Securing Wastewater Collection Systems from Accidental and Intentional Harm: a Cost-Benefit Analysis

Steve Papa, William Casper

Lockheed Martin Aeronautics, Fort Worth, TX
Southern Methodist University, Dallas, TX
{steve.m.papa,william.d.casper}@lmco.com

Tyler Moore

Southern Methodist University,
Dallas, TX
tylerm@smu.edu

Abstract. It has been widely reported that industrial control systems underpinning critical infrastructures ranging from power plants to oil refineries are vulnerable to cyberattack. A slew of countermeasures have been proposed to secure these systems, yet their adoption has been disappointingly slow according to many experts. Operators have been reluctant to spend large sums protecting against threats that have only rarely materialized as attacks. But many security countermeasures are dual-use, in that they help protect against service failures caused both by hackers and by accident. In many critical infrastructure sectors accidents caused by equipment failures and nature do happen regularly, and investment to detect and possibly prevent accidents and attacks could be more easily justified than investments detecting and preventing attacks alone. In this paper, we conduct a cost-benefit analysis of adopting security countermeasures to reduce the incidence of sewer overflows in wastewater collection systems. We estimate the expected annual losses of wastewater systems due to large overflows exceeding 10,000 gallons using publicly-available data on overflows, cleanup costs, property damage and regulatory fines. We also estimate the cost of adopting security countermeasures for wastewater facilities in eight large US cities, finding that for many even a modest 20% reduction in large overflows can make adopting countermeasures cost-effective.

Keywords: process control system security, cost-benefit analysis, wastewater collection systems, security economics

1 Introduction

Supervisory Control and Data Acquisition (SCADA) and Industrial Control Systems (ICS) are widely used to control systems including water, wastewater collection and treatment, oil refineries, oil and gas pipelines, factories, ships, and subways. These systems have evolved from direct human control to computer-based control over the last several decades. Once computer-based control became common practice, a migration from proprietary standards and interfaces to standards-based occurred. Today, many systems have adapted standard wireline and RF physical interfaces, and the TCP/IP protocol is commonly used to move command and status messages within

these systems. To ease management, the trend has been to connect these control networks to the company intranet, which are also normally connected to the internet.

Unfortunately, these systems were not designed to defend against even the simplest network attacks. Operational commands, controller software updates, and operational status messages are not authenticated [1]. As a result, these systems are vulnerable to command injection [2]and middleperson attacks [3]. A Programmable Logic Controller (PLC) attack was at the heart of the Stuxnet virus targeting Iranian facilities [12]. Effectively, Stuxnet used a middleperson attack to change the PLC logic to report normal centrifuge operation to the operator while issuing control operations that damaged centrifuges.

Existing papers on control systems security typically take for granted that an attack will occur and instead focus their efforts on adopting security countermeasures to thwart attacks. However, attacks have been so rare in practice that operators are reluctant to invest in adequate defenses. In this paper, we study one particular critical infrastructure sector – wastewater collection and treatment systems – and investigate whether the expense of security countermeasures can be justified, provided that they can also be used to prevent accidents as well as attacks. We select the wastewater sector precisely because the goal of cyberattacks is the same as a relatively common failure mode: sewer overflows. Furthermore, systems for detecting malicious overflows in wastewater systems can also detect accidental ones.

In Section 2, we outline the threat model for wastewater collection systems and explain how security countermeasures can be deployed on a representative system to detect and prevent sewer overflows. In Section 3, we present a framework for calculating the expected costs of large sewer overflows. We use detailed public data from the California Water Board to estimate the incidence of large sewer overflows. We collate reports of legal settlements to estimate the cost of property damage and consult EPA data on Clean Water Act violations to estimate the cost of regulatory fines, as well as the probability of drawing the ire of regulators. We also provide an estimate for the cost of comprehensive security countermeasures. In Section 4 we use the findings presented in Section 3 to carry out a cost-benefit analysis. We present the net expected utility by comparing costs to the benefit of experiencing fewer overflows. Because wastewater systems vary greatly in complexity, we provide a detailed analysis for eight US cities, finding that some cities are likely to find the costs of security acceptable while others will not. In Section 5 we review related work and we outline opportunities for future research in Section 6.

2 System Model

We first describe the threat model for wastewater collection systems examined in this paper. We then explain the countermeasures that have been proposed and explain how to secure a representative wastewater facility using available countermeasures.

2.1 Threat Model

The threat model we consider includes all sewage system overflow failures occurring at wastewater facilities, regardless of intent. A wide range of common failures include electrical equipment failures (sensors, pumps and control electronics), blockages and structural failures. However, an overflow can also be triggered by an actor with malicious intent. The primary methods of attack on industrial control systems include command injection, service-denial and middleperson [1][2][3]. Regardless of whether the attacker's motivation is wealth, fame, notoriety or terror, invariably the attack's aim is to disrupt the operation of the system. In this paper we do not differentiate Sanitary Sewer Overflows (SSO) from Combined Sewer Overflows (CSO) which may be caused by accidents or attacks. A CSO is a single collection system for both storm water and sanitary wastewater, and a SSO is just wastewater, but for either case we refer to these as a Sewer Overflow (SO). We note that all overflows cannot be prevented even if detected, notably those caused by by excessive storm water inflow.

For wastewater collection systems the most likely and disruptive method of attack is to trigger a sewer system overflow. A famous attack on a wastewater collection system is the Maroochy Water Service Attack [16]. In this attack a SCADA system installer injected commands to a lift station, triggering millions of liters of SOs on at least 46 separate occasions. While the incident persisted for nearly two months, we view it as a single, sustained attack rather than 46 separate attacks since it was carried out by the same perpetrator. The person responsible, Vitek Boten, was caught, sentenced to two years in prison, and was fined to help reimburse cleanup costs incurred due to the attack. His motive was an attempt to get a consulting job with the utility to stop the SO incidents. In general, the PLCs controlling the lift station pumps are the most logical targets for causing overflows. Attack methods could include turning off one or more pumps, under pumping, or repeatedly cycling power to the pumps in order to cause motor damage and malfunctions. Any of these techniques can be executed by modifying the PLC control logic, by injecting malicious control commands, or by modifying operator commands. PLCs are vulnerable to attack because these devices often have no ability to authenticate commands.

2.2 Countermeasures to Prevent Sewage Overflows

Two complementary types of countermeasures have been proposed to protect against attacks on control systems. The more proactive approach is to improve the integrity of control elements such as PLCs or RTUs in a SCADA system and the communications channels they rely on to transmit messages. For example, researchers have proposed retrofitting communications channels with devices to encrypt communications at the link level [22],[23],[24]. Alternatively, integrity can be achieved at the system level by deploying new sensors and PLCs that incorporate trusted hardware (e.g., the Trust Anchors proposed in[4]).While the approach offers a high level of protection against attacks, adding systems such as Trust Anchors are expensive and do not on their own aid in detecting system failures or attacks.

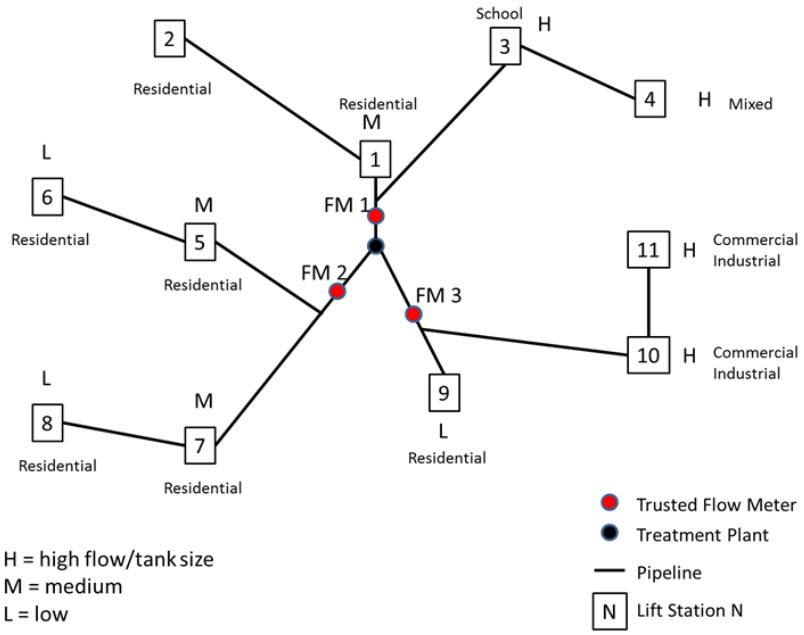

Fig. 1. Reference Wastewater Collection System (Trust Anchors placed in Flow Meters in red).

A second class of countermeasures is much more reactive. Instead of preventing attacks by improving system and communications integrity, one can also detect attacks and failures by monitoring systems for aberrant behavior. Several researchers have proposed intrusion detection systems tailored to industrial control systems ([1],[2],[3],[25],[26],[27]). These detection mechanisms typically must be individually tailored to the system being protected, based not only on the industrial sector but also on the activities of the particular plant.

Approaches that improve system integrity, such as Trust Anchors, can be leveraged as a source of unmodified signal data. For instance, in the case of wastewater collection systems, trusted flow meter data in the reference system can be used with untrusted data from PLCs (tank levels and pump control Status) along with known pump characteristics, pipeline lengths and flow speeds to model the system behavior of the reference wastewater collection system. One or more system models then predict system signals based on known PLC control logic and characteristics of the physical system. In each case failures can be detected when observed behavior deviates substantially from the estimated behavior. Crucially, detection mechanisms can identify failures caused by accident as well as by strategic attackers, often with enough lead time to mitigate widespread calamity.

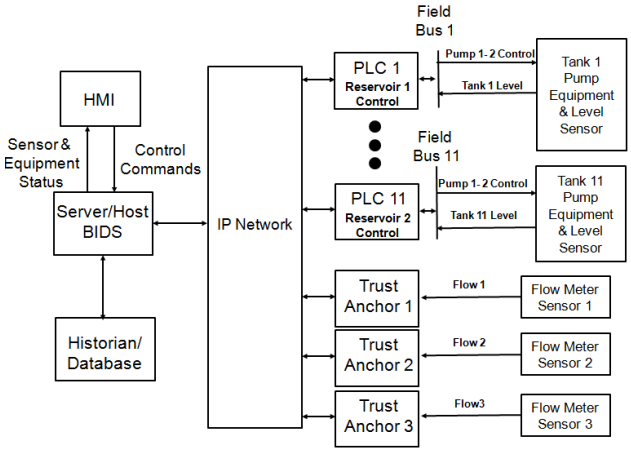

Fig. 2. Reference System Hardware Block Diagram

2.3 Securing a Representative Wastewater System

We make the discussion of available countermeasures concrete by considering how a joint deployment of Trust Anchors and a behavioral IDS (BIDS) might work in a representative wastewater collection system outlined in Figure 1. This system is modeled after a city wastewater collection system that feeds into a single wastewater treatment facility. The system consists of eleven lift stations (indicated by square nodes in the figure), each of which relies on a PLC to control the behavior of two pumps, a storage tank, and a tank-level sensor. Additionally there are three flow meter sensors monitoring the pipelines entering the treatment plant.

Fig. 2 shows the wastewater collection system from a different perspective, using a block diagram of the hardware elements and connections. This diagram also indicates how the security countermeasures can be integrated into the system. We can see from the top right of the figure that each of the eleven PLCs relies on a tank-level signal from a sensor to automatically determine when to turn on and off the two pumps under its control. Rules for turning pumps on and off can vary by lift station PLC, and can be modified by the operator. Additionally, three flow meters transmit back to the operator information on flow rates to the treatment plant. These flow meters can be augmented with Trust Anchors to ensure that the flow rate is not manipulated.

The PLC and Trust Anchors transmit readings to the operator and BIDS via the IP network. The BIDS automatically inspects for suspicious readings that indicate an overflow may occur. The operator, based on either manual inspection of readings or alerts from BIDS, can decide to override PLC commands and turn pumps on or off to react to conditions in the larger system context, or dispatch maintenance crews if he suspects failures. As part of normal operation the system saves operational status from flow meters, PLCs and operator commands to a historian database for future system analysis and as input to the BIDS.

How might an overflow be detected using BIDS and Trust Anchors?

For instance, the pump flows from lift station 1 and 3 combined with the flow between the lift stations and flow meter 1 can be used to estimate flow for comparison to the trusted flow meter 1 measurement. If the estimated and measured flow fall within an acceptable range, then the two lift stations (1&3) and the pipelines that feed flow meter 1 (FM1) are deemed to have not failed or become compromised. After this initial assessment, the flows from lift station 2 based on pump status of LS 1 and 2 can be used to estimate the level of LS 1's tank. The integrity of the pipelines and lift station 2 is verified if the reported level and estimated level are acceptably close. Lift station 4 and its pipelines to lift station 3 can be verified in a similar fashion. This process is repeated for the collection segments monitored by FM2 and FM3.

When the estimated signal (flow or tank levels) deviates substantially from the reported signal values, then a blockage, pump failure, PLC failure or attack on the lift station could be to blame. BIDS can raise an alarm and notify the operator. Furthermore, BIDS can often determine which component has failed. For instance, if the error in predicted versus measured flow is approximately equal to the output flow of a feeding pump, then that pump must have failed. If instead all four pumps report similar output rates, then we can only conclude that one of the four pumps has failed without pinpointing which one. When the error is not proportional to a pump's output, then there may be a blockage or structural failure between the lift station and the flow meter, or between lift stations.

We have verified that BIDS is effective at detecting failures by running Matlab simulations of the reference system with three trusted flow meters. We also verified that BIDS can detect simulated PLC attacks and pump failures, giving us confidence that a pump failure or other medium-to-large flow problem can be detected, isolated and reported.

While the combination of BIDS and Trust Anchors offers a powerful way to detect failures and cyberattacks early, the question remains whether it is economically feasible to deploy these countermeasures. We set out to answer that question next.

3 Empirical Estimation of Expected Sewage Overflow Costs

3.1 Framework for Sewage Overflow Costs

When a utility experiences a sewer overflow a number of costs are incurred. We follow the approach of Anderson et al. 2012 [28] and divide these costs according to direct losses experienced by the utility, indirect losses imposed on society, and defense costs that mitigate SOs. Direct losses associated with an overflow incident include cleanup costs, collateral property damage (buildings/environmental/property), regulatory fines and penalties, and adverse health impacts sustained by the victims. Additional indirect losses associated with an incident include lost business, environmental impacts and distress on the people suffering from the overflow.

Table 1 enumerates the different types of costs and assigns them to the appropriate category. In the following subsections we present appropriate data sources in order to

Table 1. Cost breakdown of sewage overflows, along with a note whether data is available to estimate costs (✓: available, ×: not available) and the corresponding variable used in the model.

Cost Category	Data?	Variable
Direct losses		
Cleanup costs	✓	c_{cln}
Property damage	✓	c_{dam}
Regulatory costs (e.g., fines, settlements)	✓	c_{EPA}
Lost business for victims	×	
Victim health costs	×	
Indirect losses		
Lost business for non-victims	×	
Broader environmental impact	×	
Psychological distress	×	
Defense costs		
Integrity protection (e.g., Trust Anchors),	✓	c_{sec}
incident detection (Behavioral BIDS)		

estimate the costs when possible. We use the derived data to calculate the expected annual cost of SOs using the following formula:

$$ALE_0 = E(n_{SO}) \times (c_{cln} + p_{dam}c_{dam} + p_{EPA}c_{EPA}) \quad (1)$$

Here E(n_{SO}) represents the expected number of SOs exceeding 10,000 gallons for a utility per year, which is computed in Section 3.2. For each expected overflow, we tally the cost of cleanup (c_{cln}), the expected cost of property damage ($p_{dam}c_{dam}$), and the expected cost of regulatory penalties ($p_{EPA}c_{EPA}$), which are computed in Section 3.3.

3.2 Estimating the Incidence of Sewage Overflows

We exhaustively searched public sources in order to estimate the historical probability of cyberattacks targeting wastewater collection systems. However, the search turned up only one well-publicized attack on a wastewater system worldwide [16]. The Repository for Industrial Security Incidents (RISI) provides reports around 17additional water and wastewater system incidents between 2000 and 2009. Based on the limited publicly available information, these incidents appear to have been triggered by a mix of software and hardware equipment failures, system failures, network failures, sabotage, and operator or maintainer errors. Notably, there is no indication that the reported incidents were actually cyberattacks on wastewater collection systems [7][34]. Consequently, we conclude that although wastewater systems are vulnerable to attack, the empirical probability of a cyberattack is extremely low based on its past incidence. This is consistent with the finding that cyberattacks on SCADA systems in general have also been very rare to date, even if those attacks that have been executed,

such as Stuxnet, have attracted significant notoriety [7][12]. The extremely low incidence of published SCADA cyberattacks in general points to a similarly low probability that wastewater treatment facilities in particular would be targeted. Of course, the absence of attacks in the past is no guarantee that such attacks could never happen in the future, but it does mean that investment in systems to ward off attacks are unlikely to be justified on a cost-benefit basis of preventing malicious attacks alone. To that end, we instead empirically examine the probability of a non-malicious sewage overflow in order to see if countermeasures that protect against malicious and accidental overflows could be economically justified.

The most reliable and comprehensive data on overflows in the US comes from the California Water Board, which reports that 4,738 SOs occurred in the state during FY 2011 [20]. However, overflow size varies greatly. Just 2% of overflow incidents – those exceeding 10,000 gallons – account for 84% of the total volume of spilled sewage. Of course, not all overflows could be prevented by early detection. The California Water Board distinguishes between four broad categories of overflows: operational, condition, structural, and other. Operational overflows arise from acts of nature such as debris clogging systems, while condition failures can be caused by outdated infrastructure. These types of overflows are unlikely to be prevented by early detection. However, structural overflows where components such as pump stations fail should be detectable before an overflow exceeds 10,000 gallons. Likewise, most failures in the "other" category are detectable, such as operator and maintenance errors.

We choose to only track the incidence of SOs exceeding 10,000 gallons. As stated above, these account for the vast bulk of overflow volume, and so are the best targets for prevention. California experienced 96 SOs exceeding 10,000 gallons in FY2011, approximately 46 of which should be detectable by a system like BIDS. A 10,000 gallon SO is also a reasonable threshold for detecting failures with systems such as BIDS because sensor measurements (flow meter and tank level) need to be filtered to reduce false alarms and instantaneous measurement errors.

While there are approximately 16,000 sewer systems in the US, they vary widely in size, and therefore, in their likelihood of experiencing a large overflow. Fortunately, the California Water Board reports that there are 110,593 total miles of sewer lines. We can use this to compute an estimate of the number of detectable overflows in SO exceeding 10,000 gallons using the following formula as a function of sewer line miles m:

$$n_{SO}(m) = \frac{46\ detectable\ large\ SOs}{110{,}593\ sewer\ miles} = 4.16 \times 10^{-4} \times m \qquad (2)$$

We searched public records in order to determine the average number of miles of sewer mains for cities over 100K population. We started with US Census data available from the 2010 census [33] to obtain a list of 273 cities that met the criteria. For each city we then started with the city's specific website. Often the cities have sewer information categorized in a water, wastewater, or public works section of their websites. These sections were usually associated with departments or public utilities portions of the websites. Searches were made in each city's website using keywords of "sewer", "sewerage", "sewer mains", "sewer miles", and "wastewater treatment". Frequently there were statistical summaries on the amount of water lines and sewer lines in miles for each city area. Some cities have relegated all of their water and wa-

ter treatment efforts to separate water and wastewater utility organizations. These were independently controlled entities with their own websites, and for cities utilizing these entities the amount of sewer mains serviced were often documented there. This results in data for 135 of the 278 cities, with an average of 1300 miles of sewer mains per city over 100K population reported for these 135 cities. Using this value of 1300 miles in the equation for number of detectable overflows previously discussed yields the following: $n_{SO}(1300) = 0.541 = n_{SO}$.

3.3 Estimating the Cost of Sewage Overflows

We now review available data on each of the types of costs associated with sewage overflows in order to derive robust estimates. We follow the structure set out in Table 1: direct losses resulting from the overflow, including cleanup costs, property damage, and regulatory fines; and indirect losses affecting others not directly impacted by the overflows.

Direct Losses – Cleanup

For spills that only require cleanup (no property damage) the best public data available is from an EPA report in 2000 [18]. Based on this EPA report the average cost of cleanup (labor and materials) was approximately $2130 per event in 2000. Adjusting for inflation from 2000 to 2012 the average cost per event is estimated to be $2,854 [19]. However, this is not an appropriate estimate of the cleanup costs for large SOs, since larger events are much more likely to cause property damage.

The California Sanitation Risk Management Association (CSRMA) is a large non-profit organization that includes forty city and regional wastewater utilities and it provides credible wastewater cleanup information. According to CSRMA the average cleanup cost is $22,554 per SO event, based on 133 claims made in 2011 [17]. Even so, this data is also likely an underestimate since the CSRMA is a cooperative of mostly smaller wastewater districts, but since this is the best data we have, we set c_{cln}=$22,554.

Direct Losses – Property Damage

Spills that result in property damage have costs that vary widely depending on the location and volume of the SO. The majority of all property damage claims are paid by insurance, and the remaining claims are settled in lawsuits. Many large wastewater utilities are city or county government owned and self-insured. Smaller utilities can either purchase insurance or join a larger pool to distribute their individual risk of a large loss.

The last category of costs is legal and non-insurance settlements. After an exhaustive internet search including a review of lawsuit data associated with SOs, public utility records, and media reports for property damage settlements only 19 claims were found. To find these incidents a significant number of key word searches on google.com and bing.com were performed and approximately the highest 100-400 web sites for each search term were reviewed for related loss data. Search terms in-

Table 2. Non-insurance settlement costs.

Cases	Property Damage Estimates from Non-Insurance Settlements			
#	minimum	median	mean	maximum
19	$11,331	$151,000	$1,403,345	$11,600,000

cluded "sewer overflow", "wastewater overflow", SSO, CSO, along with "property" and "property damage". All unique instances of reported property damage settlements were used to create a database from these data sources. These incidents included damages sustained by single homes, multiple homes, an apartment complex, and a business, and this database is summarized in Table 2. The average cost due to legal and non-insurance settlement is $1,403,345.

How often does property damage occur? 9 of the 19 cases with confirmed property damage were reported in 2011-2012. We can establish a lower bound for the probability of property damage occurring as follows:

$$p_{dam} \leq \frac{9\ SOs\ with\ damage\ in\ US\ in\ 2011-12}{\frac{46\ SOs\ in\ CA\ in\ 2012}{12.1\%\ US\ pop.in\ CA} \times 2\ yrs} = 1.18\% \quad (3)$$

Of course, this likely a gross underestimate, since a 10,000 gallon sewage spill is very likely to damage property. However, many property damage claims are settled out of court in such a way that does not attract news coverage. Consequently, the true value for p_{dam} is somewhere between 1.18% and 100%. While we lack data to support it, a more plausible conservative estimate for the probability of incurring property damage in a large overflow is for property damage is $p_{dam} = 25\%$.

We also note that the RISI incident information provides independent cost estimates associated with six incidents. Three incidents cost less than $10K, two incidents cost between $10K and $100K and one incident cost over $10M [7][34]. This is fairly consistent with the data we collected.

Direct Losses – Regulatory Fines

The Environmental Protection Agency (EPA) levies fines on cities, municipalities, and special utility districts for violations of the 1972 Clean Water Act (CWA), which includes sewage overflow incidents. In addition to federal penalties levied by EPA, the legal settlement of these violations often includes the offender agreeing to environmental enhancement projects, referred to as Supplemental Environmental Policy (SEP) agreements. SEP costs are typically a small fraction of the total estimated cost to comply with the EPA's recommendations for system upgrades. SEPs can be thought of a way to compensate society for the harm imposed by the CWA violation.

The EPA maintains a database of the CWA infractions with associated costs for resolution of violations from 2001-2012 [8]. Table 3 summarizes the penalty, SEP, and compliance cost data associated with CWA violations. The table presents summary statistics from 46 of the 85 total CWA violations specifically caused by SOs. We obtained this subset by filtering the search results from the EPA database on two parameters. First, we set the Facility Characteristics SIC Code to 4952, which is the

Table 3. Regulatory fines & compliance costs for Clean Water Act violations arising from SOs.

46 Incidents 2001-2012	**Regulatory Fines and Compliance Costs**			
	minimum	median	mean	Maximum
Federal Penalties	$0	$122,500	$341,205	$2,200,000
SEP Value	$10,800	$305,000	$2,546,344	$42,000,000
Compliance Costs	$0	$72,000,000	$679,600,000	$4,700,000,000

Standard Industrial Classification (SIC) code for Sewerage Systems. Second, we set the Case Attribute Primary Law to CWA – Clean Water Act. All other parameters were left to their default values. Additionally, we validated many of the EPA fines by inspecting independent legal settlement information [10].

The average EPA penalty assessed was $341,205, and average SEP was $2,546,344. We can safely assume that EPA fines would only occur in the case of large overflows exceeding 10,000 gallons. Consequently, we can estimate the probability that a large overflow receives an EPA fine as follows:

$$p_{EPA} = \frac{46\ EPA\ violations\ in\ US\ in\ 2000{-}11}{\frac{46\ SOs\ in\ CA\ in\ 2012}{12.1\%\ US\ pop.in\ CA} \times 12\ yrs} = 1.01\% \quad (4)$$

Indirect Losses

The primary indirect loss associated with SOs is pollution of the environment. For business that are impacted by a SO, there may be loss of business due to cleanup, lost customers due to health concerns, and damaged inventory or equipment. Additionally there is a significant amount of psychological distress on the victims of a SO due to the overflow or concerns of future overflows.

Overflows often reach rivers, water sheds or the ocean causing additional environmental losses and health hazards. Lakes, rivers, beaches and shell fisheries can become contaminated from SOs. A SO can cause contamination of water that is well described in a 2004 report to the US congress. The list of contaminants includes a long list of microbial pathogens, viruses, parasites, metals, synthetic organic chemicals, toxins and bacteria [21]. All of these indirect costs are difficult to quantify, but a small percentage of these costs are likely captured in the lawsuit loss and regulatory fine data provided with the direct loss data.

Finally, it is worth noting that typical home insurance policies do not normally cover sewer overflow damage. An additional rider is usually needed to provide this coverage, and the cost of this rider is typically $40-$60/year [14][15].This low cost indicates that the insurance companies have determined there is low risk for sewer overflow damages that are the home owner's responsibility.

3.4 Estimating Defense Costs

Table 4 provides a summary of the estimates of the non-recurring development and annual reoccurring costs associated with the addition of Trust Anchors and Behavioral Intrusion Detection Systems (BIDS) hardware and software to the reference system

Table 4. Cost estimate for reference system development and installation of TA and BIDS.

Non-reoccurring Costs	TA HW, SW, BIDS and Training	$79,095
Reoccurring (annual) Costs	TA HW, SW, Training, and Key Updates	$ 7,850

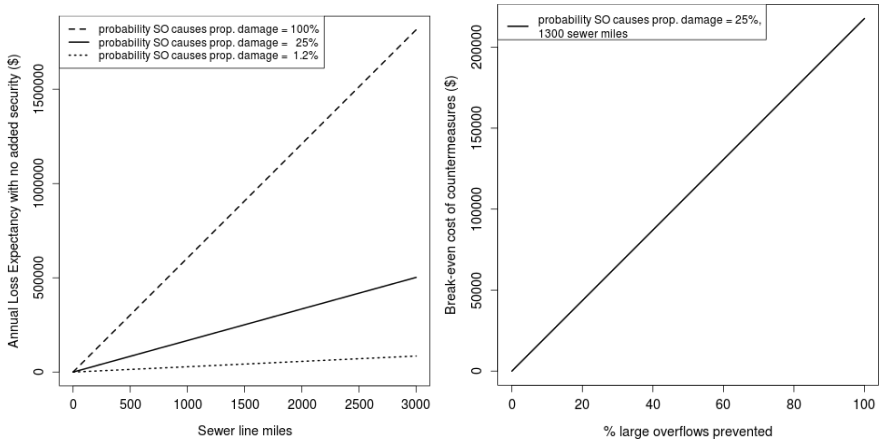

Fig. 3. The expected loss ALE_0 for a utility expressed as a function of the length of sewer lines managed, shown for varying probabilities of experiencing property damage (left); the break-even costs of countermeasures expressed as a function of the percentage of large overflows prevented through improved detection (right).

outlined in Section 2.3. The reoccurring expenses include TA hardware, system specific TA and BIDS software, and initial operator training. Based on Matlab system simulation that the minimal number of Trust Anchors required to support the BIDS is three (flow meters). With these three trusted flow measurements and data from each PLC controlling a lift station the BIDS can detect failures in operation due to pump failures of operational inconsistencies in data due to a cyberattack on the PLC.

4 Cost-Benefit Analysis

4.1 Security Metrics

Having estimated various costs associated with sewer overflows and associated countermeasures, we are now in a position to evaluate the effectiveness of the countermeasures. We can calculate today's expected annual loss by inserting the empirically derived estimates into Equation (1). Doing so, we find that $ALE_0 = \$217{,}675$ for a utility managing 1300 miles of sewer lines, the average for large US cities.

Figure 3 (left) plots the expected annual loss ALE_0 as a function of the number of miles of sewer pipe managed by a utility. Three lines are included in the plot for different probabilities of experiencing property damage p_{dam}. The dotted line uses the lower bound estimate of 1.18%, showing expected losses ranging from $3,000 for a

Table 5. Estimated costs of countermeasures for wastewater systems in 8 cities plus the reference system shown in Figure 1. The annual expected loss is shown for each city, as well as the expected net benefit of countermeasures based on a 40% risk reduction and $p_{dam} = 0.25$.

System	Pump. Stns	# WWTP	Lift/ WWTP	Min #TA	Cost Fac.	Initial Cost	Op. Cost (20yr)	Cost/ year	Sewer Miles	ALE_0	$ENBIS$
Ref.	11	1	11	3	1	$79K	$157K	$20K			
City 1	16	4	4	5	2	$158K	$314K	$39K	2125	$356K	$103K
City 2	25	1	25	7	3	$237K	$471K	$59K	1800	$301K	$62K
City 3	56	2	28	16	6	$475K	$942K	$118K	993	$166K	($51K)
City 4	83	2	42	23	8	$633K	$1,256K	$157K	1600	$268K	($50K)
City 5	93	14	7	26	9	$712K	$1,413K	$177K	6000	$100K	$225K
City 6	116	1	116	32	11	$870K	$1,727K	$216K	3100	$519K	($9K)
City 7	154	11	14	42	14	$1,107K	$2,198K	$275K	6700	$1,122K	$174K
City 8	200	4	50	55	19	$1,502K	$2,983K	$373K	895	$150K	($314K)

town with 100 miles of sewer lines to $85,000 for a city with 3000 miles of sewer lines. By contrast, with the upper bound of probability $p_{dam} = 1$, a city with 3000 miles of sewer faces an expected annual loss of $1.8 million. The loss falls to $500 thousand when there is only a 25% chance that an overflow causes property damage.

The benefits of security are notoriously difficult to measure. We follow the practice in the security economics literature by measuring security benefits as the amount of loss avoided. We can calculate the expected annual loss when investing in countermeasures as:

$$ALE_S = (1 - r) \times ALE_0 \tag{5}$$

Here r represents the fraction of large overflows that can be prevented through early detection. We have no reliable data on what this rate should be, since few utilities have adopted failure detection systems. A value of $r = 0.4$ seems plausible; we examine the impact of varying r on the viability of countermeasures below.

But first, we must also include the cost of countermeasures in determining whether a countermeasure is worth adopting. To that end, we must calculate the expected net benefit of security:

$$ENBIS = ALE_0 - ALE_S - c_{sec} \tag{6}$$

A standard approach in cost-benefit analysis is to identify the cost at which a countermeasure breaks even. To do that, we can set $ENBIS = 0$ and solve for c_{sec}. Substituting from Equation (5), security countermeasures break even when $c_{sec} = r\, ALE_0$. Figure 3 (right) plots the break-even cost for a city with average sewer lines and a 25% chance of property damage as a function of the overflow prevention rate r. For a 40% prevention rate, the breakeven cost is just under $100,000 per year, rising linearly. Countermeasures that cost more than $200,000 require complete prevention to be economically viable for a typical utility.

But what exactly constitutes a typical utility, and how do the estimated costs of countermeasures stack up? We next examine the expected costs facing U.S. cities.

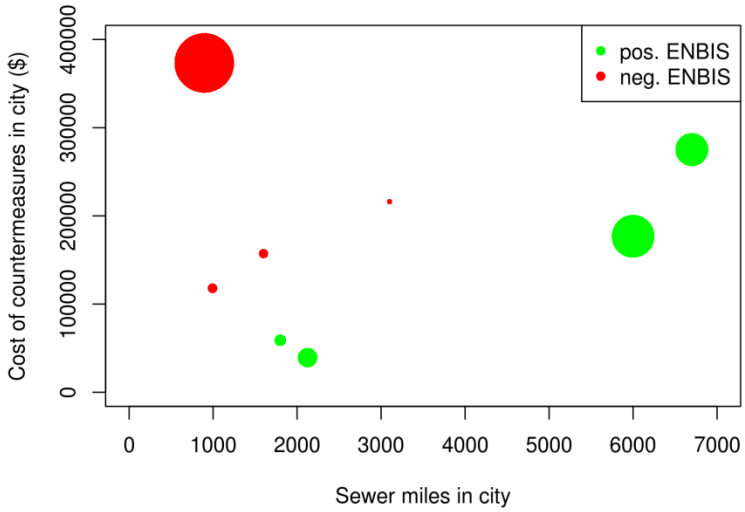

Fig. 4. Comparing the number of sewer miles to the costs of deploying countermeasures for the 8 representative cities. The size of the point indicates the magnitude of the net benefit (or cost) of deploying countermeasures.

4.2 Cost-Benefit Analysis for Eight Large US Cities

To determine the viability of adding protection measures it is necessary to estimate the costs of adding this hardware and software to the systems of several cities (Atlanta, Baltimore, Los Angles, New Orleans, New York, Orlando, San Francisco, and Washington D.C.). We selected these eight major US cities in order to estimate the cost of protection for a diverse range of system layouts. We compared each city's system layout to the reference system and scaled the protection equipment costs appropriately. This scaling was based on the number of standalone lift stations and combined lift stations/Wastewater Treatment Plants (WWTPs) identified. Because SCADA systems have 20-30 year hardware replacement lifecycles, we adopted a 20-year lifecycle to calculate total cost and average operational cost per year of the countermeasures.

Table 5 presents the results. The annual costs varied considerably, ranging from $20,000 to $373,000. The table also shows the sewer miles managed by each city. While positively correlated, some cities have far more pumping stations and lifts than the length of sewer lines might suggest. This invariably has to do with local geography and the age of the city. One consequence of this, however, is that protective countermeasures are much more economically feasible for some cities than for others. Table 5 also includes an estimate of the expected annual loss without added protection (ALE_0), which is tied to the sewer line length in each city. Finally, we can compute whether the protection mechanisms are viable using $ENBIS$. For a 40% risk reduction and 25% probability of suffering property damage, four of the cities should invest and four should not.

Table 6. Return on security investment (ROSI) for 8 cities varying the percentage of overflows prevented using countermeasures (r). Positive numbers indicate a viable investment.

	10%	**20%**	**30%**	**40%**	**50%**	**60%**	**70%**	**80%**	**90%**	**100%**
City 1	-9	81	81	262	353	443	534	624	715	805
City 2	-49	2	2	104	156	207	258	309	360	411
City 3	-86	-72	-72	-44	-30	-15	-1	13	27	41
City 4	-83	-66	-66	-32	-15	2	19	36	53	70
City 5	-43	14	14	127	184	241	298	354	411	468
City 6	-76	-52	-52	-4	20	44	68	92	116	140
City 7	-59	-18	-18	63	104	145	185	226	267	308
City 8	-96	-92	-92	-84	-80	-76	-72	-68	-64	-60

Figure 4 explores the relationship between sewer miles, cost of protection, and $ENBIS$ visually using a scatter plot. Each point represents a city; points to the right indicate the city has more sewer miles and points towards the top have more expensive costs of protection. The points are scaled according to the size of the gain (or loss) from investing in protection. We can see that points in the upper left quadrant fare worst, which makes sense since their costs are high but their risk of large overflows are lower. But the figure also shows that that even relatively expensive protection mechanisms can be economically viable if the risk of a large overflow is substantial, as is the case for cities 5 and 7. Individual utilities can of course inspect their own overflow history in order to determine if the risk of overflows is substantial enough to invest in better detection.

As mentioned above, it is not known how effective the protection mechanisms will be in preventing overflows. We can deal with this by measuring the viability of countermeasures for a wide range of detection rates. Because each city faces different costs and benefits, it can be helpful to normalize the benefit measure. We use the standard metric called $ROSI$ (return on security investment), defined as:

$$ROSI = \frac{ENBIS}{c_{sec}} = \frac{r \times ALE_0 - c_{sec}}{c_{sec}} \tag{7}$$

Table 6 plots $ROSI$ for detection rates ranging from 10% to 100% effectiveness for each of the 8 cities. Positive percentages indicate that investing in protection is worthwhile, while negative numbers suggest that the added protection costs too much compared to the reduction in risk. We can see that if the countermeasures reduce large overflows by 10%, none of the cities will find the protection cost-effective. Once 20% of overflows are prevented, the protection is viable for cities 1,2 and 4. Notice that city 8's infrastructure is so complex for its size that even preventing *all* sewer overflows would not make the countermeasures cost-effective. Consequently, we can safely conclude that protection mechanisms may be reasonable for some, but never all, wastewater utilities.

5 Related Work

Critical infrastructures are susceptible to disruption. While failures triggered by accidents or acts of nature have long presented a challenge, during the past decade researchers and practitioners have begun warning of the threat from malicious parties to exploit vulnerabilities in the computer systems that control operations [29]. The vulnerabilities in process control systems affect a broad range of industries, including electric utilities, oil refineries, and wastewater collection and treatment systems.

Computer scientists have proposed two main ways to protect against threats arising from these vulnerabilities. First, proposals to improve the integrity of the systems and communications channels have been made, ranging from less-expensive retrofits (e.g., [22],[23],[24]) to more comprehensive replacement solutions (e.g., Trust Anchors [4]). A second approach has been to build systems that can detect attacks and hopefully stop them from succeeding, borrowing ideas from intrusion detection systems used in IP networks [1],[2],[3],[25],[26],[27]. When applied to SCADA, intrusion detection systems can identify unauthenticated command injections, response injections and denial-of-service attacks [2]. Papa et al. devised a risk assessment methodology to determine the most vulnerable assets within a system [5], and then used it to recommend the least disruptive and most cost-effective configurations of Trust Anchors and intrusion detection systems to detect attacks [6]. We leveraged this approach in estimating the configuration required to secure wastewater systems for the eight representative cities in Section 4.2.

The literature referenced above has argued that vulnerabilities in process control systems, once found, must be fixed. Given our society's reliance upon critical infrastructures to function, this is an understandable position. However, the proposed countermeasures come at substantial cost, and operators within the industry have pushed back, arguing that attacks exploiting these vulnerabilities are an exceedingly rare occurrence, if they happen at all. Recent research in security economics can shed light on this problem in two ways. First, researchers have argued that insecurity is a form of negative externality, which suggests that firms often lack an appropriate incentive to improve security [30]. In the context of our study, operators may not wish to invest enough in protecting against insecurity when the harmful consequences of an attack are primarily borne by society. The second way security economics can help is to quantify the costs of insecurity, as well as the benefits of improved security. In Section 3 we apply the cost framework used in [27] in the context of cybercrime to estimate the costs associated with sewage overflows. We also used loss expectancy and return on security investment metrics consistent with [31].

Cost-benefit analysis has been applied in the context of combating terrorism. For example, Stewart and Mueller conduct cost-benefit analysis for securing bridges from terrorist attack [13]. Similar to our study, this threat has been realized very rarely if at all; unlike wastewater systems, the authors could not rely upon a non-malicious threat of failure to estimate probabilities. The same authors also examined aviation security countermeasures [32], comparing different techniques for effectiveness. We adopt a similar approach in piecing together empirical estimates of probabilities and costs using public data sources.

6 Concluding Remarks

Detecting non-malicious failures could make security countermeasures economically viable for wastewater collection systems. Absent improved failure detection mechanisms, we estimate that the expected annual loss due to sewer overflows exceeding 10,000 gallons is approximately $200,000 for US cities with populations exceeding 100,000. Each utility's cost depends on the complexity of its collection system. Some will find that investing in security mechanisms that improve early detection of overflows is justified.

There are several limitations to the current study that could be improved. For one, the robustness of the data may be improved. We are limited by what can be obtained through public resources. Notably, we remain uncertain of the probability that large overflows will cause property damage, and our estimates for cleanup costs are likely understated. We have chosen to use average values of cost estimates, even though the distribution of losses is highly skewed. Using median values instead would have biased the estimates downward, but in the end we decided that using mean values was more appropriate given the risk aversion many operators exhibit. Finally, while we did not account for attacks that triggered overflows due to their historical rarity, we would like to be able to derive some measure of their expected cost, which may be substantial.

We are optimistic that the same approach we have taken in this paper – justifying security improvements by quantifying their ability to prevent accidents – can be usefully applied to other critical infrastructure sectors. For instance, higher value assets such as petroleum refineries or power plants seem like promising targets for cost-benefit analysis. The number of these facilities is more limited (144 refineries and 6,313 electrical power plants in the U.S. [9]) as compared to over 16,000 wastewater treatment facilities. Available incident data on these industries show a higher incident rate and higher cost per incident which may justify investments that prevent failures regardless of intent.

References

[1] VERBA, J. AND Milvich, M. 2008. Idaho National Laboratory Supervisory Control and Data Acquisition Intrusion Detection System (SCADA IDS), *2008 IEEE Conference on Technologies for Homeland Security*, Waltham, MA, USA, May 2008, IEEE, 469-473

[2] GAO, W., MORRIS, T., REAVES, B., RICHEY, D. 2010. On SCADA Control System Command and Response Injection and Intrusion Detection, *IEEE eCrime Researchers Summit (ecrime)*, Dallas, TX, USA, October 2010, IEEE, 1-9

[3] PAPA, S., CASPER, W., NAIR, S. 2012. A Transfer Function based Intrusion Detection System for SCADA Systems, *IEEE International Conference on Technologies for Homeland Security (IEEE HST)*, Waltham, MA, USA, November 2012, IEEE, 93-98

[4] PAPA, S., CASPER, W., NAIR, S., 2011. Placement of Trust Anchors in Embedded Computer Systems, *IEEE Hardware Oriented Security and Trust (HOST)*, San Diego, CA, USA, June 2011, IEEE, 111-116

[5] PAPA, S., CASPER, W., NAIR, S. 2011. Availability Based Risk Analysis for SCADA Embedded Computer Systems, *The 2011 World Congress in Computer Science, Computer Engineering, and Applied Computing (WorldComp11)*, Las Vegas, NV, USA, July 2011, WORLD ACADEMY OF SCIENCE, 541-547

[6] CASPER, W., PAPA, S., NAIR, S. 2012, Security Fusion Implementation and Optimization in SCADA Systems, *IEEE International Conference on Technologies for Homeland Security (IEEE HST)*, Waltham, MA, USA, November 2012, IEEE, 620-625

[7] TUDOR, Z. AND FABRO, M. 2010. What Went Wrong? A Study of Actual Industrial Cyber Security Incidents, SRI International, *Industrial Control Systems Joint Working* Group (ICSJWG), San Antonio, TX, USA, April 2010, The Department of Homeland Security United States Computer Emergency Readiness Team (US-CERT),
http://www.us-cert.gov/control_systems/icsjwg/presentations/spring2010/02%20-%20Zach%20Tudor.pdf

[8] ENVIRONMENTAL PROTECTION AGENCY, 2012.
http://www.epa-echo.gov/echo/compliance_report_sep.html

[9] US ENERGY INFORMATION ADMINISTRATION, 2012.
http://www.eia.gov/dnav/pet/pet_pnp_cap1_dcu_nus_a.htm
http://www.eia.gov/electricity/annual/html/table5.1.cfm

[10] LEXOLOGY, 2012.
http://www.lexology.com/library/MoreLikeThis.aspx?g=7183714f-d1b3-439c-bbbc-259f0c1d5074&SameJurisdiction=1
LAWYERS AND SETTLEMENTS, 2012.
http://www.lawyersandsettlements.com/settlements/09802/sewer-system-improvements.html
LAWYERS AND SETTLEMENTS, 2012.
http://www.lawyersandsettlements.com/settlements/02648/sewage.html
LAWYERS AND SETTLEMENTS, 2012.
http://www.lawyersandsettlements.com/settlements/02648/sewage.html

[11] NIST Summary of Maroochy Incidents:
Nist.gov, 2008. http://csrc.nist.gov/groups/SMA/fisma/ics/documents/Maroochy-Water-Services-Case-Study_report.pdf

[12] New Scientist Stuxnet Article:
New Scientist.com, 2011. http://www.newscientist.com/article/dn20298-stuxnet-analysis-finds-more-holes-in-critical-software.html

[13] STEWART, M. and MUELLER, J. 2011, Assessing the Risks, Costs, and Benefits of Counter-Terrorism Protective Measures for Infrastructure, The CIP Report, *Center for Infrastructure Protection and Homeland Security*, pp 3-5, 31, November 2011

[14] Sewer overflow insurance rider cost source:
Timeshare User's Group, 2005.
http://www.tugbbs.com/forums/showthread.php?t=145439

[15] Sewer backup insurance rider source: Insure.com, 2008.
http://www.insure.com/articles/generalinsurance/sewer-backup.html
[16] Maroochy Water Service Attack:
http://csrc.nist.gov/groups/SMA/fisma/ics/documents/Maroochy-Water-Services-Case-Study_report.pdf
[17] CSRMA.org:
http://www.csrma.org/docs/meeting-agendas/Agenda-PL-110512.pdf
http://csrma.org/template/members.asp?id=366
[18] EPA Report:
http://archive.nacwa.org/getfile05d1.pdf?fn=ra01-4f.pdf
[19] Internet Source:
http://www.usinflationcalculator.com/
[20] California Water Boards, **Statewide Sanitary Sewer Overflow Reduction Program Annual Compliance Report FISCAL YEAR 2011 – 2012, 2012**
http://www.waterboards.ca.gov/water_issues/programs/sso/docs/compliance_report_fy1112.pdf
[21] 2004 EPA report to Congress:
http://www.epa.gov/npdes/pubs/csossoRTC2004_chapter06.pdf
[22] Tsang, Patrick, and Sean Smith. "YASIR: A low-latency, high-integrity security retrofit for legacy SCADA systems." In Proceedings of the IFIP TC-11 23rd International Information Security Conference, pp. 445-459. Springer Boston, 2008.
[23] Wright, Andrew, John Kinast, and Joe McCarty. "Low-latency cryptographic protection for SCADA communications." In Applied Cryptography and Network Security, pp. 263-277. Springer Berlin/Heidelberg, 2004.
[24] Majdalawieh, Munir, Francesco Parisi-Presicce, and Duminda Wijesekera. "DNPSec: Distributed network protocol version 3 (DNP3) security framework." Advances in Computer, Information, and Systems Sciences, and Engineering (2006): 227-234.
[25] Yang, Dayu, Alexander Usynin, and J. Wesley Hines. "Anomaly-based intrusion detection for SCADA systems." In 5th Intl. Topical Meeting on Nuclear Plant Instrumentation, Control and Human Machine Interface Technologies (NPIC&HMIT 05), pp. 12-16. 2006.
[26] Cheung, Steven, Bruno Dutertre, Martin Fong, Ulf Lindqvist, Keith Skinner, and Alfonso Valdes. "Using model-based intrusion detection for SCADA networks." In Proceedings of the SCADA Security Scientific Symposium, pp. 127-134. 2007.
[27] Linda, Ondrej, Todd Vollmer, and Milos Manic. "Neural network based intrusion detection system for critical infrastructures." In Neural Networks, 2009. IJCNN 2009. International Joint Conference on, pp. 1827-1834. IEEE, 2009.
[28] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michael van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. In Workshop on the Economics of Information Security (WEIS), 2012.
[29] Vinay M. Igure, Sean A. Laughter, Ronald D. Williams, Security issues in SCADA networks, Computers & Security, Volume 25, Issue 7, October 2006, pp. 498-506.
[30] Ross Anderson and Tyler Moore. The economics of information security. Science, 314(5799):610-613, 2006.

[31] Rainer Böhme and Thomas Nowey: Economic Security Metrics. In I. Eusgeld, F.C. Freiling, and R. Reussner (Eds.): Dependability Metrics, LNCS 4909, pp. 176–187, 2008.
[32] Stewart, M. G. and Mueller, J. (2012), Terrorism Risks and Cost-Benefit Analysis of Aviation Security. Risk Analysis.
[33] Internet Source:
http://www.census.gov/population/www/
[34] Repository for Industrial Security Incidents (RISI). Quarterly Report on Cyber Security Incidents and Trends Affecting Industrial Control Systems (Sample Copy), 30 November 2009.