

Towards an Optimal Network Survivability Reporting Threshold

Andrew P. Snow and Shweta Agarwal

McClure School of Communication Systems Management

Ohio University

Abstract – This paper considers the appropriateness of a network survivability reporting threshold by investigating the optimality of different thresholds, and threshold types. The threshold in question was set by the Federal Communications Commission (FCC) in the U.S. where a telecommunication outage affecting 30,000 or more users for at least 30 minutes must be reported. Previous work has shown that this threshold masks considerable communications loss. Here, we dub this type threshold an AND threshold, and we investigate the optimality of a variety of different AND thresholds. In addition, we define here a PRODUCT threshold where the product of the outage size and the duration is a constant. The investigation is empirical, using over 19,000 network outage events above and below the FCC reporting threshold. We wish to search for an optimal threshold, which minimizes the number of reports, and maximizes the amount of communication loss exposed by the reports. As a consequence we define here a threshold figure of merit, which is the ratio of the communications loss and the number outages exposed relative to the FCC threshold, and seek to maximize this metric. Findings are that the PRODUCT threshold is superior, and a specific threshold is identified as a candidate for further consideration. The FCC threshold is found to be 33% from the optimal threshold, and to censor significant outliers. Based upon this work, the FCC may want to consider modifying the current reporting threshold to a PRODUCT threshold.

1. Introduction

There are three different perspectives of network dependability dealing with outages, namely reliability, availability, and survivability. Reliability deals with the probability a network or network service will be acceptable and not fail over some specified period of time. Instantaneous availability deals with the probability that the network service or network will be ready at any instant, while average availability is the percentage of time the network service or network is up or down. Although network reliability and availability can be applied to a single user, a group of users, or all users, survivability is a measure of the resiliency of entire network

to absorb failures without unduly affecting the majority of network users. Survivability deals with the frequency of the failures (reliability), the duration of the resulting outage due to the failure (availability), and the number or percentage of subscribers affected by the outage. A non-survivable event is an outage which affected “too many” subscribers for “too long” a time period. Too many for too long is a survivability threshold.

So where does one set the threshold? A threshold set too high might censor important happenings in a network, while a threshold set too low could result in excessive reporting costs and many nuisance reports.

1.1 A Survivability Threshold in the U.S.

In the U.S., the FCC set a reporting threshold for telecommunication outages in the Public Switched Telephone Network. One of the reporting criteria is for switch outages that affect, or potentially affect, at least 30,000 users for at least 30 minutes [1]. This is an “and” threshold in that both the size of the outage and the duration of the outage must each exceed its respective threshold in order to be a reportable event. These outage reports have been filed since 1992, and are used to assess the wellbeing of the network. Over 150 reports a year are made by telecommunications carriers. How many outage events and how much communications loss is censored by this particular threshold? Previous work has shown that there is significant loss below the current FCC threshold [3]. This work addresses the optimality of report thresholds, which takes into consideration the amount of communications loss exposed by the threshold, yet is mindful of the number of reports necessary to expose this loss.

1.2 Two Types of Thresholds

Some obvious problems with the FCC AND threshold is apparent in that the following outages do not exceed the threshold:

- 5,000,000 users affected for 29 minutes
- 29,000 users affected for 24 hours

A simple way to envision the impact of an outage is to convert the outage to lost line-hours (the product of the size of the outage and the duration of the outage). For instance, a (30-minute, 30,000 affected) outage is equivalent to a 15,000 lost line hour (LLH) outage. Likewise, the

outage examples above (which do not exceed the FCC threshold) are 2,416,667, and 696,000 LLH outages, respectively. This leads to a natural question as to whether a PRODUCT threshold is an attractive alternative to the AND threshold. For instance, in a (Duration, Size) xy plot, we can express the AND threshold as:

$T_{AND} = (30, 30K)$, bound by two loci:

$D=30$ ($D \geq 30$), where D is duration in minutes, and

$S=30K$ ($S \geq 30K$), where S is size of the outage (number of users affected by the outage)

Likewise, we can express the PRODUCT threshold as:

$T_{PRODUCT} = 15K$, bound by a hyperbolic locus:

$SD = 15K$, expressed in constant LLH

Each type threshold is “left and bottom censored” in that outages to the left of the threshold, and below the threshold are masked from view. The AND threshold is shown in Figure 1, while the PRODUCT threshold concept is in Figure 2.

1.3 Communications Loss Metric

Some thresholds result in revealing more outage events, while others result in revealing more communication loss. A communication loss metric is needed to assess the impact of each outage. As outage data above and below the threshold is available for local switch outages (discussed in Section 1.4), the communication loss metric must be a reasonable measure of impact for this type outage. LLH has limitations as a loss metric because the number switch lines in actual use depends on the time of day and the day of the week that the switch failed. This shortcoming can be overcome by multiplying LLH times a Time Factor (ranging from 0.1 to 1.0) which accounts for percentage of the lines that would have otherwise been in use, had it not been for the switch outage. The prime of the business day (Monday through Friday) is 8 am to 5pm, and the Time Factor is 1.0, according to Committee T1, an ANSI accredited standards body [2]. Here, we apply a communication loss metric created by multiplying the appropriate Time Factor times the LLH , resulting in an equivalent prime lost line hours ($PLLH$) for each switch outage. Total communication loss is the sum of the $PLLH$ for all the switch outages [3]. In this paper,

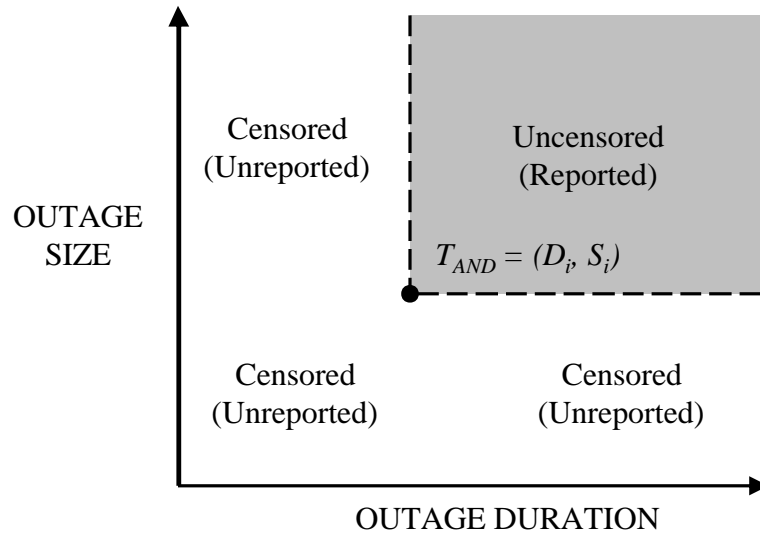


Figure 1. AND Threshold

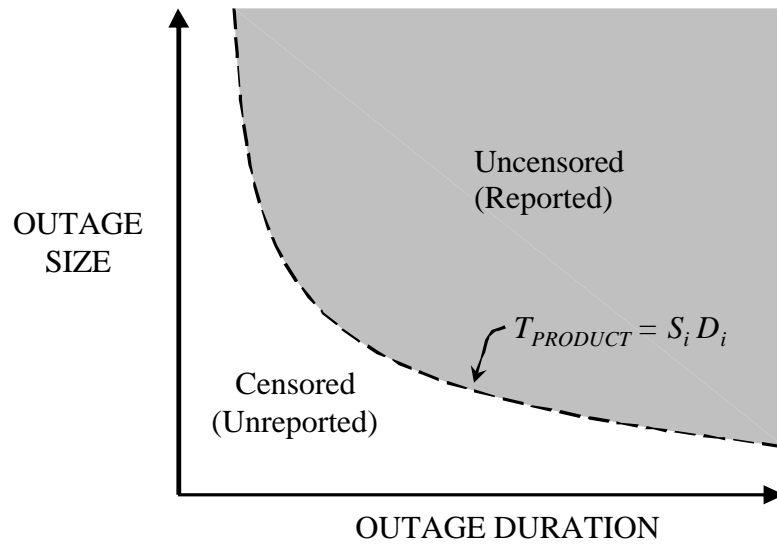


Figure 2. PRODUCT Threshold

we use the *PLLH* as our communications loss metric rather than the Committee T1 outage index because previous work shows use of this un-validated metric in threshold examinations is not defensible [3].

1.4 Local Telecommunication Switch Outage Data

Because of the criticality of local telecommunication switches to customers, the FCC also requires, in a separate reporting venue, reports on any local switch outage of 2-minutes or more, irrespective of the switch size (lines)¹. This offers an opportunity to examine the large-scale FCC threshold of 30,000 customers affected and 30 minutes, above and below the threshold – at least for switch outages. The hope is that switch outages may be a proxy for all type outages, or at least offer insights into the efficacy of the FCC threshold.

2. Research Questions

The objective is to search for an efficient threshold, so we want a “local optimization” which is “near” the current number of events exposed by the FCC threshold, but which maximizes the exposed communication loss. Three questions are investigated:

- R1: Is the current FCC AND threshold near optimal?
- R2: Are there better AND thresholds?
- R3: Are PRODUCT thresholds better than AND thresholds?

¹Automated Reporting Management Information System (ARMIS) filings by carriers include financial, infrastructure, and quality of service information. See <http://www.fcc.gov/wcb/armis/orders.html>. ARMIS was initiated by the FCC in the late 1980s in order to collect financial and operational data from the largest incumbent local exchange carriers. In 1991 additional ARMIS reports were required to collect service quality and network infrastructure information from local exchange carriers subject to price cap regulations. The data from ARMIS used in this study was obtained from <http://www.fcc.gov/wcb/armis/>

3. Methodology

Over 19,000² local switch outage events characterized by lines, duration, date/time of start were obtained from the FCC. These outages are at least 2-minutes in duration and occurred over a 10-year period (1993-2002). Prime Lost Line Hours for each outage event was then calculated. The number of events and *PLLH* was then aggregated above and below the particular candidate threshold. Similar sub-optimization strategies were adopted for each threshold type, which are described in 3.1 and 3.2. After a sub-optimization for each threshold type, results are compared and an local optimal solution identified.

3.1 AND Threshold Sub-optimization

The ranges of lines and duration were examined, and then all possible AND thresholds were created starting from (10,000 lines, 10-minutes) to (150,000 lines, 12,570-minutes)³, incrementing lines by 10,000 and duration by 10-minutes. The reason for choosing these intervals is a practical one as those reporting would find such thresholds as 55 minutes and 17,000 lines awkward. A total of 18,885 different AND thresholds were formed and the number of events and *PLLH* for each threshold was calculated. The thresholds having the same impact and the same number of events were disregarded except for one, leaving 324 thresholds for further consideration. Next, we searched for suboptimal AND thresholds which maximized communication loss and minimized number of events, subject to the following constraints:

- Noted that above the FCC threshold there are 231 events and a total communication loss of 24.2M *PLLH*
- Considered only those thresholds which have outage duration less than 60 minutes
- Took into consideration only those thresholds whose uncensored loss is greater than or equal to FCC threshold loss (24.2M *PLLH*)

The duration constraint of 60 minutes is chosen so that long duration outages are unmasked by the threshold. The additional constraint of exposing at least as much communication loss as the current FCC threshold is added so as not to degrade the value of the current threshold in exposing communication loss. This resulted in 14 local suboptimal candidate thresholds for

² 19,101 to be exact

³ These values represent the largest switch experiencing an outage and longest duration switch outage in the dataset.

further consideration. A plot of these 14 thresholds is shown in Figure 3. The communication loss and event range for these thresholds are:

Communication Loss [24.7M, 34.7M]

Events [253, 1404]




Figure 3. Suboptimal AND Thresholds

3.2 Product Threshold Suboptimization

Different PRODUCT thresholds are formed by incrementing the Size-Duration (Line-Hours) product by 5000 starting at 10,000 Line-Hrs. These increments were chosen for the same practical reporting issue identified earlier. The suboptimal PRODUCT thresholds were searched which maximized communication loss and minimized number of events. This, sub-optimization was subjected to the following constraints:

- Took into consideration only those thresholds whose uncensored loss is greater than or equal to FCC threshold loss (24.2M PLLH)
- Investigated product thresholds less than or equal to 60,000 Lines-Hrs

The second constraint was added because it became apparent that too many significant outage events would be excluded if a product threshold were higher. Using these results, we identified what initially appeared to be a local suboptimal threshold. The results from the investigation of the initial product thresholds are shown in Table 1.

Table 1. Initial PRODUCT Threshold Examination

Product Threshold	Events	Communication Loss (M)
60,000	216	30.0
55,000	233	30.5
50,000	264	31.0
45,000	296	31.7
40,000	324	32.1
35,000	377	32.9
30,000	441	33.8
25,000	536	34.8
20,000	658	35.9
15,000	855	37.2
10,000	1220	38.9

4. Optimal Threshold Investigation

Our objective is to search for an optimal threshold which minimizes the number of reports, and maximizes the amount of communication loss exposed by the reports. As a consequence we define here a threshold figure of merit (*FOM*), and seek to maximize this metric. To create this metric, we compare the ratio of the communication loss and the number of events exposed by the threshold, relative to the communication loss and events exposed by the FCC threshold. *FOM* is given by:

$$FOM_i = \frac{Loss_i / Events_i}{Loss_{FCC} / Events_{FCC}}$$

The *FOM* maximizes loss exposed by a threshold and minimizes the number of events that need to be reported. This definition of *FOM* provides a quantitative measure of how much more efficient the threshold is than the FCC threshold. If the *FOM* of a threshold is 1.0, it is equivalent to the FCC threshold. For the AND threshold, we examined Figure 5 to identify the top three thresholds which exposed the most loss and the least number of events. We identified the

following three AND thresholds, in this format (Min, K-Lines): (60, 20), (50, 20) and (20, 30). For the PRODUCT threshold, we use a similar constraint as used for AND. We observe that the best three thresholds, in the format (K-Line-Hrs) are: 50, 55, and 60.

The AND and PRODUCT threshold comparisons are in Table 2. Here, we also show the figure of merit for each threshold, keeping in mind that the FCC threshold exposed 24.2 M in loss, with 231 events.

Table 2. Threshold Comparisons

Threshold (Min, K) or K-Line-Hours	Threshold Type	Events (#)	Loss (PLLH-M)	Figure of Merit
(50, 20)	AND	206	23.2	0.91
(60, 20)	AND	231	26.6	1.04
(20, 30)	AND	194	26.8	0.72
50	Product	264	31.0	1.12
55	Product	233	30.5	1.25
60	Product	216	30.0	1.33

4.1 AND Threshold Results

The (60, 20) threshold has the best figure of merit of the AND thresholds. This threshold is compared to the FCC threshold in Figure 4. The (60, 20) threshold reveals 3.1M more communication loss than the FCC threshold. The (60, 20) threshold reveals significant communication loss for events of longer duration and smaller number of lines (Region II of Figure 4).

4.2 PRODUCT Threshold Results

The 60K-Line-Hrs threshold has the best figure of merit of the PRODUCT thresholds. This threshold is compared to the FCC threshold in Figure 5. This PRODUCT threshold reveals 5.8M

more communication loss than the FCC threshold. The PRODUCT threshold reveals significant communication loss for events of longer duration and smaller number of lines (Region I) at the expense of Region III, in Figure 5.

4.3 Threshold Comparison By Type

Although the 60K Line-Hrs PRODUCT threshold has the best figure of merit, it is instructive to compare it to the best AND threshold, seen in Figure 6. Here we see that the PRODUCT threshold reveals 3.4M more communication loss than the best AND threshold found.

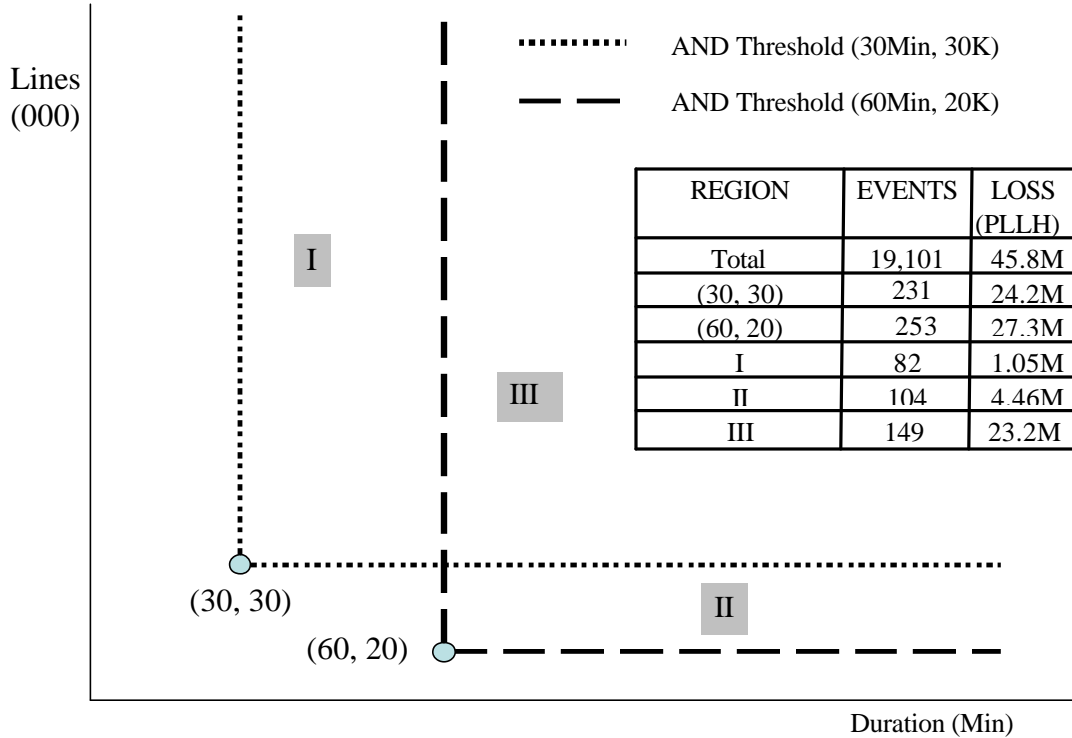


Figure 4. Best AND Threshold Compared to FCC Threshold

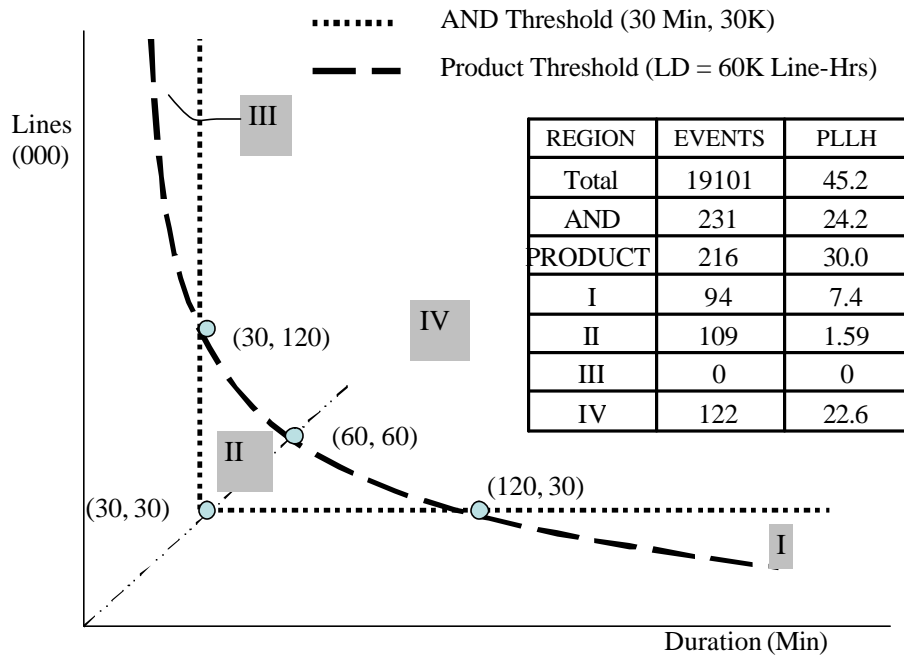


Figure 5. Best PRODUCT Threshold Compared to FCC Threshold

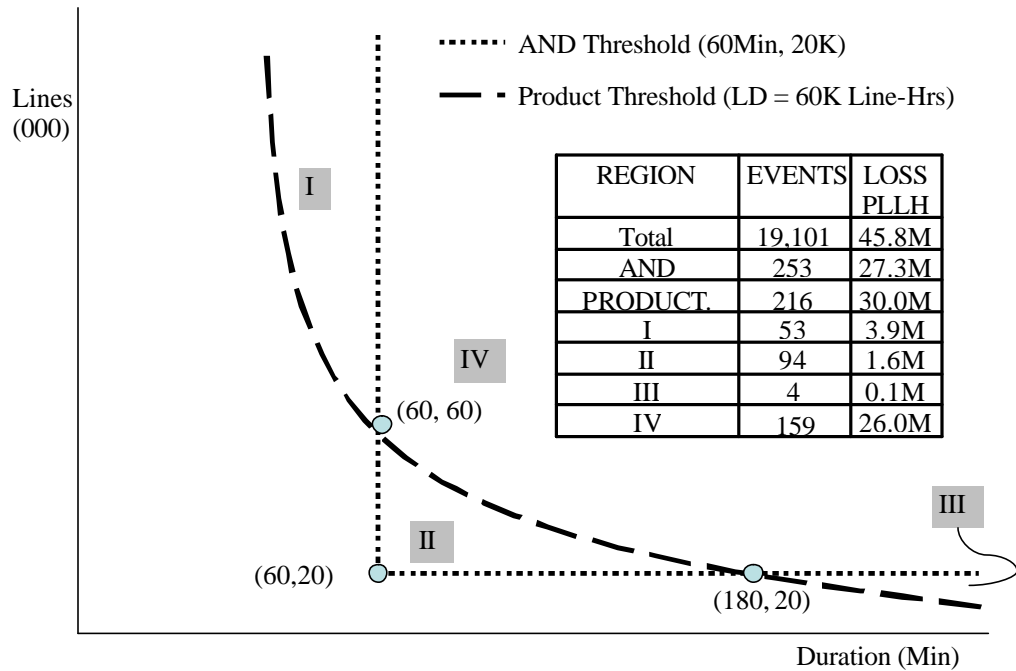


Figure 6. Best AND & PRODUCT Threshold Comparison

5. Conclusions

Conclusions regarding the research questions are:

- R1: Is the current AND threshold near optimal? On a figure of merit basis, compared to the optimal PRODUCT threshold, the FCC threshold is about 33% less optimal.
- R2: Are there better AND thresholds? Although we identified (60,20) AND threshold to be more optimal than the FCC threshold, it is very susceptible to missing very large, but short outages.
- R3: Are PRODUCT thresholds better than AND thresholds? Overall, the PRODUCT threshold is not as susceptible to missing significant outliers, and has higher figures of merit than AND thresholds.

The 60K-Line-Hrs PRODUCT threshold is the most optimal threshold found. In addition, the PRODUCT threshold is superior to the AND threshold in that important outliers (small-duration/large-lines and long-duration/small-lines) are less apt to be missed by such a threshold. Also, the 60K-Line-Hrs PRODUCT threshold requires slightly less reports than the current FCC threshold.

We have some limitations to our results:

1. Switch outages are not totally representative of all outage types. For instance, a fiber optic cut can result in millions affected, while the largest local switch outage is 150,000 lines. However, consideration of switch outages is one of convenience because we have switch outages below the FCC threshold in which to do this analysis.
2. While determining figure of merit we have given equal importance to both, the exposed communications loss and number of reports (events).

Although switch outages may not be totally characteristic of all outages, results indicate that the FCC may want to reconsider their current threshold. Large outages like the fiber optic cut affecting 1 million, if recovered in 20 minutes through use of an alternate fiber transmission path, would show up in Region III of Figure 5 (and missed by the FCC threshold). The equal weight to communication loss and number of events deserves additional comment, and points to further research. Not only should cost of reporting be considered, but also the cost benefit to society by revealed communication loss. This is an economic optimization and is the next step

in this research. However, this work has moved us towards a better understanding of network survivability threshold optimization.

Optimal threshold reporting is an important issue if the goal is to assess the state of a national or international telecommunication infrastructure, be it a telephone system, the Internet, or other convergent network platform. An optimal threshold investigation must consider not only the amount of communication loss exposed, but also the quantity of reports that must be made to expose that loss. This investigation is highly suggestive that the current FCC reporting threshold is not optimal, and regulators may want to reexamine the threshold issue.

6. References

- [1] U.S. 47 CFR § 63.100, Notification of Common Carrier Service Disruptions
- [2] Standards Committee T1 Telecommunications, <http://www.t1.org>
- [3] Snow, A.P., “The Failure of a Regulatory Threshold and a Carrier Standard in Recognizing Significant Communication Loss”, The 31st Research Conference on Communication, Information and Internet Policy, Arlington, Virginia. September 19-21, 2003 .