SOURCE INVERSION WITH UNCERTAIN SENSOR MEASUREMENTS

J. Hill
Sandia National Laboratories *
Albuquerque, NM
jhill@sandia.gov

B. van Bloemen Waanders
Sandia National Laboratories
Albuquerque, NM
bartv@sandia.gov

C. Laird
Carnegie Mellon University
Pittsburgh, PA
claird@andrew.cmu.edu

Abstract

Considerable effort is underway to develop mathematical algorithms in support of contamination events in water distribution systems. In particular, least squares formulations are applied to early warning systems, minimizing the difference between measured and simulated concentration predictions, in an attempt to locate the original injection location along with the complete time profile and magnitude of the injection. Previous algorithmic development assumed sensors could provide unbiased measurements of the concentration. However, based on the latest in sensor technology, it is likely that Boolean type sensors would be more appropriate. These sensors only indicate the presence of contaminant, or more correctly, if the contaminant concentration is larger than some detection threshold or not. In this paper, we demonstrate the robustness of the inversion algorithm by comparing solutions that use both continuous readings and Boolean measurements. Although the inversion solution reports compromised magnitudes, the location can be determined with surprisingly high quality when Boolean values are used in the least squares formulation. Essentially, while the Boolean measurement formulation fails to predict the injection magnitude, sufficient spatial information is provided to recover the injection location. This is due to the uncoupled behavior between the network time delays (which are determined by the network topology and the flow patterns) and the magnitude of the concentrations (determined by the water quality calculations). Furthermore, the presented approach reliably determines the injection location with different sensor threshold values and we show the inversion solution for both continuous and Boolean measurements.

Keywords

Water quality simulation, source detection, inversion, imperfect sensors,

1 INTRODUCTION

The risk of contamination of water distribution systems through accidental or intentional attacks has motivated the investigation of alternative protection methods, including the use of numerical simulation and

*Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000

optimization algorithms. Drinking water networks are especially vulnerable to biological and chemical attack due to the large land area encompassed by the network and the number of access points. Any water outlet, such as a fire hydrant or even a household water faucet, can be an access point for backflow contamination into the network. As an alternative to physical security alone, sensors could be installed in the network to detect contaminants. Because sensors are costly to purchase and maintain, and can only be installed in accessible parts of a network, complete detection coverage will not be possible. Furthermore, the number of possible injection locations is very large (in this work we assume that every network node is a candidate injection location). Therefore an early warning system consisting of a sparse set of sensors will have limited capabilities to accurately characterize the effected area. In an attempt to improve this characterization, mathematical algorithms could be applied to this problem and one possibility is to use optimization to reconcile the difference between observed concentration data and simulated predicted values. The general idea is to manipulate source terms in a network model until calculated predictions match measurements. Although optimization algorithms can be used to automate this process, the large number of solution possibilities combined with relatively small number of measurements cause several technical difficulties including non-uniqueness in the solution and large computational requirements. These issues have been addressed in recent work in which regularized least squares formulations constrained by chemical transport dynamics were applied to identify source locations from sparse sensor measurements [6, 13]. The constrained least squares techniques demonstrate a robust and computationally efficient inversion capability by inverting complex attack scenarios such as short injection pulses, simultaneous intrusions and the use of small numbers of sensors. As such, source inversion could be a viable and practical technique to augment early warning systems.

To enable this technology in real distribution systems, however, several sources of uncertainty must be addressed including model and solution inaccuracies, measurement errors, calibration problems, and discretization inexactness. As well, there is inherent uncertainty in the true injection location due to the network topology and the sparseness of the sensor grid. Non-unique solutions are unavoidable when the limited sensor information is insufficient to distinguish neighboring nodes. To compensate for the sparsity of observations and the difficulty of non-unique solutions, a regularization term is added which results in a single unique solution. This inversion solution finds the region of interest and typically includes nodes neighboring the true injection as part of the solution. Several researchers have addressed uncertainty in the context of inversion problems for water distribution systems. McKenna et al. tested the robustness of source inversion algorithms by introducing uncertainty in a temporal aggregation of concentrations [8]. Even with coarse time discretizations, the inversion algorithm was able to produce accurate determinations of the injection source. An optimization-under-uncertainty approach has been used to solve an inverse problem that considered uncertain network flow patterns [2]. Laird et al. also developed a mixed integer post-processing approach to refine the regularized solution and remove solution profiles that had only a small effect on the least squares objective [4]. Not only can some of the solution uncertainty be reduced by this technique, but the approach was effective in uniquely identifying the location of simultaneous attack scenarios. These previous studies suggested a general robustness quality of the source inversion formulation and provided motivation to further expose the source inversion algorithm to other forms of uncertainty, in particular in the context of imperfect sensor information. Thus far, inversion algorithms have assumed continuous and accurate concentration measurements, whereas in reality current sensor technology has not reached this level of maturity. At best, affordable sensors for water distribution systems are capable of measuring water quality parameters and, through special algorithms using changes in multiple water quality parameters, can report the presence of contaminant at specific spatial locations. In this paper, we investigate the ability of the source inversion algorithms to invert for source terms with the most basic sensor functionality capable of

only reporting a Boolean value indicating the presence of contaminant (above a certain detection threshold).

This paper is organized as follows. Section 2 presents additional background information, outlines the optimization formulation of the inversion problem and describes the implementation strategy. Section 3 demonstrates the effectiveness of this formulation for a illustrative municipal water distribution network and characterizes the effects of the uncertainty of the sensor errors on the inverted solution. Section 4 summarizes our findings and discusses potential future work.

2 BACKGROUND, PROBLEM FORMULATION AND IMPLEMENTATION

Inverse problems are a well studied subject and because these problems are inexact in nature, theory and methods have naturally evolved in the statistical and probability communities [1, 12, 14], in which conditional probability methods, Bayes theorem, and filtering algorithms have been developed. These methods offer rigorous approaches to handling general uncertainty for inverse problems. Here, we focus on evaluating the robustness of the source inversion algorithm rather than attempting to characterize the uncertainty associated with the output (i.e. inversion solution). The efficient solution of inverse problems for complex dynamics poses numerous technical challenges, especially when a large number of inversion parameters (the unknown source terms) are considered and when the system can only provide sparse sensor data. First, the inversion space is very large, producing an optimization problem with a large number of degrees of freedom, which requires the implementation of large-scale constrained optimization algorithms. In these methods, the state (e.g. concentrations) and inversion variables are solved simultaneously for feasibility and optimality. Second, a regularization term needs to be added to the objective function to enforce solution uniqueness. This is an underdetermined problem because of the sparse sensor observations in comparison to the solution space and consequently many solutions are possible. The regularization forces a single unique solution to the problem, that includes a family of candidate injection profiles. This set of candidate injection profiles was effectively managed by a post-processing step [4]. Also, the value of the regularization parameter may have an effect on the conditioning of the resulting linear system. Third, inequality constraints typically need to be imposed on the inversion parameters to ensure a physical solution. In the case of our source inversion problem, positive injection concentrations are required. The solution of a optimization problem with inequality constraints requires the use of special algorithms such as active set strategies, which can become expensive with a large number of variable bounds or interior point methods. Fourth, computational performance is critical for inversion of contamination events. It is important to solve this problem sufficiently fast so that, in practice, decisions can be made on an appropriate time scale. The computational issues associated with large numbers of inversion parameters and complex dynamics requires a special simultaneous solution process whereby the forward problem is included as constraints in the optimization problem. With this approach, the forward problem is converged only once, at the solution of the optimization algorithm. In addition, a special sub-domain technique was developed to handle large datasets [5]. Finally, most inversion problems need to address multiple minima, which can be managed with continuation and multi-level techniques. In our case, the regularized problem forms a convex quadratic programming problem with a single global solution. The unregularized problem is a convex (but not strictly convex) optimization problem, implying a single optimal objective value, but the possibility of non-unique solutions in the variable space.

These technical issues have been addressed in previous work and numerical studies have demonstrated our ability to handle large inversion spaces, refinement of regularized solutions, and many inequality constraints using computationally efficient solution techniques. The next sections discuss the general mathe-

matical formulation that was used in previous studies to implementation strategy. The reader is referred to previous papers for a comprehensive review of all the numerical details [3–6, 13].

2.1 Mathematical Formulation

Water distribution systems are often described by a network of links and nodes, where links represent pipes, pumps, or valves, and nodes represent sources, tanks, or junctions [9]. When simulating these systems, it is assumed that the low concentrations of secondary species in the network do not dramatically affect fluid flow properties. Thus, the hydraulic calculations and water quality calculations can be decoupled; following dynamic hydraulic calculations, the resulting flow profiles are specified as known inputs to the water quality model. The water quality solution can then be determined by a variety of existing techniques [7, 10, 11]. Assuming that contaminant can be injected at any network node, the goal of this source inversion approach is to use information from a limited number of sensors to calculate injection times and locations. Identifying the contamination source enables security or utilities personnel to stop the contamination and propose effective containment and cleanup strategies.

This problem can be formulated as a dynamic optimization problem which seeks to find the complete time profiles of the unknown injections (at every network node) that minimize the least-squares error between the measured and calculated network concentrations. For the water quality model, the pipes are modeled with partial differential equations in time and space, and the network nodes are modeled with dynamic mass balances that assume complete mixing. This produces an infinite dimensional optimization problem that can be discretized to form a large-scale algebraic problem. A naive approach requires a discretization in both time and space due to the constraints from the pipe model. Instead, following the common assumption of plug flow in the main distribution lines, we developed an *origin tracking algorithm* which uses the network flow values to precalculate the time delays across each of the network pipes. This leads to a set of algebraic expressions that describe the time-dependent relationship between the ends of each of the pipes. The entire infinite dimensional optimization problem is then discretized using a direct simultaneous approach. This produces a large-scale algebraic nonlinear program that is solved efficiently using the interior point solver, IPOPT [15]. In previous work, this large-scale source inversion problem had over 210000 variables and 45000 degrees of freedom. Nevertheless, solutions were possible in under 2 CPU minutes on a 1.8 GHz Pentium 4 machine and the approach was very successful at determining the contamination source [6].

The completely discretized dynamic optimization problem can be written as follows:

$$\min_{\bar{c},c,m} \quad f = \frac{1}{2}\left[c - c^{\star}\right]^T W \left[c - c^{\star}\right] + \frac{\rho}{2} m^T m \tag{1}$$

$$\text{s.t.} \qquad \bar{c} - Pc = 0, \tag{2}$$

$$\bar{N}\bar{c} + Nc + Mm = 0, \tag{3}$$

where $\bar{c}=\left[\cdots \bar{c}^I_{ij}, \bar{c}^O_{ij} \cdots\right], \forall i \in \mathcal{P}, j \in \Theta$ is a vector of pipe concentrations for the inlet ($\mathcal{I}$) and outlet ($\mathcal{O}$) of every pipe, discretized in time, $c=\left[\cdots c_{ij} \cdots\right], \forall i \in \mathcal{N}, j \in \Theta$ is the vector of calculated concentrations for every node at every time discretization, and $m=\left[\cdots m_{ij} \cdots\right], \forall i \in \mathcal{N}, j \in \Theta$ is the vector of unknown contaminant mass injections for every node at every time discretization. The matrix P is defined by the origin tracking algorithm, and the matrices $\bar{N}$, N, and M are the Jacobians of the discretized mass balance constraints (for junctions and storage tanks) with respect to the pipe concentrations, node concentrations,

and unknown injections, respectively. The function f is a regularized weighted least squares objective where $c^\star$ are the measured concentrations and the diagonal matrix W is a flow-based weighting matrix for the least squares errors, with nonzero elements only for sensor nodes at sample times. The details of this discretization are presented in [6]. It is important to note that this is not a purely topological formulation and includes the effect of time-varying network flows. The concentration and injection vectors are discretized for each node and each timestep. Differential tank equations are present in discretized form.

The second term in the objective is a regularization term that forces an unique solution to the problem. The QP solution is a global minimum for the objective value, but the minimizer is not necessarily unique. This non-uniqueness is due to the limited number of sensors. It should be noted that in previous studies, positive inequality constraints were imposed on the inversion parameters. In this study we have removed these constraints for computational efficiency reasons but acknowledge that we have ignored some negligible negative concentration values in a small percentage of our solution vector.

A range of relatively simple solution techniques can be used to solve this optimization problem. Here, we use IPOPT [15], the primal-dual interior point solver for nonlinear problems used in previous studies. This solver is distributed through the COIN-OR foundation and project details can be found at `http://projects.coin-or.org/Ipopt`.

To test the effects of Boolean sensor measurements on the quality of the inversion solution, the least squares formulation is altered to include a translated Heaviside function of the perfect sensor measurements, or

$$\min_{\bar{c},c,m} \quad f = \frac{1}{2}\left[c - \mathcal{H}(c^\star - \gamma)\right]^T W \left[c - \mathcal{H}(c^\star - \gamma)\right] + \frac{\rho}{2} m^T m \tag{4}$$

$$\text{s.t.} \qquad \bar{c} - Pc = 0, \tag{5}$$

$$\bar{N}\bar{c} + Nc + Mm = 0, \tag{6}$$

where γ is the activation concentration threshold for the Boolean sensors. The Heaviside function $\mathcal{H}$ transforms the perfect sensor measurements into Boolean readings, or

$$\mathcal{H}(c_{ij}^\star - \gamma) = \begin{cases} 0 & \text{if } c_{ij}^\star < \gamma \\ 1 & \text{otherwise} \end{cases} \forall i \in \mathcal{N}, j \in \Theta \tag{7}$$

With this formulation, we have complete control over the accuracy of the concentration measurements and can test the inversion capabilities with varying degrees of uncertainty in these measurements.

3 NUMERICAL RESULTS

We demonstrate the effectiveness of the source detection formulation for imperfect, or Boolean, sensors on a municipal water network which contains 534 pipes and 397 nodes. Table 1 summarizes the network properties, hydraulic and water quality parameters. An intentional contamination event will most likely consist of a significant amount of harmful material injected over an extended period of time. With larger injections and additional spatial and temporal sensor information, the quality of the inversion solution should improve. In practice however, the goal will always be to minimize detection time so that containment procedures can be considered. Therefore in our numerical experiments, a difficult detection scenario is emulated by exposing the distribution system to small and time-limited pulse injections that consequently

Item	Description	Item	Description
Number of Junctions	395	Maximum Trials	40
Number of Tanks	1	Quality Analysis	Chemical
Number of Pipes	534	Total Duration	12 hrs
Number of Pumps	1	Water Quality Time Step	5 min
Number of Valves	4	Water Quality Tolerance	0.01 mg/L
Headloss Formula	Hazen-Williams	Specific Gravity	0.998
Hydraulic Timestep	15 min	Relative Kinematic Viscosity	0.98
Hydraulic Accuracy	0.001	Relative Chemical Diffusivity	1.00

Table 1: Summary of the network properties, hydraulic and chemical assumptions for the municipal water network used in the numerical studies.

are more challenging to invert. A twelve-hour window is monitored with a contamination scenario in which a one hour pulse travels through the system. Figure 1 illustrates the attack scenario and the evolution dynamics of the contaminant concentration. The links and nodes are colored coded to indicate the velocity and concentration intensity, respectively. Obviously a continuous injection of contaminant would produce a more compelling visualization of the spreading dynamics. Also, other and potentially more complicated scenarios can be devised, but we defer such variations to future work.

3.1 Comparison Between Boolean and Perfect Sensors

For fifty sensors randomly positioned in our network, we compare the source location solution predicted using both the perfect and the Boolean measurements. In these comparisons, we assume the Boolean sensors require a concentration of 10 mg/L or greater to detect the presence of a contaminant. Qualitatively, the source inversion using Boolean measurements results in a solution of similar quality when compared to a solution derived with perfect sensor measurements. Table 2 compares the inversion results for perfect concentration measurements and Boolean measurements. Reported is the percentage of the total injection that the source detection algorithm allocates to each node. In both cases, the inversion identifies the true injection location along with neighboring nodes. As well, in both cases we could refine the solution further by applying the mixed integer linear programming technique [4]. Figure 2 graphically illustrates the predicted source locations in the network for both sensor types. The intensity of the nodal coloring indicates the relative magnitude of the injection solution at that node.

There is little qualitative difference between the Boolean and perfect sensors, and only a small quantitative difference between the two. In both cases, the predictions of the primary injection nodes are within two links of the actual injection source. Comparing the time profiles of the solution at the injection node, we see that the source detection algorithm accurately characterizes the time of the injection using both Boolean sensors and perfect sensors (see Figure 3). Observing that the magnitude of the contaminant injection is not completely recovered in the source inversion, even with perfect concentration measurements, we recall that for a sparse sensor network, multiple source nodes are identified and the regularized formulation allocates a portion of the overall injection to each of these nodes.

It is very encouraging that the formulation is able to accurately identify the location and the time profile for the true contaminant injection using only Boolean measurements. As expected, the Boolean formulation

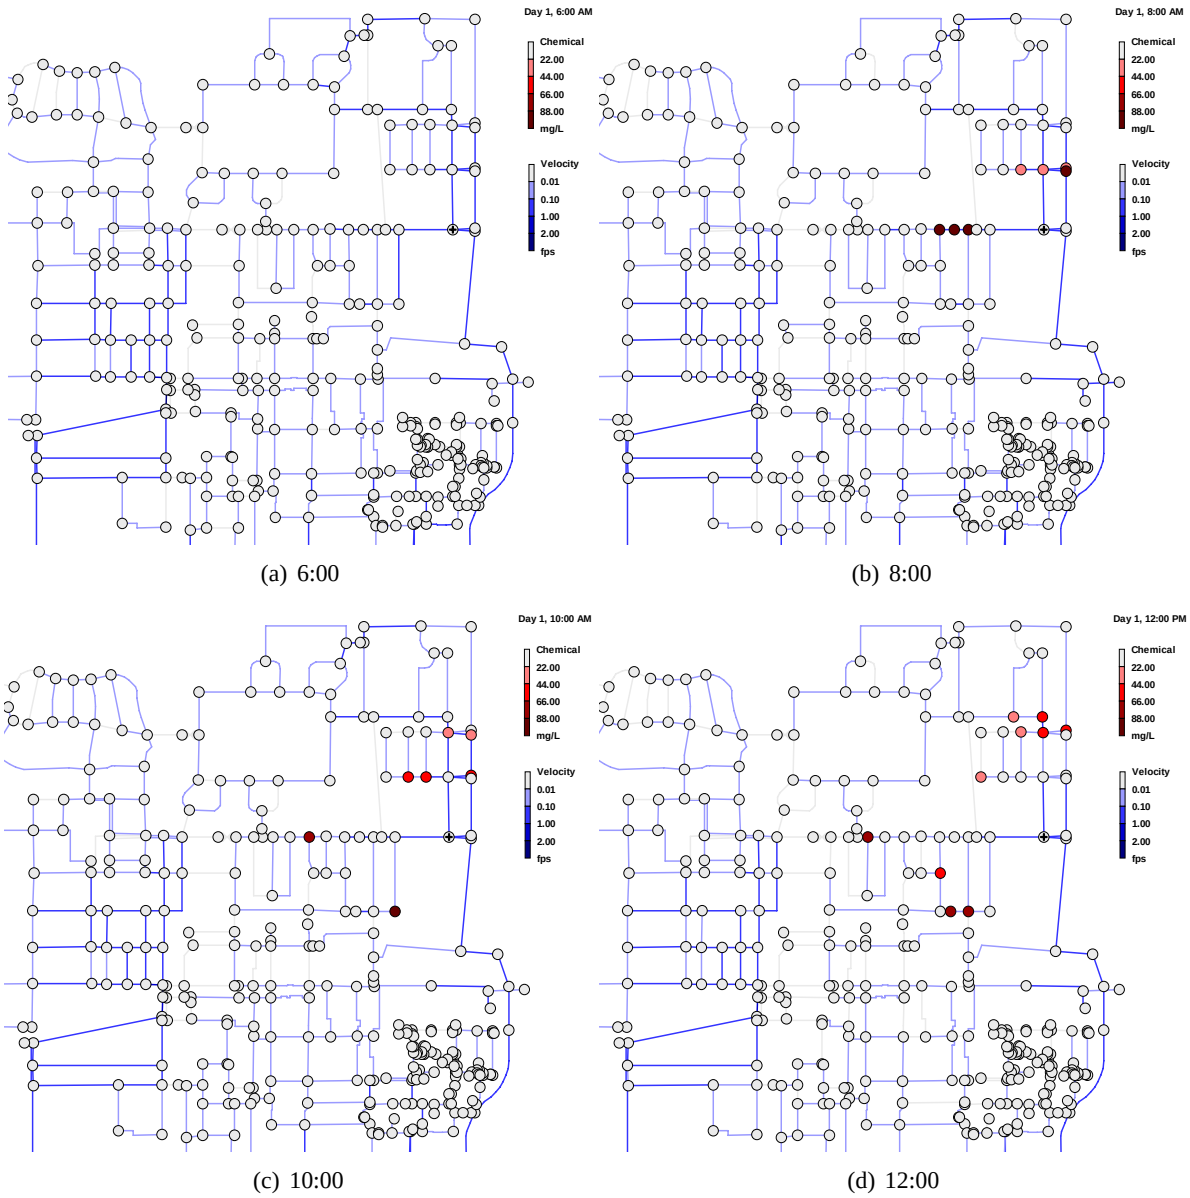

(a) 6:00 (b) 8:00

(c) 10:00 (d) 12:00

Figure 1: The municipal water distribution network that is used for all numerical experiments. This network contains 534 pipes, 395 junctions, 1 reservoir, 1 tank, 1 pump and 4 valves. A contaminant is injected into the network at the node indicated by $\oplus$ at 6:00 for one hour. The evolution of the pipe velocity profiles and junction concentrations following the injection is shown at (a) 6:00, (b) 8:00, (c) 10:00 and (d) 12:00.

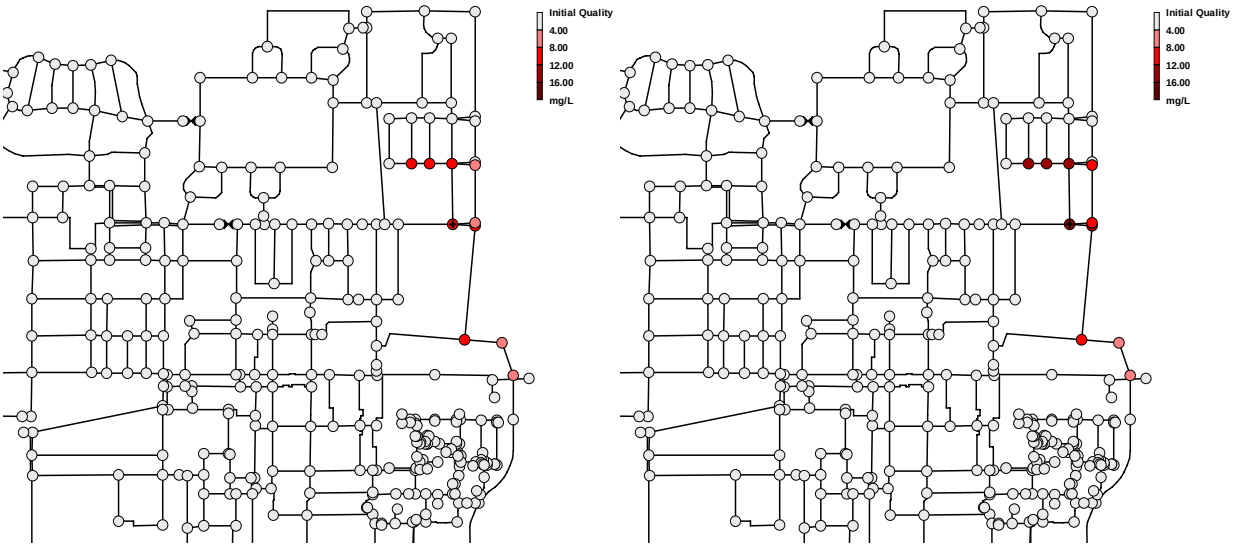

(a) Predicted Source Location with Perfect Sensors (b) Predicted Source Location with Boolean Sensors

Figure 2: Predicted source location for (a) perfect sensors and (b) Boolean sensors for a water distribution network with fifty sensors. The injection node is marked with a $\oplus$. The intensity of the nodal coloring indicates the likelihood that the node is the injection node, with greater intensity indicating a higher probability.

Node ID	Perfect Sensors	Boolean Sensors
2201	14.9 %	14.7 %
INJ	14.9 %	14.7 %
431	9.4 %	11.1 %
432	7.8 %	8.4 %
433	7.7 %	8.2 %
2401	6.3 %	7.5 %
2301	6.3 %	7.4 %
439	8.2 %	5.5 %
437	7.1 %	4.8 %
578	5.3 %	3.5 %
430	1.8 %	2.9 %

Table 2: Comparison of the predicted source location for a network of fifty sensors, both perfect and Boolean. Reported is the percentage of the total injection that the source detection algorithm allocates to each node. **INJ** indicates the actual injection node for the simulated attack.

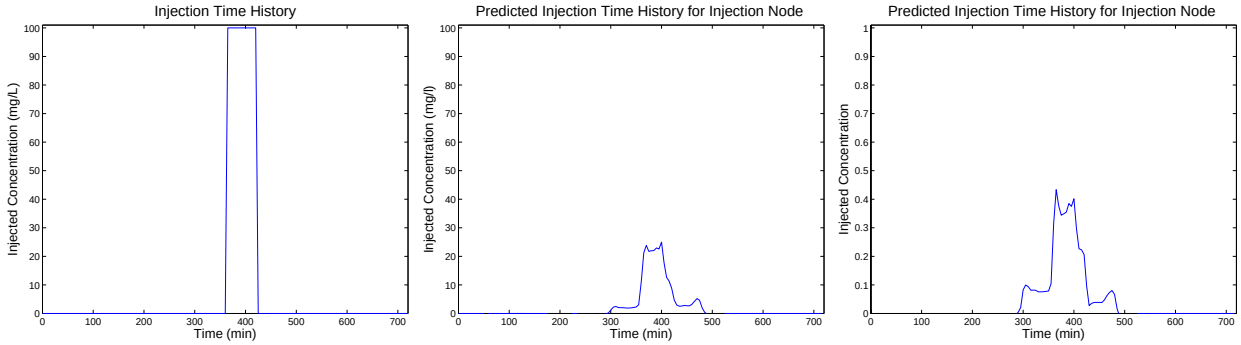

(a) Actual Contamination Event (b) Predicted Event with Perfect Sensors (c) Predicted Event with Boolean Sensors

Figure 3: The time history of a simulated contamination event at the injection node is illustrated in (a). (b) and (c) illustrate the predicted contaminant at the injection node in time for a sensor network containing fifty sensors for (b) perfect sensors and (c) Boolean sensors. As expected, the Boolean sensors cannot predict the contaminant concentration, but they do characterize the time of the contamination event compared with the perfect sensors.

is not able to accurately predict the magnitude of the injection. In addition to applying a threshold detection limit, using Boolean measurements also scales the problem from the true concentrations to values between zero and one. Furthermore, since the hydraulic calculations are decoupled from the water quality calculations, the network time-varying flow patterns are assumed to be known inputs to the water quality model. These network flows, coupled with the topology of the network, completely determine the time-delays associated with individual paths through the network. Therefore, the network time-delays are independent of the magnitude of the contaminant concentration. This helps explain why the Boolean formulation is able to correctly identify the node locations and the contaminant injection time, independent of the true magnitude of the injection.

We also demonstrate the effectiveness of the Boolean sensors for varying sensor densities. Figure 4 illustrates the improvement in the quality of the source detection as the sensor network becomes more dense. The characterization of the predicted source location changes between a sensor network of twenty-five sensors and a network with fifty sensors. We see some nodes that are identified with twenty-five sensors that are ruled out as potential candidates when we use fifty sensors. As well, two additional nodes are added as candidates. With the denser sensor network, more of the distribution system is being monitored. The contaminant presence in this portion of the network was not captured until the additional sensors were included. It is important to note through this discussion, however, that these discrepancies are only on nodes where the injection solution is a low percentage of the total injection.

Figure 5 illustrates the predicted time history of the source injection at the injection node as the number of sensors is increased. As expected, the accuracy of the profile shape and the contaminant injection time is improved with increased sensor density.

Next, we evaluate the solution quality as a function of varying the threshold detection limit. Table 3 shows the source inversion results for threshold values ranging from five percent to forty percent of the maximum network concentration. Of course, in an actual system with Boolean sensors the threshold value would be determined by the sensor design, but this allows us to predict the effect of that threshold on the

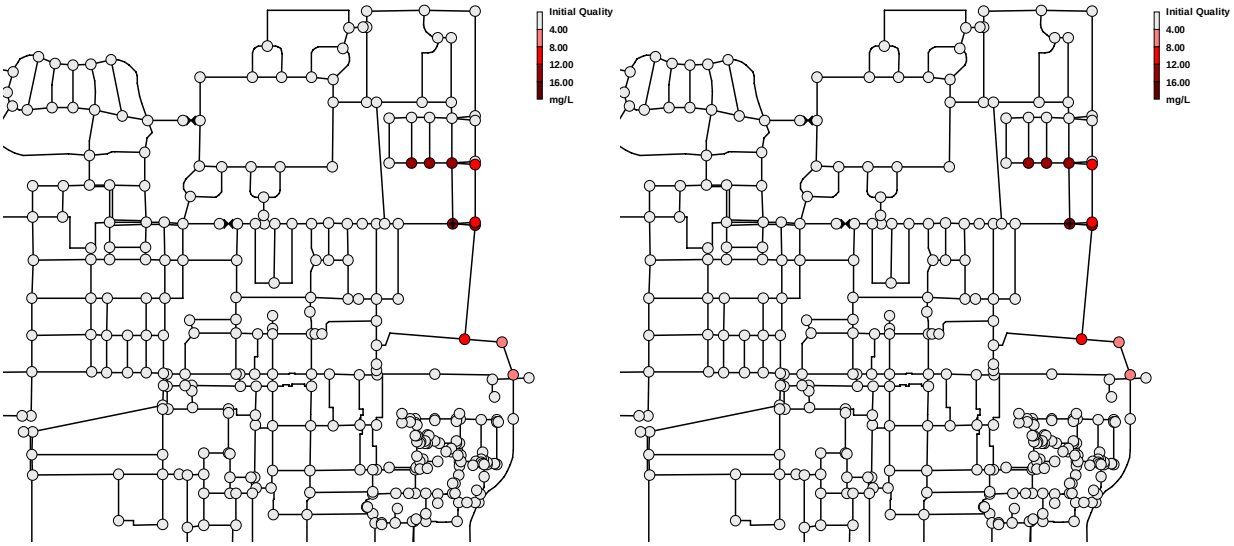

(a) Prediction with 13 Boolean Sensors

(b) Prediction with 25 Boolean Sensors

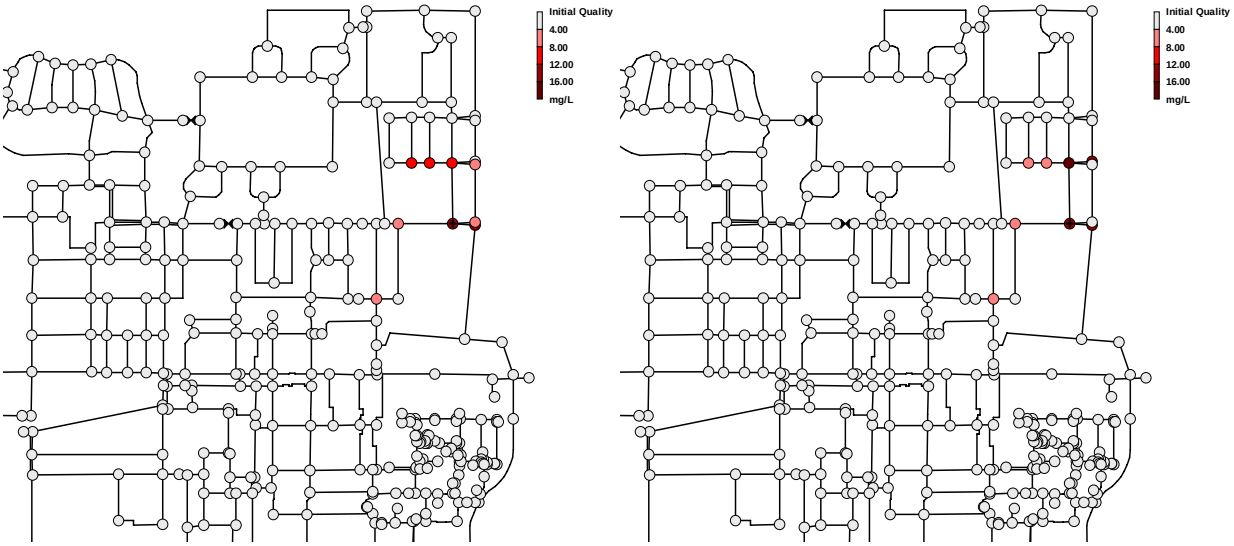

(c) Prediction with 50 Boolean Sensors

(d) Prediction with 100 Boolean Sensors

Figure 4: Qualitative comparison of the quality of the source detection based on the (a) 13 Boolean sensor, (b) 25 Boolean sensors, (c) 50 Boolean sensors, and (d) 100 Boolean sensors. The injection node is marked with a $\oplus$. The intensity of the nodal coloring indicates the likelihood that the node is the injection node, with greater intensity indicating a higher probability. As expected, as the number of sensors increases, the quality of the source detection improves.

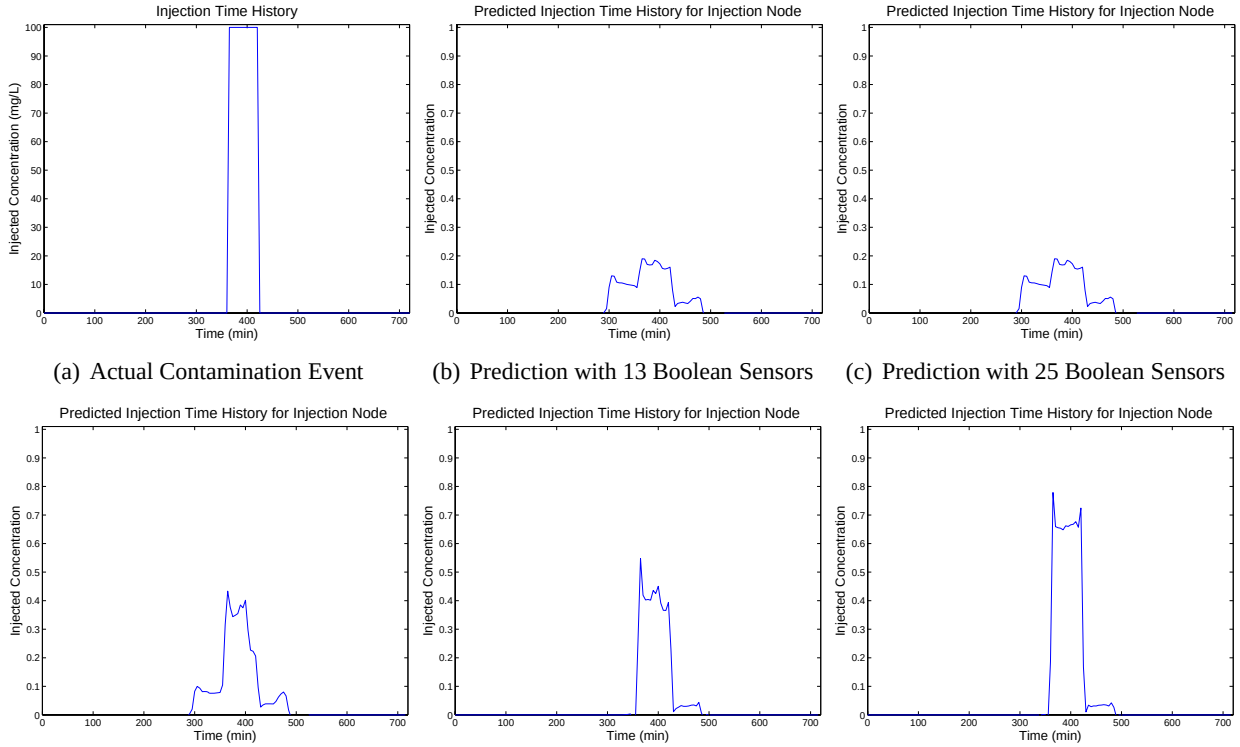

(a) Actual Contamination Event (b) Prediction with 13 Boolean Sensors (c) Prediction with 25 Boolean Sensors

(d) Prediction with 50 Boolean Sensors (e) Prediction with 100 Boolean Sensors (f) Prediction with 199 Boolean Sensors

Figure 5: The time history of a simulated contamination event at the injection node is illustrated in (a). We illustrate the predicted contaminant at the injection node in time for sensor networks containing (b) 13 Boolean sensors, (c) 25 Boolean sensors, (d) 50 Boolean sensors, (e) 100 Boolean sensors and (f) 199 Boolean sensors. As the sensor density increases, the quality of the source prediction more closely approximates the actual injection at the source node.

NodeID	5 %	10 %	20 %	40 %
2201	12.6 %	12.7 %	12.6 %	11.8 %
433	12.6 %	12.6 %	12.5 %	11.8 %
432	12.6 %	12.6 %	12.5 %	11.8 %
431	12.6 %	12.6 %	12.5 %	11.8 %
INJ	12.6 %	12.6 %	12.5 %	11.8 %
577	9.5 %	9.8 %	10.0 %	9.1 %
2401	8.4 %	8.5 %	8.4 %	7.9 %
2301	8.4 %	8.5 %	8.4 %	7.9 %
807	7.7 %	7.8 %	7.9 %	9.6 %
809	5.4 %	5.4 %	5.6 %	8.1 %

Table 3: Comparison of predicted source location as a function of threshold percentage for Boolean sensors. Reported is the percentage of the total injection that the source detection algorithm allocates to each node. **INJ** indicates the actual injection node for the simulated attack.

source inversion. The inversion approach is able to effectively identify the region of interest over the entire threshold range studied. In this particular injection scenario a square pulse was injected which inherently has an *on-off* type behavior that is captured well with the Boolean measurements. We anticipate that the inversion approach would be equally successful with a constant injection scenario as long as the concentration was significant when compared to the detection limit. However, the inversion approach may encounter difficulty if the injection produces concentrations that fluctuate around the sensor threshold value.

4 CONCLUSIONS AND FUTURE WORK

The source inversion algorithm has successfully inverted for the location of source terms using sensor information with a significant amount of uncertainty. This uncertainty however does not penalize spatial information but instead influences the magnitude. This means that current water quality sensors, capable of identifying the presence of contaminant by combining the signature of multiple parameters, can be used to support early warning systems. Furthermore, these sensor designs appear to provide sufficient information for the source inversion problem.

Our numerical experiments show that the source inversion approach is remarkably reliable over a large range of threshold detection limits. It is clear, however, that increasing the sensor density has a positive effect on the ability to invert for the injection source, as well as providing a more effective early warning system. While further studies are necessary, this study seems to indicate that it may be more effective to install a dense network of lower cost Boolean sensors rather than a very sparse network of highly accurate concentration sensors. Of course, this strategy is only effective if the Boolean sensors are sufficiently sensitive to detect dangerous concentrations of the contaminant. Future work will investigate this trade-off further.

Other future work will consist of further understanding uncertainties associated with the source inversion algorithm. Although this paper is focused on evaluating specific type of sensor uncertainty, a range of parameters needs to be evaluated in the context of this uncertainty, such as number of sensors, diffusion

coefficient of the contaminant, different demand patterns and larger datasets.

BIBLIOGRAPHY

[1] B. ANDERSON AND J. MOORE, *Optimal Filtering*, Prentice Hall, 1979.

[2] C. LAIRD AND L. BIEGLER, *Nlp algorithm for multi-scenario optimization*, in 3rd International Conference on High Performance Scientific Computing, Hanoi, March 6-10, 2006, To appear.

[3] C. LAIRD, L. BIEGLER, AND B. VAN BLOEMEN WAANDERS, *Time dependent contaminant source determination: A network subdomain approach for very large networks*, in World Water and Environmental Resources Congress 2004, Salt Lake City, Utah, June 2004.

[4] ——, *A mixed integer approach for obtaining unique solutions in source inversion of drinking water networks*, in World Water and Environmental Resources Congress 2005, Anchorage, Alaska, May 2005.

[5] ——, *Real-time, large scale optimization of water network systems using a subdomain approach*, in Second CSRI Conference on PDE-Constrained Optimization, To appear.

[6] C. LAIRD, L. BIEGLER, B. VAN BLOEMEN WAANDERS, AND R. BARTLETT, *Contamination source determination for water networks*, A.S.C.E. Journal of Water Resources Planning and Management, 131 (2005), pp. 125–134.

[7] C. LIOU AND J. KROON, *Modeling the propagation of water-borne substances in distribution water networks*, Journal of American Water Works Association, 79 (1987), pp. 54–58.

[8] S. A. MCKENNA, B. G. VAN BLOEMEN WAANDERS, C. D. LAIRD, , S. G. BUCHBERGER, Z. LI, AND R. JANKE, *Source location inversion and the effect of stochastically varying demand*, World Water and Environmental Resources Congress, Anchorage, AL, 2005.

[9] L. ROSSMAN, *EPANET 2.0 User's Manual*, EPA, Cincinnati, Ohio, 2000.

[10] L. ROSSMAN AND P. BOULOS, *Numerical methods for modeling water quality in distribution systems: A comparison*, Journal of Water Resources Planning and Management, 122 (1996), pp. 137–146.

[11] L. ROSSMAN, P. BOULOS, AND T. ALTMAN, *Discrete volume-element method for network water-quality models*, Journal of Water Resources Planning and Management, 119 (1993).

[12] A. TARANTOLA, *Inverse Problem Theory*, SIAM, 2005.

[13] B. VAN BLOEMEN WAANDERS, R. BARTLETT, L. BIEGLER, AND C. LAIRD, *Nonlinear programming strategies for source dection of municipal water networks*, in EWRI World Water and Environmental Congress, Philadelphia, Pennsylvania, June 2003.

[14] C. VOGEL, *Computational Methods for Inverse Problems*, SIAM, 2002.

[15] A. WÄCHTER AND L. T. BIEGLER, *On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming*, To Appear in Mathematical Programming, (2005).