Nonparametric Statistical Methods to Analyze the Internet Connectivity Reliability

Dimitri Papadimitriou[1], Davide Careglio[2]

[1]Bell Labs
Antwerp, Belgium

[2]Technical University of Barcelona
Barcelona, Spain

IEEE CQR'15

Outline

Outline

1. Motivations and Objectives
2. Model and Preliminaries
3. Nonparametric Statistical Methods
4. Case Study
5. Conclusion

Motivation and Objectives (1)

Motivations

- Transient but frequent changes in spatio-temporal properties of routing paths may affect performance of corresponding forwarding paths (connectivity)
- Frequent instabilities when observed for same (subset of) path(s) and attributable to spatially localized portion(s) of the Internet may reveal failure-prone physical topology
  - Stability: spatial changes affecting sequence of (abstract) nodes and edges of paths
- Provide for longer term prediction of Internet routing-forwarding system performance using well-proven statistical analysis accounting for recurrence of events and correlation between various events

Objectives

- Characterize the dynamic properties (in particular, stability properties) of the Internet routing and corresponding forwarding path(s).
- Quantitative evaluate the reliability of the Internet connectivity (also referred to as reachability in computer networking)
- Predict its evolution over time without requiring to infer the parameters of the distribution characterizing network failure probability and rate.

Motivation and Objectives (2)

Generalized Weibull distribution

- Modeling reliability of engineered systems and their components (e.g., physical links) by their failure probability at specific time as well as their failure rate variation over time
- Applicable when modeling simultaneous and/or correlated failures together with their probability distribution, including node failures and common resource failures which lead to the failure of numerous network paths
- When each link failure rate depends on different parameters, leads to consider multi-variate joint distributions (NOT simple product of individual distributions)

Nonparametric statistical methods

- Instead of model-driven statistical methods, apply data-driven nonparametric statistical methods (Kaplan-Meier survival probability estimator and Mean Cumulative Function)
  - Unlike parametric statistics, makes no assumptions about the probability distributions of the variables under study
  - Requires few or no assumptions about the populations from which data are obtained
  - Includes both descriptive and inferential statistics not relying on the estimation of parameters (e.g., mean or variance) describing distribution of variable(s) of interest in the population

Outline

Model (1)

Network topology

- Modeled as undirected graph $G = (V, E)$ where,
- V: finite vertex set ($|V| = n$)
- E: finite arc/edge set ($|E| = m$)

Network paths

- $\mathcal{P}(u, v) \triangleq$ set of all paths $p(u, v)$ from vertex u to v
- Loop-free path $p(u, v) \in \mathcal{P}$ from u to $v \triangleq$ finite sequence $[v_0(= u), v_1, \ldots, v_{i-1}, v_i, \ldots, v_p(= v)] \mid v_{i-1}$ adjacent to $v_i, \forall (v_{i-1}, v_i)_{(i=1,\ldots,p) \in E}$
- Path instability (or perturbation) characterized by a change (or interruption if detectable) in the sequence of vertices along path $p(u, v)$

Distance vs. Length

- Length of path $p(u, v) \in \mathcal{P}$: number of edges the path $p(u, v)$ traverses from u to v
- Distance $d(u, v)$ between vertex u and v: minimum length path from u to v

Model (2)

Distinction between topological, forwarding and routing paths

Topological path $p(u, v)$

Path from u to v as computed from the topology graph G
$\Rightarrow$ Topological distance defined by the shortest distance (topological) path on G

Routing path $r(u, v)$

Not necessarily min.length path from u to v as produced at u by the distributed routing algorithm using as input information on G
$\Rightarrow$ Routing topology $\triangleq$ sub-graph H of G representing actual nodes and links along the paths as selected/computed by the routing algorithm

Forwarding path $f(u, v)$

Path followed by the traffic directed from u to v derived at u from local routing tables
$\Rightarrow$ Forwarding topology $\triangleq$ sub-graph H' of G representing actual nodes and links as selected by router's forwarding decision
Note: sub-graphs H and H' not required to be identical (a routing table entry may exist without a corresponding forwarding entry)

Elements of Reliability Theory (1)

T: continuous random variable representing the failure time (lifetime) of a physical system

Cumulative probability distribution function $F(t)$

- Probability that the system will fail by time t

$$F(t) = P[T \leq t] = \int_0^t f(x)\,dx \tag{1}$$

- Failure probability density function (p.d.f.) $f(t) \triangleq$ expected number of failures experienced in a given time interval

$$f(t) = \frac{dF(t)}{d(t)} \tag{2}$$

Reliability function $R(t)$ (or survival function $S(t)$)

Probability that the system survives at least until time t

$$R(t) = P[T > t] = 1 - F(t) = \int_t^\infty f(x)\,dx \tag{3}$$

$$R'(t) = \frac{d}{dt}R(t) = -f(t) \tag{4}$$

Elements of Reliability Theory (2)

Average failure rate $\lambda(t)$

If system survived up to time t (no failure event before time t) and failure occurs during $[t, t + \Delta t]$; then average failure rate $\lambda(t)$ during time interval Δt:

$$\lambda(t) = \frac{F(t + \Delta t) - F(t)}{\Delta t R(t)} = \frac{R(t) - R(t + \Delta t)}{\Delta t R(t)} \tag{5}$$

Hazard (a.k.a. failure rate) function $h(t)$

Instantaneous failure rate at time t, given that the system survived until time t:

$$h(t) = \lim_{\Delta t \to 0} \frac{F(t + \Delta t) - F(t)}{\Delta t R(t)} = \lim_{\Delta t \to 0} \frac{R(t) - R(t + \Delta t)}{\Delta t R(t)} \tag{6}$$

$$= \frac{f(t)}{R(t)} = \frac{-R'(t)}{R(t)} = -\frac{d(\ln R(t))}{dt} \tag{7}$$

Cumulative hazard or failure rate function $H(t)$

Accumulation of the hazard over time (quantifies number of times one would expect to observe failure event in a given time period, if the event was repeatable): $H(t) = \int_0^t h(x)\,dx$

Example: 2-parameter Weibull distribution

Scale parameter $b > 0$ (or slope), and shape parameter $c > 0$ (or characteristic life):

$$\text{Probability density function: } f(t) = \frac{c}{b}(\frac{t}{b})^{c-1} \exp{(-\frac{t}{b})^c} \tag{8}$$

$$\text{Probability distribution function: } F(t) = \int_0^t f(x)\, dx = 1 - \exp{(-\frac{t}{b})^c} \tag{9}$$

$$\text{Probability that system survives until time t: } R(t) = \int_t^\infty f(x)\, dx = \exp{(-\frac{t}{b})^c} \tag{10}$$

$$\text{Instantaneous failure rate: } h(t) = \frac{c}{b}(\frac{t}{b})^{c-1} \tag{11}$$

Practice

- Commonly used in survival analysis and reliability engineering
- But unable to capture the behavior of a lifetime data set that has a non-monotonic failure rate function

Elements of Reliability Theory: Generalized Weibull Distribution

- Fundamental relationship: $R(t) = exp(-aH(t))$ with $a > 0$
  $\Rightarrow$ With suitable choice of $H(t)$ one can obtain a bathtub shaped failure rate distribution
- Generalized Weibull distributions: reference for modeling reliability of engineered systems and their components (e.g., physical links) by their instantaneous failure probability, failure rate variation over time, etc.

Practice

- Link failure rate depends on different shape and scale parameters
  $\Rightarrow$ Multi-variate joint distributions that is not simply the product of the individual distributions when modeling simultaneous and/or correlated failures, e.g., node failures, or common resource failures
- Additional parameters (coupling effects $\nu > 0$ and time thresholds $\tau \geq 0$) required to model the joint survival distribution $R_{\mathcal{K}}(t)$ of set $\mathcal{K}$ comprising k components with individual failure rate λ_k: $R_{\mathcal{K}}(t) = \exp\left(\tau_k^\nu - \left[\tau_k + \sum_{i=1}^{k}(\lambda_i t_i^{c_i})\right]^\nu\right)$
- Assuming that parameters can be inferred out of observations, resulting model is computationally intractable without simplifying assumptions

Conclusion: instead of modeling and analyzing reliability of Internet connectivity starting from joint failure probability of different components underlying the network topology (bottom-up structural approach), top-down statistical perspective

Outline

Kaplan-Meier Survival Probability Estimation (1)

Assumptions and Notation

- k distinct event times: $t_1 < \ldots < t_i < \ldots < t_k$
- Number of failures (or deaths) at event times: $d_1 < \ldots < d_i < \ldots < d_k$
- At t_i, risk set n_i:
  - Without censoring: n_i = number of survivors prior to time t_i (original sample − all those that experienced the event before time t_i)
  - With censoring: n_i = number of survivors − number of losses/withdrawals (censored cases)

Kaplan-Meier estimator $\hat{S}(t)$

- Probability of surviving beyond time t_{k+1}: $S(t_{k+1}) = P[T > t_{k+1}]$
- $S(t_{k+1})$ depends conditionally on the probability of surviving beyond time $t_k : P[T > t_k] = S(t_k) \Rightarrow S(t_{k+1}) = S(t_k) \times P[T > t_{k+1} | T > t_k]$
- **Goal**: build iteratively a numerical estimate $\hat{S}(t)$ of true survival function $S(t)$

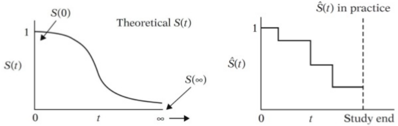

Kaplan-Meier Survival Probability Estimation (2)

Procedure

- Conditional probability formula $P(A \cap B) = P(A) \times P(B|A)$, $\forall t \in [t_i, t_{i+1})$
  - A: event to survive to time t_i
  - B: event to survive from time t_i up to some time t before t_{i+1}
  - $A \cup B$: event to survive to beyond time t before t_{i+1}
- Estimated probability $P[T > t]$ to survive at $t \in [t_k, t_{k+1})$ given by the Kaplan-Meier (piecewise) estimator $\hat{S}(t)$:

$$\hat{S}(t) = (1 - \frac{d_1}{n_1})(1 - \frac{d_2}{n_2}) \cdots (1 - \frac{d_k}{n_k}) = \prod_{i=1:t_i \leq t}^{k} [1 - \frac{d_i}{n_i}] \tag{12}$$

  - Proportion that failed at event time t_i: $\frac{d_i}{n_i}$
  - Proportion that survived event time t_i: $1 - \frac{d_i}{n_i}$

$$\hat{S}(t) = (1 - \frac{d_1}{n_1})(1 - \frac{d_2}{n_2}) \cdots (1 - \frac{d_k}{n_k}) = \prod_{i=1:t_i \leq t}^{k} [\frac{n_i - d_i}{n_i}] \tag{13}$$

  - Number of surviving entities after event time t: $n_i - d_i$
  - Number of surviving entities at risk in the interval just prior to time t: n_i

Kaplan-Meier Survival Probability Estimation (3)

Cumulative Hazard Function Estimation

- Instantaneous hazard function $h(t)$ (a.k.a. hazard rate, conditional failure rate) $\triangleq$ event rate at time t conditional on surviving up to or beyond time t

Cumulative hazard function $H\hat(t)$

- Integral of the instantaneous hazard rates from time 0 to t: $H(t) = \int_0^t h(x)\,dx$
- Represents accumulation of hazard over time
- Estimated Peterson method (1977):

$$\hat{H}(t) = -\ln(\hat{S}(t)) \tag{14}$$

Kaplan-Meier Survival Probability Curve

Kaplan-Meier survival probability curve

- Probability of surviving in a given length of time while considering time in many small intervals
- Kaplan-Meier curve
  - Can take into account some types of censored data (right-censoring) if an entity is withdrawn from a study, i.e. lost from (original) sample before the final outcome is observed
  - Small vertical tick-marks indicate losses, where a entity's survival time has been right-censored.
  - When no truncation or censoring occurs (which is the case for repairable systems) Kaplan-Meier curve is the complement of the empirical distribution function.

Recurrent Event Data Analysis (RDA) (1)

Recurrent event data

- Multiple events occurrence (across observation periods) with possible correlation
  - Example: node failure observed through the occurrence of multiple link failures
- Temporal trajectories of observed data often very complex to determine and their statistical modeling leads to generalized multivariate distributions (difficult to use in practice for analytic or predictive purposes)

Consequences

- Parametric statistical models may not be flexible enough to capture their main features and (stochastic) temporal networks where the sequence of activation times is a stochastic model that preserves the observed inter-event distribution difficult to apply
- Nonparametric (or semi-parametric) statistical models: mean structures are modeled non-parametrically (or semi-parametrically) and distributional assumptions are non-parametric

Recurrent Event Data Analysis (RDA) (2)

RDA vs. LDA

- RDA commonly used in various engineering fields and is particularly useful when performing reliability analysis of repairable systems
- Compared to Life Data Analysis (LDA)
  - Focuses on time to event occurrence data
  - Assumes that events (failures) are independent and identically distributed (i.i.d) whereas in certain situations, the events are dependent and not identically distributed (common property of repairable system data)

Non-parametric RDA

- **Goal**: model number of occurrences of events over time rather than length of time prior to first event occurrence (compared to LDA)
- Non-parametric RDA
  - Provides a non-parametric graphical estimate of mean cumulative number of recurrences of events versus time
  - Non-parametric analysis method relies on Mean Cumulative Function (MCF)

MCF Estimation Procedure

To compute MCF $M(t_i)$ (estimate MCF at t_i):

- At each observation time t_i, the number of events n_i that occurred since the previous observation time t_{i-1} is recorded
  - Recurrent events (non-fatal), e.g., failure events followed by restoration
  - Events assumed to occur randomly
- Number of events n_i divided by number ρ_{i-1} of pairs observable at time t_{i-1} (with ρ_1 set to the total number of initially observable number of entities)
- Compute MCF estimate using the formula:

$$M(t_i) = \frac{n_i}{\rho_{i-1}} + M(t_{i-1}) \tag{15}$$

$$M(t_1) = \frac{n_1}{\rho_0} \tag{16}$$

- All observation intervals: nonrandom, identical for all observations, of equal length

Recurrent Event Data Analysis (RDA) (4)

MCF Plot

- The shape of the MCF plot can reveal several important properties about the behavior of the recurrent events under consideration in a reliability study
- The MCF vs. time (age) curve can be numerically differentiated to obtain the slope, called recurrence rate

MCF Plot Interpretation

- Analyze shape of MCF plot: derive interpretation assuming that instability events induce transient changes in paths properties that affect their performance and operating conditions (hence, their reliability):
  - Constant recurrence rate: MCF plot increases monotonically, slope remains constant → events under consideration occur at constant rate
  - Increasing recurrence rate: MCF plot is convex (slope increases), recurrence rate increases over time → system performance degradation over time
  - Decreasing recurrence rate: MCF plot is concave (slope decreases), recurrence rate decreases over time → maintenance improvement over time (decreasing repair rate)

Outline

- Characterize and analyze dynamic properties of the forwarding and the routing paths as well as their relationships
- Determine whether routing paths follow mainly the perturbations experienced by the forwarding paths or vice versa (causality effect or not ?)
- Method: detection and identification of perturbation events following the methods and procedures documented in [1] based on the stability criteria and metrics introduced in [2] [3]

Case Study: Data

Routing Paths

- Routing path information extracted from BGP datasets provided by RouteViews project
- Datasets collected over 50 days from monitored BGP routers:
  - complete Routing Information Base (RIB) entries (updated every two hours)
  - received BGP routing updates received from peering AS (separated in files recorded every 15 minutes)

Forwarding Paths

- Extracted from data recorded by RADAR tool
  - Measurements: traceroute-like probes initiated from a set of monitoring nodes
  - Targets: large set of IP address prefixes distributed across the Internet
- Based on these measurements, RADAR builds ego-centered views of forwarding topology, i.e., initiating router collects traces along forwarding paths that it probes
- Subset of forwarding paths traced by RADAR corresponds to the routes obtained from RouteViews dataset → subset of monitored routing paths also monitored by RADAR

Note: In total, analyzed dataset includes 1000 forwarding - routing path pairs over 50 days

Case study: Kaplan Meier Survival Probability Estimation (1)

- Estimated probability of survival (number of forwarding paths surviving beyond time t_i) drops quickly

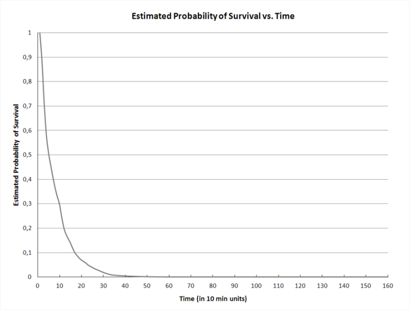

- Perturbations affecting directly (or indirectly) forwarding paths are common events leading rapidly (within 10 hour) to complete connectivity unavailability if no connectivity restoration action
- The probability of survival for the routing paths follows the same type of curve than the one observed for the forwarding paths but with a time shift

Case study: Kaplan Meier Survival Probability Estimation (2)

- First 10 hours: estimated survival probability slowly decreases to reach about 90%
- During this period, estimation for forwarding paths exponentially decreases (less than 1%): second order effect on survival probability estimation for routing paths
- Major routing perturbation events at 100 and 700 (X-axis) seriously affecting survival probability of routing paths whereas (time interval between events, probability decreases rather slowly (< 0.001))

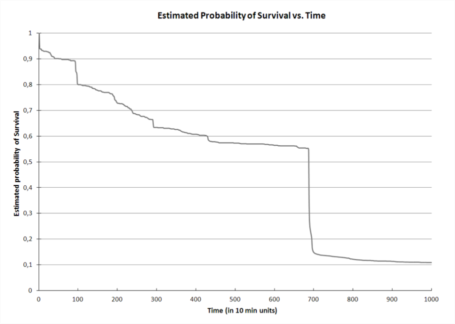

Case Study: Mean Cumulative Function (MCF)

- MCF plot indicates that both forwarding and routing paths experience instabilities at constant recurrence rate
- Rate experienced by forwarding paths is about 10x higher than the rate experienced by routing paths: main source of perturbation affecting reliability of Internet connectivity would be caused by the forwarding plane

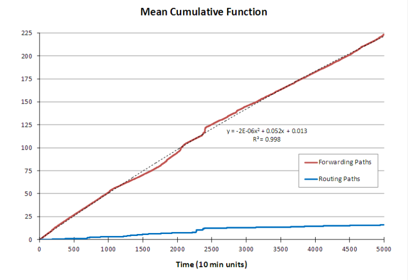

- Number of Internet active paths 512k (July 2014): representative sample size of total population to obtain confidence level of 95% and confidence interval of ± 3 (approx. 16k pairs to reach a confidence level of 99% with a confidence interval of ± 1)
  $\Rightarrow$ Main limit to generalization comes from nonrandom selection of location where datasets have been obtained

Outline

Conclusion

- Application of nonparametric statistical methods (Kaplan-Meier survival probability estimator and mean cumulative function) to characterize dynamic properties of Internet routing paths and corresponding forwarding path(s) directly related to reliability of Internet connectivity (a.k.a. reachability in computer networking)
    1. Simple causality effects between forwarding and routing paths unavailability as identified about 10 years ago not verified anymore: our analysis determines that main source of connectivity perturbations caused by forwarding plane instabilities (which experiences 10x higher recurrence rate compared to routing plane)
    2. Corroborates assumption that dynamic properties of routing system mainly driven by its adaptation to forwarding system: causality effect does not find anymore a simple explanation as forwarding paths become dominant source of instability affecting reliability of Internet connectivity
    3. Reproducing nonparametric statistical procedures on similar datasets obtained from randomly selected locations would enable to further generalize outcomes of this study
- Future work: extend method to characterize intra-AS perturbations