IT Lessons Learnt from Real Time Dike Monitoring

Matthijs Vonder and Bram van der Waaij
TNO, Eemsgolaan 3, Groningen, The Netherlands

Keywords: Dike Monitoring, Sensor Data, Scalable, Cloud, Use Case, Anomaly Detection, Real-time, Lessons Learnt.

Abstract: The Dutch lowlands are protected by many kilometres of dikes. Currently these dikes are visually inspected on a regular basis. During heavy weather this frequency is raised, up to 24/7 in very extreme situations. After a dike failure at the Dutch town Wilnis in 2003, the question was raised whether modern sensor technology could be used to assess extra information on dike conditions. To answer this question, different experiments have been conducted in order to gain more knowledge about dike failure mechanisms and to validate real time sensor dike monitoring in existing dikes. Based on these use cases, this paper presents several IT lessons learnt and future IT challenges concerning data storage, anomaly detection and dike stability models in relation to CPU power usage for small, medium and large scale dike monitoring.

1 INTRODUCTION

The Dutch lowlands are protected by many kilometres of dikes. Despite the fact that building dikes started in the late Middle Ages, today designing, constructing and maintaining dikes still involves a lot of empiricism (Van, 2009). During high water conditions the information on the actual strength of a dike is usually obtained by visual inspection. Questions about the time to failure or the maximum load increase that a specific dike location can withstand are hard to answer. Modern sensor technology is used to obtain (sub)soil information. After a dike failure at the Dutch town Wilnis in 2003 (Bezuijen, 2005), the question was raised whether modern sensor technology could be used to obtain extra information on dike conditions.

Currently dikes are visually inspected on a regular basis, e.g. every month. During heavy weather this frequency is raised, up to 24/7 in very extreme situations. Next to the visual inspections some temporal experiments are performed using sensors (e.g. pore pressure) with batch processing afterwards. Special theoretical models are used in this batch processing, which are based on lab experiments and hind sight analyses of the real disasters.

Van (2009) states that sensor technology could be used as an early warning system: when a monitored parameter reaches a certain value, people are warned and action can be taken. When using modern sensor technology for an early warning system, it should be known which parameter should be monitored at which interval in time and space and at which location in the cross-section, but also at which point an action should be taken and what time frame is available. In order to find an answer to these questions the idea of a fieldlab IJkdijk was born. In the past years a couple of experiments have been done. In this paper we focus on the IT-related aspects while using sensor systems to monitor a dike.

2 FIELDLAB "IJKDIJK"

In 2005 the idea of a fieldlab IJkdijk was born (Vries, 2010). It is pronounced as 'Ike-dike' and is Dutch for 'calibration dike'. In 2008 the IJkdijk Foundation was established, which is an initiative of TNO and Deltares, STOWA, NOM and IDL (Pals, 2009). It is an initiative where knowledge on dikes and sensor technology comes together. The plan has been emerged to build test dikes to enable the systematic testing of existing and new theoretical models using various types of new sensors and communication technologies, both during construction and on the entire lifetime of a dike.

As part of the IJkdijk program, in Booneschans, the Netherlands, a number of dikes is built at full scale and brought to failure with two explicit goals: to increase the knowledge on dike behaviour and to

Vonder M. and van der Waaij B..
IT Lessons Learnt from Real Time Dike Monitoring.
DOI: 10.5220/0004202700730079
In *Proceedings of the 2nd International Conference on Sensor Networks* (SENSORNETS-2013), pages 73-79
ISBN: 978-989-8565-45-7

develop and test new sensor technologies for flood early warning systems under field conditions (Van, 2009). This should increase both the quality of the dike inspection and monitoring process and the safety assessment of dikes.

The IJkdijk project functions as an open innovation platform for testing sensor techniques, dike monitoring systems and to improve dike technology, providing benchmarking to all contributors (Pals, 2009). By conducting experiments in a controlled environment and under pre-determined conditions, knowledge about the failure mechanisms of dikes can be improved and the value of new technologies can be demonstrated. At present, more than forty companies and institutions from five different countries cooperate in this initiative.

Between 2007 and 2010 ground breaking experiments were conducted in dike monitoring with the aid of sensor technology. 'Overtopping' - where water slushes over the dike - was explored in 2007 (Meijer, 2008), macro stability in September 2008 and 'backward piping erosion' – where a kind of tunnel arises through seepage underneath a dike - in September-December 2009 (Kruiver, 2010). All dikes were stressed to a point where they failed.

2.1 IJKDIJK Environment

Figure 1: IJKDIJK environment (Weijers, 2009).

The IJKDIJK location is an area of 800m x 120m and has the advantage that it was already surrounded with its own dikes. So when something might go wrong during the experiments the water will stay in the polder. It is also located next to a 30m wide canal, which supplies the necessary water for the experiments (Meijer, 2008). Figure 1 gives an impression of the IJkdijk environment (while preparing the macro stability experiment).

For failure mechanisms mentioned earlier (overtopping, macro stability and piping) several test dikes (i.e. IJKDIJK's) were built under supervision of Deltares. These test dikes had real-life dimensions concerning height and width, but with a limited length (i.e., up to 100 meters for macro stability). For the experiments focusing on macro stability and piping, other companies were invited to become partner who brought their own sensor systems equipment. This resulted in a list of 13 different sensor systems for the macro stability experiment (Weijers, 2009) and 9 for the piping experiment (Koelewijn, 2010). Figure 2 gives an instrumentation overview for the macro stability experiment, while the dike breach is shown in figure 3. The sensor systems varied from known pore pressure sensors from different suppliers to fibre optic temperature sensors, infrared cameras and even experimental "listening tube with hydrophones" (Meijer, 2008). In the design phase all supplied sensor systems were taken into account and they were placed during the building of the test dikes (and not put in the dike afterwards, as would be the case for existing dikes).

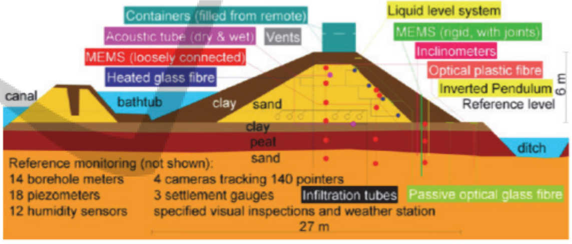

Figure 2: Instrumentation overview macro stability experiment (Van, 2009).

Figure 3: Dike breach of macro stability experiment (photo: TNO 2008).

For each experiment a dike is furnished with in-situ intra-dike and extra-dike sensors. From these sensors dedicated short-distance communication lines run to small communication hubs, where the data is aggregated and transferred onto other longer-distance communication lines. This data is then received at a central data aggregation point where it is stored in such a way that it is easily accessible by computer applications that use models to analyse the data and have specific need for a certain sensor, chronological, geometrical and/or geographical point of view, i.e. a subset of all measured data.

2.2 Lessons Learnt

Based on the experiments at the IJKDIJK facility, we have learnt several lessons in the domain of applying information technology for monitoring dikes, which will be elaborated further in this section:

1. Models need to be converted from batch to (near) real time.
2. Development of dike stability models for new sensor types is necessary.
3. When there are no dike stability models, generic anomaly detection techniques can be an alternative.
4. There is a very little syntactical standardization in dike monitoring.
5. There is a very little semantic standardization in dike monitoring.

1. *Models Need to be Converted from Batch to (Near) Real Time.* As described above, while performing dike inspections using sensors in real situations, this is done mostly by using data loggers. After a certain period the collected data is processed using batch-oriented models. For the IJkdijk experiments the sensor data was available in (near) real time, but the available dike stability models were not well-suited for that. A first workaround was a semi real-time analysis by feeding small batches to these models. As described by Langius (2009) and Kruiver (2010), the FEWS-DAM-model (FEWS-DAM stands for Flood Early Warning System extended with Dike stability Analysis Module) for slope stability and piping was adjusted to be able to cope with the real-time sensor values of pore pressure.

2. *Development of Dike Stability Models for New Sensor Types is Necessary.* The available theoretical models are based on parameters like pore-pressure in the dike (at several places) and water height on both sides of the dike. During the experiments companies also wanted to test their innovative solutions on sensors, measuring other parameters, like displacements with in situ sensors (e.g. geobeads, fibre optics) and infrared cameras. Other examples are changes in temperature (e.g. fibre optics) and specific sound (hydrophones). However, for the failure mechanisms under investigation, there were no proven models that took these aspects into account: the models contained no parameters for geotechnical displacements, temperature or sound. Therefore additional (long term) research is needed in future to develop new proven dike stability models for these sensor types (Langius, 2009).

3. *When there are No Dike Stability Models, Generic Anomaly Detection Techniques can be an Alternative.* For those sensor types that have no proven dike stability models, anomaly detection techniques can be of great assistance. The dike stability models are modelling the known failures and provide the probability of actual occurrence of these failures. Anomaly detection techniques can be used to model the normal "good" behaviour of the dikes and demonstrate a probability of how much deviation from that normal behaviour is occurring (see also Kruiver 2010 and look for "trendspotting"). Krzhizhanovskaya (2011) with the Neural cloud from Lang (2008) and Mititelu (2011) with change detection have already experimented with anomaly detection on dike monitoring.

4. *There is a very Little Syntactical Standardization in Dike Monitoring.* Each sensor system partner in the IJkdijk fieldlab had its own syntactical interface to export sensor data which resulted in the rather poor integration of all solutions. There needs to be more confirmation towards the existing international syntactical standards (Weijers, 2009).

5. *There is a very Little Semantic Standardization in Dike Monitoring.* During the fieldlab IJKDIJK it became clear that dike monitoring with sensors is still a very immature area. Often there was no mutual understanding of important dike monitoring parameters such as: pore pressure, dike stability, etc. There is need for more international semantic standardization. Within the Netherlands a logical institute (AQUO) (www.aquo.nl) is founded to develop these standards called the information desk water standards.

3 SMALL SCALE: LIVEDIKES

After the successful IJKDIJK experiments, the water board authorities were confident that monitoring could be safely tested on real dikes. The first real dike equipped with a real time dike monitor system was the harbour dike in the Eemshaven in the north of the Netherlands (see figure 4). Since October 15th, 2009 a segment of 800m of this sea dike are continuously monitored using 208 sensors and a fibre optic cable (Kolk, 2011).

Figure 4: Location of Livedike Eemshaven (blue line). Background image Google Earth (Kruiver, 2010).

Based on a successful first year of this trail, several other "livedike" locations have been developed. Within the Netherlands: Stammerdijk, Vechtkade, Vlaardingse kade (see also www.Livedijk.nl). Since 2010 also several international dikes are being monitored, i.e. in Australia (Brisbane) and United Kingdom (Boston) (see also www.urbanflood.eu).

The goal of the Livedike Eemshaven trail is to gain insight in relevance and usefulness of the use of sensor technology for dike monitoring (Kolk, 2011). Within this paper, the following three derived goals are investigated:

- To learn about placement of sensors at a larger scale, to get already a feeling for large scale deployment.
- To learn about operational issues that occur during monitoring for several years.
- To learn how to set up a proper IT dike monitoring infrastructure.

3.1 Operational Issues

Based on the experience of the last two years of livedike, the following operational issues came to light:

- *Getting Power and Internet on the Location of the Dike is not Trivial.* The livedikes are located in non-urban areas, which is typical for most dikes. Almost never a power line and/or Internet infrastructure is nearby. Even wireless Internet coverage and bandwidth is not always guaranteed in remote dike locations (rural areas). For sensor dike inspection these are non-trivial issues.

- *Some Sensor Systems were more Sensitive to Lightning than Expected.* When lightning stroke at the vicinity of this dike, a lot of the electrical sensor system sensors died (Langius, 2011). The hypothesis is that when the lightning stroke land, the salty water on the other side of the dike was more "attractive" than the parts deeper in the ground. This resulted in a large current through the dike, which was too much for the electrical sensors. However, such type of sensors as the fibre optic sensors gave no problems.

3.2 Lessons Learnt

Based on the experience of the last two years of livedikes, the following IT lessons learnt can already be presented, which will be elaborated further in this section:

1. Adding new sensors into the monitoring system should be automated as much as possible.
2. It should be possible to "correct" measurements.
3. It should be possible to view long periods of measurements.
4. Sample rate increase during flood conditions should be handled in a right way.
5. Simulations are desired.
6. Use of noSQL databases for robustness

1. *Adding New Sensors into the Monitoring System should be Automated as much as Possible.* Most dike monitoring sensor systems use complex sensors which are a combination of many simple sensors. This is especially true for the fibre optics, which can contain hundreds of sensors per kilometre of cable. In order to deploy fast new complex sensors in the database, a template mechanism is needed to pre-specify the layout of that complex sensor previously of installation. That template can then be used to easily instantiate many fibre optic cables, which can contain the measurements.

2. *It should be Possible to "Correct" Measurements.* During the trial the need for multiple versions of the measurements arose. The technique, the pore pressure sensor uses to perform its measurement, is related to the outside air pressure. The sensor itself does not make this correction. To have easier use of the pore pressure, for instance, in the models, it is desired to have both the air pressure corrected version of the pore pressure measurements and the original version (Kruiver, 2010 and Kolk, 2011). In other situations, only the corrections on limited timestamps are needed, i.e., in the case of correction of the measurements of a temporal faulty sensor.

3. *It should be Possible to View Long Periods of Measurements.* To let people interact with the livedikes, a multi touch table was used to display the

sensors within the dikes. The user could easily browse through all the measurements (Krzhizhanovskaya, 2011). During the trial, the time frame of the measurements became larger and larger. Users often zoomed out to the entire period of the trial, more than a year at that time. At some point the used linear sample technique to collect a visualization dataset of 1000 points, was no longer generating representative graphs.

To facilitate the large timeframe viewing request, a different aggregation technique needs to be developed to be able to show a representative graph for each possible timeframe (Kruiver, 2010). Traditional average techniques filter too many details and spikes away. At the moment wavelets are under investigation as a promising aggregation technique, based on (Li, 2002).

4. *Sample Rate Increase during Flood Conditions should be Handled in a Right Way.* The sample rate for the sensors depends on the time scale at which the phenomenon that is to be measured occurs. Other considerations are (Kruiver, 2010): data storage limitations, computer power for data analysis (for larger data sets data analysis takes longer), how important remediation measures are and in which time frame we need to take those measures. And Kruiver also considers if it is useful to transiently adjust the measurement frequency: in periods of higher risk of vulnerabilities, during storms or high water for instance, it might be very useful to increase the measurement frequency, so more data will be gathered and better information about vulnerabilities can be given to the local authorities. For different failure mechanisms different measurement intervals are suggested (e.g. for macro stability: 10 min maximum during flood conditions, otherwise once every hour.)

Kolk (2011) states that (for the specific situation in Livedike Eemshaven) a measurement frequency of e.g. once per hour will be sufficient, but with an automatic increase up to once per 5 minutes as wind speeds of 7 Bft and higher.

5. *Simulations are Desired.* In case of an upcoming crisis, looking at the sensor data and the dike stability models gives not insight that is always enough. One way of attaining additional insight is through the use of simulations where the effect of certain changes in the physical situation (e.g. due to remediation measures) is measured within a simulation; situations which cannot be created for real on an actual dike under threat (Kruiver, 2010).

6. *Use noSQL Database for Robustness.* The data storage of the measurements must be constructed in a very robust and flexible scalable manner. For the IJkdijk experiments and the livedike trails, a traditional SQL database (postgreSQL) is used. For large scale implementations the database partitioning is a realistic problem. In case of a disaster it is not unthinkable that a datacentre actually falls out. The CAP (Consistent, Available & Partitioning tolerant) theorem (Brewer, 2000; Gilbert, 2002 and Langius, 2011) shows that current SQL databases are not robust against partitioning. NoSQL databases make the choice to be available and partition tolerant and are therefore better suited for large scale critical sensor applications (Veen, 2012).

4 LARGE SCALE IT CHALLENGES

Based on the successful trials with IJkdijk and livedikes in the Netherlands, development has been started on a so called Dike Data Service Centre (DDSC) (www.ijkdijk.nl/en/ddsc). The DDSC will become the national (near) real time dike monitoring centre, with facilities for storage as while as providing knowledge. Its first tasks will be the continuation of the operational monitoring of the livedikes and subsequently a mid-scale dike monitoring system.

When looking at the step towards large scale monitoring up to many thousands of kilometres dike, at least the following two IT challenges have to be solved:

1. *How to manage a large number of dike stability model instances for inspection and simulation during normal and crisis situations?* From the lessons learnt of the livedikes, the adaptation of sample rates and the desire for simulations make a dynamic need for CPU power. Temporal increase of sample rate results in a temporal increase of analysis power and each simulation has its own temporal additional need for CPU power. It is, on a cost base, undesirable to scale the dike monitoring datacentre to be able to deal with all of this dynamic CPU power need.

Therefore it is suggested to scale the dike monitoring datacentre based on the baseline CPU power. To deal with sample rate increases and simulations, cloud based CPU power should be requested (asked for). A first attempt in this direction is a cloud model management system developed to set up and configure model instances upon request (Meijer, 2010). In addition, also anomaly detection models as described in the

IJKDIJK section should be treated in the same way as these dike stability models.

2. How to avoid unnecessary running of the CPU intensive dike stability models? Dike stability models can be complex and even in non-critical situations demand a lot of CPU power. Running these models continuously for thousands of kilometres dikes can therefore be quite costly.

It is suggested by Langius (2011) to use simple anomaly detection techniques as a trigger for running the dike stability models. This results in a reduction of using the CPU power most of the time. Only during potentially critical situations additional CPU power can be required in the cloud to run the dike stability models to get insight into the changes of particular dike failure mechanisms.

5 CONCLUSIONS

Based on the work presented in this paper we can state that monitoring dikes using sensor systems in combination with information and communication technology is possible, on a small scale. Based on the lessons learnt we advise:

- to adapt or develop dike stability models to deal with new types of sensors used in the dikes;
- to use anomaly detection techniques when there are no dike stability models available;
- to use noSQL databases to realize a robust (highly available & partitioning tolerant) sensor data storage;
- to work on standardization and semantics for a more mature market where integration of different components is less costly in terms of time and money;
- to apply innovative aggregation techniques to enable viewing data from a large timeframe.

Scaling towards mid and large scale dike monitoring requires solutions to be found for bringing power and Internet (i.e. communication) to the (rural) dike locations.

For large-scale dike monitoring two major IT-challenges have been identified. To cope with these challenges we advise to work in the following two directions.

- To use cloud technology to deal with the dynamic CPU power needs due to sample rate changes and/or simulations.
- In order to reduce the costs of CPU power in non-critical situations, also use anomaly detection techniques to avoid continuous usage of computationally intensive dike stability models.

Finally we want to state that to address these IT lessons learnt, we are developing and combining suitable technologies. Our goal is to make them as generic as possible in order to be useable also in other domains. At the moment we are already involved in projects concerning the monitoring of cracks in steel bridges, ground movement of gas pipes and dairy farming.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the IJkdijk Foundation for making the fieldlab IJkdijk possible and the water board Noorderzijlvest for the Livedike Eemshaven location. And finally the national Flood Control 2015 project and FP7 UrbanFlood project for the research opportunities.

REFERENCES

Bezuijen, A., Kruse, G. A. M.,Van, M. A., 2005. Failure of peat Dikes in the Netherlands, *Proc. 16th Int. Conf. On Soil Mechanics and Geotechnical Engineering*, Osaka, Rotterdam, Millpress, 2005, Vol.3, pp. 1857-1860.

Brewer, E. A., 2000. Towards robust distributed systems, (Invited talk) *Principles of distributed computering*, Portland July 2000.

Gilbert, S., Lynch, N. A., 2002. Brewer's Conjecture and the Feasibility of Consistent, Available: *Partition-Tolerant Web Services*, News, June, 2002.

Koelewijn, A., Pals, N., Sas, M. J., Zomer, W. S., 2010. *IJkdijk Pipingexperiment - Validatie van sensor- en meettechnologie voor detective van optreden van piping in waterkeringen,* ISBN 978.90.5773.485.4, 2010-26 PIW.

Kolk, B. J. van der, 2011. *Livedijk Eemshaven - analyse metingen*, 1200888-001-GEO-0006.

Kruiver, P., Hopman, V., Research team, 2010. Models and analysis for flood control systems, *Flood Control 2015,* 2010.02.02.3.

Krzhizhanovskaya, V. V. et al, 2011. Flood early warning system: design, implementation and computational modules, Procedia Computer Science 4 (2011) 106–115, *International Conference on Computational Science, ICCS 2011*.

Lang, B., Poppe, T., Minin, A., Mokhov, I., Kuperin, Y., Mekler, A., Liapakina, I., 2008. Neural Clouds for Monitoring of Complex Systems. *Optical Memory and Neural Networks (Information Optics)*. 2008. V. 17. N 3, pp. 183–192.

Langius, E. A. F., et al, 2009. *Robust monitoring – Added value of sensor streams in dike monitoring systems*, 2009.02.01.1.

Langius, E. A. F., Meulenhoff, P., Oostveen, J., Vonder, M. R., Waaij, B. D. van der, 2011. Large Scale Sensor Data Analysis - D2.3 rapportage over timelines, trendspotting en geoinformatiestandaarden, *Flood Control 2015*, 2011.07.

Li, T., Li, Q., Zhu, S., Ogihara, M., 2002, A Survey on Wavelet Appliacations in Data Mining, *SIGKDD Explorations*, Volume 4, Issue 2, pp. 49 - 68.

Meijer, R. J., Koelewijn, A. R., 2008. The Development Of An Early Warning System For Dike Failures, Presented at the *Waterside Security Conference 2008,* Copenhagen, Denmark.

Meijer, R. J., 2010. *Functionality & Architecture of Internet Based EWS and EWS Hosting Platforms*, UrbanFlood - D2.2, UFD2.2v7TNO.

Mititelu, G., 2011. *Detecting Change in Levee Sensor Data*, MSc Thesis, VU Amsterdam, Faculty of Sciences, Department of Artificial Intelligence.

Pals, N., Vries, A. de, Jong, A. de, Boertjes, E., 2009. *Remote and in situ sensing for dike monitoring: The IJkdijk experience*, GeoSpatial Visual Analytics NATO Science for Peace and Security Series C: Environmental Security, 2009, Part 7, 465-475, DOI: 10.1007/978-90-481-2899-0_37.

Van, M. A., Zwanenburg, C., Koelewijn, A.R., Lottum, H. van, 2009. *Evaluation of Full Scale Levee Stability Tests at Booneschans and Corresponding Centrifuge Tests*, Proceedings of the 17th International Conference on Soil Mechanics and Geotechnical Engineering (IOS Press).

Veen, J. S. van der, Waaij, B. D. van der, 2012. Sensor Data Storage Performance: SQL or noSQL, Physical or Virtual, to be published in *IEEE Cloud 2012 conference proceedings*.

Vries, G. de, Koelewijn, A. R., Hopman, V., 2010. IJKDIJK Full Scale Underseepage Erosion (Piping) Test: Evaluation of Innovative Sensor Technology; Scour and Erosion: *Proceedings of the Fifth International Conference on Scour and Erosion*.

Weijers, J., Elbersen, G. T., Koelewijn, A. R., Pals, N., 2009. *Macrostabiliteit IJKDIJK: Sensor- en meettechnologie*, ISBN: 978.90.5733.432.8, VIW 2009-19.