

Probabilistic Response-Time-Aware Search for Transient Astrophysical Phenomena



Daisy Wang* Marion Sudvarg* Filip Marković† Jeremy Buhler* Sanjoy Baruah* Gregory Kehne*

*Washington University in St. Louis, United States †University of Southampton, United Kingdom

(w.yanwang, msudvarg, jbuhler, sbaruah, kehne)@wustl.edu f.markovic@soton.ac.uk

Abstract—Timely observation of transient astrophysical phenomena (TAP) is of crucial importance for our understanding of the universe and the laws of physics, as recognized by the National Academies in the Astro2020 decadal survey. Ultimately, the goal is to observe TAPs as early as possible using optical telescopes. This is non-trivial due to the probabilistic nature of the search problem, where multiple potential sky locations for a TAP, each with an associated probability, must be scheduled for observation before successful localization. The problem lies at the intersection of several research disciplines, including real-time systems, cyber-physical systems, astrophysics, and operations research, motivating the need for a unified modeling framework.

To this end, we introduce the first formal stochastic, response-time-aware model for search planning toward detection and localization of TAPs. We consider the problem of maximizing expected utility of early localization and show that it is reducible to the Orienteering Problem. Building on this formulation, we develop the real-time-capable Greedy-Christofides Pathfinding (GCP) algorithm. An evaluation on 37 probability maps from LIGO demonstrates that GCP consistently achieves high solution quality and computational efficiency across diverse search scenarios. GCP achieves $\leq 0.5\%$ deviation from the ILP-computed optimal solution on tractable problem instances while running within a second, on average, for larger inputs.

I. INTRODUCTION

The U.S. National Academies of Sciences, Engineering, and Medicine, in collaboration with other scientific organizations (NASA, NSF, DOE), released the Astro2020 decadal survey [59] to identify scientific challenges for astronomy and astrophysics in the next decade. The survey highlighted the “space-based **time-domain** and multi-messenger program” as the highest-priority sustaining activity in space. The goal of this program is to coordinate concurrent observations of astrophysical phenomena by instruments with different observational modalities (e.g., gravity waves, neutrinos, cosmic rays, and across the electromagnetic spectrum).

A key component of the broader multi-messenger program is the *real-time* coordination of astrophysical observations to achieve *timely* optical detection of *transient astrophysical phenomena* (TAPs), such as kilonovae and supernovae.

A typical example of such coordination is presented in Fig. 1 in four stages (a–d). First, a TAP emits prompt messenger signals, e.g., a gamma-ray burst (GRB) (stage a), which are subsequently detected by a wide-field, non-optical instrument (stage b). Observational data from that instrument are used to construct a map that encodes the spatial probability distribution of the source location. This map is transmitted to an optical telescope (stage c) for follow-up observations.

A key challenge is that many non-optical instruments cannot constrain the TAP’s origin to within a region smaller than




Fig. 1. **I and II:** Overview of the real-time search pipeline for transient astrophysical phenomena (TAP). **III:** Temporal decay of the modeled energy flux ε_{obs} across different energy bands [9], overlaid with deadlines D_1 , D_2 , and D_3 corresponding to scientifically significant observation phases.

the narrow field-of-view (FoV) of most optical telescopes, as illustrated in Fig. 1(I); the likely region for a TAP often spans an area equal to numerous FoVs (e.g., [7, 29, 66]). To address this challenge, the optical telescope conducts a time-consuming *search* (stage d in Fig. 1(II)) across multiple FoVs to identify the TAP’s precise location. Effective search strategies are therefore important to ensure that early features of the TAP’s evolution across the broader electromagnetic spectrum are observed (see Fig. 1(III)).

The scientific motivation for prioritizing early observations stems from the fleeting nature of several physical processes involved in many TAPs, which evolve on timescales ranging from milliseconds to hours and vary across different types of transients (e.g., [14, 15, 68, 75]). Several theoretical models describing the early evolution of various TAPs remain unconfirmed and are continuously refined through prompt optical observations (e.g., [35, 60, 80]). Whether these processes can be captured depends on the temporal properties of the event and the response-time capabilities of the telescopes. We elaborate the scientific objectives in Section II.

To jointly capture the *time-sensitive* nature of the problem and the varying *scientific value* of localization at different stages of TAP evolution, we adopt the figure of merit (FoM), which is a quantitative measure of mission success, and is a notion

instantiated broadly as a *science merit function* [19] in space research [65]. Ultimately, an FoM is assigned to each of several deadlines for observation, corresponding to the timescales at which physical processes of interest to the mission are expected to remain observable. For example, Fig. 1_(III) shows simulated light curves from supernova afterglow [9]. These may be used to inform deadline selection; e.g., D_1 corresponds to the time of peak emission across most energy bands, D_2 to the onset of the plateau phase and emergence of 1 eV (near-infrared) emission, and D_3 to the start of the late-time decay.

The Search Problem. In this paper, we consider the search problem (stage **d** in Fig. 1_(II)) that seeks to *maximize the expected FoM* for optical observations of a single detected TAP. In particular, we consider the *entire* mission-planning problem before the telescope starts to move. We conduct our efforts within the following framework. A TAP’s spatial probability distribution is discretized for an optical telescope by projecting its FoV as a set of *tiles* providing complete coverage of the sky (see Fig. 1_(I)); each tile is associated with the probability that the TAP originated from a direction within its boundary. We consider the following problem to model the detection of a TAP by a searching optical telescope.

Given: a set of tiles, each characterized with ① the probability that the TAP originated within its boundary and ② an upper bound on the dwell time required for the optical telescope to determine whether or not the TAP is present in that tile. Each pair of tiles is characterized with ③ the time it takes for the telescope to move between them. And ④ a set of deadlines, each assigned an FoM associated with meeting it. **Produce:** the sequence of tiles to search so as to maximize the expected FoM.

We assume in this paper that the times to move between tiles satisfy the triangle inequality, a restriction justified by the mechanical properties of many real telescope mounts.

Contributions. This paper presents, to our knowledge, the first formal probabilistic, time-aware model for search planning (defined in Section III) that (i) considers the time both to move between search locations, and to observe the current location to detect the target, and (ii) incorporates multiple deadlines by explicitly encoding the expected FoM as the optimization objective. Although the problem is stochastic, we show in Sections V and VI that it can be solved deterministically. In particular, it is solved optimally by finding a traversal pattern that maximizes the cumulative detection probability amassed within each deadline, weighted by the corresponding FoMs.

With even a single deadline, we observe that our problem inherits hardness of approximation from orienteering (Corollary 3) and is APX-hard. However, we also demonstrate that the problem can be reduced to the Orienteering Problem (OP), such that any α -approximation to OP becomes an α -approximation to our problem (Corollary 4). Further, an α -approximation for the single-deadline case may be extended to optimize the FoM over multiple deadlines while maintaining a polynomial bound on the increase in approximation ratio (Section VIII).

From these observations, we propose the response-time-aware Greedy-Christofides Pathfinding (GCP) heuristic for maximizing the expected FoM in time-constrained search




Fig. 2. LIGO spatial probability map from gravitational-wave signal GW230529_181500 [2] with inset showing a tiled region.

scenarios.¹ We evaluate GCP on 37 probability maps from gravitational-wave events detected by LIGO [2, 4] using realistic models of optical telescopes participating in multi-messenger follow-up campaigns. We compare it to a Genetic algorithm and a greedy heuristic (Section IX). Results indicate that GCP consistently achieves high solution quality, closely matching the optimal solution of an ILP formulation on small problem instances, while remaining computationally efficient across complex and diverse search scenarios.

II. BACKGROUND

Scientific Objectives. Time-domain and multi-messenger astrophysics provide insights arising from temporally aligned “information carried by photons, gravitational waves, neutrinos and cosmic rays” [58]. The Astro2020 survey [59], released as a joint effort among U.S. scientific agencies, identifies the field as its highest priority among space-based programs.

Many sources of these “messenger signals” are transient astrophysical phenomena (TAPs), which are often initially detected by prompt, high-energy emission signatures, e.g. gamma-ray bursts (GRBs). Theoretical models describing the TAP’s early evolution may be confirmed or refined through additional optical observations. For example, open questions remain about prompt optical counterparts concurrent with the initial short GRBs emitted by binary neutron star mergers; shock breakout and cooling effects which occur in the first seconds to a fraction of an hour of supernovae [81]; and thermal and radioactive decay mechanisms for which important data may be obtained from the first minutes to hours of early afterglow light curves. To capture these processes, observations must be made at the appropriate timescales, implying a *deadline* by which the TAP must be localized by a searching telescope. The associated challenges are longstanding problems but, if modeled appropriately, may be particularly suitable to the collective expertise of the real-time systems community.

Given that the target TAP’s location is not known a priori, unless the complete search region can be covered, even an optimal strategy cannot guarantee success within the deadline. A deadline miss, however, should not be considered a mission failure: subsequent observations may still be useful. Indeed, the late optical afterglows of GRBs often last several days (see Fig. 1_(III)). This suggests some *time utility* of the follow-up observations, which decreases the longer they are delayed during the search. Borrowing terminology from NASA for

¹<https://github.com/Daisy0419/Telescope-Searching-Problem>

space missions [65], we quantify this utility as a figure of merit (FoM) associated with meeting successive deadlines (Fig. 1_(III)) imposed by the different processes we want to observe.

TAP Detection and Localization. After messenger signals **(a)** emitted by a TAP reach instruments on (or in orbit around) the Earth, the several stages depicted in Fig. 1 must execute to produce the inputs to the search problem considered in this paper. A high-energy detector, e.g., for gamma rays (Fermi [11], COSI [76], APT [22]), neutrinos (IceCube [1], ANTARES [6]), or gravitational waves (LIGO [3], VIRGO [16], LISA [31]), rapidly detects the TAP with an almost omnidirectional FoV.

Observed data from the TAP are then analyzed **(b)** using a parametric model of the high-energy instrument’s detector response, which does *not* depend on prior probabilities of TAP appearances, to evaluate the likelihood that the TAP originated from each spatial direction, e.g., with the BAYESTAR [71] software for gravitational-wave signals (see Fig. 2) or the `cosipy` library [23, 57] for GRBs. The resulting probability map is typically discretized using HEALPix [39], an equal-area tiling of the sphere, for **(c)** transmission to optical telescopes. Ongoing efforts to estimate the TAP’s source direction *in real time* aboard high-energy detectors [73, 74] are particularly important for space-based instruments like APT, for which limited communication bandwidth and high speed-of-light latency to Earth may induce unwanted delays in stages **(b–c)**.

Properties of Optical Telescopes. Upon receiving a probability map, an optical telescope projects its FoV into a set of *tiles* to completely cover the constrained search region, as illustrated in the inset in Fig. 2. Each tile is assigned a probability based on its intersection with the HEALPix regions in the probability map, e.g., using the methods of [37]. Tiles are then observed in sequence. Current methods adopted by the astronomy community to determine the order, e.g., those outlined in Section X, do not attempt to solve the formal response-time and FoM-driven problem explored in this paper.

Captured images are analyzed to identify whether the TAP is present in the observed tile. Standard analytical techniques, e.g., those outlined in [8], include astrometric alignment [49] with cataloged images [24, 82], image subtraction [45], and the use of ML classifiers to identify whether features in the difference image indicate the presence of the TAP [21].

Two key properties of the optical telescope must be considered when defining a search order: the time it takes to move between tiles, and the “dwell time” to collect enough light from a single tile to ensure TAP identification. Many optical telescopes participating in multi-messenger observational campaigns, e.g., RAPTOR [18] and the 7DT [48], use direct-drive mounts that enable geodesic movement among tiles with traversal time proportional to the angular difference between the tile centers. The set of tiles and the movement times among them thus form a *metric space* (distances are non-negative, symmetric, and satisfy the triangle inequality), which is important for the problem reduction in Section VI. Movement costs for mounts using separate elevation and azimuthal motors are history-dependent due to degeneracy of the spherical coordinate system

TABLE I. TERMINOLOGY

Symbol	Explanation
τ	The set of tiles in a telescope’s field of view.
$\tau_i \in \tau$	A tile from τ .
$m(\tau_i, \tau_j) \in \mathbb{R}_{\geq 0}$	Upper bound on the transition time needed to refocus the telescope from tile τ_i to tile τ_j (assumed metric, i.e., symmetric and satisfies the triangle inequality).
$C_i \in \mathbb{R}_{\geq 0}$	Upper bound on dwell time for successful detection.
$\mathbb{D} \triangleq (D_1, \dots, D_{z-1})$	Sequence of deadlines corresponding to stages of interest in the TAP emission.
$\mathbb{M} \triangleq (M_1, \dots, M_{z-1})$	Figures of Merit (FoMs) quantifying the scientific merit of TAP detection and localization within the corresponding deadline.
$\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$	A function encoding a spatial probability distribution derived from the non-optical telescope.
$p_i \in [0, 1]$	Probability of TAP originating in τ_i .

at the poles; considering these is therefore left to future work.

The telescope’s optical sensitivity and estimates of the TAP’s optical brightness based on the flux observed by the high-energy instrument are used to define an upper bound on the dwell time. Atmospheric opacity, e.g., as computed by the SMARTS model [41], also impacts dwell times for ground-based optical telescopes. This implies a non-uniform dwell time across tiles; for example, a telescope must look through a greater mass of atmosphere at angles closer to the horizon [47].

We note that methods for *deriving* the parameters discussed in this section are not the focus of this paper; these parameters are assumed to be given as inputs to the considered problem.

III. SYSTEM MODEL AND PROBLEM STATEMENT

Tiling. Let $\tau \triangleq \{\tau_1, \dots, \tau_n\}$ be a set of tiles projected with respect to the FoV for the telescope of interest. The upper bound on the time needed for the telescope to shift its focus from τ_i to τ_j is $m(\tau_i, \tau_j)$. We assume m is metric, i.e., $m(\tau_i, \tau_j) = m(\tau_j, \tau_i)$, and for any three tiles τ_i , τ_j , and τ_k , $m(\tau_i, \tau_k) \leq m(\tau_i, \tau_j) + m(\tau_j, \tau_k)$ always holds. Upon commencing optical detection (hereafter referred to as *detection*), on tile τ_i , it is assumed that the time to determine the TAP’s presence is upper-bounded by a known value C_i .

Deadlines. Let $\mathbb{D} \triangleq (D_1, \dots, D_{z-1})$ be an increasing sequence of mission-specified deadlines corresponding to specific stages of interest in the TAP emission (see Fig. 1_(III)), relative to the time D_0 (=c in Fig. 1_(II)) at which the optical telescope receives the spatial distribution.

Figure of Merit (FoM). Let $\mathbb{M} \triangleq (M_1, \dots, M_{z-1})$ be a sequence of figures of merit for each deadline. An FoM is a number representing a measure of success in space missions [65]. In this context, $M_k \in \mathbb{M}$ is a mission-specified scalar quantifying the *scientific merit* [19] of TAP detection within the time interval $(D_{k-1}, D_k]$. For the sake of explainability [13], we define $M_1 \triangleq 1$ and impose the ordering $M_1 \geq M_2 \geq \dots \geq M_z$, reflecting the principle that successful detection before D_k implies success at all subsequent deadlines. This setup enables an intuitive comparison of values. For instance, if $M_1 = 1$ and $M_2 = 0.5$, then detection before D_1




Fig. 3. An example with four tiles and corresponding parameters from Table I. **Note:** $p_1 = 0$ and $C_1 = 1$, while for $\tau_* \in \tau \setminus \{\tau_1\}$, $m(\tau_1, \tau_*) = 1$.

is considered twice as valuable as detection occurring after D_1 but not after D_2 . For completeness, we define $D_z \triangleq +\infty$ and $M_z \triangleq 0$, indicating that the mission is considered unsuccessful beyond the last mission-specified deadline D_{z-1} .

Tile-Based Probability Distribution. Formally, we consider the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let Ω denote the sample space over outcomes $\omega \in \Omega$. For the considered problem, an outcome represents the origination of a single TAP in a specific tile τ_i ; we define a function $o : \Omega \rightarrow \tau$ that maps each outcome ω to its corresponding tile τ_i . Then $\mathcal{F} \subseteq 2^\Omega$ is the event space, where each event is a subset of outcomes in Ω . Finally, $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$ is a function that maps each event to its corresponding probability. Here, $\mathbb{P}$ encodes a single TAP's spatial probability distribution inferred from observational data taken by a non-optical telescope, and discretized to the FoV-based tiling (see, e.g., the probability map in Fig. 2).

Problem. Given the parameters defined above, the primary objective is to maximize the mission's FoM. Ideally the TAP would always be detected at or before D_1 , resulting in the maximum possible mission FoM $M_1 = 1$. However, several factors make this problem inherently complex:

- A telescope may be unable to observe all tiles within a deadline, potentially resulting in a lower figure of merit due to the problem instance.
- From the perspective of search planning, the resulting figure of merit is a random variable, reflecting the stochastic nature of the search process over a probability space where successful TAP detection is uncertain yet open to probabilistic quantification.
- Optimizing the search strategy for a specific deadline, e.g., D_1 , may lead to inefficiencies in long-term mission planning, as it can cause significant divergence from strategies optimal for later, more feasible deadlines.

We note that, although optical afterglows from multiple TAPs might be visible simultaneously, the classes of high-energy TAP signals we consider are observed only about once per day globally; their initial bursts would almost never coincide. We therefore assume that the probability map being searched originates from a single TAP. Searching for multiple TAPs within distinct probability maps is left to future work.

IV. MOTIVATING EXAMPLE

To illustrate the above concerns, we present a motivating (and running) example highlighting the importance of incorporating response-time awareness into the problem statement.




Fig. 4. Cumulative distribution function of a successful TAP detection for six search orders over time (up to 17 TUs) considering two mission deadlines.

Consider a configuration of four tiles $\tau = \{\tau_1, \tau_2, \tau_3, \tau_4\}$ as illustrated in Fig. 3, with probabilities $p_1 = 0$, $p_2 = 0.1$, $p_3 = 0.1$ and $p_4 = 0.8$. Assume that the optical telescope is initially focused on τ_1 ; since $p_1 = 0$ it is thus excluded from further consideration. Let us consider two deadlines $D_1 = 4$ and $D_2 = 12$ for two TAP emission phases of interest. Six complete search orders (SO) starting from τ_1 visit each tile:
 SO_1 : $\tau_1 \rightarrow \tau_2 \rightarrow \tau_3 \rightarrow \tau_4$, SO_2 : $\tau_1 \rightarrow \tau_3 \rightarrow \tau_2 \rightarrow \tau_4$
 SO_3 : $\tau_1 \rightarrow \tau_2 \rightarrow \tau_4 \rightarrow \tau_3$, SO_4 : $\tau_1 \rightarrow \tau_3 \rightarrow \tau_4 \rightarrow \tau_2$
 SO_5 : $\tau_1 \rightarrow \tau_4 \rightarrow \tau_2 \rightarrow \tau_3$, SO_6 : $\tau_1 \rightarrow \tau_4 \rightarrow \tau_3 \rightarrow \tau_2$
but the corresponding response-time distributions for successful GRB detection vary significantly (see Fig. 4).

Interpretation. Fig. 4 depicts the cumulative distribution function (CDF) of a successful TAP detection up to time t (expressed in generic time units on the x-axis). For instance, the probability that SO_1 , SO_2 , SO_3 , or SO_4 successfully detect the TAP up to time $t = 4$ is 0.1, while for SO_5 or SO_6 it is 0. This means that for the mission-specified deadline $D_1 = 4$, search orders SO_1 – SO_4 are preferred over the others. However, for $D_2 = 12$, SO_5 and SO_6 are better than the alternatives, as the probability of a successful TAP detection would be 0.8 compared to 0.2 (SO_1 or SO_2) and 0.1 (SO_3 or SO_4).

Let us now introduce the figures of merit. By definition, $M_1 = 1$. Suppose $M_2 = 0.5$, suggesting that the contribution to the overall mission value of detecting a TAP after D_1 but not after D_2 is half that of detecting it no later than D_1 . The FoM of the entire search mission is a random variable depending on the search order as follows:

- For SO_1 and SO_2 , the FoM takes the values 1 (Pr 0.1), 0.5 (Pr 0.1), and 0 (Pr 0.8). The expected FoM in this case is $\mathbb{E}[\text{FoM}] = 1 \cdot 0.1 + 0.5 \cdot 0.1 + 0 \cdot 0.8 = 0.15$.
- For SO_3 and SO_4 , the FoM takes the values 1 (Pr 0.1), 0.5 (Pr 0), and 0 (Pr 0.9). The expected FoM in this case is $\mathbb{E}[\text{FoM}] = 1 \cdot 0.1 + 0.5 \cdot 0 + 0 \cdot 0.9 = 0.1$.
- For SO_5 and SO_6 , the FoM takes the values 1 (Pr 0), 0.5 (Pr 0.8), and 0 (Pr 0.2). The expected FoM in this case is $\mathbb{E}[\text{FoM}] = 1 \cdot 0 + 0.5 \cdot 0.8 + 0 \cdot 0.2 = 0.4$.

From the analysis above, SO_5 and SO_6 are the preferred search orders, yielding a higher expected FoM than alternatives.

In conclusion, an efficient algorithm that explores a broad search space is essential. Overly local strategies (e.g., optimising for D_1) risk mission failure, while computationally heavy algorithms may fail to produce a timely search order.

In Section V, we formally define a notion of the search order, its response time, and its expected FoM. In Section VI,

we show that the simplest instance of the problem is APX-hard and can be modeled as a variant of the Orienteering Problem (OP) [77], while Section VII presents the proposed algorithm.

V. RESPONSE TIME OF AN OPTICAL DETECTION

While this section provides a rigorous probabilistic formulation of the problem, readers primarily interested in its deterministic resolution may skip ahead to Section VI, where we show that the problem reduces to a deterministic optimization, namely *the orienteering problem*.

To define the response time for detecting a TAP, we first define the concept of a **search order**. Let $k \leq n$ and let $\pi : \{0, 1, \dots, k\} \rightarrow \tau$ be a search order function, where for $1 \leq i \leq k$, $\pi(i)$ represents the i -th tile allocated for observation while $\pi(0)$ is the start tile, *i.e.*, the tile of focus at time D_0 according to the optical telescope's current pointing when it receives the probability map. We require that π is injective on the subset $\{1, \dots, k\}$. That is, each tile is allocated only once for TAP detection, as allowing revisits could inflate the worst-case response time: early exits followed by revisits may lead to overly inefficient searches. Let $\pi^{-1}(\cdot)$ be the inverse function of $\pi(\cdot)$ and let $\pi^\#(\cdot)$ be the function returning the tile index. Finally, let $\mathbb{S}$ be the set of all possible orders.

The **response time** of a successful TAP detection depends on the assigned search order and the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We consider its upper bound to be a random variable, denoted by $\mathcal{R} : \mathbb{S} \times \Omega \rightarrow \mathbb{R}_{\geq 0}$. Its cumulative distribution function (CDF), for a given search order π is $F_{\mathcal{R}_\pi}(t) \triangleq \mathbb{P}[\mathcal{R}_\pi \leq t]$. $\mathcal{R}_\pi$ is formally defined as follows.

Def. 1. The upper bound on the response time of a successful TAP detection in sample $\omega \in \Omega$ with search order π is:

$$\mathcal{R}_\pi(\omega) \triangleq \sum_{i=1}^{\pi^{-1}(o(\omega))} m(\pi(i-1), \pi(i)) + C_{\pi^\#(i)}.$$

Example 1. Consider $\pi = SO_6$ from the running example, which starts at τ_1 and assesses the TAP's origin at $\tau_4 \rightarrow \tau_3 \rightarrow \tau_2$. Suppose that ω denotes the outcome corresponding to the TAP originating in τ_3 , that is, $o(\omega) = \tau_3$. Since τ_3 is the second tile considered in SO_6 , $\pi^{-1}(o(\omega)) = \pi^{-1}(3) = 2$. The corresponding response time is:

$$\begin{aligned} \mathcal{R}_\pi(\omega) &= \sum_{i=1}^2 m(\pi(i-1), \pi(i)) + C_{\pi^\#(i)} \\ &= m(\pi(0), \pi(1)) + C_{\pi^\#(1)} + m(\pi(1), \pi(2)) + C_{\pi^\#(2)} \\ &= m(\tau_1, \tau_4) + C_4 + m(\tau_4, \tau_3) + C_3 = 1 + 10 + 2 + 1 = 14. \end{aligned}$$

$\mathcal{R}_\pi(\omega)$ is an upper bound rather than the exact response time since we acknowledge that detection may occur earlier. However, we conservatively assume that the detection time on each tile $\tau_i \in \tau$ is given by C_i , which represents an upper bound on the execution time required for successful TAP detection. The transition times are also treated conservatively.

Similarly, the lower bound on the mission's FoM is also a random variable $\mathcal{M} : \mathbb{S} \times \Omega \rightarrow [0, 1]$ defined as follows.

Def. 2. The lower bound on the mission's figure of merit in sample $\omega \in \Omega$ with search order π is:

$$\mathcal{M}_\pi(\omega) \triangleq \min \{M_i \mid \mathcal{R}_\pi(\omega) \leq D_i\}.$$

Example 2. Building upon Example 1, $\mathcal{M}_\pi$ is computed as:

$$\mathcal{M}_\pi(\omega) = \min \{M_i \mid 14 \leq D_i\} = M_3 = 0,$$

implying that search order $\pi = SO_3$ is unsuccessful in event $\{\omega\}$ (TAP originating in τ_3) because $\mathcal{R}_\pi(\omega) = 14 > D_2 = 12$. Since $D_2 = 12$ is the last mission-defined deadline, FoM evaluates to M_3 , which corresponds to the auxiliary deadline $D_3 = +\infty$, indicating an unsuccessful outcome ($M_3 = 0$).

Conversely to the response time, the mission's FoM can be higher if the TAP is localized earlier, but here we once again seek a conservative estimate by assuming that detection happens only after conservative $\mathcal{R}_\pi(\omega)$ units of time.

Finally, we utilize the following lemma to efficiently compute the upper bound on the expected value of $\mathcal{M}_\pi$.

Lemma 1. Let $\mathcal{M}_\pi$ be the random variable of Definition 2 and for legibility, let $\phi_{\pi,i} \triangleq F_{\mathcal{R}_\pi}(D_i)$. Then, its expectation is:

$$\mathbb{E}[\mathcal{M}_\pi] = \sum_{i=1}^{z-1} M_i \cdot (\phi_{\pi,i} - \phi_{\pi,i-1}).$$

Proof. $\mathbb{E}[\mathcal{M}_\pi]$

$$\begin{aligned} &\stackrel{(i)}{=} \sum_{i=1}^z \sum_{\omega \in \Omega \mid D_{i-1} < \mathcal{R}_\pi(\omega) \leq D_i} \min \{M_i \mid \mathcal{R}_\pi(\omega) \leq D_i\} \cdot \mathbb{P}(\{\omega\}) \\ &\stackrel{(ii)}{=} \sum_{i=1}^z M_i \cdot \mathbb{P}[D_{i-1} < \mathcal{R}_\pi \leq D_i] \\ &\stackrel{(iii)}{=} \sum_{i=1}^z M_i \cdot (\phi_{\pi,i} - \phi_{\pi,i-1}) \stackrel{(iv)}{=} \sum_{i=1}^{z-1} M_i \cdot (\phi_{\pi,i} - \phi_{\pi,i-1}). \end{aligned}$$

Step (i) follows by definition of expectation $\mathbb{E}[\mathcal{M}_\pi] \triangleq \sum_{\omega \in \Omega} \mathcal{M}_\pi(\omega) \cdot \mathbb{P}(\{\omega\})$, Definition 2, and by partitioning the sample space Ω according to the successful detection within $(D_{i-1}, D_i]$ intervals. Step (ii) follows by definition of min in the interval $D_{i-1} < \mathcal{R}_\pi(\omega) \leq D_i$ where the inner expression simplifies to M_i , by factoring out the constant M_i from the inner sum and by recognising the inner sum as the probability $\mathbb{P}[D_{i-1} < \mathcal{R}_\pi \leq D_i]$. Step (iii) follows by the CDF identity $\mathbb{P}[D_{i-1} < \mathcal{R}_\pi \leq D_i] = F_{\mathcal{R}_\pi}(D_i) - F_{\mathcal{R}_\pi}(D_{i-1})$ (*e.g.*, from Equation 4.14 in [61]). Finally, Step (iv) follows since $M_z = 0$ by definition from Section III, which concludes the proof. $\square$

Lemma 1 is convenient because it expresses the expected FoM entirely in terms of the cumulative distribution function evaluated at a finite number of deadlines. Consequently, **comparing different search orders reduces to comparing their respective CDFs at the given deadlines**, thus simplifying both analysis and optimization.

This observation motivates the following theorem, which formalizes the principle that a search order whose CDF dominates that of others across all deadlines necessarily achieves a greater or equal expected FoM.

Theorem 1. If $\pi, \pi' \in \mathbb{S}$ and for all $i \in \{1, \dots, z-1\}$, $F_{\mathcal{R}_\pi}(D_i) \geq F_{\mathcal{R}_{\pi'}}(D_i)$, then $\mathbb{E}[\mathcal{M}_\pi] \geq \mathbb{E}[\mathcal{M}_{\pi'}]$.

Proof. We prove $\mathbb{E}[\mathcal{M}_\pi] - \mathbb{E}[\mathcal{M}_{\pi'}] \geq 0$. For legibility, let $\phi_{\pi^*, i} \triangleq F_{\mathcal{R}_{\pi^*}}(D_i)$ and let us start from $\mathbb{E}[\mathcal{M}_\pi] - \mathbb{E}[\mathcal{M}_{\pi'}]$

$$\begin{aligned} & \stackrel{(i)}{=} \sum_{k=0}^{z-2} M_{k+1} \cdot (\phi_{\pi, k+1} - \phi_{\pi, k}) - \sum_{k=0}^{z-2} M_{k+1} \cdot (\phi_{\pi', k+1} - \phi_{\pi', k}) \\ & \stackrel{(ii)}{=} \sum_{k=0}^{z-2} \phi_{\pi, k+1} \cdot (M_{k+1} - M_{k+2}) - \sum_{k=0}^{z-2} \phi_{\pi', k+1} \cdot (M_{k+1} - M_{k+2}) \\ & \stackrel{(iii)}{=} \sum_{k=0}^{z-2} (\phi_{\pi, k+1} - \phi_{\pi', k+1}) \cdot (M_{k+1} - M_{k+2}) \stackrel{(iv)}{\geq} 0. \end{aligned}$$

Step (i) follows from Lemma 1 and from applying reindexing $k = i - 1$, implying $i = k + 1$. Step (ii) follows from Abel's summation by parts lemma (e.g., Equation 2.56 in [40]): $\sum_{k=0}^{n-1} b_k(a_{k+1} - a_k) = a_n b_n - a_0 b_0 + \sum_{k=0}^{n-1} a_{k+1}(b_k - b_{k+1})$ applied with $n = z-1$, $b_k = M_{k+1}$ and $a_k = \phi_{\pi^*, k}$ for $\pi^* \in \mathbb{S}$. Terms corresponding to $a_n b_n$ and $a_0 b_0$ equal zero due to model assumptions ($b_n = b_{z-1} = M_z = 0$ and $a_0 = F_{\mathcal{R}_{\pi^*}}(D_0) = 0$). Step (iii) follows from applying the linearity of summation and the distributive law of multiplication. Step (iv) follows by theorem assumption, i.e., for $0 \leq k \leq z-2$, $\phi_{\pi, k+1} - \phi_{\pi', k+1} = F_{\mathcal{R}_\pi}(D_{k+1}) - F_{\mathcal{R}_{\pi'}}(D_{k+1}) \geq 0$, and by system model assumptions on FoM: $M_{k+1} - M_{k+2} \geq 0$. $\square$

By iteratively applying Theorem 1, we conclude that if there exists a search order $\pi \in \mathbb{S}$ such that $F_{\mathcal{R}_\pi}(D_i) \geq F_{\mathcal{R}_{\pi'}}(D_i)$ for all $i \in \{1, \dots, z\}$ and all $\pi' \in \mathbb{S}$, then π achieves the maximum expected FoM among all search orders.

Corollary 1. Let $\pi \in \mathbb{S}$ be such that for all $\pi' \in \mathbb{S}$ and all $i \in \{1, \dots, z-1\}$, it holds that $F_{\mathcal{R}_\pi}(D_i) \geq F_{\mathcal{R}_{\pi'}}(D_i)$. Then,

$$\max_{\pi' \in \mathbb{S}} \mathbb{E}[\mathcal{M}_{\pi'}] = \mathbb{E}[\mathcal{M}_\pi].$$

Proof. By induction on the cardinality of $\mathbb{S}$, using Theorem 1 at each step to compare π to every other search order. $\square$

In the following section, we show that the problem of finding the optimal search order is APX-hard, and that a special case ($z = 2$) reduces to the orienteering problem.

VI. REDUCTION TO THE ORIENTEERING PROBLEM

Let us consider the special case $z = 2$, where there is only one mission-specified deadline D_1 with $M_1 = 1$ denoting a successful mission, and an auxiliary deadline $D_2 = +\infty$ with $M_2 = 0$ denoting an unsuccessful mission. We begin by showing that the optimal search order is the order achieving the highest probability of detection before the mission deadline.

Corollary 2. If $|\mathbb{D}| = 2$, $\pi \in \mathbb{S}$, and for all $\pi' \in \mathbb{S}$, the CDF $F_{\mathcal{R}_\pi}(D_1) \geq F_{\mathcal{R}_{\pi'}}(D_1)$, then

$$\max_{\pi' \in \mathbb{S}} \mathbb{E}[\mathcal{M}_{\pi'}] = \max_{\pi' \in \mathbb{S}} F_{\mathcal{R}_{\pi'}}(D_1) = F_{\mathcal{R}_\pi}(D_1).$$

Proof. Follows by Lemma 1 and corollary assumptions as for $D = 2$, $\forall \pi' \in \mathbb{S}$, $\mathbb{E}[\mathcal{M}_{\pi'}] = M_1 \cdot (F_{\mathcal{R}_{\pi'}}(D_1) - F_{\mathcal{R}_{\pi'}}(D_0)) = 1 \cdot F_{\mathcal{R}_{\pi'}}(D_1) - 0 = F_{\mathcal{R}_{\pi'}}(D_1)$. $\square$

The importance of Corollary 2 lies in the fact that when $z = 2$, the expected FoM $\mathbb{E}[\mathcal{M}_\pi]$ is fully determined by the CDF $F_{\mathcal{R}_\pi}(D_1)$. Therefore, the optimal search order is the one that maximizes CDF at D_1 . This allows us to show that such an optimization can be formulated as a **probability collection problem under budget constraints**.

According to the definition by Golden et al. [38], the orienteering problem is defined as follows.

Def. 3 Orienteering problem (OP). Consider a set of nodes $N = \{1, \dots, |N|\}$, where each node $i \in N$ is associated with non-negative reward r_i . Two nodes $s, t \in N$ are designated as start and end nodes, respectively. If an edge exists between nodes i and j , the non-negative travel cost between i and j is $d(i, j)$. The goal of the OP is to determine a path, limited by a given time budget $T_{\max}$, that visits a subset of N and maximizes the total collected reward. It is assumed that rewards can be added across nodes and that, while a node may be visited more than once, its reward is obtained only once.

Reduction. In the one-deadline ($z = 2$) case, the search problem, in which search must start from some designated tile τ_0 , can be reduced to the OP as follows.

- 1) Create a node i in N for each tile $\tau_i \in \tau$, and add artificial nodes s, t that will serve as start and end nodes for the OP.
- 2) Assign node i reward p_i where $p_i \triangleq \mathbb{P}[\{\omega\}]$ s.t. $o(\omega) = \tau_i$, i.e., the probability that the TAP originates at τ_i . Artificial nodes s, t are assigned zero reward.
- 3) Create edges between every pair of nodes arising from tiles, and between all such nodes and both s and t .
- 4) Define the travel cost $d(i, j)$ as follows:

$$d(i, j) \triangleq \begin{cases} m(\tau_i, \tau_j) + \frac{1}{2}C_i + \frac{1}{2}C_j, & \text{if } i, j \notin \{s, t\}, \\ m(\tau_0, \tau_j) + \frac{1}{2}C_j + B, & \text{if } i = s, \\ \frac{1}{2}C_i + B, & \text{if } j = t, \end{cases}$$

where B is a constant chosen to be both $> \max_{i,j} m(\tau_i, \tau_j)$ and $> \frac{D_1}{2}$.

- 5) Set the time budget $T_{\max} = 2B + D_1$.

Intuitively, the reduction distributes the dwell-time of a tile onto its incident edges, so that a path passing through the corresponding node once pays the dwell time. The large value B enforces the triangle inequality when s and t are added (our heuristics and known approximations for the OP require metric distances) and prevents paths that meet the time budget from visiting these nodes multiple times or using them as shortcuts.

Lemma 2. The constructed graph for the orienteering problem in Section VI defines a metric distance on its edges.

Proof. Let u, v, w be any three vertices of this graph that are connected in a triangle. We show the triangle inequality holds.

If none of u, v, w are s or t , we have that $d(u, v) + d(v, w)$

$$\begin{aligned} & = m(\tau_u, \tau_v) + m(\tau_v, \tau_w) + \frac{1}{2}C_u + C_v + \frac{1}{2}C_w \\ & \geq m(\tau_u, \tau_w) + \frac{1}{2}C_u + \frac{1}{2}C_w = d(u, w). \end{aligned}$$

If (WLOG) the triple is s, u, v , neither u nor v can be t . The edges involving s each have weight at least B , so $d(u, s) + d(s, v) > d(u, v)$. Moreover, $d(u, v) + d(v, s)$

$$\begin{aligned} &= m(\tau_u, \tau_v) + m(\tau_v, \tau_0) + \frac{1}{2}C_u + C_v + B \\ &\geq m(\tau_u, \tau_0) + \frac{1}{2}C_u + B = d(u, s). \end{aligned}$$

If (WLOG) the triple is u, v, t , neither u nor v can be s . The edges involving t each have weight at least B , so $d(u, t) + d(t, v) > d(u, v)$. Moreover, $d(u, v) + d(v, t)$

$$= m(\tau_u, \tau_v) + \frac{1}{2}C_u + C_v + B \geq \frac{1}{2}C_u + B = d(u, t) \quad \square$$

Example 3. To illustrate the reduction, we return to the tile configuration shown in Fig. 3. Assume this time that the optical telescope is initially focused on tile τ_2 . Since the probability $p_1 = 0$, we exclude tile τ_1 from consideration.




Fig. 5. Extending the example from Fig. 3 to illustrate the reduction to OP, assuming the optical telescope starts at τ_2 .

A **start node** s is added; its distance to **node 2** is $d(s, 2) = C_2/2 + B = 0.5 + B$ since no actual movement is required to reach tile τ_2 , but still captures its dwell time. Distances to **nodes 3 and 4** do capture the telescope's movement times: $d(s, 3) = m(\tau_2, \tau_3) + C_3/2 + B = 2 + 1/2 + B = 2.5 + B$, and similarly for $d(s, 4)$. Distances among **nodes 2, 3, and 4** capture the dwell times of both corresponding tiles, e.g., $d(2, 4) = m(\tau_2, \tau_4) + C_2/2 + C_4/2 = 2 + 1/2 + 10/2 = 7.5$. An **artificial end node** t is also added. Its distances from **nodes 2, 3, and 4** are computed as $d(i, t) = C_i/2 + B$, e.g., $d(4, t) = 10/2 + B$, which allows the dwell time of tile τ_4 to be captured.

Equivalence. Let I_s be an instance of the search problem, and let I_o be its constructed OP instance. There is a close connection between s - t paths in I_o and search orders in I_s :

Lemma 3. I_s has a search order π that begins at τ_0 and achieves probability P of detection by deadline D_1 iff I_o has a path that finishes by $T_{\max}$ and achieves total reward P .

Proof. For any valid search order π in I_s , there is a unique corresponding *simple* orienteering path Π in I_o obtained by dropping τ_0 from π and adding endpoints s, t to the corresponding path of nodes. Let $c(\pi)$ be the total cost (traversal plus dwell time) associated with π , and let $\ell(\Pi)$ be the total length of Π . Finally, let $r(\Pi)$ be the total reward of all nodes in an orienteering path, and or equivalently the total probability of all tiles (except τ_0) in a search order.

Observation 1: $\ell(\Pi) = c(\pi) + 2B$. Indeed, we know that

$$c(\pi) = \sum_{(\tau_u, \tau_v) \in \pi} m(\tau_u, \tau_v) + \sum_{\tau_v \in \pi} C_{\tau_v}.$$

Because Π has no repeated nodes and incurs the added length B only at its first and last edges,

$$\ell(\Pi) = \sum_{(u,v) \in \Pi} d(u, v) = \sum_{(\tau_u, \tau_v) \in \pi} m(\tau_u, \tau_v) + \sum_{\tau_v \in \pi} C_{\tau_v} + 2B.$$

Observation 2: A path in I_o from s to t of length $\leq 2B + D$ never visits s or t except at its endpoints. Indeed, if path Π visits node $z \in \{s, t\}$ other than at its endpoints, the two edges to and from z have length at least $2B$, in addition to the length of at least $2B$ required to leave s initially and arrive at t finally. Hence, the total length of the path must be at least $4B$, which is greater than $2B + D$ by our choice of B .

($\rightarrow$) Given search order π in I_s that starts with τ_0 and ends at τ_v by time D , set Π to $[s, \pi - \{\tau_0\}, t]$. By Observation 1, $\ell(\Pi) = c(\pi) + 2B$, and so $\ell(\Pi) \leq D + 2B$. Moreover, $r(\Pi) = r(\pi)$, since the nodes of nonzero reward touched by Π exactly match the tiles of nonzero probability touched by π .

($\leftarrow$) Given Π in I_o of length $\leq 2B + D$, by Observation 2, Π has s and t only at its endpoints. Let Π' be the simple path obtained by enumerating nodes of Π in order of their first appearance. Π' exists in I_o because its graph is complete; by the triangle inequality, $\ell(\Pi') \leq \ell(\Pi) \leq 2B + D$. Now form π in I_s by removing the endpoints of Π' and prepending τ_0 ; by Observation 1, $c(\pi) = \ell(\Pi) - 2B \leq D$. We have $r(\Pi) = r(\pi)$ by correspondence of the two paths' node and tile sets. $\square$

Hardness of Search Problem. Blum et al. [17, Theorem 7.2] show that OP (with fixed start but not fixed end node) is APX-hard. Our search problem generalizes this formulation of OP; hence, we have the following:

Corollary 3. The search problem (Section I) is APX-hard, even if there is only one deadline and all dwell times $C_i = 0$.

This result indicates that exact methods are generally too time-consuming for practical applications, requiring heuristic approaches. Building on Section IV, the next section proposes GCP, a real-time-capable heuristic for the search problem.

Approximability of Search Problem. The reduction from search to OP is approximation-preserving, so approximation algorithms for OP give approximation algorithms for search.

Corollary 4. If A is an algorithm that finds an α -approximation to the max-reward path for the OP, then A can be used to construct an α -approximation for the search problem.

Proof. Based on the correspondence established in Lemma 3, given an optimal path Π for orienteering instance I_o built from search instance I_s , the corresponding search order π in I_s must also be optimal. Moreover, if A finds a feasible path for I_o with reward P , the corresponding search order also has total probability P . Thus, A 's approximation ratio vs the optimum in I_o also applies to the induced search order in I_s . $\square$

VII. GREEDY-CHRISTOFIDES PATHFINDING ALGORITHM

Chekuri et al. [26] give a $(2+\varepsilon)$ -approximation algorithm for the orienteering problem (OP) that employs a recursive dynamic programming framework with tree decompositions. While this




Fig. 6. Probability map for GRB from COSI instrument simulations [28].

method offers strong theoretical guarantees via Corollary 4, its time complexity is impractical for real-time decision-making, such as telescope follow-up for transient events, where a path must be computed within seconds of an alert.

To better balance computational speed with solution quality, we propose an efficient heuristic, Greedy-Christofides Pathfinding (GCP), outlined in Algorithm 1. Our approach decomposes the planning process into two stages: (i) subset selection and (ii) deadline-aware s - t path construction.

Subset Selection. We apply two parallel strategies to generate candidate subsets of tiles: a *probability-based strategy* that selects the top k tiles ranked by probability, and a *spatially adaptive strategy* that incrementally grows the subset by greedily adding tiles with the highest probability-to-distance ratio. Distance here is measured from the nearest already selected tile. For a given candidate subset of tiles, GCP will construct an s - t path that touches the corresponding set of nodes **in the constructed OP instance**.

These complementary strategies are motivated by spatial patterns commonly observed in gravitational-wave probability maps. The probability-based approach is well-suited to compact, high-density confidence regions (e.g., Fig. 6). In contrast, the spatially adaptive strategy is more effective when probability mass is dispersed across multiple clusters, especially when these clusters are far apart (e.g., Fig. 2). We run our heuristic with both strategies in parallel and keep the solution with highest utility, thereby allowing it to adapt dynamically to different topological structures in the sky map.

Path Construction. Given a candidate subset of tiles, we construct an s - t path (corresponding to a search order) visiting every tile in this subset using Hooogeveen’s modification of Christofides’ minimum-spanning-tree heuristic [44], which yields a $\frac{5}{3}$ -approximation for the minimum-weight undirected s - t path spanning the chosen subset.

Two-Phase Search. The heuristic uses a two-phase search to determine the largest feasible subset that yields a valid path within the time budget. The search is performed separately, in parallel, for the two subset selection strategies. In the initial *doubling phase*, the subset size k is incrementally doubled from a small initial value. For each k , an MST is computed over the tiles selected by the considered strategy. Since the MST cost provides a lower bound on any feasible s - t path, if it exceeds the budget, further expansion is halted.

The subsequent *binary search phase* searches between the last feasible and first infeasible subset sizes. For each size k' , full s - t paths are constructed using the modified Christofides procedure on the selected tiles. If the path cost lies within the

Algorithm 1 Greedy-Christofides Pathfinding (GCP)

Require: Set of tiles τ , rewards $\{p_i : \tau_i \in \tau\}$, travel cost function $d(i, j)$, time budget T , start node s , end node t

- 1: best path $\pi_{\text{best}} \leftarrow \emptyset$, best reward sum $P_{\text{best}} \leftarrow 0$
- 2: **for** each tile-sorting strategy **do**
- 3: subset size $k \leftarrow 2$ $\triangleright$ *Doubling phase*
- 4: **while** true **do**
- 5: Select top k tiles
- 6: Construct MST over selected tiles using $d(\tau_i, \tau_j)$
- 7: **if** MST cost exceeds T **then**
- 8: **break**
- 9: $k \leftarrow 2k$
- 10: $k_{\min} \leftarrow \lfloor k/2 \rfloor$, $k_{\max} \leftarrow k$ $\triangleright$ *Binary search phase*
- 11: **while** $k_{\max} - k_{\min} > 1$ **do**
- 12: $k' \leftarrow \lfloor (k_{\min} + k_{\max})/2 \rfloor$
- 13: Select top k' tiles
- 14: Construct MST over selected tiles
- 15: **if** MST cost exceeds T **then**
- 16: $k_{\max} \leftarrow k'$
- 17: **else**
- 18: Construct s - t path using Hooogeveen’s algorithm
- 19: Compute path cost C and total reward P
- 20: **if** $C \leq T$ **then**
- 21: update $\pi_{\text{best}}, P_{\text{best}}$
- 22: $k_{\min} \leftarrow k'$
- 23: **else**
- 24: $k_{\max} \leftarrow k'$
- 25: **return** $\pi_{\text{best}}, P_{\text{best}}$

time budget, the best solution is updated, and the algorithm explores larger subsets; otherwise, it searches smaller ones. This continues until a largest feasible subset is identified.

For each of the two subset selection strategies, this two-phase search procedure produces a candidate s - t path that satisfies the time budget constraint. The final output is the path with the higher expected reward among the two.

The heuristic is designed to run efficiently while adapting to real-time constraints. To further improve path quality without significantly increasing computation time, we apply two lightweight post-processing steps. First, a single-pass 2-opt heuristic is applied to remove potential edge crossings and reduce travel cost. Then, we perform a greedy insertion of unvisited nodes before the terminal node t , selecting the highest-reward node that can be added without violating the time budget. This allows the algorithm to collect additional rewards within the remaining budget opportunistically.

Time Complexity. GCP’s runtime is bounded by $\mathcal{O}(n^3 \log n)$. This guarantee makes GCP particularly suited for real-time search problems. Moreover, GCP achieves strong empirical performance across realistic TAP follow-up scenarios consistently identifying high-quality search orders, as shown in Section IX.

VIII. MULTI-DEADLINE APPROXIMATIONS

We further leverage the orienteering problem (OP) and Corollary 4 to exhibit the polynomial-time tractability of our search problem when extended to the two-deadline and k -deadline cases (where here $k \triangleq z - 1$ is the number of mission FoMs, for z as defined in Section III). The theoretical approach to leveraging single-deadline algorithms to give k -deadline

guarantees is also efficient enough to inform heuristics for k -deadline instances.

Here we let SP_k denote our TAP search problem with k deadlines, and let A be an algorithm for SP_1 with approximation ratio α . Corollary 4 together with prior work on the OP [26] provides a P-time A with $\alpha = 2 + \varepsilon$ for any fixed $\varepsilon > 0$. We also let the *merit aspect ratio* of an instance be $\beta \triangleq M_1/M_k \geq 1$ and let $r(\pi)$ denote the total reward of search order π .

The Two-Deadline Search Problem. We start with a simple reduction from SP_2 to SP_1 .

Intuitively, the following theorem shows that an approximation can be obtained by solving two instances of the single-deadline problem, one for each deadline D_j , with rewards scaled by M_j , then taking the path which gives the better collected (scaled) reward.

Theorem 2. *Given an instance I_s of SP_2 , let I_s^1 and I_s^2 be derived by setting $M_2 \leftarrow 0$ in I_s and by setting $M_1 \leftarrow M_2$ in I_s , respectively, and let $\pi^i = A(I_s^i)$ for $i \in \{1, 2\}$. Then the π^i maximizing $r(\pi^i)$ is a 2α -approximation for SP_2 .*

Proof. Let the optimal two-stage solution to I_s be $\pi^* = \pi_1^* \sqcup \pi_2^*$, where π_1^* is the sequence of nodes visited (and rewards collected) by OPT before deadline D_1 , and π_2^* is the nodes visited and rewards collected by OPT between D_1 and D_2 . Note that $r(\pi^*) = r_1(\pi_1^*) + r_2(\pi_2^*)$, where $r_1(\cdot)$ and $r_2(\cdot)$ denote the rewards of paths in instances I_s^1 and I_s^2 , respectively.

Note also that π_1^* is feasible for I_s^1 and $\pi_1^* \sqcup \pi_2^*$ is feasible for I_s^2 . Then on the one hand, by our SP_1 algorithm we have

$$\alpha \cdot r(\pi_1) = \alpha \cdot r_1(\pi_1) \geq r_1(\text{OPT}(I_s^1)) \geq r_1(\pi_1^*).$$

On the other hand, applying A to I_s^2 gives

$$\alpha \cdot r(\pi_2) \geq \alpha \cdot r_2(\pi_2) \geq r_2(\text{OPT}(I_s^2)) \geq r_2(\pi_2^*).$$

Combining, we have $2\alpha \cdot \max(r(\pi_1), r(\pi_2)) \geq r(\pi^*)$. $\square$

We can also provide a β -dependent improvement:

Theorem 3. *Given an instance I_s of SP_2 , let I_s^2 be defined as in Theorem 2. Then $A(I_s^2)$ is an $\alpha \cdot \beta$ -approximation.*

Proof. Let the optimal two-stage path on I_s be $\pi^* = \pi_1^* \sqcup \pi_2^*$. Furthermore let $\pi = A(I_s^2)$, and let $\text{OPT}(I_s^2)$ denote the optimal one-stage path. By the guarantee for A on I_s^2 ,

$$\alpha \cdot r(\pi) \geq \alpha \cdot r_2(\pi) \geq r_2(\text{OPT}(I_s^2)) \geq r_2(\pi^*),$$

while, recalling that $\beta \triangleq M_1/M_2$,

$$\begin{aligned} \beta \cdot r_2(\pi^*) &= \beta \cdot r_2(\pi_1^*) + \beta \cdot r_2(\pi_2^*) \\ &\geq \beta \cdot r_2(\pi_1^*) + r_2(\pi_2^*) = r(\pi^*). \end{aligned}$$

Combining these inequalities gives the stated claim. $\square$

Taken together, this implies that the algorithm proposed in Theorem 2 is an $\alpha \cdot \min(2, \beta)$ -approximation for SP_2 .

The k -Deadline Search Problem. A natural approach for the many-deadline case is to reduce to *orienteering with time windows* [12], wherein node rewards arrive and depart at

specified times. The closest analog is the common-start, few-deadline variant, for which Chekuri and Kumar [25] claim a polynomial-time $(\alpha + 1)$ -approximation for fixed k (and polynomial discretization of time). However, it is unclear if the approximation-preserving reduction to OP (Corollary 4) generalizes to multiple deadlines. Nevertheless, we provide a k -deadline guarantee by generalising Theorem 2.

Intuitively, the following theorem suggests solving an instance of the single-deadline problem for each of the k deadlines, with the rewards scaled by the corresponding FoM, then taking the best such path.

Theorem 4. *Given an instance I_s of SP_k , for each $i \in [k]$ let I_s^i be the derived SP_1 instance with deadline D_i and FoM M_i . Let each $\pi_i = A(I_s^i)$ be an α -approximation to the SP_1 instance I_s^i . Then the π_i maximizing $r(\pi_i)$ is an $\alpha \cdot \min(k, 2 \cdot \log_2 \lceil \beta \rceil)$ approximation.*

Proof. Let the optimal k -stage solution to I_s be $\pi^* = \pi_1^* \sqcup \dots \sqcup \pi_k^*$, where π_i^* is the sequence of tiles visited (and rewards collected) by π^* between deadlines D_{i-1} and D_i . Let $r_i(\cdot)$ be the rewards of paths in instances I_s^i ; then $r(\pi^*) = \sum_i r_i(\pi_i^*)$.

If $\pi_{\leq i}^*$ denotes the prefix of π^* before deadline D_i , then note also that each π_i^* is feasible for I_s^i . Then on the one hand, by our SP_1 algorithm A and since FoMs are decreasing,

$$\alpha \cdot r(\pi_i) \geq \alpha \cdot r_1(\pi_i) \geq r_i(\text{OPT}(I_s^i)) \geq r_i(\pi_{\leq i}^*) \geq r_i(\pi_i^*).$$

Summing over all $i \in [k]$ then yields

$$\alpha \cdot \sum_{i \in [k]} r(\pi_i) \geq \sum_{i \in [k]} r_i(\pi_i^*) = r(\pi^*).$$

Since the maximum upper-bounds the average, we have

$$\alpha \cdot k \cdot \max_{i \in [k]} (r(\pi_i)) \geq \alpha \cdot k \cdot \frac{1}{k} \sum_{i \in [k]} r(\pi_i) \geq r(\pi^*).$$

This demonstrates the proposed algorithm is an $\alpha \cdot k$ -approximation. The $\log \beta$ guarantee follows directly by bucketing together phases that have FoMs within a factor of two. $\square$

IX. EVALUATION

We compare both the *solution quality* (measured in terms of the expected FoM $\mathbb{E}[\mathcal{M}]$) and the *computational efficiency* of the proposed **GCP** (Algorithm 1) to the following baselines:

- **Genetic**: a stochastic baseline that explores the solution space using evolutionary search.
- **Greedy**: a standard heuristic in the astrophysics literature (e.g., [66, 67]) that at each step selects the unvisited tile with the highest probability.
- **ILP**: the optimal solution obtained from an exact integer-linear-programming formulation following [78, Section 2.3.1], solved using Gurobi [43], and used as the ground truth for evaluating approximation quality.

The experiments are conducted on a collection of 37 sky probability maps derived from real LIGO events [2, 4]. They were run on a server with $2 \times$ Intel Xeon Gold 5118 12-core CPUs with HyperThreading enabled and 196 GB of RAM.

A. Experiment Setup

The dataset includes 13 small- and 24 large-scale instances, curated to reflect diverse probability distributions derived from real non-optical instruments. We use the methods in [37] to define tiles by projecting the telescope’s FoV over HEALPix representations of the LIGO probability maps, fixing the pixel with the highest probability directly above the telescope, then considering just those pixels within the visible hemisphere.

We define a small instance as a probability map in which the top 100 tiles, ranked by probability, collectively account for at least 80% of the total probability mass. To determine this, we extract the top 100 (or all tiles if fewer than 100 exist), compute their cumulative probability, and use them as the test set if the cumulative mass exceeds 80%. Otherwise, the instance is discarded. This setup ensures computational tractability for comparison with the ILP baseline and enables a detailed assessment of solution quality. To generate a sufficient number of small instances, we project a $10 \times 10^\circ$ FoV; optical observatories that are equipped with multiple telescopes can be configured to provide such a wide FoV by arranging the individual telescopes so that their FoVs are adjacent.

We define large-scale probability maps as instances where the probability is broadly distributed across the sky. Specifically, we retain the full 99% containment region for each selected event, which typically comprises 1,000–6,000 tiles when projecting a smaller $2.5 \times 2.5^\circ$ FoV. We keep any instances where 1,000–10,000 tiles are retained. These instances are intended to emulate realistic operational scenarios in telescope scheduling, where sky localization regions are large and resources are significantly constrained.

Each tile is associated with a detection probability, and the pairwise traversal cost between tiles incorporates both angular slew time and dwell time. The start tile for each instance was taken to be the tile of highest probability so as to avoid measuring the confounding, instance-dependent effect of an initial slew from an arbitrary starting position. For each instance, the objective is to compute a search order that maximizes the expected FoM as per Lemma 1.

Dwell Times. A complete derivation of the dwell time C_i for each tile τ_i requires in-depth knowledge of the participating telescope’s optical properties, the TAP’s flux as indicated by the high-energy detecting instrument, and atmospheric opacity in the wavelengths corresponding to the telescope’s spectral sensitivity. As a simple approximation to provide sufficient realism for illustrative purposes in our evaluation, we consider that the dwell time increases as the telescope approaches the horizon due to the additional atmosphere through which it must look. We compute air mass for a given zenith angle θ in radians using the Kasten & Young model [47]:

$$AM = \left[\cos \theta + 0.50572 (96.07995 - 180\theta/\pi)^{-1.6364} \right]^{-1}$$

We then use the Simple Model of the Atmospheric Radiative Transfer of Sunshine (SMARTS) [41] to estimate the rela-




Fig. 7. Average percentage deviation from ILP baseline across deadlines. Shaded region denotes observed deviation ranges.

tive irradiance of 780 nm light² as a function of air mass: $s = 1.1129e^{-0.107AM}$. Under the assumptions that (i) a target signal-to-noise (SNR) ratio in the optical data is required to identify the TAP in a given tile, (ii) signal scales with exposure time, and (iii) noise scales with the square root of exposure time, dwell time is $C_i = C_0/s^2$, where C_0 is the dwell time at the zenith; for our experiments, we assume C_0 is 1 second.

Move Times. The time required for the telescope to move between two sky tiles depends on the angular distance between them and the telescope’s slew rate. Each tile τ_i is associated with a sky position specified in equatorial coordinates (α_i, δ_i) . Given two tiles τ_i and τ_j , the angular separation θ_{ij} between them is computed using the spherical law of cosines:

$$\theta_{ij} = \arccos [\sin(\delta_i) \sin(\delta_j) + \cos(\delta_i) \cos(\delta_j) \cos(\alpha_i - \alpha_j)]$$

In our experiments, we adopt a constant slew rate of $50^\circ/\text{s}$, consistent with the fast-slewing direct-drive Planewave mounts used by recent and upcoming observatories, e.g., 7DT [48].

B. Expected FoM Comparison on Small Maps

Goal and setup: This experiment evaluates the approximation quality of **GCP**, **Genetic**, and **Greedy** relative to the **ILP** baseline on small-scale search instances, where exact solutions are computationally feasible. The objective is to maximize the expected FoM $\mathbb{E}[\mathcal{M}]$ under a single-deadline setup across deadlines ranging from 10 to 90 seconds in increments of 10. The reported metric is the percentage deviation from the **ILP**-optimal $\mathbb{E}[\mathcal{M}]$, averaged over 13 probability maps.

Observations: As shown in Fig. 7, **GCP** consistently achieves deviations within 0.05% of the optimal $\mathbb{E}[\mathcal{M}]$ across all deadlines, exhibiting the best performance overall. The **Genetic** method follows closely, with slightly higher deviations. In contrast, the **Greedy** baseline shows significantly worse performance, especially under tight deadlines (e.g., 10–50 seconds), where its average deviation reaches up to 7%. The shaded area (indicating 95% confidence intervals) around the Greedy curve shows substantial variability, while **GCP** and **Genetic** results remain consistently stable. In all these trials, **GCP** returned a search order in around 0.01 s, compared to around 1 s for the **Genetic** method and $<10^{-4}$ s for **Greedy**.

Conclusions: These results confirm that GCP offers robust and near-optimal performance on small-scale instances. It out-

²780 nm red light is at the end of the visible spectrum, and is captured by the near-infrared filters used in several sky surveys [24, 53, 64, 72, 82].




Fig. 8. **Left:** Percentage deviation in FoM from the GCP baseline. **Right:** Computation time vs Deadline presented in log-scaled plot.

performs both heuristic baselines in accuracy while maintaining significantly better runtime efficiency over the **Genetic** method. Its minimal deviation from ILP warrants its use in time-sensitive settings where exact solutions are computationally infeasible.

C. Expected FoM and Efficiency Comparison on Large Maps

Goal and setup: Similarly to the previous experiment, this experiment’s goal is to evaluate the scalability and solution quality of **GCP** and the baselines. This time, we consider the large-scale search problem instances, where exact methods such as **ILP** are computationally *infeasible*. As before, the objective is to maximize $\mathbb{E}[\mathcal{M}]$ across mission deadlines, this time ranging from 10 to 1200 seconds in the set $\{10, 100, 200, \dots, 1200\}$. The reported metric is the percentage deviation from the **GCP**-resulting $\mathbb{E}[\mathcal{M}]$, averaged over the 24 probability maps.

Observations: Fig. 8_{Left} shows that **GCP** consistently yields the highest $\mathbb{E}[\mathcal{M}]$ result across all deadlines. **Genetic** follows closely (within at least 1.6% of the **GCP**’s $\mathbb{E}[\mathcal{M}]$) but consistently remains below **GCP** across all deadlines, with larger deviations observed under shorter ones (100 to 500 seconds). **Greedy** shows several orders of magnitude lower performance, and is thus not shown in the figure.

Fig. 8_{Right} reports computation times with the shaded area indicating 95% confidence interval. **Greedy** is the fastest method, with execution times consistently below 10⁻⁴ seconds, but it also yields the lowest $\mathbb{E}[\mathcal{M}]$. **Genetic** is the slowest and most variable, with runtimes occasionally approaching 100 seconds. **GCP** strikes a favorable balance, exhibiting predictable scaling with deadlines and maintaining execution times below 10 seconds, with an average of 1 second across all evaluated horizons. It is important to note that **GCP** is implemented as a *dual-threaded* algorithm, whereas **Genetic** parallelizes across all 24 physical cores.

Conclusions: These results reconfirm **GCP**’s advantage in achieving a robust trade-off between computation time and solution quality even in realistic large-scale search scenarios. It consistently achieves nearly the highest $\mathbb{E}[\mathcal{M}]$ results while maintaining predictable computation time on limited cores, making it a desirable choice for real-time telescope scheduling.

D. Impact of Computation Time on the Expected FoM

Goal and setup: In real-time optical follow-up missions, the utility of a TAP search strategy depends not only on the quality of the computed path but also on how quickly it can be produced. Since planning itself consumes time, algorithms must




Fig. 9. Results after accounting for MOET. **Left:** Average Figure of Merit vs Deadline. **Right:** FoM percentage deviation from GCP.

account for this delay when operating under mission deadlines. Let D_k be an arbitrary deadline and WCET_m the worst-case computation time of method $m \in \{\mathbf{GCP}, \mathbf{Genetic}, \mathbf{Greedy}\}$. Then the effective planning horizon is $D'_k = D_k - \text{WCET}_m$ as the generated search plan can begin execution only once computation has completed. This experiment evaluates the resulting effect on overall expected FoM for **GCP**, **Genetic**, and **Greedy**. For each combination of method, deadline, and sky map, we ran five trials and recorded the maximum observed execution time (MOET); these values were then subtracted from the nominal deadlines before computing the search orders.

Observations: The MOET was 67 seconds for **Genetic**, 197 ms for **GCP**, and 108 μ s for **Greedy**. Fig. 8_{Left} shows the average $\mathbb{E}[\mathcal{M}]$ of each method across deadlines, while Fig. 8_{Right} reports the corresponding $\mathbb{E}[\mathcal{M}]$ deviation from **GCP** under the adjusted deadlines D'_k . Both plots suggest that when $\text{MOET}_{\mathbf{Genetic}}$ is accounted for, **Genetic** suffers a significant $\mathbb{E}[\mathcal{M}]$ degradation, particularly at short deadlines where the effective planning window is severely reduced. And even accounting for its MOET, **GCP** retains an $\mathbb{E}[\mathcal{M}]$ at least as high as **Greedy**’s, even at short deadlines where **Greedy**’s short MOET allows a larger search budget.

Conclusions: The results demonstrate that accounting for computation time is essential in time-sensitive missions. Methods like **GCP** that maintain low and predictable runtime are better suited for real-time applications, as they preserve more of the mission deadline for actual execution. In contrast, slower methods like **Genetic** may lose significant utility, as the solver consumes a non-negligible portion of the planning horizon.

E. A Closer Look at Response-Time Awareness

Goal and setup: This experiment evaluates **GCP** in a multi-deadline planning context, where the utility of a successful detection depends on its timing. As formalized in Section III, each deadline D_i is associated with an FoM M_i , representing decreasing scientific value for later detections. Once again, the objective is to find a search order that maximizes $\mathbb{E}[\mathcal{M}]$. We consider a timeline with deadlines $\{D_1, D_2, D_3\} = \{100, 200, 500\}$ seconds and associated FoMs $\{M_1, M_2, M_3\} = \{1.0, 0.5, 0.2\}$. For each deadline, we report the cumulative value $E(D_k)$ of the partial value of $\mathbb{E}[\mathcal{M}]$ collected up to D_k , i.e., more formally defined $E(D_k) \triangleq \sum_{i=1}^k M_i \cdot (F_{\mathcal{R}_\pi}(D_i) - F_{\mathcal{R}_\pi}(D_{i-1}))$. This is done for three search orders (**SO1**, **SO2** and **SO3**), each constructed according to the principles formalized in Theorem 4:




Fig. 10. Expected FoM progression over the designated deadlines for skymaps GW200322_091133 (**Top**) and GW200216_220804 (**Bottom**). Probability maps shown are limited to the telescope’s visible hemisphere.

- **SO1**: $GCP(D_1)$ followed by $GCP(D_3 - D_1)$
- **SO2**: $GCP(D_2)$ followed by $GCP(D_3 - D_2)$
- **SO3**: $GCP(D_3)$ (single long path)

The results are reported for two illustrated maps (see Fig. 10) that have significantly different probability distributions.

Observations: The results show that the concerns laid down in Section IV remain valid even for real spatial probability maps, namely *an algorithm should explore a broad search space as local strategies may risk hindering mission’s performance*. For instance, in Fig. 10^{top}, **SO1** is preferred for $D_1 = 100$ while **SO3** has the lowest partial FoM score. However, through deadline D_3 this ranking is completely reversed. A similar inversion occurs between **SO2** and **SO1** in Fig. 10^{Bottom}.

Conclusion: Notably, the application of Theorem 4 using **GCP** enables adaptive planning without succumbing to short-sighted, locally optimal choices that may jeopardize success.

X. RELATED WORK

Timing aspects of space-related applications have long been recognized within the real-time systems community (e.g., [50–52, 54, 74]). To model them in our problem (Section V), we adapted the well-known concept of response-time analysis [46]. To consider the probabilistic counterpart, we adapted the multi-evolution semantics developed in recent works on probabilistic analysis [20, 55, 56]. Furthermore, Theorem 1 utilizes the temporal aspects of stochastic dominance between response-time distributions, a topic extensively studied in real-time systems (e.g., [20, 32–34]). For a broader overview, we refer the reader to the surveys by Davis and Cucu-Grosjean [32, 33].

In Section VI, we showed that the considered search problem reduces to the orienteering problem (OP). The OP, introduced in the seminal work by Golden et al. [38], has myriad variants, including the s - t orienteering problem, team orienteering, and time-dependent orienteering. For comprehensive overviews, we refer the reader to surveys by Gunawan et al. [42], Saller et al. [69], Vansteenwegen et al. [79]. Related to OP is the budgeted prize-collecting traveling salesperson problem, which aims to maximize the number of visited nodes under one or more time

constraints [62, 63]. Prior work on heuristic and exact methods has also considered dwell times (i.e. wait costs, service times, processing times) [10, 27, 70]; but the single-deadline reduction to OP (Section VI) is to our knowledge novel.

The astrophysics community has developed a rich body of work focused on search strategies for transient astrophysical phenomena (TAPs), such as kilonovae. Early efforts focused on prompt follow-up observations using optical telescopes based on alerts from wide-field gamma-ray detectors (e.g., [36]). More recent work has explored coordinated tiling strategies to optimize sky coverage under visibility and telescope constraints [5, 7, 30, 66]. These methods primarily focus on HEALPix-based probability maps, attempting to balance latency, sensitivity, detection accuracy, and FoV limitations.

The key distinction of our work is that prior approaches do not explicitly model expected scientific utility as a stochastic, time-aware optimization problem, nor do they incorporate a strict, deadline-driven timing model as we propose in this paper.

XI. CONCLUSIONS AND FUTURE WORK

The problem of planning efficient telescope observations for transient astrophysical phenomena (TAP) lies at the intersection of real-time systems, operations research, astrophysics and cyber-physical systems. Addressing this challenge requires both theoretical rigor and practical efficiency, particularly under stringent timing constraints. In this work, we introduced the first formal stochastic, response-time and scientific-merit aware model for TAP search planning, capturing the temporal and probabilistic nature of the problem. We showed that the problem is APX-hard and reduces to a variant of the Orienteering Problem, motivating the design of a novel, real-time-capable heuristic: *Greedy-Christofides Pathfinding (GCP)*. GCP achieves a strong trade-off between computational efficiency and solution quality by combining probability-based and spatially adaptive strategies with deadline-aware path planning. For multiple deadlines, we established worst-case performance guarantees by proving an $\alpha \cdot \min(k, 2 \log_2 \lceil \beta \rceil)$ -approximation bound.

Empirical results on real LIGO gravitational-wave probability maps demonstrate that: (i) for ILP-tractable problem instances, GCP yields near-optimal solutions with predictable runtime; (ii) it consistently outperforms baseline heuristics in solution quality; (iii) it remains robust across diverse probability-map scenarios, maintaining high solution quality under realistic operational conditions. Finally, (iv) results suggest that GCP is an effective planning method for real-time TAP localization, enabling broad and adaptive search without succumbing to short-sighted strategies, constituting a step toward reliable scheduling pipelines for next-generation time-domain astronomy.

Opportunities remain to consider generalizations of our model in future work. For example, we might consider that a TAP’s optical luminosity changes as it evolves, giving rise to non-constant dwell times with increasing detection probability over longer observations. Furthermore, the well-separatedness of the LIGO probability map in Figure 2 suggests that *multiple* optical telescopes could be deployed for coordinated search.

ACKNOWLEDGMENT

This work was supported by NSF grants CNS-2141256, CNS-2229290 (CPS), and CNS-2502855 (CPS), NASA award 80NSSC21K1741, and a Washington University OVCR seed grant. We would like to thank the reviewers for their thorough and insightful feedback. We would also like to acknowledge the members of the ADAPT collaboration, which inspired this work. To Matthew Andrew, Blake Bal, Elisabetta E. Bissaldi, Richard G. Bose, Dana Braun, James H. Buckley, Eric Burns, Marco M. Cecca, Davide Cerasole, Roger D. Chamberlain, Wenlei Chen, Michael L. Cherry, Federica F. Cuna, Gaia G. De Palma, Davide D. Depalo, Riccardo R. Di Tria, Leonardo L. Di Venere, Jeffrey Dumonthier, Manel Errando, Stefan Funk, Fabio F. Gargano, Priya Ghosh, Francesco F. Giordano, Jonah Hoffman, Aldana A. Holzmann Airasca, Ye Htet, Zachary Hughes, Aera Jung, Patrick L. Kelly, John F. Krizmanic, Makiko Kuwahara, Calvin Lee, Francesco F. Licciulli, Antonio A. Liguori, Gang Liu, Pierpaolo P. Loizzo, Leonarda L. Lorusso, Mario Nicola Mazziotta, John Grant Mitchell, John W. Mitchell, Boris Murmann, Georgia A. de Nolfo, Jennifer Ott, Giuliana G. Panzarini, Richard Peschke, Riccardo Paoletti, Roberta R. Pillera, Brian Rauch, Davide D. Serini, Garry Simburger, George Suarez, Teresa Tatoli, Gary S. Varner, Eric A. Wulf, Adrian Zink, and Wolfgang V. Zober, we extend our thanks.

REFERENCES

- [1] M. G. Aartsen *et al.*, “The IceCube Neutrino Observatory: instrumentation and online systems,” *Journal of Instrumentation*, vol. 12, no. 03, p. P03012, 2017.
- [2] A. G. Abac, R. Abbott, I. Abouelfettouh, F. Acernese *et al.*, “Observation of gravitational waves from the coalescence of a 2.5–4.5 solar mass compact object and a neutron star,” *The Astrophysical Journal Letters*, vol. 970, no. 2, p. L34, Jul. 2024. [Online]. Available: <http://dx.doi.org/10.3847/2041-8213/ad5beb>
- [3] B. Abbott *et al.*, “LIGO: The laser interferometer gravitational-wave observatory,” *Reports on Progress in Physics*, vol. 72, no. 7, p. 076901, 2009.
- [4] R. Abbott *et al.*, “Open data from the third observing run of LIGO, Virgo, KAGRA, and GEO,” *The Astrophysical Journal Supplement Series*, vol. 267, no. 2, p. 29, jul 2023. [Online]. Available: <https://dx.doi.org/10.3847/1538-4365/acdc9f>
- [5] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “Multi-messenger observations of a binary neutron star merger,” *The Astrophysical Journal Letters*, vol. 848, no. 2, p. L12, 2017.
- [6] M. Ageron *et al.*, “ANTARES: the first undersea neutrino telescope,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 656, no. 1, pp. 11–38, 2011.
- [7] T. Ahumada, S. Anand, M. W. Coughlin, I. Andreoni, E. C. Kool, H. Kumar, S. Reusch, A. Sagués-Carracedo, R. Stein, S. B. Cenko *et al.*, “In search of short gamma-ray burst optical counterparts with the Zwicky Transient Facility,” *The Astrophysical Journal*, vol. 932, no. 1, p. 40, 2022.
- [8] R. Artola *et al.*, “TOROS optical follow-up of the advanced LIGO–VIRGO O2 second observational campaign,” *Monthly Notices of the Royal Astronomical Society*, vol. 493, no. 2, pp. 2207–2214, 2020.
- [9] K. Asano, “Early gamma-ray afterglow from gamma-ray bursts,” in *Proceedings of 38th International Cosmic Ray Conference — PoS(ICRC2023)*, vol. 444, 2023, p. 633.
- [10] N. Ascheuer, M. Fischetti, and M. Grötschel, “Solving the asymmetric travelling salesman problem with time windows by branch-and-cut,” *Mathematical programming*, vol. 90, pp. 475–506, 2001.
- [11] W. Atwood, A. A. Abdo, M. Ackermann, W. Althouse, B. Anderson, M. Axelsson, L. Baldini, J. Ballet, D. Band, G. Barbiellini *et al.*, “The large area telescope on the Fermi gamma-ray space telescope mission,” *The Astrophysical Journal*, vol. 697, no. 2, p. 1071, 2009.
- [12] N. Bansal, A. Blum, S. Chawla, and A. Meyerson, “Approximation algorithms for deadline-TSP and vehicle routing with time-windows,” in *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, 2004, pp. 166–174.
- [13] S. Baruah and P. Ekberg, “Towards efficient explainability of schedulability properties in real-time systems,” in *35th Euromicro Conference on Real-Time Systems (ECRTS 2023)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023, pp. 2–1.
- [14] E. Berger, “Short-duration gamma-ray bursts,” *Annual review of Astronomy and Astrophysics*, vol. 52, no. 1, pp. 43–105, 2014.
- [15] E. Berger, C. Leibler, R. Chornock, A. Rest, R. Foley, A. M. Soderberg, P. Price, W. Burgett, K. Chambers, H. Flewelling *et al.*, “A search for fast optical transients in the Pan-STARRS1 medium-deep survey: M-dwarf flares, asteroids, limits on extragalactic rates, and implications for LSST,” *The Astrophysical Journal*, vol. 779, no. 1, p. 18, 2013.
- [16] D. Bersanetti *et al.*, “Advanced Virgo: Status of the detector, latest results and future prospects,” *Universe*, vol. 7, no. 9, p. 322, 2021.
- [17] A. Blum, S. Chawla, D. R. Karger, T. Lane, A. Meyerson, and M. Minkoff, “Approximation algorithms for orienteering and discounted-reward TSP,” *SIAM Journal on Computing*, vol. 37, no. 2, pp. 653–670, 2007.
- [18] K. N. Borozdin, S. P. Brumby, M. C. Galassi, K. McGowan, D. Starr, T. Vestrand, R. White, P. Wozniak, and J. A. Wren, “Real-time detection of optical transients with RAPTOR,” in *Astronomical Data Analysis II*, J.-L. Starck and F. D. Murtagh, Eds., vol. 4847, International Society for Optics and Photonics. SPIE, 2002, pp. 344 – 353. [Online]. Available: <https://doi.org/10.1117/12.461102>
- [19] W. J. Borucki, J. M. Jenkins, and R. M. Duren, “Science merit function for the Kepler mission,” *Journal of Astronomical Telescopes, Instruments, and Systems*, vol. 6, no. 4, pp. 044 003–044 003, 2020.
- [20] S. Bozhko, F. Marković, G. von der Brüggen, and B. B. Brandenburg, “What really is pWCET? A rigorous axiomatic proposal,” in *2023 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2023, pp. 13–26.
- [21] D. M. Bramich, “A new algorithm for difference image analysis,” *Monthly Notices of the Royal Astronomical Society: Letters*, vol. 386, no. 1, pp. L77–L81, 05 2008. [Online]. Available: <https://doi.org/10.1111/j.1745-3933.2008.00464.x>
- [22] J. Buckley *et al.*, “The Advanced Particle–astrophysics Telescope (APT) project status,” in *Proc. of 37th Int’l Cosmic Ray Conference*, vol. 395, Jul. 2021, pp. 655:1–655:9.
- [23] J. Buhler and M. Sudvarg, “Real-time Likelihood Map Generation to Localize Short-duration Gamma-ray Transients,” in *Proc. 39th Int’l Cosmic Ray Conf.*, vol. 501, Jul. 2025, pp. 587:1–587:9.
- [24] K. Chambers *et al.*, “The Pan-STARRS1 surveys,” 2019. [Online]. Available: <https://arxiv.org/abs/1612.05560>
- [25] C. Chekuri and A. Kumar, “Maximum coverage problem with group budget constraints and applications,” in *International Workshop on Randomization and Approximation Techniques in Computer Science*. Springer, 2004, pp. 72–83.
- [26] C. Chekuri, N. Korula, and M. Pál, “Improved algorithms for orienteering and related problems,” *ACM Transactions on Algorithms (TALG)*, vol. 8, no. 3, pp. 1–27, 2012.

- [27] N. Christofides, A. Mingozi, and P. Toth, "State-space relaxation procedures for the computation of bounds to routing problems," *Networks*, vol. 11, no. 2, pp. 145–164, 1981.
- [28] Compton Spectrometer and Imager (COSI) Collaboration, "COSI Data Challenges," Apr. 2025. [Online]. Available: <https://doi.org/10.5281/zenodo.15126188>
- [29] M. W. Coughlin, T. Ahumada, S. B. Cenko, V. Cunningham, S. Ghosh, L. P. Singer, E. C. Bellm, E. Burns, K. De, A. Goldstein *et al.*, "2900 square degree search for the optical counterpart of short gamma-ray burst GRB 180523b with the Zwicky Transient Facility," *Publications of the Astronomical Society of the Pacific*, vol. 131, no. 998, p. 048001, 2019.
- [30] M. W. Coughlin *et al.*, "Efficient targeting of gravitational wave sky localizations with Zwicky Transient Facility," *Publications of the Astronomical Society of the Pacific*, vol. 131, no. 1005, p. 048001, 2019.
- [31] K. Danzmann, LISA Study Team *et al.*, "LISA: laser interferometer space antenna for gravitational wave measurements," *Classical and Quantum Gravity*, vol. 13, no. 11A, p. A247, 1996.
- [32] R. I. Davis and L. Cucu-Grosjean, "A survey of probabilistic timing analysis techniques for real-time systems," *Leibniz Transactions on Embedded Systems*, vol. 6, no. 1, pp. 03–1–03:60, 2019.
- [33] —, "A survey of probabilistic schedulability analysis techniques for real-time systems," *Leibniz Transactions on Embedded Systems*, vol. 6, no. 1, pp. 04:1–04:53, 2019.
- [34] J. L. Diaz, J. M. Lopez, M. Garcia, A. M. Campos, K. Kim, and L. L. Bello, "Pessimism in the stochastic analysis of real-time systems: Concept and applications," in *25th IEEE International Real-Time Systems Symposium*. IEEE, 2004, pp. 197–207.
- [35] M. M. Fausnaugh, R. Jayaraman, R. Vanderspek, G. R. Ricker, C. J. Burke, K. D. Colón, S. W. Fleming, H. M. Lewis, S. Mullally, A. Youngblood *et al.*, "Observations of GRB 230307A by TESS," *Research Notes of the AAS*, vol. 7, no. 3, p. 56, 2023.
- [36] N. Gehrels, E. Ramirez-Ruiz, and D. B. Fox, "The swift gamma-ray burst mission," *Annual Review of Astronomy and Astrophysics*, vol. 47, pp. 567–617, 2009.
- [37] S. Ghosh, S. Bloemen, G. Nelemans, P. J. Groot, and L. R. Price, "Tiling strategies for optical follow-up of gravitational-wave triggers by telescopes with a wide field of view," *Astronomy & Astrophysics*, vol. 592, p. A82, Aug. 2016. [Online]. Available: <http://dx.doi.org/10.1051/0004-6361/201527712>
- [38] B. L. Golden, L. Levy, and R. Vohra, "The orienteering problem," *Naval Research Logistics (NRL)*, vol. 34, no. 3, pp. 307–318, 1987.
- [39] K. M. Gorski *et al.*, "HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere," *ApJ*, vol. 622, no. 2, p. 759, 2005.
- [40] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics: A Foundation for Computer Science*, 2nd ed. Reading, Massachusetts: Addison-Wesley, 1994.
- [41] C. Gueymard *et al.*, *SMARTS2: a simple model of the atmospheric radiative transfer of sunshine: algorithms and performance assessment*. Florida Solar Energy Center Cocoa, FL, USA, 1995, vol. 1.
- [42] A. Gunawan, H. C. Lau, and P. Vansteenwegen, "Orienteering problem: A survey of recent variants, solution approaches and applications," *European Journal of Operational Research*, vol. 255, no. 2, pp. 315–332, 2016.
- [43] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2024. [Online]. Available: <https://www.gurobi.com>
- [44] J. Hoogeveen, "Analysis of Christofides' heuristic: Some paths are more difficult than cycles," *Operations Research Letters*, vol. 10, no. 5, pp. 291–295, 1991.
- [45] L. Hu, L. Wang, X. Chen, and J. Yang, "Image subtraction in Fourier space," *The Astrophysical Journal*, vol. 936, no. 2, p. 157, 2022.
- [46] M. Joseph and P. Pandya, "Finding response times in a real-time system," *The Computer Journal*, vol. 29, no. 5, pp. 390–395, 1986.
- [47] F. Kasten and A. T. Young, "Revised optical air mass tables and approximation formula," *Appl. Opt.*, vol. 28, no. 22, pp. 4735–4738, Nov 1989. [Online]. Available: <https://opg.optica.org/ao/abstract.cfm?URI=ao-28-22-4735>
- [48] J. H. Kim, M. Im, H. M. Lee, S.-W. Chang, H. Choi, and G. S. H. Paek, "Introduction to the 7-Dimensional Telescope: Commissioning procedures and data characteristics," 2024. [Online]. Available: <https://arxiv.org/abs/2406.16462>
- [49] D. Lang, D. W. Hogg, K. Mierle, M. Blanton, and S. Roweis, "Astrometry.net: Blind astrometric calibration of arbitrary astronomical images," *The astronomical journal*, vol. 139, no. 5, p. 1782, 2010.
- [50] Q. Li, S. Wang, C. Xu, X. Ma, M. Xu, A. Zhou, R. Xing, B. Yang, Z. Zhu, Y. Zhang *et al.*, "Exploring real-time satellite computing: From energy and thermal perspectives," in *2024 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2024, pp. 161–173.
- [51] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of the ACM (JACM)*, vol. 20, no. 1, pp. 46–61, 1973.
- [52] A. Loveless, R. Dreslinski, B. Kasicki, and L. T. X. Phan, "Igor: Accelerating byzantine fault tolerance for real-time systems with eager execution," in *2021 IEEE 27th Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 2021, pp. 360–373.
- [53] lsst, "lsst/throughputs: Lsst simulations repository for baseline evaluation information," <https://github.com/lsst/throughputs/>, accessed: 2025-08-08.
- [54] M. Lv, X. Peng, W. Xie, and N. Guan, "Task allocation for real-time earth observation service with LEO satellites," in *2022 IEEE Real-Time Systems Symposium (RTSS)*, 2022, pp. 14–26.
- [55] F. Marković, P. Roux, S. Bozhko, A. V. Papadopoulos, and B. B. Brandenburg, "CTA: A correlation-tolerant analysis of the deadline-failure probability of dependent tasks," in *2023 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2023, pp. 317–330.
- [56] F. Marković, G. von der Brüggen, M. Günzel, J.-J. Chen, and B. B. Brandenburg, "A distribution-agnostic and correlation-aware analysis of periodic tasks," in *2024 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2024, pp. 215–228.
- [57] I. Martinez *et al.*, "The cosipy library: COSI's high-level analysis software," in *Proc. of 38th Int'l Cosmic Ray Conf.*, vol. 444, 2023, pp. 858:1–858:8.
- [58] P. Mészáros, D. B. Fox, C. Hanna, and K. Murase, "Multi-messenger astrophysics," *Nature Reviews Physics*, vol. 1, no. 10, pp. 585–599, Oct 2019. [Online]. Available: <https://doi.org/10.1038/s42254-019-0101-z>
- [59] National Academies of Sciences Engineering and Medicine, *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*. Washington, DC, USA: The National Academies Press, 2023. [Online]. Available: <https://nap.nationalacademies.org/catalog/26141/pathways-to-discovery-in-astronomy-and-astrophysics-for-the-2020s>
- [60] G. Oganesyan, L. Nava, G. Ghirlanda, A. Melandri, and A. Celotti, "Prompt optical emission as a signature of synchrotron radiation in gamma-ray bursts," *Astronomy & Astrophysics*, vol. 628, p. A59, 2019.
- [61] K. I. Park, M. Park *et al.*, *Fundamentals of probability and stochastic processes with applications to communications*. Springer, 2018.
- [62] A. Paul, D. Freund, A. Ferber, D. B. Shmoys, and D. P. Williamson, "Budgeted prize-collecting traveling salesman and

- minimum spanning tree problems,” *Mathematics of Operations Research*, vol. 45, no. 2, pp. 576–590, 2020.
- [63] —, “Erratum to “budgeted prize-collecting traveling salesman and minimum spanning tree problems”,” *Mathematics of Operations Research*, vol. 48, no. 4, pp. 2304–2307, 2023.
- [64] “Photometric system — filters used,” Wikipedia, https://en.wikipedia.org/wiki/Photometric_system#Filters_used, February 2025, last edited on 27 February 2025; retrieved 8 August 2025.
- [65] B. Preiss, T. Pan, and K. Ramohalli, “Development of a figure-of-merit for space missions,” *NASA Space Engineering Research Center for Utilization of Local Planetary Resources*, 1991.
- [66] J. Rana, A. Singhal, B. Gadre, V. Bhalerao, and S. Bose, “An enhanced method for scheduling observations of large sky error regions for finding optical counterparts to transients,” *The Astrophysical Journal*, vol. 838, no. 2, p. 108, 2017.
- [67] J. Rana, S. Anand, and S. Bose, “Optimal search strategy for finding transients in large-sky error regions under realistic constraints,” *The Astrophysical Journal*, vol. 876, no. 2, p. 104, 2019.
- [68] M. W. Richmond, M. Tanaka, T. Morokuma, S. Sako, R. Ohsawa, N. Arima, N. Tominaga, M. Doi, T. Aoki, K. Arimatsu *et al.*, “An optical search for transients lasting a few seconds,” *Publications of the Astronomical Society of Japan*, vol. 72, no. 1, p. 3, 2020.
- [69] S. Saller, J. Koehler, and A. Karrenbauer, “A systematic review of approximability results for traveling salesman problems leveraging the TSP-T3CO definition scheme,” *arXiv preprint arXiv:2311.00604*, 2023.
- [70] M. W. Savelsbergh, “Local search in routing problems with time windows,” *Annals of Operations research*, vol. 4, pp. 285–305, 1985.
- [71] L. P. Singer and L. R. Price, “Rapid bayesian position reconstruction for gravitational-wave transients,” *Phys. Rev. D*, vol. 93, p. 024013, Jan 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.93.024013>
- [72] M. Sirianni, M. J. Jee, N. Benítez, J. P. Blakeslee, A. R. Martel, G. Meurer, M. Clampin, G. De Marchi, H. C. Ford, R. Gilliland, G. F. Hartig, G. D. Illingworth, J. Mack, and W. J. McCann, “The photometric performance and calibration of the hubble space telescope advanced camera for surveys,” *Publications of the Astronomical Society of the Pacific*, vol. 117, no. 836, p. 1049, sep 2005. [Online]. Available: <https://dx.doi.org/10.1086/444553>
- [73] M. Sudvarg *et al.*, “A Fast GRB Source Localization Pipeline for the Advanced Particle-physics Telescope,” in *Proc. of 37th Int’l Cosmic Ray Conf.*, vol. 395, Jul. 2021, pp. 588:1–588:9.
- [74] —, “Parameterized workload adaptation for fork-join tasks with dynamic workloads and deadlines,” in *Proc. of 29th Int’l Conf. on Embedded and Real-Time Computing Systems and Applications*, 2023, pp. 232–242.
- [75] S. Tingay and W. Joubert, “High cadence optical transient searches using drift scan imaging ii: Event rate upper limits on optical transients of duration < 21 ms and magnitude < 6.6 ,” *Publications of the Astronomical Society of Australia*, vol. 38, p. e001, 2021.
- [76] J. A. Tomsick, S. E. Boggs, A. Zoglauer, D. Hartmann *et al.*, “The Compton Spectrometer and Imager,” 2023. [Online]. Available: <https://arxiv.org/abs/2308.12362>
- [77] P. Vansteenwegen and A. Gunawan, “Orienteering problems,” *EURO advanced tutorials on operational research*, vol. 1, 2019.
- [78] P. Vansteenwegen, W. Souffriau, and D. Van Oudheusden, “Orienteering problem variants: Time windows, service times, and multiple visits,” in *Orienteering Problems: Models and Algorithms for Vehicle Routing Problems with Profits*. Berlin, Heidelberg: Springer, 2011, pp. 23–25.
- [79] —, “The orienteering problem: A survey,” *European Journal of Operational Research*, vol. 209, no. 1, pp. 1–10, 2011.
- [80] V. A. Villar, E. Berger, B. D. Metzger, and J. Guillochon, “Theoretical models of optical transients. i. a broad exploration of the duration–luminosity phase space,” *The Astrophysical Journal*, vol. 849, no. 1, p. 70, 2017.
- [81] E. Waxman and B. Katz, *Shock Breakout Theory*. Cham: Springer International Publishing, 2017, pp. 967–1015. [Online]. Available: https://doi.org/10.1007/978-3-319-21846-5_33
- [82] D. G. York, J. Adelman, J. E. Anderson Jr, S. F. Anderson, J. Annis, N. A. Bahcall, J. Bakken, R. Barkhouser, S. Bastian, E. Berman *et al.*, “The Sloan Digital Sky Survey: Technical summary,” *The Astronomical Journal*, vol. 120, no. 3, p. 1579, 2000.