Erik Hollnagel, Jeanette Hounsgaard & Lacey Colligan

FRAM – the Functional Resonance Analysis Method

– a handbook for the practical use of the method

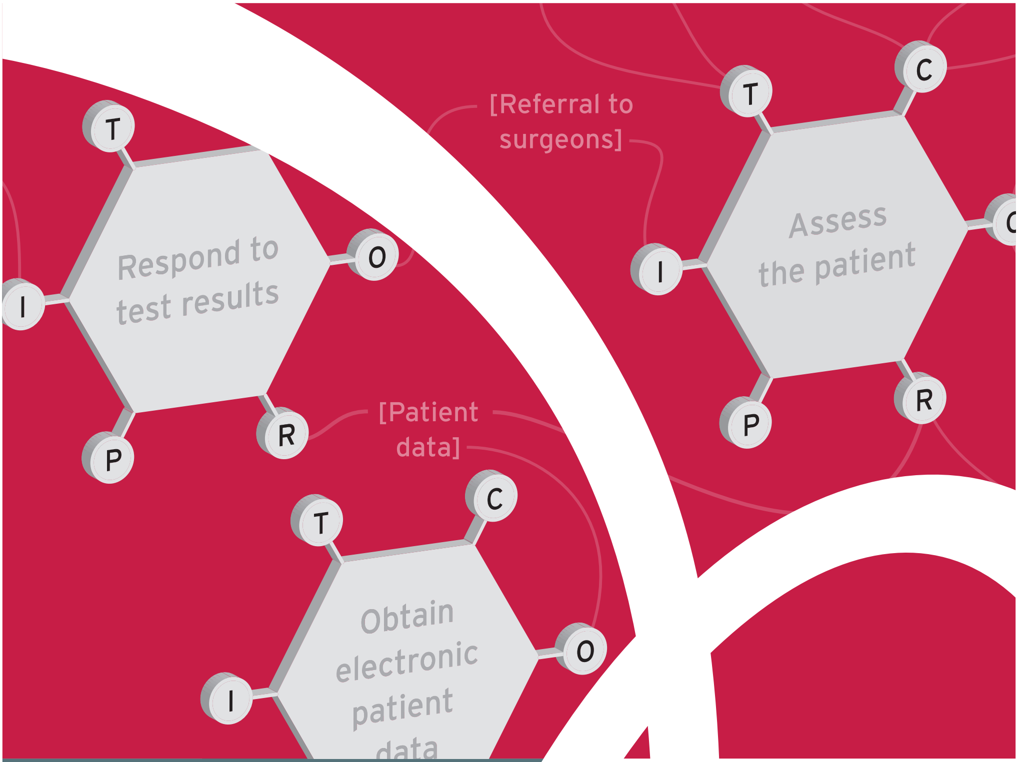

First edition, june 2014

Centre for Quality

Region of Southern Denmark

Hollnagel, E; Hounsgaard, J; Colligan, L: FRAM – the Functional Resonance Analysis Method – a handbook for the practical use of the method

The handbook is developed, published and financed by Centre for Quality,
in the Southern Region of Denmark.

ISBN 978-87-997477-0-2

Graphic production: ProGrafisk ApS – www.prografisk.dk

Other handbooks developed, published and financed by Centre for Quality:

Interne survey i Region Syddanmark (only in Danish)
Center for Kvalitet og Region Syddanmark 2014

Global Trigger Tool (only in Danish)
Center for Kvalitet, Region Syddanmark og OUH, 2008

SPC – Statistisk Proces Kontrol (only in Danish)
Center for Kvalitet, Region Syddanmark m.fl., 2008

For further information:
www.centerforkvalitet.dk

Foreword

This is the first handbook that Centre for Quality has published in the English language. It is a translation of a corresponding handbook in Danish, used to train clinical staff in the FRAM method.

Three authors made this possible: Lacey Colligan, Jeanette Hounsgaard and Erik Hollnagel. They are introduced in the beginning of this book.

The handbook is based on Centre for Quality's experience using the FRAM in hospitals in the Southern Region of Denmark. The FRAM has been used for both retrospective and prospective analyses and to improve the design of systems so that quality and patient safety is ensured.

We want to share our experiences with the FRAM outside Denmark and hope that the handbook will inspire others to contribute to quality and safety all over the world.

Middelfart, Denmark, 15 June 2014

Arne Poulstrup
CEO Centre for Quality

Table of contents

Presentation of the authors

Lacey Colligan MD MSc FAAP is an American pediatrician with expertise in neonatology and intensive care. She obtained a Masters in Patient Safety from University College London and has continued her academic work in human factors and safety research. She worked on the translation of the Danish Handbook to English as a visiting scholar to the Centre for Quality.

Jeanette Hounsgaard, BSc in Mechanical Engineering, diploma in Human Resource Development, IRCA Lead Auditor and lecturer at the University of Southern Denmark, is deputy manager and developer at the Centre of Quality, Region of Southern Denmark. For more than 30 years she has worked with Quality and Risk Management – mainly focusing on external environment and safety of workers – in off-shore, natural gas and chemical industries. Her professional field has been quality and patient safety at the public hospitals and in Psychiatry, Region of Southern Denmark, for the last 7 years. Since 2011 she has worked together with Erik Hollnagel, using the FRAM in the clinical practice to analyze adverse events, to prepare risk assessments and to improve the quality by changing the design of systems.

Erik Hollnagel, PhD, is chief consultant at the Centre of Quality, Region of Southern Denmark, professor at the University of Southern Denmark and emeritus professor at the University of Linköping in Sweden. He has worked at universities, research centers, and industries in Denmark, England, Norway, Sweden and France, focusing on problems from many domains, including nuclear power generation, aerospace and aviation, software engineering, land-based traffic, and healthcare. His professional interests include industrial safety, resilience engineering, patient safety, accident investigation, and understanding large-scale socio-technical systems. Dr Hollnagel had published widely and is the author/editor of 20 books – including four books on resilience engineering – as well as a large number of papers and book chapters.

Reading guide: How to get started with this book?

The purpose of this handbook is to provide practical guidance in the use of the FRAM. It is *not* intended to be a complete description of the method, and it is strongly advised that it is not used unless the analyst is familiar with the theoretical basis for the FRAM. (Hollnagel 2012).

Introduction

When something unexpected happens, a natural reaction is to try to find an explanation. This is especially true if the unexpected is undesirable or if an expected result is not achieved. The need to find an explanation is present both in activities that are complicated and take a long time, such as building a bridge or renovating a hospital, and in activities that are simpler and of shorter duration, such as an accident at work or the hospitalisation and treatment of a patient.

Finding an explanation for the unexpected can serve multiple purposes, such as relieving uncertainty and setting the mind at ease, assigning responsibility (or blame), taking steps to prevent something from going wrong again, etc. In many cases the search is for a cause. But finding a cause and finding an explanation are not quite the same. While a cause usually is part of an explanation, it is possible to develop an explanation without pointing to a cause. We may, for instance, have understood what has happened without being able to explain *why* it happened. Unfortunately a clear distinction between causes and explanations is not always maintained in practice, and this may sometimes make communication difficult.

Looking for causes

We try to find causes of adverse events because we take for granted that knowledge of causes is necessary to prevent things from going wrong in the future. This is based on the dictum that if there is an effect then there must also be a cause. This idea was already proposed by Leucippus of Miletus (ca. 480 – ca. 420 BCE), who is supposed to have said that "nothing happens in vain, but everything from reason and necessity". In practice causes are, however, not found but constructed by mutual agreement. A cynical definition of a cause is that it represents the mutual determination of a number of aspects or factors that are considered to be necessary and sufficient to explain why something happened. It should be possible to link a cause to a known part or function of a system (people, components, procedures, etc.). It should also be possible to do something about it to eliminate or reduce the effects – with a reasonable consumption of funding and time, of course. Finally a cause should be consistent with the generally accepted notion of how something happens.

Causes, in other words, are relative rather than absolute. The causes we use reflect the prevailing technology and current understanding of how the world works. For a long time, from the late 1700s to mid-1900s, it was assumed that causes should be found in the world of technology; accident analyses historically began by looking for failures of technical components and functions. More recently, i.e., since the middle of the 20th century, causes have almost always been linked to humans. This is partly related to the realisation that technology does not work by itself, but that it always is part of a human-machine system or a socio-technical system. The understanding of the human role in human-machine systems (particularly in relation to accidents) began with a concept of the human as a possible risk or as a "fallible" machine – e.g., something which was prone to accidents, a term that was used as early as 1914. This focus on the human factor intensified slowly over the following decades and was given a further boost by the accident at the Three Mile Island nuclear power plant in Harrisburg PA in 1979. During the 1990s, however, it gradually became clear that using "human error" as a simple and unambiguous explanation was not tenable. Instead, we began to understand how different situation factors could affect human activities and thus directly or indirectly force "human errors". (See also the following section on Safety-I and Safety-II.)

New types of accidents have generally speaking been met with suggestions of new types of causes, without questioning the basic assumption of causality. For centuries, we have become accustomed to explaining accidents by means of cause-effect chains – either simple or complex. We stubbornly hold on to that tradition, even though it becomes less and less suitable as time goes by.

Since the mid-1990s it has, however, gradually been realised that it is a mistake to work with one type of causes or explanations for actions that go well and another for actions that go badly. It is far more reasonable to assume that both in principle are much the same. A practical consequence of this assumption is that we should stop looking for causes and instead try to understand how work is adjusted to the situation and what the consequences – both positive and negative – of such adjustments may be. The FRAM (Functional Resonance Analysis Method) has been developed with this purpose in mind. The method consists of four steps:

- Identify and describe the important system functions and characterise each function using the six basic characteristics (called aspects). Together, the functions constitute a FRAM model.
- Characterise the potential variability of the functions in the FRAM model, as well as the possible actual variability in one or more implementations of the model.
- Determine the possibility of functional resonance based on dependencies / couplings among functions given their potential / actual variability.
- Develop recommendations on how to monitor and influence the variability, either by attenuating the variability that can lead to undesirable results, or by enhancing the variability that can lead to desired results.

The purpose of this guide is to show how the method can be used in practice. For an expansion of the FRAM, including the theoretical background, the reader is referred to Hollnagel (2012).

Chapter 1: An Example of a FRAM Analysis

This chapter provides a brief description of a FRAM analysis carried out in the Danish health service. The example will be referred to throughout the handbook to illustrate central concepts and processes of the FRAM.

The example describes a Patient with a Spinal Fracture, drawn from the Danish Patient Safety Database (DPSD) where it was reported as inappropriate treatment of a patient. The patient was unknowingly part of four parallel clinical investigations that were mostly independent of each other. The patient's clinical pathway involved seven independent parties including one General Practitioner (GP) office and three hospital units:

- General Practitioner, ambulatory office (Party 1)
- Hospital 1: an outpatient osteoporosis clinic; a X-ray department; and a surgical cancer unit (with its own CT) (Parties 2-4)
- Hospital 2: an out-patient Spine Center that is comprised of a triage, or screening unit, and medical spine treatment department (Parties 5 and 6)
- Hospital 3: a unit for spinal surgery (Party 7). This hospital was not involved in the incident report.

Three imaging modalities were involved in the patient evaluation:

- DXA scan (Dual X-ray absorptiometry) at Hospital 1's osteoporosis clinic [Pathway 1]
- an X-ray at Hospital 1 [Pathway 2]
- a CT-scan (computed tomography) at Hospital 1 [Pathway 3].

The Incident Report in the DPSD was filed by the medical spine treatment department at Hospital 2:

Incident report

"The patient had an unstable fracture of L1 seen on X-ray [date] and CT scan [date]. Despite this finding, the patient was referred for general medical assessment of his back pain.

Impact

The reading on the x-ray and CT were not recognised and the patient was dismissed while still in considerable pain.

Recommendation for Prevention

The triage evaluation must be more thorough."

Patient History in Chronological Order

12/02/11	Patient falls in the bathtub and hurts his back.
12/29/11	The Patient is referred from the GP to Hospital 1 for suspicion of osteoporosis. The DXA scan confirms osteoporosis and also collapse of T12 (Thoracic level 12). Note: this referral was not related to the Patient's report of the bathtub fall [Pathway 1].
Precise date unknown	The Patient sees the GP and complains of back pain after his bathtub fall. The GP refers the Patient to X-ray at Hospital 1 specifically for suspicion of traumatic spinal fracture. [Pathway 2].
1/06/12	Report of Patient's X-Ray from Hospital 1 indicates compression of T12. There is no sign of an unstable fracture of T12. The report is sent to the GP.
3/19/12	After another visit, the GP refers the Patient to the Surgical Department of Hospital 1 for evaluation for suspected cancer. The Surgical department at Hospital 1 requests a CT scan of the thorax and abdomen as a standard part of a work-up. [Pathway 3]. The CT scan confirms the patient does not have cancer but the radiologist notices an incidental finding of an unstable compression fracture of L1 (Lumbar 1). The report is sent to the GP but the GP only notices that there is no cancer (it is "ruled out"). The GP does not notice the incidental finding of the unstable fracture at the bottom of the report.
3/29/12	The Patient returns to the GP because he is still experiencing back pain. The GP is not comfortable managing this pain on his own and refers the Patient for specialised evaluation of back pain at the out-patient Spine Center at Hospital 2 [Pathway 4]. The Patient is screened and referred to the medical spine treatment department of the out-patient Spine Center at Hospital 2.

4/24/12	In the medical treatment department at the out-patient Spine Center at Hospital 2, the patient is interviewed and reports that he has increasing gait unsteadiness and "foot-drop" (a sign of a spine-related neurological problem) since December 2011.
	The medical treatment department at the out-patient Spine Center Hospital 2 recognises the Patient's critical spinal problem and urgently refers the Patient for specialised in-patient orthopedic surgery at Hospital 3. [Pathway 5].

This representation of the patient's pathway as a timeline is common but misleading because the chronological order does not make the dependencies clear. The patient was treated through five different pathways:

- suspected osteoporosis [Pathway 1]
- suspected fracture [Pathway 2]
- suspected cancer [Pathway 3]
- out-patient evaluation for back pain [Pathway 4] and
- confirmed surgical spinal injury [Pathway 5].

Only the first four pathways were mentioned in the DPSD Incident Report. They are also the ones included in the FRAM analysis.

Patient History by Pathway

Pathway 1: Suspected osteoporosis.
On December 29, 2011, the GP refers the patient for a DXA scan at Hospital 1 for a suspicion of osteoporosis, unrelated to the fall in the bathtub. The study confirms that the Patient has osteoporosis and compression of T12.

Pathway 2: Suspected fracture.
On December 2, 2011, the Patient falls and hurts his back in the bathtub. On an (unknown) date, the GP refers the Patient to X-ray at Hospital 1 and the study indicates full compression of T12, but there is no sign of an unstable fracture.

Pathway 3: Suspected cancer.
The GP refers the patient to the surgical department of Hospital 1 for suspicion of gastro-intestinal cancer on March 19, 2012. The CT scan of the abdomen and thorax rules out cancer. However, the compression fracture of L1 (as opposed to T12 on the X-Ray and DXA scan) is reported as an incidental finding.

Pathway 4: Out-patient evaluation for back pain.

The Patient remains in pain and is persistent in his return to the GP office. The GP refers him to the out-patient Spine Center at Hospital 2 on March 29, 2012, for general evaluation of back pain.

Patient's History in terms of what actually happened

If you only look at the Patient's case chronologically, you would think that the unstable compression fracture was present all along. The purpose of the analysis would therefore be to understand why this important finding was missed through the many patient encounters.

If you only consider the Patient's case by means of the pathways, you might think that no one had an overview of the Patient's medical history and condition. The purpose of the investigation would therefore be to understand how a GP could misunderstand the work-up results for the Patient.

However, you can also consider the Patient's case by focusing on the processes, rather than the results. This kind of analysis acknowledges that in most cases this type of patient evaluation goes right. Instead of asking why things went wrong in the specific case, we should ask why it did not go right as usual. In consequence of this, one must begin by describing and understanding how these patient pathways or treatment usually works. This description will provide the basis of understanding *how* this Patient's course differed from what usually happens. Later on in this handbook, we will show how the use of the FRAM provides this type of description.

Chapter 2: The FRAM and Safety Thinking

This chapter introduces *Resilience Engineering* and *Safety-II*, a new way of thinking about safety and safety science. A system is said to be resilient if it can adjust its functioning prior to, during, or following events (changes, opportunities, and disturbances) and thereby sustain required operations under both expected and unexpected conditions. The focus of resilience engineering is therefore both situations where things go wrong and where things go right.

The FRAM is an analysis tool that reflects Resilience Engineering and Safety-II thinking. The first comprehensive description of the basic principles of the FRAM is found in the book *Barriers and Accident Prevention* (Hollnagel, 2004). The scientific background for these principles stems from general systems theory, the psychological and cybernetic research of the 1950s and 1960s, and the principles of structured software development from the end of the same period. The development of the FRAM coincided with the development of Resilience Engineering as an alternative to traditional safety thinking and the FRAM can be seen as a tool for this new way of looking at safety.

Safety-I and Safety-II

To most people safety means the absence of unwanted outcomes such as incidents or accidents. Safety is generically defined as the system quality that is necessary and sufficient to ensure that the number of events that can be harmful to workers, the public, or the environment is acceptably low – or even zero.

The starting point for safety concerns has historically been the occurrence of accidents or recognised risks (actual or potential adverse outcomes). Adverse outcomes – things that go wrong – have been explained by pointing to their causes, and the response has been to either eliminate or contain these. New types of accidents have been adressed by introduction of new types of causes. Because this has been effective in providing short-term solutions, we have become so accustomed to explaining accidents in terms of cause-effect relations, that we no longer notice it. And we cling tenaciously to this tradition, although it has become increasingly difficult to reconcile with reality.

The current state of affairs reflects a common understanding of safety, Safety-I, that defines safety as a condition where the number of adverse outcomes is as low as possible. Since the purpose of safety management is to achieve and maintain that condition, safety goals are defined in terms of a reduction of the measured outcomes over a given period of time.

An important underlying assumption of Safety-I is that we can understand our systems by looking at their components and how these function or malfunction (as failures, errors). Another important assumption is that the reasons why things go wrong are different than the reasons why things go right. If that was not so, a reduction of the causes of accidents would also reduce the chance of things going right. This is called the *hypothesis of different causes.*

During the last 50-60 years, socio-technical systems in our society – including health care – have become increasingly complex and difficult to understand and control. Methods built on the assumption that we can understand our systems can therefore no longer achieve the desired freedom from accidents. Although some parts still are linear and understandable, complexity has made it impossible to fully understand or control *all* parts of our systems. Our inability to fully understand a system and its parts completely undermines the Safety-I assumptions.

Safety-II accepts that we cannot understand all of our system, nor understand them entirely. Instead of looking only at adverse events, Safety-II focuses on everyday work and situations where things go right. This shift in focus towards things going right is analogous to the World Health Organisation's (WHO's) argument that "health is a state of complete physical, mental, and social well-being and not merely the absence of disease or infirmity". Safety-II does not define safety as the absence of adverse events, but as the presence of successful everyday functioning.

While Safety-II emphasizes the importance of understanding how everyday work succeeds, it is, of course, still necessary to spend time understanding failures. However, it is naïve to think that our systems work because they are well-engineered and that workers function the same way in each situation. When we try to understand what is really happening, it becomes apparent that the main reason that systems succeed is that people and organisations constantly adjust to what they do in a situation (see Chapter 3.1 on Approximate Adjustments).

Reactive and Proactive Safety Management

From a Safety-I perspective, the purpose of risk management is to minimise the number of accidents and adverse events. This leads to a reactive approach to safety management where the organisation responds when an unacceptable incident has occurred, on the assumption that if there is no incident then everything is fine. Root cause analysis (RCA) is an example of a reactive analysis. Reactive risk management strategies can in principle work if the frequency of adverse events is so low that they do not disrupt daily function. However, when adverse events happen so often that daily function suffers, then reactive risk management strategies are insufficient. Faced with too many events that disrupt daily work, the quality of reactive risk management is eroded and analyses and reactions to what happens will be delayed and degraded. Control of safety is thereby eventually lost.

Risk management is in principle *proactive*, because it it studies things that could potentially happen, that pose a risk, but that have not happened. Healthcare Failure Modes and Effects Analysis (HFMEA) is an example of that. However, risk management is also reactive because a risk analysis rarely is performed more than once, and practically never performed continuously.

From a Safety-II perspective, risk management should be proactive so that mitigation or prevention can take place before an injury occurs. A proactive safety management will make it possible to identify risks early enough to prevent an adverse event from occurring. Timely, anticipatory interventions are usually more straightforward and less onerous than interventions required when damage has occured and harmed the system.

Tractability of Systems

Classical safety thinking and methodologies are based on a number of assumptions rooted in regulatory requirements for major technology and industrial systems developed in the late 1950s. These assumptions are deeply entrenched in safety practices and are rarely made explicit. The four main assumptions are:

- Systems can be divided into meaningful elements (typical components or activities). This is also called the principle of decomposition.
- System functioning is bimodal. This means that components are seen as being in one of two states, either working (functioning) or not working (malfunctioning).
- The order or sequence of events is predictable or predetermined. This can be seen as a kind of weak – or wishful – determinism, where the convenient assumption that events develop as planned overrules how events develop in reality.
- Combinations of events are logical and understandable. Thus, combinations of events can be described by logical relations (conjunctions, disjunctions, and negations) so that one always can determine what the outcome is.

The last assumption is the foundation of the concept of comprehensibility or tractability. The degree of tractability is a measure of how well the system's "inner" functioning is known. That is, whether we can clearly describe or specify components, so that we can understand how they interact. It becomes more clear if we for a moment assume that the opposite is true. If we do *not* have a clear description of a system and / or if we do *not* know how it works, then it is impossible to conduct a risk assessment or to manage it effectively.

It is important for almost all forms of human activity to make our actions and our surroundings as comprehensible as possible – and thus also as predictable as possible. This is especially the case for the manufacturing sector, where the basic ideas of quality control were developed. But the larger systems become, and the more integrated they become, the more difficult it is to understand them. Most large socio-technical systems in our daily lives (transport, communications, finance, distribution – not to mention the entire health care system) are so complicated that tractability is inadequate. The differences between systems with high and low tractability are summarised in Table 1.

Table 1. High and low tractability

	High, adequate tractability	**Low, inadequate tractability**
Description	Simple descriptions that require minimal detail.	Comprehensive descriptions that require extensive detail.
Knowledge about the system	Operational principles are known.	Operational principles are only partly known.
Stability	The system changes slowly and it is possible to foresee most situations.	The system changes so quickly that the description never fits reality because it has changed. It is impossible to foresee all situations.
Inter-relatedness to other systems	The system can operate independently.	The system cannot operate independently.

Traditional methods and approaches to risk management tacitly assume and therefore indirectly require that systems are tractable. In the 1960-70s, when these risk methodologies were developed, this assumption was not unreasonable although it was not universally true. However, information technology (in particular) has transformed the complexity of our systems at an exponential rate and the established methodologies are no longer suitable for low tractability or intractable systems. Thus, we need a new approach and new tools.

Chapter 3: The FRAM: The key concepts

This chapter first presents the basic assumptions of the FRAM method. Following that, the chapter gives an introduction to the basic concepts used to build a FRAM model of a system or set of functions. This includes a short introduction to the main concepts that are needed to use a FRAM model as part of a wider analysis.

3.1 The basic principles of the FRAM

The FRAM is based on four principles or assumptions about how things happen. Some of these principles have already been alluded to above, but are described here in detail. The four principles are:

1. **The principle of equivalence** (of successes and failures): this is the assumption that different kinds of consequences do not necessarily require different kinds of explanations causes, but that the same explanation can be used in most – if not all – cases.
2. **The principle of approximate adjustments:** this is the assumption that people continuously adjust what they do so that the actions match the conditions.
3. **The principle of emergence:** this is the acknowledgement that not all results can be explained as having a specific, identifiable cause.
4. **The principle of resonance:** in cases where it is neither possible – nor reasonable – to base explanations on the cause-effect principle (causality), functional resonance can be used instead to describe and explain non-linear interactions and outcomes.

1. The principle of equivalence, or that successes and mistakes arise from the same processes

Explanations of why something has gone wrong typically rely on taking the system apart into parts or components, an approach known as *decomposition*. In both physical and socio-technical systems there are "natural" parts, such as people, machines, etc. In the description of a process or an event sequence, the decomposition results in a description of individual actions or steps in the process, for example, as they can be found in an instruction or a step-by-step guideline. The explanation of the adverse outcome relies on finding one or more parts or components in the situation that have failed or malfunctioned, or a step that somehow has been performed incorrectly. The consequence is in this way explained as a result of the malfunctions of a system and / or its components. This view is based on the *hypothesis of different causes*, which is the assumption that things that go right and things that go wrong do so for different reasons, i.e., that the two kinds of outcomes have completely different causes. The FRAM – and Resilience Engineering – takes another approach, namely that things that go right and and things that go wrong happen in much the same way. (This is also reflected in the definition of Safety-II, see above.) The fact that the results in the two cases are different does not mean that the underlying causes also must be different. The principle of approximate adjustments explains why this is so.

2. The principle of approximate adjustments

As described above, many socio-technical systems can only be partially understood. The actual conditions for work or for carrying out an action are therefore never completely in agreement with what was expected or prescribed. In order to carry out work it is therefore necessary to constantly adjust performance to fit the existing conditions (resources, time, tools, information, requirements, opportunities, conflicts, interruptions). These adjustment are made by individuals, by groups and by organisations and take place at all levels, from the performance of a specific task to planning and management. However, since resources (time, labour, information, etc.) almost always are limited, such adjustments will inevitably be approximate rather than precise. In most cases this is good enough because the situations rarely differ much from what usually is the case. The approximate adjustment are the reason why things mostly go right, but also the reason why they occasionally go wrong (see Figure 1).

Figure 1: The principles of approximate adjustment

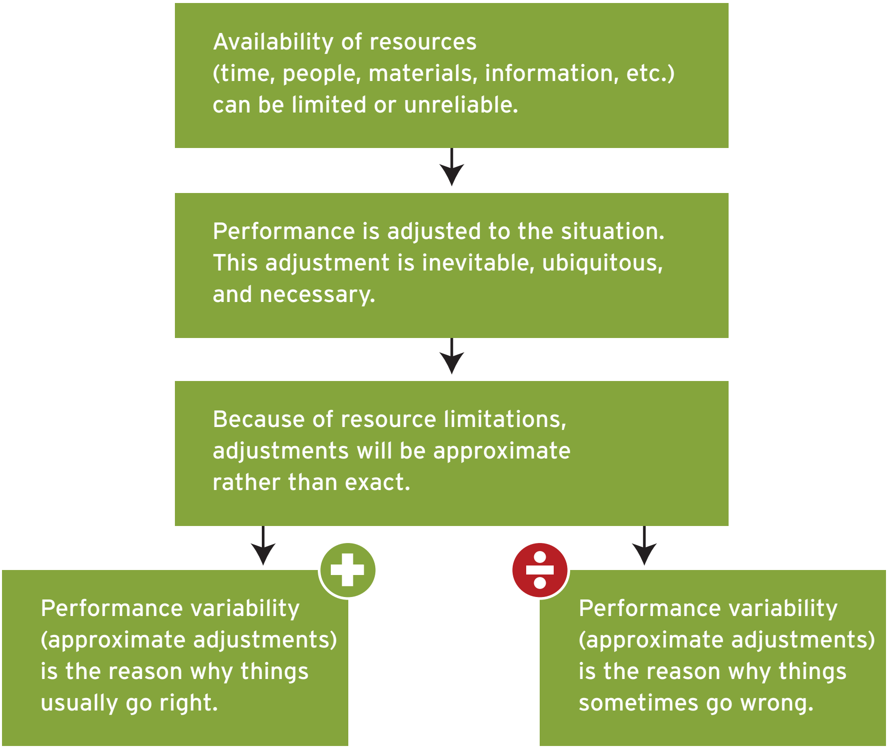

3. The principle of emergent outcomes

The variability that is due to the everyday adjustments is rarely large enough in itself to serve as the cause of something going wrong or large enough to be described as a failure. But if it is impossible to find a specific cause that explains the observed consequence, then one cannot rightly claim that the results were caused by certain conditions or situations (see Table 2). On the one hand there will in such situations be nothing that really went wrong, i.e., nothing that was so different from what commonly happens that it could justifiably be characterised as a failure or deviation. The performance adjustments were the ones commonly made, precisely because they had shown their value in the past. They were neither documented, nor can they be "found" again or reconstructed. The variability of multiple functions may on the other hand coincide and mutually affect each other in unexpected ways, thereby leading to unexpected and disproportionate impacts – both in a negative and a positive sense.[1] This way of describing or explaining how consequences arise is technically called non-linear.[2] Both errors (that which turns out badly) and everyday work (that which turns out well) can be explained as *emerging* from variability rather than as a *result* of it (or caused by it). In practice this means that the effects no longer can be attributed to or explained by referring only to malfunctions or defects in specific components or parts.

Table 2: Examples of resultant and emergent outcomes

	The outcome as resulting	**The outcome as emerging**
Finance	Interest rates and yields	Stock quotes / share values
The beach	Coastal erosion: every time a wave hits the coast, some sand is taken. (a result of the waves continuous exposure)	Wave patterns in the sand: whenever a wave hits the shore, sand ripples emerge, but we can neither explain their appearance, nor predict it.
Winter weather	Ice formation: when the temperature drops below 32F or 0C, water will freeze as a result.	Snowflake formation: when water crystals fall through freezing temperatures, we do not know what form the snowflake will take.
Nature	Young birds (or eggs?) – a result of natural reproduction.	The pattern of a flock of migratory birds flying (but may be not geese and swans in formation)

---

[1] The technical name for this phenomenon is coupling.

[2] A non-linear effect has two characteristics. Firstly, there is no proportionality between "cause" and "effect". Secondly, the effect cannot be explained by causal (linear) thinking.

Figure 2 shows resultant outcomes can be seen as the determinated one or several identifiable causes. The outcom itself is a (relatively) permanent change to a system state. Causes are here considered as components or functions that in some way failed or did not work correctly. Since many of these causes correspond to semi-permanent system conditions, they can be found by tracing backwards from the consequence. This can either be done when the event is analysed, or their existence can be deduced and possibly verified afterwards.

Figure 2: Resultant outcome

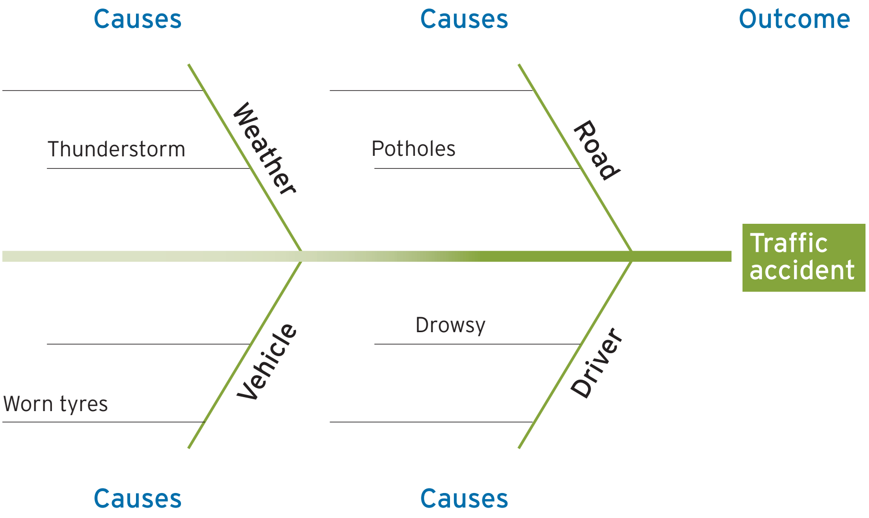

Figure 3 shows how the emerging outcome can be seen as being produced by unstable (short-term) combinations of states and events. The consequences can however not be explained as an effect of specific components or functions. Instead the incident occurs because of conditions that are transient or temporary.

Figure 3: Emergent outcomes

The »causes« are configurations or combinations of states and events that existed at a certain point of time. Their existence may be inferred, but they cannot be found.

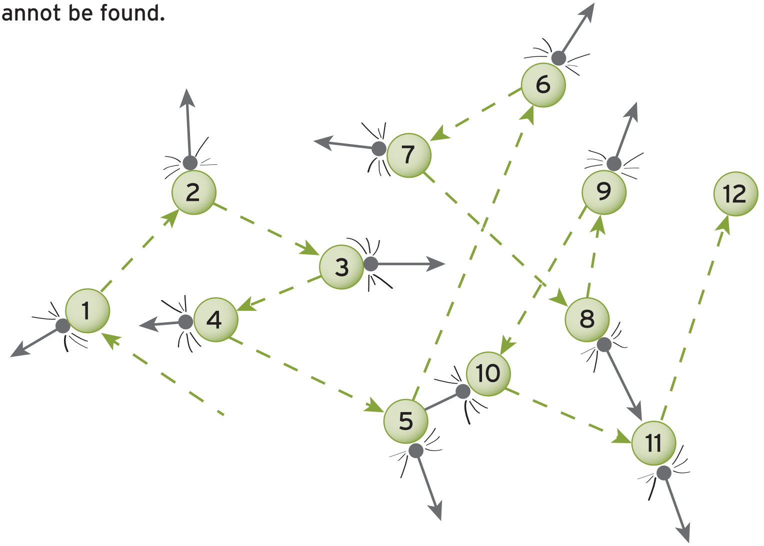

The outcome is a (relatively) stable change in the system or its parts.

4. The principle of resonance

The variability of a number of functions may sometimes coincide, i.e., they may mutually influence each other. This can cause the amplitude of one or more functions to become unusually large (leading to either positive or negative outcomes). The consequences of such increased variability may spread to other functions in analogy with the phenomenon of resonance. It is thus no longer a question of single or multiple cause-effect chains, because that would imply that one could speak of one or more specific and recognisable causes (something that went wrong). The impacts instead emerge, i.e., they appear in a manner that cannot be explained by or reduced to linear causality.

There are three types of resonance. In physical systems, *classical resonance* is the phenomenon that a system can oscillate with larger amplitude at some frequencies than at others (see Figure 4). These are known as the system's resonant (or resonance) frequencies. At these frequencies even small external forces that are applied repeatedly can produce large amplitude oscillations, which may seriously damage or even destroy the system. Classical resonance has been known at least since ancient Greece. A more recent form is *stochastic resonance*, has been understood since the early 1980s (see Figur 5). In stochastic resonance there is no forcing function, but rather random noise, which every now and then pushes a subliminal signal over the detection threshold. Stochastic resonance can be defined as the enhanced sensitivity of a device to a weak signal that occurs when random noise is added to the mix. The outcome of stochastic resonance is non-linear, which simply means that the output is not directly proportional to the input. The outcome can also occur instantaneously, unlike classical resonance which must be built-up over time.

A third type of resonance, which is central to the FRAM, is called *functional resonance*. As a phenomenon, functional resonance describes the noticeable performance variability in a socio-technical system that is the outcome of the multiple approximate adjustments that are the basis for everyday work activities. Functional resonance is the detectable outcome (or signal) that emerges from the unintended interaction of the everyday variability of multiple signals. The approximate adjustments comprise a small number of recognisable short-cuts or heuristics, which means that performance variability is semi-orderly and therefore also partly predictable. There is a regularity in how people behave and in how they respond to unexpected situations – including those that arise from how other people behave. The resonance effects that occur can be seen as a consequence of the ways in which the system functions, and the phenomenon is therefore called *functional resonance* rather than stochastic resonance. Functional resonance offers a way to understand outcomes that are both non-causal (emergent) and non-linear (disproportionate) in a way that makes both predictability and control possible.

Figure 4: Classical resonance

- A child swinging
- A guitar (acoustic or electric that is amplified)
- The balance wheel of a mechanical watch

Figure 5.

The Tacoma Narrows Bridge was subjected to strong winds that caused oscillations and the resulting resonance and swaying caused its eventual collapse.

3.2 Basic concepts in developing a FRAM model

The FRAM is a systematic approach to create a description or representation of how an activity (a piece of work, a sequence of actions) usually takes place. This representation is called a FRAM model. The selected event or performance is described in terms of the functions that are necessary to carry out the activity, the potential couplings between the functions, and the typical variability of the functions. The purpose of the FRAM is to provide a concise and systematic description of "normal" work, as it typically takes place.[3] What follows is a presentation of the basic concepts that you need to build or construct a FRAM model to describe how work typically is done.

3.2.1 Functions and aspects

What is a function?

A function in the FRAM represents the means that are necessary to achieve a goal. More generally, a function represents the acts or activities – simple or complex – that are needed to produce a certain result.

- A function typically describe what people – individually or collectively – have to do to perform a specific task and thus achieve a specific goal, for example, triage a patient or carry out medication reconciliation.
- A function can also refer to something that an organisation does: for example it is the function of an emergency department to treat incoming patients.

---

[3] The meaning of normal is not "normative", but simply that it is what you would expect to see in everyday – or normal – conditions.

- A function can refer to what a technical system does either by itself (an automated function, such as a dialysis machine) or together with one or more people (an interactive or socio-technical function, like an electronic health record).

A general rule in the use of the FRAM is that a function should be described by a verb (verb) if it is a single word, or a verb phrase – in both cases using the infinitive form. For instance "(to) diagnose a patient" rather than "diagnosing a patient" or "(to) order medication" rather than "ordering medication".[4]

How to find functions (analysis of relationships and dependencies)?

The following two examples demonstrates how to recognise the functions from a general description of a piece of work.[5]

The first example is a simplified description of a workflow. One part describes the delivery of medication to a department, while the other describes how the medication is dispensed to patients. This is the formal job description, hence neither a description of how the work is done in practice nor the description of a specific incident.

Example 1. Pharmacy

The pharmacy has an assistant, who oversees the electronic ordering and distribution of medications. Medications are delivered to the pharmacy section in sealed boxes, a delivery confirmation is signed, and the sealed boxes are placed in a locked room.

When medication is prepared for a specific patient, it is placed in a pill container with a patient ID label attached. Before the medication is administered, the label is scanned to make sure it is for the correct patient.

In order to receive their medication, patients go to, but do not enter, the medication room. The Clerk makes sure all patients receive their medication. The Nurse actually administers the medication.

If the Nurse leaves the medication room to take care of the patient during the process, they close the door and lock the medication room.

---

[4] In other words, use the infinitive form but leave out the "to".

[5] The example is taken from a report written by Karen Ørnebjerg and Benedicte Schou and is used with permission of the authors.

This work description includes many different actions. Each action corresponds to a function, and is typically characterised by a verb phrase. The main actions / functions are highlighted below. (Try to see if you can find more.)

> *The pharmacy has an assistant, who* ***oversees*** *the supply of medications. Medications are* ***delivered*** *to the pharmacy section in sealed boxes, a delivery confirmation is* ***signed****, and the sealed boxes are* ***placed*** *in a locked medication room.*
>
> *When the medication is prepared for a specific patient, it is* ***placed in pill container*** *with a patient ID* ***label attached.*** *Before the medication is administered, the* ***label is scanned*** *to make sure it is for the correct patient.*
>
> *In order to receive their medication, patients* ***go to,*** *but do not enter, the medication room. The Clerk* ***makes sure*** *all patients receive their medication. The Nurse actually* ***administers*** *the medication.*
>
> *If the Nurse* ***leaves*** *the medication room* ***to take care*** *of a patient during the process, the door to the medication room must be* ***closed*** *and* ***locked.***

After finding the tasks, we can describe them as activities (using verbs). This differs from a description of a state, such as "system ready", or a change of state, such as "system not ready". States and changes of states are described with nouns.

<oversee supply of medications>
<deliver medications>
<acknowledge (or sign) the delivery of medication>
<place the medications in the room>
<attach a label to the pill container>
<place medication in patient specific pill container with label attached>
<scan label to confirm right medication right patient>
<make sure all patients get medication>
<administer the medication to each patient>
<discontinue medication administration temporarily to attend to other>
<close door>
<lock door>

These twelve functions can be used as a basis for describing how the work is done in practice. This description may in turn be used to analyse a specific pathway or event. This is done by describing each function in detail, including who performs it and as many of the six aspects (see below) that are relevant to understand how the function can be performed.

Example 2. Patient discharge

The second example is based on a description of a number of conditions that have been found in practice.[6] *It therefore does not refer to a single process, but is rather a conglomerate of what has been found in a number of cases.*

A patient comes home from the hospital with medication dispensed in dosage boxes for three days and an accompanying medication list. The number of tablets in the boxes does not match the number of tablets on the medication list.

The cardiac medications that the patient usually receives do not appear on the medication list. The patient was prescribed pain meds for a fall during the hospitalisation, but they were not included in the discharge medication list. There is no follow-up for the medications after discharge.

To begin a FRAM analysis, one must first describe what should happen when a patient is discharged.

Before the discharge, the patient's medication requirements must be determined.
The patient should receive the prescribed medication in the correct quantity.
A list should be prepared of medications that correspond to each drug and given to the patient.

The next step is to describe the key actions in the same way as in the first example. This could look like this:

---

[6] The example is taken from a report written by Ann-Kathrine Fog, Merete Larsen and Birte Fester and used with permission of the authors.

Before the ***discharge,*** *it is necessary to* ***determine*** *the patient's medication requirements.*

The patient must ***receive*** *the medications* ***packaged in the correct quantity.*** *The hospital should also* ***prepare a list of medications*** *that correspond to each drug.*

From this the following functions can be described in the same way as before, i.e., as activities (using verbs) rather than states.

<discharge a patient>
<determine patient's medication requirements>
<package medication in the correct quantity>
<prepare medication list for discharge>

These four functions provide the basis for a description of how work is done in practice. This description can in turn be used to analyse a specific pathway or event. This is done by describing each function in detail, including who performs it and as many of the six aspects (see below) that are relevant to understand how the function can be performed.

What are the FRAM aspects?

In the FRAM, functions are described by means of six aspects:[7] Input, Output, Requirements, Resources, Control, and Time.[8] The general rule of the FRAM is that an aspect should be described when it is seen as necessary or appropriate by the analysis team, provided there is sufficient information or experience to do so. It is thus not necessary to describe all six aspects of every function, and it can indeed sometimes be either impossible or unreasonable to do so.

The guidelines for how and when to describe aspects are given in the following. As a minimum, at least one Input and one Output must be described for all foreground functions. (The meaning of a foreground function is also explained in the following.) Note, however, that the FRAM model is reduced to an ordinary flow chart or network

---

[7] The aspects can be thought of as features or dimensions of a function.

[8] In the FRAM, the six aspects are written with capital initials when they refer to aspects as part of a functional description (e.g., Time or Control), but without the capital initial when used in other contexts (e.g., timely or effective control).

diagram if only the Input and Output aspects are described. A general rule in the use of the FRAM is that an aspect is described with a noun (noun), or a noun phrase. In other words, an aspect is described as a state or as a result of something – but not as an activity.

A brief description of the six aspects

Input. The Input to a function is traditionally defined as that which is used or transformed by the function to produce the Output. The Input can represent matter, energy, or information. There is, however, another sense of the term Input that is just as important for the FRAM, namely the Input as that which activates or starts a function.[9] The Input in this sense may be a clearance or an instruction to begin doing something that must be detected and recognised by the function. Input can be seen as a form of data or information, or more generally as a state change that is recognised by a function as a signal to begin. Formally, an Input is always the result of a change in the condition of something, whether it is energy, information, or position. It is for that reason that the description of the Input always is a noun or a noun phrase.

In the FRAM, designated foreground functions must have defined Inputs, while designated background functions need not have. The difference between foreground and background functions will be explained in the following.

Output. The Output of a function describes the result of what the function does, for example, the result of processing the Input. The Output may therefore represent material, energy, or information – an example of the latter would be a permission or clearance, or the result of a decision. The Output describes a change of state – of the system or of one or more output parameters. The Output may, for example, be the signal to start a function. The description of the Output should be a noun or a noun phrase.

Precondition. In many cases a function cannot begin before one or more Preconditions have been established. These Preconditions can be understood as system states that must be [True], or as conditions that ought to be verified before a function is carried out. A Precondition does, however, not itself constitute the signal that starts the function. An Input, on the other hand, can activate a function. This simple rule can be used to determine whether something should be described as an Input or as a Precondition. It is not critical for a FRAM analysis whether something is labelled Input or Precondition, as long as it is included in the model in one way or another. A Precondition must always be an Output from another function. The description of a Precondition should be a noun or a noun phrase.

---

[9] In such cases the Input functions are that which gives "green light" to start the function.

Resource (or Execution Condition). A Resource is something that is needed or consumed while a function is carried out. A Resource can represent matter, energy, information, competence, software, tools, manpower, etc. Time can, in principle, also be considered as a Resource, but since Time has a special status it is treated as a separate aspect.

Since some Resources are consumed while the function is carried out and others are not, it is useful to distinguish between (proper) Resources on the one hand and Execution Conditions on the other. A (proper) Resource is consumed by a function and will therefore be reduced as time goes by; an Execution Condition only needs to be available or exist while a function is active. (The difference between a Precondition and an Execution Condition is that the former is only required before the function starts, but not while it is carried out.) The description of a Resource (an Execution Condition) should be a noun or a noun phrase.

Two examples are blood plasma during surgery and competence in relation to an operation. Blood plasma is a (proper) Resource and the competence is an Execution Condition. After surgery, some of the blood plasma has been consumed, while competence after an operation will be the same – if not actually increased.

Control. Control, or control input, is that which supervises or regulates a function so that it produces the desired Output. Control can be a plan, a schedule, a procedure, a set of guidelines or instructions, a program (an algorithm), a “measure and correct ”functionality, etc. Another, less formal type of control is social control or expectations to how the work should be done. Social control can be external, such as the expectations of others (management, organisation, co-workers) or a person’s own expectations. Social control can also be internal, for example, when we plan a job and mentally go through when and how to do it, or when we imagine what others expect of us. The description of Controls should be a noun or a noun phrase.

Time. This aspect represents the various ways in which Time can affect how a function is carried out. Time, or rather temporal relations, could be seen as a form of Control, as when Time represents the sequencing conditions. A function may, for instance, have to be carried out (or be completed) before another function, after another function, or overlapping with – parallel to – another function. Time may also relate to a function alone, seen in relation to either clock time or elapsed time.

Time could also be interpreted as a Resource, such as when something must be completed before a certain point in time, or within a certain duration. Time could, of course, also be interpreted as a Precondition, e.g., that a function must not begin before a certain time or that it must not begin before another functions has been completed. Yet rather than having Time as a part of either of the other three aspects of a function it seems reasonable to acknowledge its special status by having it as an aspect in its own right. The description of Time should be a noun or a noun phrase.

Figure 6: The six aspects of a function or activity

TIME
Temporal aspects that affect how the function is carried out (constraint, resource).

CONTROL
That which supervises or regulates the function, e.g. plans, procedures, guidelines or other functions.

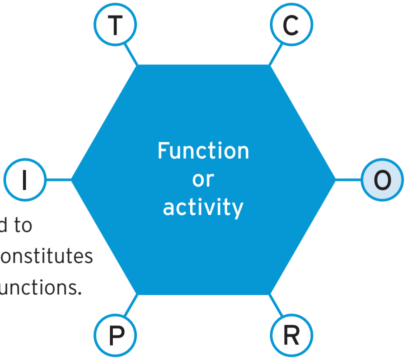

INPUT
That which activates the function and/or is used or transformed to produce the output. Constitutes the link to upstream functions.

OUTPUT
That which is the result of the function. Constitutes the links to downstream functions.

PRECONDITION
System conditions that must be fulfilled before a function can be carried out.

RESOURCES (execution conditions)
That which is needed or consumed by the function when it is active (matter, energy, competence, software, manpower).

3.2.2 Relationship between functions

All the functions of a FRAM model are characterised by the six aspects. If the same values (names) are assigned to aspects of different functions – for instance the Output of one function and the Precondition of another – then there is a potential dependency or coupling between the functions.

Couplings

The basis for traditional event and risk analyses is a description of relations. It may be the structure of sub-tasks in a hierarchical task analysis, or the order of (sub-)events on a timeline. Relations are often described by means of route or network diagrams (“boxes and arrows”), of which there are many different forms. The starting point for a traditional analysis is the specific relations between elements such as cause and effect, part-whole, goals-means, etc.

The basis of the FRAM is the description of the functions that make up an activity or a process. The description starts by the functions themselves, and not by how they are ordered or related. The relationship is not described directly by a graphical rendering, such as an event tree, but indirectly as relations defined by aspects of functions. The common technical term for such relations is coupling. If, for example, the Output of a function A also is defined as a Precondition for another function B, then the two functions are potentially coupled. One can also say that the function B uses the Output of function A as a Precondition. The couplings in a FRAM model are generally n-to-n (or many-to-many) rather than 1-to-1. For example, the Output from function A may also be Input to a function E. Function B can also have another Precondition that is Output from a function H.

Figure 7: Couplings for Function A

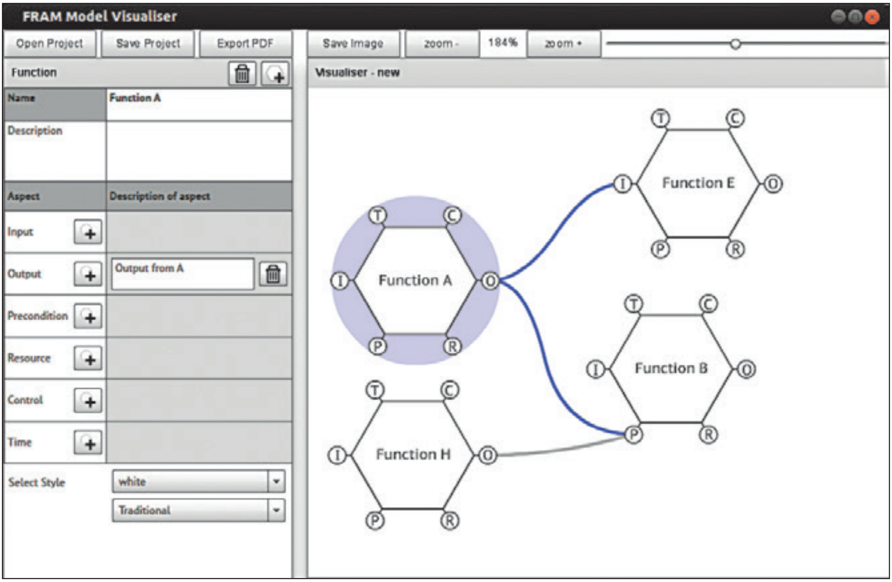

Name: Function A

Output: Output from A

Select Style: white

Traditional

Figure 8: Couplings for Function B

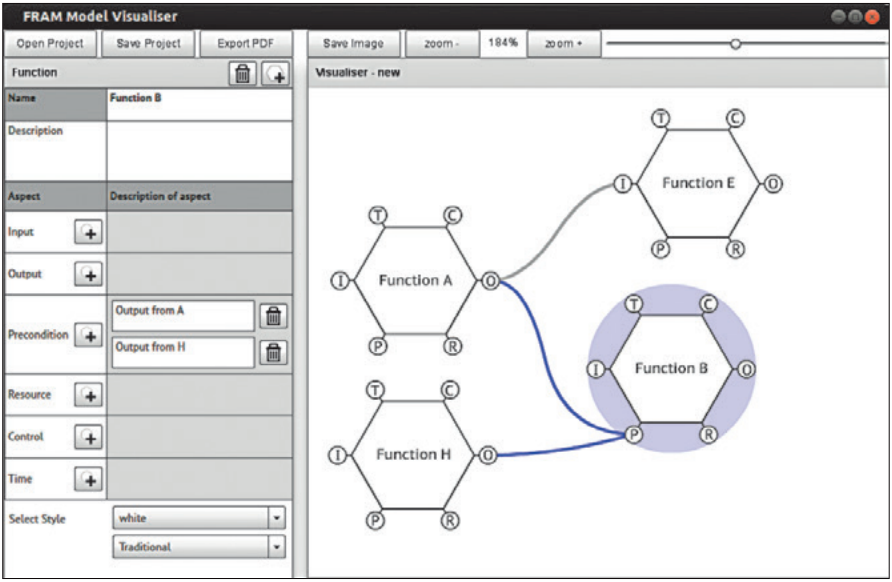

Name: Function B

Description:

Precondition: Output from A

Precondition: Output from H

Select Style: white

Traditional

Figure 9: Couplings for Function E

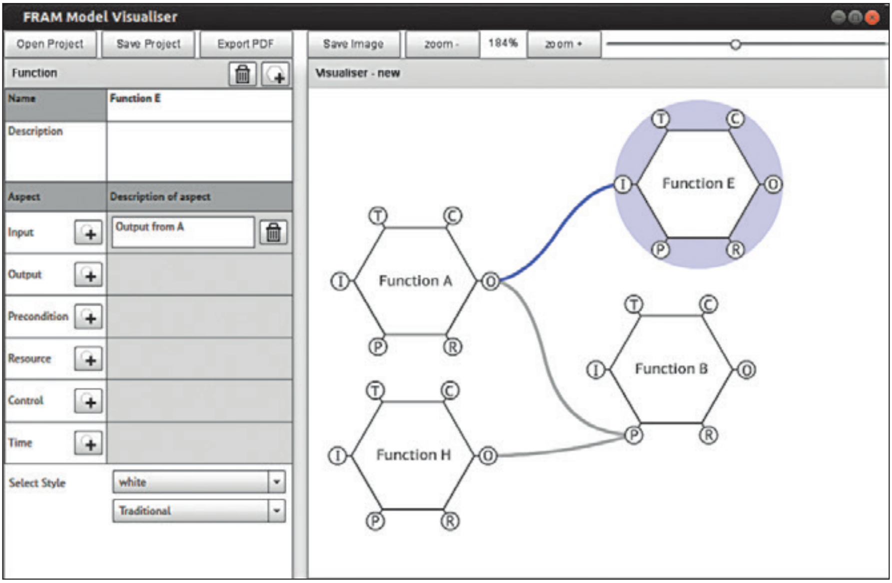

Name: Function E

Description:

Input: Output from A

Select Style: white

Traditional

Figure 10: Couplings for Function H

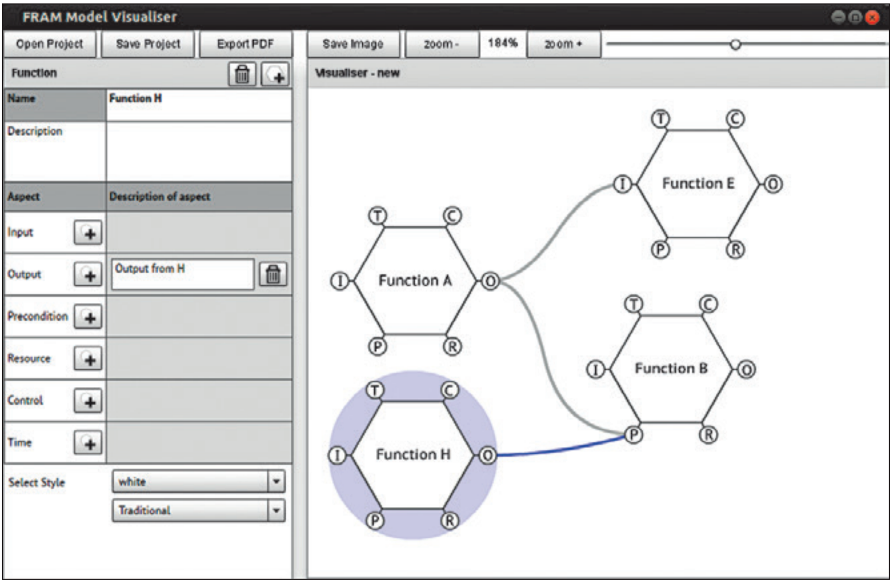

Name: Function H

Output: Output from H

Select Style: white

Traditional

The couplings that are described in a FRAM model, i.e., the dependencies that are the consequence of shared aspect attributes, are called potential couplings. This is because a FRAM model describes the potential or possible relationships or dependencies between functions without referring to any particular situation. In an instantiation of a FRAM model, only a subset of the potential couplings can be realised; these represent the actual couplings or dependencies that have occurred or are expected to occur in a particular situation or a particular scenario.

Figure 11 shows how the functions **<Respond to test result>** and **<Assess the patient>** are interconnected. One coupling is that the Output from **<Respond to test results>** serves as Input to **<Assess the patient>**. This means that the variability of the Output from the **<Respond to test result>** can affect the performance of **<Assess the patient>.** Another connection is that both functions have **[Patient data]** as a Resource. This means that the variability of the Output from **<Obtain electronic patient data>** can affect the performance of the **<Respond to test result>** and **<Assess the patient>.**

Figure 11: Illustration of couplings between functions.

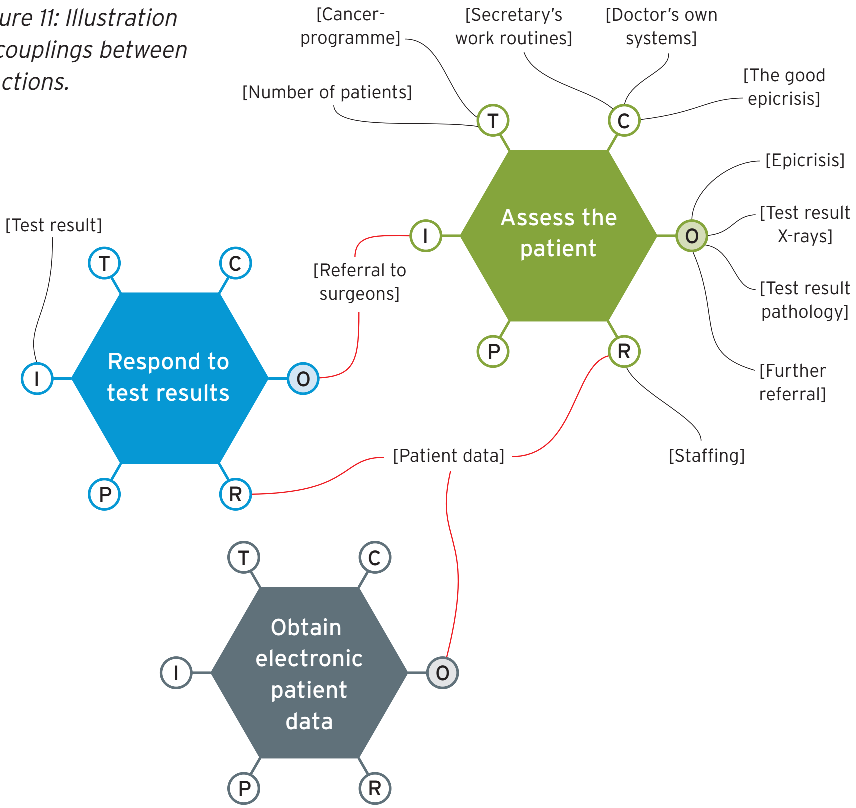

An instantiation of a FRAM model represents how a subset of functions can be mutually coupled under given conditions or within a given time frame. The couplings realised for a specific instantiation do not vary but are "fixed" or "frozen" for the assumed conditions. For an event analysis the instantiation will typically cover the entire event pathway and the couplings that existed at the time.

Foreground and background functions

Functions in the FRAM can be described either as foreground functions or background functions. The terms have nothing to do with the type of functions that are involved, but with the role of the function in a particular model – and of course also for an instantiation of the model. A function is considered a foreground function if it is part of the study focus, which in practice means if the variability of the function may have consequences for the outcome of the event or process being examined. A background function is similarly a function which can be assumed not to vary, i.e. which can be assumed to be constant – during the incident, or the process under study.

The terms foreground – background function thus refer to the relative importance of a function in the model and not to function as such. If the study focus changes, a function may change from being a foreground function to becoming a background function, and vice versa.

Background functions typically represent something that is used by foreground functions, but which are assumed to be stable in the situation under consideration. It could, for example, be a Resource (the right level of staffing or the competence of the staff) or an instruction (Control). A person's competence is generally assumed to be stable (not varying) during the execution of a task, just as an instruction also must be assumed to be stable. This does not mean that competence is sufficient or that the instruction is correct, but only that they should be regarded as stable during the time it takes to perform the task.

While the execution of an instruction may vary, the instruction itself only changes in case it is corrected or modified. The instruction is therefore only variable when considered over a longer time span, which is typically many times longer than the duration of the event. In this case the focus would be on the writing and maintenance of the instructions, which means that this becomes the function. The relationship between foreground and background functions can be shown as in Figure 12.

Figure 12: Relationship foreground and background-functions

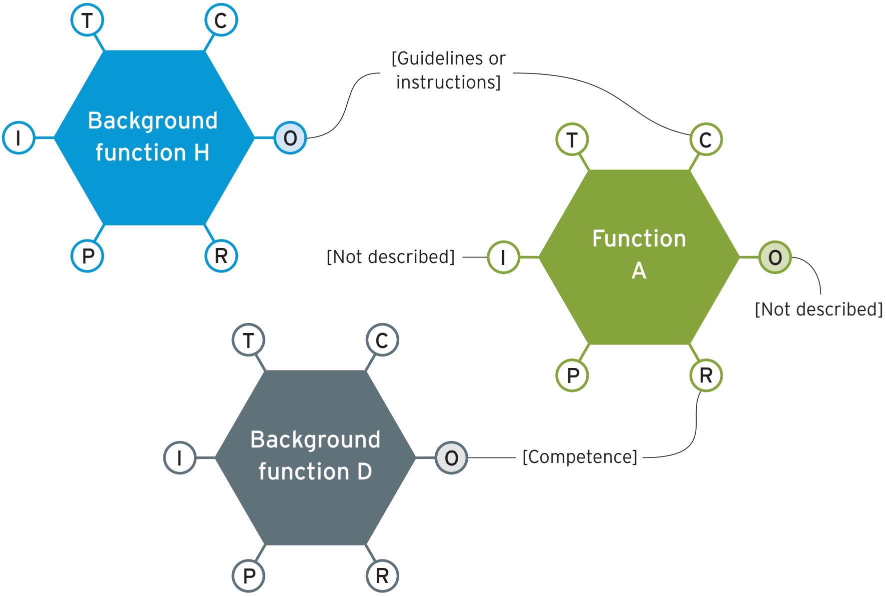

Performance shaping factors as background functions

It has been common practice in reliability analysis, especially human reliability analysis, to assume that the situation (context) or the working conditions affect work performance and how events develop. This is often described in terms of a number of performance shaping factors. But instead of proposing the existence of particular performance shaping factors, a decisive impact on the situation may just as well be seen as coming from background functions. In other words "background" functions do not vary during the time frame of the study, but they affect how an event progresses. In this way background functions represent everything that can affect the foreground functions being examined.

A foreground function should be described in as much detail as possible, i.e., so that all relevant aspects are included. The description of a background function is somewhat simpler, since usually only the Output is required. This means that the development of a model FRAM stops whenever a background function is reached (see Figure 13).

In the FRAM there is a special kind of background function that is called a drain. A drain is a function that receives an Input, and which represents the further process or the continuation of the event. A drain is a placeholder for downstream functions that are not included in the analysis; of course they can be developed in more detail if the analysis requires. For example, if a function is called **‹Discharge patient›** for which the Output is **[Patient discharged]**, then there must also be a function that "receives" the discharged patient or has this as an Input, but which is not further described. Such a function could be **‹Receive patient in nursing home›**, as shown in Figure 13.

Figure 13: Example of a drain

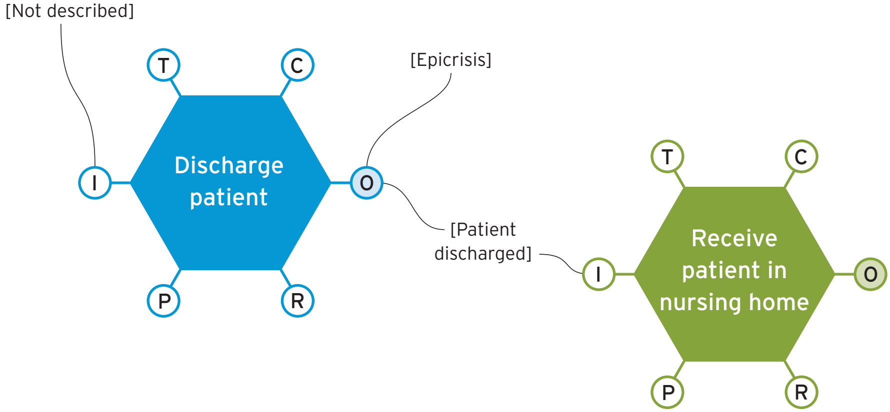

Upstream and downstream functions

While the terms foreground and background represent a function's role in a model, the terms upstream and downstream are used to describe the temporal relationship between a function that is in focus and the other functions. The analysis of the FRAM model takes place by following the potential couplings between the functions step by step. This means that there will always be one or more functions that are in focus, i.e., whose variability is being considered. The functions that have been in focus before, which means functions that already have been performed, are referred to as upstream functions. Similarly, the functions that follow the function that is in focus, are called downstream functions. During the implementation of an analysis, any function can change status from being downstream, to come into focus, and to become an upstream function.

A FRAM model describes the functions and their potential couplings for a *typical* situation, but not for a *specific* situation. It is therefore not possible to say with certainty whether a function always will be performed before or after another function. That can only be determined when the model is instantiated. By contrast, the labels foreground function and background function are valid both for the FRAM model as its instantiations. An instantiation of the model uses detailed information about a particular situation or scenario to create an instance or a specific example of the model. This corresponds to a temporal organisation of functions that reflects the order in which they will take place in the scenario, depending on how much variability there is. An upstream function is a function that for a given instantiation is performed *before* others, and therefore may affect them. Functions that – in the instantiation – take place *after* other functions and therefore can be affected by them, are called downstream functions. The terms upstream and downstream function are thus relative and not absolute.

Graphical representation of a FRAM model

As explained above, a FRAM model represents a system's functions (the union of the foreground and background functions). The model also describes the potential couplings between the functions that can be derived from the functions' aspects. A graphical representation of a FRAM model uses hexagons to represent functions. The representation is, however, usually without lines or connections between the functions (hexagons).[10] The graphical representation therefore does not define a default orientation or ordering of the hexagons (such as from left to right or from top to bottom).

---

[10] In the current version of the FRAM Model Visualiser (FMV), the connections are included by default.

An instantiation of the FRAM model shows how a subset of functions can be mutually coupled under given conditions or within a given time frame. The couplings contained in a specific instantiation are assumed to be stable during the scenario. For an event analysis the instantiation will typically correspond to the duration of the whole event and the couplings that existed at that time. A risk assessment will usually include a set of instantiations, where each instantiation represents the couplings between upstream and downstream functions at a particular time or for given conditions.

The FRAM Model Visualiser is not described in this handbook. The FMV can be found at www.functionalresonance.com from which the current version, as well as a brief set of instructions, can be downloaded.

Chapter 4: Development of a FRAM Model

In the rest of this training handbook, we will use the example of Patient with a Spinal Fracture that was described in Chapter 1. The key functions are:

- Respond to test result (X-ray, DXA and CT) (GP)
- Evaluate CT scan (Surgery department Hospital 1)
- Evaluate X-ray (Radiology, Hospital 1)
- Evaluate DXA scan (Radiology, Hospital 1)
- Triage patient at outpatient spine center for medical or surgical treatment (Hospital 2).

The five functions can be found in many places throughout the healthcare system, although not always in a predictable or related sequence. The functions occur daily, and usually several times a day.

The FRAM describes each function, as well as typical everyday variability and adjustments. This is used to understand how the variability and adjustments can affect other functions, i.e., how the functions are coupled to each other. Functions can mutually dampen each other (absorb variability), so that the situation can become stabilised. Functions can also mutually reinforce each other (amplify variability) so that the situation becomes unstable and leads to an unwanted result. The description of the potential couplings can be used to explain how functions together can succeed or fail and also how such developments can be monitored and managed.

4.1 How to prepare for data collection

Information needed

The FRAM is a tool to describe or represent how an activity usually is carried out. The selected activity is described in terms of the functions necessary for carrying it out, the potential coupling between the functions, and the typical variability. If we assume that the result of an activity is normally distributed (as in Figure 14), the purpose of the FRAM is not to describe the events represented by the left tail, 0.6% serious accidents and errors, nor the events represented by the right tail, 0.6% conspicuous successes. Instead, the FRAM aims to describe the 98.8% that falls in between – the “gray” area of everyday activity. These are the activities that function as they should – without being extraordinarily good or bad.

Figure 14: Normal distribution

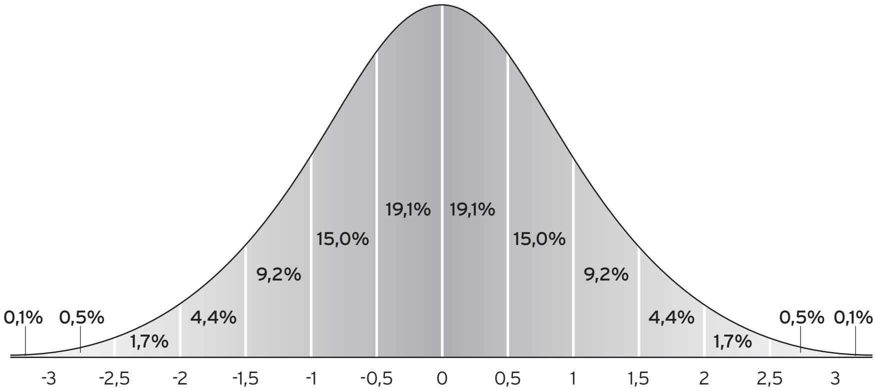

How to obtain this information

The best sources of information about activities of interest is the people who actually perform the work. They can either be the people who are in the workplace under consideration, or people who work in a similar workplace. Interviews are the primary tool of investigation, but may be supplemented by field observations and document review.

Preparing for Interviews

This training handbook will introduce the interview technique, and offer some advice and examples.

For those performing the interviews, it is important to think through the purpose of the study: how much information is needed and how will that information be helpful? It is essential to prepare as much as possible before going into the field, for instance by consulting all available sources of information such as rules and regulations, documents, protocols, job descriptions, etc. Data on turnover of personnel, equipment, procedures and organisation and major events or changes to the function can also be very valuable. This information will form the basis of the set of questions that should be asked during the interview.

The interview questions should focus on daily activities and practices, and their characteristic variabilities. So, instead of asking about successes or failures, questions should focus on the daily routines and habits – things that might be taken for granted or passed over – or even suppressed – in discussion of an adverse event.

In addition, it is important to find out as much as possible about the physical and environmental conditions of the workplace. This may require examination of the architectural drawings, photos or videos etc.

Examples of Possible Questions

- When do you start this activity? What "signals" that you can begin?
- How do you adjust the activity to different conditions? How do you determine how and when to adjust?
- How do you respond if something unexpected happens? For example, an interruption, a pause required by a more urgent task that takes priority, a missing resource, etc.
- How stable is staffing? Is staff allocation permanently assigned or adjusted daily? What happens if staffing is short?
- How stable is the environment? Supplies? Resources? Demands? Etc
- Are there often undesirable conditions that you have to tolerate or get used to?
- What preconditions are usually met?
- Are there factors that are taken for granted?
- How do you prepare for your work (documents, instructions, colleagues, etc)? What do you do if these resources are not available?
- What information do you need (equipment, services, etc)? What do you do if this is not available?
- How does time pressure affect your work?
- What skills do you need? Does everyone performing this work have these skills?
- What is the optimal way to perform this work? Is there an optimal way?
- How often do you change or adjust your work?

The Interview

If possible, the interviews should occur at the place of work, or where the event took place. A tour of the workplace is often valuable to get a feeling for the environment. The interviewer may bring a valuable set of "new" eyes to things that workers may have become "blind" to.

It is also important to prepare the interviewees for the process. They must first and foremost agree to participate in the interview and understand the purpose of the interview.

It can be useful if two interviewers conduct the interviews together: one can then concentrate on the dialogue, and the other on taking notes. One of the interviewers needs to be drawn from the work domain, but they must be aware of and control their own biases. This interviewer should not be a supervisor or manager (they only know "work as imagined" and the interviewees may not feel free to speak openly in front of their supervisor). It can be helpful to record the interview if the interviewees explicitly agree.

4.2 How to begin analysis and synthesis of the data

The first step is to type the notes from the interview and consider these along with information previously obtained during the preparation. The team needs to identify the important functions and arrange the material so that the information is sorted by functions. If possible, foreground and background functions can be identified already at this stage.

For each foreground function, one should try to identify as many of the six aspects as possible: Input; Output; Preconditions; Resources; Control; Time. Information on Input and Output represent the bare minimum required.

The Output, with its expected variability, should be described in detail with respect to *time* and *precision*. For time, one must determine whether the Output varies by coming too early, too late, on time, or not at all. For precision, one must determine if the Output is likely to be imprecise, acceptable, or precise.

How to document the Interview

For each function, you should identify the following:

Function Name: it is important to find a short and clear name that describes the activity. This should be a verb or verb phrase.

Description of the function: describe the function in as much details as necessary (free text description). Initially, most functions should be considered as background functions unless further analysis suggests otherwise. The description should include who performs the function – not the specific individual but the person in the correct organisation role. This description can be as long or as short as you wish.

The naming and description of a function is followed by a characterisation of some or all of the six aspects. Each of these is described to the extent possible with the information available and to the extent necessary to best describe the function. This helps you figure out when to stop the analysis (set the scope). A function can have more than one Input, Output, Precondition, etc.

The description of the functions and aspects can be produced by the FRAM Model Visualiser (FMV). This is a simple software tool that makes helps structure the information and also provides some useful functions to check the completeness of the model. An example of the FMV (without data) is shown in Figure 15.

Figure 15: FMV data input for functions and aspects

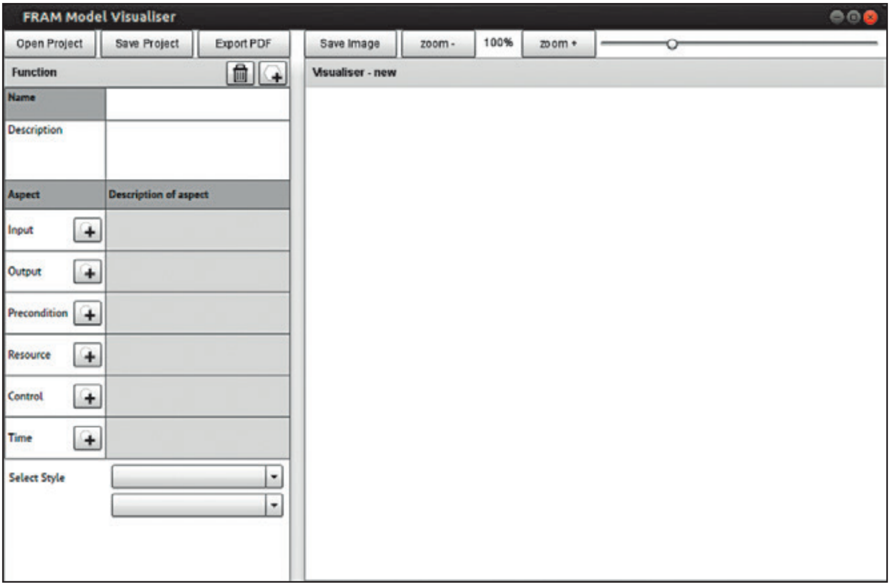

Finding the first function

So where should the model building begin? In principle, a FRAM model can start from any function. This is because the method itself ensures that the model will be complete, regardless of the starting point. Nonetheless, it may be a good idea to start with a function that is central to the activity being analysed.

In the example, Patient with a Spinal Fracture, the function **<Respond to test result>** is a central function and therefore a possible starting point of the model. The description is as follows:

The GP receives and evaluates oral or written information about the patient. The information has emerged on the basis of tests, surveys, or interviews conducted outside practice. Based on the GP's assessment of this information, the GP will decide what to do next.

Description of the six aspects and variabilities

The following table reflects the results of the interview with the GP. The details were matched to the six aspects of the function **<Response to test result>**. This work can form the basis for a more formal description of the function.

*Table 3: Worksheet for **<Response to test result>***

Input	**Test results** • Test results go directly to the Clinic's electronic email "in-box". • The email system routes the message to the GP who last saw the patient. This GP may be someone other than the patient's usual doctor, or the person who ordered the test. • The GP is busy and has a routine for reading the test result. They usually only read the top line of the test result. • If the wrong GP receives the result, they have to reroute the email to the right GP. • If a GP is on holiday, there may be a delay in rerouting.
	Requested (Urgent) Tests • When the test that is ordered is urgent, or requires action, the GP adds the test to a personal reminder list so that he can track the result. • The GP can make a note about the urgent test (a sticky note) and give it to the Clinic secretary to keep on a reminder list.
	Call from hospital • For critical results, the Hospital may call the Clinic directly.
	Follow-up visit by patient • For certain results, the hospital may direct the patient to arrange a follow-up visit with their GP in order to be informed about the test result. • The GP can arrange for the patient to have a follow-up visit in order to be informed about the test result.
Output	**Schedule patient visit** • If something serious or urgent is detected, the patient can be contacted directly phoned (or letter) by the GP. • If the result is indeterminate, contradictory or requires further investigations, the GP requests that the patient has a follow-up visit to discuss further.
	Referral • If the GP isn't sure what to do next, the GP can arrange referral for a consult.
	Results given to patient at the next visit • Results that are not serious or urgent are conveyed at the next routine visit.
	Reroute test result • If the wrong GP receives the result, they reroute the result to the right GP.

Variability of Output	This interview did not include further discussions of the variability of output. However, the interviewer can ask about variability related to time or quality. It is often helpful to ask workers in downstream functions about output variability.
Pre-condition	Electronic email system is functioning properly.
Resources	Electronic patient record system provides other information about the patient. VISINFOSYD (an open information platform for the Danish healthcare system. This website provides general information about how the system is organised). Other knowledge about the patient (the GP and the patient have a common history).
Control	Email inbox empty by Friday afternoon. Reminder from patient. Sometimes the patient is asked to book a follow-up visit. When they return, the GP will find the result of the test and share with the patient. Coordination of GP emails during absence (holidays, time-off). GPs agree to monitor each other's patients when they are away. The away GP's emails are rerouted automatically and the covering GP directs the results as appropriate. The personal reminder lists that the GP made to track urgent results. General reminders (that the GP makes and gives to the secretary to track urgent results). Systematic reading of test results. 1. The GP looks for the diagnosis and whether the result rules in or out the diagnosis. 2. The rest of the report is read or skimmed based on the findings of task 1. The results are read in priority, such that the results from the personal and secretary's reminder lists are read first. Results from other studies not requested by the GP are read last. When the GPs receive too many test results, they prioritise their review of the test results. They also rely on their own, or the secretary's, reminder lists to find the urgent and critical results first.
Time	Competing activities influence the number of test results that can be managed and the thoroughness with which the GPs read the results.

When the FRAM model of the case Patient with a Spinal Fracture is constructed, it indicates that the GP performed his job according to his normal routine.
When this information has been entered into the FMV, it looks like this. (Note that the screen dump cannot show all the aspects.)

Figure 16: FMV data input for ***<Response to test results>***

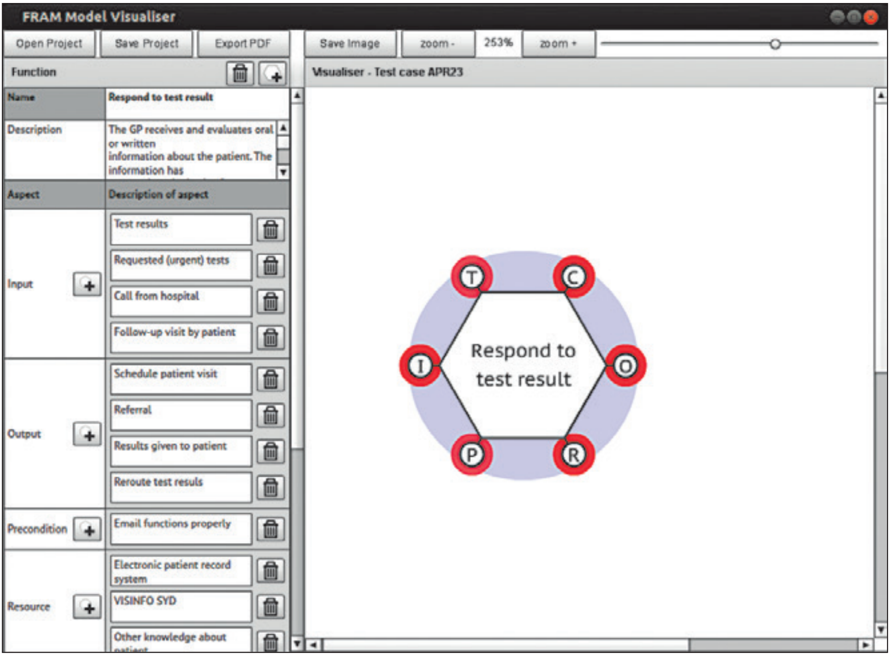

Aspect	Description of aspect
Name	Respond to test result
Description	The GP receives and evaluates oral or written information about the patient. The information has
Input	Test results Requested (urgent) tests Call from hospital Follow-up visit by patient
Output	Schedule patient visit Referral Results given to patient Reroute test resuls
Precondition	Email functions properly
Resource	Electronic patient record system VISINFO SYD Other knowledge about

The description of this function shows that a number of additional functions need to be defined in order to understand how **<Response to test result>** operates and may vary. The basic principle of the FRAM tool is that each input (Input, Precondition, Resource, Control, Time) to a function is an Output of another function.

If we first look at the Input to **<Respond to test result>,** we can infer the following:

- there must be a function whose output is **<Test result>**
- there must be a function whose output is **<Requested (urgent) Tests>**
- there must be a function whose output is **<Call from hospital>**
- there must be a function whose output is **<Hospital calls clinic>**
- there must be a function whose output is **<Follow visit by patient>**

If we examine the other aspects in the same manner, we see that all aspects of **<Respond to test result>** must be the Output of another function. This means that we may need to define a maximum of 18 other functions (including the five mentioned above).

Breadth before depth

When a FRAM model is developed it is advisable to describe all the important functions first, before looking at the functions that are "needed" by the aspects. In the example used here, this means the following four functions:

- Evaluate CT scan
- Evaluate X-ray
- Evaluate DXA scan
- Triage patient at outpatient spine center

This prevents the analysis team from becoming distracted by the details of the investigation at hand, especially if it is an event that has happened.
In the example, the next function to be described could be **<Evaluate CT scan (Surgery department Hospital 1)>.** As is clear from the description, this study is performed at Hospital 1. This function is described in the same way as the previous one. The result is shown in Figure 17.

The remaining core functions in this example are:
<Evaluate X-ray>
<Evaluate DXA scan>
<Triage patient at outpatient spine center>

Figure 17: FMV data input for ***<Evaluate CT scan>***

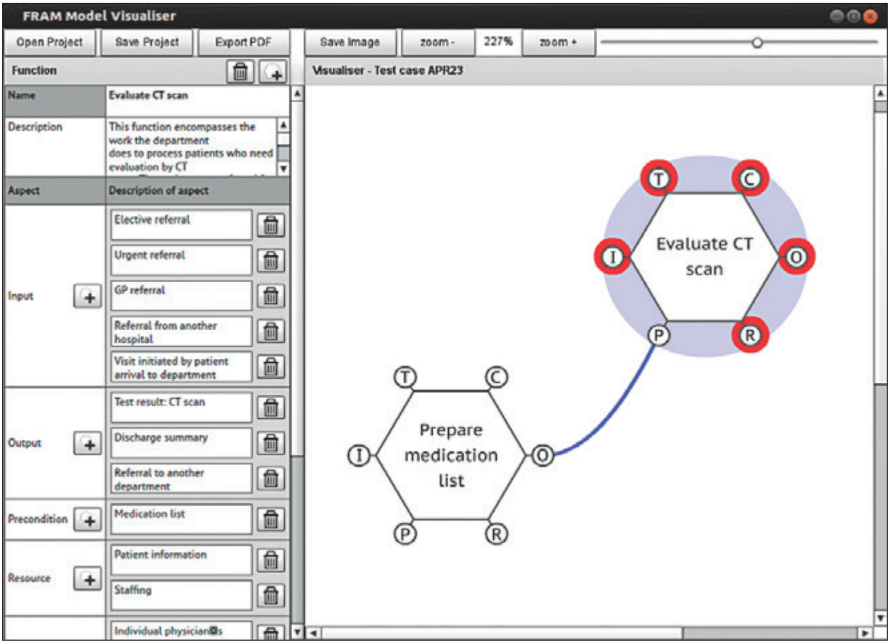

Name: Evaluate CT scan

Description: This function encompasses the work the department does to process patients who need evaluation by CT

Aspect	Description of aspect
Input	Elective referral Urgent referral GP referral Referral from another hospital Visit initiated by patient arrival to department
Output	Test result: CT scan Discharge summary Referral to another department
Precondition	Medication list
Resource	Patient information Staffing
	Individual physician's

4.3 How to proceed?

Once you have described the functions that were initially designated as the important functions, the next steps are to make sure all aspects are completely specified, using the following rules:

- Every *aspect* has to be an Output of another function.
- Every Output must be a *non-Output aspect* of another function (Input, Precondition, Resource, Control, Time).

For example, the aspect **[Medication list]** is a Precondition of the function **<Evaluate CT scan>** and must therefore also be the Output of another function. This could be a background function **<Prepare medication list>**. In this way the **[Medication list]** provides a coupling between the two functions.

It may seem daunting to describe so many functions especially because some of the new functions will have aspects that in turn require more functions. This is where the importance of distinguishing between foreground and background functions becomes critical. Since background functions only need the Output to be defined, the expansion stops there.

In practice, the new functions also often "share" aspects with functions that already have been defined. The FRAM models also generally contain more background functions than foreground functions.

The process of checking whether all aspects have been defined is built into the FMV. In the graphical rendering, aspects are marked with a red circle if they have been incompletely defined.

4.4 When is the model finished?

The FRAM model is complete when there are no "loose" aspects, as shown by the FMV. As explained above, the FRAM model can in principle be started with any function because the method itself ensures that all necessary functions are described.

In practice, the FRAM model is built through several iterations. These iterations ensure that the model is complete and consistent. When the model is used, some details may become more important and then background functions may need to become foreground functions – and vice versa. Through its development, the FRAM model will have many versions, with each version presenting an improvement over the previous.

Chapter 5: Application of a FRAM model?

The model that results from the FRAM can be the basis for an event analysis, a risk analysis, or a design evaluation. Two things are important for any type of analysis: a description of the variability and a description of how functions can depend on each other.

How to describe the variability?

The characterisation of variability in a FRAM model is the starting point for understanding how functions can be interconnected and how this can lead to unexpected results. The analysis focuses more on the variability of the Output of functions than the variability of functions as such. The reason is simply that if the performance of a function is variable without this showing up in Output, then the variability is in principle not important. However, if the Output of a function is variable, then the variability of the function is of interest because it is that which determines the characteristics hence the quality of the Output.

In principle, there are three different reasons why the Output of a function is variable.

- The variability of the Output can be a result of variability of the function itself, i.e., a consequence of the function's uniqueness or character. This is called *internal* or *endogenous variability.*
- The variability of Output may be due to variability of the work environment, i.e., the conditions under which the function is performed. This is called external or *exogenous variability.*
- The variability of Output may finally be due to variability of the Output from upstream functions, since this provides the Input, Requirement, Resource, Control, or Time for downstream functions. This type of coupling is the basis of functional resonance and is called *functional upstream-downstream coupling.*

The variability of a function can, of course, also be due to a combination of the three conditions, i.e., internal variability, external variability and upstream-downstream couplings.

The variability of the different types of functions

The FRAM assume that there are characteristic differences in the variability of technological functions, of human functions, and of organisational functions. (This three-way split corresponds to the traditional distinction between M, T, and O – huMan, Technology and Organisation.)

Technological functions are performed by different types of "machines", which in most cases include information technology. As technological functions are designed to be both predictable and reliable, the default assumption in the FRAM is that they have no significant variability. This does not exclude that technological functions can be variable; it only means that the FRAM considers them to be stable unless there is reason to believe otherwise.

Human functions are carried out either by individuals or small groups (formal or informal). A FRAM analysis assumes that human functions are variable with high frequency and large amplitude. The high frequency means that performance can change quickly or even abruptly. People react promptly to changes, especially in response to other people. The large amplitude means that differences in performance can be large, sometimes dramatically so – for better or worse. The variability depends on many different things, including the working conditions. One purpose of FRAM is to provide a clear and comprehensive description of such dependencies.

Organisational functions are performed by a group, sometimes very large groups, where activities are explicitly organised. Although organisations clearly consists of people, organisational functions differ from human functions and are usually described – and defined – on another level. They are thus functions of the system itself, rather than of the people working in the system. A FRAM analysis assumes that the frequency of organisational variability typically is low, but that the amplitude is large. The low frequency means that organisational performance changes slowly, but that the differences in the results, i.e., amplitude, can be very large.

Endogenous variability

The endogenous variability can influence human, technological and organisational functions alike.

- Technological functions can vary because the "inner workings" often are so complicated that it is not really known how the technology functions – it is intractable or underspecified. This may be true for pure "mechanical" systems, and is clearly the case for software systems. Variability can also be caused by the inevitable degradation of the mechanical components. Apart from these, there are no other significant sources of internal variability of the technological functions.
- Human functions can vary due to physiological and psychological factors. (Social factors are treated as an external source of variability in the FRAM.) Fatigue and stress (workload) are probably the most studied of the several physiological factors. Other factors are diurnal rhythm, well-being (or illness), various physiological needs, temporary disabilities, etc. There are also many different psychological factors that may affect the performance of a task, such as personality traits, cognitive style, bias in assessment and decision-making, etc.
- Organisational functions can vary for several reasons, such as how effective communication is, the authority gradient, confidence, organisational culture, organisational memory, etc.

Exogenous variability

The exogenous variability can also influence technological, human, and organisational functions alike.

- Technological functions can vary due to improper maintenance, inappropriate operating conditions – especially if they exceed the design specifications, faults in measuring instruments and sensors, overloading, misuse, etc.
- Human functions can vary due to social factors (peer pressure, unspoken norms and expectations, and so on), and because of organisational factors such as expectations, standards, requirements, commercial considerations, political considerations, etc.
- For organisations, the greatest external influence comes from the environment, the physical as well as the legal (regulations) and the commercial. The environment includes customer requirements or expectations, the availability of resources and spare parts, the regulatory environment, commercial pressures, supervisors, as well as weather and other forces of nature.

Even this brief discussion explains why functions vary when they are carried out, and that no type of function is immune to variability. For simplicity, the default assumption of the FRAM is that technological functions are relatively stable, that human functions vary with high frequency and high amplitude, and that organisational functions vary with low frequency but high amplitude. As a result of this, the variability of human and organisational functions is of most interest, whether it is the potential or actual variability.

Manifestations of variability

Once the likely possible internal and external sources of variability have been identified, the next step is to describe how the variability will appear in the function's output – how it will manifest itself, we also call this the phenomenology of variability. This step is important both because it clarifies the basis for observing or detecting variability, and because it gives an idea of how variability can affect downstream functions. The manifestations of variability can in principle be described in two different ways, one simple and one more detailed. The simple way is efficient, but not as thorough as the detailed – which in turn is more thorough, but not as efficient. In practice, it is recommended to start by the simple way, and then later go into more detail if needed.

The simple description characterises the variability of a function's Output in terms of *time* and *precision*.

In terms of time, an Output can occur too early, on time, too late, or not at all. (The last category, "not at all", can be seen as an extreme version of "too late". The consequence may be that Output either never occurred or that it occurred so late that it was useless.) An Output that is not available on time can affect the variability of downstream functions in several different ways.

Example:

<Radiology evaluation of X-ray> In case of an acute response to the GP after hours, contact the attending physician and send a message to the GP. This could mean that the GP receives the reply too late.

Example:

<Radiology evaluation of DXA scan> Results that must be discussed with other doctors are saved for the next day in order to enable a common response. This may lead to a delay in the outcome of the evaluation.

In terms of precision, an Output can be precise, acceptable, or imprecise. (More detailed characterisations can also be used.) Because the Output provides the coupling between upstream and downstream functions, the meaning of precision is relative rather than absolute.

- A precise Output meets the needs of a downstream function. A precise Output will therefore not increase the variability of downstream functions, and may potentially even reduce it.
- An *acceptable* Output may be used by a downstream function, but will require some adjustment or variability of the receiving function. An acceptable Output may therefore increase the variability of downstream functions.
- An *imprecise* Output is incomplete, incorrect, ambiguous or otherwise misleading. An imprecise Output cannot be used as it is, but requires interpretation, verification, comparison with other data or with the situation as such. These are all things that can increase the variability in the receiving function, typically by consuming resources and time that could and should have been used for other purposes.

Example:

<Triage of patient> Whether "additional findings" from the X-ray evaluation are considered depends on what the radiologists write and whether the doctor thinks it is necessary to respond. The Output may be imprecise.

Example:

<Radiology evaluation of X-ray> The Output ends with a conclusion – varying from two lines to one A4 page (8½″ × 11″). The Output may thus be imprecise and / or incomplete.

The detailed description of Output variability can be with respect to time (too early, too late) and duration (short, long), strength (weak to strong), distance (too long, too short) and direction (wrong direction), object or target (wrong item, wrong recipient), and finally with regard to sequence or order (of two or more sub-activities).

Potential and actual variability

The potential variability describes what might happen under different conditions. The actual variability describes what should realistically be expected to happen under given conditions (existing assumptions about demands, opportunities, and resources), i.e., for an instantiation of the model. For technological functions the FRAM assumes that the potential variability is not realised as long as the operating conditions corresponding to the nominal conditions. For both human and organisational functions it is assumed that the potential variability will become realised as actual variability, unless working conditions are absolutely ideal. How the actual variability will express itself depends to a very high degree on the level of detail of the instantiation, i.e., the situation for which the model is analysed. For an event analysis there will often be quite detailed information, corresponding to the facts of the event that took place. For a risk analysis, it depends on the degree of detail of the assumptions made for that scenario.

It is important to distinguish between the potential and actual variability in the description of how variability – or the consequences of variability – can propagate, specifically how variability can affect downstream functions. An analysis of coupling and resonance should only be carried out for an instantiation of a model, and thus for the actual variability – which always will be a subset of the potential variability. For this reason it is wise to begin by describing the potential variability, in order to avoid being unduly biased by having a specific scenario.

Dependence between the functions

Work involves multiple tasks and sub-tasks, here called functions, a collaboration between different people, and a coordination of their work. Each activity, i.e., the functions a person performs, must be adjusted to the conditions as described previously. But each function also forms part of the conditions of other functions.
This means that the adjustments a person makes at a certain time becomes part of the variability (of the environment) for other (downstream) functions, whether per-

formed by the same person or someone else. For a downstream function, the adjustments of the upstream functions is not known with certainty, although it rarely completely unknown or unexpected. It can usually be assumed that the upstream (previous) functions have been performed in accordance with established practice. The adjustment to upstream functions thus constitutes a variability that affects the adjustments of subsequent (downstream) functions. In a stable working environment with limited organisational variability (such as changes in demands, resources, personnel, etc.), the variability and adjustments will match each other eventually and thus provide the basis for effective everyday performance. In a working environment that is not stable, unexpected and unwanted situations may easily arise.

The previous sections have briefly described the internal (endogenous) and external (exogenous) variability, both how they occur and how this can affect downstream functions – and thus either enhance or suppress other variability. However the main reason for variability is the result of upstream-downstream couplings, particularly for the actual contexts described by an instantiation of the model.

Given that the variability is described for the Output of a function, there can in principle be five different upstream-downstream couplings: between Output and Preconditions, Output and Resources, Output and Control, Output and Time, and finally between Output and Input. Each of these is briefly described below. For each type the potential impact on the downstream function is shown in a table, using the following symbolism:

[V↑] means that the variability is most likely to increase.
[V↓] means that the variability is most likely to be reduced.
[V↔] means that the variability is most likely to remain unchanged.

Coupling between Output and Preconditions

The Preconditions describe conditions that must be present before a function can be carried out. Preconditions are provided as Output from one or more upstream functions. In cases where a function cannot be carried out unless some Preconditions are established, the first step of the function is often to verify whether the Preconditions are met. This check or verification may itself be variable or possibly skipped completely because it is assumed that the situation is the same as it used to be.

If the Output providing the Precondition is variable, this may increase the variability of downstream functions. A downstream function may possibly have to wait until the Precondition has been established (loss of time), to verify or interpret the status of the Precondition (waste of time and / or resources), or even to start before the Precondition has been established, for instance because of demands from other functions (see Table 4).

Table 4: Effects on Preconditions from upstream variability

Variability of the Output from an upstream function		Possible effects via a downstream function's Preconditions
Timing	Too early	Incorrect start because prerequisite overlooked **[V↑]**
	On time	Possible dampening of variability **[V↓]**
	Too late	Possible loss of time **[V↑]**
	Omission (not at all)	Increased improvisation, possible loss of time **[V↑]**
Precision	Imprecise	Possible loss of time (clarification); possibility of misunderstanding **[V↑]**
	Acceptable	No change **[V↔]**
	Precise	Possible dampening **[V↓]**

Coupling between Output and Resources

Resources and Execution Conditions must be present when the function is carried out. Resources are consumed during the performance, and must therefore be replenished or renewed as Output from an upstream function, which typically is described as a foreground function. With regard to Execution Conditions, it can generally be assumed that they are stable while the function is performed, i.e., that they can be described as the Output of a background function.

In some cases a lack of Resources or Execution Conditions will lead to a search for alternatives. This is an unforeseen activity that may lead to a delay in the Output of the function. If an alternative does not fit completely (e.g.. in quantity or functionality), this may also lead to increased variability of the function's Output. Conversely, fully sufficient Resources or Execution Conditions could dampen the variability of the function. These relations are described in Table 5.

Table 5: Effects on Resources or Execution Conditions from upstream variability

Variability of the Output from an upstream function		Possible effects via a downstream function's Resources or Execution Conditions
Timing	Too early	No effect **[V↔]** or possibly damping **[V↓]**
	On time	Possible dampening of variability **[V↓]**
	Too late	Possible loss of time **[V↑]**
	Omission (not at all)	Alternative solutions may be used if possible; improvisation **[V↑]**
Precision	Imprecise	Insufficient or reduced functionality **[V↑]**
	Acceptable	No effect **[V↔]**
	Precise	Possible dampening **[V↓]**

Coupling between Output and Control

Control represents that which guides or steers how a function is carried out. Variability of the Output of an upstream function that provides Control of a downstream function will obviously lead to greater variability of its Output. If Control is standardised, for example as an instruction or a procedure, it is reasonable to describe the Control as a result (Output) of a background function. However, if Control is more active and flexible, it would be more appropriate to describe it as a result (Output) of a foreground function.

The temporal aspects of Control influences whether it is available when needed. This applies, for example, for instructions that are provided verbally rather than in writing. Similarly, the contents of the Control can be imprecise or insufficient. This can lead to an increased use of resources – especially time – to confirm the contents, or get more information. These relations are described in Table 6.

Table 6: Effects on Control from upstream variability

Variability of the Output from an upstream function		Possible effects via a downstream function's Control
Timing	Too early	Instructions/guidance may be missed **[V↑]**
	On time	Possible dampening of variability **[V↓]**
	Too late	Default or *ad hoc* control is used instead **[V↑]**
	Omission (not at all)	Alternative control is sought, if possible **[V↑]**
Precision	Imprecise	Delays and trade-off with regard to accuracy and precision **[V↑]**
	Acceptable	No effect **[V↔]**
	Precise	Possible dampening **[V↓]**

Coupling between Output and Time

Time represents the various conditions that may affect how a function is carried out. It may, for example, be the time available, the time when a function can begin or be completed, requirements for synchronisation with other functions, etc. The difference between Time and Control is that Time relates to *when* a function is performed while Control relates to *how* the function is carried out.

Time is also important relative to how functions can be coupled and how they may depend on each other. (The terms upstream and downstream, for example, are defined in terms of time.) Time can affect how a function is carried out, and thus the variability of its Output. Too little time because a function has started too late, because it must be ended prematurely, or because it takes longer than usual – will typically lead to compromise in the performance of the function (see Table 7). This can be described by means of the ETTO principle (Hollnagel, 2009).

Table 7: Effects on Time from upstream variability

Variability of the Output from an upstream function		Possible effects via a downstream function's Time.
Timing	Too early	Too early start of function, wrong timing **[V↑]**
	On time	Possible dampening of variability **[V↓]**
	Too late	Delays; problems with coordination; use of shortcuts; lack of synchronisation [**[V↑]**
	Omission (not at all)	Mistimed or incorrect start or end of the function **[V↑]**
Precision	Imprecise	Increased variability **[V↑]**
	Acceptable	No effect **[V↔]**
	Precise	Possible dampening **[V↓]**

Coupling between Output and Input

The Input represents what starts a function and what is used or changed by the function. In the former case, variability of the Output from upstream may mean that the function starts too early or too late, which in turn may lead to the synchronisation and coordination problems. For example, the performance of a function may be "trimmed" to save a little time. Such "trimming" may absorb a delay, but may also increase the variability of the Output, for example, in terms of precision.

If the Input is used or processed by function, then variability in precision is more important than variability in timing. If an Input is imprecise, it may both cause delays and lead to incorrect results.

Table 8: Effects on input from upstream variability

Variability of the Output from an upstream function		Possible effects via a downstream function's Input
Timing	Too early	Premature start; Input may be overlooked **[V↑]** No effect or possible dampening **[V↓]**
	On time	No effect or possible dampening of variability **[V↓]**
	Too late	Delays which can lead to increasing use of shortcuts **[V↑]**
	Omission (not at all)	The function is not performed or is significantly delayed **[V↑]**
Precision	Imprecise	The loss of time, loss of accuracy, misunderstandings **[V↑]**
	Acceptable	No effect **[V↔]**
	Precise	Possible dampening **[V↓]**

Couplings between upstream and downstream functions help identify the variability of the Output of a function can affect other functions without reliance on linear cause-effect relationships. This reflects the concept that the way in which an event takes place depends on how the situation develops. A FRAM model describes the potential couplings and the instantiations of the model describe the actual couplings. Thus, the FRAM explains how everyday approximate adjustments can lead to unexpected results, and also how the non-linear outcomes can occur. The method is more qualitative than quantitative, because there are no generally accepted methods of expressing variability numerically.

Chapter 6: Application of the FRAM method

The purpose of a FRAM analysis is to describe how a system should function to achieve its goals (i.e., everyday work) and to understand how the possible variability of functions, alone or in combination, may affect how this happens – either by preventing it from happening or by enhancing the functionality. To do this it is necessary first to build a FRAM model of the system and then to analyse a number of scenarios or instantiations of this model.

There are many different methods for event analysis and risk assessment. Well-known examples of incident analysis methods are root cause analysis (RCA), MTO analysis, and the London Protocol. Similar examples for risk assessment are Failure Mode and Effects Analysis (FMEA), Fault Tree Analysis, and Hazard and operability Studies (HAZOP).

The FRAM differs from these methods because its purpose is to produce a description of the way in which a system functions; this description is called a FRAM model. The analysis of a past or future event uses a FRAM model to understand how something happened (an event), to assess how something may happen (a risk analysis), or to assess the impact of changes / improvements (design). The FRAM can therefore be used as part of event analysis, as part of risk assessment, or as part of the design process, but is strictly speaking neither an accident analysis method, a risk assessment method, or a design method. The three potential applications of the FRAM are illustrated in Figure 18.

Figure 18. Three applications of the FRAM.

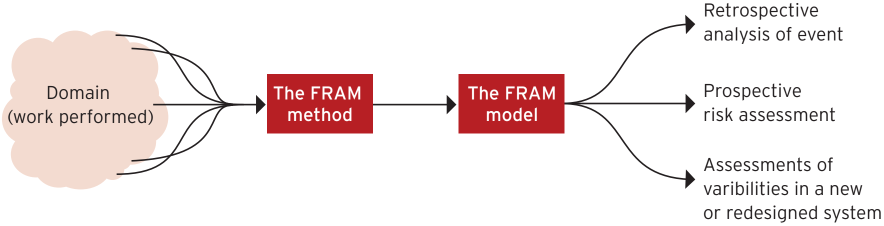

The FRAM as a basis for event analysis (retrospective analysis)

The purpose of a traditional event analysis is to propose a set of acceptable observed outcomes or consequences. These can either be the immediate causes, as in a root cause analysis, or the underlying causes, as in an MTO (huMan, Technology, Organisation) analysis. An accident analysis typically starts with the observed (adverse) outcome, and goes backwards step by step to find something that did not work, failed or malfunctioned. The purpose of the analysis is to identify the conditions and prior events that were necessary and sufficient for the actual outcome to occur. Often the reason is expressed that "If X had not happened, then this (accident) would not have happened. Therefore X is the cause of the accident".

The purpose of a FRAM analysis is not to find a cause, but to describe what should have happened for the work to succeed, i.e., what happens in everyday conditions when nothing goes wrong. This description is the starting point for finding how the combined variability of the functions can explain the actual (adverse) outcome. A FRAM model thus describes a system's features and the potential couplings or dependencies between functions, but is not a description of a sequence of individual steps of an event, such as in an accident scenario.

An accident scenario represents an instantiation of the model. A FRAM model describes the potential or possible couplings without referring to any particular situation. An instantiation of a FRAM model describes the specific couplings in a particular situation or a particular scenario. For an incident analysis, the specific context will be the actual context, to the extent this is known. For a risk assessment, the specific contexts will be the realistic expectations of what may happen.

The analysis of FRAM model instantiation includes the following steps:

- Describe the variability in the actual event of both the foreground functions and the background functions, where the latter account for the context. Consider whether the actual variability corresponded to what one might expect (i.e., whether it was typical), or whether it was unusually large (atypical). The variability is described by how the output of each function can vary, using either the simple or detailed description.
- Identify the dynamic dependencies or couplings (functional resonance) that played a role in the described event. This represents the instantiation of the model, which can be used to understand why the event evolved as it did. Compared to a traditional accident analysis, this instantiation explains what happened without the need to identify specific causes or postulate errors or malfunctions. The explanation may, for example, be a description of the typical dependencies between the variability of the functions.
- Suggest ways to monitor and mitigate or dampen performance variability (through indicators, barriers, design / modification, etc.). For unexpected positive results one should, of course, look for (controlled) ways to amplify the variability rather than ways to dampen it.

The FRAM as prospective risk assessment

A risk assessment is the most common form of proactive analysis. The purpose of a risk assessment is to identify the hazards or risks that may make it difficult or impossible for the system to fulfil its purpose, and possibly further to calculate the probability that this may happen. The risk is usually defined by the probability that something will happen, and the consequences of the result.[11] This relationship between the probabilities forms the basis of the traditional risk matrix.

A risk assessment is usually based on a specific representation of how an event may develop, such as a fault tree or an event tree. A FRAM analysis differs by being based on a functional model of how the system is intended to operate rather than of how it fails. The considerations that are the basis for the choice of representative scenarios are essentially the same in both forms of analysis, and require deep domain knowledge.

A risk assessment using the FRAM comprises the following steps:

- Characterise the possible or expected actual variability of the chosen instantiation of the model. Consider whether the actual variability will correspond to what one would expect (i.e., whether it will be typical), or whether it will be unusually large (atypical).
- Identify the dynamic couplings (functional resonance) that are expected to play a role in an event or in a particular situation. These comprise an instantiation of the model to predict how an event will develop and how control may be lost. Compared to the traditional risk assessment, this identifies what might happen without describing unique or specific results. The explanation is based on couplings and recognition of variability in daily work, rather than the presence of failures, malfunctions and their consequences.
- Suggest ways to detect and mitigate variability (indicators, barriers, design / modification, etc.). In the event of unexpected positive results, of course, search for ways to control or amplify the variability rather than for ways to alleviate it.

The FRAM for assessment of variability in a new or redesigned system

Innovation or redesign of a system can benefit from proactive assessment of variability. The aim is to identify the conditions or factors that affect whether something new (an apparatus, a change to working procedures, an organisational change, etc.) will succeed or fail. The FRAM can be used to determine how combinations of variability in the preconditions and / or resources may affect a design, or to which extent a loss of control or time constraints may hamper the performance of a specific function. The use of the FRAM for design evaluation is not described in further detail here.

---

[11] This is normally expressed as risk = probability * consequences.

Bibliography

Hollnagel, E. (2004) *Barriers and Accident Prevention*. Farnham, UK: Ashgate.

Hollnagel, E. (2009). The ETTO principle: *Efficiency-thoroughness trade-off. Why things that go right sometimes go wrong.* Farnham, UK: Ashgate.

Hollnagel, E. (2012). *FRAM: the Functional Resonance Analysis Method: Modelling Complex Socio-technical Systems.* Farnham, UK: Ashgate.

Hollnagel, E., Woods, D. D. & Leveson, N. G. (2006). *Resilience engineering: Concepts and precepts.* Aldershot, UK: Ashgate.

Glossary

Aspect (of a FRAM function): Each FRAM function is described by six aspects, namely: Input, Output, Precondition, Resource, Control, and Time. An aspect should be described if it deemed reasonable by the analysis team, and if there is sufficient information or experience to do it. But it is not necessary to describe all six aspects for all functions; indeed, sometimes it may be either impossible or unreasonable to do so. For foreground functions, it is necessary to describe at least Input and Output. For background functions that represent a source of something, it may be sufficient to describe the Output. Similarly, for background functions that represent a sink, it may be sufficient to describe the Input. Note, however, that if only Inputs and Outputs are defined, the FRAM model regresses to a simple flowchart or network.

Approximate adjustments: When working conditions are underspecified or when time or resources are limited, workers adjust performance to match the conditions. This is a main reason for performance variability. But the very conditions that make performance adjustments necessary also result in the adjustments bring approximate rather than perfect. The approximations are under most conditions good enough to ensure successful performance.

Background function: Foreground and background refer to the relative importance of a function in the model. It has been common practice in human reliability assessment to invoke the presence of so-called Performance Shaping Factors (PSFs) to describe the conditions that influence the events being studied. Instead of doing that, the FRAM represents the PSFs as background functions. The background functions denote things that affects the foreground functions being studied. In other words, the background functions constitute the context or working environment. These background functions should be described in exactly the same way as other functions, although possibly with less detail. A designated background function may become a foreground function if the focus of the analysis changes.

Bimodal principle: Technological components and systems function in a bimodal manner. This means that for every element of a system, whether a component or the system itself, the element will either function or it will not. In the latter case the element is said to have failed. The bimodal principle does, however, not apply to humans and organizations. Humans and organizations are instead multi-modal, in the sense that their performance is variable – sometimes better and sometimes worse but never failing completely. A human "component" does not stop functioning and cannot be replaced in the same way a technological component can.
The principle of bimodal functioning may become blurred in the case of systems with a large number of components, and/or systems that depend on software. In such sys-

tems there may be intermittent functions, sudden freezes of performance, and/or slow drift, e.g., in sensor measurements. However, the principle of bimodal functioning is true even in these cases, since the components are bimodal. It is just that the systems are intractable and that the ability adequately to describe and understand what is going on therefore is limited.

Boundary: An important problem in any analysis is setting the scope; that is how far should the analysis continue or when should it stop (the stopping rule). An analysis may continue beyond the boundaries of the system as initially defined. This may be because the description of the aspects of a function makes it necessary to include additional functions in the model. It may also be because some background functions vary and thereby affect designated foreground functions, in this case they should be treated as foreground functions and the boundary scope extended correspondingly. The semi-explicit stopping rule of the FRAM is that the analysis should continue until there is no unexplained (or unexplainable) variability of functions – which is the same as saying that the analysis has reached a set of functions which can be assumed to be stable rather than variable. But the boundary is relative rather than absolute, and refers to functional characteristics rather than physical characteristics.

Control (as aspect): Control, or control input, supervises or regulates a function so that it results in the desired Output. Control can be a plan, a schedule, a procedure, a set of guidelines or instructions, an algorithm, a "measure and correct" functionality, etc. A different type of Control is social control and/or expectations. Social control can be external, e.g., the expectations of others, for instance the company or management, or internal, such as when we plan do to some work and make clear to ourselves when and how to do it. Social control can also be internal, as a kind of self regulation. External social control will typically be assigned to a *background function. The description of a Control should be a noun, if it is a single word, or begin with a noun if it is a short sentence.*

Coupling: In the book "Normal accidents," Perrow (1984) proposed that systems in general could be described by two dimensions called *coupling* and *interactiveness.* Coupling describes the degree to which sub-systems, functions, and components are connected or depend upon each other; the degree of coupling can range from loose to tight. The FRAM makes a distinction between the potential couplings that are defined by a FRAM model, and the actual couplings that can realistically be assumed to exist for a given set of conditions (an instantiation). While the actual couplings always will be a subset of the potential couplings, they may differ from the couplings that were intended by the system design.

Drain: A drain is a function that receives an Input, and it represents a further process or continuation of the event. A drain is a placeholder for downstream functions that are not included in the analysis; these can of course, be developed in more detail if the analysis requires.

Downstream functions: A FRAM model describes the functions and their potential couplings, rather than the organisation of the functions for specific conditions. Therefore it is not possible to be certain that one function will always be performed out prior to or after another function. In an instantiation of the model, detailed information about a specific situation or scenario is used to create an instance or a specific example of the model. This establishes a temporal organisation of the functions as they are likely to unfold (or become activated) in the scenario, and it is only possible to consider functions in their temporal and causal relations when an instantiation has been produced. Functions that – in the instantiation – occur after other functions, and which may be affected by them, are called downstream functions. The notion of a downstream function is relative rather than absolute.

Equivalence of successes and failures: Whenever something is done, the intention is always to do it right and never to do it wrong. For each action, the choice of what to do is determined by many different things, including competence, understanding of the situation, experience, habit, demands, available resources, and expectations about how the situation may develop – not least about what others may do. If the expected outcome is obtained, the next action is taken, and so on. But if the outcome is unexpected, the preceding action is re-evaluated and classified as wrong rather than right. Thus, an error is seen as a mistake, by the common but fallacious *post hoc ergo propter hoc* argument. With hindsight, we see what should have been done, if only people had made the necessary effort at the time. The entire argument is, however, unreasonable because the action was chosen based on the expected rather than the actual outcome. Failures and successes are equivalent in the sense that we can only say whether the preceding action was right or wrong after the outcome is known. That changes the judgement of the action, but not the action itself.

Execution Condition (as aspect). An Execution Condition needs to be available or to exist when a function is active. However, it is not consumed during a function like a (proper) Resource; that is, it does not diminish over the performance of the function. The difference between a Pre-condition and an Execution Condition is that the former is only required before the function starts, but not while it is carried out. *The description of a Resource (or an Execution Condition) should be a noun, if it is a single word, or begin with a noun if it is a short sentence.*

Foreground functions: Foreground and background refer to the relative importance of a function in the model. The foreground functions denote that which is being analysed or assessed, i.e., the focus of the investigation. A designated foreground function may become a background function if the focus of the analysis changes.

Function: In the FRAM, a function represents the means that are necessary to achieve a goal. More generally, a function refers to the activities – or set of activities – that are required to produce a certain outcome. A function describes what people – individually or collectively – have to do in order to achieve a specific aim. A function can also refer to what an organisation does: for example, the function of an emergency room is to treat incoming patients. A function can finally refer to what a technological system does either by itself (an automated function) or in collaboration with one or more humans (an interactive function or co-agency). *The description of a Function should be a verb, if it is a single word, or begin with a verb if it is a short sentence.*

Functional resonance: Functional resonance is defined as the detectable signal that emerges from the unintended interaction of the everyday variability of multiple signals. The signals are usually subliminal, comprised of both the "target" signal and the remaining signals that constitute the noise. But the variability of the signals is subject to certain regularities that characterize different types of functions; hence these variabilities are not random or stochastic. Since the resonance effects result from the ways in which the system functions, the phenomenon is called functional resonance rather than stochastic resonance.

Hexagons (the FRAM): In the usual graphic rendering of a FRAM model, or of an instantiation, a function is represented by a hexagon. The reason is obviously that a hexagon has six corners corresponding to the six aspects.

Input (as aspect). The Input to a function is traditionally defined as that which is used or transformed by the function to produce the Output. The Input can be matter, energy, or information. This definition corresponds to use of the term in flowcharts, Process- and Instrumentation Diagrams, process maps, logical circuits, etc. However, there is another sense of the term Input that is just as important for the FRAM; namely that the Input activates or starts a function. The Input in this sense may be a clearance or an instruction to begin doing something, which in turn requires that the Input is detected and recognised by the function. While this nominally can be seen as being data, it is more important that the Input serves as a signal that a function can begin. Technically speaking, the Input represents a change in the state of the environment, just as if the Input was matter or energy. *The description of an Input should be a noun, if it is a single word, or begin with a noun if it is a short sentence.*

Instantiation: An instantiation of a FRAM model is a "map" of how a subset of functions are mutually coupled under given conditions (favourable or unfavourable) or for a given time-frame. The couplings that are represented by a specific instantiation do not vary but are "fixed" or "frozen" for the assumed conditions. To represent the result of performance variability of a function – or of several functions – therefore requires an instantiation. For an event (accident) investigation, the instantiation will typically cover the time-frame of the accident and represent the couplings that existed at the time. For a risk assessment, it is appropriate to work with a set of instantiations; each instantiation represents the couplings between upstream and downstream functions at a given time or for given conditions. The instantiation may be represented graphically, although this is not necessary. (See *Performance Variability*.)

Intractable systems: Systems are intractable if it is difficult or impossible to understand how they function. This typically means that the performance is irregular, and that descriptions are complicated in terms of parts and relations. Intractable systems are also underspecified, meaning that it is impossible to provide a complete specification of how work should be carried out for a sufficiently large set of situations.

Model: A FRAM model describes a system's functions (the union of the sets of foreground functions and background functions). The potential couplings among functions are defined by how the aspects of the functions are described. The FRAM model, however, does not describe the actual couplings that may exist under given conditions (see: *Instantiation*). A graphic representation of a FRAM model will be a set of hexagons, where each hexagon stands for a function, but without any lines or connections among functions.

Output (as aspect): The Output from a function is the result of what the function does, for instance by processing the Input. The Output can represent matter, energy, or information – the latter as a command issued or the outcome of a decision. The Output can represent a change of state – of the system or of one or more output parameters. The Output can also represent the signal that starts a downstream function. *The description of an Output should be a noun, if it is a single word, or begin with a noun if it is a short sentence.*

Performance variability: Traditionally, the study of risk and accidents has focused on how failures or malfunctions of components or elements (technological, human, organisational) occured and how the effects propagated through the system. This reflects a *bimodal view* of functions and performance. The FRAM is based on the principle of equivalence of successes and failures and the principle of approximate adjustments. Consequently, in practise performance is always variable. The performance variability of upstream functions may affect the performance variability of downstream functions, and thereby lead to non-linear effects (functional resonance).

Resonance: In physical systems, classical (or mechanical) resonance refers to the phenomenon that a system can oscillate with larger amplitude at some frequencies than at others. These are the system's resonant (or resonance) frequencies. At these frequencies even small external forces applied repeatedly can produce large amplitude oscillations, which may seriously damage or even destroy the system.

Pre-condition (as aspect): In many cases, a function cannot begin before one or more Pre-conditions have been established. These Pre-conditions are system states that must be true, or conditions that must be verified before a function is performed. Although a Pre-condition is a state that must be true before a function is carried out, it does not itself constitute a signal that starts the function. An Input, on the other hand, can start a function. This simple rule can be used to determine whether something should be described as an Input or as a Pre-condition. *The description of a Pre-condition should be a noun, if it is a single word, or begin with a noun if it is a short sentence.*

Resilience: A system is said to be resilient if it can adjust its functioning prior to, during, or following changes and disturbances, and thereby sustain required operations under both expected and unexpected conditions.

Resilience engineering: The scientific discipline that focuses on developing the principles and practices that are necessary to enable resilience in systems.

Resource (as aspect): A Resource is something that is needed or consumed while a function is performed. A Resource can represent matter, energy, information, competence, software, tools, manpower, etc. In principle, time can also be considered a Resource, but Time has a special status so it is treated as a separate aspect. Since some Resources are consumed while the function is carried out and others are not, it is useful to distinguish between *Resources* and *Execution Conditions.* A Resource is consumed by a function, so that there will be less of it as time goes by, but an Execution Condition needs to be available or exist while a function is active and is not consumed. The difference between a Pre-condition and an Execution Condition is that the former is only required before the function starts, but not while it is carried out. *The description of a Resource (or an Execution Condition) should be a noun, if it is a single word, or begin with a noun if it is a short sentence.*

Safety-I: The goal of Safety-I is that the number of adverse outcomes (accidents / incidents / near misses) are as low as possible. Safety-I is achieved by trying to ensure that things do not go wrong, either by eliminating the causes of malfunctions and hazards, or by containing their effects.

Safety-II: The goal of Safety-II is that the number of successful outcomes is as high as possible. Safety-II seeks system strength through the the ability to succeed under varying conditions. Safety-II is achieved by trying to make sure that things go right, rather than by preventing them from going wrong (see also *Work-as-Done*).

Stochastic resonance: Stochastic resonance is the enhanced sensitivity of a device to a weak signal that occurs when random noise is added. The outcome of stochastic resonance is non-linear, which means that the output is not directly proportional to the input. The outcome can also occur – or emerge – instantaneously; this differs from classical resonance which must be built-up over time.

System: Systems are usually defined with reference to their structure, i.e., in terms of their parts and how they are connected or put together. Common definitions emphasise that the system is a whole, and that it is composed of independent parts or objects that are interrelated. Definitions of this type rely on the principle of decomposition, and explain the overall functioning of the system in terms of the functioning of the components or parts – keeping in mind that the whole is larger than the sum of the parts. However, it is entirely possible to define a system in terms of how it functions rather than in terms of what the components are and how they are combined. From this perspective, *a system is a set of coupled or mutually dependent functions.* This means that the characteristic performance of the system – of the set of functions – cannot be understood unless all the functions i.e., the set as a whole are described. Thus, the system is not delimited by its structure or relations among components (the system architecture). An organisation, for instance, should not be characterised by what it is but by what it does. Neither should it be characterised by the people who are in a given place (on the organisation chart or in reality) but by the functions they performs. One consequence of a functional perspective is that the distinction between a system and its environment, and thereby also the system boundary, becomes less important, cf., the distinction between foreground and background functions.

Time (as aspect): The Time aspect of a function represents the various ways in which Time can affect how a function is carried out. Time, or rather temporal relations, can be seen as a form of Control. One example is when Time represents a sequencing condition. For instance a function may have to be performed (or completed) before another function, after another function, or overlapping with – parallel to – another function. Time may also relate to a function alone, seen in relation to either clock time or elapsed time. Time can also represent a Resource, such as when something must be completed before a certain point in time, or within a certain duration (as when a report must be produced in less than a week). Time can also be seen as a Pre-condition, e.g., that a function must not begin before midnight, or that it must not begin before another functions has been completed. Yet rather than including Time as a part of other of the three aspects of a function – or even four since it could also be considered as an Input – the FRAM recognises its special status by having it as an aspect in its own right. *Time as an aspect should be a noun*, if it is a single word, or begin with a noun if it is a short sentence. *The description of a Time aspect could also be one of the commonly used descriptors of a temporal relationship [Allen, J. F. (1983). Maintaining knowledge about temporal intervals, Communications of the ACM, 26(11), 832-843].*

Tractable systems: Systems are called tractable if it is possible to understand how they function. This means that the performance is highly regular, and that descriptions are relatively simple in terms of parts and relations.

Upstream functions: An upstream function is a similar way to a *downstream function*. It is a function that – in a given instantiation – happens before other functions, and which may affect them. The notion of an upstream (or downstream) function is relative rather than absolute.

Variability: See *Performance Variability.*

Work-as-Imagined/ Work-as-Done: Because performance adjustments are always necessary, Work-as-Done (WAD) is always different than Work-as-Imagined (WAI). These two terms explain how proximal, or sharp-end factors, in combination with distal or blunt-end factors can lead to accidents. Workers at the sharp end accept that WAD is, and must be, different from WAI. For them it is no surprise that descriptions based on WAI cannot be used in practice and that actual work is different from prescribed work. Unlike Safety-I, Safety-II considers WAD at the sharp end and the blunt end on equal terms, and recognizes that performance adjustments are required in both cases. Safety-II acknowledges that WAI and WAD are different, and focuses analysis on WAD.

Notes

Notes

Notes

Notes

Centre for Quality
P. V. Tuxens Vej 5, 1
5500 Middelfart
Denmark

centerforkvalitet.dk

Region of
Southern Denmark