Heuristic Approach for Operational Response to Drinking Water Contamination

Annie Poulin, M.ASCE[1]; Alain Mailhot[2]; Patrice Grondin[3]; Louis Delorme[4]; Nathalie Periche[5]; and Jean-Pierre Villeneuve[6]

Abstract: This paper introduces a simple topological approach to systematize the isolation of contaminated areas within the pressure zones of drinking water distribution systems (DWDSs). Assuming optimal location of contaminant detectors and known flow conditions, a heuristic procedure delineates the area to be isolated and identifies the valves to be closed by response teams sent in the field, taking into account a response delay from the time of first detection. As a first step leading to the development of a more comprehensive algorithmic application, the approach was elaborated and validated from a pragmatic perspective using two real-world DWDSs. Depending on each network's design, configuration, and assumed flow conditions, application of the isolation procedure will result in different isolation strategies (extent of isolated areas, number of required isolation valves). The current approach is based on a set of simplifying assumptions, and needs to be further validated with other networks. The proposed methodology can be used to assess the required emergency response capabilities and even possible network design improvements.

DOI: 10.1061/(ASCE)0733-9496(2008)134:5(457)

CE Database subject headings: Potable Water; Water distribution systems; Security; Risk management; Water pollution.

Introduction

Because of their potential adverse impacts on public health, drinking water contaminants occurring within distribution systems are highly worrying. Pathogens, chemical contaminants, and radionuclides can be introduced in a drinking water distribution system (DWDS) either accidentally, through system deficiencies (e.g. cross connection and back-siphonage, Craun and Calderon 2001), or intentionally, as a result of malevolent acts. Over the past 4 decades, many accidental contamination events associated with DWDS deficiencies have caused illness and death (Craun et al. 2006; Craun and Calderon 2001). Since the events of September 11, 2001, concern about intentional attacks against DWDSs has also grown considerably. A multiple-barrier approach can be implemented by water utilities to reduce health risks associated with contaminants introduced in DWDSs. But these barriers are unlikely to effectively mitigate most of the high-impact health risks (Allman and Carlson 2005).

The vulnerability of DWDSs naturally leads to the desire to enhance their security. Many water utilities now recognize in their mission statement the need to focus on "new" security concerns. Specifically, DWDSs should be operated in a way to protect against, detect, and respond to human-caused and natural hazards (Bell et al. 2004). To meet these goals, new protective technologies and risk mitigation measures are needed. Such resources will only make water systems safer if they are supported by well organized emergency response and recovery plans (Herrick et al. 2006). Therefore, new simulation tools and procedures need to be developed and made available to water utilities, to help detect and manage contamination events in practical applications.

Background

In the last decades, computerized simulation tools have become standard practice for the design, operation, management, and analysis of DWDSs (Walski et al. 2003). Nowadays, with growing security concerns, development of new simulation tools and procedures that are better adapted to this specific context is an active research area. Recent works have first focused on coupling existing modeling software with mathematical programming methods to support the design of contamination warning systems (CWSs). A CWS is an integrated system for continuously collecting, integrating, analyzing, and communicating information to provide a timely warning of potential water contamination incidents (USEPA 2005). When designing a CWS, a well thought out placement of sensors throughout a DWDS is a key element since accessibility and budget constraints preclude complete coverage.

[1]Ph.D. Student, Institut National de la Recherche Scientifique, Centre Eau Terre & Environment, 490 rue de la Couronne, Quebec QC, Canada G1K 9A9 (corresponding author). E-mail: annie.poulin@ete.inrs.ca

[2]Professor, Institut National de la Recherche Scientifique, Centre Eau Terre & Environment, 490 rue de la Couronne, Quebec QC, Canada G1K 9A9. E-mail: alain.mailhot@ete.inrs.ca

[3]Senior Engineer, Teknika-HBA, 150 rue de Vimy, Sherbrooke QC, Canada J1J 3M7. E-mail: patrice.grondin@teknika-hba.com

[4]Researcher, Hydro-Québec Institut de Recherche, 1800 Boulevard Lionel-Boulet, Varennes QC, Canada J3X 1S1. E-mail: delorme.louis@ireq.ca

[5]Senior Engineer, Aqua Data Inc., 95 5e Ave., Pincourt QC, Canada J7V 5K8. E-mail: nperiche@aquadata.com

[6]Professor, Institut National de la Recherche Scientifique, Centre Eau Terre & Environment, 490 rue de la Couronne, Quebec QC, Canada G1K 9A9. E-mail: jpv@ete.inrs.ca

Note. Discussion open until February 1, 2009. Separate discussions must be submitted for individual papers. The manuscript for this paper was submitted for review and possible publication on December 18, 2006; approved on November 7, 2007. This paper is part of the ***Journal of Water Resources Planning and Management***, Vol. 134, No. 5, September 1, 2008. ©ASCE, ISSN 0733-9496/2008/5-457–465/$25.00.

Various models address the problem of optimal sensor location, to increase protection against a predefined set of contamination events. From a modeling point of view, these models either assume steady-state flow conditions (e.g., Kessler et al. 1998; Berry et al. 2005) or unsteady flow conditions (e.g., Ostfeld and Salomons 2004; Berry et al. 2006; Propato 2006). Models using time-variable flow can more accurately analyze the response of DWDSs to contamination but they also require extensive water quality simulation computing, as opposed to steady-state models which are based on hydraulic simulation only. Various sensor placement objectives have been considered, generally seeking to maximize the detection likelihood or to minimize the impact of contamination events based on surrogate indicators (e.g., time to detection, volume of contaminated water consumed prior to detection, extent of contamination prior to detection) for a fixed number of sensors.

Other related works have addressed the inverse problem of contamination source identification, assuming the existence of CWSs (e.g., Laird et al. 2006; Guan et al. 2006; Propato et al. 2006; Tryby et al. 2007). An important issue about the source inversion problem is the tradeoff between the inversion feasibility and the number of sensor measurements to be collected under known flow conditions. As the number of measurements increases over time, the problem is better defined but contaminant spread and public exposure also increase (Propato et al. 2006).

All of the above cited works assumed perfect sensors (no failure, contaminant specific) and instant alarm triggers. Efforts have been undertaken to make the sensor placement and source inversion models more realistic (e.g., Ostfeld and Salomons 2005; Hart et al. 2007; Hill et al. 2006; Preis et al. 2007b). However, research fields related to critical issues such as contaminant detection and accurate impact assessment on public health still remain open.

Contamination Consequence Management

To better guide water utilities in preparing for and responding to drinking water contamination, the U.S. Environmental Protection Agency (USEPA) recently provided the Response Protocol Toolbox (RPTB) (USEPA 2003). According to the RPTB, the occurrence of contamination, once suggested by a particular warning, should trigger a series of actions aiming to limit the impact of potential contamination on public health and to return to a normal operation of the water network. As threat credibility becomes clearer, response actions to be implemented are more precisely defined. From an operational point of view, an efficient response strategy may consist of: (1) appropriate valve closures to isolate the contaminated water; (2) hydrant flushing to evacuate the contaminated water from a previously isolated network area and, in some cases, combining such field actions with valve manipulations to reach high velocities in unidirectional pipe sections and remove adhering contaminants; and (3) injection of chemical oxidants or cleaning agents to enhance the removal of adhering contaminants (USEPA 2004a,b). Contaminated water removed from a DWDS may be discharged to a nearby source or to a wastewater collection system if it does not represent a hazard to the environment. Otherwise, it should be pretreated prior to proper disposal (on site or off site).

The computerized tools presented previously increase protection against drinking water contamination incidents and may help better assess the consequences of such events (e.g., contamination spread, affected users). However, these do not give indications regarding the management of those consequences, to effectively prevent further exposure to contaminants once detection has occurred, and to eliminate the contaminants. In fact, few analytical tools or systematic procedures are available to support the response and recovery processes. Baranowski (2007) and Preis et al. (2007a) have suggested, as a first response, optimal combinations of hydrant openings (or demand altering) and valve shutoffs to reduce contaminant concentration in water networks. Such approaches are mostly useful when the contaminant concentration profile throughout the network is precisely known at every time step of the event (i.e., assuming complete knowledge of water consumption). If not, it may be difficult to prevent contaminant dissemination to unaffected network areas. Moreover, it seems logical that purging contaminated water from a system be deferred until the completion of thorough investigation/confirmation and, in cases where the contaminated water is deemed hazardous, until equipment is in place to contain or treat the flushed water. In such situations, preventing contaiminant spread prior to early hydrant opening may be warranted. Poulin et al. (2006) examined the single issue of contaminant isolation through targeted valve closings as a first operational response to detected contamination. Isolation prevents contaminant spreading, hence reducing the risk of contaminated water consumption, even when additional delays are required before contaminated water evacuation (flushing) takes place. Physical containment of contaminated water also permits better management of subsequent flushing activities.

The works of Poulin et al. (2006) are further extended in this study, which introduces a heuristic procedure to systematize the isolation of contaminants within the pressure zones of DWDSs, taking into account operational and modeling issues. Potentially contaminated areas are delineated based on first detection information provided by optimally located sensors. Through an iterative set of heuristic rules, an isolation procedure (IP) then guides the elaboration of isolation strategies and identifies related operations in the field. The methodology was developed and validated using two small real-world DWDSs located in the province of Québec, both having different designs and topological configurations. The goal was to elaborate a first version of a pragmatic methodology that could eventually be applied to other DWDSs with various configurations and sizes, and evolve to become a more comprehensive computerized application. The procedure is part of ongoing research aimed at elaborating an operational response strategy that includes isolation of contaminated areas and flushing contaminated water from DWDSs, to limit health risks and efficiently return to a normal level of service.

Methodology and Assumptions

Isolation of contaminants requires effective response actions and, at first glance, seems well suited to optimization. To this end, many relevant performance objectives could be considered, similar to those proposed by Baranowski (2007) and Preis et al. (2007a) though in this study, no general optimization scheme is utilized. Instead, a heuristic approach was elaborated to define isolation strategies based on a set of rules that seek, at the same time, to limit the operational constraints related to subsequent unidirectional flushing feasibility, limit the extent of isolated areas within pressure zones, and as much as possible the number of valve closings, while integrating security margins and ensuring that isolation operations in the field can be executed as quickly as possible. Taking all such issues into account within a single problem may greatly increase optimization complexity, especially when one considers that every real-world DWDS is unique.

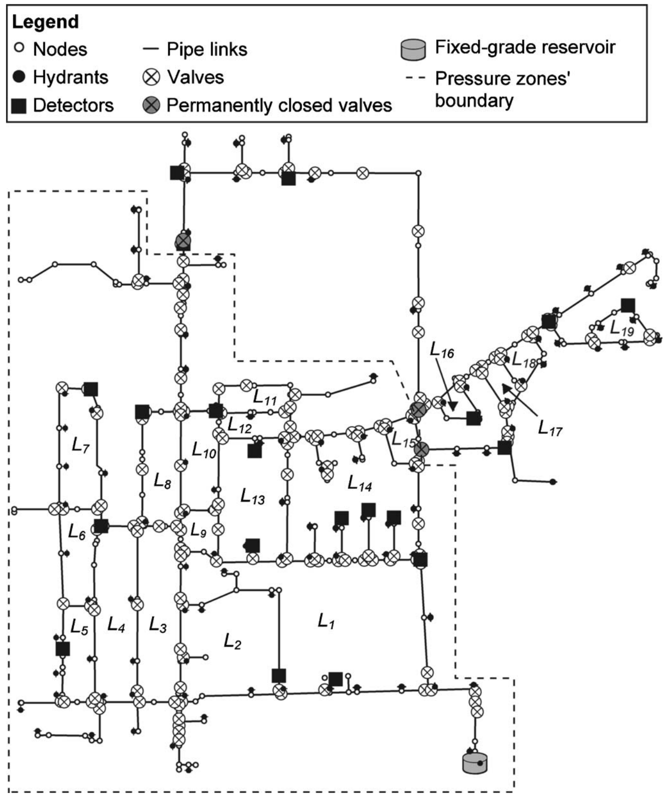

Fig. 1. Valcourt network, with optimal location of 20 contamination detectors

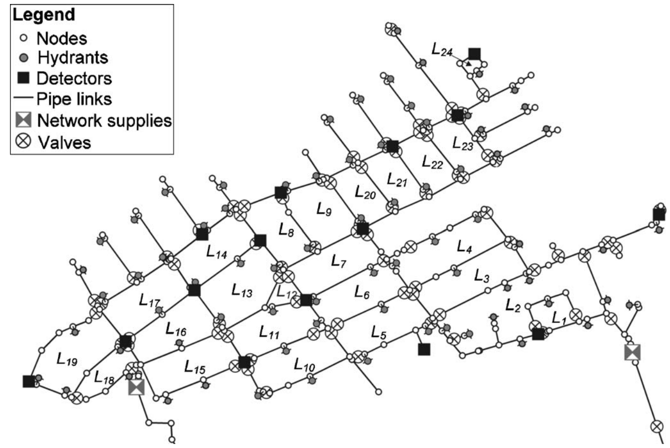

Fig. 2. TV network, with optimal location of 15 contamination detectors

Experimental DWDSs

Two real-world DWDSs with different configurations and designs served to elaborate the IP. The Valcourt network (Fig. 1) was used to build a first version of the procedure. This network is located about 114 km east of Montreal. It is mainly looped and partially branched, it is supplied by a single surface water source (fixed-grade reservoir), and contains 303 nodes and 324 pipe links. The Terrasse–Vaudreuil (TV) network (Fig. 2) was used to validate and refine the IP. This second network is located about 40 km west of Montreal. It is more densely looped and this characteristic added complexity to the analyses. The TV network has two water supplies form a neighboring utility, and contains 206 nodes and 229 pipe links. The DWDSs cover areas of about 4 and 1.5 km^2 and have total pipe lengths of 20.4 and 12.7 km, respectively. Both networks mainly serve residential water consumers.

The network models were created with Aqua Cad Suite software (Aqua Data 2005) and calibrated in steady-state mode. They are all-pipes models and include every valve and hydrant as required by Aqua Cad Suite's, AquaGeo, which is specifically designed for DWDS operation, management, and analysis. In the models, valves are pipe link attributes, whereas hydrants are node attributes. The Valcourt network has two pressure zones (Fig. 1) and TV has only one.

Implementation of Contamination Warning System

We assume in this study that a CWS comprised of optimally located detectors raises the first alarm, should a contamination occur in a DWDS. Provided that, ideally, consumption of contaminated water would be completely prevented throughout a DWDS, detectors were located according to the objective of minimizing the volume of contaminated water consumed networkwide prior to detection. A static sensor-placement model proposed by Watson et al. (2004) was used and is based on the following assumptions: (1) a single static flow pattern is considered and represents the mean daily flow conditions in the experimental DWDSs (over 24 h time horizon); (2) contaminant transport is based on water travel times in pipe links, and consequently the impact of contamination is only estimated by tracking the presence or absence of contaminants at network points; (3) the time at which contamination occurs is not taken into account since flow conditions are time invariant (they represent mean daily conditions); (4) failure of the CWS is ignored; (5) the CWS can detect a wide range of contaminants at any concentration level and an alarm is raised as soon as contamination is detected; (6) each contamination event occurs at a single node in the network; (7) contaminant growth or decay in DWDSs is ignored; (8) contaminated water represents a danger for public health immediately as it reaches a node for the first time (i.e., by the shortest path in terms of travel time from a contamination source); (9) contaminant discharge is continuous until isolation is completed (remember that isolation is seen as a first operational response; in the worst-case scenario, contaminant release would be maintained long after isolation but at least spreading would be contained; in every case, however, flushing would only be implemented upon confirmation that discharge has ended); (10) detectors are located at network nodes; and (11) the assumed detector budget is not restricted by the cost of available technologies.

Steady-state sensor placement models are usually run with contamination scenarios that occur in several consecutive non-overlapping flow patterns over 1 day. Although such models do not simulate transitions between each flow pattern, they at least take into account the possible flow inversions in separate patterns. In this paper, locating contaminant detectors based on only one flow pattern was favored in the context of a first validation of the methodology and since the experimental DWDSs were calibrated based on this single flow pattern. It also is in accordance with the idea of a simple modeling framework that allows, upon detection, the rapid identification of a set of potential contamination sources from where contaminant spread may be tracked through time. This is made possible because the sensor placement model assigns every contamination scenario either to a detection site or to "no detection" (for further details see Watson et al. 2004). Each detector then determines a potentially contaminated area, considering all possible contamination sources (see the next section).

Heuristic Isolation Procedure

In terms of operations, the isolation strategy adopted consists of: (1) shutting down appropriate valves to physically contain the contaminated water while (2) leaving one pipe "open" as a clean water supply to the isolated area. The assumptions underlying the IP are as follows:

1. As this is the aim of this study, we assume that contaminant isolation is possible within pressure zones, i.e., inside boundaries delimited by remotely controlled pressure reducing valves or pumps, and closed valves;
2. After a first detection occurs, a *response delay* is assumed. This delay is set to take into account the time required for operational response planning (assuming previous training and good coordination of response personnel), for the organization of public notification, and ideally for contamination confirmation or at least undertaking an investigation. The end of this delay marks the point in time at which isolation operations are set in motion;
3. Consumption conditions remain unchanged across a DWDS during isolation operations, meaning that compliance with public notification begins after isolation is completed. Based on this study, it is fair to believe that the time required for isolation is significantly shorter than the response delay, at the end of which public notification broadcast should be initiated;
4. Operational response planning begins at the time of first detection. The approach would need to be adapted if subsequent detections were taken into account, as this additional information would modify the assumption about contaminant location;
5. City roads follow the exact same layout as DWDSs;
6. Every DWDS has a "drinking water alert station" (DWAS), from where response operations are initiated;
7. Each valve in a DWDS is assumed watertight, accessible, and in maneuverable state, although a constraint on valve operational condition could easily be included in the IP;
8. Each response team consists of one specialized vehicle and two operators;
9. A mean velocity of 30 km/h is assumed for response teams' displacements along city roads, considering the shortest path between every starting and ending point;
10. Valve operations required for isolation are executed simultaneously. This implies that an unlimited number of response teams are available, and that field operations are only executed once every team is in place;
11. A fixed delay is assumed for valve shutoff. This *manipulation* delay is mostly dependent on valve diameter. To avoid water hammers, larger diameters require longer shutting times. An *operational delay* refers to the total time required to reach a valve (variable *displacement delay*) and shut it down (fixed *manipulation delay*); and
12. Dead-end lines without terminating hydrants can be purged, for instance through blowoff valves, since contaminated water that reaches these areas needs to be evacuated.

The basic idea of the IP is to delineate and isolate areas from which the maximum volume of (and ideally all) the contaminated water will be readily flushed, without having to further modify their extent. Topologically, areas to be isolated are comprised of a set of natural loops and adjacent linear or branched pipe sections.

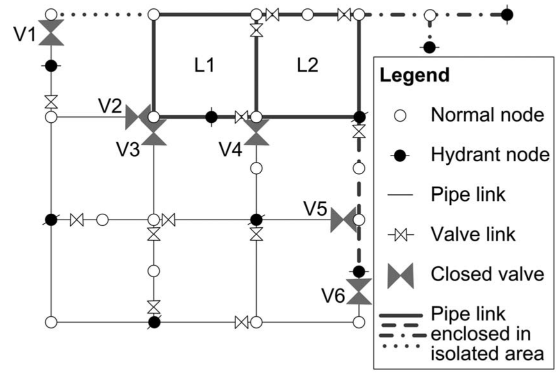

Fig. 3. Explicative representation of natural loops, supply links, and linear or branched pipe sections

A natural loop is a closed pipe circuit in a network with no additional interior pipe circuits (Wood and Rays 1981) and, pragmatically, would most likely be picked visually by a designer (Rahal 1995) (e.g., L1 and L2 in Fig. 3). Epp and Fowler (1970) introduced an algorithm for natural loop data generation from the pipe-node connectivity data. The linear or branched pipe sections may be comprised of: (1) existing networks' dead-end lines or stems (dash-dot lines in Fig. 3); (2) "artificial" dead-end lines created by shutting down valves on natural loops, where their terminating valve can either be categorized as a clean water supply to the isolated area (doted line in Fig. 1) or not; in the latter case, the valve must be connected to a hydrant enclosed in the isolated area (dashed lines in Fig. 3); and (3) a combination of (1) and (2). After isolation is completed, unidirectional flushing is intended and will proceed through the clean water supply links, using the hydrants located on the natural loop pipes and at the end nodes of the linear or branched pipe sections.

The general form of the heuristic IP introduced in this study is represented by the flowchart in Fig. 4. The procedure is triggered by first contaminant detection, at a given detector. Next, assuming a 2 h response delay, the "worst-case" contaminated area is delineated. In this study, worst case means taking into account all the presumed contamination sources at the time of first detection, as mentioned in the "Implementation of Contamination Warning System" subsection. Aggregated contaminant spread from all these sources is obtained using the results from the steady-state hydraulic analysis of the network and is subject to the same simplifying assumptions as in the sensor placement model. So clearly, the "worst-case contaminant spread" is dependent upon the simulation framework. The potentially contaminated areas unavoidably contain uncertainty, unless a unique source is associated to the first-detection site, which is unlikely. But, as already mentioned, in this paper a quick response is advocated as soon as possible after initial detection, making the best use of the limited information available at this time. Steps 2–9 of the IP (Fig. 4) lead to the identification of the final isolation strategy and are detailed in the "Specific Results" subsection, using two application examples.

Results

Optimal Location of Contaminant Detectors

Using the GLPK package (GNU 2006), 20 and 15 contaminant detectors were optimally located on the Valcourt and TV net-

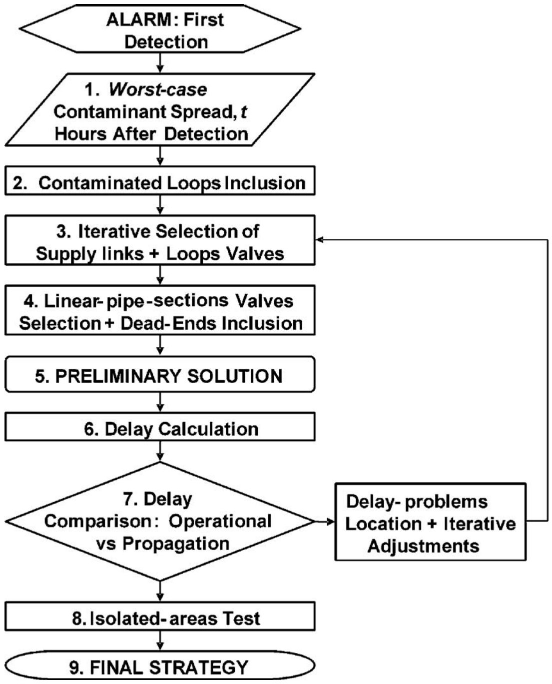

Fig. 4. Flowchart of IP

works, respectively, representing approximately 7% of the total number of nodes in both networks. In accordance with the main objective of the current study, these detector budgets allowed the delineation of contaminated areas of a fairly limited extent at first detection. Hydraulic simulation results were computed using AquaGeo (Aqua Data 2005).

Excluding the supply nodes (assumed to be protected by existing sensors), 206 nodes of the TV network and 166 nodes of the Valcourt network were selected as contamination sites, representing in each case a set of equally probable contamination scenarios. Valcourt contamination sites included hydrants and pipe intersections that were not located in dead ends (dead-end lines, dead-end stems or dead-end loops), and consumption points. This selection reduced the size of the optimization problem, while retaining a realistic set of scenarios. Owing to the smaller number of nodes of the TV network, size reduction was judged worthless.

Results for detector locations are indicated by the squares in Figs. 1 and 2. One particular aspect of the Valcourt network's results is that 12 out of 20 detectors (60%) are located in dead ends. This is an artefact of the way residential water uses were assigned to the model's nodes. As these demands were not metered, the total water use was inferred from the mean daily water production. Aggregated values were then assigned to the central node of corresponding tributary areas. In spite of that, sensor placement results still reflect protection of water consumers, as detectors are located in areas where consumption (mainly residential) is of greater magnitude.

In both experimental DWDSs, some of the assumed contamination scenarios remain undetected by the sensor networks, either because no detector is reachable from these nodes or because accessible detectors cannot be reached within 24 h. For the Valcourt and TV networks, respectively, 11 and 59 scenarios are undetected, leaving 2 and 14% off the mean daily consumptions unprotected. Among the uncovered scenarios for the TV network, 52 (88%) would occur at dead-end nodes, of which 26 have no assigned consumption (such scenarios were not included in Valcourt detector location optimization) and 25 have relatively low consumptions.

Table 1. Contaminated and Isolated Areas' Characteristics

	Pipe length (percent of total length)			
Instances	Contaminated (%)	Isolated (%)	Increase factor	Valves
Valcourt				
D-18	7.6	11.3	1.5	2
D-31	27.2	54.0	2.0	3
D-51	4.3	14.8	3.4	7
D-105	6.4	13.4	2.1	4
D-112	17.5	27.2	1.6	5
D-131	10.7	27.8	2.6	7
D-133	11.7	25.5	2.2	6
D-142	10.5	25.5	2.4	6
D-159	30.0	55.2	1.8	9
D-279	13.6	27.8	2.0	6
D-363	10.0	18.8	1.9	5
D-368	24.2	40.0	1.7	5
D-375	11.7	25.7	2.2	6
Mean	14.3	28.2	2.1	5.5
TV				
D-030	8.3	21.6	2.6	4
D-033	3.4	10.6	3.1	4
D-039	8.5	31.0	3.6	7
D-046	7.5	24.4	3.3	11
D-048	9.0	52.6	5.9	10
D-049	7.8	44.6	5.8	10
D-051	13.2	59.4	4.5	8
D-057	27.6	80.5	2.9	5
D-060	3.0	27.8	9.3	7
D-067	15.4	24.3	1.6	4
D-102	16.9	40.3	2.4	11
D-192	9.8	24.9	2.5	2
D-201	14.6	46.7	3.2	10
Mean	11.2	37.6	3.9	7.2

Application of Isolation Procedure

General Results

In each experimental network the number of potentially contaminated area instances was equal to the number of located detectors. Thirteen instances were analyzed for each network, as in the remaining cases the contamination extent was restricted to dead-end areas. The extent of contaminated and isolated areas in terms of pipe length is presented in Table 1 for each one of the 26 analyzed instances, as well as the number of valves required for isolation. An increase factor is also presented in Table 1, as the ratio of isolated pipe length relative to the corresponding contaminated pipe length.

Specific Results: Description of IP

To help the reader better understand the application of the IP as exposed in this subsection, some definitions are given hereafter.

1. Contaminated area or contamination: all the pipes enclosed

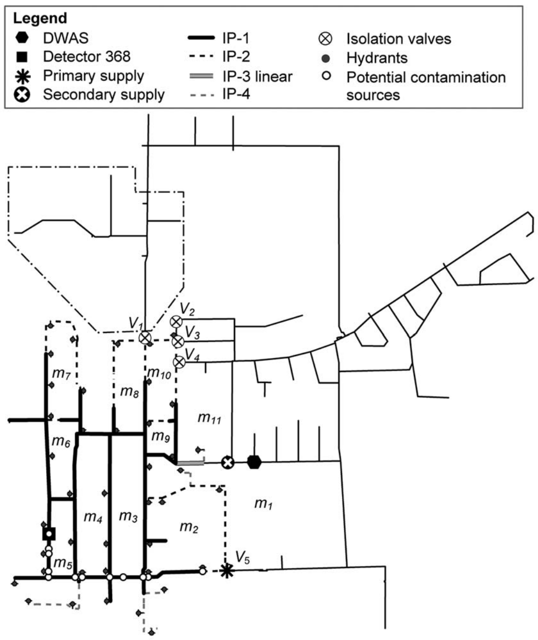

Fig. 5. Application of IP, first detection at D-368, Valcourt network

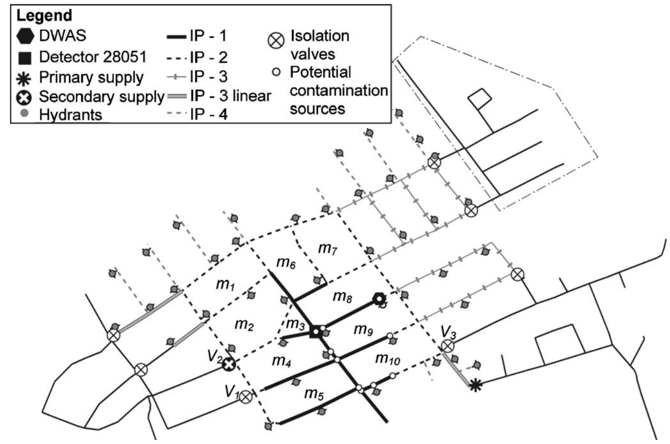

Fig. 6. Application of IP, first detection at D-051, TV network

in "worst-case" contaminated areas 2 h following detection (assumed response delay);

2. Pipe links' state: as pipe links are added to the area to be isolated, their state is changed from "not included" to "included." The same applies to the valve links selected for isolation;
3. Adjacent links: the IP is based on adjacency between links, i.e., connected to a common node;
4. Security margin: pipe "distance" left between a potentially contaminated link and a valve selected for isolation since: (1) contamination keeps spreading while response teams are on their way to shut down isolations valves; and (2) whether the flow is directed towards a valve selected for isolation or not, a safety factor is added. The security margin is expressed in terms of a pipe link or segment in which the travel time must be at least the value of the manipulation delay ("Heuristic Isolation Procedure" section). In our examples, this delay was set to 5 min, as the diameters of the valve links to likely be selected for isolation did not exceed 200 mm. Pipe links that comprise the security margin are added to the isolated area;
5. Natural loops set: the number of natural loops $\boldsymbol{l}$ in a DWDS graph representation is given by $\boldsymbol{l}=p-n+1$, with p the total number of pipe links excluding closed valves, and n the total number of nodes (including reservoirs). $\{C\}$ denotes the entire natural loop set of a given DWDS and $\{M\}$ a particular set of marked loops (step 2 of the IP). Then $|\{C\}|=\boldsymbol{l}$ and $\{M\}\subseteq\{C\}$. In Figs. 1 and 2, the natural loops are numbered according to L_i ($i=\{1,\ldots,\boldsymbol{l}\}$); and
6. Primary and secondary supplies: valve links selected as clean water supplies to the isolated area, at step 3 of the IP. The first one is left open as isolation takes place and is likely to be used as the primary water supply for subsequent flushing operations. The second one is closed and retained as a possible backup water source to also be used during flushing. Selection of these links is based on two main criteria: they should be located on maximum diameter pipes adjacent to the isolated area or as close as possible to a network water source.

Two examples were chosen to illustrate detailed applications of the heuristic IP, one for each experimental network (Figs. 5 and 6). The notation IP-x in the figures' legends refers to the steps of the IP in Fig. 4. Delineation of the area to be isolated uses as an input the "worst-case" extent of contamination 2 h following detection, which is represented by the IP-1 lines in Figs. 5 and 6. At step 2 of the procedure, the natural loops in which at least one constituting node is connected to a contaminated link are first identified. Those loops (and their constituting pipe links) are marked as part of $\{M\}$ (notation m_i, $i\in\{1,\ldots,\boldsymbol{l}\}$ in Figs. 5 and 6) and are likely to be included in the isolated area if: (1) more than 50% of their pipe length is contaminated (in such cases the inclusion is systematic, see loops m_3, m_4, m_5, m_6, m_9 in Fig. 5 and m_3, m_4, m_5,m_8, m_9, m_{10} in Fig. 6); (2) or if based on the previous loop inclusions, and taking into account the security margin, the valves located on these loops would not allow enclosing the contaminant within less than 50% of their pipe length, or worse, would not allow complete enclosing of the contaminant (see loops m_2, m_7, m_8, m_{10} in Fig. 5 and m_2, m_6 in Fig. 6); (3) or if still based on the previous loop inclusions, more than 50% of their pipe length is either contaminated or already included (see loops m_1, m_7 in Fig. 6). After each loop inclusion, pipe links state information is updated. For the sake of limiting the loop inclusion process and hence the extent of isolated areas, at this stage only the loops in $\{M\}$ may be included if they meet any one of the three conditions mentioned above. Outcome from step 2 of the procedure (i.e., included pipe links) is represented by lines IP-2 in Figs. 5 and 6. As can be seen from these figures, all the loops in $\{M\}$ were included in the TV example, but not in the Valcourt example (where m_a, and m_{11} were not included.

Step 3 of the IP seeks to identify valves required to isolate the looped network area delineated at step 2. This is accomplished by a progression of rules to which the security margin still applies. Starting from each network junction between the included loops and adjacent pipe links not yet included (without paying attention to dead-end lines or stems which are treated at step 4): (1) find valve links that are connected to ideally all or at least some of these junctions (see valves V_1, V_2, V_3, V_4, V_5, in Fig. 5 and V_1, V_2, V_3 in Fig. 6); (2) at every remaining junction try to find a single valve that would create an adjacent linear pipes section (as represented by the dashed line in Fig. 3); (3) based on the previously mentioned criteria, find two supply links adjacent to the included area and identify, if required, additional valves necessary to create a linear pipe section (as represented by the doted line in Fig. 3). The supply links may be among the valve links already identified in (1) and (2). In Figs. 5 and 6, the pipes included after applying rules (2) and (3) are represented by lines IP-3 linear. In the Valcourt example, the first and second supply links were selected based on a diameter of 400 mm and proximity to a network sourse, respectively (Fig. 5). In the TV example, both supply links were selected because of their diameters of 250 and 200 mm (Fig. 6). (4) at this stage, as was the case for the TV example (Fig. 6), there may still remain junctions where no isolation valve was found. When the links adjacent to these junctions belong to natural loops (that obviously were not yet completely included), this indicates that the valve configuration of these loops does not allow meeting the conditions for rules (1)–(3). Taking these loops one by one, when a junction without an isolation valve is found, two adjacent valve links are sought in an ultimate attempt to create a linear pipe section. If the search in unsuccessful, then as indicated by lines IP-3 in Fig. 6, the loops have to be included. This fourth rule of step 3 is iterative and after each resulting loop inclusion, rules (1)–(3) are applied again. In the example shown in Fig. 6, three iterations were required. After each one, in that specific case, the supply links remained the same.

Step 4 of the IP is divided in two substages. To this point, the IP has focused on finding two supply links and the valves required to isolate a looped network area. Although this is not the case in the examples of Figs. 5 and 6, some linear or branched pipe sections adjacent to the previously delineated area may still require isolation. Taking into account the security margin, the first substage of step 4 then finds the required isolation valves. In the second substage, existing dead-end lines or stems (see dash-dot line in Fig. 3) adjacent to the included pipe links are added (line IP-4 in Figs. 5 and 6). Isolating these dead ends from water networks does not pragmatically justify additional valve closures, and besides, as shown in Figs. 1 and 2, some of these areas do not even contain a single valve. A condition could be added to the procedure, though, to require isolation of dead ends in which critical users such as hospitals or schools would be located.

At step 5 of the IP, a preliminary isolation solution is obtained, and the required number of response teams is computed. In the examples of Figs. 5 and 6, respectively, five and eight response teams would have to be sent in the field.

Steps 6 and 7 serve to determine whether response teams can reach the isolation valves and manipulate these simultaneously before the contamination actually reaches the valves. If not, the IP is applied iteratively until this condition is met. In the examples of Figs. 5 and 6, the delay condition was met at first iteration. Among all the 26 instances analyzed in this study, only Valcourt D-51 required a second iteration and this essentially has to do with the value of the security margin. When this value is set accordingly, clearly the number of required iterations is reduced. Owing to the relatively small sizes of the Valcourt and TV networks, the operational delays are almost always dominated by their fixed-manipulation part. Even at a mean velocity of 30 km/h, the displacement times are mostly shorter than the 5-min manipulation delay. Poulin et al. (2006) presented the case of Valcourt D-159 and showed it was constrained by the maximum operational delay, which is associated with the farthest isolation valve from the DWAS. In this previous paper, the security margin was set to a distance of one pipe link, independently of water travel times. A second iteration was then needed to obtain the final solution. Setting the security margin to the value of the manipulation delay (expressed in terms of water travel time in pipes), in that case, produces the same result after only one iteration.

Before the final solution is obtained, step 8 of the IP identifies uncontaminated areas where water supply is interrupted as a consequence of isolation (dash-dot lines in Figs. 5 and 6). At least consumers located in these areas cannot be reached by contamination. Having this information readily accessible may guide utility managers in deciding how to treat these out-of-supply areas.

Discussion

Modeling Assumptions

The methodology and results presented in this paper are dependent on a set of simplifying assumptions regarding the modeling framework and the number of contaminant detectors.

1. The assumed static behavior of the DWDSs is a strong simplification. The assumption of a single mean daily flow pattern most importantly affects the definition of the "worst-case contaminant spread" at step 1 of the IP and the subsequent spread during isolation operations, as temporal variations of flow are not taken into account. A first straightforward alternative would be to consider a set of at least four consecutive flow patterns occurring over 1 day and still use Watson et al. (2004)'s steady-state sensor placement model. Isolation strategies would then be defined for each flow pattern. The contaminant spread would be delineated considering all the potential contamination sources in the most realistic flow pattern active at detection. To a certain extent, the single flow pattern assumed in our study could be seen as one of several possible static patterns in a day. Even with such modifications, the basic structure of the IP would remain valid;
2. Considering the characteristics of currently available technologies, locating 15–20 contaminant detectors in DWDSs would certainly be cost prohibitive (for example, thousands of dollars per customer). But in the near future, more affordable and reliable technologies will likely and hopefully become available. Smaller (and more realistic) sensor budgets (for instance two or three) would have most likely generated contaminated areas so large that it would have been impossible to isolate the contamination within pressure zones or, worse, at all;
3. Valves reliability (operability, watertightness) is another important issue. As opposed to the assumption made in this study, reliability is generally smaller than 100%. Various sets of constraints on valves availability could be added to the IP in order to assess how the extent of isolated areas would be modified. The approach introduced in this paper may help identify a set of critical valves requiring closer maintenance; and
4. The current version of the IP assumes all-pipe network models in which the location of every valve and hydrant is available. To treat larger-scale skeletonized models, the IP would need to be adapted. In such cases, for instance, the IP could produce output that includes virtual valves to be mapped by utilities to the best available real response actions.

The above mentioned assumptions allowed a first validation of the IP, in a simple and suitable modeling framework. All the related crucial issues must be further examined in future works.

Delays and Emergency Response Capabilities

The response delay refers to the time between first contaminant detection and initiation of isolation operations. It was set to 2 h assuming this time period is allowed for proper organization of operational responses (and for parallel contamination investigation/confirmation). As for the operational delay, it refers to the total time needed for a response team to reach a valve from the DWAS (displacement time) and shut it down (manipulation time). In this study, a velocity of 30 km/h is assumed for the response teams' displacements. Simultaneous field operations are also assumed and, as mentioned in the "Specific Results" section, delineation of an isolation area is constrained by the maximum operational delay (farthest valve from DWAS), which in fact represents the total time for all operations to be completed. The assumptions related to both response and operational delays contribute the effectiveness of the IP.

First, the choice of the response delay value may vary from one application instance to another. Of course, as indicated by Murray et al. (2006), longer response times reduce the effectiveness of CWS technologies and response strategies. Utilities should seek to minimize this delay by enhancing emergency response capabilities, including sound maintenance practices, response team training, and coordination. Second, the assumption of simultaneous manipulations implies that a sufficient number of response teams are available. This reduces the operational time but, on the other hand, may seem rather optimistic, especially when it comes to larger utilities. The numbers of isolation valves to be closed remain reasonably low in all Valcourt and TV instances (Table 1), and this of course has to do with the small sizes of these networks. In much larger networks, with for instance more than 10,000 pipes and many parallel pipes, the number of valves required for isolation may become too important for simultaneous shutoffs. Conducting valve closures in separate steps may result in creating high velocity contaminant transport paths, which is undesirable. Our approach is not currently adapted for such applications. Eventually, it could rely on the installation of remotely controlled valves for easier simultaneous shutoffs.

The delay and response-capabilities-related issues will need to be examined, to further validate the approaches' feasibility in other DWDSs. At least, applying the strategy as proposed in this paper can help assess required emergency response needs and upgrades.

Isolated Areas' Characteristics

As can be seen from Table 1, the number of valves required for isolation is generally higher for TV instances, as compared to Valcourt. This is due to the higher density of natural loops per pipe length unit in the TV network. In most TV instances, there were a higher number of pipes adjacent to the included loops and where valve shutoff was required, unless the contaminated area affects peripheral loops of the network.

Table 1 also reveals that the increase factor is generally much higher for TV instances than for Valcourt instances, even though the mean extent of contaminated areas (% of total pipe length) is lower for the TV network. This has to do with the valve pattern of each network. As a first observation, the total number of valves relative to the total number of pipe links is lower for the TV network (0.36) compared to the Valcourt network (0.53). Also, the general rule of thumb where every n-pipe intersection should contain at least $n-1$ valves is less frequently observed in the TV network. These two characteristics of TV's valve configuration contribute to extending the isolated areas (see lines IP-3 Fig. 6). Ideally, the number of natural loops included in an isolated area should be smaller than or equal to the number of elements in $\{M\}$ (lines IP-2 in Figs. 5 and 6). This was observed in all 13 Valcourt instances, but in only four of the 13 TV instances. From a design point of view, results similar to those obtained for the TV network could be an indicator of deficient valve patterns. Again, such results could justify further improvements on water networks, to reduce health risks related to contamination. Based on the characteristics discussed above, overall Valcourt results appear to be more interesting than TV results and this is also reflected in Figs. 5 and 6. Nonetheless, as stated in ASCE (2004), no water system will be perfect in regards to isolating and flushing contaminated water without further impact to users that would not be directly affected by contamination.

Fire Protection

The IP presented in this paper does not allude to fire protection but obviously this issue cannot be completely ignored. Should a fire occur during a contamination event, the simultaneous occurrence of both incidents will most often result in worsening the emergency situation, whether part of a DWDS is isolated or not. First, as already mentioned, if the contaminant is not properly isolated, hydrant opening may contribute to disseminating it to unaffected areas of a DWDS. Second, if a fire occurs close to or inside the boundaries of a contaminated area, whether isolation has taken place or not, relying on the DWDS for extinction implies that contaminated water will likely have to be used. As a possible solution, utilities should be aware of that issue and plan for alternative fire protection as part of a more global emergency response plan.

Summary and Conclusions

In this paper heuristic procedure was introduced systematize the isolation of contaminated areas within the pressure zones of a DWDS, based on a set of pragmatic and operational rules. The procedure is motivated by the idea that physical containment of detected contamination in a DWDS should be undertaken as a first operational response, as quickly as possible, to prevent further spreading and hence better protect water consumers. Once the contaminant is contained, subsequent operations seeking to evacuate it (through hydrant flushing) may be better managed. The approach relies on an optimally located CWS which allows identification of a set of possible contamination sources upon fist detection and delineation of a "worst-case" potentially contaminated area. A response delay is taken into account, during which the contaminant spread is computed. The IP identifies a set of valves and ensures that these can be shut down by response teams sent in the field as soon as possible and within contaminant spreading delays.

As a first step towards a more comprehensive computerized application, the IP was developed and validated using two different real-world DWDSs. Depending on each network's design and configuration, application of the IP will result in different isolation strategies. In a less densely looped and more appropriately valved network the number of required isolation valves tends to be lower and the extent of isolated areas relative to the corresponding contaminated areas is generally smaller.

The approach should be further validated using other network instances with different configurations and under dynamic flow conditions. Despite strong assumptions as to the required sensor densities and simultaneous valve shutoffs, the IP as it stands can help assess necessary resources to better protect drinking water consumers and reduce health risks, from an operational response point of view and even from a design improvement point of view.

Acknowledgments

This work was funded by a scholarship from National Science and Engineering Research Council (NSERC) of Canada and Aqua Data Inc., a consulting firm in computerized diagnosis of water distribution and wastewater collection systems. The writers also wish to thank the reviewers for their insightful and valuable comments.

References

Allman, T. P., and Carlson, K. H. (2005). "Modeling intentional distribution system contamination and detection." *J. Am. Water Works Assoc.*, 97(1), 58–91.

Aqua Data. (2005). "Aqua Cad Suite software." ⟨http://www.aquadata.com/⟩ (December 10, 2007).

ASCE. (2004). *Interim voluntary guidelines for designing an online contaminant monitoring system*, Reston, Va.

Baranowski, T. M. (2007). "Development of consequence management strategies for water distribution systems." Ph.D. thesis, Vanderbilt Univ., Nashville, Tenn.

Bell, R. N., et al. (2004). "Interim voluntary security guidance for water utilities." American Water Works Association, Denver.

Berry, J., Fleisher, L., Hart, W. E., Phillips, C. A., and Watson, J.-P. (2005). "Sensor placement in municipal water networks." *J. Water Resour. Plann. Manage.*, 131(3), 237–243.

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 methods." *J. Water Resour. Plann. Manage.*, 132(4), 218–224.

Craun, G. F., and Calderon, R. L. (2001). "Waterborne disease outbreaks caused by distribution system deficiencies." *J. Am. Water Works Assoc.*, 93(9), 64–75.

Craun, M. F., Craun, G. F., Calderon, R. L., and Beach, M. J. (2006). "Waterborne outbreaks reported in the United States." *J. Water Health*, 4(2), 19–30.

Epp, R., and Fowler, A. G. (1970). "Efficient code for steady-state flows in networks." *J. Hydr. Div.*, 96(1), 43–56.

GNU–GNU's Not Unix. (2006). "GLPK–GNU linear programming kit." ⟨http://www.gnu.org/software/glpk/⟩ (December 10, 2007).

Guan, J., Aral, M., Maslia, M. L., and Grayman, W. M. (2006). "Identification of contaminant sources in water distribution systems using simulation–optimization method: Case study." *J. Water Resour. Plann. Manage.*, 132(4), 252–262.

Hart, W. E., Berry, J. W., Murray, R., Phillips, C. A., Riesen, L. A., and Watson, J.-P. (2007). "SPOT: A sensor placement optimization toolkit for drinking water contamination warning system design." *Proc., World Environmental and Water Resources Congress: Restoring Our Natural Habitat*, Environment and Water Resources Institute, Tampa, Fla.

Herrick, C., Pratt, J., Raucher, R., and Kalas-Adams, N. (2006). "Emergency response and recovery planning for water systems: A kit of tools." *Research Rep. No. 91097f*, American Water Works Association Research Foundation, Denver.

Hill, J., van Bloemen Waanders, B., and Laird, C. (2006). "Source inversion with uncertain sensor measurements." *Proc., Water Distribution System Analysis Symp.*, ASCE, Cincinnati.

Kessler, A., Ostfeld, A., and Sinai, G. (1998). "Detecting accidental contaminations in municipal water networks." *J. Water Resour. Plann. Manage.*, 124(4), 192–198.

Laird, C. D., Biegler, L. T., and van Bloemen Waanders, B. G. (2006). "Mixed-integer approach for obtaining unique solutions in source inversion of water networks." *J. Water Resour. Plann. Manage.*, 132(4), 242–251.

Murray, R., Uber, J., and Janke, R. (2006). "Model for estimating acute health impacts from consumption of contaminated drinking water." *J. Water Resour. Plann. Manage.*, 132(4), 293–299.

Ostfeld, A., and Salomons, E. (2004). "Optimal layout of early warning detection stations for water distribution systems security." *J. Water Resour. Plann. Manage.*, 130(5), 377–385.

Ostfeld, A., and Salomons, E. (2005). "Securing water distribution systems using online contamination monitoring." *J. Water Resour. Plann. Manage.*, 131(5), 402–405.

Poulin, A., Mailhot, A., Grondin, P., Delorme, L., and Villeneuve, J.-P. (2006). "Optimization of operational response to contamination in water networks." *Proc., Water Distribution System Analysis Symp.*, ASCE, Cincinnati.

Preis, A., Mayorchik, Y., and Ostfeld, A. (2007a). "Multiobjective contaminant response model." *Proc., World Environmental and Water Resources Congress: Restoring Our Natural Habitat*, Environment and Water Resources Institute, Tampa, Fla.

Preis, A., Ostfeld, A., and Perelman, L. (2007b). "Contamination source detection with fuzzy sensors data." *Proc., World Environmental and Water Resources Congress: Restoring Our Natural Habitat*, Environment and Water Resources Institute, Tampa, Fla.

Propato, M. (2006). "Contamination warning in water networks: General mixed-integer linear models for sensor location design." *J. Water Resour. Plann. Manage.*, 132(4), 225–233.

Propato, M., Cheung, P. B., and Piller, O. (2006). "Sensor location design for contaminant source identification in water distribution systems." *Proc., Water Distribution System Analysis Symp.*, ASCE, Reston. Va.

Rahal, H. (1995). "A cotree flows formulation for steady state in water distribution networks." *Adv. Eng. Software*, 22(3), 169–178.

Tryby, M. E., Propato, M., and Ranjithan, R. (2007). "Monitoring sensor network design for water distribution source inversion problems." *Proc., World Environmental and Water Resources Congress: Restoring Our Natural Habitat*, Environment and Water Resources Institute, Tampa, Fla.

USEPA. (2003). "Response protocol toolbox: Planning for and responding to drinking water contamination threats and incidents—Overview and application." *Report EPA 817-03-007 Interim Final—December 2003*, Washington, D.C.

USEPA. (2004a). "Response protocol toolbox: Planning for and responding to drinking water contamination threats and incidents—Module 5: Public health response guide." *Report EPA 817-03-005 Interim Final—April 2004*, Washington, D.C.

USEPA. (2004b). "Response protocol toolbox: Planning for and responding to drinking water contamination threats and incidents—Module 6: Remediation and recovery guide." *Report EPA 817-03-006 Interim final—April 2004*, Washington, D.C.

USEPA. (2005). "Watersentinel system architecture." *Draft Report 12105*, Washington, D.C.

Walski, T. M., Chase, D. V., Savic, D. A., Grayman, W. M., Bockwith, S., and Koelle, E. (2003). *Advanced distribution system modeling and management*, Haestad Press, Waterbury, Conn.

Watson, J.-P., Greenberg, H. J., and Hart, W. E. (2004). "A multiple-objective analysis of sensor placement optimization in water networks." *Proc., World Water and Environmental Resources Congress*, Environment and Water Resources Institute, Salt Lake City, Utah.

Wood, D. J., and Rayes, A. M. (1981). "Reliability of algorithms for pipe network analysis." *J. Hydr. Div.*, 107(10), 1145–1161.