Math Matters

Apply It.

The math behind...
Alarm Redundancy Detection

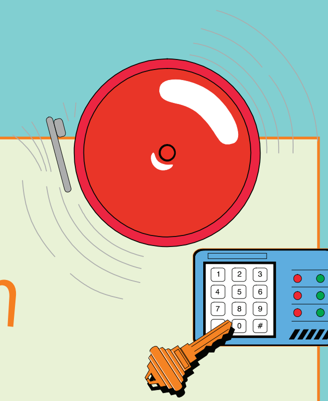

Technical terms used:

Directed weighted graph, adjacency matrix, transition probability matrix

Uses and applications:

It is common for almost every physical component in an industrial plant to be continuously monitored. When a component of the plant is operating outside of normal parameters—for example, the pressure in a valve exceeds a prescribed limit—an alarm will inform plant operators of the issue. However, a single physical anomaly can trigger multiple alarms, which makes it difficult to isolate and handle the original problem.

It is therefore useful to rank alarms by the amount of new information they contribute. Those that do not contribute much are called redundant. The goal is to identify and remove the alarms with the highest redundancy, which would have the effect of reducing the operator load without the risk of disabling safety critical alarms.

How it works:

An alarm system can be visualized as a directed graph, with nodes representing alarms and edges representing correlations between alarms. In particular, if alarm A is followed by alarm B within a fixed period of time (for instance, 30 seconds), the alarm system graph contains an edge from alarm A to alarm B, which indicates that alarm B is potentially redundant. However, even in a small industrial plant there could be thousands of alarms, and the interconnection between any two alarms within this large network becomes unclear.

Industrial plants keep detailed records of alarm activity, and using this data it is possible to find how frequently a particular alarm preceded other alarms within a short period of time. We can restate these frequencies as probabilities, which are used as the edge weights for the alarm system graph.

The adjacency matrix of this weighted graph is a transition probability matrix; weighted sums of powers of this transition matrix, along with an average alarm state vector, are used to find the redundancy of each alarm. By disabling alarms with high redundancy, the strain on the plant operators is reduced, and dangerous accidents are made less likely.

Interesting fact:

In 1979, the Three Mile Island Nuclear Generating Station underwent a partial nuclear meltdown. According to the plant operators' reports, at the beginning of the incident the initial alarm was followed by an alarm cascade (100 alarms within minutes). This made interpreting the alarms significantly harder. One operator, Craig Faust, said, "I would have liked to have thrown away the alarm panel. It wasn't giving us any useful information."

Reference:

T. Butters, S. Guettel, and J. L. Shapiro. Detecting and reducing redundancy in alarm networks. 2015 IEEE International Conference on Automation Science and Engineering (CASE), 1224-1229 (Aug 2015).

M. Rogovin et al. Three Mile Island: a report to the Commissioners and to the public, U.S. Nuclear Regulatory Commission, Special Inquiry Group (1979).

Submitted by: Mario Berljafa, Massimiliano Fasi, Matthew Gwynne*, Goran Malic, Mante Zemaityte, Weijian Zhang, The University of Manchester, UK, Second place, Math Matters, Apply it! contest February 2016