QUICKEST DETECTION OF SUDDEN TRAFFIC CHANGES IN NETWORKS

Ravi R. Mazumdar, Catherine Rosenberg and Edward Coyle

School of ECE

Purdue University

Joint work with P. Dube (INRIA, France)

CERIAS Symposium, Purdue, April 27, 2001

Overview

- Motivation
- Issues
- Framework
- Model
- Main results
- Discussion of results
- Simulation results
- Conclusion and future issues

Motivation

The objective of this research is to see whether signatures for network anomalies can be found in traffic measurements. Our goal is to be able develop network monitoring algorithms based on local traffic measurements in order to take preventive measures as soon as possible.

Problem

Traffic anomalies arise due to many causes:

1. Concerted attacks on a particular location or even a subnet: security

2. Failure of network elements: routers etc. which can involve rerouting

3. Occurrence of a "hotspot".

4. Presence of large users

Effects of sudden changes

Network congestion: Affects latency, causes throughput to be drastically reduced. This can have a multiplicative effect making some parts of the network inaccessible and routers to be overloaded.

This is particularly important because TCP works based on packet loss and thus the rate is reduced.

Issues

1. Detect as soon as possible to take preventive action
2. Network anomaly can often be only perceived through secondary effects (local information)
3. Need to differentiate between anomaly and "normal" statistical variation: Reduce false alarms
4. Devise algorithms which do not need global network information and can be easily distributed

Illustrative example

Consider traffic passing through a router which is associated with a particular part of the network. Normal traffic on the average is λ_1 bits per sec.

At some random time the rate changes to λ_2 per sec.

How can we detect this based on measurements?

Naive approach: Take empirical averages over certain intervals. How long should they be?

Can we guarantee that we will not exceed a certain probability of false alarm? This is to differentiate between normal fluctuations and a genuine change.

The key is that we must have a statistical model for the traffic otherwise we cannot address the false alarm issue

The naive approach needs too long a window to achieve a given false alarm.

However there is a very nice framework available which can give algorithms which can detect changes in the minimal possible time (on average)

The theory is called Optimal Stopping Theory

Problem formulation

Consider a random process $\{X_k\}$. Suppose that before a random time θ the probability distribution of X_k is given by $P_{1,k}$ and after θ the distribution of X_k is given by $P_{2,k}$ i.e.:

$$P_k = P_{1,k}\mathbf{1}_{(k<\theta)} + P_{2,k}\mathbf{1}_{(k\geq\theta)}$$

We observe a history $\mathcal{F}_k$. This history gives the desired probabilistic information about X_k.

Problem: based on observing $\mathcal{F}_k$, determine whether θ has occurred or not with the minimal possible delay subject to the probability of false alarm being $\leq \alpha$, for $\alpha << 1$.

Problem formulation (contd)

Stated in mathematical terms the problem is:

Find a decision rule τ^* (called an optimal stopping rule) such that:

$$\begin{aligned} \tau^* \quad &= \quad argmin E[(\tau - \theta)^+] \\ subject\ to \quad & P(\tau < \theta) \leq \alpha \end{aligned}$$

The term $E[(\tau - \theta)^+]$ is called the detection delay.

$P(\tau < \theta)$ is the probability of false alarm

Main result

Let $\pi_n = P(\theta \leq n | \mathcal{F}_n)$ be the aposteriori probability of the random time θ occurring before n.

Theorem:

If π_n is a transitive statistic w.r.t. $\mathcal{F}_n$ i.e. $\pi_n = f_n(\pi_{n-1}, T_n)$ where $\mathcal{F}_n = \sigma\{T_m, m \leq n\}$ then there exists a constant $A(\alpha)$ such that the stopping time:

$$\tau^* = \inf\{n : \pi_n \geq A(\alpha)\}$$

is optimal in the class of stopping times τ s.t. $P(\tau < \theta) \leq \alpha)$

It can be shown that $A(\alpha) \approx 1 - \alpha$.

Application to traffic models

A good traffic model is the following: Consider N heterogeneous, independent sources,each of which is on-off. When the ith source is ON it transmits at rate $h_i(n)$,

The N sources enter the buffer of a router. Let at any discrete time n, X_n be the random variable denoting the aggregate number of packets entering the node.
At time n, let S_n^i be the state of the source. $S_n^i = 1(0)$ if the source is on(off) respectively. Then we can write

$$X_n = \sum_{i=1}^{N} h_i(n) S_n^i$$

and

$$h_i(n) = h_1(n)\mathbf{1}_{(n<\theta)} + h_2(n)\mathbf{1}_{(n\geq\theta)}$$

If X_n is a so-called Markov modulated fluid and an apriori geometric distribution with parameter p is given for θ, then it can be shown that:

$$\pi_{n+1}^{\pi,x} = \frac{\pi_n^{\pi,x} P^1(X_{n+1}|X_n) + (1-\pi_n^{\pi,x}) p P^0(X_{n+1}|X_n)}{\pi_n^{\pi,x} P^1(X_{n+1}|X_n) + (1-\pi_n^{\pi}) p P^0(X_{n+1}|X_n) + (1-\pi_n^{\pi,x})(1-p) P^0(X_{n+1}|X_n)}$$

or π_n is a transitive statistic with respect to $\mathcal{F}_n = \sigma\{X_m, m \leq n\}$

What this result says is that it is possible to detect changes optimally if we measure total rates

Concluding remarks and future work

- A framework for developing detection rules has been given
- Further work will involve sensitivity to priors (this is not an issue if algorithm is re started after a window of length N). What is the appropriate size of N?
- What if the new parameters are unknown: approach via robust statistics will be studied
- Other information flows will be studied including delay information
- Decentralized detection: multiple distributed sensors.

These algorithms will always beat algorithms based on empirical means even though we can get priors wrong.