Mitigating Risk of Cascading Failures in Storage Area Networks

Guixiang Lyu (PhD Advisor: Prof. Liudong Xing)

Electrical and Computer Engineering Department, University of Massachusetts Dartmouth

I. Overview

- Storage area networks (SANs) are a widely used and dependable solution for data storage. Nevertheless, the occurrence of cascading failures caused by overloading has emerged as a significant risk to the reliability of SANs, impeding the delivery of the desired quality of service to users. This poster presents static and dynamic load-triggered reallocation strategies to alleviate the cascading failure risk during the specified mission time. Based on the SAN component reliability assessed using the accelerated failure-time model, the SAN reliability is evaluated using binary decision diagrams. A case study of a mesh SAN is provided to demonstrate and compare the performance of different cascading failure mitigation schemes in terms of the SAN reliability improvement ratio and resulting SAN reliability.

II. An Illustrative Example

- Node degree (d): The number of connections/links that a switch has to other switches in the SAN. For example, d_{sw1}=3, d_{sw2}=4, d_{sw3}=3, d_{sw4}=4, and d_{sw5}=4.
- Parameters:

Switch	Failure rate λ (per hour)	Initial load applied L_0
Sw_1	3.0e-6	15
Sw_2	5.0e-6	50
Sw_3	3.0e-5	5
Sw_4	3.0e-6	1
Sw_5	3.5e-6	8

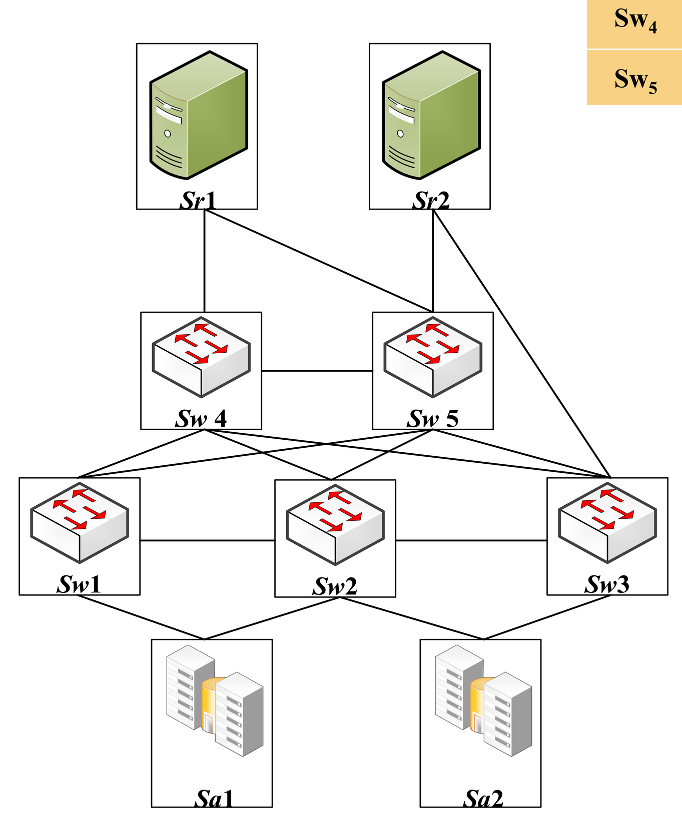

- Fault Tree Modeling:

The example SAN functions as long as at least one server can communicate with at least one storage array.

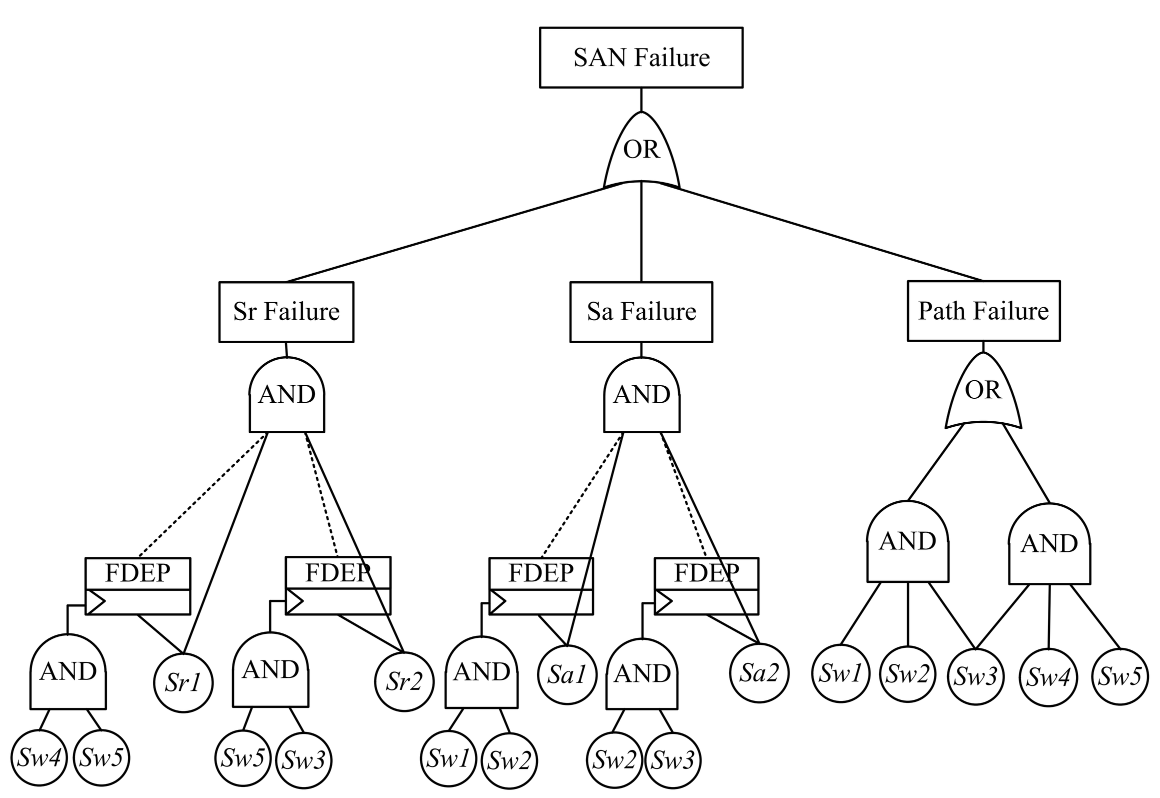

III. Switch Reliability Evaluation

- Accelerated failure-time model (AFTM): Modeling the load-failure rate relationship of a SAN device.
- The reliability of a switch under loading L at mission time t

$$R(t;L) = R_0(t\phi(L))$$

where

$$\phi(L) = \begin{cases} L^{\alpha}, & Power\ law \\ e^{L\alpha}, & Exponential\ law \end{cases}$$

- Illustration (power law) when the baseline time-to-failure distribution is exponential

$$R(t;L_0) = e^{-\lambda(t\phi(L_0))} = e^{-\lambda t L_0{}^{\alpha}} \quad F(t;L) = 1 - e^{-\lambda t L^{\alpha}};\ \lambda(L) = \frac{F'(t;L_0)}{R(t;L_0)} = L^{\alpha}\lambda$$

IV. SAN Reliability Evaluation

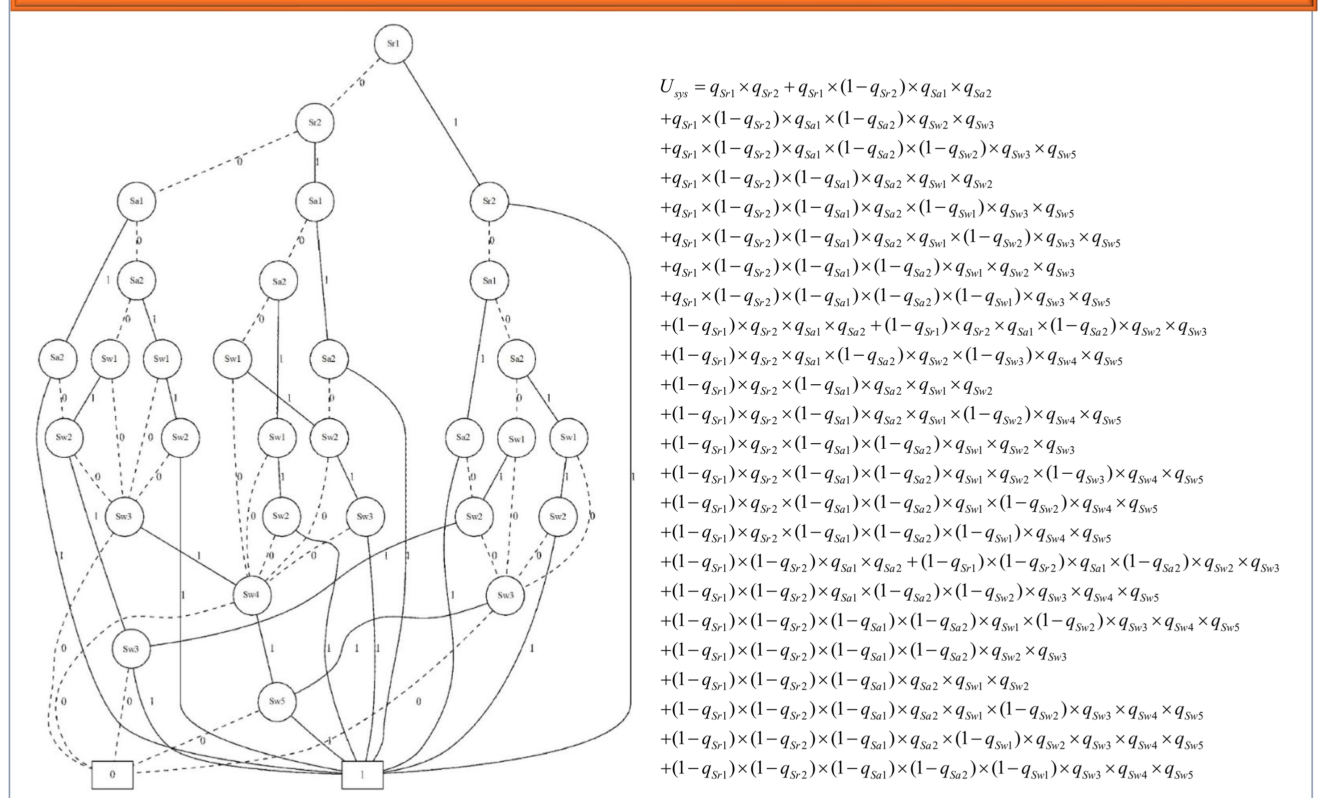

BDD model of the example mesh SAN

- The ordering of component variables:

$Sr_1 < Sr_2 < Sa_1 < Sa_2 < Sw_1 < Sw_2 < Sw_3 < Sw_4 < Sw_5$.

V. Threshold-Triggered Mitigation Strategies

- Load redistribution: Redistributing the workload of selected nodes to other nodes, enhancing the overall SAN system reliability and further mitigating the high risk of cascading failures
- Node Selection Rules: load-sensitive and reliability-sensitive
- Redistribution Rules:

Node degree-based proportional rule

$$\Pi_j = \begin{cases} \dfrac{d_j^{\beta}}{\sum_{m\in N_i} d_m^{\beta}}, & for\ any\ j \in N_i \\ 0, & for\ any\ j \notin N_i \end{cases}$$

Load updating

$$\Delta L_{jk} = L_k \Pi_j$$

$$\bar{L}_k = \Delta L_{kk}$$

$$\bar{L}_k = \sum_{y\in\Phi} \Delta L_{yk}$$

$$\bar{L}_j = L_j + \sum_{k\in\Phi} \Delta L_{jk}$$

- Mitigation Strategies:

	Reliability-Sensitive	Load-Sensitive
Static Threshold	Scheme 1	Scheme 2
Dynamic Threshold	Scheme 3	Scheme 4

VI. Analysis Results

- Performance Comparisons [x-axis: t; y-axis: SAN reliability (left), load of Sw2 (right): $L_{Sw_2} = 0.05 \times t$]

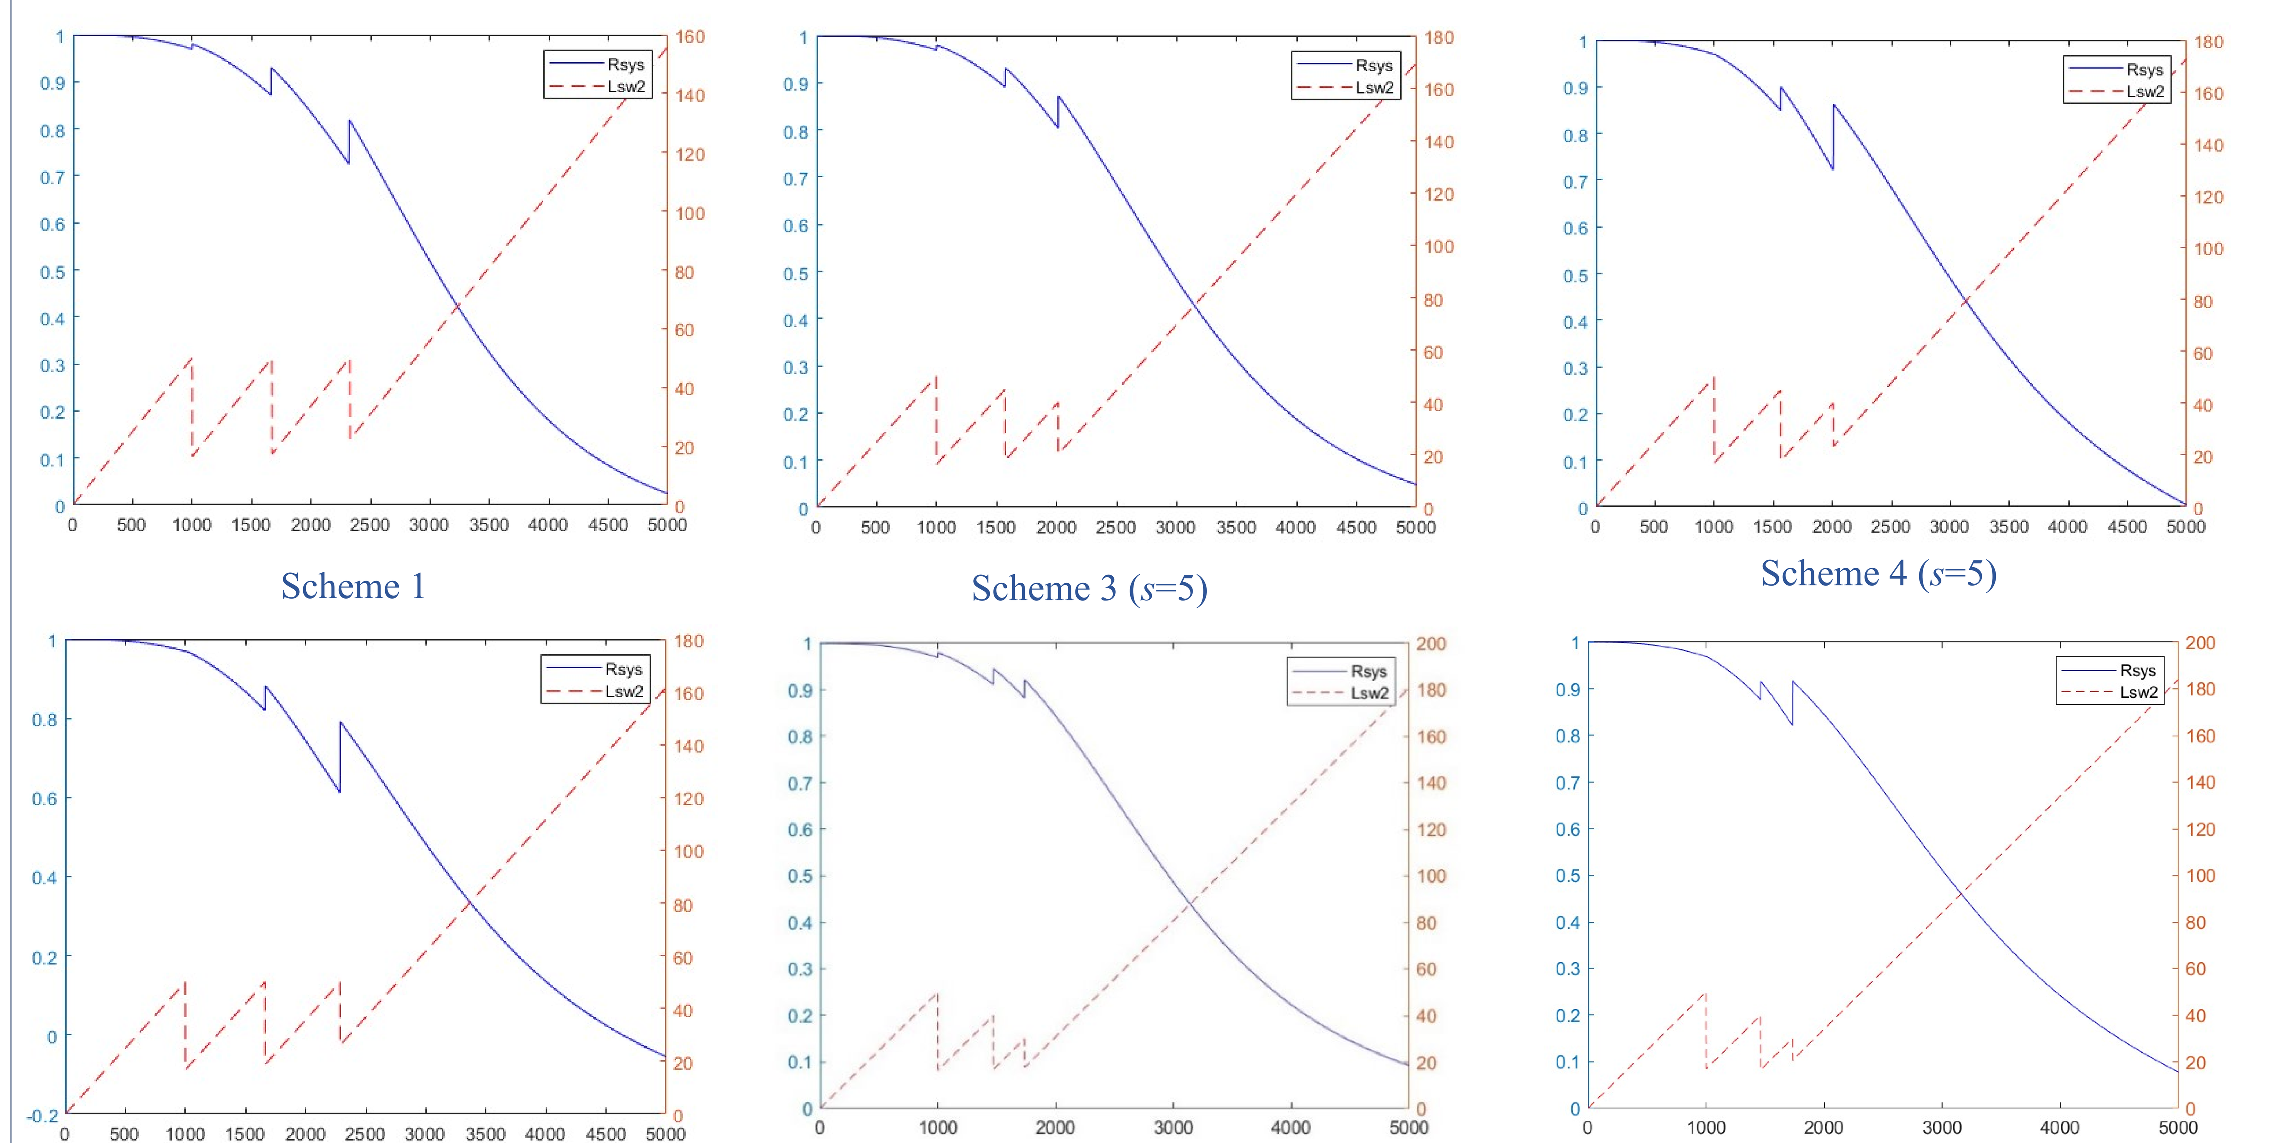

Scheme 1 | Scheme 3 (s=5) | Scheme 4 (s=5)

Scheme 2 | Scheme 3 (s=10) | Scheme 4 (s=10)

- The average SAN reliability improvement ratio (IR):

$$IR_{\text{average}} = \frac{1}{n}\sum_{i=1}^{n} IR_i = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{AR_i - BR_i}{BR_i}\right]$$

	Scheme 1	Scheme 2	Scheme 3		Scheme 4	
	s=5	s=5	s=5	s=10	s=5	s=10
IR_1	0.0108	0.0010	0.0108	0.0109	0.0010	0.0011
IR_2	0.0550	0.0772	0.0462	0.0377	0.0598	0.0453
IR_3	0.1595	0.2960	0.0861	0.0442	0.1962	0.1162
$IR_{average}$	0.0751	0.1247	0.0476	0.0310	0.0857	0.0542

- Performance Comparisons using resulting SAN reliability
  - The SAN reliability at t=2013h under Scheme 1 (0.8104713) and Scheme 3 (0.8731286), and at t=2008h under Scheme 2 (0.7396682) and Scheme 4 (0.8633420).
  - The SAN reliability at t=2286h under Scheme 1 (0.7112488) and Scheme 2 (0.7932698), and at t=2008h under Scheme 3 (0.8055559) and Scheme 4 (0.8633420).
  - The SAN reliability at t=1735 for s=5 (0.8897228) and s=10 (0.9215826) under Scheme 3, and at t=1730 for s=5 (0.8388160) and s=10 (0.9173340) under Scheme 4.

VII. Conclusion

- In terms of the average reliability improvement ratio, the mitigation scheme using the static threshold always outperforms the mitigation scheme using the dynamic threshold regardless of the node selection rule adopted
- In terms of the resulting SAN reliability, the mitigation scheme using the dynamic threshold outperforms the mitigation scheme using the static threshold regardless of the node selection rule adopted
- The load-sensitive selection rule always outperforms the reliability-sensitive selection rule.
- As the step value increases, the dynamic mitigation schemes become less effective in terms of the average reliability improvement ratio but more effective in terms of the resulting SAN reliability.