Sample paper review

Paper: **A Data Mining Analysis of RTID Alarms**

Reviewer: **XXXXX**

RATINGS OF PAPER

[Please rate the following by entering a score between -3 to 3 with 0 being the average based on the following guidelines:

3: *Strong Accept* (As good as any top paper in reputable journals)
2: *Accept* (Comparable to good papers in reputable journals)
1: *Weak Accept* (I vote acceptance, but won't argue for it)
0: *Neutral* (I don't like it, but I won't object if others like it)
-1: *Weak Reject* (I would rather not see this paper accepted)
-2: *Reject* (I would argue to reject this paper)
-3: *Strong Reject* (Definitely detrimental to the journal quality if accepted)]

Originality of the paper: [2]
Technical soundness: [2]
Significance: [2]
Clarity of presentation: [1]
Relevance: [3]

LENGTH (relative to the useful contents of the paper)

About right: [X]
Should be shortened:
Should be extended:

If the paper should be shortened, please indicate an expected number of pages (in its submission format) to be removed:

OVERALL RECOMMENDATION

Accept with no changes:
Accept subject to minor revisions: [X]
Author should prepare a major revision for a second review:
(Not applicable if the paper in question is itself a major revision of a previously reviewed paper.)
Reject:

If the paper is accepted, it should be published as a

regular paper: [X]

short paper:

A SHORT SUMMARY OF THE RATIONALE FOR YOUR RECOMMENDATION *(5 lines expected)*

The paper mainly presents two key contributions:

1. The idea of using volumes of alerts generated by a misuse detection system to achieve anomaly detection by characterizing "normal" alarm patterns.
2. The notion of using alarm contexts and histories to achieve better anomaly detection.

The authors also present some experimental results on real sensor data to show that the methodology is effective and can be improved. The paper, however, seems to be unclear along several lines making the technical and presentation aspects appear weak. With suggested improvements along these lines, the paper can make a good case for the integration of both *misuse* and *anomaly* detection systems for network intrusion detection. This integrated approach, I think, has good practical implications, particularly since misuse detection system sensors can generate very high volumes of alarms, analysis of which becomes a daunting task. Segmentation of customers based on the sensors they use, for assisting service strategies is a nice idea. Comments and suggestion are listed in the next section.

DETAILED COMMENTS FOR AUTHOR(S)

(Please enter your comments for the author(s) giving your reasons for accepting or rejecting the paper by considering the following suggested points: (1) main contributions; (2) positive aspects; and (3) negative aspects. In particular, you are encouraged to substantiate negative comments. If you claim that the work is not original, please give specific references to the earlier allegedly similar work.)

The paper addresses the issue of effectively utilizing alarm sets generated by a misuse detection system for detecting alarm anomalies. The novel contribution is the idea of using contexts and histories of alarms in detecting anomalies.

The use of "*alarm burst*" seems to be a novel idea and will properly fit the need of a real-time system, which the authors intend to pursue. An alarm burst has been said to correspond to a transaction, but since we are dealing with alarm data generated by the RTID sensors, this correspondence appears a little fuzzy and I think it requires some more clarification. For example, how do we view a stream of alarms, which may be independent, as a coherent group as suggested by the notion of a transaction? The importance of the clear notion of an "*alarm burst*" is even more heightened by the fact that the basic notion of the "*normal*" alarm stream is related to frequent itemsets within an alarm burst.

As mentioned (in page 4), "*the alarm bursts are identified by larger than average inter-alarm times at their start times and end points*". A small example with a figure will help better explain this. As the alarm streams are generated by the RTID, the *frequent episodes* technique proposed by Lee *et. al.* (in Wenke Lee, Sal Stolfo, and Kui Mok. "Mining Audit Data to Build Intrusion Detection Models"*In Proceedings of the Fourth International Conference on Knowledge*

Discovery and Data Mining (KDD '98), New York, NY, August 1998) that uses a window of temporal interval to generate sequential patterns of alarms, may be a good technique to experiment on. Yet another technique to characterize a "normal" alarm may be to simply record the temporal information (e.g., occurrence times) of the alarm over an extended period of time.

How the algorithm incorporates the contextual and history information is not clear either. For example the authors write: "*This is one type of anomaly where the context in which alarms occur is of importanc*e", but there is no clear indication of what a context is in the whole paragraph. Is the alarm burst, as is suggested by its use in the algorithm, capturing the contextual information? Furthermore, there is no indication as to how the history information is being used in the anomaly detection algorithm.

At one point the authors state that "*the minimum support and confidence thresholds were chosen empirically; more on this later*". However, the issue is not addressed later. The authors also suddenly introduce the notion of "*interestingness*", but do not indicate how and where it has been used.

In the anomaly detection algorithm, we do not exactly know how the modules `check_rules`(*R*) or `report_infrequent`(*B, M, D*) work. For example, should not it be `check_rules`(*R, s*)?

Descriptions of the related works have been injected within discussions at various places. I suggest putting them together before section 2, in a separate section "Ralated Work". This will give an opportunity to describe the content of the last two paragraphs of section 2 and some other referenced works (such as in the beginning of section 2) at one place. Furthermore, brief introductions of **NetRanger**, which at present is done somewhat late, and other systems used for the experiment can be put together in this section. I believe this will increase readability.

Section 3 seems to be much clearer in its discussion. I think the idea of sensor profiling and extracting distinct segment alarm behavior is novel and highly useful for, as nicely stated also in the paper, servicing strategies at the segment-level rather than the individual sensor level. Some minor points of confusion in this section are:

> In the fourth paragraph, it says "(i) *alarm volumes and rates by priority level and across levels*"; it is not clear what *levels* refer to in the phrase "*across levels*".
>
> The following sentence referring to the *Condorset Criteria* is confusing and needs improvement, *"In other words, it favors clusters that contain similar vectors when, at the same time, vectors assigned to different clusters are dissimilar*". Intuitively, vectors in a cluster will be similar and those assigned to different clusters will be dissimilar – thus the sentence is just stating an apparently obvious fact; if something else is meant, it is not clear at present.
>
> In the second sentence of the first paragraph – "*Can sensor behaviors <u>can</u> be assigned* ..?"has a redundant *can.*

CONFIDENTIAL COMMENTS TO THE EDITORIAL BOARD

(*As many lines as you like.)*

The paper makes a good effort in trying to use data mining techniques over a huge volume of data produced by a set of misuse detection systems. The critical components involved in this technique are:

- Characterizing properly the alarm bursts from the incoming streams of alarms from the RTID sensors.
- Capturing and using context and history information to better detect the anomalous behaviors of alarm streams

The paper clearly indicates that it is reporting preliminary results towards building *anomaly* detection system over the *misuse* detection systems. Efforts to develop empirically justified way of characterizing "*alarm burst*" cannot be seen in the paper. Information gathering, in section 3, is also somewhat ad hoc and seems to have many attributes related to each alarm. Although they may capture a lot of information, dimensionality may be a problem, because the RTID sensors generate huge volumes of alarms. The paper, however, presents a practically viable technique, which I think can be very useful, provided authors can establish an empirically sound sets of "*alarm burst*" characteristics that can be used in different kinds of networks. I think, at its present form, with improvements suggested in the previous section, the paper makes an acceptable case for publication. I think this will generate research interests in improving the two components mentioned above.