NextDrop: Using Human Observations to Track Water Distribution

Emily Kumpel
Dept. of Civil and Environmental Engineering
University of California, Berkeley

Anurag Sridharan, Thejo Kote, Ari Olmos
Next Drop

Tapan S. Parikh
School of Information
University of California, Berkeley

Abstract

Households around the world access water through an intermittent piped water supply; however, delivery of water in these systems is often unpredictable, creating a burden for households waiting to collect water and utilities managing its distribution. We present NextDrop, a system that allows water operators to report information through existing networks of mobile phones and use that information to relay water delivery times to households. NextDrop has been deployed for six months in Hubli, a mid-sized city in India. We describe the NextDrop system, evaluate the accuracy of data collection, and present data collected through NextDrop. The experience of NextDrop shows that crowd-sourced observations can be used to monitor the operation of a complex infrastructure system in a low-resource setting.

1 Introduction

Residents of every city in India, and many others throughout Asia, Africa, and Latin America, access water through a piped water supply that provides water intermittently [8]. In addition to limited durations of water supply, households often receive water at irregular times.

Water utilities providing intermittent supply are often unable to track water timings and durations in their highly decentralized distribution systems. In many continuous piped water distribution systems, utilities instrument the network with sensors and use Supervisory Control And Data Acquisition (SCADA) systems to monitor and control the flow of water. However, utilities with intermittent supply lack the resources for a highly instrumented system, and SCADA systems that do exist were designed for continuously pressurized pipes and would be ineffective for intermittent operation.

NextDrop is a system that collects and disseminates information similar to that provided by a highly instrumented sensor network, but appropriate for the context of intermittent piped water supplies. NextDrop uses existing mobile phone networks to collect observable data about the distribution system from water system operators, and uses this information to send text or voice message alerts about water delivery to households. Mobile phones are already in widespread use among water consumers in urban India: wireless teledensity, expressed the number of lines for every 100 people, is 159 [7].

We present the results of deploying NextDrop in Hubli, Karnataka, India. This is a novel demonstration of how humans can use mobile phones to record easily observable data to improve management of and access to a complex physical infrastructure service. While information may not directly reduce unreliability of the infrastructure, the hope is that NextDrop will improve households' ability to manage unreliable water delivery and help water utilities respond more quickly to problems in the distribution system. Data from NextDrop's system offer a complete picture of the complex reality of intermittent supply distribution system operations and demonstrate the degree to which households currently experience unreliable water supply.

2 Background

Hubli-Dharwad are twin cities with a population of one million, making it the second most populous urban area in the state of Karnataka and one of over 400 cities in India with populations of between 100,000-1,000,000 [1]. The northern Karnataka Urban Water Supply and Drainage Board (KUWSDB) operates the water distribution system on behalf of the Hubli-Dharwad Municipal Corporation. Surface water from two sources are treated at water treatment plants and conveyed by pumps or gravity to service reservoirs throughout the cities. Water is then supplied to consumers intermittently, with supply rotating throughout the city on a cycle that ranges from every other day to once in a week. While for most of the year there is enough water at the source, the capacity (e.g. size of treatment plants and service reservoirs), reliability (e.g. interrupted electricity for pumps), and condition (e.g. excessive leakages in aging pipes) of the infrastructure cannot meet the demand for water by all households simultaneously.

Households with intermittent supply must store water to meet their needs between supply hours. While those with high socio-economic status may have large underground or roof tanks that automatically store water as it is delivered, the majority in Hubli-Dharwad must have a member present to collect and store water in a multitude of storage containers, share tap use with neighbors, or perform chores such as washing while water is flowing. Household members may forgo work, school, or leisure activities while waiting for water delivery. Currently their only way to know when they will receive water is though day-old newspaper announcements.

From the perspective of a household, the supply of water is unreliable at several levels: days until the next supply, time that supply will start, and the supply's duration. The causes of unreliability are complex. The utility provides water rotationally to small geographic areas that receive water at the same time known as *supply areas*. Operators known as *valvemen* control the flow of water to these supply areas. All pipes and reservoirs upstream of a supply area must have a sufficient quantity of water; a chain of utility employees manage valves and pumping operations for these upstream components throughout the city.

Over 800 valves are operated during a single supply cycle, which is the time it takes to supply water rotationally to the whole city. Communication among utility employees, who each control small portions of the distribution system, enables water delivery. Currently, most information about the state of system components is relayed by utility employees and valve operators through ad hoc mobile phone calls and field visits. Some information about supply times and reservoir levels are recorded in logbooks; however, these are often written post hoc with varying levels of accuracy. In addition to unreliability in water delivery caused by operating such a complex rotational system, further unpredictability is introduced by unanticipated pipe breaks, electricity outages that affect pumping, and human error.

There is an opportunity to help households and utility employees better manage the unreliability inherent in intermittent water supply by crowd-sourcing the currently fragmented information about distribution system components. The NextDrop system, which addresses this need, was designed to meet the design criteria of intermittent supply in Hubli-Dharwad; while other cities' water supplies will vary, we believe the key operational elements, such as rotational supply controlled by networks of valve operators, are typical of those in other similarly-sized cities in India.

3 The System

Most utility employees and valvemen in Hubli-Dharwad already own simple mobile phones they regularly use to make and receive calls relaying information about opening and closing valves and repairing broken pipes. Households commonly know the valveman responsible for providing water to their area and will call him frequently to find out about water delivery. Given their already ubiquitous use for relaying information about water availability, mobile phones are a natural fit for NextDrop's system.

NextDrop uses this existing network to collect, process, and distribute more accurate and useful data. Valvemen use mobile phones to report information to NextDrop when they operate valves. NextDrop uses this information to send an alert about water delivery through a text or voice message to households associated with that valve.

3.1 Architecture

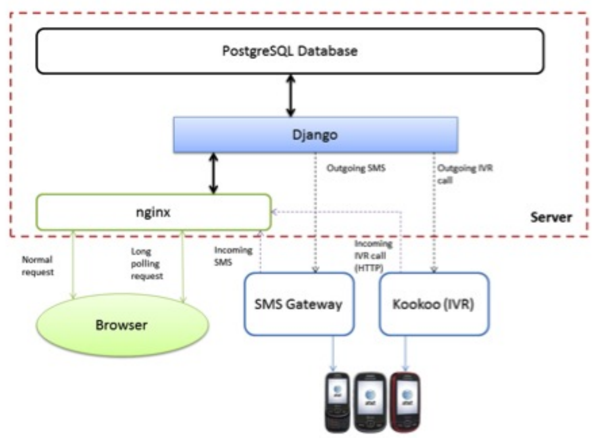

Figure 1: NextDrop System Architecture

NextDrop was built with a number of open source software components. It was primarily developed using Django, a web application development framework for the Python programming language. As Figure 1 shows, the core application receives requests and sends responses through the nginx web server and uses the PostgreSQL relational database to store data. End users interact with the NextDrop application either through a web browser (NextDrop admins and water utility employees) or via SMS and IVR systems (residents of the city). NextDrop depends on external services for SMS and voice based interaction with end users.

The creation of SMS and voice based services has grown progressively easier with the launch of services which abstract away the complexities involved. When we first developed NextDrop, we set up our own voice based telephony infrastructure using open source software (e.g. Asterisk); eventually, we shifted to using Kookoo, a service provider in India, which simplifies the creation of voice services by exposing HTTP APIs for the major operations of making and receiving calls. Access to a SMS gateway over HTTP has been available for some time. Using existing services allowed us to focus on areas where we could add more value: studying the causes of unreliability in the water supply and developing methods for tracking water supply operations and providing accurate information to households.

3.2 Valve Operators

Over 50 valvemen in Hubli turn more than 800 valves on and off throughout the city. Thirty-three of them currently send in updates about 112 water supply areas to NextDrop. Participating valvemen are first trained on their responsibilities and accompany a NextDrop employee to map the geographic borders of supply areas associated with each valve.

The valvemen provide four different types of information to NextDrop: 1) advanced notice; 2) valve opening; 3) valve closing; and 4) last minute changes. They report this information by using their mobile phones to call NextDrop's interactive voice response (IVR) system, which prompts them to enter the type of update and the valve number. Valvemen report advanced notice when they know they will open a valve within an hour, while valvemen send opening and closing notifications when they perform those actions. Valvemen report last minute changes when a supply that has already started is interrupted due to pipe breaks or pump failures.

To evaluate the accuracy and reliability of the information provided by valvemen, NextDrop employees manually tracked missed information from 2/8/2012 to 3/19/2012 by searching for inconsistencies in reported data and following up with phone calls (Figure 2). Valve open notifications, the emphasis of NextDrop's data collection, was reported for 94% of possible valve open occurrences during this time period. Missed notifications are concentrated mostly in certain areas whose valvemen reported forgetting to provide updates.

Advanced notice and valve closure updates were more difficult to reliably collect. Valvemen may not be able to send an advanced notice, as they need to know the state of upstream distribution system components outside of their control to know at what time they will turn on a valve. Valve closures have not been collected as reliability as the other notification types, likely because training emphasizes valve open notifications, as that is the information needed to relay delivery times households. There may also be a disincentive to report closures, as there is an opportunity for corruption by providing longer durations.

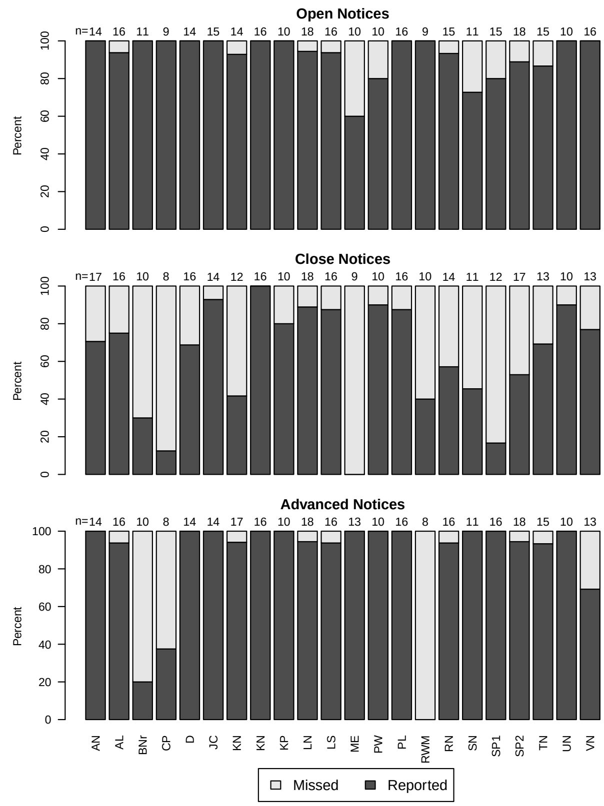

Figure 2: Accuracy of notices from valvemen by area from 2/8/12-3/19/12

3.3 Households

After a valveman reports an opened valve or advanced notice to NextDrop, the NextDrop system sends the appropriate message to enrolled households, who opt to receive it through text message or an automated voice call. NextDrop enumerators recruited customers by going door-to-door, which served to personally introduce households to the service, collect phone numbers, and assign each house to a supply area based on the maps developed in collaboration with valvemen.

More than 7,000 households in 80 supply areas have enrolled and received updates from NextDrop. One month of free updates are provided initially, after which customers can continue the service by paying 10 Rs.

($0.196 USD) for text or 15 Rs. ($0.294 USD) for voice calls per month. 96% of customers chose text message. To date, most of NextDrop's service areas have been in middle- or high-income areas, where residents are more likely to be able to read a text message; as NextDrop offers the service to more low-income areas, households with lower literacy levels may choose IVR messages.

NextDrop verified the accuracy of notifications by randomly calling three households in a supply area at the end of each supply day to check whether they received advanced notice before water delivery and whether water arrived soon after the valve open notification. Three weeks of monitoring showed 100% accuracy, suggesting missed updates are likely to be more of an issue than misreported updates.

3.4 Water Utility

The water utility is currently unable to systematically and reliably track water delivery. NextDrop has been using the collected data to provide reports on water supply timings and durations to utility employees. Currently, only valvemen report information to NextDrop; however, NextDrop has piloted data collection from utility employees regarding the state of other distribution system components, including pipe breaks and service reservoir water levels, which can help utilities respond more quickly to unanticipated problems and allow NextDrop to offer water delivery predictions.

4 Results

NextDrop collected data from valve operations in 112 areas, phased in between 11/1/2011-3/19/2012. Missed alerts - events known to happen but not reported to NextDrop - were tracked manually for a portion of this time as described previously. These data are used to interpret three types of unreliability for a supply area: the day water arrives (supply cycle), the time it arrives (supply timing), and how long it lasts (duration).

4.1 Supply Cycles

The number of days between supplies vary within and between areas. Over six weeks, supply cycles varied from 1-10 days (Figure 3). Starting in October 2011, after increasing capacity at the water treatment plant, transmission lines, and storage reservoirs, the utility changed its official cycle policy from 4-6 days to 2-3 days. However, this reverted back to an (unofficial) 4-6 day cycle in March 2012 due to dry summer season weather. This transition is observed in the data, with supply every 2 days in most areas in February, and a lengthening of the supply cycle starting March 9, 2012.

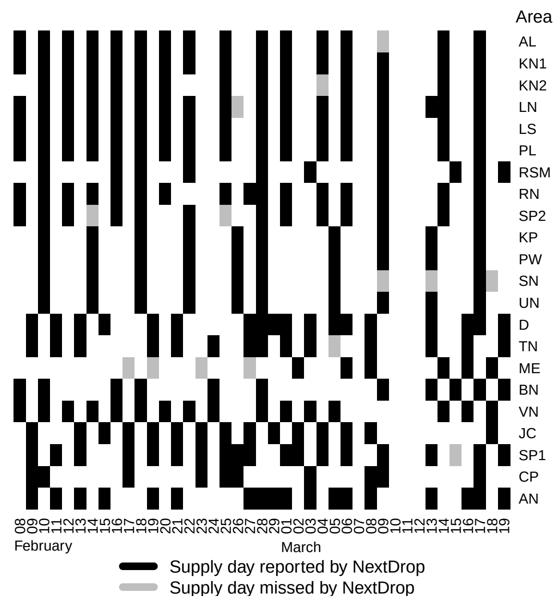

Figure 3: Reported water supply by date in each area, showing variations in the days between supply cycles.

NextDrop employees verified delays of more than three days in valveman-provided updates. While the causes of delays are not formally tracked, these verifications provided anecdotal information on pipe breaks, empty reservoirs, and electricity outages. The variation in cycle lengths between supply areas is also evident, as some supply areas regularly have longer or more erratic cycles than others. There are several possible explanations for this irregularity: distribution system layout, abnormal events such as pipe breaks that require weeks to trace and repair, or re-scheduling of areas with smaller populations due to system-wide supply disruptions.

4.2 Supply Timing

The time that water starts flowing varies with each supply cycle, as shown in Figure 4, which displays the start time for two adjacent supply areas scheduled to begin supply between 5:30-6:30am (indicated by the shaded grey rectangle). Slight variations (within an hour) are likely the result of the time it takes for a valveman to travel between several areas whose valves are opened at the same time. Delays on the order of hours are likely caused by delays in providing water to upstream pipes and reservoirs outside of the valveman's control. Another explanation is that changes in supply duration (discussed in the next section) can affect the start time of nearby areas, as often one valve needs to be closed before another can be opened. Delays of an entire day (discussed in the previous section) may throw off the rotation city-wide,

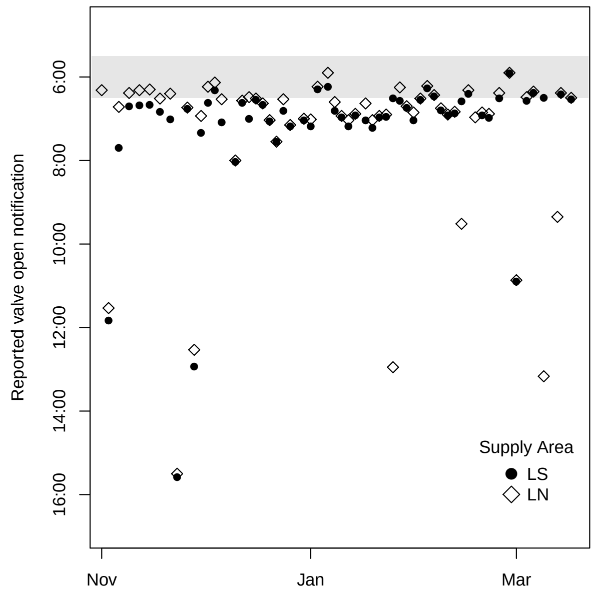

Figure 4: Time that water supply started in two adjacent supply areas. Shaded grey marks the scheduled start times.

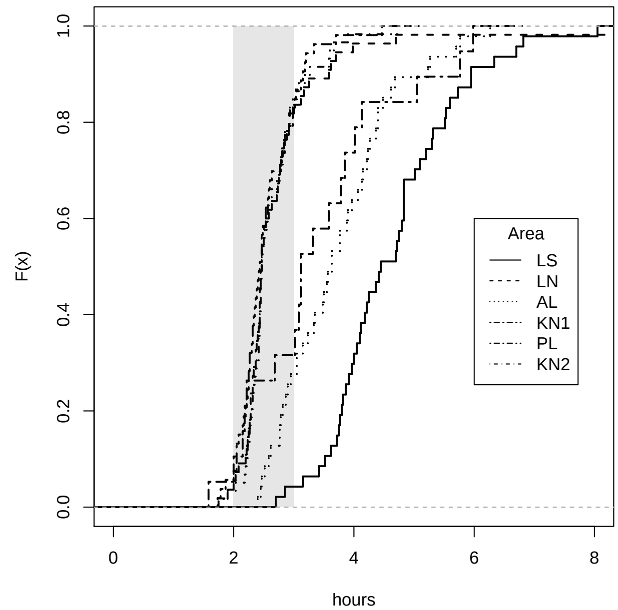

Figure 5: CDF of the duration of supply to adjacent areas, 11/1/2011-3/19/2012. Shaded grey marks the scheduled duration.

altering the start time by half a day or more.

4.3 Duration

The utility has a schedule for the duration that water should be supplied to each area, however, this is often designed without knowledge of a supply area's context, such as the number of households drawing water, available water pressure, or condition of pipes (which affect flow and pressure). Therefore, valvemen are, in practice, given the flexibility to change supply durations. For example, they may supply for more hours if pressure appears to be low on a particular day. Other supply areas regularly receive longer durations because they serve as a main line connecting water to downstream supply areas.

Figure 5 shows variations in supply durations between areas scheduled to receive 2-3 hours of supply (indicated by the grey rectangle). For the reasons discussed above, providing an adequate and equitable amount of water to households is difficult given that the estimated demand and available water is different for each area. Information on existing durations, provided through NextDrop's data, can help the utility adapt the schedule to reflect the context of supply areas.

4.4 Updates to Water Consumers

The only way households currently know when to expect water delivery is through prior experience (which, as shown through Fig. 3 and 4, is of varying reliability) or through announcements published in local newspapers the day before water is to be supplied. A typical announcement in the newspaper simply lists area names, with no timing information and a disclaimer that supply will change if there is a problem with electricity or emergency repairs.

NextDrop sent updates to enrolled households every time valvemen reported a valve open or advanced notice notification. Advanced notice was provided within an hour of the valve open notice, with most valvemen able to provide the notification close to the targeted 30 minutes before start of supply (Figure 6). Households receiving these advanced notice updates reported arranging to return home or have family or neighbors collect water.

5 Related Work

SCADA systems have been in use for several decades by water utilities to gather information as well as to automate operations and link to predictions from hydraulic models. While numerous hardware and software SCADA components are available, most require instrumentation with constant electricity and communications capabilities [2, 3], constraints that limit their applicability for many utilities with intermittent water supply. However, the prevalence of these systems demonstrate a strong case for the value of this data to utilities.

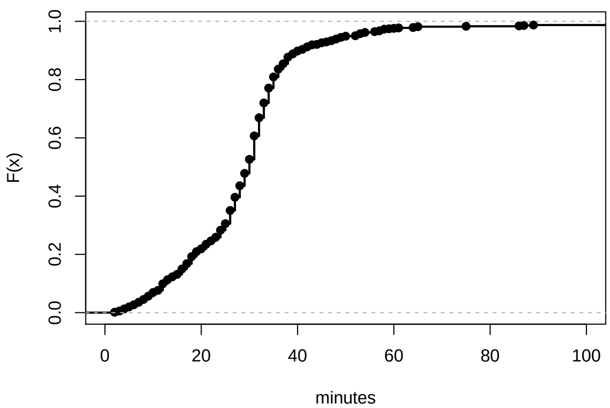

Figure 6: CDF plot of the time between advanced notice and valve open notifications

Several systems have used cellular networks to monitor remote rural water supplies. The Fontes Foundation developed a solar-powered data collection device linked to a cellular network to enable rural water committees to communicate sales and production information to project managers [4]. Water for People deployed an Android smart phone application, the Field Level Operations Watch (FLOW) [9], to monitor whether rural water supplies are functional, and Aquatest, a low-cost test for water quality, uses a mobile phone application on camera phones to assist reading and reporting results from the field [5]. These systems use humans to input the data, however, they also rely on access to technology that may not already be in place (e.g. smartphone or camera phone). Platforms like Ushahidi [6] have demonstrated the usefulness of real-time information crowd-sourced through text messages sent on any device.

NextDrop builds on these concepts, simplifying data collection by using technology already owned and operated by people (simple mobile phones) who can provide data to create a low-cost smart-grid to monitor water distribution operations.

6 Conclusion and Future Work

In this paper we presented the use of a technically simple, low-cost system that uses human observations to provide information through existing mobile phone networks to monitor complex infrastructure operations. We demonstrated that valve operators in Hubli-Dharwad can reliably and accurately report information about intermittent water delivery through NextDrop, and that this information can be relayed to households using mobile phones.

The next step is to identify other data that water utility employees can report to NextDrop to improve distribution system operations, as well as test methods for predicting water availability using NextDrop's historical data. We will also explore how to leverage the capabilities of smartphones to improve the system as they become affordable and available to the target population. While the NextDrop system was designed for intermittent water supplies, the core concept of the system - using reported human observations sent through existing mobile phone networks - could be used by utilities in other sectors, such as electricity or irrigation water, where service may also be irregular and unreliable.

7 Acknowledgments

Thank you to NextDrop staff Madhusudhan B., Fanus Pasha, Avinash Anigol, Adity Paithankar, and Manjunath Bendigiri, as well as to Mr. K. P. Jayaram (Executive Engineer, Hubli, KUWSDB), Dr. K.V. Thrilok Chandra (Comissioner, Hubli-Dharwad Municipal Corporation), and to funding for NextDrop provided by the Knight News Foundation.

References

[1] India population and housing census 2001. Tech. rep., Office of the Registrar General and Census Commissioner (India), New Delhi, India, 2001.

[2] Coulbeck, B., and Orr, C. Essential considerations in the computer control of water distribution systems. *Reliability Engineering & System Safety 42*, 1 (1993), 55–64.

[3] EPA. Water distribution system analysis: Field studies, modeling and management. Tech. rep., U. S. Environmental Protection Agency, 2006.

[4] Koestler, M. A. Live monitoring of rural drinking water schemes using mobile phone infrastructure. *Water, sanitation and hygiene: sustainable development and multisectoral approaches. Proceedings of the 34th WEDC International Conference, United Nations Conference Centre, Addis Ababa, Ethiopia, 18-22 May 2009* (2009). CABI:20103078455.

[5] Loudon, M., Ajmal, T., Rivett, U., De Jager, D., Bain, R., Matthews, R., and Gundry, S. A 'Human-in-the-Loop' mobile image recognition application for rapid scanning of water quality test results.

[6] Okolloh, O. Ushahidi, or 'testimony': Web 2.0 tools for crowdsourcing crisis information. *Participatory Learning and Action 59*, 1 (2009), 65–70.

[7] TRAI. Indian telecom services performance indicator report for the quarter ending september 2011. Tech. rep., Telecom Regulatory Authority of India, 2012.

[8] van den Berg, C., and Danilenko, A. *The IBNET Water Supply and Sanitation Performance Blue Book: The International Benchmarking Network for Water and Sanitation Utilities Databook.* World Bank, Washington, DC, 2011.

[9] Water for People. FLOW: field level operations watch. http://www.waterforpeople.org/programs/field-level-operations-watch.html, 2012.