Sensor Placement Guidance for Small Utilities

Stacey Schal[1], L. Sebastian Bryson[2], Lindell Ormsbee[3], Srinivasa Lingireddy[4]

[1]Research Assistant, Department of Civil Engineering, University of Kentucky, 161 Raymond Building, Lexington, KY 40506, email: slscha3@g.uky.edu
[2]Associate Professor, Department of Civil Engineering, University of Kentucky, 161 Raymond Building, Lexington, KY 40506, ph: (859) 257-3247, email: sebastian.bryson@uky.edu
[3]Director, Kentucky Water Resources Research Inst., University of Kentucky, 233 Mining & Minerals Resource Building, Lexington, KY 40506, ph: (859) 257-1299, email: lormsbee@engr.uky.edu
[4]Fluid Hammer Consultancy Services Ltd, Hyderabad, India, email: srini@fluidhammer.com

Abstract

Contamination warning systems (CWS) are strategies to lessen the effects of contamination in water distribution systems by delivering early indication of events. A critical component of CWS, online quality monitoring, involves a network of sensors that assess water quality and alert an operator of contamination. Utilities developing these monitoring systems are faced with the decision of what locations are optimal for deployment of sensors. The TEVA-SPOT software was developed to analyze the vulnerability of systems and aid utilities in designing sensor networks. However, many small utilities do not have the technical or financial resources needed to effectively use TEVA-SPOT. As a result, a sensor placement algorithm was developed and implemented in a commercial network distribution model (i.e. KYPIPE) as a simple tool to aid small utilities in sensor placement. The developed tool was validated using 12 small distribution system models and multiple contamination scenarios for the placement of one and two sensors. Results were compared with data from identical simulations runs in TEVA-SPOT to verify the effectiveness of the proposed algorithm. Because the proposed algorithm uses a simple complete enumeration strategy for sensor placement, the algorithm was able to select the same or superior nodes to those selected by TEVA-SPOT.

Introduction

Water distribution systems are an integral part of society, and the availability of a clean and dependable supply of water influences both the socioeconomic status and health of a populace. In recent years, protecting the nation's critical infrastructure from terrorist attacks has become a priority, and efforts to protect the water infrastructure are included in this goal. Water distribution systems are considered to be vulnerable to intentional, along with accidental, contamination because they have a large spatial distribution and multiple points of access (Hart & Murray, 2010). In an effort to mitigate the risks from contamination of the water supply, contamination warning systems (CWS) have been proposed as a cost-effective and reliable strategy.

Contamination warning systems are proactive strategies to reduce public health impacts and economic loss from contamination in distribution systems by providing an early indication of an intentional or accidental contamination event (Janke et al., 2006). A CWS includes deployment and operation of online sensors, which involves a network of sensors that can assess water quality in the system and alert an operator of a potential contamination event. The challenge involved with developing these water quality monitoring systems is determining which locations are best suited for deployment of sensors. Budget constraints will limit the number of sensors a utility can deploy, and they must be placed in locations that maximize their ability to detect contaminates and provide the greatest protection of human health (McKenna et al., 2006).

To date, there is no applicable federal or state guidance to assist utilities in the deployment of water quality sensors. Technological advancements in sensor placement optimization software may help solve the problem of sensor placement issues for some utilities. The TEVA-SPOT software (Threat Ensemble Vulnerability Assessment Sensor Placement Optimization Tool) has been developed to analyze the vulnerability of drinking water distribution networks and aid utilities in the design of sensor networks (Berry et al., 2010). While TEVA-SPOT uses public domain software (e.g. EPANET) along with a sophisticated optimization algorithm to evaluate the ability of different sensor combinations to detect contamination events, the software can be intimidating to use for medium to small utilities. As a result, application of the software has largely been limited to large utilities or research studies. This paper proposes the use of a fairly simple enumeration method coupled with a widely used commercial network distribution model (i.e. KYPIPE) for applications of sensor placement to small or medium sized utilities. For the purposes of this discussion, KYPIPE was primarily used in order to facilitate access and generation of a dataset of network models from a statewide database. The proposed heuristic can be easily adapted for use with EPANET directly or with any other commercial software.

In order to evaluate the proposed heuristic, the model is applied to 12 different water distribution systems associated with small water utilities in Kentucky (Jolly et al., 2013). The model is executed with a variety of contamination scenarios for all systems, placing a number of sensors reasonable for the budget of a small utility. The results of these applications are then compared to the results obtained from applying TEVA-SPOT to the same 12 systems. In each case, sensor locations are selected by minimizing the time to detect the contaminant.

Current Sensor Placement Software

The Threat Ensemble Vulnerability Assessment Sensor Placement Optimization Tool (TEVA-SPOT) Program was developed as a probabilistic framework for analyzing the vulnerability of drinking water distribution networks (Murray et al., 2004). This collection of software tools to aid utilities in the design of sensor networks was developed by researchers from the Environmental Protection Agency (EPA), Sandia National Laboratories, the University of

Cincinnati, and Argonne National Laboratory (Murray et al., 2010). TEVA-SPOT creates a threat ensemble, consisting of a set of contamination scenarios, and the vulnerability of the network is assessed using the entire threat ensemble (Murray et al., 2004).

TEVA-SPOT contains three main software modules. The first module simulates the set of incidents in the threat ensemble, the second module calculates the potential consequences of each incident, and the third module optimizes for sensor placement. The design basis threat consists of the set of incidents for the sensor network to detect. The consequences are calculated based on one or more of the performance objectives that include the number of people who become ill as a result of exposure, percentage of incidents detected, time to detection, and length of pipe contaminated. When TEVA places sensors, the mean consequence for a given objective is minimized. This is equivalent to assuming that each contamination scenario is equally likely to occur and that each is important when selecting sensor locations. The user is able to specify weights to put more importance on locations with a higher likelihood of being contaminated (Murray et al., 2010). A flow chart of the TEVA-SPOT software is shown in Figure 1.

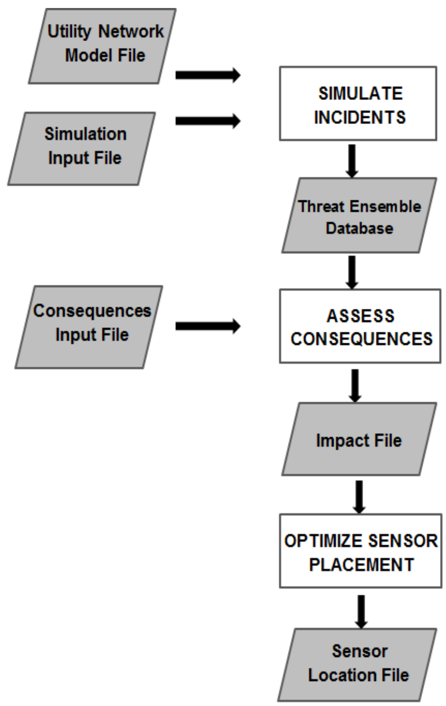

Figure 1. Flowchart of TEVA-SPOT Software (Murray et al., 2010).

TEVA uses simulation and optimization models to select optimal placement of sensors for a CWS by implementing two main steps: modeling and decision-making. The modeling process first involves creating a network model for a hydraulic and water quality analysis. TEVA-SPOT uses EPANET (U.S. EPA, 2000) to perform these analyses. An EPANET INP file is used to describe the physical characteristics of the system, and this file is built using the EPANET user interface. Using models for the purpose of contamination warning systems requires a high degree of accuracy.

The design basis threat describes the type of threat that the utility wants to protect against when designing a CWS. Contamination incidents are described by the specific contaminant, the quantity and duration of injection of the contaminant, and the locations where the contaminant is introduced. The program understands that these conditions cannot be known before an incident occurs, so the modeling process takes this uncertainty into account. For example, probabilities can be assigned to each node in a system to specify the likelihood that contamination would occur at that location. An ensemble of incidents is then simulated, and sensor network designs are chosen based off how they perform for the entire ensemble of incidents (Murray et al., 2008).

TEVA measures performance of sensor network designs based on minimizing certain performance objectives including the number of failed detections, extent of contamination, mass/volume consumed, and time to detection. If a utility has several important priorities for the performance of their sensor design, multiple objectives can be considered by assigning a relative importance, or weight, to each objective. Modeling the utility response time, the time between initial detection of the contaminant and effective warning of the population, is another important aspect of the modeling process (Murray et al., 2010). The user must also input information defining all potential sensor locations in the system. When selecting nodes for potential sensor locations, certain requirements are needed such as accessibility, security, and protection from the environment. A longer list of feasible sensor sites results in a sensor design that is more likely to perform well. The benefits of using sites that need some adaptation to meet requirements may be worth the additional costs, but sensor placement can be restricted to locations preferred by the utility. Characteristics of the sensor behavior are also needed to measure performance of a CWS, so the sensor type, detection limit, and accuracy should be included (Murray et al., 2008).

The second main step in the TEVA sensor placement framework is the decision process. The decision making process uses an optimization method and evaluates sensor placement by analyzing trade-offs and comparing a set of designs to account for modeling and data uncertainties. The goal of this step is to aid utilities in understanding the public health and cost tradeoffs between different sensor placement designs and ultimately help them choose the sensor design that will best meet their needs. This is accomplished by using an incremental approach for applying optimization to generate a set of sensor placement designs that can be compared and contrasted (Murray et al., 2008).

The first sensor placement design is found under ideal conditions with simplifying conditions. The assumptions are then removed one at a time to make the designs more realistic. For example, simplified conditions would assume all nodes in the system as potential sensor locations, instantaneous response time, and perfect sensors. More realistic conditions would assume delayed response time and would force sensors to be placed at utility-owned or public locations (Murray et al., 2006). The performance of the new sensor design is compared with the previous designs and baseline case with no sensors, quantitatively and visually, to understand what has been gained or lost with each assumption. The tradeoffs can be analyzed in terms of the desired performance objective, such as the percent reduction in the number of illnesses with each design (Murray et al., 2010).

TEVA-SPOT provides three optimization method options in order to develop a sensor design: mixed-integer programming (MIP), a greedy randomized adaptive search procedure (GRASP) heuristic, and a Lagrangian relaxation method. The standard formulation used to evaluate impacts is a MIP formulation, which optimizes linear objective functions by maximizing or minimizing the function subject to constraints. The MIP solver minimizes the predicted impact of an ensemble of contamination events using the specified sensor set size. This solver is exact and will guarantee to find the optimal solution (Murray et al., 2010). MIP technology can be used if the problem instances are not very large and if sufficient power is available. However, heuristic methods are commonly used when working with large problem instances because the number of constraints and variables will grow rapidly as problem size increases (Berry et al., 2006).

The GRASP algorithm finds solutions by systematically exploring the space of possible sensor layouts, and it typically produces solutions as effective as results from a MIP in less time. This solver was utilized to collect data in TEVA-SPOT for this study. The GRASP randomly creates a set of starting points using greedy bias to make these reasonable approximations. It then explores ways to move one sensor to a new location that will improve the objective function, making these swaps until a better solution doesn't exist. The only limitation of this method is that it still has a fairly large memory requirement (Murray et al., 2010).

The Lagrangian method removes a set of "difficult" constraints, resulting in a problem that is easier to solve. Penalties are then added to the objective function to satisfy the relaxed constraints. Penalty weights are manipulated and an iterative algorithm drives the solution to feasibility (Murray et al., 2010).

KYPIPE Sensor Placement Tool

KYPIPE was first developed in the 1970s to calculate steady state flows and pressures in a water distribution system (Wood, 2010). The program can complete an analysis for any configuration of pipes including hydraulic components such as pumps, valves, fittings with significant head losses, and storage tanks. The program also has the capabilities to execute an extended period simulation (EPS) that can account for the variation in storage tank levels over time (Wood, 2010). KYPIPE performs hydraulic analyses using the KYPIPE hydraulic engine which is based on a nonlinear solution of the network loop energy equations (Wood, 1981).

The Water Quality Sensor Placement Tool has been developed to work with the existing KYPIPE graphical user interface. The goal is to provide a simple tool to aid utility managers in the optimal placement of sensors in their distribution systems. The simplicity and ease of use of the tool makes it attractive for use in small utilities. The sensor placement tool recommends optimal sensor placement, regardless of how many sensors are implemented, based on minimizing time to detection. The tool considers detection events at nodes throughout the entire system, and recommends optimal sensor placement based on the locations that can detect contamination events the fastest.

The KYPIPE sensor placement routine utilizes four different input files. The first input consists of a normal KYPIPE input file. This file is used to describe the physical parameters of the network and to specify the parameters for the required extended period simulations. The second file is an EPANET INP file that is generated internally within KYPIPE using data from the normal KYPIPE data file. Some adjustments are made to accommodate differences in the way the two programs handle certain components such as pumps (i.e. nodes vs. links). The third file is the travel time matrix, which is generated using hydraulic and water quality calculations. The fourth file is used to prescribe the parameters associated with the sensor placement algorithm. Figure 2 displays the entire procedure executed by the KYPIPE WQ Sensor Placement Tool, followed by further explanation of the process.

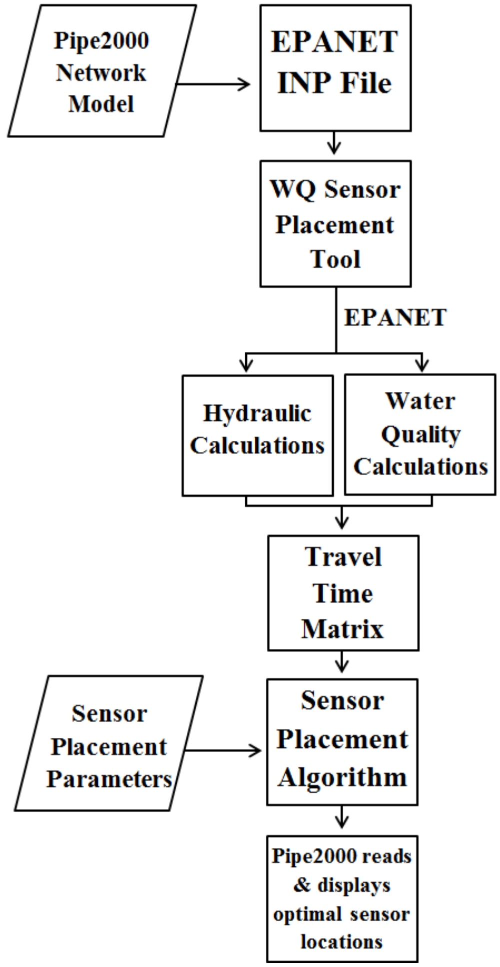

Figure 2. Flowchart of KYPIPE Sensor Placement Methodology.

When performing a sensor placement analysis, KYPIPE utilizes the EPANET engine for both hydraulic and water quality calculations. Therefore, use of the sensor placement tool will first require KYPIPE to export the network model into EPANET format (INP file) that contains all hydraulic and water quality data. KYPIPE then makes calls to the EPANET engine to perform the hydraulic and water quality analyses, and result files are generated to the hard drive. The tool also uses this data to generate the travel time matrix, which is used to perform the optimal sensor placement calculations.

The optimal sensor location information is written to a file on the hard drive, and KYPIPE is able to read the file and display the chosen sensor locations on a graphical representation of the water distribution network. KYPIPE is also able to read the results files and display data from the hydraulic and water quality analyses. Figure 3 shows the steps to be carried out by the user in order to execute the sensor placement tool in KYPIPE.

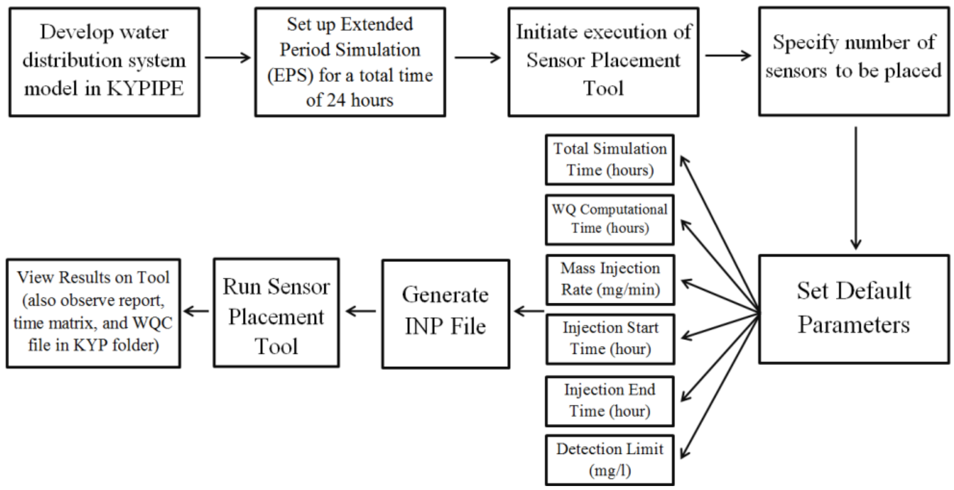

Figure 3. Sensor Placement Tool Flowchart.

For the case where only one sensor is being considered, the algorithm simply iterates through all combinations of injection points and possible sensor locations. This results in a matrix of travel times between the injection locations and possible sensor nodes. The algorithm then simply enumerates through the set of solutions to find the "optimal" sensor location (with the lowest average time to detection). Because the algorithm computes the travel times for all sensor locations, it is also possible to list the average time to detection for each sensor location, thus providing the user with alternative options in the event the "optimal" node turns out to be impractical for other secondary conditions (e.g. physical access, communication or power limitations, etc.).

For situations involving multiple sensors, the algorithm calculates the average travel time for each set of sensor locations. As with the single case, these results are then stored in a matrix of travel times for each combination of multiple sensors. The methodology for determining the average travel time for two sensors is illustrated in Figure 4. The contaminant is "injected" at the first possible injection site, and the travel time for the contaminant to reach each of the sensors in the first sensor combination is determined. The values for T_1 and T_2 represent the travel times from each injection node to sensor 1 and sensor 2, respectively. The travel time assigned to this particular set of sensor locations and injection node will be the minimum of the two travel times, since the contaminant is considered to be detected when it reaches the first sensor. This process is repeated for all possible injection nodes in the system. The average travel time for the particular set of sensor locations is calculated by averaging the minimum travel times from all injection sites.

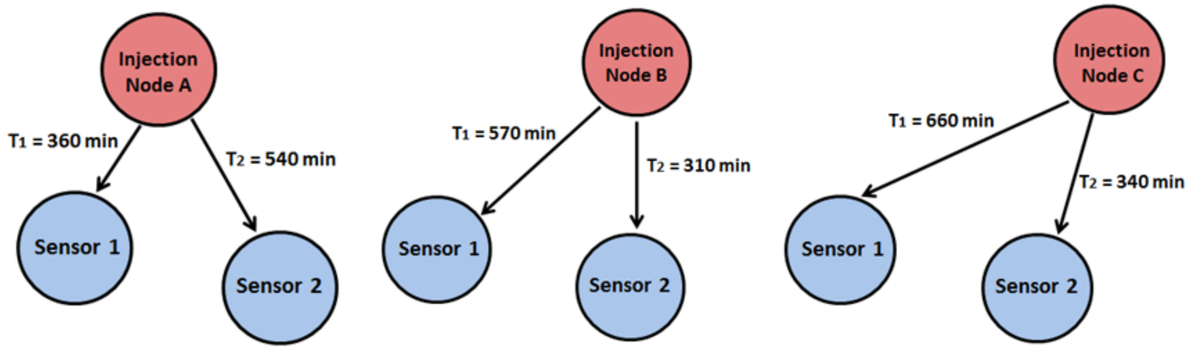

$$Travel\ Time = \frac{\sum Minimum\ Travel\ Times}{\#\ Injection\ Nodes} = \frac{360 + 310 + 340}{3} = 336.67\ minutes$$

Figure 4. Sensor Placement Tool Theory.

This process is then repeated for every set of possible sensor locations, resulting in an average travel time for all combinations of two sensors in the system. The sensor combination with the lowest average travel time will be considered the optimal sensor location. As with the single sensor case, the user may display the travel times for all the sensor location combinations if desired, thereby providing information on inferior solutions as well.

It should be recognized that use of an optimization algorithm based on complete enumeration will result in an exponential increase in the evaluation of combinations of travel times as one considers more than one sensor. However, many small and even medium sized systems may not have the financial resources to place a large number of sensors across the network. Thus the ability to site a smaller number of sensors (e.g. two to three) may be sufficient to provide adequate coverage for such systems. As system become larger, it is likely that they may contain multiple pressure zones. In such cases the proposed algorithm may be then be applied to each pressure zone individually, thereby limiting the computational burden while providing for more sensors. Finally, the ability to obtain and display the complete solution space for the problem may allow the utility to identify sub-optimal sites that may still be able to provide for expanded protection for the utility.

In order to maximize the potential coverage of the sensor locations, the KYPIPE Sensor Placement Tool considers all nodes to be potential sensor locations (i.e. including tanks, pumps, reservoirs, and junctions). However, the algorithm excludes all dead-end nodes as possible sensor locations. The average travel time to dead-end nodes will generally be much higher, skewing the average times to detection. Possible injection sites are considered to be all non-zero demand nodes, excluding dead-end nodes. Dead-end nodes are considered to be consumption nodes, so any contaminant injected at these nodes is assumed to be consumed immediately and the contaminant will not be able to travel further in the system.

The contamination detection limit for each sensor is entered in the default parameters menu for the program (a detection limit of 0.01 mg/l was used in this study). When the concentration of the contaminant reaches the detection limit at the particular sensor node, the contaminant is considered to be detected. The tool requires input for the total simulation time and considers this the maximum travel time. Any travel time past the total simulation time (24 hours in this study) will be considered this maximum time for calculation purposes.

As indicated previously, TEVA-SPOT provides the user with several different design options (i.e. minimize number of failed detections, extend of contamination, mass/volume consumed, or time to detection). In order to minimize the computational burden and maximize use of the algorithm from small to medium sized systems, the KYPIPE Sensor Placement Tool considers minimum time to detect as the sole operational objective.

Water Distribution System Models

In order to demonstrate the utility of the proposed approach to the sensor placement problem, the sensor placement software was evaluated using a database of 12 small water distribution system models (Jolly et al., 2013). While each of the models represent a real distribution system in Kentucky, all models were given a generic name in the form KY #. All identifying information for the actual systems represented by the models was removed, such as names of tanks, to protect the security of the utilities. Model names were grouped by physical configuration type (Von Huben, 2005). The first four models, KY 1 – KY 4, are characterized as grid systems. Models KY 5 – KY 8 are classified as loop systems, while the remaining models, KY 9 – KY 12, are characterized as branch systems. The layout of one system representing each of the three configurations is displayed in Figure 5.

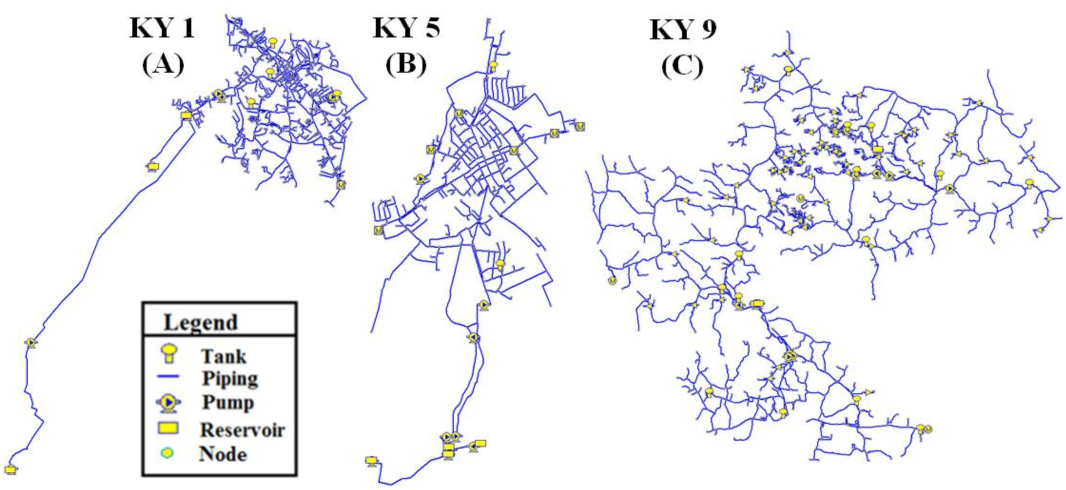

Figure 5. System Models (A) Grid; (B) Loop; (C) Branch.

Verification Studies

Contamination Scenarios

In order to evaluate the performance of the KYIPE Sensor Placement Tool and to compare its performance against TEVA-SPOT, the tool was executed on the 12 hydraulic models for 15 different contamination scenarios. The contamination scenario is determined by both the rate of injection of the contaminant (in mg/min) and the total injection time (in hours). Contamination scenarios were created for three different general scenarios: fixed amount, fixed rate, and fixed time. Each general scenario is comprised of five specific sets of an injection rate with a total injection time. The 15 contamination scenarios performed on each model are displayed in Table 1.

Table 1. Contamination Scenarios.

	Injection Rate (mg/min)	Injection Time (hours)	Total Contaminant Injected (g)
Fixed Amount (Vary Time)	4000	1	240
	2000	2	240
	1000	4	240
	500	8	240
	250	16	240
Fixed Rate (Vary Amount)	1000	1	60
	1000	2	120
	1000	4	240
	1000	8	480
	1000	16	960
Fixed Time (Vary Rate)	600	4	144
	800	4	192
	1000	4	240
	1200	4	288
	1400	4	336

All 15 contamination scenarios were executed on each of the 12 water distribution system models using both TEVA-SPOT and KYPIPE. It was desired to compare the sensor placement results between KYPIPE and TEVA-SPOT for a variety of scenarios. To be able to directly compare results, it was required that all parameters matched between the programs. First, the general network models used in each comparison were identical. The TEVA-SPOT program uses a model input from EPANET. Even though minor differences exist between KYPIPE and EPANET, all major system components and characteristics of these components matched between the two programs. An example of a difference between KYPIPE and EPANET is that KYPIPE allows tanks to be measured as a total volume or fixed diameter, while EPANET only allows a fixed diameter as input for tank size. To make the models as similar as possible, all tanks in both KYPIPE and EPANET were set as fixed diameters.

Parameters used in the sensor placement tool in KYPIPE and TEVA-SPOT were also standardized. The WQ computational time (labeled as the hydraulic timestep in TEVA-SPOT) was set to 60 seconds, and the total simulation time was set to 24 hours. The detection limit for both programs was also set to 0.01 mg/l. This ensured one program would not detect the contaminant faster than the other simply because it had a lower detection limit. As mentioned, the KYPIPE sensor placement tool utilizes the EPANET engine to perform hydraulic and water quality calculations. This further reduces any differences in the programs prior to sensor placement optimization.

Comparison of Times to Detection

The baseline contamination scenario (considered the baseline case because it was present in all three general contamination scenarios) injected a contaminant at 1000 mg/min for four hours.

The comparison of time to detection between the two programs using the baseline condition for the placement of one sensor is shown in Figure 6. Because the time differences between the nodes selected by TEVA-SPOT and KYPIPE are minimal, the percent differences in times are also displayed.

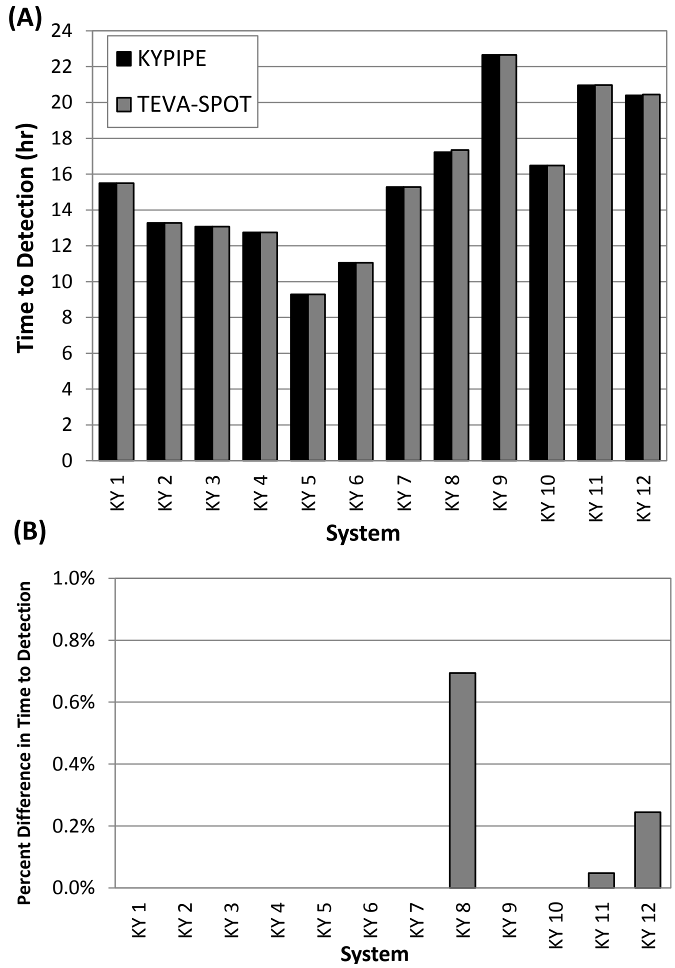

Figure 6. Comparison of Nodes Selected by KYPIPE and TEVA-SPOT for Baseline Conditions and Placement of One Sensor: (A) Times to Detection; (B) Percent Difference in Times to Detection.

Figure 6 shows that sensors selected by KYPIPE for the baseline contamination scenario resulted in times to detection either equal to that of nodes selected by TEVA-SPOT or slightly lower for all system models. Similar results were found for all 15 of the contamination scenarios. The KYPIPE optimization method resulted in slightly lower or equal times to detection than TEVA-SPOT for all contamination scenarios in all 12 systems, thus confirming that the complete enumeration strategy employed by KYPIPE was either equal to or superior to the performance of the GRASP algorithm used by TEVA-SPOT. Similar results were obtained for the two sensor solutions as well. Although sensors selected by KYPIPE were always superior or equal to those chosen by TEVA-SPOT, the differences in average times to detection were minimal. The time to detection between the nodes selected by each program for the same system and contamination scenario were fairly similar (if not equal), as shown by the average percent difference in time to detection (averaged over the 15 contamination scenarios) displayed in Figure 7. The average percent difference in times is relatively low for all systems. It should be noted that the few systems with 0 percent average time differences did not all have matching sensor selection for all 15 scenarios, but the time differences between the selected nodes in these cases were negligible.

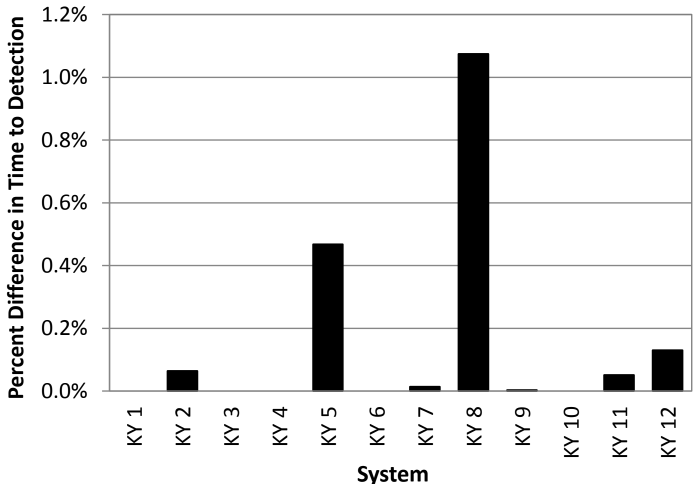

Figure 7. Average Percent Difference in Time to Detection between Nodes Selected by TEVA-SPOT and KYPIPE for Baseline Scenario and Placement of One Sensor.

While the TEVA-SPOT algorithm generally required less computation time than KYPIPE, the differences in times were not that significant. The largest system (i.e. KY 12) required 45 minutes for KYPIPE and 13 minutes by TEVA-SPOT for the placement of one sensor. When placing two sensors, KYPIPE required 1 hour and 20 minutes while TEVA-SPOT used 13.5 minutes. Differences in computational times were not as significant for the remainder of systems in the model database.

In addition to comparing the times to detection for nodes selected by KYIPIPE and TEVA-SPOT, the times to detection obtained by KYIPIPE averaged over the contamination scenarios for both one and two sensor systems were also compared (see Figure 8). As can be seen from the figure, addition of a second sensor, at least for the 12 systems examined, did not seem to add a significant amount of benefit. All but three of the systems resulted in less than 15 percent decrease in average time to detection when the second sensor was added. Thus, for the small systems analyzed, one might argue that one sensor might be adequate. Additional analyses would be required to confirm such a hypothesis.

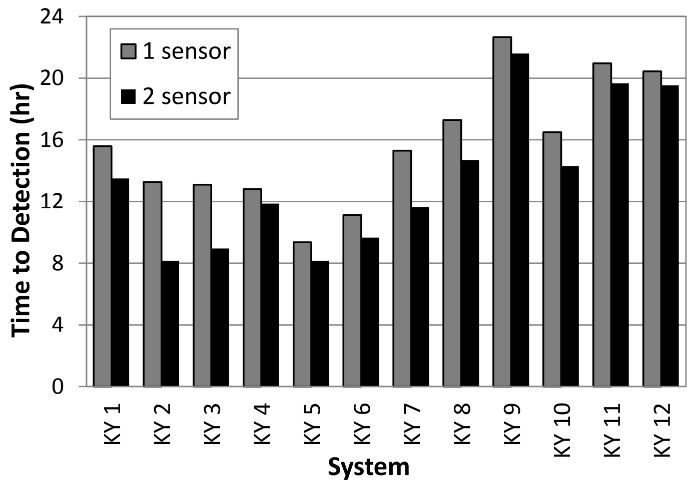

Figure 8. Comparison of Performance between the Placement of One and Two Sensors.

Comparison of Identical Sensor Placement

Along with a comparison of the times to detection of nodes selected by KYPIPE and TEVA-SPOT, the location of nodes chosen as optimal sensors locations by both programs were also compared. Some contamination scenarios for the same system resulted in TEVA-SPOT and KYPIPE selecting the same nodes as the optimal sensor locations. The percentage of the 15 contamination scenarios that resulted in identical sensor selection between KYPIPE and TEVA-SPOT for each system is summarized in Figure 9 and Figure 10.

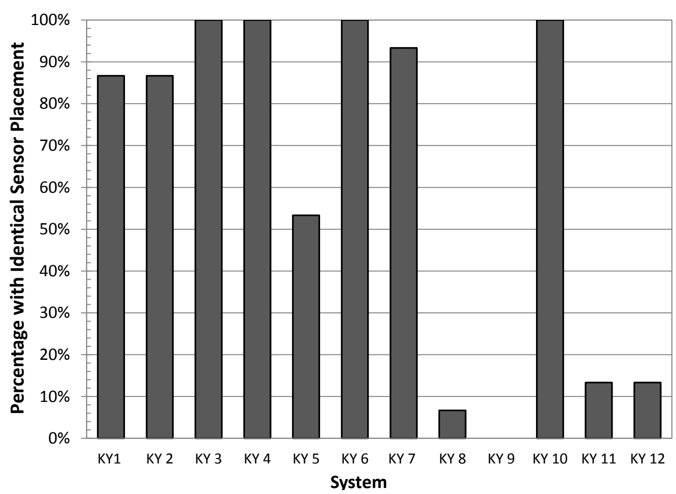

Figure 9. Percentage of Contamination Scenarios with Identical Sensor Selection between KYPIPE and TEVA-SPOT (1 sensor).

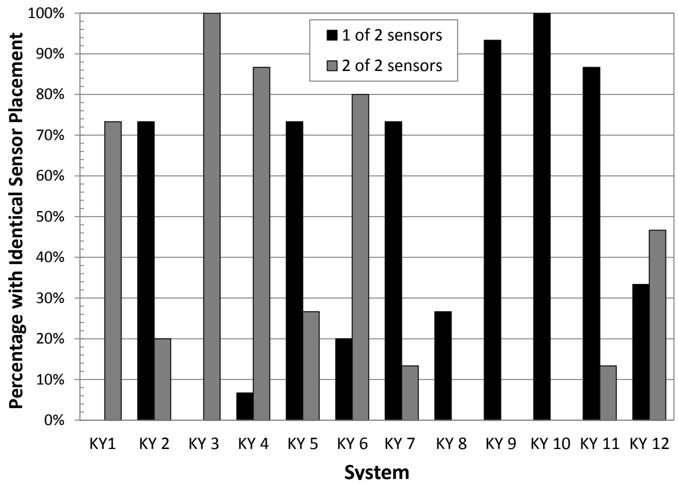

Figure 10. Percentage of Contamination Scenarios with Identical Sensor Selection between KYPIPE and TEVA-SPOT (2 sensors).

Figure 9 shows that some systems resulted in matching sensor selection between KYPIPE and TEVA-SPOT for all 15 contamination scenarios, while other networks did not have any matching sensor placement among scenarios. For the placement of one sensor, four of the 12 systems had matching optimal sensor nodes for all 15 scenarios, and eight of the 12 models had identical placement for over 50 percent of scenarios. On average, 9.4 out of the 15 scenarios (63%) resulted in identical placement of sensors for all systems. There were four systems that had less than 20 percent matching sensor nodes between KYPIPE and TEVA-SPOT, and one system (KY 9) did not have any matching sensor selection. In these systems with very few matching sensors, further investigation revealed that the vast majority of these sensors were still in very close proximity to each other. Only six contamination scenarios (out of the 15 scenarios performed on 12 systems for a total of 180 simulations) led to sensor locations that were considered to be far away from each other in the system, and all of these cases occurred in KY 8. Although these few scenarios resulted in sensor selection that was considered to have significant spatial variation, the nodes were still located in the same general region of the system. This concept is shown in Figure 11. The top portion shows sensor selection that is considered to have significant spatial variation, and the bottom portion illustrates sensor selection in close proximity. As mentioned, the vast majority of scenarios with different sensor selection between KYPIPE and TEVA-SPOT resulted in nodes in very close proximity to each other for the placement of one sensor.

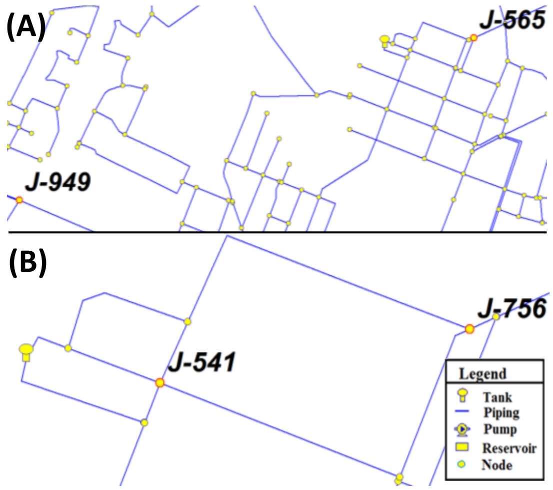

Figure 11. Spatial Variation in Sensor Selection (KY 8): (A) Significant Spatial Variation; (B) Close Proximity

Similar trends were present when placing two sensors as with the one sensor scenario. Figure 10 shows that some systems resulted in both sensors matching between KYPIPE and TEVA-SPOT for many of the 15 contamination scenarios, while other systems matched one out of two sensors placed for certain scenarios. KY 3 was the only system to result in identical sensor placement for both sensors between KYPIPE and TEVA-SPOT for all 15 contamination scenarios. KY 1, KY 3, KY 4, and KY 6 resulted in two out of two matching sensor node selection for over 50 percent of the contamination scenarios, showing that KYPIPE and TEVA-SPOT produced very similar results in these systems. KY 8, KY 9, and KY 10 did not have any scenarios that matched both sensors, but these systems did have one out of two identical sensor nodes for several scenarios. KY 2, KY 5, KY 7, KY 9, KY 10, and KY 11 resulted in one out of two matching sensors for over 70 percent of the contamination scenarios. An average of 5.8 scenarios (38%) of the 15 contamination scenarios averaged over all systems resulted in matching sensor selection for both sensors, while 7.3 out of the 15 scenarios (49%) resulted in identical sensor placement for one of the two sensors, averaged over all 12 networks.

The cases where KYPIPE and TEVA-SPOT recommended different sensor nodes were investigated. As with the placement of one sensor, the vast majority of these cases resulted in placement of sensors that were in close proximity to each other. Only a few cases produced results where the sensors recommended by KYPIPE and TEVA-SPOT were significantly far away from each other. Specifically, 20 cases (out of the 15 simulations run on all 12 systems) led to sensor selection between the two programs that varied considerably spatially. In all of these cases, only one of the two sensors placed showed significant spatial variation between the two programs. In 18 of these cases, placement of the other sensor was identical, and the other two scenarios placed the remaining sensor in very close proximity. 15 of these 20 cases of significant spatial variation of one sensor occurred in KY 10, three were present in KY 6, and the final two cases occurred in KY 12.

Because several cases occurred where KYPIPE and TEVA-SPOT selected different sensor nodes, the ranking of the nodes chosen by TEVA-SPOT were investigated using the average times to detection for all possible sensor locations generated by KYIPE. The data in Table 2 displays the percentage of the 15 contamination scenarios that resulted in differing sensor node selection between KYPIPE and TEVA-SPOT for the placement of one sensor (also shown in Figure 9). Table 2 also shows the ranking of the node selected by TEVA-SPOT (averaged over all 15 contamination scenarios) based on times to detection generated by KYPIPE, the total number of possible sensor locations, and the percentage of sensor locations with higher rankings (lower times to detection) than the location selected by TEVA. The values reported for average ranking of the node and percent of sensor locations with lower times to detect include cases where KYPIPE and TEVA-SPOT selected the same sensor.

Table 2. Ranking of Sensor Nodes Selected by TEVA-SPOT (1 sensor).

System	Possible Sensor Locations	Percentage of Scenarios Resulting in Different Sensor Selection	Average Ranking (by KYPIPE) of Node Selected by TEVA-SPOT	Average Percent of Sensor Locations with Lower Times to Detect
KY 1	509	13.3%	1.13	0.03%
KY 2	597	13.3%	1.13	0.02%
KY 3	224	0.0%	1.00	0.00%
KY 4	683	0.0%	1.00	0.00%
KY 5	305	46.7%	1.87	0.28%
KY 6	386	0.0%	1.00	0.00%
KY 7	378	6.7%	1.20	0.05%
KY 8	926	93.3%	7.27	0.68%
KY 9	836	100.0%	2.53	0.18%
KY 10	673	0.0%	1.00	0.00%
KY 11	547	86.7%	2.73	0.32%
KY 12	1908	86.7%	8.13	0.37%

If the node selected by TEVA-SPOT is ranked high (and thus the percent of sensor combinations with lower times to detection is low), this shows the results provided by the two programs are similar. Even when KYPIPE and TEVA-SPOT select different ideal sensor nodes, similar times to detection and high rankings of the node chosen by TEVA-SPOT show that both programs are effective in providing sensor locations that will detect contaminants quickly. Table 2 shows that all systems averaged less than 1 percent of possible sensor combinations that are considered better than the sensor locations selected by TEVA-SPOT in terms of low time to detection. Even though KY 8 and KY 12 had slightly higher values for average ranking of the node selected by TEVA (7.27 and 8.13, respectively), they still had very low percentages of sensor locations with faster times to detection.

Analysis and Discussion

Results of the verification study demonstrate the effectiveness of the KYPIPE sensor placement optimization method for small systems. The slightly faster times to detection (as compared to nodes selected by TEVA-SPOT) for the non-common selected sensor locations using the

KYPIPE algorithm show that KYPIPE is producing slightly superior sensor placement results to TEVA-SPOT (utilizing the GRASP heuristic method). However, the relatively low differences in average times to detection of the selected sensor nodes show that results between the two programs are very similar. The programs did select identical sensor locations in many cases, but the KYPIPE algorithm will always produce superior (or the same) sensor placement over TEVA-SPOT. The average percent difference in time to detection between nodes selected by TEVA-SPOT and KYPIPE for the placement of one sensor was 0.15 percent. This value takes into account the numerous scenarios where TEVA-SPOT and KYPIPE selected the same sensor node. The maximum percent difference in average time to detection between differing nodes selected by the two programs was 2.9 percent.

Because the models used in both TEVA-SPOT and KYPIPE software are uncalibrated models of real distribution systems, the hydraulic/water quality analyses and subsequent times to detection for all possible sensor locations are only estimates. However, the data should be fairly similar to results that would occur through tracer studies or other field testing. Because KYPIPE utilizes the EPANET engine to perform hydraulic/water quality analyses, the occasional slightly faster times to detection can be attributed to the optimization method utilized by KYPIPE, which will always produce superior or equal results to TEVA-SPOT. KYPIPE utilizes an enumeration method that calculates travel times for the entire solution space while a GRASP heuristic method was utilized in TEVA-SPOT for this study. Therefore, it is logical that the times between the programs are similar but KYPIPE consistently produces slightly faster times.

Summary and Conclusions

TEVA-SPOT has been developed to analyze the vulnerability of drinking water distribution networks to contamination and recommend locations to deploy water quality sensors as a component of contamination warning systems. However, the software may not be appropriate for small utilities in terms of the simplicity and ease of use. A Water Quality Sensor Placement Tool was developed with KYPIPE to accomplish the objective of providing a simple tool to aid small utilities in sensor placement. The KYPIPE tool uses a complete enumeration optimization scheme along with EPANET for both hydraulic and water quality analysis. Both TEVA-SPOT and KYPIPE were used to locate either one or two water quality sensors for 12 different water distribution systems. 15 different contamination scenarios were evaluated for each system. The results of this analysis have provided the following conclusions:

1) The KYPIPE sensor placement tool provides sensor locations equal to or superior to those provided by TEVA-SPOT when using the GRASP optimization option along with an objective to minimize the time to detection.

2) The TEVA-SPOT algorithm was able to converge to a solution in less time than the KYPIPE sensor tool. The relative difference in computational times was greater for the two sensor solutions than the one sensor solutions. Such an observation is to be expected given the fact that the KYPIPE sensor tool uses a complete enumeration algorithm. Despite this fact, the actual computational requirements for the KYPIPE tool were not excessive, even for the largest system analyzed (i.e. KY 12 which had nearly 2000 potential sensor locations).

3) While the KYPIPE sensor placement tool was superior to the TEVA-SPOT algorithm, the results were generally not that significant. Thus, the GRASP algorithm used by TEVA-SPOT would appear to be fairly efficient, even providing the global optimal solution in eight of the 12 systems analyzed (for the baseline contamination scenario when placing one sensor).

4) A comparison of the times to detection for both the placement of one and two sensors in the systems revealed that the average time to detect did not significantly decrease with the addition of a second sensor. Thus, for small systems, use of a single sensor might be adequate to provide an acceptable level of protection for utilities with limited financial resources. Additional analyses with an increased number of sensors should be performed to validate this hypothesis.

5) If one or two sensors will provide sufficient coverage for many small systems, some general rules or guidelines could be developed for sensor placement that might be able to avoid the requirement of a calibrated network model as currently required by TEVA-SPOT and KYPIPE. One possible such methodology for single sensor systems has been proposed by Schal et al. (2013). Such methodologies could prove to be especially helpful for small utilities that may not have the technical or financial resources to employ computer model based approaches such as TEVA-SPOT.

Acknowledgement

Funding for this research was provided by the U.S. Department of Homeland Security, Science and Technology Directorate, through a technology development and deployment program managed by The National Institute for Hometown Security, under an Other Transactions Agreement, OTA #HSHQDC-07-3-00005, Subcontract #02-10-UK.

References

Berry, J., Hart, W. E., Phillips, C. A., Uber, J. G., and Watson, J-P. (2006). Sensor Placement in Municipal Water Networks with Temporal Integer Programming Models. *Journal of Water Resources Planning and Management*, 132(4), 218–224.

Berry, J. W., Boman, E., Riesen, L. A., Hart, W. E., Phillips, C. A., and Watson, J.-P. (2010). *User's Manual: TEVA-SPOT Graphical User Interface User's Manual.* EPA-600-R-08-147, U.S. Environmental Protection Agency, Office of Research and Development, National Homeland Security Research Center, Cincinnati, OH.

Hart, W. E., & Murray, R. (2010). Review of Sensor Placement Strategies for Contamination Warning Systems in Drinking Water Distribution Systems. *Journal of Water Resources Planning and Management*, 136 (6), 611-619.

Janke, R., Murray, R., Uber, J., & Taxon, T. (2006). Comparison of Physical Sampling and Real-Time Monitoring Strategies for Designing a Contamination Warning System in a Drinking Water Distribution System. *Journal of Water Resources Planning and Management*, 132 (4), 310-313.

Jolly, M.D., Lothes, A.D., Bryson, L.S., & Ormsbee, L. (2013). Research Database of Water

Distribution System Models. *Journal of Water Resources Planning and Management,* Accepted.

McKenna, S. A., Hart, D. B., & Yarrington, L. (2006). Impact of Sensor Detection Limits on Protecting Water Distribution Systems from Contamination Events. *Journal of Water Resources Planning and Management*, 132 (4), 305-309.

Murray, R., Janke, R., & Uber, J. (2004). The Threat Ensemble Vulnerability Assessment (TEVA) Program for Drinking Water Distribution System Security. *Critical Transitions in Water and Environmental Resources Management* (p. 8). Salt Lake City, UT: World Water Congress.

Murray, R., Hart, W., & Berry, J. (2006). Sensor Network Design for Contamination Warning Systems: Tools and Applications. *American Water Works Association* (p. 20). Washington, D.C.: Water Security Congress Conference.

Murray, R., Janke, R., Hart, W. E., Berry, J. W., Taxon, T., & Uber, J. (2008). Sensor network design of contamination warning systems: A decision framework. *American Water Works Association*, 100 (11), 97-109.

Murray, R., Haxton, T., & Janke, R. (2010). *Sensor Network Design for Drinking Water Contamination Warning Systems.* Cincinnati, OH: National Homeland Security Research Center, Office of Research and Development, U.S. Environmental Protection Agency.

Schal, S. L., 2013. Water Quality Sensor Placement Guidance for Small Water Distribution Systems. Master's thesis, Department of Civil Engineering, University of Kentucky, Lexington, KY.

U.S. Environmental Protection Agency. (2008). *Water Quality in Small Community Distribution Systems - A Reference Guide for Operators.* Cincinnati, OH: Office of Research and Development.

Von Huben, H. (2005). Water Distribution Operator Training Handbook. American Water Works Association. Denver, CO.

Wood, D. J. (1981). *Algorithms for Pipe Network Analysis and Their Reliability.* Lexington, KY: University of Kentucky Water Resources Research Institute.

Wood, D. J. (2010). *KYPipe Reference Manual.* KYPipe, LLC