University at Buffalo
The State University of New York

Department of Industrial and Systems Engineering
School of Engineering and Applied Sciences

December 12, 2014

Dear 2015 MIT Sloan Sports Analytics Conference Review Committee,

Attached, please find our research paper "Alleviating Competitive Imbalance in NFL Schedules: An Integer-Programming Approach" for exclusive consideration as a candidate for the final stage of the research paper competition. Also enclosed, please find a separate, extended version of the paper as an appendix including all mathematical formulations and analytical details, and a supplementary document highlighting the complaints on imbalanced features of the last five years' NFL schedules from players, coaches, managers and owners along with the links to the media they appeared. Excluding me, the other co-authors of the paper are Dr. Mark Karwan, Mr. Niraj Kumar Pandey and Mr. Kyle Cunningham.

Our paper proposes a novel integer-programming approach to an emerging but vexing sports scheduling application which has not been considered from an optimization point of view earlier. To alleviate competitive disadvantages mainly from the differences in teams' rest durations between games, we develop integer programming formulations for the scheduling of the National Football League's regular season games. Our formulations also consider balancing several other features of the schedule such as the distribution of more important games across the season and the elimination of long undesirable streaks of home-away games. We devise a decomposition-based solution approach to create a desirable volume of more competitively balanced schedules in reasonable amount of time and empirically validate its stability in generating competitively more appealing schedules on past and numerous possible future NFL seasons. As an appendix to the paper we highlight the major imbalanced features of the 2014 NFL schedule and attached one of the alternative 2014 schedules that our method creates.

As the anomalies that have been observed in current and past NFL schedules lead to dissatisfaction from several franchises which receive national media's attention frequently, our paper should be of interest to a broad audience including but not limited to the league's scheduling department, professionals from the sports industry, leading sports writers, researchers working at the intersection of optimization and sports analytics, and football fans.

We confirm that our paper is an original study conducted at the State University of New York at Buffalo.

Please address all correspondence concerning this submission to me at the University at Buffalo and do not hesitate to contact me at the e-mail address or phone number below.

Thank you for your time and consideration.

Sincerely,

Murat Kurt
Assistant Professor
Department of Industrial and Systems Engineering
The State University of New York at Buffalo
309 Bell Hall, North Campus, Buffalo, NY-14260
Phone: (716) 645 4716
E-mail: muratkur@buffalo.edu

Alleviating Competitive Imbalance in NFL Schedules: An Integer-Programming Approach

Abstract

The National Football League (NFL) uses numerous complex rules in scheduling regular season games to maintain fairness, attractiveness and its wide appeal to all fans and franchises. While these rules balance a majority of the features, they are not robust in spacing games to avoid competitive imbalance. We consider the scheduling of NFL regular season games and formulate a mixed-integer linear program (MILP) to alleviate competitive disadvantages originating from the assignment of bye-weeks, Thursday games and streaks of home-away games among various other sources. We propose a two-phase heuristic approach to seek solutions to the resulting large-scale MILP and conduct computational experiments to illustrate how past NFL schedules could have been improved for fairness. We also demonstrate the efficiency and stability of our approach by creating balanced schedules on an extensive set of simulated possible future NFL seasons. Our experiments show that the heuristic can quickly create a large pool of schedules that are completely free of disadvantages due to scheduling of bye-weeks and well-balanced in preparation time differences due to Thursday games.

1. Introduction

The NFL is currently the highest revenue generating sports league in the world and grossed more than $9.5 billion revenue in 2013 [1]. During a regular season (or "season" in the remainder of the paper), weekly NFL games are primarily scheduled on Sundays with the exception of two or three additional games on Thursday and Monday nights of the same week. In a seventeen-week season each NFL team plays sixteen games in total and rests one week, which is commonly referred to as a "bye-week". A Thursday game provides a team with three/four days of extra rest and practice time while a bye-week provides the team a whole extra week to prepare for the following week's opponent.

Although the NFL's scheduling routine has become sophisticated in terms of its ability to optimize complex logistics issues and consider the fairness of the schedules to all of the participant franchises, every year when the precise dates of the upcoming season games are released, several teams believe their schedules lack competitive balance. Recently, after the release of the 2013 schedule, there was a great debate about the fairness in the assignments of bye-weeks and Thursday games.

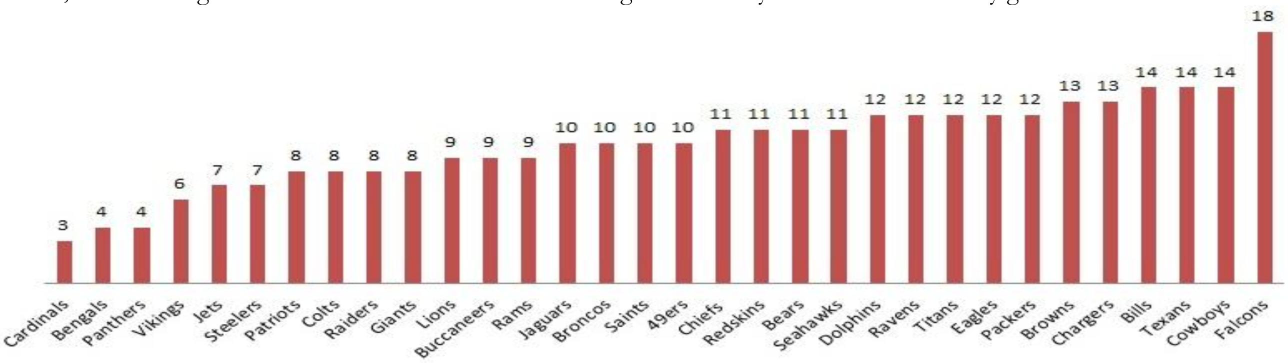

Figure 1 - Total number of games played by each NFL team against teams coming off a bye-week between 2003 and 2012

The 2013 season schedule was particularly imbalanced in the number of games that each team was playing against more-rested opponents. In 2013, the Buffalo Bills led the NFL with 5 games against teams either coming off a Thursday game or a bye from the previous week. Dissatisfied by their 2013 season schedule, the Bills raised a complaint on their official website stating "It's very difficult to call the NFL a league of parity when there's one team with half of their division games against clubs with extra time to rest and prepare, while another in the same division has none. The league simply has to do better" [2]. The number of games against more-rested opponents was not all that different for the Bills in 2012 either, as they had 4 such games in a span of 5 weeks and lost 3 of those 4 games. Following the Bills, the Atlanta Falcons were the second in the league in 2013 with 4 games against more-rested opponents. While the Bills and the Falcons had multiple games against more-rested opponents, there were fourteen teams which had only 1 such game and 3 teams had none during the 2013 season. When questioned about this scheduling quirk, the NFL claimed that such an anomaly does not occur frequently, and over many years of scheduling occurrences with similar extreme disadvantages are equally distributed among all NFL teams. Figure 1 depicts the total number of games that each team played against teams with an advantage of extra rest due to bye-weeks from 2003 to 2012 during which a total of 320 such games were played. The numbers in Figure 1 do not only conflict with the NFL's claim on the uniformity of the distribution of games against more-rested opponents across the league over the years, but with the figures in Table 1, they also provide empirical evidence that they affect the outcomes of the games.

In Table 1, the center column represents the average of seasonal win percentages of teams which played against more-rested opponents. The right column represents the average of the win percentages of such teams only against their opponents which were coming off extra rest from the previous week. While the average win percentage goes down by 3.77% when facing a team with extra rest, over these five years, particularly in years 2010 and 2013, the negative effects of playing against more-rested opponents were more dramatic and adversely affected the league's competitive balance. The competitive imbalance of the schedule and complexity of season scheduling calls for advanced scheduling techniques.

Year	Average win % against all opponents	Average win % against rested opponents
2009	44.9	46.9
2010	**51.9**	**37.5**
2011	50.3	50.0
2012	47.3	48.4
2013	**47.9**	**40.6**
Average	48.5	44.7

Table 1 - Comparison of win percentages of teams which played against teams coming off their bye-weeks with extra rest for seasons 2009 through 2013

The existing research on NFL is fledgling compared to the vast body of research on scheduling in other sports and attempts to answer various questions other than refining the league's scheduling problems, including the effects of draft system on team performance, realignment of divisions for shorter travel distances, West Coast teams' disadvantages when playing on the East Coast, physiology of the athletes and learning/in-game strategy [3, 4, 5, 6]. The only work that proposes scheduling methods for NFL games is Dilkina and Havens (2004) [7], in which the authors first describe the highly constrained problem faced by the NFL to schedule television broadcasts and then develop a prototype scheduling system using constraint programming. Our work differs from Dilkina and Havens (2004) in that our objective is to improve the competitive balance and fairness of the NFL's season schedules, whereas Dilkina and Havens (2004) focuses primarily on the assignment of games to television broadcasts.

In this paper, we study an emerging but vexing sports scheduling problem, particularly the improvement of the competitive balance of the NFL's season schedules, which has not previously been considered from an optimization point of view. With an objective of minimizing the maximum number of games that any team plays against more-rested opponents, we develop a mixed-integer linear programming (MILP) model under restrictions pertinent to bye-week and Thursday game assignments, streaks of home and away games, and the distribution of divisional games across weeks, among several others. We further minimize the number of games in which a team is at a disadvantage of playing with less rest against an opponent that has a Thursday game in the previous week and improve the schedule by avoiding any long streaks or undesirable patterns of home and away games, and distributing divisional games across weeks towards a competitively appealing schedule.

2. The Optimization Model and Solution Approach

In this section we describe the problem of scheduling the games of an NFL season. First, we explain the NFL's schedule format, which has been in use since the division realignment in 2002. Then, we present a list of scheduling rules that we formulate by means of mixed-integer linear inequalities and devise a heuristic which seeks solutions to the resulting optimization model.

2.1. The Schedule Format: The NFL is comprised of thirty two team franchises in two conferences: The American Football Conference (AFC) and the National Football Conference (NFC). Each conference is further broken into four separate divisions as East, West, North and South, containing four teams each. In a season each team plays two games against each of the three teams from its own division, which are referred to as the "divisional opponents" of that particular team. The remaining ten games in each team's schedule are against teams from other divisions. Specifically, each team from the same division plays one game against all the teams in a particular division from the same conference and one game against all the teams in a particular division from the other conference. The pairs of divisions from the same (different) conference which are scheduled to play against each other are called Intra-Conference divisions (IaCD) (Inter-Conference divisions or shortly IeCD). Each team plays its remaining two games against two teams which finished the previous season at the same standing in their respective divisions as the team did in its own division. We refer to these teams as the "standing-based opponents". Because each team plays against all teams in two divisions from its own conference (its own division and its IaCD), the standing-based opponents are chosen from the other two divisions in that team's conference. Although the NFL's scheduling format specifies the distribution of games among the opponents of each team, it does not reveal the particular weeks and the dates of games. In Table 2, we present the list of NFL teams with their respective divisions and conferences, 2012 divisional standings, and the breakdown of the Indianapolis Colts' 2013 season games among their opponents.

2.2. Scheduling Rules: In this section, we present the list of rules that are currently imposed by the NFL along with additional rules added by us to the model to make schedules fairer and more attractive than the current NFL schedules. The rules we additionally impose on the schedule are based on the literature and our analyses of the previous years' schedules. We specify the extra rules we impose by *.

Conference	2012 Standing	Divisions: East	North	South	West
AFC	1	Patriots	Ravens	Texans	Broncos
	2	Dolphins	Bengals	Colts	Chargers
	3	Jets	Steelers	Titans	Raiders
	4	Bills	Browns	Jaguars	Chiefs
NFC	1	Redskins	Packers	Falcons	49ers
	2	Giants	Vikings	Panthers	Seahawks
	3	Cowboys	Bears	Saints	Rams
	4	Eagles	Lions	Buccaneers	Cardinals

Legend
Division Opponent
IaCD Opponent
IeCD Opponent
Standing-based Opponent

Table 2 - The list of NFL teams with their respective divisions highlighting the 2013 season opponents of the Colts

R1: Every team should play each of its divisional opponents exactly twice in a season.
R2: Every team should play each of its non-divisional opponents exactly once in a season.
R3: Each team should have a bye during a week within the NFL-mandated bye-week window (weeks 4 through 12, both inclusive).
R4: For any given week within the bye-week window, there should be no more than six and no less than two teams on bye.
R5*-Bye-Week Fairness Rule: To eliminate all competitive imbalances due to bye-weeks, for any game, each of the participating opponents should either play the previous week or have a bye.

Divisional games are vital towards the season-record of a team as they are primary tie-breakers for a berth in playoffs or a division title against its divisional opponents. In late-season games, teams clinching playoff spots may bench their key players for the playoffs rather than playing them to win, compromising the attractiveness and integrity of the NFL games. Resting key players in late-season games has been a major concern for the NFL Competition Committee. In 2010, Roger Goodell, the commissioner of the NFL, has suggested scheduling as many divisional games as possible in the last three weeks of the season [8]. In line with the commissioner's suggestions, the NFL has been scheduling the last week of the season exclusively as divisional matchups since 2011. We impose rules R6-R8 to keep the competition among the divisional opponents alive throughout the season.

R6*: Games between divisional opponents should be separated by at least two weeks.
R7*: A team should play each of its divisional opponents at least once in the second half of the season, i.e., between weeks 9 and 17, inclusive.
R8*: Dividing the season into a set of predetermined stretches of weeks, each team should play at least a certain number of divisional games in each of these stretches. For instance, consistent with the NFL, one of these stretches may consist of only the last week of the season, while other stretches may deal with the spread of the divisional games in prior weeks.

Based on the NFL's current scheduling format, rules R9-R13 specify the distribution of home and away games to each team.

R9: Each team should play each of its divisional opponents once at home and once on the road.
R10: Each team should play a game against each of the four teams from its IaCD. Two of these games should be at home and the other two should be on the road. Similarly, each team should also play a game against each of the four teams from its IeCD. Again, two of these games should be at home and the other two should be on the road.
R11: Each team should play one of its standing-based opponents at home and the other on the road.
R12-West Coast Modification Rule: The Patriots and the Jets had to make four cross-country trips to the West Coast during the 2008 season when they were scheduled to play all AFC West and NFC West teams. Starting with the 2010 season, to alleviate the affliction of such trips, the NFL modified its scheduling format in that the teams that are scheduled to play the AFC West (NFC West) as their IaCD or IeCD should face either the Raiders or the Chargers (the 49ers or the Seahawks) on the road.
R13: The Kickoff game of the season should be hosted by the reigning Super Bowl champions.

The disadvantages of away games are amplified when teams are assigned to streaks of two or three road games. Such streaks are considered undesirable widely by the teams, media and fans as they add extra burden to teams' schedules. Between the 1990 and 2010 NFL seasons, there have been 110 occurrences of a streak of three road games and in 65% of them teams that were on such a streak lost two or all of their games. Our analyses of earlier schedules show that no team had more than a single occurrence of a streak of three road games or more than three separate occurrences of a streak of two road games in a single season. In accord with our analyses, rules R14 and R15 aim to avoid the aforementioned undesirable road game patterns. In the remainder of the paper we let RS_2 and RS_3 refer to streaks of two and three road games, respectively.

R14*: No team should have more than three occurrences of possibly overlapping RS_2s. Note that based on the definition of an RS_2, an RS_3 starting from a week is equivalent to two overlapping RS_2s: One starting from that particular week and the other starting from the week after. Therefore, this rule avoids multiple occurrences of an RS_2 along with an RS_3 as well.
R15*: No team should have more than a single occurrence of an RS_3.

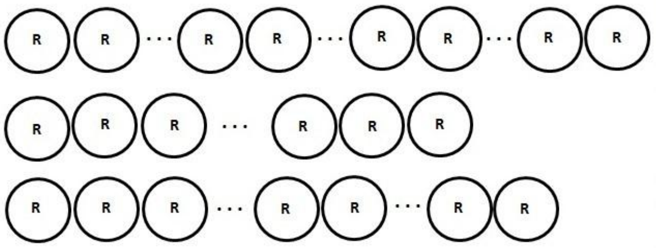

Figure 2 - Disallowed major road game patterns

R16*: To reduce any disadvantage due to long road game streaks at the end of a season, no team should end the season with an RS_3. Likewise, as home games attract higher attendance in early season, no team should begin the season with an RS_3.
R17*: To avoid any long streaks of home games each team should play at least one road game in any four consecutive weeks. Likewise, so as not to have any team far from its home fan base for a long period of time, each team should play at least one home game in any four consecutive weeks.

Starting with the 2012 season, the league instated a rule that each team should play a Thursday game during a season giving rise to the number of Thursday games. We formally state this change by rules R18 and R19 below.

R18: In each of the NFL-mandated Thursday game weeks (weeks 1 through 15 both inclusive), except for Thanksgiving week, there should be one Thursday game.
R19: The Kickoff game should be scheduled to Thursday and excluding the Kickoff game, each team should play one Thursday game during the season. As the reigning Super Bowl champions and their first week opponents open the season by a Thursday game, their second Thursday game should not be scheduled to the second week.
R20: By tradition, both the Lions and the Cowboys should host a Thursday game in Thanksgiving week.
R21*: Since the 2006 season, the NFL has been scheduling a third game on a Thanksgiving Thursday with no fixed opponents. Although not required by the NFL, for even representation of each conference on Thanksgiving Thursday we impose the participating teams of these three games should comprise of three AFC and three NFC teams.

Besides reduced preparation and rest time, the effects of Thursday games are intensified by the long flights that a team has to take for a Thursday game on the road after a Sunday matchup. In particular, teams are referred to as being on a "short week" when they play a Thursday game following a non-Thursday game. The assignment of opponents to short week matchups has been a source of concern and dissatisfaction among teams and players [9, 10, 11].

R22*: To avoid any long trip in a short week, no team's Thursday game should require more than 3.5 hours of flight time.
R23*: Although the NFL tries not to send any team to back-to-back road games in a short week, in 2013, the Steelers had to travel for their Thursday game after coming from a road game on Sunday. Therefore to lessen the burden of a short week, we avoid two travel obligations in a short week, i.e., no team should play twice on the road in a short week.

Similar to the NFL, we pair the teams that are in close proximity as "same area teams". In particular, we refer to the Jets and the Giants as "New York City teams", and the Raiders and the 49ers as "Bay Area teams". Rules R24-R27 below enforce restrictions on the games of the same area teams due to coverage of their fan base and possible conflicts in their shared media markets.

R24: No pair of the same area teams should simultaneously have a bye in the same week.
R25*: Because the Giants and the Jets share the same stadium for their games, when they simultaneously host other teams in the same week, one of the games should be on Thursday to give stadium management enough preparation time for the later game.
R26*: By rule R19 each team typically plays a single Thursday game in a season. Therefore, the New York City teams can simultaneously host other teams on at most two occasions but only within the NFL-mandated Thursday game weeks in a season.
R27*: Based on our analyses of recent schedules, for the regularity of having a game in the area, the number of occasions the Bay Area teams simultaneously play on the road should be three or less in a season.

We develop an MILP formulation of the scheduling rules we describe and refer it to as the Full Model (FM) (the details of the formulation are provided in the full paper Appendix). One of the main challenges in highly constrained sport scheduling models is the scarcity of feasible solutions and the difficulty in their characterization whereas the actual schedules that are implemented by the decision-makers are mostly chosen from a variety of options after weighing different criteria. Therefore, creating a pool of alternatives (as we do below with a two phase model) is more essential to the nature of the schedule selection process as it enables the incorporation of secondary objectives into the selection process and allows the schedules to be compared on a diverse set of characteristics. Our model balances several features of an NFL season schedule by hard-constraints. On top of these constraints, we consider the minimization of the following four main criteria:

(C1): Across the whole league, the maximum number of games that a team plays against more-rested opponents. By the bye-week fairness constraints, a team may play against a more-rested opponent only if the opponent comes off a Thursday game from the previous week while the team itself does not. Therefore, for each team, we call the number of games that is played against a more-rested opponent a "Thursday effect" and abbreviate the "league's maximum Thursday effect" across all teams as LMTE.
(C2): The number of teams with the LMTE.
(C3): The number of teams playing three straight road games, i.e. teams with an RS_3.
(C4): The number of teams playing three separate blocks of back-to-back road games, i.e. teams with three separate RS_2.

While an ideal schedule is desired to minimize each of the four criteria above, due to their conflicting nature the trade-offs among them do not render their simultaneous minimization possible. Our objective is to generate as many feasible schedules as possible and refine them based on a lexicographic priority order of the four main criteria above as (C1), (C2) and (C3)-(C4) (with no priority of (C3) over (C4) and vice-versa).

Another desirable feature of a scheduling process is its scalability to generate alternatives in a reasonable amount of time. Faster generation of schedules does not only save computational resources but also enables the schedule-makers to design more robust schedules by facilitating a faster analysis of the sensitivity of their solutions with respect to critical parameters of the process.

Due to the very large number of binary variables and constraints, our computational experiments indicate that when the FM was run for the 2012-2014 NFL seasons with no explicit objective it was able to produce only a few feasible schedules after 100 hours of computation time. To overcome the intractability due to scale of the FM we devise a heuristic, which we refer to as the two-phase method (or model) (shortly as TPM). The TPM explores solutions to the NFL's season scheduling problem by partitioning the FM into two simpler mathematical programs and solving them in sequence. Specifically, the first phase of the TPM assigns the games to weeks with no specification on their venues and identifies the bye-weeks for each team. Then the second phase determines the host of each game subject to rules on R9-R17 and the set(s) of teams for each particular week that will play the Thursday game(s) subject to rules R18-R23 among others (a detailed breakdown of the FM into two phases with explanations and interpretations of variables and constraints are provided in the full paper Appendix).

Splitting the FM into two phases brings computational efficiency due to reduction in the size of the model and the interaction between the two phases. Because each phase of the TPM is concerned with only a subset of rules and the second phase exploits the outputs of the first phase as parameters, the decision variables of the mathematical programs in the TPM do not need to carry as much information as those of the FM and the constraints in each phase can be described with fewer decision variables. Due to fixed weekly matchups from the first-phase, certain constraints are needed only for weeks in which two teams are scheduled to play against each other whereas the FM requires those constraints for all weeks as it also seeks weeks for each particular matchup. Hence, the mathematical programs in each phase include substantially less binary variables and constraints than the FM. Specifically, the FM has 15689 constraints and 8546 variables, 7520 of which are binary. On the other hand, the first phase of the TPM has 5451 constraints and 3536 binary variables, and the second phase consists of 3514 constraints, 1024 binary variables and 1122 continuous variables. Since each phase includes substantially less binary variables and constraints than the FM, the TPM enhances the computational efficiency of the solution process which in turn helps achieving a diverse set of schedules.

Based on the priority order of criteria (C1), (C2) and (C3)-(C4), after solving the second phase MILP for the minimum possible LMTE, the set of optimal solutions can be refined by solving another MILP which minimizes the number of teams with the resulting value of LMTE. However, based on the volume of solutions generated in the first-phase, this refinement procedure may negatively affect the computational performance and the appeal of the TPM as it requires optimizing an additional second-phase MILP for each schedule that has the smallest value for LMTE. Therefore, we develop side relations which we make use of to show that the optimal value of LMTE and the minimum number of teams with the optimal value of the LMTE can be simultaneously achieved without having to solve the second-phase MILP twice (a detailed explanation of the side relations and a mathematical proof of how criteria (C1) and (C2) are simultaneously minimized with the help of those side relations are provided in the full paper Appendix).

3. Implementation of the Two Phase Model (TPM)

In this section we present computational results from the implementation of the TPM on past NFL seasons. We also evaluate the stability of the results and the computational efficiency of the TPM on an extensive set of problem instances by simulating possible future NFL seasons. In our experiments, for the distribution of divisional games across weeks, we divide a season into four stretches of weeks as $\{1,\ldots,5\}, \{1,\ldots,8\}, \{13,\ldots,16\}$ and $\{17\}$. We require each team to play at least one divisional game in stretches 1, 3 and 4 and at least two divisional games in stretch 2. We solved all mathematical models using IBM-CPLEX 12.5 on an Intel Xeon PC with 2.0 GHz CPU and 32 GB RAM.

For each actual and simulated season we followed a series of steps to generate a list of schedules with the minimal competitive imbalance due to Thursday effects and the fewest number of teams with undesirable road game patterns. We employ an objective which maximizes the total number of divisional games in the second half of the season, we iteratively invoke CPLEX's populate procedure to create a first-phase solution pool and gather all distinct optimal solutions in the first-phase solution pool. For each solution coming out of the first-phase, we ran the second-phase MILP so as to generate schedules with the minimum possible LMTE and the minimum number of teams that are exposed to the resulting optimal value of the LMTE and gathered the resulting schedules in the second-phase solution pool. Since the second-phase generates final schedules that are optimal with respect to the fixed weekly matchups from the first-phase, different first-phase solutions may lead to schedules differing in LMTE and the number of teams with the resulting LMTE. Therefore, we refined the second-phase solution pool in three stages with respect to the lexicographic order of the criteria (C1)-(C4). First, we built a secondary pool consisting of final schedules with the least value for the LMTE. Then, we filtered the schedules of this secondary pool that have the minimum number of teams with the LMTE into a third pool. Because the schedules in the third pool differ only in criteria (C3)-(C4) and those with fewer values for each of these criteria are more desirable, we formed a final pool of schedules that are non-dominated with respect to criteria (C3)-(C4) in the third pool. We provide a schematic representation of this refinement process in Figure 3 and in the rest of this section we will present the major characteristics of the schedules in the final pool.

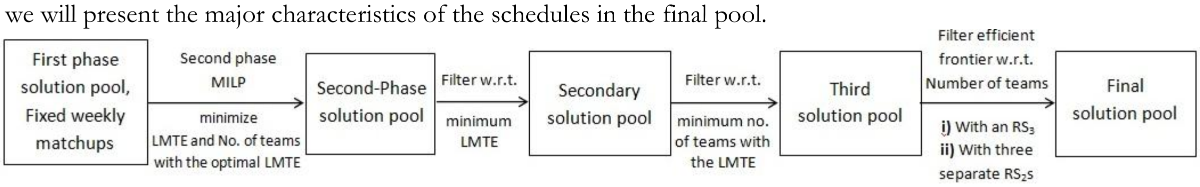

Figure 3 - Demonstration of the solution filtering process

3.1. Comparison with the Actual NFL Schedules: Given the NFL has been scheduling a Thursday game for each team since the 2012 season, using fixed, season-specific opponent lists for each team, we illustrate how the TPM could balance the schedules that are actually employed by the NFL for the 2012-2014 seasons. Our experiments demonstrate that the schedules generated by the TPM outperform the actual NFL schedules in distributing games against more-rested opponents to teams. We compare the performances of the actual NFL schedules in the four major criteria to those from the TPM in the table below.

	Season	LMTE	Number of Teams with LMTE	Max Bye-week Effect	Corresponding Number of Teams	Max Combined Effect	Corresponding Number of Teams	Total RS_3	Total Triple RS_2s
NFL	2012	3	2	3	1	5	1	2	8
	2013	3	1	3	2	5	1	2	7
	2014	2	5	2	5	4	1	4	3
	Season								
TPM	2012	2	4	-	-	2	4	6	14
	2013	2	5	-	-	2	5	2	10
	2014	2	6	-	-	2	6	2	6

Table 3 - Major features of the actual NFL schedules and a sample of schedules from the TPM for 2012-2014 seasons

Similar to the definition of a Thursday effect, we call each game that is played against an opponent coming off a bye a "bye-week effect" on the less-rested team. In Table 3, the fourth column represents the maximum bye-week effect across the whole league and the fifth column represents the number of teams that are exposed to that extreme effect. The sixth column represents the maximum combined effect, i.e., maximum number of games played against more-rested opponents (those coming off a bye or a Thursday game) and the seventh column denotes the corresponding number of teams with the maximum combined effect.

Unlike the actual schedules, the schedules achieved by the TPM are devoid of any bye-week effects and lower the LMTE to almost the minimum attainable value. While the maximum combined effect across all teams was no less than 4 in the actual NFL schedules, our approach was able to generate schedules where no team plays more than 2 games against more-rested opponents. Moreover, a bye-week effect is more disadvantageous than a Thursday effect for a team by providing longer preparation time for the opponent. Due to absence of the bye-week effects in our schedules, our combined effects consist of only Thursday effects and hence provide better balanced schedules.

We also investigated the flexibility of the schedules released by the NFL and sought more balanced Thursday game assignments for the 2012-2014 seasons using a subset of the second-phase constraints, that are mainly associated with the Thursday game assignments. Given the fixed weekly matchups and their host-visitor information from the NFL, there was no assignment of Thursday games to weeks satisfying all constraints on scheduling of Thursday and short week games. From Table 3 as the TPM could clearly generate weekly matchups which admit feasible Thursday game assignments for the 2012-2104 seasons, the inflexibility of the actual NFL schedules reveals the importance of the assignment of games to weeks for avoiding disadvantages due to short week-related travel obligations and rest-durations. It also verifies the amenability of a two-phase approach such as ours in seeking solutions to the problem.

3.2. Sensitivity Analyses: Unlike a round-robin format, the NFL determines the non-divisional matchups between teams through a formula using teams' previous year standings, and their IaCDs and IeCDs. For each division because there are three possible IaCDs and four possible IeCDs to play against, the IaCD and IeCD of a division change every season based on a rotating cycle of three and four years, respectively, to guarantee each team to play against every other team in the league every 12 years. We tested the robustness of the schedules and the stability of the computational performance of the TPM on 100 randomly generated end-of-season divisional standings for each of the twelve scheduling patterns based on which the NFL assigns IaCD and IeCD of a division for a season (the twelve scheduling patterns for the IaCD and IeCD assignments, and a sample cycle of twelve seasons that the scheduling patterns correspond to are presented in the full paper Appendix). Along with the NFL's rotating cycle of IaCD and IeCD assignments, each of the standings that we simulated describes the set of opponents for each team's schedule in a season. Therefore, in the remainder of the paper we will use the terms standing and season interchangeably.

Pattern	Avg Number of Schedules from the first-phase	Max Number of Schedules from the first-phase	25th Percentile of Number of Schedules from the second-phase	Avg Number of Schedules from the second-phase	Max Number of Schedules from the second-phase	Avg. Total Time per Standing (in mins)
1	324	758	42	94	283	24.9
2	324	694	40	89	254	26.3
3	316	692	25	56	155	21.4
4	286	788	27	78	273	22.8
5	317	715	29	76	185	24.0
6	347	669	43	106	277	26.0
7	327	711	36	92	193	18.2
8	385	714	33	82	185	23.2
9	347	667	59	106	260	24.7
10	300	672	46	91	236	21.5
11	285	722	28	79	196	18.4
12	321	752	27	94	242	23.3

Table 4 - Number of schedules generated by the TPM and the associated computation times

The process of populating a pool of first-phase solutions and solving the second-phase MILP for each of them was taking less than half an hour on average per season. In Table 4, we demonstrate the computational efficiency and the performance of the TPM in producing a vast number of solutions from both phases. In all patterns except 4 and 11, the first-phase was able to generate a pool of more than 300 solutions per season. On the other hand, in patterns 4 and 11, the number of solutions from the first-phase was not far from 300. While the maximum number of schedules generated by the first-phase could be in six and seven hundreds, the second-phase was able to find feasible home-away and Thursday game assignments for only about one third of these schedules. The average number of schedules that the TPM was able to generate ranged from 56 to 106 and was highly dependent on the scheduling pattern. The response of the TPM in the number of feasible schedules was also varying significantly with respect to season in that the volume of the schedules created for a season could be at least twice as much as the average number of schedules per season under the same pattern. For instance, in pattern 1, while the TPM generated an average of 94 schedules for a season, the maximum number of schedules it returned across all seasons was 283.

Scheduling Pattern	Number of Teams: 1	2	3	4	5	6	7	8	9	10	11	12
1	0	2	18	54	84	95	96	98	99	99	99	99
2	0	1	16	65	94	96	99	100	100	100	100	100
3	0	1	8	34	71	87	94	95	96	96	96	97
4	0	2	4	28	67	87	92	98	98	98	98	98
5	0	1	5	30	61	82	94	95	97	98	99	99
6	0	2	21	80	96	99	100	100	100	100	100	100
7	0	0	0	26	64	97	99	99	99	99	99	99
8	0	0	0	18	55	65	85	93	97	97	97	97
9	0	2	11	64	90	100	100	100	100	100	100	100
10	0	5	40	80	97	99	99	99	99	99	99	99
11	0	0	1	9	41	77	84	90	98	98	98	98
12	0	0	3	48	81	88	96	100	100	100	100	100

Table 5 - The cumulative number of seasons with respect to the number of teams with LMTE = 2

In Table 5, we demonstrate the performance of the schedules from the TPM in terms of the cumulative distribution of the number of teams with two Thursday effects as the league's maximum. For instance, under scheduling pattern 2, in 94 standings the LMTE was equal to two and the number of teams with the LMTE was less than or equal to five.

In our experiments, the minimum value for the LMTE was no less than two under any standing in any pattern. However, it was possible to generate a schedule with an LMTE = 2 under all standings in four of the patterns. In the remaining eight patterns, the best value of the LMTE was not more than three under any standing. In those patterns, the best value of the LMTE was three in at most three standings but in those standings there was only a single team with three Thursday effects as the league's maximum. When the best possible value of the LMTE was two, the number of teams with the LMTE varied between two and twelve across patterns. In half of the patterns, the number of teams with two Thursday effects as the league's maximum was no more than four in 48% - 80% of the 100 randomly generated standings. Moreover, the number of teams with two Thursday effects as the league's maximum was no more than six in at least 65% of the standings in all patterns, and no more than six in at least 80% of the standings in all patterns except 8 and 11. Throughout all patterns, at least 90% of the standings had schedules that had eight or less teams with two Thursday effects as the league's maximum.

Pattern	Avg RS_3 (Overall)	Avg Triple RS_2s (Overall)	Avg Min RS_3	Corresponding Avg Triple RS_2s	Avg Min Triple RS_2s	Corresponding Avg RS_3
1	5.49	10.32	5.01	10.90	9.73	6.01
2	5.28	10.62	4.64	11.33	9.89	5.90
3	5.42	10.19	4.99	10.73	9.61	5.90
4	5.15	10.20	4.70	10.67	9.76	5.60
5	5.10	10.30	4.74	10.75	9.85	5.47
6	5.35	9.92	4.84	10.40	9.44	5.88
7	5.16	10.94	4.33	12.13	9.75	6.00
8	5.42	10.98	4.75	11.79	10.17	6.08
9	5.19	10.29	4.68	11.04	9.56	5.68
10	5.07	10.64	4.76	10.88	10.32	5.36
11	5.09	10.46	4.64	11.08	9.84	5.52
12	5.34	10.76	5.04	11.20	10.32	5.64

Table 6 - Average performances of schedules in the final pool in terms of major undesirable road game streaks

In Table 6, we present the average performances of the schedules across all seasons with respect to criteria (C3) and (C4). The second and third columns represent the average number of teams with triple RS_2s and average number of teams with an RS_3 in a season, respectively. The fourth column denotes the average of the minimum number of teams with triple RS_2s in a season and the fifth column denotes the corresponding average number of teams with an RS_3 in that season. Likewise, the sixth column denotes the average of the minimum number of teams with an RS_3 and the last column denotes the corresponding average number of teams with triple RS_2s in that season. In all patterns, roughly five teams were scheduled to play three straight road games and about ten teams were to make three separate back-to-back road game trips on average per season. When criteria (C3) and (C4) are prioritized over each other for a further filtering of the schedules in the final pool these averages did not change significantly for either of them. For instance, in pattern 1, the average number of teams playing three straight road games was 5.49 across all schedules, whereas this average was only 5.01 across the schedules with the minimum number of teams that have to play three straight road games and only 6.01 across the schedules that had the minimum number of teams with three separate back-to-back road games.

4. Conclusions

The flaws in the NFL schedules including the assignment of bye-weeks and Thursday games, undesirable home-away game patterns, and travel times in short weeks have been largely unaddressed. Our modeling approach and computational experiments indicate that it is possible to create more balanced schedules in various dimensions while completely eliminating disadvantages due to scheduling of bye-weeks and keeping those due to Thursday games reasonably low for each team. Experiments on past and possible future NFL schedules also demonstrate a desirable volume of high-quality solutions in practical time limits.

As explained in Section 2.1, each team plays all of its IaCD and IeCD opponents in its schedule exactly once during the season on a rotating basis. With an emphasis on cross-season fairness among teams, the NFL also rotates the venues of these games on a separate multi-year basis. More explicitly, the scheduled host of a game between two IaCD and IeCD opponents flips every three and four years, respectively. In contrast to the NFL, our model places greater emphasis on in-season fairness of the schedule by arranging the venues of the non-divisional games flexibly and independent of their past assignments.

Our results elucidate the effects of the NFL's cross-season home-away game scheduling format on schedules' in-season competitive balance, an immediate future research direction is to incorporate the cross-season home-away scheduling scheme on top of the restrictions we propose in this paper. Although our model is flexible enough to accommodate any additional requirements, extra requirements may limit the feasibility and the efficiency of our two-phase approach. Therefore, another promising future research direction is to seek scalable algorithms to cope with the enormity of the resulting model.

5. References

[1] C. Isidore (2013). Why football is still a money machine.
http://money.cnn.com/2013/02/01/news/companies/nfl-money-super-bowl/

[2] C. Brown (2013). NFL scheduling in need of regulations.
http://www.buffalobills.com/news/article-2/NFL-scheduling-in-need-of-regulations/908bf9be-181f-46a6-9b76-c632ea671dbc

[3] V. Billat, G. Py, P. Sirvent, J. Koralsztein, and J. Mercier (2003). The concept of maximal lactate steady state. *Sports Medicine 33* (6), 407-426.

[4] M. Cardinale and C. Bosco (2003). The use of vibration as an exercise intervention. Exercise and Sport Sciences Reviews 31 (1), 3-7.

[5] D. Etzion (2003). Annual vacation: Duration of relief from job stressors and burnout. Anxiety, Stress, and Coping 16 (2), 213-226.

[6] B. Macdonald and W. Pulleyblank (2010). Realignment in the NHL, MLB, the NFL, and the NBA. Technical report, United States Military Academy.

[7] B. N. Dilkina and W. S. Havens (2004). The US National Football League scheduling problem. In Association for the Advancement of Artificial Intelligence, pp. 814-819.

[8] M. Florio (2010). Weeks 16, 17 could feature division games only.
http://profootballtalk.nbcsports.com/2010/03/24/weeks-16-17-could-feature-division-games-only/.

[9] M. Florio (2010). Giants' owner doesn't like Thanksgiving schedule.
http://profootballtalk.nbcsports.com/2009/11/17/giants-owner-doesnt-like-thanksgiving-schedule/

[10] F. Concha (2013). Thursday Night Football: Bad Games, Players At Risk. . . So Here's The Solution.
http://www.sportsgrid.com/nfl/thursday-night-football-bad-games-players-at-risk-heres-solution/.

[11] M. Sando (2011). Jim Harbaugh: 49ers drew short straw.
http://espn.go.com/nfl/story/_/id/7268293/jim-harbaugh-says-49ers-drew-short-straw-game-vs-john-harbaugh-baltimore-ravens

Appendix A: Scheduling Quirks for the 2014 NFL Regular Season

The 2014 schedule released by the NFL in April 2014 has several undesirable features all of which can be eliminated by our model. We list those quirks below based on the family of scheduling rules of our model that they violate.

Violated Rules on Streaks of Home-Away Games

The Steelers have to make 4 separate trips of two-consecutive road games, which is the maximum number of such trips that any team could have during a season as each team plays 8 road games in total. The schedule has 4 teams playing three straight road games. Among these teams, the Buccaneers are also scheduled to play 2 separate streaks of two road games which in turn leave them with only 4 home games in a twelve-week window. While there is no team playing 4 straight home or road games, 3 teams, the Colts, the Patriots and the Raiders, do not have any travel obligations for 4 consecutive weeks due to a combination of a bye and three home games.

Violated Bye-Week Fairness Rules

Though it is impossible to eliminate all disadvantages arising from facing opponents with extra rest due to Thursday games, the schedule has significant disparities in teams' rest durations between games. For instance, there are 6 teams with 3 or more games against more-rested opponents and 21 teams play against opponents coming off their bye-weeks. In particular, both the Chargers and the Bears face two opponents that are coming off their byes, but making these games more critical for these two teams is that they are back-to-back against divisional opponents.

Violated Rules on Spacing and Distribution of Divisional Games

The schedule shows that pushing more divisional games to the last weeks of the season does not leave a fair balance of such games in early and mid-season. During the season, 3 teams, the Cowboys, the Seahawks and the Jets, do not face any of their divisional opponents until week 7. Moreover, because of this early season imbalance in the number of games against divisional opponents, the Seahawks face a major end of season challenge as all but one of their last 6 games are against their divisional opponents. Similar to the Seahawks, the 49ers play half of their divisional games in the last 5 weeks of the season although they do not face the Rams in the second half of the season. Their only non-divisional opponents during this five-week window are the Raiders and the Chargers. On top of this, with the home game in week 12, the 49ers do not leave the West Coast in the last 6 weeks of the season. One possible source of this tidbit for the 49ers is that they play against the Seahawks twice in 3 of the last 6 weeks, which is disallowed in our model. The 49ers and the Seahawks are not the only pair of divisional opponents whose matchups are separated by only a week in that the Cowboys and the Eagles also face each other twice in a span of 3 weeks. In contrast to teams with an end of season schedule mostly loaded by divisional games, there are also teams, e.g. the Packers and the Bills, with only a single divisional game in the last 5 weeks of the season.

Among all divisions, the AFC North has the most imbalanced distribution of divisional games across weeks with a possible influence on an early finish in the divisional supremacy race. The Steelers play both of their divisional games against the Browns by week 6 and the Bengals finish both of their matchups against the Ravens by week 8. Again in one of the most intense rivalries of the NFL, the Steelers and the Ravens do not face each other after week 9. Moreover, the Ravens do not play any divisional game after week 9 until the closing week of the season.

Violated Rules on Thursday/Short Week Games

Despite the flexibility the NFL has gained by adding a third game with no fixed opponents to the Thanksgiving Thursday, unlike previous years the 2014 schedule features only NFC teams in all Thanksgiving Thursday games and lacks the traditional representation of AFC teams. In fact, 2 of those participating teams, the Bears and the Packers, play back-to-back Thursday games which is happening for the first time since the league expanded the limited capacity of the Thursday games. Another flaw of the schedule concerning the competition between divisional opponents is that the Cardinals face the Rams on a Thursday of a short week for which the Cardinals have to fly for more than 4 hours after their Sunday game.

Violated Rules on Same Area Teams

For the regularity of having a game in an area hosting multiple teams there is no explicit rule by the NFL dictating the number of occasions that they simultaneously play on the road. However, the 2014 schedule is sending both of the Bay Area Teams (the 49ers and the Raiders) simultaneously for a road game in 4 weeks which is not only more than the limit we impose but also tied for the most (together with 2010 and 2012 season figures) since the division realignment of the NFL in 2002.

Appendix B: 2014 NFL Regular Season Schedule by the TPM

	Week																
Team	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**
NYJ	@MIA	KC	@ DEN	NE	@ PIT	MIN	@ BUF	CHI	@ SD	X	@ GB	TEN	@ DET	OAK	MIA	BUF	@ NE
NE	@ OAK	@ CHI	MIA	@ NYJ	@ BUF	SD	DET	@ IND	GB	X	BUF	@ KC	@ MIA	DEN	CIN	@ MIN	NYJ
BUF	@ DEN	CLE	OAK	@ MIA	NE	@ DET	NYJ	@ SD	KC	X	@ NE	CHI	@ HOU	MIN	@ GB	@ NYJ	MIA
MIA	NYJ	GB	@ NE	BUF	@ OAK	@KC	SD	@ BAL	DEN	@ CHI	X	@ MIN	NE	JAX	@ NYJ	DET	@ BUF
BAL	TEN	CAR	@ JAX	@ NO	@ CLE	HOU	PIT	MIA	@ IND	X	CLE	@ PIT	CIN	@ TB	ATL	@ SD	@ CIN
CLE	NO	@ BUF	CIN	TB	BAL	@ CAR	@ HOU	TEN	@ ATL	X	@ BAL	JAX	PIT	@ CIN	@ IND	OAK	@ PIT
PIT	@ ATL	JAX	@ KC	@ CIN	NYJ	@ TB	@ BAL	X	NO	CIN	@TEN	BAL	@CLE	@HOU	CAR	IND	CLE
CIN	@ JAX	DEN	@ CLE	PIT	@ CAR	@ NO	IND	ATL	@ TEN	@PIT	X	TB	@BAL	CLE	@ NE	HOU	BAL
HOU	PHI	@ IND	@ NYG	WAS	@ JAX	@ BAL	CLE	X	@ OAK	TEN	IND	@DAL	BUF	PIT	JAX	@ CIN	@ TEN
JAX	CIN	@ PIT	BAL	@ PHI	HOU	X	@ TEN	@ WAS	DAL	@IND	SD	@CLE	TEN	@MIA	@ HOU	NYG	IND
IND	@ NYG	HOU	@ TEN	X	DAL	WAS	@ CIN	NE	BAL	JAX	@HOU	@DEN	@PHI	TEN	CLE	@ PIT	@ JAX
TEN	@ BAL	@ DAL	IND	NYG	@ WAS	X	JAX	@ CLE	CIN	@HOU	PIT	@NYJ	@JAX	@IND	KC	PHI	HOU
SD	KC	@ SEA	@ STL	X	DEN	@NE	@ MIA	BUF	NYJ	OAK	@JAX	AZ	@OAK	@KC	SF	BAL	@ DEN
DEN	BUF	@ CIN	NYJ	X	@ SD	OAK	@ SF	AZ	@ MIA	@STL	SEA	IND	@KC	@NE	@ OAK	KC	SD
OAK	NE	STL	@ BUF	KC	MIA	@DEN	@ AZ	X	HOU	@ SD	@SF	SEA	SD	@NYJ	DEN	@ CLE	@ KC
KC	@ SD	@ NYJ	PIT	@ OAK	STL	MIA	X	@SEA	@ BUF	SF	@AZ	NE	DEN	SD	@ TEN	@ DEN	OAK
DAL	@ SEA	TEN	@ PHI	X	@ IND	AZ	@ WAS	SF	@ JAX	PHI	@CHI	HOU	WAS	NYG	@ STL	NO	@ NYG
WAS	@ SF	NYG	@ AZ	@ HOU	TEN	@ IND	DAL	JAX	SEA	@ NYG	X	STL	@DAL	PHI	MIN	@ TB	@ PHI
PHI	@ HOU	AZ	DAL	JAX	@ NYG	@ STL	GB	X	@ CAR	@ DAL	NYG	@SF	IND	@WAS	SEA	@ TEN	WAS
NYG	IND	@ WAS	HOU	@ TEN	PHI	@ SEA	STL	@ DET	SF	WAS	@PHI	X	ATL	@DAL	@ AZ	@ JAX	DAL
DET	@ GB	@ TB	CH	@ MIN	ATL	BUF	@ NE	NYG	X	@ AZ	NO	GB	NYJ	@CAR	@ CHI	@ MIA	MIN
CHI	MIN	NE	@ DET	CAR	X	@ SF	TB	@ NYJ	@MIN	MIA	DAL	@BUF	@NO	@ GB	DET	@ ATL	GB
MIN	@ CHI	ATL	@ TB	DET	@ GB	@ NYJ	CAR	STL	CHI	@ NO	X	MIA	GB	@ BUF	@ WAS	NE	@ DET
GB	DET	@ MIA	NO	SEA	MIN	@ ATL	@ PHI	TB	@ NE	X	NYJ	@DET	@MIN	CHI	BUF	@ CAR	@ CHI
TB	@ CAR	DET	MIN	@ CLE	NO	PIT	@ CHI	@ GB	@ STL	ATL	X	@CIN	CAR	BAL	@ NO	WAS	@ ATL
ATL	PIT	@ MIN	@ CAR	AZ	@ DET	GB	@ NO	@ CIN	CLE	@ TB	CAR	X	@NYG	NO	@ BAL	CHI	TB
CAR	TB	@ BAL	ATL	@ CHI	CIN	CLE	@ MIN	X	PHI	@ SEA	@ ATL	NO	@TB	DET	@ PIT	GB	@ NO
NO	@ CLE	SF	@ GB	BAL	@ TB	CIN	ATL	X	@ PIT	MIN	@ DET	@CAR	CHI	@ATL	TB	@ DAL	CAR
AZ	STL	@ PHI	WAS	@ ATL	SEA	@ DAL	OAK	@ DEN	X	DET	KC	@SD	@ SF	@ STL	NYG	@ SEA	SF
SF	WAS	@ NO	SEA	@ STL	X	CHI	DEN	@ DAL	@ NYG	@KC	OAK	PHI	AZ	@ SEA	@ SD	STL	@ AZ
SEA	DAL	SD	@ SF	@ GB	@ AZ	NYG	X	KC	@ WAS	CAR	@ DEN	@ OAK	STL	SF	@ PHI	AZ	@ STL
STL	@ AZ	@ OAK	SD	SF	@ KC	PHI	@ NYG	@ MIN	TB	DEN	X	@ WAS	@ SEA	AZ	DAL	@ SF	SEA

Legend	Meaning
@	Road Game
X	BYE
(black)	Divisional Game
(gray)	Thursday Game