$

$$= 1 - \int_{-\infty}^{t} f(t)dt \quad = \quad \int_{t}^{\infty} f(t)dt$$

Conditional Reliability

- The previous equation assumes that we started using the component at t = 0 (new). But sometimes we want to evaluate the reliability of a component, given that it has worked until time T (i.e., used components)

$$R(t \mid T) = R(T + t) / R(T)$$

Conditional Reliability (contd..)

- Assume that a component does not age over time. In other words, its survival probability over time (y + t) is independent of its present age t.

$R(y + t) = R(y)\, R(t)$

$(R(t + y) - R(t)) / t = R(t)\, [R(y) - 1] / t$

Taking limit as $t \rightarrow 0$ and as $R(0) = 1$, we get

$$R'(y) = R'(0)\, R(y)$$

Solving the differential equation, $R(y) = e^{\,yR'(0)}$

We set $R'(0) = -\lambda$, then $\mathbf{R(y) = e^{-\lambda y}, y > 0}$

Exponential Distribution

This yields the famous exponential distribution

$F(x) =$ **$1 - e^{-\lambda x}$, $x > 0$**

0 , otherwise

$f(x) = \lambda e^{-\lambda x}$, $x > 0$

0 , otherwise

λ is called the failure-rate in the context of reliability

Exponential Distribution -2

- f(t)

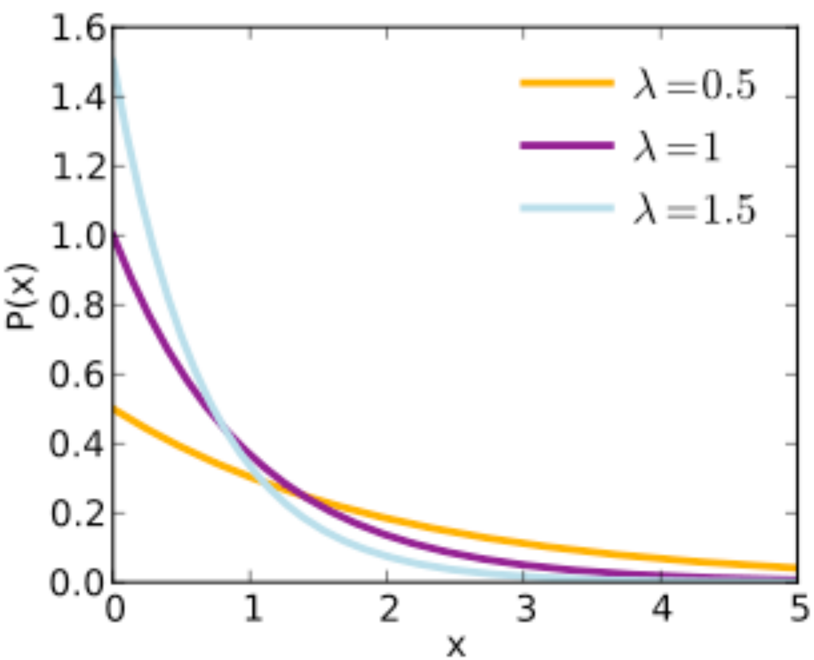

- F(t)

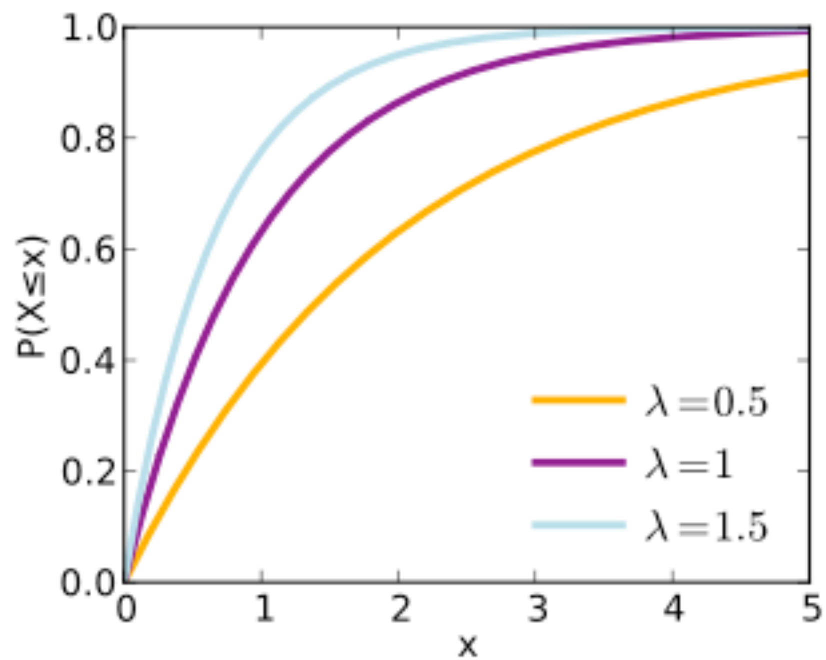

Exponential Distribution -3

Some properties of exponential distributions:

1. $P(X >= t) = \mathbf{e}^{-\lambda t}$
2. $P(a <= X <= b) = \mathbf{e}^{-\lambda a} - \mathbf{e}^{-\lambda b}$
3. Mean time to failure (MTTF) = Mean = $1 / \lambda$
4. **Memory-less property** (we started with this):

   $P(T > y + t \mid T > y) = P(T > t)$

Example: $P(T > 40 \mid T > 30) = P(T > 10)$

Does NOT mean: $P(T > 40 \mid T > 30) = P(T > 40)$

Why is this useful for Reliability ?

- A used component is as good as new, so no need to replace components that are working fine
- In calculating MTTF, reliability etc. we do not need to keep track of history of the system
  - Especially useful for Markov models (later)
- Makes it very simple to reason about reliability as failure rate is a constant (series/parallel systems)

Exponential Failure Rate: Series

- Consider a series system in which each component has exponentially distributed and independent lifetimes, with rates $\lambda_1, \lambda_2, \lambda_3, \ldots, \lambda_n$

R_s = Prob that system works = $R_1.R_2.R_3....R_n$

$= (1 - F_1)(1 - F_2)(1 - F_3) \ldots. (1 - F_n)$

$= e^{-\lambda_1 t} e^{-\lambda_2 t} e^{-\lambda_3 t} \ldots\ldots e^{-\lambda_n t}$

$= e^{-(\lambda_1 + \lambda_2 + \lambda_3 + \lambda_n)t}$ → exponentially distributed

$\lambda_{sys} = (\lambda_1 + \lambda_2 + \lambda_3 + \lambda_n)$

Exponential Failure Rate: Parallel

- Consider a parallel system in which each component has exponentially distributed and independent lifetimes with rates $\lambda_1, \ldots, \lambda_n$

$R_p = 1 - \Pi(1 - R_i) = 1 - \Pi F_i = \mathbf{1 - \Pi(1 - e^{-\lambda_i t})}$

NOTE: The corresponding failure distribution is not exponential, but is a function of its age.

When $\lambda_1 = \lambda_2 = \ldots = \lambda_n$, $\mathbf{R(t) = 1 - (1 - e^{-\lambda t})^n}$

Failures in practice: Bathtub curve

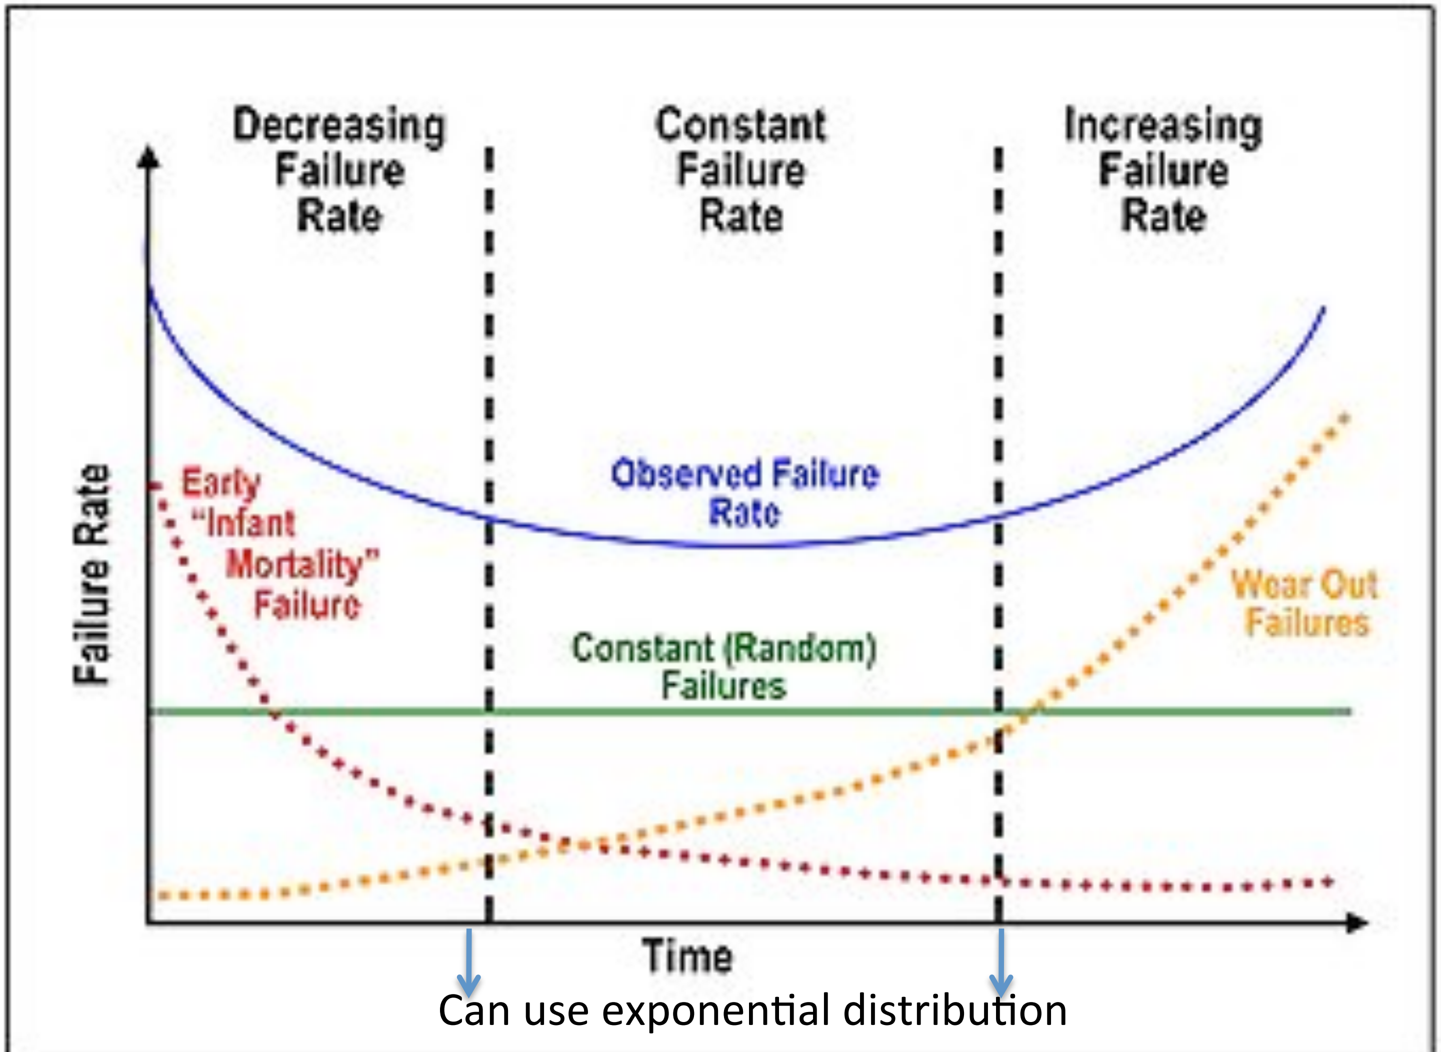

Other Distributions

- Hyper-exponential
- Weibull
- Log-normal
- Normal

Hypo-exponential distribution

- Sometimes you need to add two random variables, X and Y, each of which is exponentially distributed (i.e., Z = X + Y)
  - Z follows a distribution called **Hypo-exponential**

$$f_z(t) = \lambda_1\lambda_2/(\lambda_1 - \lambda_2)[\ e^{-\lambda_2 t} - e^{-\lambda_1 t}\ ]$$

$$F(t) = 1 - [\ \lambda_2/(\lambda_2 - \lambda_1)e^{-\lambda_1 t} - \lambda_1/(\lambda_2 - \lambda_1)\ e^{-\lambda_2 t}]$$

Learning Objectives

- At the end of this lecture, you will be able to
  - Define combinatorial models of reliability
  - Evaluate the reliability of series, parallel systems
  - Evaluate reliability of non-series, parallel systems
  - Evaluate standby redundancy schemes
  - Model failures using the exponential distribution
  - Evaluate the reliability of TMR and TMR Simplex
  - Understand the pitfalls of single voter in TMR

M-out-of-N system

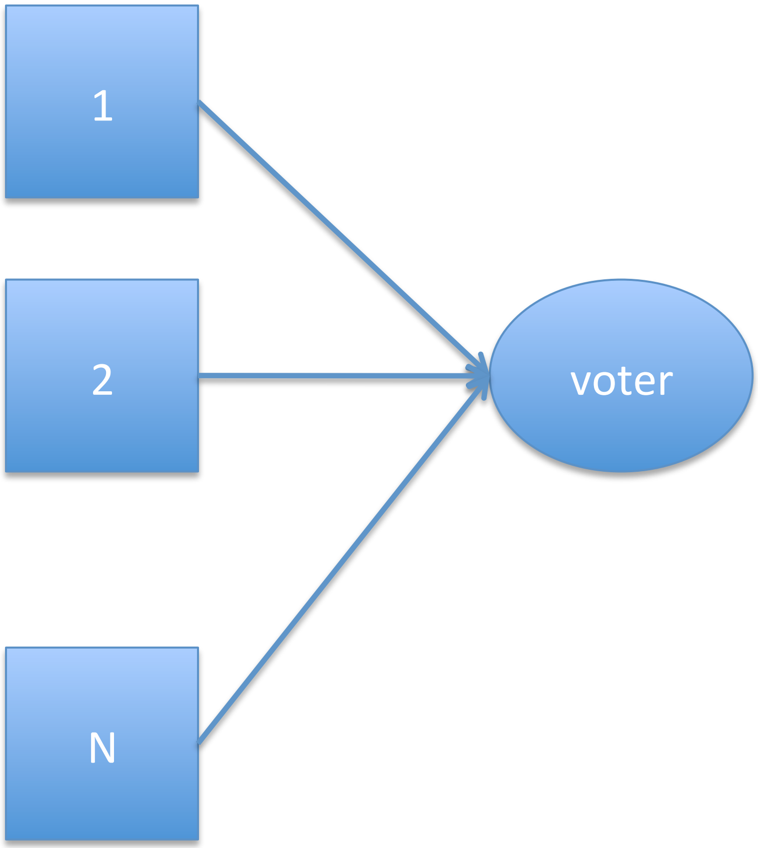

Consider a system with 'N' components such that at least 'M' should work for the system to work.

Component lifetimes are I.I.E.D

$$R_{sys} = \sum_{j=M}^{j=N} \binom{n}{j} R^{j} (1-R)^{n-j}$$

TMR system

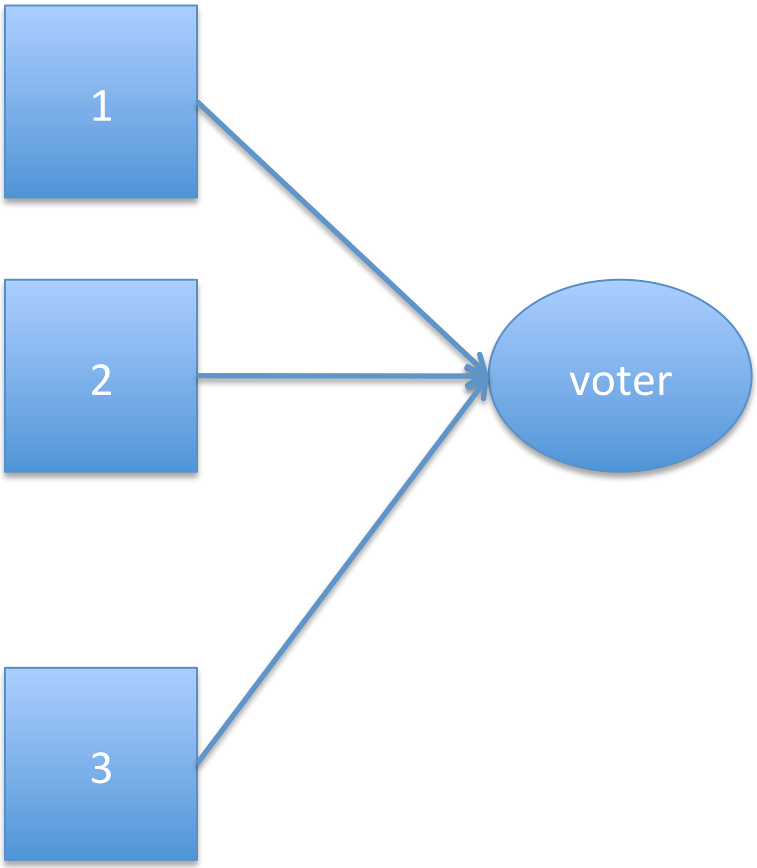

- Special case of NMR where N=3, M=2

$$R_{TMR} = 3R^2(1 - R) + R^3$$
$$= 3R^2 - 2R^3$$

Let $R(t) = e^{-\lambda t}$, then

$$R_{TMR} = 3e^{-2\lambda t} - 2e^{-3\lambda t}$$

$$R_{simplex} = e^{-\lambda t}$$

Reliability of TMR

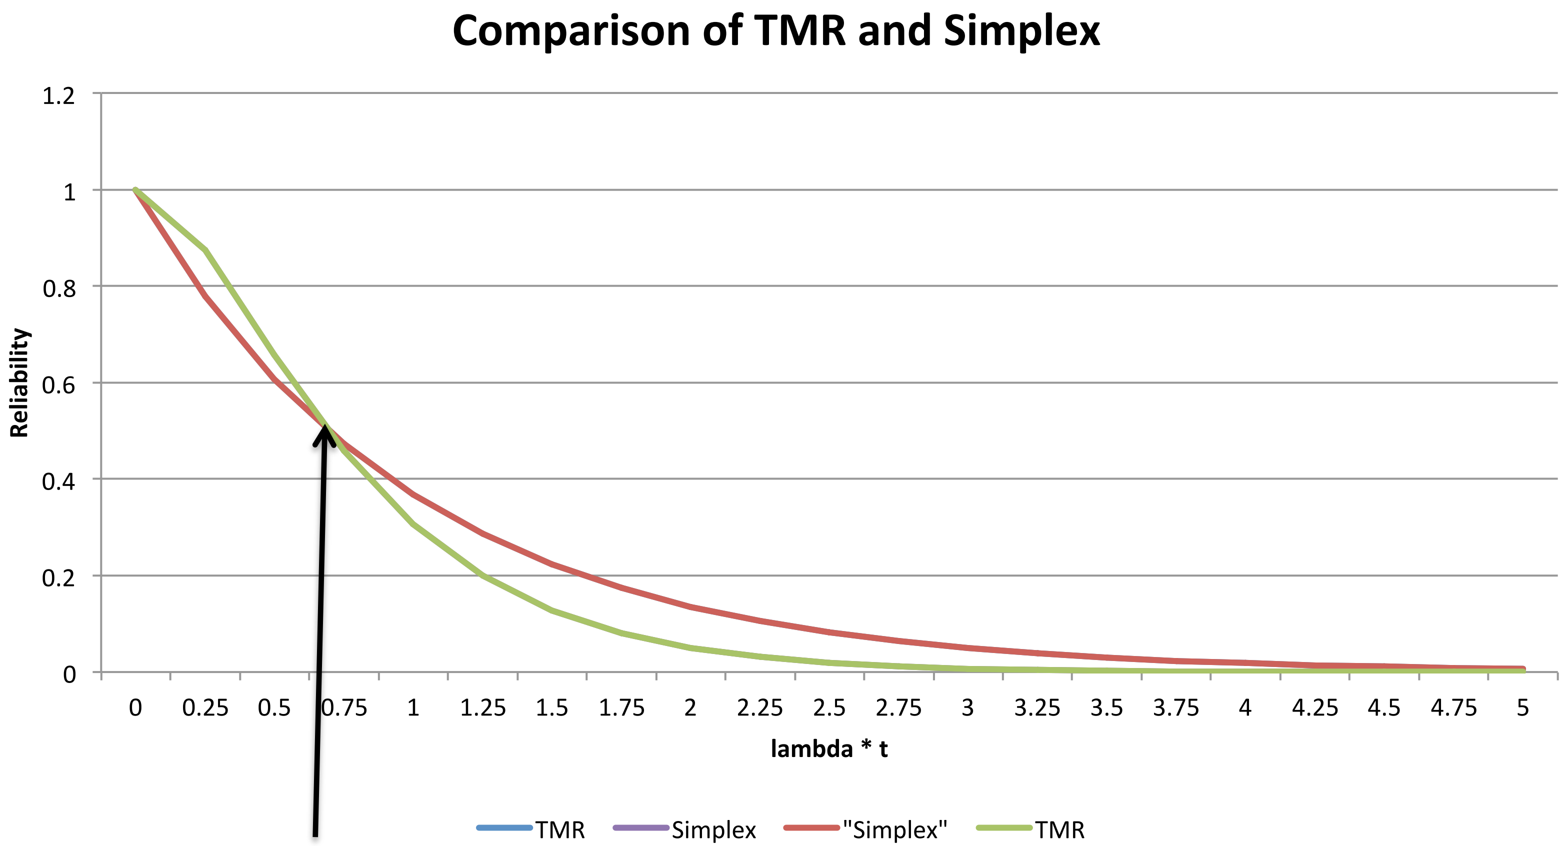

$T_o = \ln 2 / \lambda$
$= 0.7 / \lambda$

For $T > T_0$: $R_{TMR} < R_{simplex}$ (Why ?)

TMR Versus Simplex

- TMR offers substantially better reliability than Simplex for short mission times ($T < T_0$)
  - Used in systems such as airplanes where mission times are typically short (< component lifetime)
  - Not suitable for long missions such as space systems
- After the first failure, the TMR is equivalent to a system of 2 components in series
  - Failure rate is double that of a single system
  - Second component does not provide any benefit

TMR-Simplex System

- Can we combine the advantages of TMR and Simplex in the same system ?
  - After one system fails in a TMR, we switch to a simplex configuration by discarding a component. So this means we discard a good component

$$R(t) = 1 – [\ 1 - \ 3/(3 – 1)\ e^{-\lambda t} + 1/(3 – 1)\ e^{-3\lambda t}\ ]$$

$$= (3/2)e^{-\lambda t} – (1/2)e^{-3\lambda t} \qquad \text{(hypo-exponential)}$$

Reliability of TMR-Simplex

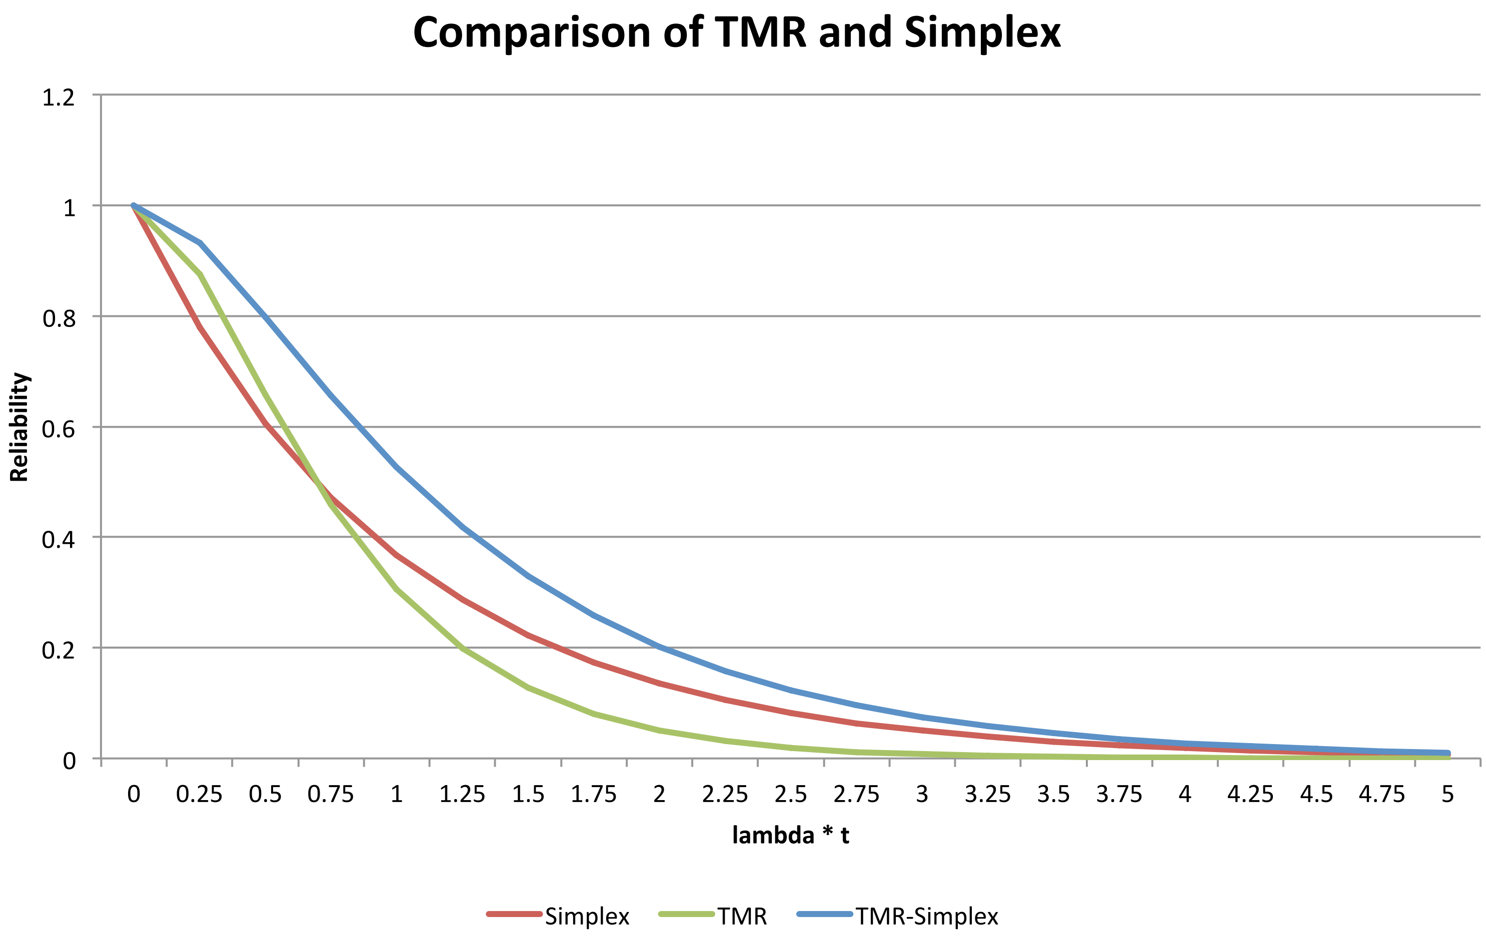

TMR versus TMR-Simplex

- TMR Simplex achieves much higher MTTF than TMR – can be used in long-term missions
- However, the reliability benefits provided by TMR are not available after the first failure
- Also, false-alarms possible if wrong detection
  - May degrade reliability considerably

Learning Objectives

- At the end of this lecture, you will be able to
  - Define combinatorial models of reliability
  - Evaluate the reliability of series, parallel systems
  - Evaluate reliability of non-series, parallel systems
  - Evaluate standby redundancy schemes
  - Model failures using the exponential distribution
  - Evaluate the reliability of TMR and TMR Simplex
  - Understand the pitfalls of single voter in TMR

TMR – Voter Reliability

Voter is a single point of failure in TMR systems (equivalent to a series system with Voter)

$$\mathbf{R_{TMR} = R_V[R_M^3 + 3R_M^2(1 - R_M)\ ]}$$

Reliability is only as good as reliability of Voter

1. Voter can fail silently and discard correct outcomes -> switch to Simplex in worst case
2. Voter can prevent faulty outcomes from being suppressed – much more serious kind of error

TMR with redundant voters

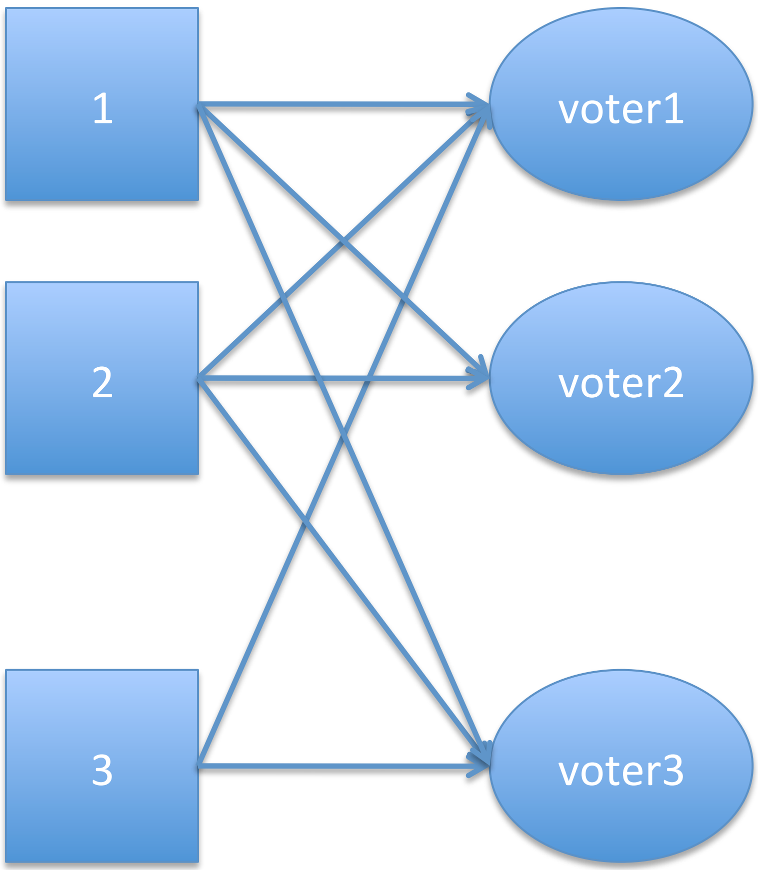

- Use a TMR configuration for the Voter as well.

$R_V = R_V^3 + 3R_V^2(1 - R_V)$

However, there needs to be a single voter somewhere !!!

TMR Voting: Practical issues

- Voter also introduces a performance delay due to variations in clock-speeds/network delays
- What is the right granularity of voting ?
  - instruction-level, module-level, syscall boundaries
- How to handle non-determinism in voting ?
  - Ensuring determinism among replicas is hard
  - Discard non-deterministic state during voting

Learning Objectives

- At the end of this lecture, you will be able to
  - Define combinatorial models of reliability
  - Evaluate the reliability of series, parallel systems
  - Evaluate reliability of non-series, parallel systems
  - Evaluate standby redundancy schemes
  - Model failures using the exponential distribution
  - Evaluate the reliability of TMR and TMR Simplex
  - Understand the pitfalls of single voter in TMR