A LOW POWER, DIGITAL TRANSCEIVER FOR WIRELESS SENSOR NETWORKS

L. H. Crockett, N. C. MacEwen, E. Pfann, R. W. Stewart

DSP Enabled Communications Group, University of Strathclyde, Glasgow, Scotland, UK. G1 1XW.
specknet@eee.strath.ac.uk

Keywords: sensor networks, low power radio, hardware implementation, synchronisation, pulse shaping

Abstract

This paper introduces a new, ultra-small ad-hoc sensor network technology, and reviews the requirements on a DSP physical layer to support RF transmissions between the severely power-limited nodes therein. Specifically, we present results from design and implementation work on a low power, digital transceiver which uses Manchester encoding to reduce the synchronisation overhead. With power consumption in mind, we consider the effect of the physical layer design on clock oscillator specification, and the addition of pulse shaping filters to enhance the efficiency of inter-speck transmissions. These evaluations are made in the context of 125mm^3 nodes transmitting over channels of up to 10cm, and with limited energy storage capacity.

1 Introduction

The Speckled Computing Consortium[1] is a collaboration of research groups in five universities focusing on the research and the realisation of a wireless ad hoc sensor network comprising tiny nodes (or "specks"). Each of these specks will be autonomous, and capable of communicating via both optical and radio frequency (RF) links. Individually, the specks must have modest resources due to their physical size (the initial target volume is 125mm^3), but when their collective power is harnessed in a "SpeckNet" – which will typically comprise several hundreds or even thousands of specks - significant processing capability will be realised.

Specknets are envisioned as sensor networks and as such will process and relay sensor data. Early versions are intended to support data rates of a few kilobits per second, with this figure increasing in the longer term.

Figure 1 shows an overview of the Speckled Computing research interests (reproduced from [11]).

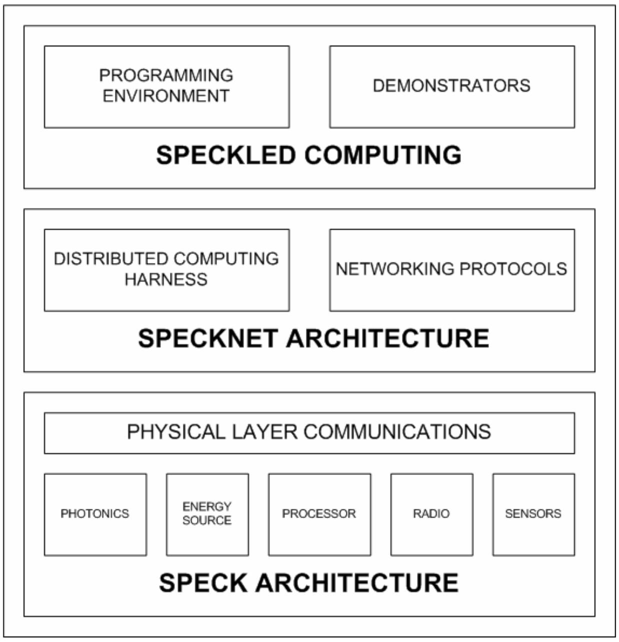

Figure 1: Overview of Speckled Computing Research Interests (reproduced from [11])

A first generation Specknet will be developed over the next few years, comprising specks of approximate size 5 x 5 x 5 mm. Work is currently underway on the specks, which will be capable of simple processing using a commercially available microcontroller [2], and RF communication. Details of this work will be presented in Section 4.

Each speck will be equipped with its own energy source within the stated volume and as volume is severely restricted, battery dimensions must be kept to a minimum. At present, lithium ion battery technology provides the highest volumetric energy density [9]. For the first generation specks the smallest currently commercially available rechargeable lithium batteries [4] will be suitable, but nevertheless a transceiver design exhibiting low power consumption is vital in order to provide a useful lifetime. Therefore minimising the power consumption of the design - both in terms of expenditure on the speck and transmissions between specks - is the critical factor.

Section 2 provides a background for the particular issues involved in enabling RF communications in Specknets. Section 3 focuses on the digital transceiver, and in particular describes the adopted strategy of Manchester encoding to reduce the synchronisation overhead. Section 4 provides details of work on the prototype to date, Section 5 discusses pulse shaping as a potential improvement to the system, and

Section 6 outlines intended future work. Finally, the paper is concluded in Section 7.

2 Inter Speck Communications

Wireless sensor networks place particular demands upon the physical layer [10]. With energy storage capacity tightly limited, the need to minimise power consumption – not just in the DSP physical layer, but in all hardware design, and also the communications protocol - is crucial.

In the following subsections, we examine briefly some of the key aspects of the communications system, primarily from a physical layer perspective.

2.1 The Channel

In the first instance, the Specknet Consortium has elected to use the unlicensed band at 2.4GHz. In future versions a higher frequency band may be used; however analysis in this paper refers to the 2.4GHz band.

It is shown in [5] that for a wireless sensor network with robustness to single-link deep fades, equalisation is not required at symbol rates up to 200ksps. Note that the specks will communicate over very short ranges (initially ~10cm), at which delay spread will be much smaller than that reported in [5]. Early studies show that data rates of up to 400ksps will be possible without equalisation [6].

2.2 RF Section

The RF section is often the part of a transceiver which consumes most power, particularly in carrier based systems employing coherent demodulation. In this type of design, active components such as mixers and phase locked loops (PLLs) add significantly to the total power consumed. The design of an RF circuit for Speckled Computing therefore must, for energy conservation reasons, use an alternative method which involves as few active components as possible.

Possible alternatives include carrier-less transmission such as pulse position modulation ultra wideband (PPM-UWB); and non-coherently demodulated carrier-based systems. As PPM presents complex synchronisation problems, the latter strategy has been chosen.

Clearly the choice of RF modulation and detection methods impacts on the requirements for the DSP physical layer design. Details of the RF design will be discussed in Section 3.

2.3 Modulation

Again with the aim of reducing the complexity of the design, and therefore the power consumed, the modulation scheme must be chosen carefully. Data rates are envisaged to be comparatively low, i.e. at most tens of kbps. Bandwidth constraints are thus not a problem and consequently the simplest possible modulation scheme, On-Off Keying (OOK) [8], can be employed. Using this scheme is also beneficial from an RF design perspective as it only requires the RF oscillator to be turned on and off using the information signal. This also removes the need to include a mixer in the RF transmitter circuit.

However, a particular drawback of OOK is the sharp transitions present in the signal: these cause its bandwidth to spread outside the intended frequency band. This unwanted spreading negatively affects spectral efficiency and detection and also, depending on the frequency band, has implications for adherence to spectral usage regulations.

In Section 5 we describe simulation work on pulse shaping, which has been undertaken to identify an appropriate filter for incorporation into the transceiver in subsequent design iterations. Doing so will mitigate the undesirable effects described above.

2.4 Access Technique

In order for a Specknet to function successfully, many specks must coexist in the same radio space, and thus some form of medium access must be implemented. In the long term it is believed that a code division multiple access (CDMA) technique will provide the best solution, as it will take advantage of the individual identity given to each speck, and avoid the undesirable issues involved with time and frequency division multiple access (TDMA/FDMA) (i.e. the need for time synchronisation or antennas accurately tuned to different frequencies).

The first generation Specknet will, primarily in the interests of creating a low complexity transceiver, employ a Carrier Sense Multiple Access (CSMA) strategy, allowing random access of a single frequency channel.

3 Transceiver

Figure 2 shows a model of the basic transceiver used in the first generation speck. In order to ensure that the implementation is as simple and consumes as little power as possible, an OOK modulation scheme has been selected for RF radio communication. A local oscillator directly modulates the baseband data and an envelope detector is used to perform non-coherent down-conversion at the receive side. This system removes the need for any active components such as power-hungry mixers and PLLs in the receiver.

Once the RF signal has been demodulated by the envelope detector, it is sliced using an analogue comparator, and the recovered baseband data is input directly to the digital receiver. The slicer employs a hysteresis region in order to ensure that low level noise is not presented as data to the digital part of the receiver. At present design, manual gain

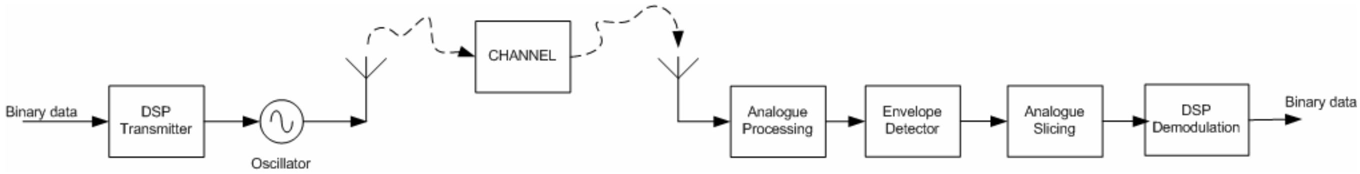

Figure 2: Basic transceiver block diagram for specks

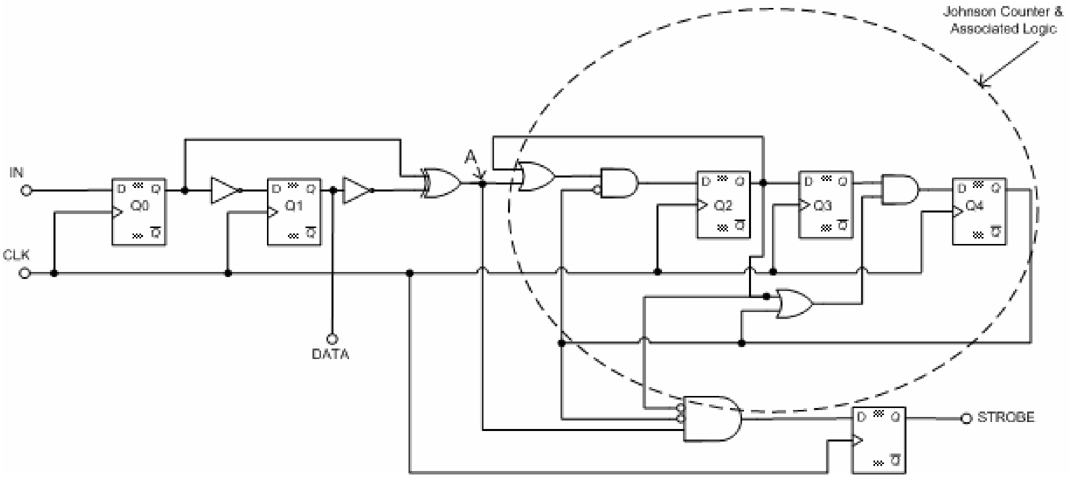

Figure 3: Logic circuit diagram of the Manchester decoder implementation

control is employed, wherein the threshold level of the comparator is determined at a higher level in the protocol stack, and communicated to the physical layer of the transceiver by the microcontroller. Future designs of the specks will incorporate automatic gain control (AGC).

The digital transceiver therefore must include baseband modulation and demodulation, and frame and symbol synchronisation. A simple method of symbol synchronisation is to embed the transmitter clock in the transmitted data and then extract it at the receiver. This avoids expensive digital synchronisation circuitry. For the first generation speck, a Manchester-encoded transceiver has been implemented.

Manchester encoding is a form of bi-phase pulse code modulation (PCM) which produces a transition at the centre of each bit period. This transition can hence be used to extract the transmitter clock, and results in an extremely simple receiver structure. The immediate drawback is that the transmitter RF oscillator must switch at twice the rate, and the transmitted bandwidth is increased correspondingly.

Manchester decoding involves edge rather than level detection, allowing the received signal to be AC-coupled, thus facilitating use of the analogue comparator, which has no need for an averaging circuit to find a threshold.

4 First Generation Implementation

A digital transceiver has been designed to act as a DSP physical layer, interfacing with both the commercially available microcontroller and the custom-designed OOK RF transceiver described previously. The DSP transceiver communicates with the microcontroller via a Serial Port Interface (SPI).

The transceiver consists of a Manchester encoder, Manchester decoder, SPI interface logic, control logic and various buffers and registers. In total, the un-optimised circuit comprises around 11000 gates. A programme of ASIC design iterations has recently commenced.

The logic circuit diagram of the Manchester decoder implementation is shown in Figure 3 [3]. It consists of a transition detector and a Johnson counter which is used to block between-bit transition detection. Between-bit transitions occur when two consecutive bits have the same value, as shown in Figure 4 where two consecutive '1' bits are transmitted. This creates the possibility that the between-bit transition may be detected at point A. The data must be oversampled and the design of the decoder results in a finite range of oversampling ratios that can be used, which is fixed directly by the counter's maximum count. This range of oversampling ratios results in the decoder being tolerant to clock inaccuracies.

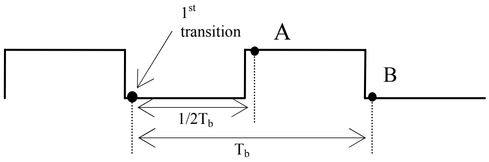

Figure 4: Limitations of the Manchester decoder oversampling ratio

Figure 4 shows how the decoder ignores between-bit transitions. Once a valid transition has been detected, the counter commences, and transitions are ignored until the counter has completed one entire cycle. Referring to the points marked in Figure 4, in order to both miss a between-bit transition (*A*), and to detect correctly the next valid transition (*B*), the count must complete between half a bit period, and a whole bit period, after the detection of a valid transition. This can be described by the following inequality:

$$\frac{1}{2}T_b < NT_s \leq T_b \qquad (1)$$

Where T_b is the bit period, N is the maximum count of the counter, and T_s is the sampling period. Resolving the inequality for the sampling and bit rates, Equation (2) gives the bounds on the sampling frequency.

$$NR_b \leq R_s < 2NR_b \qquad (2)$$

Where R_b is the bit rate, and R_s is the sampling rate. The oversampling ratio must therefore be between N and $2N-1$ inclusive.

The permitted frequency tolerance for the oscillators in the specks can be found by analysing the two worst cases of clock error, i.e. transmitter and receiver clocks at the minimum and maximum possible frequencies. The maximum oscillator variations can be found from the following two inequalities.

$$N[R_b + \delta R_b] \leq OSR[R_b - \delta R_b] \leq (2N-1)[R_b + \delta R_b] \qquad (3)$$

$$N[R_b - \delta R_b] \leq OSR[R_b + \delta R_b] \leq (2N-1)[R_b - \delta R_b] \qquad (4)$$

Where OSR is the chosen oversampling ratio and δ is the maximum acceptable fractional tolerance. Resolving the inequalities gives Equation (5), which expresses the bounds of permissible frequency error:

$$-\frac{OSR-N}{OSR+N} \leq \delta \leq \frac{OSR-N}{OSR+N} \qquad (5)$$

A counter with a maximum count of $N = 6$ gives the best compromise between an easily generated power of 2 oversampling ratio and a wide tolerance. The implemented Manchester decoder operates at an oversampling ratio of 8 and is therefore tolerant to clock inaccuracies including central frequency error and jitter to the extent of +-14.29%. This again enables circuit simplification as the use of an inaccurate, low-cost and low-power clock source is permissible.

5 Pulse Shaping

Pulse shaping [7] is used in wireless communications to define the spectrum occupied by the transmitted RF signal, specifically increasing the proportion of power in the main lobe, and decreasing the power in the sidelobes. Naturally reductions in both the bandwidth of the signal, and the magnitude of out-of-band emissions, will increase the signal to noise ratio (SNR) at the receiver for a given transmit power.

Matched pulse shaping filters are almost invariably used, meaning that a filter is present at the receiver mirroring that at the transmitter. This correlates the received signal with the transmitted pulse shape, and enhances SNR by acting as a low-pass filter. The drawback is the increased receiver cost in terms of area and power consumption.

Including these filters may have the effect of introducing distortion and inter-symbol interference (ISI), the nature of which relies on the response of the filters, and the channel through which the signal passes. However, careful filter design allows these effects to be avoided or controlled.

5.1 Pulse Shaping Filter Design

Of the several different pulse shapes used in RF communications, the basic speck design described in Sections 2 and 3 uses a rectangular filter – analogous to no filtering at all – which in general is seldom used due to the unwanted spectral leakage it creates. Other filters include root raised cosine, half sine, Gaussian, and Butterworth filters. In each of these cases, a non-trivial analogue or digital filter is incorporated into the transmitter, and usually the receiver as well.

Other parameters of the pulse shaping filter, all of which need to be considered for an actual implementation, include:

- Oversampling ratio
- Filter length (in weights, or symbol periods)
- Roll-off rate (root-raised cosine); BT (Gaussian) etc.
- Resolution of filter coefficients

5.2 Pulse Shaping in SpeckNet

In the context of Speckled Computing the performance increase using pulse shaping filters needs to be carefully traded off against any cost in additional circuitry. Given the low power budget, and the short transmission distances and simple channels involved, it is likely that simple, low cost, filters will be most appropriate.

Additionally, it is important to consider any consequences of pulse shaping to other sections of the transceiver. In particular, it has to be considered that the successful operation of the Manchester Decoder requires binary data to be easily and successfully extracted from the received signal as discussed in Section 6. This makes root-raised cosine filtering, for example, an inappropriate choice, given the extra circuitry required to accurately interpret the matched filtered signal.

The advantages of omitting a matched filter at the receiver are the DSP hardware savings, and that the use of a multi-bit analogue-to-digital converter (ADC) can potentially be avoided. This latter device is comparably complex and power-hungry. Instead, an analogue slicing component may be used, as in the initial speck design (see Figure 2).

As a solution to the problems described, we propose a method of half-sine pulse shaping at the transmitter, with no matched filter at the receive side. For this type of filter, the input data is upsampled to the oversampling ratio, and the half-sine filter can be implemented as a lookup table. This has the benefit of reducing its cost and power consumption.

The architecture of the system considered is shown in Figure 5 below. An analogue band-pass filter rejects unwanted frequency components prior to detection using an envelope detector. It is desirable to create the band-pass filter cheaply (i.e. specify a filter of low order) in order to minimise analogue hardware cost.

Figure 5 – Receiver architecture

The effect of introducing a digital, half-sine pulse shaping filter was simulated using an oversampling ratio of 8, and 6 bits resolution. The simulation assumed an additive white Gaussian noise (AWGN) channel. For the analogue bandpass filter, a simple 2nd order Butterworth filter was used.

Simulation results show that introducing this half sine filter increased the proportion of signal energy retained in the bandwidth of interest from 23.46% to 45.00%. This is illustrated in Figure 6: note the larger main lobe, and lower sidebands, of the half-sine shaped waveform compared to the rectangular shaped signal.

Correspondingly, applying the filter results in an increase in SNR at the receiver (i.e. after the analogue band pass filter) of 2.83dB over the unshaped version, as shown in Table 1.

This represents an improvement in the integrity of the transmitted data, or alternatively implies that the power emitted by the communicating node could be reduced by the same ratio while maintaining the original receiver SNR.

The cost of introducing this filter, implemented (without optimisation) as a lookup table, is approximately 300 gates. As a guide, the overall gate count of the rest of the transceiver design is currently around 11000 gates. Therefore it may be concluded that the inclusion of the half sine filter would add very little (less than 3%) to the overall area and power consumed by the design.

Table 1 – Performance comparison of half sine and rectangular pulse shapes (SNR values measured after the analogue band pass receiver filter).

Filter Type	Taps	SNR at receiver (dB)	Improvement (dB)
Rectangular	0	11.3393	-
Half Sine	8	14.1689	2.8296

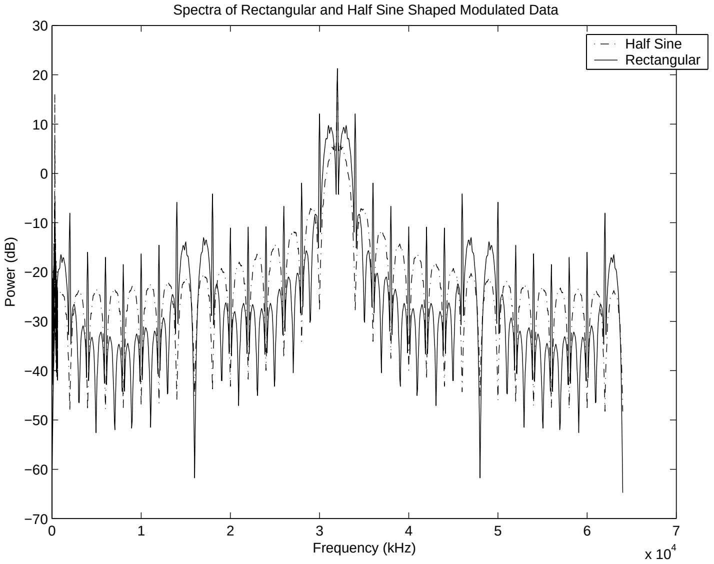

Figure 6 – Spectral comparison of half sine shaped and rectangular (unshaped) modulated signals.

Additionally, a digital to analogue converter (DAC) of at least 6 bits would be required in place of the 1-bit device currently used. This would have an associated power cost, the exact value of which would depend on ASIC implementation, to be offset against the SNR improvement when evaluating the net benefit of introducing the filter.

These results show that incorporating a half-sine pulse shaping filter into the Manchester encoded, on-off keyed system, produces a net benefit in terms of the power/performance trade-off.

In the specific case simulated, a 2.83dB improvement in SNR was observed when such a filter was inserted into the system. Even allowing for the power consumed by the filter and the DAC, this is likely to represent a real and useful development. The SNR improvement could be used either to make the transmission more robust to noise and interference, or to reduce the transmit power while maintaining the same performance. A design could also incorporate a combination of these two effects.

6 Future Work

One potential drawback of introducing half-sine pulse shaping, yet to be fully evaluated, is that it may not be possible to extract the sharp transitions of the Manchester encoded binary waveform as reliably, especially in the presence of high noise levels. Spurious transitions in the signal may be experienced when received data is oversampled by the receiver, especially in relatively high noise conditions. In particular, it is anticipated that samples of '1' symbols at symbol period edges may be detected as '0's. Further investigations will focus on techniques to reduce the impact of this effect. Thus other pulse shapes, alternative implementations of Manchester decoding, and associated implications for the tolerance of the oscillator will be considered.

7 Conclusions

Speckled Computing presents novel challenges for the DSP physical layer. In particular all design must be undertaken with low-power as a first priority, due to the limitations of speck size, and current battery technology.

A simple transceiver has been developed that enables communication across an OOK RF link using Manchester encoding. An added advantage to the design is that it enables clocks with a wide tolerance to be used while still providing resilience to clock jitter. The main disadvantage to the Manchester encoded signal is that it increases the transmitted bandwidth. The current version of this transceiver uses only around 11000 gates.

Pulse shaping can be used to better concentrate the energy of the transmitter signal within the bandwidth of interest. The case of half sine pulse shaping was considered, and shown to increase the energy from 23.46% to 45.00% - representing a 2.83dB increase in SNR - without adding significantly to the overall cost of the transceiver.

Further work will concentrate on low cost implementations of alternative pulse shaping filters, and the associated effects on the performance and oscillator requirements of the Manchester decoding strategy described in this paper.

8 Acknowledgements

This work is supported by the Scottish Higher Education Funding Council (SHEFC), and the Engineering and Physical Sciences Research Council (EPSRC).

The authors would like to thank members of the Speckled Computing Consortium for their input.

9 References

[1] www.specknet.org

[2] EM Microelectronic EM6812 http://www.emmicroelectronic.com/Products.asp?IdProduct=209

[3] "Manchester Decoder in 3 CLBs". In *Xilinx XCell Journal 17, Q2 1995.*

[4] Sanyo ML421 (http://sanyo.wslogic.com/pdf/pdfs/ML421.pdf) and Sanyo ML414 (http://sanyo.wslogic.com/pdf/pdfs/ML414.pdf)

[5] M. Josie Ammer, Jan Rabaey, "Low Power Synchronization for Wireless Sensor Network Modems." In *Proc. 2005 IEEE Wireless Communications and Networking Conf*, pp 670- 675.

[6] F.Darbari, I. McGregor, G. Whyte, R.W. Stewart, I. Thayne, "Channel Estimation for Short Range Wireless Sensor Network," *To be presented at the IEE Conference on DSPenabledRadio, Southampton, UK, 2005.*

[7] K. Gentile - "The Care and Feeding of Digital, Pulse-Shaping Filters", *RF Design magazine, April 2002, pages 50 - 61.*

[8] B .P .Otis, Y. H. Chee, R. Lu, N. M. Pletcher, and J. M. Rabaey, "An Ultra-Low Power MEMS-Based, Two-Channel Transceiver for Wireless Sensor Networks," *Symposium on VLSI Circuits,12-19th June 2004, Digest of Technical Papers, pages 20 -23.*

[9] Donald R. Sadoway, Anne M. Mayes, "Portable Power: Advanced Rechargeable Lithium Batteries", *MRS Bulletin, August 2002.*

[10] K. Daniel Wong, "Physical Layer Considerations for Wireless Sensor Networks," *Proceedings of the 2004 IEEE Conference on Networking, Sensing and Control, pages 1201-1206, Vol. 2.*

[11] D.K Arvind, "Speckled Computing – an Overview", *2nd Workshop in Speckled Computing, Edinburgh, 2004.* http://www.specknet.org/publications/20904workshop/