

A General Model for Wildfire Suppression Problems

Gustavo Delazeri

Universidade Federal do Rio Grande do Sul
Porto Alegre, RS - Brazil
gustavo.delazeri@inf.ufrgs.br

Marcus Ritt

Universidade Federal do Rio Grande do Sul
Porto Alegre, RS - Brazil
marcus.ritt@inf.ufrgs.br

ABSTRACT

Wildfires are an important problem of our time, causing substantial economic, ecological, and infrastructural damage. The operations research community has long been involved in developing decision support models for fire suppression, and modeling approaches vary greatly. An important class of models is based on encoding a landscape using a graph and integrating fire propagation and suppression strategies within the optimization process. Several works in the literature describe graph-based models, which differ in the assumptions about fire propagation, the objective function, and the features considered. In this work, we propose a new mixed-integer programming (MIP) model that includes various features from previous research. In experiments, we compare this model against a more specialized model using benchmark instances from the literature. The results indicate that the proposed model is easier to solve and performs competitively against more sophisticated algorithms.

KEYWORDS. Wildfire suppression. Mixed-integer programming. Exact algorithms.

Mathematical programming. Combinatorial Optimization.

1. Introduction

Wildfires are increasingly destructive due to climate change, causing severe environmental and economic damage, as seen recently in Brazil and Greece (Paddison, 2024; Kuebler, 2024). The cost of fighting these fires is substantial, often exceeding billions of dollars annually (Joint Economic Committee, U.S. Senate, 2023). Decision support models help optimize firefighting strategies by improving resource allocation and reducing costs in wildfire management.

The operations research (OR) community has been interested in the combat and prevention of forest fires for a long time (Martell, 1982), and several OR models have been proposed to support a series of decisions relevant to forest-fire management. Modeling techniques vary greatly. In demand-based models, for example, decisions are related to where to locate fire suppression resources to cover a set of demand points or how to sequence fire suppression operations to minimize the value lost in the fire. In fireline-based models, a predefined fire perimeter evolves over time, and the objective is to allocate resources to construct a fireline that contains the perimeter.

Graph-based (or grid-based) models are a class of models that try to represent the space where a fire is spreading. The idea is to discretize the space into basic units and represent adjacency between basic units using a graph. The arcs connecting these basic units are labeled with fire propagation times, which can be estimated by using Geographic Information System data and an appropriate fire spread model (Rothermel, 1972). An advantage of grid-based models is their ability to simultaneously model the interaction between fire spread and suppression activities within the optimization process. This integrated approach may lead to more realistic and effective fire management strategies. There are several works in the literature proposing graph-based models, which differ in the assumptions about fire propagation, the objective function, and the features considered, such as resource availability and safety constraints regarding resource allocation. In this work, we propose a new, comprehensive MIP model that unifies features from previous research. In experiments, we compare the model against a more specialized model on a set of benchmark instances of the literature. Our results indicate that our model is easier to solve.

We close this section with an overview of the paper. Section 2 provides a review on OR methods for fire suppression and gives a more precise description of the type of scenario we are trying to model. In Section 3, we describe the proposed mathematical model. Section 4 presents experiments comparing the proposed model with previous research. Section 5 concludes the work.

2. Literature Review

2.1. Fire suppression

Jewell (1963) is one of the first works to propose the application of OR methods to forest fire problems. Since then, much work has been done on decision support models, such as the coordination of fire crews and the dispatch of fire suppression resources. Dimopoulou and Giannikos (2001), for example, model fire suppression as a demand-covering problem, where we have a set of locations, a set of resources (e.g. firefighting teams) that can be assigned to these locations, and a set of demand points (e.g. potential fire events). Resources must be assigned to locations to cover the demand points. The goal is to find an assignment of resources to locations that maximizes the number of covered points. Rachaniotis and Pappis (2010) model fire suppression as a scheduling problem. Forest fires are treated as jobs processed by a single fire suppression resource. The processing time of a job (i.e. the time to contain the fire) increases with time, and the goal is to find a schedule that maximizes the total value of the unburnt areas remaining after the end of the operation. We refer the reader to Granda et al. (2023) for a detailed survey.

In graph-based models, the idea is to rasterize a landscape and encode it as a directed graph. Each vertex of the graph represents a basic unit of the rasterized landscape, and arcs connecting the vertices are labeled with fire travel times that can be used to model fire propagation. As




Figure 1: Example of an instance. The graph has five vertices and the ignition set is $I = \{s\}$. We have two resources available at times 2 and 3. The graph on the right-hand side depicts a possible allocation, where vertices v_1 and v_3 receive the available resources (vertex v_2 cannot receive a resource since it burns at time 1 and the first resource is released at time 2). Due to the allocation, a delay δ is added to the outgoing arcs of v_1 and v_3 , which will increase the fire arrival time at vertex v_4 . If the objective were to minimize the number of burnt vertices at time 4 and $\delta = 2$, our allocation would have an objective value of 4 (since vertices s, v_1, v_2 , and v_3 burn before time 4).

time progresses, fire suppression resources arrive and can be allocated to the vertices of the graph. Allocating a resource to a vertex delays fire propagation on its outgoing arcs, i.e. we have a local slowdown in fire propagation. The goal is finding an allocation of resources to the vertices of the graph that optimizes some objective, such as the number of unburnt basic units. In the next section, we present a more formal definition of the problem and a review of the literature on this class of models.

2.2. Graph-based models

We are given a directed graph $G = (V, A)$ and a set of ignition vertices $I \subseteq V$. Each arc $uv \in A$ is labeled with a fire propagation time t_{uv} . Graph G and vertices I define a fire propagation process, i.e. a shortest path graph under propagation times t_{uv} . We are also given a set R of fire suppression resources, and each resource $i \in R$ can be allocated to a vertex of the graph to cause a slowdown in fire propagation. Allocating resource i to a vertex $v \in V$ adds a delay δ_i to the outgoing arcs of v . Each resource $i \in R$ is released at time t_i and cannot be allocated to vertices that burned before t_i . The goal is to find a feasible allocation of resources to the vertices of the graph that optimizes a given objective. A common objective, relevant to this work, is minimizing the number of vertices that burn before a target instant (or horizon) H . Figure 1 shows an example.

The previous paragraph described the general setting, and what makes the models different is the objective function, the restrictions imposed on the allocation of resources, and the assumptions on how fire propagates. In the following paragraphs, we summarize some of the models that can be found in the literature.

Hof et al. (2000) introduces a linear programming model in which all resources are available from the beginning, at time 0 (i.e., $t_i = 0$ for $i \in R$), and the goal is to find an allocation that maximizes the earliest fire arrival to a set of target units. In this model, there is a finite number of resources, and it is possible to allocate only a fraction of a resource to a vertex $v \in V$, which will result in only a fraction of the delay to the outgoing arcs of v . Unlike the other work discussed below, the delay introduced by a resource depends on the vertex to which the resource is allocated.

Wei et al. (2011) propose a MIP model where the set of resources R is finite, the delay caused by a resource is the same for all resources, and all resources are available from the beginning. Each vertex $v \in V$ has an associated value w_v which is lost if the vertex burns. The goal is to find an allocation of resources to vertices that minimizes the total loss of value at a target instant H . The authors also consider a version of the problem where the resource release times t_i may be greater than zero. They do not extend the model to account for that, but instead propose two iterative heuristic approaches.

Belval et al. (2015) propose a MIP model where the number of resources is unlimited and allocating a resource to a vertex $v \in V$ completely blocks fire propagation through v ($\delta_i = \infty$ for $i \in R$). All resources are available from the beginning at time 0, but allocating a resource to a vertex incurs a cost that is later added to the objective function (otherwise the optimal solution is trivial). Besides the fire arrival time, the model also computes the fireline intensity at each vertex, which is a measure of how much heat a fire releases. Each vertex has an associated fireline intensity threshold; if the fireline intensity in the vertex exceeds the threshold, the fire is considered harmful, otherwise it is beneficial. Each vertex has an associated value which is lost if the vertex burns. This value is different depending on if the vertex was burnt by a beneficial or harmful fire. Multiple objective functions are studied, which are basically the minimization of different combinations of vertex attributes (e.g. the cost of all allocated resources plus the value of all burnt vertices whose fireline intensity is above its intensity threshold) evaluated at a target instant H .

In Alvelos (2018), the number of resources is finite, and the delay introduced by a resource is always the same. Differently from Hof et al. (2000), Wei et al. (2011) and Belval et al. (2015), the release time t_i of each resource i may be greater than zero. The author proposes several objectives, but relevant to this work is the minimization of the number of burnt vertices at a target instant H .

Avci et al. (2024) propose a MIP model similar to a discrete time process. They consider a finite sequence of time instants $T = (0, 1, \dots, H)$ and, for each $t \in T$ and $v \in V$, a set $C_{tv} \subseteq V$ defines the vertices which will burn at time t if vertex v burns at time $t - 1$. An advantage of this approach is the possibility to model a change in wind direction and intensity as time progresses. In their model, the number of available resources is finite, and the release times may be non-zero (similar to Alvelos (2018)) but, differently from previous works, resources are not homogeneous. They consider three types of resources inspired by real-life fire suppression practice. A *firefighting team* consists of a group of firefighters and a water tanker. During fire combat, the water tanker is positioned at a designated basic unit, while firefighters advance into the forest with hoses to contain the fire at one or more basic units. The position of the water tanker is restricted to basic units that are part of roads and fire can only be suppressed in basic units that are close to the water tanker. Both the water tanker and the firefighters must maintain a safety distance to the fire. A *firefighting helicopter* flies over the landscape and drops water on one or more basic units. This type of resource is similar to the one modeled by Alvelos (2018) in that it does not impose constraints related to safety or spatial locality. Finally, a *firebreak construction team* consists of a group of bulldozers that work on razing basic units to remove their flammable content. Bulldozers can only work on basic units that are adjacent to roads or to previously razed basic units. Like firefighting teams, firebreak construction teams must respect a safety distance to the fire. Each fire suppression resource $i \in R$ is available at time t_i and each basic unit has a value w_v . The goal is to minimize the total loss of value at time H .

2.3. Solution Methods

The model of Alvelos (2018) has received considerable attention recently. Mendes and Alvelos (2023) propose a heuristic algorithm based on iterated local search. The authors introduced a new set of instances and showed that the iterated local search obtains better results than a commer-

cial solver. Harris et al. (2023) noticed that once we define an allocation of resources, the rest of the model can be easily solved by using a shortest path algorithm. They use this observation to propose an exact algorithm based on logic-based Benders decomposition. Their exact algorithm is able to solve all instances used in the work of Mendes and Alvelos (2023) in a few seconds. They propose a more challenging set of instances, and show that the proposed algorithm obtains better results than a commercial solver and the heuristic of Mendes and Alvelos (2023). Finally, Delazeri and Ritt (2024) propose a heuristic algorithm based on beam search and compare it against all the previous works on the instances proposed by Harris et al. (2023). Their results indicate that the beam search algorithm can consistently find the optimal solution in less time than previous methods.

Avci et al. (2024) propose a reformulation of their original MIP using combinatorial Benders cuts. They basically construct a slightly different formulation that has an exponential number of constraints, which are dropped before passing the model to the solver. These constraints are incrementally added to the model whenever the solver finds an integer solution that violates one of them.

3. Mathematical model

In this section we describe the proposed mathematical model. The model combines the minimum travel time approach of previous works such as Alvelos (2018) and Belval et al. (2015), with the heterogeneity of resources, safety constraints, and spatial locality introduced by Avci et al. (2024). Tables 1 and 2 summarize the parameters and the variables of the model. In the subsequent subsections we define the different aspects of the model.

Table 1: Parameters of the model.

Parameter	Description
V	Set of vertices
A	Set of arcs
N_v^+	Successor set of $v \in V$
N_v^-	Predecessor set of $v \in V$
t_{uv}	Fire travel time between $u \in V$ and $v \in N_u^+$
H	Time horizon
$I \subseteq V$	Ignition vertices
R	Set of resources
V_i	Set of vertices that can be the base location of resource $i \in R$
$d(v, u)$	Distance between vertices v and u , $u, v \in V$.
w_v	Value of vertex $v \in V$
t_i	Release time of resource $i \in R$
c_i	Number of vertices that can be protected by resource $i \in R$
r_i	Range of resource $i \in R$ (i.e. maximum distance)
z_i^b	Safety time for the base location of resource $i \in R$
z_i^p	Safety time for the vertices protected by resource $i \in R$
d_i^b	Safety distance for the base location of resource $i \in R$
d_i^p	Safety distance for vertices protected by resource $i \in R$
δ_i	Delay caused by resource $i \in R$

Table 2: Variables of the model.

Variable	Type	Description
a_v	Continuous	Fire arrival time at vertex $v \in V$
Δ_v	Continuous	Delay added to the outgoing arcs of vertex v
y_v	Binary	Indicates if vertex $v \in V$ burns before time H
s_v^i	Binary	Indicates if vertex $v \in V_i$ is the base location of resource $i \in R$
r_v^i	Binary	Indicates if fire in vertex $v \in V$ is being suppressed by resource $i \in R$

3.1. Fire spread

Let a_v be the earliest time vertex v can be reached by fire. Then we have simply, for predecessor set N_v^-

$$a_v \leq a_u + t_{uv} + \Delta_u, \quad v \in V, u \in N_v^-, \quad (1)$$

where Δ_u is added to the outgoing arcs of u by a resource. The fire arrival time is set to $a_u = 0$ for all the ignition vertices $u \in I$. It is known that the constraint set above does not ensure exact fire arrival times, but only upper bounds. For our model, however, this is sufficient.

3.2. Fire suppression resources

One aspect in which the model of Avci et al. (2024) differs from previous work is the enforcement of spatial locality. A firefighting team, for example, can only suppress fire in vertices that are spatially close and close to the position of the water tanker. Inspired by that, we require that each resource $i \in R$ has a fixed position in the landscape, its *base location*, as indicated by the binary variable s_v^i . Different types of resources have different types of constraints regarding their base location. The base location of a firefighting team of Avci et al. (2024), for example, must be a vertex that is part of a road. A firefighting helicopter, on the other hand, may be located at any vertex. The set of vertices that can be the base location of a resource is a parameter $V_i \subseteq V$, for each $i \in R$. Naturally, only one vertex can be the base location of a given resource

$$\sum_{v \in V_i} s_v^i \leq 1, \quad i \in R. \quad (2)$$

Binary variable r_v^i says whether resource $i \in R$ was allocated to vertex $v \in V$. Fire suppression resources have a limited range of action, i.e. they can only protect vertices that are not too distant from their base location. Consequently, if vertex v is protected by resource i , then the distance between v and the base location of the resource is limited by a constant r_i . The constraint below implements this logic, where $d(v, u)$ denotes the distance between any two vertices $u, v \in V$

$$r_v^i \leq \sum_{u \in V_i: d(v, u) \leq r_i} s_u^i, \quad v \in V, i \in R. \quad (3)$$

Each resource has a finite capacity, and we denote by c_i the maximum number of vertices that it can protect, and obtain

$$\sum_{v \in V} r_v^i \leq c_i, \quad i \in R. \quad (4)$$

Following previous work, we also require that each vertex can be protected by at most one resource

$$\sum_{i \in R} r_v^i \leq 1, \quad v \in V. \quad (5)$$

Each resource $i \in R$ also has a release time t_i and can only be allocated to vertices that are not burnt at t_i . Previous work considers safety constraints during resource allocation. Avci et al. (2024), for example, require that suppression activities are conducted at a safety distance from the fire. Our model imposes safety constraints on the base location s_v^i and on the protected vertices r_v^i considering both time and space. More specifically, a vertex $v \in V_i$ can be the base location of a resource $i \in R$ only if the fire arrival time to v and to nearby vertices is greater or equal to $t_i + z_i^b$, where z_i^b is a parameter defining a safety time for the resource. The vertices nearby v are all the vertices (including v) whose distance to v is less than a predefined constant d_i^b . Therefore, we require

$$(t_i + z_i^b)s_v^i \leq a_u, \quad i \in R, v \in V_i, u \in V : d(u, v) \leq d_i^b. \quad (6)$$

We impose similar constraints to the variables r_v^i using parameters z_i^p and d_i^p :

$$(t_i + z_i^p)r_v^i \leq a_u, \quad i \in R, v, u \in V : d(u, v) \leq d_i^p. \quad (7)$$

Each resource $i \in R$ has an associated delay value δ_i , which is added to the outgoing arcs of a protected vertex. The delay Δ_v in a vertex $v \in V$ is then

$$\Delta_v = \sum_{i \in R} \delta_i r_v^i, \quad v \in V. \quad (8)$$

3.2.1. Objective

Finally, binary variable y_v defines whether vertex $v \in V$ is burnt or not at time H

$$y_v \geq 1 - a_v/H, \quad v \in V. \quad (9)$$

Each vertex $v \in V$ has an associated value w_v , and the objective is to minimize the accumulated loss of value at time H

$$\min. \quad \sum_{v \in V} w_v y_v. \quad (10)$$

3.2.2. Limitations & Extensions

The model described is more general than the model of Alvelos (2018). Let R be the set of resources for an instance of their model and T be the set of all time instants where one or more resources are released. We simply set safety times and safety distances to zero and, for every $t \in T$, we create a resource with capacity equal to the number of resources released at t and with a range big enough to not impose any spatial restrictions. Since in the model of Alvelos (2018) there is no concept of base location, we set $V_i = V$ for each resource i . The objective is to minimize the number of burnt vertices at a given time horizon H , so we assign to each vertex a value of one.

As mentioned earlier, our model does not ensure that each variable a_v holds the correct fire arrival time. Following Belval et al. (2015), we could fix that by forcing each non-ignition vertex $u \in V \setminus I$ to have exactly one predecessor $v \in N_u^-$ indicated by a binary variable p_{uv} . The following constraints implement this logic

$$a_v \geq a_u + t_{uv} + \Delta_u - M(1 - p_{uv}), \quad v \in V \setminus I, u \in N_v^-, \quad (11)$$

$$\sum_{u \in N_v^-} p_{uv} = 1, \quad v \in V \setminus I. \quad (12)$$

Table 3: Instances used in the experiments.

Group	Resources per time instant						H	δ
	10	20	30	40	50	60		
LA	3	3	3	3	0	0	70	50
LB	3	3	3	3	3	3	70	30

where M is an upper bound for the fire arrival time at a vertex. Using the constraints above we could easily integrate to our model the notion of fireline intensity, as introduced by Belval et al. (2015).

Fire combat usually happens through direct attacks (e.g. throwing water at the flames) and indirect attacks (e.g. construction of firebreak lines). Fire suppression at a basic unit v via a direct attack requires the fire to be imminent (i.e. it does not make sense to throw water long before the fire arrives). Our model does not distinguish between direct and indirect attacks, such as Avci et al. (2024) does. To account for that, we could assign to each resource $i \in R$ an expiration time e_i , which represents the maximum interval between an attack at time t and the fire arrival time at the basic unit. When the expiration time of a resource is greater than zero, we have no guarantees that there is an optimal solution that allocates i as soon as it is available (at time t_i). As a result, to introduce expiration times we also have to introduce a continuous variable d_i , for $i \in R$, that defines the time at which resource i is allocated. In summary, we replace constraints (6) and (7) by

$$\begin{aligned}
 t_i &\leq d_i \leq H, & i &\in R, \\
 -(H + z_v^b)(1 - s_v^i) + d_i + z_i^b &\leq a_u, & i &\in R, v \in V_i, u \in V : d(u, v) \leq d_i^b, \\
 -(H + z_v^p)(1 - r_v^i) + d_i + z_i^p &\leq a_u, & i &\in R, v, u \in V : d(u, v) \leq d_i^p, \\
 M(1 - r_v^i) + d_i + e_i &\geq a_v, & i &\in R, v \in V.
 \end{aligned}$$

4. Experiments

In this section we present some computational experiments. All the experiments were done on a platform with a 3.5 GHz AMD Ryzen 9 3900X 12-Core processor, 32 GB of main memory, running Ubuntu Linux 20.04 LTS. All jobs were executed on a single thread.

4.1. Instances

In our experiments, we consider a set of instances for the model of Alvelos (2018). We selected these instances because our model is more general than that of Alvelos (2018) (see 3.2.2) and there is previous work proposing solution methods for this model (see 2.3). This set consists of 16 instances and each instance is a 20×20 grid graph. The set of ignition nodes consist of a single vertex located at the center of the grid and the time horizon H is 70. The 16 instances are divided into two groups of 8 instances each, based on the delay caused by a resource (i.e. the value of δ) and the quantity of resources released at each time instant. These values are shown in Table 3. The optimal solution of all 16 instances is known, so in the sections below we report algorithm performance in terms of the absolute deviation from the optimal objective value.

4.2. Comparison with the literature

In this subsection we compare our model against the model of Alvelos (2018) using the instances described in the last section. In our comparison, we use a strengthened formulation of the model of Alvelos (2018) proposed by Harris et al. (2023). We also consider the logic-based Benders decomposition of Harris et al. (2023) as a baseline. The algorithms rely on a solver to execute, and in our experiments we used Gurobi 10.0.1. In Harris et al. (2023), the decomposition-based algorithm

Table 4: Comparison between M_h , M_p , and M_a . The first group of columns (Abs. gap) shows the absolute difference between the obtained solution and the optimal solution. The second group shows the relative gap between the best obtained solution and the best obtained lower bound. Each line in the first two groups of columns presents the average and the standard deviation of 20 replications on the same instance, and the last line shows the mean of the column. In all 16 instances, M_a failed to produce a feasible solution in at least one replication, so the third column of the first group shows the number of replications in which a feasible solution was found. Each line in the third group of columns shows the percentage of the 20 executions in which the optimal solution was found, while the last group shows the percentage of the executions in which optimality was proven. The last line in these two groups show the percentage of all executions in which the optimal solution was found (third group) or the optimality was proven (fourth group).

	Abs. gap			MIP gap (%)			OPT (%)			POPT (%)		
	M_h	M_p	M_a	M_h	M_p	M_a	M_h	M_p	M_a	M_h	M_p	M_a
LA0	0.00 ± 0.00	0.80 ± 1.54	(17)	0.00 ± 0.00	12.90 ± 1.07	-	100	60	0	100	0	0
LA1	0.00 ± 0.00	0.35 ± 1.18	(19)	0.00 ± 0.00	4.75 ± 3.37	-	100	90	0	100	25	0
LA2	0.00 ± 0.00	0.55 ± 0.60	(16)	0.00 ± 0.00	11.95 ± 0.89	-	100	50	0	100	0	0
LA3	0.00 ± 0.00	0.05 ± 0.22	(19)	0.00 ± 0.00	5.75 ± 0.85	-	100	95	0	100	0	0
LA4	0.00 ± 0.00	0.30 ± 0.92	(10)	0.00 ± 0.00	10.05 ± 1.00	-	100	90	0	100	0	0
LA5	0.00 ± 0.00	0.80 ± 0.95	(16)	0.00 ± 0.00	8.35 ± 1.50	-	100	55	0	100	0	0
LA6	0.00 ± 0.00	0.00 ± 0.00	(11)	0.00 ± 0.00	9.55 ± 0.83	-	100	100	0	100	0	0
LA7	0.00 ± 0.00	0.65 ± 2.08	(9)	0.00 ± 0.00	7.55 ± 1.00	-	100	85	0	100	0	0
LB0	5.40 ± 3.28	4.40 ± 1.70	(9)	10.25 ± 2.63	23.25 ± 1.33	-	15	10	0	0	0	0
LB1	0.00 ± 0.00	1.25 ± 1.83	(11)	0.00 ± 0.00	14.05 ± 0.89	-	100	60	0	100	0	0
LB2	3.10 ± 3.29	3.15 ± 2.03	(10)	8.30 ± 2.70	21.25 ± 1.16	-	20	15	0	0	0	0
LB3	1.15 ± 1.93	0.00 ± 0.00	(19)	5.15 ± 2.78	12.80 ± 0.62	-	65	100	0	10	0	0
LB4	6.90 ± 3.75	5.90 ± 5.51	(6)	8.85 ± 2.46	19.05 ± 2.28	-	5	20	0	0	0	0
LB5	2.95 ± 2.96	6.80 ± 2.65	(13)	6.65 ± 2.01	17.40 ± 1.05	-	20	0	0	0	0	0
LB6	5.30 ± 4.09	7.05 ± 3.62	(12)	7.75 ± 2.31	18.55 ± 1.67	-	20	5	0	0	0	0
LB7	5.10 ± 3.23	0.65 ± 0.67	(9)	9.25 ± 2.12	14.15 ± 0.81	-	5	45	0	0	0	0
Mean	1.87 ± 1.41	2.04 ± 1.59	(12.88)	3.51 ± 1.06	13.21 ± 1.27	-	65.62	55.00	0.00	56.88	1.56	0.00

starts with an initial solution generated by a greedy algorithm. To make a fair comparison, we do not provide initial solutions for any of the algorithms. All three algorithms (i.e. the two models and the Benders decomposition) were executed 20 times on each instance with a time limit of 600 seconds. From now on, we refer to our model as M_p , to the model of Alvelos (2018) as M_a , and to the algorithm of Harris et al. (2023) as M_h . Table 4 summarizes the results.

Several observations can be made. M_h obtains the best results, both in terms of objective value and optimality gap. The proposed model comes in second, finding the optimal solution in about 55% of the time, followed by M_a . Notably, M_a failed to find a feasible solution in 114 out of the 320 executions and was not able to find the optimal solution of a single instance. M_h and M_p have a similar performance in terms of solution quality, but in terms of proving optimality, however, M_h performs considerably better.

The results show that the solver is able to find better solutions to M_p than to M_a . We take a closer look into that by comparing the average size of the two models across the 16 instances. Our comparison is based on analysing the number of variables, the number of constraints, and the number of non-zero coefficients in the model. Most modern solvers apply a preprocessing step to the model, known as presolve, which aims to make it smaller and tighter. Table 5 shows a comparison between M_p and M_a considering the original model and the model obtained after the presolve step.

As we can see, M_p before presolve is larger than M_a (this can be seen in the last column, which shows the ratio of non-zero coefficients). After the presolve step, however, M_p becomes smaller, which may indicate why it obtains better results.

Since M_a struggles to find feasible solutions, and we can find in the literature heuristic methods for the model of Alvelos (2018) (see Section 2.3), we repeated the experiment of Table 4 but this time all algorithms received a high-quality initial solution. Instead of giving a time limit of 600 seconds to each algorithm, we run a heuristic for 100 seconds to obtain an initial solution, and

Table 5: Model size comparison. For each metric s (e.g. number of variables), we present the average of s_{M_p}/s_{M_a} across the 16 instances, where s_{M_p} and s_{M_a} are the value of the metric for M_p and M_a , respectively. The considered metrics were (in order), number of variables, number of binary variables, number of constraints, and the number of non-zero coefficients.

Variables		Bin. Variables		Constraints		Non-zeros	
Original	Presolve	Original	Presolve	Original	Presolve	Original	Presolve
0.71	0.43	1.12	0.71	1.13	0.42	21.11	0.39

Table 6: Comparison between M_h , M_p , and M_a when an initial solution is provided. The notation is the same from Table 4. Due to space constraints, we just present the mean across the 16 instances.

	Abs. gap			MIP gap (%)			OPT (%)			POPT (%)		
	M_h	M_p	M_a	M_h	M_p	M_a	M_h	M_p	M_a	M_h	M_p	M_a
Mean	0.07 ± 0.08	0.09 ± 0.08	0.14 ± 0.16	3.78 ± 0.53	13.21 ± 0.89	39.34 ± 1.52	92.81	90.94	86.56	55.62	1.88	0.00

then we execute the algorithm for 500 seconds. We selected the beam search of Delazeri and Ritt (2024) as the heuristic algorithm, since, in their study, it obtained the best results among the iterated local search of Mendes and Alvelos (2023) and the greedy heuristic of Harris et al. (2023). As in the first experiment, we do 20 replications for each instance. Table 6 shows the results.

As we can see in the table, a strong initial solution improves overall solution quality, but the solver still had difficulties proving optimality. A possible explanation are the weak relaxations of M_p and M_a . Table 7 shows the mean and the standard deviation, across the 16 instances, of the relative gap (in %) between the optimal value of the instance and the relaxation (which is equal for both models). We also show the relative gap of the root relaxation (the relaxation of the presolved model). The relaxation values of the new model are significantly far from the optimal objective values, which suggests the need to explore new decomposition algorithms or valid inequalities when a proof of optimality is desired.

Lastly, it is interesting to note that the MIP gap of M_h increased when compared to Table 4. This happens because, for some instances, M_h can quickly find a solution as good as beam search in less than 100 seconds and then spend the rest of the time improving the lower bound.

5. Conclusions

In this work, we proposed a new model for fire propagation and suppression that integrates several features from previous research. To the best of our knowledge, this is the first model to combine minimum travel times with constraints on resource allocation related to safety and spatial locality. Even though the proposed model is more general, our experiments show that, when evaluated on a set of benchmark instances, it is competitive with more sophisticated algorithms and outperforms more specialized models in terms of solution quality. Our results also reveal that existing benchmark instances in the literature do not present a challenge to current methods. A more diverse and challenging set of benchmark instances would allow more interesting analyses. Additionally, it would be interesting to see whether the performance gains over the model of Alvelos

Table 7: Relative gap (in %) between the relaxation value and the optimal objective value. We report the mean and the standard deviation across the 16 instances.

Relaxation (%)	Root relaxation (%)	
	M_p	M_a
96.22 ± 0.56	73.7 ± 7.16	77.24 ± 5.17

(2018) continue even if the extensions proposed in Section 3.2.2 are considered. Finally, there is no exact or heuristic algorithm in the literature that can solve the proposed model, so this would be another interesting direction for future work.

References

- Alvelos, Filipe Pereira e (2018). “Mixed Integer Programming Models for Fire Fighting”. In: *Computational Science and Its Applications – ICCSA 2018*. Ed. by Osvaldo Gervasi et al. Springer International Publishing, pp. 637–652. ISBN: 978-3-319-95165-2.
- Avci, Mualla Gonca et al. (2024). “The wildfire suppression problem with multiple types of resources”. In: *European Journal of Operational Research* 316.2, pp. 488–502. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2024.03.005.
- Belval, Erin J., Yu Wei, and Michael Bevers (2015). “A mixed integer program to model spatial wildfire behavior and suppression placement decisions”. In: *Canadian Journal of Forest Research* 45.4, pp. 384–393.
- Delazeri, Gustavo and Marcus Ritt (2024). “Iterated Beam Search for Wildland Fire Suppression”. In: *Applications of Evolutionary Computation*. Ed. by Stephen Smith, João Correia, and Christian Cintrano. Springer Nature Switzerland, pp. 273–286. ISBN: 978-3-031-56852-7.
- Dimopoulou, Maria and Ioannis Giannikos (2001). “Spatial Optimization of Resources Deployment for Forest-Fire Management”. In: *International Transactions in Operational Research* 8.5, pp. 523–534. DOI: 10.1111/1475-3995.00330.
- Granda, Bibiana et al. (2023). “Decision Support Models and Methodologies for Fire Suppression”. In: *Fire* 6.2. ISSN: 2571-6255. DOI: 10.3390/fire6020037.
- Harris, Mitchell G., Michael A. Forbes, and Thomas Taimre (2023). “Logic-based benders decomposition for wildfire suppression”. In: *Computers & Operations Research* 160, p. 106392. ISSN: 0305-0548.
- Hof, John et al. (Aug. 2000). “A Timing-Oriented Approach to Spatial Allocation of Fire Management Effort”. In: *Forest Science* 46.3, pp. 442–451. ISSN: 0015-749X. DOI: 10.1093/forestscience/46.3.442.
- Jewell, William S. (Oct. 1963). “Forest Fire Problems—A Progress Report”. In: *Operations Research* 11.5, pp. 678–692.
- Joint Economic Committee, U.S. Senate (2023). *Climate-exacerbated wildfires cost the U.S. between \$394 to \$893 billion each year in economic costs and damages*. <https://web.archive.org/web/20231114011935/https://www.jec.senate.gov/public/index.cfm/democrats/reports?id=E31AF93E-34C7-4C35-A416-533FF796369B>, Accessed: November 13, 2023.
- Kuebler, Martin (Aug. 2024). *Brazil wildfires worsened by ‘mega drought,’ extreme heat*. DW. URL: <https://www.dw.com/en/brazil-wildfires-made-worse-by-drought-extreme-heat/a-70088581>.
- Martell, D. L. (1982). “A review of operational research studies in forest fire management”. In: *Canadian Journal of Forest Research* 12.2, pp. 119–140.
- Mendes, André Bergsten and Filipe Pereira e Alvelos (2023). “Iterated local search for the placement of wildland fire suppression resources”. In: *European Journal of Operational Research* 304.3, pp. 887–900. ISSN: 0377-2217.
- Paddison, Laura (Aug. 2024). *Photos from space show just how close Greece wildfires got to burning whole towns*. CNN. URL: <https://edition.cnn.com/2024/08/16/climate/greece-wildfires-athens-photos-climate-intl/index.html>.

- Rachaniotis, Nikos P. and Costas P. Pappis (May 2010). “Minimizing the total weighted tardiness in wildfire suppression”. In: *Operational Research* 11.1, pp. 113–120. ISSN: 1866-1505. DOI: 10.1007/s12351-010-0079-x.
- Rothermel, Richard C. (1972). *A mathematical model for predicting fire spread in wildland fuels*. Research Paper INT-115. Ogden, UT: U.S. Department of Agriculture, Intermountain Forest and Range Experiment Station, p. 40.
- Wei, Yu, Douglas B. Rideout, and Thomas B. Hall (Oct. 2011). “Toward Efficient Management of Large Fires: A Mixed Integer Programming Model and Two Iterative Approaches”. In: *Forest Science* 57.5, pp. 435–447. ISSN: 0015-749X. DOI: 10.1093/forestscience/57.5.435.