# The Best Rest

BY DAVID HELLERSTRÖM, HANS ERIKSSON, EMMA ROMIG AND TOMAS KLEMETS

**A comparison of differing regulatory efforts to control pilot fatigue.**

In commercial aviation, crew schedules are regulated by duty time limits, flight time limits, minimum rest rules and other constraints. These rules and limits, collectively referred to as flight time limitations (FTLs), originally were conceived as a simple scheme for limiting fatigue among flight crewmembers.

Over time, FTLs have evolved, driven by industrial pressures or new scientific data, or to cope with changing aircraft capabilities. Today, there are major differences among FTL schemes in different parts of the world affecting crew productivity, crew alertness — and airline competitiveness.

With the results of new research on sleep and work-related fatigue in hand, it becomes useful to compare existing regulations with the new findings.

FTLs are relatively straightforward, and, combined with labor agreements and other safeguards, they do a reasonable job of protecting alertness under most circumstances. Unfortunately, FTLs tend to be extremely rigid and limit operational flexibility and efficiency. But by far the most troublesome aspect of FTLs is the illusion of safety that they create — suggesting that to fly within the limits is inherently safe, while flying outside the limits is inherently unsafe.

In recent years, considerable effort has been directed toward increasing scientific knowledge of fatigue and

© Monika Wisniewska/iStockphoto

alertness. By combining new knowledge of fatigue with safety and risk management processes, the concept of the fatigue risk management system (FRMS) was created. In previous work, we have demonstrated that a properly implemented and managed FRMS can be vastly superior to FTLs in managing alertness while maintaining or improving productivity.[1] Whereas FTLs are not feedback-driven and often lack a scientific basis, an FRMS is by definition intended to be a closed-loop, data-driven process. In addition to the stronger scientific basis of an FRMS, an added benefit is increased operational flexibility.

FRMSs are built around predictive tools including, but not necessarily limited to, mathematical models of fatigue and alertness. Models predict crew alertness from planned and actual schedules and inferred sleep and wake history. Models also consider known physiological phenomena, such as circadian rhythms and sleep propensity, and make predictions based on these considerations. Unfortunately, while models have been developed and validated in a laboratory environment, more work is required to validate the models in a commercial aviation environment. Without validation and other checks, the use of any specific model on FRMS in scheduling is ill advised.

Thus, we are faced with a dilemma. FTLs are imperfect, but well understood and easy to apply. An FRMS is better for managing fatigue-related risk but must be developed and validated to be trusted. Until FRMSs are widely proved and implemented, the goal must be to refine FTLs to be as close as possible to an FRMS-based approach. A refined FTL should strive to guarantee an equivalent or better level of flight safety while allowing airlines to efficiently and flexibly operate their businesses.

For this article, we analyzed three different sets of FTLs for productivity and alertness. We compared these regulatory formulations to a model-based FRMS. The analysis used a fatigue model within crew scheduling optimization software on the timetables of three short-haul airline fleets. Finally, we demonstrated our suggested alternative for improving FTLs.

## Analytical Methods

To build the schedules for comparing FTLs, we used the system illustrated in Figure 1. Our system centers on an "optimizer," which considers an airline's timetable and a set of rules and objectives to build crew schedules. In each of our FTL comparisons, we created a schedule using one airline's timetable and one of the FTL sets as a constraint. To simulate an FRMS, we created schedules without the constraint of an FTL set, instead using the predictions of our alertness model.

The FTL sets used were EU-OPS with Subpart Q — abbreviated as Joint Aviation Requirements (JARs); U.S. Federal Aviation Regulations (FARs) Part 121; and China Civil Aviation Regulations (CCARs) 121 Rev 3. Each FTL scheme has a different focus:

- JARs focus on duty-time limitations with reduced daily limits based on the number of legs and time of day. Duty time can be extended twice in seven days. Minimum rest between duty periods is 10 hours. There may be no more than seven days of work between rest periods of at least 36 hours.

**Tools and Conditions Used in Schedule Creation**

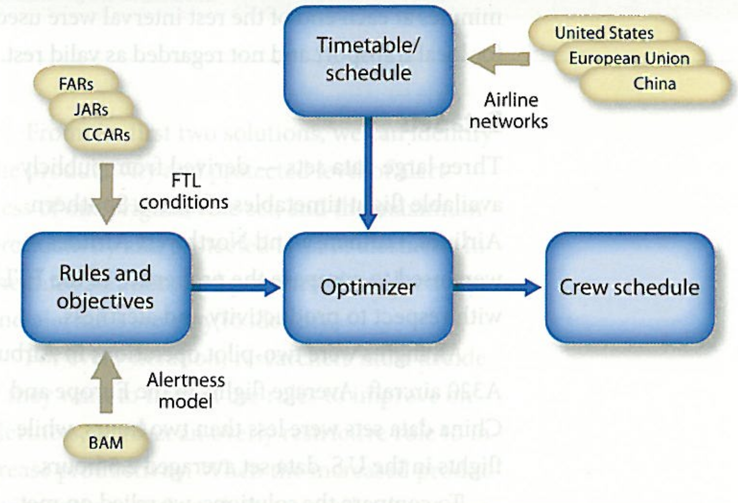

BAM = Boeing Alertness Model; CCARs = China Civil Aviation Regulations; FARs = U.S. Federal Aviation Regulations; FTLs = Flight time limitations; JARs = Joint Aviation Requirements (referring to EU-OPS with Subpart Q)

Source: David Hellerström, Hans Eriksson, Emma Romig and Tomas Klemets

**Figure 1**

- FARs limit block time and lack real duty-time limits. Minimum rest between duty periods is eight hours. There must be weekly rest of at least 24 hours in every seven-day period.
- CCARs address both block time and duty time limits. Minimum rest between duty periods is 10 hours. The weekly rest requirement is 48 hours in any seven-day period.

In addition to the three FTL sets, we created an "FRMS" rule set based on a model's predicted alertness. The rule set was created using the Boeing Alertness Model (BAM), a bio-mathematical model of alertness.[2,3] In this rule set, there were no rules on flight time, duty time, or rest time; instead, an alertness limit was set, under which no flights would be scheduled. Alertness is predicted on a scale from zero (least alert) to 10,000 (most alert), which we call the Common Alertness Scale.[4]

**It is possible to build solutions that protect against fatigue without sacrificing productivity.**

JARs and CCARs consider duty time to include briefing and debriefing; for this analysis, we set the parameters for briefing time to 45 minutes before active duty and 30 minutes before passive duty.[5] Debriefing time was set to 15 minutes. CCARs define "rest at rest location" as being rest at a hotel, rather than at an airport; therefore, 20 minutes at each end of the rest interval were used for local transport and not regarded as valid rest.

## Data Sets

Three large data sets — derived from publicly available flight timetables of China Southern Airlines, Lufthansa and Northwest Airlines — were used to compare the properties of the FTLs with respect to productivity and alertness.

All flights were two-pilot operations in Airbus A320 aircraft. Average flights in the Europe and China data sets were less than two hours, while flights in the U.S. data set averaged 2.5 hours.

To compare the solutions, we relied on metrics representing the resources needed to implement a solution for each flight and the predicted alertness level of flight crewmembers. A low level of predicted alertness on a flight is associated with higher risk. The alertness properties in the solutions were hard to map to a single descriptive value or statistical measure; therefore, we chose to report and compare the lowest level of predicted alertness, as well as the average alertness value of the lowest 1 percent, 5 percent and 10 percent of flights within the schedule.

To quantify the relative productivity of the solutions, we created a composite measure of productivity called the "Resource Index (RI)." RI values are a measure of how much less efficient a solution is than a theoretically "perfect" solution. Using all three airline data sets, we observed the same trend in the RI: The FARs were the most flexible and most efficient of the FTL schemes, followed by the CCARs and finally the JARs. The flexibility of the FARs comes primarily from the lack of duty-time limits and the possibility of a rest period as short as eight hours. However, the BAM outperformed all three FTL sets in terms of the resource index.

When we considered average block time per duty day — another measure of productivity — we saw similar performance on predicted alertness from BAM and the FTLs. Only when applied to the Chinese data set did the FARs generate a solution more efficient than that created by BAM.

Under the U.S. airline operating conditions, with relatively fewer legs and legs of longer duration, the JARs outperformed the CCARs in terms of crew productivity per day; in all other cases, the JARs were the least efficient of the FTL schemes. The performance shortfall on the other FTL sets probably stemmed from the reduction in duty-time limits for many sectors under the JARs. We also noted that the FARs — without any real duty-time limit — consumed much more duty time than the other FTL schemes.

Figure 2 shows the level of fatigue is highly dependent on the data set because legs scheduled very early or very late always cause low alertness. As shown in the figure, the FARs provided the least protection against fatigue; the CCARs and JARs were comparable to each other, but the JARs provided somewhat better protection. The solutions produced by BAM were better at protecting against fatigue — not surprising because when constructing a schedule with BAM, predicted alertness is a primary objective. The BAM solutions were interesting because they showed that it is possible to build solutions that protect against fatigue without sacrificing productivity.

Worth noting is that many of the FTL-permissible flights associated with low alertness would not be allowed under BAM-based rules.

**Alertness Comparison**

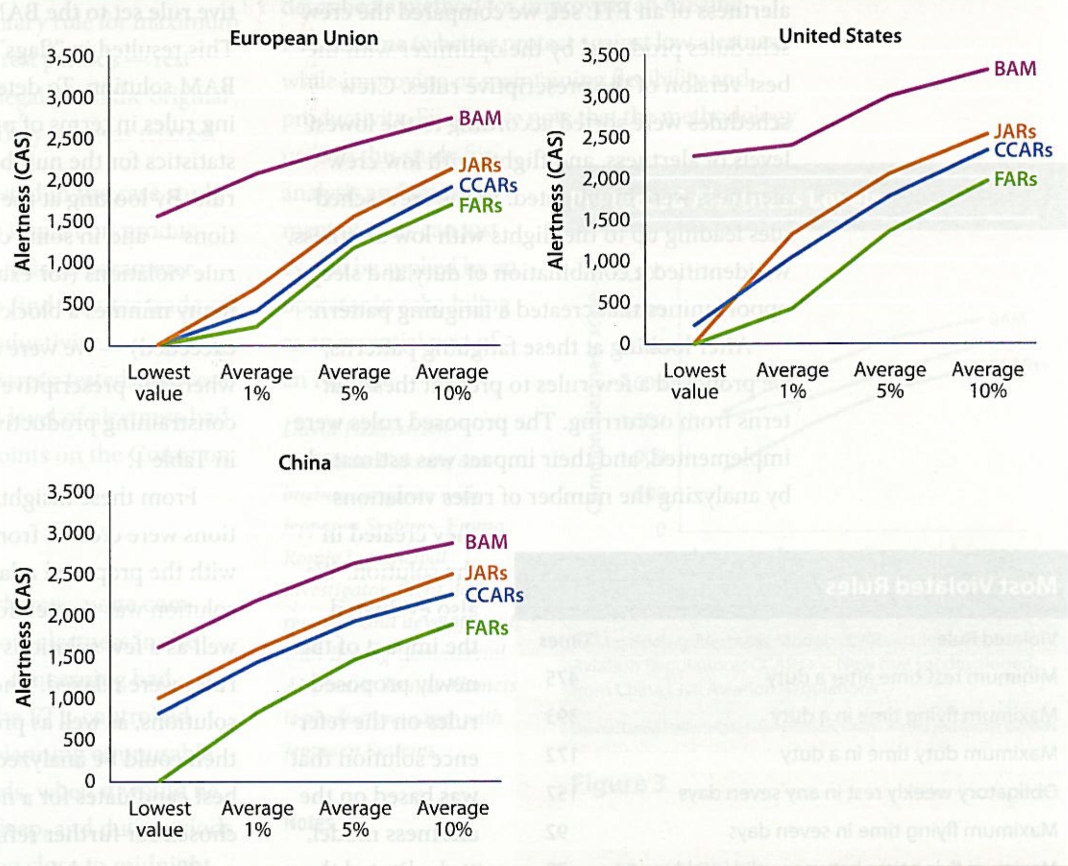

BAM = Boeing Alertness Model; CAS = Common Alertness Scale; CCARs = China Civil Aviation Regulations; FARs = U.S. Federal Aviation Regulations; JARs = Joint Aviation Requirements (referring to EU-OPS with Subpart Q)

Source: David Hellerström, Hans Eriksson, Emma Romig and Tomas Klemets

**Figure 2**

## Improving the Rules

The tools used for this productivity and alertness comparison can be extended into a framework to improve prescriptive rules, such as an FTL scheme, to help the FTL scheme provide better protection against low alertness while also maintaining or improving productivity. In this application, the optimizer can be used to analyze the properties, including productivity and alertness, of an evolving rule set. The method identifies overly restrictive rules and loopholes in the existing rule set.

The improvement begins with the creation of three reference solutions. One solution is based solely on the alertness model with no other limiting rules. The second solution is based on the limits in the prescriptive rules. The third solution is a stress test solution, also based on the limits in the prescriptive rules. In the stress test, the researcher activates an incentive so that the optimizer will produce the most tiring solutions allowed under an FTL.

From the first two solutions, we can identify the productivity and protected level of alertness of our original rule set, and the maximum productivity and protected level of alertness. In the third solution, bad patterns of productivity and alertness are easy to identify.

For every iteration, researchers must decide if they want to tighten the rules to improve on alertness, or relax an overly restrictive rule to increase productivity. When the increased productivity option is selected, the revised rule set also changes the alertness outcome — probably for the worse. Likewise, when alertness is improved, the rule set usually causes loss of productivity. Changes that improve productivity or alertness — without one affecting the other — are ideal.

## Improving Alertness

In our effort to improve the protected level of alertness of an FTL set, we compared the crew schedules produced by the optimizer with the best version of the prescriptive rules. Crew schedules were sorted according to the lowest levels of alertness, and flights with low crew alertness were highlighted. In the crew schedules leading up to the flights with low alertness, we identified a combination of duty and sleep opportunities that created a fatiguing pattern.

After looking at these fatiguing patterns, we proposed a few rules to prevent these patterns from occurring. The proposed rules were implemented, and their impact was estimated by analyzing the number of rules violations they created in the solution. We also evaluated the impact of the newly proposed rules on the reference solution that was based on the alertness model, and adjusted the proposed rules as warranted.

The final impact of new rules was then analyzed by generating a set of new solutions from the prescriptive rule set and the newly proposed rules. One new solution for each added rule, and a few solutions using combinations of new rules, were generated. The productivity and level of alertness were analyzed for each solution, and the data were plotted on a chart. One rule, or a few rules that collectively improved alertness, were chosen to move forward.

**Most Violated Rules**

| Violated Rule | Times |
|---|---|
| Minimum rest time after a duty | 475 |
| Maximum flying time in a duty | 393 |
| Maximum duty time in a duty | 172 |
| Obligatory weekly rest in any seven days | 157 |
| Maximum flying time in seven days | 92 |
| Maximum flying time between valid weekly rest | 28 |

Source: David Hellerström, Hans Eriksson, Emma Romig and Tomas Klemets

**Table 1**

## Improving Productivity

To examine possibilities for improving productivity, the BAM reference solution became the starting point. As noted, this solution had no constraints other than maintaining a protected level of alertness. Theoretically, then, it should be the most productive solution possible, unless all protection of alertness is sacrificed. In our system, it was possible to apply the prescriptive rule set to the BAM reference solution. This resulted in "flags" of rule violations in the BAM solution. To determine the most limiting rules in terms of productivity, we compiled statistics for the number of violations of each rule. By looking at the frequency of rule violations — and in some cases the degree of the rule violations (for example, looking at by how many minutes a block or duty time limit was exceeded) — we were able to gain insight into where the prescriptive rules were unnecessarily constraining productivity. Examples are shown in Table 1.

From these insights, new FTL-scheme solutions were created from the prescriptive rules, with the proposed relaxations added. One new solution was created for each relaxed rule, as well as a few solutions in which combinations of rules were relaxed. The productivity of the new solutions, as well as protected level of alertness, then could be analyzed. One or several of the best candidates for a new rule set then could be chosen for further refinement.

The research has validated the methodology by applying it to the CCARs rule set and the data set representing the Chinese airline.

In three iterations, nine rule changes were tried and five rule changes were introduced. The final result was a rule set in which the average block time per day was increased by 6 percent from 5 hours 59 minutes to 6 hours 21 minutes and alertness was improved between 250 and 700 points on the Common Alertness Scale. Differences in alertness are compared in Figure 3, where the new rule set is named CCARs+. The new solution's resource index also dropped 8.5 percent.

The following rule changes were introduced:

- Prohibiting pilots from being asked to report for duty more than once in a 24-hour day;
- Reducing the maximum duty time for duty periods that fall partly within 2300 to 0330;
- Relaxing the rule governing maximum block time in a duty period;

- Relaxing the rule governing minimum rest after duty; and,
- Adding a complementary rule for maximum duty time after short rest periods — rest periods that became legal when the original minimum rest-after-duty rule was relaxed.

The parameter changes tested in the case study were large and had a large impact on productivity and alertness. More refined parameter changes could be tested to find a better trade-off between alertness and productivity.

The final rule set was stress-tested. The test showed that the protected level of alertness had increased by 250 to 450 points on the Common Alertness Scale.

## Conclusions

Of the three tested FTL schemes, none completely protected against low alertness in the crew schedules. The most concerning bad patterns encountered in the FTL-controlled crew schedules were the planning of unusable rest during daytime periods, when it would be difficult for the pilots to sleep, and duty periods of maximum length ending close to midnight. These situations are legal and appeared in solutions generated from all FTL schemes.

The JARs and CCARs rule sets are comparable in many aspects, both in productivity and in the protection against low alertness. The JARs FTLs are slightly better at protecting against fatigue but less productive if there are many legs in the average duty. The FARs FTLs are the most efficient of the three FTLs but allowed for very long duty times. FARs FTLs also performed worst in protecting against low alertness.

The levels of alertness predicted by BAM for the FTLs should be viewed with caution because the model is not yet fully validated in airline operations. When the model is shown to be valid, the safety and business case for FRMS will be further strengthened. Our results indicate that FTLs do not appear to protect well against low alertness — and within an airline's FRMS, model-based scheduling should be both safer and more productive.

In the meantime, assuming that current FTL schemes are to be moved toward FRMS, we have described a method for improving an existing FTL scheme to better protect against low alertness while improving or maintaining flexibility and productivity. Finally, we note that the methodology used in this study for analysis and improvement of rules can just as well be applied by an operator in scheduling as an essential part of an FRMS.

*David Hellerström and Hans Eriksson are business analysts with Jeppesen Systems. Emma Romig is principal investigator, flight deck research and development, with Boeing Commercial Airplanes. Tomas Klemets is product manager with Jeppesen Systems.*

**Protected Alertness Levels**

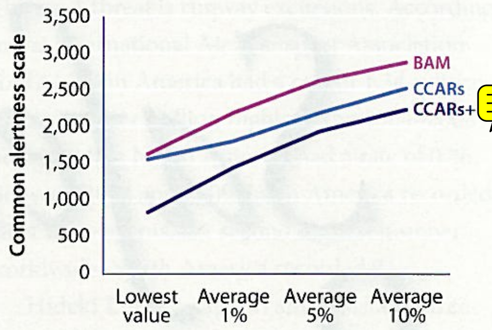

BAM = Boeing Alertness Model; CCARs = China Civil Aviation Regulations; CCARs+ = New rule set developed from China Civil Aviation Regulations

Source: David Hellerström, Hans Eriksson, Emma Romig and Tomas Klemets

**Figure 3**

### Notes

1. Romig, Emma; Klemets, Tomas. "Fatigue Risk Management in Flight Crew Scheduling." *Aviation, Space, and Environmental Medicine* Volume 80 (December 2009): 1073–1074(2).
2. Åkerstedt, T.; Folkard, S. "The Three-Process Model of Alertness and its Extension to Performance, Sleep Latency, and Sleep Length." *Chronobiology International* 14(2), 115–123, 1997.
3. Åkerstedt, T.; Folkard, S.; Portin, C. "Predictions From the Three-Process Model of Alertness." *Aviation, Space, and Environmental Medicine* 2004; 75(3, Suppl.): A75-83.
4. This scale has been anchored to the Karolinska Sleepiness Scale, a widely used scale for rating sleepiness. The Common Alertness Scale, and the interface by which BAM connects with the scheduling software, has been formalized in a document shared with leading fatigue modelers, with the intention that other commercial models can be integrated into crew scheduling as BAM has been.
5. Passive duty is duty time during which the crewmember flies as a passenger to be positioned for further duty.