Master of Science Thesis in Vehicular Systems
Department of Electrical Engineering, Linköping University, 2016

Supervision of the air loop in the Columbus Module

of the International Space Station

Jasper Germeys

LINKÖPING UNIVERSITY

Master of Science Thesis in Vehicular Systems

Supervision of the air loop in the Columbus Module of the International Space Station

Jasper Germeys

LiTH-ISY-EX--16/5014--SE

Supervisor: **Daniel Jung**
ISY, Linköpings University

Examiner: **Erik Frisk**
ISY, Linköpings University

Department of Vehicular Systems
Department of Electrical Engineering
Linköping University
SE-581 83 Linköping, Sweden

Abstract

Failure detection and isolation (FDI) is essential for reliable operations of complex autonomous systems or other systems where continuous observation or maintenance thereof is either very costly or for any other reason not easily accessible.

Beneficial for the model based FDI is that there is no need for fault data to detect and isolate a fault in contrary to design by data clustering. However, it is limited by the accuracy and complexity of the model used. As models grow more complex, or have multiple interconnections, problems with the traditional methods for FDI emerge.

The main objective of this thesis is to utilise the automated methodology presented in [Svärd, 2012] to create a model based FDI system for the Columbus air loop. A small but crucial part of the life support on board the European space laboratory Columbus.

The process of creating a model based FDI, from creation of the model equations, validation thereof to the design of residuals, test quantities and evaluation logic is handled in this work. Although the latter parts only briefly which leaves room for future work.

This work indicate that the methodology presented is capable to create quite decent model based FDI systems even with poor sensor placement and limited information of the actual design.

Carl Svärd. *Methods for Automated Design of Fault Detection and Isolation Systems with Automotive Applications*. PhD thesis, Linköping University, Vehicular Systems, The Institute of Technology, 2012

Acknowledgements

This thesis has taken some time to finish and has defined quite a good portion of my being and by that does it not only mark the end of a great time but also a promise of future developments. This is a huge milestone for me and thus could these acknowledgements never fully contain the depth of my gratitudes.

First and foremost must I acknowledge Erik Frisk and the whole Vehicular Systems department at Linköping university and Airbus DS for allowing me to work upon such an interesting system, even though it has given me plenty of grief at times. Daniel Jung for always keep trusting in me but also Enrico Noack and Mikael Persson for the time during the preparation work. Not to forget my opponent Gustav Romeling.

Special thanks to Bengt Göransson, 千葉大奈, 李森 and 仇隽挺 for helping me to get back on track at a time where motivation was a great shortcoming, for this will you not be forgotten.

Alla mina vänner under studietiden, speciellt Björn Holm, Dan Adolfsson, Fredrik Henriksen och Mikael Göransson samt familjerna LeMoine och Lockwood. Men även Stefan Persson som lyckats stå ut med mig under så lång tid.

Carl Forden, Jonas Frossmed och Fritiof Olsson för simpel men pålitlig vänskap. Anna för alla dessa underbara konversationer som skapar glädje i vardagen. Familjen Asp för ett evigt familjeband samt David, Matte och Peter i Nässjö.

Joakim Råberg då även små människor kan göra stordåd.

The supporting friends in the Mahjong community like David Clarke, Matthias Köhler, Senechal and Gemma Sakamoto among others, you know who you are.

Folket på Gata/Park i Nässjö kommun för er hjälp att komma ut ur arbetslöshetens lömska klor men även för er förståelse när jag valde att lämna er för högre utbildning, det tog sin tid, men nu är det klart och utan er hade detta inte skett.

Familjen Borg för deras outgrundliga vänlighet och stöd. Familjerna Hjelm, Kvist, Nylander samt Pieters i Kärrgruvan, Norberg samt familjen Högström för allt ni betytt för mig.

Solbritt Sundvall för sitt gedigna engagemang att försöka reda ut en karl av mig.

En natuurlijk mijn familie, die altijd voor mijn zaak blijft staan.
Ik hou van jullie allemaal.

Linköping, December 22, 2016
Jasper RK Germeys

Contents

Abbreviations

AFS air mass flow sensor. 12, 16, 18, 22, 23, 25

Airbus DS Airbus Defence and Space. 1

Airbus SE Airbus Group SE. 1

CFA cabin fan. 8, 15, 19, 25, 32

CHX condensate heat exchanger. 8, 22–24

CHXFA condensate heat exchanger fan. 12

COL-CC COLumbus Control Centre. 2

COL-DMS COLumbus Data Management System. 2

CTCU cabin temperature control unit. 8, 13, 20

CUSUM cumulative sum. 5, 35

CWSA condensate water separator assembly. 13, 16

DAG directed acyclic graph. 30, 31

EADS European Aeronautic Defence and Space company NV. 1

EU electronic unit. 10

FDI failure detection and isolation. 3–5, 45–47

GLR generalized likelihood ratio. 5

HEPA high-efficiency particulate arrestance. 1, 8, 40

HS humidity sensor. 13, 19, 20

IMV inter module ventilation. v, vi, 8, 10

IRFA inter module ventilation (IMV) return fan. 8, 10, 15, 19, 25

ISFA IMV supply fan. 8, 15, 19, 22, 23, 25

ISS International Space Station. 1

KH total dynamic head. 11, 24

LCOS liquid carry over sensor. 13

MAE% mean absolute error percentage. 16, 18–20, 23–25

MSE mean square error. 16, 18–20, 23–25

MSO minimal structurally overdetermined set. 5, 27–30, 46

NaN not-a-number. 32, 35

SCC strongly connected components. 27

TPS air pressure sensor. 13, 17

Units

I	Current $[A]$
N	Revolutions per minute $[min^{-1}]$
Q	Volume flow $[m^3/s]$ most of the time in $[m^3/h]$
R	Gas constant $[J/kgK]$
T	Temperature $[^oC]$
W	Watt $[W]$
ϕ	Relative Humidity [-]
ρ	Density $[kg/m^3]$
p	Pressure $[kPa]$
q	Mass flow $[kg/s]$ exclusively used as $[kg/h]$ in this document

1

Introduction

The European space laboratory Columbus is designed and built by European Aeronautic Defence and Space company NV (EADS), now Airbus Group SE (Airbus SE), affiliate Astrium, now Airbus Defence and Space (Airbus DS) [Airbus Defense and Space, 2013a], and is a part of the International Space Station (ISS). It is equipped with a range of experimental facilities and basic life support for up to three astronauts. Together with the microgravity environment, Columbus has enabled numeral extraordinary experiments previously not possible in many scientific fields such as physics, material science, biology, medicine, and human physiology [Airbus Defense and Space, 2013b,c].

1.1 The Columbus air loop

The Columbus air loop main function is to provide a forced air flow to avoid dead air pockets which is safety critical for the crew. The forced air flow also enables fire detection and removes heat from air cooled equipment. As the Columbus Module has no O_2 or CO_2 control, the inter module ventilation has to provide fresh air to support life on board.

The air loop consists of four fans, three with check valves, one without, a high-efficiency particulate arrestance (HEPA) filter and a condensate heat exchanger. A simple overview of the system is shown in Figure 1.1. The system is designed to be operating in 8 stable modes and additionally 43 interim modes intended for air loop reconfiguration.

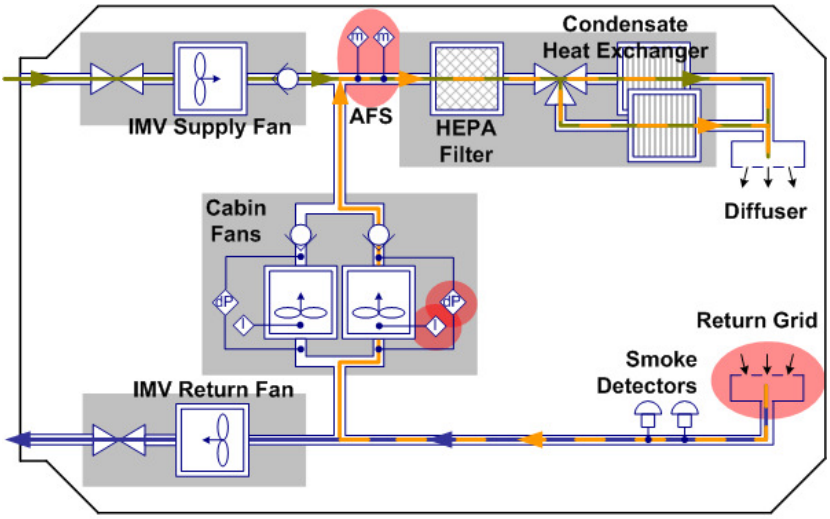

Figure 1.1: *A schematic overview of the Columbus air loop, [Source: Airbus DS, 2013]*

1.2 Columbus failure management system

The failure management system on Columbus consist of multiple parts; First is the on board COLumbus Data Management System (COL-DMS) that collects and logs all measured signals. Being fully automatic it can quickly initiate required responses but due to sparse resources, the diagnosis part of COL-DMS is limited to only detect time critical failures like fire outbreaks and other environmental hazards.

The second part is the COLumbus Control Centre (COL-CC) which is the primary ground based unit. It receives data from COL-DMS in near realtime through a downlink and performs data analysis, both manually and automatically. The failure management part of COL-CC has a similar purpose as the onboard unit and focuses primarily on detection of critical failures and short term responses. This is supported by an Engineering Support Centre that does offline analyses of selected data to find and direct long term corrective actions for COL-CC and an Assembly, Integration and Test facility for testing and development purposes.

An onboard monitoring experiment, the ERNO BOX, has been added as a third part to the failure detection arsenal. This system runs in parallel to COL-DMS, monitoring the same data but with no means for interaction. Instead the responses are submitted to COL-CC for comparison with the decisions made by COL-DMS [Noack et al., 2012].

There has been some recent developments in the COL-CC infrastructure too improve the failure management efficiency [Sabath et al., 2012, 2014].

1.3 Background

Since launch in February 2008, Columbus has encountered multiple failures and abnormalities. The failures have either been minor or been recovered before any serious damage to the system has occurred. It has emerged that the original failure detection system has weak robustness and poor fault sensitivity [Noack et al., 2010].

Upon request from Astrium, European universities have been engaged in development and improvement of the diagnosis system on Columbus. BTU Cottbus is using a data mining approach and Linköping University by means of model based failure detection. [Noack et al., 2011, 2012, Noack and Schmitt, 2012]

Beneficial for the model based failure detection and isolation (FDI) is that there is no need for fault data to detect and isolate a fault in contrary to data clustering. However, it is limited by the accuracy and complexity of the model used. As models grow more complex, or have multiple interconnections, problems with the traditional methods for FDI emerge. In [Svärd, 2012] an automated methodology is presented for design of an FDI system for complex systems operating in multiple modes.

This thesis focuses on the design of a model based FDI system for the Columbus air loop. A small but critical part of the Columbus life support.

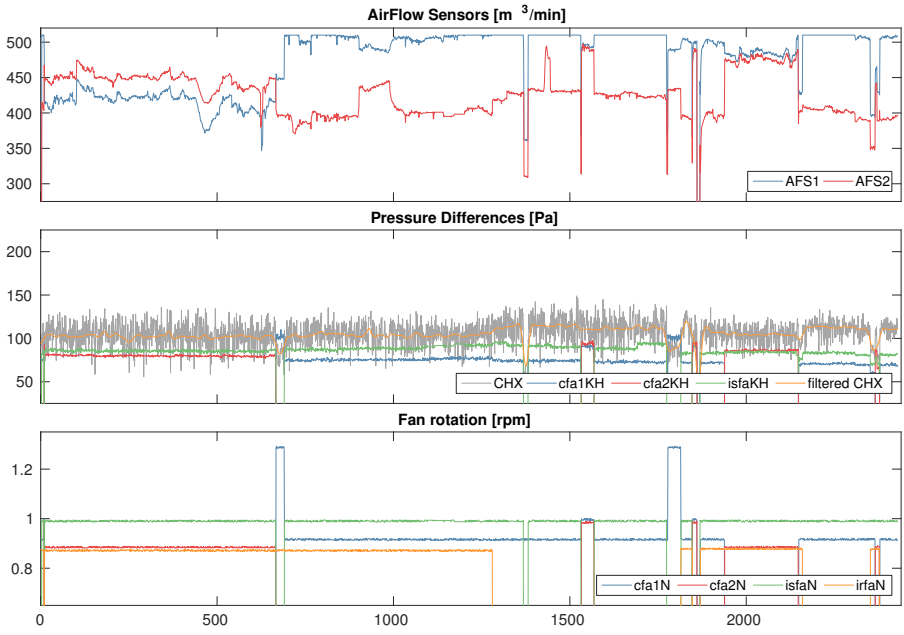

Figure 1.2: *Subsampled unproccessed data*

1.4 Model and data

A model of the Columbus air loop was proposed during the summer of 2012 by Jasper Germeys and Mikael Persson. This model, which is based on physical relations and measured data from the time span of 2008 to 2010 provided by Astrium, has been used as basis for the model in this thesis. A second batch of data containing a more detailed selection of all the more interesting parts from the initial set was later received.

In this thesis only the second data set and an additional third data set, where the fans' rotation are increasing in steps, are used. This third data set has been used as training data for the parameterisation of the model. A selection of the signals in the second and third data sets are displayed in Figure 1.2 and Figure 3.1 respectively. In Tables 2.1 and 2.2 the sensors contained in the data sets are listed. Additionally, Figure 1.1 and Figures 2.1-2.4 are taken from these documentations.

1.5 Problem formulation

The purpose of this thesis is to develop a diagnosis system that will not only detect but isolate the failures included in the model. The system suffers from noisy sensors and fault propagation, hence will simple means for fault isolation not suffice in aspect to minimise the risk of false alarms. This includes detection of the fans' current work mode and determine actual mass flow as it is this central state that has poor observability. This will pose a problem in unknown modes and when clogging or leaks occur as those faults are highly linked.

The FDI system should;

- Be designed based on a model of the air loop.
- Be able to detect which mode the fans operate in.
- Detect and isolate single faults such as fan loss, leaks, wear, (partial) clogging, and sensor loss.

1.6 Related research

A model-based design of a diagnosis system can be divided into three phases; A modelling phase, a residual generation phase and finally an evaluation phase where the rules for fault detection and isolation are set. These phases entwine and will be done iteratively.

In the modelling phase the full model equation set will be extended with the desired detectable faults and modes. Structural fault detectability and isolability will be analysed by using the Dulmage-Mendelsohn decomposition [Frisk et al., 2012]. The parameters of the equations will then be estimated from sample data with no faults.

The residual generation phase consist of extracting the minimal structurally overdetermined set (MSO) to generate the residuals [Krysander et al., 2008]. As it is unknown how many alternative sets the model allows, initially a greedy selection algorithm described in [Svärd et al., 2013] will be applied.

The evaluation phase consist of creating the rules for detection and isolation. There are many methods available for detection, where constant thresholds being the simplest, but also more robust methods as cumulative sum (CUSUM), likelihood functions, adaptive thresholds [Sneider and Frank, 1996], or combinations thereof [Meinguet et al., 2012] have proven effective. As the system requires robustness and will be operating in multiple modes, a relaxed generalized likelihood ratio (GLR), as proposed in [Svärd et al., 2014], was to be considered.

1.7 Expected results

The expected results of this work is an FDI system that detects and isolates faults in the air loop system. The system should be able to work in real time and detect faults within reasonable time. It is also expected that the system could measure wear on the fans and partial clogging.

An evaluation shall be done if the methodology presented in [Svärd, 2012] is a working methodology for this system or not. Also an analysis if the solution can be improved by using Kullback-Leiber divergence for fault isolation as suggested by [Eriksson et al., 2011].

Additionally an evaluation on how well the FDI system works in comparison to the uncertainties in the model and if simpler methods could suffice.

1.8 Outline of the report

The plan for the report outline is as following;
Chapter 1 introduces the problem and the system, the chosen methodology and related research as well as the expected results and this outline of the rest of the report.

In Chapter 2 the details on the model and the equations sets are listed. The work of setting the model parameters are presented in Chapter 3 together with an estimation of the expected noise while in no fault node. Chapters 2 and 3 will together form the modelling phase of the work.

Chapter 4 displays the full model including modes and faults and continues to the residual generation and what sets the methodology generates. These are then evaluated and the result are presented. How well the design works against collected data is presented in Chapter 5.

Chapter 6 contains the evaluation of the work, made together with ideas for potential future work.

2

Model description

The Columbus air loop consist of a small set of flows with some sensors and fans. This chapter will describe these compontents and the models used to represent them. A 3D model of the air loop has been provided by Astrium which is shown in Figure 2.1. Even though the 3D model is trusted to be an accurate representation of the air loop, no detailed data has been gathered from this figure.

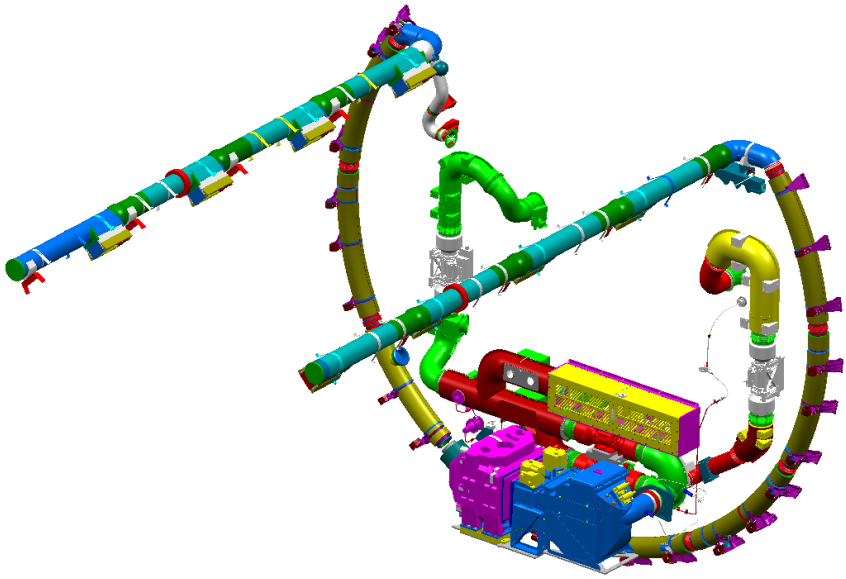

Figure 2.1: *A 3D model of the Columbus air loop, [Source: Airbus DS, 2013]*

2.1 The ducts

The ducts have been split into the different regions based on the air flow they are expected to contain for modelling purposes as shown in Figure 2.2. Below follows a description of the different sections. The actual dimensions of the ducts are not known but it can be fair to assume they all are static.

Harmony module: The module supplying Columbus with air and receiving the returned air. This module is not considered a part of this system and is thus not modelled. We can not assume there are any correlation between the two flows from this section due to unknown internal structure.

supply: Transfers air from the Harmony module. This region has many unknowns as the only sensors are those from the inter module ventilation (IMV) supply fan (ISFA). The input flow may have any temperature or humidity even if it is assumed to be constant during normal operation. There is both a check valve and a manually controllable valve.

outlet: From the junction of the supply and fan{1,2} to the multiple outlets to the cabin. In this section the air gets filtered through a HEPA filter, regulated in a controllable cabin temperature control unit (CTCU) with a condensate heat exchanger (CHX) and finally distributed to the cabin. Apart from the cabin this is the only section where the properties of the air is expected to change. The pressure loss is expected to be quite high. There is a sensor measuring the pressure difference over the heat exchanger. It is unknown if the water purging system is contained between this sensor or not.

fanJ: Small region where both cabin fans share the ducts, it is unknown if the majority of this section is before or after the actual fans but should not be relevant for modelling purposes as it should only contain the air flow that is produced by the cabin fans.

fan{1,2}: Small region where only the cabin fan{1,2} (CFA) affects the flow, drawing air from the inlet region to the supply and outlet junction. Both fans have check valves and should virtually be identical. These sections are subsections of the fanJ section.

inlet: Small region that is transferring air from the cabin to the fanJ and return section. There should be an air filter that can be clogged here. There is also smoke detectors mounted in this section.

return: The section where air returns to the Harmony module. Just as for the supply section, only the sensors provided by the IMV return fan (IRFA) are known, however, this section should not affect the system that much on a whole. There is no check valve in this section but there is a controllable valve.

cabin: The actual cabin. Some of the sensors, like pressure and cabin temperature, should be mounted in this section. This section is not a duct and can

not be considered sealed. It is however, fair to assume that there is some correlation between the in and out flow from this section.

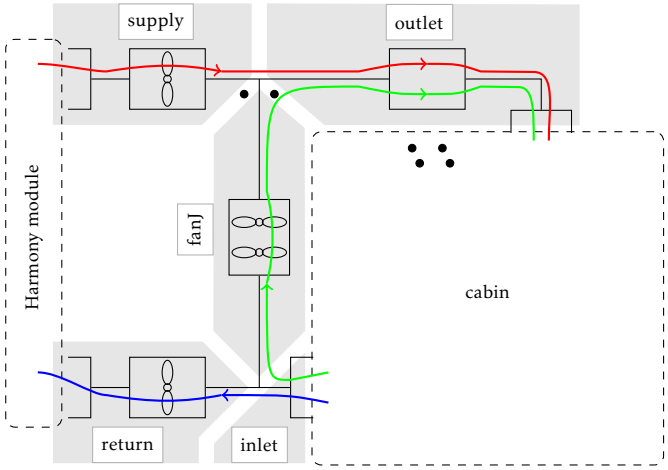

Figure 2.2: *A schematic overview of the ducts and air flows*

2.1.1 Mass conservation

Under the assumption that there are no leaks in the system, the local pressure will variate based on the mass flow in and out of each node as a function of in and out flow.

$$\dot{p} = f(Q_{in}, Q_{out}) \tag{2.1}$$

Additionally, it is assumed that the local pressures in the system are constant in the time frame giving the following relations between the flows:

$$Q_{outlet} = Q_{supply} + Q_{fanJ} \tag{2.2a}$$
$$Q_{fanJ} = Q_{fan1} + Q_{fan2} \tag{2.2b}$$
$$Q_{inlet} = Q_{return} + Q_{fanJ} \tag{2.2c}$$

The sections are described in Section 2.1.

2.1.2 Passive flow

Sections without a fan will have a flow driven by the pressure difference. This can be modelled by using the Bernoulli equation but as there are too many unknowns in the system, as duct size, number of bends and if the flow is laminar or not, a simplified model is used instead given by

$$\Delta p_i = K_{f,i}(Q_i K_{c,i} - K_{d,i})^2 \tag{2.3}$$

which translate to that most terms in Bernoullis equation are considered constant, neglectable or flow dependant.

2.2 Fans

Each fan has a certain number of sensors as listed in Table 2.1. All four fans are of the same type with an attached electronic unit (EU). Thus, one model structure should apply to them all. There is a check valve connected to the fan output to prevent reverse flow for all fans, except the IMV return fan. A schematic overview of a fan is displayed in Figure 2.3.

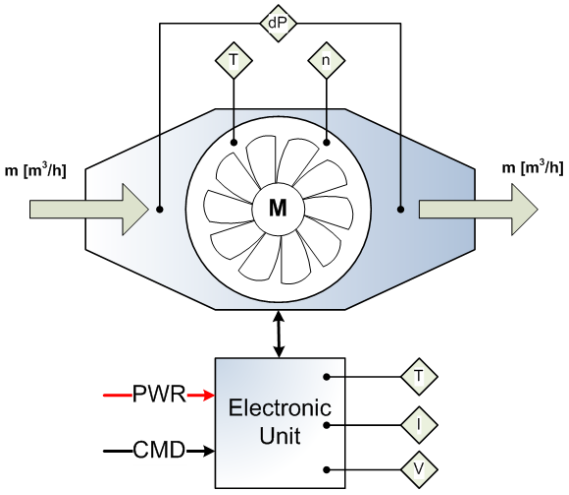

Figure 2.3: *A schematic overview of a fan. [Source: Airbus DS, 2013]*

Table 2.1: *Fan signals*

Measurements		
part	signal	unit
CFA{1,2}_/ I{S,R}FA_	Delta_P_	kPa
	Fan_Speed_	min^{-1}
	Input_Current_	A
	EU_Temp_	^{o}C
	Fan_Temp_	^{o}C
	Input_Voltage_	V
Switches/Analyses		
CFA{1,2}_/ I{S,R}FA_	Pwr_Stat_	ON/OFF
	Avail_Stat_	(UN)AVAIL
CFA_	Redun_Stat_	(UN)AVAIL
I{S,R}SOV_	Vlv_Open_Stat_	X/OPEN
	Vlv_Closed_Stat_	X/CLOSED

2.2.1 Fan curve

It is common practice to utilise fan curves when designing ventilation systems to ensure that the fans are working in their efficient load regions by calculating the different operation points. Figure 2.4 shows the fan curve for the current fans. This curve is normally generated in ideal (dry) conditions and to scale this for other fan speeds (N) and air densities (ρ) the affinity equations

$$\frac{dp_1}{dp_2} = \frac{\rho_1}{\rho_2}\left(\frac{N_1}{N_2}\right)^2, \quad \frac{Q_1}{Q_2} = \left(\frac{N_1}{N_2}\right), \quad \frac{W_{f1}}{W_{f2}} = \frac{\rho_1}{\rho_2}\left(\frac{N_1}{N_2}\right)^3 \tag{2.4}$$

are utilised. A map has been made to calculate flow (Q) as a function of total dynamic head (KH) (dp), i.e., the fan's workload. However, there are two possible flows, and the fluid effect (W_f) the fan can generate depends on where on the curve the fan is operating as illustrated in Figure 2.4. The lower flow output, the so called stall region, is considered a fault as it is not only less efficient but also induces more wear on the fan. There are multiple reasons for a fan to enter stall mode where the most common ones are from the design stage like overdimensioning of the fans or competing parallel flows. It is assumed that this fan curve is valid for 8500 rpm and with an air density of 1.2041 kg/m^3 (dry air at 20^oC).

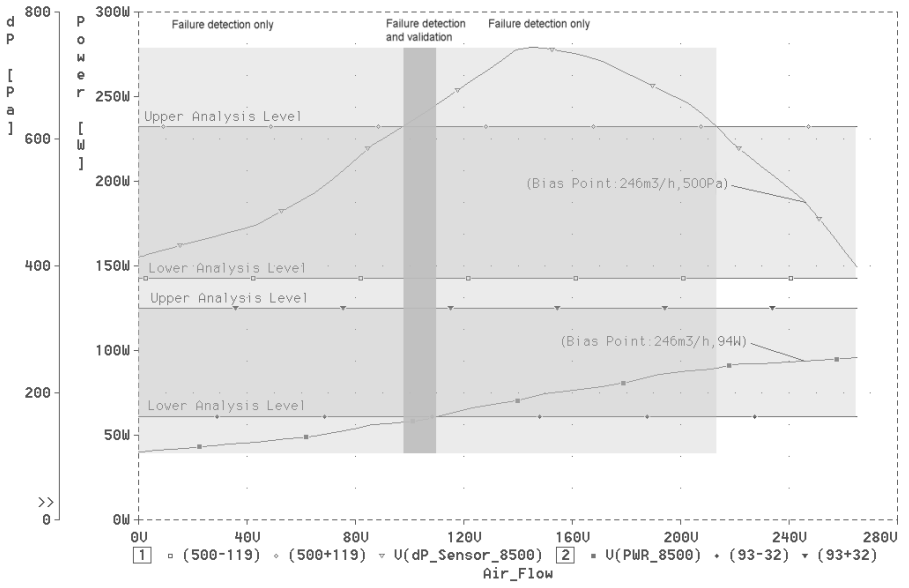

Figure 2.4: *Characteristics of the fans, [Source: Airbus DS, 2013]*

2.2.2 Fluid work

All attempts to utilise the fan curve to determine an operational point by measured power consumption has been very unreliable. The system loss is not insignificant. Analysis indicates that the effect in the fan curve is measured (W_m) and not fluid effect (W_f), this implies the scaling law for W does not apply directly. This observation is strongly supported by data as well.

$$W_m = W_f + W_{loss}$$

To determine W_{loss} in the ideal conditions is not feasible and there is no basis to assume W_{loss} would scale as nicely. Instead, a model which is given by

$$W = K_p \Delta p Q + K_q Q^2 + K_d \tag{2.5}$$

has been estimated from the fan curves map by linear regression where K_p would be the fan's efficiency in building pressure, K_q flow work efficiency and K_d any constant losses. This approach is accurate enough to determine if the fan is working in the desired mode or in the faulty stall mode. This is later validated in Section 3.4.

2.2.3 Fan flow by power consumption

An already parameterised model has been received from Astrium which is a direct correlation between the fan power consumption and the flow that the fans produce. This model is done by linear regression of collected and, possibly, preprocessed data.

$$W = K_f Q + K_d \tag{2.6}$$

The validation in Section 3.6 conclude that this model is accurate within certain intervals which should be enough for residual generation.

2.3 Sensors

A complete listing of the available sensors is shown in Table 2.1 together with Table 2.2. Some of these sensors have been concluded to be of no use to the current model and are consequently not included.

AFS : The air mass flow sensors, positioned in the junction between supply, fan{1,2} and outlet. Will be abbreviated as AFS{1,2} in most parts of this document. The sensors are validated in Section 3.3 but are often saturated and due to local vertices does rarely produce redundancy. An improved model (2.7) has been made in combination with (2.2a) which is validated in Section 3.5.

$$AFS_i = \sum K_{j,i} Q_j + K_0 \tag{2.7}$$

CHXFA : The pressure sensors over the condensate heat exchanger fan. Are in used in combination with (2.3) to produce an estimation of Q_{outlet}, val-

idated in Section 3.5 and directly in (3.2) which also is validated in Section 3.5.

CTCU : Multiple sensors and switches, only the cabin temperature signals are used in the current model and these are validated in Section 3.3.

CWSA : Condensate water separator assembly. None of these signals are used in the current model.

HS : Air humidity sensor, measuring percentual water saturation in the air, used to calculate current air density (2.11) and is validated in Section 3.3.

TPS : Air pressure sensor, measured in $mmHg$ and used in the system and is validated in Section 3.3. Will be abbreviated as TPS(1-4).

LCOS : Liquid carry over sensor. Used to detect if the CWSA fails which is outside the scope of this thesis.

Cabin : Smoke detectors and a low airflow state, none of these are currently in use by this model.

Table 2.2: *Non fan signals*

Measurements		
part	signal	unit
AFS{1,2}_	Cab_Air_Massflow_	kg/h
CHXFA_	Delta_P{1,2}_	Pa
CTCU{1,2}_	Cabin_Temp{1-3}_	^oC
	Avg_Cabin_Temp_	^oC
	TCV_Posn_	%
CWSA{1,2}_	Input_Current_	A
	Delta_P_Air_	kPa
	Delta_P_Water_	kPa
HS{1,2}_	Air_Humidity_	$\%rH$
TPS{1-4}_	Air_Press_	$mmHg$
LCOS{1,2}_	Level_	V
Cabin_	SD{1,2}_Obscuration_	V
	SD{1,2}_Scatter_	V
Switches/Analyses		
CTCU{1,2}_	Pwr_Stat_	ON/OFF
	Kick_Posn_Up_Stat_	(IN)ACTIVE
	Kick_Posn_Down_Stat_	(IN)ACTIVE
	Kick_TCV_Hold_Posn_	(IN)ACTIVE
	Dryout_Stat_	DISABLED
	Cntl_Loop_Stat_	(DIS)ABLED
CWSA{1,2}_	Pwr_Stat_	ON/OFF
Cabin_	Low_Airflow_Stat_	NOM/LOW

2.4 Other equations

There are some other equations used in the model that are of a more generic type, like the pressure difference

$$\Delta p = p_a - p_b \tag{2.8}$$

and conversion from mass flow (q) to volume flow (Q)

$$q = \rho Q \tag{2.9}$$

Ohm's Law for calculating power

$$W = IV \tag{2.10}$$

used to calculate the fans' power consumption where W equals Watt, I the supplied current and V the voltage.

To calculate the air density (ρ) as a function of pressure (p), temperature (T) and relative humidity (ϕ) the ideal gas law is used

$$\rho = f(T, p, \phi) \tag{2.11}$$

$$\rho = \frac{p_{air}}{R_{air}T} + \frac{p_{water}}{R_{water}T}, \quad p_{water} = p_{sat}\phi, \quad p_{air} = p - p_{water}$$

where p_x are partial pressures and R_x being the specific gas constants.

3

Model parametrisation and validation

This chapter contains verification and parametrisation of the models listed in Chapter 2. Starting with an overview of the data sets used throughout the chapter and a description of how the model error is measured. This is followed by verification of the system redundancy and the fan curve assumptions. Finally the parametric models are both estimated and verified.

3.1 Data sets

For parametrisation of the models, a data set where the fans are being tested in different modes has been used. This training data is not normal operation but made upon request for assisting in modelling the system. For validation some data sets with fairly normal operation has been selected and a data set which is a subsampled version of all available data not used as training data. Below follows a short description of each data set and an overview of the training data set is shown in Figure 3.1.

Training data: Data set from 20100[1-2]* which is a testing run to gather information of how the system change in the different modes. Both IRFA and ISFA is enabled with open valves. The temperature, cabin pressure and humidity is either constant or within normal variation.

n20080319: Validation data set with low sample rate on some signals, CFA2 is doing the cabin flow and all fan modes and valves are kept at a constant rate.

n2008070: Validation data set going from normal operation to internal mode where the supply fan gets disabled and both cabin fans are run at high

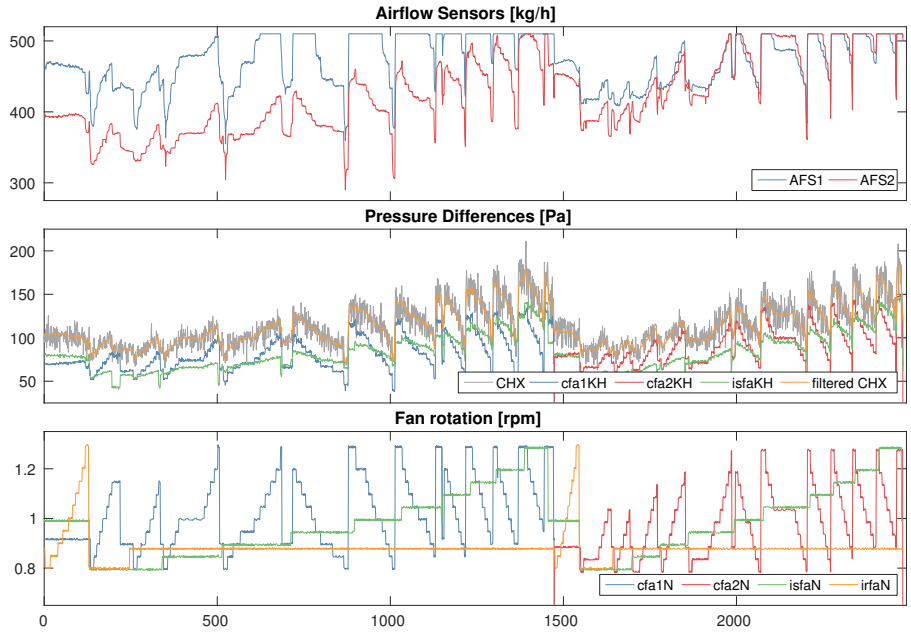

Figure 3.1: *Raw sensor data from the training data set*

speed. The return fan is disabled the whole set but its valve is opened in the internal mode. The supply valve is not closed. There is a slight peak in humidity when the valves close.

20091129: Validation data set in normal operation with a swap of which cabin fan is running, followed by the return fan disabled and its valve closes. There is also noticeable drops in the AFS sensors due to the CWSA working.

subsampled data: All available data is subsampled for quicker processing and used for validation. The last registered value is used when data is missing.

3.2 Model error measurement

To measure the model error, the model equations are evaluated on the validation data sets. Mean square error (MSE) and mean absolute error percentage (MAE%) have been used to determine if an equation seem to hold true as

$$MSE : \frac{1}{N}\sum_{i=1}^{N}(x_i - \hat{x}_i)^2 \qquad MAE\% : \frac{1}{100N}\sum_{i=1}^{N}\frac{|x_i - \hat{x}_i|}{x_i} \tag{3.1}$$

in addition to a manual behavioural check.

3.3 Validation of sensor redundancy

Validation of the simplest form of redundancy where multiple sensors ought to monitor the same information. This is done to determine if the sensors truly are redundant and, if not, how much their output differs. As there is no need for estimation in this section, the training data set is also used as validation data.

Air pressure sensors

The air pressure sensors are quantified to steps of $0.4mmHg$. The residual ends up being within one, occasionally two, quantification steps after adjusting offset. If this is due to constant local pressure differences or due to bad calibration cannot be determined. The only contradiction to this can be found in the subsampled data set as shown in Figure 3.2. The biggest offset is between TPS1 and TPS4 of 9 steps being $0.48kPa$ ($3.6mmHg$).

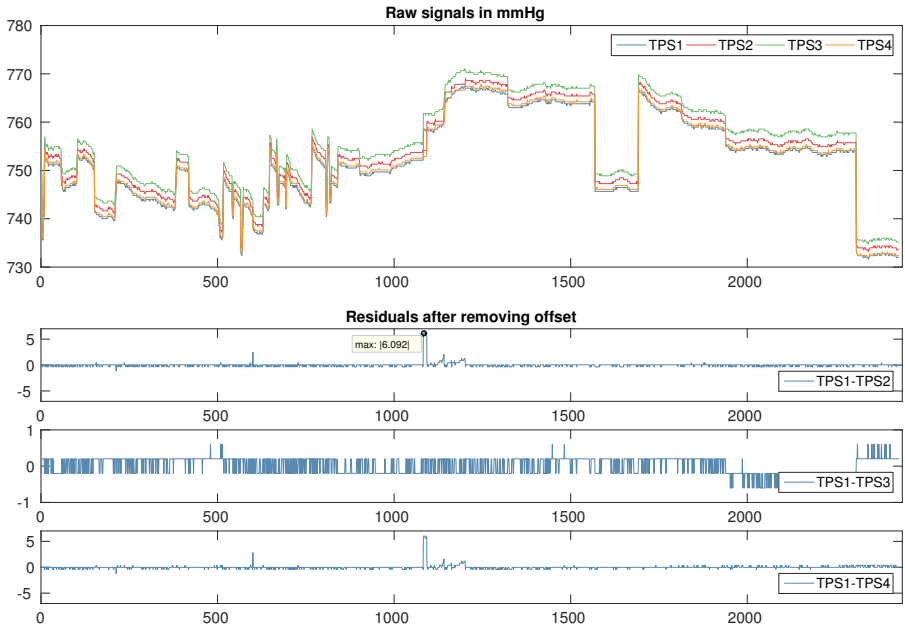

Figure 3.2: *Validation of the air pressure sensors on the subsampled data set*

Air flow sensors

At the junction with the air mass flow sensors (AFS) there are multiple models that could be used to determine the actual flow in the region. As shown in Figure 3.1 the AFS sensors do not give a good indication of the true flow in the region. In Table 3.1, the results are listed together with how many percent of the data is excluded due to saturated sensor. The difference between the sensors is often bigger when the sensor is capped.

Table 3.1: *Air mass flow sensors redundancy error*

data	MSE	MAE%	Capped%
training	3847.4	11.42	33.5
n20080319	1088.7	8.22	0.0
n2008070	214.78	1.79	70.7
n20091129	182.27	2.48	21.1
subsampled	4407.9	11.29	32.3

Cabin fan pressure sensors

It is only when both fans are enabled at the same time that we can utilize this redundancy. It is not known if this is never recorded or if it is discarded. Among the selected validation data set there is only redundancy in the n2008070 data set and the subsampled data set. There seems to be an offset between these and we can assume part of it originates from the differences in duct design. However it does not seem to be constant which could be due to clogging in the system. In the n2008070 data set the predicted offset is as large as 10%. The MAE% is 5.34% in the subsampled data set while 3.79% in the n2008070 data set.

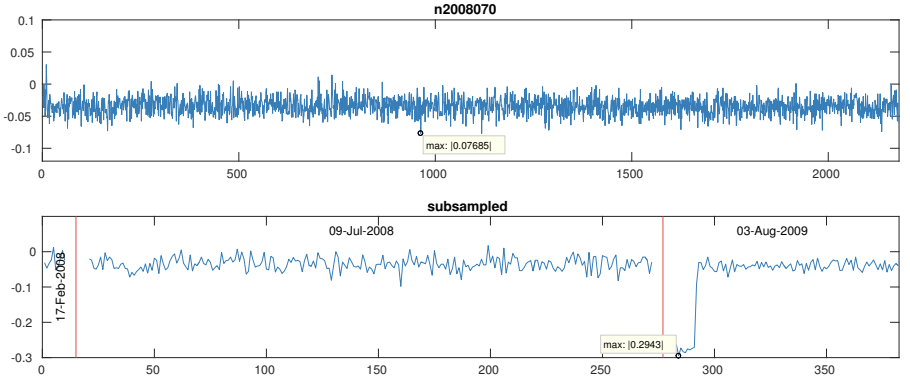

Figure 3.3: *Residuals of the cabin fan pressure sensors*

At the beginning of the 3rd August 2009 section of the subsampled residual, as shown in Figure 3.3, there is a clear disturbance that is believed to originate from that the just starting fan is having problems leaving stall mode. During this interval both fans are set to run at high speed while the supply fan was abruptly

disabled indicating that there was something wrong with the supply and a desire to vent out the module was present. However, it is believed that this made one of the fans not able to ensure a positive flow.

Fan input voltage

The input voltage to the fans are often varying and the signal is depending on the modes the fan is working in. This variation is fairly small and the variation is normally within one percent. In Table 3.2 the MSE and MAE% of the different fans are shown when comparing to the supply fan.

Table 3.2: *Fan voltage redundancy error*

ISFA vs	IRFA		CFA1		CFA2	
data	MSE	MAE%	MSE	MAE%	MSE	MAE%
training	0.964	0.79	0.191	0.34	0.859	0.74
n20080319*	2.189	1.20	-	-	0.953	0.79
n2008070*	-	-	0.044	0.17	-	-
n20091129*	1.141	0.86	0.134	0.30	0.978	0.80
subsampled	1.945	1.12	0.094	0.24	0.963	0.79

Humidity sensors

As shown in Figure 3.4 are the humidity sensors displaying slightly different dynamics indicating that they must have different positions in the system. The MSE and MAE% are displayed in Table 3.3.

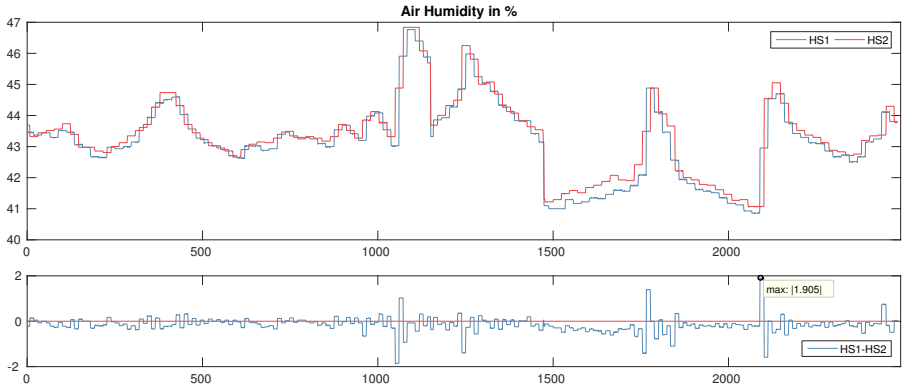

Figure 3.4: *Humidity sensors from the training data set*

Table 3.3: *Humidity sensors redundancy error*

data	MSE	MAE%
training	1865.67	5.94
n20080319	1849.63	5.17
n2008070	1178.48	2.16
n20091129	1848.76	3.76
subsampled	1641.13	4.87

Temperature sensors

When comparing the temperature sensors, it is concluded that they stay within a quantification step between each other. The variation of the average temperature signal indicates that it is calculated from more detailed data. CTCU2 is most of the time not in use at all. No offset or drifting has been detected.

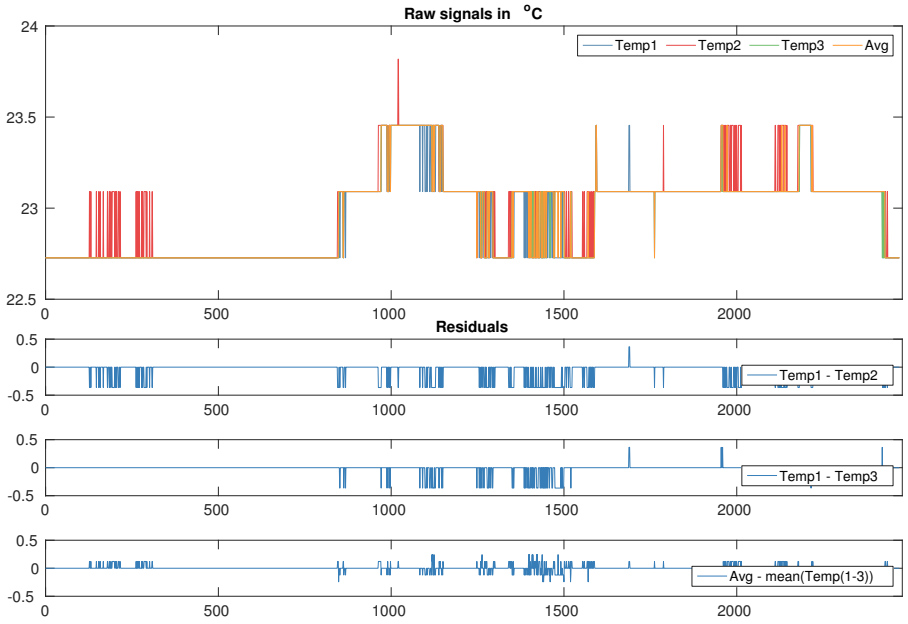

Figure 3.5: *Validation of the CTCU1 temperature sensors*

3.4 Validation of the fan curves

As no true flow is known in the system, it is assumed that the fan curve together with the affinity laws generate a flow that could be utilised instead. The equations displayed in Section 2.2.1 are part of the affinity laws and will not be validated. Since the estimations involving flows produced by this method yield better results than the other methods tested indicate that the method is functional.

Detection of fan modes

To utilise the fan curves the fan mode must be determined. Initially this was done manually on the training data set to generate estimations but as the generalisation, described in Section 2.2.2, produced an overall better result, all estimations were re-generated using this. A comparison of the manual and automated mode detection is shown in Figure 3.6.

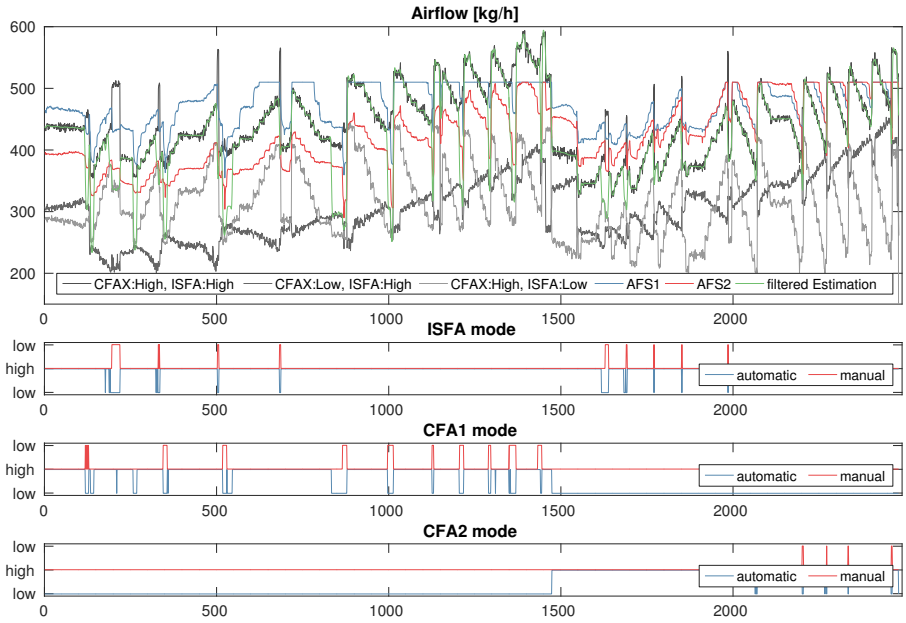

Figure 3.6: *Estimation of modes*

Decoupling of air density

Early in the construction of the model it was realised that the air density dependency often was not properly defined among the equations and input signals. In addition it was desired to decouple the density from the fan curves. The conclusions however, was that the fan curve and mode detection could not be satisfactory decoupled from the air density on a modelling basis. As fan mode is required to determine the produced flow, air density has not been decoupled from the fan curve model.

3.5 Model parametrisation and validation

The model is parametrised using linear regression in the case of (2.7) and (3.2) and curve fitting (Levenberg-Marquardt) in the case of (2.3). The parametrisation is done using the whole training data set except in the case of (2.7) where the samples with capped AFS signals were discarded.

$$AFS_i - Q: \qquad AFS_i = \sum K_{j,i} Q_j + K_0 \tag{2.7}$$

$$CHXdp - Q: \qquad \Delta p_i = K_{f,i}(Q_i K_{c,i} - K_{d,i})^2 \tag{2.3}$$

$$kh - CHX: \qquad kh = K_v \Delta p + K_0 \tag{3.2}$$

AFS-Q equation

A complete overview of the parametrisation of (2.7) is shown in Table 3.4. Overall yields the AFS2 sensor a very good result while the AFS1 parametrisation appears to be clearly faulty. An observation of the estimated parameters show that AFS2 uses about 60% of ISFA and 30% of the cabin fans. While the AFS1 parametrisation uses 30% of them all. An re-parametrisation weighting ISFA higher for AFS1 might give better results, but as the AFS1 sensor is often saturated, this might never be possible.

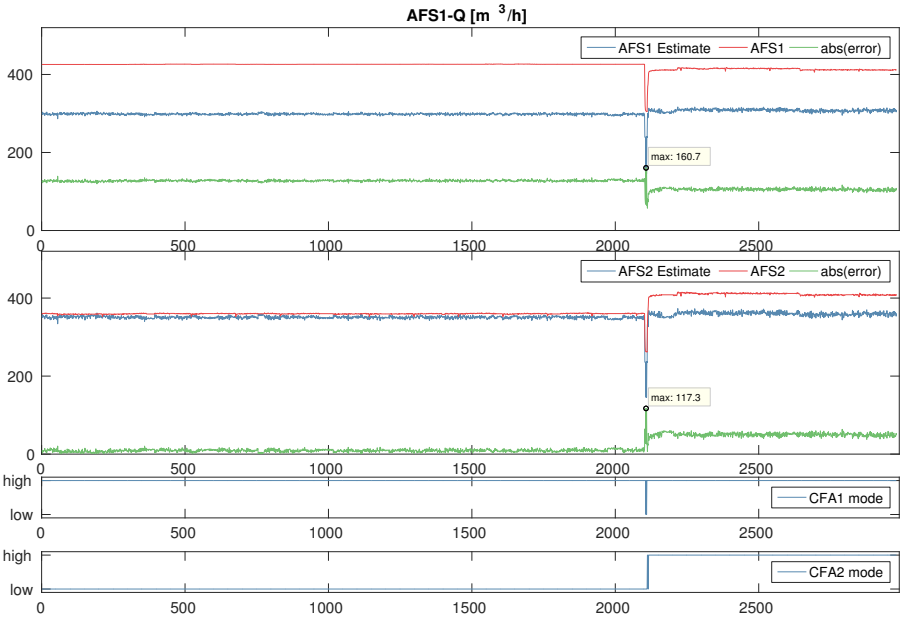

Figure 3.7: *Validation of the AFS-Q parametrisation (n2008070)*

In the n2008070 data set, displayed in Figure 3.7, the AFS2 parametrisation turned worse when the ISFA valve is opened. This could indicate that either the ISFA check valve is not operational or that there is a passive flow from the

supply module. This behaviour can also be seen in the n20091129 data set when the ISFA valve is closed.

Table 3.4: *AFS-Q parametrisation error*

data	AFS1 MSE	AFS1 MAE%	AFS2 MSE	AFS2 MAE%
training	12689	27.14	385.4	4.09
n20080319	4001	18.27	34.2	1.15
n2008070	14738	28.62	806.9	5.38
n20091129	15400	29.91	419.1	5.04
subsampled	13915	28.43	576.2	5.09

CHX-Q equation

It can clearly be seen that the parametrisation of (2.3) is not perfect when viewing its match to the training data set but seems to be holding quite well overall even if it appears to be suffering from an offset. It has to be analysed if this is drifting and thereby could be caused by clogging or any similar fault. In Figure 3.8 can the worst case of offset among the selected validation data sets be seen. Interesting is how the offset change sign when the working cabin fan is changed about sample 1700. Additionally the fan is having problems leaving the faulty stall mode which it finally manages once the return fan gets disabled just before sample 2000.

Table 3.5: *CHX-Q parametrisation error*

data	MSE	MAE%	MSE filtered	MAE% filtered
training	246.0	11.03	151.7	8.90
n20080319	228.5	12.45	58.9	6.08
n2008070	207.0	10.80	9.5	2.20
n20091129	156.0	9.62	52.5	5.71
subsampled	218.1	16.23	69.0	6.47

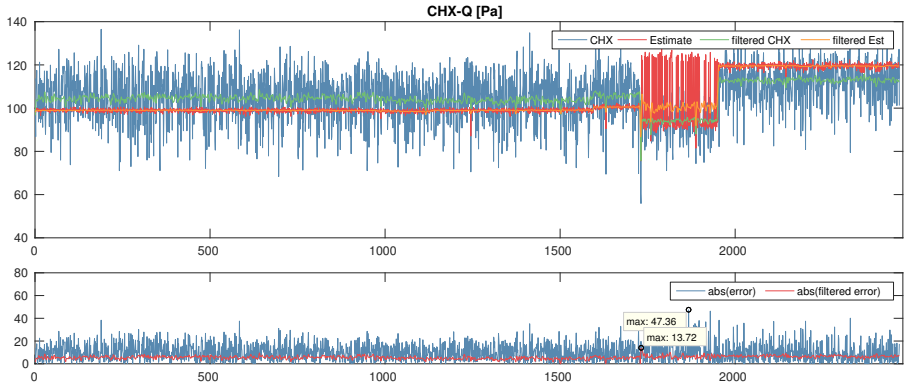

Figure 3.8: *Validation of the CHX-Q parametrisation (n20091129)*

KH-CHX equation

The results of the parametrisation of (3.2) is shown in Table 3.6. It can be noted that the parametrisation seem to overshoot when the return valve is closed. This happens in the first part of the n2008070 data set as seen in Figure 3.9 and in the latter part of n20091129 data set. This could indicate a back flow through the return fan or that there is a forced exhaust flow even when the fan is disabled.

Table 3.6: *KH-CHX parametrisation error*

data	MSE	MAE%	MSE filtered	MAE% filtered
training	0.012	10.41	0.0089	8.53
n20080319	0.011	10.36	0.0022	4.00
n2008070	0.023	16.27	0.013	14.19
n20091129	0.014	11.81	0.0084	10.36
subsampled	0.018	13.85	0.010	9.49

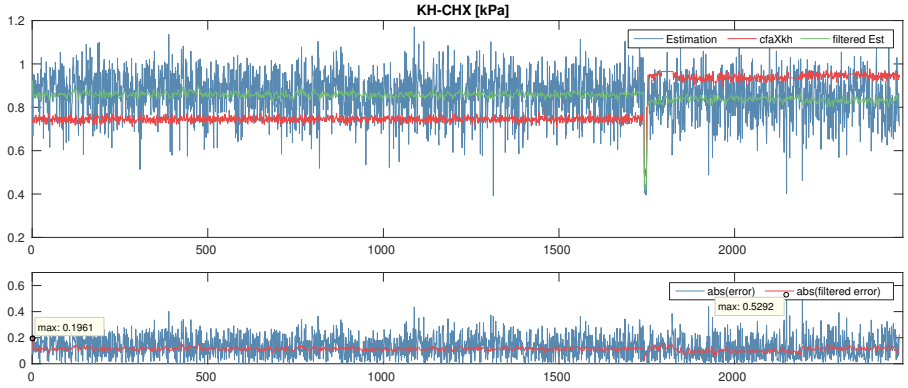

Figure 3.9: *Validation of KH-CHX parameterisation (n2008070)*

3.6 Validation of fan flow by power consumption

This model (2.6) has been verified against both the air flow sensors (AFS against $f(W)$) as displayed in Table 3.7 and against the air flow generated from the fan curves (W against $f(Q)$) (2.4) as shown in Table 3.8 with varying results. The model seems to hold fairly well when the fans operate in certain regions, but as this is quite often not the case the usability might be limited. Note that all fans are identical while the fan model (2.6) has different parameters for each fan. This could be a case of parametrisation for that fan's default working region rather than an overall. When using the parameters for CFA2 on the CFA1 flow an MAE% improvement of about 10% is achieved. Doing a similar attempt for IRFA does not yield any improved results, however as the IRFA flow is unknown this method can not be determined incorrect.

Table 3.7: *AFS estimation by power consumption error*

data	AFS1 MSE	AFS1 MAE%	AFS2 MSE	AFS2 MAE%
training	57562	37.9	56518	37.2
n20080319	4288.7	16.2	1095.0	7.6
n2008070	3360.6	10.8	3377.9	11.6
n20091129	2782.3	9.6	569.6	4.9
subsampled	6730.6	15.6	3361.3	8.1

Table 3.8: *Power consumption estimation by fan flow error*

	ISFA		IRFA		CFA1		CFA2	
data	MSE	MAE%	MSE	MAE%	MSE	MAE%	MSE	MAE%
training	940.2	20.9	3645.9	58.9	1092.2	28.0	1810.0	18.7
n20080319	3.3	1.2	2891.0	40.7	-	-	37.4	6.5
n2008070	14.0	2.6	-	-	432.6	19.7	523.1	17.7
n20091129	22.2	3.7	5.6	1.3	928.9	33.7	10.5	3.6
subsampled	14.0	2.6	1684.0	25.5	1161.7	32.4	61.2	6.9

4

Design of fault detection system

This chapter will explain how the faults are implemented, how the residuals are selected and finally tested. The desired detectable faults are added to the model where it then is partitioned into minimal structurally overdetermined sets (MSOs) by using Dulmage-Mendelsohn decomposition analysing the strongly connected components (SCC) [Frisk et al., 2012, Krysander et al., 2008]. Only when this is done is the greedy approach used to select the residuals.

4.1 The model

The complete model has 78 equations generating 55 states from 36 input signals. 32 errors are handled and below follows a description on how these are included in the modelled. The whole system is displayed in Figure 4.1.

Sensor Faults

Sensor faults (F_s) have been added into the model as an additive parameter to the measured sensor (S_s) as

$$S_t = S_s + F_s$$

As the original signal is known during simulation can multiplicative faults also be simulated by letting F_s be a function of S_s. This is added to all input signals except the enabled and open/closed valve signals.

Duct clogging

A simulation of the complete model is required to determine how clogging really affects the system. This is however not done and a heavily simplified fault is in-

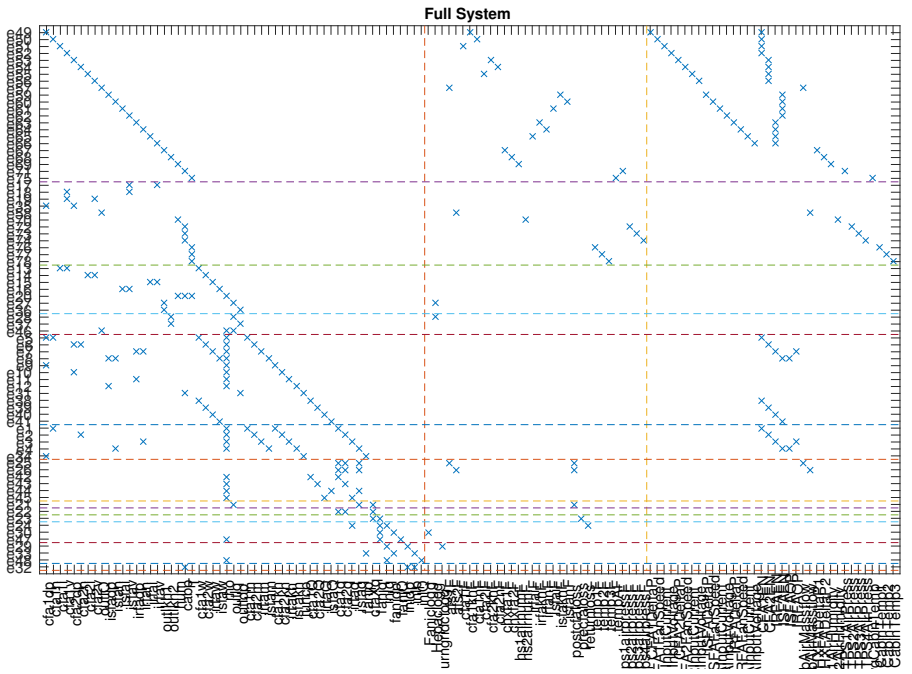

Figure 4.1: *Dulmage-Mendelsohn decomposition of the full model*

stead simulated that could be seen as clogging. The fault is added as a parametic fault to (2.3) as additional duct flow resistance.

$$\Delta p_i = K_{f,i}(Q_i(K_{c,i} + F_{c,i}) - K_{d,i})^2$$

Air leaks

Like with the clogging, a complete simulation is required to truly know how this type of fault is affecting the system. The simplified method here is that the negative fault (F_q) is added to the flows in (2.2) as

$$Q_{out} = Q_{in} + F_q$$

simulating that air has taken route elsewhere than the monitored connections.

4.2 Generating the MSOs

Finding the MSOs for a system this size is a very heavy task and has therefore been split into simpler sub tasks. For the chosen method the system is first sorted using Dulmage-Mendelsohn decomposition as shown in Figure 4.1. An upper subset, A, is arbitrarily selected that is considered fast enough to be completely searched while still having overdetermined components. After which the rest, B, of the system is detached. This detached subsystem is searched for all existing

MSOs and for each found MSO the dependant equations and states from A is added to form a new matrix C that is again searched for all MSOs.

Consider the system shown in Figure 4.2. After selecting an A, which has at least an overdetermined part (e5 and e6) the rest B is selected. Analysis of B yields two MSOs that forms $C1$ and $C2$ as illustrated. Observe that for $C2$ has e2 been completely removed from the final search.

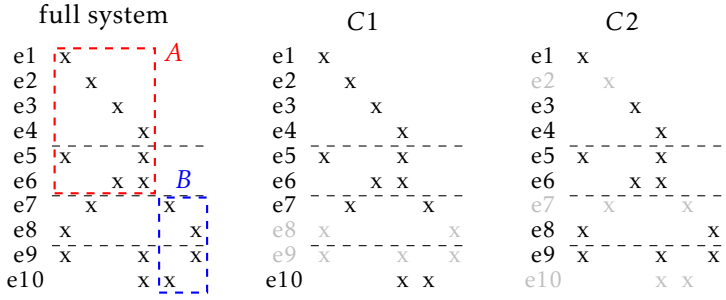

Figure 4.2: *Illustration of the MSO generation process*

This method should find most MSOs of the system but can not guarantee a complete set as this would require that also the exactly determined sets in B have to be considered for generation of additional C sets.

These MSOs are then sorted according to which structural faults they could detect in preparation for the residual selection. For this particular system the method has produced 1.8 (1776385) million MSOs covering 63604 different combinations of structural fault detectability.

4.3 Selecting the residuals

To select a sufficient subset of MSOs, a greedy search algorithm is applied [Svärd et al., 2013]. The greedy search works such that a residual is added to the solution set that fulfils as many new fault isolation requirements as possible. To do this each of the residual is tested and the best candidate is included in the solution set. This whole process is repeated until no further improvements are to be found.

4.3.1 System modes

As the approach above were selecting residuals purely based on their structural fault detectability, except when the residuals were tested, it was clear that some residuals were only valid in certain modes. This is primarily the case for which cabin fan currently is in operation. To cope with this all the selected residuals were tested and classified.

Residuals that appeared to be invalid in any mode were completely discarded and the search operation was re-run from that position onward. This is also done

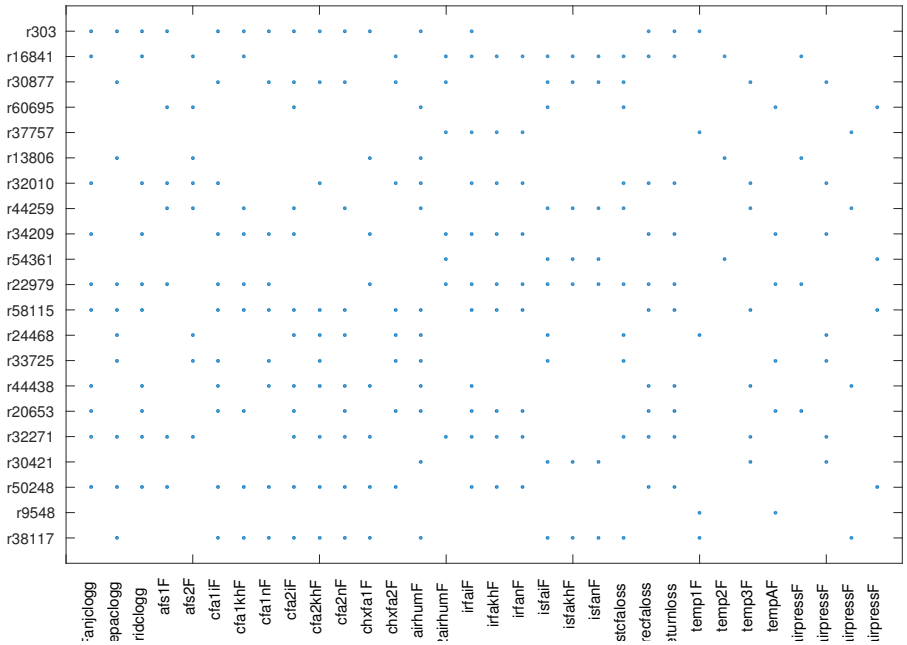

Figure 4.3: *Final structural detectability matrix*

on-the-fly when a new residual has been selected. In the case where there exists multiple MSOs with the same set of structural detectability has only the second one been tested if the first has been marked as invalid.

Once a general set was generated with ideal isolability without regards of the different modes, the specific mode dependant isolation matrices were analysed and the set was expanded to improve the currently worst matrix using the same method. This was iterated until no improvements to any of the isolation matrices could be found. The final structural detectability matrix is shown in Figure 4.3.

4.4 Generating the residual code

During the classification process the residuals have been generated and tested. The residual itself is processed into executable code using a function that will go through the directed acyclic graph (DAG) and return the required equations to calculate each state. A limitation of this function is that it cannot handle when an equation set has to be solved to generate the next needed state, requiring manual supervision of the output before execution.

As many of the residuals calculate the same states using the same equation sets, another function was made to merge residuals by creating new states and equations when there is a conflict in method to calculate an otherwise shared state. This allows the generation of one function containing all residuals.

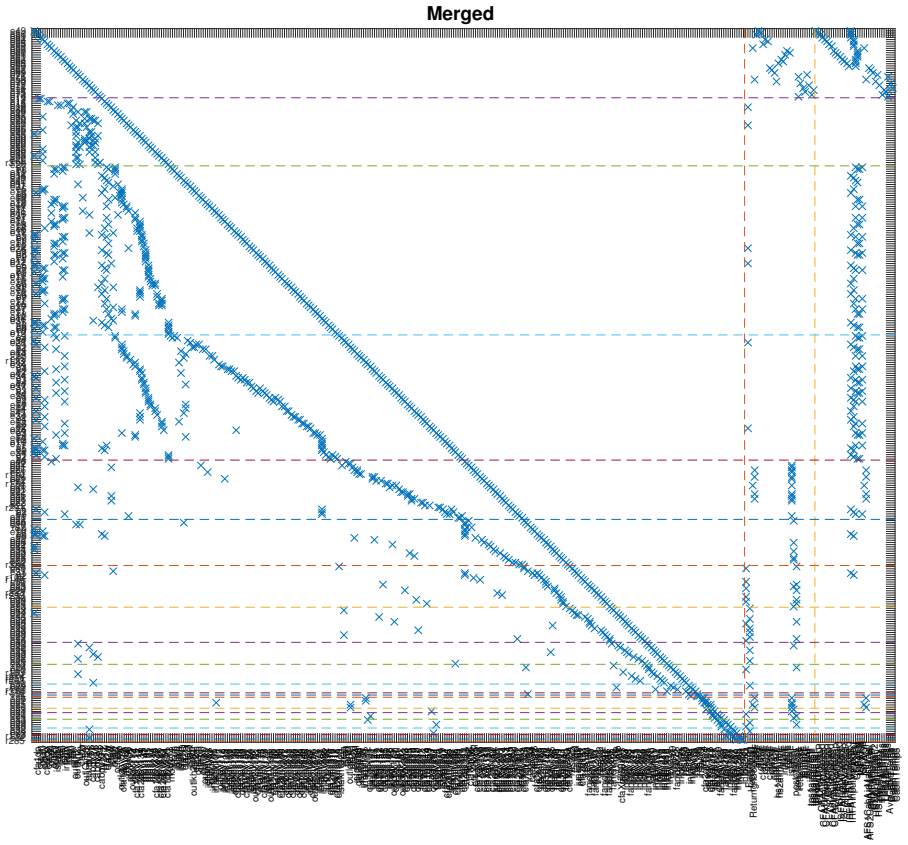

Figure 4.4: *Dulmage-Mendelsohn decomposition of the merged residuals*

Even though the merging function is far from ideal and can still produce multiple identical steps the produced output is remarkably faster than running the residuals separately. The Dulmage-Mendelsohn decomposition of the merged residuals is displayed in Figure 4.4 with horizontal lines marking the splits of the DAG, i.e., each block of equation can only be calculated by using the states from previous blocks.

4.5 Validation of the residuals

During the classification process a heavily subsampled selection of both training data and the whole data set was used as described in Section 3.1. A fault profile, that was repeatedly alternating and increasing in magnitude was injected. This was very effective for the purpose of detecting what modes the residual required but the results could be very hard to interpret.

Due to that a more simple data set have been selected for validation and a single percentual fault has been injected. The fault appears at sample 100 as a step of 50% of original signal and stays constant for a few samples to then decrease to zero. At sample 300 is the same procedure done in reverse. For easy reference, the fault profile is added to all the figures.

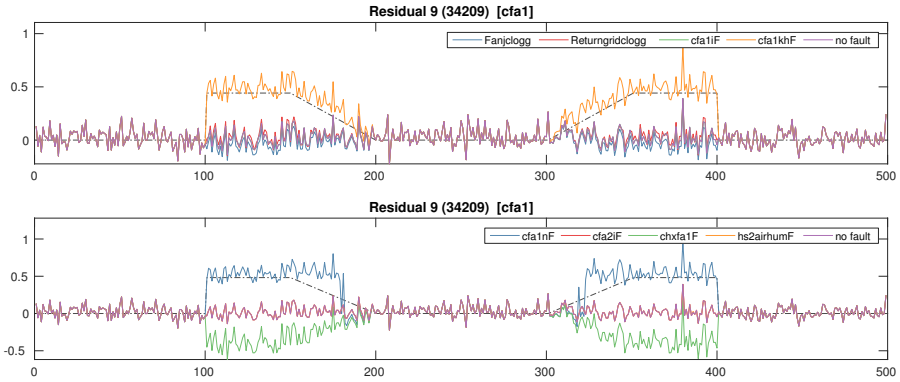

Figure 4.5: *Validation of residual 9*

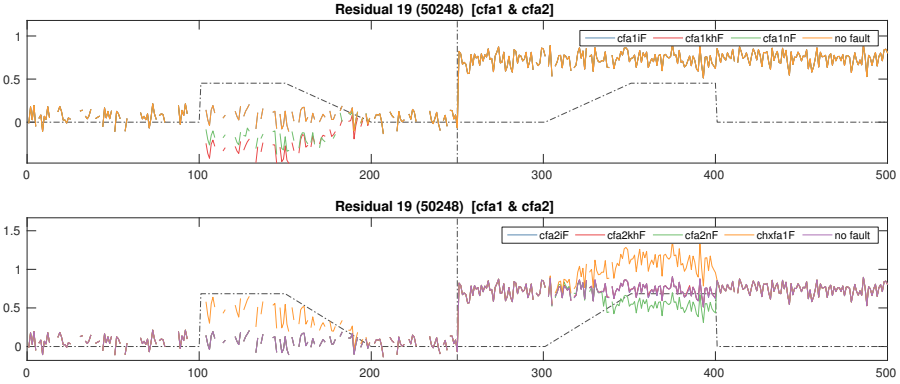

Figure 4.6: *Validation of residual 19*

The data sets n20090803 and n20091127 have been used respectively for the residuals that require either CFA1 or CFA2 enabled. Regarding residuals working with both or neither fans have part of both data sets been used, switching at sample 250.

The results are varying quite much among the residuals, some are noisy and only detects parts of their indicated faults, like residual 9 displayed in Figure 4.5, while others are clean and produce easily detectable injected faults.

A few of them are not centred around the zero axis, like residual 7 and 8 displayed in Figure 4.7, while are still reacting to the injected fault in a detectable way.

Note that residual 7 produces not-a-number (NaN) with an injection on the TPS2 air pressure sensor and could be considered triggering. In the case of residual 19 however, should a NaN not be considered an alarm which seems to be requiring a separate handling depending on which of the fans are currently running. This can be seen in Figure 4.6.

Generally, most of the residuals are prone to the fault that they structurally are able to detect, but they all need unique methods of detection, and in some cases, a unique method per fault is required.

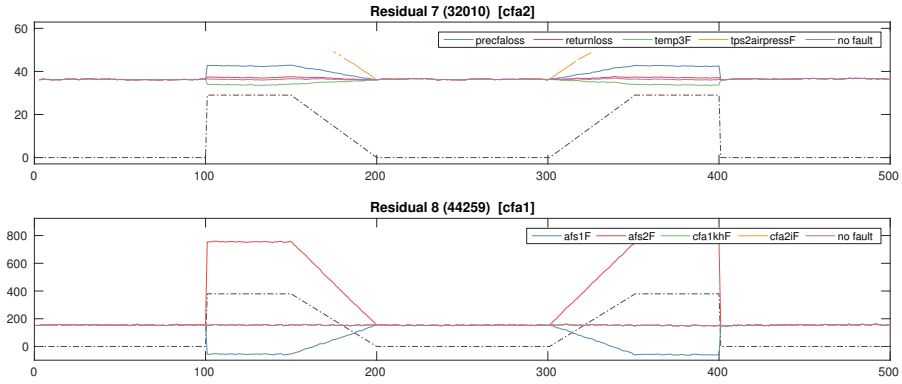

Figure 4.7: *Validation of residual 7 and 8*

5
Results

In this chapter the results of the design are presented both when detecting on injected faults and when analysing the real data.

5.1 Fault detection and isolation with injected faults

Analysis of the structural fault detectability and isolability of the final system yield that 20 of 32 faults are fully isolable. The structural fault isolation matrix is shown in Figure 5.1a.

A manual detection on the residuals with injected single faults cannot with confidence detect any air humidity faults. Neither can it detect the cabin fan 1 current when the system is in a cabin fan 2 only mode, which would be expected, but this brings the question on why this is detected for the opposite case. The rest of the faults are all detectable and about half of the faults are possible to isolate by manual analysis of the residuals. The isolation matrix when doing a manual analysis on injected single faults is presented in Figure 5.1b and show a fair reduction of isolability which indicate that the residuals selected might not be ideal and further iterations of the residual picking process may be needed.

Simple test quantities were produced for automated detection on the real data as manual detection on the residuals is a time consuming task. The test quantities are using either absolute value, cumulative sum (CUSUM) or a windowed mean value depending on best detectability on the injected fault data. The thresholds were then determined automatically using the test validation data so no false alarms were triggering. For better isolation were the set completed with a few not-a-number test quantities. In Figure 5.1c is the final isolation matrix presented. This indicate that the issue with detecting the air humidity faults were not a

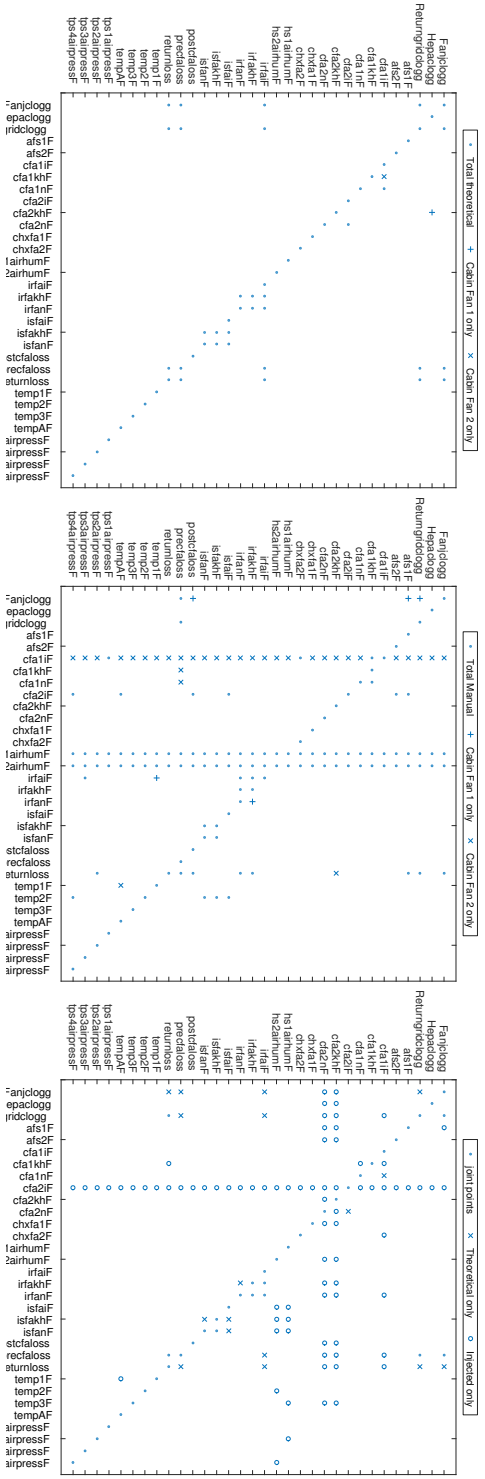

(a) Structural *(b) Manually detected* *(c) Automatically detected*

Figure 5.1: *Isolation matrices*

residual problem. This automated design can not detect the cabin fan 2 current fault and a few other pairs can not be isolated from each other.

5.2 Detection and isolation on real data

Analysing the test quantities described in Section 5.1 on the real data yield that many of the test quantities triggers and many of the faults are detected but it is not often a single fault could be isolated with certainty at any given time. Almost all faults isolated are indicated in short time intervals which makes it hard to determine if these are proper or false alarms. There are however a few cases of prolonged repeated patterns presented below. Noted can be that less than 3% of the time no test quantities are triggering giving us a total of only about 11% of no fault data including the cases where the data is not properly defined.

The figures throughout this chapters displays the alarms as how many test quantities are triggering for this fault in single fault isolation certainty. That means that if an alarm is lower than 1 then there is at least one conflict in isolation. This is calculated by count of alarming candidates of detected candidates for the specific test quantities triggering. No faults are alarming if only one test quantity is triggering nor if there are contradiction among the test quantities triggering. This is a simple form of isolation by column matching. One shortcoming with this is that in case of multiple faults could one fault prevent another faults alarm to go off which is also the case for false triggering of test quantities. All faults not covered in the figures are not alarming during the displayed data set, not even briefly.

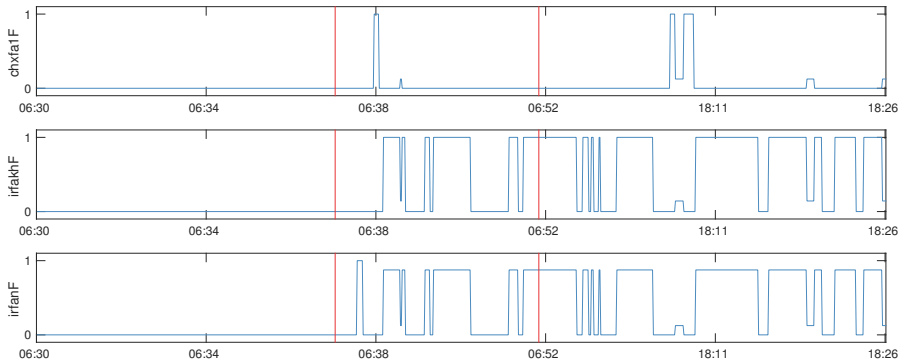

Figure 5.2: *First noted alarm in the outset data set. The vertical lines indicate time lapses with missing data in the data set.*

Return fan faults

Figure 5.3a displays the alarm signal for fault on the pressure and rotation speed on the return fan on real data from 11th Mars 2008. It cannot isolate the two faults from each other but it is fairly certain there is a fault on either of those

two sensors. Also to be considered is that this is not a constant indication which could imply that these all are false alarms. There are occasionally some other faults alarming, but those are generally brief in nature. Figure 5.2 displays the first occurrence of this pattern.

Clogging faults

Another pattern that is repeating over quite large portions of the real data processed is where there is alarms on multiple parametric faults and pressure sensors. These faults are confirmed in Figure 5.1c to be hard to isolate from each other. A sample of this real data have been selected and is displayed in Figure 5.3b.

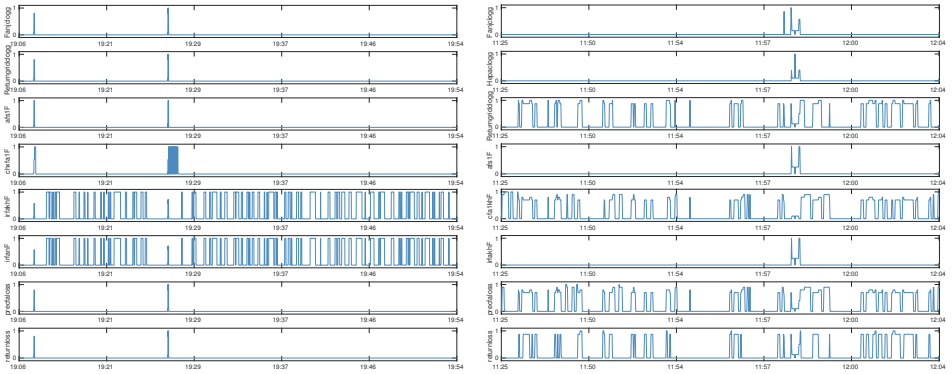

(a) *11th of Mars 2008* ***(b)*** *21th Mars 2008*

Figure 5.3: *Alarms during selected parts of the outset data set*

Outset data set

Both of the above mentioned patterns are part of the a data set, which is the longest set, containing more than half of the available real data. The name is arbitrary selected due to this set being the first continuous data available for analysis. This whole set has been marked as return fan loss but there is also documented unexpected clogging of the return grid. Additionally this data have a fairly low and inconsistent sampling.

However, the above detected patterns and the documented faults are not entirely coherent as only half of the data is detected to have return fan problems whilst the clogging seems to continue further than what has been reported. An overview of the alarms during this whole set is shown in Figure 5.4 with markings on where the part of unexpected clogging is documented.

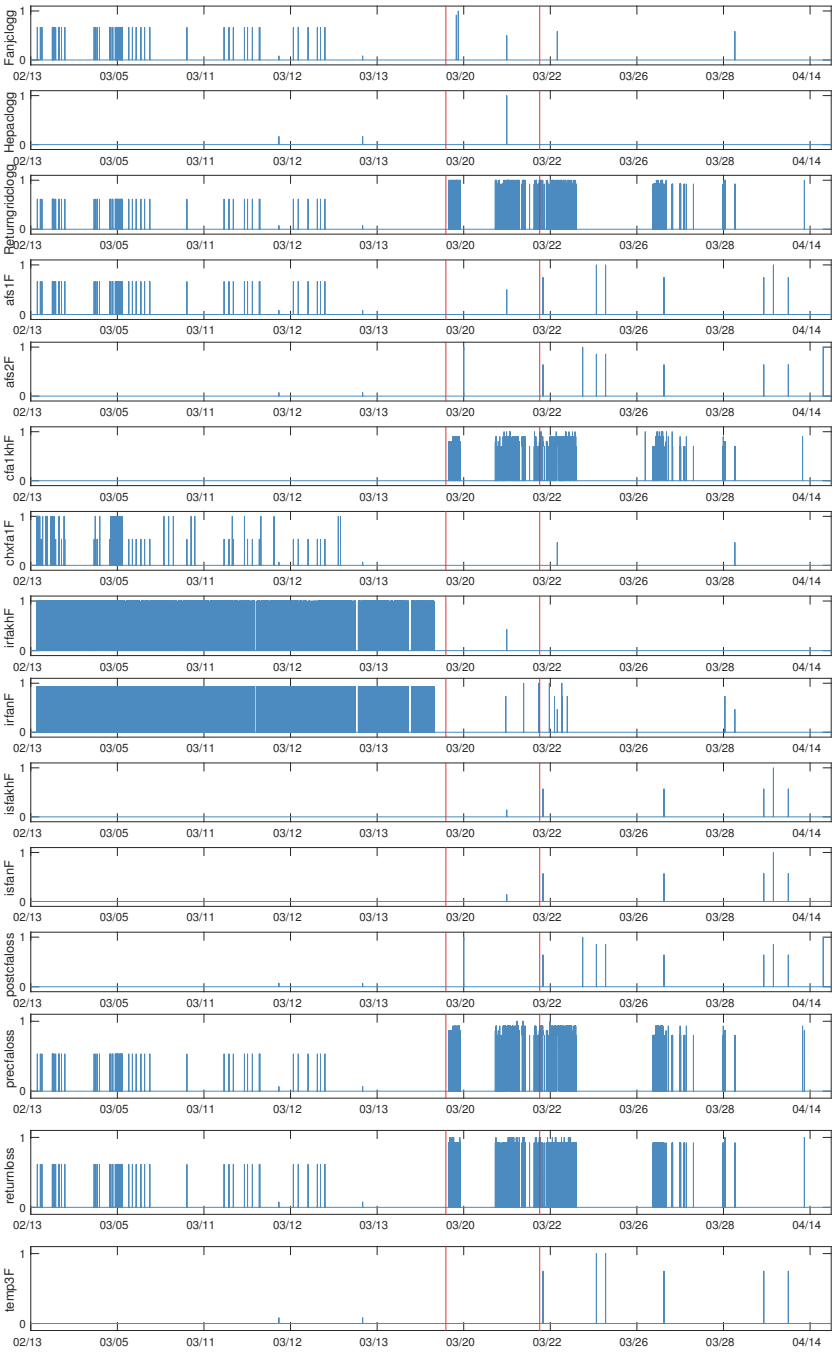

Figure 5.4: *Full overview of the alarms in the outset data set. The vertical lines indicate the start and end of the time period where within also unexpected clogging of the return grid has been reported.*

HEPA filter clogging

Another data set, covering only the 3rd of August 2009, begins with a fairly clear indication of issues in the post cabin region followed by quite certain alarms on the HEPA filter. This then stops about 2 am with hardly any alarms until almost 7 hours later where yet again issues are indicated in the post cabin fan region and finally uncertain alarms on quite many of the components in the system. The alarms over this set can be examined in Figure 5.6.

The data set itself is flagged as unexpected clogging of the return grid, which could be one of multiple faults available in this data set, but this is not clearly isolated and if this design is not producing large quantities of false alarms, is not the only fault.

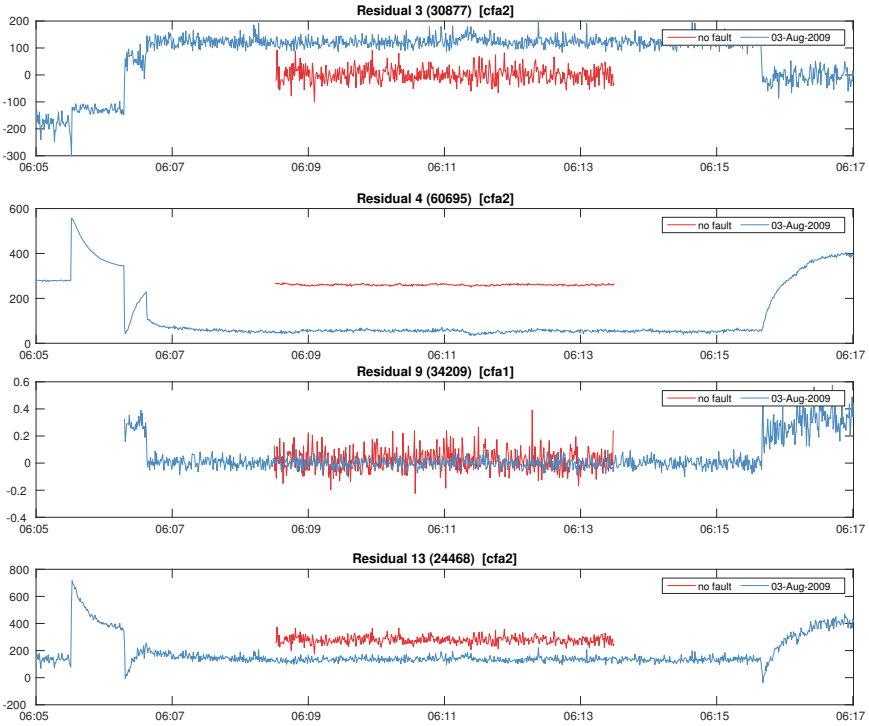

Figure 5.5: *Selection of residuals during the 3rd of August 2009.*

There are quite a lot of the test quantities triggering in the section where no faults are being isolated. Figure 5.5 displays a selection of the residuals during a portion of this time frame and clearly displays that something is detected and that this is a matter of evaluation rather than detection. Included in the graphs is also the no fault data used in the validation for easier comparation. This could indicate a passage with multiple faults. One possible explanation for this could be something faulty with the supply air, like for example smoke, that then propagates and cloggs most of the system. This has however not been confirmed.

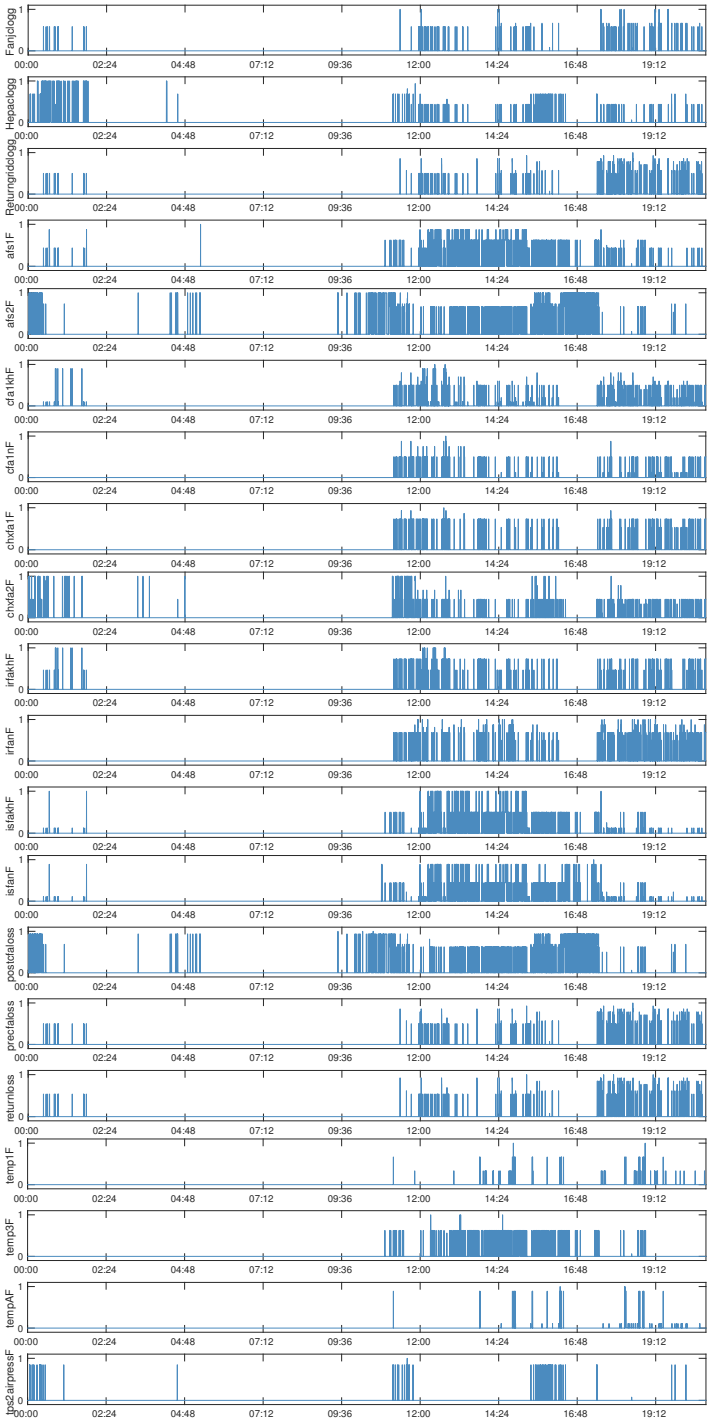

Figure 5.6: *Full overview of the alarms during the 3rd of August 2009.*

Another return fan failure

One data set, covering the 18th of January 2010, is documented as another unexpected clogging of the return grid. However this is not detected by this design. It does however, briefly, detect but not isolate faults of the parametric type before it quite certainly detects faults with the return fan.

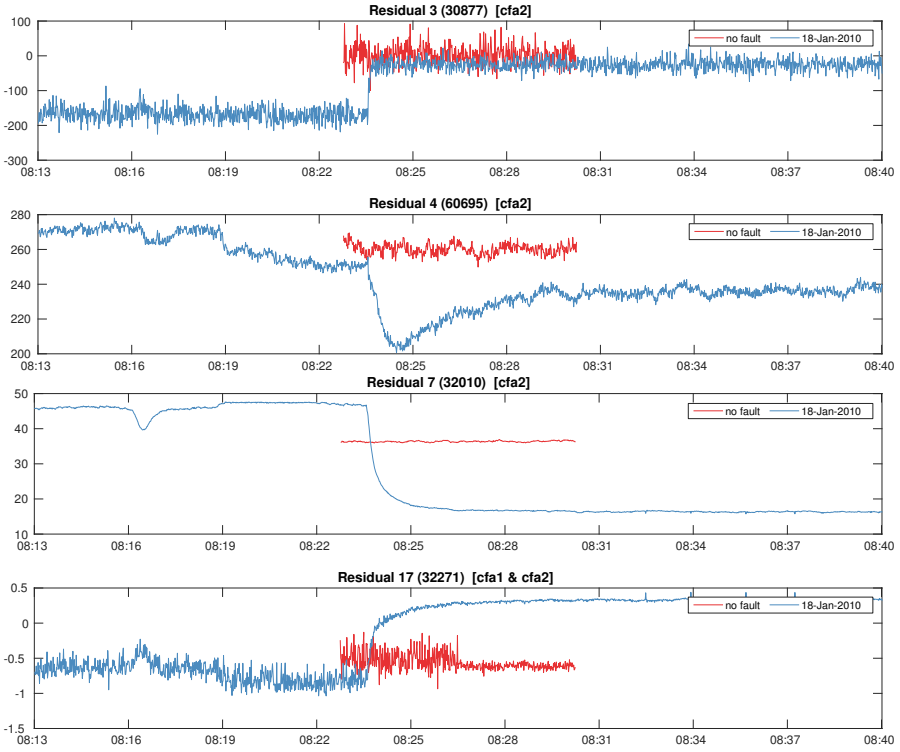

Figure 5.7: *Selection of residuals during the 18th of January 2010.*

When analysing the test quantities during the silent section of the alarms which are displayed in Figure 5.8 it is observed that many test quantities trigger about 8:23. A selection of the residuals highlighting this event is displayed in Figure 5.7. This indicate that there is something happening that the current evaluation cannot isolate which means that the alarming faults that follow could be a result of the measurements taken towards this fault.

Unknown data set

One of the data sets has no documentation of what type of fault could be present, if any. The analysis indicate return fan issues as the alarms displayed in Figure 5.9 tell. However, in general few test quantities trigger which could indicate that all of these could be false alarms.

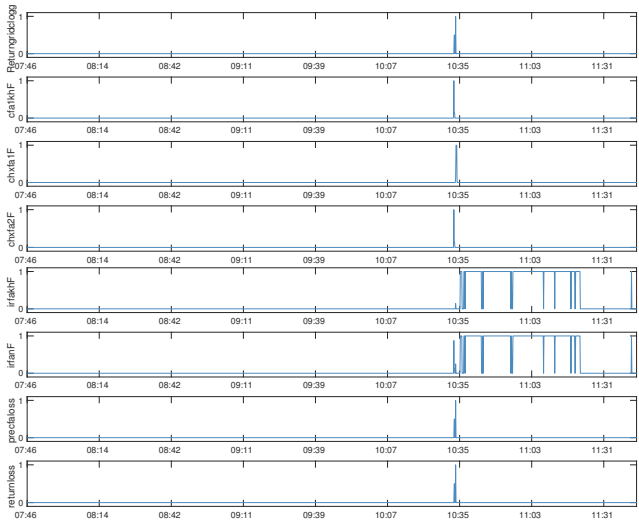

Figure 5.8: *Alarms from part of 18th January 2010.*

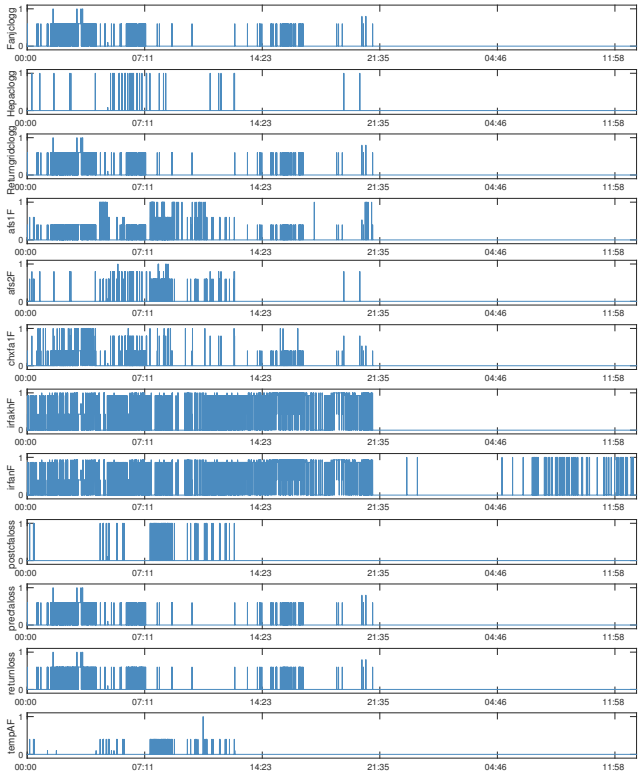

Figure 5.9: *Data set with unknown fault.*

No alarms

There was additionally a few sets where the design could not isolate any faults. They are all very similar and only triggers alarms very briefly, mostly on the airflow sensors and post cabin fan loss.

Two of the sets are marked as unexpected return grid clogging, one has documented return fan loss while the last one is supposed to have loss of the supply fan. Neither of these faults are isolated even though all set are indicated to only have one simultaneous fault.

However, many of the test quantities are either constantly on or triggering on and off repeatedly which could indicate that this section is actually suffering of multiple faults which is outside of the scope of the current design.

6
Conclusion

In this chapter an overall evaluation of the work will be presented followed by reflections and suggestions for eventual future work.

6.1 Overall evaluation

The expected results were to have an FDI system that detects and isolates faults in the air loop system that could work in real time and detect the faults within reasonable time.

On injected data the system is able to detect and isolate most of the modelled faults. However the models of leaks and cloggings are all parametric faults that, in practice, could be either leakage, clogging or wear.

The system can detect some of the indicated faults in the real data and even though it cannot isolate the faults perfectly, it can indicate the nature of the fault which, in practice, is sufficient to know which physical components that requires service. The proposed solution shows promising results but additional work could improve fault detection and isolation performance by, for example, consider multiple simultaneous faults during the evaluation process.

The system is based on a model of the air loop and is with fairly high certainty able to detect which mode the fans are operating in even though that information currently not is propagated to the fault detection and isolation level.

The methodology presented in [Svärd, 2012] was partially used and is considered to be working for this system as a whole. Further development following this methodology is believed to improve the current results when applicable.

The results on real data indicate that the methods used, even though they are yielding results, are not really sufficient for the design of an FDI solution for this particular system. This implies that there is a great potential for further improvement if more advanced methods for fault detection and isolation are explored.

6.2 Future work

This section will present suggested future work divided by task in a top-down fashion as it is suggested that future improvements of this design should be made in such an order. This as the underlying model itself is to be considered of desired quality and without need for immediate improvements.

Fault detection and isolation

The algorithms for detection and isolation should be extended so that multiple simultaneous faults are handled as there are indications that such exist in some, if not most, sets of real data. It is also recommended that more advanced test quantities are to be generated from the current selection of residuals.

One possible method to improve this could be fault detection and isolation by analysing the Kullback-Leibler divergence as was originally suggested for this work, and should be considered as one of the earlier paths for design improvement [Eriksson et al., 2011].

It could also be of interest to make the set of test quantities more dynamic as this model has different operation modes. The current solution is fairly basic and the method described in [Eriksson et al., 2012] could be considered for integration.

Residual generation and selection

For a model with a fair amount of equations, but more importantly, plenty of interconnections between them, the MSO count is quickly increasing to sizes that makes the residual generation and selection a very time consuming task. While this particular model is yet within reasonable size for modern computers, only a few more additions and it might not be.

Due to this are alternative, preferable automatic, methods of determining which MSOs to consider for further analysis. One way might be by analysing the residuals distinguishability [Eriksson et al., 2013], however, as the task to find the MSOs is close to the upper limits of feasibility this analysis must be brought closer to the model equations. My suggestion would be to analyse the signal (and fault) to noise ratio of the equations themselves and create weights that are to be used in the MSO search. This would need to be a directional analysis similar to how causality is handled in [Svärd and Nyberg, 2010] and is particular interesting in the quantified or otherwise non-linear equations. Due to this is it believed that this model could be an interesting set for this kind of analysis even if a (almost) complete MSO set already has been extracted.

This should not only assist in quicken up the search for MSOs but also should be able to assist in selecting how to structure the MSOs when designing residuals.

Modelling

In the current model the fan modes are estimated on the fly within the model by a simple test quantity, this is needed for making the fan curve equations valid. However, this test quantity is not propagated to the fault detection and isolation layer of the system. Stall mode is considered a fault and should be implemented and the detection of it ought to be improved. Multiple faults must be handled before this is propagated as stall mode occurs fairly often in the live data.

Additionally, in some sets of live data the model has a tough time deciding which mode the fans are in which in turn gives a very inconsistent output flow for the fans. Analysis indicate that this might not actually be a fault in the model as it not only fits the data fairly well but also fits that parallel fans easily can end up racing each other. In the real data available the sample rate is too low for a deterministic conclusion. However, due to this might it be interesting to smooth out the calculated air flows when problems to determine fan mode occur.

Other improvements of the model could be modelling of the dynamic properties and the parametric faults. There has also been some experimentation on making a Simulink model of the system with fairly decent results, but such an implementation could rapidly make the system infeasible in real time. However it is not suggested to focus on the underlying models unless no further improvements can be found in the actual detection and isolation layer of the FDI system.

Bibliography

Airbus Defense and Space. Airbus defense and space homepage, 2013a. URL https://airbusdefenceandspace.com/. Online; accessed 01/10/2016. Cited on page 1.

Airbus Defense and Space. Columbus celebrates fifth anniversary in space, 2013b. URL http://www.space-airbusds.com/en/news2/columbus-celebrates-fifth-anniversary-in-space.html. Online; accessed 01/03/2013. Cited on page 1.

Airbus Defense and Space. Utilisation and operation of the ISS, 2013c. URL http://www.space-airbusds.com/en/programmes/utilisation-and-operation-of-the-iss.html. Online; accessed 01/10/2016. Cited on page 1.

Daniel Eriksson, Mattias Krysander, and Erik Frisk. Quantitative stochastic fault diagnosability analysis. In *50th IEEE Conference on Decision and Control*, Orlando, Florida, USA, 2011. Cited on pages 5 and 46.

Daniel Eriksson, Erik Frisk, and Mattias Krysander. A sequential test selection algorithm for fault isolation. In *10th European Workshop on Advanced Control and Diagnosis*, Copenhagen, Denmark, 2012. Cited on page 46.

Daniel Eriksson, Erik Frisk, and Mattias Krysander. A method for quantitative fault diagnosability analysis of stochastic linear descriptor models. *Automatica*, 49(6):1591–1600, June 2013. Cited on page 46.

Erik Frisk, Anibal Bregon, Jan Åslund, Mattias Krysander, Belarmino Pulido, and Gautam Biswas. Diagnosability analysis considering causal interpretations for differential constraints. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, 42(5):1216–1229, September 2012. Cited on pages 4 and 27.

M. Krysander, J. Aslund, and M. Nyberg. An efficient algorithm for finding minimal overconstrained subsystems for model-based diagnosis. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 38(1): 197–206, 2008. Cited on pages 5 and 27.

F. Meinguet, A. SANDULESCU, X. Kestelyn, and E. Semail. A method for fault detection and isolation based on the processing of multiple diagnostic indices: Application to inverter faults in ac drives. *Vehicular Technology, IEEE Transactions on*, PP(99):1, 2012. ISSN 0018-9545. doi: 10.1109/TVT.2012.2234157. Cited on page 5.

E. Noack, W. Belau, R. Wohlgemuth, R. Müller, S. Palumberi, P. Parodi, and F. Burzagli. Efficiency of the columbus failure management system. In *AIAA 40th International Conference on Environmental Systems*, 2010. Cited on page 3.

E. Noack, T. Noack, V. Patel, I. Schmitt, M. Richters, J. Stamminger, and S. Sievi. Failure management for cost-effective and efficient spacecraft operation. In *Adaptive Hardware and Systems (AHS), 2011 NASA/ESA Conference on*, pages 137–144. IEEE, 2011. Cited on page 3.

E. Noack, A. Luedtke, I. Schmitt, T. Noack, E. Schaumlöffel, E. Hauke, J. Stamminger, and E. Frisk. The columbus module as a technology demonstrator for innovative failure management. Deutcher Luft- und Raumfahrtkongress, Berlin, Germany, 2012. Cited on pages 2 and 3.

T. Noack and I. Schmitt. A cyclic process model for monitoring mobile cyber-physical systems. 2012. Cited on page 3.

Dieter Sabath, Gerd Söllner, and Dirk Schulze-Varnholt. Development and implementation of a new columbus operations setup. 2012. Cited on page 2.

Dieter Sabath, Thomas Kuch, Gerd Soellner, and Thomas Müller. The future of columbus operations. In *SpaceOps 2014 Conference*, page 1618, 2014. Cited on page 2.

H. Sneider and P.M. Frank. Observer-based supervision and fault detection in robots using nonlinear and fuzzy logic residual evaluation. *Control Systems Technology, IEEE Transactions on*, 4(3):274–282, 1996. Cited on page 5.

Source: Airbus DS. Source: Airbus ds, 2013. URL `https://airbusdefenceandspace.com/`. Cited on pages 2, 7, 10, and 11.

Carl Svärd. *Methods for Automated Design of Fault Detection and Isolation Systems with Automotive Applications*. PhD thesis, Linköping University, Vehicular Systems, The Institute of Technology, 2012. Cited on pages 3, 5, and 45.

Carl Svärd and Mattias Nyberg. Residual generators for fault diagnosis using computation sequences with mixed causality applied to automotive systems. *IEEE TRANSACTIONS ON SYSTEMS MAN AND CYBERNETICS PART A-SYSTEMS AND HUMANS*, 40(6):1310–1328, 2010. Cited on page 46.

Carl Svärd, Mattias Nyberg, and Erik Frisk. Realizability constrained selection of residual generators for fault diagnosis with an automotive engine application. *IEEE Transactions on Systems, Man and Cybernetics: Systems*, 43(6): 1354–1369, 2013. Cited on pages 5 and 29.

Carl Svärd, Mattias Nyberg, Erik Frisk, and Mattias Krysander. Data-driven and adaptive statistical residual evaluation for fault detection with an automotive application. *Mechanical systems and signal processing*, 45(1):170–192, 2014. Cited on page 5.