A FRAM Primer

This Appendix provides a quick introduction, a primer, to the FRAM. More extended descriptions can be found on www.functionalresonance.com and of course in Hollnagel (2012).

The purpose of the Functional Resonance Analysis Method is to analyse how something is done or could be done in order to produce a representation of it. This representation is effectively a model of the activity in the sense of being a compact description that captures the essential features of how something is done. In the case of the FRAM, the essential features are the functions that necessary and sufficient to describe the activity and the way in which they are coupled or mutually dependent.

Unlike all other safety analysis methods the FRAM does not include or refer to an existing model. Consider, for instance these examples.

- Root Cause Analysis (RCA) represents adverse outcomes and the events that lead to them by single or multiple cause-effect chains starting with the (root) cause and ending with the observed outcomes. The underlying model of accident causation is the Domino Model that was introduced by Heinrich (1931).
- The Swiss cheese model (Reason, reference) represents adverse outcomes, and therefore also adverse events, by a combination of latent conditions and active failures. The former are the famous holes in the 'planes' or slices of cheese, while the latter are the dangers or hazards that can penetrate the organisation. The purpose of the associated method, for instance TRIPOD, is to identify the underlying factors of accidents, incidents, and near misses.
- The AcciMap approach (Rasmussen & Svedung, reference) represents accidents in complex socio-technical system by mapping the possible causes onto six system levels. The map is a network of connections from the physical sequence of events and activities right up to the causes at the governmental, regulatory and societal levels. The same principles is used by STAMP (Leveson, 200x), which explains accidents as a result of inadequate control or inadequate enforcement of safety-related constraints. Both are linear causal analysis methods based on assumptions about how systems generally are structured.

Since each of these methods refer to a pre-defined model or description about how the world is structured, the resulting descriptions essentially map the events onto the model. The FRAM, however, is not a method to analyse an event or activity in terms of a model, but a method to analyse an activity in order to produce a model (description). Instead of having a pre-defined model, structure, or organisation of the world, the FRAM is based on some assumptions about how things happen.

The basic principles of the FRAM

The assumptions on which the FRAM is based can be expressed by four basic principles.

The principle of equivalence

Explanations of how something happens typically rely on decomposing a system or an event into parts, either physical parts such as people and machines, or the segments of an activity such as individual actions or steps in a process. Outcomes are explained by linear cause-effect relations among the parts and adverse outcomes are attributed to malfunctions or failures of parts. This implies that things that go right and things that go wrong have completely different causes. The FRAM – and Resilience Engineering – takes different approach, namely that things that go right and and things that go wrong happen in much the same way. The fact that the results are different does not mean that the underlying causes also must be different. The principle of approximate adjustments explains why this is so.

The principle of approximate adjustments

Most socio-technical systems cannot be specified in every detail. Effective work therefore requires that performance continuously is adjusted to the existing conditions (resources, time, tools, information, requirements, opportunities, conflicts, interruptions). 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. Since resources (time, materials, information, etc.) almost always are limited, the adjustments will typically be approximate rather than precise. In most cases this does not matter since people will know what to expect and be able to compensate for that. The approximate adjustment are the reason why things mostly go right, but also the reason why they occasionally go wrong.

The principle of emergent outcomes

The variability of individual functions is rarely large enough to serve as the cause of something going wrong or to be described as a failure. The variability of multiple functions may on the other hand combine in unpredictable (non-linear) ways that can lead to unexpected and disproportionate outcomes – negative as well as positive. Acceptable and unacceptable outcomes can both be explained as *emerging* from variability due to the everyday adjustments rather than as a *result* of single or multiple cause-effect chains starting from malfunctions or failures of specific components or parts.

The principle of functional resonance

As an alternative to linear causality, the FRAM proposes that the variability of two or more functions can coincide and either dampen each other or amplify each other to produce an outcome, or output variability, that is disproportionally large. In the latter case the consequences may spread to affect other functions in analogy with the phenomenon of resonance.

There are three types of resonance. In classical resonance, a regular external force can increase the amplitude of an oscillating system at a specific preferential frequency. Outcomes build-up gradually and are proportional to inputs. In stochastic resonance the external force is replaced by random noise, which every now and then can push a subliminal signal over the detection threshold. Outcomes can occur instantaneously and are non-linear, i.e., not directly proportional to the input. Functional resonance, finally, describes the noticeable performance variability in a socio-technical system that can happen when multiple approximate adjustments coincide. Performance variability is not random because the approximate adjustments comprise a small set of recognisable short-cuts or heuristics. 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. Functional resonance offers a systematic way to understand outcomes that are both non-causal (emergent) and non-linear (disproportionate).

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 work as it typically takes place.

The Meaning of Functions in the FRAM

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 complicated – 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 guide an approaching aircraft.
- A function can also refer to something that an organisation does: for example it is the function of a railway line to transport people and goods.
- A function can finally refer to what a technical system does either by itself (an automated function, such as a robot) or together with one or more people (an interactive or socio-technical function, like a check-in kiosk in an airport).

To emphasise that functions represent activities or something that is done, it is recommended to describe FRAM functions by a verb or a verb phrase. For instance "(to) diagnose a patient" rather than "diagnosing a patient" or "(to) request information" rather than "requesting information".

The meaning of Aspects in the FRAM

In the FRAM, functions can be described in terms of six aspects: Input, Output, Requirements, Resources, Control, and Time. The general rule of the FRAM is that a function's aspects should be described the analysis team things it is appropriate, 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. But as a minimum, at least one Input and one Output must be described for all foreground functions, cf., below. Note, however, that a FRAM model is reduced to an ordinary flow chart or network diagram if only the Input and Output aspects are described. The FRAM recommends that an aspect is described with a 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.

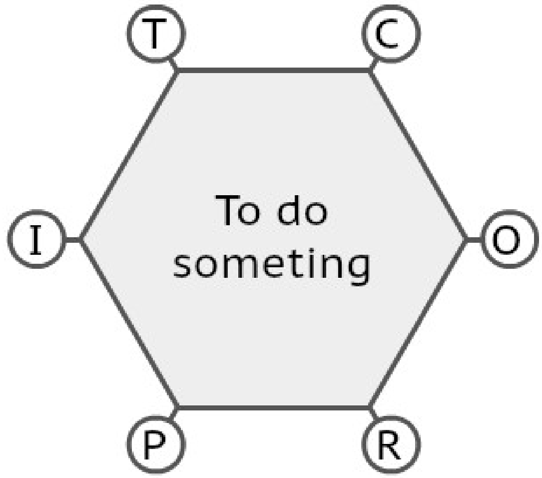

Figure 1: A FRAM function

- The **Input** (**I**) to a function is traditionally that which is used or transformed by the function to produce the Output. The Input can represent matter, energy, or information. But an Input can also be that which activates or starts a function, such as a clearance or an instruction to begin doing something. Input can be seen as a form of data or information, or more generally as a state change a function interprets 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. When a Input aspect is described for one function, it must also be described as an Output from another function.
- The **Output** (**O**) of a function describes the result of what the function does, for example, the result of processing the Input. The Output may 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. When an Output aspect is described for one function, it must also be defined as an Input, Precondition, Resource, Control, or Time aspect for another function.
- A function can sometimes not begin before one or more **Preconditions** (**P**) have been established. These Preconditions can be understood as system states that must exist, or as conditions that ought to be verified before a function is carried out. A Precondition

does, however, not itself constitute a signal that can activate a function. This simple rule can be used to determine whether something should be described as an Input or as a Precondition. The description of a Precondition should be a noun or a noun phrase. When a Precondition aspect is defined for one function, it must also be defined as an Output from another function.

- A **Resource** (**R**) 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. 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. When a Resource aspect is defined for one function, it must also be defined as an Output from another function.
- **Control** (**C**) 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. When a Control aspect is defined for one function, it must also be defined as an Output from another function.
- **Time** (**T**) represents the various ways in which time can affect the performance of a function. Time, or rather temporal relations, can be seen as a form of Control. 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 can also be seen as representing 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, of course, also be seen as a Precondition, 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 having Time as a part of either of the three aspects of a function it is reasonable to recognise its special status by having it as an aspect in its own right. The description of a Time aspect should be a noun, if it is a single word, or begin with a noun if it is a short sentence. When a Time aspect is defined for one function, it must also be defined as an Output from another function.

Couplings

Couplings describe how functions are connected or depend upon each other. Formally, two functions are said to be coupled if that have an aspect in common. 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).

The couplings that are described in a FRAM model, i.e., the dependencies that are the due to common aspects, are called **potential** couplings because a FRAM model describes the possible relationships or dependencies between functions without referring to a particular situation. An instantiation of a FRAM model represents how a subset of functions can **actually** become coupled under given conditions or within a given time frame. The subset represents the actual couplings or dependencies that have occurred or are expected to occur in a particular situation or a particular scenario. The couplings described for 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.

While the actual couplings always will be a subset of the potential couplings, they may be different from the couplings that were intended by the system design. The couplings in a FRAM model are generally n-to-n (or many-to-many) rather than 1-to-1.

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 as 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 event, 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 become a background mode, and vice versa.

Background functions typically represent something that is used by foreground functions, but which is 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 must 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 such cases the focus would be on the writing and maintenance of the instructions, which means that this becomes the function.

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 a background function that serves as placeholder for downstream functions not included in the analysis represent (a drain), it may be sufficient to describe the Input. This means that the development of a model FRAM stops whenever a background function is reached.

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.

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 and also shows the potential connections between the functions. (The FRAM Model Visualiser or FMV is not described here. The FMV can be found at www.functionalresonance.com from which the current version, as well as a brief set of instructions, always can be downloaded.) The graphical representation does not define a default orientation or ordering of the hexagons (such as from left to right or from top to bottom).

An instantiation of 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 'logic' of the analysis is as follows:

The purpose of the FRAM is to describe or represent 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 couplings due to common aspects of functions.

- Describe the event, using either an existing description (for instance an event report), a future scenario (for instance a safety case), a proposed change (design specifications, functional specifications) or any other available material.
- Try to describe what should happen (meaning Work-as-Done, not Work-as-Imagined). This will require data from the daily work place, from people who have expert knowledge about the activity, rather than (post mortem) data from an accident investigation. The description is basically the set of functions that are required for everyday performance to succeed.
- The outcome is the FRAM model. This is the basis for characterising the expected (potential) variability of the activity as carried out in the everyday work environment. The characterisation of this variability also provides the basis for describing the potential couplings.
- Use more specific information (e.g., an event or accident analysis, or a safety case scenario) to produce one or more instantiations of the model.
- Analyse these instantiations, either to find an explanation why something happened, or a plausible scenario of what may happen (as in risk assessment).