Paving the Way for More Accessible Cancer Care in Low-Income Countries with Optimization

Abel Sapirstein*, Lauren N. Steimle, and Mathieu Dahan

H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA.

*Corresponding author(s). E-mail(s): asapirstein3@gatech.edu;
Contributing authors: steimle@gatech.edu; mdahan3@gatech.edu;

Abstract

Cancers are a growing cause of morbidity and mortality in low-income countries. Geographic access plays a key role in both timely diagnosis and successful treatment. In areas lacking well-developed road networks, seasonal weather events can lengthen already long travel times to access care. Expanding facilities to offer cancer care is expensive and requires staffing by skilled medical professionals, which are often in short supply. In this article, we propose a mathematical model to improve geographic access to cancer care by jointly considering expansions to care facilities and improvements to the road network. We model this as a multi-period stochastic facility location network design problem. In each period, a decision maker must simultaneously choose a set of facilities at which to add tertiary cancer services and a set of roads to improve while facing demand and travel time uncertainty. Once demand for cancer care and weather events are realized, patients observe road conditions and use the transportation network to travel towards the closed facility with available cancer services.

We create a new path-based formulation of this problem and develop a new branch-price-and-cut algorithm with acceleration techniques that take advantage of this formulation's structure. We demonstrate our approach using Rwanda as a case study and show that the reductions in travel time to cancer care that are directly attributable to the road network improvements can be as high as 1 hour.

Keywords: Healthcare access, facility location, network design, cancer care, geospatial analysis, low-income countries

Highlights

- Tertiary cancer care is expensive and often not widely available.
- Accessing this care is particularly challenging in regions with poor road conditions, especially during rainy seasons.
- We develop a multi-period stochastic optimization model that plans both the expansion of cancer care and improvements to the road network to improve access.
- We design a tailored branch-price-and-cut algorithm with custom acceleration strategies to solve large-scale real-world instances of the model.
- In a case study of Rwanda, we find that jointly improving roads and expanding care provides the biggest benefits for patients currently farthest from care.

1 Introduction

1.1 Motivation

Cancer is a leading cause of morbidity and mortality worldwide, but the impacts of cancer are not felt equally. It is estimated that 9.6 million people die from cancer each year, and that 70% of these deaths occur in low- and middle-income countries [1]. The majority of new cancer cases occur in these countries and have higher mortality than in high-income countries [2]. For instance, 6% of cancer cases globally are estimated to occur in sub-Saharan Africa, but this region accounts for 15% of cancer-attributable deaths [2]. Compared to high-income countries, there is a higher proportion of cancer patients in low-income countries (LICs) who have already progressed to an advanced stage of the disease by the time they are diagnosed [3], and consequently there is a higher case fatality rate [4]. Accordingly, early diagnosis and access to treatment play an important role in the cancer control strategy for many LICs [5].

Geographic accessibility plays a key role in early diagnosis and treatment in LICs. It is well-documented that patients with worse geographic access receive care at later stages of disease [2, 6]. However, the uneven distribution of cancer care facilities and the quality of the road network can result in extremely long travel times to access care, especially for those living in rural areas. For instance, more than 25% of Rwandans and Nigerians live more than 4 hours from cancer care [6–8]. Moreover, people living in rural regions are more likely to experience disruptions in access due to extreme weather events [9]. Often, roads in rural areas tend to be gravel or dirt, which makes them particularly susceptible to being washed out by heavy rainfalls [10]. Meanwhile, the available data suggest that rural regions in sub-Saharan Africa have higher rates of cancer-attributable mortality than their urban counterparts [6, 11]. Altogether, rural populations in LICs are especially impacted by a lack of geographic access to cancer care.

Given the importance of geographic access on early diagnosis and access to treatment, identifying solutions to improve access to cancer care is a crucial step in improving health outcomes and mitigating health disparities on national and global scales. Indeed, many ministries of health and non-profit organizations are looking to make investments to expand access to care [7, 12–14]. However, it remains unclear how service and infrastructure improvements can best improve access to care, particularly for rural populations when roads are severely impacted by weather events.

This lead us to the following research question: *How should health system planners allocate resources to jointly improve cancer care services and road infrastructure for maximizing access care under demand and travel time uncertainty?* We hypothesized that enhancing the roads on which patients travel in concert with expanding services could significantly improve access to cancer care, especially for currently under-served patients.

1.2 Contributions

In this article, we describe a new approach for planning the expansion of healthcare and transportation infrastructure to improve geographic accessibility of cancer care.

We present a new Facility Location Network Design Problem (FLNDP) in which a central planner can make improvements to existing healthcare facilities and to the existing road network over a multi-period planning horizon by allocating monetary and staffing budgets. During each period, some uncertainty is realized that impacts the demand for healthcare services from population centers and the travel time across the road network.

To solve this model, we formulate a new path-based mixed-integer program (MIP). Our formulation has an exponentially large number of decision variables, which represent the possible paths from population centers to facilities in the network. To handle the scale of the problem, we first derive a branch-and-price algorithm. Leveraging the structure of the problem, we are able to efficiently solve the pricing problems using shortest-path subroutines. We further enhance the scalability of the algorithm by deriving acceleration techniques based on tailored heuristics, branching rules, and valid inequalities. Our final proposed algorithm is an enhanced branch-price-and-cut algorithm. We test the algorithm on synthetic instances and observe that our acceleration techniques increase the rate at which optimal solutions are found on small instances and the likelihood that an optimal solution is found for larger instances.

Finally, we apply our method to explore potential ways to expand access to tertiary cancer care in Rwanda over a five-year planning horizon. To do so, we create a high-fidelity network representation of the road network in Rwanda using Open Street Maps [15] and gather data on cancer care services from the Rwanda Ministry of Health and Partners in Health, a non-governmental organization. We also consider the impacts of oncologist scarcity and weather events in our planning. In a representative scenario, our results reveal that cancer care is, on average, 30 minutes closer when investments in both healthcare and transportation infrastructure are considered, compared to improvements in healthcare infrastructure alone. Moreover, our findings suggest that the recommendations from our FLNDP model could result in a 4-hour reduction in travel time to care for the patients who currently must travel the longest to access care, with 1 hour of this reduction being directly attributable to road network improvements. Our results demonstrate the value of jointly considering both network improvements and facility expansions, providing managerial insights to healthcare planners seeking to improve care access in LICs.

1.3 Background and Related Literature

With the growing interest in understanding and addressing barriers to geographic access to care, modeling is starting to emerge as a powerful tool to help inform policy planning in LICs. While there has been a dearth of studies examining geographic access in LICs historically, recently there has been an emergence of the use of Geographical Information Systems (GIS) modeling to understand current levels of geographic access and how plans to expand cancer care would change geographic access in Rwanda [7], Ghana [16], and Nigeria [6].

While GIS modeling can be helpful to examine current levels of access and evaluate different expansion policies, these methods are limited in that they can only evaluate pre-defined plans, and there is no guarantee that these plans would make the best use of available resources. *Facility location modeling* extends these GIS approaches by

using mathematical optimization to select the best facilities to improve a pre-specified measure of geographic access given limited resources. Despite its utility, very few studies have used facility location modeling for planning the expansion of cancer care services in LICs. Among those few studies are Sapirstein et al. [8] who examine the expansion of tertiary cancer care in Rwanda and de Campos et al. [17] who study the optimal location of mammography units in Brazil [17]. However, these existing studies do not consider improvements to the road network as a means to improve access to cancer care.

The operations research community has developed models and methods devoted to solving problems in which a planner must simultaneously choose where to locate facilities that can provide a service of interest and a transportation network that can be used to route demand to the selected location. Daskin et al. [18] provided the first formulation and coined such problems as FLNDPs. Shortly thereafter, Melkote and Daskin [19] demonstrated that an FLNDP can be reduced to a pure network design problem. Since then, considerable progress has been made on the ability to solve FLNDPs [20–23], which has enabled their use in various application areas, including healthcare.

Our work adds to a relatively recent stream of literature focusing on using FLNDPs to improve healthcare access in LICs. To the best of our knowledge, Cocking et al. [24] were the first to propose an FLNDP approach to improve access to healthcare through both facility expansions and road network improvements. They explore a single-period planning problem in which a decision-maker can build healthcare facilities and new roads to improve access to healthcare. They demonstrate their approach on an instance of an FLNDP representing Burkina Faso's Nouna District, which they can solve using a commercial solver in a matter of hours. Ghaderi and Jabalameli [25] use an FLNDP approach to simultaneously design the capacity of healthcare facilities and the transfer network among these healthcare facilities over multiple time periods. The authors propose a fix-and-optimize approach to heuristically solve large instances of their problem, such as their case study of Iran's Ilam Province. In contrast to our proposed model, Ghaderi and Jabalameli [25]'s model considers distinct monetary budgets for the construction of facilities and the construction of the transfer network. Shishebori et al. [26] also use an FLNDP model with distinct network and facility budgets, adopting a heuristic approach and considering health-system capacity constraints. More recently, Pourrezaie-Khaligh et al. [27] proposed an improvement on the heuristic techniques of Ghaderi [22] while explicitly incorporating equity in their optimization model. While these studies have demonstrated the benefits of FLNDP as a means to improve healthcare access, the existing literature in this area has been limited to the solution of relatively small instances with a commercial solver or the solution of more realistic large-scale instances using heuristic approaches that do not provide any guarantees on the optimality gap. Therefore, there is a need for new optimization methods that can solve larger-scale instances to optimality or provide guarantees on the optimality gap. Our solution approach based on branch-price-and-cut seeks to fill this gap.

The methods proposed in this article build upon previous efforts to solve large-scale FLNDPs. More specifically, our work builds on two streams of research in large-scale

network design problems related to "edge-based" formulations for FLNDPs and "path-based" formulations for pure network design problems. Edge-based formulations use binary variables to represent the decisions of whether or not to construct an edge, continuous variables to represent the amount of demand from a demand point flowing along a given edge, and constraints to ensure that demand flows from a demand point along only constructed edges to reach a facility. Pearce and Forbes [23] were one of the first to propose an exact method for solving large-scale instances of edge-based formulations of FLNDPs. They use a Benders decomposition approach, adapted from a similar approach originally for pure network design problems [28], in which they fix the binary decision variables corresponding to facility and edge construction in a main problem. Then, they solve a subproblem that can be disaggregated by demand point and time period. Similar methods have been used in a variety of FLNDP contexts including hazardous waste transport [29], urban planning [30], and traffic flow [21]. Such edge-based formulations offer nice benefits because they lend themselves to decomposition techniques, such as Benders decomposition, which can improve solution times on large-scale instances. However, these formulations oftentimes struggle to converge on large instances: They still require solving many (potentially challenging) subproblems, leading to potentially substantial optimality gaps on large instances.

Path-based formulations are becoming popular in other network design problems. In these formulations, a path typically corresponds to a set of edges along which demand flows from the originating demand point of the path to a facility that satisfies this demand, which is the end point of the path. These formulations typically require fewer constraints than edge-based formulations, but this comes at the cost of requiring an enormous number of decision variables. Thus, these path-based formulations tend to be mixed-integer programs with a huge number of decision variables. Accordingly, existing studies have leveraged branch-and-price [31] as a way to solve these formulations of pure network design problems [32, 33]. As a result, path-based formulations have been effective for generating solutions to many complex logistics, routing, and scheduling problems [34, 35]. To our knowledge, the only work on path-based formulations for FLNDPs has been limited to Contreras et al. [36]. In their deterministic single-period setting, they use a min-max objective function to minimize the maximum travel time across their network and demonstrate that a path-based formulation results in improved solution times compared to edge-based formulations. However, their approach relies on the generation of valid inequalities that result from the need to only consider the longest path. Thus, these valid inequalities do not generalize to the objective function considered in our modeling framework in which we consider population-average travel time, accounting for multiple improvement decision periods as well as demand and travel time uncertainty.

Despite the use of path-based approaches to solve a variety of pure network design problems and the strong connections between FLNDPs and pure network design problems, to our knowledge, the use of path-based formulations for solving FLNDPs has been extremely limited. In this article, we seek to create a solution methodology for our proposed FLNDP that combines the structure of path-based formulations with decomposition techniques more frequently seen in edge-based formations.

1.4 Organization of this Paper

The remainder of this article is organized as follows. In §2, we define our problem and formulate a two-stage stochastic MIP that can be used to solve it. In §3, we propose a new branch-price-and-cut (BP&C) framework for solving the previously formulated MIP. In §4, we show the efficacy of our approach via computational experiments. In §5, we present our motivating case study and discuss the important findings. Finally, in §6, we conclude and identify future avenues of research.

2 Problem Description

2.1 Problem Statement

We consider a health system planner seeking to improve geographic access to tertiary cancer care for the regional population. Given a budget allocated over multiple time periods, the planner aims to invest in both expanding cancer services in existing facilities and improving road infrastructure. Considering uncertainty in the demand for cancer services and the impact of weather on road travel times, the planner's objective is to minimize the expected travel time for the population to access cancer care.

Formally, we consider a set $\mathcal{O}$ of population centers with patients in need of cancer care. We denote by $\mathcal{F}$ the set of candidate healthcare facilities that can provide such care with some additional investment from the health system planner. Population centers and facilities are connected by a road network, which the planner seeks to improve. Road surfaces (e.g., dirt, gravel, paved) can be categorized into an ordered set $\mathcal{S} = \{1, \ldots, S\}$ of S types, such that surface type $s \in \mathcal{S}$ is of lower quality than surface type $s' \in \mathcal{S}$ when $s < s'$. Next, we define as $\mathcal{G} \coloneqq (\mathcal{N}, \mathcal{A})$ the undirected multigraph representing the potential improved road networks achievable after investment. Its set of nodes $\mathcal{N}$ comprises population centers, facilities, and road junctions, which are connected by a set of type-dependent arcs $\mathcal{A}$. Each arc $a = (i, j, s) \in \mathcal{A}$ reflects a potential road between nodes $i \in \mathcal{N}$ and $j \in \mathcal{N}$ with a surface of type $s \in \mathcal{S}$. Without loss of generality, we assume that each arc is associated with a single surface type; otherwise, a road can be equivalently represented by multiple smaller arcs of potentially different types.

Over the course of a planning horizon, discretized into a set $\mathcal{T} \coloneqq \{1, \ldots, T\}$ of T (monthly/yearly) decision time periods, the planner can invest in the addition of cancer services to candidate healthcare facilities and in the improvement of the road network. Adding cancer services to a facility $f \in \mathcal{F}$ at the start of time period $\tau \in \mathcal{T}$ incurs both a one-time monetary cost c_f^τ and a human capital cost h_f^τ, while improving the road between nodes i and j from surface type $s-1$ to type s incurs a one-time monetary cost c_{ijs}^τ. We assume that road improvement costs are additive, namely, the cost of improving the road between i and j from surface type s_1 to type $s_2 > s_1$ at the beginning of time period τ is given by $\sum_{s=s_1+1}^{s_2} c_{ijs}^\tau$. For the sake of exposition, we neglect operations and maintenance costs, and we assume cancer services and road surfaces invested in at the beginning of a time period are available for the population during that time period. In each time period $\tau \in \mathcal{T}$, the planner receives a monetary

budget C^τ and a staffing budget H^τ. Any leftover budget that is not used by the end of a time period can be rolled over for use in the subsequent time period.

After the planner has made investments in a particular time period, the patients seeking cancer care travel across the road network to reach facilities offering cancer services. However the planner faces uncertainty regarding the demand for such services. Furthermore, the condition of some roads (e.g., gravel or dirt) may significantly worsen during inclement weather (e.g., extreme rainfall), negatively impacting patients' travel times. Thus, we consider a set Ω^τ of stochastic scenarios for each time period $\tau \in \mathcal{T}$. Each scenario $\omega \in \Omega^\tau$ occurs with a given probability $p^{\tau\omega}$ and is realized independently of scenarios that are realized in other time periods. Each scenario ω^τ is characterized by the demand $d_o^{\tau\omega}$ for cancer services from each population center $o \in \mathcal{O}$, as well as the travel time $\ell_a^{\tau\omega}$ across each type-dependent arc $a = (i, j, s) \in \mathcal{A}$. We naturally assume that in every scenario $\omega \in \Omega^\tau$ occurring at a time period $\tau \in \mathcal{T}$, arcs connecting two nodes i and j in the graph $\mathcal{G}$ satisfy $\ell_{ijs_1}^{\tau\omega} > \ell_{ijs_2}^{\tau\omega}$ for every $s_1 < s_2 \in \mathcal{S}$, i.e., the travel time to traverse a road of a given length will be higher if it is a lower-quality surface type than if it is a higher-quality surface type. We note that Ω^τ can also account for additional factors such as potential armed conflicts or the effectiveness of cancer screening programs that can influence the likelihood of patients seeking care.

During each time period and each scenario, we assume that patients can accurately estimate the road network's travel times, and select a path (i.e., route) that minimizes their travel time to a facility that provides cancer services. We define a path $\lambda \in \Lambda$ as a subset of adjacent type-dependent arcs in $\mathcal{A}$ that connect a population center $o_\lambda \in \mathcal{O}$ to a facility $f_\lambda \in \mathcal{F}$. Under the realized scenario $\omega \in \Omega^\tau$ during a time period $\tau \in \mathcal{T}$, the total travel time along each path $\lambda \in \Lambda$ is given by $\ell_\lambda^{\tau\omega} \coloneqq \sum_{a\in\lambda} \ell_a^{\tau\omega}$.

We then model the health system planner's decisions as a two-stage stochastic multi-period facility location network design problem with uncapacitated facilities and uncapacitated arcs. In the first stage, the planner selects for each time period the subset of facilities to invest in the addition of cancer services and the subset of type-dependent arcs to construct, given the monetary and staffing budgets. Then, in the second stage, for a given realized scenario $\omega \in \Omega^\tau$ during each time period $\tau \in \mathcal{T}$, the patients from each population center $o \in \mathcal{O}$ select a shortest path $\hat{\lambda}_o^{\tau\omega}$ that crosses arcs that are constructed by time period τ to seek care in a facility that received investments by that same time period. Thus, the planner's objective is to minimize the population's expected travel time, given by

$$\sum_{\tau\in\mathcal{T}} \sum_{\omega\in\Omega^\tau} p^{\tau\omega} \left(\sum_{o\in\mathcal{O}} d_o^{\tau\omega} \ell_{\hat{\lambda}_o^{\tau\omega}}^{\tau\omega} \right).$$

We illustrate the decision and uncertainty timeline in Figure 1 and summarize the model notations in Table 1. Note that we interchangeably use the subscripts ijs and a when expressing parameters and variables that depend on a type-dependent arc $a = (i, j, s) \in \mathcal{A}$.

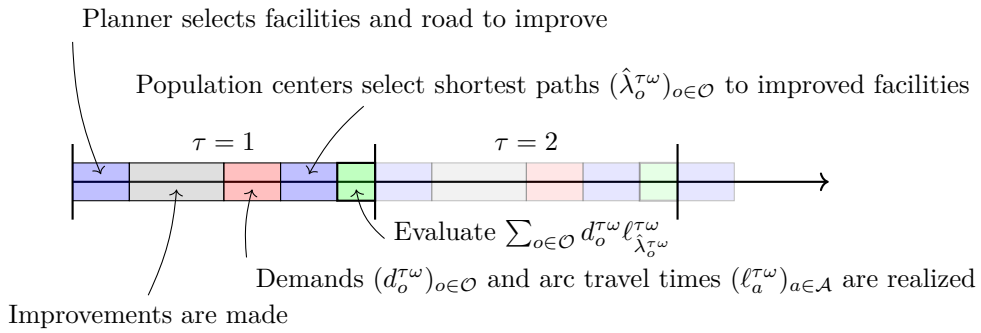

Fig. 1: Decision and uncertainty timeline of the two-stage stochastic multi-period facility location network design problem.

Table 1: Notation used to describe the facility location network design problem.

Symbol	Description
$\mathcal{N}$	Set of network nodes
$\mathcal{F}$	Set of facility locations, $\mathcal{F} \subseteq \mathcal{N}$
$\mathcal{O}$	Set of population centers $\mathcal{O} \subseteq \mathcal{N}$
$\mathcal{S}$	Set of road surface types, $\mathcal{S} = \{1, \ldots, S\}$
$\mathcal{A}$	Set of type-dependent arcs
$\mathcal{G}$	Graph of potential improved road networks, $\mathcal{G} = (\mathcal{N}, \mathcal{A})$
Λ	Set of paths in graph $\mathcal{G}$ connecting a population center to a facility
$\mathcal{T}$	Set of time periods, $\mathcal{T} = \{1, \ldots, T\}$
Ω^τ	Set of stochastic scenarios in time period $\tau \in \mathcal{T}$
c^τ_{ijs}	Monetary cost to construct arc $(i, j, s) \in \mathcal{A}$ at the start of period $\tau \in \mathcal{T}$
c^τ_f	Monetary cost of adding services to facility $f \in \mathcal{F}$ at the start of period $\tau \in \mathcal{T}$
h^τ_f	Human capital cost of adding services to $f \in \mathcal{F}$ at the start of period $\tau \in \mathcal{T}$
C^τ	Monetary budget received at the start of period $\tau \in \mathcal{T}$
H^τ	Staffing budget received at the start of period $\tau \in \mathcal{T}$
o_λ	Population center from where path $\lambda \in \Lambda$ originates
f_λ	Facility at which path $\lambda \in \Lambda$ terminates
$p^{\tau\omega}$	Probability of scenario $\omega \in \Omega^\tau$ occurring during period $\tau \in \mathcal{T}$
$d_o^{\tau\omega}$	Demand at population center $o \in \mathcal{O}$ during period $\tau \in \mathcal{T}$ under scenario $\omega \in \Omega^\tau$
$\ell_a^{\tau\omega}$	Travel time along arc $a \in \mathcal{A}$ during period $\tau \in \mathcal{T}$ under scenario $\omega \in \Omega^\tau$
$\ell_\lambda^{\tau\omega}$	Travel time along path $\lambda \in \Lambda$ during period $\tau \in \mathcal{T}$ under scenario $\omega \in \Omega^\tau$

2.2 Model Formulation

We formulate the stochastic facility location network design problem described above as a mixed-integer program (MIP). The planner's first-stage decisions can be described using four sets of binary decision variables: For each time period $\tau \in \mathcal{T}$, we use u^τ_f to indicate if the planner adds cancer services to facility $f \in \mathcal{F}$ at the start of time period τ. We also use x^τ_f to indicate if services are available at facility f at the start of time

period τ, including the investment decision taken at the start of that time period. Similarly, v_a^τ (resp. y_a^τ) determines if arc $a \in \mathcal{A}$ is constructed (resp. is available) at the start of time period τ. For each stochastic scenario $\omega \in \Omega^\tau$ that is realized during a time period $\tau \in \mathcal{T}$, we define for every path $\lambda \in \Lambda$ the second-stage continuous variable $z_\lambda^{\tau\omega} \in [0,1]$ that represents the fraction of demand from population center o_λ (i.e., the origin of the path) traversing along λ to reach facility f_λ (i.e., the end of the path).

We formulate the problem in (1) and summarize the decision variables in Table 2.

$$\underset{u,\,v,\,x,\,y,\,z}{\text{minimize}} \quad \sum_{\tau\in\mathcal{T}} \sum_{\omega\in\Omega^\tau} p^{\tau\omega} \sum_{\lambda\in\Lambda} d_{o_\lambda}^{\tau\omega} \ell_\lambda^{\tau\omega} z_\lambda^{\tau\omega} \tag{1a}$$

$$\text{subject to} \quad \sum_{\tau=1}^{\tau'} \left(\sum_{f\in\mathcal{F}} c_f^\tau u_f^\tau + \sum_{a\in\mathcal{A}} c_a^\tau v_a^\tau \right) \le \sum_{\tau=1}^{\tau'} C^\tau, \quad \forall\, \tau' \in \mathcal{T}, \tag{1b}$$

$$\sum_{\tau=1}^{\tau'} \sum_{f\in\mathcal{F}} h_f^\tau u_f^\tau \le \sum_{\tau=1}^{\tau'} H^\tau, \quad \forall\, \tau' \in \mathcal{T}, \tag{1c}$$

$$x_f^\tau - x_f^{\tau-1} = u_f^\tau, \quad \forall\, f \in \mathcal{F},\ \forall\, \tau \in \{2,\dots,T\}, \tag{1d}$$

$$y_a^\tau - y_a^{\tau-1} = v_a^\tau, \quad \forall\, a \in \mathcal{A},\ \forall\, \tau \in \{2,\dots,T\}, \tag{1e}$$

$$x_f^1 = u_f^1, \quad \forall\, f \in \mathcal{F}, \tag{1f}$$

$$y_a^1 = v_a^1, \quad \forall\, a \in \mathcal{A}, \tag{1g}$$

$$y_{ijs}^\tau \le y_{ij(s-1)}^\tau, \quad \forall\, (i,j,s) \in \mathcal{A} \text{ s.t. } s \ge 2,\ \forall\, \tau \in \mathcal{T}, \tag{1h}$$

$$\sum_{\{\lambda\in\Lambda\,:\,o_\lambda = o\}} z_\lambda^{\tau\omega} = 1, \quad \forall\, o \in \mathcal{O},\ \forall\, \omega \in \Omega^\tau,\ \forall\, \tau \in \mathcal{T}, \tag{1i}$$

$$\sum_{\{\lambda\in\Lambda\,:\,o_\lambda = o,\, a\in\lambda\}} z_\lambda^{\tau\omega} \le y_a^\tau, \quad \forall\, a \in \mathcal{A},\ \forall\, o \in \mathcal{O},\ \forall\, \omega \in \Omega^\tau,\ \forall\, \tau \in \mathcal{T}, \tag{1j}$$

$$\sum_{\{\lambda\in\Lambda\,:\,o_\lambda = o,\, f_\lambda = f\}} z_\lambda^{\tau\omega} \le x_f^\tau, \quad \forall\, f \in \mathcal{F},\ \forall\, o \in \mathcal{O},\ \forall\, \omega \in \Omega^\tau,\ \forall\, \tau \in \mathcal{T}, \tag{1k}$$

$$z_\lambda^{\tau\omega} \ge 0, \quad \forall\, \lambda \in \Lambda,\ \forall\, \omega \in \Omega^\tau,\ \forall\, \tau \in \mathcal{T}, \tag{1l}$$

$$x_f^\tau,\ u_f^\tau \in \{0,1\}, \quad \forall\, f \in \mathcal{F},\ \forall\, \tau \in \mathcal{T}, \tag{1m}$$

$$y_a^\tau,\ v_a^\tau \in \{0,1\}, \quad \forall\, a \in \mathcal{A},\ \forall\, \tau \in \mathcal{T}. \tag{1n}$$

The objective function (1a) computes the cumulative expected travel time over the entire planning horizon. Constraints (1b)-(1c) ensure that at each time period, the planner does not exceed the cumulative monetary and staffing budgets received by then. The linking constraints (1d)-(1g) guarantee that the variables x and y respectively represent the availability of facilities and arcs at the start of each time period, as a result of the investment decision variables u and v. By enforcing that an arc of type s can only be constructed if lesser arc types have also been constructed, constraints (1h) model the additive nature of road improvement costs. Constraints (1i)

Table 2: Decision variables used in the MIP formulation of the two-stage stochastic multi-period facility location network design problem.

Variable	Description
$u_f^\tau \in \{0,1\}$	1 if services are added to facility $f \in \mathcal{F}$ at the start of period $\tau \in \mathcal{T}$
$v_a^\tau \in \{0,1\}$	1 if arc $a \in \mathcal{A}$ is constructed at the start of period $\tau \in \mathcal{T}$
$x_f^\tau \in \{0,1\}$	1 if services are available at facility $f \in \mathcal{F}$ at the start of period $\tau \in \mathcal{T}$
$y_a^\tau \in \{0,1\}$	1 if arc $a \in \mathcal{A}$ is available at the start of period $\tau \in \mathcal{T}$
$z_\lambda^{\tau\omega} \in [0,1]$	Fraction of demand from population center $o_\lambda \in \mathcal{O}$ using path $\lambda \in \Lambda$ in scenario $\omega \in \Omega^\tau$ during time period $\tau \in \mathcal{T}$

account for the fact that all patients seeking care in each scenario travel to a facility. Constraints (1j)-(1k) ensure that in each time period and each possible scenario, any path selected by a patient only traverses available type-dependent arcs to reach a facility with available services. We note that constraints (1d)-(1g) can be adjusted to account for construction lead times. Furthermore, monetary and human capital costs can be set to zero for any facility that already provides services. Similarly, if an arc $(i,j,s) \in \mathcal{A}$ is already constructed, then we set $c_{ijs'}^\tau = 0$ for every time period τ and every type $s' \in \{1,\ldots,s\}$. Finally, we observe that at optimality of the MIP (1), every patient will select in each scenario a shortest path that traverses improved roads to reach a facility with available services.

The resulting formulation requires $2T(|\mathcal{F}| + |\mathcal{A}|)$ binary first-stage decision variables and $T|\Omega||\Lambda|$ continuous second-stage variables. Since the number of paths is exponential, (1) cannot directly be solved using commercial solvers and requires new methodologies based on decomposition algorithms.

3 Solution Methodology

3.1 Algorithm Overview

In this section, we first develop an algorithm based on branch-and-price for solving the large-scale MIP (1). This tailored algorithm iteratively solves the linear programming (LP) relaxation of (1) using column generation and recursively splits the search space using a branch-and-bound algorithm to find the best integer-feasible solution.

To solve the LP relaxation of (1), referred to as the master problem, we iteratively compute an optimal solution to a restricted version that considers all first-stage binary variables, but only a small subset of second-stage path variables. We then solve the pricing subproblem to determine the variable in the original master problem with the lowest reduced cost. If the lowest reduced cost is negative, the corresponding variable is added to the restricted master problem, which is subsequently solved. Column generation terminates when all reduced costs are nonnegative, guaranteeing that the optimal solution to the restricted master problem is optimal for the original master problem.

If the resulting solution to the master problem satisfies the integrality constraints (1m)-(1n), then the branch-and-price algorithm terminates with an optimal solution to (1). Otherwise, the algorithm branches on a first-stage variable with a fractional value, creating new problems to be solved using column generation. Each problem is represented as a node of a branch-and-bound tree, and the algorithm terminates when all tree nodes have been processed. The algorithm updates the incumbent solution as it obtains integer-feasible solutions with lower objective values. At optimality, the incumbent solution is optimal for (1). We note that if the LP relaxation at a given tree node has a fractional optimal solution with an objective value that is no less than that of the incumbent solution, then the algorithm does not branch on that fractional solution, as it cannot lead to a better incumbent solution. Furthermore, column generation at each tree node is warm-started using all columns that have been generated at previously processed nodes.

To improve scalability of the algorithm we leverage the problem's structure to develop acceleration techniques. The first innovative element of our approach consists of designing a heuristic procedure to initialize the branch-and-price algorithm with a good-quality incumbent solution, leading to a smaller number of problems to solve. The heuristic first solves a facility location model without considering any road improvements and then solves another problem that uses the remaining budget to improve roads. In the second acceleration technique, we leverage the multi-period structure to adjust the branching rule using special ordered sets. Finally, we develop valid inequalities based on optimal dual solutions of the second-stage problem to reduce the number of iterations of column generation in subsequent tree nodes. We add these valid inequalities when the algorithm finds an integer-feasible solution, making our tailored approach a branch-price-and-cut algorithm.

3.2 Column Generation

We solve every LP relaxation at a node of the branch-and-bound tree using the column generation algorithm. Given a tree node, let $\mathcal{B}$ denote the set of constraints obtained by branching on first-stage variables at its ancestor nodes. For instance, $\mathcal{B}$ may enforce services to be added to a subset of facilities at the start of some periods or require services to be unavailable at the start of some periods. Similarly, this set encapsulates constraints that require arcs to be available at the start of some periods, or that prevent arcs from being constructed at the start of some periods. The master problem we aim to solve at that tree node can be formulated as the following linear program:

$$
\begin{aligned}
\mathcal{L}(\mathcal{B},\Lambda): \quad & \underset{u,v,x,y,z}{\text{minimize}} && \sum_{\tau\in\mathcal{T}}\sum_{\omega\in\Omega^\tau} p^{\tau\omega}\sum_{\lambda\in\Lambda} d^{\tau\omega}_{o_\lambda}\ell^{\tau\omega}_{\lambda} z^{\tau\omega}_{\lambda} && && \text{(2a)}\\
& \text{subject to} && (1b)-(1k), && && \\
& && z^{\tau\omega}_{\lambda}\geq 0, && \forall\ \lambda\in\Lambda,\ \forall\ \omega\in\Omega^\tau,\ \forall\ \tau\in\mathcal{T}, && \text{(2b)}\\
& && 0\leq x^\tau_f,\ u^\tau_f\leq 1, && \forall\ f\in\mathcal{F},\ \forall\ \tau\in\mathcal{T}, && \text{(2c)}\\
& && 0\leq y^\tau_a,\ v^\tau_a\leq 1, && \forall\ a\in\mathcal{A},\ \forall\ \tau\in\mathcal{T}. && \text{(2d)}\\
& && (u,v,x,y)\in\mathcal{B}. && && \text{(2e)}
\end{aligned}
$$

The master problem $\mathcal{L}(\mathcal{B}, \Lambda)$ cannot be solved directly using commercial solvers due to the extremely large number of path variables. Instead, we develop a column generation algorithm that iteratively introduces subsets of variables. Specifically, we solve the restricted master problem $\mathcal{L}(\mathcal{B}, \overline{\Lambda})$, obtained by only considering a subset of the path variables $\overline{\Lambda} \subseteq \Lambda$. We denote by $(\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z})$ the resulting optimal solution to the restricted master problem. To determine whether or not this solution is optimal for the original master problem $\mathcal{L}(\mathcal{B}, \Lambda)$, we check whether any excluded path variable $z_\lambda^{\tau\omega}$ for $\lambda \in \Lambda \setminus \overline{\Lambda}$ has a negative reduced cost.

For every $o \in \mathcal{O}$, $f \in \mathcal{F}$, $a \in \mathcal{A}$, $\tau \in \mathcal{T}$, and $\omega \in \Omega^\tau$, we let $\hat{\alpha}_o^{\tau\omega}$, $\hat{\beta}_{oa}^{\tau\omega}$, and $\hat{\gamma}_{of}^{\tau\omega}$ respectively represent the dual variables associated with (1i), (1j), and (1k) at optimality of $\mathcal{L}(\mathcal{B}, \overline{\Lambda})$. Then, the reduced cost of *any* path variable $z_\lambda^{\tau\omega}$ is given by:

$$\bar{c}(z_\lambda^{\tau\omega}) \coloneqq p^{\tau\omega} d_{o_\lambda}^{\tau\omega} \ell_\lambda^{\tau\omega} - \hat{\alpha}_{o_\lambda}^{\tau\omega} - \sum_{a \in \lambda} \hat{\beta}_{o_\lambda a}^{\tau\omega} - \hat{\gamma}_{o_\lambda f_\lambda}^{\tau\omega}, \ \forall\, \lambda \in \Lambda, \ \forall\, \omega \in \Omega^\tau, \ \forall\, \tau \in \mathcal{T}. \quad (3)$$

To avoid complete enumeration of the exponentially large number of such reduced costs, we instead formulate the pricing problem, which aims to find the variable with the lowest reduced cost. It is critical to efficiently solve this pricing problem to enable scalability of the column generation algorithm. By leveraging the network structure of the problem, we show in the next proposition that the pricing problem can be solved in polynomial time using shortest-path subroutines:

Proposition 1 *Given a population center $o \in \mathcal{O}$, a time period $\tau \in \mathcal{T}$, and a scenario $\omega \in \Omega^\tau$, the lowest reduced cost of a variable $z_\lambda^{\tau\omega}$ corresponding to path λ that originates at location o is given by*

$$\min_{\{\lambda \in \Lambda \,:\, o_\lambda = o\}} \sum_{a \in \lambda} \left(p^{\tau\omega} d_o^{\tau\omega} \ell_a^{\tau\omega} - \hat{\beta}_{oa}^{\tau\omega} \right) - \hat{\alpha}_o^{\tau\omega} - \hat{\gamma}_{of_\lambda}^{\tau\omega}, \quad (4)$$

which can be solved as a shortest-path problem in $O(|\mathcal{N}||\mathcal{A}|)$ time.

Therefore, computing the variable in $\mathcal{L}(\mathcal{B}, \Lambda)$ with lowest reduced cost can be achieved in $O(|\mathcal{N}|^2 |\mathcal{A}| \sum_{\tau \in \mathcal{T}} |\Omega^\tau|)$ time.

Proof of Proposition 1. Let $o \in \mathcal{O}$, $\tau \in \mathcal{T}$, and $\omega \in \Omega^\tau$. We create a graph $\mathcal{G}_o^{\tau\omega}$ as a copy of the graph of potential improved road networks $\mathcal{G} = (\mathcal{N}, \mathcal{A})$. Next, we augment $\mathcal{G}_o^{\tau\omega}$ by adding a dummy sink node $\bar{f}$ and connecting it to each facility $f \in \mathcal{F}$ with a single arc. The resulting graph is then given by $\mathcal{G}_o^{\tau\omega} = (\mathcal{N} \cup \{\bar{f}\}, \mathcal{A} \cup \{(f, \bar{f}),\ \forall\, f \in \mathcal{F}\})$. We set the length l_a of each arc a in $\mathcal{G}_o^{\tau\omega}$ as follows:

$$l_a = \begin{cases} p^{\tau\omega} d_o^{\tau\omega} \ell_a^{\tau\omega} - \hat{\beta}_{oa}^{\tau\omega} & \text{if } a \in \mathcal{A} \\ -\hat{\alpha}_o^{\tau\omega} - \hat{\gamma}_{of}^{\tau\omega} & \text{if } a = (f, \bar{f}) \text{ for some } f \in \mathcal{F}. \end{cases} \quad (5)$$

An example of such a graph is provided in Appendix A.1. We note that all arcs $a \in \mathcal{A}$ have nonnegative lengths since $\hat{\beta}_{oa}^{\tau\omega} \leq 0$.

In the graph $\mathcal{G}_o^{\tau\omega}$, we observe that the length of every path originating at o and traversing facility f to reach $\bar{f}$ has a length of $\sum_{a \in \lambda} \left(p^{\tau\omega} d_o^{\tau\omega} \ell_a^{\tau\omega} - \hat{\beta}_{oa}^{\tau\omega} \right) - \hat{\alpha}_o^{\tau\omega} - \hat{\gamma}_{of}^{\tau\omega}$. Therefore, Problem (4) can be solved by determining the shortest path in $\mathcal{G}_o^{\tau\omega}$ from o to $\bar{f}$. The Bellman-Ford algorithm finds such a path in $O((|\mathcal{N}| + 1)(|\mathcal{A}| + |\mathcal{F}|)) = O(|\mathcal{N}||\mathcal{A}|)$ time.

By repeating this process for every $o \in \mathcal{O}$, $\tau \in \mathcal{T}$, and $\omega \in \Omega^\tau$, we can find the variable $z_\lambda^{\tau\omega}$ with the lowest reduced cost in $O(|\mathcal{N}|^2|\mathcal{A}|\sum_{\tau\in\mathcal{T}}|\Omega^\tau|)$ time. □

From Proposition 1, we obtain that the pricing problem can be efficiently solved by running $|\mathcal{O}|\sum_{\tau\in\mathcal{T}}|\Omega^\tau|$ shortest-path algorithms. Let $\hat{\lambda}_o^{\tau\omega} \in \Lambda$ represent an optimal solution of (4) for every $o \in \mathcal{O}$, $\tau \in \mathcal{T}$, and $\omega \in \Omega^\tau$. If $\bar{c}(z^{\tau\omega}_{\hat{\lambda}_o^{\tau\omega}}) \geq 0$ for every $o \in \mathcal{O}$, $\tau \in \mathcal{T}$, and $\omega \in \Omega^\tau$, then $(\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z})$ is an optimal solution of the master problem $\mathcal{L}(\mathcal{B}, \Lambda)$. Otherwise, we augment Λ by including the paths we have found with negative reduced costs and solve the new restricted master problem.

We note that our approach deviates from a traditional implementation of column generation in that if a path λ is identified such that $\bar{c}(z_\lambda^{\tau'\omega'}) < 0$ for some $\tau' \in \mathcal{T}$ and $\omega' \in \Omega^{\tau'}$, then the corresponding variables $z_\lambda^{\tau\omega}$ are added to the new restricted master problem for *all* $\tau \in \mathcal{T}$ and $\omega \in \Omega^\tau$. Due to the structure of the problem, the same path is likely to be used by patients across many time periods and scenarios, thus adding all corresponding variables reduces the number of iterations of the column generation algorithm. We detail our implementation of column generation in Algorithm 1.

3.3 Branch-and-Price Algorithm

We now formally describe in Algorithm 2 our custom branch-and-price algorithm for solving (1).

3.4 Acceleration Techniques

To further enhance the scalability of the branch-and-price algorithm, we devise additional acceleration techniques that leverage the structure of the problem.

3.4.1 Heuristic Solution

We design a two-step optimization-based method for generating good-quality feasible solutions of (1), providing a strong incumbent solution at the onset of Algorithm 2. In this method, we first solve a simpler version of (1) by assuming that the planner can only invest in adding services to facilities. Then, we solve a pure network design problem that uses the budget that is left over after expanding facilities to improve the road network.

To solve the first problem, we consider a dummy sink node $\bar{f}$, and connect it to each facility $f \in \mathcal{F}$ with a dummy arc of length equal to zero. We also let $\mathcal{A}^0 \subseteq \mathcal{A} \cup \{(f, \bar{f}),\ \forall\ f \in \mathcal{F}\}$ denote the set of arcs already constructed in the road network (including the dummy arcs). For the sake of exposition, we assume that all arcs in $\mathcal{A}^0$ are of surface type 1. Then, we consider the same variables u (resp. x) to indicate when cancer services are added (resp. available) at each facility. However, we replace the second-stage decisions with continuous variables $w_a^{\tau\omega}$, representing the flow of patients (originating at any population center) traveling along arc $a \in \mathcal{A}^0$ during scenario $\omega \in \Omega^\tau$ at time period $\tau \in \mathcal{T}$. Since arcs in $\mathcal{A}^0$ are originally undirected, we arbitrarily direct them. Consequently, $w_a^{\tau\omega} < 0$ signifies that $-w_a^{\tau\omega}$ patients are traveling arc a in its opposite direction. We also introduce auxiliary variables $\theta_a^{\tau\omega}$ (for $a \in \mathcal{A}$, $\omega \in \Omega^\tau$, $\tau \in \mathcal{T}$) to represent the absolute value of $w_a^{\tau\omega}$, which corresponds to the nonnegative

Algorithm 1: Column Generation for Solving $\mathcal{L}(\mathcal{B}, \Lambda)$

Input : Tree node with added constraints $\mathcal{B}$, subset of paths $\overline{\Lambda}$
Output: Optimal solution $(\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z})$, generated path set $\bar{\Lambda}$, optimal value $\hat{\xi}$

do
  Solve the restricted master problem $\mathcal{L}(\mathcal{B}, \bar{\Lambda})$:
  $(\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z}) \leftarrow$ optimal solution, $\hat{\xi} \leftarrow$ optimal value
  $(\hat{\alpha}_o^{\tau\omega}), (\hat{\beta}_{oa}^{\tau\omega}), (\hat{\gamma}_{of}^{\tau\omega}) \leftarrow$ optimal dual variable values associated with (1i), (1j), and (1k).
  $\Lambda' \leftarrow \emptyset$
  for every $o \in \mathcal{O},\ \tau \in \mathcal{T},\ \omega \in \Omega^\tau$ **do**
    $\mathcal{G}_o^{\tau\omega} \leftarrow (\mathcal{N} \cup \{\bar{f}\}, \mathcal{A} \cup \{(f, \bar{f}),\ \forall\ f \in \mathcal{F}\})$
    for every $a \in \mathcal{A}$ **do**
      Set length of a in $\mathcal{G}_o^{\tau\omega}$ to $p^{\tau\omega} d_o^{\tau\omega} \ell_{oa}^{\tau\omega} - \hat{\beta}_{oa}^{\tau\omega}$
    end
    for every $f \in \mathcal{F}$ **do**
      Set length of $(f, \bar{f})$ in $\mathcal{G}_o^{\tau\omega}$ to $-\hat{\alpha}_o^{\tau\omega} - \hat{\gamma}_{of}^{\tau\omega}$
    end
    $\hat{\lambda}_o^{\tau\omega} \leftarrow$ optimal solution of (4) by computing the shortest path from o to $\bar{f}$ in $\mathcal{G}_o^{\tau\omega}$ using Bellman-Ford algorithm
    if $\bar{c}(z^{\tau\omega}_{\hat{\lambda}_o^{\tau\omega}}) < 0$ **then**
      $\Lambda' \leftarrow \Lambda' \cup \{\hat{\lambda}_o^{\tau\omega}\}$
    end
  end
  $\bar{\Lambda} \leftarrow \bar{\Lambda} \cup \Lambda'$
while $\Lambda' \neq \emptyset$
return $((\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z}), \bar{\Lambda}, \hat{\xi})$

number of patients traveling along arc a. For every node $i \in \mathcal{N} \cup \{\bar{f}\}$, we denote by $\delta^+(i) \subseteq \mathcal{A}^0$ (resp. $\delta^-(i) \subseteq \mathcal{A}^0$) its subset of outgoing arcs (resp. incoming arcs). The problem can then be formulated as the following MIP:

$$\underset{u,x,w,\theta}{\text{minimize}} \quad \sum_{\tau\in\mathcal{T}} \sum_{\omega\in\Omega^\tau} p^{\tau\omega} \sum_{a\in\mathcal{A}^0} \ell_a^{\tau\omega} \theta_a^{\tau\omega} \tag{6a}$$

$$\text{subject to} \quad \sum_{\tau=1}^{\tau'} \sum_{f\in\mathcal{F}} c_f^\tau u_f^\tau \leq \sum_{\tau=1}^{\tau'} C^\tau,\ \forall\ \tau' \in \mathcal{T}, \tag{6b}$$

$$\sum_{\tau=1}^{\tau'} \sum_{f\in\mathcal{F}} h_f^\tau u_f^\tau \leq \sum_{\tau=1}^{\tau'} H^\tau,\ \forall\ \tau' \in \mathcal{T}, \tag{6c}$$

$$x_f^\tau - x_f^{\tau-1} = u_f^\tau,\ \forall\ f \in \mathcal{F},\ \forall\ \tau \in \{2, \ldots, T\}, \tag{6d}$$

Algorithm 2: Branch-and-Price Algorithm for Solving (1)

Input: Instance of (1)
Output: Optimal solution $(u^*, v^*, x^*, y^*, z^*)$ and optimal value ξ^*

Initialize the incumbent objective value $\xi^* \leftarrow +\infty$, a restricted path set $\overline{\Lambda}$ to guarantee feasibility, a set of branching constraints $\mathcal{B}_0 \leftarrow \emptyset$, a list of pending nodes $\mathcal{Q} \leftarrow (\mathcal{B}_0)$
while $\mathcal{Q} \neq \emptyset$ **do**
  Remove a pending node (with corresponding set $\mathcal{B}$) from $\mathcal{Q}$
  Solve $\mathcal{L}(\mathcal{B}, \Lambda)$ using column generation, warm-started with $\overline{\Lambda}$:
  $((\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z}), \bar{\Lambda}, \hat{\xi}) \leftarrow$ output of Algorithm 1
  if $\hat{\xi} < \xi^*$ **then**
    if $(\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z})$ is integer feasible **then**
      $(u^*, v^*, x^*, y^*, z^*) \leftarrow (\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z}), \quad \xi^* \leftarrow \hat{\xi}$
    else
      Select a fractional variable on which to branch
      for each new branching constraint **do**
        Create a new set $\mathcal{B}'$ by appending the branching constraint to $\mathcal{B}$, and add the corresponding pending node to $\mathcal{Q}$
      end
    end
  end
end
return $((u^*, v^*, x^*, y^*, z^*), \xi^*)$

$$x_f^1 = u_f^1, \ \forall\ f \in \mathcal{F}, \tag{6e}$$

$$-\theta_a^{\tau\omega} \leq w_a^{\tau\omega} \leq \theta_a^{\tau\omega}, \ \forall\ a \in \mathcal{A}, \ \forall\ \omega \in \Omega^\tau, \ \forall\ \tau \in \mathcal{T}, \tag{6f}$$

$$\sum_{a \in \delta^+(i)} w_a^{\tau\omega} - \sum_{a \in \delta^-(i)} w_a^{\tau\omega} = \begin{cases} d_i^{\tau\omega} & \text{if } i \in \mathcal{O} \\ -\sum_{o \in \mathcal{O}} d_o^{\tau\omega} & \text{if } i = \bar{f} \\ 0 & \text{otherwise,} \end{cases}$$
$$\forall\ i \in \mathcal{N} \cup \{\bar{f}\}, \ \forall\ \omega \in \Omega^\tau, \ \forall \tau \in \mathcal{T}, \tag{6g}$$

$$-\sum_{o \in \mathcal{O}} d_o^{\tau\omega} x_f^\tau \leq w_{f\bar{f}}^{\tau\omega} \leq \sum_{o \in \mathcal{O}} d_o^{\tau\omega} x_f^\tau, \ \forall\ f \in \mathcal{F}, \ \forall \omega \in \Omega^\tau, \ \forall \tau \in \mathcal{T}, \tag{6h}$$

$$x_f^\tau, \ u_f^\tau \in \{0, 1\}, \quad \forall\ f \in \mathcal{F}, \ \forall\ \tau \in \mathcal{T}. \tag{6i}$$

In this MIP, the flow conservation constraints (6g) ensure that patients travel from their population centers to the dummy sink node (after going through one facility). Furthermore, constraints (6h) prevent patients from traveling to a facility where services are not available.

We solve the smaller MIP (6) directly with commercial solvers; we denote its optimal solution by $(\hat{u}, \hat{x}, \hat{w}, \hat{\theta})$. In the next step, we allocate the remaining monetary

budget to improve the road network: We develop a second optimization problem to minimize the patients' expected travel times, assuming that they will travel according to $\hat{\theta}$ towards facilities with service availability governed by $\hat{x}_f^\tau$ at optimality of (6). Given the same variables v and y to indicate when arcs are constructed and available, we formulate the following MIP:

$$\underset{v,\,y}{\text{minimize}} \quad \sum_{\tau\in\mathcal{T}}\sum_{\omega\in\Omega^\tau} p^{\tau\omega} \sum_{(i,j,s)\in\mathcal{A}} \hat{\theta}_{ij1}^{\tau\omega}(\ell_{ijs}^{\tau\omega} - \ell_{ij(s-1)}^{\tau\omega})y_{ijs}^\tau \tag{7a}$$

$$\text{subject to} \quad \sum_{\tau=1}^{\tau'}\sum_{a\in\mathcal{A}} c_a^\tau v_a^\tau \le \sum_{\tau=1}^{\tau'}\left(C^\tau - \sum_{f\in\mathcal{F}} c_f^\tau \hat{u}_f^\tau\right), \quad \forall\ \tau'\in\mathcal{T}, \tag{7b}$$

$$y_a^\tau - y_a^{\tau-1} = v_a^\tau, \quad \forall\ a\in\mathcal{A},\ \forall\ \tau\in\{2,\dots,T\}, \tag{7c}$$

$$y_a^1 = v_a^1, \quad \forall\ a\in\mathcal{A}, \tag{7d}$$

$$y_{ijs}^\tau \le y_{ij(s-1)}^\tau, \quad \forall\ (i,j,s)\in\mathcal{A}\ \text{ s.t. } s\ge 2,\ \forall\ \tau\in\mathcal{T}, \tag{7e}$$

$$y_a^\tau,\ v_a^\tau \in \{0,1\}, \quad \forall\ a\in\mathcal{A},\ \forall\ \tau\in\mathcal{T}, \tag{7f}$$

with the convention that $\ell_{ij0}^{\tau\omega} \coloneqq 0$. We note that the inner summation in the objective function (7a) indeed represents the travel time incurred by the $\hat{\theta}_{ij1}^{\tau\omega}$ patients traveling between i and j on the best surface type available during time period τ.

We solve the small MIP (7) using commercial solvers and denote by $(\hat{v},\hat{y})$ (resp. $\hat{\xi}$) its optimal solution (resp. optimal value). From this approach, we obtain feasible first-stage variables $\hat{u},\hat{v},\hat{x},\hat{y}$ of (1) and an incumbent objective value to warm-start our branch-and-price algorithm.

3.4.2 Branching Rules

Conventional bifurcated branching in branch-and-bound divides the search space at a given node into two subproblems by enforcing that a fractional variable at the node takes on a value of either 0 or 1. These subproblems are added as child nodes in the tree. For instance, if column generation outputs an optimal solution to the LP relaxation $\mathcal{L}(\mathcal{B},\Lambda)$ at a tree node with a fractional first-stage variable $\hat{u}_f^\tau$ for some $f\in\mathcal{F}$ and $\tau\in\mathcal{T}$, then conventional branching will create two child nodes, where in the first (resp. second) node, the constraint $u_f^\tau = 0$ (resp. $u_f^\tau = 1$) is added to $\mathcal{B}$.

From our problem's structure, there is at most one period at the start of which cancer services are added to a facility (or similarly, an arc is constructed). Thus, the improvement decision variables u and v form a special ordered set with at most one non-zero entry. We use this structure in our branching strategy. For instance, if $\hat{u}_f^\tau$ is fractional at optimality of $\mathcal{L}(\mathcal{B},\Lambda)$ at a tree node for some $f\in\mathcal{F}$ and $\tau\in\mathcal{T}$, we implement this special wide branching as follows. Instead of branching on the u_f^τ variables corresponding to facility f separately, we branch on the set $(u_f^1,\dots,u_f^T)$ to identify when (if ever) facility f will be expanded. This process creates $T+1$ subproblems, one for each $\tau\in\mathcal{T}$ to represent that $u_f^\tau = 1$ and $u_f^{\tau'} = 0$ for $\tau'\neq\tau$, and a final subproblem corresponding to $u_f^\tau = 0,\ \forall\ \tau\in\mathcal{T}$. The process is similar for arc construction variables v.

Additionally, sequential decisions by the planner imply that fixing first-stage decisions in earlier periods reduces the number of feasible decisions in later periods. Therefore, we prioritize branching on fractional variables in earlier periods before considering those in later periods. We found that this branching strategy accelerates the convergence of our algorithm.

3.4.3 Valid Inequalities

As we process nodes of the branch-and-bound tree, we aim to incorporate information from the optimal solutions of the associated LP relaxations into the remaining tree nodes. To this end, we investigate valid inequalities, as described in the following proposition:

Proposition 2 *Given a population center $o \in \mathcal{O}$, a time period $\tau \in \mathcal{T}$ and a scenario $\omega \in \Omega^\tau$, valid inequalities to the MIP* (1) *are given by*

$$\sum_{\{\lambda \in \Lambda \,:\, o_\lambda = o\}} \ell_\lambda^{\tau\omega} z_\lambda^{\tau\omega} \geq \alpha_o^{\tau\omega} + \sum_{a \in \mathcal{A}} \beta_{oa}^{\tau\omega} x_a^\tau + \sum_{f \in \mathcal{F}} \gamma_{of}^{\tau\omega} w_f^\tau, \tag{8}$$

for any α, β, γ satisfying

$$\alpha_o^{\tau\omega} + \sum_{a \in \lambda} \beta_{oa}^{\tau\omega} + \gamma_{of_\lambda}^{\tau\omega} \leq \ell_\lambda^{\tau\omega}, \quad \forall\, \lambda \in \Lambda \,:\, o_\lambda = o, \tag{9}$$

$$\beta_{oa}^{\tau\omega} \leq 0, \quad \forall\, a \in \mathcal{A}, \tag{10}$$

$$\gamma_{of}^{\tau\omega} \leq 0, \quad \forall\, f \in \mathcal{F}. \tag{11}$$

Proof of Proposition 2. We consider a feasible solution $(\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z})$ to (1), a population center $o \in \mathcal{O}$, a time period $\tau \in \mathcal{T}$, and a scenario $\omega \in \Omega^\tau$. By definition of (1), we obtain that:

$$\begin{aligned}
\sum_{\{\lambda \in \Lambda \,:\, o_\lambda = o\}} \ell_\lambda^{\tau\omega} \hat{z}_\lambda^{\tau\omega} \geq \underset{z}{\text{minimize}} \quad & \sum_{\{\lambda \in \Lambda \,:\, o_\lambda = o\}} \ell_\lambda^{\tau\omega} z_\lambda^{\tau\omega} \\
\text{subject to} \quad & \sum_{\{\lambda \in \Lambda \,:\, o_\lambda = o\}} z_\lambda^{\tau\omega} = 1, \\
& \sum_{\{\lambda \in \Lambda \,:\, o_\lambda = o,\, a \in \lambda\}} z_\lambda^{\tau\omega} \leq \hat{y}_a^\tau, \quad \forall\, a \in \mathcal{A}, \\
& \sum_{\{\lambda \in \Lambda \,:\, o_\lambda = o,\, f_\lambda = f\}} z_\lambda^{\tau\omega} \leq \hat{x}_f^\tau, \quad \forall\, f \in \mathcal{F}, \\
& z_\lambda^{\tau\omega} \geq 0 \quad \forall\, \lambda \in \Lambda \,:\, o_\lambda = o.
\end{aligned}$$

By duality, we obtain the following equivalent inequality:

$$\begin{aligned}
\sum_{\{\lambda \in \Lambda \,:\, o_\lambda = o\}} \ell_\lambda^{\tau\omega} \hat{z}_\lambda^{\tau\omega} \geq \underset{\alpha, \beta, \gamma}{\text{maximize}} \quad & \alpha_o^{\tau\omega} + \sum_{a \in \mathcal{A}} \beta_{oa}^{\tau\omega} \hat{y}_a^\tau + \sum_{f \in \mathcal{F}} \gamma_{of}^{\tau\omega} \hat{x}_f^\tau \\
\text{subject to} \quad & \alpha_o^{\tau\omega} + \sum_{a \in \lambda} \beta_{oa}^{\tau\omega} + \gamma_{of_\lambda}^{\tau\omega} \leq \ell_\lambda^{\tau\omega}, \ \forall\, \lambda \in \Lambda \,:\, o_\lambda = o,
\end{aligned}$$

$$\beta_{oa}^{\tau\omega} \leq 0, \ \forall \ a \in \mathcal{A},$$
$$\gamma_{of}^{\tau\omega} \leq 0, \ \forall \ f \in \mathcal{F}.$$

Thus, we obtain that for any (α, β, γ) satisfying (9)-(11), the following inequality holds:

$$\sum_{\{\lambda \in \Lambda \,:\, o_\lambda = o\}} \ell_\lambda^{\tau\omega} \hat{z}_\lambda^{\tau\omega} \geq \alpha_o^{\tau\omega} + \sum_{a \in \mathcal{A}} \beta_{oa}^{\tau\omega} \hat{y}_a^\tau + \sum_{f \in \mathcal{F}} \gamma_{of}^{\tau\omega} \hat{x}_f^\tau.$$

Since constraints (9)-(11) do not depend on $(\hat{u}, \hat{v}, \hat{x}, \hat{y}, \hat{z})$, we deduce that (8) is a valid inequality for the MIP (1). □

From Proposition 2, we can generate these decomposition-based valid inequalities (or equivalently cuts) to speed up the column generation algorithm when processing tree nodes. Specifically, within the branch-and-price algorithm (Algorithm 2), we introduce a global cutset $\mathcal{C}$ that is added to the master problem $\mathcal{L}(\mathcal{B}, \Lambda)$ at each tree node. Every time we process a node and the column generation algorithm terminates with an integer-feasible solution, we extract the optimal dual variables $\hat{\alpha}_o^{\tau\omega}$, $\hat{\beta}_{oa}^{\tau\omega}$, and $\hat{\gamma}_{of}^{\tau\omega}$ respectively associated with (1i), (1j), and (1k), for every $o \in \mathcal{O}$, $f \in \mathcal{F}$, $a \in \mathcal{A}$, $\tau \in \mathcal{T}$, and $\omega \in \Omega^\tau$. Such optimal dual variables naturally satisfy (9)-(11) and we add the corresponding decomposition-based cuts (8) to the cutset $\mathcal{C}$, which will be applied to all subsequent tree nodes.

We note that adding the cutset $\mathcal{C}$ to each master problem impacts reduced cost computations during column generation. Nonetheless, we still find that the pricing problems can be solved using shortest-path algorithms. In Appendix A.2, we discuss refinements to the process of generating cuts. First, we consider that cuts of the form (8) may be dominated by previously generated cuts. Therefore, we provide an algorithm to generate a set of non-dominated cuts to be added to the master problem at each tree node. Second, adding cuts early in the algorithm can further accelerate its convergence. Thus, we aim to derive cuts from the integer feasible solution identified via the heuristic from §3.4.1. However, dual variables cannot be computed using complementary slackness and are not directly available. In Appendix A.2, we show how to analytically derive cuts from only first-stage variables.

4 Computational Study

In this section, we analyze the computational performance of our proposed algorithm on synthetic test cases. The primary purpose of these experiments is to evaluate the efficacy of our approach as a function of network size and to quantify the benefit of the valid inequalities derived in §3.4.3. Our secondary purpose is to identify the parameters of (1) to which the performance of the proposed approach is sensitive.

4.1 Test Instances

We generate a portfolio of random test instances based on methods used in prior literature [22, 23, 25]. First we fix the number of nodes, $|\mathcal{N}|$, number of candidate facilities, $|\mathcal{F}|$, and number of time periods, T. Then, we randomly distribute $|\mathcal{N}|$ nodes across a square grid and sample $|\mathcal{F}|$ of these nodes to be the set of candidate facilities. We assume that $\mathcal{O} = \mathcal{N}$, namely, patients can originate from any node. To generate

the road network, we first construct a random adjacency matrix such that nodes that are closer to each other are more likely to have an arc connecting them than nodes that are farther away from each other. We then randomly assign the surface type for each road from a set of $|\mathcal{S}| = 5$ surface types, with lower-quality types being more likely than higher-quality types. In so doing, we construct the road network available at the beginning of the planning horizon. Note that the graph $\mathcal{G} = (\mathcal{N}, \mathcal{A})$ of potential improved road networks can be obtained by making a copy of each road with different surface types, thus generating the set of type-dependent arcs $\mathcal{A}$. We also generate $|\Omega^\tau| = 2$ event scenarios for each time period τ. For brevity, we defer additional details regarding the random generation of the adjacency matrix, surface types, and other model parameters to Appendix B.

In total, we construct a set of 360 test instances. These 360 test instances result from varying the the number of nodes in the network to 4 different values ($|\mathcal{N}| \in \{10, 20, 50, 75\}$) and the number of candidate facilities to 3 different values ($|\mathcal{F}| \in \{\lfloor \frac{1}{4} \cdot |\mathcal{N}| \rfloor, \lfloor \frac{1}{2} \cdot |\mathcal{N}| \rfloor, \lfloor \frac{3}{4} \cdot |\mathcal{N}| \rfloor\}$). Then, for each of the 12 $(|\mathcal{N}|, |\mathcal{F}|)$ combinations, we generate 30 different instances according to the procedure above. We fix $T = 5$ for all instances.

For each instance, we calculate the total cost of expanding all candidate facilities during the first time period, $\bar{C} \coloneqq \sum_{f \in \mathcal{F}} c_f^1$. Then, we assume that a certain proportion, κ, of $\bar{C}$ would be divided equally across all time periods so that $C^\tau = \kappa \cdot \frac{\bar{C}}{T}$ for all $\tau \in \mathcal{T}$. Similarly, we calculate $\bar{H} \coloneqq \sum_{f \in \mathcal{F}} h_f^1$ and set $H^\tau = \kappa \cdot \frac{\bar{H}}{T}$ for all $\tau \in \mathcal{T}$.

We conduct two sets of experiments. In the first set of experiments, we explore the impact of problem size on solution times, considering all 360 random instances for a fixed value of $\kappa = 50\%$. In the second set of experiments, we only consider 15 random instances for which $|\mathcal{N}| = 50$ and vary the budgets by considering $\kappa \in \{10\%, 20\%, \ldots, 100\%\}$. We hypothesize κ may impact the performance of our solution method.

In all experiments, we compare the performance of our solution methodology to a branch-and-price solver. That is, we compare our complete branch-price-and-cut (BP&C) algorithm with a branch-and-price (B&P) algorithm that uses the column generation approach discussed in §3.2 but does not add the cuts discussed in §3.4.3 (hereafter, referred to as "B&P"). We assume that both algorithms employ the heuristically generated warm start from §3.4.1 and the special ordered set branching rule from §3.4.2. We compare the performance of these two algorithms to evaluate the potential benefits of adding the valid inequalities into a B&P scheme.

All algorithms are implemented in Python 3.11. We solve all optimization problems using Gurobi v.11.0.2, running at 3.0 Ghz (1 core) with 32 GB of assigned RAM and Red Hat Linux v.8.10. We specify a time limit of 120 minutes for all instances.

4.2 Results

Figure 2 presents the results from our first set of experiments. When the number of nodes in the network is relatively small (i.e., $|\mathcal{N}| \in \{10, 20\}$), we observe that the BP&C algorithm is able to solve all instances to optimality within 28 minutes. Meanwhile, the B&P algorithm is able to solve most, but not all, of the instances within

the time limit. For the larger network sizes that we tested (i.e., $|\mathcal{N}| \in \{50, 75\}$) we observe that the BP&C algorithm solves at least as many instances as the B&P algorithm within the time limit for all values of $|\mathcal{F}|$. This trend is evident in Figure 2 as the solid lines (representing BP&C) frequently appear above the dashed lines (representing B&P), indicating a higher fraction of instances solved. For example, with 50 nodes, BP&C solves 26 (87%) of the 30 instances with the fewest number of facilities, but B&P only solves 14 (47%) of the same instances (see the blue lines in Figure 2 for $|\mathcal{N}| = 50$). These findings demonstrate that there is value to adding the cuts into our B&P algorithm.

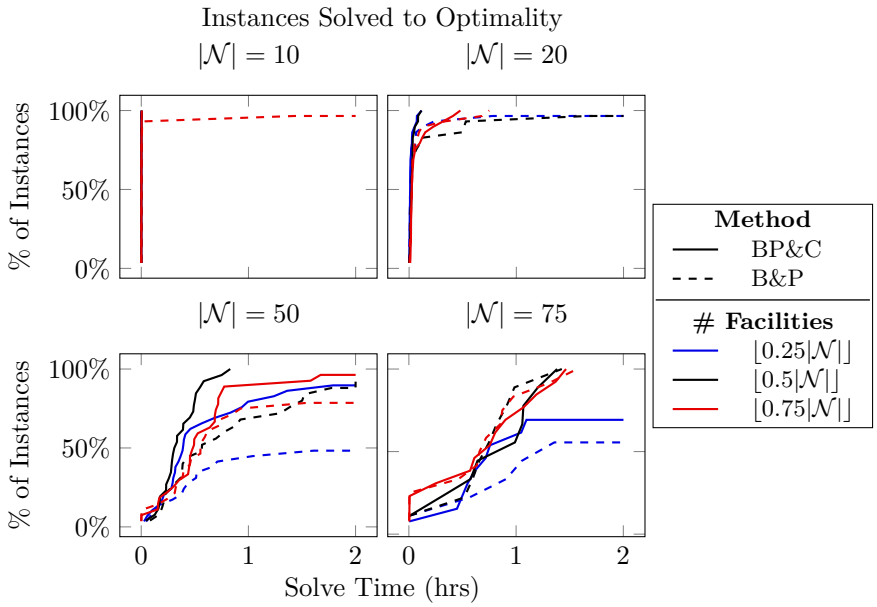

Fig. 2: The cumulative percentage of instances solved to optimality within the given time for our BP&C algorithm (solid lines) compared to the B&P approach described in §3.2 without the cuts described in §3.4.3 (dashed line). Line color is based on the fraction of nodes that are have facilities. The maximum allocated time was 120 minutes. The results are faceted by the number of nodes in the instances, $|\mathcal{N}|$. Each line is composed of the outcomes of 30 separate instances with the same $(|\mathcal{N}|, |\mathcal{F}|)$ parameters. The same instances are tested across solution methods.

Figure 3 illustrates the results of our second set of experiments that examine the sensitivity of the performance of our algorithms to the values of our budget parameters, C and H. For the smallest tested values of $\kappa = 10\%$, which correspond to the smallest tested monetary and staffing budgets for each instance, both the BP&C and the B&P struggle to obtain an optimality gap under 50%. When we consider κ in the range of 30%-50%, there are noticeable differences in performance between BP&C and B&P. For example, when $\kappa = 30\%$, the BP&C algorithm solves 4 of the 15 instances

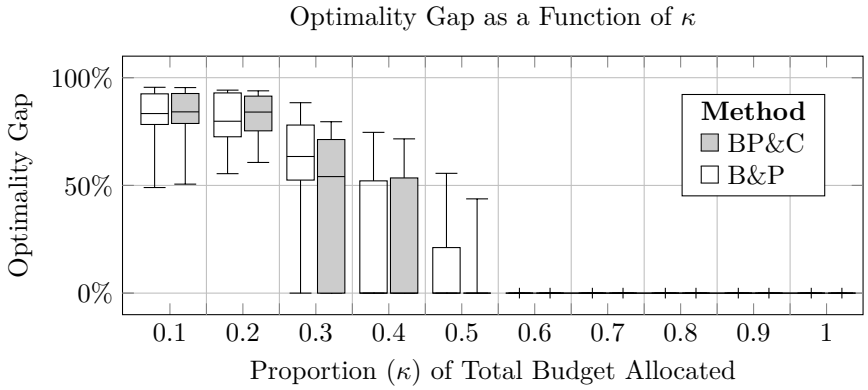

Fig. 3: Among 15 instances with $|\mathcal{N}| = 50$ we vary κ, the fraction of the total construction and staffing costs allotted for C and H, from 10% to 100%. Each bar consists of 15 instances which are each allocated 120 minutes of compute time.

to optimality, but the B&P can only solve one instance to optimality within the time limit. When $\kappa \geq 60\%$, both methods are able to solve all instances to optimality. These results suggest that solving versions of the Facility Location Network Design Problem (FLNDP) with tight budgets is challenging, as they require a more complex budget allocation to achieve optimality. On the other hand, both solution methods perform well when budgets are less restrictive. Importantly, the BP&C algorithm outperforms the B&P algorithm for moderate budget levels.

In summary, we find that BP&C outperforms B&P and that the difference in performance is most pronounced in our test cases with larger network sizes and moderate budgets. We also observe that instances where $|\mathcal{F}| \ll |\mathcal{N}|$ or $\kappa < 30\%$ are the most challenging for both solution methods. Conversely, the instances with the smallest networks and those with the highest values of κ are quickly solved regardless of the number of facilities considered.

5 Case Study: Geographic Access to Cancer Care in Rwanda

In this case study, we explore how both expansions of facilities and improvements to the road network might improve access to cancer care in Rwanda. We begin by providing some context about the current state of cancer care in Rwanda, a low-income country in East Africa, which in many ways exemplifies the challenges in geographic access to cancer care. Then, we will describe the data used to parameterize our model, the policies we compared, and the results of using our model to examine how cancer care in Rwanda might be improved under different monetary and human capital budgets.

In 2022, it was estimated that more than 7,000 people died from cancer in Rwanda, a country that only has 13.6 million people [2]. Although cancer data is limited, these data indicate that Rwanda primarily treats late-stage cancers [11] and that patients

located furthest from care are more than two times as likely to have a late-stage diagnosis compared to their peers who live close to care [37].

Rwanda's tertiary cancer care services are not distributed uniformly throughout the country. While basic pathology and radiology diagnostics are available at all district and provincial hospitals in Rwanda, there are only three hospitals that offer tertiary cancer care services. Two of these hospitals are public hospitals—the Rwandan Military Hospital (RMH) and the Butaro Cancer Center of Excellence (BCCOE)—and the third is a private hospital—the King Faisal Hospital. RMH is located in Rwanda's capital, a city of approximately 1.8 million people that is located in the middle of the country. Cancer services at RMH are provided by the Rwandan Government. BCCOE is located in the far north of the country's Northern Province, approximately 120 kilometers by road from Kigali. BCCOE is owned and operated by Partners in Health, an American non-governmental organization whose mission is to deliver health care "to the world's poorest places." The location of BCCOE was meant to provide excellent cancer care to a primarily rural population. We focus on access to public hospitals offering tertiary cancer care. All Rwandans participate in a collective insurance scheme. While some work by Rubagumya et al. [38] has detailed that the cost of care likely presents a burden, the collective insurance scheme makes care at public hospitals more financially accessible than private ones. The geographic location of the public hospitals offering tertiary cancer care is shown in Figure 4.

Traveling to these tertiary cancer care services can be challenging. Rwanda is known as "the land of a thousand hills" due to its hilly and mountainous terrain, especially in its central and western regions. Moreover, Rwanda has two three-month rainy seasons, during which the rain can be heavy and persistent. Some regions in Rwanda receive over 1500 mm (approximately, 59 inches) of annual rainfall concentrated in a few months [39]. Flooding and landslides during the rainy season are common, can be fatal, and frequently damage infrastructure [40–42]. Together, this makes travel along unpaved roads quite challenging, and certain parts of the road network can require long travel times to traverse.

Because of the locations of the existing tertiary cancer care services, there is significant variation in how long it takes Rwandans to access this type of cancer care. Less than 5% of the population within the capital province have to travel more than two hours to reach a comprehensive cancer center. Meanwhile, more than 95% of the population in the rural Western Province has to travel more than two hours to access this type of care [7]. These disparities in access are associated with disparities in the receipt of cancer care: rural regions in Rwanda have later diagnoses and experience longer diagnostic and treatment delays [11].

Previous studies have explored potential plans to expand tertiary cancer care to reference hospitals in Rwanda. Fadelu et al. [7] evaluated a policy that expanded tertiary cancer care at all reference hospitals and showed that this plan would increase the percentage of rural Rwandans living within 1 hour of care from 8% to 33%. Sapirstein et al. [8] used an optimization approach and demonstrated that other selections of reference and provincial hospitals to expand care could result in better access at lower cost. However, even with the highest levels of investment in cancer care, they demonstrate that around 20% of Rwandans would still live more than 2 hours from

tertiary cancer care. In our case study, we examine the extent to which simultaneously expanding healthcare facilities and improving the road network could improve access to cancer care in Rwanda, given these barriers to geographic access.

5.1 Model Parameterization

We construct an instance of (1) that could be used to help the Rwandan Government and Ministry of Health (MOH) plan its expansion of tertiary cancer care over a five-year planning horizon to be consistent with its five-year cancer control plans [12]. A summary of our data sources is provided in Table 3.

Table 3: Summary of the estimates of model parameters and their corresponding data sources for the case study.

Set/ Parameter	**Description of estimate(s)**	**Source**
$\mathcal{N}$	All population centers, road junctions, and health facilities in Rwanda	[15, 43]
$\mathcal{F}$	Existing district, provincial, and reference hospitals	[7, 43]
$\mathcal{O}$	The spatial distribution of Rwanda's population aggregated to the closest community health worker (CHW)	[8, 43]
$\mathcal{S}$	The five road surfaces as defined in Table 4	[15]
$\mathcal{A}$	Rwanda's road network extracted from OpenStreetMaps	[15]
$\mathcal{T}$	$\{1, 2, 3, 4, 5\}$; a five-year planning period corresponding to the MOH's cancer planning periods	[12]
Ω^τ	{dry, rainy}; the dry and rainy season in each time period τ	Assumption
c^τ_{ijs}	The cost to improve a road depends on its current surface type; estimated from the Rwandan Transport Ministry	[44]
c^τ_f	\$12.9 million United States dollar (USD) if f is currently a reference or provincial hospital; \$26.5 million USD if f is a district hospital	[45]
h^τ_f	3 oncologist teams if f is either a reference and provincial hospital; 6 oncologist teams if f is a district hospital	[8]
C^τ	Varies from \$0 to \$25 million in test cases	
H^τ	Varies from 2 to 6 in test cases	
$p^{\tau\omega}$	30% probability of the rainy scenario in each year (3-4 months of heavy rainfall per year)	[39]
$d^{\tau\omega}_o$	Proportional to the number of people in population center o, for all τ and ω	[2, 8, 46]
$\ell^{\tau\omega}_a$	Based on the estimated speeds outlined in Table 4	[7, 44, 47, 48]

First, we construct the set of population centers in our model and their associated demand for cancer care services. We use the methodology described in Sapirstein et al. [8] to estimate population density from satellite data. This process results in a continuous spatial distribution of the population across Rwanda. To create a representative discrete set of population centers, $\mathcal{O}$, we assigning population density to its closest community health worker (CHW) post in Rwanda. We assume that the prevalence of cancer is uniformly distributed across the population of Rwanda. Thus, the total demand for cancer care from a given population center o is proportional to its population. We also assume that the demand for cancer care, $d_o^{\tau\omega}$, is constant across all time periods $\tau \in \mathcal{T}$ and weather scenarios $\omega \in \Omega^\tau$ (which we define later).

We consider that the Rwandan MOH may wish to expand tertiary cancer care to hospitals that currently offer one of three levels of care (district, reference, and provincial). These existing hospitals form our candidate set of facilities, $\mathcal{F}$. We obtain the locations of these candidate facilities from Rwanda's MOH through Rwanda's Africageoportal website [43]. Consistent with prior literature, we estimate that the cost to expand tertiary cancer care at a reference or provincial hospital is \$12.9 million USD for reference and that the cost to expand this care to a district hospital is \$26.5 million USD [45]. These costs represent initial costs associated with adding cancer services. Similarly, we estimate that adding services to reference and provincial hospitals would require 3 additional oncologist teams and that district hospitals would require 6 such teams. That is, $h_f^1 = 3$ if f is either a reference or provincial hospital and $h_f^1 = 6$ otherwise.

We also consider that the Rwandan government may wish to improve its road network to improve access to care. We obtain data to construct a representation of the road network from OpenStreetMaps (OSM) as of March 2025. The existing surfaces of these roads are labeled to be consistent with OSM's classification. This process results in 5 road surface types, which are described in Table 4. To estimate the cost to improve a road between i to j from surface s to $s+1$, we consider that there is a per-kilometer cost to upgrade and that the total cost is proportional to the total distance between i to j. These surface-specific per-kilometer road construction costs are estimated from the Rwandan Ministry of Transport [44].

Table 4: The types of road surfaces considered in our case study extracted from OpenStreetMaps (OSM) [15] with their corresponding speeds during the dry and rainy seasons [7, 47, 48].

s	**Road Surface Description**	**OSM Classification**	**Speed (km/hr)**		**Cost***
			Dry Season	Rainy Season	(\$1,000/km)
1	All other roads	Unclassified	10	6	-
2	Residential Streets	Tertiary + Residential	18	11	97
3	Minor Highway	Secondary	27	16	97
4	Highways	Primary	40	37	151
5	Major Highways	Trunk + Motorway	60	54	793

*Costs are estimated from [44] and converted into United States dollars from Rwandan Francs

We consider that, each year, we have two distinct weather scenarios: the rainy seasons and the dry seasons. The time to travel along each road surface depends on the season. During the rainy seasons, the speed that a patient can travel along a road is significantly slower than the speed achievable during the dry season. The gap between the speed during the rainy season and dry season is more pronounced for lower-quality roads, such as gravel and dirt roads due to mud and washout (see Table 4, [47, 48]).

Using the instance described above, we consider several different test cases that vary in terms of the monetary budget C^τ and the staffing budget H^τ for each period τ. We vary the annual monetary budget for Year 1 from \$0 to \$25 million USD in \$5 million dollar increments. Independently, we also consider three different annual staffing budgets, H^τ: 2, 4, and 6 new oncologist teams for each $\tau \in \mathcal{T}$. The monetary cost estimates and budgets described above for Year 1 are adjusted using a 5% annual inflation rate to obtain the corresponding cost estimates in future time periods (i.e., $\tau = 2, \ldots, 5$).

5.2 Policy Comparisons and Outcome Measures

Using the instances of (1) describe above, we compare two policy approaches to expand access to cancer care in Rwanda. The first approach is to consider only improvements to healthcare infrastructure. To optimize planning using this approach, we solve (6), the Facility Location Problem (FLP) in which only facility expansions are considered; no improvements to the road network are considered. The second approach is to consider improvements to both the healthcare infrastructure and the transportation infrastructure using our proposed FLNDP model. We solve this large-scale FLNDP using our BP&C algorithm described in §3. For brevity, hereafter we refer to these two approaches as the "FLP policy" and the "FLNDP policy," respectively.

To compare these two policies, we calculate the travel time required to access tertiary cancer care on average and by quartile of current travel time, in both the rainy and the dry season. We also visualize the changes in travel time to better understand which populations see the benefits of improved access under these different policies. Further, we investigate when an FLNDP policy offers the most benefit relative to an FLP policy. To quantify these benefits, we describe the reduction in travel times that is directly attributable to road network improvements. To do so, we compare the travel times to the facilities expanded under the FLNDP in the current road network to the travel times to the facilities expanded under the FLNDP policy in the FLNDP policy's improved road network. This comparison enables us to determine how much of the reduction in travel time is attributable to facility expansions versus road network improvements.

5.3 Results

Using the instances of (1) described above, we evaluate the performance of the two policies for expanding access to cancer care in Rwanda. The processes described above result in instances consisting of a network with $|\mathcal{N}| = 3{,}597$ nodes, $|\mathcal{F}| = 48$ candidate facilities, $|\mathcal{O}| = 1{,}066$ population centers, $|\mathcal{S}| = 5$ surface types, $|\mathcal{A}| = 7{,}644$ type-dependent arcs, $T = 5$ time periods, and $|\Omega^\tau| = 2$ stochastic scenarios per period

$\tau \in \mathcal{T}$. We are able to solve the FLP to optimality across all test cases, and we are able to solve the FLNDP to within $4.6\% - 14.7\%$ of optimality for all test cases.

Table 5 presents the average travel time required to access tertiary cancer care under the FLNDP policy and the FLP policy in each of our test cases. These travel times are reported for the rainy season, dry season, and on average. With the current facilities and road network, the travel time to tertiary cancer care is approximately 3.5 hours on average. However, the current travel time to care is more than 6.5 hours in the rainy season. For each tested combination of monetary and staffing budget, the FLNDP policy results in lower average travel times, driven in large part by reductions in the travel times during the rainy season. For example, when the annual monetary budget is \$5 million and in all staffing budget levels, the FLNDP policy achieves an average travel time of 151.7 minutes compared to the FLP policy's average travel time of 168.2 minutes. This difference is due to a 11.2 minute difference during the dry season (104.6 minutes under the FLNDP policy compared to 116.0 minutes under the FLP policy) and a 28.6 minute difference in the rainy season (261.5 minutes under the FLNDP policy compared to 290.1 under the FLP policy). These trends persist as the monetary and staffing budgets increase, but are most pronounced for high monetary budgets and low staffing budgets. For example, for our highest monetary budget considered (\$25 million) and the lowest staffing budget of (2 oncologist teams per year), the FLNDP policy achieves an average travel time that is 24 minutes shorter than the FLP policy, and this improvement is in large part because the FLNDP policy has an average travel time that is 41.4 minutes shorter than the FLP policy during the rainy season.

We now present detailed results for a representative test case: test case 2. We select this test case to explore in detail because the monetary budget of \$5 million USD closely resembles historical trends in cancer care spending [45] and 4 new oncologist teams per year is an optimistic estimate of the impact that new training programs may have [49]. Figure 4 illustrates the optimal facility expansions over the planning horizon in this representative test case. In this test case, the FLNDP and FLP policies recommend the same facility expansions during the same time periods (as illustrated in the top of Figure 4). In Year 2, both policies recommend expansion of the University Teaching Hospital of Butare (located in the Southern Province) and, in Year 3, they both recommend expansion of the Kibuye Referral Hospital (located in the Western Province). Due to budget constraints, neither policy is able to recommend opening hospitals during Year 4. In Year 5, both policies recommend that the Musanze Referral Hospital (located in the Northern Province) be expanded. Under the FLNDP policy, many graveled minor highways in the Southern Province are improved to be paved highways during Year 1. Later, we observe that the FLNDP policy recommends more roads in the Eastern Province to be improved. Interestingly, we also observe that the FLNDP policy recommends the improvement of roads in the regions where future hospital expansions will occur. For example, the FLNDP policy recommends significant investment in the roads near Butare in Years 1 and 2, ahead of the facility expansion in Butare which is planned to occur in Year 3. We observed that the FLNDP policies favored improving roads that are minor highways rather than unclassified and residential streets. In the Eastern province road improvements made on the

Table 5: The average travel time required to access tertiary cancer care under the FLNDP policy and the FLP policy in each of our test cases, across a five-year planning horizon. Each test case varies in terms of monetary and staffing budget. The travel times are reported as a five-year average for the rainy season, the dry season, and overall average (in minutes).

Test Case	**Budget**		**Travel Time by Season (minutes)**					
	Monetary	**Human**	**FLNDP**			**FLP**		
	(mil. USD)	**Capital**	Dry	Rainy	Average	Dry	Rainy	Average
0*	-	-	162.8	407.0	210.7	162.8	407.0	236.0
1	5	2	104.6	261.5	151.7	116.0	290.1	168.2
2		4	104.6	261.5	151.7	116.0	290.1	168.2
3		6	104.6	261.5	151.7	116.0	290.1	168.2
4	10	2	101.7	254.3	147.5	116.0	290.1	168.2
5		4	81.3	203.1	117.8	90.5	226.3	131.3
6		6	81.3	203.2	117.8	89.9	226.3	130.4
7	15	2	100.8	252.0	146.2	116.0	290.1	168.2
8		4	79.8	199.4	115.7	90.7	226.8	131.5
9		6	69.1	172.8	100.2	76.6	191.6	111.1
10	20	2	100.1	250.2	145.1	116.0	290.1	168.2
11		4	79.2	197.9	114.8	90.8	227.1	131.7
12		6	66.1	165.1	95.8	73.9	184.8	107.2
13	25	2	99.5	248.7	144.2	116.0	290.1	168.2
14		4	78.3	195.6	113.5	90.5	226.3	131.3
15		6	66.1	165.3	95.9	75.3	188.3	109.2

*: Test case 0 represents the current state of access
FLNDP: Facility Location Network Design Problem
FLP: Facility Location Problem
USD: United States dollar

way to Kigali (small, central province) have a cascading impact on travel times, with patients located further from existing facilities having larger reductions in travel time, even when no facilities are built.

Figure 4 also illustrates the travel time required to access these facilities when road network improvements in the FLNDP policy are also made. Under both the FLNDP policy and the FLP policy, the southwest and northeast regions of the country experience the longest travel times in Year 1, before any facilities are expanded. However, we observe that under the FLNDP policy these regions have large reductions in travel times that are directly attributable to the improvements to the road network at this time. For example, some populations see their travel times reduced by close to 1 hour due to the improvements to the road network alone in Year 1. In the rural northeast of the country, neither the FLNDP nor FLP recommend a facility expansion in the Eastern Province. However, some people in this province see their rainy season travel times reduce by 90 minutes, which is attributable to road network improvements alone.

Figure 5 illustrates how the distribution of the population's travel time during the rainy seasons changes over the planning horizon in our representative test case under both the FLP and FLNDP policies. We observe that the FLNDP policy improves

access more quickly than the FLP policy. For instance, in Year 1, the FLP policy results in 20% of Rwandans traveling more than 10.8 hours to access care during the rainy season compared to 9.2 hours under the FLNDP policy. By Year 4, the FLNDP ensures that 80% of the population can access care in less than 5.7 hours during the rainy season. The FLP can achieve this level of access as well, but a year later, in Year 5.

Figure 6 illustrates the distribution of the reductions in travel time for each population center for each year of the planning horizon, relative to the current travel time for the respective population center. These distributions are shown for both the FLNDP policy and the FLP policy, stratified by quartile of current travel time. The Bottom 25% quartile is comprised of the quarter of Rwandans who currently travel the least to access cancer care, while the Top 25% quartile is comprised of the quarter of Rwandans who currently travel the farthest to access care. Figure 6 demonstrates that the FLNDP policy consistently provides higher reductions in travel time for each year and each quartile. In addition, we observe that Rwandans who currently have the worst access (the Top 25% quartile) receive the largest reductions in travel time under both the FLP and FLNDP policies. Importantly, the FLNDP policy provides this group with larger travel time reductions, and these reductions are observed earlier in the planning horizon. For example, the FLNDP policy reduces the travel time for this quartile by 36 minutes on average in Year 1, compared to no reduction in travel time under the FLP policy in Year 1. In Years 2 through 5, the FLNDP policy consistently provides this quartile with larger travel time reductions than the FLP policy.

In summary, the FLNDP policy results in meaningful improvements in access compared to the FLP policy, especially for Rwandans who currently travel the farthest to access tertiary cancer care.

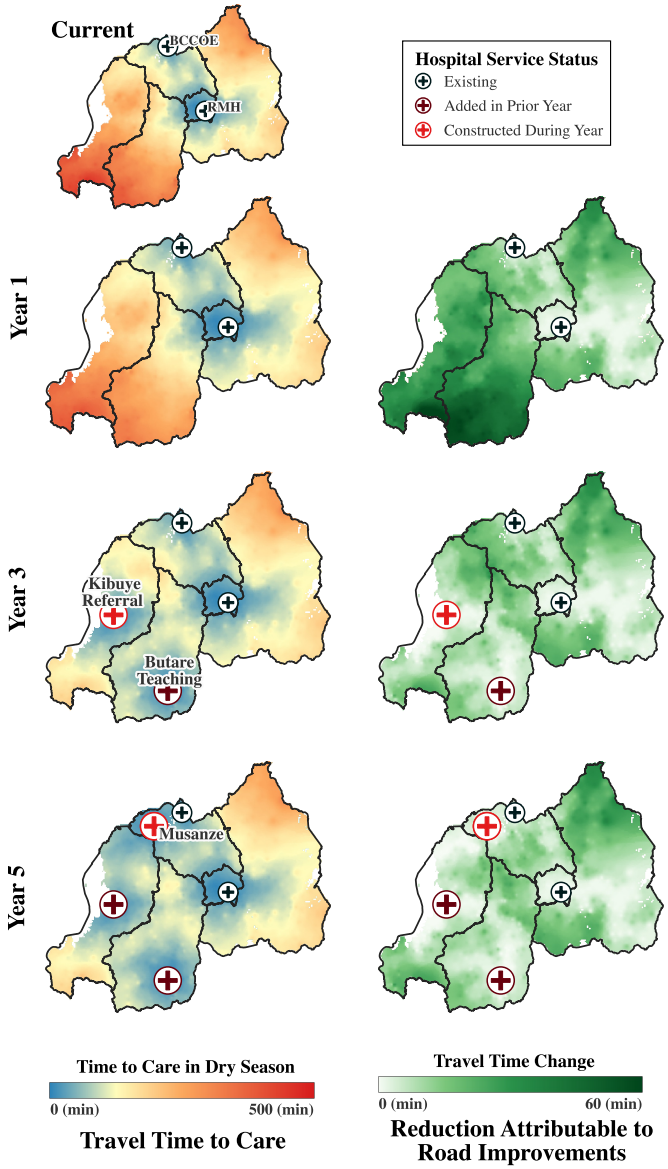

Fig. 4: An illustration of how the optimal expansion of facilities and road network improvements change travel times to access tertiary cancer care in Rwanda over the five-year planning horizon. Facilities are indicated with the + symbol. Existing facilities are colored black, facilities expanded during the corresponding year are colored red, and facilities expanding during a previous year are colored Maroon. (Left) The travel time required to access care by population center by year. (Right) The reduction in travel time that is directly attributable to road network improvements, by year.

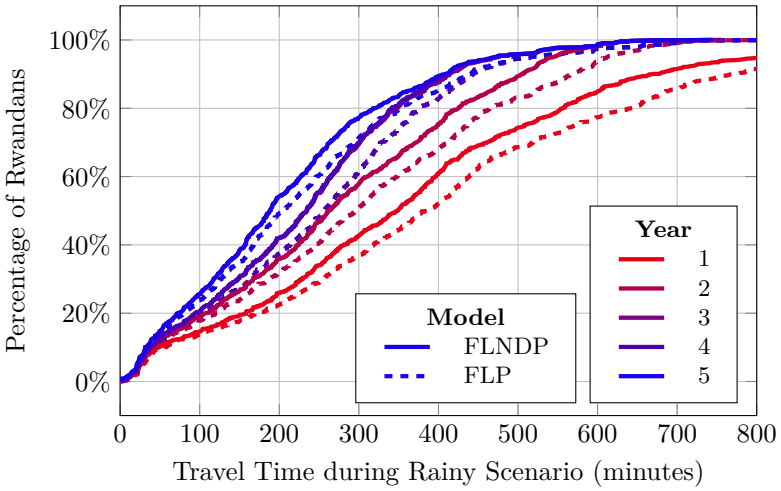

Fig. 5: The distribution of travel time (in minutes) across the population over the planning horizon under both FLP and FLNDP policies (in representative test case 2)

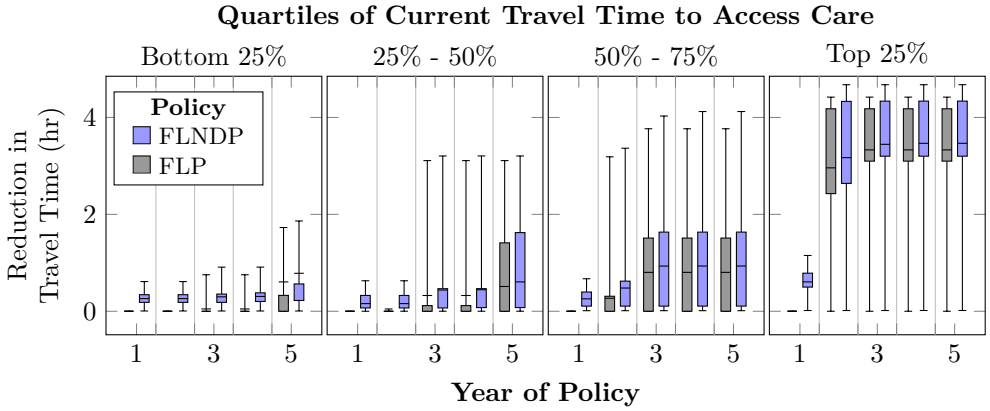

Fig. 6: The distribution of reduction in travel time from baseline over the planning horizon under each policy in our representative test case, stratified by quartile of current travel times during the dry season. The Bottom 25% quartile is comprised of the quarter of Rwandans who currently have the shortest travel times to access cancer care, while the Top 25% quartile is comprised of the quarter of Rwandans who have the longest travel times to access care.

6 Conclusions

In this article, we consider the problem of limited geographic access to cancer care in a low-income context and the extent to which improvements to transportation infrastructure could be used in concert with facility expansions to improve access.

To address this challenge, we formulate a multi-period Facility Location Network Design Problem (FLNDP) with stochastic demand and travel times. During each period in our model, a central planner decides which facilities to expand and which roads to improve. Then patients will select the shortest path to a facility offering cancer services given the road conditions due to weather. This formulation requires many fewer constraints than existing formulations in the literature. However, because there are a large number of possible paths, the resulting formulation requires a huge number of decision variables. To handle this, we propose a branch-price-and-cut (BP&C) approach.

Within the BP&C algorithm, we use column generation to solve the linear programming relaxation at each node. In this step, we start by considering only a subset of path variables and iteratively generate new paths as needed by solving a pricing problem. We leverage the problem's structure to efficiently solve the pricing problems by expressing them as shortest-path problems. Then we tailor several modifications to accelerate our method, including heuristics and special ordered set branching rules that take advantage of the sequential decision periods. We also derive decomposition-based cuts that can be added at integer feasible solutions in the branch-and-bound tree. In a computational study, we compare our complete BP&C algorithm with a branch-and-price (B&P) algorithm that uses our proposed column generation approach, but does not include any cuts. We find that the BP&C is able to solve more test instances than the B&P approach, suggesting that our proposed cuts provide meaningful improvements to the algorithm. We also find that our proposed model is most challenging to solve in instances with relatively tight monetary budgets.

In our case study, we explore geographic access to cancer care in Rwanda and compare two different approaches for expanding care: an FLP policy that only considers expansions to the healthcare infrastructure and an FLNDP policy that considers improvements to both the existing healthcare and the transportation infrastructure. The FLNDP consistently outperforms the FLP policy by further reducing travel times to access care due to its proposed road network improvements. In a representative test case, we observe that reductions in travel time that are directly attributable to road network improvements can be as high as 1 hour for some population centers. Moreover, we see that the patients who have the largest reductions in travel time due to road network improvements are the patients who currently travel the farthest to access care.

There are open opportunities for future work that builds on this work. Future work could further improve our BP&C approach. For instance, such work could strengthen our cuts via mixed-integer rounding, consider alternative branching schemes, or further improve the described cuts by considering the dual variables associated with existing cuts. Another opportunity for future work is to apply this modeling framework to expand geographic access to other diseases, such as HIV or diabetes, which

require diagnosis and treatment. Yet another opportunity for future work is the consideration of how increased geographic access to diagnosis could improve the rates of early diagnosis, which in turn may influence the demand for treatment. Finally, in our case study, we used a simple model of weather events. However, future work could use more sophisticated forecasting techniques to quantify the risks and impacts of various weather events.

In summary, this article introduced a new method for expanding both healthcare and transportation infrastructure as a means to improve geographic access to cancer care. Our work demonstrates that improvements to the road network could meaningfully improve geographic access to cancer care, especially among cancer patients who are currently the most underserved. Our analysis provide valuable managerial insights to healthcare planners seeking to improve care access in low-income countries.

Acknowledgments and Funding

This work was supported by the Watson Foundation, National Defense Science and Engineering Graduate Fellowship program, the George Family Fellowship at the Georgia Institute of Technology, and the Nash Endowment to the Georgia Institute of Technology. A.S. was supported by the Watson Foundation, is a George Family Fellow, and is a National Defense Science and Engineering Graduate Fellow through the Office of Naval Research. Additionally, this work would not have been possible without the support of the Institute of Global Health Equity Research and the University of Global Health Equity, who graciously hosted A.S. in the fall of 2022. L.N.S was supported by the Harold R. and Mary Anne Nash Endowment Fund from the H. Milton Stewart School of Industrial and Systems Engineering.

References

[1] Brand, N.R., Qu, L.G., Chao, A., Ilbawi, A.M.: Delays and barriers to cancer care in low-and middle-income countries: a systematic review. The oncologist **24**(12), 1371–1380 (2019)

[2] Bray, F., Laversanne, M., Sung, H., Ferlay, J., Siegel, R.L., Soerjomataram, I., Jemal, A.: Global cancer statistics 2022: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. CA: A Cancer Journal for Clinicians **74**(3), 229–263 (2024) https://doi.org/10.3322/caac.21834

[3] Bizuayehu, H.M., Ahmed, K.Y., Kibret, G.D., Dadi, A.F., Belachew, S.A., Bagade, T., Tegegne, T.K., Venchiarutti, R.L., Kibret, K.T., Hailegebireal, A.H., Assefa, Y., Khan, M.N., Abajobir, A., Alene, K.A., Mengesha, Z., Erku, D., Enquobahrie, D.A., Minas, T.Z., Misgan, E., Ross, A.G.: Global disparities of cancer and its projected burden in 2050. JAMA Network Open **7**(11), 2443198 (2024) https://doi.org/10.1001/jamanetworkopen.2024.43198

[4] Unger-Saldaña, K.: Challenges to the early diagnosis and treatment of breast cancer in developing countries. World journal of clinical oncology **5**(3), 465 (2014)

[5] Stefan, D.C., Tang, S.: Addressing cancer care in low-to middle-income countries: a call for sustainable innovations and impactful research. BMC cancer **23**(1), 756 (2023)

[6] Knapp, G.C., Tansley, G., Olasehinde, O., Wuraola, F., Adisa, A., Arowolo, O., Olawole, M.O., Romanoff, A.M., Quan, M.L., Bouchard-Fortier, A., Alatise, O.I., Kingham, T.P.: Geospatial access predicts cancer stage at presentation and outcomes for patients with breast cancer in southwest nigeria: A population-based study. Cancer **127**(9), 1432–1438 (2021) https://doi.org/10.1002/cncr.33394

[7] Fadelu, T., Nadella, P., Iyer, H.S., Uwikindi, F., Shyirambere, C., Manirakiza, A., Triedman, S.A., Rebbeck, T.R., Shulman, L.N.: Toward Equitable Access to Tertiary Cancer Care in Rwanda: A Geospatial Analysis. JCO Global Oncology (8), 2100395 (2022) https://doi.org/10.1200/GO.21.00395

[8] Sapirstein, A., Steimle, L.N., Stefan, D.C.: Toward expanded access to cancer care with cost awareness: An optimization modeling analysis of rwanda. JCO Global Oncology (10), 2400022 (2024) https://doi.org/10.1200/GO.24.00022

[9] Dewi, S.P., Kasim, R., Sutarsa, I.N., Dykgraaf, S.H.: A scoping review of the impact of extreme weather events on health outcomes and healthcare utilization in rural and remote areas. BMC Health Services Research **24**, 1333 (2024) https://doi.org/10.1186/s12913-024-11695-5

[10] Mikou, M., Rozenberg, J., Koks, E., Fox, C., Quiros, T.P.: Assessing rural accessibility and rural roads investment needs using open source data (2019)

[11] Pace, L.E., Mpunga, T., Hategekimana, V., Dusengimana, J.-M.V., Habineza, H., Bigirimana, J.B., Mutumbira, C., Mpanumusingo, E., Ngiruwera, J.P., Tapela, N., Amoroso, C., Shulman, L.N., Keating, N.L.: Delays in breast cancer presentation and diagnosis at two rural cancer referral centers in rwanda. The Oncologist **20**(7), 780–788 (2015) https://doi.org/10.1634/theoncologist.2014-0493

[12] Nsanzimana, S., NgamiJe, D.: Rwandan National Cancer Control Plan. Non-Communicable Diseases Division, Rwanda Biomedical Centre, Ministry of Health, Rwanda (2020)

[13] Shulman, L.N., Mpunga, T., Tapela, N., Wagner, C.M., Fadelu, T., Binagwaho, A.: Bringing cancer care to the poor: Experiences from Rwanda. Nature Reviews Cancer **14**(12), 815–821 (2014) https://doi.org/10.1038/nrc3848

[14] Manirakiza, A.V.C., Rubagumya, F., Mushonga, M., Mutebi, M., Lasebikan, N., Kochbati, L., Gwayali, B., Booth, C.M., Stefan, D.C.: The current status of national cancer control plans in africa: Data from 32 countries. Journal of Cancer Policy **37**, 100430 (2023) https://doi.org/10.1016/j.jcpo.2023.100430

[15] OpenStreetMap contributors: Planet dump retrieved from

https://planet.osm.org . https://www.openstreetmap.org (2017)

[16] Scott, A.A., Polo, A., Zubizarreta, E., Akoto-Aidoo, C., Edusa, C., Osei-Bonsu, E., Yarney, J., Dwobeng, B., Milosevic, M., Rodin, D.: Geographic accessibility and availability of radiotherapy in ghana. JAMA Network Open **5**(8), 2226319 (2022) https://doi.org/10.1001/jamanetworkopen.2022.26319

[17] Campos, M.V.A., Assis, R.d.S.L., Souza, M.J.F., Siqueira, E.C., Silva, M.A.L., Souza, S.R.: Multi-objective mammography unit location–allocation problem: A case study. Operations Research for Health Care **41**, 100430 (2024)

[18] Daskin, M., Hurter, A., Van Buer, M.: Toward an integrated model of facility location and transportation network design. the transportation center. Technical report, Northwestern University, Working Paper. Evanston, IL, USA (1993)

[19] Melkote, S., Daskin, M.S.: An integrated model of facility location and transportation network design. Transportation Research Part A: Policy and Practice **35**(6), 515–538 (2001) https://doi.org/10.1016/S0965-8564(00)00005-7

[20] Contreras, I., Fernández, E., Reinelt, G.: Minimizing the maximum travel time in a combined model of facility location and network design. Omega **40**(6), 847–860 (2012) https://doi.org/10.1016/j.omega.2012.01.006

[21] Drezner, Z., Wesolowsky, G.O.: Network design: Selection and design of links and facility location. Transportation Research Part A: Policy and Practice **37**(3), 241–256 (2003) https://doi.org/10.1016/S0965-8564(02)00014-9

[22] Ghaderi, A.: Heuristic Algorithms for Solving an Integrated Dynamic Center Facility Location - Network Design Model. Netw Spat Econ **15**(1), 43–69 (2015) https://doi.org/10.1007/s11067-014-9269-z

[23] Pearce, R.H., Forbes, M.: Disaggregated Benders decomposition and branch-and-cut for solving the budget-constrained dynamic uncapacitated facility location and network design problem. European Journal of Operational Research **270**(1), 78–88 (2018) https://doi.org/10.1016/j.ejor.2018.03.021

[24] Cocking, C., Flessa, S., Reinelt, G.: Improving access to health facilities in nouna district, burkina faso. Socio-Economic Planning Sciences **46**(22), 164–172 (2012) https://doi.org/10.1016/j.seps.2011.12.004

[25] Ghaderi, A., Jabalameli, M.S.: Modeling the budget-constrained dynamic uncapacitated facility location–network design problem and solving it via two efficient heuristics: A case study of health care. Mathematical and Computer Modelling **57**(3), 382–400 (2013) https://doi.org/10.1016/j.mcm.2012.06.017

[26] Shishebori, D., Jabalameli, M.S., Jabbarzadeh, A.: Facility Location-Network Design Problem: Reliability and Investment Budget Constraint. Journal of Urban

Planning and Development **140**(3), 04014005 (2014) https://doi.org/10.1061/(ASCE)UP.1943-5444.0000187

[27] Pourrezaie-Khaligh, P., Bozorgi-Amiri, A., Yousefi-Babadi, A., Moon, I.: Fix-and-optimize approach for a healthcare facility location/network design problem considering equity and accessibility: A case study. Applied Mathematical Modelling **102**, 243–267 (2022) https://doi.org/10.1016/j.apm.2021.09.022

[28] Magnanti, T.L., Mireault, P., Wong, R.T.: Tailoring Benders decomposition for uncapacitated network design. In: Gallo, G., Sandi, C. (eds.) Netflow at Pisa. Mathematical Programming Studies, pp. 112–154. Springer, Berlin, Heidelberg (1986). https://doi.org/10.1007/BFb0121090

[29] Liu, X., Kwon, C.: Exact robust solutions for the combined facility location and network design problem in hazardous materials transportation. IISE Transactions **52**(1010), 1156–1172 (2020) https://doi.org/10.1080/24725854.2019.1697017

[30] Bigotte, J.F., Krass, D., Antunes, A.P., Berman, O.: Integrated modeling of urban hierarchy and transportation network planning. Transportation Research Part A: Policy and Practice **44**(7), 506–522 (2010)

[31] Barnhart, C., Johnson, E.L., Nemhauser, G.L., Savelsbergh, M.W.P., Vance, P.H.: Branch-and-Price: Column Generation for Solving Huge Integer Programs. Operations Research **46**(3), 316–329 (1998) https://doi.org/10.1287/opre.46.3.316

[32] Li, X., Aneja, Y., Huo, J.: Using branch-and-price approach to solve the directed network design problem with relays. Omega **40**(5), 672–679 (2012)

[33] Kulkarni, O., Dahan, M., Montreuil, B.: Resilient Relay Logistics Network Design: A k-Shortest Path Approach. *Preprint* (2023)

[34] Greening, L.M., Dahan, M., Erera, A.L.: Lead-time-constrained middle-mile consolidation network design with fixed origins and destinations. Transportation Research Part B: Methodological **174**, 102782 (2023) https://doi.org/10.1016/j.trb.2023.102782

[35] Rothenbächer, A.-K., Drexl, M., Irnich, S.: Branch-and-price-and-cut for the truck-and-trailer routing problem with time windows. Transportation Science **52**(5), 1174–1190 (2018) https://doi.org/10.1287/trsc.2017.0765

[36] Contreras, I., Fernández, E., Reinelt, G.: Minimizing the maximum travel time in a combined model of facility location and network design. Omega **40**(6), 847–860 (2012) https://doi.org/10.1016/j.omega.2012.01.006

[37] Bhangdia, K., Natarajan, A., Rudolfson, N., Verguet, S., Castro, M.C., Dusengimana, J.-M.V., Shyirambere, C., Schleimer, L.E., Shulman, L.N., Umwizerwa,

A., Kigonya, C., Butonzi, J., MacDuffie, E., Fadelu, T., O'Neil, D.S., Nguyen, C., Mpunga, T., Keating, N.L., Pace, L.E.: The association of travel distance and other patient characteristics with breast cancer stage at diagnosis and treatment completion at a rural rwandan cancer facility. BMC Cancer **25**(1), 146 (2025) https://doi.org/10.1186/s12885-025-13489-2

[38] Rubagumya, F., Wilson, B., Manirakiza, A., Mutabazi, E., A Ndoli, D., Rudakemwa, E., Chamberlin, M.D., Hopman, W.M., Booth, C.M.: Financial toxicity: Unveiling the burden of cancer care on patients in rwanda. The Oncologist **29**(3), 345–350 (2024) https://doi.org/10.1093/oncolo/oyad291

[39] Bank, W.: Climate Change Knowledge Portal (N.D.). https://climateknowledgeportal.worldbank.org/

[40] Nkurunziza, M.: Rwanda: Two killed, dozens of houses damaged by floods, landslides. The New Times (2025)

[41] Idukunda, C., Michellier, C., De Longueville, F., Twarabamenye, E., Henry, S.: Assessing community vulnerability to landslide and flood in northwestern rwanda. International Journal of Disaster Risk Reduction **123**, 105329 (2025) https://doi.org/10.1016/j.ijdrr.2025.105329

[42] Li, L., Mind'je, R.: In: Li, L., Mind'je, R. (eds.) Conclusion, pp. 103–104. Springer, Singapore (2023). https://doi.org/10.1007/978-981-99-1751-8_9

[43] Health, R.M.: Health Facilities. https://rwanda.africageoportal.com/ (2022)

[44] Sibomana, S., Ngendahimana, P.: Unit Price Study. Local Administrative Entities Development Agency, P.O. Box: 7305 Kigali, Rwanda (2021). https://www.loda.gov.rw/index.php?eID=dumpFile&t=f&f=39777&token=01f9d91a190dec25fa5ac55aed2bd2ebd7844806

[45] Neal, C., Rusangwa, C., Borg, R., Tapela, N., Mugunga, J.C., Pritchett, N., Shyirambere, C., Ntakirutimana, E., Park, P.H., Shulman, L.N., Mpunga, T.: Cost of providing quality cancer care at the butaro cancer center of excellence in rwanda. Journal of Global Oncology (4), 1–7 (2018) https://doi.org/10.1200/JGO.17.00003

[46] Ray, N., Ebener, S.: Accessmod 3.0: computing geographic coverage and accessibility to health care services using anisotropic movement of patients. International Journal of Health Geographics **7**(1), 63 (2008) https://doi.org/10.1186/1476-072X-7-63

[47] Blanford, J.I., Kumar, S., Luo, W., MacEachren, A.M.: It's a long, long walk: accessibility to hospitals, maternity and integrated health centers in niger. International Journal of Health Geographics **11**(1), 24 (2012) https://doi.org/10.1186/1476-072X-11-24

[48] Ouma, P.O., Maina, J., Thuranira, P.N., Macharia, P.M., Alegana, V.A., English, M., Okiro, E.A., Snow, R.W.: Access to emergency hospital care provided by the public sector in sub-saharan africa in 2015: a geocoded inventory and spatial analysis. The Lancet Global Health **6**(3), 342–350 (2018) https://doi.org/10.1016/S2214-109X(17)30488-6

[49] Manirakiza, A., Rubagumya, F., Fehr, A.E., Triedman, A.S., Greenberg, L., Mbabazi, G., Ntacyabukura, B., Nyagabona, S., Maniragaba, T., Longombe, A.N., Ndoli, D.A., Makori, K., Kiugha, M., Rulisa, S., Hammad, N.: Oncology training in rwanda: Challenges and opportunities for undergraduate medical students (the educan project). Journal of Cancer Education **35**(2), 359–365 (2020) https://doi.org/10.1007/s13187-019-1473-6

Appendix A Additional Details on Branch-Price-and-Cut

A.1 Pricing Problem Example

In this section, we provide an example on how to generate the graph $\mathcal{G}_o^{\tau\omega}$ used in the pricing problems corresponding to population center o in time period τ and scenario ω (as described in §3).

Figure A1 provides an example of a small graph $\mathcal{G}$ with nodes $\mathcal{N} = \{o_1, o_2, j, f_1, f_2\}$, population centers $\mathcal{O} = \{o_1, o_2\}$, and candidate facilities $\mathcal{F} = \{f_1, f_2\}$. The corresponding travel time along the arc between node i and node k of surface type s for this particular time period τ and scenario ω is indicated on the corresponding type-dependent arc $a = (i, k, s)$ as $\ell_a^{\tau,\omega} = \ell_{(i,k,s)}^{\tau,\omega}$. Recall that $d_o^{\tau\omega}$ is the demand from population center o during time period τ and scenario ω, and $p^{\tau\omega}$ is the probability of scenario ω in time period τ.

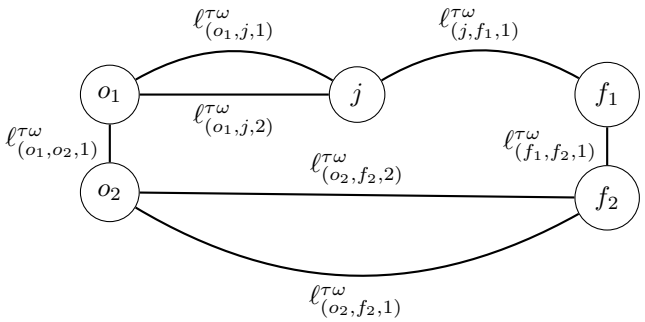

Fig. A1: An example of a graph $\mathcal{G}$ with nodes $\mathcal{N} = \{o_1, o_2, j, f_1, f_2\}$, population centers $\mathcal{O} = \{o_1, o_2\}$, candidate facilities $\mathcal{F} = \{f_1, f_2\}$. The type-dependent arcs $\mathcal{A}$ are shown with their corresponding travel times.

We now demonstrate how to construct the graph $\mathcal{G}_o^{\tau\omega}$ that is used to solve the shortest path problem corresponding to each population center o in time period τ and scenario ω. First, we create $\mathcal{G}_o^{\tau\omega}$ as a copy of $\mathcal{G}$ described above. We augment $\mathcal{G}_o^{\tau\omega}$ by adding a dummy sink node $\bar{f}$ and connecting it with each facility f. Then, we assign the length of each type-dependent arc in graph $\mathcal{G}_o^{\tau\omega}$ according to (5):

$$l_a = \begin{cases} p^{\tau\omega} d_o^{\tau\omega} \ell_a^{\tau\omega} - \hat{\beta}_{oa}^{\tau\omega} & \text{if } a \in \mathcal{A} \\ -\hat{\alpha}_o^{\tau\omega} - \hat{\gamma}_{of}^{\tau\omega} & \text{if } a = (f, \bar{f}) \text{ for some } f \in \mathcal{F}, \end{cases}$$

where $\hat{\alpha}, \hat{\beta}$, and $\hat{\gamma}$ are obtained from solving the restricted master problem.

Figure A2 provides an example of such a graph for two different populations, o_1 and o_2, during time period τ and scenario ω. Figure A2a illustrates the graph to be

used in the shortest path problem for population o_1. Figure A2b illustrates how a similar process would be used to generate the graph corresponding to population o_2.

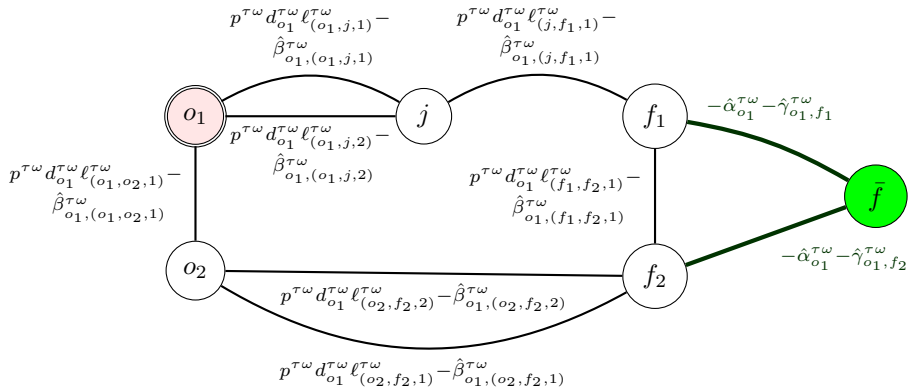

(a) Graph $\mathcal{G}^{\tau\omega}_{o_1}$ used to solve the pricing problem from o_1 during τ and ω

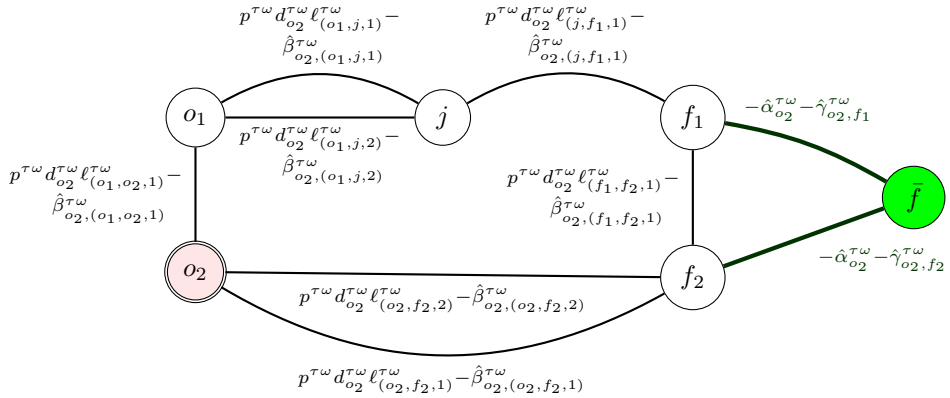

(b) Graph $\mathcal{G}^{\tau\omega}_{o_2}$ used to solve the pricing problem from o_2 during τ and ω

Fig. A2: Examples of graphs to be used to solve the shortest-path problems corresponding to time period τ and scenario ω for two different population centers o, as described in §3.2. Figure A2a (resp. Figure A2b) illustrates the graph to be used to solve the shortest path problem for $o = o_1$ (resp. $o = o_2$). In each graph, $\hat{\alpha}, \hat{\beta}$, and $\hat{\gamma}$ denote optimal dual values of the restricted master problem.

A.2 Decomposition-Based Cuts

A.2.1 Non-Dominated Cuts

Here, we consider that some of the cuts of the form (8) may be dominated by previously generated cuts. Therefore, we provide an algorithm to generate a set of non-dominated cuts to be added to the master problem solved at each tree node. If the node which we are currently processing with column generation yields an integer solution, from which we can construct the cuts (8), we want to ensure that these cuts are not dominated by any pre-existing cuts to avoid slowing down the branch-price-and-cut algorithm. Thus, we compare the new cut, associated with $(\alpha_o^{\tau\omega}, (\beta_{oa}^{\tau\omega})_{a\in\mathcal{A}}, (\gamma_{of}^{\tau\omega})_{f\in\mathcal{F}})$ with prior cuts associated with $(\alpha_o^{\tau\omega'}, (\beta_{oa}^{\tau\omega'})_{a\in\mathcal{A}}, (\gamma_{of}^{\tau\omega'})_{f\in\mathcal{F}})$. The new cut is *dominated* by the old cut if

$$\min\{\alpha_o^{\tau\omega'} - \alpha_o^{\tau\omega}, \min_{a\in\mathcal{A}}\{\beta_{oa}^{\tau\omega'} - \beta_{oa}^{\tau\omega}\}, \min_{f\in\mathcal{F}}\{\gamma_{of}^{\tau\omega'} - \gamma_{of}^{\tau\omega}\}\} \geq 0.$$

If the new cut is not dominated by any of the existing cuts, then we add it to the cut set $\mathcal{C}$.

A.2.2 Generating Cuts from Heuristic Integer Solutions

We now describe our process for generating valid inequalities (8) when we only have access to integer-feasible first-stage solutions $(\hat{u}, \hat{v}, \hat{x}, \hat{y})$. This would permit us to generate cuts from the heuristic integer solution in §3.4.1 without having access to dual variables at optimality of the master problem. Recall that given binary first-stage variables, all second-stage decisions result from solving shortest-path problems. Thus, this problem bears structural similarity to that of Magnanti et al. [28] and Pearce and Forbes [23], which we now extend to consider our specific FLNDP in which there are multiple surface types, $\mathcal{S}$. The dual variables have the following natural interpretations for a fixed τ and fixed ω. $\gamma_{of}^{\tau\omega}$ is the change in travel time during ω for o if f were to have been constructed at the start of τ. Similarly, $\beta_{o(i,k,s)}^{\tau\omega}$ denotes the change in travel time during ω for o if (i,k,s) were to have been constructed at the start of τ. Lastly, $\alpha_o^{\tau\omega}$ is the travel time currently incurred by o in τ under ω.

Since $\hat{u}, \hat{v}, \hat{x}, \hat{y}$ are integer, we create a subgraph $\overline{\mathcal{G}}^{\tau\omega} := (\mathcal{N}, \overline{\mathcal{A}}^{\tau\omega})$ for each $\tau \in \mathcal{T}$ and $\omega \in \Omega$ by including only the arcs that are available at the start of period τ, namely, $\overline{\mathcal{A}}^{\tau\omega} := \{(i,k,s) \in \mathcal{A} \text{ s.t. } \hat{y}_{ijs}^{\tau} = 1\}$. Then, for $o \in \mathcal{O}$, we compute the travel time to all other nodes in $\overline{\mathcal{G}}^{\tau\omega}$ using $\ell^{\tau\omega}$ as arc lengths. Following the convention of Pearce and Forbes [23], let D_{on} denote the shortest travel time to any node $n \in \mathcal{N}$ from o. Then $\alpha_o^{\tau\omega} = \min\{D_{of},\ \forall\ f \in \mathcal{F} \text{ s.t. } \hat{x}_f^{\tau} = 1\}$. Let $\Delta_n := \min\{0, D_{on} - \alpha_o^{\tau\omega}\}$. Note that Δ_n is equal to 0 when n is located further than the closest existing facility and is negative when n is closer than such a facility. Thus we can define the remaining dual variables based on their intuited definitions:

$$\gamma_{of}^{\tau\omega} = \Delta_f,\ \forall\ f \in \mathcal{F}$$

$$\beta^{\tau\omega}_{o(i,k,s)} = \min\{0, \underbrace{\min\{\Delta_i, \Delta_k\}}_{\text{close endpoint}} - \underbrace{\max\{\Delta_i, \Delta_k\}}_{\text{far endpoint}} + \ell^{\tau\omega}_{(i,k,s)}\}$$

Briefly, $\beta^{\tau\omega}_{o(i,k,s)}$ is obtained by considering the difference in Δ between the endpoint of (i,k,s) located closest to o and the far endpoint of (i,k,s). This difference expresses the reduction in travel time if no time was required to traverse (i,k,s). To this difference, we add the travel time $\ell^{\tau\omega}_{ijs}$ required to traverse. Without altering $\overline{\mathcal{G}}^{\tau\omega}$, we can repeat the process for each $o \in \mathcal{O}$. To solve for all dual variables, we must construct a new $\overline{\mathcal{G}}^{\tau\omega}$ for each $\tau \in \mathcal{T}$ and $\omega \in \Omega$. We use these dual variables to construct cuts of the form (8).

Appendix B Additional Details of Computational Experiments

In this appendix, we provide more details about the randomly-generated instances we used to test our solution algorithms. We generate the nodes $\mathcal{N}$ and the candidate facilities as described in §4. We also define the time periods $\mathcal{T}$ and scenarios per time period Ω^τ for each τ. Now, we provide more details on how each instance is generated given $\mathcal{N}, \mathcal{F}, \mathcal{T}$, and Ω.

Road Network Construction

First, we describe how we constructed the road networks available at the beginning of the planning horizon. Our first step is to calculate the Euclidean distance between each pair of nodes in the network. Then, we select a random node in the network, $i \in \mathcal{N}$. For this node i, we rank the other nodes in the network according to their distance from i. Therefore, for node i, we can label the other nodes such that i_1 is the closest node to i and $i_{|\mathcal{N}|-1}$ is the farthest node from i. Finally, we randomly generate an arc between i and i_n with probability $\frac{1}{n}$ for all $|\mathcal{N}|$. We then ignore node i when considering the generation of future arcs and select a new node i'. We repeat this process where we now rank the $|\mathcal{N}|-1$ nodes in $\mathcal{N} \setminus \{i\}$ and generate arcs between the remaining nodes with the probability described above (i.e., the arc between i' and i'_n is generated with probability $\frac{1}{n}$). This process guarantees we will construct a connected graph. In addition, this process generates graphs in which nodes that are closer to each other are more likely to have an arc between them, which is characteristic of road networks. Figure B3 illustrates an example of a randomly generated graph using this process.

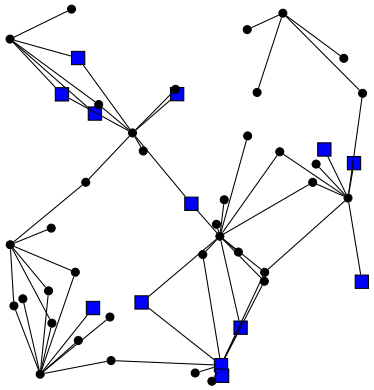

Fig. B3: An example instance of a randomly constructed instance. Candidate facilities are shown as blue rectangles, with population centers shown as black circles. In this instance $|\mathcal{N}| = 50$ and $|\mathcal{F}| = 12$.

Next, we define the type of surface for each arc (i, j). We consider $|\mathcal{S}| = 5$ different surface types. The probability that an arc (i, j) will take on surface type s is given a

truncated, shifted geometric distribution:

$$p(\text{surface}(i,j)=s) = \begin{cases} 1 - \frac{1}{|\mathcal{S}|}\sum_{s'=2}^{|\mathcal{S}|}\left(1-\frac{1}{|\mathcal{S}|}\right)^{s'-1}, & \text{if } s=1, \\ \frac{1}{|\mathcal{S}|}\left(1-\frac{1}{|\mathcal{S}|}\right)^{s-1}, & \text{if } s \in \{2,\ldots,|\mathcal{S}|\}. \end{cases}$$

At this point, we have now generated the nodes $\mathcal{N}$, candidate facilities $\mathcal{F}$, time periods $\mathcal{T}$, scenarios for each time period Ω^τ, and the existing road network. We now describe how we generate the remainder of the parameters for each instance.

Monetary and Human Capital Cost Parameters

We assume a uniform monetary cost of constructing arcs $c_{ijs}^{\tau\omega}$ and adding services c_o^τ across time periods $\tau \in \mathcal{T}$. The cost associated with constructing arc (i,j,s) is a function of $\|i-j\|_2$, the Euclidean distance from i to j, the surface condition s, as well as the already existing surface condition s', the existing surface connecting i and j. We set that

$$c_{ijs}^{\tau\omega} = \begin{cases} \|i-j\|_2 \cdot \frac{1}{10}s, & s > s', \\ 0, & s \leq s'. \end{cases}$$

The cost associated with a single facility c_o^1 was distributed according to a truncated normal distribution centered at 75 with a mean $\mu = 5$ truncated to the interval $(50, 100)$. These costs are assumed to be consistent across the horizon, i.e., $c_o^\tau = c_o^1$, $\forall \tau \in \mathcal{T}$. Staffing costs for each hospital were also assumed to be consistent across time periods, with $h_o^1 \sim \text{Unif}(2,6)$ and $h_o^\tau = h_o^1$, $\forall \tau \in \mathcal{T}$.

Demand Parameters

We consider that all nodes are population centers. The demand for a fixed $o \in \mathcal{O}$ was constant across $\omega \in \Omega_\tau$ but varied across time periods $\tau \in \mathcal{T}$. Demand during the first time period $d_o^{1\omega} \sim \text{Unif}(50,100)$. Demand in all other time periods, i.e., $\tau \neq 1$, is generated such that $d_o^{\tau\omega} \sim d_o^{(\tau-1)\omega} + \text{Unif}(25,50)$.

Travel Time Parameters

We specify the travel time $\ell_{ijs}^{\tau\omega}$ for type-dependent arc $(i,j,s) \in \mathcal{A}$ in each of the two scenarios, $\omega^\tau \in \{1,2\}$, to be the same in all time periods $\tau \in \mathcal{T}$. We specify $\ell_{ijs}^{\tau\omega}$ as follows:

$$\ell_{ijs}^{\tau\omega} = \frac{\|i-j\|_2}{s+1} \cdot \left(1 + \frac{1}{2}\mathbb{1}\{\omega = 2\}\right)$$

where $\|i-j\|_2$ is the ℓ_2 norm between the arcs endpoints and $\mathbb{1}$ is an indicator function.