ROI 2010-0222-1576

Flight Time Limitations and Fatigue Risk Management: A comparison of three regulatory approaches

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 or FTLs, were intended to be a simple method of limiting and accounting for fatigue among flight crew members. Today there are major differences among FTLs formulations in different parts of the world affecting crew productivity and the level of alertness for the crew. In recent years, an alternative to FTLs has emerged – Fatigue Risk Management Systems or FRMSs.

In this paper we present a comparison between three FTL formulations based on crew schedules for typical carriers from three regions: Europe, the US, and China. For each set of FTLs, crew schedules are created using typical scheduling incentives to simulate realistic conditions within an airline.

The crew schedules are compared in terms of a number of key indicators of productivity and predicted alertness. The analysis identifies loopholes and productivity limitations built into each FTL scheme. In addition, we will compare these crew schedules with schedules built with an alertness model under an envisioned Fatigue Risk Management System.

Finally we present a methodology for improving a set of prescriptive rules to increase the alertness while maintaining, or improving crew productivity in the solutions.

Keywords: *crew scheduling, crew rostering, fatigue, BAM, alertness, flight time limitations, FRM, Fatigue Risk Management*

DAVID HELLERSTRÖM david.hellerstrom@jeppesen.com
HANS ERIKSSON hans.eriksson@jeppesen.com
EMMA ROMIG emma.romig@boeing.com
TOMAS KLEMETS tomas.klemets@jeppesen.com

1. INTRODUCTION

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 or FTLs, are intended to be a simple method of limiting and accounting for fatigue among flight crew members, as part of overall safety concerns and objectives.

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

With new research on sleep and work-related fatigue it becomes useful to compare existing regulations with new findings.

FTLs are relatively straightforward to understand and apply to crew scheduling. Combined with labor agreements and other safeguards, FTLs 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, there has been considerable effort spent on increasing our scientific knowledge in the areas of fatigue and 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 should be vastly superior to FTLs in managing alertness [1] while maintaining or improving productivity. Whereas FTLs are not feedback-driven, and often lack a scientific basis, FRMS is by definition intended to be a closed-loop, data-driven process for managing fatigue. In addition to the stronger scientific basis from FRM, there is an added benefit of 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 from which the model can infer or predict sleep and wake history. Models consider known phenomena, such as the circadian rhythms and sleep propensity, and make predictions based on these considerations. Unfortunately, while models have been developed and validated in a laboratory environment, there is still effort required to validate the models in a commercial aviation environment. Without validation and vetting, the use of models and FRM in scheduling is ill-advised.

Thus, we are faced with a dilemma. FTLs are imperfect, but well-understood and easy to apply. FRM is a better system for managing fatigue related risk, but needs development and validation to be trusted. Until FRMSs are widely proven and implemented, the goal must be to refine FTLs to be as close to an FRMS-based approach as possible. The goal of a revised FTL should be to guarantee an equivalent or better level of flight safety but at the same time allow for the airlines to efficiently and flexibly operate their business.

In this paper, we present an analysis of three different FTLs from a productivity and alertness point-of-view. We compare these regulatory formulations against a model-based FRM approach. The analysis is carried out using a fatigue model within crew scheduling optimization software on the time tables of three short-haul airline fleets. Finally, we demonstrate an iterative approach for improving FTLs.

2. METHODS

Analytical Conditions

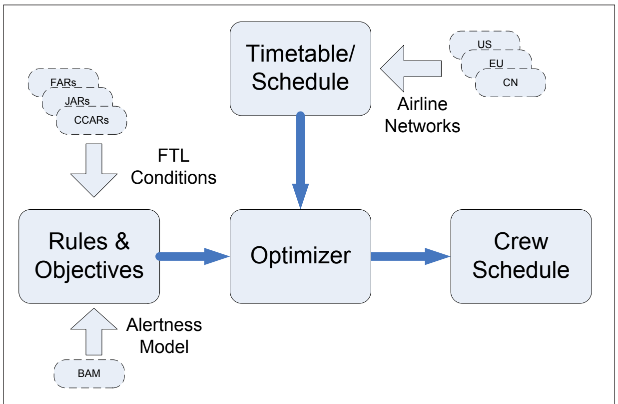

Figure 1. Schematic representation of the tools and conditions used in schedule creation

In order to the build the schedules upon which the FTLs are compared, we utilized the system illustrated schematically in Figure 1. Our system is centered around an optimizer, which considers an airline's timetable and a set of rules and objectives in order to build crew schedules. In each of our FTL comparison runs, we create a schedule using one airline's timetable and one FTL set as a constraint. To simulate an FRMS approach, we also create schedules without the constraint of FTLs, instead using the predictions of our alertness model.

Alertness Model

For prediction of alertness the Boeing Alertness Model (BAM) was used. BAM is a commercially available alertness model based on the Three Process Model of Alertness (TPMA) [2][3] with some extended capabilities and modifications. Perhaps the most significant modifications relates to the model's sleep prediction. When predicting alertness the sleep strategy is of great importance. The model was configured to use sleep prediction according to TPMA [4] with the addition of a "late sleep" strategy as well as the possibility of an "afternoon nap". In addition, the parameters for sleep opportunity, the period during which the crew may chose to sleep, were set so that the sleep opportunity occurs two hours after a flight arrival and ends two hours before the next departure[1].

The alertness model was configured to reset its homeostatic component when encountering sleep opportunities longer than 24 hours, so chronic fatigue is not considered in this study.

BAM interfaces to the crew planning environment through the Common Alertness Prediction Interface (CAPI), a draft protocol for integrating alertness models with crew scheduling software, and thus predicts alertness on the Common Alertness Scale (CAS) in the range 0 to 10 000. CAS is linearly mapped against the Karolinska Sleepiness Scale (KSS) where a KSS value of 9 maps to 0 CAS points and KSS of 1 maps to 10 000 CAS points.

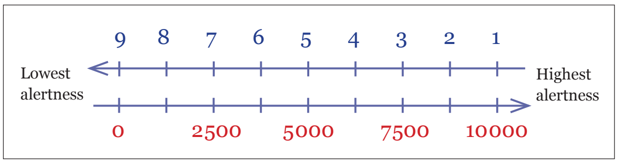

Figure 2. The KSS scale on the top with the corresponding Common Alertness Scale (CAS) used by CAPI below.

In order to capture the fatigue associated with work activity[2] BAM supports a task-load concept based on activities in the schedule. Task-load has been applied at deliberately low levels[3] using a setting of 30 CAS points per flight and 30 CAS points per duty hour. Furthermore a load of 100 CAS points per each preceding consecutive duty day is applied at the start of duties. The task-load for flights and duty hours are only applied on the current duty while the

[1] BAM allows for individual adaption of these opportunities, for example to be used for commuting crew, but throughout this study the mentioned defaults were used.
[2] Fatigue from work stress as opposed to physiological fatigue.
[3] Even though the concept of task-load is acknowledged there is no scientific consensus on how task-load affects alertness and fatigue. The values used here will have a small effect on predicted alertness relative to the circadian and homeostatic effects.

task-load for consecutive duty days carry over and accumulates to the next day.

It is important to point out that even while BAM is based on the validated TPMA model, BAM including extensions to TPMA as described[4], was at the time for this study only validated at a cursory level with two airlines.

Compared Flight Time Limitations

This study was carried out on three regulatory FTLs; JAR (EU-Ops with Subpart Q), FAR (Part 121 rules as of February 2010) and CCAR (CCAR 121 Rev 3); from Europe, the US, and China respectively. Qualitatively speaking, each FTL scheme appears to have developed with a slightly different focus:

JAR – JAR focuses on duty time limitations with reduced daily limits based on number of legs and time-of-day. JAR addresses circadian rhythm with rules around Window of Circadian Low (WOCL). Duty time can be extended twice in seven days. Minimum rest between duties is at least 10 hours. There must be at most seven days work between rest-periods of at least 36 hours.

FAR – FAR works by limiting block time and there is no real duty time limit. Minimum rest between duties is at least 8 hours. There must be weekly rest of at least 24 hours in every seven-day period.

CCAR – CCAR addresses both block time and duty time limits. Minimum rest between duties is at least 10 hours. The weekly rest requirement is 48 hours in any seven-day period.

BAM – In addition to the three FTLs above, an FRM-based rule set based on predicted alertness was created. In this rule set, identified as “BAM”, there were no rules on flight time, duty time, or rest time but a limit for lowest allowed alertness was set to 1500 CAS points. The limit of 1500 was chosen from a preliminary investigation of the existing FTLs looking at to what extent they protected from low levels of alertness.

JAR and CCAR consider duty-time which includes time for briefing and debriefing. For this study we set parameters for briefing time to be 45 minutes before active duty, 30 minutes before passive duty. Debriefing time was set to 15 minutes. CCAR also defines 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 thus were not regarded as valid rest.

The programmatic implementations used for JAR and FAR are validated and in use by many airlines. The CCAR implementation was however new, created for this study, and only cursory validated.

[4] The extensions to TPMA are the task-load parameters and sleep prediction methods.

Optimization

To construct crew work schedules, an optimizer widely used in commercial aviation[5] was used. The constructed crew schedules were work periods, bound by weekly rest periods at homebase. Rest at homebase within the work-period was allowed. The maximum working period length is different depending on the FTL scheme used.

The optimizer solves a set-covering [5] [6] problem where a set of crew schedules are constructed to exactly cover, or assign staff to, the set of flights in the time-table. All constructed crew schedules adhere to a rule set, which is typically configured through a rules engine [7]. Each crew schedule also carries an associated cost, or objective function. The optimizer minimizes the cost for the whole solution, i.e. the sum of the cost of all crew schedules.

FTL	Max length of work-period
JAR	168 hours
FAR	144 hours
CCAR	120 hours
BAM	6 calendar days

Figure 3. The maximum lengths of a working period for the respective FTLs.

The solution constructed by the optimizer can be investigated in detail as well as on an aggregated level.

Data sets

Three large data sets, i.e. flight time tables, were constructed for the study. For each region an airline with a large short-haul fleet was selected – they were Lufthansa, Northwest Airlines and China Southern Airlines for Europe, US and China – respectively. While the time tables used are actual time tables from these airlines this study says nothing about the productivity or safety about these airlines’ operations. The study merely uses their time tables to compare the properties of the FTLs.

Data set	# Flights	Block hrs	Bases
Europe (EU)	3922	6853	FRA, MUC, HAM, TXL, DUS
China (CN)	4314	8099	PEK, SHA, CAN
USA (US)	2898	7491	MSP, DTW, MEM

Figure 4. An overview of the data sets. Number of flights and block hours are from a weekly schedule.

All flights are two-pilot operations with A320 aircraft. The average flight in the EU and CN data-sets are below 2 hours duration while flights in the US data-set are on average 2.5 hours.

To construct the data sets, flight schedules were

[5] Jeppesen’s Carmen Crew Pairing optimizer.

extracted from the publicly available OAG schedule. Flights are from the week 14-20th September 2009 and A320 only. The airline's other fleets (non-A320) were imported as positioning (passive transfer) candidates. Aircraft rotations were reconstructed with a first-in first-out algorithm. Crew bases were chosen for each data set, matching the network topology.

Cost structure

When using an optimizer to construct crew schedules, an objective function for the optimization must be defined. The objective function is dependent on the cost structure of the airline. A major factor in the cost structure is the pilot pay structure, which varies between regions and airlines. In the US pay is often tightly related to block time or actual duty time. Other regions have monthly salary combined with overtime.

A mixed cost structure was used in the optimization to give a realistic objective for an arbitrary operator. The highest priority was given to crew productivity, i.e. the active block hours per calendar day should be high. To reflect possible overtime implications and crew acceptance, there was also a incentive added to keep the ratio between duty time and active block time low avoiding unnecessarily long connections between legs.

The cost structure also considered costs associated with layovers, deadheading and minimizing number of aircraft changes[6].

Comparison metrics

To compare the solutions we rely on four different metrics.

Productivity metrics

The *Resource Index* describes how much resources are needed to implement a solution. The index is a normalized value of the monetary costs in the optimizer's objective function. Simplified, an airline operating with a resource index of 4.0 requires four times the staffing of an airline operating on 1.0.

The cost is normalized against a perfect solution for a perfect time table assuming an ideal duty period including 8 hours block time and 9 hours 45 minutes duty time[7]. There are no deadheads in this solution and all duties start and end at the crew home base. If all the block time in the time table had been covered by duties such as these, we would have the solution requiring the fewest possible resources to staff. This is our reference solution with a resource index value of 1.0. The 1.0 solution is idealized and cannot be achieved due to the nature of the time table as well as constraining rules, but it provides a basis for comparison.

The greatest weight within the resource index is given to the cost for each crew production day and to total duty time. We therefore present values for *average block time per working calendar day* and the *ratio between duty time and block time*.

Alertness evaluation

A low level of predicted alertness on a flight is associated with higher risk. The alertness properties in the solutions are hard to map to a single descriptive value or statistical measure. We chose to report and compare the lowest level of predicted alertness experienced by crew during each flight in the schedule. This gives an alertness distribution of the solution with as many data-points as flights (illustrated in Figure 5). It was deemed most important to focus on the flights in the low-alertness "tail" of the alertness distribution and we further chose to compare the average alertness for the worst 1%, 5% and 10% flights, as well as reporting the lowest alertness value reported on any flight within the solution. We use these three values to compare the *protected level of alertness*, i.e. to what extent a set of prescriptive rules ensures fatiguing situations are avoided.

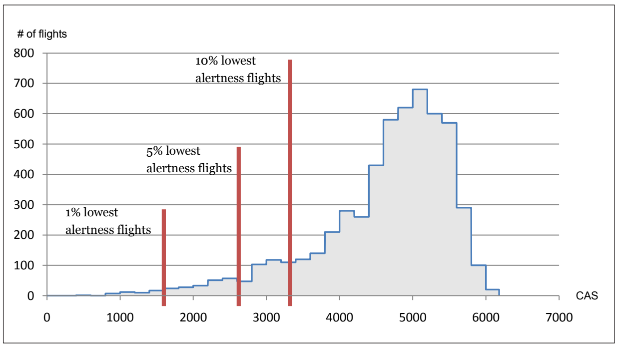

Figure 5. An example distribution of the worst predicted alertness experienced by the pilots in a crew schedule comprised of 5518 flights. The Y-axis depicts the number of flights in each alertness interval. The X-axis represents the lowest predicted alertness experienced on the flight, CAS-scale. The histogram bins are 200 CAS-points wide. Dividers for the 1, 5 and 10% lowest alertness flights are marked with vertical lines. To compare the protected level of alertness we compared the average alertness experienced among these flights.

[6] An aircraft change is a connection where crew could, but does not, follow the aircraft to the next flight, resulting in an undesirable instability in the solution making it more sensitive to delays.
[7] The additional 1h45m on top of the block time is explained by briefing, debriefing and one turnaround.

3. COMPARISON OF FTLs ON PRODUCTIVITY AND ALERTNESS

Productivity comparison

The table below summarizes the results of comparison of the three FTL formulations. Comparing the Resource Index (the overall productivity metric) across the various networks, we can see that there are inherent differences in the time-table and network of the airline which affect their productivity – regardless of the FTL formulation applied.

Under all three network conditions (EU, US, CN), we observe the same trend in Resource Index among the FTLs. In all cases, the FARs are the most flexible and thus most efficient of the FTL formulations, followed by the CCARs and finally the JARs. The flexibility of FARs comes primarily from the lack of duty time limits and the possibility to have a rest period of as few as eight hours. Under all network conditions, the BAM conditions outperform the prescriptive rule sets in terms of Resource Index.

Delving deeper into the secondary metrics of productivity, there are a few noteworthy results. When we consider average block time per duty day, we see similar performance from BAM relative to the FTLs. However, on the Chinese data set the FARs actually produce a solution more efficient than BAM.

Only under the US network conditions, where there are fewer and longer duration legs, do the JARs outperform the CCARs in terms of crew productivity per day, in all other cases the JARs are again the least efficient of the FTL formulations. The short-fall on other data sets is probably connected to the reduction in duty time limits for many sectors in JARs. We can also note that FARs without any real duty time limit consumes much more duty time than the other FTLs.

Data set	FTL	Resource index	Avg. resource index	Avg. block-time / day	Duty-time /block-time
EU	JAR	3.71	3.36	5:31	1.62
	FAR	3.11		6:23	1.79
	CCAR	3.26		6:07	1.75
	BAM	3.04		6:51	1.66
US	JAR	3.97	3.94	6:20	1.50
	FAR	3.73		6:43	1.60
	CCAR	4.12		6:01	1.56
	BAM	3.35		7:55	1.50
CN	JAR	4.10	3.77	5:24	1.56
	FAR	3.56		6:37	1.74
	CCAR	4.01		5:59	1.67
	BAM	3.42		6:28	1.56

Figure 6. The table lists the achieved productivity from the four rule-sets in each data-set.

Alertness comparison

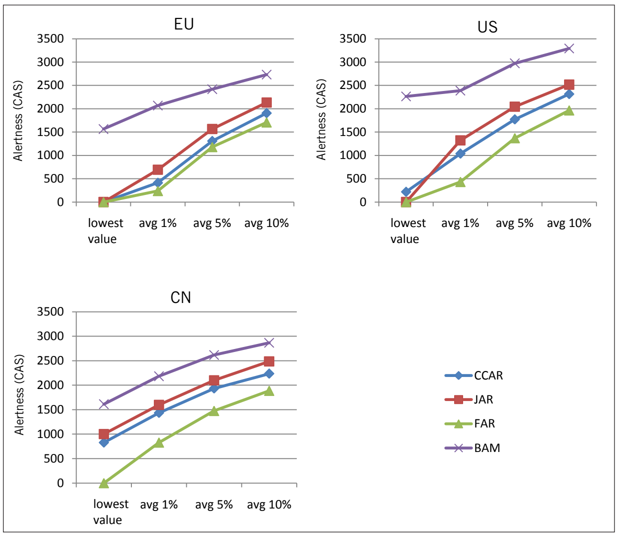

Figure 7. Protected level of alertness for the rule sets on the three data sets, Europe top-left, USA top-right and China bottom-left. Average predicted alertness among the 1, 5, and 10% flights with lowest alertness is compared.

In the graphs above, the protected level of alertness is plotted for each data set and each FTL. The most left-hand point is the alertness of the flight with lowest alertness followed by average alertness among the 1%, 5% and 10% flights with lowest predicted alertness in the solution.

The level of fatigue is highly dependent on the data set. Some legs are scheduled very early or late and will always cause low alertness. The performance of the FTLs is the same in all three networks. FAR protects the least from fatigue. CCAR and JAR are comparable but JAR is somewhat better at protecting against fatigue. The solutions produced by BAM are much better at protecting against fatigue which are not surprising since they were optimized with BAM, where predicted alertness is part of the objective function. The BAM solutions are interesting because they show that it is possible to build solutions that protect against fatigue without sacrificing productivity.

Worth noting is the fact that many of the worst flights allowed by FTLs, would not be allowed with the BAM based rule set – still without sacrificing crew productivity. On the US data set for example, the average alertness on the 10% of the worst flights allowed to be operated by FAR is lower than the flight with lowest alertness operated under the use of BAM.

Stress test comparison

We also have the possibility to "stress test" the FTLs. This is done by adding an incentive to the optimizer to deliberately produce schedules with low alertness. The same cost structure was still used but the objective to reduce duty time was removed. The first priority in these runs was still to build efficient and realistic schedules. The new incentive will however guide the optimizer to pick bad connections, creating the most fatiguing situations allowed under each FTL.

As can be seen in Figure 8 below, the FTLs still rank in the same order as in the production runs. The FARs offer the least protection from low levels of alertness while CCARs and JARs are comparable. This stress test highlights that no FTL scheme manages to protect against low levels of alertness. The BAM-based rule set was deliberately not stress tested as it cannot (by definition), deflect down below 1500 on any flight.

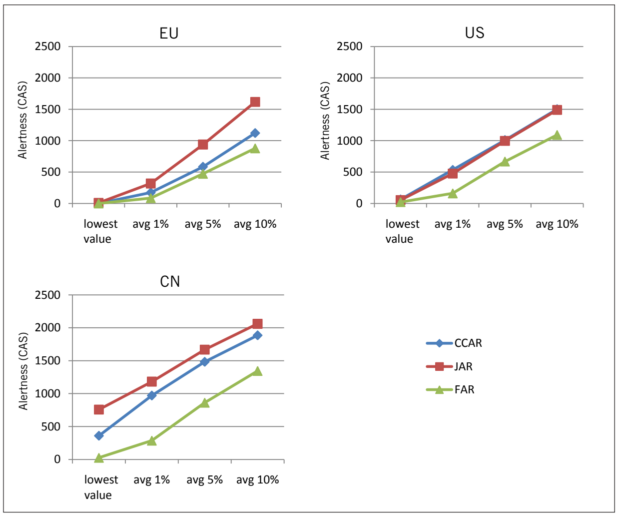

Figure 8. Protected level of alertness from the stress test runs. The stress test incentive reduces overall alertness. The rule sets still rank in the same order as in the productivity runs.

4. ITERATIVE IMPROVEMENT OF A SET OF PRESCRIPTIVE RULES

Introduction

The tools used for the productivity and alertness comparison can easily be extended into a framework to iteratively improve a set of prescriptive rules, such as an FTL. The goal of the proposed method is to improve the set of rules so that they protect better against low alertness while maintaining or improving the productivity in the solutions. The optimizer is 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 quantitative multi objective approach in this methodology, taking both fatigue and crew productivity into account, thus differs dramatically from a qualitative, fatigue-only focused approach for improving rules taken in e.g. the work done for EASA reported on in the Moebus report [8].

Method description

The method starts with the creation of three reference solutions. One solution is based solely on the alertness model with no other limiting rules[8]. The second solution is created towards the limits in the prescriptive rules to be improved. The third solution is a stress test solution created towards the limits in the prescriptive rules to be improved. In this stress test solution the incentive for the optimizer to produces tiring solutions is activated.

All runs are constructed with the same cost structure as in the previous chapter, but if the method is to be applied to a specific airline or region, a cost structure matching the airline or region shall be used.

From the first two solutions, we identify the productivity and protected level of alertness of our original rule set and the possible productivity and protected level of alertness we can achieve. The third solution, where the incentive to produce arduous patterns is active, is used as an eye-opener. By applying the alertness model to this solution and investigating the crew schedules with worst alertness, bad patterns are quite easy to identify.

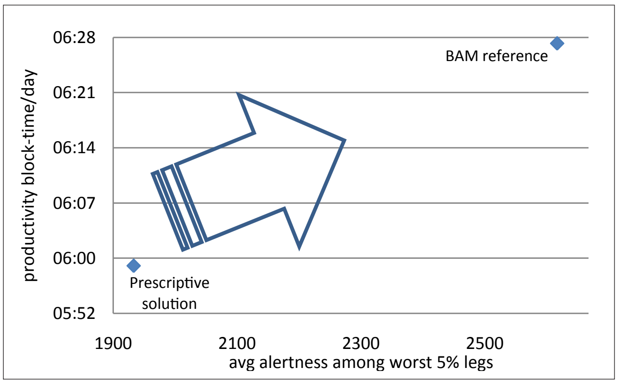

Figure 9. The reference solution governed only by the alertness model is both more efficient and has, according to the model, higher level of protected alertness. The method will transform a set of prescriptive rules such as it moves closer to the properties in the reference solution.

Iterations

For every iteration, it must be decided if we want to tighten the rules to improve on alertness, or if we want to increase productivity by identifying an overly restrictive rule to relax. When productivity is improved the new rule set also changes its alertness properties, most likely for the worse. Likewise, when alertness is improved the rule set usually loses some productivity. Changes that improve productivity or alertness without affecting the other are naturally ideal.

[8] Additional bounding rules on maximum duty-time of 16hrs, maximum flying-time of 12hrs and max 7 legs per duty were used to speed up optimization. These rules were chosen to be so relaxed as not to limit the flexibility of the generated reference solution.

Improving alertness

To improve the protected level of alertness of a rule set, we investigate the crew schedules produced by the optimizer with the best version of the prescriptive rules. This is the version of the rule set we want to improve further. Crew schedules are sorted according to worst alertness and flights with low crew alertness are highlighted. In the crew schedules leading up to the flights with low alertness, there is a combination of duty and sleep opportunities creating a fatiguing pattern.

An attempt is made to identify a common pattern and propose a couple of loose rules that will capture the fatigue cases, see figure 11. The rules are then implemented in a parameterized fashion. The impact of the new rules can be estimated by analyzing the number of rule violations created by them. Preferably the new rule will only trigger in situations where low alertness exist in your current crew solution.

The reference solution, created with the alertness model, is investigated. The new rule set with the newly added rules, is loaded. The new rules are supposed to only capture situations with low alertness, thus few rule violations are to be expected in the reference solution. If there are too many rule violations in the reference solutions the rule may be overly restrictive. If that is the case the rule should be relaxed.

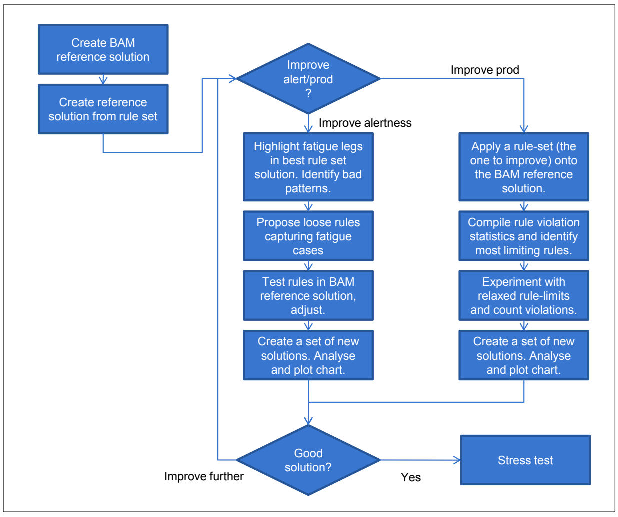

Figure 10. An overview of the iterative method. For every iteration it is decided to either improve alertness or productivity.

A set of new solutions from the prescriptive rule-set are created with the newly proposed rules added. One new solution for each added rule, and a few solutions where combinations of new rules are active, are created. If there is a parameterized limit in the rule, a couple of values can be tried where one solution is created per parameter setting. The new solutions' productivity and level of alertness values are analyzed, and the solutions' statistics are plotted on the chart. One, or a few of the best candidates for new rule sets to continue from, are chosen.

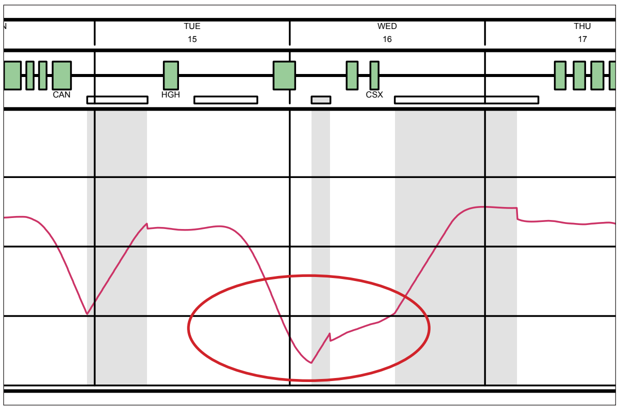

Figure 11. A low alertness pattern is visualized. The boxes on top are flights and sleep opportunities. There are two long rest-periods planned but crew can only sleep in the first. Level of alertness drops below 1 000 CAS-points and crew operates three flights with less than 2 000 CAS-points.

Improving productivity

When improving productivity, the reference solution, produced by the rule-set only governed by the alertness model, is addressed first. Since this solution had no other constraints than maintaining a good protected level of alertness it should be the most productive solution possible, unless alertness is sacrificed.

The BAM reference solution is loaded and the version of the prescriptive rules to be improved is applied. This is typically the last version of the rule set created in the last iteration. The prescriptive rules and alertness model do not necessarily agree on what is an unwanted crew schedule, so the crew schedules created in the reference solution will violate several of the prescriptive rules. Rule violation statistics are compiled for the number of violations of each rule, typical limit of violated rules and typical overshoots of violated rules, example in figure 12.

From the rule violation statistics, the most limiting rules can be identified. It is possible to experiment with the rule limits to decrease the number of rule violations. Sometimes it is necessary to replace the limiting rule with a new rule in the spirit of the old rule, in order to have parameterized limits to adjust.

A set of new solutions from the prescriptive rules are created, with the newly proposed relaxations added. One new solution is created for each relaxed rule, as well as a few solutions where combinations of rules are relaxed. The productivity of the new solutions, as well as protected level of alertness, can be analyzed and data points plotted on the chart. One or several best solution candidates for new rule set can be chosen, for continuing work.

Violated rule	times	avg. limit	avg. fail value
Min rest time after a duty	475	10:06	8:12
Max flying (incl. DH) time in a duty	393	8:00	9:15
Max duty time in a duty	172	14:00	14:42
Obligatory weekly rest in any seven days	157	48:00	18:17
Max flying time in seven days	92	40:00	43:41
Max flying time between valid week rest	28	40:00	42:11

Figure 12. A report summarizing the most violated rules. This report is generated after the prescriptive rules to be improved have been applied to the solution governed only by the alertness model. We would expect the most violated rules to be those that are overly restrictive.

A case study

To validate the method described, it has been applied to the CCAR rule set on the CN data-set.

In three iterations, 9 different rule changes were tried and five rule changes were introduced. The final result was a rule set where the average block time per day was increased by 6% from 5h59m to 6h21m and alertness was improved with between 250-700 CAS-points. The alertness is compared in Figure 13 where the new rule set is named CCAR+. The new solution's resource index dropped 8.5%.

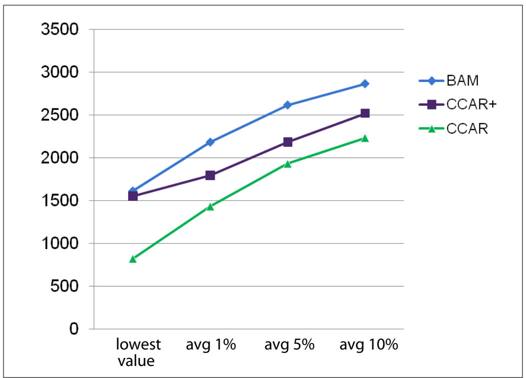

Figure 13. A comparison of the protected level of alertness between CCAR, the evolved CCAR+ and the BAM reference solution.

Introduced changes

The following changes were introduced to the rule set. The properties of the resulting rule sets were plotted in the graph in Figure 14.

a) Added a rule limiting the number of check-ins for duty on the same day (cut at 03:59 crew homebase time)

b) Reduced the maximum duty time for duties partly falling within 23:00 to 03:30

c) Relaxed rule governing maximum block-time in a duty

d) Relaxed rule governing minimum rest after duty

e) Added a complementing rule for maximum duty time for duties after short rest periods, i.e. rest periods that became legal when the original minimum rest after duty rule was relaxed.

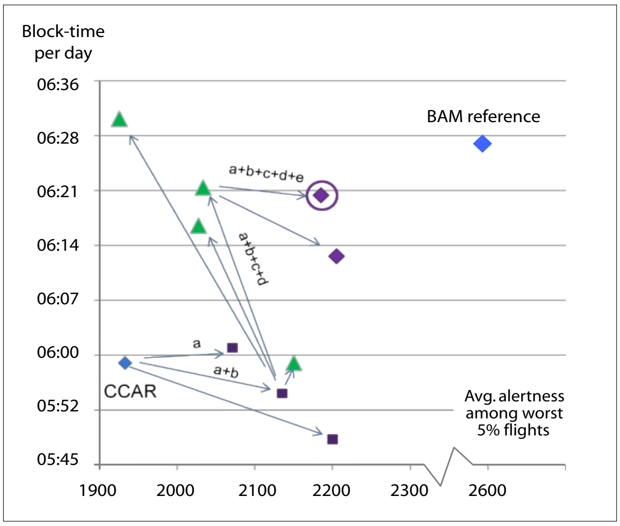

Figure 14. In our case study 9 changes were tried to CCAR and the properties of the resulting rule-set were plotted.

As can be seen in the figure above 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 good alertness and productivity.

The final rule set was stress tested. The test showed that the protected level of alertness had increased with 250 to 450 CAS-points in the different measure points.

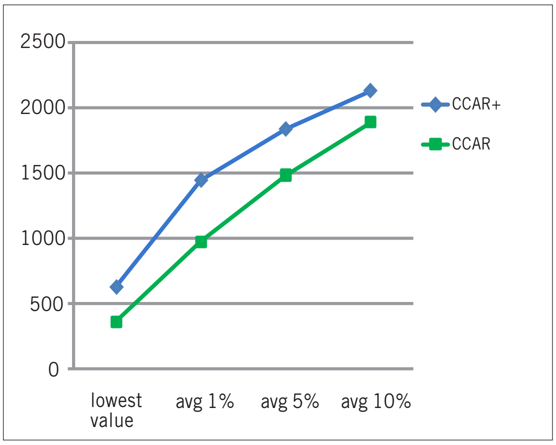

Figure 15. Protected level of alertness from stress tested CCAR and the new rule set, depicted CCAR+.

5. CONCLUSIONS

Of the three tested FTL schemes, none managed to completely protect against low alertness in the crew schedules. The most common bad patterns encountered in the crew schedules were the planning of unusable rest during day time (periods of low sleep propensity where it would be difficult for the pilots to sleep) or duties of maximum length ending close to midnight. These situations appeared in solutions generated from all FTLs. Often, but not always, these cases are caught by labor agreements or other planning rules, however at day of operations such rules are usually set aside and operations are managed by FTL only.

The JAR and CCAR rule sets are comparable in many aspects, both in productivity and in the protection against low alertness. JAR is slightly better at protecting against fatigue but also less productive if there are many legs in the average duty. FAR is the most efficient of the three FTLs but at the expense of using much duty time[9]. FAR also performs worst when it comes to protecting against low alertness.

The levels of alertness predicted by BAM for the FTLs should be viewed with caution since the model is not yet fully validated in airline operations. However, assuming the model is shown to be valid, the safety and business case for FRMS is strengthened since we see that FTLs do not protect well against alertness and that a model based scheduling is both safer and more productive.

Assuming current FTLs are to be improved as part of a move toward FRMS, we have described a method for iteratively improving an FTL, while assaying the impacts on productivity and alertness. The method presented can improve a set of rules relatively easily, to better protect against low alertness while improving or maintaining flexibility and productivity. The method can be used with any alertness model connected to an optimizer and could just as well be applied to labor agreements to ensure that alertness, productivity, and quality of life can be maintained as labor agreements are improved.

Note that the functionality used in this study purely for analysis and improvement of rules, can just as well be applied by an operator as an essential part of an FRMS.

6. REFERENCES

1. **Romig, Emma and Klemets, Tomas**. Fatigue Risk Management in Flight Crew Scheduling. *Aviation, Space, and Environmental Medicine, Volume 80, Number 12, December 2009 , pp. 1073-1074(2).*

2. **Åkerstedt, T and 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 and Portin, C**. Predictions from the three-process model of alertness. *Aviation, Space, and Environmental Medicine 2004; 75(3, Suppl.):A75-83.*

4. **Åkerstedt, T and Folkard, S**. Predicting duration of sleep from the three process model of regulation of alertness. *Occup Environ Med. 1996 February; 53(2): 136–141.*

5. **Hjorring, Curt and Hansen, Jesper**. Column Generation with a Rule Modelling Language for Airline Crew Pairing. *In Proceedings of the 34th Annual Conference of the Operational Research Society of New Zealand, December 10-11, 1999, Hamilton, New Zealand.*

6. **Bengtsson, Lennart, et al**. Railway Crew Pairing Optimization. *Carmen Research and Technology Report CRTR-0408, November 2004.*

7. **Andersson, Erik, et al**. Problem Solving in Airline Operations. *OR/MS Today, pages 36-40, April 2005.*

8. **Moebus, P**. Scientific and Medical Evaluation of Flight Time Limitations. *EASA, TS.EASA.2007.* OP.08, Issue 1.0.

[9] The high usage of duty time is an effect of the cost structure used. FAR is a very flexible rule set and with another cost structure duty time over block time can be reduced at the expense of productivity per day.