

Optimized Pilot Training for Combat Aircraft

**Guro K. Svendsen, Idar Ulvestad Haugstuen,
Ole Martin Mevassvik
Norwegian Defence Research Establishment (FFI)
Kjeller, Norway
Guro.Svendsen@ffi.no, Idar.Haugstuen@ffi.no,
Ole-Martin.Mevassvik@ffi.no**

**Henning Rorvik
Norwegian F-35 Program,
Norwegian Defence Materiel Agency
Kjeller, Norway
hrorvik@f35.mil.no**

ABSTRACT

Training of combat aircraft pilots is a complex task, with available time being one of the main challenges. The pilots train on a set of missions according to a comprehensive training programme, with formations of various sizes. Required training currency for the different missions implies that it is challenging for the pilots to train sufficiently within available time. In addition to live training in the aircraft, simulator training will also be an important part of the future Norwegian training concept for combat aircraft pilots and is integrated in the training programme.

To support the Norwegian F-35 Program, the Norwegian Defence Research Establishment (FFI) performs analyses on the training system planned for the F-35. Here, we present the simulation tool TREFF2. The tool has been developed to analyse how realizations of possible training concepts affect training efficiency. One important factor is the qualification of the pilots, which limits the possible positions they can fill in a formation. Another important factor is the distinction between planned pilot and aircraft availability, and the actual resources in use. The training system must have enough flexibility to handle unexpected events like pilot absence and sortie cancellations. TREFF2 simulates scheduling of pilots to training missions for a given planning horizon, and also the daily training utilizing the actual available resources.

TREFF2 can be used to analyse the training system and identify possible bottlenecks, taking into consideration the qualifications of the pilots in the squadron and the available number of aircraft and simulator sorties. TREFF2 provides insight into the number of sorties needed to complete the training program and whether or not there is sufficient time available for the pilots to complete their training.

ABOUT THE AUTHORS

Guro K. Svendsen is a researcher at the Norwegian Defence Research Establishment (FFI). Her research is directed towards analysis of training systems through modelling and simulation. She holds a Ph.D. in Electronics and Telecommunication (2013) and a Master of Science (MS) in Applied Physics (2006) from the Norwegian University of Science and Technology (NTNU).

Idar Ulvestad Haugstuen is a researcher at FFI. He works with simulation analysis of training systems. In 2010 he received his MS from the Department of Engineering Cybernetics, NTNU.

Ole Martin Mevassvik is a principal scientist at FFI. His main research interests are systems architecture and simulation interoperability with the focus on Command and Control (C2) to Simulation (C2SIM) interoperability. Ole Martin has participated in a number of national research projects and international activities on Modeling and Simulation. He has acted as a consultant for the Norwegian Armed Forces in several simulation acquisition projects.

Major Henning Rorvik is a deputy program manager for F-35 simulators at the Norwegian Defence Materiel Agency, F-35 Program. He has experience with the F-16, with 2000+ flight hours. He also holds a Cand. Mag. in Informatics from the University of Oslo (UiO).

Optimized Pilot Training for Combat Aircraft

**Guro K. Svendsen, Idar Ulvestad Haugstuen,
Ole Martin Mevassvik
Norwegian Defence Research Establishment (FFI)
Kjeller, Norway
Guro.Svendsen@ffi.no, Idar.Haugstuen@ffi.no,
Ole-Martin.Mevassvik@ffi.no**

**Henning Rorvik
Norwegian F-35 Program,
Norwegian Defence Materiel Agency
Kjeller, Norway
hrorvik@f35.mil.no**

INTRODUCTION

Simulation-support for military training has great potential to improve training while keeping costs acceptable. In order for such training to reach its potential, it is important that simulation-based training is an integrated part of the training programmes.

A substantial amount of high-fidelity simulator training is included in the future Norwegian training programme for F-35 pilots. Incorporating full mission simulators into the F-35 training system can increase the time and cost efficiency of the training. This is also very important for the training quality. With a 5th generation aircraft, several training aspects are much better suited for simulators. In order to fully exploit the potential of simulators, the simulator training is integrated into the pilots' daily training schedule together with training in the aircraft. A discussion on the balance of live and virtual training can be found in a 2009 study conducted by the RAND Corporation (Schank et al., 2009). The training system is also designed to be prepared for a future LVC (Live, Virtual, Constructive) capability. This will allow pilots in real aircraft to train together with pilots in simulators and constructive forces, and has a potential for reduction of the need for live Red Air¹ (Ausink et al., 2011).

The pilots train according to a comprehensive training programme with requirements on the currency for the various missions. It is challenging for the pilots to train sufficiently within the available time. Future training will consist of an increased number of training sorties, distributed on two sets of training media, i.e., aircraft and simulators. The scheduling of pilots to perform specific training missions according to the training programme is an especially complicated task. The missions require formations of various sizes, and the pilots fill different positions in the formation from mission to mission. Possible positions are limited by their qualification.

FFI performs analyses of the future training system of F-35 in support of the Norwegian F-35 Program. Models of different training concepts and training systems have been developed in collaboration with the Norwegian F-35 Program. The models are based on how training is performed at current combat aircraft squadrons and on plans for the future system. The question of whether or not sufficient training is possible with the planned training system has been investigated earlier using a simulation tool called TREFF, developed at FFI (Skogsrud and Mevassvik, 2011). One of the findings of that work was that efficient scheduling of pilots to training missions is a key element to success. In the follow-up of the TREFF analysis, we addressed the question of how this is best achieved by investigating scheduling bottlenecks. The different pilot qualifications, and the distinction between planned and actual pilot availability are important factors in this regard. A new simulation tool, TREFF2, has been developed for that purpose. A scheduling module, SOFT (Schedule Optimized Fighter Training), has been developed to handle scheduling in a given planning horizon. This planned schedule is then the basis for a simulation of the daily training utilizing the actual available resources.

In this paper we address the different aspects related to actual and planned availability of pilots in TREFF2, and how the tool can be used to identify possible bottlenecks. The training outcome is sensitive to the availability of training resources and to training conditions like weather and daylight. There is also a strong degree of interdependency among pilots, as all positions in a formation must be filled with a pilot of sufficient qualification. The training will therefore not always follow the scheduled plan. Accordingly, it is not really necessary to find a schedule that is guaranteed to be optimal. The goal of SOFT is to find a solution that is sufficiently robust and distributes the

¹ Red Air: when the pilots play the role of adversary.

training sorties evenly among the pilots. In the resulting analysis, challenging elements can be identified, allowing for a more optimal training system.

The scheduling algorithm in SOFT and the simulated conduct of the daily training in TREFF2 are specifically tailored to pilot training. However, the approach is general, and can be transferred to other domains. It is especially suitable for systems with a rotation of roles such that one trainee may fill different roles in a team, or when there is a strong dependency between the trainees such that unexpected events affecting one trainee or group of trainees have a significant impact on the other trainees. One other area where we have identified such dependencies is within training and education of an army brigade, where the cooperation between different battalions and different areas of expertise is important for the overall level of the brigade.

We will first present the overall training system, and an overview of the TREFF2 simulation tool, introducing the scheduler SOFT as the main component. The scheduling problem is then further explained with a high-level description of SOFT, which is implemented using a Constrained Programming (Tsang, 1993) algorithm. Deviations from the scheduled training are handled by a simulation of the daily training in TREFF2. The outline of this model is briefly explained before a simplified example case is presented.

MODEL OVERVIEW

The annual training programme for pilots consists of a set of missions of different categories, with different levels of complexity, currency requirements, and needs for coordination with other units (e.g., from the army). The order in which the missions are trained is therefore not arbitrary. To account for this, the annual training programme is divided into sessions with dedicated theme-periods focusing on specific missions (e.g., Air-to-Air missions). A dedicated constraint-based algorithm has been developed to ensure that the missions are distributed in a way that ensures the required currency, while distributing the workload evenly over the year.

The planning horizon for the scheduler SOFT is around 1-2 weeks, which is typically shorter than the length of the theme-periods. The planning horizon of SOFT provides predictability for the pilots and for the logistics organization, so that they can plan their activities. During this period, one attempts to minimize the deviations from the plan. For the daily training, the working hours are such that each pilot at most can perform two activities per day. An activity can be a training mission in an aircraft or in a simulator. The smallest time horizon in the model is thus the length of the daily training slots. An overview of the time horizons is given in Figure 1.

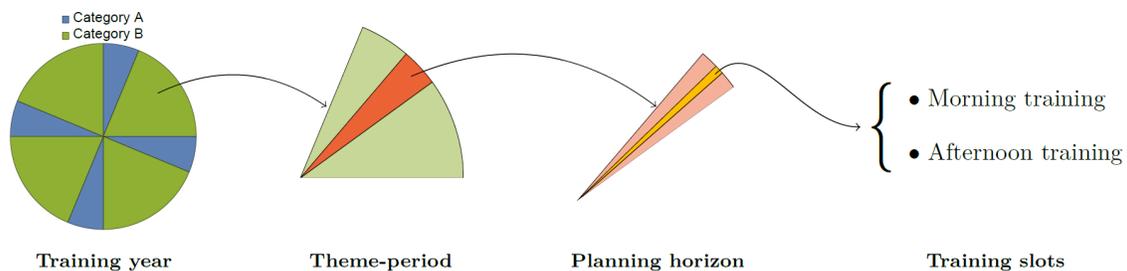


Figure 1. Relations between time frames in TREFF2. The **training year** is divided into **theme-periods**, each belonging to a certain category of training (e.g., Air-to-Air missions). An example periodization is depicted by the blue and green sectors. Within each theme-period, detailed training plans are created by SOFT for a given **planning horizon**, marked by the red sector. TREFF2 simulates the daily training according to the predefined plan and handles unexpected events. The daily training (orange sector) is performed in aircraft and/or simulators. These activities are restricted to a set of **training slots**.

TREFF2 is a discrete event simulation tool (Banks, 2010). The simulation is performed in two main steps: 1) planning of the subsequent planning horizon (scheduling); and 2) simulation of the execution of the plan. A flow diagram of how TREFF2 simulates pilot training is shown in Figure 2. The plan is developed by the scheduler SOFT, a deterministic module calculating a schedule for a given planning horizon. The schedule is based on the training progress of the individual pilots, expected training device availability, and on planned absences of pilots such as vacations. During the next step, the execution of the plan is simulated. This step includes stochastic factors like the number of available training devices, snag at start-up, weather cancellations, and unforeseen absence of

pilots. It is also possible to bypass the daily training simulation in TREFF2 and run SOFT as a standalone schedule planner.

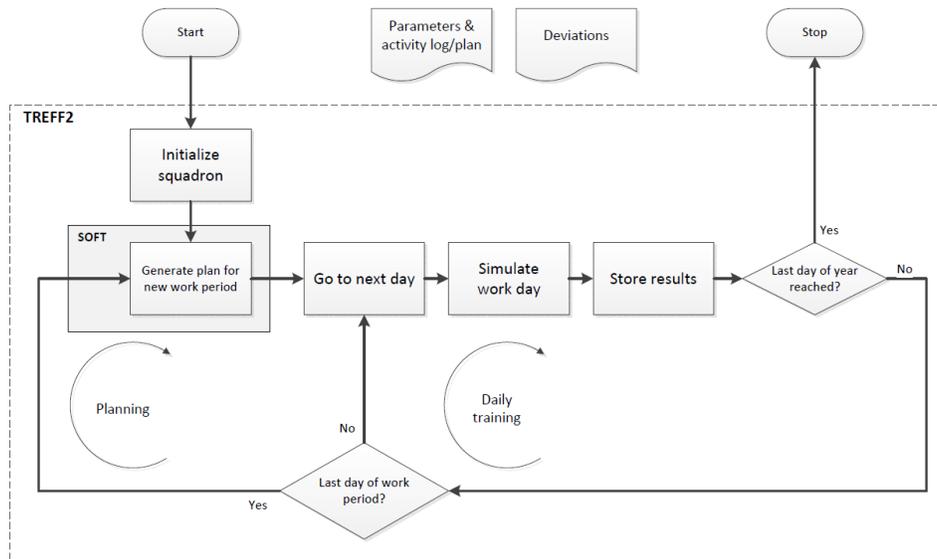


Figure 2. TREFF2 flow diagram. The planned training schedule is provided by SOFT, and used as a basis for the simulation of daily training. The actual training may deviate from the plan due to a set of disrupting events. Results are stored in two files, keeping track of both the accomplished activities, and how they deviate from the original plan.

SOFT

Previous TREFF-analyses (Skogsrud and Mevassvik, 2011) investigated the time available for the pilots to train in live aircraft and simulators, considering available work-hours and capabilities of the logistics/aircraft maintenance. The result was a set of possible combinations of daily activities, with the main activities being training in simulators or aircraft. Each day, this training is conducted in specific time-slots where training resources are available. The scheduling is therefore reduced to matching the pilots and missions to the available training resources in each of the training slots.

The challenge at hand deviates from a classical scheduling problem (Minton et al., 1993) (Frost and Dechter, 1999) in that the exact times for the available training slots are predefined. In terms of classical operational analysis, it resembles more the “job-shop” problem, where it is a matter of performing the job (mission) best contributing to minimizing the cost efficiency ratio (Fox, 1983). The problem at hand is restricted due to the qualifications of the pilots and constraints on each pilot's activities within a day and for subsequent days. It should be expected that deviations from the planned SOFT schedule will occur, due to unpredicted events. This implies that a schedule should contribute to a good foundation for later scheduling periods and for handling deviations. As a consequence, the aim is not really to find the optimal solution for one scheduling period, but to find a sufficiently good solution in a fair amount of time. The pilot training scheduling is characterized by a very big possibility space, a multifaceted goal function without the need to guarantee an optimal solution, and a set of constraints. Constrained Satisfaction Programming (CSP) has therefore been chosen as the means to obtain a schedule (Brailsford et al., 1999) (Tsang, 1993). The algorithm is implemented using Google OR-Tools (Google Inc., n.d).

The CSP method was developed to handle large search trees by efficiently propagating a set of constraints to prune the search tree. A general constrained satisfaction problem is formulated in terms of variables with certain domains, subject to a set of constraints. The core to solving a CSP problem is to assign valid values to all variables. The variable-value pairs are called *labels*, and are the basis for defining constraints and solutions. A *tuple* is a composite label combining several simultaneously valid labels. The synthesis of a CSP consists of searching for valid labels and combining them to build tuples.

The synthesis is visualized as a search tree in Figure 3a. Each node corresponds to a decision, and the chain of consistent decisions corresponds to a branch. As the domains in general are of different size, the ordering of the variables within the search tree will significantly alter the efficiency of the search. If the synthesis reaches a level where there are no more valid labels, it is necessary to do backtracking, i.e. changing the label in previous level (for the previous variable) to another valid label (try another value from the domain), and retry. The efficiency of backtracking is highly dependent on the geometry of the search tree in relation to the constraints. If the number of variables is large, the variable-assignment should be such that those variables with a high degree of connectivity are close to each other. This will reduce the need for backtracking in multiple steps, as illustrated in Figure 3b.

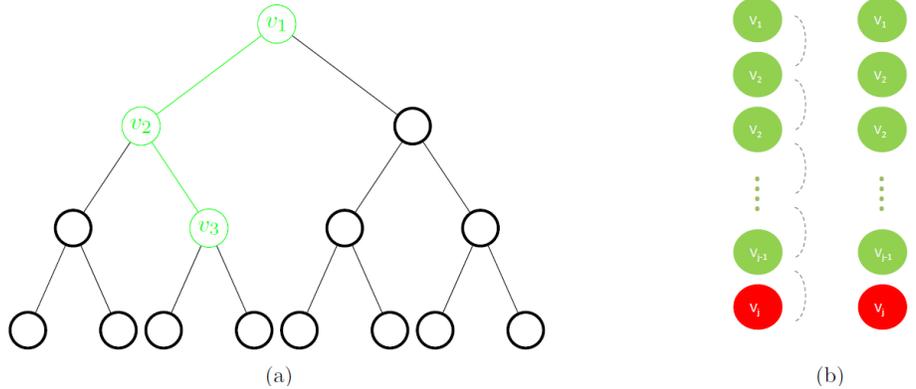


Figure 3. (a) Tuple of order 3, corresponding to a branch in the search tree. (b) A constraint causes a failure at variable v_j . In the example to the left, this constraint connects v_{j-1} and v_j and it may be sufficient to backtrack one step. On the right side, the constraint connects v_j and v_1 , and it will be necessary to backtrack through all variables.

In our problem at hand, the pilot p has a requirement to train mission m a number r_m^j times using training device j . The available resources are aircraft for live training, $j = l$, and simulators, $j = s$. Let $n_{m,p}^j$ be the number of completed missions. The end goal is then to achieve $n_{m,p}^j \geq r_m^j \forall j, m, r, p$. For a given day, the aim of SOFT is to assign missions and pilots to the available training devices in each of the training slots. The approach chosen in SOFT in transforming the scheduling problem into a CSP is:

- Assign pilot p (or no-one) to the mission m , for each day d , and all training slots s .

The above formulation is chosen as it groups variables with a high degree of connectivity close to each other. There is a strong connectivity between the pilots of a given formation, training on one mission, so assigning the pilots to one mission at a time speeds up backtracking (Figure 3b).

The search tree in SOFT is large, and finding a true optimum would be time-consuming. As mentioned this is not really necessary as there will most likely be deviations from the scheduled plan. The goal of SOFT is to find a sufficiently good solution as a basis to understand the limiting factors and bottlenecks of the training system.

Filling all positions in a formation may require pilots to train more missions than required from the training program. Pilots must also fill the Red Air formations. These kinds of activities will also provide training value, but do not give credits in terms of the training program. The search is performed in two steps, first assigning the training that contributes to the training programme progression, and then assigning pilots to fill formations.

There are a number of constraints on the schedule, such as:

- Training for the given planning horizon:
 - Valid training is restricted to the training programme of the current theme-period.
 - The training sorties should be evenly distributed among the pilots when possible.
 - Some pilots are prevented from training due to planned absence such as vacation.
- Training for a given day:
 - Allowed schedule based on shift (morning/evening).
 - Allowed combination of activities during one day.

- Training for a given slot:
 - Sufficient number of pilots to fill the formation.
 - Sufficient qualifications of the pilots.
 - Sufficient number of training devices.
 - No cross-scheduling of pilots from different squadrons in one formation.

From an operational point of view, it is an advantage to schedule the missions in a certain order. The CSP algorithm will attempt to assign missions in rising order according to the specific mission (mission ID).

TREFF2

TREFF2 is a discrete event simulation of the execution of F-35 training, taking into account the stochastic variations that may affect the training. The model is implemented using the simulation tool AnyLogic (The AnyLogic Company, n.d.). The randomness of the model can be divided into two groups: 1) the distribution of planned absence such as vacation and off-squadron activities; and 2) unexpected events like absence of pilots, failures of aircraft at start-up (“snag”), problems with simulators, and weather-conditions preventing training. Planned events are given as input to SOFT. Whenever there is an event preventing the planned training, the daily training simulation in TREFF2 will attempt to find solutions, such as finding an available substitute pilot, allowing the mission to be trained with a reduced formation size, or assigning a backup mission.

EXAMPLE CASE

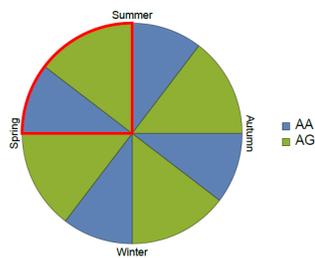
To exemplify the TREFF2 approach, we consider a simplified use case for which the parameters are summarized in Table 1. The parameters in this example are for demonstration only, and are not descriptive of the Norwegian squadrons. The pilots are separated into two equal squadrons; squadron *A* and squadron *B*. They follow a shift rotation such that squadron *A* starts earlier on even-numbered weeks, and squadron *B* has the early shift on odd-numbered weeks. The pilots are categorized according to their qualification, describing what formation size they are able to lead. The qualifications of relevance here are *wingman* (*w*), *two-ship lead* (*t*) and *four-ship lead* (*f*).

Table 1. Parameters for the example use case.

Description	Value
<i>Deterministic parameters</i>	
Days to complete training (Figure 4a)	$N_d = 65$ (week days)
Total number of pilots	$N_p = 32$
Number of (equal) squadrons	$N_{sqn} = 2$
Ratio of squadron with given qualification	$\eta(w) = 1/3, \eta(t) = 1/3, \eta(f) = 1/3$
Number of aircraft (max. available for training)	$N_{AC} = 10$
Number of simulators	$N_{sim} = 6$
National holidays (month/day)	4/13, 4/14, 4/17, 5/1, 5/17, 5/25, 6/5
SOFT	10 days (2 weeks)
Vacation per pilot	4 weeks + 5 days (dates are drawn from distribution)
<i>Stochastic parameters (TREFF2)</i>	
Probability of pilot being sick	$p_{sick} = 0.065$
Probability of weather cancellation	$p_w = 0$
Availability per simulator	$a_{sim} = 1$
Probability of aircraft snag	$p_{snag} = 0$

Planned pilot absence is divided into individual absence (vacation) and non-working days due to official holidays (e.g., Easter and national holidays). The absence is drawn from a probability distribution, and provided as input to the SOFT scheduler. The distribution for vacation is such that it is most likely for a pilot to be on vacation during the summer. The stochastic parameters for weather, sickness, and resource availability cause deviations from the training plan. In this example we only consider deviations related to pilot availability, represented as a probability for sickness.

The pilots train according to an annual training programme with Air-to-Air (AA) missions, and Air-to-Ground (AG) missions. The missions are further categorized according to their (Blue Air) formation size, and the number of aircraft required for Red Air. One calendar year is divided into eight theme-periods, four periods with AA missions and four AG periods (Figure 4a). This number of periods is based on the currency requirements of the missions. The periods are aligned with the spring equinox, such that it is possible to train night missions during two periods in each category. In this example use case, we only consider results from one quarter of the year, corresponding to one theme-period for each category. The training programme for these theme-periods is indicated in Figure 4b.



Mission name	Min blue air	Requested blue air	# Live sorties in period	# Sim sorties in period	Min red air	Requested red air	Night mission?
AA1	1	1	1	0	0	0	FALSE
AA2	2	2	2	1	0	0	FALSE
AA3	2	2	1	1	1	1	FALSE
AA4	2	2	1	1	1	1	FALSE
AA5	2	4	1	1	1	2	FALSE
AA6	2	4	1	1	0	0	FALSE
AA7	4	4	1	1	0	1	FALSE
AA8	4	4	1	1	1	1	FALSE
AG9	1	1	1	1	0	0	FALSE
AG10	1	4	2	1	0	1	FALSE
AG11	1	2	2	0	1	1	FALSE
AG12	2	2	1	1	0	0	FALSE
AG13	2	4	1	1	1	1	FALSE
AG14	4	4	1	1	0	2	FALSE
AG15	4	4	2	0	2	2	FALSE

(a)

(b)

Figure 4. The two theme-periods under consideration cover one quarter of the year (4a). Their corresponding training programme is shown in (4b). Formation sizes below those requested are only permitted in case of deviations.

Two-week Schedule

A two-week schedule for the squadrons is shown in Figure 5. The schedule shows activities that were actually performed and includes missions flown as well as planned and unplanned absence. Disruptions from the original plan are marked. The workload per pilot seems to be reasonably distributed.

Pilot	Mon	Tue	Wed	Thu	Fri	Mon	Tue	Wed	Thu	Fri
1 (w)	S S	3 ²	1 ¹	4 ²	2 ²	2 ²	3 ²	V V	3 ²	
2 (t)	2 ²	3 ²	1 ¹	4 ²	S S	S S		3 ²	R	R
3 (w)	2 ²	3 ²	1 ¹	S S	2 ²	4 ²		3 ²		
4 (t)	2 ²	3 ²	1 ¹	2 ²	2 ²	4 ²	R	3 ²	3 ²	
5 (w)	2 ² 3 ²	1 ¹	4 ²	2 ²	2 ²	3 ²	3 ²		R	4 ²
6 (f)	2 ² 3 ²	1 ¹	4 ²	2 ²	S S	R	3 ²	R		4 ²
7 (w)	2 ² 2 ²	1 ¹	3 ² 4 ²	2 ²	R	R	3 ²	R	R	4 ²
8 (f)	2 ² 2 ²	1 ¹	3 ² 4 ²	2 ²	2 ²		3 ²	5 ⁴		4 ²
9 (w)	2 ² 2 ²	1 ¹	3 ² 4 ²	2 ²	5 ⁴	3 ²	3 ²			S S
10 (f)	2 ² 2 ²	S S	3 ² 4 ²	2 ²	5 ⁴		3 ²			4 ²
11 (w)	1 ¹	2 ² 3 ²	2 ²	2 ²	4 ² 6 ⁴	3 ²	R		4 ²	
12 (f)	1 ¹	2 ² 3 ²	2 ²	2 ²	4 ² 6 ⁴	3 ²	R		4 ²	
13 (w)	1 ¹	2 ² 2 ²	2 ²	3 ² 4 ²	5 ⁴ 6 ⁴	3 ²	3 ²	R	4 ²	4 ²
14 (f)	1 ¹	2 ² 2 ²	2 ²	3 ² 4 ²	5 ⁴ 6 ⁴	3 ²		3 ²	4 ²	
15 (t)	1 ¹	2 ² 2 ²	2 ²	S S	3 ²	2 ²	S S			R
16 (f)	1 ¹	2 ² 2 ²	2 ²	3 ² 4 ²	3 ²			4 ²		R

Live Simulator (V)vacation/(S)sick Other Deviation

(a) Squadron A

Pilot	Mon	Tue	Wed	Thu	Fri	Mon	Tue	Wed	Thu	Fri
17 (w)	2 ²	1 ¹	1 ¹	2 ²	2 ²	2 ²	3 ²	3 ²	R	4 ²
18 (t)		1 ¹	1 ¹	S S	2 ²	R	3 ² 4 ²	3 ²	3 ² 3 ²	
19 (w)		2 ²	1 ¹	2 ²	2 ²	R	3 ² 4 ²	3 ²	4 ²	
20 (t)		2 ²	1 ¹	2 ²	2 ²	4 ²	3 ² 3 ²	2 ²		4 ²
21 (w)	2 ²	1 ¹	2 ²	2 ²	2 ²	R	S S		R	4 ²
22 (f)	2 ²	1 ¹	2 ²	2 ²	2 ²	3 ² 4 ²	3 ²	4 ²	R	R
23 (w)	2 ² 2 ²	1 ¹		2 ²	R	S S	3 ²	R		4 ²
24 (f)	2 ² 2 ²	1 ¹	3 ²	2 ²	2 ²	3 ²	2 ²	3 ²	4 ²	4 ²
25 (w)	2 ² 2 ²	S S	3 ²	2 ²		3 ²	R	4 ²	S S	4 ²
26 (f)	2 ² 2 ²	1 ¹	2 ²	2 ²	3 ²	3 ²	2 ²	4 ² 5 ⁴		4 ²
27 (w)	1 ¹	V V	V V	V V	V V	V V	2 ² 2 ²	2 ²	3 ² 3 ²	4 ²
28 (f)	1 ¹		S S	2 ² 2 ²	3 ²	3 ²		S S		4 ²
29 (w)	1 ¹	2 ² 2 ²	2 ²	3 ²	4 ²	S S	R	R	4 ²	R
30 (f)	1 ¹	2 ² 2 ²	2 ²	3 ²	S S	3 ²			4 ²	4 ²
31 (t)	1 ¹	2 ² 2 ²	2 ²	3 ²				4 ² 5 ⁴	4 ²	4 ²
32 (f)	1 ¹	2 ² 2 ²	2 ²	3 ²	3 ²	2 ²	3 ²	4 ² 5 ⁴	4 ²	R

(b) Squadron B

Figure 5. Extracted two-week (3/20 – 3/31) schedule for the two squadrons. The numbers in the training slots refer to the ID of the missions flown, with the corresponding formation size as superscripts. Absence due to sickness is denoted with an “S”. Deviations are marked with a red border. The pilot qualification is also indicated (wingman, two-ship lead and four-ship lead).

The two-week schedule shows the alternating shift rotation for the two squadrons. In the first week (Mon-Fri), members of squadron *A* are scheduled to perform live missions in the first daily training slot. The same week, squadron *B* is scheduled with live missions in the second daily training slot. The shift is alternated in the second week, and squadron *B* starts early. There is not sufficient time for one pilot to train more than one live sortie per day; so the schedule shows no double flight sorties. Simulator training takes less time, and this restriction does not apply. Hence, several pilots train in two simulator sorties in one day. Combinations of live and simulator training are only possible when training matching missions, in which case the time needed for briefing is reduced.

An overview of the deviations in this two-week period is shown in Figure 6. When there is a need for a pilot to fill in for an absent pilot, all vacant pilots are considered (i.e., pilots can also fill in for pilots from the other squadron). The number of sorties with deviations is often larger than the number of absent pilots, as they may have training scheduled for more than one sortie. For example; on the fourth day, there were 3 absent pilots, with a total of 5 sorties planned. Only one pilot was free to fill in.

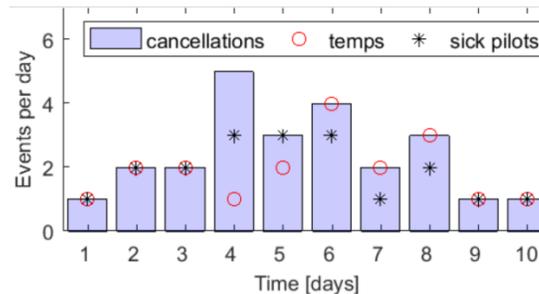


Figure 6. Chart of the deviations per day (i.e., sorties where training is not conducted according to the schedule). The period corresponds to Figure 5. The number of sorties affected by deviations per day is shown as bars, along with the number of absent pilots and the number of pilots filling in as temps.

Training Evolution in the Theme-periods

Figure 7 shows the activity level as a function of time in each of the theme-periods. In the beginning, the pilots train simpler missions. These often have small formation sizes, and it is easy to create schedules fully utilizing the available resources. As formation sizes increase, and pilots start to complete some missions, the activity level drops. This is an indication that the number of resources, the number of pilots, and the training programme could be better aligned. The fraction of pilots to the number of training media is so small in this example case that it is difficult to fully exploit the given resources. Towards the end of the AG-period, the activity level decreases due to the number of pilots on vacation. A chart of the deviations occurring in the two periods is given in Figure 8. Note that we have omitted the effects of aircraft snag, simulator snag, and weather cancellation.

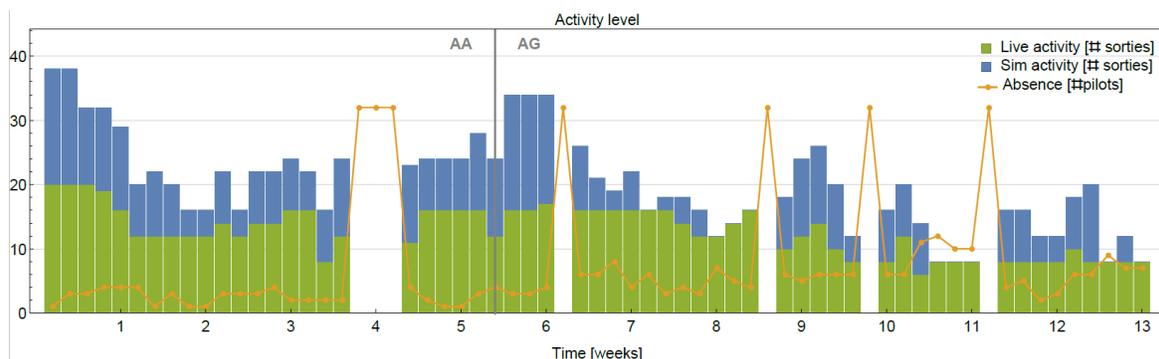


Figure 7. (a) Daily activity level. The plot shows the total number of training sorties each day, separated by in aircraft (live) and simulator sorties. Absences are also shown, consisting of; vacations, holidays, and absent pilots. Note that the number of pilots on vacation rises towards the end of the AG-period, as summer approaches, resulting in a decreased activity level.

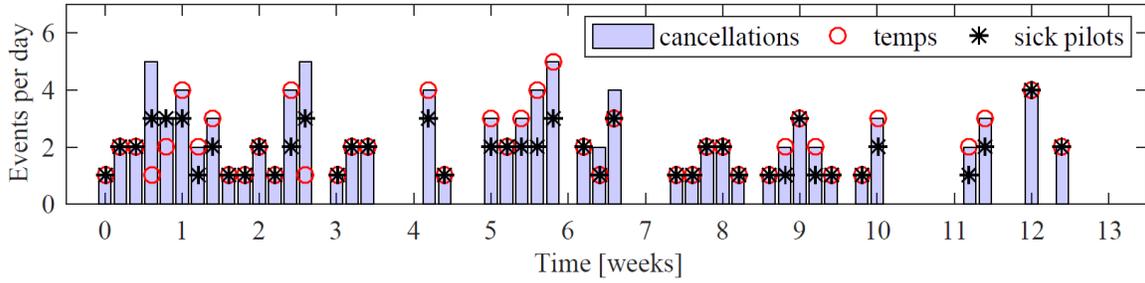


Figure 8. Deviations from planned training (all weeks). As in Figure 6, the bars show the number of cancellations for the pilots per day. The majority of these cancellations is covered by other pilots as shown by the red circles. The black stars indicate the number of absent pilots per day. Note that the number of cancellation events in some cases is larger than the number of absent pilots, which is mainly due to the absent pilots having more than one event planned. Other effects may also occur, such as pilots having to cancel their mission because an important member of the formation could not be replaced.

In order to evaluate the pilot training progress, we define two time-dependent training progress indicators. The first indicator, $TP_{p,full}$, gives a long term view on a pilot’s training progress (Equation 1), while the second indicator, $TP_{p,tp}$, gives a shorter term picture for one given theme-period (Equation 2).

$$TP_{p,full}(d) = \min\left(1, \frac{b_p(d)}{b_{ref}(d)}\right) \tag{1}$$

$$TP_{p,tp}(d) = \min\left(1, \frac{b_p(d) - b_p(d_0)}{b_{ref}(d) - b_{ref}(d_0)}\right), \text{ where } d_0 = \begin{cases} 0 & \text{for period 1} \\ 27 & \text{for period 2} \end{cases} \tag{2}$$

They indicators both compare a pilot’s registered training (b_p) against a period-wise linear progress reference (b_{ref}) based on the mission contents of the training programme of each theme-period. Per definition, the indicators are identical for the first theme-period. An indicator of unity at the end of the training implies that all required missions have been trained. Figure 9(a) shows the simulated training progress averaged over all pilots. At the beginning of a new period, the indicator will be very sensitive, as the number of missions in the progress reference is small. The dip towards the end of the AA period is of greater importance. This dip is partly due to the Easter vacation, and partly due to the fact that there are more missions requiring large formation sizes. A training progress level of unity is not achieved in any of the periods. As a reference, we also show in Figure 9(b), a similar figure based on a simulation where all training is accomplished according to the initial plan.

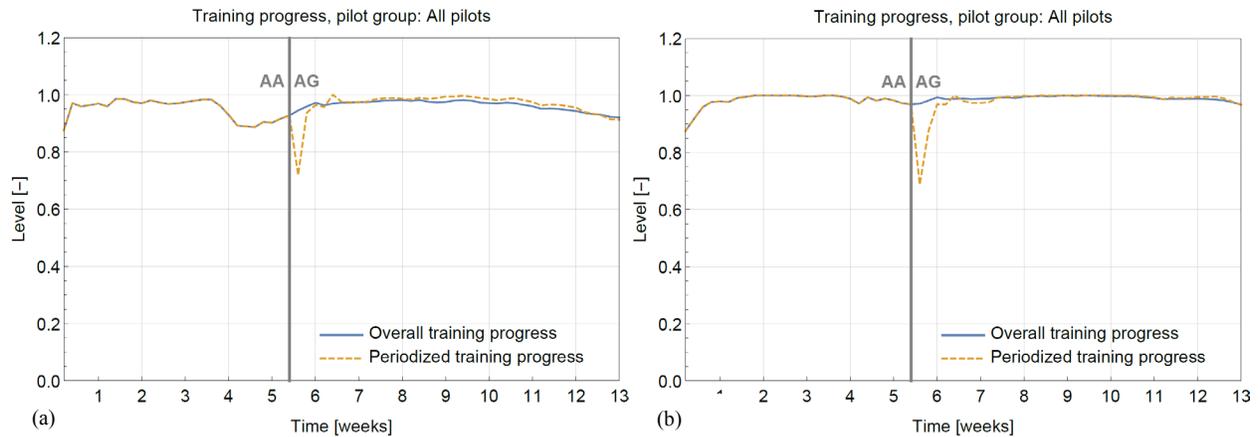


Figure 9. Training progress indicators averaged over all pilots. Subfigure (a) is from a simulation with deviations due to unexpected events, and subfigure (b) show the corresponding data in the absence of deviations.

It is important that a sufficient number of pilots is qualified to lead larger formations. For example, if there are too few four-ship leads, this may lead to a lot of excess flights for that group. Figure 10 shows the training progress for that group alone, which is relatively similar to that of the whole group of pilots. Excessive sorties do not seem to cause a reduction in the training progress of that group.

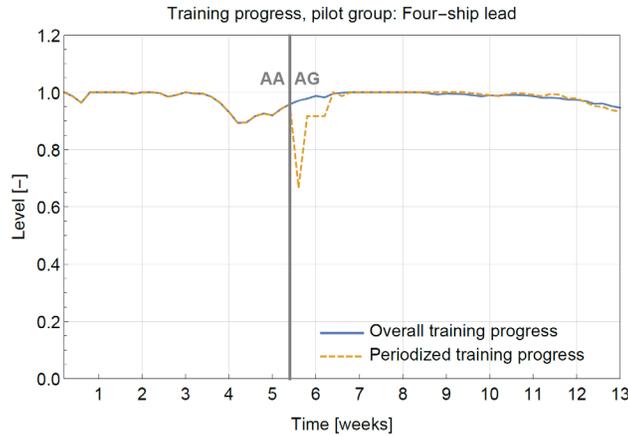


Figure 10 Training progress indicators for four-ship leads.

Overall metrics

Some metrics from the overall training per pilot are shown in Figure 11. The workload is distributed fairly well on all pilots, but there is a variation. One can also see that there are a lot of pilots missing some missions.

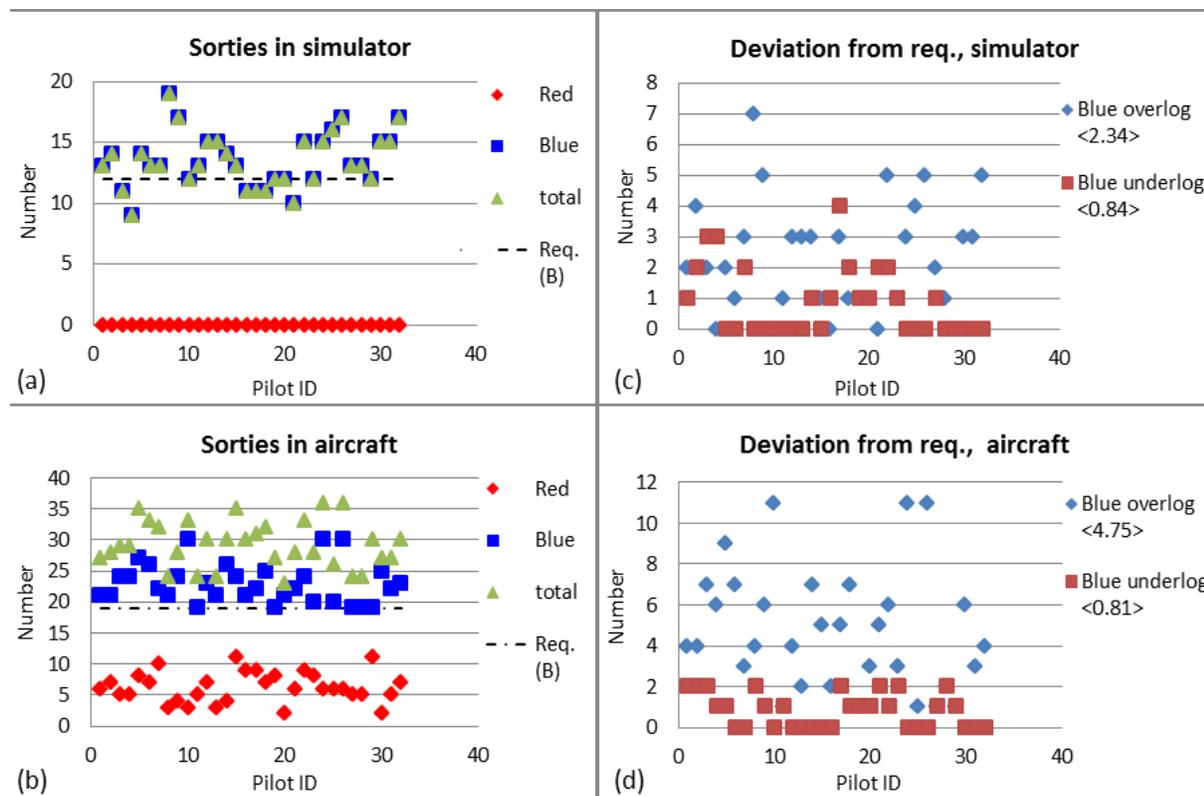


Figure 11. Overall metrics for completed training per pilot, including stochastic deviations. The upper row concerns the training conducted in simulator, and the lower row concerns live training in aircraft. Subfigures (a) and (b) display the number of sorties trained per pilot. The total number of sorties is also split into sorties flown with Blue and Red tactics. For comparison, the required number of (Blue Air) sorties is also included. To the right, subfigures (c) and (d) look more closely into the deviations between sorties flown per pilot, and the requirement from the training program. The mean values of the deviations for all pilots are also included in the legend.

To identify the cause of the insufficient training, some metrics are repeated in Figure 12 for the case without disrupting pilot absence. When comparing the figures with and without pilot absence, one sees that most of the insufficient training is due to absence. However, it is clear that the margins are too small if the ambition is to have two fully trained squadrons.

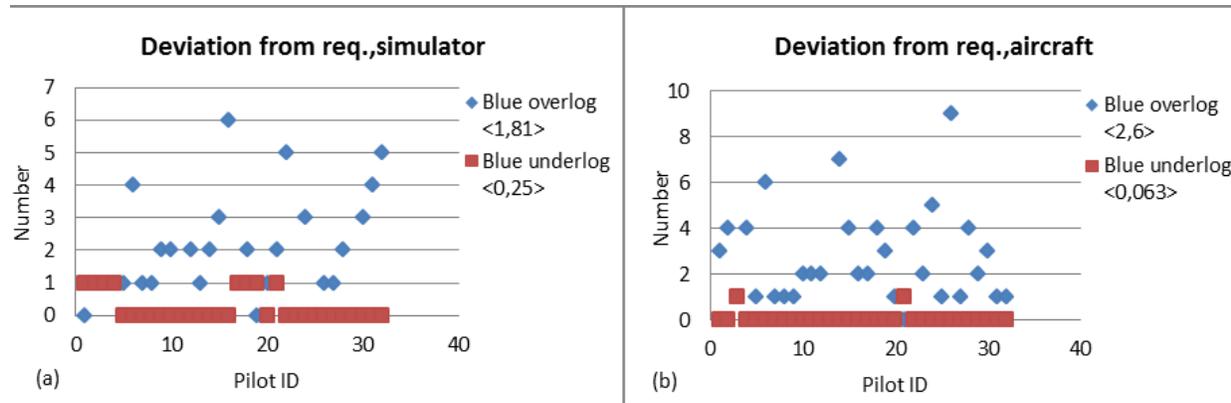


Figure 12 Overall metrics for completed training, in the absence of stochastic deviations.

In Figure 13, the deviations for the live training are plotted as a function of missions, rather than pilots. It is clear that mission 15 is problematic. This is in part because it is the last mission, but also because mission 15 has a total group size of 6, which is a difficult size with 10 available aircraft.

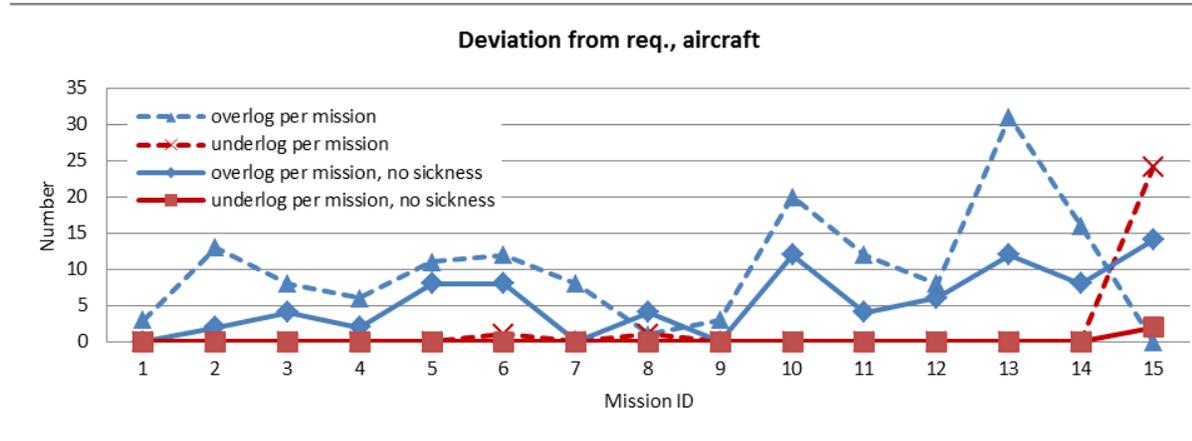


Figure 13. Deviation from training programme requirement, shown as a function of missions. Data is included both for the case with and without unexpected pilot absence.

FURTHER WORK AND CONCLUSIONS

FFI has supported the Norwegian F-35 Program since 2010 with analysis of training systems for new combat aircraft. The TREFF2 discrete event simulation tool has been developed to analyse bottlenecks and challenges complicating the completion of a training programme for combat aircraft pilots. A two-step approach is used to simulate pilot training. In the first step, a schedule is created for a given planning horizon. A constraint satisfaction algorithm is used to distribute the training sorties among the pilots. In the second step, the conduct of training is simulated to address unexpected factors. TREFF2 provides a tool for analysis of training systems, to ensure that the system of training devices and pilots is well balanced in order to fulfil training requirements.

An example case is presented to demonstrate the tool. Note that the parameters for this case are for demonstration only, and are not descriptive of the Norwegian squadrons. By using TREFF2, we could estimate the pilots' workload, and to which degree the given training programme could be completed within the time available. The results show that the training programme is a bit too ambitious. There are missions that were not sufficiently trained,

even in the case where all training goes according to plan. With deviations from the plan, this number increases significantly. For the training in aircraft, this was mainly due to missions with a formation size that is poorly matched with the maximum available number of aircraft. Also, scheduling becomes ineffective when the number of available pilots decrease in low activity periods. Designating smaller missions to such periods, or merging of the squadrons during low-activity periods are possible solutions.

In further work, we will analyse the stochastic effects systematically by running a Monte Carlo experiment with a number of replications. In support of the Norwegian F-35 Program, TREFF2 will be used to evaluate realizations of alternative training concepts for the Norwegian F-35 squadrons to ensure that the pilots receive the required amount of training.

A number of experiments can be conducted to support the development of training concepts for F-35. Some relevant factors to evaluate are: the number of missions and currency requirements in the training programme, the number of aircraft available for training, the length of the planning horizon for scheduling, segmentation of the calendar year into theme-periods, and the distribution of different pilot qualifications.

Evaluation of the possible alternatives can help decision makers design the training with reduced risk of insufficiently trained pilots. This also makes it possible to establish a connection between training outcome and costs for training devices, maintenance and number of trainees.

REFERENCES

Ausink J.A. et al. (2011). *Investment Strategies for Improving Fifth-Generation Fighter Training (TR-871-AF)*, California: RAND Corporation.

Banks J. (2010). *Discrete-Event System Simulation, Fifth Edition*, Pearson Education Inc.

Brailsford S.C. et al. (1999). Constraint satisfaction problems: Algorithms and applications, *European Journal of Operational Research*, vol. 119, pp 557-581.

Fox M. S. (1983). *Constraint-directed search: a case study of job-shop scheduling*, Ph.D. dissertation, Carnegie Mellon University.

Frost D. & Dechter R. (1999). Maintenance scheduling problems as benchmarks for constraint algorithms, *Annals of Mathematics and Artificial Intelligence*, vol. 26, pp. 149–170.

Google Inc. (n.d.). Google OR-Tools [Computer software], Available [Online]: <http://code.google.com/p/or-tools/>

Minton S. et al. (1993). Minimizing conflicts: A heuristic repair method for constraint-satisfaction and scheduling problems, *Journal of Artificial Intelligence Research*, vol. 1, pp. 25-32.

Schank J.F. et al. (2009). *RAAF F-35 Operational Pilot Training, Balancing Simulator and Live Training*, California: RAND Corporation.

Skogsrud G. & Mevassvik O. M. (2011). Simulation of a Combat Training Program for Norwegian Pilots, *Enhance or Replace: Finding the Right Live vs. Synthetic Balance (RTO-MP-MSG-087)*, pp. 20-1 – 20-8.

The AnyLogic Company (n.d.). AnyLogic [Computer software], Available [Online]: <http://www.anylogic.com/>

Tsang E. (1993). *Foundations of constraint satisfaction*, London and San Diego: Academic Press.