

Using High-Risk Pipe Clusters to Guide Capital Planning

Thomas Ying-Jeh Chen and Craig Michael Daly

Key Takeaways

Water utilities typically replace aggregated clusters of pipelines rather than individual pipes to minimize repeated disruption and mobilization cost.

Risk assessments that report high-risk pipes at the resolution of individual segments fall short of the clustering approach.

A postprocessing method improves risk assessment efficacy by scaling risk calculations from individual pipe segments to a more meaningful and actionable cluster resolution.

Layout imagery by Oleg Elkov/Shutterstock.com. Map sources: Esri, HERE, Garmin, USGS, Intermap, INCREMENT P, NRCan, Esri Japan, METI, Esri China (Hong Kong), Esri Korea, Esri (Thailand), NGCC, (c) OpenStreetMap contributors and the GIS User Community.



As good practice, utilities plan to replace aggregated areas to minimize repeated disruption rather than replace individual pipes. However, current risk assessment practices do not directly support this approach. Specifically, risk assessments are typically performed at the resolution of the individual pipe (using geographic information system [GIS] segmentation), and a selection of the highest-risk pipes is used to guide capital planning efforts. However, this approach has limited usefulness when the high-risk segments are scattered across the distribution network.

High-Risk Clusters Versus High-Risk Pipes

Capital replacement programs are typically time- and labor-intensive because they require service to be stopped, roads to be closed, and expensive excavation (Hastings 2015). Mobilization, a key part of the cost, includes developing staging areas and delivering piping material and technicians to and from the project site. Mobilization costs vary depending on location and project type, but because they can total tens of thousands of dollars, utilities lower their expenses by reducing the number of mobilizations. If the riskiest pipes are scattered widely across a system, meaning that most of the individual segments are not connected to each other, a greater number of mobilizations are required to address these critical assets. This can potentially drive up the cost of any replacement program as more mobilizations require more planning, design, contractor awards, and contractor mobilizations.

A more actionable, cost-effective alternative is to identify clusters in the distribution network that include as many high-risk pipes as possible. A cluster here is defined as a series of connected pipes that groups individual pipes into a larger area, where cluster configurations vary and cross existing political boundaries (e.g., neighborhoods). Targeting a group of clusters instead of individual pipelines is more attractive from a decision-making standpoint as it minimizes setup costs. Furthermore, reducing the number of mobilizations also limits public disruptions.

In this article, we present an algorithmic approach for identifying high-risk clusters. This approach first searches for candidate clusters then selects a subset of them on the basis of risk. This method is applied to three actual distribution networks to select high-risk clusters, and the outputs are contrasted against a direct high-to-low ranking approach without regard to asset location. Finally, we estimated the potential cost savings for planning capital upgrades based on clustered areas instead of the individual pipe.

Simplified Example

In most cases, a water utility has a spatial map of its distribution network in which the system is represented as a collection of line segments. Each line segment represents a pipe segment as drawn by the utility's computer-aided design (CAD) or GIS team. These segments typically extend from street to street or from valve to valve. No industry standard exists for line segmentation when representing infrastructure features in GIS.

To demonstrate the issues caused when reporting pipe risk only at the segment resolution, consider a risk analysis for an actual water distribution network to report the pipe's probability of failure (POF). Using a pipeline GIS shapefile and pipe break records spanning 1998–2019 from the utility, the break locations were combined with the pipeline attribute information to train a machine learning model to forecast each pipeline segment's POF.

Leaving aside more sophisticated quantifications of risk, consider POF as the risk score for each segment. Focusing on a small subset of the distribution network for easier visualization and comparison, risk assessment aims to identify the highest-risk pipes such that the utility can best prioritize and address them. Figure 1 shows a side-by-side comparison of two alternative selections of high-risk pipe for a roughly 94-mile subset of the utility's system. The left image highlights the top 10% of the network by length that has the highest POF. Note that the individual segments can have varying length, and this discretization is due to GIS segmentation. On the right half, the individual pipes are aggregated into larger clusters, and the highlighted portion accounts for the top 10% of the network by length on the basis of POF.

Focusing on the individual pipe, 68 disconnected pipeline fragments are shown in this example, and the average length of each disconnected component is 350 feet. It would not be feasible to plan a capital replacement program solely from this analysis because it would result in a high number of excavations and societal disruptions from road closures and service outages. Even if less intrusive pipeline rehabilitation techniques were used, such a disconnected effort would be unnecessarily expensive and time-consuming to plan and manage. In contrast, these pipelines can be grouped into 10 clusters that contain much of the same high-risk pipes, where each individual cluster contains between 0.8 and 1 mile of total pipe. The reduction in disparate assets represents a significant cost reduction if pipe replacements are planned to follow the cluster selection.

In practice, many factors are important to consider when selecting pipes for replacement, but this example shows how a risk analysis that focuses only on the individual pipeline has limited decision support capability.

Top-Ranked Risk for Pipes (Left) and Clusters (Right)



Figure 1


In these cases, if a high-risk asset is separate from others and relatively short in length, it doesn't warrant replacement because of practical considerations. We argue that the usability of a risk assessment is greatly improved when a selection of high-risk clusters is presented to the decision makers. The emphasis of this article is to present a methodology for identifying these clusters; to our knowledge, no previous work in the pipe risk literature has addressed this issue.

Finding Critical Clusters

A general framework for cluster selection is presented in Figure 2. The cluster aggregation method takes the output from a risk assessment, which reports risk scores for individual segments and finds a set of contiguous pipes that contain as much high-risk pipe as possible. The required inputs are

- the network data (pipe length and configuration), obtainable from GIS;

Cluster Selection Process Workflow



GIS—geographic information system

Figure 2

- the outputs from the risk assessment;
- the size limits for the total length of a cluster (upper and lower bounds); and
- the reporting threshold for cluster selection (e.g., the top 10% riskiest areas).

The risk assessment input here can be either quantitative, semi-quantitative, or qualitative. For simplicity, we will assume that risk here is the POF as reported by a statistical model. This ensures that all risk measures are bound between 0 and 1, and higher values represent riskier assets. It is possible to extend the methodology to include more sophisticated measures of risk, but this is beyond the scope of this article.

As shown in Figure 2, this approach requires two steps: (1) a cluster scan followed by (2) cluster selection, and each is explained in more detail in the following sections.

Identifying Candidate Clusters

To identify candidate clusters, aggregated areas in the distribution network configuration are evaluated on the basis of the risk scores of the individual pipes they contain. Eq 1 shows the length-weighted average risk metric that can be used to evaluate a candidate cluster; this approach normalizes the effects of the nonhomogeneous pipe segment lengths.

$$\text{Cluster Score} = \frac{\sum_{\text{pipes in cluster}} \text{Length of Pipe} \times \text{Risk Score of Pipe}}{\sum_{\text{pipes in cluster}} \text{Length of Pipe}} \quad (1)$$

The size of a cluster is defined as the total pipe mileage contained in the area, and the size of any candidate cluster can be restricted to fall within a certain range. For better usability, the size thresholds should correspond to a utility's capacity to replace and repair pipes within its planning horizon. To provide complete coverage across the pipeline network, a graphical search technique was used to enumerate all clusters falling within the size range and score every pipe using the average risk measure from Eq 1.

Prioritizing Clusters

The second step of the process is to sort the clusters and their average risk scores to determine the best combinations up to the cutoff threshold. The cutoff threshold is the proportion of the total system length for which the utility can feasibly budget for capital upgrades within the next five to 10 years. Utilities select the combination of clusters that maximizes the average risk score from the total selection of pipes while ensuring the clusters that are selected are physically disconnected to maintain their size. If any two selected clusters can touch and join,

they merge into one larger area that can break the condition placed on the size threshold of any area.

For this example, we used a greedy algorithm to identify the best cluster with the highest average risk and iteratively consider the next-best cluster on the basis of average risk that does not overlap with any of the previously selected areas. This process continues until the total mileage of selected pipe meets the cutoff threshold. A greedy approach is simple to implement for this kind of selection problem, and it can be easily adapted to account for other evaluation criteria (e.g., to avoid other underground utility work).

Implementation

Pipe break risk analysis and high-risk cluster selection were performed for three real distribution systems: Long Beach (Calif.), Denver (Colo.), and a midsized utility located in the US mid-Atlantic region (Utility A). Each utility provided a digitized map of their system along with historical records of break locations. Details for each system are reported in Table 1 along with the chronological separation of break records that were used to build and test the machine learning models implemented to estimate pipe break risk. To reiterate, the scope of the risk analysis in these examples only considers the POF and leaves aside consequences.

A machine learning model can combine information on the time and location of historical breaks with the pipeline attributes to interpolate statistical patterns that drive failures. These relationships can be used to forecast which pipe segments are most likely to fail in the future.

The most recent three years' worth of break records were set aside for performance evaluation to validate the accuracy of our machine learning methods by comparing the forecasted riskiest pipelines against those that experienced failure. For example, for the system data set of Utility A, we developed the machine learning model with breaks in 1998–2016 and forecasted the highest-risk pipes in 2017–2019. These forecasted pipes are compared against those pipes with actual breaks during the same window (2017–2019).

High-Risk Pipe Segments Versus Clusters

The machine learning model provides an estimate of the POF for each pipe segment. To better assess the practical effectiveness of this tool to support decision-planning, we focus only on the top 5% of network length with the highest POF. The aggregation method was also used to find high-risk clusters, which were limited to be between 2.5 and 3 miles in total length. Cluster size varies depending on a utility's budget, but these were kept constant across each example used here.

Summary of Test Distribution Systems

Utility Name	Total Size		Break Records for Risk Analysis	Break Records to Test Outputs
	Miles	GIS Segments		
Utility A (mid-Atlantic region)	1,381	36,047	1998–2016	2017–2019
Long Beach (Calif.)	888	42,339	1998–2015	2016–2018
Denver (Colo.)	2,672	150,254	2008–2015	2016–2018


GIS—geographic information system

Table 1

The goal was to find the top 5% by system length as a collection of a clusters that, although possibly scattered throughout the network, has potentially many fewer

individual components. Figure 3 shows the highest-risk pipes (left) and the highest-risk clusters (right) in Utility A’s system. The pipes of interest are marked in red and

Top-Ranked Risk for Pipes (Left) and Clusters (Right) for Utility A



To keep the utility anonymous, the map projection has been rotated and distorted.

Figure 3

Top-Ranked Risk for Pipes (Left) and Clusters (Right) for Long Beach, Calif.




Figure 4

account for 5% of the total system length: 397 separate components when considering only the riskiest pipe (left), and in contrast, just 26 components (right) when these pipes are aggregated into clusters. Similarly, Figure 4 compares the highest-risk pipes and the highest-risk clusters in the Long Beach system: 445 separate components when considering only the individual pipe, and just 17 components when aggregating into clusters. Finally, Figure 5 shows the comparison between the highest-risk pipe versus clusters in the Denver system: 1,034 separate components when considering only the individual pipe, and just 52 components when aggregating into clusters.

Targeting a group of clusters instead of individual pipelines is more attractive from a decision-making standpoint as it minimizes setup costs.

If we relied only on the risk reporting of individual pipes to plan asset renewals, many of the recommended segments might be ignored because of excessive mobilization costs. In contrast, the cluster aggregation presents only a limited number of high-risk clusters, and each of them are sized appropriately (2.5–3 miles) for contractor work.

For asset management, these results better support decision-making; however, a tradeoff may be needed to address the number of breaks in the targeted areas. The machine learning models forecast break locations in the three most recent years to show how many breaks are included in the two different spatial resolutions. Table 2 shows the difference in break capture along with the count of individual components.

In each case, fewer breaks were captured when targeting aggregated clusters instead of the individual pipes. However, the sizable reduction in the number of components in the cluster approach implies a large savings in mobilization costs. Using estimates from a 2017 Water Research Foundation survey (Raucher 2017) on pipe break and repair costs, Table 3 summarizes the tradeoffs in each of the three scenarios had the utilities addressed

Top-Ranked Risk for Pipes (Left) and Clusters (Right) for Denver, Colo.




Figure 5

their planning using our cluster analysis. The extra incurred cost is due to unaddressed pipe breaks, and in another column, we report the cost savings achieved as a result of reduced mobilizations. As previously stated, these cost savings are for illustrative purposes only; a

utility typically would not mobilize crews to replace every high-risk pipe scattered throughout the network. Our emphasis is to show how aggregating risk results into clusters can help utilities better target their mobilizations while still addressing critical pipes.

Performance Comparison Between Cluster and Pipe-Level Resolution

Utility Name	Top-Ranked Pipe Performance	Top-Ranked Cluster Performance	Percentage Difference
Utility A (mid-Atlantic region)	397 components 78 breaks	26 components 65 breaks	95% less 16% less
Long Beach (Calif.)	445 components 58 breaks	17 components 35 breaks	94% less 39% less
Denver (Colo.)	1,034 components 211 breaks	52 components 151 breaks	95% less 28% less

Table 2

Economic Analysis for Aggregated Clusters Versus Pipe Segments

Utility Name	Opportunity Cost From Uncaptured Pipe Breaks US\$ 2020 thousands	Cost Savings From Reduced Mobilization US\$ 2020 thousands	Magnitude Difference
Utility A (mid-Atlantic region)	27–111	2,078–9,238	75–84 times
Long Beach (Calif.)	48–197	2,397–10,657	49–54 times
Denver (Colo.)	128–513	5,499–24,452	43–47 times

Table 3

Table 3 demonstrates the significant potential cost savings when targeting clusters for upgrades. The extra costs incurred as a result of pipe breaks is on the order of hundreds of thousands of dollars, whereas the reduction in mobilization costs is at a magnitude of millions of dollars. In every case, the cost savings are between about 40 and 80 times higher than the opportunity cost from uncaptured pipe breaks. From an asset management standpoint, it is far more effective to incur the cost of extra breaks in exchange for more practical planning of capital works.

It's important to note that we used the POF as a simple measure of risk, but this kind of modeling most often incorporates more sophisticated measures and other practical planning considerations in the search of clusters, such as avoiding paving projects and underground utility works.

Implications for Risk Management

Current risk assessments report risk at the resolution of individual pipeline segments and recommend targeting those deemed most critical, but this approach is improved by considering the spatial location of the assets. This article presents a simple method for grouping risky pipes into risky clusters and demonstrates the cost-saving potential of this approach for multiple systems. By using high-risk pipe clusters to guide capital planning, utility managers can better protect the health of their critical underground systems. 💧

Disclaimer

The opinions and views expressed are those of the authors and do not necessarily reflect those of the partnering utilities.

Acknowledgment

We would like to thank our partnering utilities for allowing us to present and discuss their data in this article: Long Beach Water Department and Denver Water. We would also like to

thank Tavon Ellis, Jamie Martin, and Kavya Murhali for assisting with the clustering implementation.

About the Authors



Thomas Ying-Jeh Chen is a data scientist at Xylem Inc., Columbia, Md.; Thomas.Chen@xyleminc.com.

Craig Michael Daly is a chief engineer at Xylem Inc., Columbia, Md.

<https://doi.org/10.1002/awwa.1669>

References

- Gordian. 2020 (34th ed.). *Heavy Construction Costs With RSMeans Data*. Gordian, Rockland, Mass.
- Hastings NAJ. 2015. *Physical Asset Management: With an Introduction to ISO55000*. Springer, Cham, Switzerland.
- Walski TM. 1985. *Cost of Water Distribution System Infrastructure Rehabilitation, Repair, and Replacement*. US Army Corps of Engineers, Washington.
- Raucher R. 2017. *Managing Infrastructure Risk: The Consequence of Failure for Buried Assets*. Project 4451. Water Research Foundation, Denver.

AWWA Resources

- Is It Real or Just Apparent? A Case Study in Determining Water Loss. Tromble E, Mach ME, Bendinelli J. 2020. *Journal AWWA*. 112:5:56. <https://doi.org/10.1002/awwa.1502>
- Pressure-Transient Monitoring Supports Asset Management. Ginn T, Smither B. 2020. *Opflow*. 46:12:28. <https://doi.org/10.1002/opfl.1473>
- Asset Management Best Practices: Keeping Ahead of the Curve. Fodchuk M, Garaci M. 2019. *Journal AWWA*. 111:8:46. <https://doi.org/10.1002/awwa.1342>

These resources have been supplied by *Journal AWWA* staff. For information on these and other AWWA resources, visit www.awwa.org.