

Robustness of Market-based Task Allocation in a Distributed Satellite System

Johannes van der Horst, Jason Noble, and Adrian Tatnall

University of Southampton, Southampton, SO17 1BJ, United Kingdom
jgvdh08r@ecs.soton.ac.uk

Abstract. The perceived robustness of multi-agent systems is claimed to be one of the great benefits of distributed control, but centralised control dominates in space applications. We propose the use of market-based control to allocate tasks in a distributed satellite system. The use of an artificial currency allows us to take the capabilities, energy levels and location of individual satellites, as well as significant communication costs into account. Simulation is used to compare this approach to centralised allocation. We find the market-based system is more efficient and more robust to satellite failure, due to the adaptive allocation of tasks.

1 Introduction

The robustness and scalability of multi-agent systems are articles of faith in the artificial life community. Researchers observe that, in social insects, in human organisations and in markets, the behaviour of the whole is not critically dependent on the functioning of any single component. In recent years these types of social and biological networks have inspired a host of technological systems in which robust, cooperative behaviour is achieved across a group of autonomous agents. However, the view that distributed systems should be controlled from the centre still dominates in other engineering fields. A notable example is the engineering of spacecraft where the central control paradigm has been extremely successful. This can be traced back to a history of monolithic spacecraft, incremental development philosophies and high mission costs. In these systems, reliability is usually achieved through redundancy, fault detection and error correction. The centralised and distributed control approaches are rarely compared head to head, because their respective applications are often completely different.

The idea of using multiple coordinated spacecraft to perform the function of a single larger vehicle has recently been proposed[1]. Distributing functionality between the component spacecraft allows larger structures to be constructed in orbit, while also benefiting from the commoditisation of spacecraft, rapid deployment and mission flexibility. The increased complexity requires that the system handles component failure transparently, while also abstracting the management of components in the system. The control of such a system lies at the intersection of the spacecraft engineering and multi-agent systems fields, with neither approach assuredly superior. In this paper we investigate the suitability

of a distributed multi-agent system solution by comparing its performance to a centralised control implementation. Using a simplified task allocation model also allows us to verify the widely assumed increased robustness offered by a multi-agent system, which is one of the primary motivations for using this approach.

In [2] a multi-objective evolutionary algorithm is used to find a trade-off between signal delay and transmission power cost for communication in a multi-satellite system. Power consumption is decreased by using multi-hop routing, with the option of long distance transmissions to meet communication time constraints. An artificial potential field is used in [3] to position spacecraft in a lattice formation. By using a bottom-up approach it provides scalability and robustness. Self-organisation is also used in the proposed in-orbit assembly of large structures[4]. None of the above addresses the problem of task distribution in a group of spacecraft.

From a computational point of view, a group of spacecraft can be seen as a multi-robot system, where coordination in a hostile and noisy environment is a challenge. Alternatively, the multi-satellite system can be abstracted as a distributed computing network, where tasks need to be efficiently allocated to different nodes. Natural systems have provided useful metaphors for obtaining desired emergent behaviours in multi-robot contexts; for example, aggregations of robots through mimicking cockroaches[5], and divisions of labour[6] and foraging strategies [7] inspired by social insects. From the distributed computing perspective marketplaces have been popular mechanisms for allocating tasks in a distributed fashion. Typically, agents bid based on their suitability to perform a task and an auctioneer selects the best agent based on the size of the bid. The use of an artificial currency allows agents to make adaptive local decisions in a system where global information is incomplete and out of date. When viewed from a system-level perspective, these market-driven systems display collective social adaptive behaviour[8].

In [9] a market is used to allocate computing tasks in a heterogeneous network. The system self-organises to distribute loads fairly, by mapping idle resources into currency. The trade-off between sensing and data routing in wireless sensor networks (WSNs), is managed using a market mechanism in [10]. Nodes decide which role to fulfil, based on the payment they receive. By maximising their own revenue, the system performs close to optimal, dealing with changing sensor numbers and extending network lifetime. [11] combines the market-based allocation with robotics to efficiently explore a unknown territory. Robots compare their own cost of visiting a waypoint with trading it with a potentially better-situated vehicle. This maximises information retrieved while minimizing the system cost. Target allocation in unmanned miniature aerial vehicles (MAVs) using a market is presented in [12]. A distributed auction scheme takes the kinematic and sensing constraints of MAVs account.

The above examples demonstrate the utility of market-based control (MBC) in task allocation in distributed systems. The limited power and high communication costs in a distributed satellite mission are, however, not fully addressed. While WSNs share these constraints, their computation and communication de-

mands are quite different. We will use a distributed market-based control approach that encapsulates the energy of individual nodes, their capabilities and their location in the network to achieve robust, adaptive allocation.

Although several multi-spacecraft missions have been proposed, no actual implementations have been flown yet. We therefore propose the following reference mission: A group of small, low-cost satellites, numbering in the tens to hundreds are positioned in close proximity to each other in low earth orbit. Due to power constraints, communication is limited to only take place between neighbours, forming a network of autonomous, yet highly interdependent agents. The agents are specialised, with different classes displaying different skills: some are equipped to communicate with the ground station, while others carry remote observation cameras. The different payloads can operate independently, or be merged to provide synthesized data. For example, cameras can either be operated independently for low-resolution coverage of a wide area, or be combined to provide high-resolution images of a particular area of interest. Ideally, the group of satellites should be addressable as a single virtual spacecraft, with the detailed management of individuals left to the system. This reference mission is used to construct a simulation model that can be used to compare the robustness and efficiency of the centralised and distributed control methodologies directly.

2 Model

The desired multi-satellite system has no single point of failure and the goal is to allocate tasks in a manner that maximises the total amount of work done. We represent the system as a network of agents, with connections reflecting reliable wireless links between satellites. We currently have two types of agents: uplink nodes that communicate with the ground station; and worker nodes that perform tasks. A batch of tasks is uploaded to the uplink nodes at regular intervals, similar to a satellite having periodic contact with the ground station. Considered from the point of view of an uplink node, uploaded tasks are sequentially allocated to worker nodes using a sealed-bid reverse auction.

The uplink node acts as auctioneer and announces the auction by flooding a request describing the task through the system. If a node has enough energy and the required infrastructure to complete a job, it places a bid. The request message is repeated to its neighbours, who perform the same process. The value of the bid (B) is dependent on the ratio of maximum (e_{max}) to remaining energy (e_{rem}) of the node, the size s of the task and a scaling factor α that reflects the actual energy cost of performing the type of task: $B = \alpha s \frac{e_{max}}{e_{rem}}$. Returning bids are routed along the path of the original request message, with intermediate nodes adding a constant percentage commission to the bid. The uplink node selects the lowest bidder to assign the task to. The winning agent takes responsibility for the task, decreases its remaining energy (e_{rem}) by the cost of performing the task (αs) and receives payment. All the nodes in the communication path also receive their commission. As commission is multiplicative, it encourages the local allocation of tasks, because distant nodes appear more expensive to the




Fig. 1. Task allocation shifts to reflect changes in the network. In (a) the dotted line shows the area in which the majority of tasks sold by each auctioneer (shaded node) is allocated to. The nodes intersected by the dotted line receives tasks from both auctioneers. If a node fails, we lose its capacity to complete jobs and its routing functionality. The relative distances (measured in number of hops between nodes) and loading of nodes will change, resulting in a new distribution of labour in the network (b).

auctioneer. This bidding mechanism results in inexpensive bids from nodes that are under-utilised (large e_{rem}), while nodes that receive more allocations (lower e_{rem}) will increase their bids, making them less likely to be allocated a task. Agents do not try to win tasks by underbidding others — the quoted price is an honest reflection of the internal state of the node. Communication uses a significant portion of the total energy budget, in contrast to many multi-robot systems where it can safely be treated as negligible when compared to other energy uses. On every transmission e_{rem} is decreased by a value corresponding to the packet size. The commission parameter is responsible for minimising this communication expenditure.

When this process is executed concurrently across the network, the balance between the distributive effect of energy dependent bids and the localising effect of commission results in an allocation policy that is sensitive to levels of utilisation and changes in topology. Note that any node can act as an auctioneer, but for the simplified task structure above only uplink nodes allocate tasks.

Node failure is implemented as a uniformly distributed probability of failure per node per time step. Failed nodes cannot perform any work and have no communication links, thereby altering the network topology, as shown in Fig. 1.

This model has only one type of task and a relatively simple allocation problem. Although it is a simplified scenario, we believe it still captures enough of the dynamics of the system to allow a fair evaluation of control approaches. The design does, however, allow for future expansion: multiple task types can be used and recursive auctions can take place, which means any node can subcontract another node to perform part of a task. While the bid value is currently largely determined by the remaining energy in a node, it can easily be extended to include other system resources, such as bandwidth or memory.

Note the similarities between this abstraction and networked computing systems, where the best node for a task is determined by available CPU cycles, memory and disk space[9]. It also applies to wireless sensor networks, where node utility is dependent on remaining energy.

3 Results

The market-based control allocation scheme is used in a network with 100 nodes arranged in a 10 by 10 square lattice formation. Nine of the nodes act as uplink nodes; the remaining 91 are worker nodes. The duration of a day is 100 time steps. Our focus is on task allocation, so we assume formation flight is managed by a system similar to that presented in [3]. Tasks are introduced to the system at a constant rate (9 tasks per day), while nodes have failure probability of 0.001 per day per node. Approximately half the nodes usually fail after 800 days. The energy costs of tasks are generated from a Weibull distribution with shape $k = 2$ and scale $\lambda = 2$. The recharging of batteries from solar panels is implemented by increasing nodes' energy by 0.15 units per day, up to a maximum of 10 units per node. Transmission cost is set to 0.001 units per packet for negotiation and 0.1 units for allocation packets. As we currently only have a single task type, α is set to 1 for all nodes. Commission is set at 20%.

We compare our system to three other idealised control strategies to quantify its performance and robustness. Note that the following systems are unrealisable in a real world, but they do provide useful measures for comparison.


The *ideal* case represents the best possible performance. This assumes the controller has perfect knowledge of the network and can communicate cost-free with any node, without being constrained by network topology. Allocation is treated as a bin packing problem: for every task, the controller finds the worker node with the most remaining energy and assigns the task to it. The controller is considered immune against failure.

In the *centralised* approach, we have an intelligent mother ship that controls a network of simpler worker spacecraft. The level of realism is increased by reintroducing the network topology and transmission cost. The single controller node is positioned in the centre of the same lattice used by the distributed controller. The remaining 99 nodes are workers, as opposed to 91 workers for the MBC case. The controller has perfect information about the energy levels of nodes in the network, as well as the topology of the network — we ignore the cost of maintaining this information. Tasks are again assigned as in the ideal case, with the additional constraint that to allocate a job to a node, a valid path must exist between the central controller and the selected worker node. A path is valid if all nodes along it are active and have enough energy for transmission.


As the controller node in the centralised approach is a single point of failure, we assume that in a real mission scenario, it would incorporate redundancy to decrease its vulnerability. We therefore model this node as being immune to failure in the *centralised with immunity* (CI) case. All other variables are the same as used in the centralised approach.

3.1 Performance

The system was allowed to settle into steady state behaviour, before enabling the failure of nodes. The number of tasks successfully assigned over a period of



(a)



(b)

Fig. 2. Performance and robustness of different allocation strategies. In (a), network performance is measured by the total number of tasks allocated, normalised with respect to the steady-state performance of the ideal system. The amount of work stays constant, while nodes fail with a uniform probability. The ideal case, marked with triangles, provides an upper bound on the performance because it does not take communication cost or network topology into account. The market-based control allocation scheme (diamonds) deteriorates faster than the ideal case, but performs more efficiently than the centralised approaches (circles and squares). Robustness is defined as the ability to maintain performance despite satellite failures. In (b), the performance data is normalised with respect to the steady-state performance of the respective allocation strategies. The ideal case (triangles) shows the theoretical maximum obtainable, if topology and transmission cost have no influence. Market-based allocation (diamonds) shows a more gradual deterioration than either of the centralised approaches. The vulnerability of the network due to failure of the controller node is clearly visible when comparing centralised case (squares) to centralised with a controller that is immune to failure (circles). The solid horizontal line indicates 50% of the initial throughput.

100 days was measured and normalised with respect to the steady-state performance of the ideal allocator. This was repeated 100 times to obtain an average behaviour; the resulting performance is shown in Fig. 2(a).

The ideal system can be seen to form an upper bound on the allocation success. It slowly deteriorates over time as the number of failed nodes increases and the system’s capacity to complete jobs decreases. The MBC approach displays lower initial performance: due to the energy cost of communication it accommodates only 76.9% of the tasks the ideal case does. Steady-state performance drops to 68.9% for both centralised control schemes, because of the larger portion of the energy budget spent on communication (the average path length when allocation tasks is greater than with MBC). Progressive node failure decreases the total capacity of the network: allocation paths becomes longer and use more energy, and the network is fragmented when all routes to functioning nodes are cut. This is reflected in the steep slope of the centralised, CI and MBC data. The sensitivity of the network to failure of the central controller is significant, as can be seen when comparing the centralised approach to CI.

Additional experiments confirmed that the behaviour of the system is robust to variations in parameter values. The qualitative observations still hold, although some quantitative changes occur. For smaller networks, the centralised and MBC results converge. If the ratio of transmission cost to task size changes, the performance will increase (for smaller packets) or decrease (for larger packets) accordingly.

3.2 Robustness

We define robustness as the ability of the system to maintain steady-state performance despite satellite failures. To compare the robustness of the different systems, the results from Sect. 3.1 are normalised with regards to their respective steady-state values (Fig. 2(b)).

The ideal case again provides an upper bound. The centralised case deteriorates rapidly, largely due to the whole network collapsing if the controller node fails. While the CI approach performs better, it is still subject to network fragmentation. The MBC approach is more robust than both centralised systems. To express these results in terms of mission reliability, we define a mission as operational while it delivers more than 50% of its initial throughput. The centralised system reaches this limit at 360 days, the immune-centralised at 525 days and the MBC approach at 605, making it the most reliable of the three. In spite of having fewer worker nodes, the performance of the MBC system is superior. This is not only related to efficiency, but also to robustness. In particular, this is a result of having multiple uplink nodes which are able to adapt their allocation to changes in the network topology and node utilisation.

4 Discussion

We have shown that a market-based task allocation system completes more jobs and is more robust than a centralised approach, irrespective of whether the

central controller is subject to failure. The improved performance is a result of lower system-level communication costs when assigning jobs and improved robustness, due to the distributed and adaptive nature of the control system. These results are promising for distributed space applications. Launch mass will always be the dominant factor in total mission cost and, assuming a given launch mass and spacecraft of equal size, our results show that more work can be done more robustly using an MBC approach.

The results are also applicable to robustly controlling systems with similar constraints, such as WSNs and distributed computing systems, by using emergent behaviour. The work presented here is the first step towards a task-allocation mechanism for a distributed satellite system. Future work will look at composite tasks, requiring cooperation; adding temporal constraints to tasks; optimising the energy cost of transmissions to match communication distances; optimal composition of the satellite types in the group; and enhanced physics to provide more realism.

References

1. Brown, O., Eremenko, P.: Fractionated space architectures: A vision for responsive space. In: *Responsive Space 4*. (2006)
2. Wu, X., Vladimirova, T., Sidibeh, K.: Signal routing in a satellite sensor network using optimisation algorithms. In: *IEEE Aerospace Conference*. (2008) 1–9
3. Pinciroli, C., Birattari, M., Tuci, E., Dorigo, M., Marco, Vinko, T., Izzo, D.: Self-organizing and scalable shape formation for a swarm of pico satellites. In: *NASA/ESA Conference on Adaptive Hardware and Systems*. (2008) 57–61
4. Izzo, D., Pettazzi, L., Ayre, M.: Mission concept for autonomous on orbit assembly of a large reflector in space. In: *56th International Astronautical Congress*. (2005)
5. Garnier, S., Jost, C., Jeanson, R., Gautrais, J., Asadpour, M., Caprari, G., Theraulaz, G.: Aggregation behaviour as a source of collective decision in a group of cockroach-like-robots. In: *Advances in Artificial Life*. (2005) 169–178
6. Groß, R., Nouyan, S., Bonani, M., Mondada, F., Dorigo, M.: Division of labour in self-organised groups. In: *Simulation of Adaptive Behaviour*. Volume 5040. (2008) 426–436
7. Campo, A., Dorigo, M.: Efficient multi-foraging in swarm robotics. In: *Advances in Artificial Life*. (2007) 696–705
8. Cliff, D., Bruten, J.: Animat market — trading interactions as collective social adaptive behavior. *Adaptive Behavior* **7**(3-4) (January 1999) 385–414
9. Waldspurger, C.A., Hogg, T., Huberman, B.A., Kephart, J.O., Stornetta, S.W.: Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering* **18**(2) (February 1992) 103–117
10. Rogers, A., David, E., Jennings, N.R.: Self-organized routing for wireless microsensor networks. *Systems, Man and Cybernetics, Part A, IEEE Transactions on* **35**(3) (2005) 349–359
11. Zlot, R., Stentz, A., Dias, M.B., Thayer, S.: Multi-robot exploration controlled by a market economy. In: *IEEE International Conference on Robotics and Automation*. Volume 3. (2002) 3016–3023
12. Sujit, P.B., Beard, R.: Multiple MAV task allocation using distributed auctions. In: *AIAA Guidance, Navigation and Control Conference and Exhibit*. (August 2007)