

Time-Ordered Ad-hoc Resource Sharing for Independent Robotic Agents

Arjo Chakravarty¹, Michael X. Grey², M. A. Viraj J. Muthugala³ and Mohan Rajesh Elara³

Abstract—Resource sharing is a crucial part of a multi-robot system. We propose a Boolean satisfiability problem (SAT) based approach to resource sharing. Our key contributions are an algorithm for converting any constrained assignment to a weighted-SAT based optimization. We propose a theorem that allows optimal resource assignment problems to be solved via repeated application of a SAT solver. Additionally we show a way to encode continuous time ordering constraints using Conjunctive Normal Form (CNF). We benchmark our new algorithms and show that they can be used in an ad-hoc setting. We test our algorithms on a fleet of simulated and real world robots and show that the algorithms are able to handle real world situations. Our algorithms and test harnesses are open source and build on Open-RMF’s fleet management system.

I. INTRODUCTION

Multi-robot systems is a well studied field with a variety of sub problems. There has been extensive study of robot communications, joint exploration strategies [1], multi-agent path finding [2] [3] and task allocation [4]. Open-RMF is an open source framework that facilitates multi-fleet orchestration [5]. It comes with tools for handling problems such as traffic management, map alignment and optimal task allocation [4]. Robotics researchers have spent much effort on optimizing multi-robot path planning [2] [6], task assignment [4] [7] and multi-robot localization [8]. However, as we move these systems to production one problem that arises is ad hoc (on-demand) physical resource contention¹. Suppose two or more robots intend to use the same item like a parking spot, charger, or an elevator. There needs to be a way for their use of the resources to be orchestrated. This paper proposes algorithms for a reservation system built upon the notion of optimizing the assignment of resources while taking into account individual costs for robots using a Boolean satisfy-ability solver (SAT). When combined with traffic deconfliction, it allows lifelong operation of a multi-robot system. Our primary contributions in this paper are:

This research is supported by Intrinsic llc., the National Robotics Programme under its National Robotics Programme (NRP) BAU, Ermine III: Deployable Reconfigurable Robots, Award No. M22NBK0054, A*STAR under its “RIE2025 IAF-PP Advanced ROS2-native Platform Technologies for Cross sectorial Robotics Adoption (M21K1a0104)” programme, and also supported by SUTD Growth Plan (SGP) Grant, Grant Ref. No. PIE-SGP-DZ-2023-01.

¹Arjo Chakravarty is with Intrinsic LLC and Singapore University of Technology and Design, Singapore. arjoc@intrinsic.ai, arjo.chakravarty@mymail.sutd.edu.sg

²Michael X. Grey is with Intrinsic LLC mxgrey@intrinsic.ai

³M. A. Viraj J. Muthugala and Mohan Rajesh Elara are with Singapore University of Technology and Design, Singapore. {viraj-jagathpriya, rajeshelara}@sutd.edu.sg

¹<https://github.com/open-rmf/rmf/discussions/83#discussioncomment-1123844>

- A Conjunctive Normal Form (CNF) formulation of the robot resource allocation problem
- A greedy algorithm for optimizing resource assignment
- An algorithm for converting any assignment problem and its SAT formulation into a cost optimization problem.
- A CNF encoding of the scheduling problem
- An open source package integrated with ROS 2 and Open-RMF that handles resource assignment

The rest of this paper is organised as follows: Section II covers a background literature review. Section III introduces the resource allocation problem for fixed time ranges. Section IV compares a weighted SAT based approach to a greedy approach. Section V proposes and compares two algorithms: one based on SAT and another greedy method. Section VI extends the CNF formulation introduced in Section III to support time ranges instead of fixed time and shows that it is faster than discretization. Finally, Section VII describes an experimental validation done in simulation and on physical robots.

II. BACKGROUND

One straightforward method to solve the resource assignment problem is to check if the robot’s destination is currently occupied and wait for it to become available. This overly simplistic approach easily leads to deadlock, for example, if two robots each want to go to the location currently occupied by the other, then both are stuck waiting on the other to move.

The classic algorithm for resource assignment is the Kuhn-Munkres method [9]. This method assumes there are n resources and m requesters. Each requester assigns a cost for a given resource, e.g. the distance the requester would need to travel from their current location to arrive at each resource. The algorithm will converge on the optimal set of resource assignments in a matter of $O(n^3)$ time [9]. This method however does not account for scenarios where it is necessary for multiple requesters to take turns with one resource as is often the case in a realistic deployment. Additionally it is not possible to encode constraints into the Kuhn-Munkres algorithm. This makes it unsuitable for our use as we often need to encode constraints like “This robot must have access to a charger within the next two hours”.

It can be shown that the robot resource assignment problem has equivalent complexity to the Travelling Salesman Problem (TSP) making it NP-hard [7]. Several approaches for solving such problems exist such as using Linear Programming (LP) [7] or Constraint Programming (CP). Integer

linear programming (ILP) is heavily favoured by many TSP solvers [10], however fast ILP solvers such as Gurobi [11] and C-PLEX [12] remain proprietary. Google’s open source OR-tools use constraint programming solvers for such NP-hard optimizations [13].

The approach we favour is similar to the approach used by OR-tools as there are many very mature open source SAT solvers available for use [14]. The Conflict Driven Clause Learning algorithm allows SAT solvers to learn clauses which are impossible enabling us to limit the search space drastically [15] [16] [17]. It is easy to encode new constraints using CNF as compared to LP. This means that we know if a set of requested inputs is impossible even before trying to optimize the constraints. Similar approaches have been used by others, for instance Imeson and Smith use SAT and TSP solvers to solve integrated task and motion planning [17]. However, our approach is much simpler.

A form of the SAT problem which can encode optimization is the Weighted MaxSAT problem [18]. In the MaxSAT problem there are a few hard and a few soft constraints. The hard constraints are ones which cannot be violated, while the cost of the soft constraints is minimized with weights [18]. This is not unlike the problem we have at hand—in fact the formulation we present is one that could potentially be solved using MaxSAT solvers. However, for our formulation we do not have soft constraints thus significantly reducing the complexity and eliminating the need for MaxSAT.

Often CP-SAT solvers encode the costs of individual assignments themselves as CNF formulas [13]. This leads to the conundrum where floating point numbers cannot be represented and only integers can be used. In most cases this is rarely a problem as we can just multiply the costs by an arbitrarily large number [13]. For our problem however, thanks to the fact that only one alternative per request needs to be assigned (see section III for definition of alternative), we can perform some simple tricks to speed up the search and limit our search space when dealing with assignment type problems. Additionally, we show that it is possible to express such floating point constraints in terms of total orders, thus enabling conventional SAT solvers to reason about them.

III. PROBLEM STATEMENT

The problem we wish to solve is one where a group of heterogeneous robots request access to a certain set of resources. For instance, robots may need to use charging stations, parking spaces, or elevators. When a robot needs to utilize some kind of resource, it proposes a set of acceptable alternative choices along with the cost of each alternative. For example if a robot needs to charge its battery and is compatible with three different charging stations, it would request an assignment to one of the three stations, listing the cost of going to each station from its current location (e.g. time or distance to travel).

A. Definitions

Given a set of resource requests, we assign resources across those requests based on the available alternatives in a way that minimizes the overall cost without having any assignments that overlap in time.

- Each request can be defined as R_i .
- A request has n alternatives $\alpha_{i,j} = (s_{i,j}, r_{i,j}, d_{i,j}, c_{i,j})$ where $s_{i,j}, d_{i,j}, c_{i,j} \in \mathbb{R}$. Each alternative is made of a start time s , a resource r , a cost c and a duration d .
- An alternative $\alpha_{i,j}$ conflicts with another alternative $\alpha_{k,m}$ if it shares the same resource (i.e. $r_{i,j} = r_{k,m}$) and it overlaps in time with the other alternative.

B. Simplification to CNF

The proposed problem can be simplified into a CNF. Suppose the variable $x_{i,j}$ represents the fact that the alternative $\alpha_{i,j}$ is awarded. We have two constraints:

First only one alternative from a request R_i with n alternatives can be awarded. This gives us the clauses shown in (1) and (2)

$$(x_{i,0} \vee x_{i,1} \vee x_{i,2} \dots \vee x_{i,n}) \quad (1)$$

$$\bigwedge_{j,k} (\neg x_{i,j} \vee \neg x_{i,k}), \text{ where } j < k \quad (2)$$

Secondly, no items with conflicts can be assigned. So for every pair of conflicts $\alpha_{i,j}, \alpha_{i,m}$ we get the clause (3)

$$(\neg x_{i,j} \vee \neg x_{i,m}) \quad (3)$$

IV. COST OPTIMIZATION

A. Weighted-SAT Based Approach

While SAT can provide a feasible solution, it does not provide an easy way to calculate cost. Repeated application of SAT is a common technique used when dealing with such optimizations [18]. In our case we can formulate our optimization in terms of the boolean variables as shown in (4) where $x_{i,j} \in \{0, 1\}$ and $c_{i,j} \in \mathbb{R}$.

$$\arg \min_{x_{i,j}} \sum_{i,j} x_{i,j} c_{i,j} \quad (4)$$

One of the simplest ways to solve this problem is to ask the SAT solver for SAT assignment, then negate the assignment and add it as a clause to the original. We keep doing this until the SAT solver returns "unsatisfiable" (UNSAT). At this point we can confirm that there are no more alternative assignments. Such a naive brute-force approach does not scale very well given a scenario where there are many feasible solutions. Even when we come across the optimal solution, we would need to continue performing an exhaustive search to eliminate all possible alternatives before we can prove that the candidate solution is optimal.

There is additional information in the problem domain that we can exploit to define necessary and sufficient conditions for determining whether a feasible solution is optimal without an exhaustive search. This motivates the following Lemma:

Lemma 4.1: Given an assignment $A_1 = (x_{0,j}, x_{1,l}, \dots, x_{n,z})$ if a cheaper solution A_2 exists, then $\exists x_{i,j} \in A_2$ such that given $x_{i,k} \in A_1$, $c_{i,j} < c_{i,k}$.

Proof: (by contradiction)

Suppose not. Given that there exist two solutions A_2 and A_1 where A_2 is cheaper than A_1 . Then from the contradiction it follows that $\forall i, i < n$, given $x_{i,j} \in A_2$ and $x_{i,k} \in A_1$, $c_{i,j} \geq c_{i,k}$. This leads to the total cost of the solution $cost(A_2) = c_{0,j_1} + c_{1,j_2} + \dots + c_{n,j_n} \geq c_{0,k_1} + c_{1,k_2} + \dots + c_{n,k_n} = cost(A_1)$. But this is a contradiction because we had defined A_2 to be cheaper than A_1 . ■

Algorithm 1 uses Lemma 4.1 to formulate a new clause every time a solution is found. Given a SAT solution we add a clause to the formula that requires at least one of the cheaper alternatives to be true. This guides the solver towards the optimal solution faster than a brute force search. Additionally, as a result of Lemma 4.1 we immediately know if a solution is optimal. This is because if a cheaper solution exists it must satisfy Lemma 4.1. Thus if we get an UNSAT solution from the SAT solver, we know that a cheaper solution cannot exist and our previous assignment is optimal.

Fig. 1 shows how this heuristic greatly reduces computation time for problems. Both a Naive brute-force approach and Algorithm 1 were made to solve 100 problems with a large number of feasible alternatives. The heuristic based approach is able to handle larger problems relatively well compared to the naive SAT approach which quickly becomes too slow. One of the nice features of incremental optimization methods with SAT is that we are able to quickly obtain solutions which are sub-optimal but feasible, which is useful in real world scenarios when having an effective solution quickly is more important than obtaining the optimal solution eventually.

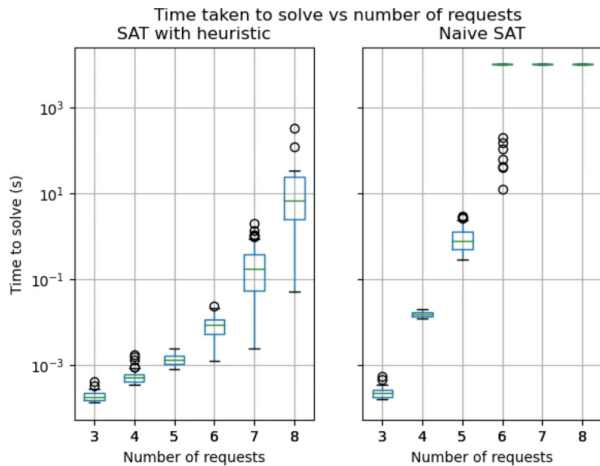


Fig. 1. Box-plot of Naive SAT based search vs heuristic driven search. Both solvers were given 100 hard optimization problems of different sizes. Benchmark was run on a 24-core GCP instance with code written in Rust. Each request has 10 alternatives.

Algorithm 1 SAT Cost Optimizer with Assignment Heuristic

Require: An incremental SAT solver

$literals \leftarrow$ Generate SAT literals from requests

$clauses \leftarrow \{\}$

$best_cost \leftarrow \infty$

$best_solution \leftarrow \emptyset$

$OPEN \leftarrow \{clauses\}$

while $OPEN$ is not empty **do**

$current_clauses \leftarrow$ pop an element in $OPEN$

$s \leftarrow$ Call Incremental SAT with $current_clauses$

if s is UNSAT **then**

 continue.

end if

if $cost(s) < best_cost$ **then**

$best_cost \leftarrow cost(s)$

$best_solution \leftarrow s$

$new_clause \leftarrow ()$

 // Create clause generated by Lemma 4.1

for all x_{ij} in s **do**

for all $cost(x_{ik}) < cost(x_{ij})$ **do**

$new_clause \leftarrow new_clause \vee x_{ik}$

end for

end for

$new_clauses \leftarrow clauses \wedge new_clause$

$new_clauses \leftarrow new_clauses \wedge \neg s$

else

$new_clauses \leftarrow clauses \wedge \neg s$

end if

$OPEN \leftarrow new_clauses$

end while

B. Alternate Solution to Optimization via a Greedy Conflict Driven Approach

We can also find an optimal solution using a greedy search (see Algorithm 2). The approach relies on greedily picking the lowest cost alternatives for each reservation, then when a conflict arises we branch the search based on the conflict. For instance if $\alpha_{m,n}$ and $\alpha_{j,k}$ conflict with each other we branch on the fact that either $\alpha_{m,n}$ must belong to the solution and hence eliminate $\alpha_{j,k}$, or $\alpha_{j,k}$ must belong to the solution and hence eliminate $\alpha_{m,n}$, or neither belong to the solution so both should be eliminated.

This gives us a branching factor of 3 where the frequency of branching comes from the number of conflicting alternatives. Therefore the more conflicts there are, the longer the greedy search will take. Additionally, if a solution is unsatisfied, the optimizer will examine every possible branch before it terminates. This combinatorial approach will not scale well with an increasing number of conflicting requests and alternatives.

C. Understanding the weaknesses of each algorithm

Algorithm 1 and Algorithm 2 are interesting duals of each other when we consider the worst case performance of each algorithm. Fig. 2 shows an adversarial example for Algorithm 1. Here, the solver has many alternatives to choose

Algorithm 2 Greedy Conflict-Driven Optimizer

```

 $B_0 \leftarrow ()$ 
 $S_0 =$  lowest cost assignment  $a_{i,j} \forall$  requests  $R_i$ 
 $OPEN \leftarrow \{(S_0, B_0)\}$ 
while  $OPEN$  is not empty do
   $(solution, B_0) \leftarrow$  pop lowest cost assignment in  $OPEN$ .
  check for conflicts in  $solution$ 
  if no conflicts found then
    return  $solution$ 
  end if
  for all  $c$  in conflicts do
     $(\alpha_{i,j}, \alpha_{k,m}) = c$ 
     $B_1 \leftarrow \neg\alpha_{i,j} \wedge B_0$ 
     $S_1 =$  lowest cost assignment  $a_{i,j} \forall$  requests  $R_i$ 
    such that  $B_1$  is satisfied
     $B_2 \leftarrow \neg\alpha_{k,m} \wedge B_0$ 
     $S_2 =$  lowest cost assignment  $a_{i,j} \forall$  requests  $R_i$ 
    such that  $B_2$  is satisfied
     $B_3 \leftarrow \neg\alpha_{i,j} \wedge \neg\alpha_{k,m} \wedge B_0$ 
     $S_3 =$  lowest cost assignment  $a_{i,j} \forall$  requests  $R_i$ 
    such that  $B_3$  is satisfied
     $OPEN \leftarrow OPEN \cup (S_1, B_1) \cup (S_2, B_2) \cup (S_3, B_3)$ 
  end for
end while

```

from. Depending on where the first iteration lands, the worst case scenario is that the SAT solver starts from the highest cost request and slowly works its way down through every possible feasible alternative. In contrast, Algorithm 2 would find the solution at its first search node.

Fig. 3 shows an example which would make the greedy algorithm (Algorithm 2) have trouble producing a solution since the overlapping times for the alternatives create numerous conflicts and therefore many branches.

In an extreme scenario where no feasible assignments can be found to simultaneously satisfy all requests, the greedy algorithm would end up doing an exhaustive search, branching at every conflict before determining that no solution exists. On the other hand, the SAT solver in Algorithm 1 would return UNSAT in the very first iteration.

It is also possible to construct an adversarial example with a feasible solution. One could create a problem where the only viable solutions have high costs, and then sprinkle in a bunch of lower cost alternatives that all lead to irreconcilable

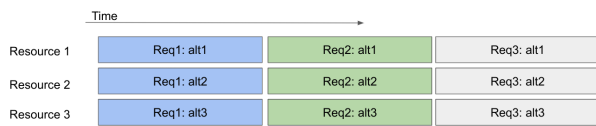


Fig. 2. Example of a scenario which elicits a worst case response in the SAT optimization process but is quickly solved by the greedy process in $O(n)$

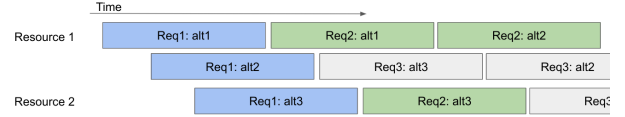


Fig. 3. Example of a scenario which elicits a worst case response in the Greedy process but is quickly solved by the SAT optimization process. Assume that alt3 has lowest cost.

conflicts. This would cause the greedy algorithm to examine the lower cost alternatives first, only to discover that they are infeasible. The SAT based solver on the other hand would quickly identify a feasible solution.

The incremental nature of Algorithm 2 makes it appealing for ad hoc request scenarios because it can rapidly find a feasible solution and then converge towards the optimal solution with however much time the overall system is willing to budget. When the multi-robot system needs to handle unknown requests coming in at unpredictable times, having any viable solution within a time frame that keeps the overall system responsive is more important than eventually finding the optimal solution. In our benchmarks, we found that the SAT method arrives at the first feasible (but sub-optimal) solution within less than a second for 40 requests with 40 alternatives. Fig. 4 shows the time to first solution of the SAT solver on a Thinkpad P50 with 64GB RAM and a 16-core 11th Gen Intel(R) Core(TM) i7-11850H @ 2.50GHz. This shows that the algorithm is well suited for ad hoc requests.

D. Experimental Benchmarks

We created multiple benchmark sets based on the scenarios described in the section IV-B. They were run single-threaded on a 24-core AMD EPYC 7B12 processor with 96GB of RAM running in Google Cloud. The algorithms were implemented in Rust [19] and the SAT solver used was the VariSAT solver [20]. The first set of benchmarks were based on a scenario where there are many feasible solutions.

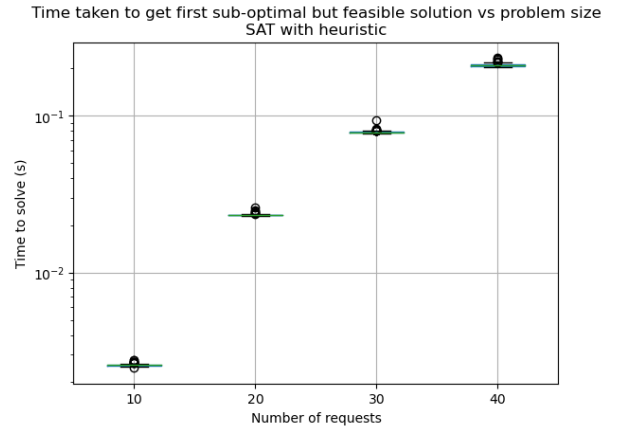


Fig. 4. Box-plot of time taken for SAT with heuristics to generate first feasible but suboptimal solution. For a request size of n , there are n alternatives. I.E for 40 requests there are 40 alternatives each.

This results in a more optimization heavy workload. In such instances, the greedy method reaches optimality faster than repeated application of SAT with heuristics. As shown in Fig. 5, as the problem size increases, the performance of the SAT method with heuristics falls far behind the greedy method for problems with simple solutions. Conversely, as the number of conflicts increases, the greedy method fairs worse even with a small number of conflicts (see Fig. 6) while the SAT based methods are able to solve large instances in a reasonable amount of time. Fig. 7 shows that the performance of the Greedy Method significantly degrades much faster than SAT with number of conflicts.

V. EXTENDING THE SYSTEM TO ENCODE TIME RANGES

In a reservation system for robots, it is important to handle scenarios that would require robots to take turns using the same resource. The simple fixed time scheme does not lend itself to solving such problems. For example if two robots each place a request to use the same resource at the same

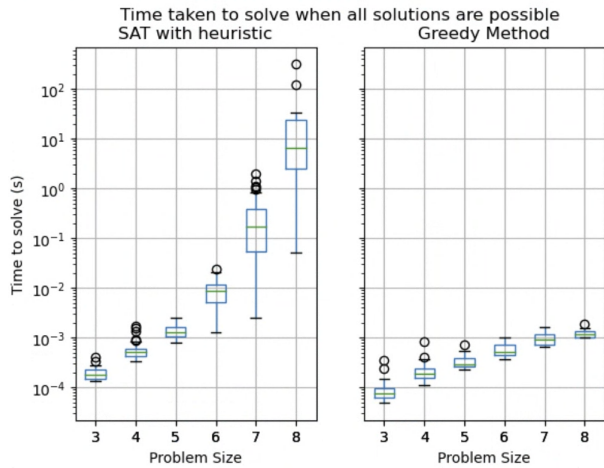


Fig. 5. Box plot showing performance comparison when many options exist and no conflicts exist. Problem size n refers to number of requests. Each request also has n alternatives

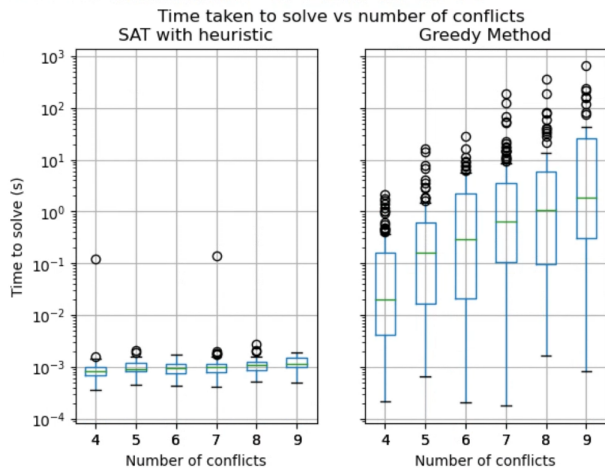


Fig. 6. Box plot of performance vs number of conflicts. Here there were 10 requests submitted each with 5 alternatives.

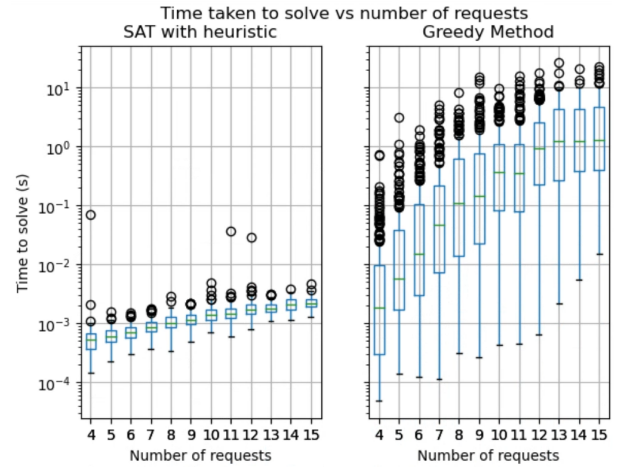


Fig. 7. Box plot of performance comparison vs requests in scenario with fixed number of conflicts. Here there were 10 alternatives per option.

time with no other alternatives, then the fixed time scheme would simply return UNSAT.

This means that we should support a time window in which each reservation alternative may begin. For this, there are two possible solutions. The first is to use the formulation in section III-A and generate a sequence of alternatives for each resource at discrete time intervals. The second option would be to encode ordering into the SAT formulation. The latter has the advantage that a sub-optimal solution can be found very quickly, and an unsatisfiable scenario can be found more quickly thanks to the fact that the starting time range not affecting the formula length (see Section V-B for more details). Additionally, there is no need to tune the time resolution to which we discretize.

A. CNF encoding of scheduling constraints to create a model

We can extend our previous formulation to add support for varying start times. Let's revise our original problem statement.

1) Definitions:

- Each request can be defined as R_i .
- A request has n alternatives $\beta_{i,j} = (s_{i,j}, l_{i,j}, r_{i,j}, d_{i,j}, c_{i,j})$ where $s_{i,j}, l_{i,j}, d_{i,j} \in \mathbb{R}$. Each alternative is made of an earliest start time s , a latest start time l , a resource r , a cost c , and a duration d .
- An alternative $\beta_{i,j}$ cannot co-exist with alternative $\beta_{k,m}$ if it shares the same resource (i.e. $r_{i,j} = r_{k,m}$) and the earliest end time of $\beta_{i,j}$ (given by $s_{i,j} + d_{i,j}$) is greater than the latest start time of $\beta_{k,m}$ (given by $l_{k,m}$).

The CNF decomposition is similar to the previous decomposition in (1) and (2), except we introduce a new set of variables X_{ijkm} to impose ordering. We say X_{ijkm} is true if $\beta_{i,j}$ starts after $\beta_{k,m}$. Based on this definition we will imply a strict total order on all X_{ijkm} within a given resource. Strict total orders generally have four properties: irreflexivity, anti-symmetry, transitivity and connectedness [21]. We do not need to consider irreflexivity as that is something that a SAT

solver already encodes. To encode anti-symmetry we start with the clause (5):

$$(x_{ij} \wedge x_{km} \wedge X_{ijkm}) \implies \neg X_{kmi j} \quad (5)$$

We can transform (5) into:

$$\neg((x_{ij} \wedge x_{km} \wedge X_{ijkm})) \vee \neg X_{kmi j} \quad (6)$$

Which by De’Morgan’s Law simplifies to:

$$\neg x_{ij} \vee \neg x_{km} \vee \neg X_{ijkm} \vee \neg X_{kmi j} \quad (7)$$

We also want to encode transitivity so we assume:

$$(X_{ijkm} \wedge X_{kmnl}) \implies X_{ijnl} \quad (8)$$

Which gets encoded into

$$\neg X_{ijkm} \vee \neg X_{kmnl} \vee X_{ijnl} \quad (9)$$

We only need connectedness if both x_{ij} and x_{km} are within the same resource and both x_{ij} and x_{km} are true. This can be given by the boolean formula below:

$$(x_{ij} \wedge x_{km}) \implies (X_{ijkm} \vee X_{kmi j}) \quad (10)$$

Which simplifies to:

$$(\neg x_{ij} \vee \neg x_{km} \vee X_{ijkm} \vee X_{kmi j}) \quad (11)$$

Finally we add all known literals of X_{ijkm} . The cases where we know X_{ijkm} are

- Iff $\beta_{i,j}$ cannot be scheduled after $\beta_{k,m}$ (i.e. $l_{i,j} < s_{k,m} + d_{k,m} \wedge s_{i,j} + d_{i,j} > s_{k,m}$), then we add the literal $\neg X_{ijkm}$.

Note that this set of constraints only provides us with a necessary condition. They are not sufficient to prove that an assignment is valid. One still needs to traverse the order determined by the SAT solver and find if the order is valid. In particular our SAT encoding does not take duration into account outside of the indefinite case. Here we can use learning to identify new clauses as defined by Algorithm 3. These new clauses will allow us to constrain the output of the SAT solver to only include assignments that satisfy the start time constraint of each alternative while taking their durations into account.

The key idea in Algorithm 3 is that certain combinations of assignments would violate the latest start time constraints of one or more alternatives. Using the X_{ijkm} values, we topologically sort the alternatives that would use the same resource within a candidate solution, and then we calculate what the concrete start time values would be for that combination of assignments. The candidate solution becomes invalid if any of the assignments has been pushed out beyond its latest start time constraint. In this case we retrace which assignments were responsible for causing the unacceptable delay and mark that combination as illegal. This allows us to further limit the search space such that we converge on a valid solution.

Algorithm 3 Conflict detection and clause learning for Resource Ordering

```

let  $G(V, E)$  be the graph formed where  $V$  is the set of all the  $x_{ij}$  assigned to be true and in resource  $r^*$ , and  $E$  is the set of all true  $X_{ijkm}$ .
sorted_assignment  $\leftarrow$  topological_sort( $G$ )
schedule  $\leftarrow$  {}
banned_clauses  $\leftarrow$  {}
for all  $i$  in sorted_assignment.len() do
   $s_i =$  sorted_assignment[ $i$ ].earliest_start
   $l_i =$  sorted_assignment[ $i$ ].latest_start
  if  $s_i >$  schedule.last.end_time() then
    schedule.insert_at( $s_i$ , sorted_assignment[ $i$ ])
  else if  $l_i >$  schedule.last.end_time() then
     $t \leftarrow$  schedule.last.end_time()
    schedule.insert_at( $t$ , sorted_assignment[ $i$ ])
  else
    banned_clause  $\leftarrow$  ()
     $j =$  schedule.len()-1
    while schedule[ $j$ ].start_time() == schedule[ $j-1$ ].end_time() do
      banned_clause  $\vee =$   $\neg$  assignment[ $i$ ]
       $j \leftarrow j - 1$ 
    end while
    banned_clauses.push(banned_clause)
  end if
end for
return banned_clauses

```

B. Comparing Discretization vs Continuous representation

An alternative to the formulation described in equations (7)-(9) is to discretize each order and check for satisfiability using the conditions described in (3). The first thing that merits discussion is the order of growth of clauses. In the case of discretization, we find that the number of literals is proportional to the interval of each alternative (t) multiplied by the number of resources to consider (n). The order of growth of the clauses is then $O(n^2 t^2)$. On the other hand, the continuous representation is independent of t . However the transitivity constraints result in $O(n^2)$ literals and $O(n^3)$ clauses. Discretization leads to the added problem that assigned reservations may have gaps between them. In practice this can be fixed with a post processing pass.

Fig. 8 shows the start time range has a significant impact on the time to solve a fixed-timestep discretization based solver, whereas the continuous solver is able to find solutions several orders of magnitudes faster and is virtually unaffected by the start time range.

Note that for cost optimization Algorithm 1 still applies with some limitations. The cost cannot depend on time when using the continuous formulation, whereas the discretized time representation can encode different costs into alternatives that start at different times.

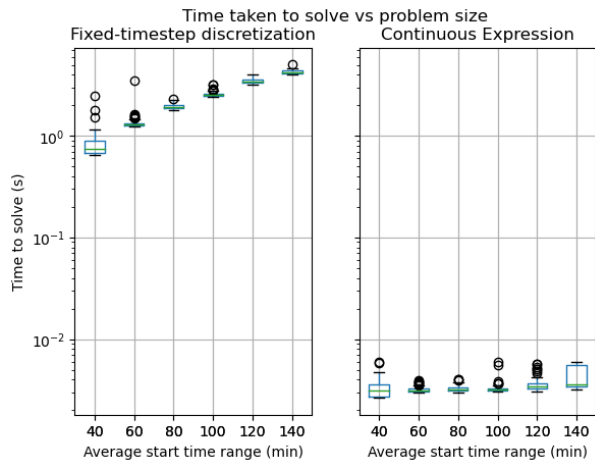


Fig. 8. Discrete vs continuous expression: 5 requests with 5 alternatives were given. The discretization time step was 10 minutes. The algorithms were run till the first feasible alternative is found

VI. VALIDATION

We ran several scenarios in Gazebo [22] for 24 hours in simulated time to check that no deadlocks occur. The first scenario involved 2 robots periodically making a request to the same location to simulate “charging”. Fig. 9 shows an example of conflicting requests made. We tried similar simulations with 3 robots in the same world. The system would automatically re-order the robots such that no conflict occurs. We use Open-RMF’s office world from the `rmf_demos` package to validate our algorithm’s behaviour. We also had scenarios where random requests were made for robots to go to different waypoints. The algorithms and test harnesses were implemented in ROS 2 [23] with Rust [19] and can be found on GitHub².

Additionally, we tested our algorithm for 30 minutes on two Smorphi robots³ shown in Fig. 10. In this test, the two robots repeatedly requested the same spot. The robot would request the resource (in this case the position in the center of the L bend), claim its assignment from the reservation

²https://github.com/open-rmf/rmf_reservation

³<https://www.wefaarobotics.com>



Fig. 9. Example of simulated world. The arrows show requested trajectory in an example scenario. The reservation system would re-order the requests such that no two robots would conflict with each other for use of the pantry.

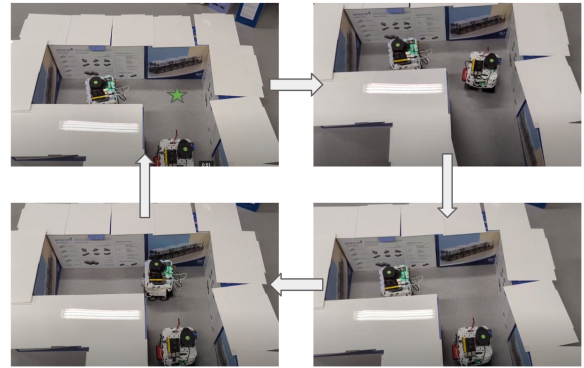


Fig. 10. Example of physical demo using two Smorphi robots. The green star shows the contested resource.

system, and then proceed to the position when permitted. In order to ensure that no other robot can move in to the same spot, the claim held by the robot must be released before any other claimants are permitted to go. The system releases the claim once the robot moves away from the claimed position. The reservation system would then give permission to the next claim. This process of requesting, claiming, waiting for permission, proceeding, and then returning to their start points cycled for 30 minutes without any deadlocks.

VII. CONCLUSIONS AND FUTURE WORK

This paper proposed a novel algorithm for an ad-hoc reservation system that optimizes the assignment of requested resources while accounting for the individual costs to the requesters. We race a SAT solver against a Greedy solver since the performance profile of each method is complementary. Additionally we have shown how to encode order into the boolean satisfiability problem and shown how the constrained assignment problem can be optimized via repeated application of a SAT solver.

Our benchmarks, simulations and real world demonstration show the feasibility of using the SAT based algorithms for resource sharing. Of particular importance are the heuristics developed for assignment and the encoding developed for total ordering.

As future work, we would like to examine how the costs of assignments produced by the algorithm in Section V can depend on time or some other arbitrary function or are we forced to use discretization to account for the cost of delaying an assignment. We also only look at discretization at fixed time steps, however we may be able to minimize computation time if we were smart about the way we sampled possible assignments. Other interesting ideas include extending the CNF formulations here to more complex planning schemes where one reservation might depend on which of the previous reservations was awarded. Finally are there ways of automatically discovering theorems that can rapidly guide optimizations.

REFERENCES

- [1] M. Dharmadhikari, H. Nguyen, F. Mascari, N. Khedekar, and K. Alexis, "Autonomous cave exploration using aerial robots," in *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2021, pp. 942–949.
- [2] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, vol. 219, pp. 40–66, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0004370214001386>
- [3] J. Kottinger, S. Almagor, and M. Lahijanian, "Conflict-based search for multi-robot motion planning with kinodynamic constraints," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 13 494–13 499.
- [4] B. Gerkey and M. Mataric, "Sold!: auction methods for multirobot coordination," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 758–768, 2002.
- [5] "Open-rmf," <https://github.com/open-rmf/rmf>.
- [6] R. Stern, N. Sturtevant, A. Felner, S. Koenig, H. Ma, T. Walker, J. Li, D. Atzmon, L. Cohen, T. Kumar, *et al.*, "Multi-agent pathfinding: Definitions, variants, and benchmarks," in *Proceedings of the International Symposium on Combinatorial Search*, vol. 10, no. 1, 2019, pp. 151–158.
- [7] C. Nam and D. A. Shell, "Assignment algorithms for modeling resource contention in multirobot task allocation," *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 3, pp. 889–900, 2015.
- [8] S. Yu, C. Fu, A. K. Gostar, and M. Hu, "A review on map-merging methods for typical map types in multiple-ground-robot slam solutions," *Sensors*, vol. 20, no. 23, p. 6988, 2020.
- [9] J. Munkres, "Algorithms for the assignment and transportation problems," *Journal of the society for industrial and applied mathematics*, vol. 5, no. 1, pp. 32–38, 1957.
- [10] W. Cook, *In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation*. Princeton University Press, 2014. [Online]. Available: <https://books.google.com.sg/books?id=UmuYDwAAQBAJ>
- [11] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023. [Online]. Available: <https://www.gurobi.com>
- [12] I. I. Cplex, "V12. 1: User's manual for cplex," *International Business Machines Corporation*, vol. 46, no. 53, p. 157, 2009.
- [13] CPAIOR, "Cpaior 2020 master class: Constraint programming," Sept. 2020. [Online]. Available: <https://youtu.be/lmy1ddn4cyw>
- [14] T. Balyo, M. Heule, M. Iser, M. Järvisalo, and M. Suda, "Proceedings of sat competition 2023: Solver, benchmark and proof checker descriptions," 2023.
- [15] J. Marques-Silva and K. Sakallah, "Grasp: a search algorithm for propositional satisfiability," *IEEE Transactions on Computers*, vol. 48, no. 5, pp. 506–521, 1999.
- [16] L. Zhang, C. F. Madigan, M. H. Moskewicz, and S. Malik, "Efficient conflict driven learning in a boolean satisfiability solver," in *IEEE/ACM International Conference on Computer Aided Design. ICCAD 2001. IEEE/ACM Digest of Technical Papers (Cat. No. 01CH37281)*. IEEE, 2001, pp. 279–285.
- [17] F. Imeson and S. L. Smith, "Multi-robot task planning and sequencing using the sat-tsp language," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 5397–5402.
- [18] F. Heras, J. Larrosa, and A. Oliveras, "Minimaxsat: An efficient weighted max-sat solver," *Journal of Artificial Intelligence Research*, vol. 31, pp. 1–32, 2008.
- [19] N. D. Matsakis and F. S. Klock II, "The rust language," in *ACM SIGAda Ada Letters*, vol. 34, no. 3. ACM, 2014, pp. 103–104.
- [20] "Varisat," <https://github.com/jix/varisat>.
- [21] S. Epp, *AISE DISCRETE MATHEMATICS WITH APPLICATIONS*. Cengage Learning, 2011.
- [22] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, vol. 3, 2004, pp. 2149–2154 vol.3.
- [23] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, "Robot operating system 2: Design, architecture, and uses in the wild," *Science Robotics*, vol. 7, no. 66, p. eabm6074, 2022. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abm6074>