CISTER
Research Center in
Real-Time & Embedded
Computing Systems

Technical Report

XDense: A Sensor Network for Extreme Dense Sensing

João Loureiro

Vikram Gupta

Nuno Pereira

Eduardo Tovar

Raghuraman Rangarajan

CISTER-TR-131203

Version:

Date: 12/6/2013

XDense: A Sensor Network for Extreme Dense Sensing

João Loureiro, Vikram Gupta, Nuno Pereira, Eduardo Tovar, Raghuraman Rangarajan

CISTER Research Unit

Polytechnic Institute of Porto (ISEP-IPP)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8340509

E-mail: joflo@isep.ipp.pt, vigup@isep.ipp.pt, nap@isep.ipp.pt, emt@isep.ipp.pt,

http://www.cister.isep.ipp.pt

Abstract

We introduce XDense, a grid topology based wired sensor network architecture tailored for high sampling rate applications and high density of sensor node deployments. The purpose of such an architecture is to monitor, detect and report features of observed dynamic phenomena, in a distributed fashion. We present preliminary simulations validating our concept.

XDense: A Sensor Network for Extreme Dense Sensing[†]

João Loureiro[‡], Vikram Gupta, Nuno Pereira, Eduardo Tovar, Raghuraman Rangarajan
CISTER/INESC-TEC, ISEP, Polytechnic Institute of Porto, Porto, Portugal
Email: {joflo, vigup, nap, emt, raghu}@isep.ipp.pt

***Abstract*—We introduce XDense, a grid topology based wired sensor network architecture tailored for high sampling rate applications and high density of sensor node deployments. The purpose of such an architecture is to monitor, detect and report features of observed dynamic phenomena, in a distributed fashion. We present preliminary simulations validating our concept.**

I. Introduction

Recent advances with respect to performance, cost and size of Micro-Electro Mechanical Systems (MEMS) have revolutionized the design of sensors, actuators, and other microsystems. This has made possible new solutions for many applications ranging from fundamental scientific research to industrial process control. For example, the measurement of shear-stress to determine flow conditions over time is of high importance in highly dynamic application scenarios like monitoring the air-flow over an aircraft wing [1].

Such applications require sensors to be smaller than the spatial granularity of the observed phenomena (of 100 μm or less) and to have high sampling rates (in excess of 10 kHz) [2]. For such scenarios, there are some limitations in the current technologies. Battery-powered wireless sensors can suffer from size limitations and transmission latency due to concurrency issues. Wired sensor networks with shared buses limit the scalability of the system and are susceptible to noise.

Our work moves away from traditional wireless and wired sensor network approaches and resembles Network-on-Chip (NoC) architectures closely, regarding some aspects like the network architecture topology, routing schemes and timing analysis [3]. On the other hand, we believe that the key differentiating features are: (a) The network is not on a single chip, but on a larger surface, which should physically interfere as less as possible with the observed phenomena, and (b) the node count is likely to be much greater than the one usually found in NoC applications.

We propose a grid network architecture for very dense sensor deployments, able to distributively detect events without the need of centralized data acquisition and processing. Moreover, it reduces number of transmissions between the nodes and the sink in order to update the sink about the overall scenario. We now describe the architecture of the proposed system and provide preliminary simulation analysis to support our proposition.

[†]This work was partially supported by National Funds through FCT (Portuguese Foundation for Science and Technology) and by ERDF (European Regional Development Fund) through COMPETE (Operational Programme Thematic Factors of Competitiveness), within projects ref. FCOMP-01-0124-FEDER-037281 (CISTER), ref. FCOMP-01-0124-FEDER-020312 (SMARTSKIN), ref. FCOMP-01-0124-FEDER-012988 (SENODS) and ref FCOMP-01-0124-FEDER-028990 (PATTERN).

[‡]João Loureiro is supported by the government of Brazil through CNPq, Conselho Nacional de Desenvolvimento Científico e Tecnológico.

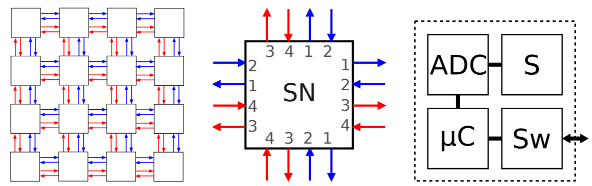

(a) (b) (c)

Fig. 1. (a) A 4×4 2D mesh sensor network, (b) is the sensor node pinout, where: 1: Tx, 2: Rx, 3: Handshake Output, and 4: Handshake Input pin. (c) shows the node architecture. The microcontroller (*μC*) is connected to the switch (*Sw*) and sensor (*S*) through the analog-to-digital converter (*ADC*).

II. System Architecture, Network Topology and Principles of Operation

The aim of our system is to detect the effects of any environmental phenomenon over a given surface. This is done by detecting spatial gradients in the distribution of various physical quantities over the surface. For example, changes in shear stress, pressure or temperature due to air flow over an aircraft wing. The architecture consists of a 2D mesh network of sensor nodes and sinks, point-to-point connected in a plane, with up to four neighboring nodes physically located in four directions. In this paper, a rectangular grid is studied for the sake of simplicity.

In addition to a sensor, each node has an ADC, a microcontroller (μC) and a switch with four full duplex serial ports, plus eight general propose input/output (GPIO) pins for handshaking. The network architecture, a single node pinout and it's hardware architecture are shown in Figure 1.

In order to detect the phenomena of interest, each node first shares its values with its neighbors, to compare its sensed value with the ones collected. The neighborhood of a node is a system defined variable which is equal to the distance from the node in number of hops, n. Intuitively, more number of hops leads to a larger neighborhood, leading to better detection whereas it also increases overhead and latency.

The sensor system has the following states of operation: *Network discovery*, *event monitoring and detection*, and *event-announcement*. In the *network discovery* state, each node discovers its neighbors and the shortest path to sink(s). In the *event monitoring and detection* state, the sensor nodes continuously sense the environment. They communicate the sensed values with their neighborhood (nodes within n hops). The value of n is selected according to the expected characteristics of the phenomena to observe. Each node by forwards data on all the four ports, propagating its own data and the data it received from its neighbors. The node then switches to an *event-announcement* state when it detects a gradient greater than a threshold based on the analysis of the collected data.

In the *event-announcement state*, a node sends the data

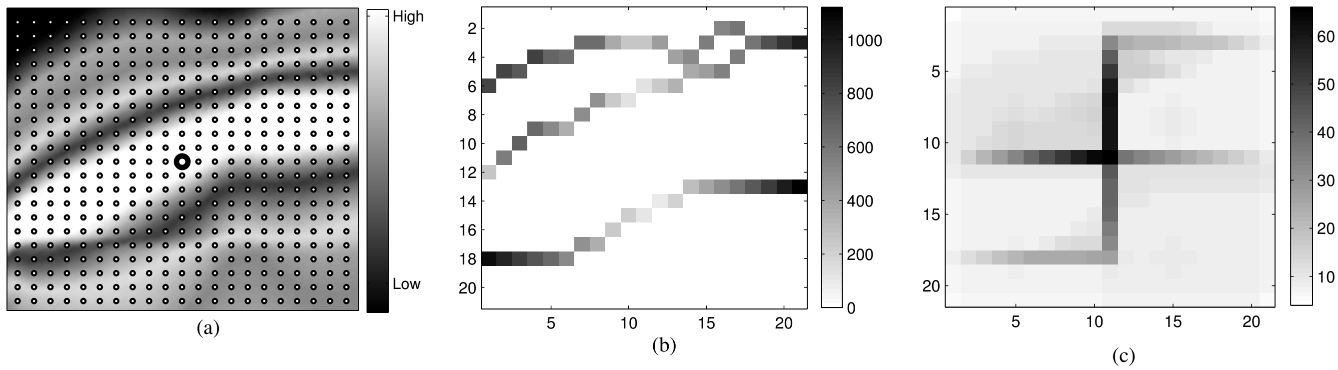

Fig. 2. Network architecture, and simulation results: (a) Grid sensor network superimposed on a pressure distribution snapshot due to airow over a surface (The sink is in the center. The wired connections are not shown for the sake of clarity); (b) Node-to-sink transmission latency after the detecting the event. It indicates the time (in μsec) required by each node in establishing a connection and sending the data to the sink. Each integral X-Y coordinate corresponds to a sensor node in the grid; (c) Total number of packets forwarded by each sensor node. It shows the route path and the bottlenecks towards the sink.

by establishing a connection path to the sink. This is done using the handshaking pin for the control signal. The adjacent node, on the path to the closest sink, gets the request in its handshaking input, and if it is not busy, it propagates the signal in the direction of the sink. Intermediate nodes do the same until the sink is reached. The sink responds with a handshake in the reverse direction granting the connection. Through this signalling process, a simple through-line circuit is established between the node and the sink. The node now sends the data to the sink over the established circuit.

In this setup, the most important metric to be analysed is the latency of the event announcement to the sink by the sensor node. This is because, latency in the network determines the speed of detection, that is, the reaction of the sensor network to the changes in the physical phenomenon. In order to study the latency metric, we simulate an example setup described in the following setup.

III. Simulation and Preliminary Results

To simulate our sensor node, we based the behaviour of our model on a low cost and low power microcontroller. [1] To the best of our knowledge, this is the best micro-controller with the necessary resources for our network design. The example scenario is a grid sensor network monitoring the pressure distribution due to airflow over a surface. We take a sample static snapshot, of pressure over a surface, to analyse the detection and transmission latency of the sensor network. The network setup is as follows: a $(21{\times}21)$ grid network of 440 sensor nodes and one sink. The sink is in the center of the grid. In this example, the feature of interest is any gradient with slope greater than a given threshold. Figure 2(a) shows the grid network superimposed on the static air flow snapshot (the sink is the node in the center).

At the time instant $t = 0$, all the nodes start the *network discovery* state. Once the network is discovered, nodes enter the *event-monitoring and detection* state. Now, sensor data is captured, sent to and received from the neighbourhood by all the sensor nodes in the network in a concurrent manner. After processing the collected data, the nodes that are positioned at high-gradient slopes trigger the *event-announcement state*. The nodes then set up a connection to the sink to send this information. The end-to-end transmission latency from all the sensor nodes, is shown in Figure 2(b) (the colormap represents the transmission latency). Farther a node is from the sink, the greater is the latency due to the number of intermediate hops. Also, the number of simultaneous detections, and hence transmissions, also determine the contention in the network.

Figure 2(c) shows the number of packets forwarded by each node in the grid. A darker node signifies congestion in the path through it. As expected, nodes close to sink forward a larger number of packets than nodes away from it. The transmission paths are also shown in the figure.

IV. Conclusions and Future Work

Low latency is an important goal in applications where the number and frequency of readings are crucial. The XDense architecture, with simple hardware and software functionalities, allows for such deployments. The proposed architecture shows some promising results for scenarios of dense deployments of sensors. We presented three concepts: Simple grid network architecture, sharing neighbourhood information allowing quick detection of events and low-latency setup of network paths to the sinks. Further work is required to examine the significance and efficacy of this approach. We intend to explore many of its aspects like routing, flow control and distributed data processing and aggregation.

References

[1] Nobuhide Kasagi, Yuji Suzuki, and Koji Fukagata. Microelectromechanical systems-based feedback control of turbulence for skin friction reduction. *Annual review of fluid mechanics*, 41:231–251, 2009.

[2] Ulrich Buder, Ralf Petz, Moritz Kittel, Wolfgang Nitsche, and Ernst Obermeier. Aeromems polyimide based wall double hot-wire sensors for flow separation detection. *Sensors and Actuators A: Physical*, 142(1):130–137, 2008.

[3] Shashi Kumar, Axel Jantsch, J-P Soininen, Martti Forsell, Mikael Millberg, Johny Oberg, Kari Tiensyrja, and Ahmed Hemani. A network on chip architecture and design methodology. In *VLSI, 2002. IEEE Computer Society Annual Symposium on*, pages 105–112. IEEE, 2002.

[4] Microchip PIC24FJ128GA310. www.microchip.com/downloads/en /devicedoc/39996f.pdf.

[1] We chose the Microchip PIC24FJ128GA310 [4] which is an extremely low powered 16-bit μC that can run at up to 16 MIPS. It has the following required peripherals: four high speed USART's, one DMA channel and one high speed ADC channel.