

An Analysis and Application of Distributed Constraint Satisfaction and Optimization Algorithms in Sensor Networks*

Weixiong Zhang^{*}, Zhao Xing, Guandong Wang and Lars Wittenburg
Department of Computer Science and Engineering
Washington University in St. Louis
St. Louis, MO 63130, U.S.A.

ABSTRACT

This research is motivated by a distributed scheduling problem in distributed sensor networks, in which computational resources are scarce. To cope with limited computational resources and restricted real-time requirement, it is imperative to apply distributed algorithms that have low overhead on resource requirement and high anytime performance. In this paper, We study distributed stochastic algorithm (DSA) and distributed breakout algorithm (DBA), two distributed algorithms developed earlier for distributed constraint satisfaction problems. We experimentally investigate their properties and compare their performance using our distributed scheduling problem as a benchmark. We first formulate the scheduling problem as a distributed multi-coloring problem. We then experimentally show that the solution quality and communication cost of DSA exhibit phase-transition or threshold behavior, in that the performance degenerates abruptly and dramatically when the degree of parallel executions of distributed agents increases beyond some critical value. The results show that when controlled properly, DSA is superior to DBA, having better or competitive solution quality and significantly smaller communication cost than DBA. Therefore, DSA is the algorithm of choice for our distributed scan scheduling problem.

Categories and Subject Descriptors

I.2.11 [Computing Methodologies]: Artificial Intelligence - Distributed AI; I.2.8 [Computing Methodologies]: Artificial Intelligence - Problem Solving, Control Methods, and Search

General Terms

Algorithms, Performance, Experimentation

*This research was funded in part by NSF Grants IIS-0196057 and IET-0111386, and in part by DARPA Cooperative Agreements F30602-00-2-0531 and F33615-01-C-1897. Thanks to Stephen Fitzpatrick for many discussions related to this research.

*Corresponding author. Email: zhang@cse.wustl.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'03, July 14–18, 2003, Melbourne, Australia.
Copyright 2003 ACM 1-58113-683-8/03/0007 ...\$5.00.

Keywords

Sensor networks, distributed constraint satisfaction and optimization, distributed coloring, distributed breakout, distributed stochastic search, object tracking, phase transitions

1. INTRODUCTION

We have been developing a large-scale, wireless sensor network controlled by a distributed multi-agent system. We place an agent atop of a sensor, which runs on a battery and has, among other things, a slow CPU, a small memory and a wireless communication unit. The sensors/agents need to be programmed in such a way that they can collaboratively carry out a task, which no single sensor/agent can do on its own. A typical application of our system is mobile object tracking in dynamic and real-time environments, where at least three sensors are required to sense at the same time in order to triangulate the location of a target.

Before the sensors start to track a mobile object, they need to detect the object. The sensors have to scan the areas that they monitor. The sensing actions of the sensors must be scheduled in such a way that the new objects can be detected as quickly as possible with as little energy usage as possible. Overall, object detection and tracking operations must be interleaved; when a sensor is not involved in a tracking task, it will scan its areas for detecting new objects. As to be described in the next section, the object detection using multiple sensors is scan scheduling problem can be cast as a distributed constraint satisfaction problem (CSP) or constraint optimization problem (COP).

Due to the real-time nature of the application and limited computation and communication resources on the physical devices, the key to large-scale, real-time sensor/agent networks are the mechanisms that can enable agents to make viable distributed decisions in restricted time with limited computation and communication resources. In this research, we are particularly interested in *low-overhead methods* for solving distributed CSPs/COPs. By *low-overhead methods* we mean those in which limited communication among agents is needed and agents make their decisions without global knowledge of the system and without knowing the actions of the other agents.

In this paper, we study distributed stochastic algorithm (DSA) [1, 3, 8] and distributed breakout algorithm (DBA) [4, 6, 7, 9], two existing algorithms for distributed constraint problems, in the context of solving our scan scheduling problem. We consider these algorithms as low-overhead methods because they indeed require limited communications and use simple mechanisms to search, distributedly, for global solutions. The main differences between these

two approaches are that DBA introduces priorities for conflict resolution among agents' actions, while DSA uses randomness to possibly avoid conflicts.

We apply DSA and DBA to our distributed scan scheduling problem and evaluate their performance. In particular, we investigate their characteristics and compare their performance on this application. The objective is to identify the right algorithm and its appropriate parameters for the application.

A detailed account of the scan scheduling problem is given in Section 2. We formulate the problem as a weighted multi-coloring problem, in which a node may have multiple colors, in Section 3. We then briefly describe DBA and DSA in Section 4. We also discuss how to extend DBA to handle weighted constraints. As reported in our previous research, DSA exhibits phase-transition or threshold behavior if its parameter is not controlled properly when applied to regular graph coloring problems. To choose the right parameters to avoid DSA's degraded performance on our scan scheduling problem, we reinvestigate the phase-transition issue on multi-coloring problems in Section 5. We compare the performance of DSA and DBA at their plausible equilibrium states as well as their anytime performance in Section 6. We evaluate their communication cost at Section 7. In Section 8, we evaluate DSA and DBA for finding the shortest scan schedule of quality exceeding a threshold. Finally, we conclude in Section 9.

2. DISTRIBUTED SCAN SCHEDULING

Detecting and tracking mobile objects in large open environments is an important topic that has many real applications in different domains, such as avionics, surveillance and robot navigation. We are developing such an object detecting and tracking system using MEMS sensors, each of which is operated under restricted energy sources, i.e., batteries, and has a small amount of memory and restricted computational power. In a typical application of our system, a collection of small Doppler sensors are scattered in an open area to detect possible foreign objects moving into the region. Each sensor can scan and detect an object within a fixed radius. However, the overall detecting area of a sensor is divided into three equal sectors, and the sensor can only operate in one sector at any given time. The sensors can communicate with one another through radio communication. The radio channel is not reliable, however, in that a signal may get lost due to, for instance, a collision of signals from multiple sensors, or distorted due to environment noises. Moreover, switching from one scanning sector to another sector and sending and receiving radio signals take time and energy. To save energy, a sensor may turn itself off occasionally or periodically if doing so does not reduce system performance.

The overall system is operated to achieve two conflicting goals, to detect foreign objects as many and quickly as possible and, at the same time, to preserve energy as much as possible so as to prolong the system's lifetime. The overall problem is thus to find a distributed scan schedule to optimize an objective function that balances the above two conflicting goals.

One of our goals here is to develop a scalable sensor network for object detection in that the system response time for detecting a new object does not degenerate when the system size increases. To this end, we place an agent atop each sensor and make the agent as autonomous as possible. This means that an agent will not rely on information of the whole system, rather the local information including its neighboring sensors or agents. These agents also try to avoid unnecessary communication so as to minimize system response time. Therefore, complex coordination and reasoning methods are not suitable.

We set forth to apply simple, low-overhead methods for solving

this distributed scheduling problem. We assume that the sensors are stationary, in that they will not move after being placed in an operation environment. We consider two distributed algorithms, distributed breakout and distributed stochastic search, that were developed earlier to solve distributed constraint satisfaction problems. Our objectives are twofold. First, we want to understand how these two algorithms compare with each other on a real-world application such as our scan scheduling problem. Second, we want to identify the right application condition and parameters of the algorithm that we will use based on the evaluation.

3. MODEL IN MULTIPLE-COLORING

Our first step to address the distributed scan scheduling problem is to model it as distributed constraint satisfaction or optimization problem. Indeed, the problem can be captured as a distributed multi-coloring problem.

To detect objects quickly, a sensor needs to scan its three sectors as often as possible. However, to save energy, two neighboring sensors should try to avoid scanning a common area covered by two overlapping sensor sectors at the same time since one sensor's sensing of an area is sufficient to detect possible objects in the area. The objective is to find a sequence of sensing actions so that each sensor sector can be scanned at least once within a minimal period of time. In other words, the objective is a cyclic scan schedule in which each sector is scanned at least once with the constraint of minimizing the energy usage.

As all agents execute the same procedure, scanning actions on different sensors take approximately the same amount of time. Assume that the sensors are synchronized¹, an assumption that can be lifted by a synchronization method [5]. Therefore, the problem can be viewed as coloring a weighted graph so that the total weight of violated constraints is minimized. Here, a node corresponds to a sensor or agent, a link between two nodes represents the constraint of a shared region between two agents, and the weight captures the area of a common region. The larger a shared region is, the more energy will be wasted if two sensors scan the shared region at the same. Moreover, each color corresponds to a time slot in which a sector is scanned and the total number of colors represents the scanning cycle length or the number of time slots required to scan each of the sectors of all sensors. A node must have at least one color so that the corresponding sector is scanned at least once, and may also have multiple colors so that the sector may be sensed multiple times to increase system's object detection rate.

We further consider a sensor as a hypernode of three nodes, each of which corresponds to a sector of the sensor. Within an hypernode, two nodes cannot share a color since two sectors cannot be scanned at the same time. In other words, the colors assigned to the nodes within an hypernode are mutually exclusive. This is a hard constraint that cannot be violated. Furthermore, not all available colors must be used as a sensor may sometimes turn itself off. A node within a hypernode may also be constrained by a node of another hypernode if the two corresponding sectors share a common region. This constraint may be violated, resulting in wasted energy. Overall, we are looking for such a coloring that satisfies all hard constraints within hypernodes and minimizes the weights of violated constraints among hypernodes so as to reduce the overall conflicts of scanning common regions among overlapping sectors.

We can model the distributed scan scheduling problem as fol-

¹The algorithms we will consider shortly, DSA and DBA, are synchronous. It has been observed that under certain conditions asynchronous actions may actually improve the performance of a distributed algorithm [2].

lows. Let $G(V, E)$ be an undirected graph with a set of hypernodes V and a set of edges E . Each hypernode $v \in V$, which represents a sensor, consists of k nodes, v_1, v_2, \dots, v_k , which model k sectors of the sensor. (In our application, we have $k = 3$.) An edge between nodes $u_i \in u$ and $v_j \in v$ is denoted as $(u_i, v_j) \in E$, and u_i and v_j are called neighbors. Two hypernodes are neighbors if any pair of their nodes are neighbors. The weight of an edge (u_i, v_j) , denoted as $w(u_i, v_j)$, is the area of a common region covered by sectors of sensors u and v . Every hypernode is given a total T available colors that can be used and a sensor activation ratio $\alpha \leq 1$. $|T|$ corresponds to a schedule length. α captures the frequency, how often a sensor should be active, and $1 - \alpha$ is the frequency determining how often a sensor should turn itself off to save energy. Given the total available colors T , the weighted multi-coloring problem is to color the nodes following these criteria. First, exactly $\lceil \alpha T \rceil$ colors are used within a hypernode. This means that a node may have more than one color if the available colors are more than the nodes in a hypernode, i.e., $\lceil \alpha T \rceil > k$. Second, every node must have at least one color and no two nodes within a hypernode can share a color. Third, minimizing conflicts between pairs of nodes in different hypernodes, so the total weight of violations going across hypernodes is minimized. When no conflict is allowed or a minimal level of allowed conflicts is specified, the overall problem is to find the minimal number of allowed colors T and such a schedule satisfying the criteria.

4. THE ALGORITHMS

The two algorithms we will discuss are synchronous [5], in that all distributed processes proceed in synchronized steps. In each step, a process sends and receives (zero or more) messages and then performs local computations, i.e., changing local state, if necessary.

4.1 Distributed breakout algorithm (DBA)

To understand DBA, we first consider the sequential breakout algorithm [4]. It is a local search method for CSP equipped with an innovative scheme of escaping local minima. Given a CSP, the algorithm first assigns a weight of one to all constraints. It then picks a value for each variable. If no constraint is violated, the algorithm terminates. Otherwise, it chooses a variable that can reduce the total weight of the unsatisfied constraints if its value is changed. If such a weight-reducing variable-value pair exists, the algorithm changes the value of a chosen variable. The algorithm continues to select a variable and change its value until no weight-reducing variable can be found. At that point, it reaches a local minimum if a constraint violation still exists. Instead of restarting from another random initial assignment, the algorithm tries to escape from the local minimum by increasing the weights of all violated constraints by one and proceeds as before. This weight change will force the algorithm to alter the values of some variables to satisfy the violated constraints.

Centralized breakout can be extended to synchronized distributed breakout algorithm (DBA) [6, 7]. Without loss of generality, we assign an agent to a variable, and assume that all agents have unique identifiers. Two agents are *neighbors* if they share a common constraint. An agent communicates only with its neighbors. At each step of DBA, an agent exchanges its current variable value with its neighbors, computes the possible weight reduction if it changes its current value, and decides if it should do so. To avoid simultaneous variable changes at neighboring agents, only the agent having the maximal weight reduction has the right to alter its current value. If ties occur, the agents break the ties based on their identifiers. The above process of DBA is sketched in Algorithm 1.

The DBA algorithm was not designed for weighted constraint

Algorithm 1 Sketch of DBA

```

set the local weights of constraints to one
value ← a random value from domain
while (no termination condition met) do
  exchange value with neighbors
  WR ← BestPossibleWeightReduction()
  send WR to neighbors and collect their WRs
  if (WR > 0) then
    if (it has the biggest improvement among neighbors) then
      value ← the value that gives WR
    end if
  else
    if (no neighbor can improve) then
      increase violated constraints' weights by one
    end if
  end if
end while

```

Algorithm 2 Sketch of DSA, executed by an agent.

```

Randomly choose a value
while (no termination condition is met) do
  if (a new value is assigned) then
    send the new value to neighbors
  end if
  collect neighbors' values, if any; compute the best possible
  conflict reduction  $\Delta$ 
  if ( $\Delta > 0$ ) or ( $\Delta = 0$  but there is a conflict) then
    change to a value giving  $\Delta$  with probability  $p$ 
  end if
end while

```

problems, in which constraints carry weights. We can extend DBA to weighted CSP as follows. A constraint weight is taken as the initial weight of the constraint for the algorithm. Whenever the weight of a constraint needs to be increased in the algorithm, the constraint weight in the original problem is added. This way, the weight increment reflects the actual constraint weight. We use this extended version of DBA in our experiments.

4.2 Distributed Stochastic Algorithm (DSA)

Distributed stochastic algorithm (DSA) is uniform [5], in that all processes are equal and have no identities to distinguish one from another. In contrast, DBA is not uniform since agents have identities or priorities for tie breaking. In general, a uniform algorithm is more general and easier to implement than a non-uniform algorithm. Therefore, if their performances are equivalent under all conditions, DSA should be preferable over DBA.

DSA is a family of algorithms [1, 3, 8]. Its basic idea is simple. With initial random values, agents go through a sequence of steps until a termination condition is met. In each step, an agent sends its current state information, i.e., its variable value in our case, to its neighboring agents if it changed its value in the previous step, and receives the state information from its neighbors. It then decides, randomly, to keep its current value or change to a new one, assuming that the neighbors will maintain their values. The objective for value changing is to possibly reduce violated constraints.

The most critical step of DSA is for an agent to decide the next value, based on its current state and its believed states of the neighboring agents. If the agent cannot find a new value in its domain to improve its current state, it will not change its current value. If there is such a value that improves or maintains state quality, the




Figure 1: Phase-transition behavior of DSA on distributed scan scheduling, $T = 6$.

agent may or may not change to the new value based on a stochastic scheme. Note that the probability p in DSA controls how frequently neighboring agents changing values, so that p is called the degree of parallel executions.

Five specific variations of DSA have been proposed and studied [1, 3, 8], among them DSA-B (or CFP in [3]) is the best on regular graph coloring problems [8]. In this research, we adopt DSA-B for our distributed scheduling, and refer to it as DSA in the rest of the discussion. A sketch of this DSA algorithm is in Algorithm 2.

5. THRESHOLD BEHAVIOR OF DSA

In our previous study, a phase-transition behavior of DSA was observed on regular coloring problems [8]. Here we consider this issue again on distributed, weighted multi-coloring problems generated from our scan scheduling application. The objective is to better understand the behavior of DSA related to the degree of parallel executions among distributed agents.

A phase transition refers to the phenomenon of a system in which a global property changes rapidly and dramatically when an order parameter goes beyond a critical value. A simple example is the melting of ice with rising temperature. When the temperature reaches above the melting point, a phase transition from ice to water occurs. For our problem here, the system property is the average weight of violated soft constraints in DSA solutions normalized by the weight of all soft constraints. The control parameter is the degree of parallel execution p .

In our experiments, we set the sensing radius of a sensor to one unit, and used a square of 10×10 units as the area to be monitored. The number of sensing sectors is set to three to match our hardware system. We randomly and uniformly placed a fixed number of sensors with arbitrary orientations in the square. We then changed the probability p from 0.1 to 0.99, with an increment of 0.01. For each p , we generated 100 instances and converted them into multi-coloring problems. We evaluated the performance of DSA when

it has reached relatively stable states. Specifically, we ran DBA to the point where its performance does not change significantly from one step to the next. On all network sizes we considered, DSA's performance seems to be stabilized after 256 steps. Similar results have been observed after 1024, 2048 and longer steps. In the rest of this paper, we report the results at 256 steps.

In our first experiment, we used 500 sensors and six allowed colors, i.e., $T = 6$. We set the sensor activation ratio $\alpha = 2/3$ so that that four colors will be used for a hypernode with three nodes in it. Following the hard constraints within a hypernode, i.e., a node must have at least one color and no two nodes can share a color, this means that one node must have two colors. Figure 1(a) shows the average performance of DSA, controlled by the parallelism probability p . The performance is measured by the ratio of the total weight of soft constraints violated in the solutions found by DSA to the weight of all soft constraints. As the figure shows, the performance of DSA experiences a sharp and dramatic transition when the parallelism probability p increases beyond 0.8. One interesting observation from this experiment, as well as other similar ones such as that in Figure 1(b), is that for a fixed sensor density, DSA reaches its best performance before the phase transition occurs.

Moreover, the phase-transition behavior occurs at a higher probability p when the sensor or graph density decreases. One example can be found by comparing the phase transition points in Figures 1(a) and 1(b), where the network for Figures 1(a) has 500 nodes while that used for Figures 1(b) has only 300 nodes. The phase-transition point in Figures 1(a) appears sooner than that in Figures 1(b). This result indicates that the more constrained the underlying graph is, the sooner the phase-transition behavior occurs. It seems to be an important, but difficult, problem to analytically determine the phase transition location given a graph. Nevertheless, our experimental results provide some qualitative information on how the degree of parallel executions should be chosen. The tighter the constraints in a given problem or the higher the connectivity of a constraint graph, the smaller the degree of parallel executions should be. Similar phase-transition patterns were observed on the same sensor networks but with more allowed colors, for instance $T = 18$.

The phase transition result here is qualitatively similar to that in [8] where regular graph coloring was considered. This implies that the threshold behavior of DSA persists in these two problems. However, while on both regular coloring and multi-coloring problems DSA's performance degenerates significantly when its degree of parallelism is too high, the algorithm pans out better on multi-coloring than on regular coloring. On regular graph coloring problems, DSA's degenerated performance may be much worse than that of a random coloring, while on multi-coloring problems the algorithm's degenerated performance is still better than that of a random coloring. For instance, on a 300-node multi-coloring problem, the average normalized solution quality of random multi-coloring is more than 0.65 while the averaged normalized quality of degenerated solution of DSA when $p = 0.99$ is less than 0.16.

6. SOLUTION QUALITY

We now directly compare DBA and DSA on multi-coloring problems. We used the same set of parameters as used for studying phase transitions in Section 5. Specifically, each sensor has three equal sectors, and its sensing radius is set to one unit. Sensors are then randomly and uniformly placed in an area of 10×10 square with arbitrary orientations. We experimented with different values of maximum allowed colors and different sensor activation ratio α . In the rest of this section, we report the results using $T = 6$ and $T = 18$ with $\alpha = 2/3$.

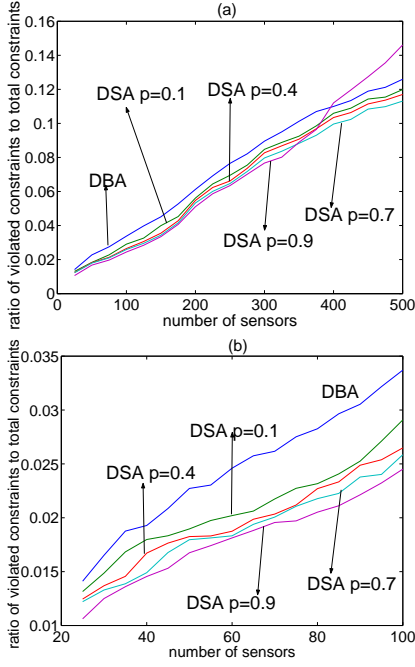


Figure 2: DSA vs. DBA in terms of number of sensors, $T = 6$.

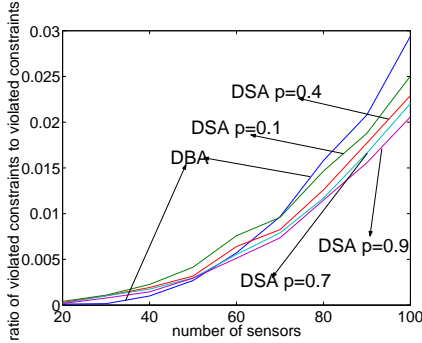


Figure 3: DSA vs. DBA in terms of number of sensors, $T = 18$.

6.1 Solution quality in terms of network sizes

Since one of our ultimate objectives is to choose DBA or DSA to solve our distributed scan scheduling problem, we need to investigate the relationship between the quality of the schedules found by these two algorithms and the properties of the underlying networks. To this end, we experimentally compared DBA and DSA on multi-coloring problems produced from sensor networks of various sizes. We changed the density of multi-coloring graphs by changing the number of sensors N . The solution quality is the total weight of violated soft constraints normalized by the total weight of soft constraints, measured at 256 steps of the algorithms' executions when their performances are relatively stable.

We run DBA and DSA with four different representative probabilities p of parallel executions, 0.1, 0.4, 0.7 and 0.9. We varied the density or number of sensors and compared the quality of the colorings that DSA and DBA produced. We changed the number of sensors from 25 to 100 with an increment of 5 sensors, and from 100 to 500 with an increment of 25 sensors. We averaged the re-

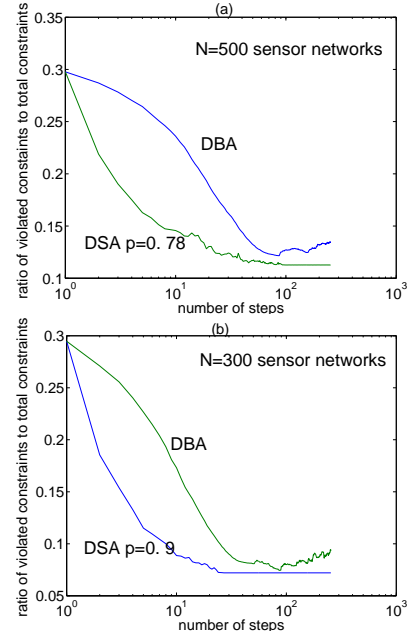


Figure 4: Anytime performance of DSA and DBA in dense sensor networks, $T = 6$.

sults over 100 random problem instances for each fixed number of sensors. Figure 2 shows the result on $T = 6$. The horizontal axes in the figures are the numbers of sensors, and the vertical axes are the normalized solution quality after 256 steps. Longer executions, such as 512 and 1024 steps, exhibit almost identical results. Figure 2(a) shows the result in the whole range of 25 to 500 sensors and Figure 2(b) expands the results of Figure 2(a) in the range of 25 to 100 sensors.

As analyzed in [8], DBA may perform better than DSA on underconstrained problems. An underconstrained scan scheduling problem may be created when more colors are available. Indeed, when we increase the number of allowed colors (targeting schedule cycle length) to eighteen ($T = 18$), DBA outperforms DSA on sparse networks with less than 50 sensors. This result is shown in Figure 3, where each data point is averaged over 100 trials.

Based on the experimental results, we can reach three conclusions. First, DBA typically performs worse than DSA when its degree of parallelism is not too high in the range of 25 to 500 sensors. Second, when the sensor density increases, the performance of DSA may degenerate, especially if its degree of parallelism p is high. For instance, DSA with $p = 0.9$ becomes the worst of all when there are more than 400 sensors (Figure 2). This means that the denser the sensor networks are, the smaller the parallel degree p should be. The degenerated performance of DSA with a large p is mainly due to its phase-transition behavior revealed in the previous section. When the sensor density increases, more constraints will be introduced into the inherited constraints of the scan scheduling problem, so that the problem becomes overconstrained. As indicated in the phase-transition section, DSA's phase-transition behavior appears sooner when overall constraints are tighter. Third, in the underconstrained region, a higher degree of parallelism is preferred to a lower degree. This result is in line with that in [8].

6.2 Anytime performance

An important feature of our targeting sensor network for object




Figure 5: Anytime performance of DSA and DBA in sparse sensor networks, $T = 6$.

detection is real-time response. High real-time performance is important, especially for systems with limited computation and communication resources in which it may be disastrous to wait for the systems to reach stable or equilibrium states. This is particularly true for our sensor-based system for object detection and mobile object tracking. Therefore, the algorithm for the distributed scan scheduling must have anytime property, i.e., the algorithm can be stopped at anytime during its execution and is able to provide a feasible solution at that point. Fortunately, DBA and DSA can both be used for this purpose because the hard constraints internal to individual sensors (a sensor cannot scan its two sectors at the same time) are always maintained.

In the rest of this section, we directly compare DBA and DSA as anytime algorithms. We use the same set of experimental conditions and parameters as in the previous sections, i.e., sensors have three sectors and are randomly and uniformly placed on a 10×10 grid, with results averaged over 100 trials.

We first consider dense networks with $N = 500$ and $N = 300$ sensors. Based on the phase-transition results in Section 5, DSA performs the best with $p = 0.9$ and $p = 0.78$ for the networks of $N = 500$ and $N = 300$ nodes, respectively. We used these parameters in our experiments. The experimental results are in Figure 4. As the results show, DSA performs much better than DBA in both anytime performance and final solution quality.

We now consider sparse sensor networks, using networks with $N = 80$ and $N = 40$ sensors as representatives. As discussed earlier, a higher degree of parallelism should be used on sparse graphs, we thus use $p = 0.9$ for our two sparse networks. The results, averaged over 100 trials, are in Figure 5. Clearly, DBA and DSA exhibits similar performance, with DSA being able to produce slightly better solutions at the end.

To complete our analysis, we also compared DSA and DBA on the same sets of instances, but with 18 available colors ($T = 18$). The results on dense and sparse sensor networks are included in Figures 6 and 7, respectively. Interestingly, DBA's anytime perfor-

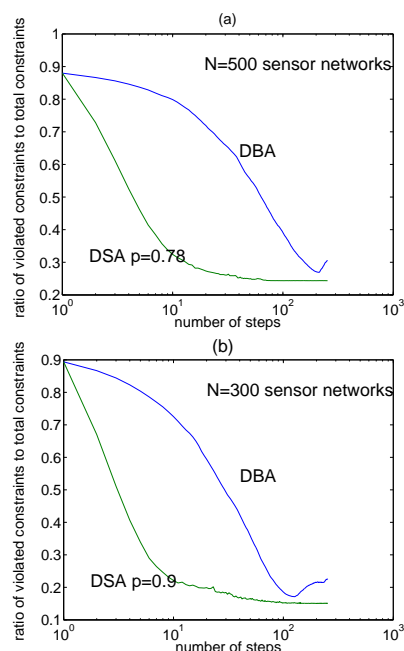


Figure 6: Anytime performance of DSA and DBA in dense sensor networks, $T = 18$.

mance degenerates, compared to that using $T = 6$.

In summary, as far as solution quality (anytime and final solutions) is concerned, our experimental results indicate that DSA should be adopted for the distributed scan scheduling problem.

7. COMMUNICATION COST

We have so far concentrated on the solution quality of DBA and DSA without paying any attention to their communication cost. As mentioned earlier, communication in a sensor network has an inherited delay and could be unreliable in most situations. Therefore, a good distributed algorithm should require a small number of message exchanges.

In each step of DBA, an agent announces its best possible conflict reduction to its neighbors and receives from the neighbors their possible weight reductions. Thus, the number of messages sent and received by an agent in each step of DBA is no less than the number of its neighbors, and the total number of messages exchanged in each step is more than a constant for a given network.

In contrast, an agent may not have to send a message in a step in DSA if it does not change its value. In an extreme case, an agent will not change its value if it is at a local minima. If solution quality of DSA improves over time, its communication cost will reduce as well. In principle, the communication cost of DSA is correlated to its solution quality. The better the current solution, the less the number of messages. In addition, the communication cost is also related to the degree of parallel executions of the agents. The higher the parallel probability p is, the higher the communication cost will be. As shown in [8], the communication cost of DSA goes hand-in-hand with its solution quality and also experiences a similar phase-transition or threshold behavior on regular coloring problems. Figure 8 shows the phase-transition behavior of DSA's communication cost on the $N = 500$ and $N = 300$ sensor networks using $T = 6$ that we studied before. Here we considered the accumulative communication cost of all 256 steps. This result

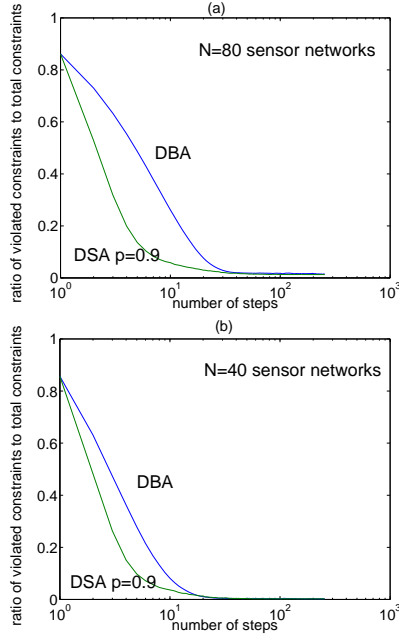


Figure 7: Anytime performance of DSA and DBA in sparse sensor networks, $T = 18$.

indicates that the degree of parallelism must be controlled properly in order to make DSA effective. Similar phase-transition patterns have been observed when we use $T = 18$.

We now compare DSA and DBA in terms of communication cost. Figure 9 shows the results evaluating DBA and DSA with probability $p = 0.78$ on $N = 500$ networks using $T = 6$ and $T = 18$, averaged over 100 trials. The figures plot the average numbers of messages exchanged in DSA and DBA at a particular step. Clearly, DSA has a significant advantage over DBA on communication cost. The large difference on communication cost between DSA and DBA will have a significant implication on how these two algorithms can be used in real sensor networks, especially when the sensors are connected through delayed, unreliable and noisy wireless communication. For our particular application and system where communication was carried out by radio frequencies, DBA's high communication cost makes it noncompetitive.

In summary, in terms of solution quality and communication cost, DSA is preferable over DBA for our distributed scan scheduling if DBA's degree of parallelism is properly controlled.

8. SOLVING SCHEDULING PROBLEM

Based on the results from Sections 5 to 7, we now apply DSA and DBA to dealing with two related problems at the same time, finding the shortest scan cycle length T and obtaining a good schedule given the shortest cycle length T .

To this end, we run DSA and DBA in iterations, starting with an initially large T . T is reduced after each iteration. Given a T in an iteration, DSA or DBA searches for a schedule of a quality better than a predefined threshold Q . The iteration stops whenever such a schedule is found within a fixed number of steps, and a new iteration may start with a smaller T .

In our simulation, we checked the quality of the current schedule after each simulated step. As soon as the quality of the current schedule exceeds the given threshold Q , we terminate the current

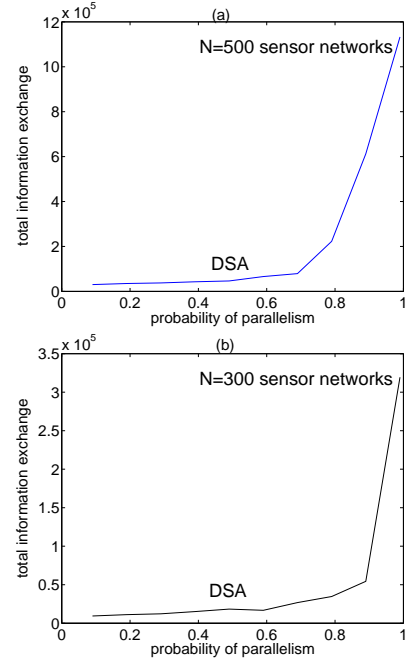


Figure 8: Communication-cost phase transitions of DSA on scan scheduling, $T = 6$.

iteration. This is equivalent to having an agent compute the global state of a distributed system, a method infeasible for our completely distributed system. We use this mechanism here simply to evaluate the performance of DSA and DBA.

Figures 10 and 11 show the results on two networks, one with $N = 80$ sensors and the other with $N = 300$. In our experiments, we fixed the sensor activation ratio at $\alpha = 2/3$, used initial $T = 36$, and reduced T by three after each iteration, which ran a maximum of 256 steps. The threshold for schedule quality was set to $Q = 0.01$ for $N = 80$ and $Q = 40$ for $N = 300$. As the results show, DSA is superior to DBA. On the $N = 80$ network (Figure 10), DSA finds a targetting schedule of length $T = 15$ in 242 steps, while DBA needs 588 steps. On the $N = 300$ network (Figure 11), DSA takes 64 steps, while DBA uses 691 steps, which is an order of magnitude difference.

In summary, our results clearly show that DSA is superior to DBA on the distributed scan scheduling problem. If communication cost is also a concern, DSA is definitely the algorithm of choice for the problem.

9. CONCLUSIONS

Motivated by real applications of multiagent systems in sensor networks, we studied low-overhead distributed algorithms for solving distributed constraint satisfaction and optimization problems. We analyzed and compared distributed stochastic algorithm (DSA) [1, 3, 8] and distributed breakout algorithm (DBA) [4, 6, 7, 9] on a distributed scheduling problem in sensor networks. We specifically investigated the relationship among the degree of parallel executions, problem constrainedness, and DSAs' behavior and performance. We showed that DSA exhibits a threshold behavior similar to phase transitions in which its performance, in terms of both solution quality and communication cost, degrades abruptly and dramatically when the degree of agents' parallel execution increases

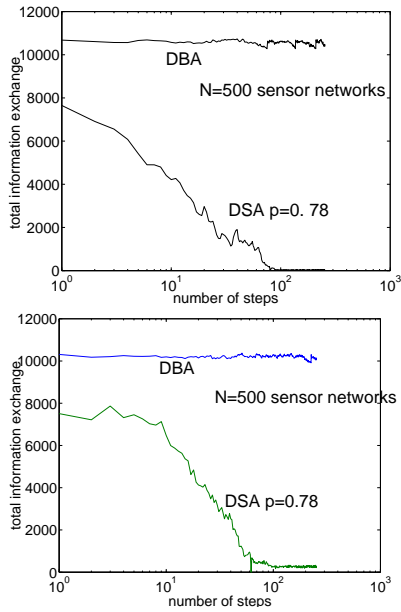


Figure 9: Communication cost of DSA and DBA, $T = 6$ (left) and $T = 18$ (right).

beyond a critical point. On the other hand, we also showed that if controlled properly, DSA is significantly superior to DBA, finding better solutions with less computational cost and communication overhead. For distributed scheduling problems such as the one considered in this paper, DSA is the algorithm of choice.

10. REFERENCES

- [1] M. Fabiunke. Parallel distributed constraint satisfaction. In *Proc. Intern. Conf. on Parallel and Distributed Processing Techniques and Applications (PDPTA-99)*, pages 1585–1591, 1999.
- [2] S. Fitzpatrick and L. Meertens. Experiments on dense graphs with a stochastic, peer-to-peer colorer. In *AAAI-02 Workshop on Probabilistic Approaches in Search*, to appear.
- [3] S. Fitzpatrick and L. Meertens. An experimental assessment of a stochastic, anytime, decentralized, soft colourer for sparse graphs. In *Proc. 1st Symp. on Stochastic Algorithms: Foundations and Applications*, pages 49–64, 2001.
- [4] P. Morris. The breakout method for escaping from local minima. In *Proc. AAAI-93*, pages 40–45.
- [5] G. Tel. *Introduction to Distributed Algorithms*. Cambridge University Press, 2000.
- [6] M. Yokoo. *Distributed Constraint Satisfaction: Foundations of Cooperation in Multi-Agent Systems*. Springer Verlag, 2001.
- [7] M. Yokoo and K. Hirayama. Distributed breakout algorithm for solving distributed constraint satisfaction problems. In *Proc. ICMAS-96*, pages 401–408, 1996.
- [8] W. Zhang, G. Wang, and L. Wittenburg. Distributed stochastic search for distributed constraint satisfaction and optimization: Parallelism, phase transitions and performance. In *Proc. AAAI-02 Workshop on Probabilistic Approaches in Search*, pages 53–59, 2002.
- [9] W. Zhang and L. Wittenburg. Distributed breakout revisited. In *Proc. AAAI-02*, pages 352–357, 2002.

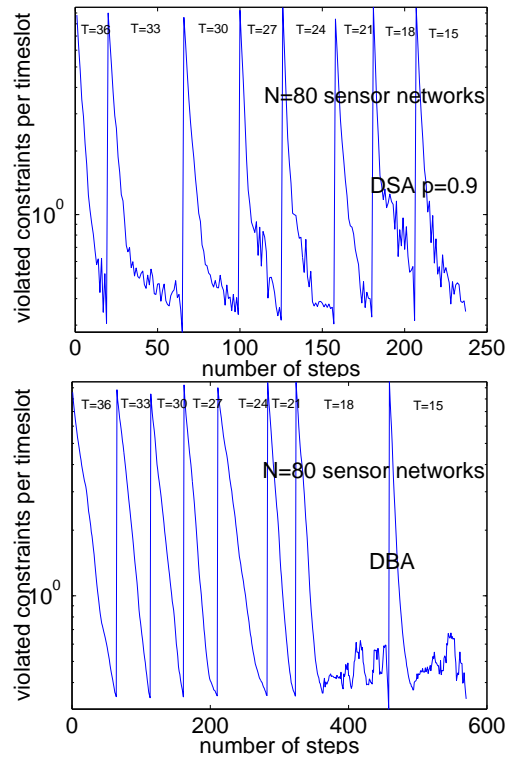


Figure 10: Finding best possible schedule using DSA (top) and DBA (bottom), $N = 80$.

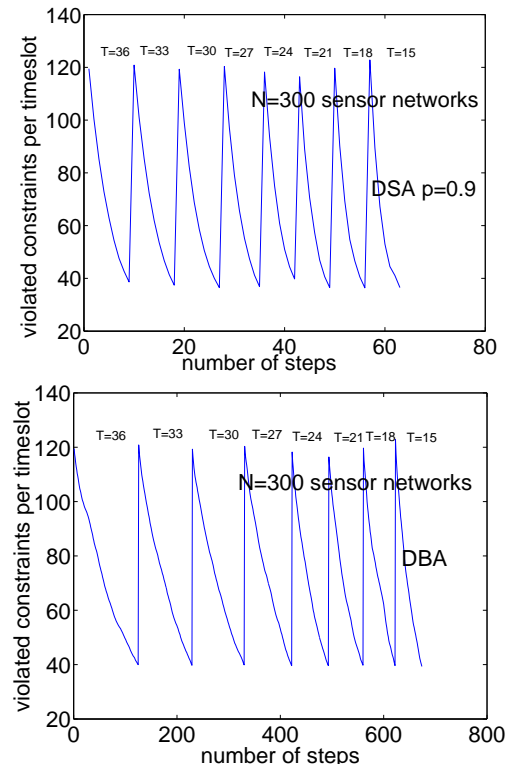


Figure 11: Finding best possible schedule using DSA (top) and DBA (bottom), $N = 300$.