Inferring Groups of Correlated Failures*

Jean Lepropre† and Guy Leduc
Research Unit in Networking
University of Liège
{lepropre, leduc}@run.montefiore.ulg.ac.be

ABSTRACT

We compare and evaluate different methods to infer groups of correlated failures. These methods try to group failure events occurring nearly simultaneously in clusters. Indeed if several failures occur nearly at the same moment in a network, it is possible that these failures have the same root cause. The input data of our algorithms are IP failure notifications that can be provided by several sources. We consider two sources: IS-IS Link State Packets (LSPs) and Syslog messages. Our first results on the Abilene and GÉANT networks show that the inference methods behave differently and that using IS-IS LSPs provides more accurate results than using Syslog messages.

Categories and Subject Descriptors: C.2.3 [Network Operations]: Network monitoring

Keywords: Inference, failures, SRLG, clustering

1. INTRODUCTION

The IP topology of a network is often very different from its physical topology, in the sense that a logical IP link may be composed of several physical segments interconnected by data link and/or physical devices.

Consider then the topology of figure 1. In this figure, Lx' is a physical optical link and Lx is a logical IP link using Lx'. If the optical equipment E fails, $L2'$ and $L4'$ will also fail. Thus $L2$ and $L4$ will not work any more though they seem to be disjoint in the IP topology. We say that $L2$ and $L4$ belong to the same Shared Risk Link Group (SRLG), they share a common risk.

Some IP operators, like GÉANT, only see the logical IP (Lx) links and have no information about the physical topology. So they do not know their SRLGs. However, it is very interesting to have this information. For example, when establishing backup Label Switched Paths (LSPs) in Multi-Protocol Label Switching (MPLS) networks, it is desirable that no link of the backup LSP belongs to the same SRLG as links of the corresponding primary LSP.

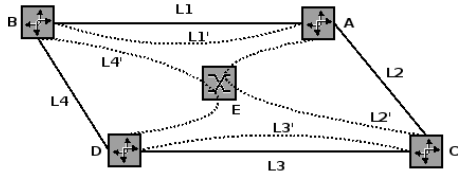

Figure 1: IP topology vs physical topology

The key point of this study is that, in figure 1, if E fails, $L2$ and $L4$ will fail nearly at the same time. Therefore we will analyse failure events of the network and observe links that fail nearly simultaneously. If we put failure events on a time line, it amounts to group events close to each other in clusters. We will say that the clustering methods we are going to develop try to infer groups of correlated failures.

It is important to note that the failure events will not occur exactly at the same time. Indeed several timers tied to different protocols are involved in the detection of a failure. Moreover if the data source is centralised, we can have a delay between the failure event and the arrival of its notification at the server. If the data source is decentralised, the clocks of the measurement devices may not be synchronised. So we have to consider a time window during which the events will be considered correlated.

2. DATA SOURCES

The input data of our algorithms are IP failure notifications. These can be provided by several sources. We consider two of them: IS-IS LSPs and Syslog messages. In the LSPs, each IS-IS router sends other routers in the domain the list of its neighbouring routers. The information is sent periodically or triggered by particular events such as "link up" or "link down". Thus, if we save in a database the neighbours list of each router, we can detect a failure (or the failure recovery) when there is a difference between the content of an LSP and the content of the database. We have two IS-IS traces: one from Abilene, recorded at Atlanta and covering the year 2005, and one from GÉANT, recorded at Geneva and going from 5th February 2004 to 30th March 2005 with some gaps.

Syslog messages are produced by different parts of a router and are sent to the Syslog process of the router. Depending on the configuration of this process, the messages are saved

*This work has been partially supported by the Walloon Region (TOTEM project) and by the European IST-FET ANA project. We thank DANTE and Chris Small for the information about GÉANT and Abilene respectively. We are also grateful to Simon Balon and Pierre François for their comments.

†Research Fellow of the Belgian National Fund for the Scientific Research (F.N.R.S)

CoNEXT'06, December 4–7, 2006, Lisbon, Portugal.

in files, displayed in console or sent to a remote Syslog process. In the last case, User Datagram Protocol (UDP) is used. We have one set of Syslog files from Abilene and covering the year 2005. In Abilene, Syslog messages are sent to a remote Syslog process and so they can be lost.

3. METHODS

We developed and compared three methods. The first is called "the fixed interval". Let $(P_1, P_2, \ldots, P_n)$ be the sequence of chronologically sorted failure events, where n is the total number of failures. The method considers P_i (initially $i = 1$) and groups in a cluster the failures $P_i, P_{i+1}, \ldots, P_j$ such that $t(P_j) \leq t(P_i) + interval$ and $t(P_{j+1}) > t(P_i) + interval$ (if $j + 1 \leq n$), where $1 \leq i \leq j \leq n$, $t(P_i)$ is the time associated with P_i and $interval$ is the size of the interval (this is a parameter of the method). After that, i takes the value $j + 1$ and the method repeats the same process until there is no more failure to consider.

The other two methods come from the field of automatic learning. One is based on the hierarchical agglomerative clustering method. It starts by creating one cluster per failure and then merges the two closest clusters recursively until the distance between the two closest clusters is greater than a given threshold value. The method accepts two parameters: the distance function and the threshold value.

The last method is based on regression trees. A regression tree aims at automatically design "if-then" rules to predict a numeric value. The goal of the building process of regression trees is to define a partition of the input space into regions where the output variable is constant or has limited variance. Thus, if we use as input and output attributes the time of failure events, the building procedure of regression trees will divide the time line into regions where events that can be considered to have occurred at the same moment are grouped into the same cluster. The implementation of regression trees we used is the one of Pepito[1]. In this implementation, there are two parameters that specify when the tree growing has to stop.

Hierarchical agglomerative clustering and regression trees are described in more details in [5]. The three methods have been implemented into the TOTEM toolbox [1].

4. RESULTS

First, we compare results obtained with Syslog messages and IS-IS LSPs. To this end we use the fixed interval method on data provided by Abilene. If the interval is smaller than one minute (it is reasonable to think that beyond one minute the events are not correlated any more), the fixed interval produces twice less clusters with Syslog than with IS-IS. For example, if the interval is 1s, 25 of the 41 clusters produced with IS-IS are missing in results produced with Syslog. None of these missing clusters are false positives because the corresponding failures involve all the links of a node (so we can reasonably think that they are node failures and that the failures have the same root cause). These clusters are missing because the failure events are missing in the Syslog traces, even though the Abilene NOC tickets[2] indicate that the failures actually occurred. Thus, in this case, it is better to use our methods on IS-IS LSPs than on Syslog messages.

Then we compared the three methods on data provided by GÉANT (IS-IS LSPs only). The fixed interval proved to be difficult to configure as the method is very unstable (the results change noticeably even if the interval changes a little bit). The hierarchical agglomerative clustering method produces more or less the same results as the fixed interval and is as difficult to configure as the fixed interval (even more as there are two parameters). Finally, the regression trees are more stable as we identified a good value for each parameter on a first data set and use the same values on the remaining data. However, the method tends to produce more false positives and some interesting clusters are missing (while these were found by preceding methods). A cluster is considered a false positive notably when one failure has already recovered while another one has not occurred yet in the same cluster. These false positives can be easily filtered.

The results suggest to use regression trees to find SRLGs of a network. Indeed, they produce quite good results (if filtered) and are quite simple to configure.

The main SRLGs identified in GÉANT are GR-UK and IT-IL (found 82 times), GR-DE2 and IT-IL (5 times), LU-FR, FR-BE and UK-NL (5 times) and FR-BE and UK-NL (3 times)[3]. The last SRLG was confirmed by DANTE, the company operating GÉANT.

5. RELATED WORK

In [3, 2], the authors also use a technique similar to the fixed interval to do event clustering. However, [3] is focused on characterisation of failures and [2] on localisation of failures. In this paper, we focused on inference of groups of correlated failures and we proposed two original methods. Moreover, we compared two different data sources. In [4], the authors are interested in SRLGs auto-discovery, a subject close to ours. However, they use location-based methods which require location information associated with active components (e.g., optical amplifiers). Here we assumed that the operator has no access to this kind of information.

6. CONCLUSION AND FUTURE WORK

We presented and compared three methods to infer groups of correlated failures. We showed that regression trees are rather stable and accurate. We also compared results obtained with two data sources and showed that results are better with IS-IS LSPs than with Syslog messages. Future work includes improving regression trees, trying other clustering methods and finding some characteristics about clusters.

7. REFERENCES

[1] S. Balon, J. Lepropre, O. Delcourt, F. Skive, and G. Leduc. Traffic Engineering an Operational Network with the TOTEM Toolbox. *IEEE eTransactions on Network and Service Management, under revision*, 2006.

[2] R. R. Kompella, J. Yates, A. Greenberg, and A. C. Snoeren. IP Fault Localization Via Risk Modeling. In *Proc. of NSDI'05*, pages 57–70, Boston, May 2005.

[3] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot. Characterization of Failures in an IP Backbone. In *Proc. of INFOCOM'04*, volume 4, pages 2307–2317, Hong Kong, March 2004.

[4] P. Sebos, J. Yates, D. Rubenstein, and A. Greenberg. Effectiveness of Shared Risk Link Group Auto-Discovery in Optical Networks. In *Proc. of OFC 2002*, pages 493–495, Anaheim, March 2002.

[5] L. Wehenkel. Applied inductive learning, October 2000. `http://www.montefiore.ulg.ac.be/~lwh/AIA/notes-aia.ps.gz`.

[1] `http://www.pepite.be/en/produits/PEPITo`

[2] `https://listserv.indiana.edu/archives/abilene-ops-l.html`

[3] GR stands for Greece, UK for United Kingdom, IT for Italy, IL for Israel, DE2 for 2nd point-of-presence of GÉANT in Germany, LU for Luxembourg, FR for France, BE for Belgium and NL for Netherlands.