Economics-based risk analysis of correlated failures

Maurizio Naldi - Giuseppe D'Acquisto

Università di Roma "Tor Vergata"
Dip. di Informatica, Sistemi e Produzione
Rome, Italy

`naldi@disp.uniroma2.it`

Econ@tel meeting
Wien · 27 April 2009

Contents

- Outage statistics
- Shortcomings of current approach
- Reliability and economic risk
- A proposal for the economic approach to reliability

Design trends and reliability

- Network Design trends
  - Cost-reduction strategies
  - Use of modular software
  - Use of COTS (Commercial off-the-self) products
  - Reduction of functional redundancy
- Side effects
  - Increased failure frequency (increase in the number of devices)
  - Wider correlation of failure events

Failure timescales

- Time resolution is
  - 1 day (top picture)
  - 1 hour
  - 5 minutes (bottom picture)
- Failures do occur at all timescales (Iannaccone, 2002)

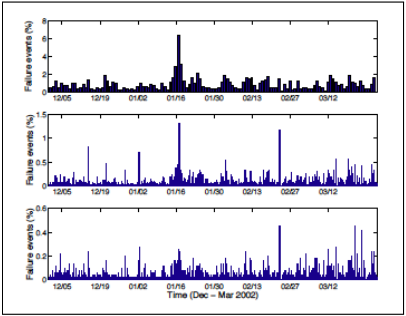

A taxonomy of failures

- Failures may be planned (preventive maintenance) or unplanned
- Failures may concern a single link (individual link failures) or two or more links (shared link failures, which represent correlation)
- Shared link failures may take place simultaneously (identical start and end times) or be overlapping (within a few seconds of one another)
- Failures may concern a router or a transmission device (optical)

Statistics on Sprint's network (Markopoulou et alii, 2008)

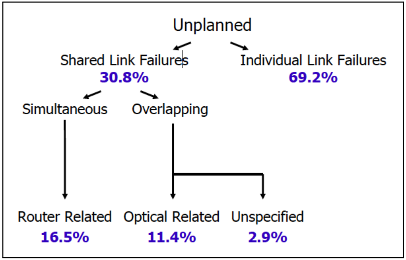

An example of failure concurrency

A case of triple protection on a base station in a mobile network

Power supply is guaranteed by a triple line of protection

- Commercial AC power
- Backup AC generators
- An 8-hour battery backup if there is a dual AC power failure

One outage involved loss of power under this scenario:

1. Lightning caused a loss of commercial AC power.
2. The same lightning strike damaged the AC generator.
3. The alarm system to this un-staffed facility was either not enabled or not tested after installation.
4. After 8-hours the entire facility went down when the batteries were depleted.

Shortcomings of the traditional approach to reliability

- Correlation between within-network failures is seldom considered
- Interdependence between networks is not considered
- Network-wide measures of reliability (e.g. connectivity) are often of the ON-OFF kind
- Deeper consequences of failures are not considered

Challenges for a novel approach

- Failures are often correlated or depend on a common cause
- Identical software may be installed on many devices
- Deeper consequences of failures should be considered
- Failures differ as to their consequences
    - Number of customers affected
    - Number of services affected
    - Degrees of severity of impairment
    - Economic consequences

An economic approach to reliability

Network failures are relevant in relation to the economic loss they cause

An economic approach to reliability

Network failures are relevant in relation to the economic loss they cause

Both direct losses and hidden costs should be considered

- Lost revenues
- Penalties for breach of SLA conditions
- Recovery costs

An economic approach to reliability

Network failures are relevant in relation to the economic loss they cause

Both direct losses and hidden costs should be considered

- Lost revenues
- Penalties for breach of SLA conditions
- Recovery costs

Greater efforts should be devoted to improve the reliability of those devices whose failure has larger economic impact

Metrics for economic risk

- Poor reliability has to be addressed in economic terms
- We need a (simple) measure of the economic loss associated to the failure risk for
  - Design decisions
  - Protection and recovery policy
  - Insurance
- The risk measure should possess some desirable properties

Coherent measures of risk

A relevant class of measures of risk $\rho(X)$ is represented by the *coherent* class with the following properties (Artzner et alii, 1999)

- Monotonicity $X_1 \geq X_2 \Longrightarrow R(X_1) \geq R(X_2)$
- Subadditivity $R(X_1 + X_2) \leq R(X_1) + R(X_2)$
- Homogeneity $R(\alpha X) = \alpha R(X)$, $\alpha \geq 0$
- Translational invariance $R(\alpha + X) = \alpha + R(X)$

A measure of risk

The Value-at-Risk (VaR) is the loss that is not exceeded with a prescribed probability

$$\texttt{VaR}(X;\beta) = F_X^{-1}(\beta)$$

Properties

- Homogeneity: Yes
- Monotonicity: Yes
- Translational invariance: Yes
- Sub-additivity: No

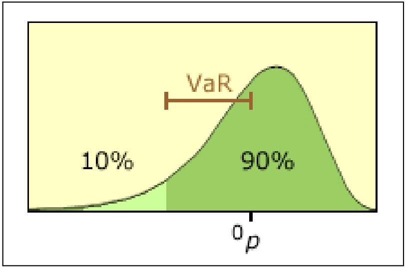

The Tail Value at Risk

Since the VaR doesn't consider the value of the losses incurred beyond the VaR itself, a better measure of risk may be the Tail Value at Risk (T-VaR), defined as the average Value at Risk

$$\texttt{T-VaR}(X;\beta) = \frac{1}{1-\beta}\int_{\beta}^{1}\texttt{VaR}(X;\xi)d\xi$$

The T-VaR is also named Expected Shortfall and is a coherent measure of risk, and is related to both the expected loss and the VaR
$\texttt{T-Var}(X;0) = \mathbb{E}[X]$
$\texttt{T-Var}(X;\beta) \geq \mathbb{E}[X]$
$\texttt{T-Var}(X;\beta) \geq \texttt{VaR}(X;\beta)$

A reliability-oriented model of the network

- The overall set of customers/services is divided into a number of service basins
- Each basin represented a homogeneous group of customers using a specific service
- Basic characters of homogeneity are the contract conditions and the level of consumption (traffic/revenues)
- In each service basin service is accomplished by a number of devices (possibly all customers in a service basin are served by the same devices)
- A device may serve multiple service basins
- The service to a service basin is disrupted if any of the basin devices fails

Mathematical formulation of the loss model

- The number of service basins is N
- The number of devices supporting the i-th basin is M_i
- The state of the j-th device in the i-th basin is represented by the binary variable Y_{ij}
- The state of the i-th basin is represented by the state variable $S_i = \max(Y_{i1}, \ldots, Y_{iM_i})$
- The loss associated to the disruption of the i-th basin is a_i
- The overall loss is $L = \sum_{i=1}^{N} a_i S_i$

A latent variable model for the single subsystem

- The state of each subsystem is determined by a continuous latent variable $Y_{ij} = \mathbb{I}(X_{ij} > b_{ij})$, where the threshoold b_{ij} is set so to match the marginal failure probability for the subsystem
- Each latent variable incorporates the effects of its individual risk factor η, a number of joint risk factors Z and a common shock factor W

$$X_{ij} = \frac{\sum_{k=1}^{D} \rho_{ik} Z_k + \alpha_{ij} \eta_{ij}}{W}$$

Two versions may be considered, depending on the characteristics of the shock factor:

- Normal Copula
- T-Copula

The Copula models

In the Normal Copula

- The shock factor is absent
- The joint risk factors are i.i.d. random variable following a standard normal distribution
- The latent variable follows a standard normal distribution

In the T-Copula

- The shock factor is the square root of a chi-square variable
- The joint risk factors are i.i.d. random variable following a standard normal distribution
- The latent variable follows a t-Student distribution

A toy model

A preliminary analysis has been performed on a toy network

- No. of service basins $N = 100$
- A single subsystem for each service basin $M_i = 1, \forall i$
- A single common risk factor $D = 1$
- No shock factor $W = 1$
- Subsystem failure probability $=0.16$
- Loss due to a service basin $a_i = 1, \forall i$

Simulation results

Impact of the weight factor for VaR=50

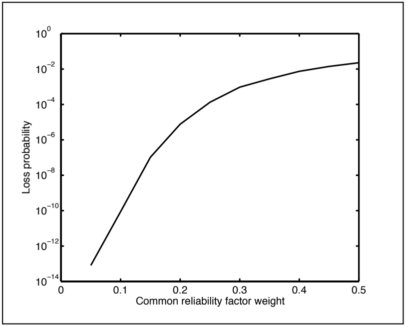

VaR for $\alpha = 0.1$

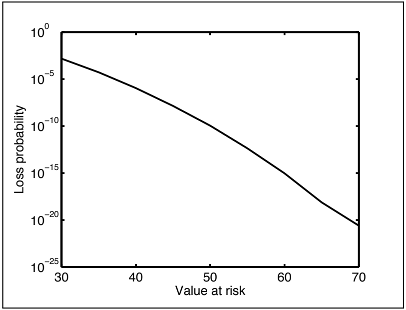

Open issues

- Instantiating the general framework on a real network
- Identifying the common risk factors
    - Too small = The system is not represented adequately
    - Too large = Too many parameters
- Setting the thresholds for the latent variables (relatively easy: inverting the latent variable distribution)
- Setting the values of the correlation parameters (calibration)