Modelling the Production and Absorption of Pilots

The Development of the Production, Absorption and Retention Simulation (PARSim)

Dr. Norman C. Corbett
Department of Mathematics & Statistics, Okanagan College

DRDC CORA TR 2013-023
March 2013

Defence R&D Canada

Centre for Operational Research and Analysis

Operational Research and Analysis Directorate
1 Canadian Air Division/
Canadian NORAD Region Headquarters
Comd 1 Cdn Air Div, A1/A4, A1 Trg

National Defence Défense nationale

Canada

Modelling the Production and Absorption of Pilots

The Development of the Production, Absorption and Retention Simulation (PARSim)

Dr. Norman C. Corbett
Department of Mathematics & Statistics, Okanagan College

Defence R & D Canada - CORA

Technical Report

DRDC CORA TR 2013-023

March 2013

Author

Original signed by

Dr. Norman C. Corbett

Approved by

Original signed by

Dr. R.E. Mitchell
Head Maritime and Air Systems OR

Approved for release by

Original signed by

Paul Comeau
Chief Scientist–DRDC CORA

The information contained herein has been derived and determined through best practise and adherence to the highest levels of ethical, scientific and engineering investigative principles. The reported results, their interpretation, and any opinions expressed therein, remain those of the authors and do not represent, or otherwise reflect, any official opinion or position of DND or the Government of Canada.

Abstract

This report is a record of the technical aspects pertaining to the development of the Pilot Production Absorption Retention Simulation (PARSim). The simulation was produced at the Operational Research and Analysis Directorate at 1 Canadian Air Division (1 Cdn Air Div) Headquarters under the sponsorship of Commander 1 Cdn Air Div, A4/A1 and A1 Trg. The report is intended to be a guide for analysts: providing a basis for the maintenance of PARSim and the development of similar simulations.

PARSim is a model of the "flow" of pilots from formal Undergraduate Pilot Training (UGPT) to the various major operational communities and is designed to give decision makers the ability to gauge the effect of factors, like attrition, production and flying rates on the overall health of the pilot system. The simulation is comprised of two main modules: a course module, which is used to build UGPT portion of the simulation and the mentor module, which is a generic model of the major operational communities. An important feature of the mentor module is that it simulates the *transfer of experience* to new pilots, by dynamically adjusting the rate at which this transfer occurs according to health of the community and the availability of resources.

Résumé

Ce document est un rapport des aspects techniques pertinents au développement de la Simulation de la production, de l'absorption et de la conservation des pilotes (PARSim). Cette simulation a été produite par la Direction d'analyse et de recherche opérationnelle (DARO) à 1re Division aérienne du Canada (1 DAC) sous la direction du Commandant de 1 DAC et de A1 Instruction. Ce rapport est visé aux analystes afin de les guider : leur offrant une base pour maintenir PARSim ainsi que pour développer des simulations semblables.

PARSim est un modèle du mouvement des pilotes de la formation officielle de premier cycle (UGPT) vers diverses communautés opérationnelles signifiantes et a été créé afin de donner aux décideurs l'habileté de mesurer l'effet de certains facteurs, tels que l'attrition, la production et l'activité aérienne sur la santé globale de l'ensemble des pilotes. La simulation comprend deux modules principaux : un module des cours, qui est utilisé afin de construire la partie UGPT de la simulation, et le module mentor, un modèle général des communautés opérationnelles principales. Une caractéristique importante du module mentor est la simulation du *transfère d'expérience* aux nouveaux pilotes. Celle-ci est atteinte en modifiant dynamiquement le taux de transfère selon la santé de la communauté et des ressources disponibles.

This page intentionally left blank.

Executive summary

Modelling the Production and Absorption of Pilots: The Development of the Production, Absorption and Retention Simulation (PARSim)

Dr. Norman C. Corbett; DRDC CORA TR 2013-023; Defence R&D Canada – CORA; March 2013.

This report is a record of the technical aspects pertaining to the development of the Pilot Production Absorption Retention Simulation (PARSim). The simulation was produced at the Operational Research and Analysis Directorate (ORAD) in 1 Canadian Air Division (1 Cdn Air Div) Headquarters and is the result of more than a year of developmental effort by the ORAD team. The project was sponsored by the Commander 1 Cdn Air Div, A4 and A1 Training. This document is intended to be a guide for analysts: providing a basis for the maintenance of PARSim and the development of similar simulations.

To this end, we give mathematical descriptions of the major components of PARSim and discuss some of their important properties. Our primary objective is to provide the analyst with an understanding of PARSim that is independent of its implementation. Our intent in doing so is to facilitate future developments that will not be constrained to specific software packages or modelling methodologies.

PARSim is a model of the movement or "flow" of pilots through formal Undergraduate Pilot Training (UGPT) to the various major operational communities. The model is designed to give decision makers the ability to gauge the effect of factors, like attrition, production and flying rates on the overall "health" of the pilot system. PARSim is essentially comprised of two distinct modules: a *course module*, which is used to build the UGPT portion of the simulation and the *mentor module*, which is a generic model of the major operational communities. An important feature of the mentor module is that it approximates the process of upgrading inexperienced pilots. This is accomplished by dynamically adjusting the rate at which inexperienced pilots fly according to the health of the community and the availability of resources. For example, if the experience levels or the flying rates at the unit decrease, then the time that it takes to upgrade ab-initio pilots increases.

The user can construct and investigate a wide variety of scenarios by specifying 162 input parameters: setting values that determine production, attrition rates, flying rates, etc. Subsequent model runs produce a variety of output information, which includes experience levels, yearly flying rate (YFR) requirements for operational squadrons and waiting time estimates for the various personnel awaiting training (PAT) pools that are distributed throughout the pilot system. The analysis of this information allows the user to estimate the effects of the prospective changes that are encoded in the input parameters.

One of the primary motivations for the development of PARSim was promulgated by the Chief of the Air Staff (CAS). In 2000 CAS directed that the pilot military occupation (MOC) was to be regenerated to meet a preferred manning level (PML) of about 1,500 pilots. However, this regeneration was complicated by the looming expiration of the Pilot Terminable Allowance (PTA) program, which was scheduled to occur in 2003.

An extreme increase in the attrition of pilots was anticipated as a result of the PTA expiration. This, and low experience levels at many of the operational communities, led decision makers at 1 Cdn Air Div to seek an impartial method to assist in balancing the production from the UGPT system with the absorption

capacity of the operational units. In the long term, the ability to manage the “pilot system” so as to maintain robustness, despite fluctuations in factors like attrition, became paramount.

Since its inception, PARSim has been used extensively for a variety of investigations. Along with the PTA analysis, the model has been used to test proposals for future fighter force structures (see [1]), in support of staff college research papers (see [2]) and as a means to validate the decisions made at the annual Pilot Production and Absorption Working Group Meetings (see [3]). The results of a study, concerned with the problem of balancing UGPT production with community absorption capacity, can be found in the [4]. Here, the output from a series of PARSim runs with varying production levels and attrition rates, are summarised.

Sommaire

Modelling the Production and Absorption of Pilots : The Development of the Production, Absorption and Retention Simulation (PARSim)

Dr. Norman C. Corbett ; DRDC CORA TR 2013-023 ; R&D pour la défense Canada – CARO ; Mars 2013.

Ce document est un rapport des aspects techniques pertinents au développement de la Simulation de la production, de l'absorption et de la conservation des pilotes (PARSim). Cette simulation a été produite par la Direction d'analyse et de recherche opérationnelle (DARO) à 1re Division aérienne du Canada (1 DAC) sous la direction du Commandant de 1 DAC et de A1 Instruction. Ce rapport est visé aux analystes afin de les guider : leur offrant une base pour maintenir PARSim ainsi que pour développer des simulations semblables.

Pour ces fins, nous offrons une description mathématique des composantes principales de PARSim et nous discutons de quelques propriétés importantes de ces composantes. Notre premier but est de donner à l'analyste une compréhension de PARSim qui est indépendante de son implémentation. En faisant ceci, nous espérons faciliter des développements futurs qui seront indépendants de logiciels spécifiques ou de méthodes de modélisations.

PARSim est un modèle du mouvement des pilotes de la formation officielle de premier cycle (UGPT) vers diverses communautés opérationnelles signifiantes. Une caractéristique importante du module mentor est son approximation de l'action de monter en grade par les pilotes inexpérimentés. Celle-là est accomplie en modifiant dynamiquement l'activité aérienne de pilotes inexpérimentés selon la santé de la communauté et des ressources disponibles. Par exemple, si le niveau d'expérience ou l'activité aérienne à l'unité diminue, le temps que prend un pilote débutant pour monter en grade augmente.

L'usager peut construire et investiguer diverses scénarios en précisant 162 paramètres d'entrée qui fixent les valeurs qui déterminent la production, les taux d'attrition, l'activité aérienne, etc. Des exécutions ultérieures produisent une variété de renseignements de sorties, incluant les niveaux d'expérience, le contingent annuel d'heures de vol (CAHV) exigé par les escadrons opérationnels et des estimations de temps pour diverses équipes de personnel en attente d'instruction (PAI) qui sont distribuées à travers le système-pilote. L'analyse de ces renseignements laisse l'usager estimer l'effet de modifications potentielles encodées dans les paramètres d'entrée.

Une des motivations principales du développement de PARSim a été promulguée par le Chef d'état-major de la Force aérienne (CEMFA). Dans l'année 2000, le CEMFA a ordonné que l'occupation militaire des pilotes doive être régénérée afin de satisfaire au niveau préférentiel de dotation (NPD) d'environ 1 500 pilotes. Malheureusement, cette régénération a été compliquée par l'expiration menaçante de l'indemnité provisoire de pilotes (IPP) en 2003.

En conséquence de l'expiration de l'IPP, une augmentation marquée à l'attrition de pilotes était anticipée. Cette augmentation ainsi que le faible niveau d'expérience à plusieurs des communautés opérationnelles ont encouragé les décideurs à 1 DAC de trouver une méthode non-biaisée qui assistera à équilibrer la production du système UGPT avec la capacité d'absorption des unités opérationnelles. A long terme, l'habileté de gérer le système-pilote afin de sauvegarder la robustesse, malgré les fluctuations de facteurs tels que l'attrition, est de la plus haute importance.

Depuis ses débuts, PARSim a été utilisé extensivement pour une variété d'investigations. Accompagné de l'analyse IPP, le modèle a été utilisé afin de tester certains futurs plans pour la structure des forces de chasse (voir référence [1]), à l'appui des articles de recherche de l'école supérieure de guerre (voir référence [2]), et comme moyen de valider les décisions commises au groupe de travail annuel de la production et de l'absorption de pilotes (voir référence [3]). Les résultats d'une étude visée vers l'équilibre de la production UGPT avec la capacité de l'absorption de la communauté se trouvent dans la référence [4]. Là, les sorties d'une série d'exécutions de PARSim sont résumées, incluant diverses niveaux de production et de taux d'attrition.

Table of contents

Tables

Figures

1 Introduction

This report is a record of the technical aspects pertaining to the development of the Pilot Production Absorption Retention Simulation (PARSim). The simulation was produced at the Operational Research and Analysis Directorate (ORAD) in 1 Canadian Air Division (1 Cdn Air Div) Headquarters and is the result of more than a year of developmental effort by the ORAD team. The project was sponsored by the Commander 1 Cdn Air Div, A4 and A1 Training. This document is, in part, intended to be a guide for analysts: providing a basis for the maintenance of PARSim and the development of other, like simulations.

To this end, we give mathematical descriptions of the major components of PARSim and discuss some of their important properties. In doing so, we intend to provide the analyst with an understanding of PARSim that is independent of its implementation. It is hoped that this approach will facilitate future developments that will not be constrained to specific software packages or modelling methodologies.

PARSim is a model of the movement or “flow” of pilots through formal Undergraduate Pilot Training (UGPT) to the various major operational communities. The model is designed to give decision makers the ability to gauge the effect of factors, like attrition, production and flying rates on the overall “health” of the pilot system. PARSim is essentially comprised of two distinct modules: a *course module*, which is used to build the UGPT portion of the simulation and the *mentor module*, which is a generic model of the major operational communities. An important feature of the mentor module is that it approximates the process of upgrading inexperienced pilots. This is accomplished by dynamically adjusting the rate at which inexperienced pilots fly according to the health of the community and the availability of resources. For example, if the experience levels or the flying rates at the unit decrease, then the time that it take to upgrade ab-initio pilots increases.

The user can construct and investigate a wide variety of scenarios by specifying 162 input parameters: setting values that determine production, attrition rates, flying rates, etc. Subsequent model runs produce a variety of output information, which includes experience levels, yearly flying rate (YFR) requirements for operational squadrons and waiting time estimates for the various personnel awaiting training (PAT) pools that are distributed throughout the pilot system. The analysis of this information allows the user to estimate the effects of the prospective changes that are encoded in the input parameters.

One of the primary motivations for the development of PARSim was promulgated by the Chief of the Air Staff (CAS). In 2000 CAS directed that the pilot military occupation (MOC) was to be regenerated to meet a preferred manning level (PML) of about 1,500 pilots. However, this regeneration was complicated by the looming expiration of the Pilot Terminable Allowance (PTA) program, which was scheduled to occur in 2003.

An extreme increase in the attrition of pilots was anticipated as a result of the PTA expiration. This, and low experience levels at many of the operational communities, led decision makers at 1 Cdn Air Div to seek an impartial method to assist in balancing the production from the UGPT system with the absorption capacity of the operational units. In the long term, the ability to manage the “pilot system” so as to maintain robustness, despite fluctuations in factors like attrition, became paramount. It is important to note, that this is something of a departure from the traditional Operational Research (OR) *modus operandi*, where analysis is performed in response to a request for advice pertaining to a specific set of questions. Here, a request to develop a specific model for continued and regular use accompanied the request for analysis.

Since its inception, PARSim has been used extensively for a variety of investigations. Along with the PTA analysis, the model has been used to test proposals for future fighter force structures (see [1]), in support of staff college research papers (see [2]) and as a means to validate the decisions made at the annual Pilot Production and Absorption Working Group Meetings (see [3]).

The results of a study, concerned with the problem of balancing UGPT production with community absorption capacity, can be found in [4]. Here, the output from a series of PARSim runs with varying production levels and attrition rates, are summarised. A supplementary analysis of this output was used to empirically postulate a number of relationships linking production, flying rates, and experience levels for the major operational communities. We shall examine relationships of this type in Section 4.

With about 13,800 pilot billets, the number of pilots in the United States Air Force (USAF) is roughly ten times the the number of pilots in the Canadian Air Force (CAF). Oddly enough, the USAF faces similar issues with regard to the pilot MOC: a shortage of pilots caused by high attrition to major airlines, low experience levels at the operational units and an inability to absorb the inexperienced pilots needed to address attrition. These issues motivated two RAND studies, documented in [5] and [6], both of which concentrated on the fighter communities.

The focus of the study, in [5], is the identification of the factors that caused, as well as the resources that could alleviate the pilot crisis. Although no explicit relationships are presented, the work presented in [5] illustrates the relationship between experience levels and the monthly flying rates of experienced and inexperienced pilots. In particular, this work shows that, as experience levels decline, experienced pilots need to fly more in order to supervise growing numbers of inexperienced pilots. On the other hand, inexperienced pilots fly less as they must compete for fewer and fewer training opportunities.

This is significant as a reduction in flying rates, for inexperienced pilots, translates into an *ageing rate deficit*. An ageing rate deficit implies that the time T_e (in months) it takes for inexperienced pilots to become experienced will be significantly longer when experience levels on the operational squadrons are low. If left unchecked, the ageing rate deficit gives rise to what the authors call the “slippery slope”. This phenomena is a progressive degradation of unit experience levels and occurs when operational squadrons are forced to absorb inexperienced pilots in excess of their capacity to do so.

Absorption capacity is the focus of the second RAND report. Here the authors offer a definition of absorption capacity and identify a number of parameters that influence it. The definition of absorption capacity, that we shall adopt, is essentially the same as that given in [6]. In particular, the *absorption capacity* $\mathscr{A}$ of a flying community is defined to be *the maximum yearly rate at which inexperienced pilots can be posted to the operational squadrons without having catastrophic negative effects on experience levels*.

In [6], the authors also present a number of steady state formulae relating the parameters that influence $\mathscr{A}$. In particular, if I represents the total number of inexperienced pilots, then the steady state absorption capacity is given by

$$\mathscr{A} = 12\frac{I}{T_e}\,. \tag{1}$$

Equation (1) tells us that the absorption capacity is the same as the yearly rate at which inexperienced pilots upgrade.

Unfortunately, the formulae given in [6] are heuristically derived and the relationship between the parameters influencing $\mathscr{A}$ and unit experience levels are never made explicit. For instance, it is demonstrated that T_e is inversely proportional to the ageing rate, which is essentially the rate at which inexperienced pilots are able to fly. However, an explicit expression relating the flying rate of inexperienced pilots to unit experience levels is not given. Our presentation of the mathematical model governing the mentor module allows us to clarify and extend the RAND approach to the estimation of a steady state absorption capacity.

Within the Operational Research division, previous efforts to model the movement of personnel through the pilot system have established a precedent of using a System Dynamics (SD) approach to these problems: or at least the associated software. The field of SD was started by Jay Forester in 1961 (see [7]). The intent was to use techniques from the fields of feedback control theory and dynamical systems in order to model and analyze problems involving complex socio-economic systems. The SD approach culminates in the development of a time-stepped computer simulation that can be used to track the evolution of dependent variables of interest. This simulation is subsequently used to investigate the real-world system under consideration.

In [8], an SD simulation was used to track the movement of pilots from one Year of Service (YOS) to the next. The objective was to assess the impact of attrition and potential changes to the restricted release policy. The modelling efforts, documented in [9], [10] focused on the UGPT system. These simulations were used to address a number of questions pertaining to PAT pools, the effectiveness of changes to the structure of the UGPT system and the prediction of future pilot intake. The UGPT portion of PARSim is based on the simulation SimPAT (see [11]) and modelling picture was completed by producing a simulation of the major operational communities (i.e. the mentor module).

1.1 System Dynamics

Since the time of its development, the SD method has been applied to problems arising in many areas, including the fields of operational research and management science (see [12] and [13]). As we have mentioned, the SD approach typically culminates in the development of a time-stepped computer simulation. This simulation can be used to track the evolution of the state or dependent variables that comprise the system under study. The state variables are generally referred to as *levels* or *stocks*. One typical interpretation ascribed to SD simulations is that of fluid flowing through a connected network of pipes, valves and reservoirs. The reader is referred to Figure 1.

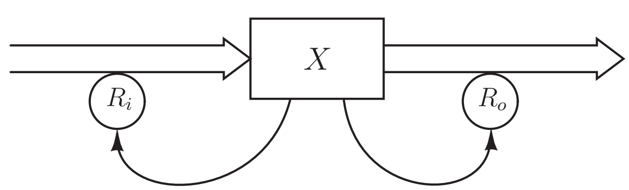

Figure 1: *Simple SD Diagram*

The level X can be thought as a reservoir. However, the amount of fluid in X can be negative. The input and output rates R_i and R_o can be likened to valves that control the flow of fluid in and out of the reservoir. In general, R_i and R_o can depend on time t and X (as well as any other state variables). In Figure 1, the dependence of the rates on X is indicated by the directed arcs emanating from X. The evolution of X is defined by the equation

$$\dot{X}(t) = R_i(X,t) - R_o(X,t)\,,$$

where the $\dot{X}$ denotes the derivative of X with respect to time t. That is, the time rate of change of X is given by the net flow or the difference between the *inflow* R_i and *outflow* R_o. The reader familiar with *compartmental analysis* will note the similarities (see [14]).

For each level in an SD simulation, an initial condition (IC) needs to be specified. The IC gives the value of the corresponding state variable at some initial time t_0, which is usually assumed to be zero. In the present case, an appropriate IC is $X(0) = x_0$, for some constant x_0. The ICs are generally not visible on an SD diagram as they are contained within the levels.

In general, SD simulations are based on coupled systems of first order ordinary differential equations (ODEs). Model development is centred on a diagrammatic representation of the system. For example, the evolution of a damped simple harmonic oscillator is described by the system of first order ODEs

$$\dot{x}(t) = y(t)\,, \tag{2a}$$

$$\dot{y}(t) = \frac{1}{m}\big(f(t) - kx(t) - \beta y(t)\big)\,, \tag{2b}$$

where x corresponds to position and y corresponds to speed. The mass, damping coefficient and spring constant are denoted by m, β and k respectively. An SD diagram that corresponds to System (2), is depicted in Figure 2. Here the constants m, β and k, as well as an exogenous[1] force $f(t)$ are represented explicitly on the diagram; however, this is not necessary.

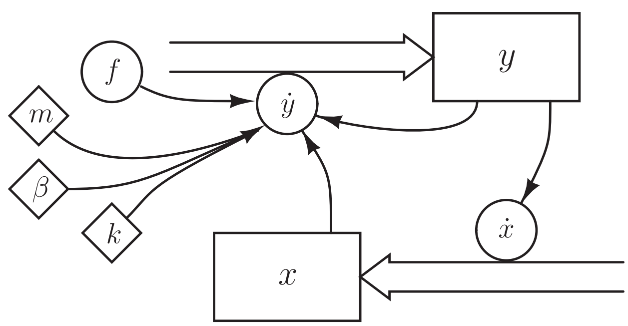

Figure 2: *SD Diagram of a Simple Harmonic Oscillator*

Typically, SD simulations are implemented in purpose-built software. For instance, the PARSim model was developed using Powersim[2]. These purpose-built software packages have a number of advantages. The user interface of these packages facilitates a graphical approach to model development by allowing the developer to “drag and drop” various predefined components onto a worksheet. As a result, the shell of the system can be quickly roughed in and the required levels and rates (auxiliary variables, constants, etc.) can subsequently be defined one at a time. As well, the resulting SD diagram can be used to give an overview of the model without the need to discuss complicated details.

[1] An influence is exogenous if it is external to the system. That is, it does not depend on any state variables.

[2] http://www.powersim.com/

When a user runs an SD simulation, the underlying system of ODEs is solved numerically. The user can select from a number of predefined numerical algorithms and this alleviates the need for the developer to manually code numerical algorithms. For instance, in Powersim, the user can choose from Euler's Method as well as second, third and fourth order Runge–Kutta Methods (see [15]).

One of the concepts which permeates the practice of SD is that of *feedback*. In SD, feedback (see [16]) refers to the case where the value of a level x has a delayed influence upon itself through a chain of cause and effect. These chains of cause and effect form, what are referred to as, *feedback loops*. Feedback loops are considered to be the fundamental building blocks of SD simulations and considerable effort has been devoted to understanding loop behaviour (see [17] and [18]). The reader is referred to Figure 3 which depicts an SD diagram of a third order feedback loop. Here the value of level x influences[3] the value of level y, which influences the value of level z. In turn, z influences the value of x. This loop corresponds to the causal chain $x \to y \to z \to x$.

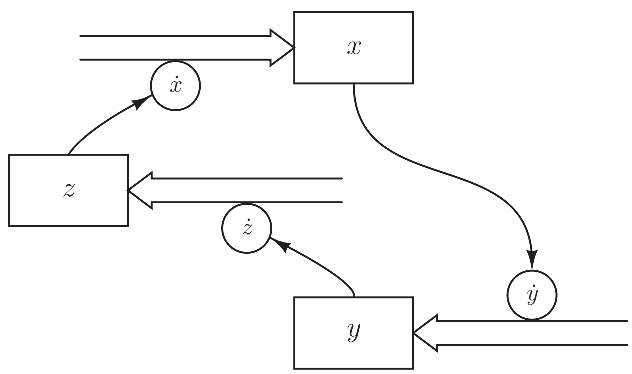

Figure 3: *Third Order Feedback Loop*

Since SD simulations are based on systems of first order ODEs, the notion of feedback espoused by the SD community is in fact the *coupling* of the equations in the underlying system; the time rate of change of one dependent variable is influenced by the value of another. For example, in the linear autonomous case, the system underlying the diagram in Figure 3 is

$$\dot{x} = a_1 z, \quad \dot{y} = a_2 x, \quad \dot{z} = a_3 y, \tag{3}$$

where a_1, a_2 and a_3 are real constants. Furthermore, unless it is explicitly modelled, the delay in the influence of a level upon itself is a product of the numerical algorithm used to solve the system of ODEs.

The SD concept of feedback is different from the one found in the field of feedback control (see [19]). Here, the term feedback refers to a situation where filtered system outputs are added to system inputs for the purpose of control. That is, the output signal of a system or *plant* that one wishes to control is first filtered by routing it through a *controller* (see Figure 4). The filtered signal is then *fed back* by adding it to the plant input signal. The controller is designed with the aim of eliciting desirable plant behaviour. The plant and the controller may or may not be described by systems of differential equations. When it is possible to model the plant with a system of ODEs, then it is comparable to an SD system. In this case, the plant input can be

[3]Each level is affected through its rate

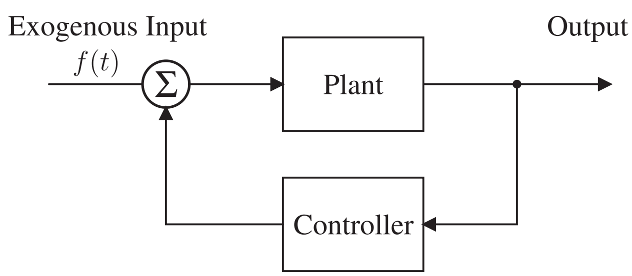

Figure 4: *Feedback Control Loop*

likened to an exogenous force, while the plant outputs are the values taken by the levels. However, without the controller to close the loop, no feedback occurs.

There are a few disadvantages to the SD modelling approach. Let us assume that our objective is to produce an accurate computational model of an observable system. We refer the reader to Figure 5. A traditional approach to model development produces, through the process of abstraction, a formal description of the model. At this point, the formal description is analyzed and, if required, revised. If an explicit solution is known, then it can be validated against experimental observations and subsequently used to make inferences about the system in question. However, explicit solutions are usually not known and the next step is to derive a computable description of the formal model. In the case of a system of ODEs, one generally uses a numerical method to produce a set of discrete equations that can be implemented on a computer. Once again, there is an opportunity for analysis and, if required, refinement of the computational and/or formal models.

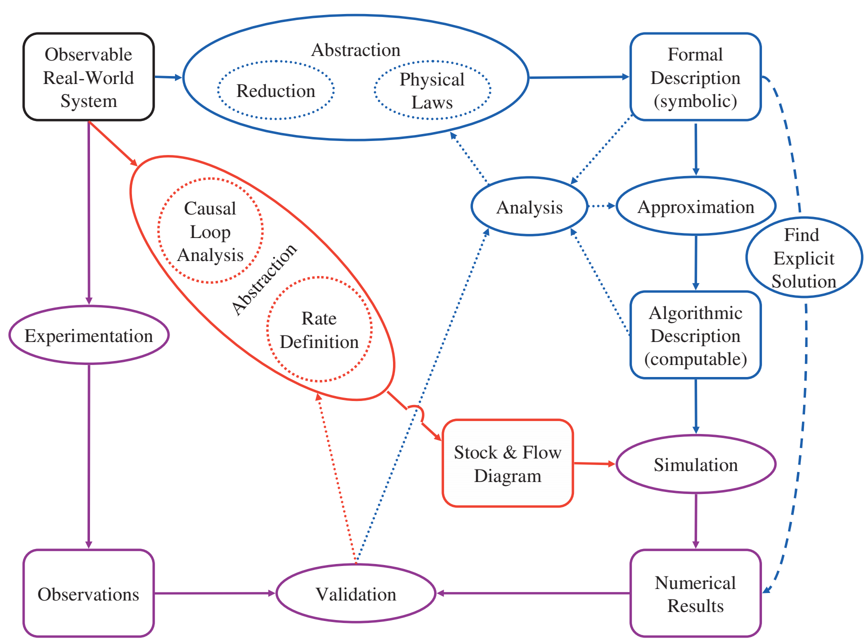

Figure 5: *Competing Modelling Processes*

When one takes an SD approach, the rates are defined piecemeal. This approach and a focus on the SD diagram means that the analyst does not usually work with a complete formal model. Although a complete set of software specific equations can be viewed, these are not amenable to analysis without translation into standard mathematical form. The effect is that the abstraction process tends to be isolated from the analysis that is typically part of traditional modelling. Also, the discrete equations forming the computational model are completely hidden and the analyst cannot easily implement an algorithm of his or her own. Consequently, the SD approach tends to jump from model development directly to simulation.

As observed in [20], bypassing an analysis of the formal model and proceeding straight to the numerics can lead to fallacious conclusions. In this paper, the author provides an example of a system of ODEs for which a Runge–Kutta algorithm gives incorrect results. In particular, underlying instability in the numerical algorithm can be confused with the onset of chaotic behaviour in a coupled nonlinear oscillator. The author then demonstrates how a relatively straightforward analysis of the formal model can uncover the errors.

In the case of observable systems, validation can uncover fallacious conclusions. However, if our sole intent is to investigate a system in a regime for which we have observational data, then why build a simulation at all? In these instances a direct interpolation of the data seems more appropriate. A primary use for a simulation is the prediction of system behaviour in regimes for which little or no data exits. Consequently complete model validation is impossible and analysis is the only persistent means by which to assess model confidence.

Lastly, we note that other modelling and simulation methodologies may be more appropriate for the pilot problem. Given the small size of the pilot MOC a discrete event approach would be a better. However, time constraints and the availability of the extant SD simulation SimPAT made the SD paradigm a compelling alternative.

The remainder of this report is organized as follows: In Section 3, we provide a mathematical description of the course module and, in Section 4, we give a detailed exposition on the mentor module. Our conclusions and recommendations for further research follow in Section 5

This page intentionally left blank

2 Model Overview

As we have mentioned, PARSim was implemented in Powersim software in order to make use of the work done on the UGPT system in [11]. Input and output are accomplished via "links" to a number of Excel workbooks. There is one workbook for the UGPT system and one for each of the operational communities that are represented in PARSim. This has the advantage of enabling a graphical presentation of the results, which can be subsequently stored with the corresponding inputs. The reader is referred to Figure 6, which depicts a typical input sheet for the mentor module. The user simply changes the value in the cell corresponding to the parameter to be modified. It should be pointed out that certain Powersim objects, like ranges, cannot be modified from the input sheet. This means that, if the dimension of any array-like object is to be changed, then the requisite modifications must be made within the simulation.

Initial Information			Other Information	
Community	E_0	44	fully manned	44
	NOP_0	61	Exp $Level_0$	64.7
	PAT OTU_0	12	T_E nominal (yrs)	2.1
	% lst/yr E	7.0	T_E nominal (wks)	109.6
	% lst/yr NOP	7.0		
	hrs/mo E	17.0	hrs/yr E	204.0
	hrs/mo I (lim)	17.0	hrs/yr I (lim)	204.0
	hrs req'd	430.0		
	SR	1.0		
	flying pos'n	68		
	min NOP	55		
	CF188 YFR	TRUE		
	UGPT IP/yr	0.0		
OTU	crs duration (wks)	16.0		
	number of crs	2		
	max loading	6		
	min loading	0		
	capacity	12		

Figure 6: *Example Input Sheet*

The use of Excel also allows many end-of-run calculations to be done in the spreadsheets rather than in the simulation. For example, estimates of the waiting times, for the PAT pools in the simulation, are calculated in Excel at the end of each simulation run. Other calculations that would normally be done in the simulation can also be offloaded to Excel. For example, the net instructor pilot (IP) requirements for the UGPT system are calculated in Excel from UGPT throughput or production values stored in the same workbook. Although this approach simplifies model development and facilitates faster simulation runtime, it does make the investigation of scenarios with highly variable production difficult.

PARSim contains a total of 650 "variables". These variables consist of levels, auxiliary variables and constants. The levels are the dependent or state variables of the underlying system of ODEs. The auxiliary variables are intermediate variables whose values depend on a number of levels, constants or other auxiliary variables. The constants represent numerical values that are static and not typically changed during a simulation run. One hundred sixty two of the 172 constants accept input from the various Excel workbooks. The output from 140 levels and auxiliary variables is collected and sent to Excel in the form of time series at the end of a run.

In particular, for a given level X, the numbers $X(k\Delta t)$, $k = 0, 1, 2, \ldots M$ are exported. Here Δt is the simulation time step and M is the number of values to be exported. If T is the duration of a run, then $M = \text{floor}(T/\Delta t)$. Since, Δt is usually set at seven days and the length of a simulation run is typically set at twenty years the typical time series for a level has 1,043 values. The exported auxiliary variables are usually sampled at a lower rate (e.g. yearly). Data from the first four or five years of a run is treated as "warm up" and discarded. Ordinarily, the transient behaviour of a system is significant. This behaviour allows one to estimate the time it takes the system to *relax* from an initial configuration or an exogenous perturbation to either a periodic (neutral) or asymptotic steady state[4].

In the context of the pilot system, the *relaxation time* represents the delay between cause and effect. For instance: How long does it take a sudden decrease in monthly flying rates to affect the experience levels in an operational community? However, in the case of PARSim, the relaxation times are influenced by the way in which attrition is modelled and the presence of state dependent nonlinear delays in the mentor module. The latter makes accurate initialisation of the simulation difficult and hence, given some initial configuration, the time it takes the simulation to reach steady state may not be realistic.

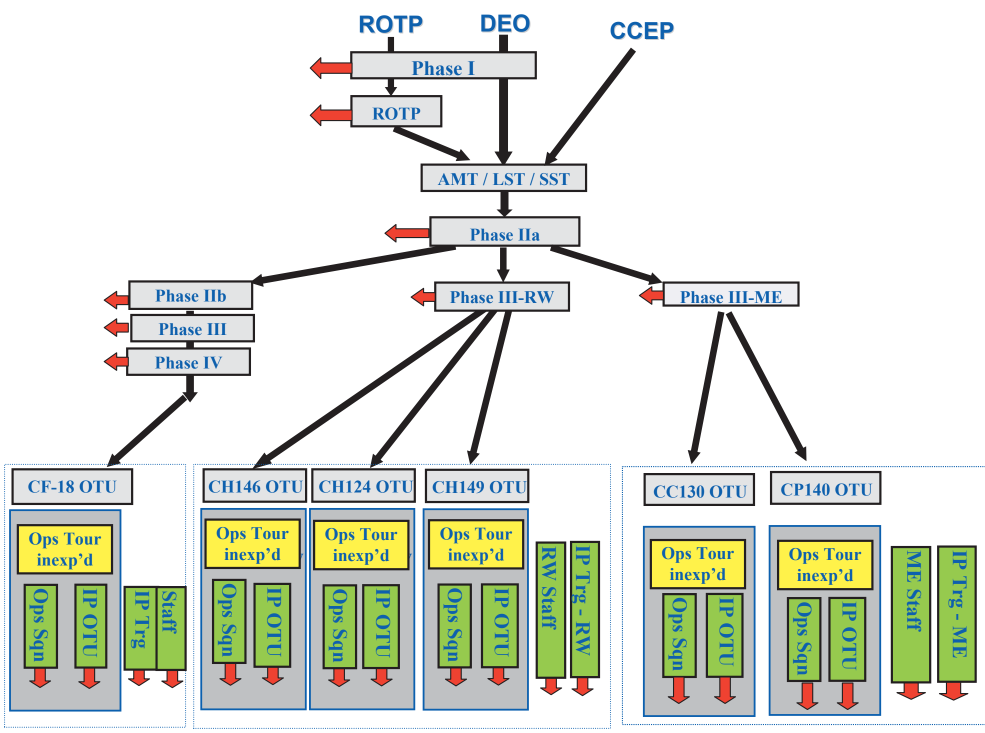

Figure 7: *Simplified Flow Diagram of the Pilot System*

The flow mapping of the UGPT system is a simplified version of the one used in Sim-PAT (see [11]). We

[4]assuming that one exists

refer the reader to Figure 7. The UGPT portion of PARSim accounts for three distinct entry plans. These are Direct Entry Officer (DEO), Royal Officer Training Program (ROTP) and the Community College Entry Plan (CCEP). Second Language Training (SLT) and the Basic Officer Training Course (BOTC) are not explicitly accounted for and, as in SimPat, Land Survival Training (LST), Sea Survival Training (SST) and Aeromedical Training (AMT) are represented with a single level. DEO and ROTP students take Phase I flying training, while CCEP students proceed directly to Phase IIa. After Phase IIa, students are assigned their respective aircraft types: jet, multi-engine or rotary wing. Jet students continue their training by taking Phase IIb, Phase III and Phase IV, while multi-engine and rotary wing students proceed to Phase III-ME and Phase III-RW respectively. After completing formal flying training the students are sent to their operational communities where they take their type course at the respective Operational Training Units (OTU).

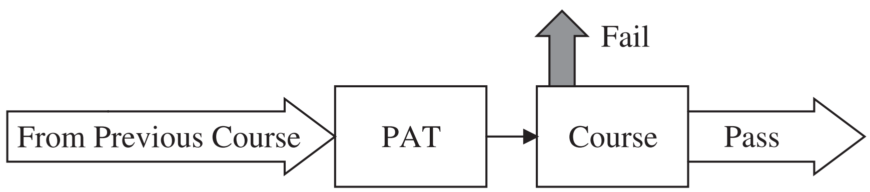

Figure 8: *Flow Diagram of the Course Module*

As in Sim-PAT, the UGPT portion of PARSim is built by connecting together a number of course modules: one for each course represented. The reader is referred to Figure 8, which depicts the major flows in the course module. Students first enter the PAT pool preceding the course and, when space becomes available, are loaded onto the course. When the course is complete, a certain number of students fail the course and leave the pilot system. The remainder flow on to the next portion of the simulation. The biggest difference between Sim-PAT and the UGPT portion of PARSim is the way in which course failure (or attrition) is modelled.

In Sim-PAT, the failure of students from a particular course is treated stochastically and, depending on the expected course loading, is modelled in one of two ways. The number of failures for courses with "large" expected loads, is calculated by assuming that the proportion of students that fail follows a normal distribution with a given mean and variance. On the other hand, the number of failures for courses with smaller expected loads is calculated according to a hard coded discrete probability distribution. Both of these approaches have certain drawbacks, which we now examine. In the case of the normal distribution we observe that there is an implicit assumption of uniformity across students. That is, the probability that any individual student fails the course does not vary over all the students loaded on the course.

In such cases, a natural choice of probability model for failures is a binomial distribution and, in certain circumstances the normal distribution can be used as an approximation. According to [21], the binomial distribution $B(n,p)$ is well approximated by the normal distribution $N(np,\sqrt{np[1-p]})$ as long as $\min(np,n[1-p]) > 15$. Here n represents the number of students loaded on the course, while p is the expected proportion of students that will fail the course. When we consider typical course loadings and failure rates (see Table 1), we see that the normal approximation of the binomial distribution is not appropriate for any of the UGPT courses.

Table 1: *Typical Course Loads and Failure Rates*

Course	Loading	Failure Rate (%)	np	$n[1-p]$
Phase I	32	12	3.8	28.2
Phase II	14	16	2.2	11.8
Phase IIb & III	22	4	0.9	21.1
Phase III RW	55	9	5	50.1
Phase III ME	25	8	2	23
Phase IV	13	2	0.26	12.7

The hard-coded discrete distribution does avoid the problems associated with small course loadings and/or small failure rates. However, the implementation of this distribution in Sim-PAT is based on the assumption that the number of students loaded on a given course offering does not vary. As the number of students loaded on an offering of a given course does vary, this approach will introduce unintended bias into the failure calculations.

In PARSim, the user can choose from either a rounded proportion or a binomial distribution to model the number of failures for each course represented in the UGPT portion of the simulation. Since Powersim software does not have a built-in binomial distribution, this distribution is implemented by defining a "large" vector of Bernoulli random variables. When a course is loaded, the appropriate number of entries in the vector are summed to obtain a sample from the desired binomial distribution (see [22]).

For each UGPT course, represented in PARSim, the user must specify:

a. The number of courses per year along with the course duration.

b. The maximum and minimum number of students that can be loaded on a course.

c. The average failure rate for the course and whether or not to use the binomial distribution for failure calculations.

At the start of a course, if there are fewer than the minimum number of students in the PAT pool, then the course is "zero-loaded" and students must wait until the next course. If the number of students in the PAT pool is between the stated minimum and maximum, then the PAT pool is emptied when the course is loaded. When there are more than the maximum number of students in the PAT pool, then the maximum number is loaded and the remaining students must wait until the next course offering.

In addition, the user must specify a number of *breakouts* which determine the proportion of students sent from:

d. Phase II Basic Flying Training (BFT) to the Advanced Flying Training (AFT) courses Phase III multi-engine, Phase III rotary wing and the follow-on jet course Phase IIb.

e. The advanced courses to the various PAT pool preceding the Operational Training Units (OTU) in the major operational communities: CF-18, CC-130, CP-140, CH-124, CH-146 and CH-149.

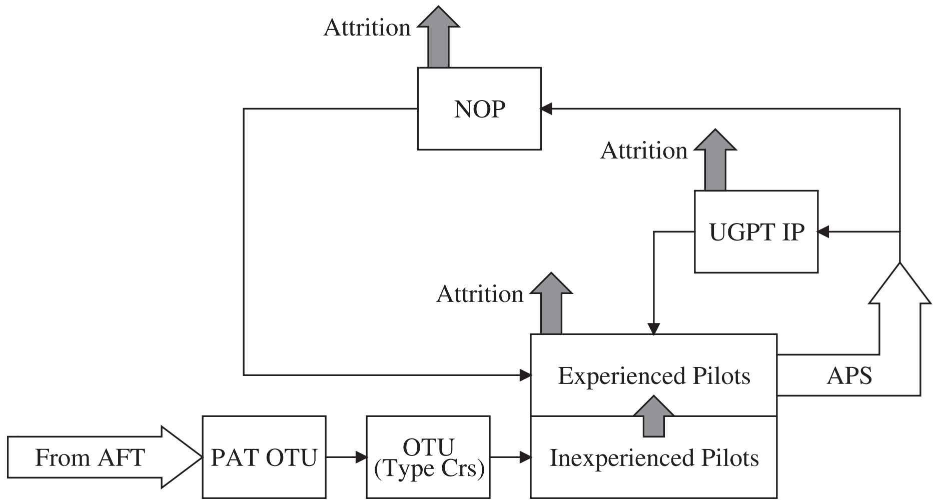

Figure 9: *Flow Diagram of the Mentor Module*

The primary flows in the mentor module are depicted in Figure 9. Newly winged pilots move into the PAT pool before the OTU, where they wait until they are loaded on the next available type course. The PAT/OTU portion of the mentor module is similar to the course module. However, as type course failures are rare, the corresponding flow is not included. After pilots complete their type course, they move onto the operational squadrons as inexperienced pilots. Here, the pilots undergo a period of "on the job training" (OJT) after which they upgrade and become experienced pilots. At this point they are available to supervise the OJT period for inexperienced pilots. It is important to note that during the OJT period, an inexperienced pilot must be supervised by an experienced pilot whenever he or she flies. This is the so-called "mentoring process".

At each annual posting season (APS), pilots can move from the operational squadrons to "non-operational" (NOP) positions and vice-versa. NOP positions include all OTU IP positions as well as staff positions at wing and higher headquarters. This flow of experienced pilots is varied in sign and magnitude so that the manning of available flying positions is tracked as closely as is possible and so that a user specified critical minimum number of NOP positions are manned. In particular if, during APS, there are more pilots in the operational community than there are flying positions, then experienced pilots are sent to NOP positions in order to free up the required number of flying positions. On the other hand, if there are unmanned flying positions, then these are filled by moving experienced pilots from NOP positions back to the operational community.

However, if the number of manned NOP positions is below the critical number, then experienced pilots are sent to fill theses regardless of the manning of flying positions in the operational community. We should point out that the most recent version of the mentor module includes a "retread flow". Specifically pilots, returning to the operational community from staff positions, are first sent to a refresher course at the OTU and then undergo an OJT period at an operational squadron before becoming experienced again. In what follows, we shall restrict our attention to the original version of the mentor module. Experienced operational pilots and those in NOP positions can be released and leave the air force. This is indicated by the attrition flows in Figure 9. Inexperienced pilots are assumed to be on restricted release and hence there is no attrition flow for these pilots.

Each of the major operational communities is represented by a copy of the mentor module. The smaller operational communities, which include the CC-115 Buffalo and the CT-142 Dash 8 and the "follow-on" communities, including the CC-150 Polaris, are not explicitly represented in PARSim. However, the net flow of pilots into these various communities is accounted for.

For each major operational community, the user must specify:

a. Attrition rates for experienced pilots and those occupying NOP positions.

b. The monthly flying rate for experienced pilots along with an upper limit on the monthly flying rate for inexperienced pilots.

c. The total number of hours that inexperienced pilots must fly on the operational squadron before they can upgrade.

d. A sortie ratio. Essentially, this parameter quantifies the proportion of proficiency missions in which an experienced pilot is accompanied by more than one inexperienced pilot.

e. The total number of flying positions available for pilots on the operational squadrons as well as the minimum number of OTU IP and staff positions that must be manned.

Although the process of upgrading pilots is performance based, in the interest of simplicity, the mentor module employs an "hours on type" marker. That is, inexperienced pilots are deemed to be experienced after flying a specified number of hours at the operational squadron. As for the course module, representing the OTU, the user must also specify:

f. The number of type courses per year along with the course duration.

g. The maximum and minimum number of students that can be loaded on a type course.

2.1 Example Runs

As a hypothetical example, let us consider two runs of the mentor module for a hypothetical operational community. In doing so, we investigate the consequences of a reduction in YFR on the ability of the community to "absorb" inexperienced pilots. Suppose that the YFR reduction results in a change from 204 to 174 hours in the yearly per pilot flying and consider whether or not 12 *ab initio* pilots per year exceeds the absorption capacity of the community. Recall that absorption capacity $\mathscr{A}$, is the rate at which inexperienced pilots can be posted to the operational squadrons without having undue negative effects on experience levels. The inputs for these two runs are given in Table 2.

In Table 2, A_e and A_s are the attrition rates for experienced pilots and pilots in NOP positions respectively. The attrition rates are simply the proportion of pilots that are expected to leave each year in the absence of all but attrition flow. The symbol R_e denotes the flying rate of experienced pilots in hours per month per pilot (hrs/mo/plt), while $R_i^{\max}$ is the limit on the flying rate of inexperience pilots. We investigate the cases where $R_e = R_i^{\max} = 204/12 = 17$ hrs/mo/plt and $R_e = R_i^{\max} = 174/12 = 14.5$ hrs/mo/plt. It is important to

Table 2: *Example Inputs*

A_e	7.0 %
A_s	7.0 %
R_e	17/14.5 hrs/mo/plt
$R_i^{\max}$	17/14.5 hrs/mo/plt
H_s	430 hrs
N_f	68 plts
ν	1
$S_{\min}$	55 plts
n_u	0 plts
M	2 crs
D	2 wks
$L_{\max}$	6 plts
$L_{\min}$	0 plts

note that the actual flying rate of the inexperienced pilots R_i, will not necessarily be equal to $R_i^{\max}$. In particular, if unit experience levels are low, then R_i will be smaller than $R_i^{\max}$ (i.e. $R_i < R_i^{\max}$). On the other hand, when the experience levels are high, $R_i = R_i^{\max}$, but R_i never exceeds $R_i^{\max}$.

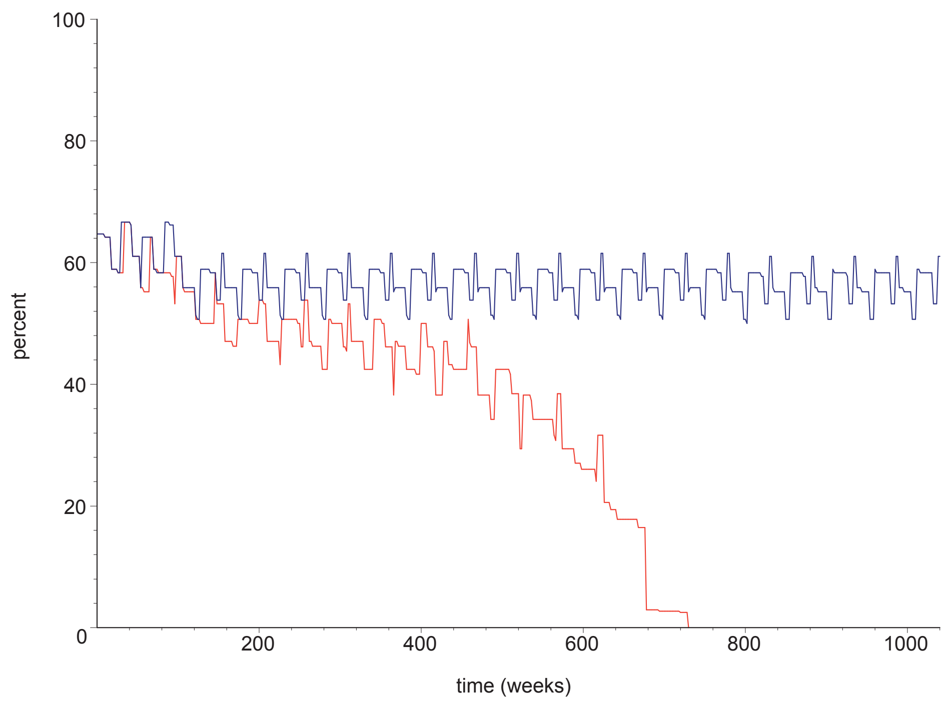

Figure 10: *Experience Levels*

The number of hours that inexperienced pilots must fly before they upgrade is denoted by H_s. The total number of flying positions on the operational squadrons is given by N_f and $S_{\min}$ is the minimum number of

NOP positions that must be filled. The so-called "sortie ratio" is represented by ν. In this case, we have set $\nu = 1$, which means that, on every proficiency sortie, one experienced pilot accompanies one inexperienced pilot. The symbol n_u is the net number of experienced pilots that are sent back into the UGPT system to act as IPs each year.

The duration of the type course at the OTU is D, while $L_{\max}$ and $L_{\min}$ denote the corresponding maximum and minimum course loadings. The number of type courses per year is M. In the current example, we have ensured that there are enough pilots in the OTU PAT pool to guarantee that all type courses are run at maximum capacity. This means that two groups of six *ab initio* pilots are sent to the operational squadrons each year.

The experience levels resulting from the inputs at Table 2, are plotted in Figure 10. The blue curve corresponds to the case $R_e = R_i^{\max} = 17$ hrs/mo/plt and the red curve corresponds to the case $R_e = R_i^{\max} = 14.5$ hrs/mo/plt. The experience level is simply the proportion of pilots on the operational squadrons that have upgraded. When we consider the blue curve, the "warm up" period is evident as the experience level gradually settles over a period of four years. From the fifth year on, the average[5] experience level is constant at approximately 56 percent. Consequently, when pilots can fly 204 hours per year, we expect that the operational squadrons should be able to sustain an absorption rate of 12 inexperienced pilots per year while maintaining an average experience level of 56 percent.

On the other hand, when $R_e = R_i^{\max} = 14.5$ hrs/mo/plt, the experienced pilots are not flying at a high enough rate to absorb the influx of inexperienced pilots. As a result, the average experience level steadily declines until the simulation stops executing after about 14 years. This happens when there are no experienced pilots left on the operational squadrons. It is important to note that this does not imply that the hypothetical flying community in our example would, in reality, fail or that the failure would happen within 14 years of beginning to fly at the reduced rate. In reality, decision makers would intervene long before such an event could occur. However, this example does illustrate that the community cannot sustain an absorption rate of 12 pilots per year when pilots fly only 174 hours per year.

One of the consequences of decaying experience levels is illustrated in Figure 11. Here the time T_e, required for inexperienced pilots to upgrade is depicted for the case $R_e = R_i^{\max} = 14.5$ hrs/mo/plt. We observe that, as the experience level decays, the time that it takes an inexperienced pilot to upgrade increases from 2.5 years to a maximum of 4.1 years just before the simulation stops executing. This is a consequence of the fact that the inexperienced pilots are not flying at the maximum rate $R_i^{\max}$. Indeed, if this were the case, then the time to upgrade would be stable at 2.5 years. According to Equation (1), the maximum absorption rate is inversely proportional to T_e.

Consequently, when a community is forced to absorb pilots at a rate which exceeds $\mathscr{A}$, the upgrade time T_e increases and this leads to a reduction in $\mathscr{A}$. As a result, T_e increases further, which leads to another reduction in $\mathscr{A}$ and so on. This is a version of the so-called "slippery slope" phenomenon discussed in [5].

We return our attention to the case $R_e = R_i^{\max} = 17$ hrs/mo/plt. Since the number of flying positions N_f is fixed, when the experience level is, on average, constant, so are the numbers of inexperienced and experienced pilots on the operational squadrons. We refer the reader to Figure 12. Here, the blue curve represents the number of experienced pilots on the operational squadrons as a function of time. The

[5] We are referring to a moving time average with a period of 52 weeks

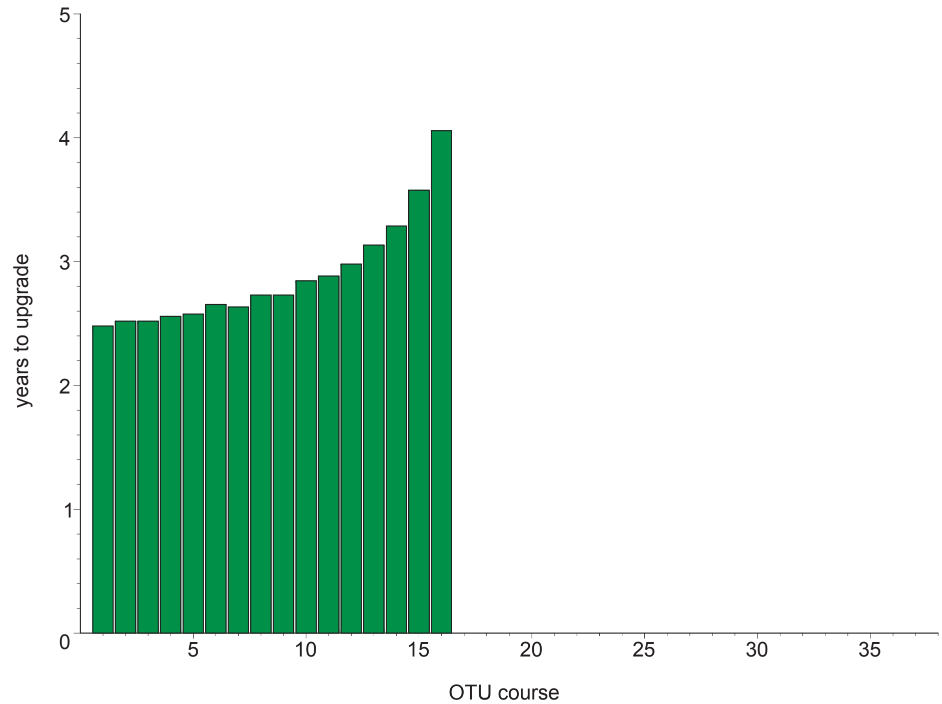

Figure 11: *Upgrade Times*

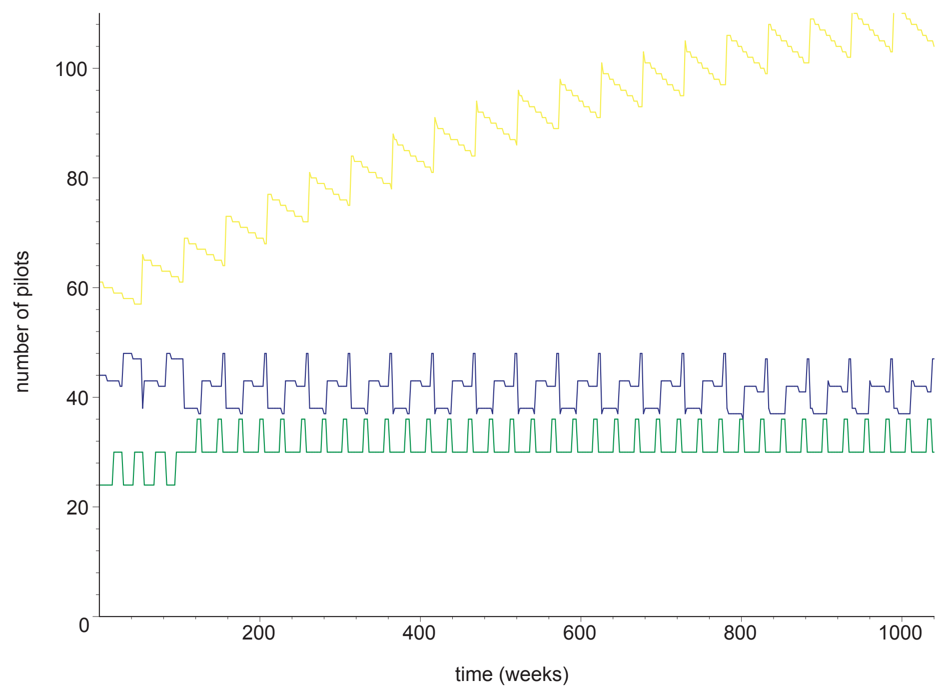

Figure 12: *Total Pilots in the Operational Community*

number of inexperienced pilots is represented by the green curve and those in NOP positions by the yellow curve. We observe that the numbers of experienced pilots is, on average, 40, while the average number of inexperienced pilots is 31. In addition, the number of pilots in NOP positions is growing. This indicates that the community is healthy enough to meet its' minimum OTU IP and staff position commitments and would likely remain healthy in the face of transient fluctuations in the attrition rates A_e and A_n. For example, since more than the minimum number of NOP positions are filled, a temporary increase in A_e would simply result in a posting of surplus pilots from NOP positions back to the operational squadrons.

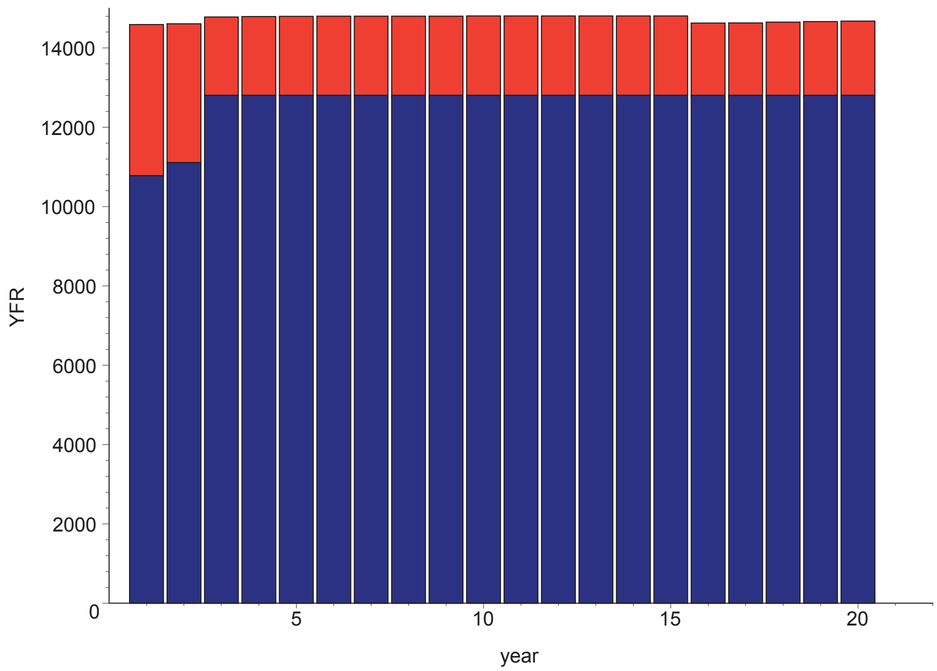

Figure 13: *YFR Requirements*

We now refer the reader to Figure 13. This figure depicts the YFR requirements for the operational squadrons in the case $R_e = R_i^{\max} = 17$ hrs/mo/plt. The YFR required for a given year is depicted with a vertical bar. The blue portion of each bar indicates the number of hours that must have force generation value (upgrading pilots). That is, these hours must count toward the accumulation of the total hours on type for inexperienced pilots. The red portion of the bar indicates hours that can be used for any purpose (e.g. Air Sovereignty Alert, Currency, etc.). We see that to maintain the flying rate of 204 hours per year for each pilot, the operational squadrons must be able to fly an average of 14,742 hours per year. Moreover, an average of 12,805 flying hours must be devoted to force generation in order to absorb 12 pilots per year and maintain an average experience level of 56 percent.

3 The Course Module and Waiting Times

In this section we present the equations which govern the behaviour of the course module and subsequently derive a formula that can be used to estimate the waiting time for students in a PAT pool. The interested reader will find the Powersim equations for the course module in Annex A. For a course of type $\mathscr{C}$ in the UGPT system, we shall use the following notation and assumptions:

a. Let $C(t)$ be the *total number of students loaded on course* $\mathscr{C}$ at time t, where t is the time in months.

b. Let M be the *number of courses* of type $\mathscr{C}$ offered per year. We assume that the courses are equally spaced in time. Denote by D, the *course duration* of $\mathscr{C}$ in months. Accordingly, the k^{th} offering or *serial* of $\mathscr{C}$ begins at time $t_s[k] = 12k/M$ and finishes at time $t_f[k] = 12k/M + D$, where $k \in \mathbb{N}_0 = \{0, 1, 2, \ldots\}$. The number of courses of type $\mathscr{C}$ that are ongoing concurrently is

$$M_{\text{con}} = \text{ceil}\left(\frac{MD}{12}\right),$$

where $\text{ceil}(x)$ is the smallest integer n satisfying $n \geq x$. Moreover, if $D > 12/M$, then $M_{\text{con}} \geq 2$.

c. We use $P(t)$ to represent the *total number of students in the PAT pool* preceding $\mathscr{C}$ at time t. That is, $P(t)$ is the total number of students waiting to be loaded on the course. Also, let $U_{\text{net}}(t)$ denote the total outflow of all courses that feed the PAT pool (i.e., courses that are directly "upstream") of $\mathscr{C}$.

d. Let $c[k]$ be the *number of students loaded on serial k*. If $\Delta t > 0$ is a "small" change in time (we can identify Δt with the simulation time-step), then $c[k]$ is given by the equation

$$c[k] = \begin{cases} 0 & P(12k/M - \Delta t) < L_{\min} \\ P(12k/M - \Delta t) & L_{\min} \leq P(12k/M - \Delta t) \leq L_{\max} \\ L_{\max} & P(12k/M - \Delta t) > L_{\max} \end{cases}, \tag{4}$$

where the numbers $L_{\max}$ and $L_{\min}$ are the *maximum and minimum loadings* for $\mathscr{C}$ respectively.

e. The *expected proportion of students failing serial k* of $\mathscr{C}$ is denoted q. The random variable $F[k]$ represents the *number of students that actually fail serial k*. We assume that $F[k] \sim B(c[k], q)$. That is, $F[k]$ is a binomial random variable with parameters $c[k]$ and q. We also assume that failing students do not leave until the end of the course in which they are enrolled.

Since, Powersim does not have an implementation of the binomial distribution, the calculation of $F[k]$ must be accomplished by summing Bernoulli random variables. In particular, if $X_i \sim \text{Bernoulli}(q)$, then $\sum_{i=1}^{n} X_i \sim B(n, q)$ (again, see [22]).

The majority of flows in PARSim are modelled with *pulse trains*. For our purposes, a *unit pulse* $\hat{\delta}$ is a positive, piecewise continuous function of t such that:

1. $\hat{\delta}$ is supported on the interval $[0, \Delta t]$, and;

2. $\int_0^{\Delta t} \hat{\delta}(t)\,\mathrm{d}t = 1$.

A pulse train Θ is simply a weighted sum of equally spaced copies of $\hat{\delta}$. For example, if $\{\theta[k] | k \in \mathbb{N}_0\}$ is a sequence of nonnegative numbers and $\Omega > 0$, then

$$\Theta(t) = \sum_{k \in \mathbb{N}_0} \theta[k] \hat{\delta}(t - k\Omega)$$

is a pulse train comprised of pulses, of height $\theta[k]$, spaced Ω apart. We refer the reader to Figure 14, which depicts an example of a typical pulse train.

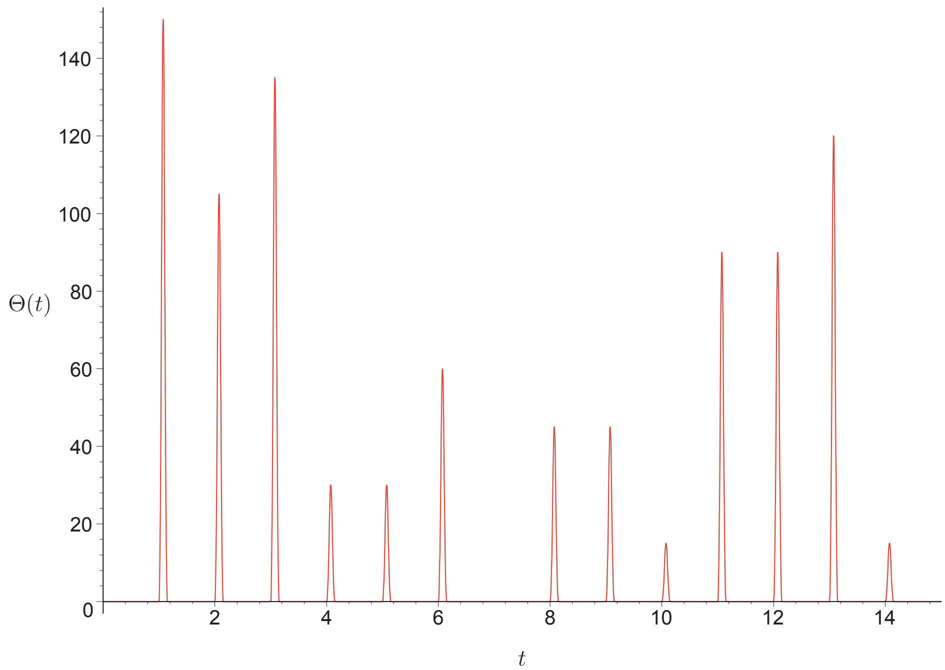

Figure 14: *Typical Pulse Train* $\Theta(t)$

In Powersim, $\hat{\delta}$ is a rectangular pulse with width equal to the simulation the time-step Δt. The number $\theta[k] \in \mathbb{N}_0$ represents a group of pilots that flows from one level to another. This flow starts at time $t = k\Omega$ and finishes at time $t = k\Omega + \Delta t$. It is important to point out that the pulse trains in PARSim are not generally periodic. In fact, $\Theta(t)$ will be $j\Omega$– periodic for some $j \in \mathbb{N}_0$ if and only if $\theta[k + j] = \theta[k]$.

The flow of students into the course $\mathscr{C}$ is given by

$$S_{\text{in}}(t) = \sum_{k \in \mathbb{N}_0} c[k] \hat{\delta}(t - 12k/M), \tag{5}$$

while the outflow of students passing the course is given by

$$S_{\text{pass}}(t) = \sum_{k \in \mathbb{N}_0} (c[k] - F[k]) \hat{\delta}(t - D - 12k/M). \tag{6}$$

The flow of students that fail the course is given by

$$S_{\text{fail}}(t) = \sum_{k \in \mathbb{N}_0} F[k] \hat{\delta}(t - D - 12k/M). \tag{7}$$

Let $C(0)$ be the number of students on course $\mathscr{C}$ at time $t = 0$. In view of Equations (5), (6) and (7), the number of students on course $\mathscr{C}$ at time t is given by the equation

$$\begin{aligned} C(t) &= C(0) + \int_0^t \left\{ S_{\text{in}}(s) - S_{\text{pass}}(s) - S_{\text{fail}}(s) \right\} \mathrm{d}s \\ &= \sum_{k \in \mathbb{N}_0} c[k] \left(\hat{\Lambda}(t - 12k/M) - \hat{\Lambda}(t - D - 12k/M) \right), \end{aligned} \tag{8}$$

where

$$\hat{\Lambda}(t) = \int_0^t \hat{\delta}(s)\, ds \tag{9}$$

is an approximation of the unit step or Heaviside function. It is important to note that the use of rectangular pulses in conjunction with a numerical method, other than Euler's method, may cause unexpected results. In PARSim, the integration method used to estimate a dependent variable can be set within the corresponding level.

The state of the PAT pool can be calculated via

$$P(t) = P(0) + \int_0^t \{U_{\text{net}}(s) - S_{\text{in}}(s)\}\, ds\,.$$

We assume that the arrival of students to the PAT pool is of the form

$$U_{\text{net}}(t) = \sum_{k \in \mathbb{N}_0} u[k]\, \hat{\delta}(t - k\Omega)$$

where $u[k]$ represents the total flow into the PAT pool at time $k\Omega$. The size of the PAT pool at time t is given by

$$P(t) = P(0) + \sum_{k \in \mathbb{N}_0} \left\{ u[k]\hat{\Lambda}(t - k\Omega) - c[k]\hat{\Lambda}(t - 12k/M) \right\}. \tag{10}$$

We can use Equation (10) to provide a rough estimate of how long students in the PAT pool must wait before they are loaded onto course $\mathscr{C}$. We assume that students are loaded on $\mathscr{C}$ in the order of their arrival and consider the waiting time $T_w(t)$ for students that enter the PAT pool at time t. The waiting time is approximately the time that it takes to clear the PAT pool.

The minimum number of serials that must be run, before students arriving at t can be loaded, is

$$M_{\text{min}} = \text{floor}\left(\frac{P(t)}{L_{\text{max}}} \right),$$

where floor(x) is the largest integer n such that $n \leq x$. The next serial to be loaded, after time t, is serial

$$m_{\text{next}} = \text{floor}\left(\frac{M\,t}{12} \right).$$

It follows that students arriving at time t wait a minimum of

$$\begin{aligned} T_w(t) &= \frac{12}{M} \{m_{\text{next}} + M_{\text{min}}\} - t \\ &= \left\{ \text{floor}\left(\frac{P(t)}{L_{\text{max}}} \right) + \text{floor}\left(\frac{M\,t}{12} \right) \right\} - t \end{aligned} \tag{11}$$

months before they are loaded on $\mathscr{C}$.

This page intentionally left blank

4 The Mentor Module

In this section, we turn our attention to the mentor module. We begin by defining the primary notation and stating the major assumptions that will be used throughout the rest of the section. Additional assumptions and notation will be introduced as required. We follow with a discussion on the "acquisition of experience" for inexperienced pilots. This amounts to a derivation of a reasonable model for the dependence of the flying rate of inexperienced pilots on quantities like the flying rate of experienced pilots, the operational flying hours required to upgrade and the experience level.

We subsequently present the formal model underlying the mentor module. As we shall see, the mentor module can be represented by a system of nonlinear delay differential equations (DDEs) with a state dependent delay. We refer the reader to [23] for additional details pertaining to delay differential equations. We complete the section with a cursory analysis of a simplified system of differential equations, which retains the nonlinearity induced by the state dependent delay. We derive a steady state solution to an averaged version of the simplified system and demonstrate how this solution can be used to quickly estimate the absorption capacity as well as the production requirements of a given community[6]. The reader, interested in the Powersim equations of the mentor module, will find them in Annex A.

4.1 Notation and Assumptions

Let $t \geq 0$ denote time in months and let $I(t)$ be the *number of inexperienced pilots* occupying flying positions in some particular operational community (i.e. CH-124, CF-188, etc.) at time t. Let $E(t)$ be the *number of experienced pilots* occupying flying positions in the same community at time t and denote by $S(t)$ the *number of pilots in NOP positions* that are attached to the community at time t. For our purposes, non-operational positions represent either staff positions at wing and higher headquarters or instructor pilot positions at the OTU.

Let N_f be the *established number of flying positions* in the community and assume that N_f is fixed. In addition, we use $N = E + I$ to denote the *operational strength* of the community. To ensure that the value of N does not exceed N_f, the number of experienced pilots is adjusted once per year so that $N(t) = N_f$ whenever possible. This is meant to simulate the Annual Posting Season (APS). In particular, if $N(t)$ exceeds N_f, then the "excess" of experienced pilots are sent to NOP positions attached to the community. On the other hand, if it is found that the $N(t)$ is below the N_f, then pilots are moved from NOP positions back to the community as experienced pilots if they are available. However, if $S(t)$ is smaller than a given critical number $S_{\min}$, then experienced pilots are moved to fill the vacant critical NOP positions regardless of the value of N.

We assume that each inexperienced pilot becomes experienced after flying a total of $H_e = H_o + H_s$ *hours on type*. Here, H_o represents *hours flown at the OTU* and H_s represents the *hours flown at an operational squadron*. We also assume that *all inexperienced pilots at the operational squadron must be "mentored" by an experienced pilot*. Specifically, an inexperienced pilot must be accompanied by an experienced pilot as the H_s hours are acquired. We call this the *mentoring assumption*. The time that it takes an inexperienced pilot to fly H_s hours will be referred to as the *time to experience* (TTE). We will use the symbol T_e to indicate this quantity.

[6] That is, without the need for simulation.

In reality, the process of upgrading pilots is performance based and the actual number of flying hours required to upgrade will vary from pilot to pilot. Generally, H_s is taken to be the minimum number of hours required for a pilot to upgrade. Hence, depending on the variability in the actual hours required, T_e may be optimistic for certain individuals[7]. On occasion, we shall refer to *upgrade time* T_u. This is simply the time at which an inexperienced pilot becomes experienced. In particular, if an inexperienced pilot arrives at an operational squadron at time t, then T_u is defined by $T_u = T_e + t$.

We shall assume that all inexperienced pilots are on *restricted release* and hence cannot leave the air force. Experienced pilots, in the operational communities and in NOP positions, can take their release from the air force. NOP positions include staff positions, at wing and higher headquarters, as well as OTU IP positions. We assume that the yearly percent attrition rates for these two groups of pilots are constant and denote them by A_e and A_s respectively.

Let

$$\beta(t) = \frac{E(t)}{E(t)+I(t)} \tag{13}$$

be the *experience level*, which is the proportion of pilots that are experienced. The experience level is a fundamental measure of effectiveness used to gauge the "health" of an operational community. We will also consider the ratio of experienced pilots to inexperienced pilots. In these cases we will use the notation $\gamma(t) = E(t)/I(t)$. We observe that

$$\gamma = \frac{\beta}{1-\beta}, \tag{14a}$$

$$\beta = \frac{\gamma}{1+\gamma}. \tag{14b}$$

Let $R_e(t)$ be the *instantaneous average flying rate for experienced pilots* (hours per month per pilot) and denote by $R_i(t)$ the *instantaneous average flying rate for inexperienced pilots* (hours per month per pilot)[8]. We assume that the flying rate of an individual experienced pilot is approximately R_e and, similarly, that the flying rate of an individual inexperienced pilot is approximately R_i.

The assumption that all the inexperienced pilots fly at the same rate implies that opportunities to fly are distributed evenly amongst these pilots. In general, this will not be the case. However, these assumptions are needed to keep the complexity of the model reasonable. The flying rates are assumed to be bounded above, which means there is a limit to how many hours a pilot can fly in a given month. As well, we assume that R_e is bounded below so that currency requirements for experienced pilots are always met. In particular, we suppose that there are positive constants $R_e^{\max}$, $R_e^{\min}$ and $R_i^{\max}$ such that

$$R_i(t) \leq R_i^{\max}, \tag{15a}$$

$$R_e^{\min} \leq R_e(t) \leq R_e^{\max}, \tag{15b}$$

for all t. We also assume that $R_i^{\max}$ and $R_e^{\max}$ are least upper bounds for the corresponding rates and set $\rho = R_i^{\max}/R_e^{\max}$.

[7]Of course, it is possible to use either the average or maximum number of flying hours required to upgrade for H_s when these quantities are known.

[8]We are referring to averages over the two groups of pilots and not over time.

Over the course of a year, experienced pilots generate a pool of hours $\mathscr{H}_e$ as they fly. These hours can be allocated to squadron tasks like proficiency training, continuation training and the performance of client missions. The mentor module is based on an assumption of the *primacy of proficiency training*. Accordingly, all hours $\mathscr{H}_e$ are assumed to be available, if needed, to upgrade inexperienced pilots. We will use the symbol ν to denote the expected ratio of inexperienced pilots to experienced pilots on a proficiency mission. We shall refer to ν as the *sortie ratio*. This value is used to account for the fact that an experienced pilot can mentor more than one inexperienced pilot concurrently.

We shall assume that, on any given proficiency mission, each experienced pilot accompanies either one or two inexperienced pilots. Let p_k, $k = 1,2$ denote the proportion missions where the ratio of inexperienced pilots to experienced pilots is $k:1$. In this case, the sortie ratio uniquely determines p_1 and p_2. Since $p_1 + p_2 = 1$ and $p_1 + 2p_2 = \nu$, we have

$$p_1 = 2 - \nu \text{ and } p_2 = \nu - 1.$$

Since the proportions must be non-negative, $1 \leq \nu \leq 2$. We observe that, as the sortie ratio is increased, the proportion of missions in which one inexperience pilot flies decreases and the proportion of missions in which two inexperienced pilots fly increases. Under the additional assumption that the duration of a proficiency mission does not vary appreciably, the pool of hours that is available to upgrade pilots is $\nu\,\mathscr{H}_e$.

The preceding analysis can be extended to the case where not all of $\mathscr{H}_e$ is available for proficiency missions. In this case, p_0 denotes the proportion of missions in which no inexperienced pilots fly and ν is now the conditional expectation

$$\nu = \mathrm{E}(k|\text{proficiency mission}) = \frac{p_1 + 2p_2}{p_1 + p_2} = \frac{p_1 + 2p_2}{1 - p_0}. \tag{16}$$

The total hours available to upgrade pilots is $\nu_a \mathscr{H}_e$ where ν_a, where $\nu_a = p_1 + 2p_2$. Equation (16) provides $\nu_a = \nu\,(1 - p_0)$. Therefore, if we assume that training directives (see [24]) fix the value of ν, then it follows that $\nu_a \mathscr{H}_e$ decreases as p_0 increases. The assumption of primacy of proficiency training is equivalent to the assumption that $p_0 = 0$ when required.

The course module, described in Section 3, is used to model the OTU of each community. In this case, we assume that no students fail their type course. The flow of inexperienced pilots from the OTU to the community is denoted by $a(t)$ and the number of OTU courses per year is denoted by the fixed positive integer $M \in \mathbb{Z}^+$. Since $a(t)$ is the OTU outflow, it does not depend on the duration D of the corresponding type course. In addition, we use $n_i[k]$ to represent the number of students on the k^{th} type course and assume that $L_{\min} \leq n_i[k] \leq L_{\max}$. Here, $L_{\min}$ and $L_{\max}$ represent the minimum and maximum course loadings respectively.

Finally we use the symbol $\mathscr{A}$ to denote the *absorption capacity* of the flying community. As previously stated, the absorption capacity will be shown to be (approximately) inversely proportional to the TTE T_e. This relationship is significant as one of the primary purposes of the mentor module is to provide a means to estimate $\mathscr{A}$. Consequently, if T_e can be estimated reliably, then so can $\mathscr{A}$.

4.2 The Acquisition of Experience

We now explore the functional relationship between the TTE T_e , the flying rates R_i and R_e, the flying hours required to become experienced H_e and the experience level β. We begin our investigation by letting

$$\mathscr{H}(t_1,t_2) = \int_{t_1}^{t_2} R_i(s)\,ds$$

be the total flying hours accumulated by each inexperienced pilot from time t_1 to time t_2, when flying at the rate R_i. Suppose that an inexperienced pilot becomes experienced at time t. That is, the pilot has accumulated H_s hours of flying on the operational squadron at time t. It follows that T_e, can be defined by the equation

$$\mathscr{H}(t-T_e,t) = H_s. \tag{17}$$

We continue by considering the dependence of R_i on the flying rate R_e, the experience level β and the sortie ratio ν. In doing so, we clarify and extend the discussion found in [5]. Let R be the overall average monthly flying rate and, as in [5], suppose that it is constant. This rate can be written in terms of the experience level and the flying rates R_i and R_e. In particular, we have

$$\begin{aligned} R &= \frac{I(t)R_i(t)+E(t)R_e(t)}{I(t)+E(t)} \\ &= (1-\beta(t))\,R_i(t)+\beta(t)R_e(t). \end{aligned} \tag{18}$$

Now, with the sortie ratio in mind, the flying rates R_i and R_e are related by the equation

$$\nu\,E(t)R_e(t) = I(t)R_i(t)$$

or, in terms of the experience level, via the equation

$$\nu\,\beta(t)R_e(t) = (1-\beta(t))\,R_i(t). \tag{19}$$

Equations (18) and (19) form a 2×2 system of equations for the unknown flying rates R_i and R_e. When these equations are written in matrix form, the system can be written as

$$\begin{bmatrix} 1-\beta(t) & \beta(t) \\ 1-\beta(t) & -\nu\,\beta(t) \end{bmatrix} \begin{bmatrix} R_i(t) \\ R_e(t) \end{bmatrix} = \begin{bmatrix} R \\ 0 \end{bmatrix}. \tag{20}$$

The solution of system (20) is

$$R_i(t) = \frac{\nu}{(1+\nu)\,(1-\beta(t))}R, \tag{21a}$$

$$R_e(t) = \frac{1}{(1+\nu)\beta(t)}R, \tag{21b}$$

which provides estimates of the flying rates for inexperience and experienced pilots in terms of β, ν and R.

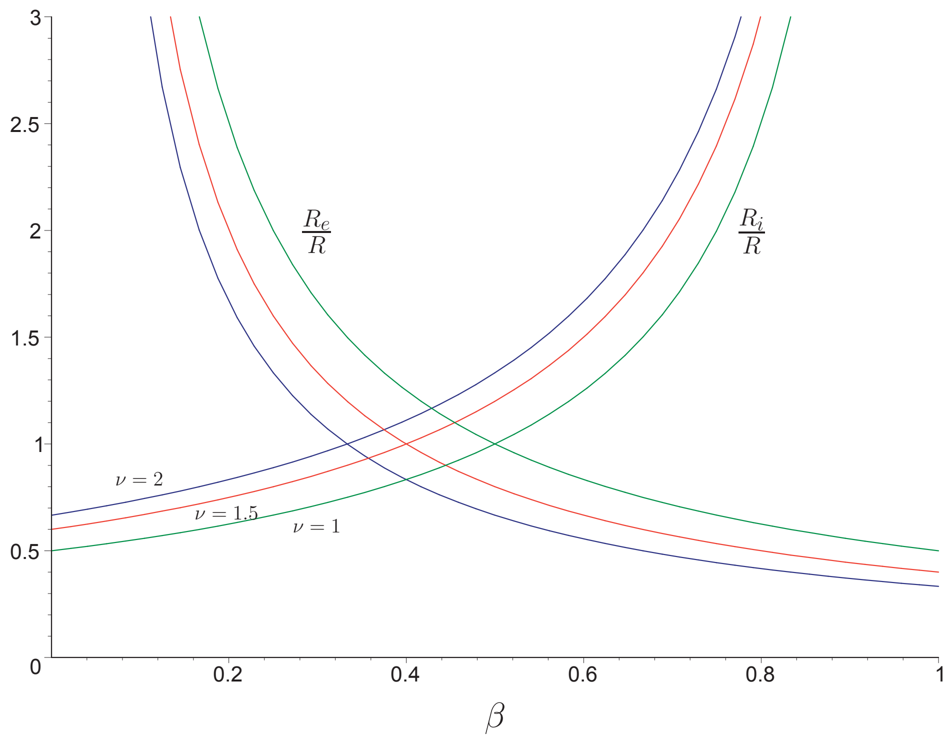

Figure 15: *Ratio of Flying Rates to Overall Average Flying Rate*

An advantage of an approach, based on Equations (21a) and (21b), is that the associated mentor module will be YFR constrained[9]. This fact follows directly from Equation (18). For example, in the fighter community (one pilot per cockpit) the total hours flow on the operational squadrons in the k^{th} year is

$$H_k = \int_{12k}^{12(k+1)} I(t)R_i(t) + E(t)R_e(t)\,dt\,.$$

Now, if $N(t) \equiv N_f$, then Equation (18) implies $H_k = 12RN_f = YFR_k$, which is constant.

We refer the reader to Figure 15. In this figure, the ratios R_i/R and R_e/R are plotted as functions of the experience level for three different values of the sortie ratio ν. The green curves correspond to the case $\nu = 1$, the red curves to the case $\nu = 1.5$ and the blue curves to the case $\nu = 2$. In all cases, we observe that, as the experience level decreases, the experienced pilots fly many times the overall average rate R, while the inexperienced pilot fly less than the average rate. This phenomena has been observed on USAF fighter squadrons (see [5] and [6]) where, as a result of decaying experience levels, combat ready pilots must fly much more than their currency requires in order to mentor inexperienced pilots.

The reader will also observe that the point at which the rates balance depends on ν. In particular, if $\nu = 1$, then $R_e = R_i$ when 50 percent of the pilots are experienced and when $\nu = 2$ the rates are equal when 40 percent of the pilots are experienced. In general, the rates will be balanced when $\beta = 1/(\nu + 1)$, which is a

[9] Here, YFR refers to the portion of yearly flying hours that has been assigned to the operational squadrons.

decreasing function of ν. Consequently, as the ability to simultaneously mentor multiple inexperienced pilots increases, the experience level required to balance the flying rates decreases. The reader should note that taking advantage of the *force multiplying* effect represented by ν may not always be possible and/or may lead to less effective "mentoring".

We can now use Equation (21a) along with Equation (17) to obtain an equation for the TTE $T_{e;1}$.[10] In particular, we have

$$R\frac{\nu}{1+\nu}\int_{t-T_{e;1}}^{t}\frac{1}{1-\beta(s)}\,ds = H_s$$

and, in a "steady state", where the experience level is (approximately) constant, we can solve the preceding equation to obtain

$$T_{e;1}(\beta) = (1-\beta)\frac{1+\nu}{\nu}\frac{H_s}{R}, \tag{22}$$

which relates $T_{e;1}$ to β, ν, H_s and R. We observe that the TTE is a decreasing function of the experience level, is proportional to H_s and is inversely proportional to R; as one might expect. However, $T_{e;1}$ has undesirable limiting properties. When $\beta \to 0^+$, $T_{e;1} \to \frac{1+\nu}{\nu}\frac{H_s}{R}$, which implies that inexperienced pilots upgrade in a finite amount of time when there are no experienced pilots in the community. This is in violation of the mentoring assumption. When $\beta \to 1^-$, the TTE tends to zero, which is absurd.

The undesirable properties of $T_{e;1}$ are caused by the rate R_i. In view of the mentoring assumption, when $\beta \to 0^+$, we expect that R_i tends to zero. This is not the case and $R_i \to \nu R/(\nu+1) \neq 0$. This leads to the finite TTE as the experience levels tend to zero. Additionally, when $\beta \to 1^-$, $R_i \to \infty$. These observations suggest that there are experience levels $0 < \beta_1 \leq \beta_2$ such that Equations (21a) and (21b) are not valid when $\beta \notin (\beta_1, \beta_2)$.

The cause of these deficiencies can be explained in terms of System (20). Let A be the coefficient matrix of the system. It is a simple matter to show that the determinant of A is

$$|A| = -\beta(1-\beta)(1+\nu)$$

and hence the system is singular whenever $\beta = 0, 1$. As a result, System (20) can no longer be simultaneously solved for R_i and R_e when experience levels are either "too high" or "too low".

As one of the primary functions of the mentor module is to provide an assessment of a community's sustainable absorption capacity, it makes sense to take an approach that "fixes" the workload of the experienced pilots. We accomplish this by treating the flying rate of experienced pilots as an input and suggest an expression for R_i in terms of R_e. Although Equation (19) provides an expression of this type, the resulting rate still becomes unbounded as $\beta \to 1^-$.

Let us assume that R_e is a positive constant. In view of Inequality (15a), we take

$$R_i(t) = \min\left(\frac{\nu\beta(t)}{1-\beta(t)}, \rho\right) R_e\,, \tag{23}$$

[10] Since we will compare two distinct formulations for the TTE, we will the use notation $T_{e;1}$ for the remainder of this section.

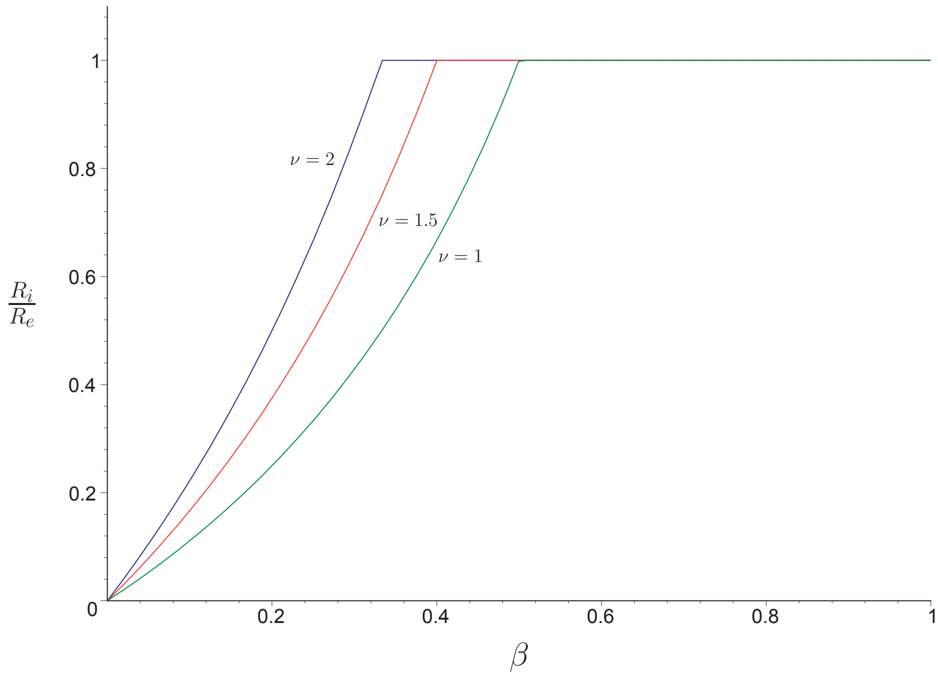

Figure 16: *Ratio of Inexperienced Flying Rate to Experienced Flying Rate*

where the constant ρ. Equation (23) guarantees $R_i(t) \leq \rho R_e^{\max} = R_i^{\max}$ as required. Observe that $\rho = R_i^{\max}/R_e^{\max}$. Plots of the ratio R_i/R_e as a function of β, for various values of ν and $\rho = 1$ appear in Figure 16. Once again, the green curves correspond to the case $\nu = 1$, the red curves to the case $\nu = 1.5$ and the blue curves to the case $\nu = 2$. Here we see that, as the experience level increases, R_i/R_e increases to a maximum of one when $\beta = 1/(1+\nu)$ and remains constant thereafter.[11] In general $R_i/Re = \rho$ for $\beta \geq \rho/(\rho+\nu)$. We also observe that, as β tends to zero, R_i tends to zero as required.

Unfortunately, an approach based on Equation (23) does not lead to a YFR constrained model. The overall average monthly flying rate is given by

$$R(\beta) = \begin{cases} \beta(1+\nu)R_e & \beta < \rho/(\rho+\nu) \\ [(1-\rho)\beta+\rho]R_e & \beta \geq \rho/(\rho+\nu) \end{cases},$$

which depends on the experience level. The reader is referred to Figure 17. In this figure, the ratio R/R_e is plotted as a function of β in the case $\rho = \nu = 1$. Since R_e satisfies Inequality (15b), $R(\beta) \to 0^+$ as $\beta \to 0^+$. Therefore, when experience levels are low, YFR goals will be difficult (if not impossible) to meet. Consequently, in the mentor module, YFR is treated as an output rather than an input. The reader is referred to Annex B, where equations for estimating YFR requirements are derived. A method, to constrain the rate R_e so that YFR goals are approximately met, is presented as well.

The TTE $T_{e;2}$, corresponding to the rate given in Equation (23), is defined by

$$\int_{t-T_{e;2}}^{t} \min\left(\frac{\nu\beta(t)}{1-\beta(t)}, \rho\right) R_e\, ds = H_s$$

[11] Recall that $\beta = 1/(1+\nu)$ is the experience level that balances the rates (21a) and (21b).

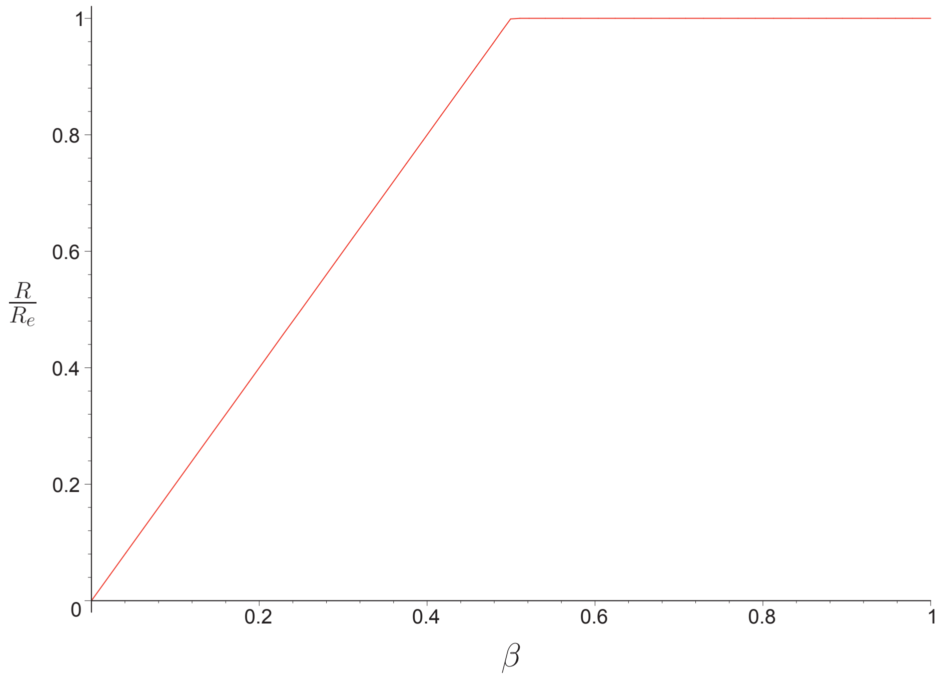

Figure 17: *Ratio of Average Flying Rate to Experienced Flying Rate*

and, in the special case where R_e and β are (approximately) constant, this equation can be solved to obtain

$$T_{e;2}(\beta) = \max\left(\frac{1-\beta}{\nu\beta}, \frac{1}{\rho}\right)\frac{H_s}{R_e}. \tag{24}$$

We have used the fact that $\left[\min(a,b)\right]^{-1} = \max(a^{-1}, b^{-1})$.

Once again, the TTE is proportional to H_s and inversely proportional to R_e. However, unlike Equation (22), Equation (24) is bounded below by $H_s/(\rho R_e)$ as $\beta \to 1^-$. When $\beta \to 0^+$, $T_{e;2} \to \infty$, which implies that an inexperienced pilot will not upgrade when there are no experienced pilots to mentor them. As an example, we compare the TTEs given by Equations (22) and (24) in the case $H_s = 425$, $R_e = 15$ and $\nu = \rho = 1$. These values are indicative of the CF-188 community.

We refer the reader to Figure 18. Here plots of the two TTEs (measured in months) are plotted as functions of the experience level β. The red curve corresponds to $T_{e;1}$, while the blue curve corresponds to $T_{e;2}$. We immediately observe that $T_{e;2}$ underestimates the TTE at all but one experience level (i.e. $\beta = 0.5$). Moreover, in contrast to $T_{e;1}$, the TTE defined by Equation (24) satisfies:

1. $\lim_{\beta \to 0^+} T_{e;2}(\beta) = \infty$ and;

2. $T_{e;2}(\beta) = H_s/(\rho R_e) \approx 28.3$ months for all $\beta \geq \rho/(\rho+\nu) = 0.5$.

We now consider the implications of the TTEs, defined by Equations (22) and (24) with respect to the absorption capacity $\mathscr{A}$. We recall the definition of the absorption capacity given by Equation (1). The

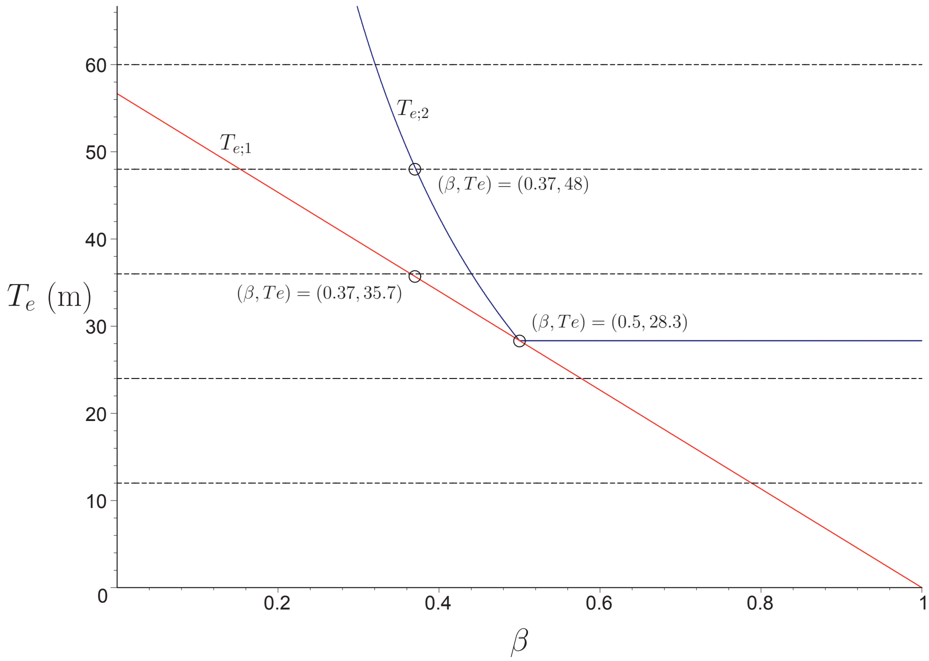

Figure 18: *Times to Experience*

steady state absorption capacity corresponding to $T_{e;1}$ is

$$\mathscr{A}_1 = 12\frac{I}{T_{e;1}} = 12\frac{\nu RI}{(1+\nu)H_s(1-\beta)}\,.$$

The number of flying positions is N_f and this constrains I. If $N \approx N_f$, then $I \approx N_f - E = (1-\beta)N_f$ and

$$\mathscr{A}_1 \approx 12\frac{\nu RN_f}{(1+\nu)H_s}\,, \tag{25}$$

which *does not* depend on β.

A derivation, based on Equation (24), yields

$$\mathscr{A}_2 = 12\min\left(\frac{\nu\beta}{1-\beta},\rho\right)\frac{R_eI}{H_s}$$

and, upon use of $I \approx N_f - E = (1-\beta)N_f$, we find that

$$\mathscr{A}_2 \approx 12\min\left(\nu\beta,\rho(1-\beta)\right)\frac{R_eN_f}{H_s}\,. \tag{26}$$

We point out that, if a unit is not at full strength, then we can simply replace N_f with N in Equations (25) and (26). We refer the reader to Figure 19 where $\mathscr{A}_1$ and $\mathscr{A}_2$ are plotted as a functions of β.

For this example, we have set $\nu = \rho = 1$, $R = R_e = 20$, $H_s = 525$ and $N_f = 71$. The absorption capacity corresponding to Equation (25) is the horizontal red line $\mathscr{A}_1(\beta) \equiv 16.22$, while the absorption capacity

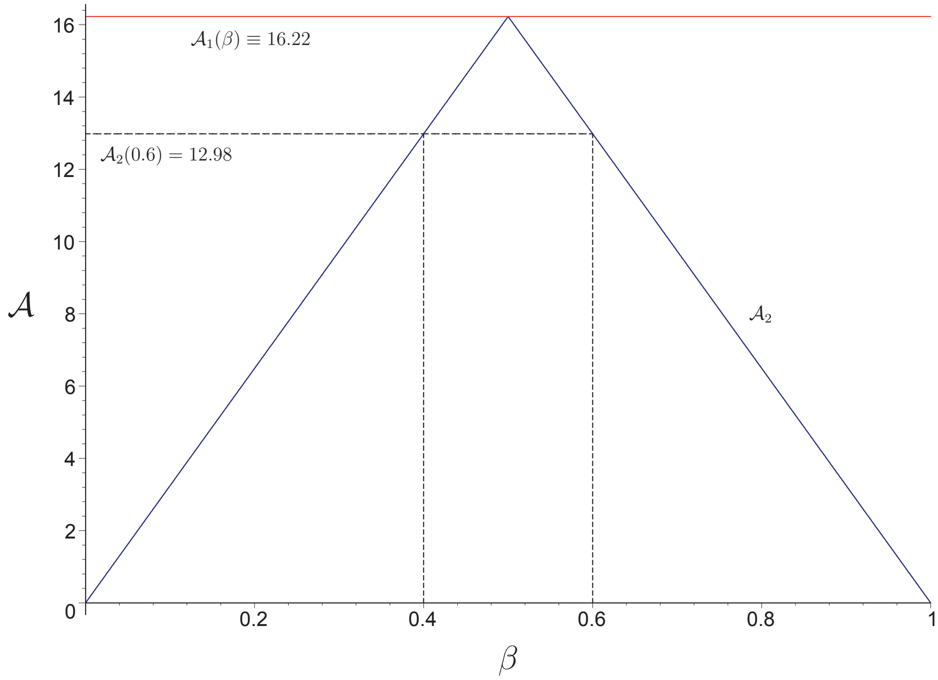

Figure 19: *Steady State Absorption Capacity*

corresponding to Equation (26) is plotted in blue. The absorption capacity $\mathscr{A}_1$ is independent of β. In other words, the same absorption capacity is predicted regardless of the experience level. This is an unrealistic property. On the other hand, $\mathscr{A}_2$ does vary with β. The two estimates agree at the maximum of $\mathscr{A}_2$, which occurs when $\beta = 0.5$. In general, the capacities will agree when $\beta = \rho/(\rho+\nu)$, which is the point where the rate R_i reaches its maximum value of ρR_e.

Some explanation regarding the interpretation of $\mathscr{A}_2$ is in order. The reader will observe that, when $\beta \neq \rho/(\rho+\nu)$, the same absorption capacity is predicted at two distinct experience levels. For instance, when $\beta = 0.4$, 0.6 the estimated absorption capacity is about 13 pilots per year. In general, when $\beta < \rho/(\rho+\nu)$, the absorption capacity is constrained by "low" experience levels. When $\beta > \rho/(\rho+\nu)$, $T_{e;2}$ is constant with respect to β. In this case, the absorption capacity is constrained by the number of available cockpit positions $(1-\beta)N_f$.

For example, assume that there is one type course per year. Upon completion of the type course, a cohort of $\mathscr{A}$ inexperienced pilots is posted to the squadrons. It will take the cohort $T_e = 26.3$ months to upgrade. As a result, there are roughly $26.3/12 \approx 2.19$ cohorts on the squadron at any given time. When $\beta = 0.6$, inexperienced pilots occupy $I = (1-\beta)N_f \approx 28.4$ of the 71 flying positions. Since the most "mature" cohort must upgrade before the next posting, $\mathscr{A} \approx 28.4/2.19 = 13.0$.

4.3 Governing Equations of the Mentor Module

We now turn our attention to the system of differential equations governing the mentor module. This particular module is meant to provide a realistic model of the transfer of experience on the operational squadrons. For our purposes, the transfer of experience and the acquisition of flying hours are equivalent. Since the equations governing the behaviour of OTU are those of the course module, we restrict our attention to the equations governing the evolution of I, E and S.

We make use of the following additional assumptions and notation:

a. The flying rate for experienced pilots R_e is a positive constant. This assumption fixes the workload of the experienced pilots.

b. We shall denote by σ, the flow of experienced pilots to and from NOP positions. "Retread" pilots, moving from NOP positions back to the operational squadrons, are assumed to have no significant impact on the absorption capacity $\mathscr{A}$. As a result, retreads do not undergo an upgrade process and enter the community as experienced pilots.

   As we have already mentioned, the number of flying positions N_f is assumed to be fixed and, during APS, experienced pilots are moved to and from NOP positions so as to track N_f, while still maintaining $S_{\min}$ NOP positions.

c. A fixed net number of experienced pilots $n_u \in \mathbb{N}$ are sent to fill UGPT instructor pilot (IP) positions each year. We will use $\upsilon(t)$ to denote the corresponding flow, which depends only on t.

d. Since inexperienced pilots are assumed to be on restricted release, the flow of pilots, from I to E, is assumed to be a delayed, scaled version of a. The delay is assumed to be T_e.

   If squadron experience levels are low and OTU PAT pool waits are long, then it may be possible that the time from wings to upgrade will exceed the current restriction period of seven years. In such cases, the assumed flow from I to E is no longer valid. However, these situations are usually followed by a "crash" where the simulation stops executing. Allowing the attrition of inexperienced pilots would slow, but not prevent this crash.

e. We shall assume that the attrition rate of experienced pilots is constant. As well, we suppose that the attrition process is adequately modelled by exponential decay. That is, in the absence of all other flows, the attrition flow of experienced pilots can be written as

$$\dot{E}(t) = -\alpha_e E(t),$$

   where $\alpha_e > 0$. Here $\dot{f}(t)$ denotes the derivative of the function f with respect to time t. The attrition process for pilots in NOP positions is modelled in an identical manner. We assume that there is a constant $\alpha_s > 0$ such that

$$\dot{S}(t) = -\alpha_s S(t)$$

   in the absence of all other flows.

   The solution of $\dot{X} = -\alpha X$ is $X(t) = e^{-\alpha t} X(0)$. Consequently, if A is the annual percent attrition

$$\frac{A}{100} = \frac{X(12k) - X(12[k+1])}{X(12k)} = 1 - \exp(-12\alpha)$$

   and the corresponding decay rate is given by

$$\alpha = -\frac{1}{12} \ln\left(1 - \frac{A}{100}\right). \tag{27}$$

In light of the preceding assumptions, the system of ordinary differential equations that governs the flow of pilots within the mentor module can be written in the form

$$\dot{I} = a(t) - a\big(t - T_e\big)\big(1 - \dot{T}_e\big)\,, \tag{28a}$$
$$\dot{E} + \alpha_e E = a\big(t - T_e\big)\big(1 - \dot{T}_e\big) - \upsilon(t) - \sigma(N,S,t)\,, \tag{28b}$$
$$\dot{S} + \alpha_s S = \sigma(N,S,t)\,, \tag{28c}$$

where $N = E + I$. The scaling factor $1 - \dot{T}_e$, in Equations (28a) and (28b), ensures that the area between $a(t)$ and the $t-$axis is conserved when $a(t)$ is delayed by T_e. In particular,

$$\int\limits_{t_1 \le t - T_e(t) \le t_2} a\big(t - T_e(t)\big)\big(1 - \dot{T}_e(t)\big)\,dt = \int_{t_1}^{t_2} a(s)\,ds,$$

as required.

Now, based on our discussion in Section 4.2, we define the TTE by the equation

$$\int\limits_{t - T_e(t)}^{t} R_i(s)\,ds = H_s, \tag{29}$$

with $R_i(t) = \min\left(\nu\gamma(t), \rho\right) R_e$. We recall that $\gamma = E/I$ and $R_e > 0$ is constant. To ensure that the System (28), is well-posed, we verify that Equation (29) can be formally solved for T_e and that the derivative of T_e exists. If we define Φ by

$$\Phi(t) = \int_0^t R_i(s)\,ds,$$

then we can rewrite Equation (29) in the form

$$\Phi(t) - \Phi(t - T_e) = H_s. \tag{30}$$

Without loss of generality, we assume that $E(t) \ge E_{\min} > 0$ and $I(t) \le I_{\max} < \infty$.[12] It follows that $R_i(t)$ is bounded below by the constant $\min(\nu E_{\min}/I_{\max}, \rho)R_e > 0$. As a result, Φ is strictly monotonic and therefore invertible. Since Φ is invertible, Equation (29) can be solved to obtain

$$T_e(t) = t - \Phi^{-1}\left(\Phi(t) - H_s\right). \tag{31}$$

We further assume that E and I are continuous. As a result, R_i is continuous as a function of t and Φ is continuously differentiable. Since Φ is invertible, Φ^{-1} is continuously differentiable as well and, from Equation (31), we conclude that $\dot{T}_e$ is continuous. In fact, with Leibniz's integral rule, we can find an "explicit" expression for $\dot{T}_e$. If we differentiate Equation (29) implicitly with respect to t, then we find that

$$\dot{T}_e = 1 - \frac{R_i(t)}{R_i(t - T_e)} = 1 - \frac{\min\big(\nu\gamma(t), \rho\big)}{\min\big(\nu\gamma(t - T_e), \rho\big)}\,. \tag{32}$$

[12] If these conditions are violated, then the simulation stops execution.

Equation (31) does not provide a practical means to compute T_e. The implementation of T_e in PARSim is indirect and is based on the upgrade time T_u. Suppose that $n_i[k]$ pilots finish their OTU type course at time t_k and are subsequently moved onto the operational squadron(s) as inexperienced pilots.

We define the sequence of functions $\{H_k(t)|\, k \in \mathbb{N}\}$ by

$$\begin{aligned} H_k(t) &= \int_0^t R_i(s)\, ds - \int_0^{t_k} R_i(s)\, ds \\ &= \Phi(t) - \Phi(t_k) \end{aligned}$$

and observe that $H_k(t)$ represents the flying hours that each pilot, in the k^{th} group, has acquired by time t. At each time-step $k\Delta t$, $H_k(t)$ is compared with H_s and, when $H_k(t) = H_s$, $n_i[k]$ pilots are subtracted from the state variable I and added to the state variable E. The solution t of the equation $\Phi(t) - \Phi(t_k) = H_s$ is the upgrade time T_u of the k^{th} group of pilots and the corresponding TTE can be recovered via $T_e = T_u - t_k$.

This implementation requires the definition of two arrays in Powersim: one for each of the sequences $\{H_k(t)|\, k \in \mathbb{N}\}$ and $\{n_i[k]|\, k \in \mathbb{N}\}$. As the length of these arrays must be "hard coded" into the simulation, the user cannot select the time-length of a simulation run from the Excel interface. Moreover, every entry of the array containing the H_k is compared with H_s at every time-step, even though the k^{th} group pilots may have already upgraded. This is inefficient and an alternative implementation should be considered.

Equation (32) forms the basis of alternative implementation of the mentor module. Indeed, one can augment System (28) with Equation (32) to form an equivalent system. The equivalent system is

$$\dot{I} = a(t) - a\big(t - T_e\big) \frac{\min\big(\nu\gamma(t), \rho\big)}{\min\big(\nu\gamma(t - T_e), \rho\big)}, \tag{34a}$$

$$\dot{E} + \alpha_e E = a\big(t - T_e\big) \frac{\min\big(\nu\gamma(t), \rho\big)}{\min\big(\nu\gamma(t - T_e), \rho\big)} - \upsilon(t) - \sigma(N, S, t), \tag{34b}$$

$$\dot{S} + \alpha_s S = \sigma(N, S, t), \tag{34c}$$

$$\dot{T_e} = 1 - \frac{\min\big(\nu\gamma(t), \rho\big)}{\min\big(\nu\gamma(t - T_e), \rho\big)}. \tag{34d}$$

We immediately observe that that the equations, governing the mentor module, form a nonlinear, non-autonomous system of DDEs with a state dependent delay T_e (see [23] and [25]).

Even the simplest delay differential equation is much more complex than its non-delay counterpart. For instance, the general solution of $\dot{X} + \alpha X = 0$ is $X(t) = Ce^{-\alpha t}$. On the other hand, if $\tau > 0$ is constant, then the characteristic equation of the linear DDE $\dot{X} + \alpha X(t - \tau) = 0$ is $\lambda + \alpha e^{\lambda \tau} = 0$. The characteristic equation has an *infinite number of solutions* given by $\lambda_k = -\frac{1}{\tau} W_k(\alpha\tau)$, where W_k, $k \in \mathbb{Z}$, is the k^{th} branch of the Lambert-W function (see [26]). The general solution of the linear DDE is of the form $X(t) = \sum_k B_k \exp(\lambda_k t)$, for "arbitrary" constants B_k.

System (34) also provides a starting point for an existence-uniqueness investigation of solutions to the equations governing the mentor module. However, System (34) does not conform to the assumptions of standard existence-uniqueness theorems for DDEs. To our knowledge, the existence and uniqueness of solutions to systems of the form (34) remains an open question. We relegate this important issue to future research.

We turn our attention to the specification of an appropriate set of initial conditions (ICs) for System (28). Ordinarily, one would simply supply values for $I(0)$, $E(0)$ and $S(0)$. However, the presence of the delays in the system complicates matters. To see how, we rewrite System (28) in integral form. A straightforward derivation yields the integral equations

$$I(t) = \int_{t-T_e(t)}^{t} a(\tau)\,d\tau, \tag{35a}$$

$$\begin{aligned} E(t) &= e^{-\alpha_e t}E(0) + e^{-\alpha_e t}\int_{-T_e(0)}^{t-T_e(t)} \exp\left(\alpha_e \Phi^{-1}(\Phi(\tau)+H_s)\right)a(\tau)\,d\tau \\ &- e^{-\alpha_e t}\int_0^t e^{\alpha_e \tau}\left(\upsilon(\tau)+\sigma(N,S,\tau)\right)d\tau, \end{aligned} \tag{35b}$$

$$S(t) = e^{-\alpha_s t}S(0) + e^{-\alpha_s t}\int_0^t e^{\alpha_s \tau}\sigma(N,S,\tau)\,d\tau, \tag{35c}$$

where we have assumed that $I(0) = \int_{-Te(0)}^{0} a(s)\,ds$.

In view of Equations (35a) and (35b), the computation of $I(t)$ and $E(t)$ for $t > 0$ requires knowledge of $a(\tau)$ for $\tau \in [-T_e(0), t)$. Consequently, the initial state of System (28) depends on $a(\tau)$ for $\tau \in [-T_e(0), 0)$. In turn, the calculation of $T_e(0)$ requires that we specify past values for $\gamma(t)$ (or equivalently $\beta(t)$) on some interval $[t_1, 0)$. However, when we try to specify $t_1 < 0$, we are led to a circular reference. To see this, consider Equation (29). When $t = 0$, Equation (29) implies that $t_1 = -T_e(0)$, which is the value we are trying to calculate!

Perhaps the simplest way to address these issues is to assume that the experience level is constant for all $t < 0$. Since E and I are assumed to be continuous, we take $\beta(t) \equiv \beta(0) = E(0)/(E(0)+I(0))$ for all $t \leq 0$. The TTE is also constant and

$$T_e(t) \equiv T_e(0) = \frac{H_s}{\min\left(\nu\beta(0)/(1-\beta(0)), \rho\right) R_e} \tag{36}$$

for $t \leq 0$. Let $n_i[k]$ be the number of pilots that finish their type course at time $t_k = 12k/M$. To define $a(t)$ for $t \in [-Te(0), 0)$ we set $n_i[k] = I(0)/m$ for $k = -m \ldots -1$, where $m = \text{floor}(MT_e(0)/12)$ and M is the number of type courses per year.

In order to implement the initial configuration in Powersim, we also need to define the functions $H_k(t)$ for $k = -m, \ldots, -1$. Since $\beta(t)$ is constant for $t \leq 0$, the flying rate $R_i(t)$ is also constant for $t \leq 0$ and we set $H_k(t) = |t_k| R_i(0)$ for $k = -m, \ldots, -1$.

It is important to note that the preceding assumptions, pertaining to β and a, are made for convenience only. These assumptions may not accurately reflect a scenario that is based on a given initial experience level $\beta(0)$. As nonlinear dynamical systems are sensitive to perturbations in ICs, the analysis of the corresponding scenario may be inaccurate. It is suggested that the mentor module undergo an analysis of sensitivity in order to assess the variability of results from configurations with a common $\beta(0)$.

We conclude this section by specifying the rates or flows used in System (28). The OTU outflow is taken to be a pulse train of the form

$$a(t) = \sum_{k=-m}^{\infty} n_i[k]\hat{\delta}(t - t_k),$$

where we assume $n_i[k]$ satisfies $L_{\min} \leq n_i[k] \leq L_{\max}$. The outflow of UGPT IP's is given by

$$\upsilon(t) = n_u \sum_{k=0}^{\infty} \hat{\delta}(t - 12k),$$

for some non-negative integer n_{u}.

The rate σ, meant to approximate the process of posting pilots, is more complicated. We begin by defining quantities that we refer to as *capacities*. In particular, we define the *operational squadron capacity* by

$$\Gamma_{\mathrm{ops}}(t) = N_f + n_u - N(t)$$

and the *NOP capacity* by

$$\Gamma_{\mathrm{nop}}(t) = S_{\min} - S(t).$$

If one of the functions Γ_{ops} or Γ_{nop} is positive, then there is room to post pilots into the corresponding positions. On the other hand, if one of the functions is negative, then this indicates a "surplus" of pilots in the corresponding positions. For example, if $\Gamma_{\mathrm{ops}} = -5$, then the operational community has five more more pilots than flying positions.

If we write the rate σ in the form

$$\sigma(N,S,t) = \sum_{k=0}^{\infty} V(12k)\hat{\delta}(t - 12k), \tag{37}$$

then, a pseudo code definition of V is given by:

IF $\Gamma_{\mathrm{ops}}(t) \geq 0$ THEN

    IF $\Gamma_{\mathrm{nop}}(t) \leq 0$ THEN
        $V(t) = -\min\left(\Gamma_{\mathrm{ops}}(t), -\Gamma_{\mathrm{nop}}(t)\right)$
    ELSE
        $V(t) = \Gamma_{\mathrm{nop}}(t)$
    END IF

ELSE
    $V(t) = \max\left(-\Gamma_{\mathrm{ops}}(t), \Gamma_{\mathrm{nop}}(t)\right)$
END IF

Alternatively, if H is the Heaviside function

$$H(t) = \begin{cases} 1 & t \geq 0 \\ 0 & t < 0 \end{cases}, \tag{38}$$

then V has the functional form

$$\begin{aligned} V(t) &= H\left(\Gamma_{\text{ops}}(t)\right)\left\{\left(1-H\left(-\Gamma_{\text{nop}}(t)\right)\right)\Gamma_{\text{nop}}(t)\right. \\ &- \left. H\left(-\Gamma_{\text{nop}}(t)\right)\min\left(\Gamma_{\text{ops}}(t),-\Gamma_{\text{nop}}(t)\right)\right\} \\ &+ \left(1-H\left(\Gamma_{\text{ops}}(t)\right)\right)\max\left(-\Gamma_{\text{ops}}(t),\Gamma_{\text{nop}}(t)\right) \end{aligned}$$

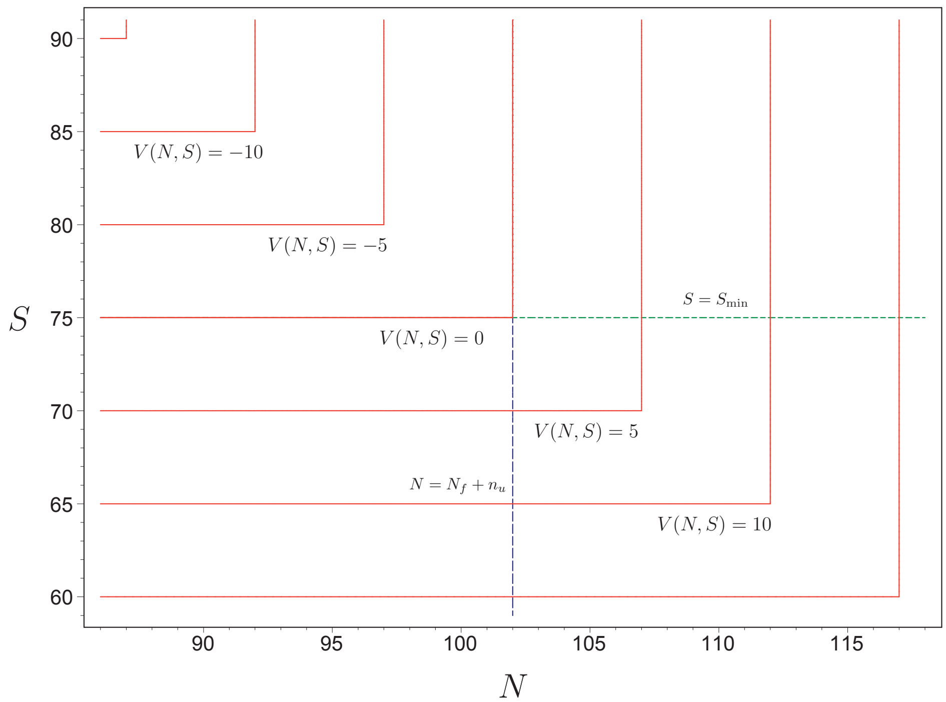

Figure 20: *Level Curves of* $V(N,S)$

We refer the reader to Figure 20. In this figure, the level curves of V are plotted as a function of N and S. Here, we have set $N_f = 100$, $n_u = 2$ and $S_{\text{min}} = 75$. The line $N = N_f + n_u$ is depicted in blue, while the line $S = S_{\text{min}}$ is depicted in green. The reader will observe that $V > 0$ whenever $S < S_{\text{min}}$ or whenever $N > N_f + n_u$ and $S > S_{\text{min}}$. These are the conditions under which there is a net flow of experienced pilots from the squadrons to NOP positions. On the other hand, when $N < N_f + n_u$ and $S > S_{\text{min}}$, $V < 0$. That is, under these conditions, there is a net flow of pilots from NOP positions to the squadrons.

4.4 An Analysis of the Equations Governing the Mentor Module

Recall that one of our goals is to estimate the sustainable absorption capacity $\mathscr{A}$ of a flying community. In Section 4.2, Equation (26) was derived under the assumption of a constant experience level β. This assumption is not realistic. For example, just before APS, a flying community is fully manned with N_f pilots and an experience level β. During APS, a cohort of n_i OTU graduates is posted to the operational squadrons as inexperienced pilots, while n_i experienced pilots are moved to staff positions. In this case, the experience level changes rapidly from β to $\beta - n_i/N_f$. This example suggests that, if a steady state exists, then β (as well as E, I and S) will vary periodically in time. As a result, estimates of $\mathscr{A}$ should be based on an analysis of the periodic (or neutral) steady states of System (28).

An analysis of the periodic solutions[13] of System (28) is beyond the scope of this report. In this section, we consider the average behaviour of E, I and S resulting from periodic rates a, σ and υ and establish conditions under which a variant of Equation (26), namely

$$\mathscr{A} \approx 12 \min(\nu\beta, \rho(1-\beta)) \frac{R_e N_f}{H_s},$$

is valid. We shall see that the right hand side of Equation (26) can be regarded as an (approximate) upper bound for the absorption capacity. We establish a lower bound for $\mathscr{A}$ as well. This lower bound captures the requirements to address losses due to attrition and to support NOP and UGPT IP positions.

We begin by simplifying System (28) to

$$\begin{aligned} \dot{I} &= a(t) - a\big(t - T_e(t)\big)\big(1 - \dot{T}_e(t)\big), & \text{(40a)} \\ \dot{E} + \alpha_e E &= a\big(t - T_e(t)\big)\big(1 - \dot{T}_e(t)\big) - \sigma(t) - \upsilon(t), & \text{(40b)} \\ \dot{S} + \alpha_s S &= \sigma(t). & \text{(40c)} \end{aligned}$$

This system retains the nonlinearity induced by T_e. However, σ is assumed to be independent of N and S. As a result, System (40) does not track the number of flying positions N_f nor the minimum number of NOP positions $S_{\min}$. These requirements will be incorporated by imposing constraints on the averages of S and $N = E + I$.

Let $f(t)$ be a twelve-periodic[14]. We define the *average* of f by

$$\langle f \rangle = \frac{1}{12} \int_b^{b+12} f(t)\, dt$$

and note that $\langle f \rangle$ is independent of the constant b. The function f can be expressed as $f = \langle f \rangle + \varepsilon_f$ with $\langle \varepsilon_f \rangle = 0$. We call the function ε_f the *variation* of f about the average $\langle f \rangle$. The constraints on the number of NOP and flying positions are expressed as

$$\begin{aligned} \langle S \rangle &\geq S_{\min}, & \text{(41a)} \\ \langle N \rangle &= N_f. & \text{(41b)} \end{aligned}$$

We assume that the rates in System (40) are twelve-periodic and seek the averages $\langle E \rangle$, $\langle I \rangle$ and $\langle S \rangle$. In particular, we suppose that a, σ and υ are given by

$$\begin{aligned} a(t) &= n_i \textstyle\sum_{k=-\infty}^{\infty} \hat{\delta}(t - 12k), & \text{(42a)} \\ \sigma(t) &= n_s \textstyle\sum_{k=-\infty}^{\infty} \hat{\delta}(t - 12k), & \text{(42b)} \\ \upsilon(t) &= n_u \textstyle\sum_{k=-\infty}^{\infty} \hat{\delta}(t - 12k), & \text{(42c)} \end{aligned}$$

[13] We are assuming that periodic solutions exist.
[14] That is, $f(t)$ is periodic with a period of twelve months.

where n_i, n_s and n_u are nonnegative integers. Recall that $\hat{\delta}$ is a positive, unimodular, piecewise continuous function that is supported on $[0,\Delta t]$. We immediately see that $n_i = 12\langle a\rangle$, $n_s = 12\langle\sigma\rangle$ and $n_u = 12\langle\upsilon\rangle$.

Implicit in Equation (42a) are the assumptions that there is only one type course per year and that each type course is loaded with n_i students. Similar assumptions are implied by Equations (42b) and (42c). For instance, since $n_s \geq 0$, we are assuming that, in steady state, the net flow of pilots is from E *to* S. The reader should contrast this with the behaviour of σ that is discussed in Section 4.3.

Now, consider Equation (40c). Upon averaging, we find that

$$\langle S\rangle = \frac{\langle\sigma\rangle}{\alpha_s}, \tag{43}$$

where we have used the fact that $\langle\dot{f}\rangle = 0$ whenever f is a differentiable, twelve-periodic function. Since $\langle\dot{I}\rangle = 0$, Equation (40a) implies $\langle a\rangle = \left\langle a\big(t - T_e\big)\big(1 - \dot{T}_e\big)\right\rangle$ and consequently, Equation (40b) yields

$$\langle E\rangle = \frac{\langle a\rangle - \langle\sigma\rangle - \langle\upsilon\rangle}{\alpha_e}. \tag{44}$$

Recall Equation (27) and suppose that A_e is the yearly percent attrition of experienced pilots. As a consequence of Equation (44), to maintain an average of $\langle E\rangle$ experienced operational pilots, the squadrons must absorb

$$n_i = \ln\left(\frac{1}{1 - A_e/100}\right)\langle E\rangle + n_s + n_u$$

inexperienced pilots per year. Unless attrition is extreme then

$$n_i \approx \frac{A_e}{100}\langle E\rangle + n_s + n_u,$$

which is expected intuitively. We may make a similar observation based on Equation (43).

To find an expression for $\langle I\rangle$, we rewrite Equation (40a) as

$$I(t) = \int_{t-T_e(t)}^{t} a(s)\,ds.$$

and substitute $a = \langle a\rangle + \varepsilon_a$. Recall that Equation (29) defines T_e uniquely. Therefore, since R_i is twelve-periodic, then T_e is twelve-periodic as well and our definition of the average makes sense for T_e. It follows from the preceding equation that

$$\langle I\rangle = \langle a\rangle\langle T_e\rangle + \left\langle \int_{t-T_e(t)}^{t} \varepsilon_a(s)\,ds \right\rangle. \tag{45}$$

Observe that, unlike S and E, the average behaviour of I depends on the variation of the relevant rate.

Now, in view of the Inequality (41a), we require

$$\langle\sigma\rangle \geq \alpha_s S_{\min}. \tag{46}$$

Since $E = \beta N$, then the assumption

$$\langle \beta N \rangle \approx \langle \beta \rangle \langle N \rangle \,, \tag{47}$$

implies that $\langle a \rangle \approx \alpha_e \langle \beta \rangle \langle N \rangle + \langle \sigma \rangle + \langle \upsilon \rangle$. We can now use Inequality (46) and Equation (41b) to derive the (approximate) lower bound

$$\mathscr{A} \gtrapprox 12\alpha_e N_f \langle \beta \rangle + 12\alpha_s S_{\min} + n_u \,, \tag{48}$$

where we have set $\langle a \rangle = \mathscr{A}/12$. Inequality (48) specifies the minimum number of pilots needed to be absorbed each year in order to sustain, on average, N_f flying position and $S_{\min}$ NOP positions.

Equation (47) is certainly valid when the experience level is constant. This assumption was made in Section 4.2. More generally, let $N = \langle N \rangle + \varepsilon_N$. If

$$\left| \frac{\varepsilon_N}{\langle N \rangle} \right| \ll 1 \tag{49}$$

then

$$\beta = \frac{1}{\langle N \rangle} \frac{E}{1 + \frac{\varepsilon_N}{\langle N \rangle}} = \frac{E}{\langle N \rangle} \left(1 - \frac{\varepsilon_N}{\langle N \rangle} \pm \cdots \right) \approx \frac{E}{\langle N \rangle} \,.$$

Hence Equation (47) is valid as long as the "relative variation" of N about its mean is not too large. Inequality (49) is equivalent to the condition $0 \ll N \ll 2\langle N \rangle$. Although we have not established this mathematically, we conjecture that this condition will hold when α_e is not too large and $n_i - n_s - n_u$ is small compared to N.

To establish an upper bound for $\mathscr{A}$, we consider Equation (45). We begin by estimating the magnitude of the last term on the right hand side. Since ε_a is twelve-periodic with $\langle \varepsilon_a \rangle = 0$, then there exists a $b_t \in [0, 12)$ such that

$$\int\limits_{t - T_e(t)}^{t} \varepsilon_a(s)\, ds = \int_0^{b_t} \varepsilon_a(s)\, ds \,.$$

Assuming $b_t \geq \Delta t$, then from Equation (42a) we obtain

$$\int_0^{b_t} \varepsilon_a(s)\, ds = (12 - b_t)\langle a \rangle$$

and consequently

$$0 < \left\langle \int\limits_{t - T_e(t)}^{t} \varepsilon_a(s)\, ds \right\rangle \leq 12 \langle a \rangle \,. \tag{50}$$

In view of Expression (50), Equation (45) yields

$$\frac{\langle I \rangle}{\langle T_e \rangle + 12} \leq \langle a \rangle < \frac{\langle I \rangle}{\langle T_e \rangle} \,.$$

Equation (47) implies $\langle I \rangle \approx (1 - \langle \beta \rangle)\langle N \rangle$ and therefore

$$\frac{(1 - \langle \beta \rangle) N_f}{\langle T_e \rangle + 12} \lessapprox \langle a \rangle \lessapprox \frac{(1 - \langle \beta \rangle) N_f}{\langle T_e \rangle} \,, \tag{51}$$

where we have enforced the constraint given by Equation (41b). Subsequently, we shall restrict our attention to the upper bound in Expression (51), but note that $\langle a\rangle \approx (1-\langle\beta\rangle)N_f/\langle T_e\rangle$ when T_e is much larger than twelve months. We continue by examining the relationship between $\langle T_e\rangle$ and $\langle\beta\rangle$.

Let $\delta_\gamma = \gamma - \frac{\langle\beta\rangle}{1-\langle\beta\rangle}$ and $\beta = \langle\beta\rangle + \varepsilon_\beta$, then

$$\frac{1-\langle\beta\rangle}{\langle\beta\rangle}\delta_\gamma = \frac{\varepsilon_\beta}{\langle\beta\rangle(1-\langle\beta\rangle)}\,.$$

We assume that

$$\left|\frac{\varepsilon_\beta}{\langle\beta\rangle(1-\langle\beta\rangle)}\right| \ll 1\,, \tag{52}$$

which is analogous to Inequality (49). Since

$$\begin{aligned}\min(\nu\gamma,\rho) &= \frac{\nu\langle\beta\rangle}{1-\langle\beta\rangle}\min\left(1+\frac{1-\langle\beta\rangle}{\langle\beta\rangle}\delta_\gamma, \frac{1-\langle\beta\rangle}{\nu\langle\beta\rangle}\rho\right)\\ &\approx \min\left(\frac{\nu\langle\beta\rangle}{1-\langle\beta\rangle},\rho\right)\,,\end{aligned}$$

then, in view of Equation (29), we obtain

$$\langle T_e\rangle \approx \frac{H_s}{R_e\min\left(\frac{\nu\langle\beta\rangle}{1-\langle\beta\rangle},\rho\right)}\,. \tag{54}$$

From Equation (54) and the upper bound in Expression (51), we find that

$$\mathscr{A} \lessapprox 12\min\left(\nu\langle\beta\rangle,\rho(1-\langle\beta\rangle)\right)\frac{R_eN_f}{H_s}\,, \tag{55}$$

which is consistent with Equation (26).

We point out that a bound on the error in Equation (54) has been neglected and should be examined in the future. However, the derivation of this error bound may be complicated by the use of the “min” function in the definition of R_i, see Equation (23). Alternatively, if R_i were a smooth function of γ, then

$$R_i(\gamma) \approx R_i\left(\frac{\langle\beta\rangle}{1-\langle\beta\rangle}\right) + {R_i}'\left(\frac{\langle\beta\rangle}{1-\langle\beta\rangle}\right)\delta_\gamma\,,$$

which may simplify the derivation. A suitable choice for a smooth R_i could be based on a logistic sigmoid function, see for example [27].

Now, since $0 < \beta, \langle\beta\rangle < 1$, then Inequality (52) is equivalent to

$$\langle\beta\rangle^2 \ll \beta \ll \langle\beta\rangle(2-\langle\beta\rangle)\,.$$

We refer the reader to Figure 21 where the preceding bounds are plotted as functions of $\langle\beta\rangle$. For reference, the line $\beta = \langle\beta\rangle$ is plotted in green. The vertical distances between the line $\beta = \langle\beta\rangle$ and either the upper or lower bounds provide bounds on ε_β. For instance, if $\langle\beta\rangle = 0.6$, then Inequality (55) may not be accurate if the experience levels vary more than 24% above or below the average.

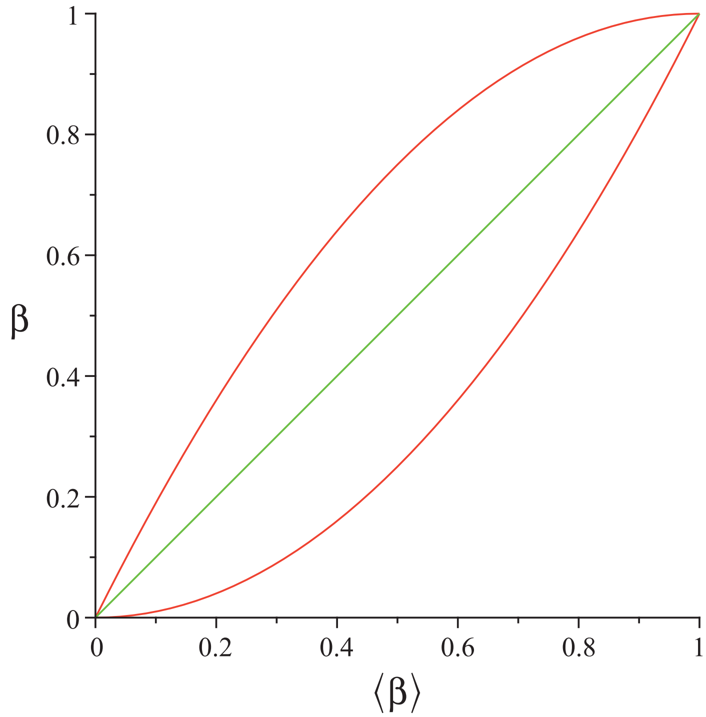

Figure 21: *Upper and Lower Bounds on the Variation* ε_β

Although additional analysis is required, the approximate Inequalities (48) and (55) have been validated by numerous PARSim runs. These inequalities form the basis of a prototypical spreadsheet *absorption estimator*. The input sheet of this estimator is depicted in Figure 22. The absorption estimator can be used to quickly validate proposed steady state absorption for a given flying community. The user supplies inputs that are analogous to those found in the PARSim input sheets (see Figure 6). The primary outputs are upper and lower bounds on the steady state $\mathscr{A}$. We relegate additional discussion to future work.

Simulation runs indicate that the upper bound, given by Inequality (55), is rarely attained. In part, this is due to the use of averaging. The source of the upper bound is Equation (45) in which averaging has smoothed out the variation of $a(t)$. In effect, the yearly allocation of *ab initio* pilots trickles onto the unit at a constant rate instead of via a small number of discrete pulses. Interestingly, this suggests that absorbing $m > 1$ cohorts of n inexperienced pilots each year is easier than absorbing one cohort of $m \times n$ pilots. This makes sense.

Suppose that a group of M *ab initio* pilots is moved to an operational squadron with N pilots and experience level β. If the new strength of the unit is $M+N$, then the relative change in experience level is

$$\Delta\beta = \frac{\frac{\beta N}{N+M} - \beta}{\beta} = -\frac{M}{N+M},$$

which decreases with M. Consequently, an OTU which runs more serials with smaller class sizes produces smaller “shocks” to the experience levels of the downstream squadrons than an OTU that runs fewer serials with larger class sizes.

Unit Absorption Estimator

	CF188	CC130	CP140	CH146	CH124	CH149				
Exp’d Attrition (%)	7.0	7.0	6.0	8.0	3.0	4.0				
Staff & OTU IP Attrition (%)	7.0	15.0	12.0	8.0	8.0	8.0				
Exp’d flying rte (hrs/mo/pilot)	12.0	22.0	30.0	25.0	20.0	20.0				
Inexp’d flying rte (hrs/mo/pilot)	12.0	22.0	30.0	20.0	20.0	20.0				
Total hrs to upgrade	500	1200	1100	800	600	500				
OTU hrs	76	91	132	60	75	0				
Sqn flying hrs to upgrade	424	1109	968	740	525	500				
Sortie ratio	1.00	1.20	1.50	1.33	1.00	1.00				
Flying pos’n (pilots)	68	110	62	181	71	46				
Min staff & OTU IP (pilots)	55	40	30	80	55	30				
UGPT IP per yr	2	3	2	5	2	0	OME	NFTC IP	Totals	Steady States
Break point (pilots/yr)	11.5	14.3	13.8	36.7	16.2	11.0	2.0	10.0	115.6	1927.0
Midpoint (pilots/yr)	10.5	14.2	11.2	29.5	12.3	7.5	2.0	10.0	97.2	1619.4
Floor point (pilots/yr)	9.4	14.1	8.6	22.3	8.3	4.0	2.0	10.0	78.7	1311.8
Proposed production (pilots/yr)	5	6	6	21	10	0	2	10	60	1000.0
Target production (pilots/yr)	10	15	11	30	6	4	2	10	88	1466.7
User Inputs										
Main Outputs										
Other Outputs										

Pan MOC Analysis

PML	Current Manning	MOC Attrition (%)	Steady State NWG
1468	1398	6.00	88.1

Rebuild Scenario I		Rebuild Scenaro II	
Specify NWG production	Yrs to Recovery	Specify Yrs to Recovery	NWG Production
90	18.7	10.0	93.0
Annual Starts	Phase II Starts	Annual Starts	Phase II Starts
127.0	107.1	131.3	110.7

Figure 22: *Absorption Estimator Input Sheet*

Denote by $U_{\mathscr{A}}$ and $L_{\mathscr{A}}$, the upper and lower bounds of $\mathscr{A}$ that are given by Inequalities (55) and (48), respectively. We refer the reader to Figure 23, where $U_{\mathscr{A}}$ and $L_{\mathscr{A}}$ are plotted as functions of the average experience level $\langle\beta\rangle$. The upper bound is plotted in red and the lower bound is plotted in blue. In this case, we have set $H_s = 430$, $R_e = 15$, $N_f = 68$, $S_{\min} = 55$, $\upsilon = 2$ and $\rho = \nu = 1$. As we have mentioned, the upper bound absorption capacity is rarely attained in simulations. We assume that the “real” absorption capacity of the community lies in between the upper and lower bounds and suggest the midpoint

$$M_{\mathscr{A}} = \begin{cases} \frac{1}{2}\big(U_{\mathscr{A}} + L_{\mathscr{A}}\big) & U_{\mathscr{A}} \geq L_{\mathscr{A}} \\ 0 & \text{otherwise} \end{cases},$$

as an estimate of $\mathscr{A}$. In Figure 23, $M_{\mathscr{A}}$ is plotted in green.

Note that the upper and lower bounds intersect at $\langle\beta\rangle \approx 0.25$ and $\langle\beta\rangle \approx 0.67$. This implies that sustained operations are not possible when experience levels are lower than 25% or greater than 67%. The maximum possible absorption occurs when the average experience level is 50%. Here the upper bound on absorption is about 14.2 pilots per year, while the lower bound on absorption is about 8.5 pilots per year. The midpoint suggests that the community should absorb about 11 pilots per year to sustain experience levels at 50%. In general, the maximum absorption occurs when

$$\langle\beta\rangle = \frac{\rho}{\rho + \nu}.$$

This observation was made in Section 4.2. On the other hand, if the preferred experience level is 60%, then the community should be absorbing about 10 pilots per year[15].

Finally, we define the distance $D_{\mathscr{A}}$ by

$$D_{\mathscr{A}} = \begin{cases} U_{\mathscr{A}} - L_{\mathscr{A}} & U_{\mathscr{A}} \geq L_{\mathscr{A}} \\ 0 & \text{otherwise} \end{cases}$$

[15]We are assuming midpoint absorption.

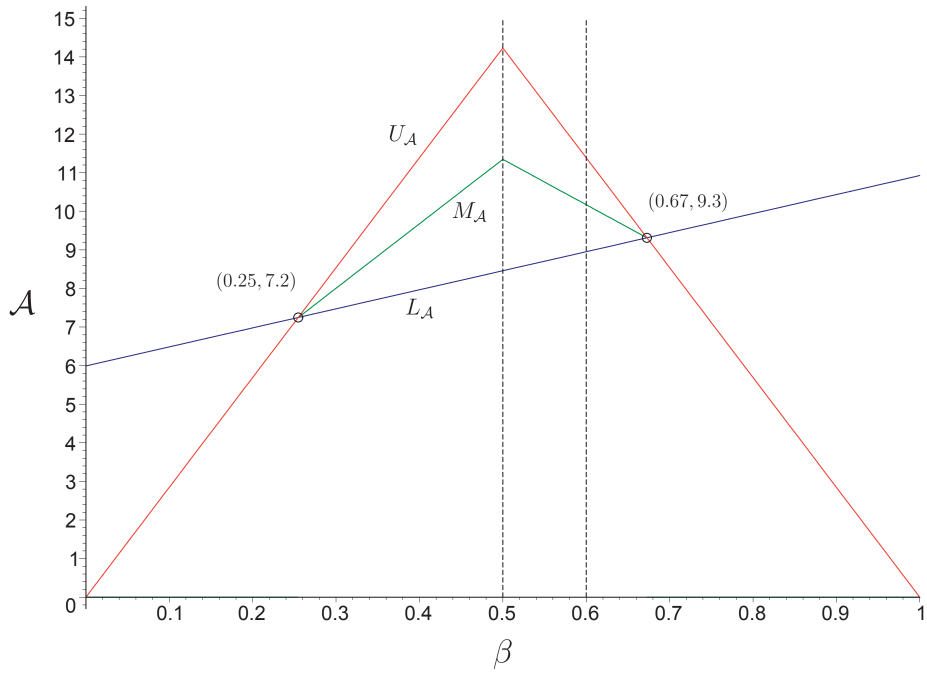

Figure 23: *Absorption Bounds*

and propose that $D_{\mathscr{A}}$ is a measure of the flexibility that a unit has in meeting absorption objectives at a given experience level. Observe that $D_{\mathscr{A}}$ takes on its absolute maximum value when $\langle\beta\rangle = 0.5$.

This page intentionally left blank

5 Conclusions and Directions for Further Research

In this report we have considered a number of the technical aspects pertaining to the development of the Pilot Production Absorption Retention Simulation (PARSim). We have presented mathematicals description of the major components of PARSim and have discussed some of their important properties. It is hoped that we have provided an implementation free understanding of PARSim that will encourage future research and development in this important area.

We have explored, extended and clarified the RAND notion of absorption and have defined a corresponding time to experience (TTE) that varies as a function of unit experience levels. Based on this new definition of the TTE, we have developed a system of nonlinear delay differential equations (DDEs) with state dependent delays. It is this system which governs the behaviour of the mentor module in PARSim. We then performed a cursory analysis of this system of DDEs. We subsequently demonstrated how this analysis provides estimates of the steady state absorption capacity and requirements of a given unit. This analysis will facilitate the creation of simple, spreadsheet based, tools that avoid the need for numerous and lengthy simulation runs.

Future work on PARSim should see the incorporation of a retread stream in communities for which re-certifying pilots requires a significant number of hours. Consequently, the analysis performed in Section 4 could be repeated for an extension of System (28). Similarly, the inclusion of flow to account for pilots changing type would enhance model fidelity for communities like CC130.

The mentor module could include the ability to specify different "learning rates" for inexperienced pilots. This enhancement would eliminate the need to use a single hours on type marker as a measure of proficiency. The flexibility of allowing inexperienced pilot to attrit could also be considered.

In any event, the analysis of System (28) is far from complete and additional work is required. Most importantly, the question of existence and uniqueness of solutions to System (28) (or equivalently System (34)) needs to be addressed. Indeed, a negative result to such an investigation would, in all likelihood, necessitate changes to System (28). As a result, the mentor module in PARSim would need to be modified.

The way in which the mentor module is initialized should be examined in greater detail. The assumptions made in Section 4.3 that pertain to the past values of a and T_e, are made for convenience. These assumptions may not accurately reflect a scenario that is based on a given initial experience level $\beta(0)$. As nonlinear dynamical systems are sensitive to perturbations in ICs, the analysis of the corresponding scenario may be inaccurate. The mentor module should undergo a sensitivity analysis in order to assess the variability of results from initial configurations with a common $\beta(0)$.

The analysis, performed in Section 4.4, is based on averages. At a minimum, this analysis should be extended to include the variation about the average. Ideally, the uncertainty in the upper and lower bound $U_{\mathscr{A}}$ and $L_{\mathscr{A}}$ should be quantified. Additionally the existence of periodic solutions was assumed. This question should be examined and, if at all possible, the stability of periodic steady states should be addressed.

In the longer term, different modelling and simulation methodologies should be considered. Due to the small size of the pilot MOC, discrete event and Monte Carlo simulation methodologies should be investigated. It certainly would be of interest to compare the results from an “aggregate last” approach to those from the “aggregate first” approach of SD.

References

1. Stroud, E. D., Lieutenant-Colonel (2002). CF18 Fighter Force – The Way Ahead. A3 Fighter, 1 Cdn Air Div Headquarters.

2. Kissman, H. P., Major (2003). The Canadian Operational Fighter Force Production Absorption Dilemma. Research Report. Air Command and Staff College.

3. Various (2003). Pilot Production and Absorption Working Group Report. A4 Training and Support, 1 CAD Headquarters.

4. Latchman, S. and Hunter, C. (2002). Preliminary Results from the Pilot Production/Absorption/Retention Simulation (PARSim) Model. (Research Note 0202). CORA.

5. Taylor, W. W. et al. (2000). The Airforce Pilot Shortage: A Crisis for Operational Units?. (Project Airforce Report MR-1204-AF). RAND. Santa Monica.

6. Taylor, W. W. et al. (2002). Absorbing Airforce Fighter Pilots: Parameters Problems and Policy Options. (Project Airforce Report MR-1550-AF). RAND. Santa Monica.

7. Forester, Jay (1962). Industrial Dynamics, Cambridge: MIT Press.

8. Desmier, P. E. (1998). Attrition and the Restricted Release Policy. (Project Report PR9816). ORD.

9. Boulet, D. P. (1995). The UPT Flow Simulator and Changes to the Undergraduate Pilot Training System. (Staff Note 93/4). ACH.

10. Boulet, D. P. (1993). The Undergraduate Pilot Training Flow Simulator and the PAT Problem. (Staff Note 94/3). ACH.

11. Iyde, Osmand, First Lieutenant (TUAF) (1998). The Simulated Pilot Attrition and Training (SimPAT) Model Phase 1 – Undergraduate Pilot Training. (Project Report PR9819). ORD.

12. Maani, K. E. and Cavana, R. Y. (2000). Systems Thinking and Modelling – Understanding Change and Complexity, Auckland: Pearson Education.

13. Sternman, J. D. (2002). Business Dynamics: Systems Thinking and Modeling for a Complex World, Boston: McGraw Hill.

14. Godfrey, K. (1983). Compartmental Models and their Application, New York: Academic Press.

15. Press, W. H. et al. (1992). Numerical Recipes in C: The Art of Scientific Computing, Cambridge: Cambridge University Press.

16. Martin, L. A. (1997). An Introduction to Feedback. (MIT System Dynamics in Education Project D-4691). Massachusetts Institute of Technology.

17. Ashford, A. C. (1995). Unexpected Behaviour in Higher-Order Feedback Loops. (MIT System Dynamics in Education Project D-4455-1). Massachusetts Institute of Technology.

18. Ford, D. N. (1999). A behavioral approach to feedback loop dominance analysis. *System Dynamics Review*, **15**(1), 3–36.

19. Doyle, J. D. et al. (1992). Feedback Control Theory, New York: Macmillan Publishing Company.

20. Skufca, J. D. (2004). Analysis Still Matters: A Surprising Instance of Failure of Runge-Kutta-Felberg ODE Solvers. *SIAM Review*, **46**(4), 729–737.

21. Bhattacharyya, G. K. and Johnson, R. A. (1977). Statistical Concepts and Methods, New York: John Wiley & Sons.

22. Law, A. M. and Kelton, W. D. (2000). Simulation Modeling and Analysis, Boston: Mc-Graw Hill.

23. Hale, J. K. and Lunel, S. M. Verduyn (1993). Introduction to Functional Differential Equations, New York: Springer-Verlag.

24. 1 Canadian Air Division Headquarters (2002). Fighter Pilot Training Directive. A3 Fighters. B-GA-432-000/FP-002.

25. Driver, R. D. (1977). Ordinary and Delay Differential Equations, New York: Springer-Verlag.

26. Corless, R. M. et al. (1996). On the Lambert W Function. *Advances in Computational Mathematics*, **5**, 329–359.

27. Bishop, C. M. (1995). Neural Networks for Pattern Recognition, Oxford: Oxford University Press.

Annex A
Powersim Equations

This annex contains the Powersim equations for the mentor and course modules.

A.1 Mentor Module Equations

```
mainmodel Mentor Module {

aux add yr {
autotype Real
autodim 0..20
def FOR(k=year|IF('YFR sw','E yr'[k]-1/(SR+1)*'upg yr'[k],
1/2*('E yr'[k]-'upg yr'[k]))) }

const AFT crs per yr {
autotype
Real
init 1 }

aux arrtm {
autotype Real
autodim 0..60
def FOR(k='OUT crs'|'crs dur'+52/'numcrs'*k) }

aux at fact E {
autotype Real
def (1/1-'pct at E'/100)^(1/52)-1)*Ex }

aux at fact N {
autotype Real
def (1/1-'pct at N'/100)^(1/52)-1)*NOP }

aux Avg hr mo {
autotype Real
def ('E hr mo'*Ex+52/12*'I hr wk'*In)/(Ex+In) }

aux cap {
autotype Real
def 'max Id'*'num crs' }

aux capacity {
autotype Real
def PML- sTES+'UGPT IP yr' }
```

```
level cm E {
autotype Real
init 1
inflow { autodef 'ins E'} }

aux cm E s {
autotype Real
autodim 0..20
def FOR(k=year|SAMPLEIF(TIMEIS(52*k),'cm E',0)) }

level cm hr I {
autotype Real
init 'hr I 0 '
inflow { autodef 'I hr wk'} }

level cm upg {
autotype
Real
init 2
inflow { autodef 'ins upg'} }

aux cm upg s {
autotype Real
autodim 0..20
def FOR(k=year|SAMPLEIF(TIMEIS(52*k),'cm upg',0)) }

aux cm yrE {
autotype
Real
autodim 0..20
def FOR(k=year|SAMPLEIF(TIMEIS(52*k),'quant E',0)) }

aux cm yrN {
autotype Real
autodim 0..20
def FOR(k=year|SAMPLEIF(TIMEIS(52*k),'quant N',0)) }

const crs dur {
autotype Real
init 3 }

aux d quant E {
autotype Real
def DELAYPPL('quant E',TIMESTEP,0) zeroorder immediate }
```

```
aux d quant N {
autotype
Real
def DELAYPPL('quant N',TIMESTEP,0) zeroorder immediate }

aux d to O {
autotype
Real
def DELAYPPL('to OTU',52,0) zeroorder immediate }

const E Hr mo {
autotype Real
init 5 }

aux E yr {
autotype Real
autodim 0..20
def FOR(k=year| IF(TIME>=52*,IF(k>0,'cm E s'[INDEX(k)]-
'cm E s'[INDEX(k-1)],'cm E s'[INDEX(k)]//When k=0//),0)) }

level Ex {
autotype Real
init 'Ex 0'
outflow { autodef 'Ist Ex' }
inflow {autodef 'fr NOP'}
outflow { autodef 'to NOP' }
inflow { autodef 'to Ex'}
outflow { autodef 'to UGPT' } }

const Ex 0 {
autotype Real
init 8 }

aux f NOP {
autotype Real
def IF(capacity>0 AND NOP-'N min'>0,MIN(capacity,NOP-'N min'),
0) }

aux fr NOP {
autotype Real
def IF(TIMECYCLE(3,52),'f NOP',0) zeroorder immediate }

aux fr OTU {
autotype Real def DELAYPPL('to OTU','crs dur',0) zeroorder
```

```
immediate }

aux hr dist {
autotype Real
autodim 0..20
def FOR(k='OTU hist'|'req hr'*k/(2*'num crs')) }

aux hr I 0 {
autotype Real
def ARRSUM('hr dist') }

aux hrs aqd {
autotype Real
autodim 0..60
def FOR(k='OTU crs'|IF(TIME>='arr tm'[k],IF('cm r I'-
nrm[k]<'req hr','cm hr I'-nrm[k],'req hr'),0)) }

aux hrs aqd hst {
autotype Real
autodim 1..6
def FOR(k='OTU hist'|IF('hr dist'[k]+'cm hr I'-'hr I 0'
<='req hr','hrdist'[k]+'cm hr I'-'hr I 0','req hr')) }

const I hr lim {
autotype Real
init 13 }

aux I hr wk {
autotype Real
def MIN(SR*(3/13)*'E hr mo'*(Ex DIVZ0 In) ,(3/13)*'I hr lim') }

level In {
autotype Real
init 'In 0'
outflow { autodef 'to Ex' }
inflow { autodef 'fr OTU'} }

aux In 0 {
autotype Real
def 'max Id'*'num crs'*2 }

aux ins E {
autotype Real
def Ex*(3/13)*'E hr mo' }
```

```
aux ins upg {
autotype Real
def IF('YFR sw',(SR+1),1/SR)*In*'I hr wk' }

aux load {
autotype Real
def IF(MIN('PAT OTU','max Id')>='min Id',MIN('PAT OTU','max Id'),
0) }

aux Ist EX {
autotype Real
def 'quand E'-'d quant E' zeroorder immediate }

aux Ist yrE {
autotype Real
def 'quant N'-'d quant N' zeroorder immediate }

aux Ist yrE {
autotype Real
autodim 0..20
def FOR(k=year|IF(TIME>=52*k,IF(k>FIRST(k),'cm yrE'[k]-
'cm yrE'[INDEX(k-1)],'cm yrE'[k]//When k=FIRST(k)//),0)) }

aux Ist yrN {
autotype Real
autodim 0..20
def FOR(k=year|IF(TIME>=52*k,IF(k>FIRST(k), 'cm yrN'[k]-
'cm yrN'[INDEX(k-1)],'cm yrN'[k]//When k=FIRST(k)//),0)) }

const max Id {
autotype Real
init 21 }

const min Id {
autotype Real
init 34 }

aux min to NOP {
autotype Real
def IF(NOP<'N min','N min'-NOP,0) }

const N min {
autotype Real
init 55 }
```

```
level NOP {
autotype Real
init 'NOP 0'
outflow { autodef 'Ist NOP' }
outflow { autodef 'fr NOP' }
inflow { autodef 'to NOP'} }

const NOP 0 {
autotype Real
init 89 }

aux nrm {
autotype Real
autodim 0..60
def FOR(k='OTU crs'| SAMPLEIF(TIMEIS('arr tm'[k]),'cm hr I',0)) }

const num crs {
autotype Real
init 144 }

level OTU {
autotype Real
init 233
outflow { autodef 'fr OTU' }
inflow { autodef 'to OTU'} }

aux OTU cls {
autotype Real
autodim 0..60
def FOR(k='OTU crs'|SAMPLEIF(TIMEIS(k*52/'num crs'),load,0)) }

level PAT OTU {
autotype Real
init 'PAT OTU 0'
outflow { autodef 'to OTU' }
inflow { autodef 'to PAT OTU'} }

const PAT OTU 0 { autotype Real init 377 }

const pct at N {
autotype Real
init 610 }

aux pct lst yr {
autotype Real a
```

```
utodim 0..20
def FOR(k=year|IF(k>FIRST(k),('lst yrE'[k]+'lst yrN'[k])/
's TES yr'[INDEX(k-1)]*100,0// When k=FIRST(k))//)))) }

const PML {
autotype Real
init 987 }

aux quant E {
autotype Real
def ROUND(INTEGRATE('at fact E'),1,0) zeroorder immediate }

aux quant N {
autotype Real
def ROUND(INTEGRATE('at fact N'),1,0) zeroorder immediate }

const req hr {
autotype Real
init 1597 }

const SR {
autotype Real
init 2584 }

aux sTES { autotype Real
def Ex+In }

aux sTES yr {
autotype Real
autodim 0..20
def FOR(k=year| SAMPLEIF(TIMEIS(k*52),sTES+NOP,0)) }

aux stop {
autotype Logical
def STOPRUNIF(Ex<=0) }

aux t Ex hst {
autotype Real
autodim 1..6
def FOR(k='OTU hist'|DERIVN(IF('hrs aqd hst'[k]>='req hr',1,
0))) }

aux t NOP {
autotype Logical
def MAX(IF(capacity<0,-capacity,0),'min to NOP') }
```

```
level thru O {
autotype
Real init 4181
inflow { autodef 'to OTU' }
inflow { autodef 'd to O'} }

aux thru O yr {
autotype Real
autodim 0..20
def FOR(k=year|SAMPLEIF(TIMEIS(52*k),'thru O',0)) }

aux to Ex {
autotype Real
def ('OTU cls' #*'t Ex pls')+'max Id'*ARRSUM('t Ex hst') }

aux to NOP {
autotype Real
def IF(TIMECYCLE(2,52),'t NOP',0) zeroorder immediate }

aux to OTU {
autotype Real
def IF(TIMECYCLE(1,52/'num crs'),load,0) zeroorder immediate }

aux to PAT OTU {
autotype Real
def IF(TIMECYCLE(0,52/'AFT crs per yr'),ROUND(cap/'AFT crs per yr'),0)
zeroorder immediate }

aux to UGPT {
autotype Real
def IF(TIMECYCLE(2,52),'UGPT IP yr',0) zeroorder immediate }

const UGPT IP yr {
autotype Real
init 6765 }

aux upg yr {
autotype Real
autodim 0..20
def FOR(k=year| IF(TIME>=52*k,IF(k>0,'cm upg s'[INDEX(k)]-
'cm upg s'[INDEX(k-1)], 'cm upgs'[INDEX(k)] //Whenk=0),0)) }

aux upgr tm {
autotype Real
```

```
autodim 0..20
def FOR(k='OTU crs'|SAMPLEIF('t Ex pls'[k]<>0,TIME-
'arr tm'[k],0)) }

const YFR sw{
autotype Logical
init TRUE }

range OTU crs {
def 0..60 }

range OTU hist {
def 1..6 }

range year {
def 0..20 } }
```

A.2 Course Module Equations

```
mainmodel Course Module {

aux Bern {
type
Real
dim binomial
def IF(generator<'pct fail',1,0) }

aux bin fail {
type Real
def Bern#*'n loaded' }

level Course {
type Real
init 10946
outflow { autodef 'fail crs'}
inflow { autodef 'to crs'}
outflow { autodef 'pass crs'} }

const crs dur {
type Real
init 17711 }

aux crs load {
type Real
def IF(MIN(PAT,'max load')>='min load', MIN(PAT,'max load'),0) }
```

```
aux d fail crs {
type Real
def DELAYPPL('fail crs',52,0) zeroorder immediate }

aux d to crs {
type
Real def DELAYPPL('to crs','crs dur',0) zeroorder immediate }

level fail {
type Real
init 28657
inflow { autodef 'fail crs'}
outflow { autodef 'd fail crs'} }

aux fail crs {
type Real
def IF('rn swith','bin fail',ROUND('pct fail'/100*end,1))
zeroorder immediate }

aux fail yr {
type Real
dim year
def FOR(k=year|SAMPLEIF(TIMEIS(k*52)fail,0)) }

aux generator { type Real dim binomial def FOR(i=binomial|RANDOM(0,100)) }

const max load {
type
Real
init 46368 }

const min load {
type Real
init 75025 }

aux n loaded {
type Real
dim binomial
def FOR(i=binomial|IF(i<=end AND i>=1,1,0)) }

const num crs {
type Real
init 121393 }
```

```
aux num to PAT {
type Real
def INTEGER(RANDOM(5,50)) }

aux pass crs {
type Real
def end-'fail crs' zeroorder immediate }

level PAT {
type Real
init 'PAT initial'
outflow { autodef 'to crs'}
inflow { autodef 'to PAT'} }

const PAT initial {
type Real
init 196418 }

const pct fail {
type Real
init 317811 }

const m switch {
type Logical
init TRUE }

level thru crs {
type Real
init 514229
outflow { autodef 'd to crs'}
inflow { autodef 'to crs'} }

aux thru crs yr{ type Real dim year def FOR(k=year|SAMPLEIF(TIMEIS(k*52),'thru
crs',0)) }

aux to crs {
type Real
def IF(TIMECYCLE(1,52/'num crs'),'crs load',0)
zeroorder immediate }

aux to PAT {
type
Real
def IF(TIMECYCLE(1,52/5),'num to PAT',0) zeroorder immediate }
```

```
range binomial {
def 1..20 }

range year { def 0..20 } }
```

Annex B
YFR Calculations

In this annex we present a derivation of the equations that form the basis for the YFR calculations performed within the mentor module. We also demonstrate how these equations could be used to build a "YFR constrained" simulation.

Let $R_e(t)$ and $R_i(t)$ denote the monthly per pilot flying rates for experienced and inexperienced pilots respectively and let ν be the expected ratio of inexperienced pilots to experienced pilots on a proficiency mission. In the k^{th} year, the experienced pilots accumulate a total of

$$\mathscr{H}_e = \int_{12k}^{12(k+1)} E(t) R_e(t)\, dt$$

flying hours and, during the same year, the inexperienced pilots accumulate a total of

$$\mathscr{H}_i = \int_{12k}^{12(k+1)} I(t) R_i(t)\, dt.$$

hours. The pool of flying hours that is available to upgrade pilots is $\mathscr{H}_e$. The pool of hours $\nu\mathscr{H}_e - \mathscr{H}_i$ could be thought of as surplus proficiency generation capacity.

The hours given by $\mathscr{H}_e$ and $\mathscr{H}_i$ do not necessarily represent YFR directly. This is because these are the hours that are logged by pilots and hence, need not be in one–to–one correspondence with the hours accumulated by a given aircraft. Indeed, when we consider the communities other than the fighter community, we must account for the fact that there will be two or three pilots in the aircraft. Accordingly, the YFR used to upgrade pilots in the k^{th} year can be estimated by

$$\begin{aligned} \text{YFR}_\text{p} &= \frac{1}{\nu}\mathscr{H}_i \\ &= \frac{1}{\nu}\int_{12k}^{12(k+1)} I(t) R_i(t)\, dt. \end{aligned}$$

Additional YFR arises from unused hours in the pool generated by experienced pilots. Here we assume that there are two experienced pilots in the cockpit and hence

$$\begin{aligned} \text{YFR}_\text{a} &= \frac{1}{2}\left\{\mathscr{H}_e - \frac{1}{\nu}\mathscr{H}_i\right\} \\ &= \frac{1}{2}\int_{12k}^{12(k+1)} E(t) R_e(t) - \frac{1}{\nu} I(t) R_i(t)\, dt. \end{aligned}$$

It follows that the total YFR requirement for the k^{th} year is given by

$$\begin{aligned} \text{YFR} &= \frac{1}{2}\left\{\mathscr{H}_e + \frac{1}{\nu}\mathscr{H}_i\right\} \qquad \text{(B.3)} \\ &= \frac{1}{2}\int_{12k}^{12(k+1)} \left\{E(t) + \min\left(E(t), \frac{\rho}{\nu} I(t)\right)\right\} R_e(t)\, dt, \end{aligned}$$

where we have used Equation (23).

As inexperienced and experienced operational fighter pilots fly in separate aircraft, the case for the CF-188 community is an exception to the preceding development. In this case, the YFR that is used for proficiency missions is given by

$$\begin{aligned}\mathrm{YFR_p} &= \left(1+\frac{1}{\nu}\right)\mathscr{H}_i \\ &= \left(1+\frac{1}{\nu}\right)\int_{12k}^{12(k+1)} I(t)R_i(t)\,dt\end{aligned}$$

and the additional YFR is given by

$$\begin{aligned}\mathrm{YFR_a} &= \mathscr{H}_e - \frac{1}{\nu}\mathscr{H}_i \\ &= \int_{12k}^{12(k+1)} E(t)R_e(t) - \frac{1}{\nu}I(t)R_i(t)\,dt\end{aligned}$$

If we use Equation (23), then $\mathrm{YFR} = \mathrm{YFR_p} + \mathrm{YFR_a}$ can be written as

$$\mathrm{YFR} = \int_{12k}^{12(k+1)} \left\{E(t) + \min\left(E(t), \frac{\rho}{\nu}I(t)\right)\right\} R_e(t)\,dt. \tag{B.7}$$

Usually $\nu = \rho = 1$ for the CF-188 community. However, the YFR formula given by Equation (B.7) allows for potential generalisations.

Equations (B.7) and (B.3) can be used to constrain the monthly flying rate $R_e(t)$ so that a target YFR[16] is approximately achieved. As an example, let us consider Equation (B.7) with $\nu = \rho = 1$ and suppose that the YFR is distributed uniformly throughout the year. In the k^{th} month, we want to set the flying rate $R_e(t)$ so that the constraint

$$\frac{\mathrm{YFR}}{12} = \int_{k}^{k+1} \{E(t) + \min(E(t), I(t))\}\, R_e(t)\,dt$$

is satisfied. Let us assume that $R_e(t)$ is constant over each month. In this case, we have

$$R_e(t) \equiv R_e[k] = \frac{\mathrm{YFR}}{12\,\langle E(t) + \min(E(t), I(t)\rangle_k},$$

for $k \le t < k+1$, where we have used the notation $\langle f(t)\rangle_k$ to denote the time-average of the function f over the k^{th} month.

If $N(t) = E(t) + I(t)$, then $E = \beta N$ and $I = (1-\beta)N$. Now, if E and I do not vary substantially over the k^{th} month, then neither do N and β. Consequently, we have

$$R_e[k] \approx \frac{R[k]}{\beta(k) + \min(\beta(k), 1-\beta(k))}, \tag{B.8}$$

where $R[k] = \mathrm{YFR}/(12N(k))$. The quantity $R[k]$ can be regarded as a target average per pilot flying rate for the k^{th} month.

[16]The reader should note that the symbol YFR refers to the portion of yearly flying hours that has been assigned to the operational squadrons. As such, YFR does not include the hours flown at the OTU etc.

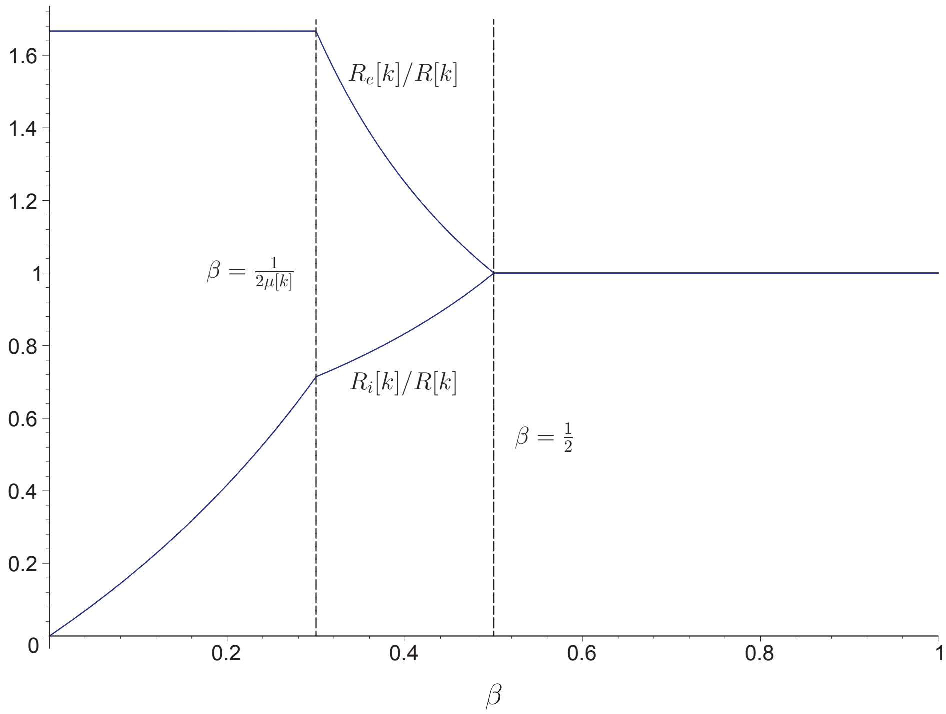

Figure B.1: *Constrained Flying Rates*

As with Equation (21b), Approximation (B.8) is not realistic when experience levels are low. This is because

$$\frac{1}{\beta(k)+\min(\beta(k),1-\beta(k))} = \frac{1}{2\beta(k)}$$

whenever $\beta \leq 1/2$ and hence $R_e[k] \to \infty$ as $\beta \to 0^+$. In view of inequality (15b), we replace (B.8) with the approximation

$$R_e[k] \approx \min\left(R_e^{\max}, \frac{R[k]}{\beta(k)+\min(\beta(k),1-\beta(k))}\right) \tag{B.9}$$

and, if we use Equation (23), then the corresponding rate for inexperienced pilots can be written as

$$R_i[k] \approx \min\left(\frac{\beta(k)}{1-\beta(k)}, 1\right) R_e[k]. \tag{B.10}$$

We refer the reader to Figure B.1, where the relative rates $R_e[k]/R[k]$ and $R_i[k]/R[k]$ are plotted as functions of the experience level β.

In this particular case we have set $R[k] = 15$ and $R_e^{\max} = 25$, which means that experienced pilots can fly approximately 67 percent more than the target rate when required. There are three distinct regimes in Figure B.1. Let $\mu[k] = R[k]/R_e^{\max}$, then the regions are:

1. $\beta \geq 1/2$: Here both of the relative rates are one, which means that all pilots fly at the target rate $R[k]$. Consequently, the YFR goal for the k^{th} month can be achieved. In addition, $R_i[k]$ does not vary with the experience level and this means that upgrade times will remain constant in this regime.

2. $1/(2\mu[k]) \leq \beta < 1/2$: As the experience levels decrease from $1/2$ to $1/(2\mu[k])$, the relative rate $R_e[k]/R[k]$increases from one to a maximum of $\mu[k]$. On the other hand, the relative rate $R_i[k]/R[k]$ decreases from one to $\mu[k]/(2\mu[k]-1)$. This regime represents the flying community's ability to "surge" when experience levels are low. Monthly YFR allotments can still be flown. However, the low experience levels suggest that the community may not be able to sustain operations indefinitely.

3. $\beta < 1/(2\mu[k])$: If the experience level falls below $\beta = 1/(2\mu[k])$, then the rate $R_e[k]$ remains fixed at $R[k]/\mu[k]$, while $R_i[k]$ continues to decrease to zero when $\beta = 0$. This indicates that when the experience levels drop below $1/(2\mu[k])$ the community will not be able to fly enough to meet a monthly YFR target.

List of Acronyms

1 Cdn Air Div	1 Canadian Air Division
CAF	Canadian Air Force
CAS	Chief of the Air Staff
CF	Canadian Forces
CORA	Centre for Operational Research and Analysis
DDE	Delay Differential Equation
DND	Department of National Defence
DRDC	Defence Research and Development Canada
IC	Initial Condition
IP	Instructor Pilot
MOC	Military Occupation
NOP	Non-Operational
ODE	Ordinary Differential Equation
ORAD	Operational Research and Analysis Directorate
OTU	Operational Training Unit
PARSim	Production Absorption and Retention Simulation
PAT	Pilot Awaiting Training
PTA	Pilot Terminable Allowance
SD	System Dynamics
UGPT	Undergraduate Pilot Training
USAF	United States Air Force
YFR	Yearly Flying Rate
YOS	Year of Service

List of Symbols

$\mathscr{A}$	Absorption capacity (pilots/yr)
$a(t)$	Flow of pilots from OTU to operational squadrons (pilots/month)
A_e	Yearly percent attrition rate of experienced pilots
A_s	Yearly percent attrition rate of pilots in NOP positions
$\beta(t)$	Experience level at time t
$C(t)$	Number of students loaded on course $\mathscr{C}$ at time t
D	Duration of course $\mathscr{C}$ (months)
$\hat{\delta}(t)$	A smooth, nonnegative pulse, of unit area
$E(t)$	Number of experienced pilots at time t
$\gamma(t)$	Ratio of experienced to inexperienced pilots at time t
$\Gamma_{\text{ops}}(t)$	Operational capacity at time t (pilots)
$\Gamma_{\text{nop}}(t)$	Non-operational capacity at time t (pilots)
H_s	Operational hours required to upgrade
$I(t)$	Number of inexperienced pilots at time t
$\hat{\Lambda}(t)$	$\hat{\Lambda}(t) = \int \hat{\delta}(t)dt$
L_{max}	Maximum loading of course $\mathscr{C}$ (students)
L_{min}	Minimum loading of course $\mathscr{C}$ (students)
M	Number of courses per year
N_f	Number of flying positions
$N(t)$	Operational strength at time t (pilots)
ν	Sortie ratio
$n_i[k]$	Number of students on the k^{th} type course
n_u	Net number of pilots sent, from the operational squadrons, to act as UGPT IPs per year
$P(t)$	Number of students in the PAT pool preceding $\mathscr{C}$ at time t
$R_e(t)$	Flying rate of experience pilots at time t (hours/month/pilot)
R_i^{max}	Maximum flying rate of inexperience pilots (hours/month/pilot)
$R_i(t)$	Flying rate of inexperience pilots at time t (hours/month/pilot)
$\sigma(t)$	Flow of pilot to and from NOP positions at time t (pilot/month)
S_{min}	Critical number of NOP positions
$S(t)$	Number of pilots in NOP positions at time t
T_e	Time to experience (months)
$\upsilon(t)$	Net flow of pilot to UGPT IP positions at time t (pilots/month)

Acknowledgements

The author would like to thank the referees for the time they took to do their reviews. Particular thanks go to a specific reviewer for supportive comments. The author would also like to thank Mr. Charles Hunter for his continued help in bringing this dated report to publication.

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.) DRDC CORA 1 Cdn Air Div/CANR HQ, Box 17000 Stn. Forces, Winnipeg MB R3J 3Y5	2. SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable). UNCLASSIFIED (NON-CONTROLLED GOODS) DMC : A Review GCEC June 2010

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title).

Modelling the Production and Absorption of Pilots

4. AUTHORS
(Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)

Corbett, Dr. Norman C.

5. DATE OF PUBLICATION (month and year of publication of document) March 2013	6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc). 82	6b. NO. OF REFS (total cited in document) 27

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered).

Technical Report

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include address).

DRDC CORA
1 Cdn Air Div/CANR HQ, Box 17000 Stn. Forces, Winnipeg MB R3J 3Y5

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Specify whether project or grant). 64529	9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written).
10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique.) DRDC CORA TR 2013-023	10b. OTHER DOCUMENT NOs. (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Defence departments and defence contractors; further distribution only as approved
( ) Defence departments and Canadian defence contractors; further distribution only as approved
( ) Government departments and agencies; further distribution only as approved
( ) Defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution beyond the audience specified in (11) is possible, a wider announcement audience may be selected).

UNLIMITED

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

This report is a record of the technical aspects pertaining to the development of the Pilot Production Absorption Retention Simulation (PARSim). The simulation was produced at the Operational Research and Analysis Directorate at 1 Canadian Air Division (1 Cdn Air Div) Headquarters under the sponsorship of Commander 1 Cdn Air Div, A4/A1 and A1 Trg. The report is intended to be a guide for analysts: providing a basis for the maintenance of PARSim and the development of similar simulations.

PARSim is a model of the "flow" of pilots from formal Undergraduate Pilot Training (UGPT) to the various major operational communities and is designed to give decision makers the ability to gauge the effect of factors, like attrition, production and flying rates on the overall health of the pilot system. The simulation is comprised of two main modules: a course module, which is used to build UGPT portion of the simulation and the mentor module, which is a generic model of the major operational communities. An important feature of the mentor module is that it simulates the *transfer of experience* to new pilots, by dynamically adjusting the rate at which this transfer occurs according to health of the community and the availability of resources.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title).

Differential Equations
Experience Levels
Flow
Force Generation
PARSim
Pilot
Simulation
System
System Dynamics
Undergraduate Pilot Training

DEFENCE DÉFENSE

DRDC CORA

www.drdc-rddc.gc.ca