Max-Sum Decentralised Coordination for Sensor Systems (Demo Paper)

W. T. L. Teacy, A. Farinelli, N. J. Grabham, P. Padhy, A. Rogers and N. R. Jennings
School of Electronics and Computer Science, University of Southampton
Southampton, SO17 1BJ, UK.
{wtlt,af2,njg,pp04r,acr,nrj}@ecs.soton.ac.uk

ABSTRACT

A key challenge for the successful deployment of systems consisting of multiple autonomous networked sensors is the development of decentralised mechanisms to coordinate the activities of these physically distributed devices in order to achieve good system-wide performance. Such mechanisms must act in the presence of local constraints (such as limited power, communication and computational resources) and dynamic environments (where the topology, constraints and utility of the sensor network may change at any time). We propose the use of message passing techniques based on the max-sum algorithm to address this challenge, and in this paper, we demonstrate its use in two different settings. We first present a software simulation where our max-sum decentralised coordination algorithm is used to coordinate sectored radar sensors tracking multiple moving targets (see the ARGUS II DARP project – `http://www.ecs.soton.ac.uk/research/projects/ARGUS`). We then present a hardware implementation of the same algorithm that performs decentralised graph colouring – an intermediate step towards deploying the algorithm to coordinate the sleep/sense cycles of a network of low-power embedded sensors (see the DIF DTC 'Adaptive Energy-Aware Sensor Network' project – `http://www.ecs.soton.ac.uk/research/projects/AEASN`).

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Multiagent systems*

General Terms

Algorithms, Design, Measurement, Experimentation

Keywords

sensor network, decentralised coordination, max-sum algorithm

1. INTRODUCTION

Sensor networks consisting of multiple autonomous sensors that seek to acquire and integrate information have received increasing attention in the computer science literature. Example applications include wide-area surveillance for security purposes and monitoring of environmental phenomena. In such domains, a fundamental challenge is to coordinate the activities of physically distributed devices in order to achieve good system-wide performance, given the specific constraints of each device (such as limited power, communication and computational resources). Additional challenges arise through the need to perform such coordination in a decentralised manner such that there is no central point of failure, no communication bottleneck, that the computation required for the coordination is shared over the distributed resources available, and that the solution scales well as the number of devices within the network increases. Furthermore, it is desirable that this coordination can be performed in dynamic environments where the topology, constraints and utility of the sensor network may change at any time.

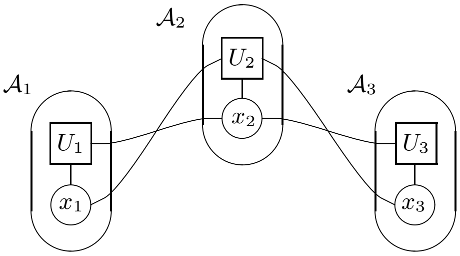

Figure 1: Diagram showing the factor graph representation of three interacting agents, $\mathcal{A}_1$, $\mathcal{A}_2$ and $\mathcal{A}_3$.

To address these challenges, we propose the use of message passing techniques based on the max-sum algorithm [1]. In doing so, we exploit the extensive evidence showing that this algorithm generates good approximate solutions when applied to loopy graphs in practical applications (e.g. iterative decoding of practical error correcting codes [3]), due to their ability to propagate information around the network such that they converge to a *neighborhood maximum*, rather than a simple local maximum [4].

2. THE MAX-SUM ALGORITHM

As described above, our goal is to solve general coordination problems, in which there are M agents whose state may be described by a discrete variable x_m. Each agent interacts locally with other agents such that the utility of an individual agent, $U_m(\mathbf{x}_m)$, is dependent on its own state and the states of these other agents (defined by the set $\mathbf{x}_m$). We assume that each agent only has knowledge of, and can directly communicate with, the few neighbouring agents on whose state its own utility depends. We make no assumptions regarding the structure of these individual utility functions. Solving the resulting coordination problem involves finding the optimal state of each agent, $\mathbf{x}^*$, such that the social welfare of the whole system (i.e. the sum of the individual agents' utilities) is maximised:

$$\mathbf{x}^* = \arg\max_{\mathbf{x}} \sum_{m=1}^{M} U_m(\mathbf{x}_m) \qquad (1)$$

Cite as: Max-Sum Decentralised Coordination for Sensor Systems (Demo Paper), W. T. L. Teacy, A. Farinelli, N. J. Grabham, P. Padhy, A. Rogers and N. R. Jennings, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons(eds.),May,12-16.,2008,Estoril,Portugal,pp. 1697-1698.

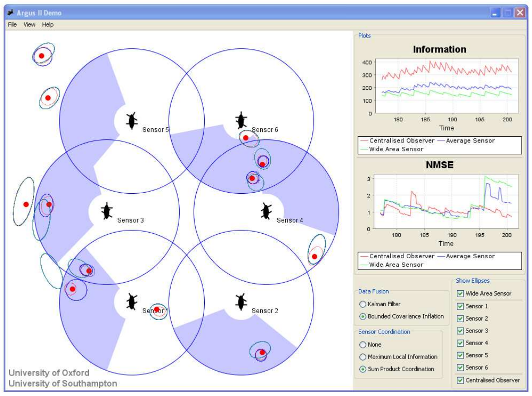

Data Fusion

○ Kalman Filter
● Bounded Covariance Inflation

Sensor Coordination

○ None
○ Maximum Local Information
● Sum Product Coordination

Show Ellipses

☑ Wide Area Sensor
☑ Sensor 1
☑ Sensor 2
☑ Sensor 3
☑ Sensor 4
☑ Sensor 5
☑ Sensor 6
☑ Centralised Observer

Figure 2: Sectored radar sensors using our max-sum decentralised coordination algorithm to track multiple targets.

Within our message passing framework, we assume that each agent is composed of a state variable and a utility function. This representation results in a loopy bi-partite factor graph composed of variable and function nodes (see Figure 1 for an example involving three agents). The max-sum algorithm provides rules for calculating the local messages that pass between the nodes of this factor graph, such that at convergence, an approximate solution to the global optimisation problem has been reached through local computation and communication [1].

3. DEMONSTRATION

We demonstrate the use of our max-sum decentralised coordination algorithm in two settings: a simulation of sectored sensors tracking multiple moving targets, and a real hardware implementation of the algorithm performing decentralised graph colouring.

3.1 Mulitple Target Tracking

The ARGUS II DARP project seeks to develop decentralised data fusion systems that can coordinate sensing actions and fuse information in the absense of centralised control. Specifically, we aim to coordinate the actions of a network of radar sensors to track multiple targets concurrently. Each sensor can make noisy range and bearing observations of multiple moving targets, and the accuracy with which the network can estimate the location of a target improves as the number of sensors tracking it increases. However, the amount of noise (and therefore accuracy) associated with an observation depends on the distance between the sensor and the target, and each sensor can only observe targets in a limited geographical area at one time. Thus, each sensor must decide which area to observe over time, so that the accuracy of each estimated target location is maximised over the network as a whole. This is achieved using the max-sum algorithm as described in Section 2 by passing messages between neighbouring sensors. Figure 2 shows a Java simulation of this system in operation, and a Java applet version is available at `http://users.ecs.soton.ac.uk/acr/argus/demo2/`.

3.2 Decentralised Graph Colouring

The DIF DTC 'Adaptive Energy-Aware Sensor Network' project seeks to develop low-power energy-harvesting wireless sensors for wide-area surveillance applications. To ensure continuous network lifetime, sensors must maintain *energy neutral operation*. That is, they must use no more energy in sensing than they can harvest from their environment. This typically results in low duty cycle operation, and to maximise the effectiveness of the entire system, individual sensors must coordinate their sleep/sense cycles with neighbouring sensors whose sensing field overlaps with their own.

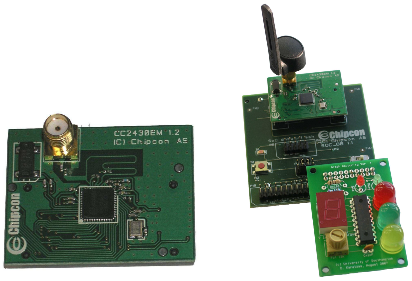

Figure 3: Hardware implementation of the max-sum decentralised coordination algorithm to perform graph colouring.

This coordination problem can be represented as an instance of a distributed graph colouring problem; a canonical benchmark problem for decentralised coordination algorithms in sensor networks [2]. More formally, agents are located at the nodes of a graph, and select their colour (i.e. the state) from a set of possible colours (i.e. $x_m \in 1, \ldots, c$) in order that they avoid conflicts (i.e. choosing the same colour) with other neighbouring agents. The utility of each agent is expressed as:

$$U_m(\mathbf{x_m}) = \sum_{i \in \mathcal{N}(m) \setminus m} x_m \otimes x_i \tag{2}$$

where:

$$x_i \otimes x_j = \begin{cases} -1 & \text{if } x_i = x_j \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

and $\mathcal{N}(m) \setminus m$ is the set of agents neighbouring m. As before, the goal is to find the state of each agent such that social welfare is maximised (i.e. minimising the number of conflicts). To prove the practical applicability of the max-sum algorithm, we have implemented it in hardware to solve the graph colouring problem with multi-coloured LEDs. Each node uses the Chipcon CC2431 System-on-Chip: a low-power device incorporating an IEEE 802.15.4 compliant RF transceiver, 8 kByte RAM, and a 32 MHz 8 bit 8051 micro-controller in a 7x7mm package, which is intended to form the core of low-power sensor nodes (see Figure 3). Videos of the graph colouring nodes in operation are available at `www.youtube.com/v/T6H1AwQ2gXw` and `www.youtube.com/v/D6vWvs3Lsj0`.

Acknowledgments

The work reported here was undertaken as part of the ARGUS II DARP (Defence and Aerospace Research Partnership), and the Data Information Fusion Defence Technology Centre (DIF DTC) Phase II 'Adaptive Energy-Aware Sensor Networks' project. The ARGUS II DARP is a collaborative project involving BAE SYSTEMS, QinetiQ, Rolls-Royce, the University of Oxford and the University of Southampton, and is funded by the industrial partners together with the EPSRC, MoD and DIUS. The DIF DTC project is joint funded by MoD and General Dynamics UK.

4. REFERENCES

[1] A. Farinelli, A. Rogers, A. Petcu, and N. R. Jennings. Decentralised coordination of low-power embedded devices using the max-sum algorithm. In *Proceedings of the Seventh Internation Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Estoril, Portugal, 2008. In press.

[2] S. Fitzpatrick and L. Meetrens. *Distributed Sensor Networks A multiagent perspective*, chapter Distributed Coordination through Anarchic Optimization, pages 257–293. Kluwer Academic, 2003.

[3] D. MacKay. Good error-correcting codes based on very sparse matrices. *IEEE Transactions on Information Theory*, 45(2):399–431, 1999.

[4] Y. Weiss and W. T. Freeman. On the optimality of solutions of the max-product belief propagation algorithm in arbitrary graphs. *IEEE Transactions on Information Theory*, 47(2):723–735, 2001.