

# An Intelligent and Adaptable Grid-based Flood Monitoring and Warning System (DRAFT)

Danny Hughes<sup>1</sup>, Phil Greenwood<sup>1</sup>, Gordon Blair<sup>1</sup>, Geoff Coulson<sup>1</sup>,  
Florian Pappenberger<sup>2</sup>, Paul Smith<sup>2</sup> and Keith Beven<sup>2</sup>

<sup>1</sup>Computing Department, Infolab21, Lancaster University, UK, LA1 4WA.

<sup>2</sup>Environmental Science Department, Lancaster University, UK, LA1 4YQ.

## Abstract

*Flooding is a growing problem in the UK. It has a significant effect on residents, businesses and commuters in flood-prone areas. The cost of damage caused by flooding correlates closely with the warning time given before a flood event, and this makes flood monitoring and prediction critical to minimizing the cost of flood damage. This paper describes a wireless sensor network for flood warning which is not only capable of integrating with remote fixed-network grids for computationally-intensive flood modeling purposes, but is also capable of performing on-site flood modeling by organising itself as a 'local grid'. The combination of these two modes of grid computation—local and remote—yields significant benefits. For example, local computation can be used to provide timely warnings to local stakeholders, and a combination of local and remote computation can inform adaptation of the sensor network to maintain optimal performance in changing environmental conditions.*

## 1. Introduction

Flooding is a growing problem in the UK and affects a large number of people financially, physically and emotionally. The problem was dramatically highlighted by the wide-spread floods of Autumn 2000, the total cost of which was estimated to be in the order of £1 billion. Following these floods, major initiatives [Environment Agency '00] have been undertaken to improve the UK's flood readiness. These include: improving flood defenses; raising public awareness; and, significantly for this project, improving flood warning systems.

Traditionally, hydrologists have approached flood prediction by deploying sensors (such as depth and flow-rate sensors) at sites prone to flooding. Data from these sensors is then collected manually or via GSM-based telemetry and used as the input to flood prediction models. Two main classes of flood prediction model are commonly used. The first, referred to as *spatial* models [Pappenberger '05], provide detailed, site-wide predictions, albeit with limited accuracy at any given point. They are computationally complex and must be executed on clusters or grids. The second class, referred to as *point prediction* models [Beven '05], provide accurate depth predictions for a single point in the flood plain. They are computationally simple so that they may be

executed in a timely fashion on standard desktop PC hardware.

In summary, then, traditional flood monitoring approaches impose a rigid separation between the on-site wireless sensor networks (WSNs) that are used to collect data, and the off-site computational grid which is used to analyze this data. Essentially, the sensor networks are computationally 'dumb', being composed of nodes that are capable only of recording and transmitting sensor data.

In order to better support timely flood warnings, we argue that more on-site 'intelligence' is required. The 'GridStix' sensor platform presented in this paper uses powerful embedded hardware, heterogeneous wireless networking technologies, and next generation grid middleware to implement an adaptable WSN that doubles as a lightweight grid, allowing nodes to not only ship data to remote fixed grids, but also to perform 'local' grid computations with significant benefits as discussed below.

This paper describes the operation of a GridStix-based flood monitoring system and specifically focuses on how local grid computation can be used to support the adaptation of the WSN to changing environmental conditions. The remainder of this paper is structured as follows. First, section 2 discusses how local computation can be exploited in WSNs. Section 3 then introduces our 'GridStix' platform, section 4 highlights potential forms of adaptation that are available

using that platform, and section 5 discusses factors that can be used to drive adaptation. Finally, section 6 discusses our ongoing deployment and evaluation work, section 7 discusses related work, and section 8 offers conclusions and outlines our plans for future work.

## 2. Exploiting Local Computation

Our prototype flood prediction system uses local grid computation to provide improved support for flood monitoring. This section discusses how local computation can be used to (i) inform system adaptation, (ii) support diverse sensors and (iii) provide timely warnings to local stakeholders.

First, local computation can be used to drive the *adaptation of WSN behaviour* based on awareness of environmental conditions such as flood data and power monitoring. For example, based on the execution of point prediction models, we can switch to a more reliable network topology (see below) at times when node failure seems more likely (i.e. when imminent flooding is predicted). Adaptation may also be informed by input from computationally intensive spatial prediction models executed in the remote fixed grid.

Second, the availability of local computation can support *richer sensor modalities* such as image-based flow prediction [Bradley '03]. Image-based flow prediction is a novel technique for measuring water flow rates that uses off-the-shelf digital cameras. It is cheaper and more convenient to deploy than the commonly-used ultrasound flow sensors, but can only be used where significant computational power is available. Flow-rate measurements are calculated based upon a series of images taken by a digital camera deployed overlooking the river. Naturally occurring tracer particles are identified on the water surface and tracked through subsequent images, from which the surface velocity of the water is inferred. The data-set used by this method, a sequence of high-resolution images, is too large for off-site transmission to be feasible using GSM or GPRS technologies and therefore the method is impractical in current sensor network deployments. However, organising computationally-capable sensor nodes into a local grid allows analysis to be performed on-site and the results of this analysis then transmitted off-site.

Finally, on site flood-modeling allows *timely flood warnings* to be distributed to local stakeholders. These flood-warnings are based

on the results of point prediction models executed by the local grid of GridStix nodes and disseminated to local stakeholders in a range of formats including on-site audio/visual warnings, a public web-site and SMS alerts. Each of these media has associated benefits and drawbacks. For example, SMS warnings are an effective method of publishing timely alerts to local stakeholders. However, SMS warnings require that users register for the service in advance and are therefore ineffective for stakeholders who might be unaware of a flood risk. Local audio/visual flood warnings may be effective without the need for stakeholders to proactively participate, however their effectiveness is dependent upon stakeholders being within audio/visual range.

## 3. The GridStix Platform

### 3.1 Overview

In order to achieve the 'local grid' functionality discussed in the previous section, a powerful and versatile sensor platform is required (in terms of both hardware and software). This section describes such a platform – GridStix. More information on the platform itself is given in [Hughes '06].

### 3.2 Hardware Platform

In order to support the proposed functionality, a sensor node device must be capable of interfacing with a variety of sensors including traditional sensors (e.g. depth sensors) and more novel sensors (e.g. the digital imaging hardware that is used to support the image-based flow prediction discussed above). A suitable device must also be capable of supporting a variety of wireless communications technologies to provide redundancy and allow sensor nodes to switch physical network technologies as conditions require. Finally, the device must have sufficient computational and storage resources to support the GridKit software platform (see below).

Sensor networks often make use of devices with extremely constrained local resources such as the Berkley Motes [Xbow '06]. This is because such devices have extremely modest power requirements and can therefore operate for long periods on small batteries. However, such constrained platforms do not offer sufficient computational power to support functionality such as on-site flood prediction, nor do they offer sufficient support for diverse networking technologies and sensor types. For this reason, more powerful embedded hardware

has been selected for use in the GridStix platform.

Each GridStix node is based on the *Gumstix* [Waysmall '06] embedded computing platform, so named as each device is roughly the same size as a pack of gum. Despite their small-size, each of these devices is equipped with a 400 MHz Intel XScale PXA255 CPU, 64Mb of RAM and 16Mb of flash memory. These hardware resources support the execution of a standard Linux kernel and Java Virtual Machine making them inherently compatible with the GridKit platform, which has been successfully deployed on comparable hardware such as PDAs [Cooper '05]. Furthermore, the PXA255, which performs comparably to a 266MHz Pentium-class CPU are capable of executing a single iteration of a point prediction model in a matter of seconds.

The Gumstix devices also provide a variety of hardware I/O mechanisms, enabling connection to a variety of sensors. For example, a network camera (for image-based flow measurement) can be connected via a standard wired Ethernet connection, while flow sensors can be connected via an on-board serial port, and depth sensors can be connected via the GPIO lines of the XScale I2C bus. In this way it is possible to connect multiple sensors to a single device. In terms of networking, each device is equipped with an onboard Bluetooth radio and Compact Flash 802.11b network hardware, which is used to provide an ad-hoc communications infrastructure. Furthermore, the devices can be equipped with GPRS modems for transmitting and receiving data from off-site.

Of course, the above capabilities come at the expense of increased power consumption: While a Berkley Mica Mote consumes only 54mW during active operation [XBow'06], our devices consume around 1W during typical operation, and thus it would not be feasible to power them for long periods using batteries alone. To address this, solar panel arrays are employed. Given aggressive power management, we have found that a single 15cm<sup>2</sup> mono-crystalline solar panel, with a maximum output of 1.9 watts combined with a 6v 10AH battery, is sufficient to continually power a device.

Finally, to minimise the effects of harsh weather conditions, flood water, vandalism, uncooperative grazing animals, etc., we have housed the devices in durable, water-tight containers that can safely be buried. In some cases burial is not possible (e.g. solar panel deployment). In such cases, the device is

situated as discreetly and securely as possible to avoid unwanted attention.

### 3.3 The GridKit Middleware Platform

The GridKit middleware platform [Coulson '05] provides the key functionality that is required to support distributed systems such as grids, peer-to-peer networks and WSNs. GridKit is based on the OpenCOM [Coulson '02] component model and the various facets of system functionality are implemented as independent component frameworks. This component-based approach allows developers to build rich support for distributed systems or conversely, to build stripped-down deployments suitable for execution on embedded hardware such as the Gumstix.

Importantly, GridKit offers rich support for *application-level overlay networks*. Its Overlay Framework [Coulson '05] supports the simultaneous deployment of multiple overlay networks and enables these to interoperate flexibly (e.g. by layering them vertically or composing them horizontally). It also supports the adaptation of overlays, allowing, e.g., one overlay to be swapped for another at run-time. The use of adaptable overlays is discussed in detail in section 4, which illustrates how overlays with different performance characteristics can be used to adapt to changing environmental conditions.

## 4. Supporting Adaptation

### 4.1 Overview

This section examines situations in which WSN adaptation is possible and then goes on to consider the factors that can be used to inform such adaptation. Three discrete classes of adaptation are identified: (i) adaptation at the level of the physical network, (ii) adaptation at the overlay network level, and (iii) adaptation of CPU performance.

### 4.2 Physical Network Adaptation

Our flood prediction system makes use of three wireless networking technologies, Bluetooth, IEEE 802.11b and GPRS, each of which has very different performance characteristics:

- The compact flash 802.11b hardware (SanDisk Connect Plus) supports speeds of 11Mbps at a range of 137 meters, 5.5Mbps at 182 meters, 2 Mbps at 243 meters and 1Mbps at 365 meters. It offers significantly better performance than Bluetooth or GPRS

and has a maximum power consumption of approximately 0.5 watts.

- The class 2 Bluetooth radio (Ericsson ROK-104-001) supports speeds of up to 768Kbps at a range of up to 25M. It offers QoS that is significantly lower than 802.11b, but significantly higher than GPRS. It consumes a maximum of 0.2 watts [Ericsson '02].
- The GPRS modem (Ambicom GPRS-CF) supports uplink speeds of up to 29kbps and downlink speeds of up to 58kbps. Range is not an issue as we assume that the entire deployment site is within the bounds of GPRS coverage. GPRS offers lower QoS than either 802.11b or Bluetooth and consumes a maximum of 2 watts. Its operating wavelength of around 30cm is longer than that of the other technologies (which operate at around 12cm) and therefore performs better under water [Ambicom '06].


Each of our three communication technologies clearly has advantages and disadvantages. For example, 802.11b offers good QoS and long range; however, it consumes significantly more power than Bluetooth. Conversely GPRS offers much poorer performance, but is not limited by range. We discuss below how these differing characteristics can best be exploited.

### 4.3 Application-level Overlay Adaptation

As previously discussed, we employ application-level overlays to provide communications support for our flood prediction system. There are a range of overlays which can be used and each has advantages and disadvantages.

Off-site data dissemination is supported by the use of *spanning tree-based overlays*. These are commonly used in WSNs to disseminate data from a large number of sensors to a small number of logging or bridging nodes which form the 'root' of the tree. Prime examples of spanning trees are Shortest Path (SP) and Fewest Hop (FH) trees. FH trees are optimised to maintain a minimum of hops between each node and the root. They minimise the data loss that occurs due to node failure, but are sub-optimal with respect to power consumption. SP trees, on the other hand, are optimised to maintain a minimum distance in edge weights from each node to the root. As a result, they tend to consume less power than FH trees, but are more vulnerable to node failure. Both forms of tree can be efficiently created using

Dijkstra's algorithm [Dijkstra '59]. Examples of SP and FH spanning trees are shown in figure 1.




**Figure 1: FH and SP Spanning Trees.**

SP and FH are just two common spanning tree types and many others are also available. We are currently investigating the performance of a range of spanning trees for off-site data dissemination. Nevertheless, these two examples serve to illustrate the trade-off that often exists between overlay performance and power consumption.

### 4.4 CPU Power Adaptation

The XScale PXA255 CPUs used in the GridStix platform support software controlled frequency scaling, which allows the CPU to be set at a variety of speeds from 100MHz to 400MHz.

Processing power increases with clock speed but at the cost of increased power consumption. Figure 2 shows the relationship between the clock frequency of the XScale CPU and the power it consumes.



**Figure 2: XScale Power Consumption.**

Of course, power consumption is also affected by the way the CPU is used by applications and therefore power management could also be implemented by explicitly controlling the manner in which local processes are scheduled (e.g. by modifying the frequency with which local point-prediction models are executed). Nevertheless, the XScale's support for software

control of clock frequency provides convenient, coarse grained adaptation of CPU power consumption.

## 5. Adaptation Scenarios

### 5.1 Overview

In the above, we have presented a number of ways in which the behaviour of GridStix nodes can be adapted. However, this is only half of the story—for adaptation to be useful and meaningful it must be suitably informed by relevant real-world triggers.

We now present three adaptation scenarios which demonstrate how awareness of real-world conditions can be used to maintain optimal system operation in changing environmental conditions. The first scenario considers situations in which sensor nodes might become immersed in water; the second considers adapting to total node failure, and the third considers adapting changes in criticality. These scenarios are not exhaustive; however they do demonstrate how local computation can be used to optimise WSN behaviour and thus produce a more useful and more robust flood prediction system.

### 5.2 Adapting to Node Immersion

In a flood-prone area, sensor nodes clearly face the risk of immersion. While effort is made to deploy nodes above the likely level of flood water, this is impossible to guarantee. As each GridStix is sealed in a waterproof enclosure, immersion, even in deep water, does not actually damage the node. However, significant immersion prevents power being produced by the solar panels and also adversely affects wireless communication technologies.

Longer wavelength wireless technologies such as GPRS are better at penetrating water than short wavelength technologies such as 802.11b or Bluetooth. Therefore, when a node becomes submerged, or the node predicts that submersion is likely, GPRS communication should be used as this is more likely to result in a sustained connection. During normal operation however, GPRS is the poorest choice for on-site communication due to its low bandwidth, low QoS and high power consumption. Therefore under normal conditions 802.11b or Bluetooth would be used.

Emergent effects of submersion may lead to the need for further adaptations. For example, a lack of power being produced by a node's solar panel might indicate a reduction of CPU clock

speed or lowering the frequency with which the flood models are executed.

### 5.3 Adapting to Node Failure

Alongside the previously-discussed risk of node immersion, sensor nodes are at significant risk of damage or destruction due to being swept away by flood water or due to collision with debris. This risk may be assessed using on-site flow rate measurements. While it is impossible for a failed node to adapt its behaviour in this instance, the impact that a node's failure has on the WSN as a whole is highly dependent on the application level overlay that is being used to disseminate data offsite.

Consider the spanning trees introduced in section 3: Shortest Path (SP) trees consume less power than Fewest Hop (FH) trees and therefore during normal operation, off-site dissemination of sensor data should be performed using an SP spanning tree. However, when on-site flow measurements indicate an increased risk of node failure, the system should switch to an FH spanning tree, which is significantly more resilient to node failure. In this way power-consumption is minimised during normal operation, whilst resilience is preserved at times of high risk.

### 5.4 Adapting to Changes in Criticality

As previously described, local point predictions are used to provide timely warnings for local stakeholders. When such a warning is in place, the computation of local flood warnings becomes more time-critical. Where initial flood warnings were accurate, local predictions can be used to show likely paths of inundation, and in the case of erroneous warnings, local predictions can be used to lift the flood warning. The latter is particularly critical as flood preparation and particularly evacuation is an expensive activity.

During normal system operation, the timely execution of flood warnings is not particularly critical and therefore, nodes can scale down their CPU speed to 200MHz to minimise CPU power consumption (to 0.3 watts). However, when stakeholder flood warnings are in place and the computation of flood warnings becomes more critical, the nodes can increase their CPU speed to the maximum of 400MHz (0.52 watts). In this way, the system conserves power during normal operation while maintaining the ability to provide timely flood warnings in critical situations.

The rate at which the sensors collect data can also be increased during these critical times.

This enables the system to provide more frequent and accurate predictions though at the cost of increased network, CPU and power resources.

Note that the performance maximisation actions taken when flood prediction becomes time-critical are in direct conflict with the power saving actions taken when battery power runs low (see above). In cases where logical adaptive actions conflict, the system must determine which action is most critical. We are currently in the process of assessing this through evaluation of system performance. Furthermore, this is by no means the only situation in which adaptations interfere with each other. In particular, there are a number of ways in which the type of physical network in use affects the optimal choice of overlay, and vice versa.

## 6. Deployment and Evaluation

### 6.1 Deployment

The system described above has already been built and tested in the lab and we are now preparing to deploy it in a real-world environment. The planned deployment site is at Cow Bridge, which is located on the River Ribble in the Yorkshire Dales. This site is prone to flooding for much of the year and thus offers good potential for evaluating the system under real-world conditions. Flooding at the site affects the nearby village of Long Preston, which thus additionally presents us with a motivation for evaluating warning systems for local stakeholders. The site is largely rural which minimises the risk to deployed hardware due to theft and/or vandalism. We anticipate initial deployment during summer 2006, when instances of flooding are relatively uncommon.

Deployment at the site will cover approximately 1km of river with an initial installation of 13 nodes. The majority of sensors deployed will be depth sensors, along with a single image-based flow sensor and a single ultrasound flow sensor.

### 6.2 Further Testing and Simulation

While the Cow Bridge deployment provides a realistic environment for evaluating system performance, it has a number of significant limitations such as a limited scale of deployment, the unpredictability of flood events, and the time required to perform tests. We are therefore concurrently pursuing lab-based and simulation-based testing of the

system to gain more insight into its generic applicability.

In particular, we are currently engaged in assessing the performance of the solar panels in terms of the power they produce in various weather conditions (hours of daylight, cloud cover etc.), and are assessing battery performance in terms of charge retention. In addition, we are assessing the performance of the physical networking hardware (802.11b, Bluetooth and GPRS) in terms of power consumption and QoS characteristics such as throughput, loss, delay and jitter. These factors are being evaluated with various usage profiles in various weather conditions and under varying levels of immersion.

Based upon these basic performance characteristics, a simulator is being constructed that will allow the testing of various application-level networks and power management strategies, using site-specific topographical information, past weather conditions and past flooding data. This will be used to prototype potential deployment technologies and investigate the performance of large-scale deployments that could not be easily tested in the real world.

## 7. Related Work

A number of grid-related projects have addressed the issues of WSN-grid integration in general and WSN-based flood prediction in particular. A prime example of the former is the Equator remote medical monitoring project [Rodden '05], and a prime example of the latter is Floodnet [DeRoure '06]. However, these systems (to the best of our knowledge) all employ a 'dumb' proxy-based approach to integrating WSNs with the grid and thus cannot take advantage of the local computational power that we employ to drive the adaptation of WSN behaviour, to support richer sensor modalities such as image-based flow prediction, and to provide timely flood warnings to local stakeholders.

## 8. Conclusions and Future Work

This paper has described a WSN that is capable of performing not only remote off-site flood modelling based on grids in the fixed network, but also local on-site flood modelling using a lightweight grid built on our GridStix platform.

The key difference between our system and existing work on WSN-grid integration is that our work aims to promote the sensors to *first class* grid entities. This allows a greater degree

of integration and flexibility than those approaches that treat sensor networks as conceptually distinct from the grid. In particular, for our flood prediction scenario it allows us to more effectively support WSN adaptation, to support richer sensor modalities, and to enable proactive behaviour such as informing local stakeholders of pending flooding.

In future research we are especially planning to work on improving our system's adaptation mechanisms. Currently, our adaptation policies are manually implemented, but we plan in the future to investigate the extent to which nodes can 'learn' appropriate adaptation behaviour. As an example, consider the performance of power management approaches. If nodes were capable of autonomously selecting appropriate power management strategies, it would significantly reduce the time-to-deployment for novel environmental monitoring applications. To accomplish this, sensor nodes could successively load different power management policies and, based on the relative success of these policies, select the most appropriate one for a given environmental monitoring scenario, or set of environmental conditions.

Currently, the information used to inform system behaviour originates exclusively from within the system itself. However, external information might also provide valuable information on which adaptation could be based. For example, local weather predictions, particularly predicted hours of sunlight, could be used to better inform battery-life models (due to fluctuations in the power captured by solar panels).

A final area of planned future work is to investigate how our WSN's functionality may be expanded from an exclusively monitoring role to additionally encompassing flood-response support. For example, real-time on-site visualisation of flood models would be useful for the emergency services who could use this data to inform the placement of sand bags and other flood defences. Similarly, the digital cameras deployed to perform image-based flow measurement could be switched to providing real-time remote imaging for flood responders. This adaptation of node roles necessitates not only modifications to local functionality, but also imposes new requirements for the supporting physical and application-level networks.

## References

- [**Environment Agency '00**] The Environment Agency "Lessons Learned, Autumn 2000 Floods", available online: <http://www.environmentagency.gov.uk/commodata/acrobat/126637>
- [**Pappenberger '05**] F. Pappenberger, K. Beven, N. Hunter 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)", published in *Hydrology and Earth System Science*, 9(4), pp381-393, 2005.
- [**Beven '05**] Beven, K J, Romanowicz, R, Pappenberger, F, Young, PC, Werner, M, "The Uncertainty Cascade in Flood Forecasting", ACTIF meeting on Flood Risk, Tromsø, in press, 2005.
- [**Bradley '03**] Creutin J. D., M. Muste, A. A. Bradley, S. C. Kim, "River Gauging using PIV Techniques: A Proof of Concept on The Iowa River", A. Kruger, *Journal of Hydrology* 277, 182-194, 2003.
- [**Hughes '06**] Hughes D., Greenwood P., Coulson G., Blair G., Pappenberger F., Smith P., Beven K., "GridStix: Supporting Flood Prediction using Embedded Hardware and Next Generation Grid Middleware", to be published in the proceedings of the 4th International Workshop on Mobile Distributed Computing (MDC'06), Niagara Falls, USA, June 2006.
- [**Coulson '05**] 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.
- [**Coulson '02**] G. Coulson, G. Blair, M. Clark et al "The Design of a Highly Configurable and Reconfigurable Middleware Platform", in the *ACM Distributed Computing Journal*, Vol. 15, No 2, pp109- 126, April 2002.
- [**Rowstron '01**] A. Rowstron, P. Druschel. "Pastry: Scalable, Decentralised Object Location and Routing for Large Scale Peer-to-Peer Systems" – Conference on Distributed Systems Platforms, Heidelberg, Germany 2001.
- [**Gnutella '00**] Limewire, "The Gnutella Protocol Specification v0.4", available online: [http://www9.limewire.com/developer/gnutella\\_protocol\\_0.4.pdf](http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf)
- [**XBow '06**] Crossbow "Mica Mote Data Sheet", available online: [www.xbow.com/Products/Product\\_pdf\\_files/Wireless\\_pdf/MICA.pdf](http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/MICA.pdf)

**[Waysmall '06]** Waysmall Computers  
"Gumstix Embedded Computing Platform  
Specifications", website:  
<http://gumstix.com/spexboards.html>.

**[Cooper '05]** "The Open Overlays  
Collaborative Workspace Environment", Chris  
Cooper, David Duce, Muhammad Younas, Wei  
Li, Musbah Sagar, Gordon Blair, Geoff  
Coulson, Paul Grace, ", In Proceedings of the  
UK E-Science All Hands Meeting 2005,  
Nottingham, September 2005.

**[Conti '05]** "A Cross-layer Optimization of  
Gnutella for Mobile Ad hoc Networks", Conti  
M., Gregori E., Turi G., Published in the  
proceedings of the 6th ACM international  
symposium on Mobile ad hoc networking and  
computing (MobiHoc '05), IL, USA, 2005.

**[DeRoure '05]** D. DeRoure, "Improving Flood  
Warning times using Pervasive and Grid  
Computing", submitted to quarterly of Royal  
Academy of Engineering, UK.

**[Rodden '05]** T. Rodden, C. Greenhalgh, D.  
DeRoure, A. Friday, "Extending GT to Support  
Remote Medical Monitoring," in the  
proceedings of UK e-Science All Hands  
Meeting, Nottingham, UK, September 2005.

**[Ericsson '02]** "Ericsson ROK 104 001 Data  
Sheet", available online:  
<http://prism2.mem.drexel.edu/~billgreen/Bluetooth/rok104001.pdf>

**[Ambicom '06]** "GPRS Compact Flash Card  
For Laptop And PDA Specifications" available  
online:  
<http://www.ambicom.com/products/gprs/gprs-feat.html>

**[Dijkstra '59]** E. W. Dijkstra: "A note on two  
problems in connection with graphs" in:  
Numerische Mathematik. 1 (1959), S. 269–271