Speckled Computing

D. K. Arvind

Research Consortium in Speckled Computing
School of Informatics, University of Edinburgh
Mayfield Road, Edinburgh EH9 3JZ, Scotland, U.K.
dka@inf.ed.ac.uk

ABSTRACT

Speckled Computing is an adventurous research programme which offers radical new ideas for information technology in the future. This will be realised by minute, autonomous semiconductor ***specks*** – each speck will encapsulate capabilities for sensing, programmable computation, and wireless networking. Computing with networks of specks (known as ***specknets)*** will enable linkages between the material and digital world with a finer degree of spatial resolution than hitherto possible, which will be a fundamental enabler for truly ubiquitous computing. The Speckled Computing Research Consortium slices across traditional academic disciplines and organisational boundaries to bring together physicists, electronic engineers and computer scientists from five Scottish universities to provide an integrated technological push towards the realisation of specks and specknets. This paper outlines the vision underlying the project, and gives an overview of some of the considerable research challenges.

Keywords: ad-hoc networks, radio, free-space optics.

1 VISION

Specks will ultimately be minute (tens of cubic millimeter) semiconductor grains that can sense, compute and communicate wirelessly. Each speck will be autonomous, with its own captive, renewable energy source. Thousands of specks scattered on surfaces or sprayed on the person, will collaborate as programmable computational networks called ***specknets***. Computing with specknets will enable linkages between the material and digital worlds with a finer degree of spatial resolution than hitherto possible – specknets will in effect be the last millimeter of the World Wide Web.

Speckled Computing is the culmination of a greater trend. As the once-separate worlds of computing and wireless communications collide, a new class of information appliances will emerge. Where once they stood proud – the PDA bulging in the pocket, or the mobile phone nestling in one's palm, the post-modern equivalent might not be explicit after all. Rather, data sensing and information processing capabilities will fragment and disappear into everyday objects and the living environment. At present information processing capability is monolithic and discrete – the computer on a desk, the PDA, laptop, mobile phone, and smart cards. In our vision of Speckled Computing, the sensing and processing of information will be highly diffused – the person, the artefacts and the surrounding space, become, at the same time, computational resources and interfaces to those resources. Surfaces, walls, floors, ceilings, articles, and clothes, when treated with specks (or "speckled"), will be invested with a "computational aura" and sensitised *post hoc* as props for rich interactions with the computational resources.

In the evolution of computer, specks represent a new generation of nano devices – the minute, sub-dollar computer which combines the capabilities to sense, process and network wirelessly.

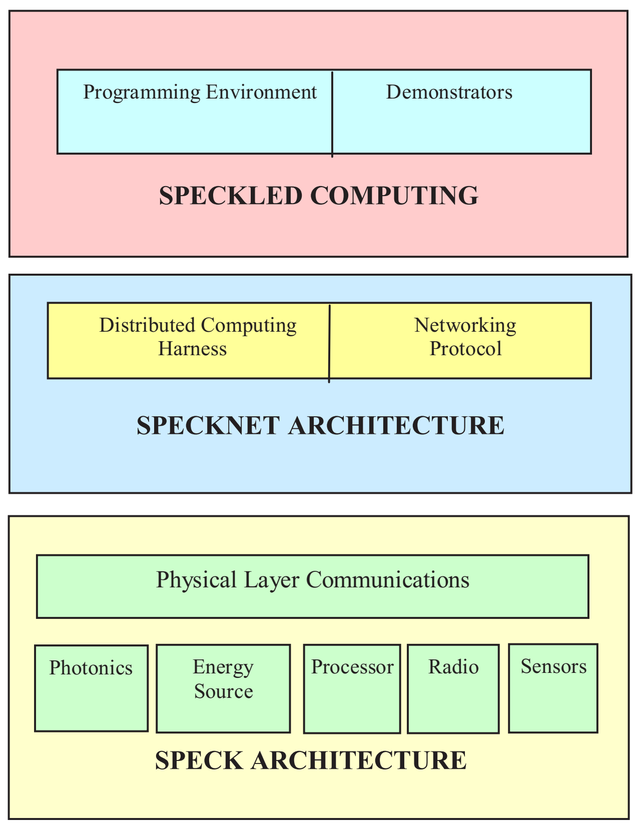

Figure 1: Speckled Computing Overview

In evolutionary terms, comparisons with the microprocessor is apt. The microprocessor integrated processing and storage on a single chip, which replaced the black box of electronics with a programmable device, and which in turn eventually enabled the PC, the PDA and the Smart Card. Similarly, the speck encapsulates sensing, processing/storage, wireless networking capabilities, together with an energy source on a single nano device, which is envisaged to be the enabler for pervasive and ubiquitous computing. And given their sizes, specks will reach computation to parts that could not be reached before, where data is required to be sensed, processed, and information extracted *in situ*.

2 SPECKS

Figure 1 gives an overview of the various components which comprise specks and specknets. We are currently investigating both free-space optics as well as radio for inter-speck communication. The key motivation for using the former is their directionality; whereas a radio transmitter radiates mainly isotropically, a laser light can be collimated and made highly directional, thus reducing the power requirements on the transmitter side. Optical sources being considered are Vertical-Cavity Surface Emitting Lasers (VCSELs) that combine low power consumption (as low as 100µA, dark current – 10nA) with high quantum efficiency and small size (the transceiver will occupy less than 100µm^2). An array of VCSELs offers the potential for electronically-controlled laser beam steering through liquid crystals, thus avoiding the expensive and voltage-hungry MEMS-based mirrors deployed in Smart Dust [3].

The design choices for the radio follow two strands: carrier-based techniques versus impulse-radio ones, and low GHz frequency regime versus millimeter-wave frequencies for both the transmission methods. A carrier-based (superheterodyne) solution results in a relatively complex radio front-end, but simplifies the processing requirements for synchronisation which eases the demodulation and data recovery issues. In contrast, a non-carrier based solution utilising pulse-position modulation simplifies the radio front-end at the expense of complex processing in the physical layer of the communication stack.

The choice of operating frequency has ramifications for both the size and energy requirements of the speck. Millimetre-wave operating frequencies will result in optimal antenna sizes in the scale proposed for the specks (a few millimeters). However, the requirement to design using distributed transmission line elements results in physically larger transceiver chip geometries. In addition, at millimeter-wave frequencies the operating DC power will be greater as the gain and noise performance of the transistor technology is inferior to that at lower frequencies. In contrast, working at a lower frequency (3-10 GHz) will allow compact lumped element designs, and the outstanding gain and noise performance of the active device technology at the lower frequency will enable transmission at reduced powers, thanks to the superior noise figures at the receiver. However, for a speck with a form factor of a few millimetre square, the antenna sizes will be suboptimal, thereby mitigating some of the gain and noise advantages mentioned earlier.

The choice of compound semiconductor for the laser and the radio front-end will result in reduced DC power, higher gain and lower noise at a given frequency in comparison to silicon-based solutions. The added attraction is the ability to integrate the source of renewable energy with the communication subsystem. Photovoltaics based on III-V heterojunctions and triple junctions based on III-V on germanium technology have demonstrated efficiencies of around 30%. Given typical outdoor operating conditions, a state-of-the-art solar cell can therefore be expected to generate around 75µWH/mm^2 (assuming a conversion efficiency of 30%, and a power-to-storage conversion efficiency of 80%) – still the highest performing power scavenging technology available. The proposed photovoltaic solution is also monolithically integrable with silicon processing, as it is based on a silicon substrate. This is the first step towards the hyperintegration of compound semiconductor radio and photonic components with standard CMOS, all on a silicon platform which would clearly ease significantly the packaging challenges.

The speck processor is being designed in CMOS technology. An optimised architecture in 0.13 µm CMOS can deliver 3 MOPs/µW and 20,000 MOPs/mm^2 (multipliers require 0.02 mm^2 and 3pJ per operation at 1V; adders and registers are an order of magnitude lower in both size and energy consumption). Although the Silicon Roadmap will continue to follow Moore's Law, primarily due to shrinkages in the minimum device feature size and increases in the clock speed, this will not significantly improve the energy efficiency (MOPs/µW). The modest reduction in the rail voltage (to 0.9V in 2006) will only result in a 10% increase in processing capability for a given power budget. For these reasons we have adopted a more radical approach towards achieving energy efficiency than the current sensor network projects. The architecture of the programmable specknet will be based on the micronet model [1], which is a network of entities that compute concurrently and communicate asynchronously. Each speck, in turn, is recursively a micronet, consisting of a network of analogue units such as the sensor and wireless interfaces, and digital units such as programmable functional units. This leads to a uniform architectural model at different levels of abstraction - at the level of a specknet and within the individual specks. In its purest form, the digital functional units within a speck will be implemented in an asynchronous style using self-timed circuits, although the model can just as well accommodate a mixture of locally-clocked modules and self-timed ones. An important side-effect of asynchrony is that the architecture is event-driven and will only consume power in those parts which are active, with the rest of the circuit only drawing quiescent current, i.e., if a speck is used only as a communication relay, then only the transmitter and receiver sections are active and the rest of the system will sleep. The asynchronous implementation also exhibits good electro-magnetic compatibility, a very desirable feature where the analogue wireless circuitry sits cheek by jowl with the digital processing platform, as it does in a speck. An asynchronous architecture will exploit both the data-dependent latencies in the operations of the functional units, and the concurrency at different levels to deliver an energy-efficient platform for the signal processing, protocol stacks and application modules.

3 SPECKNETS

The networking of sensor nets [2] has been an area of active research in recent years. But unlike traditional sensor nets, specknets will be both non-static and adaptive. Rather than being dumb conduits to sensor data, specknet programs will combine both data- and control/program-centric operations, with very limited (maximum of 2Kbytes) memory in each speck. This is networking at its extreme, which will require highly adaptive, scalable and power efficient protocols for dense and non-static networks composed of nodes with very limited energy, memory and computational speeds. New ideas have to be developed and evaluated for a bespoke Medium Access Control (MAC) layer which will address the following issues: the latency of packets is often sacrificed in favour of maximal energy efficiency (using duty cycling for instance) which may be acceptable for a data-centric sensor network, but would be less so for interactive, program- and control-centric applications intended on specknets, and research will investigate means of reducing latency without overly sacrificing energy efficiency; the proposed MAC protocol should be adaptive and function gracefully as the density of specks changes, and as network sizes (some specks may join while others leave or die) and loads vary over time. Novel light-weight, energy-conscious routing algorithms (discovery, recovery and maintenance) will be devised which take into account the modest storage and energy resources in specks. One approach to be considered is the partitioning of a specknet into Speckzones of a given size. Specks within a zone are classified as either Inner (all the immediate neighbours are within the Speckzone), or Border (otherwise). Hybrid protocols are being investigated which combine proactive routing by means of a local routing table for interspeck routing, and reactive routing by means of a foreign routing table for intraspeck routing. In the case of the latter, route discovery would be between Border specks in different Speckzones, and the foreign routing tables store the discovered routes. Should these tables fill up beyond a threshold then they are partially dumped to neighbouring Inner specks, thus creating a distributed memory model within the Speckzone, which would be necessary given the modest storage capability intended for individual specks.

Each speck will be modest in terms of processing and storage resources, but can be powerful collectively. Decentralised algorithms will be investigated for clustering specks into Speckzones of given sizes, which will be maintained as new specks join in or old ones die, and the sizes of the zones can be changed dynamically on demand depending on the computation and memory requirements. The model of distributed computation to be developed will take into account some specific features of specknets, such as unreliability of communication, a higher than normal failure rate of specks due to harsh operating environment and thanks to the very large volume manufacturing. Another feature that will be important in specknets is to relate the information that is extracted to its location within the specknet. A method for decentralized logical location discovery and maintenance has been implemented which is an improvement on existing methods based on a distributed relaxation algorithm in which each speck satisfies multiple local binary constraints based on the radio ranges from neighbouring specks [4]. In summary, we are developing a peer-to-peer style of organisation for specknets in which the resources of the specks are aggregated to provide critical system services in a highly decentralised manner.

A key distinguishing feature of specks (compared to sensor networks) is that they are intended to be overwhelmingly programmable. We intend adapting existing network-centric programming models for use in specknets, to take account of the modest storage resources. The unit of programmability will be Speckzones rather than Individual specks. The applications in the form of code modules together with dependency information will be downloaded onto a Speckzone wirelessly. A speck will accepts a module based on a probabilistic function which depends on the frequency of the computation of that particular module in the application. A scheduling algorithm based on ACO will be used to link the results of the instances of dependent modules.

4 SPECK PROTOTYPE - PROSPECKZ

The Consortium is an interdisciplinary one with key developments enabled by working not only on individual technology developments, but more importantly in defining and solving challenges spanning the traditional component layer/physical layer/architecture layer boundaries. The importance of cross-boundary research for Speckled Computing has been addressed right from the start. The Prospeckz (Programmable Specks over Zigbee Radio), as illustrated in Figure 2, has been developed using off-the-shelf components to provide academia and industry with a platform to develop MAC layer, networking protocols, programming environment and applications for this emerging technology. The underlying philosophy is to provide researchers with a means to determine the constraints and parameters in the real world which will inform the design of the semiconductor specks, and validate the simulation models of specks and specknets.

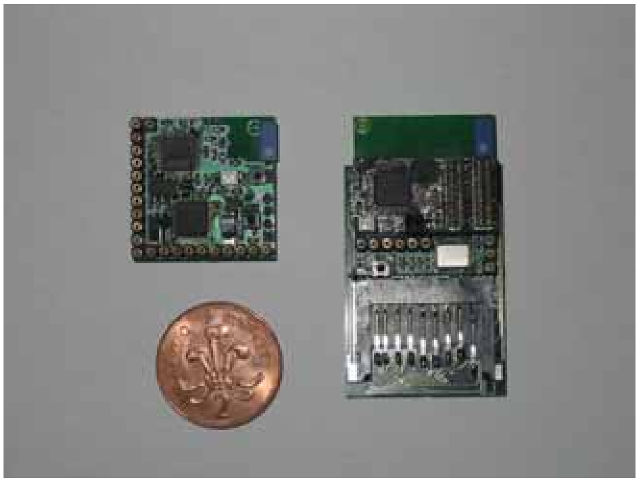

Figure 2: Speck prototype – Prospeckz II and III

Prospeckz contains a 802.11.4 compliant radio chipset (Chipcon's CC2420), 2.4GHz matched antenna and filter circuitry, and a PSoC (Cypress Semiconductor's CY8C2764) which provides soft reconfigurable analogue circuitries for easy interface to sensors, and an 8-bit microcontroller. A power-aware MAC layer has been developed to provide efficient use of the radio channel via duty-cycling and contention-free access whenever possible. A novel lightweight specknet network protocol handles the routing between specks in a unicast or multicast approach, and capabilities to transmit data over the air easily. An ultralightweight OS provides device drivers for the sensors as well as scheduling of tasks, events and commands. Prospeckz has also been incorporated by collaborators, albeit with a larger form factor, to demonstrate novel information appliances such as Smart Jewellery, Tangible Interfaces for radios, and a Stress Management System, and Smart Spaces in museums. To further develop this interdisciplinary collaboration, Prospeckz III (see Figure 2) has been designed and implemented as a stackable and interchangeable development platform, with separate boards for the wireless communication – radio or laser, CPU, Memory, Sensor, FPGA and Power. The initial realization using off-the-shelf components will be replaced progressively with bespoke ones intended for the specks as and when they become available. The FPGA board for example, will be used to experiment with the physical layer communication and the speck processor. Such an iterative approach will be used to manage the risk involved in developing optimised systems. Starting from this basis, the research programme can be envisaged as being evolutionary, with the speck volume being reduced in steps, first to 25mm^3 and smaller, as optimised components, DSP algorithms, networking and architectural concepts in the various components mature and become available for incorporation in the overall system. It is our belief that this approach will serve a number of benefits: (i) It will identify quickly both at the individual technology level, but perhaps more importantly, at the various technology interfaces, where additional effort and resource should be deployed; (ii) It will ensure that all researchers working in the consortium at all levels are fully aware of the interdisciplinarity of the project, and that collaborative links, which take time to develop and mature, are established at the beginning of the program and nurtured throughout.

5 ACKNOWLEDGEMENTS

The Research Consortium in Speckled Computing (http://www.specknet.org) has been partially funded by a Strategic Research Development Grant from the Scottish Higher Educational Funding Council since October 2003. The author wishes to thank Professors Thomas Krauss and Alan Miller of the University of St. Andrews, Professor Robert Stewart of the University of Strathclyde, Professor Iain Thayne of the University of Glasgow, and the Speckled Computing group at the University of Edinburgh for many hours of discussions.

REFERENCES

[1] D. K. Arvind and Salvador Sotelo-Salazar, "Scheduling for ILP in the "Processor-as-a-Network", in Communicating Process Architectures, Eds. Jan Broenink and Gerald Hiderink, pp 289-304, IOS Press, 2003.

[2] Anna Hac, "Wireless Sensor Network Design", John Wiley and Sons, 2003.

[3] J.M. Kahn, R.H. Katz, K.S.J. Pister, "Next Century Challenges: mobile networking for Smart Dust", Int. Conf. on Mobile Networking and Networks, Seattle, USA, 1999.

[4] R. McNally, S. Wong and D. K. Arvind, "A Distributed Algorithm for Logical Location Estimation in Specknets", Proc. IEEE Wireless Communication and Networking Conference, New Orleans, March 2005.