



National Football League | Season Scheduling

Introduction

Over the past few decades optimization has made huge strides; what were once unsolvable problems are now solved in only seconds. Of course, the response to this advancement has been to tackle ever larger and harder problems.

The US's National Football League is a perfect example of this. Four people have 10 weeks to schedule 256 games over the course of a 17-week season. To some, that may seem like a lot of planning time available for seemingly few decisions. However, when you work it out, the number of possible schedules is well into the trillions. Imagine the number one followed by 19 zeros.

A data scientist would take years to determine a good schedule out of that huge number of possibilities. Making the problem particularly hard is the addition of numerous "additional" constraints including not scheduling teams from shared markets, such as the NY Jets and Giants, to play on the same day/time, ensuring no team has more than three road games in a row, and increasing the importance of late-season games by scheduling as many divisional match-ups as possible in the final weeks.

From Boards to Computers

Until the last decade, a wooden six-foot square board was used to map out who would play whom, when and where. Hung on a wall the board consisted of a grid with 17 rows representing each week of the season and 32 columns representing each team. The original goal was simply to find a feasible schedule and accomplishing that often consumed up to 90 percent of the available planning time. The remainder was spent tweaking that feasible schedule to improve either a specific team or network's schedule.

With the switch to computers and the increasing power of optimization, a dramatic shift happened. The focus migrated from schedule creation (which by itself is far more complicated than it was only two-three years ago) toward schedule analysis. Before, the question facing the planning team was simply, "Can we just get it done?" Now, it is about finding the best possible schedule.

As Much Art as Science

This shift is important since building a great schedule can be as much art as science. For example, one goal of the scheduling process is to get the best match-ups in the TV time slots that have the widest possible audience. This is very challenging to do since it requires a subjective evaluation of the match-up while also ensuring the resulting schedule doesn't violate any constraints such as the maximum number of away games in a row for a given team. In addition, these subjective evaluations can change dramatically during the course of the scheduling process due to a variety of factors, including free agency.

The art aspect is central to the scheduling process. The planners can now evaluate any given schedule, identify what they want to change, and then see how that affects the overall schedule. If the resultant schedule is better, then it is the new leader; if not, it is dropped. Even as recently as just eight years ago planners might only have been able to look at five total schedules; using Gurobi for the 2014 regular season schedule, they examined more than one thousand feasible schedules.

From Linear to Parallel Optimization

The complexity of the NFL scheduling problem has grown exponentially over the past few years. A big part of the strategy for combatting this massive increase in difficulty was to obtain significantly more computer hardware. The NFL initially started with a single 24-core box, but recently moved to a room full of 16-core, 40-core, and larger servers.

As is often the case with new technologies, you typically won't reap the full benefits by applying existing strategies. Optimal Planning Solutions collaborated with the development team from Gurobi to rethink their approach to exploiting parallelism. This has enabled the NFL to employ a new approach, for the first time in 2014.

A key advantage of the parallelization approach is that a single problem is broken down into a finite number of smaller sub-problems, which are deployed to a pool of solving resources, all of which can share information, creating a much more coordinated search effort.

This approach has allowed the NFL to:

- do a substantially better job of sampling the whole solution space, as opposed to exploring only a small neighborhood of a few known solutions; and
- provide greater flexibility for future hardware configurations, including options to migrate away from multi-core solving boxes to smaller, less-expensive hardware configurations or even cloud-based technology.

Looking Forward

In 2014, the NFL scheduling problem was solved across a network of more than 750 cores. In future years, as more difficult scheduling requirements continue to be added, this number will surely increase and possibly move to a cloud environment. Increasing problem complexity combined with optimization tools enhancements, will make strong partnerships between experienced end-users such as the NFL and consulting firms such as Optimal Planning and focused solver developers like Gurobi even more important.

“ The team at Gurobi was great to partner with. Both the development and support personnel were very responsive and their ideas played an important part in helping deliver great results. ”

— Rick Stone, President,
Optimal Planning Solutions

“ Switching to, and working with, the team from Gurobi has been a great experience. Their solver has met our high expectations, and the support they provide has been fantastic. ”

— Michael North, Sr. Director, NFL