Tenth International Conference on Managing Fatigue: Abstract for Review

Adding Human Physiology to the Airline Crew Scheduling Process

David Karlsson, Jeppesen, david.karlsson@jeppesen.com (corresponding author)
Tomas Klemets, Jeppesen, tomas.klemets@jeppesen.com

Problem

New regulations in the airline industry are opening the door for using bio-mathematical models for evaluating schedules for alertness. As long as the airline can show an "equivalent level of safety" it is even possible to alleviate some of the regulatory rules, and thus realize some productivity gains.

There is still a lot of work to be done in finding the optimal way of integrating human physiology in the airline crew scheduling process. Even so, the early adopters of this technology are already seeing significant savings from a multitude of different fatigue risk management strategies.

Method

Airline crew scheduling has long been too complex to be done by hand. Instead it is handled by computer algorithms, so-called optimizers. These optimizers are black-boxes, striving to create the most efficient combinations given the constraints set by, for example, the regulatory rules, labor agreements, and other, self-imposed, scheduling practices. Usually, each of these domains contain examples of rules to handle crew fatigue and fatigue risk.

Traditionally the main way of handling flight patterns that are deemed fatiguing has been to introduce new constraints that force the optimizer to avoid building similar patterns. The main drawback of this method is that it introduces additional complexity in the scheduling process, and the effects of that complexity is not easily predicted. Trying to remove fatiguing patterns will often also cause the removal of some productive, non-fatiguing, ones.

Of the airlines that have begun to look into bio-mathematical modelling of fatigue, most are currently only using a model after schedule creation to find any combination of flights that are considered "too fatiguing to be flown". These patterns are broken up and re-planned, often by hand, to produce new patterns. These patterns may or may not be safer, but are always less productive.

A better approach seems to be to use the model prediction of fatigue levels directly in the optimizer and automate finding the best line of action. This can be done in two main ways. Either in terms of the predicted fatigue on each

individual flight, much like the re-active way of handling fatigue risk, or in terms of some property of the distribution of predicted fatigue for the entire operation. The latter approach has the advantage of not requiring a threshold for a "safe level of fatigue" for an individual flight.

Results

We have worked with several airlines from all over the world to make their crew schedules more aligned with human physiology. The results, so far, have been overwhelmingly positive. In most cases, just giving the optimizer access to the predictions from the fatigue model (thus giving it the possibility to manage fatigue) usually means that the fatigue levels in the resulting schedules will be reduced at a very slight cost in terms of crew productivity.

Taking the next step, to actively change the rule context based on predictions from the model, requires more work, but can also offer more rewards. In the process we use the fatigue model with the optimizer to highlight the rules that are expensive in terms of productivity, and which are not contributing to limiting fatigue in the operation. These rules can then be re-formulated in a way that is less limiting when it comes to productivity, but with an equivalent level of safety. After testing out the new rule formulation, a decision can be taken regarding the worth of re-negotiating the rule.

Only by re-formulating or removing old, non-functional rules can there be an increase in schedule efficiency. It is then up to the airline to decide what to do with this improvement potential. We predict that some airlines will take out the gains as productivity increases, but improved quality of life for crew members or decreased levels of predicted fatigue are also possible outlets.

Discussion

While a rule is conceptually easy to understand, the impact of introducing it to a complex system like the airline crew scheduling is hard to predict. The impact of a far more complex fatigue model is, paradoxically, easier to predict.

The integrated system of automatic schedule creation and fatigue model has several advantages over using the software separately. It allows the complete process to be automated and also allows for post-hoc manual fatigue analysis, without leaving the scheduling software. There are also significant schedule efficiency gains to be had by letting the optimizer dynamically handle the rebuilding of fatiguing patterns instead of doing it later by hand.

With an integrated system, there are different ways of integrating the fatigue model into the schedule creation process. The simplest approach, to just move the "too fatiguing to fly" threshold from the post-planning step into the optimizer, will un-lock some productivity that is lost in the manual planning step, but will still require a threshold below which no flight should be operated. This threshold is inherently problematic and one of the primary reasons for moving to a distribution-centric view of fatigue risk.

Even in cases where a rule based fatigue risk management system was shown to do a nice job in limiting fatigue, there are other potential gains to be had from replacing parts of the system with a fatigue model. Examples of these gains include finding productivity inhibited by the fatigue rules, but that is not particularly fatiguing, or increasing crew satisfaction with work schedules.

Summary

Using bio-mathematical models of fatigue to extend or replace (part of) a rule-based system for limiting fatigue risk has been shown to be a viable solution. It often creates better results than the reactive analyze and rebuild methods frequently in use today. It can also offer the answer to one of the most common questions in this field: “What level is safe?” by looking at the distribution of the fatigue predictions, instead of focusing on the prediction for each individual flight.