An Analysis and Design Methodology for Belief Sharing in Large Groups

Robin Glinton, Paul Scerri, David Scerri and Katia Sycara
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA

***Abstract*—Many applications require that a group of agents share a coherent distributed picture of the world given communication constraints. This paper describes an analysis and design methodology for coordination algorithms for extremely large groups of agents maintaining a distributed belief. This design methodology creates a probability distribution which relates global properties of the system to agent interaction dynamics using the tools of statistical mechanics. Using this probability distribution we show that this system undergoes a rapid phase transition between low divergence and high divergence in the distributed belief at a critical value of system *temperature*. We also show empirically that at the critical system temperature the number of messages passed and belief divergence between agents is optimal. Finally, we use this fact to develop an algorithm using system temperature as a local decision parameter for an agent.**

I. Introduction

The increasing availability of cheap, low power sensors and processors has made large scale distributed networks of coordinated autonomous systems appropriate for a range of domains including distributed tracking [11], force protection [3], and locating lost hikers[5]. In such systems agents must act cooperatively in an incompletely observable environment. To do so effectively it is critical for agents to leverage their joint sensing ability to create an accurate view of the world. Creating joint beliefs about the world is highly communication intensive, since many individual sensor readings, over time, go toward creating a view of the world. However, in all practical systems the available communication bandwidth is strictly limited.

Most approaches to distributed belief fusion, including the one presented here, allow system scalability by limiting the communication by an agent to a small number of neighbors. Joint belief is achieved when information diffuses through the system via many agent to agent interactions[8]. There are many systems both natural and artificial that operate analagously e.g. ferromagnetism in a metal [9]. Ferromagnetism, or spontaneous magnetization, occurs when all of the molecules in the metal exist in the same state or *spin*, effectively forming one large cluster of molecules with the same spin. However, the spin of a given molecule has a direct effect only on it's nearest neighbor, therefore the clustering is a result of the diffusion of *spin* information between neighboring molecules. Another example can be found in sociology where segregation, a form of clustering, occurs in a population as a result of the sharing of culture between individuals [14]. Our hypothesis is that similar system dynamics will occur in belief sharing networks and can exploited for the design of belief sharing policies.

Specifically, it is typical for such systems to exist in distinct states or *phases* dependent upon a critical value of a system property. This system property, traditionally called *temperature* because the phenomenon was first discovered in physics, is a scale factor for the probability of interaction between individuals. Below the critical temperature T_c, clustering and order is observed and information diffuses between individuals. Conversely, above T_c most individuals exist in distinct states and information flow stagnates. Furthermore, at T_c such systems become ordered with the minimal interaction between individuals. Striking similarities in the dynamics of systems where information diffuses as a result of nearest neighbor interactions have been shown through studies in a variety of fields [7, 14, 9]. The similarities in dynamics between the aforementioned systems and a large system of agents maintaining a distributed belief suggests that order/disorder phase transitions likely occur when agents use randomized information sharing policies.

For complex systems consisting of large numbers of interacting individuals or particles, the Potts model has been extensively used to statistically model the relationship between local interactions of individuals to global system properties[12]. A Potts model of a belief sharing system allows us to analytically determine randomized policies of information exchange for individual agents which are optimal in terms of minimizing both the number of messages exchanged and degree of global belief divergence across the team. It does this by allowing us to analytically determine critical values of control parameters which govern information exchange and at which an order/disorder phase transition occurs. At this critical value the number of messages exchanged and global divergence are minimized. Thus, by determining this value we can elicit system dynamics that both minimize information divergence and minimize communication.

Characterizing a system of belief sharing agents using a Potts model allows us to analytically relate the degree of allowable global divergence in beliefs across the team to policies individuals use to control interaction with their neighbors. This is important because there are applications where complete homogeneity in the beliefs of individuals is undesirable and in fact sub-optimal. For example, in a tracking application, spatially distant sensors, might only need to know that a target is present in a distant sector but not its precise location. Thus, the model provides a powerful mechanism not available in previous approaches, specifically allowing a system to be tuned to minimize communication while achieving different divergence levels.

For a system of belief sharing agents, the Potts model allows us to study the relationship between the communication network topology and the optimal operating point for local interactions between agents. This allows us to design the optimal randomized information sharing policy for an application with a given network topology. Furthermore, for applications where we are allowed to choose the network topology we can use our model to pick the one that is the best fit based on local interaction constraints.

The remainder of this paper is organized as follows. Section II formally describes the problem being addressed. Section III introduces Potts model, describing its key functions and properties. Section IV maps the belief sharing problem to Potts model, then section V describes the resulting algorithm followed by the agents. Section VI shows empirical results, illustrating that a critical point does indeed exist for the belief sharing system and comparing performance against other approaches. Sections VII and VIII give related work, conclusions and future work.

II. Problem Statement

This section formally describes the problem addressed by this paper. Agents $A = \{a_1, \ldots, a_m\}$ are a team with a joint objective in a partially observable domain. Decisions about actions by the agents are based on state variables $X(t) = \{x_1(t), \ldots, x_n(t)\}$.

The agents have uncertain sensors, thus via some filter they must determine the probability of each of the state variables. Agent a_i's probability distribution over X at time t is $P^i(X(t), t)$. For convenience, define $p^i(x_j(t), t)$.

The performance of the team will be adversely affected whenever their estimate of the state of environment varies from the actual state of the environment. The information difference (KL-divergence or similar) is $\Delta^i(X, P^i(X(t), t))$. The bigger this number, the higher the divergence. However, depending on their current activities, individual agents will not be equally effected by divergence in all variables. In general, they will only need to know precisely some values, while others can be coarsely understood or not known at all. Specifically, the cost of $\delta^i(\bullet)$ divergence to an agent a_i at a particular time is: $c(a_i, \delta^i(\bullet)) \rightarrow \mathcal{R}$.

Using their sensors, agents take sensor readings $r \in R$. A sensor reading influences $P^i(X(t), t)$ via some filter f, $P^{i'}(X(t), t) = f(P^i(X(t), t), r)$. The only assumption made about the filter is that estimates of variables improve monotonically with more sensor readings. Using the cost of information divergence and filter equations, the value of that sensor reading to a_i is $\hat{v}(s, a_i) = C(a_i, \Delta^i(X, P^i(X(t), t)), t) - C(a_i, \Delta^{i'}(X, P^i(X(t), t)), t)$, i.e., the change in cost. We assume $\hat{v}(s, a) \geq 0$. The value of s to the whole team is: $\hat{V}(s) = \sum_{a \in A} \hat{v}(s, a)$.

Thus, the overall optimization function is to minimize:

$$\sum_{a_i \in A} \int_{t=0}^{T} C(a_i, \Delta^i(\bullet)).dt + CommCost$$

where $CommCost$ is the cost of communication. Intuitively, the team should receive reward by sharing the sensor reading unless the added value exceeds the communication cost.

III. Pott's Model

A Pott's model can be defined for any system with the following characteristics:

- A group of individuals or particles, related through a neighborhood system.
- A finite set of states common to all individuals in which each individual can exist.
- Some type of interaction between neighboring individuals or particles which encourages homogeneity in the states in which they exist.

We can represent such a system formally as a graph. Define a graph $G = (V, E)$, where $V = \{v_1, \ldots, v_n\}$ is the set of vertices. These represent the individuals in the system under study. The set of edges $E = \{e_1, \ldots, e_m\}$, determine which individuals are neighbors. Each v_i can take on a label from the finite set $Q = \{q_1, \ldots, q_l\}$. The q_i correspond to the states that individuals can exist in.

For example in the case of ferromagnetism, each molecule has a certain spin. When the spins of neighboring molecules are different a magnetic force is exerted between them which encourages the molecules to switch to the same spin. In the case of belief sharing neighboring agents can share information. A strong impetus to share information, like a magnetic force between them, encourages neighboring agents to have the same belief.

Defining a Pott's model for a system with the aforementioned properties consists of defining an energy function $d(q_i, q_j)$ which gives the strength of the pull between two neighbors existing in states q_i and q_j.

The higher the energy, the greater the pull. The total energy for the system is given by the Hamiltonian function H which sums the energy between all pairs of neighboring individuals. At higher values of H the system is more likely to be disordered, i.e., there is likely to be great heterogeneity in the states in which individuals exist. This is because with neighbors constantly pulling on each other there is great impetus for any given individual to continually switch states. Conversely at lower values of H there is more likely to be homogeneity in the states in which individuals exist accross the system.

We write the set of all possible assignments of the states in Q to the vertices in V as S. Formally, $S = Q \times V$. Recall that $d(q_i, q_j)$ is a function which gives the energy that exists in a graph edge $e \in E$ which connects individuals occupying states q_i and q_j. Let $s \in S$ be a vector which gives a particular labelling of the individuals $v \in V$ in terms of states $q \in Q$. Given these definitions the total energy for the system with a particular assignment of states to individuals s is given by the Hamiltonian below:

$$H(s) = \sum_E d(q_i, q_j)$$

Using a Hamiltonian as defined above, the Potts formalism gives us the probability of the system existing with a particular assignment s of states to individuals. This distribution is a function of the Hamiltonian and system temperature T. The system temperature T serves as a scaling parameter for the distribution which is proportional to its entropy. Intuitively T dictates how sensitive an individual is to the pull of its neighbors. At higher values of T and individual is more likely to switch states in response to a strong pull from its neighbors. The critical temperature T_c is the temperature below which the system is most likely to have minimal energy. That is when most individuals in the system are most likely to occupy the same state.

Within the Potts formulation, the probability of the system existing with a given assignment of states to individuals $s \in S$ is given by:

$$P(s) = \frac{e^{-\frac{H(s)}{kT}}}{\sum_{s \in S} e^{-\frac{H(s)}{kT}}} \tag{1}$$

where k is Boltsmann's constant.

If we could evaluate Equation 1 exactly we could answer questions about the system by calculating expectations over system variables. Most important to this work, we could find the critical system temperature T_c in the following way. Recall that T_c is the temperature below which the system energy is minimal. Using Equation 1 we could calculate the expected value of H, the total energy of the system as $E(H) = \sum_s P(s)H(s)$ for a range of values of T. T_c would then be the value of T which corresponds with the minimal value of H on the resulting graph of H versus T.

However, calculating the denominator of Equation 1, which normalizes the distribution, is generally undecidable for even the smallest sytems. This is because the computational complexity of evaluating the denominator, commonly known as the partition function, is exponential in the number of possible ways to assign states to individuals. For this reason we turn to the Metropolis algorithm [10] which allows us to calculate expectations over H using samples of states s drawn according to the distribution $P(s)$.

An example of a Pott's model where each individual can exist in one of three states and a neighborhood system defined by a two dimensional lattice is show in Figure 1. This example is intended to give an intuition for how the model works. The bond energy in this case is simply a constant for bonds that link individuals occupying the same states and zero when the two states differ. In this example the divergence metric is given by, $d_{q_i,q_j} = \delta_{q_i,q_j}$ where:

$$\delta_{q_i,q_j} = \begin{cases} 1 & q_i = q_j; \\ 0 & q_i \neq q_j \end{cases}$$

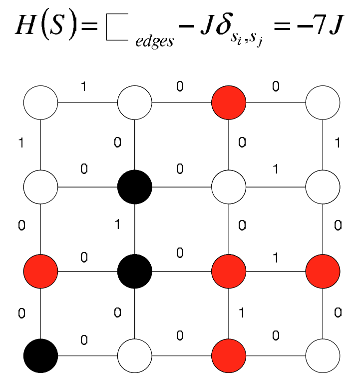

Fig. 1. Hamiltonian for a three state Pott's model with a neighborhood system defined by a two dimensional lattice. States of individuals are indicated by the color of the corresponding site in the lattice. J is a constant which gives the energy in a bond between individuals occupying the same state.

IV. Belief Sharing Mapped to Pott's Model

In this Section, we describe formally how the Pott's model maps to the belief sharing problem of interest here.

N agents maintain a distributed belief. There is a communication network defined by the set E where $\exists e_{i,j} \in E$ if a communication link exists between agents i and j The belief of agent i is represented by a vector $X^i(t) = \{x_1^i(t), \ldots, x_n^i(t)\}$ where $x_j^i(t) \in \mathcal{R}$. Ground truth is given by $X^g = \{x_1^g, \ldots, x_n^g\}$. At time t agent i takes readings $R^i(t) = \{r_1^i(t), r_2^i(t), \ldots, r_n^i(t)\}$ readings are generated by adding zero mean gaussian noise to the groundtruth according to $r_j(t) = x_i^g + N(0, \sigma_k)$. Each agent i maintains a filter f that is simply a maximum likelihood estimate of $x_j^i(t)$. That is $x_j^i(t) = \sum_{t_0}^{t} r_j^i(t)$

As a measure of information divergence we use the mean square error MSE as a metric. Formally the divergence between the belief of agent i and groundtruth is given by $\sum_{j=1}^n (x_j^i(t) - x_j^g(t))^2$ while the divergence between the beliefs of agents i and k is given by $\sum_{j=1}^n (x_j^i(t) - x_j^k(t))^2$ As a simple measure of the divergence between the beliefs of two agents i and j we use $d(i,j) = \sum_{j=1}^n (x_j^i(t) - x_j^k(t))^2$. The value $d(i,j)$ maps to $d(q_i, q_j)$ the energy in the link between neighbors as defined in Section III. Using this measure of divergence the Hamiltonian of the belief share system is given by:

$$H = \sum_{i,j} d(i,j)\delta_{i,j}$$

where

$$\delta_{i,j} = \begin{cases} 1 & \exists e_{i,j} \in E, \\ 0 & \nexists e_{i,j} \in E \end{cases}$$

and E is the set of communication links between agents.

Consequently the probability $P(H)$ of the distributed belief having total divergence or energy as defined by H above is given by:

$$P(H) = \frac{e^{-\frac{H}{kT}}}{\sum_{s \in S} e^{-\frac{H}{kT}}}$$

Inserting the belief sharing H into this function we get:

$$P(H) = \frac{e^{-\frac{\sum_{i,j} d(i,j)\delta_{i,j}}{kT}}}{\sum_{s \in S} e^{-\frac{\sum_{i,j} d(i,j)\delta_{i,j}}{kT}}}$$

V. Algorithm

T_c can be used in a belief sharing algorithm by aping the operation of the Metropolis algorithm [10] while calculating the partition function for the Potts model of the agent system. The Metropolis algorithm approximates the partition function by sampling states of the system in proportion to $P(s_i)$, the probability that the system will occupy a point $s \in S$ where S is the space of possible system configurations. These samples are then used to calculate the partition function. The algorithm achieves this by randomly choosing an agent, changing it's state, and accepting the change randomly weighted by the change in system energy. An agent can detect the change in system energy resulting from it's change in state locally since the total system energy is given by the sum of the bond energy between nearest neighbor agents.

Algorithm 1: Algorithm used by agent to decide when and with whom to share information

BeliefShare($\beta, nlist, readings$)
while true
  $i \leftarrow$ RandomInteger(Size($nlist$))
  $n \leftarrow nlist.get(i)$
  $n.belief =$ RequestBelief(n)
  $d \leftarrow$ Divergence($n.belief, belief$)
  $r \leftarrow$ UniformRandom($[0, 1]$)
  $p \leftarrow$ Exp($-\beta * d$)
  if $r < p$
    Send($readings, n$)

VI. Experiments

This section presents experimental results using the agent system, working according to the Potts model. For the experiments we have used a simulator which allows us to vary different variables in order to examine their effect. Each agent in the simulator has its own belief which is created from a Gaussian model. The simulator has four different, parameterizable network implementations. Communication between agents is simulated to allow many experiments to be performed.

In each of the graphs the model was run for 1000 phases for each value of kT and the resultant energy averaged over 10 runs. Figures 2 - 5 show the energy on the Y-axis and log(kT) on the X-axis.

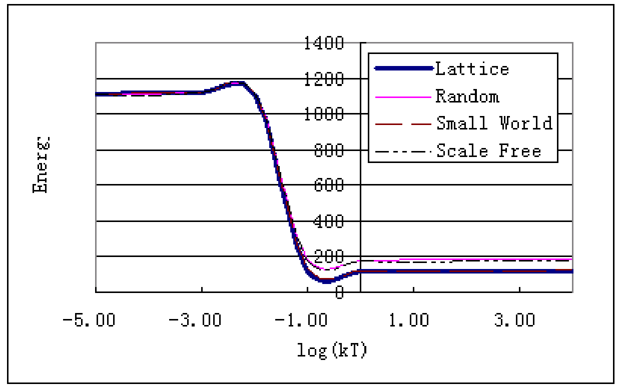

Fig. 2. Comparing performance across network types.

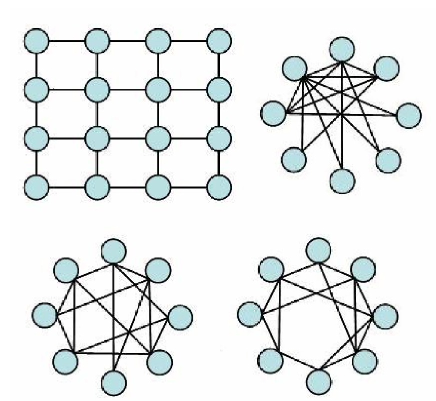

Fig. 3. Examples of the four various network types. Clockwise from top left: Lattice, Scale Free, Small World, and Random.

Figure 2 shows the results of four different network implementations. A simple lattice network was implemented whereby each agent was able to exchange beliefs with each of its 8 neighbors. A random world network was implemented by randomly pairing each agent with two other agents. The small world network involved pairing agents with close agents then randomly repairing some agents with those further away[16]. The scale free network consists of several hubs, with connections to many agents[17]. The various network types are shown in Figure 3. The key feature of this graph is that in each case there was a clear and large phase shift, occurring at the same critical point for each network. This supports the key hypothesis of this paper that the system dynamics of a Potts model apply to a belief sharing problem. The networks had close to the same performance, but both the lattice and small-worlds networks had lower divergence after the critical point, probably due to the clustering in these network types.

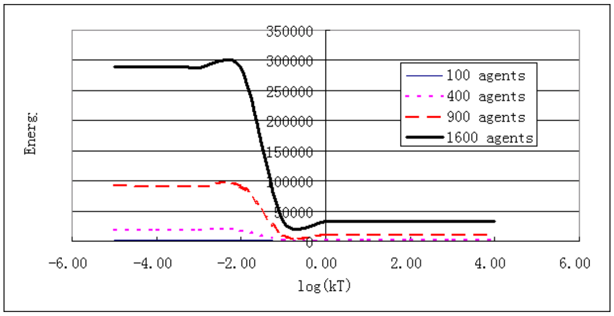

Fig. 4. Performance as noise is varied.

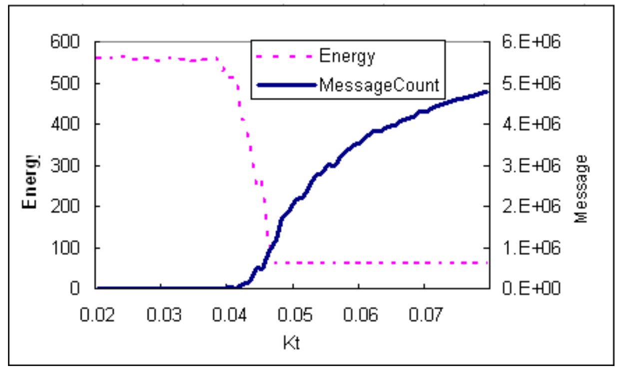

Fig. 6. System energy and messages as the temperature is varied.

Figure 4 shows the result of varying the standard deviation of the Gaussian model from which the beliefs were derived. Values of 1.0, 2.0, 4.0, 8.0 and 16.0 were used. As the noise is varied, the critical point varies, with lower noise having a lower critical temperature. This is perhaps not surprising, since one would expect that if the beliefs start closer to being the same, less interaction is required to make them the same. Importantly, even with low noise, a phase transition is still observed.

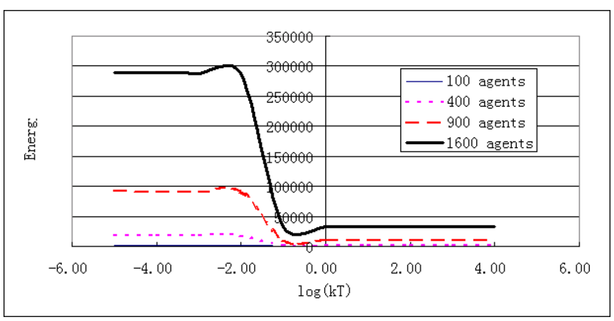

Fig. 5. Performance as team size is varied.

Figure 5 shows the results of varying the number of agents in the team. Team sizes of 100, 400, 900 and 1600 were used. The critical point, where the phase transition occurs, is at the same location for each of the team sizes. This may be somewhat surprising, since one might expect bigger teams to require more interaction to reduce divergence, but the local nature of the interactions reduces the impact of team size.

Figure 6 shows the number of messages between agents next to the resultant energy. The left Y axis shows the energy and the right Y axis shows the number of messages after 1000 phases. The x axis shows the change in kT. It is clear that the phase transition occurs at a value where a substantial number of messages are sent. This is important because predicting this number of messages, about 100,000 messages in the figure, has not been possible to estimate in the past. Even the ability to control the number of messages (and thus get this optimal behavior) has not been previously possible.

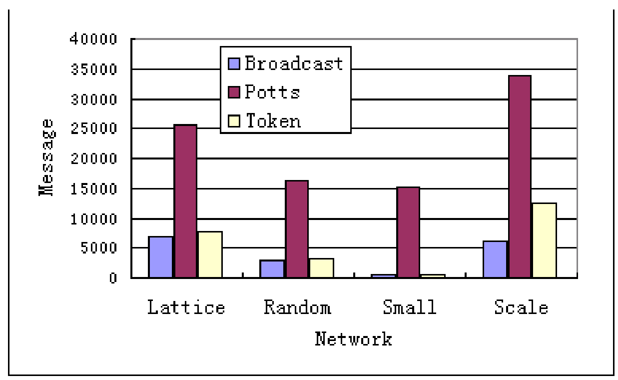

Fig. 7. Comparison with other approaches.

Figure 7 shows how the Potts derived model compares to a token based model[18] and a broadcast model. The token model was implemented by creating five tokens which contained a belief and a visit counter. The agent averaged his belief with the belief contained in the token, incremented the visit counter and passed the token on to a random neighbor. The broadcast model consisted of random agents 'broadcasting' their belief to 20% percent of the population. Agents who received a broadcast belief averaged it with their own. Notice that the Potts derived model uses more messages and is thus less efficient than the other two models. This is not because the Potts model is not optimizing performance, it is because of the constraints we arbitrarily imposed on our interaction model, constraints not imposed on the other algorithms. Relaxing these constraints, specifically giving the Potts model a richer interaction model would allow it to, presumably, achieve the same level of performance.

VII. Related Work

There has been recent interest in the use of decentralized Bayesian filters such as the ones proposed in [6, 2] to manage beliefs over a large team. Communicating these beliefs, however, is expensive, prompting several selective communications approaches. Divergence metrics such as Hellinger affinity and KL-divergence are commonly used to measure the information gain of individual communications. However, existing methods of integration such as channel managers [1] or query-based particle filters [13] face scaling issues, in these cases, dealing with redundant data and polynomial-time scaling with team size, respectively. Decision theoretic approaches have also been successfully used for small teams [15]. Consensus protocols over communication networks are another method to achieve consistent belief among agents. Multi-hop relay protocols have been demonstrated to allow for fast consensus seeking [19], but previous work has focuses on trading robustness and convergence, rather than information gain and overall network traffic

Uses of the Potts model in the literature include statistical models of ferromagnetism in metals[9], the explanation of the phenomenon of segregation in socio-economic models [14], and the study of political districting [4].

VIII. Conclusions and Future Work

This paper presented a novel approach to a key problem for distributed systems. Specifically, the Potts model from statistical mechanics was applied to the problem of analysis and design of interaction protocols for maintaining shared belief. Our hypothesis was that a critical *temperature* exists at which a phase transition occurs and belief divergence and communication are both minimized. Empirical results show that this was indeed the case, under a variety of circumstances.

However, this paper represents only a first step. The paper showed that Potts model is applicable to the design and analysis of multiagent systems, but did not leverage the huge range of analytic tools that have been built on Potts work. It is the application of these tools that will fully realize the potential of the approach presented here. One specific direction for future work is to relax the constraints on the system imposed in this paper, thus giving the Metropolis algorithm broader range in optimizing performance. We anticipate that this will lead to better performance than any competing algorithms. [1]

References

[1] F. Bourgault and H Durrant-Whyte. Communication in general decentralized filter and the coordinated search strategy. In *Proc. of FUSION'04*, 2004.

[2] F. Bourgault, T. Furukawa, and H Durrant-Whyte. Decentralized bayesian negotiation for cooperative search. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2004.

[3] L. Chaimowicz and V. Kumar. Aerial shepherds: Coordination among uavs and swarms of robots. In *in 7th International Symposium on Distributed Autonomous Robotic Systems*, 2004.

[4] C. Chou and S. Li. Taming the gerrymander, statistical physics approach to political districting problem. *Physica A Statistical Mechanics and its Applications*, 369:799–808, September 2006.

[5] J. L. Drury, J. Richer, N. Rackliffe, and M. A. Goodrich. Comparing situation awareness for two unmanned aerial vehicle human interface approaches. In *in Proceedings of the IEEE International Workshop on Safety, Security and Rescue Robotics*, 2006.

[6] B. Grocholsky. *Information-Theoretic Control of Multiple Sensor Platforms*. PhD thesis, The University of Sydney, 2002. Available from http://www.acfr.usyd.edu.au.

[7] Katsutoshi Hirayama, Makoto Yokoo, and Katia P. Sycara. The phase transition in distributed constraint satisfaction problems: First results. In *Proceedings of the 6th International Conference on Principles and Practice of Constraint Programming*, pages 515–519, 2000.

[8] Zhipu Jin. *Coordinated control for networked multi-agent systems*. PhD thesis, Electrical Engineering, California Institute of Technology, 2006.

[9] Epele L.N., Fanchiotti H., and C.A. García Canal. Ferromagnetism and antiferromagnetism in the Potts model in $(1 + 1)$ and $(2 + 1)$ dimensions. *Physical Review B*, 25:1997–1999, February 1982.

[10] N.A. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

[11] C. L. Ortiz, R. Vincent, and B. Morisset. Task inference and distributed task management in centibots robotic systems. In *AAMAS*, 2005.

[12] R. Potts. Some generalized order-disorder transformations. *Proceedings of the Cambridge Philosophical Society*, 1952.

[13] M. Rosencrantz, G. Gordon, and S. Thrun. Decentralized sensor fusion with distributed particle filters, 2003.

[14] T.S. Schelling. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1:143–186, 1971.

[15] M. Tambe. Agent architectures for flexible, practical teamwork. In *National Conference on AI (AAAI97)*, pages 22–28, 1997.

[16] Duncan Watts and Steven Strogatz. Collective dynamics of small world networks. *Nature*, 393:440–442, 1998.

[17] Yuri I. Wolf and Georgy Karevand Eugene V. Koonin. Scale-free networks in biology: new insights into the fundamentals of evolution? *BioEssays*, 2002.

[18] Y. Xu, P. Scerri, B. Yu, S. Okamoto, M. Lewis, and K. Sycara. An integrated token-based algorithm for

[1] This research has been sponsored in part by AFOSR FA9550-07-1-0039, AFOSR FA9620-01-1-0542, L3-Communications (4500257512) and NSF ITR IIS-0205526

scalable coordination. In *AAMAS'05*, 2005.

[19] Jin Zhipu. *Coordinated control for networked multi-agent systems*. PhD thesis, California Insitute of Technology, 2006.