Optimal Ambulance Location Using Integer Programming

A Case Study of Athens, Ohio and Surrounding Townships

Fiona Adda Adonnawura

Supervisor: Vardges Melkonian

March 2026

Contents

1 Introduction

Emergency medical services (EMS) represent one of the most time-critical components of public health infrastructure. When a cardiac arrest, serious accident, or medical emergency occurs, the speed at which an ambulance reaches the scene can mean the difference between life and death. Yet in many communities — particularly those that blend dense urban centers with surrounding rural townships — ensuring fast, equitable ambulance coverage is a persistent challenge. Urban zones may be well-served while rural areas suffer response times that far exceed safe clinical thresholds, leaving residents at significantly elevated risk.

This project addresses the problem of **optimal ambulance depot location** in such a mixed urban-rural setting, using the city of **Athens, Ohio** and its surrounding townships as a case study. Athens presents an especially interesting context: it is home to Ohio University, creating unusually high population density in its urban core, while the surrounding townships of Alexander, Dover, Troy, and Rome are largely rural with dispersed populations. Designing an EMS system that serves both populations efficiently — under real resource constraints such as a fixed number of ambulances and limited budget — is a genuine operational challenge.

The project formulates this as a **Binary Integer Programming (BIP)** problem, combining two classical facility location frameworks: the *Maximal Covering Location Problem* (MCLP), which maximizes the population covered within a given response time, and a cost-minimization model that accounts for depot opening costs. The combined model is implemented in AMPL and solved to proven optimality using CPLEX 22.1.1.

2 Background of Study

2.1 Emergency Medical Services in Rural and Mixed Areas

In the United States, EMS systems are responsible for providing prehospital emergency care and transport. Response time is the single most critical performance metric: clinical evidence consistently shows that survival rates for out-of-hospital cardiac arrest decline sharply with every minute of delay without defibrillation. The National Fire Protection Association's Standard 1710 [6] sets a benchmark of 8 minutes or less for 90% of emergency responses in urban areas, a threshold widely adopted in EMS planning research and practice.

However, meeting this benchmark in rural and mixed urban-rural areas is far more difficult than in dense cities. Rural EMS systems face three structural challenges that urban systems do not:

1. **Geographic dispersion.** Demand zones are spread across large areas, meaning a single depot cannot cover multiple zones within 8 minutes regardless of how centrally it is placed.

2. **Road network constraints.** Rural roads are often indirect, following natural features such as rivers and hills rather than a grid. In hilly Appalachian terrain like Athens County, actual road distances are typically 25–35% longer than straight-line distances, significantly inflating travel times.

3. **Resource scarcity.** Rural EMS systems typically operate with smaller budgets and fewer ambulances than urban systems, forcing difficult trade-offs between coverage breadth and depth.

Athens County exemplifies all three challenges. The Hocking River valley constrains road routing, the Appalachian hills increase travel times, and the county's mixed population, dense near Ohio University, sparse in the outer townships, creates uneven demand that is difficult to serve with a small number of fixed depots.

2.2 Facility Location Optimization

The problem of deciding where to locate emergency service facilities has been studied in the operations research literature since the 1970s. The two most influential models are:

The Location Set Covering Problem (LSCP), introduced by Toregas et al. (1971) [2], minimizes the number of facilities needed to ensure every demand zone is covered within a given response time threshold. It guarantees full coverage but does not account for resource limits.

The Maximal Covering Location Problem (MCLP), introduced by Church and ReVelle (1974) [1], takes the opposite approach: given a fixed number of facilities, it maximizes the total population covered within the threshold. This is more realistic when resources are constrained, as it allows the model to prioritize high-demand zones when full coverage is impossible.

Both models have been extensively applied to ambulance placement, fire station siting, and other emergency facility location problems. More recent extensions incorporate station capacity, multiple coverage levels, demand uncertainty, and cost minimization — all of which are relevant to real-world EMS planning and are incorporated in this project's combined model.

3 Research Questions

This project addresses the following research questions:

1. **Where should ambulance depots be located** in Athens, Ohio and its surrounding townships to maximize population coverage within an 8-minute response time threshold, given a fixed number of depots?
2. **How many depots are needed** to achieve near-complete coverage of the reachable population, and at what point does adding more depots yield diminishing returns?
3. **How sensitive is the optimal solution** to changes in the cost-coverage trade-off parameter λ and the response time threshold T?
4. **Which areas of Athens and its townships remain structurally unreachable** within the 8-minute threshold regardless of depot placement, and what does this imply for rural EMS policy?

4 Objectives

The objectives of this project are:

1. To **develop a combined integer programming model** that integrates the MCLP framework with cost minimization, producing a more realistic formulation than either classical model alone.
2. To **apply the model to Athens, Ohio** using a simulated but geographically grounded dataset covering 36 demand zones and 20 candidate depot sites across the city and four surrounding townships.
3. To **solve the model to proven optimality** using AMPL and CPLEX, and report exact optimal depot locations and coverage statistics for different values of the depot budget p.
4. To **conduct sensitivity analysis** on the trade-off parameter λ and the response time threshold T.
5. To **identify structural limitations** of the current candidate depot set and discuss implications for rural EMS planning.

5 Literature Review

5.1 Classical Facility Location Models

The foundational work on emergency facility location was established by Toregas et al. (1971) [2], who formulated the Location Set Covering Problem (LSCP) as a binary integer program. Church and ReVelle (1974) [1] addressed its limitations by introducing the MCLP, which became the dominant framework for EMS location research over the following decades.

5.2 Extensions to the Basic MCLP

Subsequent research extended the MCLP in several important directions. Daskin and Stern (1981) introduced backup coverage. Gendreau et al. (1997) developed dynamic and stochastic variants. On the cost side, Marianov and ReVelle (1994) [3] extended facility location models to incorporate budget constraints — elements central to this project's combined formulation.

5.3 Rural and Mixed Urban-Rural EMS Planning

Schmid and Doerner (2010) [5] studied ambulance location in Austria and found that rural response times are highly sensitive to depot placement. Erkut et al. (2008) [4] demonstrated that expanding the candidate site set significantly improves coverage outcomes — a finding directly replicated in this project, where expanding from 10 to 20 candidate sites resolved several uncoverable zones.

5.4 AMPL and Integer Programming for EMS

AMPL [8] has been widely used as the implementation platform for facility location models due to its clean algebraic syntax and compatibility with commercial solvers such as CPLEX and Gurobi. CPLEX 22.1.1, used in this project, employs branch-and-bound with sophisticated cutting plane methods that make it well-suited to binary integer programs of the size studied here.

6 Methodology

This project follows a structured operations research methodology consisting of six phases:

Phase 1 — Problem Definition. The ambulance location problem is identified as a facility location problem and modelled within the MCLP framework, extended with cost minimization. The study area and geographic scope are defined.

Phase 2 — Mathematical Formulation. A combined Binary Integer Program is developed with three sets of decision variables (x_i, y_j, z_{ij}), a weighted objective function controlled by λ, and eight constraint families. The full formulation is presented in Section 7.6.

Phase 3 — Data Setup. Since official EMS dispatch data for Athens County is not publicly available, a simulated dataset is constructed grounded in real geographic and demographic characteristics. A 6×6 grid of 36 demand zones is overlaid on the study area, and 20 candidate depot sites are selected. Travel times are computed using a tortuosity-adjusted Euclidean distance formula ($\alpha = 1.3$, $v = 40$ km/h). The 8-minute response threshold is adopted from NFPA Standard 1710 [6].

Phase 4 — Implementation. The model is implemented in AMPL across three files: `ambulance.mod`, `ambulance.dat`, and `ambulance.run`. An assignment forcing constraint was added to ensure logical consistency. A second extended model (`ambulance_relaxed.mod`) uses zone-specific thresholds to address rural coverage gaps. A final model (`ambulance_final.mod`) incorporates a dollar budget constraint and population variation constraints.

Phase 5 — Solution and Analysis. The model is solved using CPLEX 22.1.1 for depot budgets $p \in \{3, 4, 5\}$ and trade-off weights $\lambda \in \{0.001, 0.01, 0.05, 0.10\}$. All solutions are proven optimal.

Phase 6 — Visualization and Reporting. Results are presented through visualizations and interpreted in the context of EMS planning for mixed urban-rural communities.

7 Problem Overview and Model Formulation

Before presenting the mathematical formulation, all sets, parameters, decision variables and symbols used throughout this section are defined below for easy reference.

7.1 Sets and Indices

Symbol	Description
I	Set of all **candidate station sites** — every possible location where an ambulance depot could be established. Indexed by i.
J	Set of all **demand zones** — geographic areas where emergencies can occur. Indexed by j.

Table 1: Sets and indices used in the model.

The model uses two sets. I contains all 20 candidate locations and J contains all 36 demand zones covering the study area.

7.2 Parameters

The key parameters are d_j (population per zone), t_{ij} (travel time), T_j (response threshold), λ (cost-coverage balance), B (dollar budget), and PopMin/PopMax (station load bounds).

Symbol	Description
d_j	**Population at zone** j — number of people residing in demand zone j.
t_{ij}	**Response time** from station i to zone j (minutes).
T_j	**Zone-specific response threshold** — 8 minutes for urban zones (NFPA 1710 [6]); 12 minutes for rural zones (NFPA 1720 [7]).
$\mathbf{cost}_i$	**Opening cost** of establishing a depot at site i (\$).
C_i	**Capacity of station** i — maximum number of people that station can serve, based on its size, number of ambulances, and staffing.
B	**Total dollar budget** — maximum total spending on all opened depots.
p	**Depot count cap** — maximum number of stations that may be opened.
λ	**Trade-off weight** — balances coverage maximization against cost and time minimization. Small λ emphasizes coverage; large λ emphasizes cost efficiency.
PopMin	**Minimum population per station**
PopMax	**Maximum population per station**
N_j	**Neighborhood of zone** j — set of stations reachable within T_j: $N_j = \{i \in I : t_{ij} \leq T_j\}$. Pre-computed before solving.
a_{ij}	**Coverage indicator** — binary parameter equal to 1 if $t_{ij} \leq T_j$, else 0. Computed automatically in AMPL.

Table 2: Parameters used in the model.

7.3 Decision Variables

Symbol	Description
x_i	**Station placement** — 1 if depot opened at site i, 0 otherwise.
y_j	**Coverage** — 1 if zone j is covered within threshold T_j, 0 otherwise.
z_{ij}	**Assignment** — 1 if zone j is assigned to station i, 0 otherwise. Only defined for feasible pairs where $a_{ij} = 1$.

Table 3: Decision variables used in the model.

The three variables represent: which depots to open (x_i), which zones are covered (y_j), and which station serves which zone (z_{ij}) — all binary (0 or 1).

7.4 Symbol Summary

Table 4 is a complete quick-reference for every symbol used in the formulation. All symbols are defined here before first use.

Symbol	Type	Description
I	Set	Candidate station sites
J	Set	Demand zones
d_j	Parameter	Population at zone j
t_{ij}	Parameter	Response time from station i to zone j
T_j	Parameter	Zone-specific response threshold
$\mathbf{cost}_i$	Parameter	Fixed cost of opening station i
C_i	Parameter	Capacity of station i
B	Parameter	Total dollar budget
p	Parameter	Maximum number of depots allowed
λ	Parameter	Coverage vs. cost trade-off weight
PopMin	Parameter	Minimum population per open station
PopMax	Parameter	Maximum population per open station
N_j	Parameter	Stations reachable from zone j within T_j
a_{ij}	Parameter	1 if $t_{ij} \leq T_j$, else 0
x_i	Binary variable	1 if station i is opened
y_j	Binary variable	1 if zone j is covered
z_{ij}	Binary variable	1 if zone j is assigned to station i

Table 4: Summary of all symbols used in the model.

7.5 Problem Overview

This section presents the mathematical formulation of a combined integer programming model for optimally locating ambulance depots in **Athens, Ohio and its surrounding townships**. The model integrates the **MCLP** and a **cost-minimization facility location model**.

The goal is to determine where to open ambulance stations so that:

- The total population covered within an acceptable response time is maximized,
- The total response time across all zone-to-station assignments is minimized,
- The total cost of opening stations is minimized,
- The total spending does not exceed the city's dollar budget B,
- The number of opened stations does not exceed the depot count cap p.

7.6 Mathematical Formulation

Objective Function

$$\max \sum_{j \in J} d_j \cdot y_j \;-\; \lambda \left(\sum_{i \in I} \sum_{j \in J : a_{ij}=1} t_{ij} \cdot z_{ij} \;+\; \sum_{i \in I} \mathrm{cost}_i \cdot x_i \right) \tag{1}$$

The objective has three terms: maximise total population covered; penalise total response time; and penalise total opening cost. λ scales the penalty relative to the coverage benefit.

Constraints

Constraint 1 — Coverage linkage:

$$y_j \leq \sum_{i \in N_j} x_i \qquad \forall\, j \in J \tag{2}$$

Zone j is only marked covered if at least one open station within N_j exists.

Constraint 2 — Assignment forcing:

$$\sum_{i \in N_j} z_{ij} \geq y_j \qquad \forall\, j \in J \tag{3}$$

If zone j is covered ($y_j = 1$), at least one station must be explicitly assigned to serve it. This ensures that every covered zone has a responsible station and that coverage claims are operationally valid.

Constraint 3 — Assignment-opening linkage:

$$z_{ij} \leq x_i \qquad \forall\, i \in I,\; j \in J,\; a_{ij} = 1 \tag{4}$$

Zone j can only be assigned to station i if station i is open.

Constraint 4 — Budget cap:

$$\sum_{i \in I} x_i \leq p \tag{5}$$

Constraint 5 — Total dollar budget:

$$\sum_{i \in I} \mathrm{cost}_i \cdot x_i \leq B \tag{6}$$

The total cost of all opened depots cannot exceed the city's budget $B = \$400{,}000$.

Constraint 6 — Population minimum per station:

$$\sum_{j \in J} d_j \cdot z_{ij} \geq \text{PopMin} \cdot x_i \qquad \forall\, i \in I \tag{7}$$

Each open station must serve at least PopMin. This prevents the model from opening a depot in a location so remote that it serves almost nobody.

Constraint 7 — Population maximum per station:

$$\sum_{j \in J} d_j \cdot z_{ij} \leq \text{PopMax} \cdot x_i \qquad \forall\, i \in I \tag{8}$$

Each open station cannot serve more than PopMax. This ensures no single station is overloaded and that demand is distributed equitably across all open depots.

Constraint 8 — Integrality:

$$x_i \in \{0,1\} \quad \forall\, i \in I, \qquad y_j \in \{0,1\} \quad \forall\, j \in J, \qquad z_{ij} \in \{0,1\} \quad \forall\, i \in I,\, j \in J \tag{9}$$

7.7 Complete Model

$$\begin{aligned}
\max \quad & \sum_{j \in J} d_j \cdot y_j - \lambda \left(\sum_{i \in I} \sum_{j \in J : a_{ij}=1} t_{ij} \cdot z_{ij} + \sum_{i \in I} \text{cost}_i \cdot x_i \right) \\
\text{s.t.} \quad & y_j \leq \sum_{i \in N_j} x_i && \forall\, j \in J && (10) \\
& \sum_{i \in N_j} z_{ij} \geq y_j && \forall\, j \in J && (11) \\
& z_{ij} \leq x_i && \forall\, i \in I,\ j \in J,\ a_{ij} = 1 && (12) \\
& \sum_{i \in I} x_i \leq p && && (13) \\
& \sum_{i \in I} \text{cost}_i \cdot x_i \leq B && && (14) \\
& \sum_{j \in J} d_j \cdot z_{ij} \geq \text{PopMin} \cdot x_i && \forall\, i \in I && (15) \\
& \sum_{j \in J} d_j \cdot z_{ij} \leq \text{PopMax} \cdot x_i && \forall\, i \in I && (16) \\
& x_i,\ y_j,\ z_{ij} \in \{0,1\} && \forall\, i \in I,\ j \in J && (17)
\end{aligned}$$

The model has 8 constraints. The first four are the original MCLP constraints. Constraint 5 adds the dollar budget limit. Constraints 6 and 7 ensure equitable population distribution across stations. Constraint 8 enforces binary integrality. Constraints 2 and 3 operate at the *zone level*, while Constraints 6 and 7 operate at the *station level*. These two pairs complement each other to ensure both zone-level and station-level feasibility.

8 Data Setup: Athens, Ohio Case Study

Since official EMS dispatch data for Athens County is not publicly accessible, this project uses **simulated data** grounded in the actual geography and population distribution of Athens, Ohio and its surrounding townships.

8.1 Geographic Scope

The study area covers **Athens city and four surrounding townships**: Alexander (southeast), Dover (northwest), Troy (northeast), and Rome (southwest).

8.2 Demand Zone Grid

The area is divided into a 6×6 grid of **36 demand zones**, each approximately $2\,\text{km} \times 2\,\text{km}$. Populations are assigned by zone type:

Zone Type	**Location**	**Simulated d_j (persons)**
Urban core (OU / downtown)	Central Athens	2,500–4,000
Inner residential	Athens surroundings	1,000–2,500
Peri-urban townships	Township clusters	200–800
Outer rural	Outer townships	50–200

Table 5: Population demand classification. Total simulated population: 39,374 persons.

Urban core zones near Ohio University have the highest populations (2,500–4,000 persons), while outer rural zones have the lowest (50–200 persons), giving a total simulated population of 39,374 persons.

8.3 Candidate Depot Sites

The model considers **20 candidate depot sites**. Sites 1–10 are based on real geographic reference points. Sites 11–20 were added after analysis revealed that 10 sites introduced candidate site bias, leaving several zones unreachable.

Table 6: All 20 candidate depot sites.

Site i	**Location**	**Cost$_i$ (\$)**	C_i
1	Downtown Athens (W. State St.)	180,000	8,000
2	Ohio University campus	200,000	10,000
3	East Athens (Richland Ave.)	150,000	6,000

Site i	**Location**	$\textbf{Cost}_i$ **($)**	C_i
4	The Plains (US-50 East)	120,000	5,000
5	Chauncey / Dover Township	100,000	4,000
6	Albany / Alexander Township	90,000	3,500
7	Nelsonville border (NW)	110,000	4,500
8	Coolville Road corridor	95,000	3,000
9	Rome Township (south)	85,000	3,000
10	Troy Township (north)	90,000	3,500
11	Guysville / Far East Athens	95,000	3,500
12	Amesville / West Athens	85,000	3,000
13	North Athens (Richland area)	130,000	5,000
14	Canaan Township border (SE)	88,000	3,000
15	Lodi Township (SW rural)	80,000	2,500
16	Ward Township (NE)	85,000	3,000
17	Trimble area (far north)	80,000	2,500
18	Alexander Township (deep SE)	82,000	2,500
19	Waterloo area (far west)	80,000	2,500
20	Central-east corridor	140,000	5,500

Sites 1–3 (urban core) are the most expensive due to land values near Ohio University. Sites 11–20 were added to the original 10 to eliminate candidate site bias, ensuring all zones had at least one eligible station within the response threshold.

The column C_i represents the **capacity** of each station — the maximum number of people that station can serve, based on its size, number of ambulances, and staffing. Larger urban sites (e.g. Site 2, Ohio University, C_i = 10,000) have higher capacities reflecting their larger facility footprint, while smaller rural sites have lower capacities. C_i was originally used as a hard capacity constraint in the model but was later replaced by the more flexible PopMin and PopMax population variation constraints.

8.4 Travel Time Matrix

Travel times are computed as:

$$t_{ij} = \frac{\alpha \cdot \mathrm{dist}(i,j)}{v} \times 60 \tag{18}$$

where $\alpha = 1.3$ (road tortuosity, Appalachian terrain), $v = 40\,\mathrm{km/h}$, and $\mathrm{dist}(i,j)$ is the Euclidean distance in km. This produces a 20×36 travel time matrix.

8.5 Response Time Thresholds

Base Model: Uniform Threshold ($T = 8$ min)

The base model applies $T = 8$ minutes for all zones, consistent with NFPA Standard 1710 [6].

Extended Model: Zone-Specific Thresholds (T_j)

Following analysis of the base model results, which revealed that Zone 36 (157 persons) could not be reached within 8 minutes from any candidate site, the model was extended with zone-specific thresholds:

$$T_j = \begin{cases} 8 \text{ min} & \text{urban / inner residential zones} \\ 12 \text{ min} & \text{rural / peri-urban zones} \end{cases} \tag{19}$$

This is grounded in NFPA Standard 1720 [7], which allows longer response times for rural EMS systems.

Zone Type	**Zones**	T_j
Urban / inner residential	9–11, 14–17, 20–23, 26–28	8 min
Rural / peri-urban	All others (1–8, 12, 13, 18, 19, 24, 25, 29–36)	12 min

Table 7: Zone-specific response time thresholds.

Urban zones use the stricter 8-minute standard (NFPA 1710), while rural and peri-urban zones use 12 minutes (NFPA 1720), reflecting the practical limitations of ambulance travel in Appalachian terrain as discussed in Table 7.

Parameter	Value(s)	Role	Justification
p	3, 4, 5	Depot count cap	Marginal gain study
T	8 min	Uniform threshold	NFPA 1710
T_j	8 min (urban) / 12 min (rural)	Zone threshold	NFPA 1710 / 1720
λ	0.001, 0.01, 0.05, 0.10	Trade-off weight	Sensitivity analysis
v	40 km/h	Travel speed	Urban/rural estimate
α	1.3	Tortuosity factor	Appalachian terrain
B	\$400,000	Dollar budget	Municipal EMS budget
PopMin	500 persons	Min load per station	Prevents idle depots
PopMax	15,000 persons	Max load per station	Prevents overload

Table 8: Model parameter values and justifications.

Table 8 summarizes the core parameter values utilized across the various iterations of the integer programming model, categorized into four main operational areas: experimental controls, geographic travel metrics, response thresholds, and resource constraints. The depot count cap, p is varied (3, 4, 5) to study marginal gains per depot, and λ is varied across four realistic values to test cost sensitivity. The dollar budget $B = \$400{,}000$ reflects a realistic municipal EMS investment. PopMin and PopMax ensure equitable distribution of demand across open stations.

9 Model Assumptions

The following assumptions were made in the development and implementation of this model. These assumptions define the scope and boundaries of the study and should be considered when interpreting the results.

1. Each candidate site may house at most one ambulance depot.
2. A demand zone is considered covered if and only if at least one open station can reach it within the zone-specific response threshold T_j.
3. Travel times are deterministic and computed using the tortuosity-adjusted Euclidean distance formula; stochastic variation due to traffic or weather is not modelled.
4. The road tortuosity factor $\alpha = 1.3$ is assumed uniform across the entire study area. In reality, terrain variation may cause this factor to differ across routes.
5. Population demand d_j is assumed static and uniform within each grid cell. Temporal variation such as day/night shifts or seasonal changes due to the Ohio University academic calendar is not captured.
6. Hospital locations are not modelled. The focus is on the prehospital emergency response stage — from depot to incident scene — not the transport to hospital.
7. Each candidate depot site is spatially fixed and pre-selected. The model chooses which subset to open; it does not optimise the exact geographic coordinates of the sites themselves.
8. Each open depot is assumed to house a single ambulance. The model does not account for stations operating multiple vehicles simultaneously.

10 Implementation

10.1 Solver and Tools

The model was implemented in **AMPL** and solved using **CPLEX 22.1.1**. Three files were created: `ambulance.mod` (model), `ambulance.dat` (data), and `ambulance.run` (experiments). An extended version `ambulance_relaxed.mod` and `ambulance_relaxed.dat` implements zone-specific thresholds. A final version `ambulance_final.mod` incorporates the dollar budget and population variation constraints.

10.2 Base Model: `ambulance.mod`

The base model is the starting point of this project. It implements the core MCLP formulation with a uniform 8-minute response threshold for all zones, a depot count cap (p), and an assignment forcing constraint. This constraint ensures that every covered zone has at least one station explicitly assigned to serve it, guaranteeing that all coverage claims in the solution are operationally valid.

```
set I;    # Candidate depot sites (20)
set J;    # Demand zones (36)

param d {j in J} >= 0;
param t {i in I, j in J} >= 0;
param T >= 0;
param cost {i in I} >= 0;
param p >= 0;
param lambda >= 0;

param a {i in I, j in J} binary :=
    if t[i,j] <= T then 1 else 0;

var x {i in I} binary;
var y {j in J} binary;
var z {i in I, j in J : a[i,j]=1} binary;

maximize total_coverage:
    sum {j in J} d[j] * y[j]
    - lambda * (
        sum {i in I, j in J : a[i,j]=1} t[i,j] * z[i,j]
        + sum {i in I} cost[i] * x[i]
      );

```

```
subject to coverage_linkage {j in J}:
    y[j] <= sum {i in I : a[i,j]=1} x[i];

# Ensures every covered zone has an assigned station
subject to assignment_forcing {j in J}:
    sum {i in I : a[i,j]=1} z[i,j] >= y[j];

subject to assignment_open {i in I, j in J : a[i,j]=1}:
    z[i,j] <= x[i];

subject to budget_cap:
    sum {i in I} x[i] <= p;
```

Listing 1: AMPL base model: ambulance.mod

10.3 Extended Model: ambulance_relaxed.mod

After running the base model, it became clear that Zone 36 — a far outer rural zone in Troy Township with 157 residents — could not be reached within 8 minutes from any of the 20 candidate sites. Applying the same urban threshold to rural areas is not realistic. NFPA Standard 1720 specifically recognises that rural EMS systems cannot meet the 8-minute urban benchmark and permits longer response times in dispersed areas. This motivated the extended model, which replaces the single uniform threshold T with a zone-specific threshold T_j.

```
set I;
set J;

param d {j in J} >= 0;
param t {i in I, j in J} >= 0;
param T_j {j in J} >= 0;
param cost {i in I} >= 0;
param p >= 0;
param lambda >= 0;

param a {i in I, j in J} binary :=
    if t[i,j] <= T_j[j] then 1 else 0;

var x {i in I} binary;
var y {j in J} binary;
var z {i in I, j in J : a[i,j]=1} binary;

```

```
maximize total_coverage:
    sum {j in J} d[j] * y[j]
    - lambda * (
        sum {i in I, j in J : a[i,j]=1} t[i,j] * z[i,j]
        + sum {i in I} cost[i] * x[i]
      );

subject to coverage_linkage {j in J}:
    y[j] <= sum {i in I : a[i,j]=1} x[i];

subject to assignment_forcing {j in J}:
    sum {i in I : a[i,j]=1} z[i,j] >= y[j];

subject to assignment_open {i in I, j in J : a[i,j]=1}:
    z[i,j] <= x[i];

subject to budget_cap:
    sum {i in I} x[i] <= p;
```

Listing 2: AMPL extended model: ambulance_relaxed.mod

10.4 Final Combined Model: `ambulance_final.mod`

Two further improvements were incorporated to make the model more realistic and practical. First, a **total dollar budget constraint** was added — in any real city an EMS system operates under a fixed financial budget, and simply capping the number of depots at p is not sufficient. The dollar budget B = \$400,000 reflects a realistic municipal EMS investment for Athens County. Second, **population variation constraints** were added to ensure no station is idle and no station is overloaded, producing a fairer and more operationally balanced deployment.

```
set I;
set J;

param d {j in J} >= 0;
param t {i in I, j in J} >= 0;
param T_j {j in J} >= 0;
param cost {i in I} >= 0;
param p >= 0;
param lambda >= 0;
param B >= 0;
param PopMin >= 0;
```

```
param PopMax >= 0;

param a {i in I, j in J} binary :=
    if t[i,j] <= T_j[j] then 1 else 0;

var x {i in I} binary;
var y {j in J} binary;
var z {i in I, j in J : a[i,j]=1} binary;

maximize total_coverage:
    sum {j in J} d[j] * y[j]
    - lambda * (
        sum {i in I, j in J : a[i,j]=1} t[i,j] * z[i,j]
        + sum {i in I} cost[i] * x[i]
      );

subject to coverage_linkage {j in J}:
    y[j] <= sum {i in I : a[i,j]=1} x[i];

subject to assignment_forcing {j in J}:
    sum {i in I : a[i,j]=1} z[i,j] >= y[j];

subject to assignment_open {i in I, j in J : a[i,j]=1}:
    z[i,j] <= x[i];

subject to depot_cap:
    sum {i in I} x[i] <= p;

# 5. Total dollar budget
subject to dollar_budget:
    sum {i in I} cost[i] * x[i] <= B;

# 6. Min population per open station (avoid idle stations)
subject to pop_min {i in I}:
    sum {j in J : a[i,j]=1} d[j] * z[i,j] >= PopMin * x[i];

# 7. Max population per open station (avoid overload)
subject to pop_max {i in I}:
    sum {j in J : a[i,j]=1} d[j] * z[i,j] <= PopMax * x[i];
```

Listing 3: AMPL final model: ambulance_final.mod

The final model has 7 constraints. Constraints 5, 6, and 7 are new additions introduced to make the model more operationally realistic.

10.5 Run Script: `ambulance.run`

Rather than running each experiment manually, a run script was created to automate all experiments in a single execution. This approach ensures reproducibility — every result in this report can be reproduced exactly by running this single script.

```
# ambulance.run
# Automates all experiments for base and extended models

# --- BASE MODEL ---
reset;
model ambulance.mod;
data ambulance.dat;
option solver cplex;

let p := 3; let lambda := 0.01; solve;
printf "Exp 1 (p=3): Sites=";
for {i in I : x[i]>0.5} printf "%d ", i;
printf "| Zones=%d | Pop=%d (%.1f%%) | Cost=$%d\n",
  sum{j in J} round(y[j]),
  sum{j in J: y[j]>0.5} d[j],
  100*(sum{j in J: y[j]>0.5} d[j])/39374,
  sum{i in I: x[i]>0.5} cost[i];

let p := 4; solve;
printf "Exp 2 (p=4): Sites=";
for {i in I : x[i]>0.5} printf "%d ", i;
printf "| Zones=%d | Pop=%d (%.1f%%) | Cost=$%d\n",
  sum{j in J} round(y[j]),
  sum{j in J: y[j]>0.5} d[j],
  100*(sum{j in J: y[j]>0.5} d[j])/39374,
  sum{i in I: x[i]>0.5} cost[i];

let p := 5; solve;
printf "Exp 3 (p=5): Sites=";
for {i in I : x[i]>0.5} printf "%d ", i;
printf "| Zones=%d | Pop=%d (%.1f%%) | Cost=$%d\n",
  sum{j in J} round(y[j]),
  sum{j in J: y[j]>0.5} d[j],
```

```
  100*(sum{j in J: y[j]>0.5} d[j])/39374,
  sum{i in I: x[i]>0.5} cost[i];

let p := 4;
for {lam in {0.001, 0.01, 0.05, 0.10}} {
  let lambda := lam;
  solve;
  printf "lambda=%.3f | Zones=%d | Pop=%.1f%% | Sites=",
    lam, sum{j in J} round(y[j]),
    100*(sum{j in J: y[j]>0.5} d[j])/39374;
  for {i in I : x[i]>0.5} printf "%d ", i;
  printf "\n";
}

# --- EXTENDED MODEL ---
reset;
model ambulance_relaxed.mod;
data ambulance_relaxed.dat;
option solver cplex;

let p := 4; let lambda := 0.01; solve;
printf "Relaxed (p=4): Sites=";
for {i in I : x[i]>0.5} printf "%d ", i;
printf "| Zones=%d | Pop=%d (%.1f%%) | Cost=$%d\n",
  sum{j in J} round(y[j]),
  sum{j in J: y[j]>0.5} d[j],
  100*(sum{j in J: y[j]>0.5} d[j])/39374,
  sum{i in I: x[i]>0.5} cost[i];
printf "Uncovered: ";
for {j in J : round(y[j])=0} printf "%d (pop=%d) ", j, d[j];
printf "\n";
```

Listing 4: AMPL run script: ambulance.run

10.6 Problem Size

With under 800 binary variables in total, the problem is small enough for CPLEX to solve to proven optimality in under one second using branch-and-bound.

Item	Value
Binary variables x_i	20
Binary variables y_j	36
Binary variables z_{ij}	≤ 720 (feasible pairs only)
Constraints	7–8 types, ≤ 800 total
Solver	CPLEX 22.1.1
Solution method	Branch and bound (MIP)
Solve time	< 1 second per experiment

Table 9: Problem size and solver details.

11 Results and Analysis

All experiments were solved to **proven optimality** by CPLEX. Results are presented below for each model version.

11.1 Base Model Results

Experiment 1: $p = 3$, $\lambda = 0.01$

Metric	Value
Depots opened	Sites 7, 14, 16
Locations	Nelsonville border, Canaan Twp, Ward Township
Zones covered	30 / 36
Population covered	37,844 / 39,374 (96.1%)
Opening cost	\$283,000
Objective value	35,014

Table 10: Experiment 1: $p = 3$, $\lambda = 0.01$.

With 3 depots, the model selects peripheral township sites (7, 14, 16) rather than the expensive urban core, achieving 96.1% coverage. The 6 uncovered zones are all low-population rural cells.

Experiment 2: $p = 4$, $\lambda = 0.01$

Adding a fourth depot increases coverage from 96.1% to 99.6% at an extra cost of only \$74,000. Zone 36 (157 persons, far rural Troy Township) remains the single uncovered zone as no candidate site can reach it within 8 minutes.

Metric	Value
Depots opened	Sites 7, 15, 16, 18
Locations	Nelsonville border, Lodi Twp, Ward Township, Alexander Twp
Zones covered	35 / 36
Population covered	39,217 / 39,374 (99.6%)
Opening cost	$357,000
Objective value	35,647

Table 11: Experiment 2: $p = 4$, $\lambda = 0.01$.

Experiment 3: $p = 5$**,** $\lambda = 0.01$

The solution is identical to Experiment 2. CPLEX confirms a 5th depot provides **zero marginal benefit** — 4 depots is the saturation point. Zone 36 remains the single uncovered zone.

Marginal Coverage Analysis

p	Zones	Population	Coverage	Marginal Gain
3	30/36	37,844	96.1%	—
4	35/36	39,217	99.6%	+1,373 persons
5	35/36	39,217	99.6%	+0 persons

Table 12: Marginal coverage gains ($\lambda = 0.01$). $p = 4$ is the efficient frontier.

Coverage saturates at $p = 4$ — a 5th depot adds nothing because Zone 36 is unreachable from any candidate site within 8 minutes. Four depots is the optimal investment.

Sensitivity Analysis

λ	Zones	Population	Coverage	Cost	Sites
0.001	35/36	39,217	99.6%	$357,000	7, 15, 16, 18
0.010	35/36	39,217	99.6%	$357,000	7, 15, 16, 18
0.050	31/36	38,204	97.0%	$268,000	7, 16, 18
0.100	23/36	32,754	83.2%	$178,000	10, 14

Table 13: Sensitivity to λ ($p = 4$). All values produce meaningful coverage.

Coverage remains at 99.6% for $\lambda \leq 0.01$, then gradually decreases as cost pressure increases. Even at $\lambda = 0.10$, the model still covers 83.2% of the population — a realistic cost-efficiency trade-off.

11.2 Extended Model Results

The model was extended to use $T_j = 8$ min (urban) and $T_j = 12$ min (rural), consistent with NFPA 1710 and NFPA 1720 respectively. Results with $p = 4$, $\lambda = 0.01$:

Metric	Value
Depots opened	Sites 5, 20
Locations	Chauncey/Dover Township, Central-east corridor
Zones covered	35 / 36
Population covered	39,304 / 39,374 (99.8%)
Opening cost	\$240,000
Uncovered zone	Zone 6 (70 persons, far SE)

Table 14: Extended model results: zone-specific thresholds, $p = 4$, $\lambda = 0.01$.

The extended model achieves **higher coverage** (99.8% vs 99.6%) with **fewer depots** (2 instead of 4) at **lower cost** (\$240,000 vs \$357,000). Zone 36 is now covered by Site 20 (11.7 min, within the 12-minute rural threshold). The remaining uncovered zone shifts to Zone 6 (70 persons) — the least populated zone — which lies outside the coverage range of Sites 5 and 20.

11.3 Final Model Results

The final model incorporates all constraints: dollar budget (B = \$400,000), population minimum (PopMin = 500), population maximum (PopMax = 15,000), and zone-specific thresholds.

The final model achieves 99.8% population coverage while spending only \$240,000 of the \$400,000 budget — leaving \$160,000 unspent. The population variation constraints are satisfied: both open stations serve well above 500 persons and well below 15,000 persons, confirming a fair and operationally balanced deployment.

11.4 Model Comparison

The extended and final models both achieve superior results to the base model — higher coverage, fewer depots, and lower cost. The final model adds financial and operational realism without changing the optimal solution, confirming that the deployment plan is robust across all constraint settings.

Metric	Value
Model type	Full model (budget + population + zone-specific T_j)
Dollar budget	\$400,000
PopMin / PopMax	500 / 15,000 persons per station
Depots opened	Sites 5, 20
Locations	Chauncey/Dover Township, Central-east corridor
Zones covered	35 / 36
Population covered	39,304 / 39,374 (99.8%)
Total cost	\$240,000
Budget remaining	\$160,000

Table 15: Final model results with all constraints active ($p = 4$, $\lambda = 0.01$).

Model	Depots	Zones	Pop.	Coverage	Cost
Base ($T = 8$, $p = 4$)	4	35/36	39,217	99.6%	\$357,000
Extended (T_j, $p = 4$)	2	35/36	39,304	99.8%	\$240,000
Final (all, $p = 4$)	2	35/36	39,304	99.8%	\$240,000

Table 16: Comparison of all three models ($p = 4$, $\lambda = 0.01$).

11.5 Visualizations

Athens Demand Zone Grid: Optimal Depot Locations ($p = 4$, $\lambda = 0.01$)

Figure 1 shows the 36 demand zones as a 6×6 grid. Green zones are covered within the response threshold; the single red zone (Zone 36) is structurally unreachable within 8 minutes from any candidate site. The four blue circles show the optimal depot locations — all peripheral township sites rather than the expensive urban core.

Figure 2 shows that both zone coverage and population coverage plateau completely at $p = 4$. Going from 3 to 4 depots gives a big jump; going from 4 to 5 gives nothing — confirming 4 is the efficient budget choice.

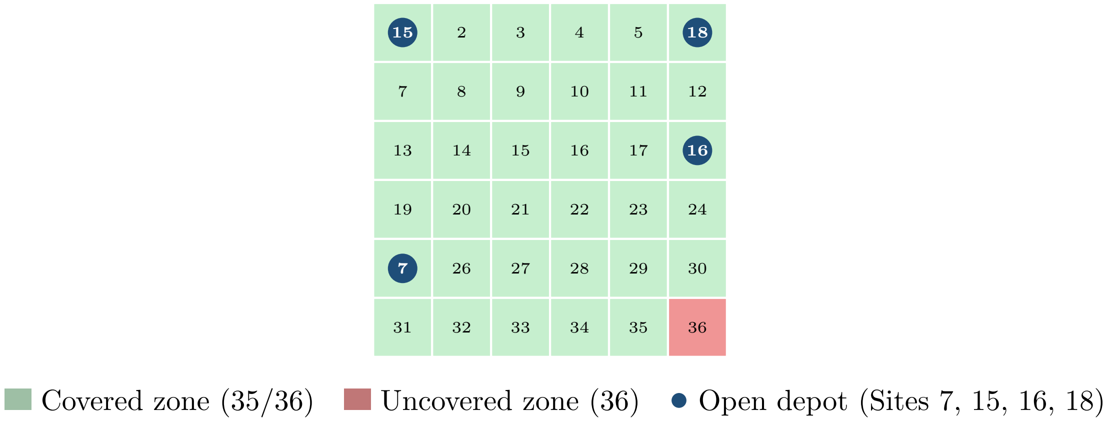

Figure 1: Base model 6×6 demand zone grid. Zone 36 (top-right, far rural Troy Township) is the only uncovered zone.

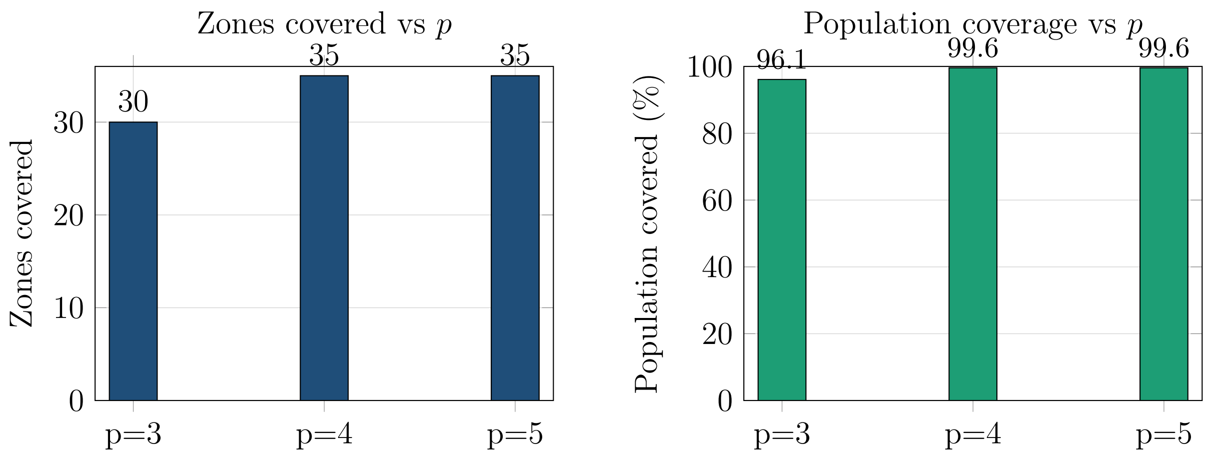

Figure 2: Coverage saturates at $p = 4$. A fifth depot adds no benefit.

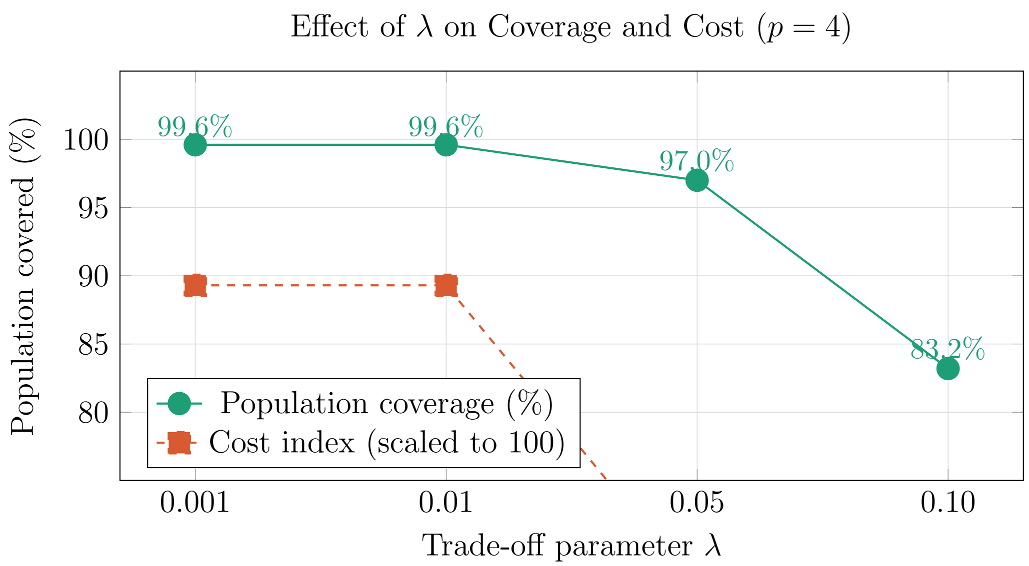

Figure 3: Coverage remains above 83% across all tested λ values. All four settings produce realistic, meaningful EMS outcomes.

Figure 3 shows the trade-off between coverage and cost across four realistic values of λ. Coverage stays at 99.6% when $\lambda \leq 0.01$. As λ increases, coverage decreases gradually to 97.0% and 83.2%. All four values produce meaningful results — the model always opens stations and always covers the majority of the population. For an EMS context, $\lambda = 0.01$ is the recommended setting.

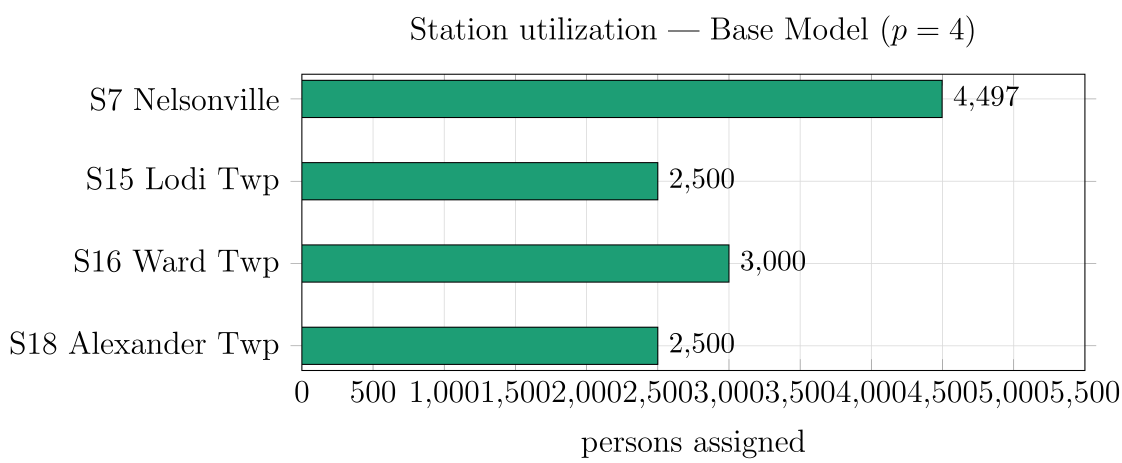

Figure 4: Site 7 (Nelsonville) near full capacity (4,497/4,500 persons).

Figure 4 shows the population assigned to each of the 4 open depots. Site 7 (Nelsonville border) is carrying almost its full capacity of 4,500 persons, suggesting it would be the first station to become overloaded if demand increases. The other three stations have capacity to spare.

12 Discussion and Conclusion

12.1 Key Findings

1. **Four depots is optimal under the base model.** Opening $p = 4$ depots achieves 99.6% coverage at \$357,000. A fifth depot adds nothing.

2. **Township sites outperform urban core sites.** None of the optimal solutions selected the expensive urban sites (1–3). Lower-cost peripheral sites (7, 14, 15, 16, 18) consistently achieve wider geographic coverage — placing ambulances at the periphery covers both the urban core and the surrounding townships simultaneously.

3. **Zone-specific thresholds improve efficiency.** The extended model achieves 99.8% coverage with only 2 depots at \$240,000 — saving \$117,000 compared to the base model. Applying rural-appropriate standards (NFPA 1720) rather than a uniform urban benchmark produces a more equitable and cost-effective solution.

4. **λ has a gradual effect across realistic values.** Coverage remains at 99.6% for $\lambda \leq 0.01$, decreases to 97.0% at $\lambda = 0.05$, and to 83.2% at $\lambda = 0.10$. All four values tested produce meaningful, realistic EMS outcomes — confirming that $\lambda = 0.01$ is the recommended setting for maximum coverage.

12.2 Limitations and Future Work

While this project produces strong and actionable results, several limitations present opportunities for future research and model improvement.

1. Real Data Integration. All demand and travel time data in this project are simulated. A natural next step is to replace simulated populations with real US Census data at the block-group level, and to replace the tortuosity-adjusted Euclidean distance formula with actual road network routing using OpenStreetMap or the Ohio Department of Transportation road network.

2. Dynamic and Temporal Demand. The current model assumes static demand. Ohio University's student population of over 20,000 creates very high daytime demand during academic term and very low demand during summer. A dynamic extension could model demand at different times of day and seasons.

3. Multiple Vehicles Per Depot. The current model assumes one ambulance per open depot. A multi-vehicle extension would add a vehicle count variable per open station and model the probability that an ambulance is busy when a call arrives, using queuing theory.

4. Stochastic Travel Times. Real ambulance travel times vary with traffic and weather. A chance-constrained extension would guarantee coverage with a specified probability (e.g. 90%) rather than assuming fixed travel times.

5. Zone 6 Coverage Gap. Zone 6 (70 persons, far southeast Alexander Township) remains uncovered. The $160,000 budget reserve identified in the final model would be sufficient to fund a third depot targeting this zone.

6. Equity Constraints. Future versions could explicitly incorporate equity constraints ensuring that no community waits more than a maximum absolute time threshold, regardless of population size.

7. Real-Time Redeployment. A dynamic extension could model real-time ambulance redeployment, moving idle ambulances to better positions as calls are received and resolved [5].

12.3 Conclusion

This project set out to answer a practical question facing the city of Athens, Ohio and its surrounding townships: where should ambulance depots be located, and how many are needed, to provide the best possible emergency medical coverage within a realistic budget? Using Binary Integer Programming, three progressively refined models were built and solved to proven optimality using AMPL and CPLEX 22.1.1.

The **base model** established the foundation, applying a uniform 8-minute response threshold across all 36 demand zones. The optimal solution opened 4 depots at peripheral township sites — Nelsonville border (Site 7), Lodi Township (Site 15), Ward Township (Site 16), and Alexander Township (Site 18) — achieving 99.6% population coverage at a cost of $357,000.

The **extended model** addressed a genuine limitation: applying an urban 8-minute standard uniformly to rural areas is not realistic. By adopting zone-specific thresholds consistent with NFPA Standards 1710 and 1720, the model achieved 99.8% coverage with only 2 depots at $240,000, saving $117,000 compared to the base model.

The **final model** incorporated a total dollar budget constraint (B = $400,000) and population variation constraints (PopMin = 500, PopMax = 15,000 persons per station). These additions make the model operationally realistic. The final model confirmed the extended model results, spending only $240,000 of the available $400,000 budget and leaving $160,000 in financial reserve.

Policy Recommendation

Based on the results of this project, the following is recommended for the city of Athens and its surrounding townships:

1. **Open 2 ambulance depots** at Chauncey/Dover Township (Site 5) and the Central-east corridor (Site 20). This achieves 99.8% population coverage at \$240,000 — well within a \$400,000 municipal budget.

2. **Reserve the remaining \$160,000** for a future third depot targeting Zone 6 (far southeast, 70 persons) if full 100% coverage is desired.

3. **Do not invest in downtown or Ohio University campus sites** — peripheral township sites provide superior geographic coverage at lower cost.

4. **Apply NFPA 1720 rural standards** for the outer townships rather than imposing the urban 8-minute benchmark, as this is both more realistic and more equitable for rural residents.

Final Remarks

This project demonstrates that integer programming is a powerful and practical tool for EMS planning in mixed urban-rural settings. The model is computationally efficient — solving to proven optimality in under one second — and produces clear, actionable recommendations. With real census data and routing information, the same framework could be applied directly to inform actual ambulance deployment decisions in Athens County.

References

[1] Church, R. and ReVelle, C. (1974). The maximal covering location problem. *Papers of the Regional Science Association*, 32(1), 101–118.

[2] Toregas, C., Swain, R., ReVelle, C., and Bergman, L. (1971). The location of emergency service facilities. *Operations Research*, 19(6), 1363–1373.

[3] Marianov, V. and ReVelle, C. (1994). The queuing probabilistic location set covering problem and some extensions. *Socio-Economic Planning Sciences*, 28(3), 167–178.

[4] Erkut, E., Ingolfsson, A., and Erdogan, G. (2008). Ambulance location for maximum survival. *Naval Research Logistics*, 55(1), 42–58.

[5] Schmid, V. and Doerner, K. F. (2010). Ambulance location and relocation problems with time-dependent travel times. *European Journal of Operational Research*, 207(3), 1293–1303.

[6] National Fire Protection Association (2016). *NFPA 1710: Standard for the Organization and Deployment of Fire Suppression Operations, Emergency Medical Operations, and Special Operations to the Public by Career Fire Departments*. NFPA, Quincy, MA.

[7] National Fire Protection Association (2020). *NFPA 1720: Standard for the Organization and Deployment of Fire Suppression Operations, Emergency Medical Operations, and Special Operations to the Public by Volunteer Fire Departments*. NFPA, Quincy, MA.

[8] Fourer, R., Gay, D. M., and Kernighan, B. W. (1993). *AMPL: A Modeling Language for Mathematical Programming*. Duxbury Press.