

Design of Early Warning Flood Detection Systems for Developing Countries

Elizabeth Basha, *Member, IEEE*, and Daniela Rus, *Member, IEEE*

Abstract—In developing countries, flooding due to natural disasters such as hurricanes and earthquakes results in massive loss of life and property. Warning communities of the incoming flood provides an effective solution to this by giving people sufficient time to evacuate and protect their property. However, the range of early warning system solutions introduces a tangle of conflicting requirements including cost and reliability, and creates several interesting problems from factors as diverse as technological, social, and political. The complexity of these systems and need for autonomy within the context of a developing country while remaining maintainable and accessible by non-technical personnel provides a challenge not often solved within developed countries, much less the developing. After describing this problem, the paper discusses a proposed solution for the problem, initial experiments in implementing the solution, and lessons learned through that work.

Index Terms—sensor network, early warning system, flood, honduras

I. INTRODUCTION

Natural disasters are a worldwide phenomenon and require significant cooperation to address. Recent hurricanes, floods, and other events have illustrated this along with the differences of the effects of disasters on developed compared to developing countries. In the recent US flooding due to storms in the Midwest, loss of life and property damage were minimized due to emergency systems available in the highly developed US, while a storm that ravaged approximately seven states caused twenty deaths and \$30 million dollars in damage with only a few left homeless or hungry [1]. On the opposite side, over a much smaller geographic area, North Korea struggled to deal with the displacement of over 300,000 people, approximately 221 deaths, and a cost of \$6 million, most to feed those made homeless by the disaster [2] that resulted in part from the lack of development of warning systems and information at the community level of the impending flooding. From this perspective, the struggle with flooding that faces developing countries presents a pressing issue that we cannot ignore while promising a solution that is globally applicable. Warning communities of the incoming flood, however, is an expensive proposal given the limited resources of the countries. Current methods add to the difficulty with the need for expensive equipment and centralized, computationally difficult flood detection schemes. This presents an opportunity to use the latest

work in information communication technology and sensor networks to solve this problem in a way that balances the minimal cost requirement and limited computational power with the need for high reliability of both the system and computation.

The problem of early warning rapidly grows in complexity upon close inspection and the addition of work within a developing country only increases that complexity. Many other requirements affect the system in addition to those listed above including those related to the devastating effect of the event in question. The problem then encompasses those requirements resulting from both very low activity times when maintenance and attention drop, and highly important times when a flood occurs and the system must continue operation. To properly work, the system also becomes not simply a technical problem, but one of cooperation between government, relief agencies, and the communities to create, maintain, and use the system. These more social and political problems define the success of the system, and ensuring their solutions involves a different approach than the technical issues.

In our work, we examine the problem of flooding on the Aguán River in north-eastern Honduras. This river basin covers a geographic area of 10,000 km² and contains at least 25 highly threatened communities of approximately 35,000 people total. The project began after the devastation caused by Hurricane Mitch in 1998 where a wall of water passing down the river during the night caused approximately 5,000 deaths with an additional 8,000 missing, and 12,000 injured [3]. While considering Mitch a significant disaster in the region, people do not view it as an isolated event. The river experiences annual flooding due to both heavy rain and hurricanes, and, within the intense hurricane period of 2005, the government declared Hurricanes Beta, Gamma, and Stan national disasters. Many lives and property could be saved if people knew the flood was coming and, after flooding occurred, could monitor the river to understand how to best focus relief efforts. With saving Honduran lives through flood warnings as our goal, this paper proposes a high-level solution to the early warning system problem. We have been working on this problem since January 2004 and have performed several prototype experiments toward our solution. This paper discusses some of those experiments and extracts a set of lessons learned from them that can aid others working on this issue along with similar large-scale technology for developing regions.

This paper describes the problem of disaster warning, a solution to the problem in the case of river flooding, a series of experiments towards this solution, and a set of lessons learned through our work in rural Honduras.

Manuscript received August 30, 2007; revised November 7, 2007.

E. Basha is with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: e.basha@mit.edu).

D. Rus is with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: rus@csail.mit.edu).

II. PREVIOUS WORK

Previous work covers a wide variety of topics including sensor networks in developing countries, sensor networks for flood detection, and current operational systems for flood detection.

Sensor Networks in Developing Countries: Juang [4], [5] proposed a system for monitoring wild zebra herds in Kenya within a 100 km² game reserve. The system measured the GPS location of each animal and communicated the information using peer-to-peer short range radios, performing little analysis of the measured data. In addition to the mobile zebra collar nodes, the base node also was mobile within the network, receiving information from whichever nodes were nearby through a long-distance radio during fixed communication windows. The test deployment of the system in 2004 consisted of 10 nodes over a 24 hour period.

De Zoysa [6] discussed a mobile sensor network for road monitoring in Sri Lanka. Actual construction and use of the proposed system had not occurred at the time of the paper for this application although tests toward that goal had occurred. The mobility of the system along with delay tolerance of the data do not make it very similar to our project, although both do share requirements due to the choice of a developing country and large geographic region to monitor.

Guy [7] created a sensor network system that has been installed in four different locations to date. In the James Reserve, a forest setting, the system measured temperature, humidity, rain, and wind using up to 27 nodes over 1.5 years. 2 nodes were installed for 1.5 years in a high-desert farm and 24 nodes in the UCLA Botanical Gardens for 3 months. Finally, a 12-node system was installed in a Bangladesh rice paddy for 2 weeks to measure nitrate, calcium, and phosphate (this experiment also described in [8]). These nodes used 433 MHz communication systems to share the data measured and a base station sent the data for off-line analysis. The goal of the researchers for the system was portability and rapid deployment, focusing on a very different set of requirements than our system.

Panchard [9] suggested a system for monitoring agricultural measurements in India. A prototype system ran for at least one year in an outdoor controlled environment, consisting of 10 nodes communicating via embedded Mica2 radios with selectable radio frequencies (exact frequency not specified in paper). Two nodes were also deployed in the field in India, but the paper did not provide further information on the experiment than node number. In addition to the system description, the paper presented results from informational surveys and outlined some human development issues.

Sensor Networks for Flood Detection: Previous work on sensor networks for flood detection is sparse with only two different examples discovered in the literature. Castillo-Effen [10] suggests an architecture for a system, but is unclear on the basin characteristics and no hardware details are suggested. Closest to our work is a paper by Hughes [11], describing a flood-predicting sensor network that uses Gumstix sensor nodes, which require significant power but allow for a Linux operating system to run on the node. As described, the system

had been tested in the lab, but no field tests were performed by time of the paper. The planned field test would consist of 13 nodes along 1 km of the river. It is unclear what flood prediction model they are using and if it is currently running on their lab test system. Given lack of information on the flood prediction side, the known details of the hardware platform dismiss it as an immediate solution to the problem introduced here as it has limited geographic range, high cost, and large power requirements.

Current Operational Systems for Flood Detection: Current operational systems for early warning of flooding cover two extremes with a new solution appearing to bridge the gap. The lack of published information on operational flood systems makes generalizations difficult, but three systems seem to summarize the approaches currently taken. The first two systems described below provide the most common examples, covering the two possible extremes. Finally, researchers and organizations have implemented a new idea, straddling the middle ground, and discussed at the end.

The one extreme involves a highly technical solution with significant resource support such as seen in the US. For this system, companies develop sensor, communication, and computation technology based on the ALERT protocol, which defines the data structure and policies of environmental monitoring systems [12]. The US Emergency Alert System provides communication of the alerts throughout the nation using television and radio channels by creating special technology and policies, requiring the installation of the technology in stations across the country along with weekly testing, and ensuring protocol compliance at all levels [13]. Implementation of specific systems trickles through each level of government: federal, state, and county. Given the large number of counties in the US, systems and policies do vary, but the majority rely on large numbers of personnel (some highly technical) and significant technical resources. Usually, counties implement the direct measurement system with help from the United States Geological Survey and create policies on how their county defines a disaster and evacuation procedures. Actual prediction usually depends on qualified hydrologists examining the data (thus removing measurement errors) and running it through a complicated physical model called the Sacramento model, which requires calibration of several unmeasurable parameters using years of historical data.

On the other end exists the system commonly seen in Central America, especially Honduras [14]. This system relies on volunteers and limited technology. Sensors to measure river state include river level markings painted on bridges and water collecting rain gages. Volunteers read the river level and rain level (also emptying the rain gage) at several intervals during a day, radioing that information to a central office run by the government. In that office, a person listens to the radio, records the values in a book, and compares them to a defined policy whereby the river level measured corresponds to a color alert. This color alert is radioed to the head office of the government branch, which then decides on the need for an evacuation alert in that region and implements some form of emergency alert procedures. Overall this system relies on very little technology and extensive policies to warn communities, working best in

small river basins where measurements indicate flooding in that area (as opposed to downstream of the measurement area). This paper discusses problems with this type of system in Section V.

Finally, a new middle ground solution exists in Bangladesh, a country regularly devastated by flooding due to its low sea level and large rivers. To combat this, the Danish Hydraulic Institute initially outfitted the country with local telemetry stations in 1995 and created a MIKE 11-based flood forecasting system [15]. However, this system experienced some of the problems seen in the low-tech approach described above and explained in Section V along with issues due to the fact that the headwaters of its major flood-causing rivers originate in India, creating complexities with monitoring. A solution to this was created by a global community of researchers and government institutions, collating all of the satellite information and forecasts generated by the US to provide short, medium, and long-term flood predictions of the major basins [16] [17]. A system called the Flood Forecasting and Warning Response System distributes the alert through reports submitted to various government agencies along with Internet, e-mail, fax, telephone, radio and TV sources [18]. This takes advantage of the ubiquity of satellite information, which looks to provide input data for flood forecasting systems of the future [19]. However, in the case of Honduras, available rainfall data from satellite only exists for satellites passing every 72 hours, which is insufficient for our work.

III. FLOOD DETECTION PROBLEM

The problem of warning communities of impending disasters quickly becomes complex due to its multifaceted nature. At the most basic level, the problem breaks down to predicting the event, communicating that prediction to the proper authorities, warning the communities affected, and evacuating those communities. Each step listed subdivides into its own set of tasks and problems, which are dealt with below. As a clarification, although this section discusses the system in the abstract sense commonly used for technology, this does not preclude people from performing any of these tasks, as Section II explains, since some early warning systems exist as a set of policies with very little technology.

First, the early warning system needs to predict the event, not simply detect the event since that would not provide enough time to evacuate. In the case of river flooding, water can travel down a large river like ours in only a couple hours, providing only enough time to alert the authorities, much less evacuate the community, so the system must predict the flood many hours in advance. Prediction entails a model of the physical system (although it could be a statistical model), an understanding of the relevant variables this model requires as input and the predicted output of the model, physical measurements of these variables, communication of this data to the computation location or locations, and a computational system to run the variables through the model. A prediction alone of the likelihood of the event or key variable defining the event does not solve the problem, however. To warn people, the system must transform the prediction of the event into

an understanding of the effect of the event, a timeline of the progression of the event, and an understanding of the uncertainties involved. To quantify this in the case of flooding, knowing the river level does not help, but knowing which regions will flood due to that river level, how long those areas have before flooding, and how likely the flood will occur at that stage provides the information that can then enable a warning.

None of this deals with the fact that the system operates within a real environment and must survive that environment. Some of the physical measurement components exist outside, requiring that they endure the elements, resist theft and damage, and operate with limited maintenance. These components will need to cover a certain percentage of the area in which the event occurs, often a large geographic area, in order to correctly measure enough data points for the model. The budget usually limits this to the minimum number allowable, demanding careful placement and creating large communication distances to aggregate the data at the computation point (or points). For the system components inside, those locations, once identified, need almost the same set of requirements: the elements and thieves should stay outside the location, the system should have protection from incidental damage, and someone should check on the system occasionally. These components also will have to cover the appropriate area and handle the communication distances necessitated by the measurement components. Plus everything needs electricity.

Once the prediction states the event will occur, authorities should know. Yet how that occurs is unclear as two things need definition. First, what communication methods inform the authorities. Should people receive a phone call, a fax, a page? Or should a siren or flashing light signal the occurrence? This leads into the second item, that of who transmits the information since either the system or local people could perform this task. At some level, the system informs someone, either by informing a local authority who then communicates with everyone else, or by informing everyone, so the questions become a matter of balancing how the communication occurs and with what levels of redundancy.

After the authorities, the communities should learn of the threat. The trickiness here, from the system perspective, develops due to the government regulations on alerts. Most governments do not want random people or systems issuing evacuation alerts, but want them to issue from the government officials only. The system can then provide information on the likelihood of the event, but must clarify that information with the warning that only the proper government channels can declare an evacuation alert. If the alert mechanism, be it a siren or light or phone call, becomes part of the technology of the system, another communication channel within the government must allow for the approval of the alert before the system triggers it.

How the system communicates any information to the communities becomes complex when dealing with some communities within developing countries. A lack of literacy and education within these communities complicates the presentation of the information, necessitating a visual or audio

method where difficulties arise in presenting the uncertainty of the information. Independent of developing or developed, communities need to understand the system to understand the message, realizing the fallibility of the technology and that this fallibility does not reduce the usefulness of the technology when the information is viewed correctly, a task difficult with non-technical people even as the world becomes more technology savvy.

Finally, when the evacuation alert occurs, the communities need to evacuate. This raises questions about where they should go and how to get there. In the case of flooding, river basins can be very flat for kilometers so the community may not have an obvious high ground point to go to or the high point could exist on the opposite side of the river, requiring that either the bridge remains clear or the community has boats. Answering these questions raises another set of problems and design points for an increasingly complicated system where failing to account for any one of the basic tasks and sub-questions renders the system ineffective.

Additionally, all this requires policy to define the more abstract questions that arise. How many hours in advance should the system predict? What margin of uncertainty of the prediction is allowable? Who needs notification and in what priority? And so on for almost every aspect of the system. Since these systems tend to cross state or county or department boundaries, assigning responsibility for answering the questions, much less implementing the solutions in some cases, involves dealing with the various levels of bureaucracy. Although the national government should resolve these issues by taking this responsibility, most often the geographic location of its capital limits its understanding of the issues and, especially in developing countries, it cannot manage the complexity of the system itself due to lack of resources, whether that be people, money or equipment.

Much of the above discussion also applies to warning systems in developed countries. However, infrastructure plays a very large role in defining the complexity of the problem as do resource availability and pervasiveness of organized government. Developed countries have all three of these, simplifying the complexity of the problem and answering many of the questions in ways that reduce the system requirements. Developing countries require much more consideration for each aspect of the system, trading off factors while planning for a variety of infrastructure, resource, and government scenarios.

IV. FLOOD DETECTION IN HONDURAS

Our work on this problem began in January 2004 during a trip to Honduras to visit with a non-governmental organization (NGO) called the Centro Técnico San Alonso Rodríguez (CTSAR). This organization, headquartered in Tocoa, consists primarily of local Hondurans and has operated in the region for over a decade. They focus on enabling the surrounding rural communities to identify and solve problems, especially those involving agriculture, while providing technical assistance with the solutions. Because of their local nature, history of reliability, and enabling behavior, the communities respect this NGO, regularly contacting them with new problems and

working together to create sustainable solutions. Additionally, recognizing that the rural nature of north-eastern Honduras limits the government presence in the area, the government and CTSAR work closely to develop necessary infrastructure.



Fig. 1. Aftermath of Hurricane Mitch in Northern Honduras [3]

In the aftermath of Hurricane Mitch, occurring in 1998, CTSAR recognized the need for a flood early warning system on the Aguán River and applied for funding for such a project. Our visit fortuitously coincided with the beginning of the project and a partnership was arranged whereby we would provide the engineering skills necessary for the prediction and notification aspects of the system, and they would work with the communities to develop their aspects of the system.

The Aguán River basin encompasses 10,000 km² of land. Situated between two mountain ranges, the extreme flatness of the basin area endangers communities along the length of the river, but especially those at the broad mouth where it meets the ocean. Communities living in that region fish for a living thereby requiring proximity to both the ocean and the river for their occupation. Additionally, the poverty, illiteracy, and racial/cultural variance from the rest of the country as descendants of African slaves hinder most chances of moving to safer locations. Mitch washed away some of the approximately 25 communities living in this region (see Figure 1) creating an unfortunate opportunity for moving, but aid, in the immediate necessity of the moment, reconstructed the communities in almost the same locations, placing them in the same danger as before.

Hurricanes Beta, Gamma, and Stan in 2005 also hurt these communities (see Figures 2). Fortunately, none were washed away, but they clearly demonstrate the continued gravity of the flooding threat.

This threat not only occurs due to hurricanes, but heavy rainfall can annually cause near-flood to flooding conditions. Figure 3 best demonstrates the potential severity of the rainfall, showing the same location on the river at three different time periods within a year. Even before the November hurricanes, the river grew significantly as seen in the August picture, Figure 3(b). By January, at the end of the rainy season, the river width had grown so much that the river banks and many trees are no longer visible compared to the August height.



Fig. 2. Aftermath of Hurricanes Beta, Gamma, and Stan in 2005



(a) March 2005



(b) August 2005



(c) January 2006

Fig. 3. River Changes Over 1 Year

In aiding these communities, our goal consists of developing a system meeting the requirements and tasks of Section III. Given a weak government presence in this area of Honduras, the system must work with as much government cooperation as possible yet function outside the government in case no cooperation occurs. This confuses the issues of responsibility and notification, but, as stated in the next section, we have worked to develop several solutions to that aspect of the system.

V. PROPOSED SOLUTION FOR FLOOD DETECTION USING SENSOR NETWORKS

In developing a solution to this problem, we divided the four tasks (event prediction, authority notification, community alert, and community evacuation) between CTSAR and ourselves. We agreed to design the event prediction system with both groups working to install the system in the river basin and some help from CTSAR in defining what constitutes a flood. Any technology needed for authority notification would fall under our purview while CTSAR would work with the Honduran government to arrange who receives the notification and what style of notification. The community alert splits similarly; we have the technology (or the hows), and they have identifying the whos and whats of the alert. For community evacuation, CTSAR agreed to work with the communities to develop evacuation policies, determine additional resource needs, and arrange for the implementation or purchase of those needs. Quite simply, we work on the technology, and they work on the people and policy issues.

For the technology, two factors define the shape of the solution. First, the geographic area involved hinders any form of volunteer-based system. The communities upstream that would need to perform the measurements and/or the communication of those measurements have almost no connection to the communities affected by the flooding. This removes any level of self-interest and peer-pressure in voluntarily performing any system tasks. Second, the nature of the problem involves measuring the river and surrounding area during heavy rains, hurricanes, and at all times of day and night. Very few volunteers would stand outside in a hurricane to perform a measurement or radio information to a central office. Neither would they perform these tasks in the middle of the night. While paying someone may allow night-time measurements, few people would remain during a hurricane, especially if that hurricane affects their own community with small-scale flooding, building leakages, or potential agriculture crop damage. Yet it is at these times that the measurements are most needed. This suggests an autonomous system and, once the system becomes autonomous, we have a sensor network.

With a sensor network solution specified, we define the following specific system requirements:

- Monitor events over large geographic regions of approximately 10,000 km²
- Measure a wide variety of variables contributing to the occurrence of the event
- Survive long-term element exposure (on the order of years)
- Recover from node losses
- Detect and predict the river flood
- Withstand the river flood
- Power system for years
- Minimize costs
- Handle significant computation requirements

Due to the distance requirement, the inability to populate the entire area with sensors, and cost limitations, a network solution cannot consist of many nodes so nodes will need to communicate over long distances on the order of 25 km. Yet the system also needs a variety of physical measurements around any one location, suggesting multiple sensor nodes, or a mini-network, at each location. Additionally any communication system cannot contain recurring costs since that limits the sustainability of the system by requiring funding for the technology beyond the initial installation. While the system may require funding for maintenance, finding occasional funding for labor to check and repair the system is easy. Finding regular funding is hard and creating a technological point of failure at which the system no longer works almost ensures the system will stop working. Other constraints affect the node design and system protocols and, taken with the above, result in Figure 4, showing an idealized overview of our system.

Based on the combination of mini-networks and long-range links, as demonstrated in Figure 4, the system communicates via two different methods. To cover the long-range communication links of approximately 25 km range, the system uses 144 MHz radios. In the US, this band covers the amateur radio operators and thereby provides many mobile radio systems for possible use. Additionally these radios cover a wide variety of applications within Central America so people who can repair these radios do exist, easing the burden of maintenance. These systems focus on voice communication, however, so we designed a modem to allow data communication thereby providing a cheap, long-range communication method without the recurring costs of a satellite or mobile telephone system. For short-range communication links within a 8 km range such as required by the mini-networks, the system operates within the 900 MHz band.

In addition to two communication methods, the system consists of four different regimes of operation: sensing, computation, government and office interface, and community interface. River flooding requires information about the state of the river, soil conditions, and meteorological conditions so we measure river level, rainfall, and air temperature at nodes powered by solar panels. (We discuss some specifics of this decision in the next section.) In order to minimize failure points, we limit the number of computation points, but ensure redundancy at the same time by using the systems communicating at 144 MHz as computation nodes, providing them with larger microprocessors and more memory. Office

nodes inform the government (along with relief agencies and local NGOs) and display more technical information about the system along with long-term data storage. This requires an office with a laptop, some form of storage system, a phone line for the communication, and, while offices tend to have grid power, a reliable backup power system to ensure the office continues to work during the disaster.

Finally, the community nodes display the river state and prediction. The least defined part of our solution, this relies greatly on our work with CTSAR to understand how the community would like to learn of the potential flood. Currently, we plan to inform a set of city members, elected for this purpose by the community, who then relay the information via centralized radio megaphones to the rest of the community. This set of city members will have access to a computer displaying the river state, but most likely will receive individual notification in their homes via a flashing light.

None of this mentions exactly how the system will predict a flood and define a flood. Ongoing research examines how to predict the future river level over several time periods extending 48 hours into the future in a computationally-light way compared to current hydrological methods by using statistical methods. Once we have a prediction of river level, we use a lookup table to determine if and where that level implies flooding. No long-term data exists for the river level either up or down stream, nor does data describe what levels in the river create floods at which communities. For this, we use local community knowledge. With the help of CTSAR, each community creates a hazard map, defining which areas always flood versus never flood, along with what flooded for known events such as Mitch. We then ask upstream communities about river levels for the known events and, using what topographic information exists, correlate the values to create a lookup table associating upstream river levels with downstream flooding.

VI. EXPERIMENTS

We have traveled nine times to Honduras over the past three and a half years of this project. The goal of these field trials has been to deploy and test the system components.

On the communication side, we verified the usability of the 144 MHz radios. We tested the various ranges necessary for the system, ensuring that they can communicate over those ranges. To communicate at these ranges reliably, the radio antennas need line-of-sight high in the air, which requires antenna towers and limits the ability to test this portion of the system in the US. With CTSAR help, we arranged access to land and built 5 meter antenna towers at two river sites where we plan to install water level sensors for 144 MHz radio communication along with 10 meter towers at the CTSAR office and the government emergency management office in Tocoa (built and staffed by CTSAR). With these towers, we verified both the communication range and the ability of our modems to communicate data over this range. Thanks to the 2005 hurricanes, we also proved that the towers and antennas will survive hurricane force winds.

At the offices, in addition to the towers, we worked to design and install secondary solar power systems. We would prefer




Fig. 4. Idealized Sensor Network Consisting of Two Communication Tiers: 144 MHz Computation, Office and Community Nodes, and 900 MHz Sensor Nodes

to use grid power if it exists, but need solar power backup for the daily fluctuations of that system along with the major outages associated with disasters. CTSAR worked with a local company to purchase panels, batteries, and a charge controller. We added an off-the-shelf inverter, a power strip, and very simple custom electronics to switch to solar at the absence of grid power. We installed these systems at both offices and are running long-term usage tests. At the government office, we also installed a permanent radio and laptop for development of that interface, using it both for longer term radio tests and exploring issues with the interface. CTSAR constructed this office for the government as part of this project and its newness has required all groups to cooperate in defining the policies and purposes of the space.

While working on all aspects of the project, another area of major focus, in addition to the communication, has been the water measuring system. This area has provided the majority of lessons learned and space to explore the various sustainability and maintainability issues. We have created five different prototypes of this system. Through these prototypes, we settled on measuring water pressure as a method of obtaining river level. Other options such as resistive water level sensors were rejected due to corrosion issues, while ultrasonic sensors were rejected due to the indirect nature of the measurement along with reduced ability in high winds. The container for these sensors also provided some interesting lessons. One prototype used locally available PVC pipes as a housing for the system. Amusingly, our collection of necessary PVC parts turned out to not all exist locally. Given that the initial reasoning for the PVC failed and the difficulties of creating reliable, secure solutions, we switched to Otter boxes, providing significantly improved reliability at the cost of local creation. Finally, these prototypes allowed us to understand the complexities of installing something in a flooding river since box movement reduces the efficacy of the measurement. Structures must hold the sensor in a fixed spot while ensuring the system does not sink in the soft ground of the river and that it is retrievable for

maintenance. We developed two different solutions allowing us to install the system on a bridge for greater reliability and also in the middle of the river when the situation necessitates. The latter uses a combination of a scooter tire, cement, and metal standoffs to ensure the box does not move and does not sink, while a metal chain both protects the cable and provides a retrieval option.

VII. LESSONS LEARNED

The technology shows great promise to aid the people of Honduras. Our early experiments confirm the feasibility of a sensor network for use in a developing country and for use in providing disaster warnings, thereby yielding an easily deployable and scalable system. We created two communication solutions and verified that both function, finalizing that portion of the system. Similarly, all three sensors work in a node structure that meets the computational requirements of current microprocessors, along with successful packaging and installation procedures. Finally, we have infrastructure needs such as power and antenna towers solved, placing us in the position where all components are functional and only connecting everything into a complete system remains. Performing these experiments, in addition to component verification, gained us constructive insights due to some unexpected issues, which, while not technical barriers to the system, furnish general lessons valuable for projects of this size within a developing country.

Partnerships: The partnership we have created with CTSAR works very well and the structure of the collaboration introduced us to one of the more successful frameworks for international development projects. Key to everything was their identification of the problem, request for our help, and securing initial funding. Having them initiate the project ensures their commitment to it through all of the setbacks and design changes such projects endure. They initiated it because of community feedback, guaranteeing the support of the community in implementing and maintaining the system.

Had we introduced the project, given its need for community input and lack of potential income, our eventual withdrawal would most likely begin the slow decay of the system, as the force behind the project would disappear.

Within the partnership, we divided the work nearly equally between both partners and divided the responsibility for providing resources. Allowing CTSAR to insulate our work from the community ensures the progress of the work and cooperation of all parties. We cannot travel to Honduras all of the time and do not have a continuous connection with the communities, but do have better access to technical resources when in the US. CTSAR has both access to and a history of work with the communities, creating a relationship that fosters cooperation, but they do not have the technical skills nor access to materials. This equal division of tasks and shared commitment through resource buy-in, while seemingly obvious, does not exist in many projects and the lack of a committed partner often causes the failure of that project.

Security: We have dealt with many security issues throughout our project. Several early sensor prototypes were stolen from their installation near the river because, while the sensor was in the river, the electronics and cable were placed at the top of bridges or near the bank for water protection and easy retrieval. In one instance, someone stole over 180 meters of cable running along the bridge connecting the electronics to the sensor. We had connected it to electricity cables running along the bridge in order to disguise it and the person, in taking those cables, took ours as well. Children climbed on our towers because they were there and, after providing more security, shot marbles at boxes on top of the towers. One marble punctured a plastic box, creating a small hole in which rainwater entered, damaging our battery and radio stored there because we had not expected marble damage.

Overall, our lesson has repeatedly been that there is no such thing as too much security and if we find accessing our equipment difficult for routine work, then it presents too much work for thieves. We have since placed all our cable within PVC pipes that we then bolt to the bridge so that people cannot cut or easily steal them. Our tower can now be climbed only with a ladder (see Figure 5) and we are upgrading our boxes to metal after the marble incident.

We did figure out some methods requiring less construction as well. For our sensors, especially the rainfall and temperature sensors, we discovered a couple of friendly families within our urban site to host the sensors thereby protecting them from thieves and minimizing the infrastructure needed. In our rural site, we located an old train bridge on which to install our sensors and camouflaged the very white rainfall sensor (see Figure 6).

Redundancy: While most agree that technology systems need redundancy, an early warning system notifying people of disasters needs significant redundancy at all levels. On the technical side, a node could fail at any time due to element exposure destroying some part of the system, people stealing or damaging some part, or the more standard death of some part. All of this can reduce or ruin the ability of the system to accurately predict flooding. Therefore, every key level of the technology needs redundancy: sensors, communication paths,



Fig. 5. 5 Meter Antenna Tower with Security



Fig. 6. Camouflaged Rainfall Sensor

data storage, batteries, computational units, and so on.

Additionally, the policies need redundancy. The government may not provide the prompt action desired in response to an alert and no community wants to watch the flood waters ascend while the government stalls on evacuation notices. Therefore, the communities and relief agencies need to define what happens in that case and how long they wait for government action before following Plan B. Maintenance may also receive delayed attention from the government so, again, the communities and NGOs need to figure out what happens in that case and who takes responsibility for ensuring the continued operation of the system. (While it may be hard to imagine a government delaying in the face of impending disaster, some governments change their entire civil service following an election which could affect the ability of the

authorities to understand the policies and system. The response of the US during Hurricane Katrina is a recent example of such a delayed response.)

Community Knowledge: Lacking historical data in which to develop flood prediction algorithms or train models, our initial plan consisted of installing sensors and gathering the data for ourselves. The realization that we could not gather all the data necessary in a reasonable timeframe led us to consider other methods than those commonly used by the hydrology community. We discovered that we could achieve the results for which we wanted the historical data, at some level, from the memories of the community members who view the river every day along with the historical effects of floods. This allows us to leapfrog the many years necessary to gather enough data for current models and, using a much simpler and cheaper method, generate a table of what constitutes a flood. We also plan to utilize this insight in providing a check on the system by allowing users to input information on the current state of the river and details on past flooding, which the system can then use to refine the models and verify sensor activity.

Testing: Every system needs testing at many different levels—most people agree on the obviousness of that statement. However, in our experience, a large-scale system such as ours that heavily relies on in-country infrastructure usually follows a test strategy whereby component testing occurs in the lab and complete system testing occurs in the field as an installation of the system in its planned location. This strategy has repeatedly failed us. On one hand, we need the components in the US for component improvements and debugging. On the other side, we need the entire system in Honduras for complete testing and we cannot be in Honduras all the time. This results in a combination of approaches. We test the components to some level, travel to Honduras, install the system as it exists, run tests for a week or so to find system problems, leave some components for longer testing, and return home with the remaining components for further development. The status of the components left, from a technical point of view, awaits our next trip to Honduras as CTSAR can tell us if they still exist and if they appear to work, but not any specifics on how they work or what may have failed without the remainder of the system. Thus we only discover long-term problems when we return to Honduras, where we can perform some debugging but need to return to the lab for further debugging and most development work. Ideally we need a system within the US, but arranging for a 10,000 km² river basin where we can install antenna towers does not seem feasible. What we have done is talk to the government organizations responsible for measuring the river locally in the US and asked for their help. This resulted in access to their cement structures for their antennas, computers, and sensors. Because we can use the existing infrastructure, we can install the complete system in a local US river and test it without constructing our own antenna towers, thus speeding up our development work.

VIII. CONCLUSION

Disasters, by their very definition, wreak havoc on countries indiscriminately. Developing countries experience much worse

havoc than developed, however, and are less equipped to deal with the aftereffects of these disasters. Foreknowledge of the disaster could thus help all, but better aid developing countries by providing time to secure property and evacuate. Developing early warning systems is a complicated problem, with many facets to the system requirements and many additional intricacies when within a developing country. This paper described these complexities, focusing on the application of river flooding, and examined a potential solution to the problem of river flooding. Through field experiments in Honduras, the problem and solution have been elaborated and refined, with many lessons learned during the process.

A complete system still does not exist within Honduras, but work continues to develop this system and create a local test platform. In the near future, we will outline a detailed system architecture, finalize testing on components of the 900 MHz sensing system, and develop a statistical prediction algorithm utilizing the community knowledge of the river. We hope to protect the people of Honduras during the rain and hurricane season of 2008, with expansion to other countries soon to follow.

ACKNOWLEDGMENT

This paper was supported in part by Microsoft Corporation, NSF Graduate Fellowships, NSF CNS-0707601, and NSF IIS-0426838. We are very grateful for this support.

REFERENCES

- [1] CBS News, "Flood clean-up under way in soggy midwest," *CBS News/AP*, August 21 2007, online at <http://www.cbsnews.com/stories/2007/08/21/national/main3188053.shtml>.
- [2] J. Herskovitz, "Bodies of North Korean flood victims float to south," *Reuters News*, August 21 2007, online at <http://uk.reuters.com/article/homepageCrisis/idUKL21191873.CH..242020070821>.
- [3] United States Geological Survey, "USGS hurricane mitch program," <http://mitchnts1.cr.usgs.gov/overview.html>.
- [4] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. S. Peh, and D. Rubenstein, "Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with ZebraNet," in *ASPLOS '02: Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems*. New York, NY, USA: ACM Press, 2002, pp. 96–107.
- [5] P. Zhang, C. M. Sadler, S. A. Lyon, and M. Martonosi, "Hardware design experiences in ZebraNet," in *SenSys '04: Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*. New York, NY, USA: ACM Press, 2004, pp. 227–238.
- [6] K. D. Zoysa, C. Keppitiyagama, G. P. Seneviratne, and W. W. A. T. Shihan, "A public transport system based sensor network for road surface condition monitoring," in *NSDR '07: ACM SIGCOMM Workshop on Networked Systems for Developing Regions*, August 2007.
- [7] R. Guy, B. Greenstein, J. Hicks, R. Kapur, N. Ramanathan, T. Schoellhammer, T. Stathopoulos, K. Weeks, K. Chang, L. Girod, and D. Estrin, "Experiences with the extensible sensing system ESS," in *Proceedings of CENS Technical Report #60*. CENS, March 2006.
- [8] N. Ramanathan, L. Balzano, D. Estrin, M. Hansen, T. Harmon, J. Jay, W. Kaiser, and G. Sukhatme, "Designing wireless sensor networks as a shared resource for sustainable development," in *ICTD '06: Proceedings of the International Conference on Information and Communication Technologies and Development*, May 2006, pp. 256–265.
- [9] J. Panchard, S. Rao, T. Prabhakar, H. Jamadagni, and J.-P. Hubaux, "Common-sense net: Improved water management for resource-poor farmers via sensor networks," in *ICTD '06: Proceedings of the International Conference on Information and Communication Technologies and Development*, May 2006, pp. 22–33.
- [10] M. Castillo-Effen, D. H. Quintela, R. Jordan, W. Westhoff, and W. Moreno, "Wireless sensor networks for flash-flood alerting," in *Proceedings of the Fifth IEEE International Caracas Conference on Devices, Circuits and Systems*. IEEE, Nov 2004, pp. 142–146.

- [11] D. Hughes, P. Greenwood, G. Blair, G. Coulson, F. Pappenberger, P. Smith, and K. Beven, "An intelligent and adaptable grid-based flood monitoring and warning system," in *Proceedings of the 5th UK eScience All Hands Meeting*, 2006.
- [12] ALERT Systems Organization, "Alert history," <http://www.alertsystems.org>.
- [13] Federal Communications Commission: Public Safety and Homeland Security Bureau, "Emergency alert system," <http://www.fcc.gov/pshs/eas/>.
- [14] E. Basha, "Interview with COPECO officials in La Masica, Honduras," January 2004.
- [15] G. H. Jørgensen and J. Høst-Madsen, "Development of a flood forecasting system in Bangladesh," in *Operational Water Management Conference*, 1997.
- [16] P. J. Webster and R. Grossman, "Forecasting river discharge into Bangladesh on short, medium and long time scales," *Climate Forecasting Applications in Bangladesh*, January 2003, online at http://cfab.eas.gatech.edu/cfab/Documents/InfoSheets/CFAB_forecast.pdf.
- [17] T. M. Hopson and P. J. Webster, "Operational short-term flood forecasting for Bangladesh: application of ECMWF ensemble precipitation forecasts," vol. 8, 2006.
- [18] R. Chowdhury, "Consensus seasonal flood forecasts and warning response system (FFWRS): an alternate for nonstructural flood management in Bangladesh," *Environmental Management*, vol. 35, pp. 716–725, May 27 2005.
- [19] F. Hossain, N. Katiyar, Y. Hong, and A. Wolf, "The emerging role of satellite rainfall data in improving the hydro-political situation of flood monitoring in the under-developed regions of the world," *Journal of Natural Hazards*, vol. 43, pp. 199–210, March 9 2007.