United States
Environmental Protection
Agency

Overview of Event Detection Systems for WaterSentinel

Draft, Version 1.0

December 12, 2005

U.S. Environmental Protection Agency
Water Security Division
Ariel Rios Building, Mail Code 4601M
1200 Pennsylvania Avenue, N. W.
Washington, DC 20460

EPA 817-D-05-001

Disclaimer

The Water Security Division, of the Office of Ground Water and Drinking Water, has reviewed and approved this draft document for publication. This document does not impose legally binding requirements on any party. The word "should" as used in this Guide is intended solely to recommend or suggest and does not connote a requirement. Neither the United States Government nor any of its employees, contractors, or their employees make any warranty, expressed or implied, or assumes any legal liability or responsibility for any third party's use of or the results of such use of any information, apparatus, product, or process discussed in this report, or represents that its use by such party would not infringe on privately owned rights. Mention of trade names or commercial products does not constitute endorsement or recommendation for use.

Questions concerning this document or its application should be addressed to:

Irwin Silverstein, Ph.D., P.E.
U.S. EPA – New England, Region I
1 Congress Street
Boston, MA 02114
(617) 918-1632
Silverstein.Irwin@epamail.epa.gov

Acknowledgements

The Water Security Division would like to recognize the following organizations and individuals for their support in the preparation of *Overview of Event Detection Systems for WaterSentinel*:

- Steve Allgeier, Office of Water – Water Security Division
- Regan Murray, Office of Research and Development – National Homeland Security Research Center
- Sean McKenna, Sandia National Laboratories
- Doron Shalvi, Computer Sciences Corporation (CSC)

Executive Summary

An effective contamination warning system (CWS) should be able to identify deviations from established baselines and system anomalies in a timely manner by integrating and analyzing information from online water quality monitoring, sampling and analysis, enhanced security monitoring, consumer complaints, and public health surveillance. The WaterSentinel (WS) approach for a CWS involves the active deployment and use of monitoring technologies/strategies that are not contaminant-specific and enhanced surveillance activities to collect, integrate, analyze, and communicate information from a variety of sources. These information streams are utilized to provide a timely warning of potential water contamination incidents and initiate response actions to minimize public health and economic impacts. As part of an effective CWS, a method for event detection should be developed such that it can quickly recognize true contamination incidents while minimizing the occurrence of false positives. An effective event detection system, coupled with the other components of the WS-CWS, should contribute to the CWS goal of reliable and timely detection of contamination to protect public health and water system infrastructure.

The event detection system component of the CWS includes methods to distinguish potential contamination incidents from the normal background variability of the baseline. Inability to identify true incidents could result in false negatives for the WS-CWS, with potentially detrimental impacts on the health and security of the public. On the other hand, the consequences of a false alarm can also be significant and consideration should be taken to minimize false positives. Hence, the event detection system should employ software that can 'train itself' on the normal variation in daily, weekly, and seasonal water quality patterns so that it can reliably recognize an anomaly that is truly indicative of a contamination incident. Algorithms, utilizing discrimination, clustering, or statistical techniques, are being used to identify anomalous patterns in the fields of public health and cyber security. Laboratory studies have determined that the use of algorithms for event detection in the environmental field is feasible. Currently, several field projects are underway using a combination of algorithmic approaches to identify anomalies in water quality patterns. Other organizations have developed software for evaluating water quality with predetermined threshold points acting as alarm triggers, as opposed to triggers based on sophisticated algorithms. Regardless of the degree of sophistication of an event detection system, human judgment and interpretation should always be an element of the credibility determination and decision-making process.

Present findings suggest that methods are available for testing and evaluating the event detection component of the WS-CWS. However, these results are merely preliminary and should be utilized with caution. Further testing of event detection within the context of a working utility distribution system is needed to validate its usefulness within an integrated system, such as the WS-CWS. Long-term testing should better ascertain the effectiveness, ability to be replicated, and sustainability of an event detection system and the entire WS-CWS system architecture.

Table of Contents

List of Figures and Tables

Section 1.0: Introduction

In response to Homeland Security Presidential Directive 9 (HSPD 9) and by its authority under section 300i-3 of the Safe Drinking Water Act (42 USC section 1434), the U.S. Environmental Protection Agency (EPA), in collaboration with other agencies, plans to build upon and expand current monitoring programs to expand current surveillance and monitoring systems that provide early detection of water contamination. Ideally, these systems would identify the presence of contamination prior to human exposures that would result in public health impacts. Currently available technologies for deployment in distribution systems are not sufficiently advanced to detect specific contaminants and provide such a timely alert. In addition, a system architecture that relies on the detection of specific contaminants faces a number of challenges for success because it is unlikely that the technologies would be able to provide contaminant-specific detection for all potential contaminants and the costs associated with deployment of multiple contaminant-specific technologies at a number of locations throughout complex distribution systems would be overwhelming. Also, water utility personnel would be severely burdened by the calibration, operation, and maintenance needs of what would likely be highly sophisticated technologies.

To meet the charge of HSPD 9, the WaterSentinel (WS) program is a demonstration project whereby EPA, in partnership with a pilot utility and laboratories, would design, deploy and evaluate a model contamination warning system (CWS). The WS approach involves the active deployment and use of monitoring technologies/strategies that are not contaminant specific and enhanced surveillance activities to collect, integrate, analyze, and communicate information from a variety of sources to provide a timely warning of potential water contamination incidents and initiate response actions to minimize public health and economic impacts. More information about WS program design is available in *WaterSentinel System Architecture* (USEPA, 2005a).

1.1 Recognizing Anomalies

The different data streams for decision-making envisioned for the WS-CWS include sensors for conventional water quality parameters (e.g., pH, residual chlorine, total organic carbon, conductivity, etc.), public health surveillance information, consumer complaints tracking, routine or triggered sampling and analysis data, and enhanced security monitoring. The fundamental challenge to the reliance on a variety of information streams as an indication of a contamination incident is a means of distinguishing anomalous patterns in these data from background signals. Statistical formulas and mathematical models that analyze data are called algorithms. While not widely deployed for analyzing water quality data and consumer complaint clusters, algorithms incorporated into event detection software also can be used to identify and 'learn from' changes in data patterns that are indicative of the introduction of a contaminant into the source water or the distribution system of a water utility, or at least indicative of a significant change from the current nature of water quality. More information about how water quality parameters change in response to the intentional or accidental introduction of contaminants into a water utility's distribution system is available in *WaterSentinel Online Water Quality as an Indicator of Drinking Water Contamination* (USEPA, 2005b).

Algorithms currently are being used to recognize anomalies in such fields as public health (Watson, W., et al., 2005), cyber security (Jackson, 2003), and for jet engine maintenance (Shroder, 2005). Algorithms can also be used to optimize the number and locations of contaminant sensors in a network (Berry, et al, 2005; Uber, et al., 2004; Ostfeld and Salomons, 2004). However, as pointed out by McKenna, et al., 2005, the majority of the work in this area has assumed contaminant-specific sensors with perfect detection characteristics.

In the public health field, detection of an anomaly that is two to three standard deviations from the baseline observation is typically the threshold for sending an alert to the appropriate public health representatives. For example, the American Association of Poison Control Centers receives poison incidence monitoring data in real-time that are first evaluated by algorithms at hourly intervals. If anomalies in the data reflecting an increased number of calls, unusually high reporting of a specific substance, and clusters of unusual health effects are cited, the results are sent to subject matter experts to make a final determination whether the data outliers are of potential public health significance. In addition, the application of event detection algorithms in the public health field is being expanded to the field of water security. In the case of the Real-time Outbreak and Disease Surveillance (RODS) project described below, anomalies regarding emergency room visits, sales of over-the-counter (OTC) drugs, and hospital admission rates can be indicative of the effect of exposure to a waterborne contaminant. This type of event detection based on syndromic surveillance is being combined with event detection for water quality with the expectation that better communication and information sharing between the public health community and the water utility will shorten the recognition time necessary for taking appropriate response actions. In addition, the Electronic Surveillance System for the Early Notification of Community-Based Epidemics (ESSENCE) project described below is a biosurveillance system being developed to evaluate data from pharmacies, hospitals, clinics, etc., with the intent to integrate a water quality data stream into an overall event detection system (EDS).

1.2 Event Detection and Consequence Management

Event detection is the trigger that sets in motion the process of consequence management that includes credibility determination, response actions, and recovery. The relationship between event detection and consequence management in the WS concept of operations is described in *WaterSentinel System Architecture* (USEPA, 2005a) and illustrated in **Figure 1-1** below. While event detection initiates the credibility determination process, it is the analysis of the information associated with this trigger that determines the set of response actions the utility should take. These actions can include sending a site characterization team into the field and requesting information from law enforcement and/or public health agencies.

While the EDS can indicate a possible contamination threat, it is only part of the credibility determination process and it alone cannot be the sole decision-making tool for initiating response actions that are warranted to protect public health. In addition, the human element is integral to the credibility determination and subsequent consequence management steps. However, it is possible to develop a tool to support officials in decision-making by guiding them through the initial stages of the evaluation process and aiding in the synthesis of information necessary to make timely and appropriate response decisions. Although ultimately the decision should always rely on human judgment and the evaluation of incomplete information, this decision tool can be a great aid in the process, and might substantially reduce the time to make critical response decisions. Therefore, event detection software that utilizes algorithms developed to evaluate water quality data are an important component of a successful CWS.

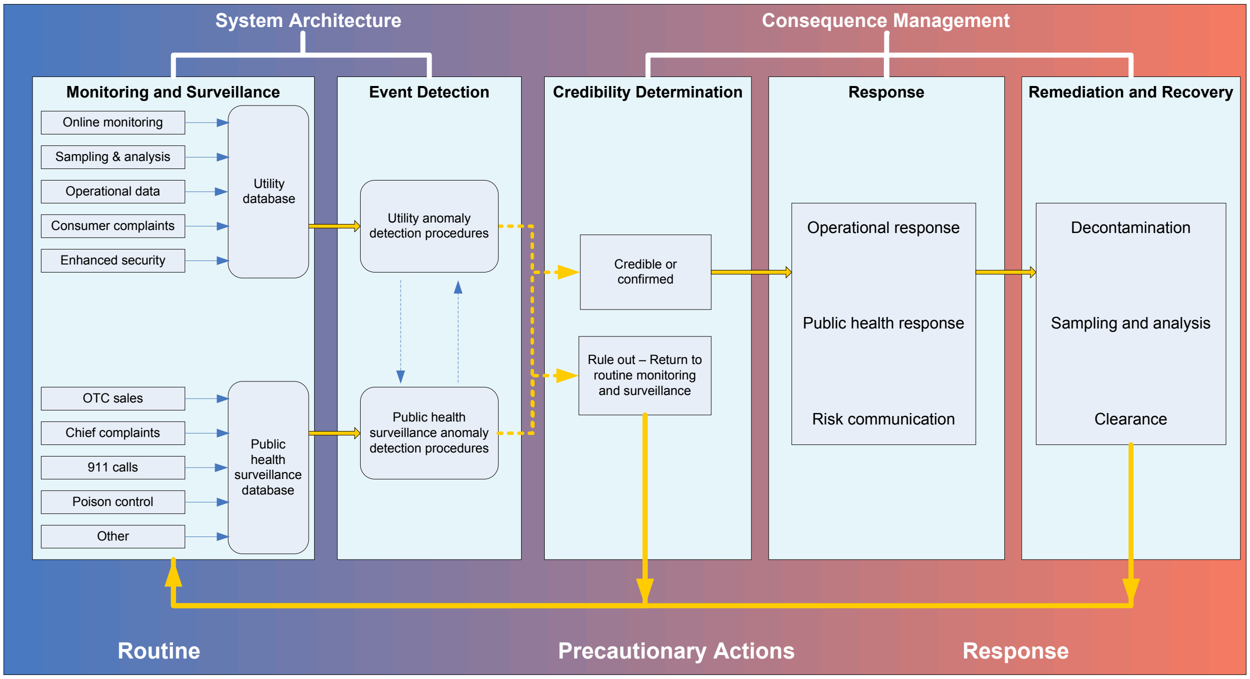

Figure 1-1. Overview of WS Concept of Operations

1.3 Objectives

The objectives of this document are as follows:

- Describe the concept of event detection algorithms to identify water quality anomalies
- Describe tools and approaches to evaluate the effectiveness of event detection
- Cite current projects and currently available event detection products being used to evaluate water quality data
- Cite other projects involved with water quality monitoring
- Discuss how the WS pilot is an appropriate exercise for building upon the preliminary results of current event detection projects

1.4 Document Organization

The remaining sections of this document describe the following aspects of the WS EDS:

- **Section 2.0: Overview.** This section presents an overview of the functions of EDSs.

- **Section 3.0: Event Detection and Water Quality Anomalies: Proof of Concept**. This section describes the different algorithm categories and ways to evaluate how they perform, to support the online water quality monitoring component of WS.

- **Section 4.0: Evaluating the Effectiveness of Event Detection Systems.** This section outlines an approach that can be used for evaluating the effectiveness of an EDS.

- **Section 5.0: Projects Currently Using Event Detection.** This section describes projects involving event detection in the water sector.

- **Section 6.0: Other Water Quality Analysis Activities.** This section provides examples of projects conducted by utilities, commercial vendors, and research organizations to improve water monitoring capabilities and data analysis and a description of companies that offer off-the-shelf tools for event detection and decision support.

- **Section 7.0: Summary and Preliminary Conclusions.** This section provides a summary of the document and some preliminary conclusions.

- **Section 8.0: References**. This section provides a bibliography of the references cited in this document.

- **Acronyms**

A complete glossary of terms related to event detection and the WS program is available in *WaterSentinel System Architecture* (USEPA, 2005a).

Section 2.0: Overview

The most important function of the event detection software is to filter out the anomalies or changes in water quality patterns that normally occur (e.g., changes in water turbidity caused by the surge effect of activation of a booster pump in the water distribution system), or which have known causes, and signal only those anomalies that are likely to be indicative of possible contamination incidents. In short, the purpose of the event detection software is to reduce the false positive rate without missing potential events. One approach to identifying an anomaly of concern employs a baseline estimator to quickly distinguish false warnings from regular fluctuations in water quality data upon system start-up (Cook, J., et al., 2005). Over time, the event detection software's programmed ability to become familiar with signal patterns that are associated with 'normal' operations (i.e., actuation of a booster pump every day at 3:00 PM) will result in a continued drop off in the number of false warnings. This concept is analogous to computer network security event detection software products that also define a baseline of normal activity and report any events that do not neatly fall within a predefined cluster. For example, a baseline profile of the network is created where all routine events are grouped into clusters representing normal activity. The software 'trains' itself to recognize any patterns outside of those clusters based on what different network users do, the resources they typically request, the types of files they transfer, etc. (Cook, J., et al., 2005). In another example of event detection software, a neural network developed by Intelligent Automation Corporation is used to interpret and learn from electrochemical signals produced by fish exposed to source water, so that anomalous patterns can be recognized that are indicative of a contamination event (Intelligent Automation Corporation, 2005).

The degree to which water quality data should be collected before a baseline is well established will depend on the inherent variations in the distribution system. Although these variations can occur daily, weekly, and seasonally, the process by which the event detection software 'learns' how to identify anomalies can become functional early on (e.g., one event detection software vendor, discussed below in Section 6.6, indicated that it would take two to three months for its software to complete an initial training period using a utility's water quality data). As additional water quality data are received by the event detection software, its ability to distinguish anomalies unrelated to a contamination incident (i.e., false positives) from true incidents will improve. Although a period of up to one year may be necessary to characterize certain seasonal trends, the EDS and online water quality monitoring system should become more reliable long before this full level of variability is characterized (i.e., as the system learns the normal degree of variability). Therefore, the pilot study should have a functioning EDS in its early stages and like all components of the CWS, improvement is anticipated as experience is gained in data interpretation and integration.

This document will describe various approaches to event detection, and characteristics of these systems that impact performance. To demonstrate the extent to which event detection is currently being employed in the field, a number of ongoing projects at water utilities, both civilian and military, will be presented. These projects involve utilities that are field testing event detection algorithms for changes in water quality and studies where the use of syndromic public health surveillance models that also rely on algorithms are being evaluated. The document will also include a description of other utility activities being planned which also incorporate event detection. The document will conclude with a discussion of how current projects involving event detection have performed and how the WS program can build upon this knowledge base during the pilot study, recognizing that all potential contamination events may not be detected by surrogate water quality measures.

Section 3.0: Event Detection and Water Quality Anomalies: Proof of Concept

While event detection algorithms have been used for a number of years to identify anomalies in a variety of other fields, they have limited use in the environmental field; however, a number of published and ongoing studies indicate a growing interest in EDSs for environmental applications, including the detection of anomalies in water quality data. Research investigations are under way with the EPA, the military, water utilities, and vendors of event detection softwares where the algorithms are being used to evaluate patterns in water quality data to identify contamination. The following sections describe the different algorithm categories and ways to evaluate how they perform. Subsequent sections will present projects that illustrate how this concept is being used in laboratory and field studies.

3.1 Categories of Event Detection Algorithms

Although many event detection algorithms exist, they generally fall into one of the following categories: discrimination, clustering, or statistical techniques.

Discrimination involves the development of models to separate samples into two or more discrete classes which are known in advance, using, for example, rules, decision trees, Artificial Neural Networks (ANNs), Bayesian Networks, and Support Vector Machines. For example, some types of ANNs are iteratively 'trained' to learn how to map input patterns to output classes. The iterative training process involves teaching (supervised learning) the ANN so that it can adjust its parameters according to whether its suggested output class(es) was correct or incorrect. Practical examples include training and using an ANN to input the data from a sensor (following some preprocessing), and select one of several possible output classes representing contaminants (e.g., strychnine, gasoline, ricin, none, etc.) which have been defined in advance. While many other supervised learning approaches exist, the general idea in discrimination is that given an input data pattern, a model is being trained to select one or more classes that have been defined in advance. After being trained, the model can be used as-is for operational scenarios, or can be further trained on-the-fly when presented with new training data. Therefore, it is an adaptive model.

Clustering involves the development of models to organize the data into clusters which are not known in advance, typically by looking at the similarity of the data according to selected measures based on multiple measured variables (i.e., multivariate analysis). This type of learning is often called unsupervised learning, since the target outputs are not known in advance, and there is no set of output classes to train to as a target. As an example, this approach may be used to organize online water quality monitoring sensor data into categories that were previously unknown. Such categories, indicated by clusters of one or more cells in a map grid, represent clusters of input pattern data similar in nature. Further examination is typically required to identify the input characteristics that led to the clustering; this examination may lead to clusters of data representing, for example, normal water quality, low levels of specific conductance, high levels of a toxin, presence of an unknown contaminant, and many others, including clusters for which the signature similarity of the cluster is not readily apparent. In this fashion, new, potentially actionable information may be obtained. This type of knowledge discovery may be more useful in a retrospective analysis than a prospective one, allowing the system to first identify the set of possible clusters, and then train a separate component of the system to discriminate new patterns into the clusters, using the discrimination techniques previously mentioned.

Statistical techniques involve the use of probabilities and sufficient population sample sizes to establish normal and abnormal conditions in a system, and to make predictions of future events based on the current state of the system being observed. A statistical model may be as simple as one that determines the mean of a single observed parameter, such as chloride measurement from a sensor, and flags any readings that fall outside of a certain number of standard deviations from the mean (assuming normal distribution). More commonly, statistical control charts may be used to monitor one variable over time, such as emergency room visits with respiratory or gastrointestinal complaints, triggering a response when certain signals (i.e., step, spike, exponential) are detected in the data (Shmueli, G., 2005). In another example, prior probabilities are used to compute future probabilities to determine whether spatially-organized counts of OTC drug sales are significantly higher than the expected baseline, triggering an alert (Neill, D.B. and Moore, A.W., 2005).

3.2 Demonstrating the Concept of Applying Event Detection Software to Water Quality Data

The previous section described the different categories of event detection algorithms. The following studies provide specific examples where event detection algorithms are being used to evaluate changes in water quality patterns. The first two studies present results based on laboratory experiments. All the contaminant data presented in the first study are for concentrations well below the literature LD_{50} values (*WaterSentinel Contaminant Fact Sheets*, USEPA, 2005c). For the second study, sodium cyanide, sodium fluoroacetate, and aldicarb were tested at concentrations below the lethal dose (which is based on the assumption that a person weighing 70 kg consumes 2 liters of water per day) and the range of testing for sodium arsenate encompassed the lethal dose for this fourth chemical (Inchem, 2005, CDC, 2005, and Wikipedia, 2005). Therefore, these concentrations were 'recognized' during the test at concentrations that would be protective of acute human health exposure. Research is underway by EPA's National Homeland Security Research Center (NHSRC) to determine the health risks associated with consuming sub-acute concentrations of WS priority contaminants. As that information becomes available it will likely be used in the design of future studies. The third study indicates how algorithm development is associated with the collection of baseline water quality data at a number of water utilities.

3.2.1 Study at EPA's Test and Evaluation (T&E) Facility

At EPA's T&E Facility in Cincinnati, Ohio, a laboratory study was undertaken to determine whether it was feasible to monitor water quality sensor data for changes that were indicative of the injection of a contaminant, and whether event detection software algorithms were capable of 'recognizing' such an anomaly (more information about the change in these water quality parameters observed during this study is available in *WaterSentinel Online Water Quality as an Indicator of Drinking Water Contamination,* USEPA, 2005b). Water quality data were obtained before, during, and after the introduction of contaminants (i.e., secondary effluent from a wastewater treatment plant, potassium ferricyanide, arsenic trioxide, nicotine, aldicarb, a malathion-containing insecticide, a glyphosate-containing herbicide, and *E. coli* K-12 strain with growth media), ranging in concentrations from about 6 milligrams per liter (mg/L)

to about 2 grams per liter, into the facility's recirculating pipe loop. Sensor data were continuously collected and electronically archived to establish stable baseline conditions and to record sensor responses to the contaminants injected.

The types of change in sensor data were dramatic and consistent enough for several parameters, (e.g., specific conductance, total organic carbon (TOC), total/free chlorine, chloride, and oxidation reduction potential (ORP)), to detect the injection of 11 of the 13 contaminants. After the pipe loop experiments were conducted, the sensor data were provided to the Hach Company of Loveland, Colorado for evaluation by its event detection software called the Event Monitor. The Event Monitor contains a library representing the characterization of over 80 contaminants of concern through laboratory scale 'beaker' experiments. By comparing water quality signature patterns to this library, the Event Monitor was able to identify most of the contaminants introduced into the pipe loop. A report summarizing these experiments is anticipated by mid November 2005 (Kroll, D. and King, K, 2005). A follow-up to this laboratory study has been initiated using chemical and biological warfare agents (Hall, J. et al., 2005).

These experiments first demonstrated the feasibility in detecting chemical and gross biological contamination based on a noticeable change in commonly used water quality sensors. Next an event detection algorithm using discrimination was able to interpret the sensor data to not only recognize that an anomaly had occurred but 'identify' what caused the anomaly. These results illustrate on a laboratory scale how water quality sensors and event detection algorithms can work together to provide timely warning of a contamination incident.

3.2.2 Decision Support Software Project (Discrimination and Clustering Algorithms)

In 2005, another laboratory study was initiated to investigate the capability of water quality monitoring and event detection softwares for 'recognizing' that a contaminant had been introduced (see Section 5.1 for more detail). The first portion of this project involved the development of a prototype Decision Support Software (i.e., event detection software) by using a laboratory flow loop to collect data about water chemistry response using combinations of different commercially available, field grade water quality sensors. During these experiments, four chemical threat substances (i.e., sodium arsenate, sodium cyanide, sodium fluoroacetate, and aldicarb) were introduced at concentrations ranging from 15 to 100 mg/L (sodium arsenate), 3 to 10 mg/L (sodium fluoroacetate), 0.5 to 10 mg/L (sodium cyanide), and 1 to 10 mg/L (aldicarb). The event detection software first organized existing data into clusters and then compared new data to the clusters to discriminate whether they fell within a normal or anomalous pattern in real time. The laboratory results were encouraging as the event detection software was able to discriminate normal water chemistry from the patterns signifying the introduction of contaminants (Byer and Carlson, 2005).

Like the T&E experiments, this study demonstrated the feasibility in detecting chemical contamination based on a noticeable change in commonly used water quality sensors. Similarly, event detection algorithms, in this case using a combination of discrimination and clustering, were able to interpret the sensor data and recognize that contamination had been introduced. A field testing phase of the study began in September 2005 and can provide an opportunity to observe how well water quality sensors and event detection algorithms work together in a distribution system. This phase, inclusive of a simulated attack using modulated concentrations of fluoride, will more closely represent a test of the WS concept of operations and therefore provide useful feedback in preparation for implementation of the WS pilot study.

3.2.3 Algorithm Development Project (Statistical Algorithms)

In a third study (see Section 5.3), actual water quality data (e.g., ORP, pH, conductivity, temperature, dissolved oxygen (DO), and chlorine) collected at a southwestern water utility and a consortium of northeast utilities, as well as from a simulated data set, are being used to develop event detection algorithms. The use of a simulated data set allows for the insertion of 'events' so that the rate of false positives and false negatives associated with the detection algorithms can be evaluated. This approach has the advantage over the real data set, which does not have any 'events' to detect. The transient characteristic of measured water quality data in the distribution system necessitates algorithms that can be updated in real time. The goal of the study is to develop a statistical model that first recognizes temporal trends in key water quality variables associated with diurnal and seasonal cycles, the influence of physical operations in the distribution system, electronic drift in water quality sensors or by dynamic chemical/biological processes within the distribution system, and then confidently recognizes a change caused by abnormal operating conditions (i.e., the introduction of a contaminant) (Srinivasan, S., et al., 2005).

The preliminary results for this study illustrated how a statistical relationship between data collected during one time period can be used to predict data that will be collected during future time periods. The study also indicated that the event detection algorithm could be 'tuned' so that that during periods of increased water security concern, the algorithm could operate at a higher probability of detection. In addition, this study is using simulated data to test algorithm performance. Like the previous two studies, the feasibility of using water quality data and event detection algorithms to detect contamination is being demonstrated. Obtaining information on algorithm tuning provides an opportunity to evaluate how this strategy might be used during the pilot to increase the probability of detection, with the corresponding realization that during this time period, more false positives would occur. Finally, the use of simulated data has value to the WS-CWS because it is an approach that can be used during the pilot study to test the performance of the EDS.

In summary, the three studies described above show that all three types of event detection algorithms (e.g., discrimination, clustering, and statistical) are being used in the environmental field to identify anomalies in water quality data. In the studies involving discrimination and clustering algorithms, event detection softwares successfully determined that contamination had been introduced into a laboratory control loop based on the signals generated by standard water quality sensors, as envisioned for the WS-CWS. Although proving this concept in the field will require further study, there are a number of current projects that are investigating the combination of water quality monitoring and EDS as a means of providing an early trigger to identify contamination in a distribution system. Many of these projects are in the early stages where baseline data are being collected to 'train' the event detection software so that it will be able to recognize anomalies and identify those associated with contamination incidents. A further discussion of these projects is provided in Section 5.0.

Section 4.0: Evaluating the Effectiveness of Event Detection Systems

Section 3.0 demonstrated that the application of event detection now includes the environmental field, in addition to the other fields like public health and cyber security. While several examples were cited where an event detection software product could fulfill the requirement for event detection in the overall WS concept of operations, a process is needed in the pre-implementation phase of this program to determine what event detection software product(s) should be used at the pilot utility. This section outlines an approach that can be used for evaluating the effectiveness of an EDS. First, general criteria are presented for evaluating event detection algorithms and tools, primarily based on information that can be provided by event detection software vendors currently engaged in relevant projects. An initial qualitative evaluation can be made based on this information. Next, a detailed discussion is presented regarding the use of receiver operating characteristic (ROC) curves, which includes some of these criteria, to relate tool performance to the occurrence of false positive and false negative responses. Simulated data, inclusive of baseline profiles and 'contamination incidents', can be provided to vendors so they can generate ROC curves and enable a more quantitative evaluation. Then, a more rigorous option, which can build upon the previous steps, is presented involving third party validation of available event detection products by means of EPA's Technology Testing and Evaluation Program (TTEP). This process is depicted in **Figure 4-1**, including a final evaluation step that the pilot study would provide.

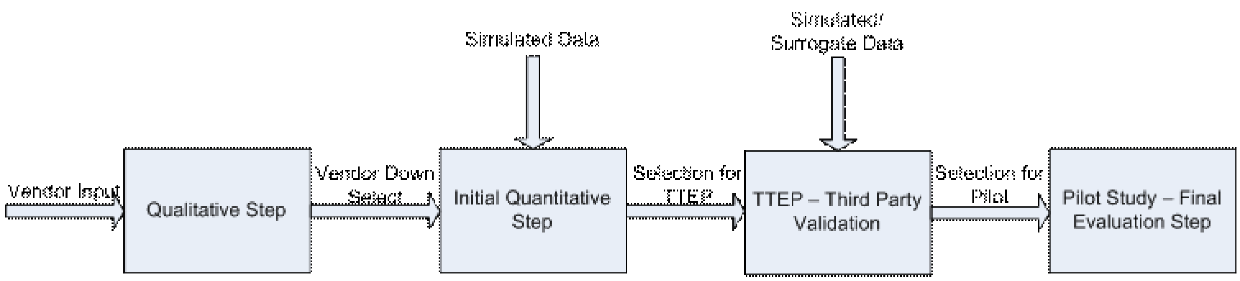

Figure 4-1. EDS Evaluation Process

4.1 Evaluation Criteria

In the context of an EDS, algorithms and tools are distinguished from one another as follows. The algorithm is the mathematical operation or statistical technique that is performed on the data for the purpose of detecting anomalies (e.g., unusual trends in water quality) and is incorporated within the event detection software or tool that interfaces with sensors, other data streams, and other utility software. The following are some measures that can be used to evaluate the effectiveness of EDS tools and/or algorithms.

Sensitivity: The sensitivity of a test is the proportion of those cases having a positive test result of all positive cases (e.g., the proportion of people diagnosed with a waterborne disease relative to the total number of people with the disease) tested; that is:

Sensitivity = (# true positives) / (# true positives + # false negatives)

In other words, the proportion of contamination incidents detected by the event detection software relative to all the contamination incidents that occurred over a given period of time.

Specificity: The specificity is equal to the proportion of true negatives of all the negative samples (e.g., the proportion of people diagnosed to be free of a waterborne disease relative to the total number of people without the disease) tested, that is:

Specificity = (# true negatives) / (# true negatives + # false positives)

In other words, the proportion of time the system is detected to be without contamination relative to the time the system is free of contamination (excluding false negatives). Specificity is also defined as 1.0 minus the false positive fraction (i.e., number of false positives divided by the sum of false positives and true negatives).

F-measure: The F-measure can be used as a single measure of performance of the test. The F-measure is the harmonic mean of sensitivity and specificity; that is:

F = (2 x sensitivity x specificity) / (sensitivity + specificity)

A perfect system, with no false positive or false negative results, would have an F-measure equal to 1.0. An F-measure between 0.7 and 0.8 would be considered a 'good' value.

Time to Detect: This measures the delay between the time the event detection software first starts receiving information about the contamination incident and the time that the system recognizes the anomalous data as indicative of contamination. The response time can be measured between the introduction of a surrogate contaminant or use of simulated data and an alert given by the event detection software. The 'time to detect' for an EDS is primarily a function of the calculation interval and the nature of the algorithm itself. It is also related to the F-measure (i.e., as the time to detect is decreased, the sensitivity and specificity may degrade through the amount of data necessary to distinguish a true 'event' from normal fluctuation in baseline data. For example, if an EDS signals an 'event' immediately when a water quality parameter rises or falls, it may provide rapid detection, but at the cost of diminished sensitivity and specificity. On the other hand, if an EDS needs ten minutes of data representing a significant change from the baseline to signal an 'event', the time to detection goes down, but the sensitivity and specificity should improve. Further discussion of the many other factors in a CWS that affect the time to detect, is provided in *WaterSentinel Contamination Incident Timeline Analysis* (USEPA, 2005d).

Ability to handle highly variable data: Water quality data are influenced by many factors (e.g., seasonal factors, source water, and treatment variables) and concentration baselines should show significant change daily, weekly, and seasonally. Event detection software should have the ability to handle these highly variable data to be effective over these various time scales, and should have the ability to relate predictable water quality changes to known causes.

Adaptivity: Can the system learn on its own, or does it need to be re-trained over time? Adaptivity is valuable in a system because it reduces the amount of off-line re-training or adjustment needed.

Resource requirements: This measure applies to the costs incurred as a result of time, labor, and consumables expended during the installation, and operation and maintenance of the event detection software, and in responding to an event trigger. This metric can also be used to track the costs associated with the execution of the event trigger protocol to determine whether the expenditures were commensurate with cause of the trigger (e.g., an event trigger that leads to a discovery by the utility that a sensor calibration problem is the cause without the implementation of a drastic response action like a 'do not use' order is indicative of good protocol because the resources expended were not excessive).

Contaminant coverage: This measure applies to the range of contaminants or contaminant groups that should be detected by the EDS. Note that contaminant coverage is also strongly dependant upon the specific water quality parameters monitored. For the purpose of evaluating the EDS, it is assumed that the sensors have the theoretical capability to detect the contaminant of interest.

The following aspects should be considered in addition to the measures discussed above when developing and executing an evaluation/testing protocol for event detection tools and algorithms. These aspects apply more in the context of how well the event detection tools and algorithms should perform when integrated with the other components of the CWS:

Calibration: Does the tool require calibration runs of various contaminants in a water matrix in order to identify or interpret events? In other words, is it sufficient for the tool to establish a baseline specific to the pilot utility so that deviations from the baseline can be evaluated effectively by the event detection software, or must training of the tool include the use of contaminants or surrogates in the pilot utility's drinking water? Additionally, does the tool have an inherent ability to recognize a contaminant by comparing water quality patterns to a pre-determined event library?

Compatibility: Is the tool compatible with existing data systems in use at the utility?

Cost: What is the cost of the tool? Is it targeted at enterprise-level (10K - 100K+), mid-range (1 - 10K), or open-source / free? What are the calibration costs, especially if live agents are required?

Customization: Typically, there should need to be some customization of whatever tool is selected. This is especially true for commercial-off-the-shelf (COTS) systems. How is this customization performed? What kinds of skills (business analyst, database analyst, software developer) are needed to perform this customization? What components (user interface, business rules, data processing) need to be customized?

Functionality: What is the ease of entering knowledge into the tool, or of the tool learning the knowledge on its own? What kind of user interface is provided in order to use the tool, for tool-builders, knowledge entry, end-users and other roles? Does the tool provide a 'user interface' builder so that an interface can easily be built or modified for end-users?

Necessary Inputs: Does it require certain water quality sensors as inputs? Does it incorporate other data streams as inputs (e.g., pressure data, OTC drug sales)? Can the algorithm work if one or more data streams are temporarily unavailable (i.e., a sensor goes off-line for one or two days)?

Neutrality: Most tools are product neutral – they should run standalone, and against most or many popular databases and other data stores. However, some event detection software products are developed to be compatible with only one product line, e.g., there are water quality event detection software products that should function only with the vendor's water quality sensors.

Performance Verification: Has the software been tested with simulated or actual baseline data to assess its performance with respect to such parameters as false negative and false positive rates (e.g., verification under TTEP)?

Scalability: Will the tool easily incorporate the addition of new sensors at a later point in time? Will the addition of another sensor negate the existing knowledge base of the tool? Will calibration need to be redone with the addition of a new sensor?

Sustainability: What are the ongoing maintenance costs of the tool? How are upgrades deployed? How do upgrades affect existing customized components?

Transparency: Is the algorithm well defined? Does it have a basis in existing theory and/or strong ties to algorithms used in other fields? Does it rely on proprietary advances that will not be available for review?

4.2 Receiver Operating Characteristic (ROC) Curves

The performance and reliability of an EDS (i.e., data collection and interpretation) depend on its ability to minimize the number of erroneous 'detections' of an event that is not a contamination incident (i.e., false positives) while avoiding the erroneous 'non-detection' of a contamination incident that has actually occurred (i.e., false negative). False negatives are associated with improper sensor selection and placement, lack of instrument sensitivity at low contaminant concentrations, interference caused by background noise, and insufficient data analysis capability. False positives are associated with oversensitive detectors that generate an indication of contamination when none exists. They can also be caused by the presence of benign substances that mimic the interaction between a target contaminant and a sensor or by inappropriate event detection software algorithms. The use of ROC analysis, first developed in the 1950s as a byproduct of research involving the filtering radio signals noise, is an important tool for determining how well the EDS should perform. Too many false positives can result in complacency during an actual contamination incident, while the occurrence of an incident that is not detected (i.e., a false negative) can have serious public health and public confidence ramifications because an incident that the system was designed to detect was missed.

The generation of ROC curves is a means of determining the likelihoods of false negatives and false positives from an EDS. These curves are produced by plotting sensitivity versus specificity. An ideal EDS would have zero false negatives (i.e., 100% sensitivity) and zero false positives (i.e., 100% specificity). In reality, such an ideal situation cannot be achieved. For example, the use of low detection limit sensors would represent a situation where the sensitivity approaches 100% (i.e., minimal false negatives because the ability to detect has been sharpened), but this heightened ability to detect increases the likelihood that a detected anomaly that is not related to a contamination incident would trigger an event detection software alert (i.e., a false positive) and as the number of false positives increase, the specificity would drop. Because the consequences are much greater if an actual event is missed (i.e., a false negative), a certain percentage of false positives should be acceptable. However, the consequences of a false alarm can be significant, particularly if they result in substantial response actions, thus the false positive rate should be minimized to the most extent possible.

Ranges of sensor performance with regard to sensitivity (false negatives) and specificity (false positives) are illustrated by the ROC curve in the **Figure 4-2**.

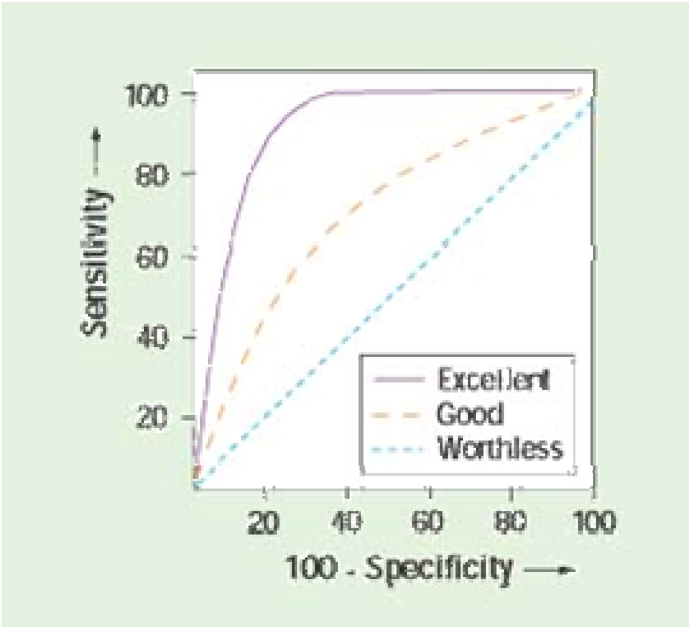

Figure 4-2. Comparative ROC Curves (Source: IVD Technology, 2005)

Another way to interpret the ROC curve is to consider that a perfect event detection algorithm for all detection threshold values would have a zero false alarm rate (‘0’ value on the x-axis corresponding to a no false positives and therefore a specificity value of 100%) and a 100% probability of detection (100 value on y-axis corresponding to no false negatives): the top left corner of the curve. For non-perfect algorithms, “the change in the threshold value at which measured values that deviate from predicted values are considered to be an ‘event’ define a continuous relationship between the rate of false positives and the probability of detection” (McKenna, et al., 2005). At relatively high detection limits, the rate of false alarms will be low and so will the probability of detecting an ‘event.’ As the detection limits are lower, more ‘events’ are detected and the number of false positives increases as indicated by top-left curve. A poorer performing algorithm is represented by the central curve, while the 45-degree line represents an algorithm that is no better than guessing (i.e., equal chance to be right or wrong) about the occurrence of an event. At a conceptual level, the ROC curve shows that the ability to detect events and the level of false alarms are inextricably linked, and have a positive and usually non-linear relationship. The construction of a ROC curve requires that a set of events exists in a form that can be used to test the event detection algorithms.

4.3 Technology and Testing Evaluation Program (TTEP)

The EPA's Office of Research and Development's (ORD) and NHSRC established the TTEP to conduct third-party performance evaluations of commercially available homeland security technologies. These evaluations incorporate the guidance of stakeholders from the water sector and other federal agencies, as well as a high degree of quality assurance oversight. Included among water security technologies are water quality sensors, field test technologies, as well as software for distribution system modeling/design and event detection. TTEP is in the process of reviewing event detection products (refer to next section) from a number of vendors to solicit their participation and is developing protocols that will be used to evaluate how effectively anomalies can be identified from variations in water quality data. The testing protocol is currently in development and may include actual field data as well as simulated data. It is anticipated that the initiation of testing and evaluation of potential event detection products for use in the WS pilot study should begin in the last quarter of 2005 and be completed by the end of the first quarter of 2006.

Section 5.0: Projects Currently Using Event Detection

There are many current projects involving event detection in the water sector. These projects typically involve collaboration between a utility and governmental, research, and commercial entities. Several of these projects are listed below.

5.1 Decision Support System for Water Distribution System Monitoring for Homeland Security

- **Utility:** South Carolina Commission of Public Works
- **Government:** U.S. Air Force
- **Research:** Colorado State University, American Water Works Association Research Foundation
- **Commercial:** Advanced Data Mining International, LLC
- **Approach:** Discrimination and Clustering
- **Status:** Prototype, Proof-of-concept

This study was initiated at Colorado State University using event detection software developed by Advanced Data Mining, LLC. Also involved in this project are the Charleston, South Carolina Commission of Public Works, the U.S. Air Force, and the American Water Works Association Research Foundation (AwwaRF, Project No. 3086). The concept behind this project involves the use of 'intelligent' software or a decision support system (DSS) and conventional water quality monitoring to determine whether a contaminant has been injected into a distribution system. The main objective of this project is to prove the concept that such a combination of sensors and software can automatically learn and remember the baseline pattern of 'normal' water quality characteristics (i.e., based on previously seen data), detect deviations from this 'normal' pattern, based on new data, that could indicate contaminant introduction, generate an alarm, and advise operators on follow-on actions. The approach is very similar to the concept used by other software products described below. It uses neural networks and multi-dimensional vectors to define normal and outlier behaviors (i.e., discrimination and clustering algorithms, as described in Section 3.2.2).

As discussed above, a combination of clustering followed by discrimination algorithms was able to distinguish pre-contaminant injection water quality sensor patterns from post-contaminant injection patterns for four chemical threat substances, in laboratory pipe loop tests. These results were presented at the AWWA Water Security Congress in Oklahoma City, Oklahoma, in April 2005 (Cook, et al, 2005). The combination of DSS and water quality sensors now are being tested in the Charleston distribution system. The main water quality parameters of interest include conductivity, pH, and UV-254 as a surrogate for TOC. Baseline water quality data are being collected that include fluoride concentrations as a prerequisite for evaluating how the EDS would respond to a simulated attack using modulated concentrations of fluoride (Cook, 2005).

5.2 Hydra Remote Monitoring System (RMS): A Case Study on Beta Test Sites and Results

- **Utility:** Greer, South Carolina water utility
- **Government:** Savannah River National Laboratory
- **Research:** Savannah River National Laboratory
- **Commercial:** PDA Technologies, Inc., Hydra RMS
- **Approach:** Statistical, Adaptive
- **Status:** In development, beta testing

In Greer, South Carolina, PDA Technologies, Inc. is working with the water utility to manage water quality data as well as to assist the utility in improving its security (e.g., biometric login and data encryption are part of the system being provided for Greer by PDA). The components of the data management system include the use of a real-time hydraulic model (e.g., EPANET or PipelinNet), collection of standard water quality parameter data using multi-probe sensors from various vendors (e.g., Hach Pipe Sonde 7 in 1) at 5 locations in the distribution system, geographic information system (GIS) data, and event detection software (Hydra RMS) that interprets the data streams with an adaptive monitoring algorithm. The manufacturer claims that public health surveillance data and consumer complaint information that would be tied to GIS, are examples of additional data streams that can be interpreted by the event detection algorithm, especially with regard to where an anomaly may have occurred. The Hydra RMS is capable of two-way communication with the sensors over secure fiber optics and includes trend analysis tools as it continuously compares real-time water quality data against benchmark water profiles. An alarm notification can be sent via cell phone, pager, or email.

Beta testing of the EDS began in Greer in November 2003. Water quality data are received from 2 Hach Pipe Sondes and are analyzed every 15 seconds. The purpose of the beta testing is to prove the concept that by continuously collecting and comparing water quality data, any deviation from an established baseline can be recognized by an event detection tool. The preliminary results of this study were presented to the Metropolitan Washington Council of Governments in early 2005. The EDS has been able to identify a range of water quality values that represent a baseline profile. The establishment of this baseline has enabled the EDS to recognize as an 'alarm event' an incident in May 2004 where a mislabeled concentration of sodium hydroxide was used during treatment (refer to **Figure 5-1**). This result is an example of the type of dual benefit that can be attained by such a system (Yang, 2005).

The beta testing is still continuing as more baseline data are being collected and adaptive algorithms are being further developed. Future plans include a cooperative research effort with Savannah River National Laboratory to develop improved water quality parameter sensors that would enable the event detection software to recognize patterns that would identify different contaminant groups. There also are plans to conduct surrogate agent exercising of the event detection capability (Page, 2005).

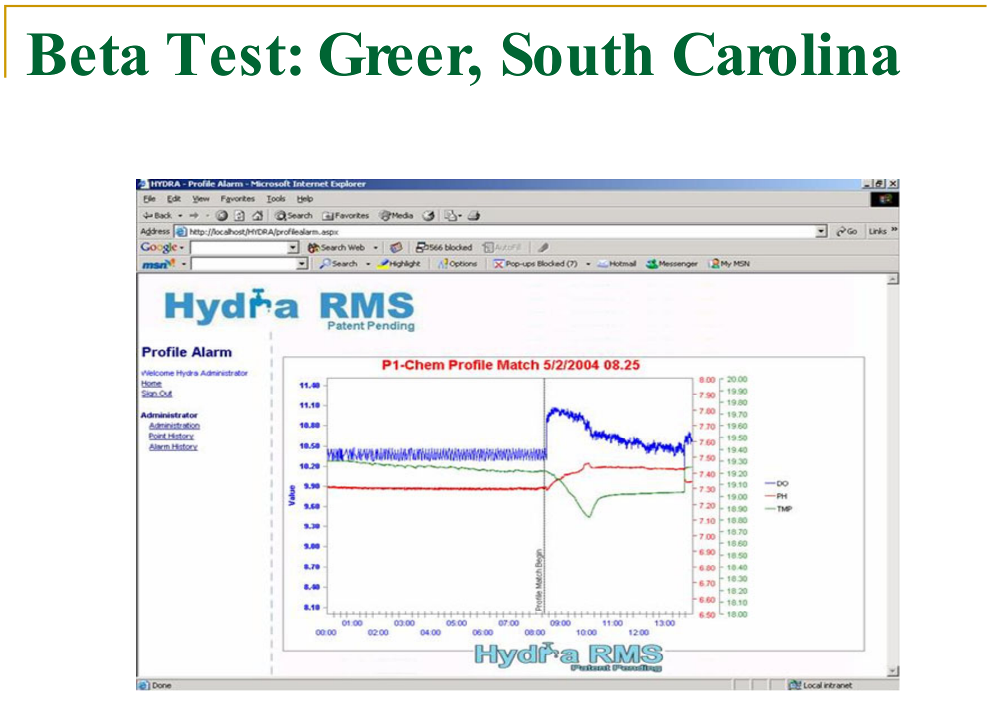

Figure 5-1 Water Quality Profile of an Alarm Event (Image was reproduced with permission from PDA Technologies, Inc.)

5.3 Water Quality Change Detection

- **Government:** USEPA ORD, U.S. Geological Survey (USGS), U.S. Dept. of Interior
- **Research:** Sandia National Laboratories
- **Approach:** Statistical
- **Status:** In development

This project (see Section 3.2.3) is being conducted by Sandia National Laboratories (SNL), EPA's NHSRC, and the USGS). Actual and simulated data are being analyzed by SNL to develop statistically-based event detection algorithms. The use of a simulated data set allows for the insertion of 'events' so that the ROC of the detection algorithms can be evaluated. To date, statistical and data summaries have been processed in detail for development of event detection algorithms (e.g., hourly measurements from 500 to 625 days) for five monitoring stations at the southwest utility for total chlorine, pH, temperature, and specific conductance. Similarly, DO, pH, specific conductance, temperature, and turbidity data collected from October 2003 to September 2004 at three stations in the northeast utility are being studied. One of the three stations is source water, which presents a matrix and set of contamination concerns that differ from the distribution system, and the other two locations are treated, chlorinated water (McKenna, et al., 2005).

Among the preliminary results for this study are the development of two event detection algorithms using measured (i.e., chlorine, temperature, pH, and conductivity) and simulated water quality data and evaluation of the change detection algorithms using ROC curves for simulated data sets. ROC curves were generated for each algorithm for each of the water quality parameters and for a fused data set consisting of all four water quality parameters. Improved ROC curves (i.e., a higher probability of detecting an ‘event’ with a corresponding lower false alarm rate) resulted from the fusing of data, the type of data input more analogous to the WS concept of using multiple water quality sensors. The study also indicated that the event detection algorithm could be ‘tuned’ such that during periods of increased water security concern, the algorithm could operate at a higher probability of detection with the corresponding realization that during this time period, more false positives would occur. Additional investigation is needed to determine how well the simulated data set corresponds to an actual contamination event that could occur in the distribution system (McKenna et al., 2005).

5.4 RODS Project

- **Utility:** Cincinnati, Pittsburgh
- **Government:** NHSRC
- **Research:** University of Pittsburgh, Carnegie Mellon University
- **Commercial:** Hach
- **Approach:** Discrimination and Clustering
- **Status:** Operational, Primarily Syndromic Health Data, OTC Sales

The RODS System is maintained by the RODS Laboratory in Pittsburgh. RODS is a collaborative effort among researchers at the University of Pittsburgh and the Auton Laboratory in Carnegie Mellon University's School of Computer Science. Drs. Wagner, Tsui, and Espino founded the RODS Laboratory in 1999 to investigate methods for real-time detection and assessment of disease outbreaks. Current research interests of the faculty include algorithm development, assessment of novel types of surveillance data, natural language processing and analyses of detectability.

Public health surveillance pilots are being conducted by NHSRC to test the concept of adding near real-time water quality parameters as a data stream into the RODS program. The demonstration project retrospectively evaluates historical water quality data/operating conditions with public health surveillance indicators, tests algorithms that correlate public health and water quality data, and shares anomalies with water utilities in addition to public health officials. Two utilities, Greater Cincinnati Water Works (GCWW) and the Pittsburgh Water and Sewer Authority (PWSA), are currently participating in NHSRC’s RODS study.

The PWSA has submitted historical data to the RODS Laboratory and is currently working with RODS to determine the best way to transfer real-time water quality data to RODS. Data proposed to be transferred are being collected using a Hach Expanded Event Monitor panel and probes, which measure pressure, temperature, conductivity, chlorine, ORP, DO, and turbidity. The GCWW is currently working with the Cincinnati RODS Steering Committee to identify the best way to share anomaly notifications. Issues being negotiated include patient confidentiality and the broad range of GCWW’s service area, which includes numerous health jurisdictions within two states.

5.5 The Electronic Surveillance System for the Early Notification of Community-based Epidemics (ESSENCE) *Project*

- **Utility:** Milwaukee, Washington Suburban Sanitary Commission (Montgomery County, MD), Albuquerque, NM
- **Government:** Walter Reed Army Institute of Research, Defense Threat Reduction Agency (DTRA), NHSRC (evaluating), Department of Defense (evaluating)

- **Research:** Applied Physics Laboratory
- **Commercial:**
- **Approach:** Statistical, Unspecified Artificial Intelligence Algorithms
- **Status:** Operational, Primarily Syndromic Health Data, OTC sales

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) developed the Electronic Surveillance System for the Early Notification of Community-based Epidemics (ESSENCE). ESSENCE collects and analyzes a variety of data sources for the early recognition of abnormal community disease patterns that could result from natural causes or terrorist activities. Similar to the RODS pilot, NHSRC is also conducting a demonstration project with ESSENCE to evaluate the feasibility and predictive value of analyzing water quality data with the health indicator data already available in ESSENCE.

NHSRC is currently negotiating a sole source contract with JHU/APL and pursuing ESSENCE Data Sharing Agreements with Milwaukee Water Works and the Washington Sanitary Sewer Commission (Maryland).

The objectives of both the RODS and ESSENCE projects are consistent with the event detection component of the WS concept of operations shown in Figure 1-1. As indicated in the figure, not only are event detection algorithms being used for water quality and public health data streams, but communication between utility and the public health sector is inherent in the project design. This feature underscores the importance of utility-public health interaction for determining how to best address any identified anomalies as the credibility determination phase of consequence management is triggered. For more information on RODS and ESSENCE, refer to *WaterSentinel System Architecture* (USEPA, 2005a).

5.6 California Space Authority Water Monitoring Project

- **Utilities:** Contra Costa and Southern California Metropolitan Water District
- **Commercial:** Frontier Technologies, Inc.
- **Approach:** Statistical
- **Status:** Planning Stage

This project is being overseen by the California Space Authority (a nonprofit corporation representing the commercial, civil, and national defense/homeland security interests of California's diverse space enterprise community) to investigate the remote transmission of water quality monitoring data using satellite technology that include standard water quality parameter data (chlorine residual, conductivity, TOC, pH, and ammonia). The intent is to evaluate the data using an approach analogous to the use of event detection algorithms developed by Frontier Technologies, Inc. to test jet engine performance. This project is still in the initial planning stages.

5.7 Data Processing and Analysis for Online Distribution System Monitoring

- **Utility (U.S.):** Philadelphia and Oklahoma City
- **Utility (Australia):** Melbourne and Newcastle
- **Research:** AwwaRF and the Commonwealth Scientific and Industrial Research Organization (CSIRO)
- **Approach:** Discrimination, Clustering, and Statistical
- **Status:** Baseline Data Collection and Initial Algorithm Development

Water quality data are being collected on daily, weekly, and seasonal time scales from U.S. and Australian utilities. Among the data being evaluated are pH, ORP, temperature, residual chlorine, conductivity, pressure, flow, and turbidity. The intent of the project is to examine data processing methods that can be used to identify anomalous patterns in online monitoring data associated with specific contamination incidents. Among the data processing methodologies being used are time series analysis and spectral analysis, state-space models, clustering and discrimination, and statistical control charts. The design of this project is consistent with the WS approach (AwwaRF, 2005).

Section 6.0: Other Water Quality Analysis Activities

There are many other activities in which utilities, commercial vendors, and research organizations are working to improve water monitoring capabilities and data analysis. In some cases the data are being evaluated with regard to predetermined set points instead of a more sophisticated approach using event detection algorithms. In other cases, data are being transmitted to commercial entities for evaluation and it is unclear whether set points or algorithms are being used. In many cases, the intent of the organizations is to eventually have the capability to evaluate water quality data with an event detection tool. Examples representative of these types of projects are summarized in **Table 6-1**. Also provided below is a description of companies that offer off-the-shelf tools for event detection and decision support.

6.1 Additional Projects

Table 6-1. Other Water Quality Monitoring Projects

Utility/Location	Research	Commercial	Project Description/Status
Langley Air Force Base/VA	No Research Organizations Involved	Hach	The U.S. Air Force has a Cooperative Research And Development Agreement (CRADA) with Hach Company that involves data collection at Langley Air Force Base. Two event detection stations have been installed and the data are sent to Hach. If funding becomes available, the Air Force would like to conduct pilot tests at 3 of their bases for the collection of water quality data and the use of event detection algorithms, possibly using off the shelf technology from two manufacturers in a parallel testing format (e.g., Hach and PureSense).
Pinellas County/FL	University of South Florida	Constellation Technologies	Pinellas County is working on a project with Dr. Daniel Lim of the University of South Florida and Constellation Technologies to develop an on-line system to detect multiple potential agents simultaneously. An initial pilot field test is anticipated during 2005. This project does not include an investigation of software to handle data output from the system at this point in time.

Utility/Location	Research	Commercial	Project Description/Status
Indianapolis Water Company/IN	No Research Organizations Involved	Clarion Sensing Systems	Clarion Sensing Systems' (Indianapolis, IN) Sentinal™ is a remote computing platform consisting of a hardware and software system for distributed data that features logical data processing at the monitoring sites and compatibility with various forms of wireless and wired data transmission. The system integrates sensor data into a single display that presents information through the internet, a local area network, or a local terminal. The data are presented in a web page format with analytical and historical data storage capability. Each monitoring site has its own Internet Provider address and serves its own web page to allow for specific site monitoring and remote configuration of the water quality profile of the site. The Sentinal™ system can be integrated into an existing system such as the supervisory control and data acquisition (SCADA), and its software is compatible with spiral development approaches, since new sensor technologies can be integrated into the system. Clarion claims to have experience within its software group to develop self-learning analytical software appropriate for event detection and evaluation. This group has developed self-learning software for monitoring water quality in distribution systems and chemical pollution in river water. Clarion is currently receiving water quality parameter data in real time from the Indianapolis Water Company, which is operated by Veolia Water. The Indianapolis system is complex with over 60 sources of ground and surface water (Harmless, 2005).
Contra Costa/CA	No Research Organizations Involved	PureSense	This water retailer has a current project with PureSense (see more information on PureSense below) involving the collection of water quality data and the wireless transmission of the data via an I-node to PureSense for their evaluation (Fowler, 2005).
Copenhagen/Denmark	No Research Organizations Involved	Unknown	An emerging technology project in Denmark is integrating information from existing data sources for a water distribution network in a utility. In the pilot project in Copenhagen, the data from real-time sensors are provided to a SCADA system. The sensor information is stored in a database allowing for analysis (the algorithm approach not known). Automated checks of the system compare against baseline measurement. Data are validated with standard modules which will flag potentially suspect or corrupt data. The project started in 2001 and is expected to be completed in 2005 (USEPA, 2005e).

6.2 Commercial-Off-The-Shelf (COTS) Products

Many companies offer COTS tools for event detection and decision support. Two such examples described above are the Hydra RMS made by PDA Technologies, Inc. (refer to Section 5.2) and the Event Monitor by Hach Company (refer to Sections 3.2.1 and 5.4). Frontier technologies also manufactures EDSs that are used to evaluate jet engine performance, and are being planned for use in a water quality monitoring project (refer to Section 5.6). There are other companies that offer tools or toolsets focused on event detection as applied to water systems; some of these companies and tools are listed below (USEPA, 2005e).

6.2.1 MIKE NET-SCADA

MIKE NET-SCADA unites EPA modeling software and SCADA systems in an effort to optimize system performance to recognize and respond to alarm conditions. The system's online module performs real-time comparisons of the measured and calculated data, automatic data pre-processing for the off-line module, and pressure/flow calculations at any point of the system. The model results are stored back into the SCADA database and the online viewer is used to display detailed model results. In addition, the online module features automatic data validation procedures, in which all measurements are automatically checked and validated with standard modules. These modules will tag questionable data and, if possible, fill in gaps in the time series. This ensures that only validated data will be transferred and used as boundary conditions in the strategic model, decreasing the potential for false alarms (USEPA, 2005e).

MIKE NET-SCADA's off-line module models IF-THEN scenarios, models system breakdowns, and predicts system behavior using demand and control rules prediction. It uses Microsoft Access to store and maintain model alternatives. Coupling the online and off-line results of MIKE NET-SCADA allows the operator to quickly detect abnormalities and help analyze ways in which the abnormality can be remedied or its impact minimized.

Analytical Technology, Inc.'s, (Collegeville, PA) Series C15 Water Quality Monitoring system allows the user to choose those parameters for which monitoring is desired and to integrate those components into a monitoring package suitable for continuous monitoring, alarming, and data collection. System components are currently available for free chlorine, combined chlorine (for chloramine treated systems), dissolved ozone, pH, ORP, conductivity, and temperature. In addition, DO and turbidity modules should be added to the system in the future (USEPA, 2005).

6.2.2 PureSense

EPA has a CRADA with another event detection software vendor (i.e., PureSense) to test its ability to mine data and detect an anomalous event. PureSense is developing a software suite that can be used across a range of water sensors from various manufacturers. The software allows data mining and an 'early warning system' approach to detect anomalous events, but the algorithms are proprietary PureSense offers services to utilities whereby the water quality data are transmitted to their Iowa facility for evaluation. The PureSense System includes the following four components to enable data transmission and analysis: 1) the iNode™ is a remote data communication device that uses cellular and Wi-Fi services to collect monitoring data and send commands to remote sensors, 2) the iWatch™ is an internet data management system that enables the integration of disparate data sets, including data from remote online sensors, 3) the iServe™ performs automated analysis of real-time data, and 4) the AlertNet™ provides automated alerts.

In November 2004, the T&E facility provided PureSense with data similar to what was provided to Hach. A preliminary review of PureSense's performance in evaluating these data was conducted in 2005 and a report is anticipated in November 2005. A prototype event detection software is anticipated in the first quarter of 2006. Puresense is also involved with arranging field trials in collaboration with a number of US water utilities.

6.2.3 AQUIS

AQUIS is a water network management system designed for both on- and off-line, real-time monitoring. The software, produced by Seven Technologies of Denmark, is used to create models to efficiently manage water resources. The models allow utility managers to minimize the impact of operational disruptions in order to maintain continuity and quality of service. The software also allows managers to explore strategies for responding to emergencies, including the introduction of contaminants and increased demands placed on the system by extensive fire-fighting or other surge demands. AQUIS is currently in use in 1,500 cities across the world.

AQUIS offers a Contingency Management Software Package that has five modules designed to establish a point of entry for contaminants, determine a method for limiting the spread of the contaminant, and determine methods to mitigate any harmful effects. The modules include a model manager for GIS data management and a hydraulic module for modeling throughout the distribution system. A water-quality module tracks the chemical composition of the water throughout the system, and a diagnostic module identifies the source of contaminants. Finally, a flushing module facilitates cleaning the distribution network.

6.2.4 Psynapse Technologies

Psynapse has developed the Checkmate Intrusion Protection System for the Technical Support Working Group (TSWG) within the U.S. Department of Defense. This product recognizes when non-typical network activity is a genuine threat. It is intended as a cyber security product by combining computer and behavioral science to conduct real-time assessment of each visitor to a network. Once determining that the behavior of the visitor indicates an attempted security breach, access is terminated. Despite the intended use of this product, the company claims that this neural net technology can be used to determine indications in any data streams that suggest a non-normal anomalous event.

6.2.5 Clarion

Clarion Sensing Systems' Sentinal™ is a remote computing platform consisting of a hardware and software system for distributed data that features logical data processing at the monitoring sites and compatibility with various forms of wireless and wired data transmission. The application of this product to the Indianapolis, Indiana water utility is described above in Section 6.1.

6.2.6 Sensicore

Sensicore, Inc., Ann Arbor, Michigan, manufactures software that can respond to a simple water quality threshold value and then sound an alarm.

6.2.7 Bristol Babcock

The Briston Babcock Company, Watertown, Connecticut, supplies water quality monitoring systems that can integrate into SCADA systems. These systems can provide online continuous monitoring for changes in water quality parameters and also provide other security information to utility operators (Elf, 2005).

Section 7.0: Summary and Preliminary Conclusions

The fundamental concept underlying the WS-CWS is the gathering, managing, analyzing, and interpreting of different information streams in a timely manner to recognize potential contamination incidents early enough to respond effectively. Event detection plays an important role in filtering out the anomalies that normally occur, or which have known causes, and signaling only those events that are likely to be possible contamination incidents. While event detection algorithms have a greater history of use in the public health and cyber security fields, as compared to the environmental field, a number of completed and ongoing laboratory and field research projects exists that involve the use of event detection tools and algorithms to evaluate data obtained from online water quality monitoring in distribution systems.

Preliminary results from the EPA T&E facility studies indicated that a variety of water quality sensors showed noticeable responses to a number of organic and inorganic chemicals, as well as gross biological contamination. A consistent pattern of sensor change was observed with the most consistent change associated with specific conductance, TOC, total/free chlorine, chloride, and ORP. In addition, the SNL study has demonstrated that statistical algorithms are capable of identifying ‘events’ associated with variations of the same types of water quality parameters monitored at the T&E facility. Similarly, the laboratory testing at Colorado State University demonstrated that clustering/classification algorithm-based decision support software could successfully characterize water chemistry based on sensor data and detect anomalies in real time when contaminants of potential concern were introduced into a closed loop system.

Studying the performance of EDSs has progressed from the laboratory to the field. The projects in Charleston and Greer, South Carolina continue to collect water quality data to establish a baseline profile and as cited above, an ‘event’ was detected in the Greer distribution system caused by the improper use of a water treatment chemical. Once sufficient baseline data have been collected in Charleston, the EDS will be evaluated on its response to a simulated attack using modulated fluoride concentrations. Baseline data collection and algorithm development are also continuing at the U.S. and Australian utilities participating in the AwwaRF-CSIRO project. Furthermore, the RODS and ESSENCE projects are examining the feasibility and predictive value of analyzing water quality data and health indicator data, whereby the water and public health sectors will communicate with one another to best inform decision-making when anomalies are identified, as envisioned in the WS concept of operations.

Although the reported results of the above projects are encouraging, they are preliminary at best. Also, while these projects are comprised of portions of the system architecture envisioned for the WS-CWS, none includes the entire set of data streams that comprise WS (i.e., water quality monitoring, consumer complaint surveillance, security monitoring, periodic water quality sampling and analysis, and public health surveillance). Long-term testing within an actual water utility’s distribution system is needed to verify how well event detection can perform both as an individual component of the WS-CWS system architecture, and as part of a system of integrated data streams. The WS pilot program presents an opportunity to conduct this type of testing and obtain knowledge that over time can be used to refine a CWS adoptable by other utilities. Therefore, a field demonstration project that includes all the WS-CWS components is needed to determine whether the system architecture can be replicated at other water utilities, whether false positive rates are kept to a minimum without sacrificing the ability to detect a real event, and whether this approach can be sustainable nationwide. It is anticipated that the capabilities of EDSs will play a critical role in determining the success of these pilots.

Section 8.0: References

AwwaRF, A Summary of AWWA Research Foundation Projects, November, 2004/2005.

Berry, Jonathan W., Hart, W.E., Phillips, C.A., Uber, J.G., and Watson, P.W., "Validation and Assessment of Integer Programming Sensor Placement Models," World Water & Environmental Resources Congress, 2005.

Byer, D. and Carlson, K.H., "Real-Time Detection of Intentional Chemical Contamination in the Distribution System," AWWA Journal, June 2005.

CDC website, 2005: http://www.cdc.gov/niosh/idlh/62748.html

Cook, J., Roehl, E., Daamen, R., Carlson, K., and Byer, D., "Decision Support System for Water Distribution System Monitoring for Homeland Security," American Water Works Association - Water Security Congress, 2005.

Hall, J., Zaffiro, A., Marx, R., Kefauver, P., Krishnan, R., Haught, R., and Herrmann, J., "Parameters for Rapid Contaminant Detection in a Water Distribution System," American Water Works Association - Water Security Congress, 2005.

IPCS Inchem, "Chemical Safety Information from Intergovernmental Organizations - Arsenic" 2005: http://www.inchem.org/documents/pims/chemical/pimg042.htm#DivisionTitle:7.2.1.1%20%20Adults

Intelligent Automation Corporation website, 2005: http://www.iac-online.com/Products/product_detail.asp?product_id=71.

IVD Technology, 2005: http://www.devicelink.com/ivdt/archive/05/03/002.html

Jackson, G., Checkmate Intrusion Protection System: Evolution or Revolution, Pysnapse Technologies. 2003: www.psynapsetech.com.

Kroll, D. and King, K.,, "Operational Validation of an On-line System for Enhancing Water Security in the Distribution System," American Water Works Association - Water Security Congress, 2005.

McKenna, S., Wilson, M., Cruz, V., Madueuke, N., and Srinivasan, S., "Status Report: Task 4, Sandia National Laboratories – EPA NHSRC Inter-Agency Agreement, FY'2005," 2005a.

McKenna, S., Hart, D., and Yarrington, L., "Impact of Sensor Detection Limits on Protecting Water Distribution Systems from Contamination Events," for submission to the Journal of Water Resources Planning and Management, 2005b.

Neil, D.B. and Moore, A. W., 2005, Methods for Detecting Spatial and Spatio-Temporal Clusters. In Wagner, et al., eds., Handbook of Biosurveillance, 2005.

Ostfeld, A. and Salomons, E., "Optimal Layout of Early Warning Detection Stations for Water Distribution System Security," Journal of Water Resources Planning and Management, 130(5), pp. 377-385, 2004.

Personal communication with Robert Fowler, Contra Costa Water Utility, July 5, 2005.

Personal communication with Martin Harmless, Clarion Sensing Systems, July 29, 2005.

Personal communication with Tom Elf, Bristol Babcock Company, August 2, 2005.

Personal communication with Dan Page, PDA Technologies, Inc., August 9, 2005.

Personal communication with Ron Shroder, Frontier Technologies, Inc., September 1, 2005.

Personal communication with John Cook, November 1, 2005.

Personal communication with Dr. Jeffrey Yang, LY International, November 4, 2005.

Srinivasan, S., Madueke, N., McKenna, S., "Literature Review of Approaches for Determining Variability in Background Water Quality Parameters and Change Detection", 2005.

Shmueli, G., "Wavelet-Based Monitoring in Modern Biosurveillance," Working Paper, Smith School of Business, University of Maryland, 2005.

Uber, J., Janke, R., Murray, R., and Meyer, P., "A Greedy Heuristics Model for Locating Water Quality Sensors in a Water Distribution System," Proceedings of the ASCE/EWRI Congress, 2004.

USEPA. *WaterSentinel System Architecture*, 2005a. For Official Use Only.

USEPA. *Online Water Quality Monitoring as an Indicator of Drinking Water Contamination*, 2005b. For Official Use Only.

USEPA. *WaterSentinel Contaminant Fact Sheets*, 2005c. SENSITIVE. For Official Use Only.

USEPA. *WaterSentinel Contamination Incident Timeline Analysis*, 2005d. SENSITIVE. For Official Use Only.

U.S. EPA, Technologies and Techniques for Early Warning Systems to Evaluate and Monitor Drinking Water Quality: A State-of-the-Art Review, 2005e.

Watson, W., Litovitz, T., Klein-Schwartz, W., Rodgers, G., Youniss, J., Reid, N., Rouse, W., Rembert, R., and Borys, D., Annual Report of the American Association of Poison Control Centers Toxic Exposure Surveillance System, Toxicology, 2003.

Wikipedia, "Cyanide" 2005: http://en.wikipedia.org/wiki/Cyanide

Appendix A: Acronym List

ANN	Artificial Neural Networks
AwwaRF	American Water Works Association Research Foundation
COTS	commercial-off-the-shelf
CRADA	Cooperative Research And Development Agreement
CSIRO	Commonwealth Scientific and Industrial Research Organization
CWS	contamination warning system
DO	dissolved oxygen
DSS	Decision Support System
DTRA	Defense Threat Reduction Agency
EDS	Event Detection System
EPA	U.S. Environmental Protection Agency
ESSENCE	Early Notification of Community-Based Epidemics
GCWW	Greater Cincinnati Water Works
GIS	geographic information system
HSPD 9	Homeland Security Presidential Directive No. 9
JHU/APL	Johns Hopkins University Applied Physics Laboratory
NHSRC	National Homeland Security Research Center
ORP	oxidation reduction potential
ORD	Office of Research and Development
OTC	over -the-counter
PWSA	Pittsburgh Water and Sewer Authority
RMS	Remote Monitoring System
ROC	receiver operator characteristic
RODS	Real-time Outbreak and Disease Surveillance
SCADA	Supervisory Control and Data Acquisition
SNL	Sandia National Laboratories
T&E	Test and Evaluation
TOC	total organic carbon
TSWG	Technical Support Working Group
TTEP	Technology Testing and Evaluation Program
USGS	U.S. Geological Survey
WS	WaterSentinel
WS-CWS	WaterSentinel Contamination Warning System
WSD	Water Security Division