Scalable Decentralised Decision Making and Optimisation in Heterogeneous Teams

George M. Mathews and Hugh F. Durrant-Whyte
ARC Centre of Excellence for Autonomous Systems
The University of Sydney, Sydney NSW 2006, Australia
Email: {g.mathews,hugh}@cas.edu.au

Mikhail Prokopenko
CSIRO ICT Centre
Locked Bag 17 North Ryde, NSW 1670, Australia
Email: Mikhail.Prokopenko@csiro.au

***Abstract*— This paper considers the scenario where multiple autonomous agents must cooperate in making decisions to minimise a common team cost function. A distributed optimisation algorithm is presented. This allows each agent to incrementally refine their decisions while intermittently receiving updates from the team. A convergence analysis provides quantitative requirements on the frequency agents must communicate that is prescribed by the problem structure. For a general multi-agent system, this solution requires every agent to have a model of every other agent. To overcome this, a specific class of systems, called Partially Separable, is defined. These systems only require each agent to have a combined summary of the rest of the system. This leads to the definition of an infinitely scalable system, which may contain an infinite number of agents while ensuring the local decisions will converge to the optimal team decision. Examples are given for reconnaissance or information gathering tasks.**

I. Introduction

This paper is part of an ongoing research program into decision making and control algorithms for large distributed active sensor networks. The vision of this work is a system of interconnected robotic agents, such as UAVs, AGVs, or simply fixed but controlable sensors, where each agent communicates and makes decisions locally while ensuring the global system goals are obtained.

This type of system has applications such as environmental monitoring, searching, surveillance, bush firefighting, target tracking, mapping and exploration, for example.

Utilising an autonomous system to perform these tasks allows human operators to be removed from possibly dangerous (or simply mundane) situations. However, the control complexity of these large distributed systems grows rapidly with the number of agents [1].

Existing approaches to this problem are either: (1) Fully centralised [2] where the system as a whole is modelled and controlled using conventional techniques. (2) Distributed or hierarchical [3] which utilise the distributed computational capacity of the multiple platforms but require a single facility to fuse information or resolve global constrains, or (3) Fully decentralised which do not require any centralised facility.

The decentralised approach can be further broken down into systems that require each platform to have global information about the group [4] and those that only require local information from a small subset. It is the latter that is of interest here.

This paper quantitatively examines what information is required by each agent such that the optimal team decision can be realised. From these requirements a specific type of system is defined, called *Partially Separable*, which allows this information to be communicated around the system in a scalable fashion, requiring only fixed inter-agent bandwidth.

Section II presents a formal definition of the distributed optimal decision problem. To solve this, a general distributed optimisation algorithm is presented in Section III for smooth and differentiable cost functions. Section IV present sufficient conditions for convergence that quantitatively specify what information is required by each agent. The scalability implications of these requirements are discussed in Section V and the Partial Separability condition defined. Conclusions are presented in Section VI.

II. Problem Definition

Consider a team of p agents, where each agent i is in charge of a local decision or control variable $u_i \in \mathcal{U}_i$. This work assumes the decision spaces are scalar and thus $\mathcal{U}_i \equiv \Re$. (For multi-dimensional decisions, it can be assumed that sub-agents control individual dimensions.)

The team members are required to select their decisions such that a given team cost function $J : \mathcal{U}_1 \times \cdots \times \mathcal{U}_p \to \Re$ is minimised

$$\mathbf{u}^* = \arg\min_{\mathbf{u}} J(\mathbf{u}) \tag{1}$$

where $\mathbf{u} = [u_1, u_2, \ldots, u_p]^T \in \mathcal{U}$ is the team decision vector and $\mathbf{u}^*$ is the desired optimal team decision. If the cost function is continuously differentiable and convex, (1) becomes

$$\frac{\partial J}{\partial u_i}(\mathbf{u}^*) = 0 \quad \forall\, i. \tag{2}$$

III. Distributed Optimisation Model

This section presents a general distributed optimisation model for smooth and continuous cost functions based on the work of Bertsekas and Tsitsiklis [5].

The minimisation in (1) is accomplished using an asynchronous gradient descent algorithm. During the optimisation procedure each agent maintains a local copy of the team decision vector. Due to the asynchronous requirements, this vector may contain outdated information about other agents and is given at discrete time t for agent i as

$$\mathbf{u}^i(t) = \big[u_1^i(t), \ldots, u_p^i(t)\big]^T \quad (3)$$
$$= \big[u_1^1\big(\tau_1^i(t)\big), \ldots, u_p^p\big(\tau_p^i(t)\big)\big]^T. \quad (4)$$

In general superscripts represent agents and subscripts represent a component of the decision vector (e.g. $u_j^i(t)$ represents agent i's local copy of agent j's component of the team decision vector). The variable $\tau_j^i(t)$ in (4) represents the time agent i's local copy $u_i^j(t)$ was generated by agent j and hence $u_j^i(t) = u_j^j\big(\tau_j^i(t)\big)$. It is assumed that $\tau_i^i(t) = t$ and thus agent i always has the latest copy of its decision variable.

A. Local Update

Each agent employs a local update rule $f_i : \mathcal{U} \to \mathcal{U}_i$ that modifies its component of the decision vector. To allow each agent to perform updates asynchronously a set of times T_U^i will be associated with each agent i that represent when the agent computes a local update. Thus

$$u_i^i(t+1) = \begin{cases} f_i\big(\mathbf{u}^i(t)\big) & \text{if } t \in T_U^i \\ u_i^i(t) & \text{else} \end{cases} \quad (5)$$

For a general gradient descent type algorithm, the update function has the form

$$f_i\big(\mathbf{u}^i(t)\big) = u_i^i(t) + \gamma_i s_i\big(\mathbf{u}^i(t)\big) \quad (6)$$

where γ_i is a constant step size and $s_i : \mathcal{U} \to \mathcal{U}_i$ is the update direction, e.g. for steepest descent

$$s_i\big(\mathbf{u}^i(t)\big) = -\frac{\partial J}{\partial u_i}\big(\mathbf{u}^i(t)\big) \quad (7)$$

B. Communication

Communication is initiated by agent i sending a message, at some time $t \in T_C^{ij}$, to agent j containing its latest decision $u_i^i(t)$. After some communication delay $b_{ij}(t)$ agent j receives it and incorporates it into its local copy of the team decision vector, thus when the message is received $u_i^j\big(t + b_{ij}(t)\big) = u_i^i(t)$ and $\tau_i^j\big(t + b_{ij}(t)\big) = t$.

IV. Convergence Results

This section presents sufficient conditions for the above algorithm to converge on the optimal team decision. These conditions quantitatively relate the communication frequency, update direction and problem structure to a maximum allowable step size for each agent.

A. General Case

Convergence is proved through ensuring cost reduction. This can be accomplished if some bounds on the cost functions Hessian are known. Each agent then bounds the effects of other agents updates on the team cost between successive communications.

If the cost is convex the algorithm will converge to the global minimum. For more general functions only a local minimum can be guaranteed.

Assumption 1 (Bounded Cost): There exists a constant E such that

$$J(\mathbf{u}) \geq E \quad \forall\, \mathbf{u} \in \mathcal{U}. \quad (8)$$

That is, the cost cannot be arbitrarily small.

Assumption 2 (Coupling): There exists positive constants K_{ij} such that

$$\left|\frac{\partial^2 J}{\partial u_i \partial u_j}(\mathbf{u})\right| \leq K_{ij} \quad \forall\, i, j, \mathbf{u} \in \mathcal{U}. \quad (9)$$

This imposes a limit on the maximum coupling between the decisions made by different agents, represented by bounds on the second derivatives of the objective function.

Assumption 3 (Local Descent): For every t and i

(i) $$s_i\big(\mathbf{u}^i(t)\big)\frac{\partial J}{\partial u_i}\big(\mathbf{u}^i(t)\big) \leq 0. \quad (10)$$

(ii) There exists positive constants C_i and D_i such that

$$C_i\left|\frac{\partial J}{\partial u_i}\big(\mathbf{u}^i(t)\big)\right| \leq \big|s_i\big(\mathbf{u}^i(t)\big)\big| \leq D_i\left|\frac{\partial J}{\partial u_i}\big(\mathbf{u}^i(t)\big)\right|. \quad (11)$$

This ensures the updates made by individual agents are in a direction of decreasing cost (based on local information), are bounded and may only be zero if the gradient of the cost function is zero.

Assumption 4 (Bounded Delays): There exists positive constants B_{ij} such that

$$t - \tau_i^j(t) \leq B_{ij} \quad \forall\, i, j, t \quad (12)$$

Thus, the time difference between agent i's local decision u_i^i, and agent j's copy of it u_i^j, is bounded. Informally, this can be relaxed such that B_{ij} represents the time difference, measured in numbers of updates computed by agent i, between $u_i^i(t)$ and $u_i^j(t)$.

Theorem 1 (General Convergence): Assumptions 1 – 4 provide sufficient conditions for the distributed optimisation algorithm defined by Equations (5) and (6) to converge with

$$\lim_{t\to\infty} \frac{\partial J}{\partial u_i}\big(\mathbf{u}^G(t)\big) = 0 \quad \forall\, i \quad (13)$$

where $\mathbf{u}^G(t) = \big[u_1^1(t), \ldots, u_p^p(t)\big]$ is the global decision vector, for all $0 < \gamma_i < \bar{\gamma}_i$, where

$$\bar{\gamma}_i = \frac{2}{D_i \sum_{j=1}^{p} K_{ij}\big(1 + B_{ij} + B_{ji}\big)} \quad (14)$$

See [6] for proof.

This general result provides a unified way to relate communication requirements and problem structure for a general cooperative multi-agent system. This is a necessary condition for a quadratic cost function with no delays (i.e. $B_{ij} = 0$ for all i and j). However, as the delays are increased it becomes only sufficient.

In general it is not known how far this diverges away from the necessary conditions. It is known however, that for a sufficiently weakly coupled system the delays can become arbitrarily large [5], [7]. This case will now be reviewed.

B. Weak Coupling

This scenario is somewhat of a special case and does not require constraints on the communication frequency. In a general sense, each agent may independently minimising the cost function in its dimension, while communicating at arbitrarily large intervals.

Convergence can be shown using a contraction mapping approach that ensures, after successive iterations, the global team decision moves closer to the optimum. The following assumptions formally define weak coupling and arbitrary delays.

Assumption 5 (Weak Coupling): For all $\mathbf{u} \in \mathcal{U}$ and all i

$$\frac{\partial^2 J}{\partial {u_i}^2}(\mathbf{u}) > \sum_{j \neq i} \left| \frac{\partial^2 J}{\partial u_i \partial u_j}(\mathbf{u}) \right|. \tag{15}$$

This means the i^{th} component of the gradient vector, $\frac{\partial J}{\partial u_i}(\mathbf{u})$, always has a greater dependence on agent i's decision than the sum of all dependencies from other agents and requires the Hessian obey a diagonal dominance condition over the full decision space $\mathcal{U}$.

Assumption 6 (Arbitrary Delays): Given any time t_1, there exists a time $t_2 > t_1$ such that

$$\tau_j^i(t) > t_1 \quad \forall\, i, j, t > t_2. \tag{16}$$

Thus the delays in the system can become arbitrarily large, while enforcing that after a sufficiently long time (t_2) all old information (from before t_1) is removed from the system.

Theorem 2 (Weak Coupling Convergence): Under Assumptions 1 – 3, 5 and 6 the distributed optimisation algorithm defined by Equations (5) and (6) will converge with

$$\lim_{t \to \infty} \frac{\partial J}{\partial u_i}\big(\mathbf{u}^G(t)\big) = 0 \quad \forall\, i \tag{17}$$

where $\mathbf{u}^G(t) = \left[u_1^1(t), \dots, u_p^p(t)\right]$ is the global decision vector, for all $0 < \gamma < \bar{\gamma}$, where

$$\bar{\gamma} = \frac{1}{D_i K_{ii}} \tag{18}$$

See [5] or [7] for a proof.

Since there is no restriction on the frequency of communications, if each agent only communicates after they have minimised the objective function in the direction of their component u_i, then the nonlinear Jacobi update is obtained

$$u_i^i(t+1) = \arg \min_{u_i^i \in \mathcal{U}_i} J\big(u_1^i(t), \dots$$
$$\dots, u_{i-1}^i(t), u_i^i, u_{i+1}^i(t), \dots, u_p^i(t)\big). \tag{19}$$

Which, under Assumptions 1 – 3, 5 and 6, is also guaranteed to converge.

C. Discussion

The convergence results of Theorem 1 and 2 provides a natural way to organise the communication network of a distributed decision making problem as they *quantitatively* describe which agents need to communicate and how frequently.

However, in the current form this optimisation procedure requires all participating agents to communicate to, and have a complete model of, potentially all other agents (i.e. so they have access to the team decision vector and can evaluate the gradient of the objective function w.r.t their decision). For large teams this poses a significant scalability issue.

V. Scalable Decentralised Optimisation

This section focuses on the team decision problem with emphasis on scalability. It considers what information each agent requires and how it should be communicated around the team.

To start, consider a general heterogeneous team of agents at a specific instant in time (i.e. for a given set of environmental and agent states) where the team is required to make a decision.

A completely general objective or cost function can be evaluated if a complete model of each agent, with their corresponding decisions, are known.

Thus, if the space of all possible agent models is $\mathcal{M}$, and the decision space for agent i is $\mathcal{U}_i$ (which may be model dependent). The team cost function can be defined as the mapping from the agents models and decisions to the real line.

$$J : \mathcal{U}_1 \times \mathcal{M} \times \cdots \times \mathcal{U}_i \times \mathcal{M} \times \cdots \times \mathcal{U}_p \times \mathcal{M} \to \Re. \tag{20}$$

Thus, in general if any agent wants to evaluate the team objective (and its derivatives) it needs the models and decisions from the rest of the team, which is not feasible for large teams.

A. Reconnaissance Example

A typical example of interest here is the reconnaissance or information gathering task. This scenario requires information to be gathered on a particular random variable $\mathbf{x} \in \mathcal{X}$. This may include the positions of targets, terrain properties of a given region, or surface information of a remote planet, for example.

In [8] this is formulated as a single step look ahead optimal control problem. This requires each agent to decide on a control action such that as a team their next set of sensor observations extracts the most information about $\mathbf{x}$.

The team is assumed to start with some prior belief $P(\mathbf{x})$. Each agent has a sensor which generates observations, $\mathbf{z}_i \in \mathcal{Z}_i$, from a conditional probability density $P(\mathbf{z}_i|\mathbf{x}, u_i)$. This describes the probability of agent i getting an observation $\mathbf{z}_i$ for a given state $\mathbf{x}$ and decision u_i. The posterior team belief, after all agents decisions and observations are made, is given by Bayes rule as

$$P(\mathbf{x}|\mathbf{z}, \mathbf{u}) = \frac{P(\mathbf{x}) \prod_{i=1}^p P(\mathbf{z}_i|\mathbf{x}, u_i)}{\int_{\mathcal{X}} P(\mathbf{x}) \prod_{i=1}^p P(\mathbf{z}_i|\mathbf{x}, u_i) d\mathbf{x}} \tag{21}$$

where $\mathbf{z} = [\mathbf{z}_1, \dots, \mathbf{z}_p]^T \in \mathcal{Z} = \mathcal{Z}_1 \times \cdots \times \mathcal{Z}_p$. The goal of the task is to minimise the uncertainty after the observations have been taken. Thus a suitable team cost function is the expected entropy [9] of the posterior, where the expectation is taken over all possible team observations (this is equivalent to the conditional entropy of $\mathbf{x}$ given $\mathbf{z}$)

$$J = -\mathrm{E}_{\mathbf{z}}\mathrm{E}_{\mathbf{x}}[\log P(\mathbf{x}|\mathbf{z},\mathbf{u})], \tag{22}$$

which can be expanded as

$$J = -\int_{\mathcal{Z}}\int_{\mathcal{X}} P(\mathbf{x})\prod_{i=1}^{p} P(\mathbf{z}_i|\mathbf{x},u_i) \times \log\frac{P(\mathbf{x})\prod_{i=1}^{p} P(\mathbf{z}_i|\mathbf{x},u_i)}{\int_{\mathcal{X}} P(\mathbf{x})\prod_{i=1}^{p} P(\mathbf{z}_i|\mathbf{x},u_i)dx} d\mathbf{x}d\mathbf{z}. \tag{23}$$

This team cost function can be cast in the form of (20) by considering the conditional density $P(\mathbf{z}_i|\mathbf{x},u_i)$ as the model for the i^{th} agent, defined as an element of the space $\mathcal{M}$ of all conditional densities defined on $\mathcal{Z}_i \times \mathcal{X} \times \mathcal{U}_i$.

It is also noted that the cost is based on the common team belief, $P(\mathbf{x})$, which is neither part of the decisions or the agent models. This is assumed part of the system state at that specific time the decisions are to be made.

Thus, if the observation models are continuous and differenable, the optimal team decision problem, defined by the minimisation of (23), can only be solved if every agent has a model of every other agent.

This scenario may be realistic for small teams, and even beneficial, since, during the optimisation process, only the decisions u_i need to be communicated (which may only be scalars). However, this requirement is unrealistic for large heterogeneous systems.

This result demonstrated that the general multi-sensor reconnaissance or information gathering problem is intractable for an arbitrarily large heterogeneous system. However, there are some scenarios (such as the cooperative control problem described in [10]) that are tractable and can to be solved in a scalable fashion. The difference between these systems and the general case will now be examined.

B. Partial Separability

In general terms, it will be shown that if the effect or impact of a group of agents on the team objective or cost function, can be described in the same manner as the impact of a single agent. Then, each agent can consider the rest of the team as a single entity which is described in the same manner as any other agent, regardless of the size the team. It will be shown this enables the team decision problem to be solved for a system of potentially infinite agents.

The key behind this simple idea, is being able to describe the impact of a group of agents in the same manner as the impact of a single agent. These types of systems will be called *Partially Separable*.

Definition 1 (Impact Function): An impact function $\Upsilon_{\mathbf{m}_i}$, of a given agent i with model $\mathbf{m}_i \in \mathcal{M}$, maps a decision u_i onto a common team impact space $\mathcal{S}$

$$\Upsilon_{\mathbf{m}_i} : \mathcal{U}_i \to \mathcal{S} \tag{24}$$

The impacts from two agents, $\alpha_i = \Upsilon_{\mathbf{m}_i}(u_i)$ and $\alpha_j = \Upsilon_{\mathbf{m}_j}(u_j)$ can be combined using a binary composition operator $*$, defined on the space $\mathcal{S}$. This allows the combined impact (or summary) to be given as $\alpha_{i,j} = \alpha_i * \alpha_j$, which remains in the same space as each individual impact. This is generalised for two non-overlapping subsets of the team.

Definition 2 (Impact Properties): For all non-empty and non-overlapping sets of agents $\mu,\nu \subset \{1,\ldots,p\}$, such that $\mu\bigcap\nu = \emptyset$ and $\sigma = \mu\bigcup\nu$, there exists

(i) A commutative and associative binary operator $*$ defined for all $\alpha_\mu,\alpha_\nu \in \mathcal{S}$ such that the combined impact α_σ is given by

$$\alpha_\sigma = \alpha_\mu * \alpha_\nu \tag{25}$$

(ii) An inverse of all elements of $\mathcal{S}$ and a unity element $\alpha_\emptyset$, such that the impact α_μ can be recovered by

$$\alpha_\sigma * \alpha_\nu^{-1} = \alpha_\mu * \alpha_\nu * \alpha_\nu^{-1} = \alpha_\mu * \alpha_\emptyset = \alpha_\mu \tag{26}$$

This formally defines an *Abelian group*.

The team objective function is equivalent to a mapping from the team impact α_T to $\Re$. This is in contrast to (20), which defines it as mapping directly from the team decisions u_i and models $\mathbf{m}_i$ to $\Re$.

Definition 3 (Generalised Cost Function): The team cost function is equivalent to the mapping $\psi : \mathcal{S} \to \Re$ and is given by

$$J(\mathbf{u}) = \psi(\alpha_T) \tag{27}$$

where $\alpha_T = \Upsilon_{\mathbf{m}_1}(u_1) * \Upsilon_{\mathbf{m}_2}(u_2) * \cdots * \Upsilon_{\mathbf{m}_p}(u_p)$.

The above definitions do not impose any restrictions on the system. This can be seen if the impact space is large enough to contain the decisions and models of all agents. Thus the largest impact space for any general system is defined as

$$\mathcal{S} = \mathcal{U}_1 \times \mathcal{M} \times \cdots \times \mathcal{U}_p \times \mathcal{M}. \tag{28}$$

The general reconnaissance example described in Section V-A can be considered to have the slightly smaller impact space of $\mathcal{S} = \mathcal{F}_1 \times \ldots \times \mathcal{F}_p$ where $\mathcal{F}_i$ is the space of all conditional densities defined on $\mathcal{Z}_i \times \mathcal{X}$. This can be seen by examining (23) and noting that if nothing is known about the agents, the cost function can be evaluated if each agents conditional probability $P(\mathbf{z}_i|\mathbf{x},u_i)$, evaluated for their chosen decisions, is known.

However, the fundamental contribution of this formulation occurs when the size of the impact space is independent of the number of agents. A system that has this property will be called *Partially Separable*.

Definition 4 (Partially Separable System): A system is partially separable if there exists an impact space for its generalised cost function whose dimensionality is not a function of the number of agents in the system.

To demonstrate this concept, the previous reconnaissance scenario will be examined for the case when all distributions are Gaussian.

C. Reconnaissance with Gaussian Uncertainty

For a detailed reference on Gaussian filtering and the extended Kalman filter see [11]. The prior team belief is assumed to have mean $\hat{\mathbf{x}}$ and covariance $\mathbf{P}$

$$P(\mathbf{x}) \sim N(\mathbf{x};\hat{\mathbf{x}},\mathbf{P}). \tag{29}$$

The i^{th} agents observation model is given in linearised form as

$$P(\mathbf{z}_i|\mathbf{x},u_i) \sim N(\mathbf{z}_i;\mathbf{h}_i(\mathbf{x},u_i),\mathbf{R}_i) \quad (30)$$

where $\mathbf{h}_i(\mathbf{x},u_i) \approx \bar{\mathbf{z}}_i + \mathbf{H}_i(\mathbf{x}-\hat{\mathbf{x}})$, $\bar{\mathbf{z}}_i = \mathbf{h}_i(\hat{\mathbf{x}},u_i)$ and $\mathbf{H}_i = \nabla_{\mathbf{x}}\mathbf{h}_i(\hat{\mathbf{x}},u_i)$.

For a given team decision $\mathbf{u}$ and observation $\mathbf{z}$, the posterior $P(\mathbf{x}|\mathbf{z},\mathbf{u})$ can be computed using Bayes rule (21). This is identical to an extended Kalman filter update, thus

$$P(\mathbf{x}|\mathbf{z},\mathbf{u}) \sim N(\mathbf{x};\hat{\mathbf{x}}_+,\mathbf{P}_+) \quad (31)$$

where the updated mean $\hat{\mathbf{x}}_+$ and covariance $\mathbf{P}_+$, given in inverse covariance form, are

$$\mathbf{P}_+^{-1} = \mathbf{P}^{-1} + \sum_{i=1}^{p} \mathbf{H}_i^T\,\mathbf{R}_i^{-1}\,\mathbf{H}_i, \quad (32)$$

$$\mathbf{P}_+^{-1}\,\hat{\mathbf{x}}_+ = \mathbf{P}^{-1}\hat{\mathbf{x}} + \sum_{i=1}^{p} \mathbf{H}_i^T\,\mathbf{R}_i^{-1}\,(\mathbf{z}_i - \bar{\mathbf{z}}_i + \mathbf{H}_i\mathbf{x}). \quad (33)$$

The cost function is given by the expectation of the posterior entropy, as defined in (23), and becomes

$$J = \frac{1}{2}\log\left((2\pi e)^{d_x}\left|\mathbf{P}_+^{-1}\right|\right). \quad (34)$$

where d_x is the dimension of the state space $\mathcal{X}$. Using (32) and noting the cost function is only required as a preference ordering over actions, (34) becomes

$$J = \left|\mathbf{P}^{-1} + \sum_{i=1}^{p}\mathbf{I}_i(u_i)\right|. \quad (35)$$

where $\mathbf{I}_i(u_i) = \mathbf{H}_i^T\,\mathbf{R}_i^{-1}\,\mathbf{H}_i$. For the above cost function to be evaluated only the matrix sum, $\sum_{i=1}^{p}\mathbf{I}_i(u_i)$, is required. Knowledge of the model parameters $\mathbf{H}_i$ and $\mathbf{R}_i$ or the individual matrices $\mathbf{I}_i(u_i)$ from each agent is redundant and is not needed to evaluate the cost of a given team decision.

This is precisely the specification for a partially separable system. This can be emphasised by defining the impact function for the i^{th} agent as $\mathbf{I}_i(u_i)$, which maps decisions u_i to an impact space $\mathcal{S}$, which contains all $d_x \times d_x$ symmetric matrices. This allows the individual models of each agent to be abstracted away and only the relevant features extracted

$$\Upsilon_{\mathbf{m}_i}(u_i) = \mathbf{I}_i(u_i) = \mathbf{H}_i^T\,\mathbf{R}_i^{-1}\,\mathbf{H}_i. \quad (36)$$

Here, the impact space, $\mathcal{S}$, is independent of the number of agents and thus, the system is Partially Separable.

With this formulation, the composition operator $*$ becomes matrix addition, an elements inverse is its negative, and the unity element becomes the zero matrix. The team cost function is redefined in Partially Separable form as a mapping from $\mathcal{S} \to \Re$ and is given by

$$J = \left|\mathbf{P}^{-1} + \alpha_T\right|. \quad (37)$$

where $\alpha_T = \Upsilon_{\mathbf{m}_1}(u_1) * \cdots * \Upsilon_{\mathbf{m}_p}(u_p)$. This formulation only requires each agent to consider the combined team impact α_T, nothing else about the team is required.

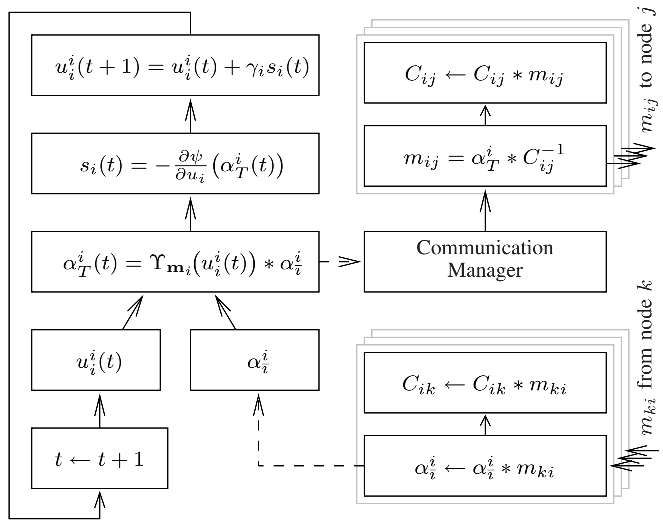

Fig. 1. Internal agent structure for steepest descent optimisation (i.e. $s_i(t) = -\frac{\partial\psi}{\partial u_i}(\alpha_T^i(t))$). The left hand side displays the optimisation loop, where local decisions are continually refined. External information regarding the rest of the teams impact $\alpha_{\bar{i}}^i$, which the refinements are based on, is continuously being updated via the communication links on the right. Dashed arrows indicate asynchronous information flows.

D. Optimisation and Partial Separability

For a partially separable system the optimal decisions are found using the distributed optimisation algorithm as described in Section III. However, instead of requiring the decisions from the team, each agent only requires the combined team impact. Thus every agent need only maintain a local copy of the team impact α_T^i.

The generic internal structure of each agent is shown in Fig. 1. The left hand side of the diagram displays the optimisation loop, where local decisions are continually refined. It is noted that the local team impact α_T^i is separated into the local impact $\alpha_i^i = \Upsilon_{\mathbf{m}_i}(u_i^i)$ and the combined impact of the rest of the team $\alpha_{\bar{i}}^i$ (where $\bar{i} \equiv \{j : \forall\, j \neq i\}$). The impact $\alpha_{\bar{i}}^i$, completely summaries the effects of the proposed decisions of all other agents in the team. The team impact can be easily recovered using $\alpha_T^i = \alpha_i^i * \alpha_{\bar{i}}^i$.

E. Communication

The proposed communication protocol (shown on the left side of Fig. 1) allows each agent to communicate locally with its neighbours within an acyclic communication network, while ensuring a copy of the global team impact is obtained. Let $n_i \subset \{1,\ldots,p\}$ be the set of neighbours agent i has in the communication network. For every neighbour $j \in n_i$ agent i has, it maintains a buffer $C_{ij} \in \mathcal{S}$ that contains the component of the team impact that is known by both agents i and j (initially this buffer contains the unity impact $\alpha_\emptyset$, representing no shared information).

Each time agent i is required to synchronise its knowledge with agent j it sends an *update message*, $m_{ij} \in \mathcal{S}$ by first removing this common component.

$$m_{ij} = \alpha_T^i * C_{ij}^{-1}. \quad (38)$$

This message contains the differences between what it has communicated before and what it now knows (by changing its own impact α_i^i, or through communications from other agents). Once sent (and known to be received), the buffer is updated

$$C_{ij} \leftarrow C_{ij} * m_{ij}. \quad (39)$$

On reception, agent j uses the message to update its impact of the rest of the team $\alpha_{\bar{j}}^j$ (since the message cannot contain information that originated from itself)

$$\alpha_{\bar{j}}^j \leftarrow \alpha_{\bar{j}}^j * m_{ij} \quad (40)$$

and similarly updates its buffer

$$C_{ji} \leftarrow C_{ji} * m_{ij}. \quad (41)$$

This process ensures consistency of each agents copy of the team impact.

For a strongly coupled system, the results of Theorem 1, provide quantitative requirements on the magnitude of the step lengths, γ_i, based on the age of the information contained within the local team impacts, α_T^i, and the degree of coupling within the system.

The results of Theorem 2 greatly lessen these requirements for weakly coupled systems. For these systems, the maximum step lengths $\bar{\gamma}_i$, are not influence by the communication frequency or the delays in the communication network.

F. Infinitely Scalable System

Consider a strongly coupled multi-agent system containing a number of agents $p \rightarrow \infty$, distributed over an area $A \rightarrow \infty$. For this type of system, it is reasonable to assume each agent is only strongly coupled to a few other *close* agents and weakly coupled to, or independent of, all others. Then, even though the communication delays will approach infinity, the system may still converge on the optimal team decision in the same manner as a ordinary finite system. This leads to the notion of infinite scalability.

Definition 5 (Infinite Scalability): If there exists a finite positive constant L such that

$$\lim_{p \rightarrow \infty} \sum_{j=1}^{p} K_{ij} \big(1 + B_{ij} + B_{ji}\big) < L \quad \forall\, i. \quad (42)$$

Then as $t \rightarrow \infty$ the system will converge on the optimal team decision and the system is infinitely scalable.

This can be proved by considering the maximum step size of each agent $\bar{\gamma}_i$. From Theorem 1

$$\bar{\gamma}_i = \frac{2}{D_i \sum_{j=1}^{p} K_{ij} \big(1 + B_{ij} + B_{ji}\big)} \quad (43)$$

and hence

$$\lim_{p \rightarrow \infty} \bar{\gamma}_i > \frac{2}{D_i L} > 0. \quad (44)$$

Thus, each agent can maintain a positive step length $\gamma_i < \bar{\gamma}_i$, which ensures the decentralised optimisation algorithm will make sufficient progress to converge.

VI. Conclusions

This paper considered the scenario of multiple autonomous agents cooperating in making decisions that minimise a common team cost function. A distributed asynchronous optimisation algorithm was presented to solve this decision problem. Based on the convergence conditions of this algorithm, quantitative communication requirements for the team were found.

However, this classical analysis assumes each agent has complete knowledge of the rest of the team; which is impractical for large heterogeneous systems. *Partially Separable* systems where introduced and describe a subset of multi-agent systems for which the above problem can be solved in a scalable fashion. These systems only require each agent to have a combined summary of the rest of the system and not a full model of all other agents.

It was shown that the general reconnaissance or information gathering problem is not separable, and requires each agents to have information about every other agent. However, when all distributions are Gaussian this problem becomes Partially Separable. In [8] an instance of the indoor mapping problem was shown to be partially separable. The multi-platform maritime search scenario, described in [12], is also believed to be partially separable.

Acknowledgments

This work is partly supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government, and by CSIRO Industrial Physics.

References

[1] D. V. Pynadath and M. Tambe, "Multiagent teamwork: Analyzing the optimality and complexity of key theories and models," in *Proc. Int. Conf. of Autonomous Agents and Multiagent Systems*, 2002.

[2] T. Furukawa, G. Dissanayake, and H. F. Durrant-Whyte, "Time-optimal cooperative control of multiple robot vehicles," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2003.

[3] W. Burgard, M. Moors, C. Stachniss, and F. Schneider, "Coordinated multi-robot exploration," *IEEE Trans. on Robotics*, vol. 21, no. 3, pp. 376–386, June 2005.

[4] S. Waslader, G. Inalhan, and C. J. Tomlin, "Decentralized optimisation via nash bargaining," in *Proc. Conf. on Cooperative Control and Optimisation*, 2003.

[5] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, 1989.

[6] G. M. Mathews, "Communication requirements for asynchronous distributed gradient based optimisation algorithms," Centre for Autonomous Systems, Tech. Rep., 2006. [Online]. Available: http://www.cas.edu.au

[7] G. M. Baudet, "Asynchronous iterative methods for multiprocessors," *Journal of the ACM*, vol. 25, no. 2, pp. 226–244, April 1978.

[8] G. M. Mathews and H. F. Durrant-Whyte, "Scalable decentralised control for muli-platform reconnaissance and information gathering tasks," in *Proc. Int. Conf. on Information Fusion*, 2006.

[9] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

[10] B. Grocholsky, "Information-theoretic control of multiple sensor platforms," Ph.D. dissertation, Univ. of Sydney, Australia, 2002.

[11] P. S. Maybeck, *Stochastic models, estimation and control*. Academic Press, New York, 1979-1982.

[12] F. Bourgault, T. Furukawa, and H. F. Durrant-Whyte, "Decentralized bayesian negotiation for cooperative search," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2004.