

Using a Grid-Enabled Wireless Sensor Network for Flood Management

**Phil Greenwood, Danny Hughes,
Barry Porter, Paul Grace, Geoff Coulson,
Gordon Blair, Francois Taiani**
Computing Department, InfoLab21,
South Drive, Lancaster University
Lancaster, LA1 4WA, UK.
danny@comp.lancs.ac.uk

**Florian Pappenberger,
Paul Smith,
Keith Beven**
Environmental Science Department,
Lancaster University,
Lancaster, LA1 4YQ.
f.pappenberger@lancaster.ac.uk

ABSTRACT

Flooding is becoming an increasing problem. As a result there is a need to deploy more sophisticated sensor networks to detect and react to flooding. This paper outlines a demonstration that illustrates our proposed solution to this problem involving embedded wireless hardware, component based middleware and overlay networks.

Keywords

Grid, overlay networks, components, sensor networks

INTRODUCTION

Flooding is an increasingly serious problem, especially as population pressures lead to increased building on flood plains. Currently, hydrologists deploy sensors at sites susceptible to flooding to record data such as water depth and rate-of-flow. The sensors are deployed statically and sensor data is fed off-site (e.g. using GSM) to grid-based computational models which predict both short and long term flooding trends [1,2]. We see considerable room for improvement in such scenarios. In particular, we propose employing a far more dynamic and autonomous sensor network organisation, and propose (selectively) shifting the execution of the prediction models to the sensor network itself which acts as a mini-grid. This more sophisticated approach promises several benefits. First, the approach can reduce the time hydrologists need to spend on site by increasing the survivability of the sensor networks in case of node or communication failures. Second, locally-computed predictions can be used to dynamically reconfigure the sensor network to optimise it for current environmental conditions (e.g., the network can employ a low power, low throughput organisation in quiescent conditions; and switch to a high power, resilient, high throughput organisation when flooding is imminent). And, third, it can be used to directly inform local stakeholders of imminent flooding.

To achieve this vision we have developed a sensor network framework based on an appropriate combination of software and hardware. Our software is based on our previously developed component-based run-time reconfigurable GridKit middleware [4]. And our hardware platform is based on the embedded Gumstix platform [3] which offers a good

trade off between computational resources (needed to run the prediction models) and power consumption. We call the combined hardware and software platform 'GridStix'.

GRIDKIT

The Lancaster developed GridKit provides the key functionality to develop Grid-applications: service-binding, resource discovery, resource management and security. It is based on our language-independent OpenCOM component model [5], with each area of Grid functionality being implemented as an independent component framework. As GridKit is component based, it is inherently configurable and extensible. A wide range of target deployments can be built by selectively combining component frameworks, either rich and complex, or basic and simple, as appropriate. In particular, minimal deployments can be developed that contain just the bare functionality necessary to perform a particular task. By removing unnecessary components in this way, and thus reducing the associated computation and storage requirements, it is possible to deploy GridKit on scarcely-resourced embedded platforms.

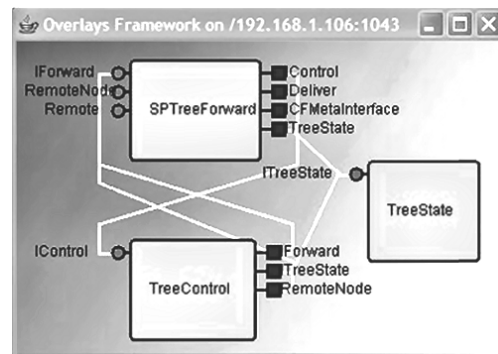


Figure 1: An Open Overlays Configuration

Of particular relevance to the proposed system is the *over-*

lay component framework. The overlay component framework is GridKit's core networking abstraction and provides a base for building application level networks (or overlays), which are typically used to implement services not provided by the underlying network, and to provide functionality that is outside of the scope of the underlying network. Classic examples include: resource discovery, content distribution event notification services, and spanning trees. Figure 1 illustrates a shortest path spanning tree implementation, which may be used to transmit sensor data offsite.

HARDWARE

As mentioned, the hardware used to host the GridKit platform is based on the Gumstix embedded computing platform—so named as each unit is roughly the same size as a pack of chewing gum. Gumstix are significantly more powerful than highly-embedded devices such as the Berkeley Motes [6]. Each unit features an Intel XScale CPU running at up to 400MHz, and comes equipped with 64MB of RAM, 16MB of flash memory and Bluetooth radio. 802.11b connectivity can also be added via a Compact Flash expansion card.

As might be expected, this computational power and flexibility comes at the expense of power consumption. During typical operation a Gumstix consumes up to 1 watt of power with a maximum theoretical power draw of up to 3 watts. This is not a significant problem in our target domain, however, as it is quite feasible for these power requirements to be met using a mid-sized solar panel and high-capacity battery.

SYSTEM OVERVIEW

As described previously, GridStix nodes are capable of maintaining both Bluetooth and 802.11b network infrastructures. This is desirable as the two networks have quite different properties and one can be selected over the other to reflect changing conditions or failures in the network. For example, a set of nodes in close proximity may be disseminating data between themselves in a multi-hop fashion via a low power consumption Bluetooth network. However, a critically placed node in this network could fail, resulting in network partition. To allow the sensor network to continue operating smoothly, our software infrastructure allows a sub-set of the nodes should then switch to 802.11b (due to the improved range) to repair the partition. Additionally, the type of data being disseminated through the sensor network can affect the decision on the network type to use. For example, to detect rate-of-flow in a rising river, we employ an image analysis algorithm to identify and track naturally occurring tracer particles on the water surface. Individual nodes are able to detect very coarse-grained changes in surface velocity; however, for more precise measurements to be performed, images need to be distributed to a number of nodes. As a consequence, the GridStix will switch network type from Bluetooth to 802.11b as Bluetooth does not have sufficient bandwidth to distribute the large image data-set in a timely fashion.

Adaptations in the sensor network can also take place at the overlay/spanning tree level, in that different overlays can be substituted at run-time depending on the environmental conditions. The GridStix nodes might initially be structured using a shortest-path tree. As 'shortest path' corresponds to the distance between nodes, it correlates with the amount of power necessary to reliably transmit data between nodes, resulting in a power-efficient topology. However, trees of this nature tend to be 'skinny' which reduces their resilience to failure. When flooding is predicted it is therefore desirable to increase resilience to failure rather than conserve power. As a result, a fewest-hop tree can be dynamically substituted, thus increasing system resilience. Fewest-hop trees tend to produce 'fat' topologies with each node having fewer children. This means fewer nodes are affected when one particular node fails. Example shortest path and fewest hop spanning trees are shown in Figure 2.




Figure 2: SP and FH Spanning Trees

SUMMARY

This paper has given a brief overview of how we have combined embedded hardware and a component-based Grid platform to offer improved support for flood prediction. This combination allows the sensor network to be adapted in rich ways to best suit the current and future environmental conditions. This increases the utility, the resilience and the performance of the network.

REFERENCES

1. F. Pappenberger, K. Beven et al., "Cascading model uncertainty from medium range weather forecasts (10 days) through a rainfall-runoff model to flood inundation predictions within the European Flood Forecasting System (EFFS)", *Hydrology and Earth System Science*, 9(4), pp381-393, 2005.
2. F. Pappenberger, P. Matgen and K. Beven, "The influence of rating curve and structural uncertainty on flood inundation predictions.", *Advances in Water Resources*. (in press) 2006.
3. Waysmall Computers "Gumstix Embedded Computing Platform Specifications", website: <http://gumstix.com/spexboards.html>.
4. G. Coulson, P. Grace, G. Blair et al, "Open Overlay Support for the Divergent Grid", in the proceedings of the UK E-Science All Hands Meeting, Nottingham, UK, September 2005.
5. G. Coulson, G. Blair et al "The Design of a Highly Configurable and Reconfigurable Middleware Platform", *ACM Distributed Comp Journal*, Vol. 15, No 2, pp109-126, April 2002.
6. Crossbow "Mica Motes and Sensors", website: <http://www.xbow.com/>