

An Efficient Algorithm For Solving Dynamic Complex DCOP Problems

Sankalp Khanna^{1, 2}
¹Institute for Integrated and
Intelligent Systems
²The Australian e-Health
Research Centre
S.Khanna@griffith.edu.au

David Hansen
The Australian e-Health
Research Centre
71/918, RBWH, Herston, QLD
4029, Australia
David.Hansen@csiro.au

Abdul Sattar
Institute for Integrated and
Intelligent Systems
Griffith University, QLD 4111,
Australia
A.Sattar@griffith.edu.au

Bela Stantic
Institute for Integrated and
Intelligent Systems
Griffith University, QLD 4111,
Australia
B.Stantic@griffith.edu.au

ABSTRACT

Multi Agent Systems and the Distributed Constraint Optimization Problem (DCOP) formalism offer several asynchronous and optimal algorithms for solving naturally distributed optimization problems efficiently. There has been good application of this technology in addressing real world problems in areas like Sensor Networks and Meeting Scheduling. Most of these algorithms however exploit static tree structures and are thus not well suited to modeling and solving problems in rapidly changing domains. Also, while in theory most DCOP algorithms can be extended to handle complex local sub-problems, we argue that this generally results in making their performance sub-optimal, and thus their application less suitable. In this paper we present new measures that emphasize the interconnectedness between each agent's local and inter-agent sub-problems and use these measures to guide dynamic agent ordering during distributed constraint reasoning. The resulting algorithm, DCDCOP, offers a robust, flexible, and efficient mechanism for modeling and solving dynamic complex problems. Experimental evaluation of the algorithm shows that DCDCOP significantly outperforms ADOPT, the gold standard in search-based DCOP algorithms.

1. INTRODUCTION

Despite being a relatively young research area, with the first asynchronous Distributed Constraint Satisfaction Problem (DisCSP) algorithm proposed in 1992 [17], and the first complete Distributed Constraint Optimization Problem (DCOP) algorithm, ADOPT, proposed in 2003 [10], the Distributed Constraint Reasoning formalism has developed rapidly to of-

fer efficient and sophisticated algorithms to model and solve a variety of naturally distributed multi-agent problems. Several notable DCOP approaches employing techniques from search (e.g. ADOPT and its several variants), dynamic programming (e.g. DPOP [12] and its several variants) and cooperative mediation (e.g. APO [8]) have emerged and are being successfully used to model and solve problems in sensor networks, meeting scheduling, etc.

This research was motivated by our effort to model and solve naturally distributed complex optimization problems in the health domain, a typical example being the scheduling of elective surgery in a large public hospital. The problem involves several departments, each with its own complex scheduling problem. The departments need to negotiate with each other to build, and maintain, the elective surgery schedule in the face of a dynamic health landscape. Allocation of airport slots to airlines, or public infrastructure to utilities companies, are examples of similar problems in relatively dynamic environments where several agents, each with complex sub-problems, are negotiating in a privacy preserving manner, to optimize a common cost function.

Working towards this aim, the first natural observation is that most current algorithms are based on tree-structures, which are static in nature and would continually need to be rebuilt in dynamic environments. Also, given the nature of the problem domain, partial centralization based strategies would not be a good fit here because of obvious privacy and decision control concerns.

We also note that the metrics used to compare algorithms are questioned by most researchers. Silaghi and Yokoo [16] have shown that it is possible to construct problems that can be exploited by algorithms such as ADOPT and DPOP to exhibit their superiority. Also, Maheswaran et al. [7] show that the performance of ADOPT in solving real world problems is significantly worse than in solving similar-sized map coloring problems.

We draw from Zhou's work [18] in the DisCSP field and gen-

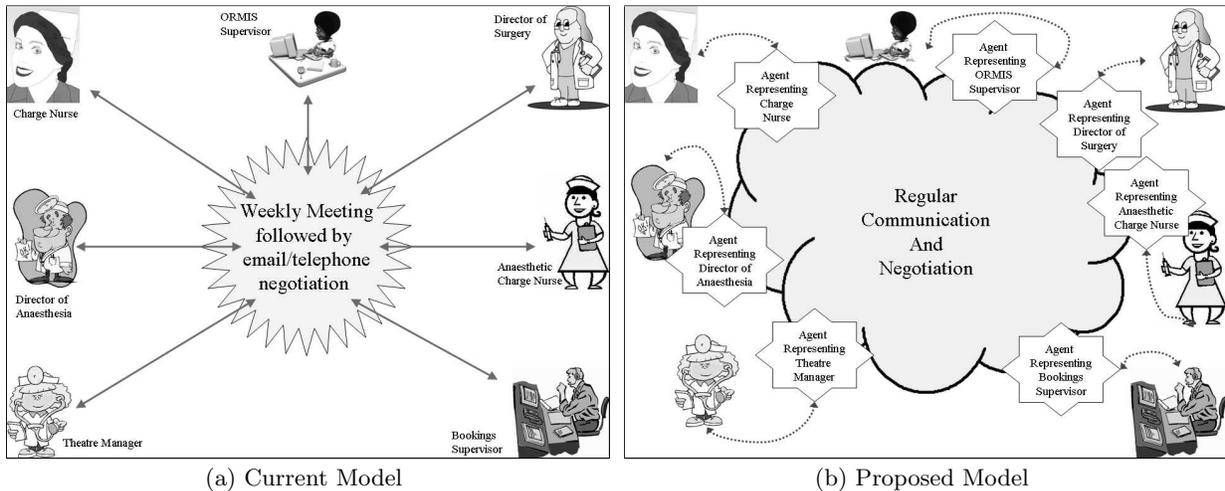


Figure 1: Scheduling Elective Surgery at the Princess Alexandra Hospital

eralize the novel measures of constraint density ¹, to introduce new DCOP measures of Dynamic Cost Density (*DCD*) and Degree of Unsatisfaction (*DU*), and then use these measures to dynamically guide agent ordering in our new Dynamic Complex Distributed Constraint Optimization Problem (DCDCOP) algorithm. We compare DCDCOP’s performance to ADOPT, the current standard in DCOP search, and show that the new algorithm offers a significant improvement over ADOPT.

The rest of this paper is organized as follows. In section 2, we describe our case study of the elective surgery scheduling problem at a large Australian public hospital. Section 3 presents our Multi-Agent System (MAS) architecture for modeling and solving this class of problems. Section 4 follows with a discussion of the DCOP formalism and the shortcomings of the current state of the art in DCOP algorithms in addressing real world dynamic complex problems. In section 5, we present the new measures of *DCD* and *DU*, and the DCDCOP algorithm, and explain these with a simple example. Section 6 reports on the empirical evaluation of the algorithm as compared to ADOPT. Lastly, in section 7, we present our main conclusions and discuss ongoing work.

2. SCHEDULING ELECTIVE SURGERY AT THE PRINCESS ALEXANDRA HOSPITAL

Elective surgery is a planned, non-emergency surgical procedure, which can be scheduled at the patient’s and surgeon’s convenience. The escalating demand for elective surgery is however compounded by a shortage of trained surgeons, anaesthetists and nurses. Recent Queensland statistics show that, as of 1 October 2008, 34,514 patients were waiting for elective surgery of which almost 21% had waited longer than a clinically desirable time [15].

We conducted an extensive study of scheduling processes followed at the Princess Alexandra Hospital (PAH), a large public hospital in Queensland’s capital city of Brisbane. PAH

offers 21 operating theaters that can be utilized by various departments for elective surgery. For the process of scheduling, the theater schedule is divided into AM and PM slots of 3.5 hours each. These slots are allocated to various doctors or departments but can be utilized for trauma or emergency if urgently required.

Each department connected (i.e. allocating staff or other resources) to the surgery carries out their individual scheduling activity. The bookings department books patients into the *Bookings Schedule* in consultation with the relevant surgical teams, and records these bookings into the *Operating Room Management Information System* (ORMIS). The different departments can access this information by looking into ORMIS or accessing the latest *Bookings Schedule* on the shared drive, where it is updated everyday at 3PM.

Every Thursday, the managers of the different departments meet and review bookings for the week ahead (Figure 1). Each session is discussed and conflicts in the departmental schedules are worked out by negotiation. Unexpected emergencies, variation in patients’ health state, sudden perturbations in staffing, and surgeon availability etc. lead to further changes being often required. However, all changes made to the system after this meeting are dealt with individually by the departments, and resolved on a case-by-case basis using conventional communication such as telephone and emails, or even by face-to-face meetings. In keeping with the dynamics of the domain, the schedule needs to be updated quickly and efficiently. This is often not possible, because of delays in inter-departmental communication, and this leads to the adoption of an easy but inefficient solution resulting in a compromised schedule. For example, if a procedure is canceled at the last minute, because new medical reports say it is no longer required, the bookings department would want to offer the slot to another procedure. They may however be unable to confirm availability of specialist staff or equipment, because the charge nurse or theater manager were temporarily unavailable, which would lead to the slot being unused.

¹Not to be confused with the traditional measure, i.e. ratio of actual number of constraints to possible number of constraints.

3. PROPOSED MAS ARCHITECTURE

Given the naturally distributed nature of problems like elective surgery scheduling, an intelligent agent based approach is a promising paradigm for modeling the environment.

We propose a methodology where intelligent agents, trained with the constraints, preferences, priorities etc. of the administrators, optimize schedule for their respective departments (Figure 1). They then negotiate in a privacy-preserving manner (i.e. without sharing more information than is essential) to resolve inter-agent constraints. The architecture of each agent (see Figure 2) incorporates an interface module to handle internal and external communication, an intelligence module to handle decision making and learning, and a DCOP engine to drive the optimization.

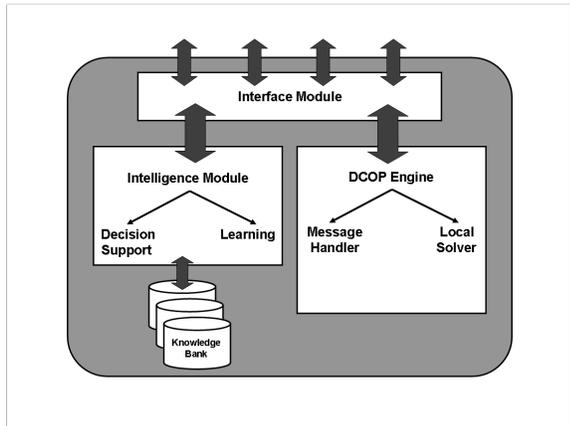


Figure 2: Proposed Architecture for Agent

The agents thus incorporate learning from domain-expert interaction, and have the ability to use intelligent reasoning to translate environmental changes and negotiation requests to constraints for the DCOP engine. Using an appropriate DCOP algorithm the agents optimize their local solution, and then collaborate with each other to resolve inter-agent constraints and align themselves.

The DCOP algorithm thus needs to be robust enough to handle the complexity of each agent’s sub-problems and perform inter-agent negotiation in a truly distributed manner, efficient enough to solve the problem in a timely and optimal manner, and flexible enough to adapt to the dynamics of the environment. Additionally if it can shield the local solver mechanism from the inter-agent negotiation process, each agent will have the ability to use a local solver strategy to suit its own need.

4. RELATED WORK

We start with discussing the DCOP formalism and some of the current approaches to solving DCOPs. We then revisit real world dynamic complex problems and discuss the shortcomings of current approaches in addressing these.

4.1 The DCOP formalism

In solving a DCOP, the goal for each agent is to assign values to its variables such that a given global objective function is minimized. The cost functions in DCOP are analogous

to constraints in DisCSP, and DCOP is thus regarded as a generalization of the DisCSP formalism. For simplicity, we use the term constraints and cost functions interchangeably. Formally, we can define a DCOP as consisting of:

1. A finite ordered set of Agents $A = \{A_1, A_2, A_3, \dots, A_n | n \in \mathbb{Z}^+\}$, where, for each Agent A there exists :
 - (a) A finite ordered set of variables $V = \{V_1, V_2, \dots, V_n | n \in \mathbb{Z}^+\}$,
 - (b) A domain set $D = \{D_1, D_2, \dots, D_n\}$, containing a finite and discrete domain D_i for each variable V_i ,
 - (c) A constraint set $C = \{C_1, C_2, \dots, C_m\}$, $m \in \mathbb{Z}^+$, where each C_j , $\forall j \in [1, m]$ is defined as a cost function on a pair of variables, $f_{i,i'} : D_i D_{i'} \rightarrow \mathbb{N}$, $\forall V_i, V_{i'} \in V$, and
 - (d) An ordered solution set $S = \{v_1, v_2, v_3, \dots, v_n | v_i \in D_i, \forall i \in [1, n]\}$ where the aggregate cost $F(A) = \sum_{(x_i, x_{i'} \in V)} f_{ii'}(d_i, d_{i'})$, $x_i \leftarrow d_i, x_{i'} \leftarrow d_{i'} \in A$.
2. The solution set of the DCOP S^* is defined as the set of the solution sets of each agent.

In keeping with the norm, all constraints are assumed to be binary, and optimization functions are assumed to be associative, commutative, and monotonic [10]. In dealing with complex DCOPs however, we do not assume one variable per agent.

4.2 Current State of the Art

Several DCOP algorithms and their variants have been introduced by recent research. Due to space constraints, we focus on the following key complete optimal algorithms, each representing a significantly different approach. Also, we do not critique each variant as they all still have the same shortcomings when applied to dynamic complex problems.

Asynchronous Distributed OPTimization, or ADOPT [10], is a complete and asynchronous DCOP algorithm. In ADOPT, agents are first prioritized into a Depth First Search (DFS) tree, whereby each agent maintains lower and upper bounds for the subtree rooted at their node. The agents then use opportunistic best-first search to assign their variables such that the lower bound is minimized. *Cost* messages propagate up the tree and *threshold* and *value* messages are sent down the tree, iteratively tightening the lower and upper bounds until the lower bound of the minimum cost solution is equal to its upper bound. If an agent detects this condition, and its parent has terminated, then an optimal solution is found and it may also terminate. The other key idea in ADOPT is to store lower bounds as a threshold and discard partial solutions before they are proven to be definitely suboptimal, thus maintaining linear space complexity at each agent. In the worst case, ADOPT may require an exponential number of messages to arrive at a solution.

Distributed Pseudotree Optimization Procedure, or DPOP [12], is a complete dynamic programming algorithm that involves a three phase process. Similar to ADOPT, the first

phase involves the formation of the DFS tree. Phase two involves calculating and propagating the utility (cost) bottom-up, i.e. from the leaves upwards to the root. Phase three involves a downward value propagation, initiated by the root node. Each agent then calculates its optimal value based on the utility message received from its subtree and the value message received from its parent. DPOP thus generates only a linear number of messages, but the message size grows with every traversal up the tree and the algorithm thus requires large amounts of memory, up to space exponential in the induced width of the problem.

Optimal Asynchronous Partial Overlay (OptAPO) [8] is an alternative approach to DCOP that utilizes partial centralization to solve difficult portions of a DCOP problem. While partial centralization offers an excellent mechanism to solve DCOPs in several scenarios, it is generally unsuitable in the kind of problems we seek to address. In most such cases, negotiating agents would normally refuse to share information or relinquish decision making control of their private sub-problems.

4.3 Solving Dynamic Complex Problems

Since both ADOPT and DPOP utilize static DFS tree structures, changes to the constraints would often result in the need for the tree to be rebuilt. Also, since ADOPT discards no-goods to maintain linear space complexity, changes to the constraints would result in bounds being discarded and the search restarted.

Both ADOPT and DPOP also offer variants to deal with dynamic environments. Modi [9] offers a formalism for mapping and solving dynamic resource allocation problems but this is applied in the DisCSP domain. This is extended to map over-constrained problems into DCOP but can handle only static problems as the author concedes to the lack of an effective DCOP algorithm for dynamic problems. Petcu et al. [11][13] propose S-DPOP and RS-DPOP, which utilize self stabilizing DFS trees to guarantee optimal solution stability in distributed continuous-time combinatorial optimization problems. Lass et al. [6] present another mechanism to deal with the *complicating factor of dynamism* by wrapping ADOPT in an *Adapter* that receives and handles dynamic event requests.

In dealing with the issue of complex sub-problems, algorithms can theoretically utilize decomposition or compilation. In practice however, decomposition results in failure to exploit the inherent benefit of domain centralization, and also blows the distributed problem size out of proportion. Burke and Brown [2] show that the compilation outperforms decomposition in case of large local sub problems but only small domain size, whereas decomposition is more appropriate when the number of inter-agent constraints and domain size is large but only for small problems. Several ADOPT variants use techniques such as decomposition [9], compilation [4], interleaving [3], and relaxation [1], to name a few, to deal with complex sub-problems. All of the above variants however still suffer from the problem of working off a static tree structure that needs rebuilding from time to time.

Further, applying decomposition to DPOP would result in a significant increase in the message sizes, while compilation

would need a novel mechanism of calculating the agent utility for different combinations of local variable assignments. It would however be interesting to evaluate the effect of utilizing our measure of DU as the utility metric on DPOP.

We thus conclude that the DCOP algorithms used in all of the above are optimal only in a static environment, and there is need for a more flexible robust algorithm, which can model the complexity and adapt better to a dynamic environment.

5. THE DCDCOP ALGORITHM

We present the new measures of DCD and DU and a new Dynamic Complex Distributed Constraint Optimization Problem (DCDCOP) algorithm and explain these measures, and algorithm execution, with a simple example.

5.1 Defining DCD and DU

Zhou [18] offers innovative measures of Dynamic Constraint Density (DCD) and Degree of Unsatisfaction (DU) (a measure of how far an agent's instantiation is from reaching a consistent state) and presents algorithms based on these, which can handle complex and dynamic constraint reasoning problems efficiently. These however are suited only to satisfaction and constraint relaxation approaches as the measure does not take varying constraint costs into account.

We generalize the definitions of Zhou to define the following new static measures of Intra-Agent Cost Density ($IACD$) and Inter-Agent Cost Density (I_ACD):

$$IACD_i = \begin{cases} 0, & \text{if } |intraV_i|=0 \\ \frac{\sum_{j=1}^{|intraC_i|} (\delta_m(intraC_i^j))}{|intraV_i|}, & \text{otherwise} \end{cases} \quad (1)$$

$$I_ACD_i = \begin{cases} 0, & \text{if } |interV_i|=0 \\ \frac{\sum_{j=1}^{|interC_i|} (\delta_m(interC_i^j) + \sum_{l=1}^{|\kappa C_i^j|} (\delta_m(intraC_i^l)))}{|interV_i|}, & \text{otherwise} \end{cases} \quad (2)$$

where $intraC_i$ is the set of intra-agent constraints for agent i , δ_m represents the maximum cost of the constraint, $intraC_i^j$ is the j^{th} intra-agent constraint of agent i , $intraV_i$ is the set of variables constrained by $intraC_i$, $interC_i$ is the set of inter-agent constraints for agent i , $interC_i^j$ is the j^{th} inter-agent constraint of agent i , κC_i^j is the set of intra-agent constraints belonging to i and connected to $interC_i^j$, and $interV_i$ is the set of variables constrained by $interC_i$ and controlled by agent i .

The measure of I_ACD takes into account the interconnectiveness of the variables that are attached to an inter-agent constraint. So the higher the cost of intra-agent constraints attached to the variable, the greater the impact of the variable on the cost density. This is quite apt as changing the value of this variable would attract a much higher effort towards optimizing the problem. Also, the measures of $IACD$ and I_ACD both equate to zero if there are no intra-agent or

inter-agent variables connected to constraints respectively. This follows from the general idea that a lower value of cost density would mean the problem is closer to the solution. Thus no variables, and a consequent cost density of 0, would imply that this component of the problem does not need to be solved further. We now define the measure of Static Cost Density (SCD):

Static Cost Density of an agent is defined as the sum of the maximum possible Intra-Agent Cost Density and the maximum possible Inter-Agent Cost Density.

$$SCD_i = IACD_i + I_ACD_i \quad (3)$$

In equations 1 and 2 we replace δ_m by δ_c , which gives us the current cost of the constraint, to get our dynamic measures of IntraUnsat (IU), and InterUnsat (I_U), which represent the dynamic intra-agent and inter-agent cost densities respectively. We utilize these to calculate the measure of Dynamic Cost Density (DCD).

$$IU_i = \begin{cases} 0, & \text{if } |IACD_i|=0 \\ \frac{\sum_{j=1}^{|intraC_i|} (\delta_c(intraC_i^j))}{|intraV_i|}, & \text{otherwise} \end{cases} \quad (4)$$

$$I_U_i = \begin{cases} 0, & \text{if } |I_ACD_i|=0 \\ \frac{\sum_{j=1}^{|interC_i|} (\delta_c(interC_i^j) + \sum_{l=1}^{|\kappa C_i^j|} (\delta_c(intraC_i^l)))}{|interV_i|}, & \text{otherwise} \end{cases} \quad (5)$$

Dynamic Cost Density of an agent is defined as the sum of the current Intra-Agent Cost Density and the current Inter-Agent Cost Density.

$$DCD_i = IU_i + I_U_i \quad (6)$$

We now redefine the measure of Degree of Unsatisfaction for agent i (DU_i) based on the above :

Degree of Unsatisfaction of an agent is defined as the ratio of the current Dynamic Cost Density to Static Cost Density.

$$DU_i = \frac{DCD_i}{SCD_i} \quad (7)$$

The DU for a problem with $|i|$ agents can similarly be calculated as the ratio of the sum of agent DU s to the number of agents $|i|$. The measure of DU provides a measure of how far away that agent's (or problem's) current instantiation is from reaching a optimal state. It does not provide a direct measure to compare the level of complexity to two agents' problems or the time it may take to solve them. However, unlike a simple summation of max cost or current cost, it attaches a higher cost to the changing of more interconnected constraints, thus providing a more realistic measure.

5.2 An Example of Calculating DU

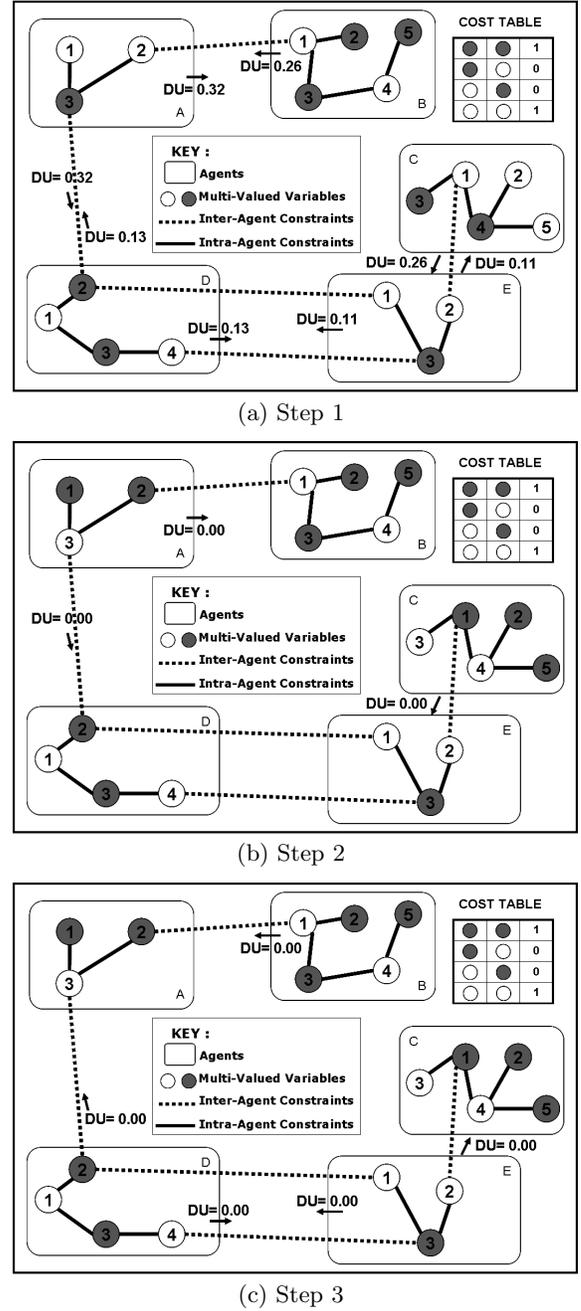


Figure 3: Example of DCDCOP Execution

To better understand the above measures, we utilize a simple example as shown in Figure 3(a). Here, we calculate the values of SCD , DCD and DU for agent D , which has four variables, three intra-agent constraints and three inter-agent constraints. The max cost of each constraint (from the cost table shown in the figure) is 1. In calculating the static measures for the problem, we have:

$$\begin{aligned}
IACD_D &= \frac{(1+1+1)}{4} = 0.75 \\
I_ACD_D &= \frac{((1+(1)) + (1+(1))) + (1+(1))}{2} = 3 \\
SCD_D &= 0.75 + 3 = 3.75
\end{aligned}$$

Now, assuming a snapshot view of the scenario, where each agent has received the communicated values of DU , we can calculate the dynamic measures:

$$\begin{aligned}
IU_D &= \frac{(0+0+0)}{4} = 0 \\
I_UD &= \frac{((1+(0)) + (0+(0))) + (0+(0))}{2} = 0.5 \\
DCD_D &= 0 + 0.5 = 0.5 \\
DU_D &= \frac{0.5}{3.75} = 0.13
\end{aligned}$$

Note that the constraint between variables 1 and 2 of agent D is counted twice in the calculation of I_ACD_D and I_UD . The calculated value of $DU = 0.13$ will be sent to neighbors A and E .

5.3 Details of Algorithm

The DCDCOP algorithm is implemented as follows :

- All agents start by calculating the values of $IACD$, I_ACD and SCD . Then the agents instantiate their local variables using a Branch and Bound algorithm [5], thus ensuring that the cost at the end of this step is the minimum as per its current context. Each agent then calculates its dynamic measures of DCD and DU and sends its DU and the related context to its neighbors (i.e. those agents with whom it shares an inter-agent constraint).
- All agents then start to receive and handle messages on incoming links. When a message is received, the context of the message is checked to ensure that it is compatible with the agent's *CurrentContext*. If not, the message is discarded and the agent continues to listen on its incoming links. If the message is compatible, the *messageContext* is added to *CurrentContext* and the *messageDU* is compared with the agent's own DU . If the agents own DU is higher, it will reassign its variables, recalculate its dynamic measures and resend messages on outgoing links. If the agent's own DU is lower than *messageDU*, it will not reassign its variables, but if relevant, it will recalculate its dynamic measures and resend messages on outgoing links. In the event that *messageDU* = DU , the agent with a higher static ordering will reassign its variables, recalculate its dynamic measures and send messages on outgoing links.
- The search stops when each agent has achieved a stable state and no more messages are transacted. In the case of a solvable problem, this equates to a situation when the agent, and all its neighbors, arrive at $DU = 0$. In

the case of an optimization problem, this equates to a situation when the agent with a higher DU does not change its local solution as doing so would raise the cost of its solution.

Algorithm 1: The DCDCOP Algorithm

```

Calculate static measures
Solve_local_problem
Calculate dynamic measures
Send message ( $DU, CurrContext$ ) to all neighbors
Receive messages
when received ( $messageDU, msgContext$ ) do
  if  $msgContext$  and  $CurrContext$  are consistent then
    add  $msgContext$  to  $CurrContext$ 
    if  $DU > msgDU$  then
      | Solve_local_problem
    end
    else if  $DU = msgDU$  and higher_order then
      | Solve_local_problem
    end
    Calculate dynamic measures
    Send message ( $DU, CurrContext$ ) to all neighbors
  end
end do
Procedure : Solve_local_problem
Branch and Bound to solve local problem

```

The pseudo code of the algorithm is shown in Algorithm 1. As agents negotiate, each negotiation is handled in one of the following ways: if the values of DU are not identical, the agent with a higher value of DU will change its value; if the values of DU are identical, the agents will follow a fixed ordering between them to decide who changes their value. When an agent with a higher DU finds no better instantiation for its local variables, it will return the same, thus reaching a steady state until it receives a context message from other agents causing it to reevaluate its cost function.

Note that in the case of a dynamic real world problem, the agent with the lower value of DU can force the agent with the higher DU to negotiate by raising the cost attached to the inter-agent constraint.

5.4 Example of DCDCOP Execution

In continuing the example from section 5.2, assuming agents A , B and C are first to receive their messages, A and C will both reassign their variables and send out new messages, as they have received DU messages less than theirs. This results in the state shown in Figure 3(b).

Now assume D and E receive both messages together. Acting on the newer messages, both D and E will attempt to reassign their variables because their last calculated DU is higher than the *messageDU* from A and C respectively, but with the updated values from A and C , both will arrive at the same local solution as being the lowest cost. Similarly, B will arrive at the same local solution as being the lowest cost and they will send out their DU messages (Figure 3(c)).

At this stage, since all agents will detect that they have a DU of 0 and their neighbors have a DU of 0, the algorithm will terminate.

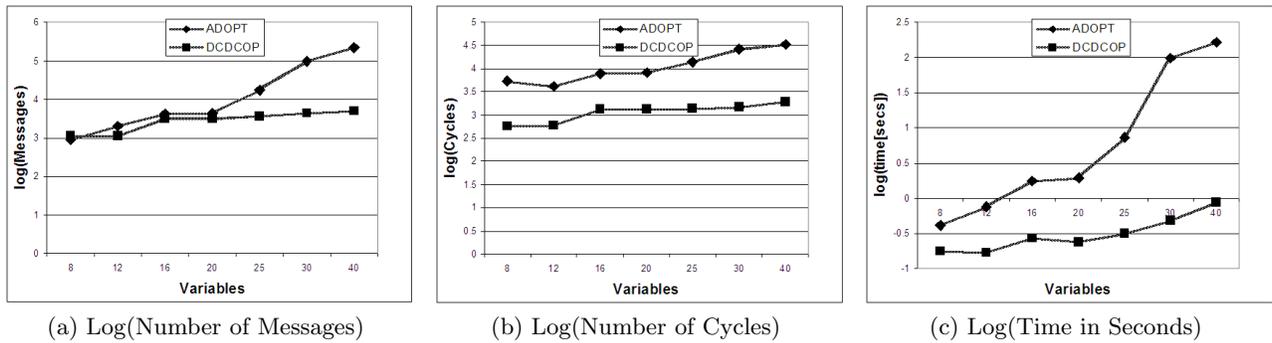


Figure 4: Performance of ADOPT vs DCDCOP ($LD = 2$)

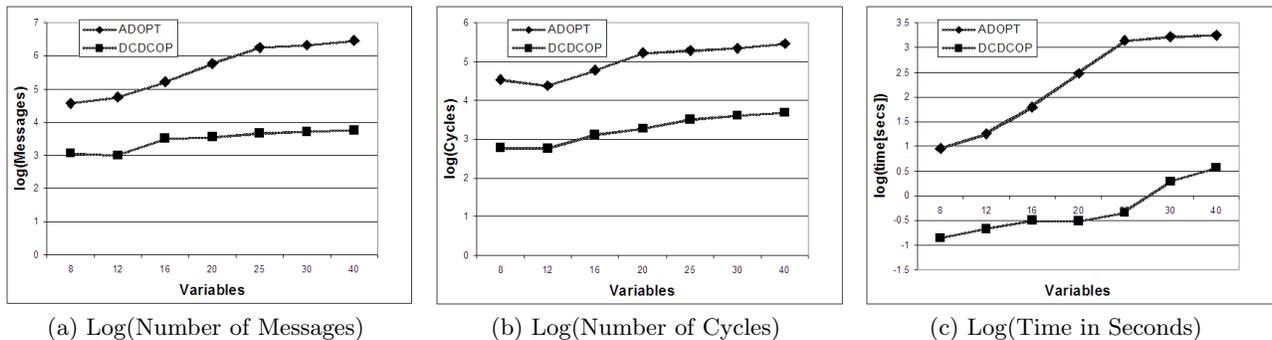


Figure 5: Performance of ADOPT vs DCDCOP ($LD = 3$)

6. EXPERIMENTAL EVALUATION

Given the need for agents to negotiate in a privacy preserving manner, otherwise optimal approaches like OptAPO, that use partial centralization, are not suitable in this context. And while DPOP and its variants are good DCOP algorithms, we choose ADOPT for our experimental evaluation because, like DCDCOP, it is based on constraint-guided search and thus provides a more realistic comparison. Further, ADOPT’s source code can handle multiple variables per agent without any modification to the algorithm, and is thus well suited to a fair and accurate comparison. To further ensure a reasonable comparison, the new algorithm is developed within the original ADOPT source code [14] and evaluated on the same graph coloring problems that were used to report ADOPT’s performance in [9] and come bundled with the source code. Also, ADOPT’s original messaging and performance evaluation procedures are utilized for the evaluation.

The original graph colouring problem data (8-40 variables) is evenly distributed between 3-5 agents. Also as in the original evaluation, we analyze the performance of both algorithms on problems with link density (LD) of 2 and 3. The performance is compared using three measures: number of messages, number of concurrent cycles, and time(secs). To prevent a large disparity between the results, the algorithms are run with a maximum time, $timeMax$ of 30 mins. Also, for a meaningful display of results, we report results on a logarithmic scale (base 10) (Figures 4 and 5).

We observe that DCDCOP outperforms ADOPT signifi-

cantly on all three scales of measurement. The speedup can be attributed partly to the algorithm exploiting domain centralization and performing each local reassignment within one cycle, and partly to the novel dynamic measures used to guide the inter-agent negotiation part of the algorithm.

The results allow for some extremely interesting observations. We observe that the difference in DCDCOP’s performance for $LD = 2$ and $LD = 3$ is not significant as in the case of ADOPT. This can be attributed to computational superiority over I/O speeds, the factor responsible for DPOPs performance gains over ADOPT [12]. Further, the performance of DCDCOP with 40 variables (distributed among 5 agents), is reasonably similar to the left side of the ADOPT charts, i.e. 8 agents with 1 variable per agent. Also, the performance of DCDCOP does not deteriorate much as we increase the problem size from 8 variables (in 3 agents) to 40 variables (in 8 agents). This further asserts the computational superiority of the centralized optimization algorithms such as Branch and Bound, and reinforces common belief that communication is the bottleneck in distributed problem solving.

We thus observe how problems that cannot be solved efficiently by DCOP algorithms can benefit greatly if there is a component of domain centralization that they can exploit. Further, a flexible robust algorithm like DCDCOP can help model the departmental centralization structure of large real world optimization problems, not only offering a natural mapping from real world problem solving structure to DCOP problem solving structure, but also exploiting this

to provide an order of magnitude improvement in performance.

7. CONCLUSION AND FUTURE WORK

Several real world optimization problems translate to agents with complex sub-problems in a dynamic environment, that need to negotiate in a manner where privacy and decision making authority is preserved. The current state of the art in DCOP deals with such problems by offering extensions to algorithms best suited for optimization of single-variable agents in a static environment and this often lead to sub-optimality.

We present a flexible, robust algorithm, DCDCOP, that is capable of exploiting the inherent domain centralization found in such problems, and uses a novel measure, DU, to dynamically guide agent ordering during optimization. Experimental evaluation shows that DCDCOP significantly outperforms ADOPT, the current state of the art DCOP search algorithm.

We are currently developing realistic and significant benchmark problems based on scheduling process information and data collected during our extensive interviews and interactions with schedulers and domain experts at the PAH. Further evaluation of DCDCOP and other DCOP algorithms on these benchmarks, and investigating the use of the measure of *DU* to guide agent based negotiation in various DCOP algorithms, are proposed as future work.

8. ACKNOWLEDGMENTS

The authors wish to thank Dr. Peter Moran and his colleagues at the PAH for their ongoing support over the last three years, for allowing us into their world, and sharing their invaluable expertise.

9. REFERENCES

- [1] D. Burke and K. Brown. Using relaxations to improve search in distributed constraint optimisation. *Artificial Intelligence Review*, 28(1):35–50, June 2007.
- [2] D. A. Burke and K. N. Brown. A comparison of approaches to handling complex local problems in DCOP. In *Sixth International Workshop on Distributed Constraint Reasoning*, page 27–33, Italy, 2006.
- [3] D. A. Burke and K. N. Brown. Interleaved search in DCOP for complex agents. In *Proceedings of the Doctoral Program, Principles and Practice of Constraint Programming - CP 2006*, Nantes, France, 2006.
- [4] J. Davin and P. J. Modi. Hierarchical variable ordering for distributed constraint optimization. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 1433–1435, Hakodate, Japan, 2006.
- [5] E. C. Freuder. Partial constraint satisfaction. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, USA*, page 278–283, 1989.
- [6] R. N. Lass, E. A. Sultanik, and W. C. Regli. Dynamic distributed constraint reasoning. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, pages 1466–1469, Chicago, 2008.
- [7] R. T. Maheswaran, M. Tambe, E. Bowring, J. P. Pearce, and P. Varakantham. Taking DCOP to the real world: Efficient complete solutions for distributed Multi-Event scheduling. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1*, pages 310–317, New York, 2004.
- [8] R. Mailler and V. Lesser. Solving Distributed Constraint Optimization Problems Using Cooperative Mediation. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1*, pages 438–445, New York, 2004.
- [9] P. J. Modi. *Distributed Constraint Optimization for Multiagent Systems*. PhD thesis, University of Southern California, USA, 2003.
- [10] P. J. Modi, W. Shen, M. Tambe, and M. Yokoo. An asynchronous complete method for distributed constraint optimization. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 161–168, Melbourne, 2003.
- [11] A. Petcu and B. Faltings. S-DPOP: superstabilizing, fault-containing multiagent combinatorial optimization. In *Proceedings of the Twentieth National Conference on Artificial Intelligence, AAAI-05*, page 449–454, Pittsburgh, Pennsylvania, USA, 2005.
- [12] A. Petcu and B. Faltings. A scalable method for multiagent constraint optimization. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 266–271, Edinburgh, Scotland, Aug. 2005.
- [13] A. Petcu and B. Faltings. Optimal solution stability in dynamic, distributed constraint optimization. In *Proceedings of the 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 321–327. IEEE Computer Society, 2007.
- [14] C. P. Portway. USC DCOP Repository. <http://teamcore.usc.edu/dcop>, 2008.
- [15] Queensland Health. Quarterly Public Hospitals Performance Report September Quarter 2008. http://www.health.qld.gov.au/surgical_access, 2008.
- [16] M. C. Silaghi and M. Yokoo. DFS Ordering in Nogood-based Asynchronous Distributed Optimization (ADOPT-ng). In *Sixth International Workshop on Distributed Constraint Reasoning*, Italy, 2006.
- [17] M. Yokoo, E. H. Durfee, T. Ishida, and K. Kuwabara. Distributed constraint satisfaction for formalizing distributed problem solving. In *International Conference on Distributed Computing Systems*, pages 614–621, 1992.
- [18] L. Zhou, J. Thornton, and A. Sattar. Dynamic Agent Ordering in Distributed Constraint Satisfaction Problems. In *Australian Conference on Artificial Intelligence*, pages 427–439, 2003.