National Basketball Association Scheduling: Location Tracking IP Approach

21-393 Final Project, Fall 2024
Animesh Ghose, Adel Heddadji, Ashwin Ramanathan, Kwaku Bimpong

1 Introduction

1.1 Back-to-Back Games

The National Basketball Association's (NBA's) extensive travel requirements create significant challenges for both player health and team performance. The physical demands of travel, particularly across long distances can strain athletes' bodies, leading to fatigue and an increased risk of injuries. The psychological toll of constant travel is equally taxing, with the disruption to daily routines and insufficient recovery periods contributing to mental exhaustion. These factors collectively undermine a teams competitiveness, impacting their ability to perform consistently over the course of the season.

Compounding the effects of travel are the back-to-back games, a fixture in the NBA schedule that forces teams to play two games on consecutive nights with little time for recovery. These high-intensity situations, where players must quickly pivot between demanding games, heighten the chances of injury and hinder overall performance. As a result, teams facing frequent back-to-back matchups often struggle with recovery and player durability, which can ultimately affect their playoff aspirations.

In light of these challenges, this project aims to explore methods of minimizing back-to-back games in the NBA schedule. By reducing the frequency of these back-to-back matchups, we can improve player health, enhance overall team performance, and reduce the financial and long-term physical costs associated with the current scheduling practices.

1.2 Travel Distance

Physical fatigue and travel-related challenges present significant concerns for NBA teams. Extended travel can severely impact players' physical condition and mental well-being, creating substantial stress for both athletes and coaching staff. The physiological and psychological toll of constant long-distance travel is not merely an inconvenience but a critical factor that can directly influence team performance.

Research indicates that extensive travel correlates with measurable performance declines. Scientific studies have demonstrated that teams subjected to more extensive travel distances experience notable decreases in competitive effectiveness. Beyond the athletic implications, these travel requirements also impose considerable financial burdens on organizations, with potential cost savings achievable through strategic travel reduction.

Moreover, the environmental impact of NBA team travel cannot be overlooked. A typical NBA team generates hundreds of metric tons of CO_2 emissions per season through air travel, contributing significantly to carbon emissions. By minimizing travel distances, teams have an opportunity to not only improve athlete performance and reduce expenses but also demonstrate meaningful environmental responsibility. This approach would allow organizations to decrease fuel consumption and actively participate in sustainable practices, potentially setting a progressive example within professional sports.

1.3 Why 82 Games?

The 82-game regular season format has been a foundation of the NBA since the 1967-1968 season, deeply embedded in the league's structure and history. Although past NBA scheduling projects from this course have reduced the number of games, we maintain the 82-game season for several reasons.

First, it allows for historical consistency, which is critical in comparisons across different eras of the league. Player achievements, team performance, and records are concrete benchmarks of NBA legacy. For example, iconic achievements, like Wilt Chamberlain's 100-point game or 1995-1996 Chicago Bulls' 72-win season, all occurred during an 82-game season. Changing the number of games carries the risk of inflating or diminishing the value of modern-day accomplishments compared to historical feats.

Another reason is that the finances surrounding the NBA are heavily dependent on the fixed structure of the 82-game schedule. Teams plan their revenue streams, such as ticket sales and merchandise, according to this established pattern. Broadcasters and advertisers also benefit from this stability. For example, marquee matchups are often scheduled to draw significant viewership, and a reduction in games could disrupt these long-standing arrangements. This would impact stakeholders across the league.

Finally, the format gives the league a robust number of games for fairness in standings and in turn, playoff qualification. A smaller schedule could lead to greater variance in outcomes, benefiting or penalizing teams due to short streaks instead of long-term, sustained performance. Although an 82-game schedule is long, it strikes a balance between showing the athletic endurance and ability of the players and obtaining an adequate number of data points to determine the best teams for the postseason.

Based on these factors, we want the 82-game schedule to remain a cornerstone of the NBA's identity. Any attempts to modify it would require addressing logistical, financial, and historical concerns.

1.4 Conferences? Divisions?

The NBA's current structure of conferences and divisions came about due to geographic considerations. Its goals were to reduce travel and foster regional rivalries. However, this structure creates several challenges as the league grows and does not completely solve travel concerns.

Each team plays 4 games against division opponents, 3 or 4 games against the rest of their conference opponents, and 2 games against out-of-conference opponents. This method of scheduling tries to create more frequent matches against geographically closer teams, but it is still inefficient. For instance, some divisions like the Northwest are very spread out, spanning from Portland to Denver to Minneapolis. This results in a disparity in travel for teams in these divisions compared to those in compact divisions like the Atlantic.

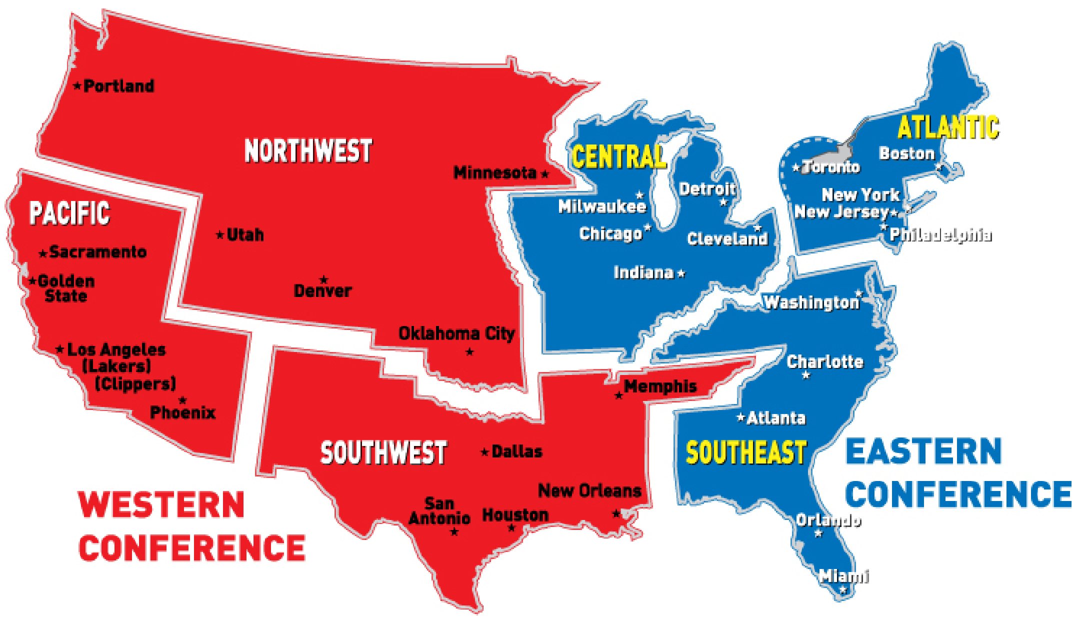

Figure 1: NBA Divisional Map with all NBA cities corresponding to the teams.

On top of this, the Western Conference as a whole spans a significantly larger geographic area than the Eastern Conference. This compounds travel-related fatigue and costs for Western Conference teams, leading to an uneven playing field. Teams in the Western Conference routinely go through cross-country trips, while their Eastern Conference peers often only have to take shorter, regional road trips. This imbalance could affect player performance, team preparation, and recovery times.

By rethinking the NBA without conferences and divisions, the league would see several benefits. Scheduling would become more equitable: Each team would play a more balanced number of games against all opponents, regardless of their stadium's location. This could create new rivalries or showcase diverse matchups with different playing styles, leading to higher overall quality of play. Additionally, our schedule can minimize travel distances with complete location-based optimization instead of somewhat arbitrary divisional boundaries.

Another advantage of removing divisions and conferences is more fair playoff seeding. Currently, teams in the weaker conference can secure playoff spots despite worse records compared to teams in stronger divisions or conferences. Even currently during the ongoing NBA season, 10 of the 16 best teams (record-wise) are in the Western conference. A no-conference structure would have seeding based solely on win-loss records so that the most deserving teams reach the playoffs.

Overall, while divisions and conferences have served to manage scheduling for decades, their usefulness in today's NBA is increasingly questionable. Advances in operations research allow us to prioritize fairness, balance, and sustainability in terms of scheduling.

2 IP Formulation

We enlist an Integer Programming (IP) approach to solve this minimization problem.

2.1 Mathematical Formulation

Sets and Indices:
T : Teams, $|T| = 30$ D : Game Days, $|D| = 190$
$\mathbf{d}_{ij}$: Distance matrix between teams i, j

Decision Variables:
$x_{dij} \in 0, 1$ (1 if home team i plays team j on day d)
$l_{dt} \in 0, \ldots, 29$ (Location of team t on day d)

Objective Function:
$\min \sum_{d \in D} \sum_{t \in T} \text{BackToBack}(t, d) + \alpha \sum_{d \in D} \sum_{t \in T} \text{TravelDistance}(t, d)$

Constraints:
1. Games per Team: $\sum_{d \in D} \sum_{j \in T, j \neq i} (x_{dij} + x_{dji}) = 82 \quad \forall i \in T$
2. Home/Away Balance: $\sum_{d \in D} \sum_{j \in T, j \neq i} x_{dij} = 41 \quad \forall i \in T$
3. One Game per Team per Day: $\sum_{j \in T, j \neq i} (x_{dij} + x_{dji}) \leq 1 \quad \forall i \in T, d \in D$
4. Inter-Team Matchups: $\sum_{d \in D} (x_{dij} + x_{dji}) \geq 2 \quad \forall i, j \in T, i \neq j$
5. Initial Location: $l_{0t} = t \quad \forall t \in T$
6. Location Update Rules:
   a. Home Game: $x_{dij} = 1 \implies l_{di} = i$
   b. Away Game: $x_{dji} = 1 \implies l_{di} = j$
   c. No Game: $l_{dt} = l_{d-1,t}$

2.2 Objective Function: Balancing Back-to-Back Games and Travel Distance

The objective function is the core of our scheduling optimization problem. It integrates our two primary goals of minimizing back-to-back games and reducing overall travel distance. Striking the right balance between these goals requires a nuanced model formulation.

Minimizing Back-to-Back Games: Back-to-back games, in which teams play consecutive days, are a significant concern for the health and performance of players. The league has already been trying to solve this problem. To address this, the objective function penalizes instances where a team is scheduled to play on both day d and day $d + 1$. This penalty makes the optimization algorithm prioritize the spacing of games whenever possible. By reducing back-to-back games, we not only mitigate the risk of injuries, but also enhance the quality of the games. Teams are better rested, leading to improved on-court performances, which in turn result in a better viewing experience for fans.

Minimizing Total Travel Distance: Travel distance is another important part of NBA scheduling. Extensive travel places physical, mental, and financial burdens on teams. The objective function incorporates the travel distance between consecutive games by using a precomputed distance matrix. This matrix holds the travel distances between any two team locations. The function then sums these distances over the entire season for all teams.

Weight Parameter (α): Balancing these two goals, reducing back-to-back games and minimizing travel, requires a weight parameter α in the objective function. A higher α prioritizes minimizing travel distance, while a lower α prioritizes reducing back-to-back games. This tunable parameter allows for flexibility in addressing the league's changing concerns. For example, if player health becomes a more pressing concern, α can be shifted to more heavily penalize back-to-back scheduling. Our current value of α is 1/3000.

3 Code Implementation

The following section describes how we actually implement our NBA scheduling optimization model. The `ortools.sat.python` package and its Constraint Programming (CP) model make up the framework of our code.

3.1 Distance Matrix

The distance matrix is an essential part of the NBA scheduling optimization model. Using the geopy package, we calculate the travel distances between the home stadiums of each team. A custom distance function is created to compute the distance between any two team stadiums based on their coordinates.

This distance function is then applied to generate a complete distance matrix that stores the distances between all pairs of teams. To integrate this into the optimization model more efficiently, the matrix is flattened, reducing the number of parameters for the integer programming (IP) solver compared to using a 2D matrix. This simplification allows for more efficient computation while still accounting for travel distance in the scheduling model.

$$\begin{pmatrix} d_{0,0} & d_{0,1} & \cdots & d_{0,29} \\ d_{1,0} & d_{1,1} & \cdots & d_{1,29} \\ \vdots & \vdots & \ddots & \vdots \\ d_{29,0} & d_{29,1} & \cdots & d_{29,29} \end{pmatrix}$$

3.2 Initialization of the Model

We begin by importing the necessary libraries and defining constants for the problem. The constants include:

- `NUM_TEAMS = 30`: Representing the 30 NBA teams.
- `NUM_GAMES = 82`: Defining the total number of games each team plays.
- `NUM_GAME_DAYS = 190`: Allocating 190 days for the season's games.
- `HOME_AWAY_RATIO = NUM_GAMES / 2`: Keeping an equal split between home and away games.

The `distance_matrix` is precomputed to store pairwise travel distances between teams' home stadiums. This matrix is flattened to optimize its integration with the CP solver.

The `cp_model.CpModel()` initializes the constraint programming model to find a solution.

3.3 Variable Definitions

Two key sets of variables are defined:

- **Schedule Variables (`schedule[day, team1, team2]`)**: Boolean variables indicating whether `team1` (home) plays against `team2` (away) on a specific `day`. These variables enforce the core game-scheduling constraints.
- **Location Variables (`locations[day, team]`)**: Integer variables tracking each team's location on a given day. These enable the calculation of travel distances and enforce location-based constraints.

3.4 Constraints

Several constraints ensure the feasibility and fairness of the schedule:

1. **Total Games Constraint**: Each team plays exactly `NUM_GAMES` games across the season. This is enforced by summing all games that a team plays, both home and away, over all days.

2. **Home/Away Balance**: Each team must play an equal number of home and away games (`HOME_AWAY_RATIO`). This is achieved by summing the `schedule` variables for home games.

3. **Daily Game Limits**: No team can play more than one game on any given day. This prevents scheduling conflicts and overloading of teams.

4. **Initial Locations**: At the start of the schedule, each team is at their home stadium.

5. **Dynamic Location Updates**: Each team's location updates based on the previous day's (possible) game. Boolean variables track whether a team played at home, away, or not at all, and the locations are updated accordingly.

6. **Team Matchup Constraints**: Each pair of teams must play against each other between 2 and 3 times over the course of the season. This constraint ensures a balanced schedule by enforcing:

$$2 \leq \sum \texttt{schedule[day, team1, team2]} + \texttt{schedule[day, team2, team1]} \leq 3,$$

   where the sum is taken over all game days.

3.5 Objective Function

The objective function combines:

- **Minimizing Back-to-Back Games**: Boolean variables track whether a team plays on consecutive days.
- **Minimizing Travel Distance**: Integer variables calculate daily travel distances based on changes in team locations. The distance matrix provides efficient lookups for these calculations.

A weight parameter scales the travel distance penalty to balance it against the back-to-back game penalty, since the total travel distance is much greater than the total number of back-to-back games. The objective is expressed as:

$$\text{Objective} = \sum \texttt{back\_to\_back} + \frac{1}{3000} \sum \texttt{travel\_distance}$$

3.6 Solution Process

The CP model solver is configured with parameters such as presolve and multi-threading to improve computational efficiency. The solver attempts to find an optimal solution within a specified time limit, providing either an optimal, feasible, or infeasible solution (no solution). However, we ran into issues with the system RAM blowing up. This was where the bulk of our work had to be done. We dealt with this with a variety of techniques: using a GPU with a higher RAM limit on Google Colab, reducing the number of variables, simplifying the constraints, splitting our optimization problem into segments of a season, and adjusting the CP model solver's parameters like `num_search_workers`.

3.7 Output and Results

The solution, if feasible, outputs the scheduled games for each day. The results include:

- Teams scheduled for each game.
- Locations of each game.
- Average number of back-to-back games and travel distance per team.

The implementation showed us how powerful constraint programming is at handling complex tasks.

4 Conclusion

We believe our methods have great promise for scheduling and reducing travel distance and back to back games optimally.

Our first approach was partially successful. Our first approach also used 170 game dates. Our model optimized back-to-back games to a point less than the current league average. Our schedule achieved an average of 6.97 back-to-back games per team compared to the current NBA average of 14.9 back-to-back games per team. However, this initial approach was unable to lessen the average total distance traveled. Our model produced an average distance traveled of 68,710 miles per team while the current NBA average is 42,435 miles traveled.

After fixing the RAM blow-up issue, we could run our final model with location tracking. The newly outputted schedule returned an average of 0.8 back-to-back games per team and an average distance traveled of 37,609.9 miles per team. This model improves upon both league averages mentioned above.

Here is a snippet of the final model's schedule:

Table 1: Comparison of Scheduling Approaches Against NBA Averages

Approach	**B2B Games/Team**	**Distance/Team (miles)**
NBA Current Average	14.9	42,435
Initial Model (no location tracking)	6.97	68,710
Final Model (with location tracking)	0.8	37,609.9

Day 131:

- Orlando Magic @ Atlanta Hawks
- Dallas Mavericks @ Brooklyn Nets
- Cleveland Cavaliers @ Chicago Bulls
- Portland Trail Blazers @ Indiana Pacers
- New York Knicks @ Los Angeles Clippers
- Detroit Pistons @ Miami Heat
- Denver Nuggets @ Minnesota Timberwolves
- Houston Rockets @ New Orleans Pelicans
- Memphis Grizzlies @ Oklahoma City Thunder
- Phoenix Suns @ Sacramento Kings
- Philadelphia 76ers @ San Antonio Spurs
- Milwaukee Bucks @ Toronto Raptors
- Boston Celtics @ Washington Wizards

Day 132:

- Portland Trail Blazers @ Charlotte Hornets
- Utah Jazz @ Oklahoma City Thunder

Day 133:

- Dallas Mavericks @ Atlanta Hawks

- Charlotte Hornets @ Brooklyn Nets
- Boston Celtics @ Cleveland Cavaliers
- Toronto Raptors @ Detroit Pistons
- Houston Rockets @ Los Angeles Lakers
- Orlando Magic @ Miami Heat
- Indiana Pacers @ Milwaukee Bucks
- Memphis Grizzlies @ Minnesota Timberwolves
- Portland Trail Blazers @ New York Knicks
- Denver Nuggets @ Phoenix Suns
- New Orleans Pelicans @ San Antonio Spurs
- Los Angeles Clippers @ Washington Wizards

Day 134:

- Utah Jazz @ Dallas Mavericks
- Chicago Bulls @ Milwaukee Bucks
- Orlando Magic @ Sacramento Kings

Day 135:

- Los Angeles Clippers @ Brooklyn Nets
- Detroit Pistons @ Charlotte Hornets
- New Orleans Pelicans @ Dallas Mavericks
- Houston Rockets @ Denver Nuggets
- Milwaukee Bucks @ Memphis Grizzlies
- Portland Trail Blazers @ Miami Heat
- Toronto Raptors @ New York Knicks
- Oklahoma City Thunder @ Washington Wizards

Our location tracking approach, although computationally more expensive than typical scheduling methods, provides accurate distance tracking as teams do not always leave from their home city and are traveling from the previous city they played their game in. Our approach also strays away from finding patterns for home and away games for each division as the divisions are not always spatially efficient to determine game schedules. Possible future directions of research include developing the schedule in smaller parts to avoid heavy computation. We could also explore eliminating back-to-back games completely by analyzing the trade off between travel distance and games played consecutively closely.

5 Appendix

Code and Schedule Solutions: https://github.com/a-ghose/NBAScheduling

6 References:

1. National Basketball Association, Wikipedia, retrived from "https://en.wikipedia.org/wiki/National_Basketba

2. Kirkwood, Chris. "NBA Standings Explained." The NBA Explained, 5 Mar. 2015, https://nbaexplained.word standings-explained/.

3. Kevin. "Analyzing the Impact of Back-to-Backs on NBA Team Performance." Medium, 6 Mar. 2024, https://medium.com/@kevhopper92/analyzing-the-impact-of-back-to-backs-on-nba-team-performance-88a4b80a5d99.

4. Uitti, Jacob. "The World Is on Fire - and the NBA Wants to Be Part of the Solution." The Guardian, 30 Apr. 2024. The Guardian, https://www.theguardian.com/sport/2024/apr/30/nba-green-climate-initiative-carbon-footprint-environment.