Failure: Analysis of an Engineering Concept

Luca Del Frate

Simon Stevin Series in the Philosophy of Technology

Failure

Analysis of an Engineering Concept

Failure

Analysis of an Engineering Concept

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op dinsdag 28 januari 2014 om 15.00 uur

door Luca DEL FRATE
Laurea in filosofia, Università degli Studi di Padova

geboren te Palmanova, Italië

Dit proefschrift is goedgekeurd door de promotor:
Prof. dr. ir. P.A. Kroes

Co-promotoren:
Dr. P.E. Vermaas
Dr. M.P.M. Franssen

Samenstelling promotiecommissie
Rector Magnificus, Technische Universiteit Delft, voorzitter
Prof. dr. ir. P.A. Kroes, Technische Universiteit Delft, promotor
Dr. P.E. Vermaas, Technische Universiteit Delft, copromotor
Dr. M.P.M. Franssen, Technische Universiteit Delft, copromotor
Prof. dr. ir. M. Boon, Universiteit Twente
Prof. dr. S.O. Hansson, Kungliga Tekniska Högskolan
Prof. dr. C.W. Johnson, University of Glasgow
Prof. dr. ir. P.H.A.J.M. van Gelder, Technische Universiteit Delft
Prof. dr. ir. I.R. van de Poel, Technische Universiteit Delft, reservelid

editors: Peter Kroes and Anthonie Meijers

ISBN: 978-90-386-3542-2
ISSN: 1574-941X

Contents

List of papers

Chapter 2

Del Frate, L., Franssen, M., and Vermaas, P. E. (2011) 'Towards a trans-disciplinary concept of failure for Integrated Product Development', in: *International Journal of Product Development* 14 (1-4): 72–95.

Chapter 3

Del Frate, L. (2013) 'Failure of Engineering Artifacts: A Life Cycle Approach', in: *Science and Engineering Ethics* 19 (3): 913–944.

Chapter 4

Del Frate, L. (2012) 'Preliminaries to a formal ontology of failure of engineering artifacts', in: Donnelly, M. and Guizzardi, G. (eds.), *Formal Ontology in Information Systems: Proceedings of the Seventh International Conference (FOIS 2012)*, IOS Press, Amsterdam: 117–130.

Chapter 5

Del Frate, L., Zwart, S. D., and Kroes, P. A. (2011) 'Root cause as a U-turn', in: *Engineering Failure Analysis* 18 (2): 747–758.

Chapter 6

A version of this chapter will be submitted to the journal *Technology and Culture*.

Maarten Franssen, Peter Kroes, Pieter Vermaas, and Sjoerd Zwart are acknowledged for granting permission to publish the co-authored papers in this dissertation.

Acknowledgements

Doing a PhD-project is a journey, a long, sometimes bumpy journey full of surprises, funny episodes, and interesting people. Admittedly, the PhD-journey analogy has been made so many times that it has become a cliché. Nevertheless, I think it is a very appropriate analogy, especially if you consider the amount of travel that working in academia today implies. Moreover, I personally associate some of the most vivid memories of this PhD with travelling. Definitely the most memorable was my second journey to Japan, in February 2012. I was going there for a conference together with Peter Kroes, my promotor, who had been invited as a keynote speaker. Our destination was Sendai, the capital of Miyagi Prefecture. Less than a year before, the 11th of March 2011, the area was struck by the massive Tohoku earthquake and the following tsunami. Indeed, Sendai is located about 100 Km north of the infamous Fukushima nuclear power plant and some of the damage was still visible around the conference venue itself in the form of long and wide cracks running along walls and staircases (everybody reassured us the buildings were totally safe, though). On the second day of our trip we were invited to visit the coastal areas to the south and see with our own eyes what happened there. It was a cold and rainy day and before us stood a vast area of complete destruction. In that location the gigantic wave reached as high as 17 meters. Only houses built above that line survived, everything else had been swept away. Such was the amount of debris that after one year of work even the super-efficient Japanese were still busy with the clean-up. Although we had already seen plenty of images of that kind on TV, walking through that deserted place and witnessing the admirable dignity of our Japanese hosts made a great impression on us.

But that was not the end of our journey and the following days we enjoyed the warmth of Japanese hospitality. We were shown around, visited beautiful temples, abundantly explored the local cuisine, and took a ride on the mighty Shinkansen, the bullet train. Even though it was a short trip, the combination of contrasting experiences and emotions made me realize that, besides being a talented philosopher (which I already knew him to be), Peter is also a wonderful travel companion. His enthusiasm is contagious and because of his genuine passion for learning there is never shortage of subjects for conversation. And he

has a gift for finding the right words for almost every occasion, both for the good days and for the less good ones. Thanks Peter for being such a great promotor.

Pieter Vermaas, co-promotor, has been my daily supervisor, but his contribution has been much greater than this description might suggest. By virtue of example, and by challenging my ideas with provoking questions he has been a key figure in my PhD, and I wish to thank him for all the support and understanding; not to mention the good laughs. Many thanks also to Maarten Franssen, co-promotor, for all the fascinating and wide-ranging conversations, which were always enlightened by his impressive philosophical acumen.

During my research, I took part in the EuJoint project, an international exchange project on engineering ontologies, and I visited two of the participating institutions. Thus, I would like to thank everyone at the Laboratory for Applied Ontology (Trento, Italy) which I visited in April 2011, particularly Nicola Guarino and Stefano Borgo for the kind hospitality and the stimulating comments on my work. Later that year, I spent a month at the Mizoguchi Lab in Osaka. That was my first visit to Japan, and I wish to thank all the researchers and staff working at the Mizoguchi Lab, especially Riichiro Mizoguchi and Yoshinobu Kitamura for the warm hospitality and the valuable feedback on my research.

Quite naturally, by travelling one gets to meet people, and, indeed, I have been very fortunate to have met many brilliant researchers with whom I collaborated, made plans for future collaboration, or just had interesting conversations. Hence, I would like to acknowledge: Gaetano Cascini (Polytechnic University of Milan) and Gualtiero Fantoni (University of Pisa) with whom I co-authored a paper, Claudia Eckert (Open University), Crispin Hales (Hales & Gooch Ltd.), and all the engineers I have met at the International Engineering Conference on Failure Analysis in 2010 and 2012 who showed interest in my research, especially Emiel Amsterdam (NLR), Richard Clegg (Queensland University of Technology), Fabrizio D'Errico (Polytechnic University of Milan), Colin Gagg and Peter Lewis (Open University), Tommaso Ghedini (ESA), and Stan Lynch (DSTO). I am grateful to Russell Wanhill (NLR) for the generosity shown in sharing his knowledge and for his ability in clarifying complex technical matters.

Michael van Tooren (TU Delft, Aerospace Engineering), Tetsuo Tomiyama (Cranfield University), Marco Ferraguti (University of Milan), and Cory Cooper (ISAF) are acknowledged for providing valuable support and insightful comments, particularly during the early stages of my research.

Admittedly, by working at TU Delft one does not need to travel in order to get in touch with different cultures and interesting people, they just happen to be there. I enjoyed wholeheartedly my time at the Philosophy department, with its friendly and yet productive atmosphere and the wonderful colleagues. I wish to express my gratitude for their support and friendship to Behnam Taebi and Christine van Burken, who accepted to show up in fancy ceremonial dresses as *paranymphs* at my defense, Malik Ahmed, Christian Detweiler, Adam Henschke, Bjørn Jespersen, David Koepsell, Filippo Santoni De Sio, Philip Serracino Inglott, Dingmar van Eck, and Sjoerd Zwart. A word of praise goes also to Diana Droog and Henneke Filiz-Piekhaar for their help in organizational and practical matters.

Many thanks to all members of the Coffee Breaks Discussion Group, whose regular meetings provided both much needed distraction and scores of insightful comments, and a special mention to the most senior members George Dafermos, Emiel Kerpershoek, Devender Maheshwari, Jop van den Hoogen and the honorary member Anish Patil.

I am at loss of words (and she knows it doesn't happen very often) to express my affection and gratitude to Barbara with whom I shared the highs and lows of this journey. Without her I wouldn't be able to travel this far. Together, we would like to thank our families which never missed to make their affection and support felt. Also, we would like to thank all our friends, whose company we hope to enjoy more often. I know I was a bore with all the ranting about finishing the PhD, it's over now, you can pick up my calls!

Finally, I cannot avoid mentioning Pen the Penguin, my dear skating teacher: thanks to his lessons I won the 2010 Philosophy Section Best Skater Award. That alone was worth the journey.

1 Introduction

It is fashionable for books about engineering failures to mention, often at the beginning, the Code of Hammurabi, a Babylonian law code dating back to about 2250 BC, see e.g. (Feld and Carper: 1997; Ratay: 2009; Grimvall et al.: 2010; Bazu and Bajenescu: 2011). In these books, the Code is presented as a stark reminder that engineers have been dealing with failures since they started realizing technical artifacts. In fact, a section of the Code deals explicitly with legal consequences of engineering failures and takes, for today's standards, a rather strong stance. For instance, law 229 is:

> If a builder builds a house for someone, and does not construct it properly, and the house which he built falls in and kills its owner, then that builder shall be put to death. (Adapted from King: 1915, 23)

Tongue-in-cheek, Feld and Carper (1997) note that this kind of legislation may have had a negative impact on engineering progress in Babylonian times by reducing the opportunities of learning from failures. On a more serious tone, the laws in the Code illustrate a duality in the concept of failure that has persisted until the present day, namely the duality between *material* and *functional* aspects of failure. The former aspect is exemplified by collapses as mentioned in law 229. The latter appears in law 235 where the legislator deals with a different branch of engineering, ship building, as thus:

> If a shipbuilder builds a boat for someone, and does not make it tight, if during that same year that boat is sent away and suffers injury, the shipbuilder shall take the boat apart and put it together tight at his own expense. The tight boat he shall give to the boat owner. (Adapted from King: 1915, 23)

In this case, the problem with the technical artifact does not reside in its structural integrity but relates to a lacking of adequate performance (e.g., its water tightness) that may obtain regardless of material or structural changes. Indeed, an artifact might be in pristine condition and yet unable to perform as expected.

The long lived duality between material and functional conceptualizations of failure is a prominent theme in this thesis. The fact that the duality has been there for such a long time does not mean, however, that the concept of failure

has not evolved since antiquity. In fact, another relevant theme of this thesis deals with changes in the concept of failure that have occurred, particularly in recent times, as a consequences of the expanding scope of engineering activity. Walter Vincenti believes a substantial increase in the scope of engineering activity unfolded within the span of his professional career:

> During my career as an engineer I have seen the scope of engineering problems also expand increasingly to include social and environmental matters. (Vincenti: 1990, 255)

By "engineering problems" Vincenti means the challenges that engineers are confronted with. Engineers, Vincenti is saying, traditionally used to deal with "purely technical" problems, like "to supply lift with the least possible drag in the case of an airfoil, or to hold two pieces of metal together in the case of a rivet"(255). Contemporary engineers, however, are expected to address a broader set of issues and requirements in their designs, including social and environmental aspects of their products. These *non-purely* technical or extra-technical considerations redefine the idea of what counts as a successful or unsuccessful product. To put it differently, besides the traditional goal of achieving adequate technical performance, there are additional extra-technical goals that engineering products are expected to achieve. Correspondingly, also the concept of failure has broadened to encompass social and environmental issues.

Failure and engineering are deeply intertwined: the moment a new technology is introduced, a new mode of failure appears. Think, for instance, of electronic devices which, in the words of failure analyst W.J. Plumbridge have opened up "new avenues for failure analysis" (Plumbridge: 2009): because of advanced materials, innovative manufacturing processes, and increasing miniaturization, new modes of failure emerge and sophisticated techniques are needed to investigate them. In turn, better understanding of failures allows engineers to improve their designs thereby making products less vulnerable to failure. This introduces a third theme, namely *learning from failures*.

Prevention of failures is a major concern for all engineers. So much so that, according to Henry Petroski, "Virtually every calculation that an engineer performs in the development of [a product] is a failure calculation" (Petroski: 1996, 89). Prevention of failure not only requires that engineers carefully scrutinize and double-check their designs in search for errors, flaws or unanticipated side-effects; prevention, one might say, begins even before actual

designing has started and continues after a design is finished. For, to be able to spot potential flaws, engineers need to learn about known modes of failure that have been diagnosed elsewhere, and they need to closely monitor how their products perform in service and investigate potential anomalies. Thus, besides lots of calculations, prevention of failure necessitates a great deal of information exchange between different disciplines and professional specializations (e.g., designers, failure analysts, manufacturers, maintenance specialists). Nevertheless, even a cursory survey of the engineering literature on the subject can reveal (as will be documented below) that this crucial notion admits of many different interpretations and has resulted in a large and partially disorganized failure terminology. Although such a plurality of interpretations has not prevented engineers from making significant progress in understanding and preventing failure phenomena, many feel that a clearer terminology would improve communications among engineers (especially in multi-disciplinary teams) and facilitate students in learning the complexities of failure as well.

Unsurprisingly then, many engineers have already proposed definitions of failure and related concepts that allegedly improve on the current situation. So far, however, these attempts have been only partially successful and have not resulted in a cumulative effort. As a result, each new proposal ends up adding to an already abundant terminology. It is also worth noting that these proposals are often motivated by the practical aim of bringing order within a certain domain after new failure phenomena have emerged thereby putting some strain on the extant terminology.

In this thesis, I take a different approach. The focus of my research lies more on the conceptual rather than practical aspects of failure. By surveying the engineering literature, I investigate how engineers *define* and *utilize* the concept of failure. The purpose of this investigation is not merely to deliver a catalog of definitions and document instances of utilization including those apparently in conflict with accepted definitions. The idea is that a conceptual analysis can deliver more than just a description of the current situation. A close inspection of the literature reveals a series of assumptions and conceptual distinctions which have not been fully spelled out and appreciated so far. From there, the analysis proceeds to delineate the preliminaries of a conceptual framework capable of rationally organizing the multiplicity of approaches retrieved from the literature. Furthermore, in line with the growing interests in sustainability and diffusion of integrated approaches to product development, this framework

somewhat expands the reach of the notion of failure and aims to take into account life cycle aspects of engineering products.

It should be stressed that, even though this work attempts to look closely at engineering language and conceptualizations, it does not do so by the traditional means of empirical studies, e.g., by interviewing engineers or recording their conversations while doing engineering work and dealing with failures. Indeed, this thesis is not an empirical study. This does not mean that, relying solely on the engineering literature, this thesis is disconnected from current engineering practice and its conclusions apply only to academic engineers whose views on failure are published in scholarly journals. In fact, the engineering literature abounds with papers, reports, case studies where the voice of practicing engineering is recorded, as it were. This part of the engineering literature can be seen as a proxy of actual utilization and provides valuable insights. Consider, for instance, the paper by Henshaw et al. (1999) which reports about an investigation into a series of failures of automobile seat belts. In Chapter 4 it is noted that the seat belt push-buttons, which play a central role in Henshaw et al.'s case, are described as failed *and yet* still functioning. The authors do not elaborate on this rather patent contradiction. Nevertheless, the fact that experienced failure investigators can entertain this problematic set of beliefs opens up an interesting perspective on the strategies that engineers may employ when confronted with the complexities surrounding the concept of failure.

By identifying such conceptually problematic areas and by showing that they can have a detrimental impact on knowledge sharing among engineers (e.g., Chapter 2), this dissertation constitutes a preparatory work for future empirical studies which could document more precisely the extend of conceptual and linguistic disagreement as well as assessing the effects on engineering practice.

A further aim of the papers collected in this dissertation is to attract the interest of philosophers of technology whom, so far, have rather neglected the study of technical failures. Of course, philosophers are fully aware that the possibility of failure is intrinsic to technical artifacts. Still, other artifact properties, notably functional properties, have received much more philosophical scrutiny. Typically, failure, or *malfunction* in the philosophical jargon, features as just an appendix to analyses of artifact functions. Many have noted, for instance, that one of the most serious shortcoming of Cummins' (1975) causal theory of functions lies in its inability to account for malfunction. In their recent monograph on technical functions, Houkes and Vermaas (2010) include the ability of

coping with malfunction as one of their *desiderata* for a sound theory of functions. Nevertheless, malfunction itself falls short of occupying the central stage.

Interestingly, the few philosophical studies that engage directly with the concept of malfunction display conflicts of intuitions not entirely dissimilar to those involving failure among engineers. Some philosophers tend to treat malfunctioning as total lack of functionality, like in malfunctioning knifes "which fail to cut, or broken corkscrews, which fail to uncork bottles" (Jespersen and Carrara: 2011, 122). Others, like Barros (Barros: 2013, 467) elect to "distinguish between mechanisms that fail (i.e., those that do not operate at all), and those that malfunction (i.e., those that operate, but do so in an unexpected way)". Thus, the investigation presented in this thesis might prompt philosophers into taking diverging intuitions about failure and malfunction more seriously. I anticipate that a better understanding of failure not only will be beneficial to philosophical explorations of technical artifacts, but will also contribute towards studies of ethical issues posed by failures and their consequences, chiefly with respect to allocation of responsibility.

1.1. Multiplicity of definitions

Conceptual disagreements about failure can dig rather deeply. *Construction Failure* (1997) is a well-known textbook on failures in the building industry written by two authorities in the field, Jacob Feld and Kenneth L. Carper. Among the first issues dealt with in the book are the causes of failures which the authors classify in a bunch of categories such as *Design errors, Construction errors*, and *Material deficiencies*. The category labels are quite self-explicatory. The interesting bit is a short remark at the beginning of the *Material deficiencies* section where Feld and Carper observe that "Some would claim that materials do not fail; people fail" (20). What they have in mind is the common sense notion of failure as breakage or rupture as exemplified by iron bars twisting and concrete pillars crumbling. However, Feld and Carper cannot help but think of a further connotation of failure, namely the idea of lack of adequate performance or, more precisely, *culpable* lack of adequate performance. But then, how can we blame an iron bar for having failed given that it was exposed to a corrosive environment that iron is not capable of withstanding. Whose fault is that? At the end of the day, iron *is* supposed to corrode (i.e., to fail) in such an environment. Instead,

given the circumstances it is the engineer who is supposed to select a different and more appropriate material.

In a short paper aptly titled *What is Failure?*, Roderick Rees contends that "It is no more than disreputable mythology to assume that failure means that something is broken" (1997, 163). In his opinion, artifacts like electrical fuses and bomb shells show that successful performance does not depend on physical integrity. If it has fulfilled its function, the spent electrical fuse should not be described as having failed, instead "it is the *function* that might be in a failed condition" (163, emphasis in the original).

To understand what failure is, then we need to answer the question: what does fail? Is it the artifact, its function, or the engineers who designed it? Clearly, intuitions about this issue are conflicting, and that is not the only area of disagreement. One of the first steps in this PhD research project was indeed to get a sense of the extent of disagreement and to this purpose a survey of engineering definitions of failure was performed which was initially published as a research paper in 2011 and now is included as Chapter 2 of this thesis. The survey shows that the most popular alternative consists in taking individual artifacts as the subjects of failure. For instance, the failure terminology presented in the *International Electrotechnical Vocabulary* (IEV) (1990) clearly presupposes that failures are predicated of individual artifacts or, to follow the terminological approach adopted in the vocabulary, *items*. Failure itself is defined as "the termination of the ability of an item to perform a required function". Thus, it is items that fail, and they do so when they *become* unable to perform a required function. Failure, the vocabulary specifies, is the event that coincides with the transition between the state or condition in which the item is able to perform its function and the state in which the ability is missing, the latter being called the *fault state* or just *fault*.

Even though the IEV definition has gained a prominent position in the literature, not everyone agrees with its proposal and, in particular, with the idea of making failure dependent on the item's *ability* to perform a required function instead of the item *actually* performing it. Consider, for instance, the attic of a house where an electric switch is installed which is utilized only infrequently. By idly sitting there, the switch gradually corrodes until it loses the ability of letting current pass through. According to the IEV definition, we should say that a failure event has occurred precisely when the accumulation of corrosion renders the switch unable to perform its function, even though its function is not re-

quired at the moment of failure nor will be anytime soon. How to describe, then, the event which will eventually occur when a user climbs to the attic, pushes the corroded switch and the light does not go on? Intuitively, *that* is the when the switch fails.

In fact, definitions of failure in terms of current performance are quite common in the literature. One example can be found in Birolini's *Reliability engineering* (2007, 3) textbook, where the IEV definition is abridged as thus: "A failure occurs when the item *stops performing* its required function" (emphasis added). Yellman (1999) explores another option consisting in a clear demarcation between, on the hand the concept of *functional failure*, and on the other hand the concept of *material failure*. The former is defined as thus, "Functional failure: Unsatisfactory *performance* (e.g., an item delivering unsatisfactory outputs) occurring during a process as operation or testing. (7, emphasis in the original); while the latter reads as follows, "Material failure: An undesired *physical condition* (e.g., an internal part of an item being damaged or broken) which is also permanent (i.e., it will persist until it is repaired). Such a condition could exist during operation or testing – or during a time there is no demand on an item to function at all" (7, emphasis in the original). So, in Yellman's view, the above mentioned switch first incurs into a material failure which, only later, becomes manifest as a functional failure.

Fuelled by conceptual disagreements of this sort, the list of failure definitions available in the literature keeps on growing. Indeed, being based on a sample of just thirty definitions, the 2011 survey does not pretend to be exhaustive of all conceivable alternatives. In fact, my personal collection of definitions is still expanding and each new entry contributes to its diversity. The reader can find these new entries in the form of a second appendix attached to Chapter 2. Again, no pretense of completeness is made. Still, the wide range of disciplines surveyed and the variety of solutions provides a good impression of the range of perspectives maintained within the engineering community.

As much revealing as they might be, definitions can reflect only partially the conceptual difficulties encountered by engineers in trying to regiment the notion of failure. Indeed, it has been a crucial characteristic of the methodology followed in this research to study the available literature in search of examples where engineers actually utilize the concept to make sense, describe, and analyze cases of failure. Therefore, textbooks on failure analysis and forensic engineering have been a primary source of information along with journals such as

Engineering Failure Analysis, *Journal of Failure Analysis and Prevention*, *Safety Science*, and others.

Valuable insights have been found elsewhere as well, particularly in publications related to Life Cycle Engineering (LCE). As pointed out by Vincenti, social and environmental concerns have broadened the set of requirements that engineers must deal with thereby expanding the reach of the concept of failure.

1.2. Life Cycle Engineering and the evolving concept of failure

LCE can be described as a "decision-making methodology that considers performance, environmental and cost requirements for the duration of a product" and which is becoming a norm in product development (Wanyama et al.: 2003, 307). According to Ishii (1995), LCE emerged as an extension of another methodology, Design for Manufacturability, which proved beneficial to many US manufacturers in improving product quality, reducing cost, and shortening development cycles. LCE extends on it by taking into account other stages in the life cycle of products besides manufacturing and by attempting to minimize environmental impacts through limitation of raw materials, energy, and emissions. In Ishii's view, "LCE seeks to maximize a product's contribution to the society while minimizing its cost to the manufacturer, the user, and the environment" (42).

Ishii's remark points out a characteristic of LCE that has significant implications for the conceptualization of failure: by emphasizing that multiple stakeholders are involved in the life cycle of a product, it challenges the predominant role that the end user and functional performance play in traditional approaches to failure such as those exemplified by the IEV definition mentioned above. Admittedly, with so many different definitions competing against each other as shown in Chapter 2, generalizations might be somewhat arbitrary. Still, Chapter 3 argues that many well-established definitions are rooted on four shared assumptions: *missing functionality*, *utilization context*, *item level*, and *negativity assumptions*.

Jointly, these assumptions define a view on failure that could be described as *event-oriented* and contrasted with a *goal-oriented* view that descends from a life cycle approach and that partially dispenses with them. Recall Vincenti's words about traditional engineering: rivets must hold tight and wings must provide lift, those are their functions and that is where failure criteria are deduced from. In

essence, that is the missing functionality assumption: failure occurs when an item stops performing its required function or, according to other formulations, it loses the ability of doing so. To put it differently, by conceptualizing item functions as measurable output (be it the force exercised by a rivet, lift generated by a wing, or flow of electrical current through a switch), failures are associated with abrupt events in the utilization stage an item. Sure enough, the physical mechanism eventually leading to the failure event can be gradual and develop over an extended period of time, for instance in case of corrosion or fatigue. Yet, the failure *event* is said to occur when the measurable output (or the ability of delivering it) has trespassed a predefined threshold.

The utilization context and item level assumptions further narrow down the domain of failure by stipulating that failure events are predicated of individual items (as opposed to groups or entire types of items) while they are deployed in their operational environment. Thus, other life cycle stages like manufacturing, servicing, or disposal are not covered. Finally, the traditional approach conceives of failures as negative or detrimental occurrences that should be avoided even when their consequences are minimal.

The goal-oriented view on failure that stems from a life cycle approach does not directly contradict the traditional approach; yet, it includes situations and events that violate one or more of the latter's assumptions except for the fourth one, negativity, which is preserved. Whereas the traditional approach focuses on the end-user's needs and the item's functional performance that is expected to satisfy them, the life cycle perspective takes into account needs and requirements of multiple stakeholders whose interests may lie in anyone of the life cycle stages, from supply of raw materials to manufacturing and recycling. Thus, product properties which are not directly related to functional performance and yet have an impact on stakeholders' interests now become relevant with respect to failure judgments.

LCE consists of a variety of methodologies targeting specific stages in the life cycle of products. Ishii (1995), for instance, distinguishes between Design for Production, Design for Assembly, Design for Service, and Design for Product Retirement. Together, these methodologies are intended to help engineers assess the life cycle implications of a candidate design and identify alternatives for improvement. Crucially, they are "most effective at the layout design stage, at which time the design is still preliminary and many decisions are uncertain" (43). It is during this stage that, among other things, designers must define the

high-level goals their products are expected to achieve and the adoption of LCE methodologies will prompt them to expand the list of goals beyond the traditional *purely technical* domain.

The adverb *purely* that I borrow from Vincenti should be clarified. The idea is that goals are purely technical when they have a technical origin and are satisfied by technical means. Having a technical origin basically means that a goal stems from the imperative of realizing a product that works. Intuitively, maximization of recycled materials after retirement is not strictly required for a product to work. Nevertheless, non-purely technical goals fall within the province of engineering because engineers have or may develop technical solutions to achieve them. The goal of maximizing recyclability, for instance, can be pursued by selecting specific materials or by means of product architectures that facilitate disassembly.

The most relevant consequence of a life cycle approach on the concept of failure is a shift from an event-oriented view to a goal-oriented view. In the latter, failure is no longer conceptualized as a discrete occurrence in the history of an item. Instead, failure judgments are based on the ability of products to achieve predefined goals that may involve anyone of the stages in the life cycle. Thus, a product may come to be regarded as a failure because it cannot achieve goals set for the manufacturing stage or for the disposal stage. Consider a personal computer whose enclosure is made of plastic and has been designed to achieve the goal of full recyclability. In a study on design for recycling of computer enclosures, Masanet and Horvath (2007, 1807) have shown that "PC enclosure components with a mass of 25 g or less would be discarded (a common practice for small plastic components)". The discarded components detract from the recycled fraction and can cause the product to miss the established goal, thus leading to a product failure in the retirement stage.

Though less common than event-oriented ones, goal-oriented definitions of failure can be found in the literature. The analysis performed in Chapter 3, however, concludes that the definition most suited to capture the concept of failure in a life cycle perspective is the one originally proposed in (Del Frate et al.: 2011), that is the survey paper featuring as Chapter 2 of this thesis. The definition advanced there claims that, from a life cycle perspective failure is:

> The inability of an engineering process, product, service or system to meet the design team's goals for which it has been developed.

This definition and the analysis supporting it are proposed to the engineering community in an attempt to foster a discussion on the concept of failure and on recent developments resulting from the widespread adoption of LCE. The intended audience of this thesis, however, is not just the engineering community. To engage philosophers of technology in dealing with conceptual issues connected to failures has also been a primary aim of this research project.

1.3. A dual audience

Formal ontology is one of the areas where the typically diverging interests of the philosophical and engineering communities can find common ground. Philosophers are attracted by the prospect of gaining clarity on fundamental conceptual issues some of which have kept philosophers busy for a very long time. Having a more pragmatic attitude, engineers see formal ontology as instrumental for the development of software tools aimed at representing and sharing engineering knowledge.

Previous research has shown the benefit of archiving knowledge about failures and making it available to designers, e.g., (Collins et al.: 1976). Recently, attempts have been made at extending available formal ontologies in order to characterize the concept of failure, e.g., (Kitamura and Mizoguchi: 1999; van der Vegte et al.: 2002; Koji et al.: 2005; Borgo and Leitão: 2007). Chapter 4 seeks to contribute to this growing body of research by building on the results of the previous two chapters. It has already been observed, e.g., by Borst (1997) and Guarino et al. (2009), that to reap the benefits of formal ontologies researchers should identify the main ontological commitments shared within the intended user community. First, if the formal ontology does not reflect these commitments users will find it hard to understand and utilize it. Second, formal ontologies should bring out "what is really shared by the community [of users] in order to enhance reuse *within* this community" (Borst: 1997, 123, emphasis in the original).

The aim of Chapter 4, which was originally published in the 2012 Proceedings of the *Formal Ontology and Information Systems* conference, consists indeed in carrying out this kind of preliminary work. It envisages a high-level formal ontology whose intended user community spans over all engineering disciplines thus requiring a very general concept of failure. As mentioned above, given the amount of alternative definitions and conceptual disagreement, finding a

common ground is highly problematic. For this reason, the paper focuses on event-oriented concepts of failure that are prevalent in the literature and have inspired some of the most influential definitions, particularly the failure definition given in the *International Electrotechnical Vocabulary*.

Even within this smaller domain, definitions of failure have been developed that result into opposing judgments. The paper distinguishes between three concepts, *function-based*, *specification-based*, and *material-based* failure. By means of an exemplary case story, the paper shows that the three concepts are mutually independent: an event that classifies as a failure given say, a function-based concept, could be classified otherwise by the other two.

Nevertheless, the paper argues that at the most abstract level these three concepts are based on the same ontological outline. The basic ingredients are constituted by the ontological categories of *occurrent* and *continuant*, and the *participation* relation. For all three concepts, failures are represented as *atomic occurrents* in which physical items *participate*. Physical items belong to the ontological category of *continuants*. States or conditions, on the other hand, belong to the *occurrent* category. More precisely, they are classified as non-atomic occurrents, because, differently from events, they have temporal parts. Two states in particular are singled out in the representation of failures. First, there are functioning states, that is to say those states in which items are performing as expected; again physical items are said to participate in functioning states. After a failure event has happened, physical items are said to participate in a second sort of states, namely, failed states. Since the three concepts of failure analyzed in the paper share this fundamental ontological structure, an engineering ontology capable of representing their mutual differences will need to deploy a set of ontological categories broader than the minimal set discussed here.

A further theme that ranks high both on the philosophical and on the engineering agenda is causality. Philosophical studies on the concept of causality are legion and date back to very origins of the discipline itself. On the other hand, the engineering literature is catching up rather quickly although, quite understandably, practical aspects tend to dominate over conceptual studies. The bulk of the literature deals with the study of causal processes responsible for failures and with methods and tools that allow engineers to ascertain causal factors from post-failure evidence. That does not mean that conceptual problems have passed unnoticed, though. Especially with respect to failures of complex systems where many factors of disparate nature are involved (e.g., organizational and technical

factors) it has become clear that intuitive notions of causality may be inadequate thereby leading engineers into drastically simplified accounts of the events.

In looking for a better understanding of causality, engineers have found that philosophical research can provide valuable insights. Lewis' (1973) theory of counterfactuals, for instance, provides the conceptual backbone to the Why-Because Analysis, an accident investigation technique developed by Ladkin (2000) with the objective of making causal investigations more rigorous. Johnson's (2003) handbook on accident reporting identifies in Mackie's (1974) *Causal Fields* and Hausman's (1998) *Causal Asymmetries* "two key theoretical ideas that must be considered when developing appropriate techniques for the analysis of adverse events" (900). Another example is Kuntz et al. (2011) work on Fault Trees, a technique utilized both by designers to prevent failures and by failure investigators to narrow down potential causal factors, which builds upon Halpern and Perl's (2005) structural-model approach to causality.

Some of these studies imply rather subtle conceptual distinctions and in some cases (e.g., Halpern and Pearl's structural-model) may lead to sophisticated logical formalisms. In contrast, the contribution presented in Chapter 5 of this dissertation relies on a relatively simple philosophical apparatus while paying considerable attention to the engineering side of the literature. The chapter itself has been previously published in *Engineering Failure Analysis*, a leading failure analysis journal, whereas a previous version was presented at the 2010 International Conference on Engineering Failure Analysis. The paper then, stems from an attempt to bridge the gap between philosophy and failure analysis by discussing a controversial engineering concept, *root cause*, and does so mainly by discussing the often overlooked distinction between backward-looking and forward-looking causality.

Understanding the causes of failure is crucial for developing corrective action and for prevention. Barring the most mundane and typically inconsequential failures which are easily explained, investigation of major failures involving complex technology is a complicated task that requires specialized skills. In fact, investigations are often carried out by multidisciplinary teams covering a wide range of disciplines. The most immediate challenge consists in reconstructing the sequence of events by collecting and analyzing material evidence, which, in some cases, could be limited due to the destruction brought about by the failure event itself. The sequence of events can be seen as the investigator's response as to the question: *What happened?* On top of that, they are also expected to answer

a further and arguably much trickier question: *Why?* Notably, the latter question is often thought to be synonymous of: *What caused it?*

The notion of causality does not appear completely out of the blue. Causal connections already start to emerge when the sequence of events is analyzed in detail. Indeed, causal connections must be identified if the sequence of events has to become a coherent whole instead of a mere series of snapshots. Thus, barring the mere chronology which is purely descriptive, causality is needed in order to tell *what* happened. Then, many engineers have assumed it is only natural that *one cause* should also provide the answer to the *why* question. From this assumption, the concept of *root cause* emerged that is to say, the idea that among all the causal factors involved in a failure it is possible to identify one which does not have antecedents therefore being "more fundamental" (Busby: 2001, 1419). Or, to put it differently, a root cause would be "the absolute beginning of the chain of events" (van Vuuren: 1999, 19). Related to root cause is the idea that causes can be ranked from the least responsible to the most responsible. Wood and Sweginnis (2006) recall that, until recently, aviation accident investigators in the United States were required to prioritize causes proportionally to their contribution to the accident. Still today, investigation reports issued by the US National Transportation Safety Board (NTSB) conclude with a "probable cause statement" singling out a few or preferably one single causal factor.

Despite the prestigious example set by the NTSB, the concept of root cause has met with criticism from many quarters. By independently going through the same path already followed by generations of philosophers, failure analysts recognized that operationalizing the concept of root cause runs into insurmountable conceptual difficulties. Also, root cause statements have been repeatedly interpreted by the public as allocation of liability, which falls outside the mandate of safety boards and is the responsibility of judiciary investigations instead. Finally, many have argued that the "root cause seduction" (Carroll: 1995) diverts investigators from their primary goal of finding lessons that can prevent reoccurrence.

Persuaded by these objections, safety agencies around the world are moving away from the "probable cause statement" and trying to distinguish their work from that of judiciary investigations by using causal terminology parsimoniously. Recently, the Australian Transportation Safety Board (ATSB) decided to expunge the term *cause* from its official accident reports altogether and deliberated that the expression *contributory safety factor* should be adopted instead. The

fact is that these recent developments seem to consist mostly of terminological adjustments which, eventually, do not challenge the assumption that causal factors, or safety factors as the ATSB would say, can somehow be prioritized based on their respective contributions to the final event.

Chapter 5 examines the concept of root cause and seeks to understand whether it is possible to reconcile the different views expressed in the engineering literature, particularly between the need to understand why a failure happened and how to prevent reoccurrence. The paper analyzes failure investigations as constituted of two sub-investigations. One is a backward looking investigation whose aim is to unearth the causal structure of events which eventually culminated into the failure event. The underlying concept of cause is deterministic and token-based, meaning that causal factors link deterministically clearly identifiable entities or events. The second sub-investigation is characterized by a probabilistic and type-based concept of cause. The causal factors identified by the backward looking investigation provide the grounds for developing potential failure scenarios that may happen in the future. The aim is to understand which factors are likely to reoccur and where corrective measures are more likely to be effective. In the forward looking perspective investigators are looking for probabilistic causal connections between types or categories of events which are based on already known causal factors.

Differently from claims about the causal connections that hold the sequence of events together, which may have strong empirical support, claims about *future* causal connections and scenarios envisaged by the forward-looking investigation are less certain and can only be expressed by means of probabilities. Still, for the investigation to achieve tangible improvements, it should motivate why a certain countermeasure (e.g., redesign of a component vs. revision of maintenance procedures) is going to be *most* beneficial in preventing reoccurrence. The factor targeted by that countermeasure *is* the root cause, which Chapter 5 proposes to conceptualize as a U-turn between the backward looking and the forward-looking investigations. The root cause of a failure, then, is that element of the factors and causes which, if corrected in future scenarios, is the most likely to prevent similar events from happening again.

Understanding the causes of failures and striving for prevention introduces the topic of the sixth chapter that concludes and to some extent summarizes this thesis, namely learning from failures.

1.4. Learning from failures and beyond

Because of their personal participation, failure analysts, safety experts, and forensic engineers are acutely aware of the amount of resources needed to effectively learn from failures as well as of the conspicuous potential benefits. The accurate study of failures and their causes not only can help engineers in preventing reoccurrence; on many occasions it has provided crucial insights eventually leading to new engineering knowledge and innovative designs. Consequently, many have come to believe that, indeed, *in engineering more is learned from failures than from successes*.

This belief has found in Henry Petroski a strong and enthusiastic advocate who has added case histories in its support coming from all epochs of engineering. Petroski reckons there is something paradoxical in claiming that more is learned from failures than from successes. At the end of the day, be it engineering, science, or literature, every student is taught to learn by looking at the masters, those who achieved remarkable success in their field. No teacher in her right mind would urge students to study a topic by following the example of those who egregiously failed. So, what concepts of failure and learning do Petroski and his sympathizers have in mind?

By looking closely at the case stories and at the arguments advanced in its support, it turns out that the paradoxical claim about learning is actually a twofold hypothesis, a *specific-learning hypothesis* and a *generic-learning hypothesis*. In both cases, failure is conceptualized from a goal-oriented perspective as the inability on the part of an engineering product to meet the goals for which it was developed. The two hypothesis, however, depend on two different interpretations of learning. According to the *specific-learning hypothesis*, the epistemic agent (i.e., the subject who learns) is either an individual engineer or a well-identifiable group of engineers (e.g., a design team or an engineering organization). The adjective *specific* indicates that the design goal facing the epistemic agent comes with clearly specified metrics for success and failure. Consider, for instance, a team of aeronautical engineers whose task is to design a landing gear for a high-performance airplane. Already in the early stages of the design process they know what sort of goals a landing gear is expected to achieve (e.g., robustness, weight, reliability) and they can specify metrics to express degrees of achievement.

In this context, learning occurs when agents utilize knowledge gained through the study of failures (either their own or somebody else's) to keep their

designs safe from those failures and, generically, to improve upon previous realizations. The engineering literature repeatedly emphasizes that merely studying failures without implementing the lessons into practice should not be regarded as actual learning. Analysis is not learning, Carroll and Fahlbruch (2011) remark, and if nothing has changed then learning has not occurred.

Many of the case studies discussed in Petroski's works and elsewhere in the literature are instances of specific learning. Sure enough, those stories show that specific learning occurs and contributes to engineering products becoming safer and more reliable. Still, stronger evidence is needed to corroborate the hypothesis that more is learned from failures than from successes. Only recently a study has been published where the specific learning hypothesis has been tested empirically. Madsen and Desai (2010, 452) claim evidence collected from the orbital launch vehicle industry allowed them "disaggregating organizational experience into failure experience and success experience and comparing the contribution of each to organizational performance". They conclude that, although organizations learn both from failures and from successes, on average more is learned from failures. According to Madsen and Desai, the reason lays in the disproportionate effect played by large-scale failures such as orbital launches dramatically falling short of achieving their goals. These events are likely to result in thorough reassessment of available knowledge and revision of current procedures with long lasting beneficial effects on future activities. Madsen and Desai acknowledge their study is not conclusive. Being based on data from a specific – and rather peculiar – industry, its results may not easily generalize over other fields. Nevertheless, it constitutes a significant step forward in the discussion about the specific-learning hypothesis that might stimulate the realization of further empirical studies.

When looking at the second learning hypothesis, the *generic-learning hypothesis*, the prospects of advancing the debate by means of empirical studies seem less straightforward. Differently from the specific, the generic hypothesis is much more ambitious and far reaching in that it aims at explaining no less than technical change on a global level. In Petroski's words:

> The failure of the Titanic contributed much more to the design of safe ocean liners than would have her success. That is the paradox of engineering and design. (Petroski: 2006, 96)

The reason is that, in his view, "the science" (Petroski: 1985, 97) of engineering structures – be it ocean liners, commercial airliners, or suspension bridges – can generally be said to have benefited more from failure events than from instances of success.

Petroski apparently believes that growth of engineering knowledge and technical change are just a direct consequence of specific learning. Since engineers or engineering organizations that learn from failures are more likely to avoid recurrence and design reliable products than those who focus mainly on examples of success, then at an aggregate level it turns out that failures are more effective than successes in shaping engineering knowledge and technical change.

Sure enough, there are many cases where the study of failures contributed decisively to the advancement of engineering knowledge and practice. The crashes of Comet airliners in the early 1950s are a prominent example and, as noted by Wanhill (2003), deserve to be considered *milestones* in the history of aircraft structural integrity:

> The Comet accidents and subsequent investigations changed fundamentally the structural fatigue design principles for commercial transport aircraft. (Wanhill: 2003, 65)

Nevertheless, Petroski's belief that generic learning simply follows from specific learning rests on questionable assumptions. In particular, it presupposes that the two phenomena involve the same concept of learning, which, I argue, is not the case. Specific learning implies a well-defined epistemic agent (either an individual or an organization) dealing with a specific engineering problem (e.g., to design a landing gear for a high performance airplane) for which there are clear, albeit qualitative, design goals. Thus, criteria can be devised to decide whether or not learning has occurred (e.g., a new landing gear has been designed which avoids a failure mode observed on previous models). Crucially, for this form of learning to occur it is not necessary that any advancement in engineering knowledge has been achieved. Learning may consist in the implementation of lessons that were already available within the engineering community's shared body of knowledge although some agents might have been unaware of it. The contribution of failure, in these cases, resides in making those agents aware that something went amiss and additional knowledge is needed. Nothing really new may have been added to the extant body of knowledge,

though. Specific learning, then, mostly consists of *diffusion* of knowledge that was already available.

The generic-learning hypothesis, however, deals with technical change and the *advancement* of engineering science which is typically regarded as a process in which new knowledge is actually *generated*. That is why historic episodes like the Comet disasters are treated as milestones. What surfaced from the investigations were fundamentally new knowledge and new design principles. In what sense could such a process be seen as a form of learning analogous to specific learning? First, the well-identified epistemic agent is no longer there and has been replaced by a diffuse entity, the loosely connected community of practitioners. Second, the community does not have a design goal of its own nor failure criteria that can be utilized to decide whether or not learning has been achieved. Nevertheless, we might be willing to say that learning has indeed occurred for the newly acquired knowledge allows the engineering community to provide society with innovative products that ostensibly outperform their predecessors on several aspects. Yet, the generation of new engineering knowledge and the improvement of technical performance are not equivalent to technical change because the latter includes a further aspect which lies beyond the engineering sphere of influence, namely what society does with technology. Whether a new technology is fully embraced by society, whether it remains confined in niche markets, or is rejected altogether, does not derive directly from its technical merits and the amount of engineering knowledge spent on it. Engineers may be able to improve technical performance of about anything, yet adoption by society does not follow automatically for economic and social factors play a crucial role in the process. The point here is the following: promoted either by failure or by success or by fundamental research, engineering knowledge may grow and allow engineers to improve existing products or create new ones. However, while the realization of a successful prototype deserves to be considered as an engineering advancement, it does not count yet as an instance of technical change. For that to occur, the prototype needs to be turned into a product which is adopted by society, at least for a while. The last part of the process, technology adoption, may well happen without any further improvement in engineering knowledge. Thus, growth of engineering knowledge and technical change, while undoubtedly linked, cannot be considered as just two faces of the same phenomenon.

Prevalent in many of Petroski's case studies are safety concerns: bridges collapsing, ships sinking, and airplanes crashing. Almost inevitably these case

studies follow a plot along these lines: failure strikes a certain technology, engineers learn the lessons and engineering knowledge grows in the wake of failure, safety is improved, and eventually the improved technology spreads. To put it differently: by focusing on safety, Petroski can easily show that technical change almost invariably follows growth in engineering knowledge. The emphasis on failure, however, may convey a distorted picture of technical change for safety ranks high among social values and it is rather unlikely that technical improvement on safety will be ignored. Sure enough, modern ships are safer than the Titanic and her sinking has a lot to do with it. Similarly, safety of modern airplanes owes greatly to the Comet's crashes and other disasters. Nevertheless, equating technical change with safety improvements would be a gross simplification. Modern engineering artifacts are not only safer, they are also more energy-efficient, less polluting and they have also become interactive and mass-customizable. Think of today's cars with their aerodynamic shape, recyclable materials, electronic gadgets, and endless lists of optional features and compare them to Ford Model T which was more or less contemporary to the Titanic. Undeniably, automotive engineering knowledge has grown substantially: today's engineers master materials, structures, and processes much better than they used to do ninety years ago. The point is that the generic-learning hypothesis assumes these improvements in knowledge *automatically* translate into technical change and society embracing engineers' latest achievements. This way it neglects one of the main lessons learned from recent historiography of technology, that is to say the role played by extra-technical factors in technical change. Social factors influenced what the automobile means for modern society: its being a means of transportation but also a status symbol, a source of pollution, a potentially dangerous device, and many other things.

In conclusion, the generic hypothesis appears to derive from an outdated *internalist* view of technical change and does not survive close scrutiny. On the other hand, by analyzing it this thesis somehow completes a full circle and finds itself one more time in agreement with Vincenti and his remark on the ever expanding scope engineering challenges and the inclusion of social and environmental matters (Vincenti: 1990, 255).

When looking at the outcomes of engineering activities from a long-run perspective, as is done in the generic-learning hypothesis, the sharp contraposition between failure and success, which works fine as far as the special-learning hypothesis is concerned, begins to fall apart. Products that initially perform well

according to the goals set by their makers may turn out to be less successful than expected or may be found responsible for unwanted social side effects. Similarly, innovations considered capable of becoming dominant have gone extinct prematurely and products thought inferior or surpassed have shown unexpected longevity. To understand technical evolution and the lessons that can be learned from it, the clear-cut contrast between failure and success should be abandoned and it should be recognized that both hold in store valuable lessons.

This reassessment of the failure-success distinction resonates well with one of the main tenets in Resilience engineering, an approach to safety in complex systems that recently has received widespread attention. Hollnagel et al.'s (2008, xi) state it very clearly at the beginning of their book as follows:

> Resilience engineering makes it clear that failures and successes are closely related phenomena and not incompatible opposites.

This interesting convergence of views between Resilience engineering and the analysis of the learning hypothesis conducted in Chapter 6 constitutes a further example of the kind of valuable insights that can be harnessed from the analysis of failure. Its investigation, however, falls outside of Chapter 6's aims and will constitute material for future work.

This thesis is an attempt to clarify a concept with which we are all familiar, engineers and non-engineers alike. It has shown that, behind the first impression of familiarity, there is a wide range of intuitions about failure which are not easily reconciled. While the ensuing ambiguities and lack of clarity may be tolerated in ordinary circumstances, engineers strive for precision and efficiency. These qualities become even more relevant given that engineering activities are increasingly being carried out by multidisciplinary and multicultural teams.

The chapters included in this thesis illustrate that pursuing conceptual clarification may result in valuable contributions to the existing literature. The identification of tacit assumptions that, so far, have gone undetected can help bringing some degree of order to discussions that have shown a tendency towards fragmentation along disciplinary boundaries. In the case of root cause, for instance, shifting the emphasis from practical matters to conceptual aspects has shown that backward looking and forward looking views typically seen as mutually exclusive actually complement each other. Critical reflection on goal-oriented concepts of failure and its ties with life cycle engineering has broadened

the reach of failure beyond products' functional performance during the utilization stage. Finally, the investigation of the learning hypothesis' conceptual underpinnings has revealed that failure provides a stimulating vantage point to approach learning in engineering and the vagaries of technical change.

Taken together, these chapters constitute the preliminaries of a conceptual framework that, once supplemented with additional engineering and philosophical contributions, may embrace the multiple facets of failure, a rather complex tangle of phenomena which, despite engineers' efforts to rein it in, is not going to disappear from the engineering agenda anytime soon.

2 Towards a Trans-disciplinary Concept of Failure for Integrated Product Development[1]

Abstract

Integrated product design approaches presuppose knowledge sharing among cross-functional teams. In this paper, such sharing is considered for failure phenomena. It aims at finding a trans-disciplinary definition of failure that facilitates the communication of knowledge about failures between the different engineering disciplines. Four criteria are given that a trans-disciplinary definition of failure should meet, and a survey of engineering proposals to define failure is presented. It is shown that none of these existing definitions meets all four criteria, and that six come close by meeting three criteria. Finally, analyzing these six definitions, a trans-disciplinary definition of failure is proposed.

2.1. Introduction

Looking for increased profitability, companies are ever more shifting from the traditional sequential approach in product design towards integrated approaches based on the establishment of cross-functional teams, i.e., Integrated Product Development, IPD. A cross-functional team gathers together people with different background and expertise which are deemed relevant for the overall project with the aim that interacting and sharing their knowledge they will come up with optimal design solutions. Members are not exclusively designers or engineers, for they may come from non-technical departments as well, e.g., marketing, finance. And members need not to be only employees, since they may come also from outside the company (e.g., suppliers, customers, subcon-

[1] This chapter has already been published as Del Frate, L., Franssen, M., and Vermaas, P. E. (2011) 'Towards a trans-disciplinary concept of failure for Integrated Product Development', in: *International Journal of Product Development* 14 (1-4): 72–95.

tractors, etc.). There is ample evidence (see Section 2.2) that this variety of perspectives is potentially beneficial for both the performance of the design process and the performance of the final product itself. For instance, it may bring to light potential conflicts among product requirements, or suggest product features that have been overlooked. The key factor is that team members are successfully integrating their efforts, which means they are sharing knowledge that is relevant for the product and not merely discussing it. This depends on the availability of an effective communication system: members should be speaking the same language especially because team members have different technical (and possibly cultural) backgrounds. Therefore, the establishment of a common terminology, the agreement on key concepts and similar communication strategies are fundamental preconditions for the successful integration in cross-functional teams.

This paper deals with the issue of knowledge sharing about failure phenomena relevant to the product which is being designed during its life cycle. It is a vital design task to anticipate and to prevent as much failures as possible, and to do so in an early stage of the design process. Because of the availability of multiple competencies and skills, cross-functional teams are well positioned to perform this task. Again, this will happen only if team members will be able to share knowledge about failure phenomena. A survey of the engineering literature shows that a number of alternative definitions of failure are available. This is hardly surprising given that the study of failure phenomena and the methods to control and prevent them are scattered through many fields and disciplines (e.g., mechanical engineering, material science, reliability engineering, safety science). However, the co-existence of multiple, sector-based and partially overlapping definitions may be a factor in creating barriers in communication and in knowledge sharing therefore undermining the aims of integrated product development. The aim of this paper is to analyse available definitions looking for the ones that are better suited to facilitate communication in cross-functional team, which may be called trans-disciplinary definitions.

The outline of the paper is as follows. In Section 2.2, IPD is compared with the sequential model of product development and it is explained, through reference to the relevant literature, why the former approach is more demanding in terms of communication needs. Section 2.3 introduces the notion and crucial features of the "IPD failure domain". In Section 2.4 the criteria used to assess the candidate definitions are explained. The definitions themselves are listed in

the Appendix. Section 2.5 discusses the results of the assessment which are summarized in table form. Section 2.6 presents a new definition of failure which is based on the results of the previous analysis and which meets all the criteria. Finally, Section 2.7 will conclude the paper summarizing the work done and suggesting directions for future work.

2.2. From the sequential model to Integrated Product Development

In the sequential model of product design, the workflow is compartmentalized. Each department or team is assigned a specific task related to a limited section of the overall project. It performs its task nearly isolated from other departments to which, when ready, it delivers the results of the assignment. The next department receives the results as a given, as the starting point for a new task. The possibility of feedback between neighboring departments is not contemplated in the model since it is assumed that they operate in virtual isolation, as if encircled by a protective wall which allows only the final results to escape. For this reason, the model is also known as the "over the wall" model. Over the years, the limitations of the sequential model have become apparent. Firstly, the various departments specialize on different components or features of the system (e.g., mechanical engineers and electric engineers) or on different stages of the life cycle (e.g., manufacturing engineers and maintenance engineers). Secondly, and related to that, it may happen that several of the many requirements and specifications that a system has to meet through its life cycle are conflicting with each other or with constraints imposed by the company's manufacturing facilities. As a consequence, it may happen that the product development process experiences multiple crises each time a conflict of requirements emerges as the development proceeds. Because of the specialized knowledge the problem cannot be anticipated, and only be tackled when it occurs. And problems encountered in late stages of the development process, especially if they demand substantial redesign, are potentially very costly.

IPD has been introduced in the hope of overcoming the limitations of the sequential model. The basic principles of IPD were made popular by Andreasen and Hein (1987) who based their model on the ideas examined earlier by Olsson (1976). As documented in (Vajna and Burchardt: 1998), the concept has subsequently been interpreted in a variety of ways and a survey of the literature has revealed that there is no single definition of IPD (Hjort et al.: 1992). Some

authors, mainly from the US, consider IPD just a synonym of Concurrent Engineering or Simultaneous Engineering, e.g., (Syan and Menon: 1994; Haque: 2003; Boyle et al.: 2006). Others, like Prasad (1996), think that IPD is one of the themes of Concurrent Engineering. Still others maintain that IPD is an extension of Concurrent Engineering (Sage: 1995) or that it includes some of the best practices in engineering design like Concurrent Engineering, customer involvement and supplier involvement (Rauniar et al.: 2008).

Vanja and Burchardt (1998) propose a general definition which takes into account the contributions of (Andreasen and Hein: 1987; Ehrlenspiel: 1995) and the results of an international workshop held in Magdeburg, September 1996, in order to promote a common understanding of IPD. The definition is the following (Vajna and Burchardt: 1998, 6):

> IPD is a human-centred procedure for developing competitive products or services of high quality, within a reasonable amount of time, and with an excellent price-performance ratio. [...] IPD describes the integrated application of holistic and multi-disciplinary methods, organization forms, and both manual- and computer-supported tools with minimized and sustainable use of production factors and resources.

As shown by the results of the Magdeburg workshop, even though some divergences still remain, several points of convergence can be found. Prominently among these, at least for the aim of this paper, is a widespread agreement that IPD promotes the establishment of cross-functional teams working together and the integration of all products and processes (people, tools and techniques) (Roe: 1996).

A substantial body of evidence documents the benefits of the application of IPD in many sectors. However, it has to be noticed that the implementation of IPD procedures is a complex process involving several costs and without a guarantee of success (Crawford: 1992; Lullies: 2000; Ford and Sterman: 2003). As pointed out by Hoopes and Postrel (1999) several investigations – e.g., (Clark: 1991) for the automotive sector; (Iansiti: 1995) for mainframe computers; (Henderson and Clark: 1990) for semiconductor photolithography – show that differences in the performances of firms can be attributed largely to the degree of integration in their development processes. Three main integrating mechanisms have been considered in the literature, namely shared knowledge, cooperation and coordination between the components of multidisciplinary teams (Hoopes and Postrel: 1999). Even though the importance of cooperation

and coordination should not be underestimated, the relevance of shared knowledge and knowledge management in product development is the subject of increased attention (Grant: 1996). This is hardly surprising considering that product development can be seen as a knowledge producing and knowledge transforming activity. In brief, these studies show that successful product development is fostered by superior organizational integration and, in turn, organizational integration is promoted by knowledge sharing among the members of cross-functional teams. Consequently, the crucial issue is how to enhance and optimize knowledge sharing. Shared knowledge is defined as facts, concepts, and propositions which are understood simultaneously by multiple agents (Hoopes and Postrel: 1999). The process of knowledge sharing is actually a subject of research among psychologists, sociologists, philosophers, management scientists and many others.

It is widely acknowledged that within a particular firm, departments can create different formalities, routines, procedures that constitute interpretative barriers (Dougherty: 1992).

In their study of the antecedents and consequences of cross-functional cooperation Pinto and co-authors (1993) argue that "divergent interests and points of view are inevitable when individuals from multiple functional areas work together" and this creates the conditions for a potential lack of cooperation. Keller's (2001) study of cross-functional project groups in new product development found evidence in support of the hypothesis that "functional diversity makes internal communication among group members more difficult owing to differing functional goals, training and orientations". Roche (2000) portrays multi-disciplinary teams for concurrent engineering, especially virtual teams, as "a Tower of Babel" because "each enterprise's actor speaks his own language, with his own terms and meanings". And this multiplicity has detrimental effects on communication since "two entities can communicate only if they agree upon on the meaning of the terms they use".

These obstacles increase when, accordingly to IPD principles, integration has to be reached with parties external to the firm (e.g., customers, suppliers, regulators, and so on) creating teams that are not only multidisciplinary and cross-functional, but also multinational and multicultural (Ramesh and Tiwana: 1999; Lang et al.: 2002; Andersen and Drejer: 2009). An appropriate illustration of these issues is a comment, quoted by Anderson et al. (2008), from a Vice President of a major supplier firm in aerospace industry. In describing the

complications that arise in working with other firms, the Vice President gave the following example:

> Boeing is an excellent aerospace company. Yet, when we work [on a product] with them we find that we speak different languages. We have different words for the same thing. [...]. Most of our procedures don't even correspond clearly to theirs.

Although there is not yet a unified solution or model about how to overcome knowledge barriers in cross-functional environments, it is commonly accepted that thorough communication among participants and the establishment of a shared common language is important. Pahl et al. (2007), for example, emphasize that working in an interdisciplinary team requires the adaptation of language and terminology; Hendlund and Nonaka (1993) argue that the mobilisation and sharing of tacit knowledge are assisted by the availability of a common language; Dougherty (1992) argues that a common language is one of the factors effective in bringing the tacit thought-world differences to conscious awareness. On the other hand, as noted by Thomas (2005), the lack of a unified language can cause a split between those who are a part of the subculture and those who are not. Many inquiries have shown that a lack of communication and shared language are detrimental to product development. Rauniar, et al. (2008) and Hoopes and Postrel (1999) present a comprehensive overview and include references to relevant literature.

One important step in the establishment of a common language is the drawing up of a vocabulary, or a set of shared definitions since, as claimed by Roche (2000), effective communication presupposes agreement on the terms used. Likewise, the work by Olsen et al. (1994) on knowledge sharing in collaborative engineering "advances the opinion that collaborators need to establish and customize sharing agreements (i.e., mutually agreed upon terminology and definitions)". Poteet et al. (2008) investigated the linguistic aspects of miscommunications related to cultural differences and came to the conclusion that "standardization of terminology seems a very important useful strategy to reduce ambiguity and thus to avoid miscommunication". Standardizing terminology introduces formalized rules and procedures in communication which has proven effective in fostering greater cooperation within cross-functional teams (Pinto et al.: 1993; Miller and Guimaraes: 2005). Definitions are both used for exchanging information about explicit knowledge, and for assisting the process of externalization of tacit knowledge.

As for the notion of failure, which is the topic of this paper, many disciplines provide a definition that allegedly captures the knowledge relevant to their specialization. Moreover, there is a vast amount of engineering literature on failure in, for instance, engineering journals like *Engineering Failure Analysis*, *Journal of Failure Analysis and Prevention*, *Journal of Loss Prevention in the Process Industries*, *Prevention Science*, *Accident Analysis and Prevention*, and *Safety Science*. Nevertheless, the result of these efforts is that the notion of failure is described from various perspectives and in a piecemeal fashion under various rubrics: Failure analysis, e.g., (Wulpi: 1999; Tawancy et al.: 2004; Affonso: 2006); Forensic engineering, e.g., (Kaminetzky: 1991; Piésold: 1991; Carper: 2001; Lewis et al.: 2003); Risk analysis, e.g., (Ale: 2009); Reliability and Maintenance engineering, e.g., (Cox: 1998; Dodson: 1999; Birolini: 2007; Yang: 2007; Daley: 2008); mechanical design, e.g., (Collins: 1993, 2003); electronic packaging, e.g., (Viswanadham and Singh: 1998); fire prevention, e.g., (Hattangadi: 2000); history of technology and engineering, e.g., (Petroski: 1985, 2006).

Because of this fragmentation it is reasonable to expect that assembling a cross-functional team will gather as many different notions – resting both on explicit and on tacit forms of knowledge – as the participants. The analysis given in the previous pages supports the conclusion that this multiplicity of sector-based and partially overlapping definitions might hamper communication and knowledge sharing, thereby undermining the benefit of gathering a cross-functional team. Yet, there have hardly been attempts to phrase a notion of failure that could help knowledge sharing in cross-functional design teams. The aim of this paper is to identify a definition – or some definitions – that could serve as an effective basis of communication and knowledge sharing about failure phenomena in cross-functional design environments, what might be called a trans-disciplinary definition.

The inspiration for this research comes from analogous efforts aiming at finding general definitions for other significant engineering notions. Rausand and Høyland (2004), for instance, in the first pages of their book on System Reliability Theory acknowledge that "there is considerable controversy concerning which is the broadest and most general concept" of reliability and yet emphasize the importance that the concept is defined in an unambiguous way. To underpin the point they declare their agreement with Kaplan (1990) who claims: "When the words are used sloppily, concepts become fuzzy, thinking is muddled, communication is ambiguous, and decisions and actions are subopti-

mal, to say the least". Another concept that has been the subject of extensive discussion is 'quality'. Garvin (1984) noted that four disciplines – philosophy, economics, marketing and operations management – provide different and contrasting definitions each based on different analytical framework and each employing its own terminology. According to Garvin, this has generated "great confusion: managers – *particularly those in different functions* – frequently fail to communicate precisely what they mean by the term" (emphasis added). Garvin, then, proceeded towards building a unified framework able to address the issues raised by the competing definitions. And many other scholars followed. Yet, more than twenty-five years later the concept is still disputed as can be seen from the table of contents of Dale's (2003) recent book about Total Quality Management. Chapter 1 begins with a section entitled "What is quality" in which it is claimed that, although quality is now a familiar word, it has a variety of interpretations and uses, and there are many definitions. This is deemed unacceptable by Dale because it may generate "misunderstanding in the communication", both within the company, and outside when the company communicates with customers and suppliers. Therefore, Dale notes, within "an organization, to prevent confusion and ensure that everyone *in each department and function* is focused on the same objectives, there should be an agreed definition of quality" (emphasis added). Coming closer to the notion of failure, a concept whose multiple definitions have been a source of concerns is that of 'defect'. Failure analysts, like Becker et al. (2005), are aware that when the term is used in a failure report, it easily gets the attention of lawyers. However, this is not the only concern and they warn that "careless use of the word *defect* can lead to communication problems during design, manufacturing, service or post service". Unfortunately, multiple definitions are being used, e.g., (Davis: 1992; Becker and Shipley: 2002), and they are not always compatible. Even worse, Becker et al. (2005) argue that some of them are internally inconsistent, as they show analysing the definition in Davis (1992) *ASM Materials Engineering Dictionary*.

The search for a trans-disciplinary definition for this paper was structured as follows. First, the engineering literature was surveyed collecting a number of different definitions. The aims of the survey were both to collect authoritative definitions and to cover a broad range of domains in order to capture the multi-disciplinarity of cross-functional teams. Therefore, first entries were definitions from widely accepted standards as (IEC 60812: 2006) and (MIL-STD-721C: 1981) and also Leonard's (1982) definition which has been adopted by the

Technical Council on Forensic Engineering of the American Society of Civil Engineers (Carper: 2001). The authoritative definitions were collected, e.g., Hubka and Eder (1996) for design science, (Birolini: 2007) for reliability engineering, and Kaminetzky (1991) for civil engineering. Other disciplines and fields covered by the candidate definitions are root cause analysis (Mobley: 1999), failure analysis (Tawancy et al.: 2004), telecommunications (Jones: 2004), structural reliability (Melchers: 1999), system reliability (Rausand and Høyland: 2004), mechanical engineering (Collins: 2003). Encompassing multiple disciplines and specializations the survey provides also indications on the varieties of failure phenomena that could be relevant for a cross-functional team (e.g., mechanical, structural, functional, process failures). This is the failure domain which features are exemplified through case stories also retrieved from the literature.

In the second stage, criteria were formulated for assessing the capability of candidate definitions to capture failure phenomena included in the failure domain. Third, candidate definitions were assessed against the criteria. The definitions retrieved from the literature are listed in the Appendix in alphabetic order. The next section summarizes the features of the cross-functional failure domain. Then, the criteria are discussed in Section 2.4.

2.3. The cross-functional failure domain

We take the cross-functional failure domain as the aggregate set of events and conditions that are classified as failure phenomena according to at least one of the notions of failure present in the cross-functional team. Given the multidisciplinary background of the team members, this cross-functional failure domain is potentially very broad and heterogeneous. Still it has some fundamental distinguishing features and boundaries. First, failure in product development must be considered a technical notion and be distinguished from the overly generic and simplistic common-sense notion which labels as failure any breaking or rupturing event. From an engineering perspective, not all material failures (e.g., fracturing, breaking) are to be considered as (technical) failures. Second, events such as accelerated degradation or plastic deformation, in which no breakage or rupture takes place, can be technical failures. Finally, the IPD failure domain is not limited to material devices for it includes also processes (e.g., in manufacturing, assembly, maintenance, usage, and so on).

2.3.1. Technical failure and material failure

The IPD failure domain deals with *technical failures*; therefore it should exclude such common sense failure events like the blowing out of a light bulb well after its expected operational life. Not every fracture or rupture event has to be considered a technical failure event. Inevitably, at some point the filament of a light bulb breaks apart, but if the event happens well after the expected operational life it does not count as a technical failure.

Fractures, cracks and ruptures are usually referred to as material, mechanical or physical failure (Collins: 1993, 2003; Becker and Shipley: 2002; O'Connor et al.: 2002). However, material failures constitute only one of the areas or contributors constituting the failure domain with which engineers have to deal in IPD. The difference between technical failure and material failure could be illustrated by considering the design of structural fuse pins connecting the aircraft's engines to the wing (see Figure 2.1).

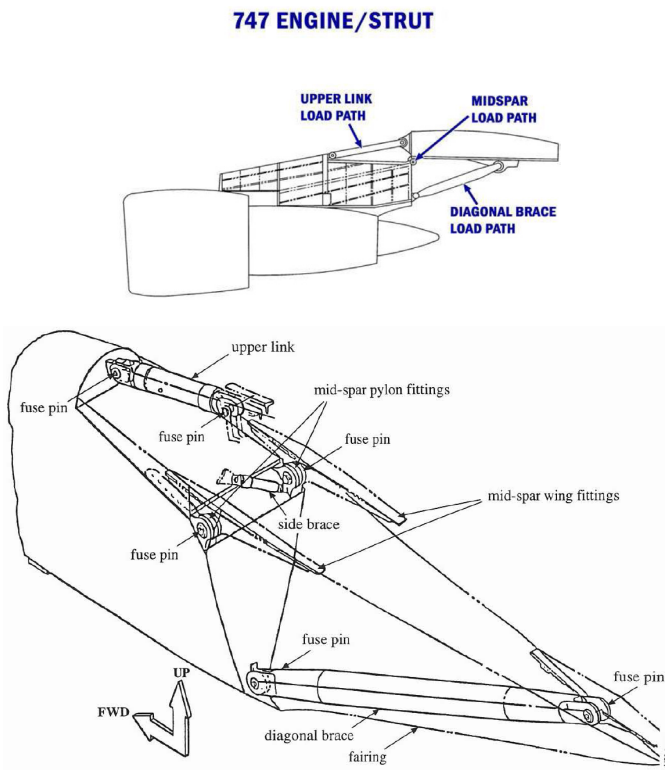

Figure 2.1: Engine-wing connection via fuse pins. From (Wanhill and Oldersma: 1997)

The fuse pins are designed so that they support the engines during normal flight operations, but they rupture and allow for clean engine detachment in case of overload or abnormal load conditions, e.g., heavy turbulence, hard landing, water landing, etc. (Wanhill and Oldersma: 1997).

Therefore, material pin rupture followed by clean engine separation is not to be considered a technical failure when the structure is overloaded. Only fuse pin rupture during normal flight operation is to be regarded as a technical failure. Since in such unfortunate events the structure behaves in aberrant ways, it could be that pin ruptures are only partial and/or not simultaneous, as they should be in case of a designed event. Therefore, instead of having a clean separation, the detached engine could tear off part of the wing or a nearby engine. According to the Netherlands Aviation Safety Board (1994) this is what happened to the El Al Flight 1862 that crashed in 1992 on the Bijlmermeer suburb of Amsterdam. After its departure, the Boeing 747 was climbing in nice weather conditions when engine number three detached from the right wing. In doing so, it damaged the leading edge of the right wing and impacted on engine number four, which then also separated from the wing. The accident investigation was able to ascertain that the fuse pins of engine number three failed because of metal fatigue (Nederlands Aviation Safety Board: 1994; Wanhill and Oldersma: 1997).

2.3.2. Degradation and deformation

On the other hand, the technical failure domain should include events where no ruptures take place, for example in case of premature degradation or excessive deformation. The structure or the component is still in place and able to perform its intended function, but its properties (mechanical, dimensional, etc.) no longer meet design specifications. Consider the following example by Tawancy, Ul-Hamid and Abbas (2004, 466–469). A contactor – a tank in the form of a welded cylinder 3 metre in internal diameter, 4 millimetres wall thickness, and 2 metre height – was part of a plant for the production of vinyl chloride monomer. After only three weeks of operation, a regular maintenance inspection found that the wall thickness, at some spots, was reduced from 4 to 0.4 millimetres. Fortunately, there was no rupture or leakage and the apparatus was disconnected for safety concerns. Inspections established that the component had been manufactured along specifications, yet the specifications for the contactor where drawn up under the assumption that most of the chloride acid should be consumed

prior to condensing in the contractor. Evidence suggested otherwise. Hence the contractor was subject to a corrosive environment it was not designed to withstand.

Even though there was no physical rupture, Tawancy, Ul-Hamid and Abbas (2004) qualify this situation as a technical failure. In doing so they are following a well-established habit in failure analysis and many other engineering disciplines.

2.3.3. Failing processes

Moreover, in IPD engineers are developing simultaneously both physical systems (and their components) and processes, by means of which those systems are manufactured, assembled, tested, maintained, and so on. A very incomplete list of examples includes processes like: heat treatments, galvanic treatments, welding, gluing, fastening, finishing processes, lubrication, cooling processes, and many more. Engineers draw up specifications both for the steps to be carried out during the process and for the expected results. The process performance is then assessed on the basis of these specifications. And, as for physical components, the possibility of a failure of the processes has to be taken into account. Therefore, failing processes are part of the failure domain. Let us consider two cases in point. The first case is presented by Tawancy, Ul-Hamid and Abbas (2004, 430–434). A steel alloy sheet was produced to have a specified creep strength (creep strain at 40 Mpa/925°C must be in excess of 15 hours). After processing, the sheet failed to pass the specification test. During creep testing, standard specimens from the sheet were ruptured after as short as 5 hours of exposure at 40 Mpa/925°C. Investigators found that the heat treatment was defective, more specifically the cooling rate was too slow. In fact, when the process was run again, this time assuring that the sheets were rapidly cooled, the product was able to pass the specification test.

The second case is about surface finishing processes and is presented in Scutti et al. (2000, 118–122). The component to be machined was the cylindrical piston shaft of a high-pressure intensifier pump. In order to avoid leakage and to improve seal life, a very smooth surface finish was required. Therefore the engineering specification was 0.4 µm, whereas the usual value is 0.8 µm. Since the piston shaft was cylindrical, the machinist was able to engage the shaft in a lathe and to do the requisite grinding and polishing during rotation to achieve

the desired surface finish (rotational polishing). However, tests carried out on the pump showed that seal life performance was unexpectedly low. After inquiry, it was possible to establish that the finishing process was failing. Small circumferential ridges were cut into the hard steel surface from the abrasive polishing grains and the rotational motion of the lathe. Even though the microscopic surface of the shaft met smoothness specifications, at the operating pressures its to-and-fro motion caused it to act like a file on the seal. Therefore, it was decided to change from rotational finishing to axial finishing which results in the microscopic ridges being oriented along the piston's motion axis instead of being along the shaft circumference (see Figure 2.2). After that modification, the operational life of the seal was extended to satisfactory times.

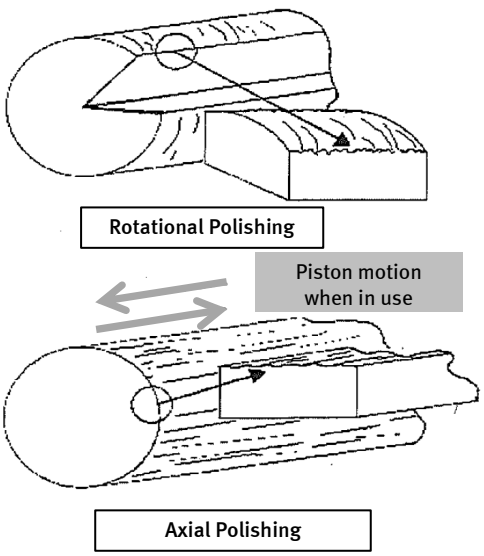

Figure 2.2: Exaggerated appearance of the piston shaft surface polished by different methods. From (Scutti and Aliya: 2000)

2.4. Criteria

This section explains the criteria that are used for establishing the capability of the candidate definitions to capture the failure phenomena included in the failure domain. The criteria are: *accuracy*, *completeness*, *flexibility* and *clarity*. The first two criteria are meant to evaluate the efficacy with which candidate definitions deal with the failure domain that results from the integration of multiple

specializations. The remaining two assess how candidate definitions manage to balance the failure domain's heterogeneity with clarity and ease of communication.

The proposed criteria are meant to be both relevant and independent. Relevance means that the criteria adequately identify the failure domain features as specified in Section 2.3. Relevance means also that they can be met. In Section 2.6, we propose a tentative trans-disciplinary definition and show that it meets all four criteria. Independence means at least that definitions can meet or fail to meet the criteria in any possible combination. A logical demonstration of complete independence of the criteria is not possible, and from the distribution of 'meets' and 'fails to meet' in Table 2.1 it could be concluded that some correlation between completeness and flexibility and between clarity and accuracy exists, but nevertheless the distribution shows the criteria to be independent to a high degree.

The definitions of both Leonards (1982) and Bhaumik (2009) were assessed not to be meeting the criterion of clarity, even to such an extent that it was not possible to properly evaluate them under the accuracy criterion.

The four criteria are motivated in the following subsections.

2.4.1. Accuracy

An *accurate* definition is one that classifies as failures only phenomena that are part of the failure domain. By means of this criterion, we require the definition to prevent false-positive claims, that is, to prevent the identification of non-failure phenomena as failures. Definitions that fall short of meeting this criterion are too broad. Kaminetzky's (1991, 2) definition, for example, adopting terms like "lack of success" and "insufficiency" explicitly includes phenomena that lie outside the technical failure's domain:

> Failure is a human act and is defined as: omission of occurrence or performance; lack of success; non performance; insufficiency; loss of strength; and cessation of proper functioning or performance.

Also the definition provided by Wahl (2006, 27) fails to meet the accuracy criterion:

> Any time that a structural system falls short of expectations or needs in even the smallest way, it may be termed a structural failure.

This definition, resting upon the notion of "even the smallest" deviation, when taken literally is overly permissive and favours misunderstandings. Mainly, during early stages of project development, or during stages like testing or initial manufacture, undue or premature claims of failure could hamper a proactive attitude towards failure prevention. For the same reason, also definitions by Stamatis (1995) – which includes problems, concerns and challenges –, Fortune and Peters (1995) – that very generically embraces everything that "has gone wrong" –, and Mobley (1999) – for which all abnormal system states count as failures –, could be said to be inaccurate.

Leonards (1982) and Bhaumik (2009) are not easily evaluated. Their definitions are not only pivoting upon a rather vague notion of "expected performance", but are framed in a metaphorical way after which failure is a "difference" or a "gap". In this case the lacking of clarity makes it difficult to assess the accuracy.

2.4.2. Completeness

A *complete* definition is one that includes all failure phenomena that are part of the broad IPD failure domain. As expected, most definitions fail under this criterion. The reason is that most of the available definitions have been designed to deal with restricted failure domains or specific failure phenomena. Several definitions explicitly confine themselves to specific aspects of failure, like Collins (1993) – "mechanical failure" –, Melchers (1999) – "structural failure" –, Hattangadi (2000) – "electrical fault" –, Wahl (2006) – "structural failure" –, Rösler's et al. (2007) – "material failure".

Other definitions appear to have a more general scope, however the formulation or terminology adopted make them unsuitable for IPD. Consider, for example, the following definition by Tawancy, Ul-Hamid and Abbas (2004, 11):

> When an engineering product ceases to perform one or more of its functions well before its expected service life, it is said to fail.

From an IPD perspective this definition is incomplete because it is unable to capture two categories of failure events. Firstly, it does not take into account premature degradation or deformation events. In such events, as discussed in Section 2.3.2, components are said to fail even if they are still performing their required function.

Secondly, this definition is limited to products that previously were working and at some point cease to perform. In doing so, it excludes events like the failure of products that did never perform before. Graça et al. (2009), for example, analyses the failure of a steel pressure vessel during hydrotest. The vessel did never perform before. While it was undergoing the mandatory hydrotest before entering service, it failed at a pressure about 20% below the maximum expected during the proof test.

Applying the completeness criterion to the candidate definitions it could be shown that all definitions pivoting on the notion of "required / intended / proper / adequate" function are vulnerable to the first criticism (Hubka and Eder: 1996; Dennies: 2002; Frawley: 2002; Collins: 2003; Jones: 2004; Rausand and Høyland: 2004; Isermann: 2005; Bauer et al.: 2006; IEC 60812: 2006; Birolini: 2007), as well as the "perform as previously specified" expression in MIL-STD-721C (1981).

There are several ways of formulating a definition that escapes this criticism.

- Affonso (2006) specifies that the intended function has to be performed "safely".
- The ASM Handbook definition (Becker and Shipley: 2002) explicitly takes into account deterioration, and connects it to safety concerns.
- Lewis, Reynolds and Gagg (2003) introduce the term "strength criteria" into their definition.
- Stamatis (1995) specifies that the inability to perform includes both "known" and "potential" cases.

Turning to the second criticism – products that never performed before –, all definitions are vulnerable that use terms like "termination" (Jones: 2004; Rausand and Høyland: 2004; IEC 60812: 2006), "stop performing" (Birolini: 2007), "interruption" (Isermann: 2005), "no longer" (Affonso: 2006; Riley et al.: 2006; Daley: 2008), "cessation" (Viswanadham and Singh: 1998; Wulpi: 1999). Frawley (2002) definition is the most blatant in this respect as far as it deals only with "previously acceptable products".

This difficulty could be overcome in different way, for example:

- Lewis, Reynolds and Gagg (2003) use the formulation "fail to meet or continue to maintain".
- MIL-STD-721C (1981) introduces the "previously specified" requirement.
- Leonards (1982) definition is framed in terms of "expected performance".

2.4.3. Flexibility

The flexibility criterion evaluates whether the candidate definitions make use of terminology that pertains to a limited domain or in some other way limits applicability. For example, Riley et al. (2006) state that:

> Failure is defined as the state or condition in which a member or structure no longer functions as intended. (264)

Clearly, their concern is about physical structures. However, as shown in Section 2.3.3, also processes are designed (e.g., manufacturing, assembling, and maintenance). Engineers define requirements and specifications, and detailed procedures are instantiated in order to monitor proper fulfillment. Therefore, it makes sense to say that a process falling short of meeting these requirements is failing. Similarly, several other definitions are specifically targeting failures of physical objects, and for that reason are vulnerable to the same criticism: the term "item" is used in MIL-STD-721C (1981), Jones (2004), IEC 60812 (2006) and Birolini (2007); the terms "part", "component", "device" and "equipment" are used by Collins (1993), Viswanadham and Singh (1998), Wulpi (1999), Hattangadi (2000), Collins (2003), Affonso (2006) and Daley (2008); and references to "products", which are usually assumed to be physical entities, are adopted in Frawley (2002), Lewis, Reynolds and Gagg (2003) and Tawancy, Ul-Hamid and Abbas (2004).

Candidate definitions show different ways around this difficulty:

- The ASM Handbook (Becker and Shipley: 2002) takes explicitly processes into account. It should be noted, however, that this is done by adding a separate section to a definition that is mainly framed on physical components. Moreover, only manufacturing processes are accounted for, while in IPD engineers design and analyze several other kinds of process (e.g., testing, installing, disassembling and recycling).
- Dennies (2002) mentions "processes" explicitly.
- Mobley (1999) and Isermann (2005) decide to use the term "system".
- Bauer et al. (2006) and Stamatis (1995) opt for the term "service".
- The definition may be phrased in terms of expected and observed performance as in Leonards (1982) and Bhaumik (2009).
- Definition's subject may be left unspecified as in Hubka and Eder (1996) and in Rausand and Høyland (2004).

A definition's flexibility could be impaired also by phrasing it in a way that presupposes a particular failure mechanism. Consider Jones (2004):

> Failure. A termination of the ability of an item to perform a required function. A failure is caused by the persistence of a defect. (209)

As discussed above for the case of early degradation of a contactor (Section 2.3.2), components may fail even though they are produced in conformance to specifications and even though they are not having material defects.

Also Becker et al. (2005) make quite a convincing case for the autonomy of the two notions of failure and defect. First, they introduce a distinction between imperfection and defect. A quench crack is an example of an imperfection created during thermal processing. Let us assume that at some point the component fractures during operation. Now, if "the crack that led to fracture did not propagate from the quench crack, the quench crack is an imperfection but not a relevant defect" (18). Next, they present an example of a failure without a defect. In one occasion, they were examining a broken component which had the rather unusual characteristic that a change in radius had been designed as a square corner leading to a dangerous stress concentration. The machine shop attempted to make the part as required by the blueprint and was able to achieve a fillet radius of 0.05 millimetres. From the design specification perspective, the component was free of defects. As summarized by Becker et al. (2005), "Just because a part fails does not imply that it contained a defect, and not all defects are a cause of failure" (16).

Jones's definition does not say that the defect that is persisting is occurring within the failing component. In the case of the corroded contactor, for example, it could be said that the defect was in the design of the chemical process that allowed HCl to condense into the contactor, and in the case presented by Becker et al. (2005) the defect was in the design of a change of radius as a square corner. However, as noted by Lewis, Reynolds and Gagg (2003): "A product may fail because it is part of a system that breaks down under some abnormal condition and unexpectedly places greater demands on one component than the design anticipated" (27-28). Hence, failure and defect are two autonomous notions.

Replacing the term 'defect' with 'fault' will not improve definitions since there is a remarkable lack of consensus on its meaning. According to MIL-STD-721C (1981), a fault is the "immediate cause of failure" (5), while IEC 60812

(2006) defines 'fault' as the state that is "the result of failure" (13). Birolini (2007) follows IEC 60812 in stating that a fault is a result state, but then he adds that a fault "can be a defect or a failure, having thus as possible cause an error (for defects or systematic failures) or a failure mechanism (for failures)" (356). Isermann (2005) agrees that the term 'fault' refers to a state of a component or a system, but interprets it as a synonym of 'defect' and finally states that "[a] failure results from one or more faults" (20).

2.4.4. Clarity

The clarity criterion combines several considerations that evaluate the usability of a definition by a group of engineers having different backgrounds and working together at the same collaborative project. First of all, they need clear terminology that prevents ambiguities. Consider, for example, the definition formulated by Leonards (1982, 108):

> Failure is an unacceptable difference between expected and observed performance.

Terms like "expectations" and "acceptable differences" may be unproblematic in everyday language, but are obscure and impractical when it comes to the multi-disciplinary engineering environment we are envisioning, where different people have different expectations and standards of acceptance.

Other definitions are vulnerable to the same criticism: Bhaumik (2009) – "the gap between the expected performance..." –; Kaminetzky (1991); Fortune and Peters (1995) – "something that has gone wrong" –; Mobley (1999) – "abnormal system state" –, Melchers (1999) – "undesirable structure response".

A definition's usability can also suffer from fragmentation and undue complications. Isermann (2005), for example, uses a twofold definition of failure: one for "permanent interruptions" and one for "intermittent irregularities" in the fulfilment of a system's functions. Daley (2008) similarly employs a twofold definition, but adopts a different criterion of allocation. In his opinion the term 'failure' should be reserved for events that involve "the critical functions for which the system or device is designed", while the term 'malfunction' more generically covers all situations in which "the service or transformation performed by a system or a device no longer meets expectations".

Finally the ASM Handbook's definition (Becker and Shipley: 2002) suffers from the same weakness. The first part of the definition is reserved for character-

izing failure in relation to physical components during operational service. The second part deals with manufacturing processes.

2.5. Definitions' assessment

The final stage towards a trans-disciplinary definition consists in taking the candidate definitions and assessing them against the proposed criteria. The results of the assessment are shown in Table 2.1 (the definitions are quoted in full in the Appendix). The ability of the definitions to meet the criteria was assessed by the authors who assessed the definitions' ability to meet each criterion on the basis of cases representing the main features of the failure domain. Given the conceptual nature of this paper, it has been decided to make use only of two possible scores: if the definition meets the criterion the score is 'yes' ('Y' value in Table 2.1); alternatively, if the definition falls short the score is 'no' ('N' value in Table 2.1). However, in two cases, namely Leonard's (1982) and Bhaumik (2009) definitions, it has been proven difficult to assess the accuracy because both are based upon an unspecified condition of "expected performance". The definitions do not clarify who possess the expectations and their legitimacy. For this reason the accuracy score of these two definitions has not been awarded ('?' in Table 2.1). The tabular presentation is helpful in providing an overview of the whole field. The outcome of the assessment is that none of the candidate definitions satisfies all the criteria.

A simple statistical analysis of the results shown in Table 2.1 reveals that Completeness and Flexibility are the more challenging criteria, 22 definitions (73%) failing to meet the former and 19 definitions (63%) failing to meet the latter. Six definitions came close to meeting all the criteria by failing only one. These definitions are the ones by (Stamatis: 1995; Hubka and Eder: 1996; Dennies: 2002; Lewis et al.: 2003; Rausand and Høyland: 2004; Bauer et al.: 2006). Stamatis (1995) is a complete, flexible and clear definition, but including such terms as problems, concerns and errors, falls short on the accuracy criterion. Hubka and Eder (1996), Dennies (2002), Rausand and Høyland (2004) and Bauer et al. (2006), focusing on lack of proper or adequate functioning, are not able to deal with products that are still performing their function but are close to rupture (for example because of accelerated degradation) and therefore do not meet the completeness criterion. The definition by Lewis, Reynolds and Gagg (2003) satisfies the completeness criterion demanding that the product achieves

not only requested performance but also requested strength. However, since the definition is designed to deal with failure in physical products, it is not flexible enough for IPD, where also systems and processes are at stake, as shown by Stamatis (1995).

These six definitions can be considered a suitable starting point for improvements towards a trans-disciplinary definition. In the next section a tentative definition is proposed which is based on the results of the assessment and which meets all four criteria. The definition is proposed mainly as a validating tool showing the consistency of the criteria and that is possible to have a reasonable definition meeting all of them.

Table 2.1: Four criteria applied on candidate definitions (AC=accuracy; CO=completeness; FL=flexibility; CL=clarity; Y=the definition meets the criterion; N=the definition does not meet the criterion; ?=not assessable). The six definitions satisfying three out of four criteria are in bold typeface.

def	AUTHOR	AC	CO	FL	CL	def	AUTHOR	AC	CO	FL	CL
1	Affonso	Y	N	N	Y	16	Jones	Y	N	N	Y
2	**Bauer et al.**	**Y**	**N**	**Y**	**Y**	17	Kaminetzky	N	Y	Y	N
3	Becker and Shipley	Y	Y	N	N	18	Leonards	?	Y	Y	N
4	Bhaumik	?	Y	Y	N	**19**	**Lewis et al.**	**Y**	**Y**	**N**	**Y**
5	Birolini	Y	N	N	Y	20	Melchers	Y	N	N	N
6	Collins (1993)	Y	N	N	Y	21	MIL-STD-721C	Y	N	N	Y
7	Collins (2003)	Y	N	N	Y	22	Mobley	N	Y	Y	N
8	Daley	Y	N	N	N	**23**	**Rausand and Høyland**	**Y**	**N**	**Y**	**Y**
9	**Dennies**	**Y**	**N**	**Y**	**Y**	24	Riley et al.	Y	N	N	Y
10	Fortune and Peters	N	Y	Y	N	25	Rösler et al.	Y	N	N	Y
11	Frawley	Y	N	N	Y	**26**	**Stamatis**	**N**	**Y**	**Y**	**Y**
12	Hattangadi	Y	N	N	Y	27	Tawancy et al.	Y	N	N	Y
13	**Hubka and Eder**	**Y**	**N**	**Y**	**Y**	28	Viswanadham and Singh	Y	N	N	Y
14	IEC 60812	Y	N	N	Y	29	Wahl	N	N	N	Y
15	Isermann	Y	N	Y	N	30	Wulpi	Y	N	N	Y

2.6. A tentative trans-disciplinary definition of failure

The six shortlisted definitions provide an interesting clue as a starting point. Four of them – namely (Stamatis: 1995; Dennies: 2002; Lewis et al.: 2003; Bauer et al.: 2006) are phrased in terms of 'inability of'. This feature is statistically remarkable since these four are the only ones out of thirty candidate definitions using it. This means that all definitions based on the notion of 'inability' made their way to the shortlist. Following this observation, we set up to investigate the reasons behind this outcome and to apply the lessons learned for designing a tentative trans-disciplinary definition. The result is the following definition:

> Failure is the inability of an engineering process, product, service or system to meet the design team's goals for which it has been developed.

In order to understand the benefits of the 'inability of' locution it should be expanded to the form 'inability of x to do y' showing that it represents a binary relation between two variables. The lessons learned during the assessment process provide valuable insights when it comes to decide what the variables x and y stand for.

The flexibility criterion, as explained in Section 2.4.3, will be met as long as x will include processes and systems besides products intended as material devices.

The completeness criterion (Section 2.4.2) rests mainly on the y variable. The trans-disciplinary definition avoids using conditions like 'to function' or 'to perform', which would result in the definition being unable to deal with cases of product degradation (e.g., the contactor case) where the device is failed even though it is still performing its function. Furthermore, the condition 'to meet' design goals is taking care of the lack-of-previous-operation cases (e.g., prototypes, safety devices) which invalidated definitions using the expression 'termination of'.

Thanks to its relational structure, the 'inability of' phrase allows to satisfy both flexibility and completeness criteria with a compact and streamlined definition. This, together with simple and appropriate terminology, allows the definition to meet also the clarity criterion (Section 2.4.4).

Finally, since the proposed trans-disciplinary definition avoids overly generic terminology and narrows down the notion of failure to the 'inability of' relation, it fulfills the accuracy criterion (Section 2.4.1).

On top of that, the trans-disciplinary definition aims to capture two related but distinct forms of failure. On the one hand, a product fails if it is unable to meet customer needs because these needs have not been properly translated into adequate requirements and specifications by the design team. The product is able to meet the technical specifications but the design falls short of meeting customer needs. To address this issue the trans-disciplinary definition is based on the assumption that developing a product is an 'engineering process' the 'goal' of which is (among others) to create a product that meets customer needs. When customer satisfaction is not achieved the goal is not fulfilled and that, according to the definition, counts as a failure. The same holds for economic considerations that may be included among the goals in case a representative of the marketing department joins the cross-functional design team.

The second form of failure deals with products, services and systems that fail because they do not work properly, that is to say they are not up to the technical specifications developed by the design team. Again, the design team's 'goal' is to create products, services and systems that deliver in conformance to specifications. Whenever a product is unable to perform as specified it falls short of meeting design team's goals and is justifiably considered to fail.

Notice that the trans-disciplinary definition makes explicit reference to the *design team's* goals. The emphasis on the team is introduced in order to prevent outsiders from claiming that this and that are also goals although not shared by the team and the product does not meet them. According to the trans-disciplinary definition, this would not be a valid claim because the only applicable goals are those agreed upon by the entire design team (it is also implicitly assumed that the design team reached agreement based on solid technical reasons and after extensive investigation).

2.7. Conclusions

The IPD approach has proved to be successful in improving product development process performance. One of the primary requirements for IPD is the establishment of effective knowledge sharing among team members taking part in cross-functional teams. This implies, at the very least, the adoption of a uniform nomenclature system, reporting format and terminology. It is important that key notions are clearly defined and their meaning agreed upon, in order to minimize misunderstandings or fruitless discussions. One of these key

notions is the notion of failure. The engineering literature offers numerous definitions, each suited to a different discipline. This multiplicity of meanings is bound to have negative repercussions on the intra-group communication. The aim of this paper is to identify a trans-disciplinary definition of failure that will help cross-functional teams involved in product development to share their knowledge and ideas for improving product development performance and reducing the risk of product failure.

Four criteria are introduced which represent the basic demands a trans-disciplinary definition should meet. These criteria are applied to assess thirty candidate definitions retrieved from the engineering literature and representing several disciplines. The outcome of the assessment is that none of the candidate definitions satisfies all criteria. However, six definitions come close by failing only on one criterion.

Based on the lessons learned from the survey and assessment processes, a tentative trans-disciplinary definition is proposed.

The conceptual analysis performed in this paper has disclosed some of the intricacies behind the notion of failure in engineering and has tried to disentangle part of it. However, the next steps towards a satisfactory trans-disciplinary definition will need the contribution of expert opinion. First, the characterization of the cross-functional failure domain will have to be strengthened through a more extensive survey of design activities performed according to the IPD approach. Second, the tentative definition should be submitted to the judgment of cross-functional design teams, for evaluating its use and for analyzing how it will assist in the communication and knowledge sharing.[2]

[2] I would like to acknowledge Stefano Borgo for valuable comments on an earlier draft of this paper. This work has been developed while taking part in the Marie Curie "EuJoint" Project (IRSES 247503).

Appendix: Failure definitions

def. 1. Failure occurs when the component or equipment no longer can perform its intended function safely. (Affonso: 2006, 3)

def. 2. Failure: The inability of an item, product or service to perform required functions on demand due to one or more defects. (Bauer et al.: 2006, 185)

def. 3. Failure. (1) A general term used to imply that a part in service (a) has become completely inoperable, (b) is still operable but is incapable of satisfactorily performing its intended function, or (c) has deteriorated seriously, to the point that it has become unreliable or unsafe for continued use. (2) Also commonly applied to manufacturing processes that produce components that do not meet specifications. (Becker and Shipley: 2002, 5)

def. 4. Failure can be defined as the gap between the expected performance and the actual performance of any component or assembly. (Bhaumik: 2009, 186)

def. 5. A failure occurs when the item stops performing its required function. (Birolini: 2007, 3)

def. 6. Mechanical failure might be defined as any change in the size, shape, or material properties of a structure, machine, or machine part that renders it incapable of satisfactorily performing its intended function. (Collins: 1993, 6)

def. 7. Improper functioning of a machine or machine part constitutes failure. (Collins: 2003, 22)

def. 8. [A failure occurs when one] of the critical functions for which the system or device is designed is no longer being done. (27)

Malfunction: the situation when the service or transformation performed by a system or a device no longer meets expectations, or when the output no longer meets requirements. (Daley: 2008, 37)

def. 9. Failure is the inability of a component, machine, or process to function properly. (Dennies: 2002, 11)

def. 10. Failure is something that has gone wrong, or not lived up to expectations. (Fortune and Peters: 1995, 20)

def. 11. Failure: An event in which a previously acceptable product does not perform one or more of its required functions within the specified limits under specified conditions. (Frawley: 2002, 33)

def. 12. Failure. The cases in which the electrical equipment fails to perform its normal function as a result of some electrical fault. (Hattangadi: 2000, 21)

def. 13. Failure is defined as the lack of adequate functioning, which may be a result of slow degradation or of a catastrophic event. (Hubka and Eder: 1996, 16)

def. 14. Failure. Termination of the ability of an item to perform a required function. (IEC 60812: 2006, 13)

def. 15. Failure. A failure is a permanent interruption of a system's ability to perform a required function under specified operating conditions.

Malfunction. A malfunction is an intermittent irregularity in the fulfilment of a system's desired function. (Isermann: 2005, 20–21)

def. 16. Failure. A termination of the ability of an item to perform a required function. A failure is caused by the persistence of a defect. (Jones: 2004, 209)

def. 17. Failure is a human act and is defined as: omission of occurrence or performance; lack of success; non performance; insufficiency; loss of strength; and cessation of proper functioning or performance. (Kaminetzky: 1991, 2)

def. 18. Failure is an unacceptable difference between expected and observed performance. (Leonards: 1982, 108)

def. 19. Failure is the inability of a product to meet or continue to maintain the performance or strength criteria in the application of which it was designed. (Lewis et al.: 2003, 27)

def. 20. Structural failure might be considered to be the occurrence of one or more types of undesirable structure response including the violation of predefined limit states. (Melchers: 1999, 51)

def. 21. Failure: the event, or inoperable state, in which any item or part of an item does not, or would not, perform as previously specified. (MIL-STD-721C: 1981, 7)

def. 22. Fault state: an abnormal system state. (Mobley: 1999, 295)

def. 23. Failure is the event when a required function is terminated (exceeding the acceptable limits). (Rausand and Høyland: 2004, 83)

def. 24. Failure is defined as the state or condition in which a member or structure no longer functions as intended. (Riley et al.: 2006, 264)

def. 25. Material failure. Plastic deformation during service is often considered as a failure criterion. One reason for this is that the deformations are usually intolerably large, another is that the yield strength is usually not small enough compared to the tensile strength so that the safety of the component is not guaranteed. A component, however, may also fail by fracture instead of plastic deformation. (Rösler et al.: 2007, 110)

def. 26. Failure. The problem, concern, error, challenge. The inability of the system, design, process, service, or subsystem to perform based on the design intent. … This inability can be defined as both known and potential. (Stamatis: 1995, 74)

def. 27. When an engineering product ceases to perform one or more of its functions well before its expected service life, it is said to fail. (Tawancy et al.: 2004, 11)

def. 28. Failure. Cessation of the ability of a component or system to perform the intended function according to the specification. (Viswanadham and Singh: 1998, 351)

def. 29. Any time that a structural system falls short of expectations or needs in even the smallest way, it may be termed a structural failure. (Wahl: 2006, 27)

def. 30. Failure. Cessation of function or usefulness of a part or assembly. The major types of failure are corrosion, distortion, fracture, and wear. (Wulpi: 1999, 271)

Appendix 2: Additional failure definitions[3]

def. 31. A service failure, often abbreviated here to failure, is an event that occurs when the delivered service deviates from correct service. A service fails either because it does not comply with the functional specification, or because this specification did not adequately describe the system function. A service failure is a transition from correct service to incorrect service, i.e., to not implementing the system function. The period of delivery of incorrect service is a service outage. (Avizienis et al.: 2004, 13)

def. 32. *Failure* is the inability of an item to perform one or more of its required functions.

Total failure, as (say) in a system "single failure point", is the inability of an item to perform *any* required function. (Badenius: 1991, 528)

def. 33. Failure is a process by which a material changes from one state of behavior to another. The more important types of failure are fracture and rupture.

Fracture is the failure process by which new surfaces in the form of cracks are formed in a material or existing crack surfaces are extended. Various stages of fracture may be visualized [1], namely, fracture initiation, fracture propagation (stable and unstable) and strength failure.

Rupture is the failure process by which a structure (e.g. a rock specimen) disintegrates into two or more pieces. (Bieniawski et al.: 1969, 323)

def. 34. When product performance falls below a desired level, the product is deemed to have failed. Failures occur in an uncertain manner and are influenced by factors such as design, manufacture or construction, maintenance, and operation. (Blischke and Murthy: 2000, 36)

def. 35. An event when machinery/equipment is not available to produce parts at specified conditions when scheduled or is not capable of producing parts or perform scheduled operations to specification. For every failure, an action is required. (Blache and Shrivastava: 1994, 69)

def. 36. Failure, in organizations and elsewhere, is deviation from expected and desired results. This includes both avoidable errors and the unavoidable negative outcomes of experiments and risk taking. (Cannon and Edmondson: 2005, 300)

def. 37. Failure: The inability of the robot or the equipment used with the robot to function normally. (Carlson and Murphy: 2005, 423)

def. 38. Failure: when the customer's expectation has not been met and/or the customer is unable to do useful work with the product. (Chillarege: 1996, 354)

def. 39. Failure can be defined as those events that lead to increased maintenance costs or reduced operating revenues (i.e., load restrictions). (Egan: 2006, 177)

---

[3] As mentioned in the Introduction, this second appendix includes engineering definitions of failure collected after the survey paper had been published.

def. 40. Failure is any malfunction or deviation from the norm that significantly detracts from performance. Excessive plastic deformation or shrinkage, wear or loss of attractive appearance may constitute failure just as much as fracture does. (Ezrin: 1996, 6)

def. 41. *Failures* are machine induced interrupts that require skilled technicians for comprehensive troubleshooting and in-depth corrective actions.

Assists may be classified as machine induced interrupts that is [sic] recovered by a machine operator. (Fashandi and Umberg: 2003, 357)

def. 42. Failure is what the structural engineer defines it to be and nothing else. For example, if the stress induced by an earthquake exceeds the yield stress of the material, it could be called failure. Alternatively, if the stress exceeds the ultimate stress of the material, it could be called failure. Failure can also be related to structure serviceability [...]. Therefore, it is fundamentally important to realize that the structural engineer defines failure and that the examples are virtually unlimited. (Hart: 1982, 118)

def. 43. We define failure as a result that was unexpected at the beginning of a project. (Hatamura: 2008, 3)

def. 44. Failure: 1) the act of falling short, being deficient, or lacking; 2) nonattainment or nonsuccess; 3) nonperformance, neglect, omission; 4) bankruptcy; and 5) loss of vigor or strength. (Hohns: 1985, 75)

def. 45. Failure: the inability of a system or component to perform its required functions within specified performance requirements. (IEEE 610.12: 1990, 32)

def. 46. *Failure*: the delivery of a service not complying with the specified service. (Laprie: 1985, 3)

def. 47. Product failure means that the product falls short of a preconceived or a predetermined specification. (Lewis: 2000, 6)

def. 48. Failure: the inability of an item to perform within previously specified limits. (NATO: 2008, 2–5)

def. 49. *Equipment (system, item) failure*: Equipment fails, if it is no longer able to carry out its intended function under the specified operational conditions for which it was designed. (54)

Mission failure: the mission fails if the specific required feasible action (=mission) cannot be carried out or completed as a result of: (a) equipment failure; (b) the inability to maintain the specified operational conditions (environment) which the equipment is designed for. (Nieuwhof: 1984, 56–57)

def. 50. Failure of a component or structure can be defined as an unacceptable gap between its expected and actual performance. It is a condition that makes the structure unable to perform its intended function safely, reliably, and economically. [...].Failures can be broadly classified into two categories: those involving fracture and those without fracture. In each category, failures can be further classified depending on whether they are caused by thermal, mechanical or chemical influences. (Ramachandran et al.: 2005, 3–5)

def. 51. The item is *failed* if it has not done what we want, and is *not failed* if it has done what we want. More exactly, it is the *function* that might be in a failed condition. [...]. Cessation of the performance of the function is failure [...]. Thus, neither hardware nor software, strictly speaking, should be described as failed; but either might be, or might have been, in a condition that can be associated with a functional failure. (Rees: 1997, 163)

def. 52. Product failure occurs when a product no longer performs its intended function in an application environment for the intended life of the product. (Thomas et al.: 2002, 641)

def. 53. Failure is defined as the incapacity of a constructed facility [...] or its components to perform as specified in the design and construction requirements. (Wardhana and Hadipriono: 2003, 152)

def. 54. Functional failure: Unsatisfactory *performance* (e.g., an item delivering unsatisfactory outputs) occurring during a process as operation or testing.

Material failure: An undesired *physical condition* (e.g., an internal part of an item being damaged or broken) which is also permanent (i.e., it will persist until it is repaired). Such a condition could exist during operation or testing - or during a time there is no demand on an item to function at all. [...]. A material failure of an item may or may not cause a subsequent functional failure. Whenever a material failure exists at a time there is no demand for an item to function, the material failure exists with no corresponding functional failure. (Yellman: 1999, 7)

3 Failure of Engineering Artifacts: A Life Cycle approach[4]

Abstract

Failure is a central notion both in ethics of engineering and in engineering practice. Engineers devote considerable resources to assure their products will not fail and considerable progress has been made in the development of tools and methods for understanding and avoiding failure. Engineering ethics, on the other hand, is concerned with the moral and social aspects related to the causes and consequences of technological failures. But what is meant by failure, and what does it mean that a failure has occurred? The subject of this paper is how engineers use and define this notion. Although a traditional definition of failure can be identified that is shared by a large part of the engineering community, the literature shows that engineers are willing to consider as failures also events and circumstance that are at odds with this traditional definition. These cases violate one or more of three assumptions made by the traditional approach to failure. An alternative approach, inspired by the notion of product life cycle, is proposed which dispenses with these assumptions. Besides being able to address the traditional cases of failure, it can deal successfully with the problematic cases. The adoption of a life cycle perspective allows the introduction of a clearer notion of failure and allows a classification of failure phenomena that takes into account the roles of stakeholders involved in the various stages of a product life cycle.

3.1. Introduction

Failure is a central notion both in ethics of engineering and in engineering practice. Apart from the most innocuous events which result in minor annoyances, failure of engineering artifacts raise a host of ethical questions related to allocation of responsibility, foreseeability of risks, prioritization of safety, and so

[4] This chapter has already been published as Del Frate, L. (2013) 'Failure of Engineering Artifacts: A Life Cycle Approach', in: *Science and Engineering Ethics* 19 (3): 913–944.

on. The possibility that failure of engineered artifacts cannot be ruled out with absolute certainty is a source of concern about the introduction of any new technology.

On the engineering side, failures are at the same time *unwanted outcomes* that should be fought with the best resources provided by engineering knowledge, and one of the *main sources* of that same knowledge (Petroski: 1985, 2006). In fact, the investigation of failures has played a crucial role in increasing the reliability and safety of aviation technology (Wood and Sweginnis: 2006). The lessons learned from the crashes of two Comet jets in 1954, for instance, were a milestone in the understanding of metal fatigue and how to improve the design of aircraft mainframes (Wanhill: 2003). Similarly to aviation, other industries and engineering specializations have achieved substantial progress by understanding the reasons behind failure events (Schlager: 1994). In fact, nowadays, the systematic analysis of potential failures has become an integral part of the design process by means of tools like computer simulations (Collins: 1993), Failure Modes and Effects Analysis (Stamatis: 1995), Fault Tree Analysis (Xing and Amari: 2008).

Moreover, the scope of the concept of failure has expanded beyond the intuitive idea of rupture and structural collapse. In aeronautics, for instance, the focus of the investigations has expanded beyond the mere "technical factors", which were prevalent in the early days of aviation, to address safety concerns related to "human factors" and to "organizational factors" (ICAO: 2009, 2–4). Concurrently, engineers put an effort in taking distance from the derogatory aspects of the concept of failure and to devise a more neutral meaning. Investigative agencies like the Dutch Safety Board or the Australian Transport Safety Bureau, for instance, make clear that it is no part of their "remit to try to establish the blame, responsibility or liability attaching to any party" (The Dutch Safety Board: 2009). Instead, their mission is to understand the causes of failures and accidents in order to prevent reoccurrences in the future.

It has to be noted, however, that better knowledge of failure, of the physical as well as of the organizational factors behind it, and increased awareness of the wider implications of failures, has not been mirrored by the creation of a unified conceptual framework shared among engineering disciplines. The urgency of finding effective measures to address failures and the complexity of failure phenomena have led to a situation of conceptual and terminological fragmentation, mostly along disciplinary divisions. Separate disciplines tend to emphasize

specific aspects and to formulate definitions tailored to particular applications. As a result, although failure is a pervasive theme in engineering and despite its importance, many different uses and definitions of the concept of failure can be found in the engineering literature (Prasad et al.: 1996; Tam and Gordon: 2009; Del Frate et al.: 2011).

The tendency towards differentiation has been somewhat balanced by unification attempts pursued by professional organizations and standardization authorities. Thanks to these efforts, a consensus, albeit partial, has coalesced around the terminology published in 1990 by the International Electrotechnical Commission (IEC) and subsequently adopted by a number of international standards. In fact, many engineering textbooks and papers dealing, one way or another, with failure include a quotation of the IEC definition:

> Failure: the termination of the ability of an item to perform a required function.

This definition can be regarded as "the traditional definition of failure" (Blache and Shrivastava: 1994, 69).

Despite its intuitive appeal and the ability to capture correctly a wide range of events, it can be shown that engineers are willing to describe as failures circumstances that do not fit this traditional definition. These problematic cases suggest that some of the basic assumptions behind the traditional approach may be unwarranted. The aim of this paper, then, is to perform a detailed analysis of the traditional approach and to explore the possibility of devising a broader notion which is able to deal with the problematic cases. In doing so, this paper aims to give ethics of engineering a more accurate and present-day understanding of failure in engineering, in support of its analysis of responsibility, liability and risks.

More specifically, it is argued that a notion of failure informed by the notion of product life cycle is a suitable answer to this quest. It is shown that it is compatible with the set of failure events constituting the basis of the traditional approach; furthermore it can account for those events and circumstances that, although in violation of the traditional assumptions, engineers are willing to consider as instances of failure. The adoption of a life cycle approach allows the introduction of a clearer notion of failure and allows a rich classification of failure phenomena that takes into account the roles of various stakeholders involved in the different stages of a product life cycle.

The paper is organized as follows: Section 3.2 starts out clarifying the basic terminology related to the traditional definition. In Section 3.3, the definition is analyzed in terms of four main assumptions, and then Section 3.4 shows how three of these assumptions may be violated in engineering practice. In Section 3.5, it is argued that the traditional approach depends on a specific interpretation of the mission of product development and that this interpretation has been challenged by recent integrated models which imply the notion of product life cycle. Section 3.6 introduces the definition of *product failure* and shows that it is compatible with a life cycle approach. The new approach and its implications are analyzed in Section 3.7. Finally, Section 3.8 summarizes and concludes the paper.

3.2. The *traditional approach* on failure

Several different definitions and characterizations of failure are available in the engineering literature (Prasad et al.: 1996; Tam and Gordon: 2009; Del Frate et al.: 2011). In a paper on failure terminology for plant asset management, Tam and Gordon (2009, 33) notice that the "looseness of terminology and often overlapping shades of meaning lead to ambiguity and confusion". Prasad et al. (1996, 14) raise the same concern from the point of view of dependable computing where the "terminology [...] is used non-uniformly by many authors and standards".

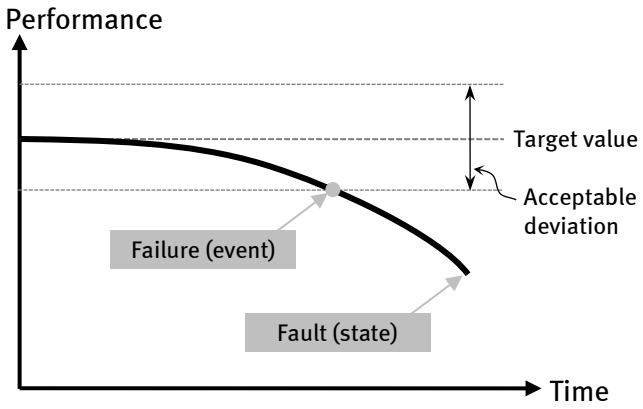

Figure 3.1: Graphical representation of the notions of *failure* and *fault* according to the IEC vocabulary. Redrawn from (Rausand and Øien: 1996)

As a paradigmatic example of this lack of uniformity let us compare two influential sources: Chapter 191 *Dependability and quality of service* of the *International Electrotechnical Vocabulary* (IEC 60050(191): 1990) published by the International Electrotechnical Commission, and the *IEEE Standard Glossary of Software Engineering Terminology* (IEEE 610.12: 1990) published by the Institute of Electrical and Electronics Engineers. The IEC vocabulary distinguishes between *failure* and *fault* which are defined as follows:

> Failure: the termination of the ability of an item to perform a required function.
>
> Fault: the state of an item characterized by inability to perform a required function, excluding the inability during preventive maintenance or other planned actions, or due to lack of external resources.

In a paper on the basic concepts of failure analysis, Rausand and Øien (1996) include a section in which they discuss the IEC terminology and clarify the distinction between failure and fault by means of the diagram that is reproduced in Figure 3.1.

The curve plots the observed level of a performance variable of an item against time. Initially, the observed performance conforms to the target value and then starts to gradually deviate from it. Failure is defined as the event in which the observed performance trespasses the acceptable limits, after which the item is said to be in a fault state. The figure makes clear that the term 'termination' in the IEC definition should not be interpreted as total lack of ability to perform, but as the trespassing of the acceptability levels. Indeed, after a failure event an item can still be capable of performing the required function, albeit at a disappointing level.

Of course, the fact that an item under investigation has failed according to a selected performance parameter does not mean it has failed with respect to all its functions. An item can still be able to perform its main required function although the ability to perform a minor function has dropped below the acceptable limits. These are the situations that the IEC vocabulary describes as "minor failures" and are classified among the "partial failures". In case the affected functions are considered of major importance, then the partial failure is considered of to be a "major" one; finally, at the end of the spectrum are "complete" failures which consist in the "complete inability to perform all required functions".

It is interesting to note that the IEC vocabulary does not make provision for gradations of fault. This means that for a given item and a given function, the notion of fault is binary: either the item performs that specific (minor or major) function within the limits or it does not. Correspondingly, there are no gradations for the "ability to perform" either. So, it does not matter how close an item is to trespassing the threshold of acceptable performance for, as far as performance is within the acceptable limits, the item is described as being in a functioning state. Instead, according to the IEC vocabulary, the notion of "gradual failure" is meant to describe the process leading up to the failure event. If the discrepancy between the observed performance and the target level builds up gradually (e.g., because the item is progressively wearing out), then the ensuing event is termed a "gradual failure"; while a sudden drop beyond the acceptable limits qualifies as a "sudden failure".

The same holds also for items performing multiple functions. Consider an item whose performance parameters are all approaching the acceptable limits but have not trespassed them yet. Then, although the overall performance is degraded, the item is still considered to be in a functioning state. The fault state will take over only when at least one of the performance parameters will trespass the boundary, at which point a partial failure will be said to have occurred. In case the item will not be removed from service and refurbished, more partial failures are likely to follow until the item will reach a state of complete failure. Therefore, in the IEC terminology the term "fault" is attributed a meaning narrower than the one prevailing in common parlance, where it can refer to defects, imperfections, or flaws that deter from the quality or value of a product which, nevertheless, may still be able to perform its required function. On the contrary, the IEC vocabulary stipulates a necessary link between the notions of fault and of lack of ability to perform one or more required functions. The discrepancies between target level and actual performance can be substantial and may affect multiple functionalities simultaneously, yet failure occurs only when at least one performance parameter has exceeded the acceptable limits.

Let us now move to the IEEE glossary (1990, 32) which also makes a distinction between the notions of failure and fault, albeit its definitions are slightly different from the definitions given in the IEC vocabulary. The two definitions read as follows:

> Failure: the inability of a system or component to perform its required functions within specified performance requirements.

> Fault: (1) A defect in a hardware device or component; for example, a short circuit or broken wire. (2) An incorrect step, process, or data definition in a computer program.

Aside from minimal wording differences, it seems that the concept which the IEC vocabulary calls "fault" is called "failure" by the IEEE glossary, while the remaining term of the couple is treated rather differently by the two standards. Adding to the confusion, engineering parlance is demonstrably divergent from the standards. In a note, the IEEE glossary itself acknowledges that the definition of fault "is used primarily by the fault tolerance discipline" but that "in common usage, the terms 'error' and 'bug' are used to express this meaning" (32). Also the IEC admits that its definitions are somewhat at odds with common usage. In the "Terms and definitions" section of the recent standard on *Analysis techniques for system reliability* (IEC 60812: 2006), the definitions of failure and fault are taken literally from the IEC vocabulary but a note is added claiming that "for historical reasons" (13) the two terms will be used interchangeably.

While, as far as the notion of failure is concerned, it is reasonable to see the differences between IEC and IEEE as mostly terminological, some authors have advocated for conceptual revisions. In discussing the notion of 'software failure', Chillarege (1996, 354) argues that a more useful definition would read as follows:

> The customer's expectation has not been met and/or the customer is unable to do useful work with the product.

In the ensuing discussion, Yellman (1999, 7) acknowledges that customer expectations should be preeminent, nonetheless he notices an important limitation in Chillarege's definition: confronted with a product that demonstrably performs below specifications, engineers are willing to call it a failure, "whether or not they [the customers] think (rightly or wrongly) their expectations have been met".

Although a few authors, like Chillarege, have proposed quite substantial reformulations, many authors have been more conservative and their proposals amount to a tailoring of the IEC definition to the needs of a specific domain. Being involved in machinery reliability and plant operations, Affonso (2006, 3) claims:

> Failure occurs when the component or equipment no longer can perform its intended function *safely* (emphasis added).

Bauer et al. (2006, 185), who write from the perspective of quality improvement, choose to define failure as follows:

> Failure: the inability of an item, product or service to perform required functions *on demand due to one or more defects* (emphasis added).

Also interesting is the variant offered by Birolini (2007, 3), which reads as follows:

> A failure occurs when the item *stops performing* its required function (emphasis added).

This characterization is noteworthy for, instead of focusing on the item's "ability to perform" as in the IEC and IEEE definitions, it is centered on the actual observable performance of the item.

The engineering literature on failure seems subject to two conflicting demands. On the one hand, there is a quest for standardization and simplification; on the other, there is the acknowledgment of the multifaceted nature of failure phenomena which appears to resist a unique characterization. As a result, the literature is characterized by a dualism between a central core occupied by the most authoritative and well-established definition and a surrounding area composed of more specialized definitions. The center of the stage is occupied by the IEC definition, which is routinely considered "the traditional definition of failure" (Blache and Shrivastava: 1994, 69). Although this definition can deal successfully with a wide range of failure events, the proliferation of amendments and variations shows that there are circumstances which engineers are willing to describe as failures but are not easily captured by the traditional definition.
In the next section it will be shown that the traditional approach is based on four assumptions. Then, in Section 3.4, it will be argued that, in certain circumstances relevant for engineering practice, these basic assumptions may not hold.

3.3. Four basic assumptions of the traditional approach

The traditional approach to failure can be analyzed in terms of four basic assumptions: *missing functionality*, *utilization context*, *item level*, and *negativity*

assumptions. An event is an instance of failure in this approach when it satisfies all four assumptions.

The *missing functionality assumption* highlights the role played by the notion of an item's function according to the traditional approach. The point is that items may behave in a number of aberrant ways and can possibly deviate from the specifications under several respects, but a failure is said to occur only when the performance of a *required function* is terminated, that is to say, it trespasses the acceptable limits.

Given the variety of function concepts available in the engineering literature (Erden et al.: 2008; Houkes and Vermaas: 2010), this assumption can be subject to multiple interpretations. Unfortunately, the IEC vocabulary does not offer much help in dispelling the ambiguity because the definition of required function it provides is circular. First, the vocabulary defines "required function" by means of the concept of "service" as follows:

> Required function: a function or a combination of functions of an item which is considered necessary to provide a given service.

Then, the concept of "service" is defined by means of the term "function":

> Service: a set of functions offered to a user by an organization.

And, to close the circle, the definition of "function" is missing.

The paper by Rausand and Øien (1996) can shed some light on the meaning of function within the traditional approach. Instead of relying on the IEC vocabulary, they endorse the approach proposed by many design methodologists of articulating item functions by means of verb-noun combinations (e.g., "transmit signal"), which express the relationship between inputs and outputs of flows of energy, materials, and signals (Stone and Wood: 2000; Pahl et al.: 2007). Then, functions can be classified in various categories like *essential*, *auxiliary*, *protective*, and so on. Although Rausand and Øien do not mention explicitly the category of *required function*, they maintain that essential functions are precisely those "required to fulfill the intended purpose of the item" and are defined as "the reasons for installing the item". Moreover, essential functions are reflected by the common names of the items themselves; for instance the essential or required function of a pump is "to pump a fluid". Finally, they require that functions of items can be split into sub-functions and organized into functional hierarchies represented by so-called "functional block diagrams" as described,

among others, in (Pahl et al.: 2007) and recommended by various technical standards as (IEC 60812: 2006) and (MIL-STD-1629A: 1980).

In order to determine whether a failure has occurred, performance parameters have to be defined for all functions as well as target levels and acceptable limits. For an item like a pump, typical performance variables are pressure and flow rate. Hence, according to the traditional approach, a pump whose output pressure gradually dropped below the acceptable limit is said to have suffered a "gradual major failure". In contrast, the gradual degradation of a sub-function that does not impair the essential function is called a "gradual minor failure".

Second is the *utilization context assumption*. Utilization is the stage in the life of a product in which the item actually delivers its required function, and in doing so it fulfills the user needs. Typically, utilization begins once the product is installed and put into operation. A pump, for instance, is installed in a chemical plant with the purpose of pumping a fluid and then, for one reason or another, it stops doing so. This assumption is made more conspicuous and easier to spot in the version of the IEC definition given by Frawley (2002, 33):

> Failure: An event in which a *previously acceptable* product does not perform one or more of its required functions within the specified limits under specified conditions (emphasis added).

Frawley's definition implies that previously, at the beginning of its useful life, the product was performing at an acceptable level, after that its performance dropped and trespassed the acceptable limits.

Although the utilization stage is, of course, the natural context in which items are expected to perform their required functions, utilization episodes may occur in another context as well, namely during testing. Tests can be considered a vicarious form of utilization of the system they represent. Certifications tests, like those conducted on aircraft engines and mainframes, are extremely severe and aim "to test a complete technological system under conditions as close as possible to actual field conditions" (Sims: 1999, 492). Smaller scale and less demanding tests are routinely run after completion of the manufacturing processes or after maintenance operations. Even though the users, whose needs are supposed to be fulfilled by the performance of the item's functions, are not actually present, their role is played by instruments and qualified personnel acting as vicarious users during testing. Therefore, as far as failure is concerned,

tests can be seen as parts of a broadly construed *utilization context* which groups together utilization stage and tests.

Third is the *item level assumption*. When looking at the traditional approach as represented by the entire set of definitions in the IEC vocabulary, it can be seen that it rests on the idea of engineering products as physical items whose behavior can be observed and measured quantitatively. As shown above, failure judgments are based on the results of a comparison between the observed physical variables of a specific item (e.g., fluid pressure or flow rate of a specific pump) and the target values and acceptable levels defined by the specifications. In fact, as users we are surrounded and are dependent on the ability of many physical items to satisfy our needs by actually performing their required functions: cars move us around, clocks tell the time, printers print paper, and so on. Many of our judgments refer to the properties and behaviors of these items with which we deal on a daily basis.

Besides these judgments based on the properties of products at the level of the physical item, it is common for engineers, as well as for lay people, to make judgments also on properties at the type level. Reliability, for instance, is a property that can be predicated both of a specific item and of a type of product. The point is that the two kinds of judgments are different and the evidence that can be used in support of a judgment at the item level may be inadequate or insufficient for a type level judgment. So, at the type level a certain product, say the Volkswagen Golf, can be judged to be more reliable than a different but comparable type, say the Ford Focus; yet it is possible to claim, without contradiction, that a specific Ford Focus item has been found to be more reliable than a specific Volkswagen Golf.

In chemical engineering, criteria have been devised for assessing the safety characteristics of processing plants at the type level. For this kind of judgments, chemical engineers refer to the notion of inherently safer design (Kletz: 1998; American Institute of Chemical Engineers: 2009). Again, the judgment that a design is inherently safer is meant to describe a property, or a group of properties, that are not fully expressed at the level of the individual item. Instead, it is a judgment about features that are shared among all items realized following the same architecture.

Unless the notions of required function and observable performance are stretched beyond the characterization given by Rausand and Øien, the traditional approach is particularly suited for failure judgments at the item level and a

broader notion of failure might be need for judgments at the type level (see Section 3.4).

Finally, there is a *negativity assumption*, that is to say, failure events are unwanted and should be avoided. The point of this assumption is that failures should be avoided *per se*, irrespectively of their consequences. Luckily only a minority of failures result in serious consequences and the vast majority cause only minor annoyances. Nonetheless, as remarked by Yellman's (1999) criticism of Chillarege (1996) definition, engineers are willing to consider an underperforming product a failure whether or not the customers complain about it.

3.4. Beyond the *traditional approach*

The view on failure embodied by the IEC definition and by the other definitions based on it has achieved a prominent status because of its undeniable agreement with a large class of events that are important in engineering practice. A great deal of engineering activity – during design, testing, manufacturing, maintenance, and so on – aims at preventing products in the field from terminating the performance of required functions. However, as it is shown by the existence of alternative definitions and by other evidence in the literature, engineers are prone to consider as failures also events that do not square with this interpretation; that is to say, events that do not satisfy the four basic assumptions.

The only assumption that most engineers will not consider challenging is *negativity*.[5] They may disagree on whether an event or a class of events constitute a failure or not, but they will maintain that, if something is a failure, then it should have been avoided. In the rest of this section it will be shown how the first three assumptions may not hold in some circumstances that are relevant in engineering practice.

---

[5] There are few exceptions to this general attitude. One exception is to maintain that some failures, like near misses, are not *totally* negative because they provide valuable learning opportunities that can be exploited for preventing more serious consequences in the future. Another exception is the approach of material scientists, some of whom maintain that failures are but physical events and processes that have measurable effects on the structure of material samples, e.g., fracture of a material occurs when the tensile stress overcomes the load bearing capacity (Dasgupta and Pecht: 1991). From this perspective, failures are not more negative than oxidative processes, fires, are for chemists.

3.4.1. Missing functionality assumption

In Section 3.3, it has been shown that, according to the traditional approach, the notion of failure is binary: given a specific item and a specific (minor or major) function, either the item is functioning or it is not, depending on whether the observed performance lies within or without the acceptable limits. For simple items which are required to perform just one function there is nothing in between functioning as required and complete failure. Once the observed performance trespasses the acceptable limits, the items are said to be in a state of complete failure. In contrast, items performing multiple required functions can suffer partial failures, which can be minor or major depending on the importance of the affected function. Anyway, for a given required function (minor or major), the difference between functioning state and fault state is straightforward and it is based only on the observed performance being within or without the acceptable limits.

However, engineers appear to make use also of a more nuanced notion of failure, namely a notion that takes into account the rate at which the observed performance is approaching the acceptable limits and the residual ability of the item to perform. Lewis et al. (2003, 27), for instance, propose the following definition:

> Failure is the inability of a product to meet or continue to maintain the performance or strength criteria in the application for which it was designed.

Therefore, an item whose observed performance is still within the acceptable limits could be classified as being in a fault state and removed from service because its performance is degrading at a rate much faster than predicted; that is to say, it is unable "to continue to maintain" the desired performance. If the same situation were to be assessed in accordance to the traditional approach, the verdict would be the opposite, for the judgment would be based on the fact that the item is still operating within the specifications.

An example of these different outcomes can be found in a case story discussed by Gagg and Lewis (2007) dealing with the failure of a swing bridge. The bridge moved over steel tracks laid in concrete, with the whole bridge riding on three castors, each one made of a steel shaft rotating within two bearing bushes manufactured from Oilon, a self-lubricating polymeric material. The bridge was in operation for six months and apparently was working as expected. Indeed, it

was only during the first *routine* inspection that evidence of wear was found on the shafts of all three castors.

Despite the fact that, until the inspection, the bridge was performing as expected, Gagg and Lewis describe the three castors as the "failed axle bearing combination" (1633). In fact, the castors were treated as failed components by removing them from service and by undertaking a detailed failure investigation. Eventually, it was determined that Oilon was not a suitable material for manufacturing the bearings and that the solution "to this 'failure' was a straightforward substitution of Oilon by Nylube, a grade of material far more capable of withstanding service loading experienced by this swing bridge castor" (1634).

Items affected by accelerated wear are the typical subjects of failure judgments made in violation of the missing functionality assumption. Another common mechanism is corrosion, like in (Suess: 1992). Although several examples can be found in the engineering literature, they are less frequent than traditional failure judgments based on clear-cut termination of a required function because the latter are more conspicuous and potentially more harmful.

3.4.2. Utilization context assumption

Although failures during utilization are usually the most worrying for engineers, they are aware that failure is lurking even before the products enter utilization and start delivering their required functions. Sudhakar and Paredes (2005) have investigated a case of "premature failure during manufacturing" (35) affecting bimetal bearings for automotive application. In a manufacturing plant, a number of bearings were found to be cracked after the sintering step during which a layer of copper alloy was soldered to the steel backing. The analysis concluded that the bearing failures were due to improper setting of the heating parameters for the sintering process. It is clear that the traditional "termination of a required function" was not the criterion adopted by Sudhakar and Paredes when they claimed that the bearings suffered a "premature failure during manufacturing" and set out to identify the "failure mechanisms".

3.4.3. Item level assumption

According to the traditional approach, the notion of failure pivots on the idea of engineering products seen as physical items and on their properties and behav-

iors that can be observed and measured. However, engineers are used also to think in terms of properties of types of products, properties that do not exist at the level of the individual item. Consider, for instance, the notion of *field return* as it is used in the automotive industry. A field return is a car component that is returned to the manufacturer after a service technician has diagnosed a failure which is covered by the manufacturer warranty. In this sense, a field return is a physical item that has failed to perform its required function. Engineers, however, generalize this property over the entire type of a product, in which case it becomes the *rate of field returns* and, as such, it does not belong to any item in particular but to the entire type.

As its counterpart at the item level, also the type level property can be used in engineering judgments over failure. In the early 1980's, Ford introduced a new type of ignition module, the so called Thick Film Ignition module, with the purpose of surpassing the reliability of Japanese cars which were invading the US market. This goal was translated into a requirement to the effect that the rate of field returns must not exceed 1.6 returns out of 100 modules installed (Pecht: 2006; Qi et al.: 2008). After a few months, the ignition modules turned out to be a substantial failure with field returns so far above the expectations that Ford decided to issue a recall. The individual failures were due to a variety of reasons, including manufacturing defects, assembly problems, and inappropriate maintenance. However, the main reason for the spate of field returns was the underestimation of the temperatures prevailing in the operational environment, that is to say, the engine compartment where the modules were installed. Unable to withstand the thermal loads, the modules behaved erratically resulting in what "may well be the most widespread intermittent failure condition ever reported" (Qi et al.: 2008, 664). Sure enough, many individual modules did operate successfully: those were the modules installed on cars operated in colder regions. The point is, however, that, as a type, the Thick Film Ignition module was unable to achieve the expected rate of field returns.

The Ford case deals with a type level property of a component, the ignition module, which is part of a larger system, the car. Engineers, however, express judgments also on the basis of type level properties of entire systems. Marks (1989) offers an analysis of the Sinclair C5, a battery-assisted tricycle which was conceived by the British inventor and entrepreneur Sir Clive Sinclair and was introduced to the market in January 1985. The sales figures were so poor that production was stopped as early as September 1985. After reviewing the main

steps in the development, launch, and failure of the C5, Marks is led to wonder: "but was the C5, in fact, a marketing failure or a technology failure?" (68). In Marks' opinion, both the product and the marketing were poor. As for the "technology", Marks notices that the product can be criticized for a series of shortcomings relating to aspects like safety (e.g., low protection in case of collisions) and usability (e.g., lack of a reverse gear). It is clear, then, that Marks' judgment does not refer to any C5 in particular for failing to perform its required functions; instead the criticisms point to crucial properties of the C5 as a type. The manufacturing quality and reliability of the physical items might have been excellent and many customers might have enjoyed riding it; yet, as a type, it was an uncontested failure.

The traditional view is not well equipped for this kind of failure judgments that are based on type level properties emerging from the interaction of multiple aspects of a product design and life cycle. Such properties like inherent safety or robustness or sustainability, especially in case of complex products, cannot be easily reduced to a specific function or small group of functions expressed as input and outputs in the functional hierarchy of a product. Still, type level characteristics are crucial for the realization of successful products and play an important role in the delineation of product requirements.

Admittedly, the distinction between functional and non-functional requirements is hotly debated in the requirements engineering community (Hull et al.: 2010). In fact, the definition of requirement given in the IEEE 1220 (2005) standard carefully avoids introducing the notion of a non-functional requirement claiming, instead, that a requirement is "a statement that identifies a product or process operational, functional, or design characteristic or constraint" (9). In turn, the notion of a design characteristic covers a wide spectrum of "performance, usability, safety, maintainability and a host of other qualities" (Hull et al.: 2010, 7).

It is worth stressing that the argument developed in this paper in support of a broader notion of failure does not depend on the availability of a shared and clear-cut distinction between functional and non-functional requirements. What is needed is the acknowledgement that the traditional notion of failure is based, among others, on the assumption that functions of engineering products should be expressed by means of input-output relations of flows of energy, materials, and signals that can be organized in functional block diagrams. Many design methodologists, e.g., (Stone and Wood: 2000; Pahl et al.: 2007), have convinc-

ingly discussed the various benefits allowed by this interpretation in support of the engineering design task, as well as in comparing, communicating and archiving design solutions. Nevertheless, nowhere is it implied that, in order to achieve these benefits, this interpretation has to cover all relevant aspects of engineering products. On the contrary, it can be expected that an overly generic notion of function would undermine the benefits that have just been mentioned.

Thus, the main purpose of this section has been to show that products can have characteristics which engineers themselves deem relevant for failure judgments and that cannot be captured easily in terms of functions and functional hierarchies. Since engineers routinely make these kinds of judgments and the traditional approach appears to struggle with them, a more efficient way of dealing with this conceptual conflict is to devise a broader notion of failure. The new notion of failure will be introduced in Section 3.6, which will also explore how the new notion connects to the life cycle perspective of product development. Before that, the next section will further investigate the conceptual background of the traditional approach to failure and its connection with the sequential model of product development.

3.5. From one customer to many stakeholders

The fact that there are problematic cases does not imply that the traditional approach is unwarranted, for it does provide an adequate representation of many kinds of failure events that engineers recognize as threats to the success of their products. Moreover, it sits well with widespread intuitions of end-users. The question then, is why the traditional approach has left out the problematic cases from the mainstream of failure events.

The analysis performed in the previous sections suggests that this exclusion depends on the fact that the traditional approach is based on a specific interpretation of the mission of engineering product development. Roughly stated, this mission can be described as follows: to develop products that provide optimal and reliable performance of required functions for the customer. Then, the traditional approach assumes that, as far as the "required functions" are concerned, the customer and the end-user are considered to be the same; that is to say, the individual or organization whose needs are satisfied by the functional performance of the product.

It has to be stressed, though, that the identification of end-user and customer may not hold for product characteristics that are not considered strictly functional. This point can be illustrated by the following historical example. In the early 1920's, during the first stages of the development of household refrigerators, General Electric engineers had to decide how to solve the problem of cooling the compressor, that is to say, the component that circulates the refrigerating fluid (Cowan: 1985). They came up with two alternatives: water cooling or air cooling. Both solutions were thought to be equally capable of contributing to the achievement of the required function of a refrigerator, that is, to keep food items fresh. The crucial difference between the two solutions was energy consumption. According to the calculations "the electric power bill of the air cooled machine would be about $ 1.30 more in six months than the water cooled machine" (209). Since electric utility companies were General Electric's most important customers, the air cooling mechanism was selected. Of course, in those days there were no standards or energy saving regulations to deter General Electric from its decision. Nevertheless, it is reasonable to assume that also within more tightly regulated markets, all else being equal a company will prioritize those non-functional characteristics that are more beneficial to its customers.

Although this interpretation of the mission statement is still extensively shared among engineers and managers alike, other interpretations (that will be discussed later in this section) have been proposed that take into account the needs (both functional and non-functional) of multiple stakeholders involved in the life cycle of a product, thus prompting more liberal views on failure.

3.5.1. The sequential model of product development

It is worth stressing that the emphasis on the end-user is far from being idiosyncratic and disconnected from practice; on the contrary, it can be considered part of the conceptual background of the so-called "sequential model" of product development (Kahn: 2005; Yang: 2007).

In this model, the overall design task is broken down into sub-tasks that are carried out by individual departments in a predefined sequence. The details may vary, but generally speaking the process starts out with the identification of a series of needs that the product is expected to address; the needs are translated into requirements, and from the requirements a design concept emerges that is progressively made more precise until it is approved for production. The up-

stream decisions taken early on in the process have the effect of freezing the main features of the final architecture and, at the same time, of dictating the boundaries or the framework within which downstream departments will have to carry out their tasks and advance the process towards completion.

The framework plays also the role of a guideline for the assessment of what kinds of product behavior count as a success or as a failure. Therefore, since the framework relies on the list of functional requirements resulting from the analysis of the user needs, it prompts for the endorsement of a notion of failure that assumes the performance of required functions as the focal criterion.

A persistent worry in the analysis of user needs is the anticipation of potential behaviors that might result in product misuse. Although, practically, products cannot be designed to survive all conceivable kinds of abuse, to a certain extent engineers can control and reduce the probability of misuse causing a product to "fail" in a dangerous manner. Here the term "fail" is between quotation marks because, as noted by Ezrin (1996), products that break apart "when someone makes [...] improper use of them, that should not be considered failure in the usual sense" (6). Still, engineers are typically aware of the duty to minimize the probability of harm also for "failures" due to misuse, and may take it into account during the design process in the form of implicit user needs.

The sequential model has been the paradigm for product design until the 1970s, when industry started switching towards alternative models based on the promotion of cross-disciplinary integration which proved to be instrumental in the commercial success of Japanese manufacturers, particularly in the automotive sector. Nowadays, the sequential model can still be effective for the development of mature products where companies have already extensively explored many kinds of different frameworks in a number of product generations (Liker et al.: 1996). Hence, downstream departments are equipped with a large repertoire of solutions that can fit with almost every possible framework.

One of the main shortcomings of the sequential model is that costs of major changes can increase exponentially as the product development proceeds to the later phases. Major changes are those demanding a revision of the framework. If, towards the end of the process, e.g., close to the manufacturing stage, a subtask cannot be completed and the design is sent back to the drawing board, then all the downstream decisions have to be checked again to assure they are compatible with the revised framework. Even in the ideal case that most of the downstream decisions can be safely retained, it is likely that delays will ensue.

Taking problems related to reliability as an illustration of the extent of these risks, Levin and Kalal (2003) have estimated that "the cost to fix a reliability problem increases an order of magnitude in each subsequent phase" (159).

Research on design methods has also pointed out that alternative models, besides curbing the cost of late stage design revisions, may allow for a host of other improvements, like more robust designs, shorter development times, and exploration of a broader range of design solutions (Henderson and Clark, 1990; Clark, 1991; Iansiti, 1995; Hoopes and Postrel, 1999).

For these reasons, ever more engineering companies are switching towards integrated approaches to product development inspired by the example of Japanese car manufacturers like Toyota (Womack et al.: 2007). Evidence for this transition has been provided by several studies (McDonough: 2000). Griffin (1997) presents the results of a five years research effort showing how new product development relies increasingly on multi-disciplinary or trans-disciplinary teams. Helper and Sako (1995) focus on the change in the customer-supplier relations and find that "where once contracts were short-term, arm's-length relationships, now contracts have increasingly become long term" (77) providing better sharing of information and cooperation.

3.5.2. Integrated models

Several alternative models have been offered, like Concurrent Engineering (Syan and Menon: 1994), Simultaneous Engineering (Kortge and Okonkwo: 1989), Integrated Product Development (Andreasen and Hein: 1987), and the literature on the subject is flourishing. One aspect these models have in common is to look at the product and the processes taking place during the product's life cycle as an integrated system and to emphasize the interdependencies both across components and across processes. Since, in one form or another, the notion of product life cycle plays a crucial role, the various integrated models on offer can be seen as interpretations of a general life cycle perspective to product realization. In the words of the recent ISO/IEC standard *Systems and Software Engineering* (ISO/IEC 15288: 2008), the aim of a life cycle perspective is to provide a "common framework to improve communication and cooperation among the parties that create, utilize and manage modern systems in order that they can work in an integrated, coherent fashion" (vi).

At the foundations of the life cycle perspective is the observation that every system has a life cycle which can be represented by means of a *life cycle model* constituted by a sequence of stages, from concept development to retirement. The number and kind of stages employed vary depending on the nature of the product to be realized and the structure of the organization. A life cycle model is not a mere chronological representation of the typical development of single items of the product, although such representation can be easily derived from the model. Instead, it is a "decision-linked conceptual segmentation" (ISO/IEC TR 24748-1: 2010, 12) used to represent and manage technical and business decisions and actions during the life of a product as a whole. The stages in a life cycle model perform two main functions: they group together homogeneous technical and managerial activities; and also provide a systematic view of the requirements that the product must be able to achieve in order to be approved for the following stages.

The key difference between the sequential model and the integrated models can be appreciated already from the early stages in the product development process. In the former, designers are in charge of translating the end-user needs into requirements and, in doing so, establishing the framework for the activities downstream. In the latter, a trans-disciplinary team is invested with the responsibility of integrating the view of the various subjects having stakes in the different stages of the life cycle. Hence, in a trans-disciplinary team, the aim of the designers to define the functions that the product is requested to perform during utilization meets with the aim of the manufacturing engineers to optimize the processes on the shop floor, and with the needs of component suppliers, and with the demands of the maintenance and support department, and so on.

Instead of being dominated by a list of requirements connected to the functions (allegedly) required by the end-user, an integrated product development process is led by a trans-disciplinary set of goals whose aim is to strike the balance between the many stakeholders involved in the whole life cycle. As product development unfolds the set of goals is increasingly refined and made more precise by exploring and expanding it into requirements and finally into technical specifications. However, in order to assure that the goals are preserved, this process is always supervised by the trans-disciplinary team.

From a life cycle perspective, the mission of engineering product development is broader than the one envisioned by the sequential model because there

is not just one customer, the end-user, to whom an optimally and reliably performing product should be provided, but there are multiple stakeholders having different – and sometimes conflicting – interests in one or more of the stages in the life of a product. Both the notion of successful product and that of failed product are affected by the life cycle approach. In fact, to be fully successful a product must satisfy end-user needs by performing its required functions during utilization. However, a product that performs well in the field but falls short of meeting the demands of other stakeholders, e.g., the manufacturing department, may result in scarce profitability and, eventually, in failure. What characterizes the life cycle approach to failure is precisely the awareness that a life cycle is an integrated whole in which the needs of many stakeholders have to be balanced in order to avoid failure.

Now, if the traditional definition of failure sits well with the sequential model, then a suitable definition should be identified that suits the needs of the life cycle approach. In the next section, a few alternative definitions will be examined and a new definition of failure will be introduced and it will be shown that, besides dealing successfully with the failure events typically addressed by the traditional approach, the new definition can also accommodate the problematic cases.

3.6. A new definition of failure

A number of attempts have been made to extend the reach of the traditional definition, but they have been only partially successful. One avenue is to link the notion of failure to the product specifications as done, for instance, by the US Military Standard (MIL-STD-721C: 1981, 7) *Definition of Terms for Reliability and Maintainability* where the following definition is given:

> Failure: the event, or inoperable state, in which any item or part of an item does not, or would not, perform as previously specified.

The problem with such kind of definitions is that product specifications may turn out to be inadequate thus leading to products that fail although they do comply with the specs. Groot Boerle (2002) describes the case of an electric wheelchair that went out of control, drove off a subway platform and badly injured the driver. The investigation established that the wheelchair was activated by a low-voltage electromagnetic field at a frequency of 1.89 GHz, which is

one of the frequencies used by digital telephone networks. During the ensuing trial, the manufacturer declined responsibility claiming that the chair met the product specifications, according to which the electrical system had to be insulated from electromagnetic fields up to 1 GHz. This line of reasoning was rejected by the court on the grounds that the specifications for a wheelchair supposedly fit for utilization in public spaces should take into account the possibility of interferences with telephone networks.

A second alternative for broadening the reach of the traditional approach consists in emphasizing the role of product requirements instead of that of product specifications. Also this fix, however, runs into the same kind of problems because it falls short of capturing situations in which product requirements themselves are misguided, as exemplified by the Thick Film Ignition module developed by Ford (see Section 3.4.3). The requirement asked for modules able to withstand continuous exposure to temperatures up to 125 °C (Pecht: 2006), which was assumed to be the hottest temperature in the area of the engine compartment where the modules were to be installed. The requirement turned out to be unrealistically low, leading to the spate of field returns that prompted the recall.

Since the amendments discussed above and other proposals available in the engineering literature have been unable to reshape the traditional definition such that it can deal with the life cycle approach, a new definition of failure is offered in this paper.

The purpose of the new definition is to capture how engineers currently use the notion of failure. The previous sections have shown that this use is broader than the range covered by the traditional approach. Moreover, the expanding range of the notion of failure appears to be correlated with the larger set of issues that are dealt with in life cycle approaches to product realization. Therefore, the new definition has been conceived by having in mind the kind of failure judgments and concerns about failure that may emerge within trans-disciplinary teams involved in product development. At the same time, the new definition has to be backward compatible and retain the ability to address the failure judgments that are the domain of the traditional approach. In order to underline the difference from the traditional definition, the new notion has been baptized *product failure* and it is defined as follows:

> Product failure is the inability of an engineering process, product, service or system to meet the design team's goals for which it has been developed.

Visibly, the structure of the new definition is remarkably similar to the traditional definition and retains the basic intuition that failure is a form of inability on the part of the product. The novelty is due to the fact that, instead of referring to the notion of required function, the intuition is anchored to the frame provided by the life cycle approach to product development by creating a connection with "goals of the design team".

While the idea of product life cycle has been a main source of inspiration for the new definition, the notion of life cycle does not appear explicitly in the text of the definition. The reason is that, although a life cycle approach to product development and, consequently, the adoption of trans-disciplinary teams is becoming increasingly popular, it is unlikely to become the exclusive approach to product design. As noted above (Section 3.5.1), the sequential model is still effective in design of mature products. Also, not all design groups are necessarily trans-disciplinary. Therefore, to make the notion of product failure as general as possible it has been decided to phrase it in terms of "design teams". When it is applied within the context of the sequential model, the new definition converges with the traditional approach where the goal of design teams is to develop products that satisfy end-user needs by performing required functions. When it is applied from a life cycle perspective, design teams become trans-disciplinary and their goals result from the integration of needs and concerns of multiple stakeholders.

It is worth noticing that trans-disciplinary teams may gather together members coming from different companies, like suppliers and contractors, in addition to representatives of the company owning the product. Hence, the set of goals on which the design team will settle can be seen also as an integration of the goals of the various companies or organizations taking part in the effort. So, the locution "design team's goals" is more general than "company's goals".

Yet, even though a life cycle approach to product realization broadens the kinds of concerns that are taken into account during the design process, the design team's goals cannot be expected to converge with the needs and priorities of society at large. Even in a life cycle approach, product development and realization remains a technical activity pursued predominantly for economical purposes. Certainly, economic viability presupposes compliance with the law and the relevant regulations, which are the most prominent expression of the needs and priorities of a society. But when engineers look at a product to assess whether it is successful or not, they think it is possible to tell apart its technical

merits and other non-technical properties. The adoption of a life cycle perspective has expanded the domain of the technical merits beyond the mere ability to perform required functions, but has not changed the fact that product development is a technical endeavor for the benefit of a well-defined set of stakeholders. Indeed, a conflict of intuitions may ensue, as it is illustrated in the following case history. As mentioned above, the development of new cars in the US automotive industry follows the integrated model and relies heavily on trans-disciplinary design teams. Nonetheless, one of the main recent innovations in the US market was the introduction of Sport Utility Vehicles (SUV) which, although extremely profitable for the industry, has raised many serious social and moral issues. Large and luxurious SUV generate more profits than any other type of car. In the early 1990's, Ford was making less than $1000 in profits on the average sedan, while the profit on large SUV models like the Explorer was nearly $8,000 (Bradsher: 2000). On the other hand, SUVs are demonstrably more polluting and more dangerous than average cars; especially in multivehicle collisions, their mass and the stiffness of their framework cause a substantial increase of the likelihood of severe injuries and casualties (Bradsher: 2002; Latin and Kasolas: 2002). These social costs, however, never played a prominent role within the definition of the design goals: as recently as 1997, Ford's Director of Vehicle Systems Engineering conceded that "crash compatibility was not an active part of design" (Bradsher: 1997). According to the definition given in this paper, SUVs like the Ford Explorer cannot be considered a "product failure", even though there might be compelling reasons to consider them examples of dangerous kinds of products.

It cannot be overstressed that satisfaction of the end-user needs through performance of the required functions is a crucial element of the set of goals of any design team. Thus, the new definition is in agreement with the traditional one in that termination of the ability to perform a required function constitutes a failure event. Moreover, it inherits the view that for failures to occur it is not sufficient that performance variables deviate from target values, but acceptable limits should be trespassed. Nonetheless, the new definition has a broader reach. Being centered on the *goals* of the design team, the concept of product failure can deal with the cases of wrong specifications and wrong requirements discussed above. Considering again the example of the electric wheel chair, it can be seen that it was indeed a case of product failure because the main goal behind the realization of an electric wheelchair is to make a product that can be safely

used for transportation on public roads. Vulnerability to interferences from telecommunication networks makes the product unsuitable for the intended application. Similarly, products may be produced which are in compliance with the requirements, but the requirements themselves may be an inadequate expression of the design goals (e.g., Ford's ignition modules). These situations are also covered by the notion of product failure.

The most important aspect of the new definition is how it deals with the traditional approach and the stance it takes towards its assumptions, i.e., missing functionality, item level, utilization stage, and negativity assumptions. It can be shown that, besides the negativity assumption, on which they are in agreement, the two approaches take different stances.

First of all, the notion of product failure does not maintain the missing functionality assumption; therefore a product that is still performing its required function but is falling short to achieve the goal for which it was developed, is classified as a failure. Thus, the notion of product failure can capture cases of failure like the swing bridge investigated by Gagg and Lewis (2007) (Section 3.4.1). Even though the bridge was still performing its required function, it was unable to meet the goals of robustness and prolonged operational life.

Secondly, the notion of product failure can accommodate cases of failure that do not comply with the utilization context assumption (Section 3.4.2). The bimetal bearings that developed cracks during manufacturing can be considered an instance of product failure because one of the design goals is to assure that products can be successfully and reliably manufactured.

Finally, since a trans-disciplinary development team may set goals for the product as a type and is not confined to item level goals, the notion of product failure does not rely on the item level assumption. Therefore, the inability of the Ford ignition modules to achieve the expected reliability goals (expressed in terms of field return rate), which proved to be problematic for the traditional approach (Section 3.4.3), is correctly classified as a type level failure.

Before the analysis moves to further aspects of the life cycle approach, it may be worthwhile summing up the main steps in the argument developed so far. It has been argued that, although the concept of failure is variously defined in engineering, a traditional approach can be identified which can deal with a large proportion of failure events that matter to engineers. However, a number of problematic cases have been shown in which the assumptions of the traditional approach are at odds with engineering intuitions and ways of speech. The

explanation of the inability to account for these violations has been identified in the relation between the traditional approach and the sequential model of product development, from which the traditional approach has inherited a bias in favor of the end-user perspective. This model has been contrasted with integrated models of product development that are based on a life cycle perspective and emphasize the role of a multiplicity of stakeholders besides the end-user. Then a new definition of failure has been proposed based on this life cycle perspective: product failure. Finally, it has been shown that, besides addressing the failure events covered by the traditional approach, the new notion can deal also with the problematic cases.

In may be worth pausing briefly to consider whether, by looking for the advantages of a broader notion of failure, this paper has ended up propounding a definition that is overly stretched. More specifically, a concern may be raised that the definition of product failure is dangerously close to becoming a synonym of design flaw, that is to say, any reason for design iteration. Although there are similarities between the two notions it can be shown that they are still clearly distinguishable. First, product failures can be addressed without recurring to design alterations. When Apple announced the new iPhone 4 back in June 2010, it claimed that the phone would have been available in two colors, black and white. However, even though the black version arrived as anticipated, the white version was delayed without further explanation, and only at the end of July Apple confessed that the white phones turned out to be "more challenging to manufacture" than expected (Apple Inc.: 2010). The company carefully avoided to specify the nature of the manufacturing challenges. Technology experts suggested that the delay was due to problems faced by the supplier of the white glass panels that constitute the front and back covers of the handset (Lai: 2010). Although the supplier was able to manufacture the prototypes according to the tight opacity and thickness specifications, when ramping up to full scale production the manufacturing yield fell dramatically with merely three panels fabricated successfully per hour (My Digital Life: 2010). The problem was not solved until, in January 2011, Apple found a new supplier able to achieve an acceptable yield while meeting the specifications. Eventually, the much awaited white models reached the shops at the end of April 2011. There was no design iteration, because the design as well as the product specifications were retained and the new supplier proved that they could be achieved. Nevertheless, the experts maintain the product was a failure because many prospective customers

postponed their purchase and decided to buy the next version of iPhone that was launched in October 2011 (Sherr: 2011).

Second, design iterations may occur that are not due to product failures. Changes in regulations (e.g., environmental protection laws) or in the financial situation may force a team to alter an otherwise sound design.

Moreover, the definition of product failure appears to be well suited for being adopted in the execution of Failure Modes and Effect Analysis (FMEA), a design procedure for the systematic identification of potential failures that is widely used in industry and recommended by international quality standards and best practices such as (ISO 9001: 2008) and (SAE J1739: 2002). In his well-known manual on FMEA, Stamatis (1995) emphasizes that the purpose of the procedure is not merely to identify and prevent failures that might occur once the product enters the utilization stage; instead, it should also be implemented to address failures located elsewhere in the life cycle. For this reason, Stamatis distinguishes four types of FMEAs as follows: Systems FMEA is based on conceptual design and focuses on potential failure modes between the functions of the system; Design FMEA is used to analyze products before they are released to manufacturing and focuses on failure modes caused by design deficiencies; Process FMEA aims at minimizing process failures during manufacturing and assembly; finally, Service FMEA is used to analyze services before they reach the customer. Even though the four types of FMEA are based on different sources and address different issues, Stamatis (1995, 74) proposes a unified definition of failure, which is the following:

> Failure: the problem, concern, error, challenge. The inability of the system, design, process, service, or subsystem to perform based on the design intent.

Since he is writing a manual for practitioners, it is not surprising that Stamatis decides to keep the conceptual analysis to the bare essentials by simply stating his definition without expanding into the motivations behind his choice. However, two crucial clues are given that highlight the similarity between his notion of failure and the one discussed in this paper. The first clue is found in the brief comment that follows the definition explaining that the design intent "usually comes from an analysis and an evaluation of the needs, wants, or expectations of the customer" (74). The second clue is a remark from the beginning of the book stressing that the customer should not be interpreted just as the end-user, instead it should be "viewed as the subsequent or downstream operation as well

as a service operation" (xxiii). Jointly, these two clauses allow for Stamatis' notion of failure to span over the same range of failure judgments that led to the notion of product failure analyzed in this paper.

3.7. The life cycle approach in action

Given the discussion above, it will come as no surprise that the term *life cycle* does not appear among the terms defined by the IEC vocabulary (IEC 60050(191): 1990). What can be found, instead, is an implicit reference to the concept of life cycle in relation to a small group of definitions dealing with the causes of failure. *Failure cause* is defined as "the circumstances during design, manufacture or use which have led to a failure". The analysis by IEC vocabulary stops here: there are no further considerations on the life cycle stages and the relations between them. And here is where the life cycle approach takes over. The first step consists in making explicit the life cycle model assumed by analysis, as it is done in Figure 3.2. Then, the next step consists in representing the three circumstances "which have led to a failure" as *failure trajectories* which connect the stage where the causal factor lies with the stage where the failure event is located.

Since Figure 3.2 has been built using the materials available within the traditional approach, only three failure trajectories can be represented, all of them point to the utilization stage, and the life cycle model stops with the utilization stage. But these limitations do not apply to the life cycle approach.

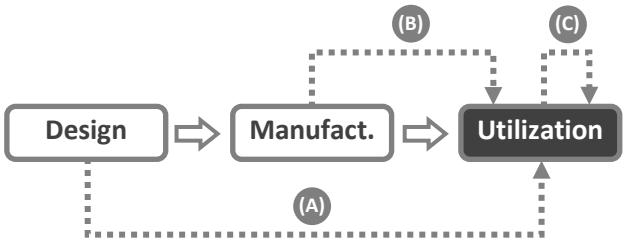

Figure 3.2: Admissible failure trajectories according to the traditional approach: arrow (A): design failure; arrow (B): manufacturing failure; arrow (C): utilization failure

Since the needs of other stakeholders besides the end-user are taken into account (e.g., those of subjects having stakes in stages like installation, shipment, maintenance, recycling, and so on), more sophisticated life cycle models can be introduced which extend beyond the utilization stage (see below).

Moreover, the withdrawal of the utilization context assumption implies that more failure trajectories can occur that are not bound to point to the utilization stage. It can now be seen that the life cycle approach can easily accommodate the case of the bimetal bearing failing during manufacturing discussed in Section 3.4.2. Although the failure did not occur during the utilization stage nor involved a termination of a required function, according to the life cycle approach it constitutes an instance of product failure because those individual items that cracked during the sintering process were unable to meet the quality characteristics which were part of the design goal. Since the failures were due to a causal factor associated to the manufacturing stage itself (i.e., inadequate setting of sintering process parameters), the corresponding failure trajectory can be represented by an arrow like (E) in Figure 3.3, which starts out from the manufacturing stage and then points to the same stage, where the failure is located.
Figure 3.3 displays also a second trajectory, arrow (D), which connects the design stage with the manufacturing stage. As the abundant literature on Design for Manufacturing and Design for Assembly shows, not all design solutions are equally easy or convenient to manufacture, and some may be unfeasible altogether. Although designs are checked for features that may hamper manufacturability, it can happen that a design is approved for production even though it contains a flaw resulting in the inability of the product to run through the entire process or, in case it can make it, in the inability to meet the acceptability requirements. Both outcomes are to be considered product failures.

Failures during manufacturing can be no less consequential than during utilization. This is especially the case in civil engineering. A notable historical case is the failure during construction of the Quebec Bridge (Quebec City, Canada) in 1907. The investigating commission established that several design shortcomings were responsible for the collapse, among which were the "use of relatively slender struts with an inefficient layout of material, calculations using methods derived from much smaller sections, inadequate lacing, [and] overstress due to inaccurate dead loads" (Collings: 2008, 25). As a result, the bridge was structurally unable to withstand its own weight.

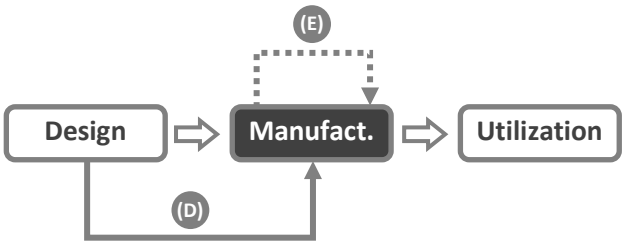

Figure 3.3: Failure trajectories of product failures located at the manufacturing stage

The introduction of more powerful engineering models and computational tools has greatly reduced the occurrence of similar failures, though not completely. In June 1970 a section of the Cleddau Bridge (Neyland, UK) collapsed during construction and the Committee of Inquiry determined that inadequate design of a pier support was a crucial factor (Merrison: 1973; Collings: 2008).
Analogous mishaps, albeit with less dramatic consequences, may occur with any kind of engineering artifact. Consider the manufacturing of tempered glass: tempering is a thermal or chemical treatment that confers to glass an increased strength compared to normal glass and it is used in the realization of products like windows, doors, tables, and building facades. The improvement of mechanical properties is due to the rapid cooling which places the surfaces of the glass in a state of high compression, and the central core in a state of compensating tension (Pfaender: 1996). As a result, tempered glass is about four times stronger than normal glass of the same thickness. Moreover, when in it breaks, tempered glass fractures into small fragments of nearly regular shape that reduce the probability of serious injury as compared to normal glass.

The downside is that, if the internal stresses between glass surface and central core are not carefully balanced, they can build up during the tempering process and shatter the glass to pieces. Therefore, the designer must be aware of a series of geometrical constraints that are dictated by the need to balance internal stresses. If holes are needed in the final product (for instance, to allow the installation of metal hinges on a glass door) their minimal diameter should be not less than the glass thickness, and the distance between the holes should be at least four times the glass thickness (Le Bourhis: 2008). Lack of compliance with these constraints will result in a product bound to fail during manufacturing because of inadequate design, arrow (D).

Figure 3.3 can also be used for presenting a third aspect that tells the life cycle approach apart from the traditional approach, which is the possibility of distinguishing from failure trajectories at the type level and at the item level. In Figure 3.3, the dashed arrow (E) represents an item level failure trajectory; the solid arrow (D) represents a type level trajectory. At first, it may be tempting to differentiate item level failures and type level failures looking at the number of items involved: just one in the former, all the items in the product population in the latter. Although this may be correct in most circumstances, it is misleading. According to this criterion, the collapse of the Quebec Bridge was just an instance of an item level failure that happened to be also a type level failure because the entire product population was constituted exactly by that precise item. However, the results of the Commission of Inquiry made it clear that there were serious problems with the bridge design such that any further attempt to rebuild it would end in a new collapse unless adequate changes were made. Precisely this link with design solutions is the aspect that characterizes type level failures. They are due to design flaws and the proper way they can be corrected or prevented is by design improvements. Certainly, changes to the manufacturing process can sometimes rescue a bad design, but those are usually mere temporary fixes and, if not properly thought through, are likely to result in other unforeseen problems in later stages. Also relevant for type level failures is the fact that, routinely, designers set goals dealing with type level properties of products. An example is the reliability goal that Ford engineers set for the ignition modules (Section 3.4.3), namely that field returns should not be in excess of 1.6 out of 100 components threshold. Clearly, this is a reliability property that cannot apply at the level of the individual item, and deals with properties pertaining to the type as a whole. Another example of type level goal is manufacturing yield, which expresses the ratio of acceptable items within the overall manufacturing output.

In contrast, item level failures are due to process variables (e.g., manufacturing, utilization, maintenance variables) that typically affect only a limited number of items, and can be prevented by means of a correct adjustment of those variables. In the case of the bimetal bearings (Section 3.4.2), for instance, the modification of the sintering timing and temperatures was successful in eradicating the problem. Since item level and type level failures are remarkably different in terms of consequences and of corrective actions, it useful to distinguish them graphically when drawing the failure trajectories in a life cycle

model. Since design flaws result in type level failures, the arrows starting out from the design stage are represented by solid arrows. The failure trajectories originating from the subsequent stages, which result in item level failures, are represented by dashed arrows.

The same line of reasoning that has been applied in the analysis of the potential failure trajectories represented in Figure 3.3 can be extended to more detailed and complete life cycle models; that is to say, models that represent sub-stages within the main stages (e.g., milling and welding within manufacturing), or models that include stages which occur after utilization. In fact, the life cycle approach introduces the possibility of analyzing failures that can occur also after a product has completed its useful life, a possibility that is not contemplated by the traditional approach. In particular, the need of minimizing the environmental impact of products, a need that is backed up by increasingly stringent regulations, implies that the design intent is informed by the requirements of stakeholders whose interests lie in the environmental performance of products during the retirement stage.

What are, then, the potential failure trajectories that point to the retirement stage of a product? Let us consider an electronic appliance for which the designers aim at minimizing the environmental impact during retirement. The product life cycle can be represented, as in Figure 3.4, by a model constituted of five stages: design, manufacture, utilization, maintenance, and retirement.

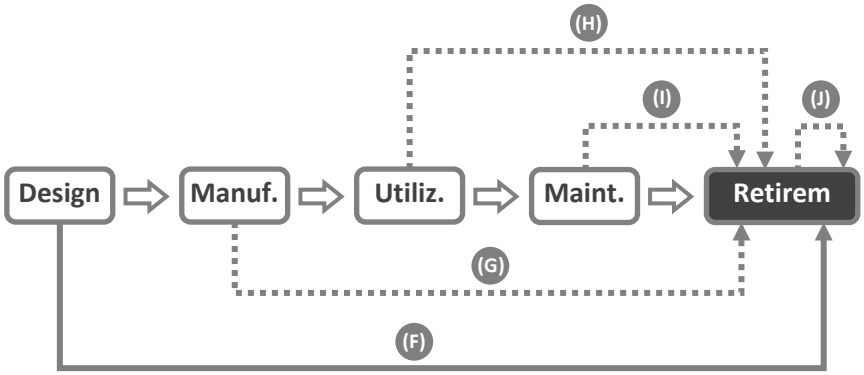

Figure 3.4: Product failure trajectories related to the retirement stage of an electronic appliance

The goal of environmental optimization evolves into a series of requirements for the retirement stage that set the acceptable limits for a range of relevant parameters like, toxic emissions, recycling of materials, reutilization of components and so on. The ability to achieve these requirements depends on decisions and actions taken during each of the stages within the life cycle. A way in which design can increase the recyclability performance is by reducing the number of small plastic components within the appliance. In a study on design for recycling of computer enclosures, Masanet and Horvath (2007, 1807) have shown that "PC enclosure components with a mass of 25 g or less would be discarded (a common practice for small plastic components)". The discarded components detract from the recycled fraction and can cause the product to miss the established goal, thus leading to a product failure during retirement due to design related factors. This failure trajectory appears as arrow (F) in Figure 3.4. Differently from the other trajectories, this failure occurs at the type level; hence it is represented by a solid arrow.

Manufacturing variables can determine item-level product failures during retirement, arrow (G). Typically, recycling of the printed circuit boards installed within electronic appliances is done through converters in which the plastic part is burned and the metals are recovered. One potential environmental risk is posed by the presence of toxic additives that are sometimes utilized in the manufacturing of the plastic part. By careful selection of materials and control of the manufacturing processes, the amount of toxic additives can be minimized. It may happen, however, that because of a mistake during manufacturing a number of plastic parts are produced containing a large amount of toxic additives that will be released during recycling, thus leading to a violation of the environmental requirements.

Arrow (H) represents product failures due to circumstances related to the utilization stage. A typical example is the failure to achieve predetermined recycling goals because users, instead of returning the products to the appropriate service centers, simply dispose of them. As documented by Behrendt et al. (1997) this is more probable for small appliances that are normally disposed of with household waste. Maintenance procedures, arrow (I), can alter a product in such a way that it becomes unsuitable for disassembly, e.g., bolted joints are replaced with welded or glued joints. Finally, arrow (J), variables related to the retirement stage itself, like incorrect temperature settings of the converter, can result in failure to achieve the design goals.

Although extremely simplified, the example above suggest the ability of the life cycle approach to identify and methodically represent a wide range of circumstances that may cause a product to miss the established design goals. Sure, the idea of product failure is intuitively more easily associated with stages like utilization or manufacturing than with the retirement stage. However, restricting failure to certain stages would undermine the life cycle approach according to which failure can be found anywhere in the life cycle.

So far, only direct failure trajectories have been examined, that is to say, trajectories originating in one stage (e.g., design) and then pointing directly where the failure is located (e.g., utilization). Engineers, however, routinely have to confront with more complicated scenarios in which multiple factors may be involved and intermediate stages also play a role. As an example, let us consider a failure case analyzed by Barella et al. (2011). A large batch of canned tuna in olive oil (1 million cans) failed during the utilization stage, approximately 6-8 months after production, when customers started complaining because the product appeared to be contaminated. The sources of contamination were oxidation products developing near the welding area of the can. Typically, tuna cans are produced by welding sheets of tin coated steel along one edge. As an additional protective measure, after welding a polymer coating, i.e., lacquer, is applied onto the can's internal surface.

As for the failed tuna cans, Barella et al. found superficial welding irregularities such that adhesion of the polymer coating was compromised. Usually, that would not constitute a problem because the oil provides a good protective environment against oxidation. Unfortunately, in this batch of cans oil with high water content (double than normal) was used, thus making the environment more corrosive and eventually leading to the contamination problems that emerged a few months after production.

To summarize, two conditions occurred: inadequate welding and oil with high water content. Each condition alone would be unable to cause the failure, but their co-occurrence determined the adverse outcome. This failure scenario can be represented by a life cycle model like the one in Figure 3.5 where the production of the tuna cans has been split into two stages: manufacturing, in which the steel sheets are transformed into lacquered cans; and assembly, in which tuna and oil are poured into the cans.

The failure trajectory starts out with the inadequate welding during manufacturing, arrow (K). Since, for unspecified reasons, the quality checks were

ineffectual, the defective cans progressed to the assembly stage, arrow (L), where oil with unusually high water content was used to fill them, arrow (M). The joint effect of (L) and (M) is represented by the final step of the failure trajectory leading to product failure during the utilization stage (N).

In Figure 3.5, the utilization stage is represented by a shaded rectangle to underline the fact that it is where the product failure is situated. Sure enough, the tuna cans failed to achieve the goal of delivering edible food to the end-users. However, the engineering and business team in charge of the project may ponder that still other goals have been missed such that the mistakes and shortcomings during manufacturing and assembly might be considered failures in their own right. Especially the problem of the defectively welded cans is likely to be taken an instance of failure during manufacturing, which might be addressed by means of corrective measures like better procedures, improved training, and so on.

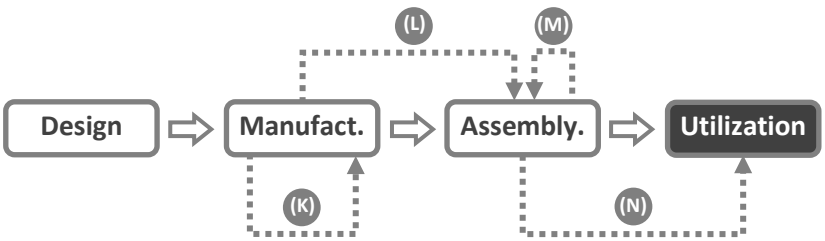

Figure 3.5: Life cycle model of the failure of tuna cans described in (Barella et al.: 2011)

3.8. Conclusion

In this paper, the traditional notion of failure has been identified as the termination of the ability of an item to perform a required function. It has been shown, however, that engineers use the notion of failure in a broader sense. Examples from the engineering literature have been given that violate three of the four assumptions on which the traditional approach is based. These assumptions, and the emphasis given to the functions required by end-users, have been analyzed as due to the traditional approach adopting the sequential model of product development as part of its conceptual background. In this model, the end-user requirements are established early on in the development process and

determine the framework within which the needs of other stakeholders are addressed in the following steps.

Actual engineering practice is progressively moving away from this sequential model and switching to integrated models in which the notion of product life cycle plays a prominent role. Inspired by the life cycle approach, a new definition for failure has been proposed, *product failure*, and it has been shown that, besides being able to deal with the traditional cases of failure, it can deal with the problematic cases as well. In conjunction with the new definition, the notion of *failure trajectory* has been introduced and it has been argued that the life cycle approach enables the introduction of new trajectories in addition to the three envisioned by the traditional approach.

In these alternative models, product development is led by a transdisciplinary team that conveys the views of the multiple stakeholders involved in the different stages of the life cycle. As a consequence, a life cycle approach to failure abandons the priority granted to the end-user and switches to a broader view that accounts for the needs of a multiplicity of stakeholders. The new approach can naturally account for failures occurring during the manufacturing stage when a product has not yet started to perform its required functions.

Especially interesting is the possibility of analyzing failure trajectories that point to stages in the life cycle occurring after utilization. The performance of products during these stages is becoming increasingly important because of the impact on sustainability. The retirement stage of a product can be almost as complex as the manufacturing stage and, from the point of view of many stakeholders, no less relevant. Also with respect to retirement a product can be a success or a failure. The life cycle approach makes the notion of failure ready for the sustainability challenges of the 21st century.[6]

---

[6] I would like to thank Pieter E. Vermaas for all the support and the insightful comments that made the writing of this paper possible, and the anonymous reviewers of *Science and Engineering Ethics* for the valuable feedback. A previous version of this paper has been presented during the International Conference on Engineering Design, ICED 2011, in Copenhagen. This work has been developed while taking part in the Marie Curie "EuJoint" Project (IRSES 247503).

4 Preliminaries to a Formal Ontology of Failure of Engineering Artifacts[7]

Abstract

The aim of this paper is to offer a conceptual analysis of the notion of failure of engineering artifacts focusing on aspects that are of import for a possible ontological formalization. Failure is a central notion in engineering, yet different taxonomies exist in the various industries and engineering domains that are not mutually compatible thereby hindering knowledge exchange. A formal definition of failure would contribute to improve knowledge exchange. However, in order to be successful such formalization should rest on shared conceptualizations. The paper analyses how the notion of failure is used in engineering, starting with the so-called "traditional definition". Then, it is shown that engineers are willing to consider as failures also events and circumstances that are at odds with this traditional definition. Therefore, it is argued that, in order to capture adequately engineering conceptualizations, three independent notions of failure should be distinguished, which are called *function-based failure*, *specification-based failure*, and *material-based failure*.

4.1. Introduction

Failure is a vital concern to engineers of all disciplines. Understanding how failures happen is crucial for prevention and also for mitigation of potential outcomes. For these reasons, tools that allow effective archiving, reuse, and exchange of data about failures are valuable to engineers. Formal ontologies have been already deployed successfully for knowledge exchange in various domains. Attempts have been made to extend formal ontologies in order to characterize

---

[7] This chapter has already been published as Del Frate, L. (2012) 'Preliminaries to a formal ontology of failure of engineering artifacts', in: Donnelly, M. and Guizzardi, G. (eds.), *Formal Ontology in Information Systems: Proceedings of the Seventh International Conference (FOIS 2012)*, IOS Press, Amsterdam: 117–130.

the notion of failure in engineering: Kitamura and Mizoguchi (1999) provide an ontological analysis of fault processes and categories of fault; van der Vegte et al. (2002), propose an ontology-based modeling of product functionality which addresses also the aspect of unintended behavior and malfunction; Koji et al. (2005) investigate the feasibility of applying ontology-based transformations to a functional model in order to create FMEA sheets; Borgo and Leitão (2007) discuss the foundations of a core ontology for manufacturing, including the concepts of disturbance and machine failure; Borgo and Vieu (2009) offer an analysis of the category of artifacts in formal ontology and outline a definition of malfunctioning artifact.

However, the analysis of conceptualizations about failure shared among engineers has played a minor role in the ontological literature so far. Indeed, as observed by Guarino et al. (2009), formal specifications of concepts do not need to be specifications of *shared* concepts. Nonetheless, Guarino et al. promptly remark that an "ontology may turn out useless if it is used in a way that runs counter to the shared ontological commitment" (14) of its stakeholders. In making this claim, they are endorsing the approach proposed by Borst (1997, 123) who argues that formal ontologies should bring out "what is really shared by the community [of users] in order to enhance reuse *within* this community" (emphasis in the original).

Therefore, an ideal starting point for a formalization of the concept of failure would be a definition which is widely shared in the engineering community and which is consistent with actual use. Unfortunately, the engineering terminology on failure and related concepts is highly fragmented and there is a lack of agreement even on the definition of failure itself. Separate disciplines tend to emphasize specific aspects of the notion of failure and to formulate definitions tailored to particular applications. As a result, conflicting definitions can be found in the engineering literature (Prasad et al.: 1996; Tam and Gordon: 2009; Del Frate et al.: 2011).

Therefore, circumstances may arise where engineering judgments about failure diverge. A paradigmatic case is failure of artifacts that have been abused, e.g., because of overloading or by exposure to environmental conditions harsher than specified. Harland and Lorenz (2005), for example, do not see any problem in classifying as failed a component which stops performing its required function because the surrounding environment has become hotter than specified. Other engineers, however, disagree and think that such events should not be

considered failures or, at least, not failures "in the usual sense" (Ezrin: 1996, 6). Disagreements may ensue also between engineers who would treat an artifact as being in a fault state because of degradation of its material properties, and those who think that a failure judgment would be unwarranted if the artifact is still functioning. Suess (1992), for instance, describes the case of a stainless steel trailer barrel used to haul various chemicals which internal surface showed evidence of severe corrosion. Even though the barrel did not develop any leakage, Suess treats the episode as a clear-cut failure, more precisely a "*failure* [which] was caused by bacteria-induced corrosion" (73, emphasis added). On the other hand, Grantham Lough et al. (2008, 473) discuss Suess' case and, by pointing out that the barrel was still able to perform its main function of storing fluids, they conclude that "the tank was still functioning properly".

The aim of this paper, then, is to perform a conceptual analysis of the notion of failure in engineering as preliminary work towards a formal definition which is informed by practitioners' intuitions and ontological commitments. The paper builds on the results of a previous survey of the engineering literature (Del Frate et al.: 2011) where it is argued that the engineering community is subject to two conflicting demands. On the one hand, there is a quest for standardization and simplification; on the other, there is an acknowledgment of the multifaceted nature of failure phenomena which stimulates the development of definitions tailored on special purposes and needs. In fact, the tendency towards unification has coalesced into the definition offered by the *International Electrotechnical Vocabulary* (IEC 60050(191): 1990) published by the International Electrotechnical Commission (IEC), where failure is defined as "the termination of the ability of an item to perform a required function". Being adopted by several international standards and influential textbooks, the IEC definition has achieved a prominent role and is often taken as "the traditional definition" (Blache and Shrivastava: 1994, 69). Nevertheless, the IEC notion has not been fully successful in superseding alternative definitions and preventing new ones being proposed. With all its merits, it has proven unable to capture relevant engineering intuitions. On the one hand, it can be shown that engineers are willing to classify as failures circumstances that do not fit the traditional definition. In this paper a proposal is made to the effect that, in order to capture engineering intuitions and to deal with the problematic cases, two additional notions should be introduced besides the traditional one. Thus, three different

notions of failure should be distinguished: *function-based failure* (i.e., the IEC notion), *specification-based failure*, and *material-based failure*.

4.2. The traditional definition: Function-based failure

The IEC vocabulary (IEC 60050(191): 1990) defines the term "failure" as follows:

> Failure: the termination of the ability of an item to perform a required function.
> NOTE 1 – After failure the item has a fault.
> NOTE 2 – Failure is an event, as distinguished from fault, which is a state.

The two notes appended to the definition make clear that, in order to understand the notion of failure, a second term should be defined as well: "fault". The IEC vocabulary definition of fault reads as follows:

> Fault: the state of an item characterized by inability to perform a required function, excluding the inability during preventive maintenance or other planned actions, or due to lack of external resources.
> NOTE 1 – A fault is often the result of a failure of the item itself, but may exist without prior failure.

Jointly, the two definitions characterize what can be considered a twofold notion of failure which in the rest of this paper will be called *Function-based failure* (FBF), in order to differentiate from the other two notions that will be discussed later on.

The motivation behind the distinction between failure events and fault states is that for engineers it is important to know when and how many times an item stopped delivering a required function (i.e., failure event), and also for how long the lack of performance persisted (i.e., fault state). The relation between the two notions has been illustrated by Rausand and Øien (1996) by means of the diagram reproduced in Figure 4.1. The curve in Figure 4.1 plots the observed level of a performance variable of an item (e.g., a pump) against time. Initially, the observed performance conforms to the target value, but later it starts gradually deviating downwards until it trespasses the acceptable limit, after which the item is still producing some output though well below the target level. The term "failure", as defined by the IEC vocabulary, refers to the instant when the observed performance trespasses the acceptable limits (or tolerance limits), after which the item is said to be in a fault state that will persist until the item is

repaired. Sure, Rausand and Øien's figure is not meant for a philosophical audience and is far from perfect.[8] Nevertheless, it has the merit of pointing out, albeit in a crude manner, some aspects relevant for the present discussion. First of all, the diagram makes clear that the item enters into a fault state immediately after exceeding the acceptable limit when it is still able to deliver some performance, even if at a disappointing level. Hence, it can be seen that the term "termination" in the IEC definition should not be interpreted as total lack of ability to perform, but as the trespassing of the acceptable levels.

Moreover, the definition of fault state does not imply that fault states are necessary permanent. In fact, fault states can be temporary, in which case the IEC terminology introduces the term "transient faults", that is to say, faults "which persists for a limited time duration following which the item recovers the ability to perform a required function without being subjected to any action of corrective maintenance". Think, for instance, of a computer that hangs because of moderate overheating; after a while the electronic components will cool down and the computer will resume operating normally.

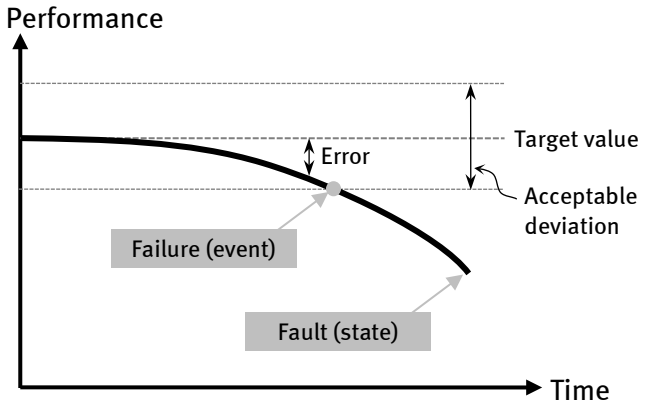

Figure 4.1: Illustration of the notions of failure event and fault state. Based on (Rausand and Øien: 1996) with minor modifications

[8] For instance the diagram plots actual performance over time even though the IEC definition refers to the item's ability to perform. For the aims of the present discussion the distinction can be safely ignored.

By analogy with the notion of "fault state" the interval preceding the failure event can be termed "functioning state", even though the term does not appear in the IEC nomenclature. According to the IEC nomenclature, the event depicted in Figure 4.1 is called a "gradual failure" for the failure event is preceded by the building up of a gap or "error" between the observed performance and the target level. Failure events where the performance departs abruptly from the target level to trespass the acceptable limits are termed "sudden failures". In these cases the gentle slope of Figure 4.1 would be replaced by a sharp turn either downwards or upwards. Another important distinction is the one between "complete failure" and "complete fault" on one hand, and "partial failure" and "partial fault" on the other. These notions are meant for items required to perform multiple functions. Therefore, a failure that results in the inability to perform some, but not all, required functions is called "partial failure"; while a failure affecting all required functions is called "complete failure".

It is interesting to note that the IEC terminology does not make provision for *gradations of fault*. This means that for a given item and a given function, the notion of fault is binary: either the item is able to perform the required function or it is not. Correspondingly, there are no gradations for the "ability to perform" either. Thus, it does not matter how close an item is to trespassing the threshold of acceptable performance for, to the extent that performance is within the acceptable limits, the item is described as being in a functioning state. Finally, the note appended to the definition of fault state (i.e., "A fault is often the result of a failure of the item itself, but may exist without prior failure") addresses a further important aspect of the relation between fault state and failure event. The phrasing of the note, though, is somewhat misleading because of the term "result" might mistakenly suggest a causal connection between failure and fault. Sure enough, for many items failure occurs after a period of satisfactory performance. Nevertheless, the failure event is just the event corresponding to the transition between the functioning state (when performance is between the acceptable limits) and the fault state. Thus, the relation between functioning, failure, and fault is one of temporal sequence and not a causal one. In fact, by saying that a fault state "may exist without prior failure", the second part of the note makes clear that causality is not required. Simply, it may happen that an item never possessed the ability to perform its required function, possibly because of a design flaw or a manufacturing defect.

Evidently, in order to understand the gist of the IEC notion of failure, the meaning of "required function" should be clarified. This is a notoriously problematic notion which is given different definitions in the engineering literature, e.g., (Erden et al.: 2008; Vermaas: 2009), and unfortunately the IEC vocabulary cannot be said to provide much clues on the issue. The vocabulary defines "required function" by means of the concept of service: "Required function: a function or a combination of functions of an item which is considered necessary to provide a given service". Then, the concept of service is defined by means of the term "function" itself: "Service: a set of functions offered to a user by an organization". However, since the definition of bare "function" is missing, one must conclude that the notion of "required function" is left undefined by the IEC terminology.

A more perspicuous discussion of the notion of function and its relation to the IEC terminology can be found in Rausand and Øien (1996), from which the diagram in Figure 4.1 has been taken. Bypassing the IEC definitions of required function, Rausand and Øien elect to endorse the approach proposed by many design methodologists of treating item functions as black boxes which perform operations, expressed by means of verb-noun combinations (e.g., "transmit signal"), on the flows of energy, materials, and signals passing through them (Stone and Wood: 2000; Pahl et al.: 2007). Rausand and Øien illustrate this approach by considering a "process shutdown gate valve" – a kind of safety valve often used in chemical plants – whose required function is to "close flow of fluid", typically in case of an emergency. In a black box model representation, the inputs are the material flow of fluid and the signal sent by the operator, and the operation consists in transforming the incoming signal in a cessation of the material flow of fluid. In normal situations, the valve is held open by a spring and the fluid can pass freely. When the need of stopping the fluid arises, the operator can send a signal and the valve performs its function by closing the flow. Thus, a failure will occur when, given that a signal has been sent by the operator, the material flow is not terminated.

Even though Rausand and Øien's characterization of the notion of function is derived from the influential work of Pahl and Beitz, other interpretations can be found in the literature – see, for instance, (Erden et al.: 2008; Vermaas: 2009) –, which could possibly result in different criteria for failure. In this paper, the decision has been made to follow Rausand and Øien's characterization because, differently from most of other works on the subject, Rausand and Øien – whose

field of expertise is reliability engineering – discuss the notion of function from the perspective of engineers dealing with failure phenomena. Moreover, a similar stance on the notion of function can be found in many other engineering publications which deal with failure phenomena and related subjects (Blischke and Murthy: 2000; Tumer and Stone: 2003; Yellman: 2006; Birolini: 2007; Grantham Lough et al.: 2008; Bellgran and Säfsten: 2010).

By looking at Figure 4.1 again, it can be seen that knowing the black-box description of the function of an item (e.g., "close flow of fluid") is not sufficient for making a failure judgment: at least one performance parameter is needed (e.g., voltage, pressure, torque, etc.), alongside with a target level and acceptable deviation. In fact, Rausand and Øien observe that in order to "identify the failure modes we have to study *the outputs* of the various functions" (76, emphasis added) performed by the items. For shutdown valves the needed output or performance parameter is given by the time it takes the valve to close the flow of fluid: if the valve closes too fast, dangerous pressure shocks may ensue; if it closes too slowly, it will be ineffective. Thus, a typical target level for shutdown valves is that they are able to close within 10 seconds, with an acceptable deviation of plus or minus 4 seconds. So, the curve in Figure 4.1 can be interpreted as representing a valve which, at the beginning, is able to close in 10 seconds; after a while the valve becomes increasingly faster such that a failure event occurs when the closing time drops below the 6 seconds threshold, after which the valve is in a fault state.

The combination of the failure-related definitions given by the IEC vocabulary together with the black-box concept of function gives rise to the *Function-based* notion of failure (FBF) which can be considered the traditional notion of failure in engineering. The main ontological commitments behind FBF could be summarized as follows. Engineering artifacts or "items" (which the IEC vocabulary defines as "any part, component, device, subsystem, functional unit, equipment or system that can be individually considered") are *continuants* characterized by the attribution of the ability to perform one or more required functions. Even though the IEC terminology offers only scant support, it can be assumed that the attribution occurs when the item completes successfully the manufacturing or construction stage and is approved by the quality checks. When the abilities which actually inhere in the items coincide with the attributed abilities, the items *participate* to functioning states. Since functioning states have temporal parts (e.g., during the first month the solar panels produced 300 kWh

of energy) they are *non-atomic occurrents*. Participation to functioning states does not imply that the items have to be actually performing their required functions. They can be in stand-by mode, like a back-up power unit waiting to be called into action, or standing on the shelves of a store. If the attributed abilities do not match the actual abilities, the items participate in fault states. Similarly to functioning states, fault states have parts (e.g., the engine was making rattling noises for a while and then it stopped completely) and are *non-atomic occurrents*. Since actual abilities of items can change in time, the IEC terminology stipulates that the transition from functioning to fault states is singled out as the failure event. Failure events are *atomic occurrents* to which items participate. It should be stressed that failure events do not need to be anything spectacular. It can be that an item, say a printer, is shipped to a dealer's store while being in a functioning state. Then, for some reason, the printer may lose the ability to print while standing idle on the shelf. Thus, a failure event has occurred, even though no one noticed.

Even though FBF can deal with a large variety of circumstances deemed relevant in engineering practice and has reached a prominent status among the community, dissenting views have emerged that will be discussed in the next section.

4.3. Specification-based failure

One of the main critiques leveled at FBF is that it does not make clear who is in charge of deciding the acceptable limits of the functional output. So, are the users allowed to decide what counts as satisfactory performance or is that the job of engineers? Chillarege (1996, 354) openly takes the side of the users by claiming that "customer expectation largely determines whether a failure has occurred or not". Other authors claim this will result in untenable judgments. On the one hand, as observed by Yellman (1999, 7), customers might be satisfied with the output they get even though the product is performing demonstrably below the specifications. In his opinion, such cases represent a clear instance of failure "whether or not any customers have explicit current expectations for the unsuccessful functionality".

On the other hand, lack of expected output might be the result of the product being abused or operated outside the stated operational conditions. Mountaineer Neal Mueller (2006), for instance, complained publicly that his iPod fell silent

while he was climbing to the top of Mount Everest. The claim, however, conflicts with the product specifications which state a maximum operating altitude of 3000 meters (Apple Inc.: 2011). Remarkably, engineers themselves frequently exploit the interpretative flexibility of FBF to describe as failed items that have been misused. Harland and Lorenz (2005), as already mentioned in the Introduction, accept that a sensor which stops working because it is operated into an overly hot environment is described as having failed. Many similar examples can be found in the engineering literature. Kieselbach (1997, 55), for example, reports the results of the investigation on the bursting of a silo and concludes that "it can be said that *failure* of this silo was caused by filling it to too high a level with liquid instead of forage" (emphasis added). Similarly, Ross et al. (2007, 961), who describe the collapse of a heavy lift crane, use the term failure even though the investigation determined that the "loads which provoked incipient *failure* [...] were almost 2–1/2 times greater" (emphasis added) than the requisite design condition.

However widespread this kind of judgments may be, many engineers think that an item which is operated outside the acceptable limits and does not perform as desired "should not be considered failure in the *usual sense*" (Ezrin: 1996, 6). Similarly, Nieuwhof (1984, 54) states that if a one-ton truck is utilized to carry a 25-ton load, then when the truck eventually collapses "we should not talk about a truck failure". Engineers like Ezrin and Nieuwhof advocate a notion of failure that looks at items within the context in which they are operated and also at the expectations that are legitimized by the intentions of the designers. In fact, Nieuwhof proposes to distinguish between two notions of failure. One, called "equipment failure", is based on the intended functions and the "specified operational conditions for which [items are] designed" (54). The other, "mission failure", is grounded on the idea of "required feasible actions" which can be assimilated to required functional output, and does not make any reference to operational conditions. Haasl (1965), urges a similar distinction, though his terms of choice are "primary failure" and "secondary failure" respectively.

As a result, in this paper a second notion of failure is proposed, i.e., *Specification-Based Failure* (SBF):

> Specification-based failure event: the termination of the ability of an item to perform as specified provided it has been operated under the stated operational environment for which it is designed.

> Specification-based fault state: the state of an item characterized by inability to perform as specified under the specified operational conditions for which it is designed, excluding (i) the inability during preventive maintenance or other planned actions, or (ii) the inability due to lack of external resources, or (iii) the inability due to previous violations of specified operational conditions.

Clearly SBF is heavily influenced by FBF from which it inherits the terminology and the main ontological assumptions. Hence, also in SBF "termination" means the trespassing of the acceptable limits as depicted in Figure 4.1. However, instead of the term "required function", the expression "perform as specified" is utilized to underline the fact that the criteria for failure are the "specifications" established by the designers of the product. Moreover, a clause has been added which requires compliance with "the stated operational environment". Thus, the new concept aims at dispelling the ambiguities which make FBF a very permissive notion and, as a result, failure judgments like those expressed by Mueller, Harland and Lorenz, and others could not be passed.

It has to be stressed that, although SBF is less a liberal notion than FBF, it can be only as precise as the set of product specifications on which it relies upon. Even if stricter regulations and threats of legal actions force manufacturers into issuing more comprehensive specifications, in practice they cannot address all potentially relevant product properties. In particular, products age by the very fact of being utilized. For instance, fuel mileage and power output of a car can be maintained within specifications only on condition that the car is periodically serviced as recommended by the manufacturer.

Many SBF are, of course, also FBF, yet in Section 4.4 it will be shown that FBF and SBF are independent concepts. In the next section a third notion of failure will be analyzed which runs parallel to the other two already discussed and which is characterized by its focus on the material properties of items.

4.4. Material-based failure

The example of the corroded trailer barrel mentioned in the Introduction has shown that engineers can arrive at contradictory evaluations about failure: based on material properties Suess (1992) described the barrel as failed while, on functional grounds, Grantham Laugh et al. (2008) pronounced it fit for purpose. Sometimes the mixing of material-based and function-based assessment can occur within the same paper. Henshaw et al. (1999, 13) analyze "*the failure* of a particular brand of automobile seat belts" (emphasis added). The failure consist-

ed in the seat belt latch assembly losing the ability to fasten properly the belt clasp, even when operated according to the specified procedures. The investigators found that small fractured pieces from the press release button (one of the components of the latch assembly) could become lodged within the assembly and interfere with its correct operation. Henshaw et al. remarked that "it is ironic that the breaking away of these small pieces *does not impede the function* of the release button itself" (17, emphasis added). Nevertheless, few sentences later, when looking closely at the offending component, they speak of "degradation *and failure* of the release button" (18, emphasis added) and conclude that "failure of the release buttons involved a combination of (1) repeated, low-level impact damage and (2) degradation of the material" (19).

Again, two rather different meanings of failure are at stake here: one based on functional grounds (latch assembly) and one relying on material properties (press button). In the previous sections, while dealing with FBF and SBF, there was no need to mention material properties for the simple reason that engineering artifacts can fail for a variety of reasons that do not involve any kind of material degradation. Take, for instance, a printer where, because of a design flaw, the rolls feeding the sheets of paper from the paper drawer exert insufficient pressure. The printer and all its components are in pristine conditions and meet all the specifications. Still, the sheets of paper get jammed in the mechanism and the printer fails to perform its required function. Another example, even more eloquent, is given by Collins and Daniewich (2006, 860) who remark that a shear pin which *does not* separate into two or more pieces upon the application of a preselected overload must be regarded as a failure, "as surely as a drive shaft has failed if it *does* separate into two pieces under normal expected operating loads" (emphasis in the original). Both events (i.e., shear pin and drive shaft) qualify as FBF and SBF, but there is a material aspect with the second that sets it apart: the material properties of the item have changed – it has fractured – such that it has lost the ability to perform its required function. Therefore, the shaft separating in two pieces counts also as a *Material-Based Failure* (MBF).

Even though fracturing and rupturing can be considered the paradigms of MBF, engineering taxonomies contain many other failure mechanisms which do not result necessarily in fracture or rupture of the affected items. In fact, as noted by Dasgupta and Pecht (1991, 531), although engineers may be tempted to think of failure in a binary manner as something being obviously fractured or not, "most real failures are more complicated than that", which means that also

non-fractured items can be said to have failed. What Dasgupta and Pecht are referring to are the numerous physical and chemical processes (i.e., the *failure mechanisms*) that result in permanent degradation of material properties. Fracturing is just one of these processes, alongside fatigue, corrosion, wear, creep, radiation damage, buckling, and so on (Collins: 1993; Tawancy et al.: 2004).

What has to be established now is whether MBF can be considered as a separate notion or just as a sub-kind of the other two notions. Indeed, the engineering literature suggests that MBF can qualify as a separate notion. The reason is that materially degraded items may be classified as failed even though they are still able to deliver their required functional output (albeit close to the acceptable limits) and do not satisfy the criteria for SBF. These cases occur when items have degraded, for whatever reason, much faster than anticipated making the items less reliable and safe to use and, ultimately, increasing the likelihood of an incoming FBF or SBF. To put it differently, considerations based on the material properties may induce engineers to declare items to be in a fault state even though considerations based on functional output would not (yet) sanction such judgments. Let us consider again the case of the stainless steel trailer barrel analyzed by Suess (1992). The investigation found that the chemical composition of the steel did comply with the requirements and that "*failure* was caused by bacteria-induced corrosion" (73, emphasis added). The most likely explanation was that water contaminated by sulphate-reducing bacteria was used to wash the barrel. Since it is known that this kind of bacteria can attack stainless steel, barrels should be dried immediately after washing. In the case at hand, the barrel had not been dried, and the material was exposed to environmental conditions for which it was not designed. Thus, it would be inappropriate to describe the event as an instance of SBF. Moreover, as noticed by Grantham Lough et al. (2008), the barrel was still able to perform its required function and FBF should be ruled out as well. Suess assessment, then, results from the observation of the negative impact of corrosion on the remaining life and residual strength of the barrel. Thus, it was an instance of MBF.

Therefore, MBF is proposed as a third notion of failure with the following definition:

> Material-based failure event: any permanent change in the values of geometrical or physicochemical properties of the materials of an item which (i) renders the item unable to perform as specified or (ii) increases substantially the likelihood that the item will become unable to perform as specified.

> Material-based fault state: the state of an item resulting from any permanent change in the values of geometrical or physicochemical properties of the materials of an item which (i) renders the item unable to perform as specified or (ii) increases substantially the likelihood that the item will become unable to perform as specified.

Here, the term "permanent" should not be interpreted in an absolute sense: changes are considered permanent when repairs are needed to restore the condition of the item. Certainly, temporary changes in geometrical properties, like reversible thermal expansion, can cause an FBF or an SBF (e.g., seizure of a valve), but are not classified as MBF because there has not been any degradation in material properties. As soon as the loads are removed, the items recover spontaneously their original conditions. On the contrary, the notion of MBF rests on the assumption that degradation processes can change permanently the abilities of items.

It is worth emphasizing that the focus of the notion of material-based failure is on the changes occurring to the properties *of materials* of which items are constituted: wear can change geometric properties without affecting physicochemical properties of materials; embrittlement and radiation damage act only on physicochemical properties; and corrosion can change both. The notion of material-based failure is not concerned with geometrical changes occurring *to the item* as a whole, like the displacement of a component within an assembly because a screw got loose. The event in which a car and one of its wheels part company because the retaining nut had not been tightened adequately counts as an FBF of the car; however there is no contextual material failure of the car (not yet, at least) nor of the retaining nut. On the other hand, if the wheel gets loose because the retaining bolt snapped, then the snapping of the bolt counts as a an MBF as well as an FBF of the bolt itself. To decide whether the snapping counts also as an SBF the operating conditions must be known: if the bolt was utilized according to the specifications, then an SBF has occurred. If the bolt was not utilized appropriately, e.g., it was not the right bolt, then no SBF has occurred.
In the next section, the trailer barrel case story will be used as a test bed for showing the mutual independence of the three notions of failure.

4.5. A case story: the mutual independence of the three notions

The case story discussed by Suess (1992) deals with a stainless steel trailer barrel which, albeit severely corroded, had not developed leaks and was still capable of

performing the required function "to store fluid". The failure investigation found evidence of bacterial attack. Stainless steel is not designed to withstand this kind of environment. Indeed, the investigation did conclude that changes in the washing procedure were to be implemented for preventing recurrence. The fact that the barrel was utilized under harsher conditions than specified implies that the barrel cannot be said to have incurred in SBF. Summing up, the original version of the case story, i.e., scenario (1), features the following combination: FBF, no; SBF, no; MBF, yes.

As observed by Suess, given the appropriate conditions bacteria-induced corrosion can be very fast and, if undetected, can result in perforation of the tank and leakage. In that case, the barrel loses the ability to perform its required function and an FBF is said to have occurred. Hence, in scenario (2) of the case story the following failures would occur: FBF, yes; SBF, no; MBF, yes.

As a third variation, let us assume that the same amount of corrosion was found on the internal surface of the barrel, i.e., no leakage, but the investigation established that nothing was wrong with the water or the washing procedure, the culprit being the defective quality of the steel. Then, even though the barrel performs the required output, an engineer would describe the situation as an instance of SBF due to the thickness of the barrel being below the specifications. Hence, scenario (3): FBF, no; SBF, yes; MBF, yes.

If the situation depicted in the previous scenario progresses until corrosion opens a hole in the barrel, FBF will occur. Therefore, in scenario (4) the barrel suffers all three kinds of failure: FBF, yes; SBF, yes; MBF, yes.

As already mentioned above, a product may be in a state of FBF even though it has not suffered any MBF. The barrel may be leaking because of a fissure resulting from a manufacturing defect, e.g., inadequate welding. Then, since the leaking violates the product specifications, also SBF is present. Summing up scenario (5): FBF, yes; SBF, yes; MBF, no.

In a further permutation, thanks to a fortunate circumstance the fissure happens to be located in the uppermost part of the barrel. Since the user does not fill up the tank until the very top, the tank is never observed leaking and is considered to be fully functional. Still it falls short of the specifications which require the tank to store fluid up to the rated capacity. Therefore, scenario (6): FBF, no; SBF, yes; MBF, no.

In the last failure scenario, the barrel has been filled above the specified limit. During transportation the fluid expands and leaks through the flanges, thus

without causing material damage. The event does not qualify as an SBF or as an MBF, hence scenario (7): FBF, yes; SBF, no; MBF, no.

To conclude, scenario (8) represents successful operation: FBF, no; SBF, no; MBF, no. The eight failure scenarios are summarized in Table 4.1.

Table 4.1: Eight failure scenarios that illustrate the mutual independence of the three notions of failure

Scenario	Function-based failure	Specification-based failure	Material-based failure
(1)	N	N	Y
(2)	Y	N	Y
(3)	N	Y	Y
(4)	Y	Y	Y
(5)	Y	Y	N
(6)	N	Y	N
(7)	Y	N	N
(8)	N	N	N

4.6. Discussion of main ontological commitments

In this paper the notion of failure as defined by the IEC vocabulary has been used as a guideline and a template for the identification and the analysis of three independent notions of failure. As a consequence, a number of conceptual aspects and ontological assumptions are shared by the three notions. At the most fundamental level is the ontological assumption that both failures and faults are *occurrents* to which engineering items participate. The term "item" recurs in all definitions and refers to physical entities characterized by a complex quality, namely the quality of being attributed the ability to perform required functions. In turn, required functions are seen as operations on flows of energy, materials, and signals.

Moreover, the notion of failure demands that the functional outputs of operations on flows of energy, materials, and signals are specified by means of appropriate target levels and acceptable limits. For the notion of FBF it is sufficient that the manufacturer of the item specifies the acceptable limits of the functional output; while, the notion of SBF demands that acceptable limits are defined for inputs, outputs, and operational environment. Let us assume that the

curve in Figure 4.1 represents the torque generated by an electrical motor which happens to be operated at an environment hotter than specified. After a while, the motor overheats and the functional output drops below the acceptable level. According to FBF, a failure event has occurred which could be further qualified as a "misuse failure" if the incorrect operational environment was due to actions or omissions on the part of the user. Thus, an FBF failure event can be defined as an atomic occurrent, to which an engineering item participates which is characterized by a transition from correct functional output to incorrect functional output. The ensuing FBF fault state will be defined as a non-atomic occurrent to which an engineering item participates which is unable to perform the required functional output. Differently from failure events, fault states are not atomic because they can have temporal parts. For instance, at the beginning of the fault state the overheated electric motor is still able to provide some amount of torque. Then, if utilization continues nevertheless, the motor can stop working altogether and perhaps for good.

The sequence of events just described does not qualify as an instance of SBF because of the violation of the product specifications. An SBF failure event can be regarded as as an atomic occurrent to which an engineering item participates. An SBF consist in the transition from compliance with specification to lack of compliance while the operational environment remains within the specifications. The ensuing SBF fault state is defined as a non-atomic occurrent to which an engineering item participates and characterized by the inability of the item to meet the specifications while the operational environment remains within the specifications. A SBF fault state can be the effect of a previous SBF event or of an FBF event; alternatively, in case of a design flaw or of a manufacturing defect, the item can find itself in a fault state from the very beginning.

So far, the discussion has dealt only with the first two notions and MBF has not been mentioned. In fact, even though the basic distinction between events and states holds also for MBF, this notion appears more challenging and complex. First, it introduces a distinction between properties of the materials that constitute an item and the item itself. Second, the changes in material properties that are relevant are only the permanent ones. Finally, the notion of material fault state depends on the previous circumstances. While an item can be in a FBF fault state from the very beginning, say because of a manufacturing defect, an item needs to go through an MBF failure event in order to enter into a MBF fault state.

4.7. Conclusion

The possibility of failure is a persistent source of concern for engineers. Failure can be subtle and minor changes in design or in manufacturing techniques can turn a robust product into an unreliable or even a dangerous one. Tools could be devised to assist engineers in archiving, retrieving, and reusing information about failure. Formal ontologies are one of the candidates. However, as argued by Borst (1997) and by Guarino et al. (2009), in order to be effective these tools need to be based on clear and shared conceptualizations. Unfortunately, the engineering literature offers a multitude of definitions partially conflicting with each other. Even the IEC definition of failure, which is often considered to be "the traditional definition", has met with critique. In this paper a conceptual analysis of the notion of failure as used by engineers has been performed. As a result, it is argued that three mutually independent notions can be identified: FBF, SBF, and MBF. The paper has examined the three notions and has sketched their main ontological assumptions.

It should be stressed that, although in this paper the analysis has been confined to the domain of engineering artifacts, the notion of failure plays a relevant role also beyond the artifactual domain. Avizienis et al. (2004), for instance, discuss the notion of failure within the context of a taxonomy of basic concepts for information systems and secure computing. Moreover, the notion of failure has strong conceptual and practical connections with the issue of human error or, more generally, of human and social factors especially within the context of complex socio-technical systems. At the most basic level, social practices such as supervision, training, and knowledge sharing have considerable influence on the likelihood of failure events. Formal ontologists are already actively investigating this area of research where technology and social factors interact closely, e.g. (Bottazzi and Ferrario: 2005; Ferrario and Guarino: 2009; Scherp et al.: 2011). Even though these studies have not addressed explicitly the notion of failure yet, it is reasonable to expect that it will attract more attention in the near future.[9] Therefore, future research might explore the possibility of expanding the conceptual analysis performed in this paper into the socio-technical domain.

---

[9] An exception is recent work by Bottazzi and Ferrario (2011) which examines the notion of "faulty institutional object"

5 Root Cause as a U-turn[10]

Abstract

Failure analysis is the process of identifying the causes and factors leading to undesired loss of functionality. Failure investigators use several kinds of notions to explain this loss. An important one is that of a root cause, but investigators still disagree about the exact meaning of this term. We maintain that two approaches to define root causes can be found in the literature. One originates in backward-looking causal analysis, which aims at determining the causes and factors accompanying a specific failure event; it is token-based and comprises mainly deterministic reasoning. The other is associated with forward-looking effects analysis, which is type-based, and sets out to find correctable factors and prevent recurrence by mainly probabilistic reasoning. Drawing on case studies from the engineering failure-analysis literature, we propose to combine the two approaches to form a new sensible notion of root cause as a U-turn.

5.1. Introduction

Several engineering disciplines and activities deal with product or component failure, such as risk assessment, safety science, reliability engineering and failure analysis. Avoiding failure is an even more important aim of engineering design. The notion of failure, however, does not have the same meaning in the various disciplines and activities. In reliability engineering, for instance, the notion of failure is mainly associated with the statistical tools used to define the failure rate of an item. In risk assessment, the effects of failure are crucially important and, multiplied with the probability of failure, are used to calculate the risk of putting the item to work. Failure analysis has two related connotations. On the one hand, the term refers to a body of knowledge, which develops scientific and engineering tools (e.g., models, methods, theories) to analyze and explain failure phenomena. As noted by Wulpi (1999), it is a very complex field

[10] This chapter has already been published as Del Frate, L., Zwart, S. D., and Kroes, P. A. (2011) 'Root cause as a U-turn', in: *Engineering Failure Analysis* 18 (2): 747–758.

based on contributions from many disciplines and from specialist fields as diverse as structural engineering, chemistry, fracture mechanics, fractography, stress analysis and metallurgy, to name but a few. On the other hand, failure analysis has to do with the *investigative process* regarding a specific failure event and to the application of the tools mentioned, and is synonymous therefore with failure investigation. To quote the definition given in the *ASM Handbook* (Aliya: 2002, 315), "failure analysis is a process performed in order to determine *the causes or factors* that have led to an undesired loss of functionality" (emphasis added). Although carrying out a failure investigation can have many reasons (e.g., to assign responsibility, to prevent recurrences or to improve productivity), it is widely accepted among practitioners that failure analysis *is* the process of finding the causes and factors that led to the failure in the first place. It is less clear what the nature of these causes and factors is. One consequence of this lack of clarity is the proliferation of terms and taxonomies, among which the distinction between causes and factors is a telling example. Other related terms that appear in the literature are: primary cause, immediate, direct, underlying, probable, latent, secondary and of course root cause that forms the topic of this paper and is probably the most controversial of all the terms.

The aim of this paper is to put forward a notion of root cause that fits into the failure analysis framework, and, while remaining credible to practicing engineers, that takes into account several reservations of the notion's opponents. It should be made clear, therefore, that this paper is not criticizing existing failure or accident models, nor is it proposing a new model for its own sake. We will argue that a definite meaning can be assigned to the notion of root cause, which is based on the distinction between backward-looking *causal analysis* and forward-looking *effects analysis*. These two directions of analysis occur in almost any in-depth failure investigation. Most of the failure and accident modeling literature does not, however, explicitly distinguish between the two, although the distinction contributes significantly to the conceptual clarity of the notion.

The paper is organized as follows. In Section 5.2 we present a range of definitions of root cause drawn from the failure and accident analysis literature. The definitions show that the notion is considered an important one, yet lacks consensus. Two main approaches appear to oppose each other. Section 5.3 provides a conceptual reconstruction of the backward-looking approach. To that end we introduce the state-sequence diagram and the failure phenomenon diagram. Section 5.4 deals with the forward-looking approach, which is charac-

terized by the emphasis on finding corrective factors to prevent recurrence. In Section 5.5, we combine our findings of Section 5.3 and 5.4, and introduce the notion of root cause as a U-turn. Finally, we draw our conclusions in Section 5.6.

5.2. Root cause

Even though it is widely accepted that the aim of failure investigation is to establish the causes of undesired and unexpected loss of functionality (Becker and Shipley: 2002), the nature of these causes and the criteria to assess whether the investigation was able to find them are open matters. It is important to note that this is not a factual issue but a conceptual one. Factual mistakes may be made during an investigation so that the correctness of its conclusions is compromised (e.g., a fracture surface is inadvertently contaminated). The point is rather that failures almost always have multiple causes which are connected to each other and to the failure event in complex ways (McKinnon: 2000; Aliya: 2002; CCPS: 2003; Bhaumik: 2009; Le May and Deckker: 2009; Ferjencik: 2010). Some of these causes are more apparent, their connection to the event is straightforward and strongly supported by available evidence. Usually, such causes are spatially and temporally proximate to the failure event and for this reasons they are usually called proximate or direct causes (also rather popular terms are physical, primary or active causes). According to Bhaumik (2009) the majority of failure investigations stop at the level of the proximate causes even when there is evidence that more remote factors played a role. This is not to say that proximate causes are irrelevant, but that they provide just a partial answer. A better answer would be able to tell what went wrong so as to create the conditions for the physical failure to occur, that is to say the root cause. Bhaumik thinks that looking for root causes makes it inevitable to deal with human and organizational factors. His characterization of root causes shows many similarities with the one defended by Scutti (2002). Scutti proposes a layered structure in which physical causes are caused by human causes that, in turn, are caused by latent causes which have roots that are organizational or procedural in nature. The view that root causes are organizational factors is well represented in the failure (and accident) analysis literature. NASA (2006) procedures for mishap reporting, for instance, explicitly define root causes as organizational factors. Similarly, other sources (Heinrich: 1980; Abdelhamid and Everett: 2000; Le

Coze: 2008; Murphy: 2008) characterize root causes as management system deficiencies.

Besides the organizational and managerial concept of root cause, however, alternative definitions have been suggested in the literature, or perhaps it would be more appropriate to say that a whole array of positions is available. Busby (2001, 1419) portrays them as the causes that do not have antecedent causes, that is to say, *the absolute beginning* of the chain of events (van Vuuren: 1999, 19). More vividly, Andersen and Fagerhaug (2006) describe root causes as "'the evil at the bottom' that sets in motion the entire cause-and-effect chain". As already noted by many ,e.g., Doerner (1980, 102), this is an attractive hypothesis because it reduces uncertainty with one stroke and encourages the feeling that things are understood and for this reason Carroll (1995) has called it "the root cause seduction". However, making sense of it has proven an insurmountable challenge because every sensible notion of causal influence proposed so far allows attaching at least one causal precursor to every conceivable event, except the origin of the universe itself (Hollnagel: 2004; Reason: 2008).

Trying to avoid this kind of criticism, Kinnersley and Roelen (2007, 33) define root causes as conditions which are necessary for an accident and stay clear of the problematic notion of *uncaused cause*. It is doubtful that their proposal amounts to significant improvement, though. On the one hand, the notion of necessary condition rather easily falls back to the notion of cause but, then again, any necessary condition has antecedents which, in turn, may be necessary conditions. On the other hand, necessary conditions may be discarded by failure analysts because of their minor explanatory value. Gravity, for instance, is by any means a necessary condition for the collapse of a building, but this is hardly an informative statement. Johnson (2003, 184) takes a slightly different stance and formulates a counterfactual definition: "if a root cause had not occurred in the singular, particular causes of an incident then the incident would not have occurred". The problem with Johnson's definition is that it seems unable to differentiate between proximate causes and root causes.

The intuition that root causes are in some respect special causes appears several times in the literature. Sheridan (2008, 421) considers them the "most responsible" among the ones appearing in the chain of causation. Wood and Sweginnis (2006, 7) recall that not so long ago aviation accident investigators in the US were required to prioritize causes to describe their contribution to the accident. Some organizations still differentiate causes and factors in terms of

their degree of connection in relation to the occurrence (Walker and Bills: 2008), again making the root cause as the predominating one. Although an intuitively attractive hypothesis, it has never been developed into a formal definition equipped with operational criteria to assess the degree of connection. If these criteria are missing then the ranking of causes becomes a subjective matter highly sensitive to the context. Admitting the lack of clear criteria, Mobley (1999, 19) concedes that root causes are often subjective (especially in relation to injury-causing accidents). Other authors have chosen to emphasize yet a different aspect of root causes, that is, the *correctability* aspect. The US Department of Energy G*uidelines for Root Cause Analysis* (1992, 1) defines root cause as the fundamental reason which, if corrected, will prevent recurrence. Others also express a comparable emphasis on correctability and prevention (Becker and Shipley: 2002; Wood and Sweginnis: 2006; Department of the Air Force: 2008; Hokstad and Rausand: 2008).

A few remarks about the short survey above are needed. First, it is not meant to be complete or cover all positions expressed by engineers, safety scientists, reliability engineers and such. The main reason for the survey is to show that the notion of root cause is considered important and controversial by at least part of the failure (and accident) analysis community. Second, it is not claimed that the various positions presented in the survey are mutually incompatible. In fact, some authors try to reconcile in one definition multiple aspects of the notion of root cause. For instance, Wood and Sweginnis (2006, 7) observe that besides its correctability, a root cause is frequently related to management issues. Paradies and Busch (1988, 479) and Marquez (2007, 127) define root cause as the most basic failure cause that can be reasonably identified and that management has the control to fix. This way it reunites aspects of root cause that were considered in isolation in other definitions (e.g., organization and management; correctability).

The survey suggests there is a tension between two perspectives on root causes. One is backward-looking and focuses on the chain of events that led to the actual failure. The idea is to look for the beginning of that chain, or to the necessary link that holds the chain together. And this search is performed going back in time, drilling down (Latino and Latino: 2006) until the root cause that set off the chain is revealed. The other perspective, while acknowledging the need for in-depth backward-looking causal analysis, holds that more backward-looking is not likely to reduce the number of contributing factors until the most

prominent or the absolutely necessary one is discovered. Drilling down into the details of a particular failure will possibly unearth local circumstances and coincidences which are specific to this exact failure but are not likely to occur again in the same way and are therefore not helpful to preventing recurrence. Instead the approach should be to learn as much as possible from the occurrence and to reason forward, applying the lessons learned in order to prevent recurrence.

This tension has come to the foreground in a debate which is currently going on among safety investigators in the aviation sector. At the 2008 seminar of the International Society of Air Safety Investigators (ISASI), Michael Walker, Senior Transport Safety Investigator at the Australian Transport Safety Bureau (ATSB) presented a paper discussing how the ATSB has approached causation as part of its investigation analysis framework (Walker: 2009). Recently ATSB has reviewed the terminology utilized in official accident reports and the term 'cause' has been dismissed and replaced by 'contributory safety factor'. According to Walker, the terminological overhaul has not just practical reasons, but also reflects a substantial conceptual change. The practical reason for dismissing the term 'cause' was that it is easily associated with legal proceedings, especially those aiming at allocating responsibility in the wake of an accident. The ATSB statute, like those of other national safety boards, makes clear that the purpose of a safety investigation is to enhance safety and not to apportion blame or liability. The conceptual reason is that the term 'contributing safety factor' is more inclusive, and can therefore provide a richer picture of the factors involved in the occurrence. These terminological and conceptual changes notwithstanding, it seems that the ATSB approach is still predominantly backward-looking. Indeed, contributory safety factors are counterfactually defined as factors that if they had not occurred then the failure or accident would probably not have happened. This is closely reminiscent of Johnson's (2003) counterfactual definition of root cause above. Moreover, the notion of contributory safety factor is connected to a link-by-link approach which is the equivalent of a traditional chain of events. It starts with a failure event (occurrence) and then, following back the various links, it attains the initial factor of the chain. ATSB has established that (ideally) a factor should have a likelihood of at least 66% to be included in the chain (Walker: 2009, 24).

MacIntosh (2010), Chief Advisor, International Safety Affairs, at the US National Transportation Safety Board, presented a dissenting view at the 2009

ISASI seminar. While MacIntosh acknowledges the worries about culpability expressed by the ATSB, he defends the appropriateness of concluding an accident report with an explicit causal statement centered on a single cause (the root cause) on the grounds that it will emphasize attention on the most crucial factor for corrective action and will provide extra momentum for overcoming financial and political resistance to safety improvement. His argument, then, exemplifies a forward-looking notion of root cause.

The aim of this paper is to make explicit the distinction between the backward-looking and forward-looking approaches to root cause and combine them in one viable concept. We argue that the two notions derive from the two directions of analysis that occur in failure investigations. One is backward-looking causal analysis, the well-established investigative process which aims at finding the causes of a failure event as described by the *ASM Handbook* quoted in Section 5.1. Bhaumik (2009, 185) has already noted that a majority of failure investigations limit themselves to this direction of analysis and stop when they have identified the main physical factors. The backward-looking notion derives from the assumption that the root cause of an occurrence can be found by pushing further in this direction until a cause stands out and proves to be the initiator of the chain of events, or the necessary initiator as exemplified by the definitions above.

The forward-looking direction of analysis is pursued less frequently and is usually applied in the wake of major failures that resulted in substantial losses. The main concern is preventing reoccurrence so in-depth knowledge of the occurrence and the circumstances that led to it is, of course, fundamental. No less important, however, is thinking about what the past tells about the future, how representative it is and what are the lessons to be learned. It would be shortsighted to try to prevent an exact replica of a failure that has already occurred. From the forward-looking perspective a detailed investigation that backtracks all the conceivable events, circumstances, and individuals that had some influence on the failure is not worth the effort. Anticipating – or controlling – the future with such detail is not feasible. Thus, the root cause will not stand out as a factor found by causal analysis alone; additional considerations about the future are needed.

To prepare the ground for our concept of root cause that combines together backward-looking and forward-looking approaches, Section 5.3 examines the

main features of the backward-looking direction of analysis and the way it deals with the issue of causality.

5.3. Backward-looking approach

It is a standard practice in failure analysis to write reports with two clearly distinguished sections. The first is mainly descriptive and the second is devoted to the analysis. The aim of the first part is to provide a factual narrative summarizing knowledge about the states which the failed item has gone through. It is based on the application of scientific and engineering methods applied to available evidence. In some instances knowledge about the evidence is rather easily transformed into knowledge about the state of the investigated item, for instance when video footage shows that at a certain time the item was operating normally, or when a maintenance log testifies that a component stopped working and was inspected. In other cases, multi-step deductions have to be performed to achieve sensible knowledge on the basis of evidence, for instance when deducing the dimension of the initial defect from a study of the final fracture surface (Janssen et al.: 2004).

The aim of the analysis is to explain the known facts by means of a causal narrative. In the case of relatively simple failure events the causal narrative consists of a series of causes and factors that explain why the item was in certain states (compared to other rationally conceivable states) and explains the transitions between consecutive states. In the case of complex failure events the causes will be arranged in branching trees and partially overlapping layers instead of simple chronological series.

Distinguishing the factual narrative from the causal narrative in the final report is not just a well-established writing strategy. It also mirrors an important difference between the fact-finding aspect of the investigation and the explanatory aspect. To account for these differences, the rest of this section offers a conceptual reconstruction of the backward-looking analysis based on two outcomes. The outcome of the descriptive process is captured by a *state-sequence diagram* and the outcome of the analysis process by the *failure phenomena diagram*.

5.3.1. The state sequence

The aim of the state-sequence diagram is to represent descriptive knowledge of the states of the failed item. Therefore it is important to provide clear definitions of the key states represented in the diagram. A crucial one is the failure state which is variously defined in the engineering literature, see (Del Frate et al.: 2011) for a recent survey. For this paper, we decided to adopt the definition provided by IEC 50(191) (1990) standard because, besides being widely used, the same standard also defines some closely related notions, like fault and error, that are sometimes confused with failure (Rausand and Øien: 1996). The definition reads as follows: "failure is the termination of the ability of an item to perform a required function". The standard also specifies that failure is an event which occurs when the item exceeds acceptable performance limits, while fault is the state in which the item is after failure. More precisely, "a fault is often the result of a failure of the item itself, but may exist without prior failure".

In their analysis of the IEC definitions, Rausand and Øien (1996) provide a useful graphical representation which is reproduced in Figure 5.1. In the diagram, the item's performance is plotted against time. Also represented are the expected performance level (or target value) and the predefined acceptable limits of performance within which the actual performance is allowed to fluctuate. For illustration purposes, Rausand and Øien assume that when the item is put into service its actual performance matches the target value and then it starts gradually diverging from it.

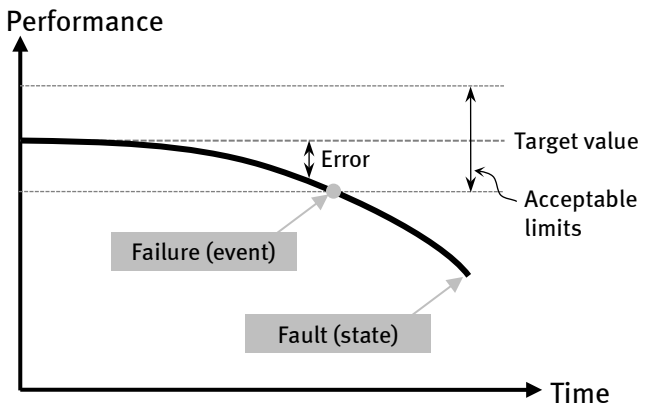

Figure 5.1: Rausand and Øien's (1996) illustration of the definitions of failure, fault and error based on IEC 50(191) (1990)

In accordance with IEC 50(191) the difference between the expected value of performance and the actual value is named 'error'. The failure event is then defined as the point in the diagram when item performance exceeds acceptable limits. It should be noted that according to the illustration after the failure event has occurred the item is still performing, although at disappointing levels.

It is easy to envisage that this would not always be the case. For instance, a catastrophic brittle fracture usually results in an abrupt downturn of the curve and the item's performance almost instantaneously drops to zero. Similarly, many other variants of the diagram may be anticipated depending on the item's failure mechanism, target performance profile, and considered time span. Rausand and Øien did not investigate these variations because their intention in drawing the diagram was just to illustrate some crucial definitions included in the IEC 50(191) standard.

We introduce state-sequence diagrams that are based on Rausand and Øien's diagram and that aim at representing these variations. These state-sequence diagrams are meant to provide a general representation of failure events and, at the same time, frameworks for the failure investigation process. Figure 5.2 shows an example of a state-sequence diagram for a hypothetical failure by wear. As in Rausand and Øien's diagram the item's performance is plotted against time. It is crucial that the state-sequence diagram is clearly associated to one exact, perhaps compound, item. Multiple diagrams can then be drawn for the same item, depending on the measured performance.

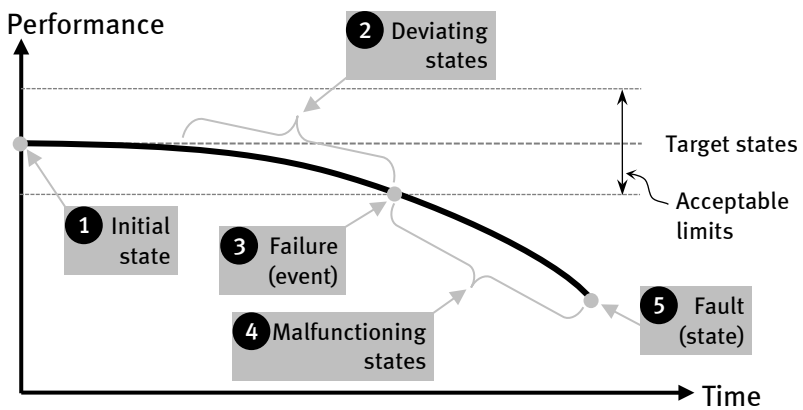

Figure 5.2: Example of a state-sequence diagram for a hypothetical failure by wear

It is also possible to split the main diagram of a compound item into a number of secondary diagrams at the component level.

Five types of states are highlighted in the diagram, together with the expected target states and the acceptable limits which are fixed by the item's specifications. *Failure event*, in accordance with the IEC 50(191) definition, is the event at which item performance exceeds acceptable limits. *Deviating states* are those in which item performance is still within limits albeit diverging from the target value (compatibly with specified tolerances). Note that, given the above definitions, objective performance criteria are provided for the identification of these states. It then depends on the availability of evidence as to whether the investigation can establish the time of occurrence and thus locate the points in the diagram.

In contrast, performance criteria are not fully provided for the fault state. According to the IEC standard, the *fault state* of an item is characterized by inability to perform a required function. This definition does not preclude that after exceeding acceptable limits the item (especially a compound item) still provides some level of performance. For instance, a car's braking performance may fall below acceptable limits but still have some braking capability. Therefore, it would be inappropriate to call all the states after failure event 'fault states'. We apply this label either to the state in which item performance is zero or to the state in which the item is declared not fit for service and removed from it. The fault state is also a prominent one because it marks the start of the backward-looking analysis. It describes the state of the item 'as received' by the investigators. On many occasions the fault state is substantially different from the state of the item corresponding to the failure event preceding it because of the effects of the failure itself. Indeed, post-failure events (e.g., an intense fire) may completely destroy important evidence about the failure event or, even worse, create misleading evidence. Of course, depending on the item's specifications and details of the event, the failure event and fault state may coincide and get represented in the diagram by the same point. Otherwise, the states occurring between the failure event and the fault state are named *malfunctioning states*.

Lastly, the *initial state* is the state located at the beginning of the sequence. It should not be confused with the item's first entrance into service, although the two states quite often coincide. The initial point of the sequence is selected by failure analysts on the basis of a number of epistemic and pragmatic considerations dependent on the scope and target of the investigation. For instance: there

is suitable amount of evidence about the state; knowledge about the state is widely accepted within the investigative team and possibly by the team's clients as well; in the causal history it can be shown that the state had a causal influence on following states (e.g., deviating states); and so on. The state sequence exemplified in Figure 5.2 depicts an item whose initial state matches the target value, but this does not always need to be the case. Actually, an item may be introduced into service without meeting expected targets. Yet, since it is within acceptable limits, it is not stopped or removed from service. An example of suitable initial state is the last maintenance check available that provided adequate evidence about an item's performance.

The state sequence represents investigators' descriptive knowledge about the failure phenomenon. Note that the state sequence is purely descriptive and avoids causal suppositions. True, capturing entirely the descriptive knowledge collected by investigators may require multiple diagrams, each related to a different performance parameter, that can then be combined in one multidimensional state sequence. Since the aim of a failure investigation is not only to come to a true description of the states but to achieve understanding of the state sequence, the description should be supplemented by an analysis that explains the phenomenon. This is the role of the causal history. It allows investigators to explain the state transitions, which are described in the state-sequence diagram.

5.3.2. Causal history

The causal history explains the state sequence by providing the causes and factors that either prevent states from changing or determine state transitions. There is a need for explanation because multiple causal histories are compatible with the same state sequence (causal under-determination of the state sequence). Thus the causal history supplies the causes and factors that were active in the main state transitions identified in the state sequence and excludes the possibility that other causes or factors played a role. Causes and factors may well have precursors of their own, possibly multiple precursors organized in diverging branches. In the case of complex failures with hundreds of causes and factors, finding out the entire causal history, even when abundant evidence is available, can be a difficult task. A number of methods have been developed to assist investigators, for instance: Fault Tree Analysis, Multilinear Event Sequencing (MES) introduced by (Benner: 1975) which has been developed into the Sequen-

tial Timed Events Plot (STEP) by (Hendrick and Benner: 1987), or Events and Causal Factor Charting (Johnson: 1980; Ferry: 1988; Buys and Clark: 1995; CCPS: 2003), and others.

Even though they have different roles, state sequence and causal history are not developed separately during the investigation. The state sequence provides the framework on which to attach the causes and factors. At the same time, growing knowledge about causes may influence the boundaries of the state sequence and demand it to be expanded further back in time. This is the case, for instance, when manufacturing defects or design flaws are assumed to have played a role. Then the initial state of the sequence has to move back to those stages of the item's life cycle.

Once investigators have completed the analysis, the identified causes and factors can be inserted in the state-sequence diagram. We call the combined failure and malfunctioning states together with their accompanying causal history the *failure phenomenon* to distinguish it from the failure event, which is only one of the states. Note that a failure phenomenon may include contributory factors that occur before the failure event. In addition we call the diagram of the entire state sequence plus all of the causal history the *failure phenomenon diagram*.

A simplified example of a failure phenomenon diagram is given in Figure 5.3. The example is restricted to the main causal factors and the state sequence has been approximated to a two-dimensional diagram in which several performance variables have been conflated in one linear variable. These limitations notwithstanding, Figure 5.3 suffices for the purposes of this paper and conveys an idea of what a failure phenomenon diagram looks like. It is based on a case history described in a recent publication by Gagg and Lewis (Gagg and Lewis: 2009). The investigated item was an estate car that crashed while being driven in a normal manner along a country road. The investigation started when the car, after the crash, was in the fault state and performance level was zero. From the driver's testimony two pieces of evidence were recognized: first the engine lost power without warning; second there was a loss of servo assistance to steering and breaking systems. This implies that at some point the car's performance exceeded the acceptable limits (failure event) and then went on operating outside the specs for a short time (malfunctioning states) before the crash occurred (fault state). It was then found that the loss of both servo assistance and power were two effects of a major fuel leak located at the fuel input connection

of the delivery rail (failure event). The fuel leak was caused by a major crack in the fuel supply pipe generated by a fatigue failure mechanism. Hence there was evidence of deviating states preceding the failure event during which the crack grew. The fatigue mechanism was initiated by a misalignment that brought about contact between the fuel pipe wall and the inner corner of a captive nut which created a stress-rising situation. A detailed examination showed that the connection between pipe and rail inlet had been broken and re-made post manufacture. At the initial state, that is, just after manufacture, the car was meeting the performance target value. At a later stage, during a service visit, the fuel rail was disassembled and then a misalignment was introduced during reassembly. From that moment onwards a fatigue mechanism was acting on the fuel pipes.

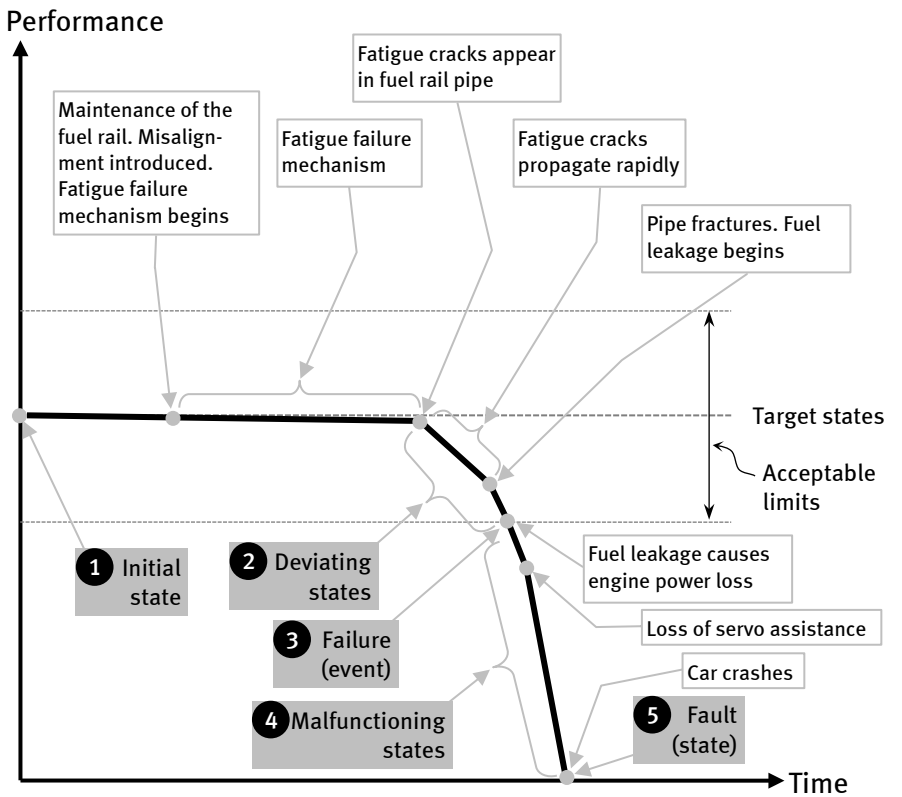

Figure 5.3: Example of (simplified) failure phenomenon diagram based on (Gagg and Lewis: 2009)

At first there were no appreciable effects on the car's performance level. Later on, once the cracks had reached a critical size, fuel leakage occurred and performance started deviating from the target values. The investigation's conclusion was that servicing was the most likely source of pipe misalignment and was therefore responsible for sensitizing the pipe to failure by fatigue. Both reconstructing the state sequence and finding out the causal history are performed by applying scientific theories and engineering methods to available evidence. Causal history, however, presents additional problems for which engineering and scientific methods can provide limited assistance. While uncertainty about the state sequence can always be reduced by additional evidence (provided that enough time and resources are available), the same does not hold for the causal history: some uncertainties are interpretative in nature and will not be reduced by additional evidence. We consider three aspects of causation generating interpretative problems: contextual sensitivity, role of negative factors, and presence of diffuse factors.

Contextual sensitivity is one important aspect proper to causal history that poses interpretative problems which do not affect the state sequence. The state sequence is descriptive and insensitive to context because it deals with measurable performance levels evolving in time. True, sometimes available evidence is not enough to provide a reliable performance assessment. Contextual sensitivity is not a problem raised by lack of evidence. The problem is that the same physical or mechanical events are classified as full-fledged causes or mere conditions depending on the context. Consider the scientific explanation of combustion. Three factors need to be present simultaneously: combustible material, oxygen and ignition source. In case a failure results in a fire, investigators will usually start looking for the ignition source and will identify it as the cause of fire. Presence of oxygen is almost always considered a mere condition not even worth mentioning. It is understandable why: in usual circumstances presence of oxygen is taken for granted and does not require explanation. In a word, it is irrelevant. There are contexts, however, in which presence of oxygen is relevant (Hart and Honoré: 1985). Consider a laboratory where experiments are carried out on flammable materials. For safety reasons, oxygen is removed from the lab. Then, if a fire breaks out it is sensible to look for the reason why oxygen was present and consider it the cause of the fire since, in that lab, the presence of flammable material and ignition sources are considered normal. Unfortunately, for most real case scenarios the choice of the appropriate context is not as easy as

in this hypothetical example. Investigators, then, have to decide what the appropriate context is.

Causal history has to deal also with the problem of *negative factors*. Again this is not a problem for the state sequence, simply because there are no 'negative states'; it is impossible to measure states which do not exist (note that a negative state is not a state for which performance is described by a negative number, e.g., temperature is -9°C). A typical example of a negative factor is absence or failure of safety barriers. The scenario is well exemplified by the catastrophic failure of a Hefler conveyor discussed in (Mobley: 1999). All the conveyor's left bars were severely bent before the conveyor could be stopped. A foreign object was established to be the cause of the blockage. Usually, these devices are provided with safety pins precisely to prevent such occurrences. Something went wrong with the shear pins, then. But for how long was the safety barrier malfunctioning? Similar conditions can be present for a long time before resulting in full-blown failure events. In order to deal with these the notion of latent or underlying failure has been introduced, e.g., (Reason: 1990). Again, in many real case scenarios the investigators have to decide if a certain condition can be interpreted as a latent failure.

Finally, some causes and factors in the causal history are not easily localized in space and time, and are therefore known as *diffuse factors* (Perrin et al.: 2006; Roelen et al.: 2011). This is the case with procedures, regulations, safety culture and other human factors which are simultaneously connected to a variety of activities within an organization and can influence the unfolding of a failure phenomenon. Yet again there are no clear cut scientific theories or engineering methods for handling such kinds of factors.

The fact that interpretative problems may affect the causal history does not imply that this must always be the case. Indeed, failure analysts can often achieve the correct causal history. These problems, however, can become an insurmountable challenge when the backward-looking direction of analysis is stretched further in the quest for the root cause. The assumption is that by digging deeper in the causal history a contributing factor will naturally emerge as the most prominent one. Prominence, as seen in Section 5.2, may be interpreted in different ways, as the absolute beginning of the causal chain, or the necessary factor that set in motion the chain, and so on. In contrast to contributing factors for which strong evidence can be provided, these attributions may be heavily dependent on investigators' interpretations. Many authors have

pointed out these difficulties and suggest taking a different approach. According to these critics, the root cause is not identified by a backward-looking analysis alone but by taking into account correctability considerations that demand forward-looking analysis. The next section will examine in more detail the alleged connection between correctability and the forward-looking direction of analysis.

5.4. Forward-looking approach

As anticipated in Section 5.2, several authors have expressed the opinion that the search for root causes is strongly associated with correctable factors. Wood and Sweginnis (2006), for instance, take a clear stance claiming that in general a statement of cause that does not contain some element of correctability is almost useless. Similar emphasis on correctability appears in Department of Energy *Root Cause Analysis Guidance Document* (1992) and Air Force *Instruction 91-204* (2008).

Intuitively, the notion of correctability seems rather clear and its relevance for failure prevention is easily understood. Wood and Sweginnis illustrate the notion by comparing two causal statements related to a hypothetical accident in the aviation sector. According to the first statement the cause of the accident is determined to be "Pilot error". In this statement, they claim, there is no element of correctability. A better statement would be "Pilot error due to lack of training on icing-detection techniques", and that would be a root cause.

It is important to note that Wood and Sweginnis' example and the notion of root cause they are defending are understood correctly only from a forward-looking perspective, according to which root causes are causal statements about future occurrences made on the basis of knowledge acquired by investigating a specific failure or accident. The investigation has proven that 'training on icing detection' was a factor. It was a factor, however, among (possibly many) other factors, another one of which was certainly 'presence of icing conditions'. Then the root cause derives from the forward-looking interpretation of the results of the investigation. And the interpretation can be phrased like this: "If appropriate corrections to the pilot's training program will be made, the probability of similar occurrences will be reduced". Again, the meaning of 'appropriate corrections' can be made more precise using the evidence provided by the investigation.

A backward-looking reading of the causal statement would be: "If the training had been better the accident would not have occurred". But this is very difficult to prove. Human beings, even the better trained ones, have always made mistakes. Moreover, it would be necessary to specify what is meant by 'better training' and how it would have determined a difference in the pilot's behavior in the past. The result is that more backward-looking investigation is required to support the interpretation and there is no guarantee that a definite cause or factor will be found. Similar objections can be raised to other backward-looking interpretations, such as: "More than 90% of the pilot's erroneous action is due to lack of training."

Although several authors share the idea that root cause should be associated with correctability, not all have clearly accepted that it implies a shift towards the forward-looking perspective. For instance, Hokstad and Rausand (2008, 623) define root cause as the most basic cause that, if corrected, would prevent recurrence. When they specify the meaning of "the most basic cause" again they refer to a backward-looking process of digging deeper into the causal history until the *fundamental* cause has been identified. A similar stance appears also in definitions by Paradies and Busch (1988, 479) and Marquez (2007, 127).

Finlow-Bates' (1998) analysis of root cause clearly connects it to the forward-looking perspective. He distinguishes between a first stage in the investigation, aimed to find the "direct physical line of cause" (11) (causal history, according to our terminology) and a second stage where the results of the first one are used to answer the question about correction and prevention: "Which of the long-term solutions on offer is the most cost-effective?" (12). This means that the root cause will not be found by looking deeper and deeper in the causal chain but reasoning about the future with the benefit of the knowledge about the past. Or, as Finlow-Bates puts it, "it is the effectiveness of the solution that finally identifies a root cause" (12).

This aspect has been acknowledged also by Dekker (2005) in his analysis of the role of accident models in accident investigations. Dekker criticizes models based on linear sequences of events because, in his opinion, they cannot account for complex interactions and emergent phenomena that are common in complex socio-technical systems. Also, the application of these models is biased because of the knowledge the investigator has of the final outcome, the so-called hindsight bias (Fischhoff: 1975). However, after closer examination he admits that when investigators are looking for the root cause they are no longer looking

backwards but are abstracting from past experience and projecting to the future. Then, searching for root cause is "more about predicting the future than about explaining the past" (82).

In the next section we will combine the backward-looking notion of root cause and the forward-looking one in a unified notion, the U-turn.

5.5. Root cause as a U-turn

In order to combine the two notions, three distinctions have to be made. First, it has to be noted that the forward-looking notion is characterized by the emphasis on suggesting effective corrective actions whereas the backward-looking one focuses on finding out the precursors of a failure event. Corrective action is a form of prevention. In doing prevention the attention goes to avoiding unwanted (usually because unsafe) system states irrespective of whether those types of states have already occurred or not. Correction, instead, is a form of prevention in which the unwanted types of system states have already occurred in the past and the aim is to avoid recurrence. Correction, therefore, can ensue only once sound knowledge about the past has been provided. Even though correctability relies heavily on knowledge about the past, it would be inappropriate to think that it simply derives from it. As pointed out by Finlow-Bates (1998) and Dekker (2005) additional considerations are required because correcting the future does not aim at preventing the exact replica of what happened. Instead it encompasses types (or classes) of states of which the past occurrence is an instantiation, or a token.

The distinction between types and tokens is the second conceptual difference between the two approaches. Distinguishing between types and tokens is something we ordinarily do when dealing with artifacts, for instance my cell phone is an exemplar, or token, of a certain model, or type, of phone. Types, therefore, are abstract concepts lacking spatio-temporal location for which tokens are exemplifications or instantiations (Wetzel: 2011). Besides artifacts, the distinction also applies to the domain of processes and events. For example, the fuel rail fatigue process described by Gagg and Lewis (Gagg and Lewis: 2009) was a token of the type fatigue process. A failure investigation is expected to start from a specific fault state and to develop, as accurately as possible, the state sequence of a specific failure phenomenon diagram. These tokens all have precise spatio-temporal locations. Similarly, the causes and factors included in the causal

history are events or conditions with a precise (at least ideally) collocation in space and time. Of course, in backward-looking causal analysis it is also important to establish the type instantiated by a certain token because this will allow the determination of the applicable laws. Knowing that a certain state is an instantiation of hydrogen embrittlement, for instance, allows the investigators to apply scientific theories and engineering methods to deduce the type of states which are precursors to it.

The next step, then, is to show that the available evidence proves that the specific preceding state was a valid instantiation of the type. In contrast, the forward-looking reasoning starts at the token level, namely the fault state, then interprets it as a type and proceeds to reason at the type level about the failure phenomenon. The starting point is the already occurred failure phenomenon which is represented by means of the state sequence and the causal history. The phenomenon is then interpreted as a token representing the type which has to be prevented as efficiently as possible. Also the key states and causal factors involved are interpreted as types. Prevention will be achieved by modifying a type of state or causal factor that will affect the unfolding of the state sequence such that either the failure event or the fault state will not occur. For instance, a revised inspection procedure may be proposed in the expectation that future inspections (tokens of the revised procedure type) will reduce the probability that deviations from expected performance will reoccur, that is, corrective action is successful. This means that the inspection procedure is reckoned as the root cause of the failure phenomenon. Otherwise, the investigators may conclude that a more efficient way to reduce the probability of recurrence is to redesign some parts of the item more robustly. In this case the future initial states will be instantiating a different type of item which will not deviate from the target values. Then the conclusion will be that the item's design is viewed as the root cause.

The crucial point is that these kinds of assessments are performed on the basis of probabilistic expectations about future types of occurrences. And this introduces the third conceptual difference between the two approaches. The backward-looking approach is interested in proving that deterministic connections about a specific state and its precursor occurred. In the forward-looking reasoning, instead, modified types of system's states are expected to result, with a certain probability, in improved types of system states. It is crucial for the

investigators to find out which modifications will provide the highest likelihood of system improvement.

Our proposal is that once these conceptual differences are acknowledged, it is easier to see the reasons of controversy between the backward-looking and forward-looking notions. The two approaches, however, instead of being incompatible are complementary. The backward-looking approach is instrumental in providing knowledge about the causal factors included in the causal history. It is however inadequate for answering questions about correction. The forward-looking approach, on the other hand, emphasizes the question of correction but needs, of course, to start from knowledge of a token that has already occurred before it is possible to generalize about future types of occurrences. The root cause of a failure phenomenon, then, is that element of factors and causes that, if corrected, is the most likely one to prevent failure phenomena similar to the one under investigation from happening again. Consequently, it is located at the U-turn between backward-looking causal analysis and the forward-looking process (Figure 5.4). In the rest of this section we will consider two failure investigations in which both the backward-looking and the forward-looking legs of the U-turn are present and a correctable factor located at the U-turn is eventually suggested as the root cause.

The first example again uses the case history presented in Gagg and Lewis (2009) paper, discussed above in Section 5.3.2 (fatigue failure in a car's fuel rail). It has to be specified that Gagg and Lewis do not use the term 'root cause' in their paper. However, relevant to our purpose is that after analyzing the causal history of the car crash (backward-looking causal analysis) they add some considerations about prevention which fall naturally within the forward-looking perspective. The crucial observation based on the investigation findings is the surprisingly short time span the crack took to propagate through the fuel pipe wall after the misalignment initiated the fatigue mechanism. As a short term corrective action, service centers should be alerted about the necessity to ensure axial alignment. However, this fix would have a rather small impact on prevention. Instead, Gagg and Lewis claim the manufacturer should take action to prevent or limit reoccurrence and should address the very reasons for the fuel rail's vulnerability to fatigue. Therefore the design of pipe restraints should be improved with the aim to minimize fluctuating stresses at the injector end of the pipe. Although Gagg and Lewis do not mention the term 'root cause' their analysis indicates that the fuel rail design, which is responsible for the predispo-

sition to fatigue failure, is the correctable factor located at the U-turn. Inadequate design is the root cause.

The second example is based on the NTSB (1990) report of an infamous aviation accident. On 19th of July 1989, United Airlines flight 232 crashed during an emergency landing in Sioux City, Iowa. The aircraft, a McDonnell Douglas DC-10, had lost hydraulic fluid in all three redundant hydraulic systems after an uncontained failure of the engine located at the base of the tail fin. The engine's first-stage fan disk, a 168 kg 80 cm diameter titanium alloy component, broke up while the plane was at cruising speed and high-energy fragments were hurled with enough force to puncture the hydraulic lines running within the horizontal empennage.

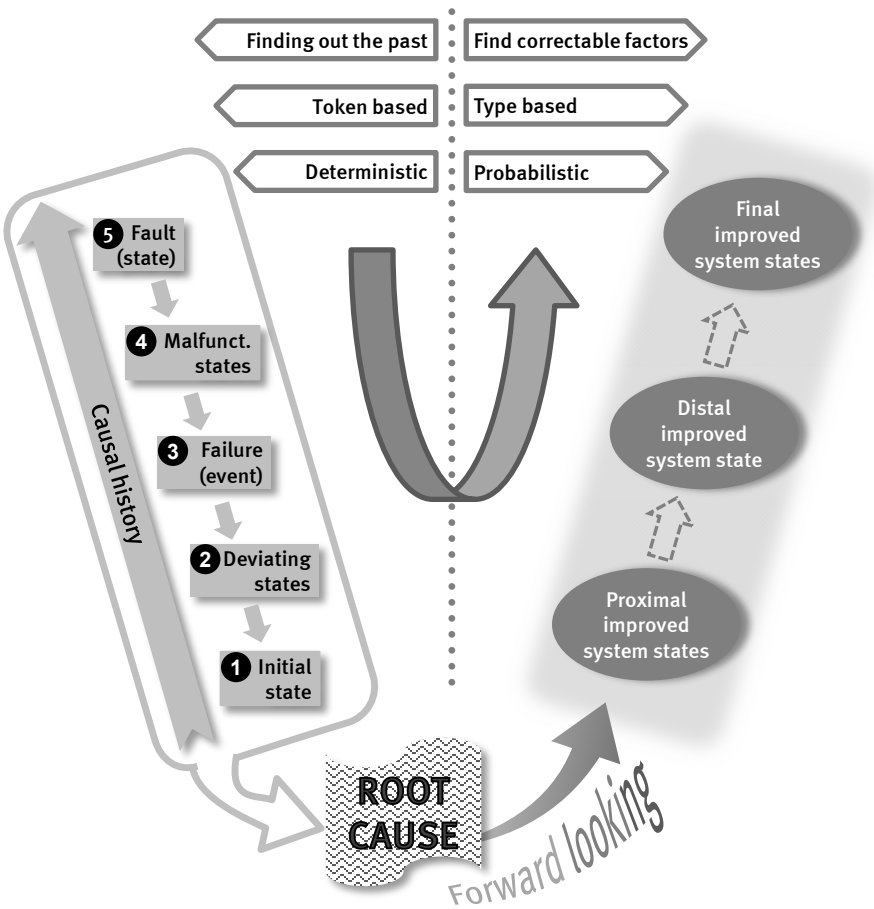

Figure 5.4: Root cause as a U-turn

Examination of the fracture surfaces of the fan disk disclosed that the primary fracture had propagated from a fatigue region on the inside diameter of the bore. Metallurgical examination revealed that the fatigue crack originated in a small cavity that formed during manufacture (Job: 1996). The disk was installed in a new engine in January 1972 and had been in operation for 17 years during which it was routinely overhauled. Its last inspection was in February 1988. The defect remained undetected through all the inspections performed during and after manufacture. However, the investigators found traces of penetrating fluorescent dye on the fracture surfaces. On top of that, calculations about crack growth speed showed that the crack was of detectable size at the time of last inspection. The NTSB investigation team concluded that "the probable cause of the accident was the inadequate consideration given to human factor limitations in the inspection and quality control procedures used by United Airlines' engine overhaul facility" (NTSB: 1990, 102). This kind of causal statement and the attached motivations make it clear that investigators were pointing at a correctable factor with positive repercussions for prevention. They were concerned about preventing a larger type of failures, not just the ones affecting engine disks. Thus they highlighted the role of human factor limitations. For this reason, although the report does not mention the term, we think that the causal statement is an example of root cause as U-turn. This becomes even clearer when the final causal statement is compared with the dissenting statement filed by a member of the investigative team. The dissenting statement lists three probable causes: the metallurgically defective fan disk, the inability of United Airlines to detect the crack, and the failure of the McDonnell Douglas design of the airframe to account for the possibility of extensive damage following a failure of the tail engine (NTSB: 1990, 108–109). Of course the dissenting statement is correct in saying these were contributing factors of the specific accident, but compared to the probable cause statement it shows a backward-looking approach.

Our analysis shows that a sensible meaning can be assigned to the notion of root cause by combining two approaches to define the notion, which, when taken in isolation, often provoke controversy. This is the case, for instance, in the debate between Walker (2009) and MacIntosh (2010) which was summarized at the end of Section 5.2. Walker is right when he argues that looking backwards is not conducive to a single unambiguous beginning of the chain of events that deserves to be called the root cause. However, his opposition to the notion seems to derive from neglecting the forward-looking concerns about correctability that

inspire MacIntosh's defense of the causal statement. At the same time, MacIntosh's position is not always clear in taking distance from interpretations of the causal history according to which the correctable factor identified as the root cause is also the 'devil at the bottom' that allows apportioning blame to the most responsible actor.

5.6. Conclusion

Failure analysis is considered the process of finding the causes and factors of undesired loss of functionality. The notion of causality, then, plays a central role in the investigative process. Yet, an important notion like root cause is still in dispute. We have argued that two main approaches to root cause are represented in the literature. One derives from what we call the backward-looking causal analysis process. It occurs in the investigative process that begins when an item is in the fault state and aims at finding the state sequence and the causal history of the failure phenomenon which ends in that very fault state. According to this backward-looking process the root cause is a factor that can be found by digging deeper into the causal history and that is more prominent than other factors either because it is at the absolute beginning of the chain of events or because it is a necessary factor.

The second approach stems from the forward-looking effects analysis. Its aim is to understand a failure phenomenon in order to discover a corrective factor that will prevent recurrence. Authors defending this approach have emphasized that finding corrective factors cannot be done by insisting on the backward-looking process alone; additional considerations about possible states of the system in the future are needed.

We have discussed the conceptual differences between the two approaches: the backward-looking one is deterministic and token-based; the forward-looking is probabilistic and type-based. Once these differences are clarified, it can be seen that the two approaches can be combined and one sensible meaning can be assigned to the notion of root cause. The root cause of a failure phenomenon, then, is that element of the factors and causes that is, if corrected, the most likely to prevent failure phenomena similar to the one under investigation from happening again. Consequently, it is located at the U-turn between the backward-looking causal analysis and the forward-looking process.

More research should be done into the instrumental value of the state-sequence diagram. First, it should be investigated whether similarities between state-sequence diagrams of different failure phenomena lead to a systematic classification of types of failure. Second, it should be found out whether a breakdown of the state-sequence diagram of a compound item into the state-sequence diagrams of its parts helps in finding the most convincing causal history of the failure phenomenon.[11]

[11] The authors are grateful for being given the opportunity of presenting the ideas developed in the present paper at the ICEFA IV conference, and for the considered responses they received. While working on this paper Luca Del Frate took part in the Marie Curie "EuJoint" Project (IRSES 247503).

6 Learning from Failure: Not so Paradoxical After All[12]

Abstract

Failures are ubiquitous and cover a spectrum of phenomena as broad as engineering itself. Engineers invest considerable resources both in the attempt of preventing failures and in the process of learning from them. Henry Petroski and many engineers with him believe that in engineering *more is learned from failures than from successes*. In this paper, I investigate this alleged paradox. After surveying the engineering literature in search of shared meanings of failure, success, and learning, I show that Petroski's arguments conceal two different hypotheses concerning respectively a *specific* mode of learning and a *generic* mode. While the former hypothesis is empirically testable, the latter rests on questionable assumptions and conflicts with currently held views of technical change.

6.1. Introduction

The second edition of the well-known engineering textbook *Corrosion for Science and Engineering* (1995) begins with a section eloquently titled *The Lessons of History*. In 1761, the Royal Navy decided to cover the hull of the frigate HMS *Alarm* with a thin copper sheathing. The purpose of the operation was to investigate experimentally a solution to two serious problems affecting wooden ships: structural damage caused by wood-boring shipworms, and increased drag due to barnacles and vegetation growing on the hull. The experiment was based on already proven toxic properties of copper which were expected to curtail both negative phenomena. The copper-sheathed frigate set sail for the West Indies were it was deployed for two years; then, in 1763, it was docked and inspected. Although the copper sheathing did succeed in providing protection from the two known problems, a third unexpected one was found: the sheathing itself had

[12] A version of this chapter will be submitted to the journal *Technology and Culture*.

become detached in many places because of corrosion of the iron nails used to fasten the copper plates to the hull.

At first, inspectors were at a loss to explain the phenomenon. Closer inspection, however, revealed a crucial hint: those iron nails which did not corrode where insulated from the copper because of pieces of brown paper trapped under the nail heads. Those pieces of paper were the remnants of the wrapping into which the copper plates were originally delivered to the yard and which was not removed prior to installation. So, through a lucky circumstance, investigators were able to draw a valuable lesson from the HMS *Alarm* copper-sheathing failure, namely that "iron should not be allowed direct contact with copper in a sea-water environment if severe corrosion is to be avoided" (Trethewey and Chamberlain: 1995, 1). Nowadays, the process that led to the sheathing failure is known as galvanic corrosion or, more precisely, bimetallic corrosion.

After such a promising beginning, one would expect Trethewey and Chamberlain to continue their narrative by illustrating how, from such fortuitous origins, knowledge about corrosion and the proper ways of preventing it were disseminated, thereby leading to the design of better ships and other structures alike. As a matter of fact, they take a different direction and recount numerous episodes where neglecting the lesson of HMS *Alarm* resulted in serious galvanic corrosion and eventually in failure. Ironically enough, some of these episodes involve ships or other equipment owned by the Royal Navy, the same organization that first encountered galvanic corrosion and, allegedly, learned how to deal with it. Especially striking was the failure of a sea-water evaporator in a submarine which happened in 1962, a sort of reenactment of the bicentennial antecedent. The copper end-plate of the evaporator had been fixed in place by uninsulated steel bolts that, like the iron nails of HMS *Alarm*, effectively dissolved by galvanic corrosion causing the separation of the end-plate itself.

In the rest of the chapter, Trethewey and Chamberlain offer a brief summary of economic and social costs of failures due to various kinds of corrosion processes (which are discussed in detail in the rest of the book) like lost production, loss of quality, and pollution. At this point, the reader of the book will not be surprised in learning that since the initial failure of HMS *Alarm* "bimetallic corrosion and many other forms of corrosion have continued to cause service failures, despite their apparently well-publicized effects". Indeed, after reading of so many repeated episodes of failure, one might start wondering why so little is learned from experience. Yet, Trethewey and Chamberlain decide to conclude

the chapter on a positive note and to draw attention to the many cases where engineers have been able to achieve successful control of corrosion: bridges and buildings do stay up most of the times, air travel has never been so safe, and today's cars are more reliable than ever. Hence, in spite of many examples pointing in the opposite direction, they claim that engineers do "*learn* more from failures. In order to build upon this success into the new millennium, we must continually mull over the reasons for past failures" (19, emphasis in the original).

Such a conclusion may appear somewhat paradoxical given the abundance of occasions where lessons from failure went egregiously unheard. Nevertheless, Trethewey and Chamberlain are by no means an exception and many authors from different disciplines have advocated the fundamental role of learning from engineering failures. Philosopher Gary Gutting (1984, 63), for instance, claims that "The mere fact that a system fails to perform properly in certain circumstances in itself constitutes a piece of knowledge essential to the technological enterprise". Aeronautical engineer and historian Walter Vincenti quotes Gutting's claim approvingly in his renowned book on engineering knowledge (1990). Civil engineer and failure analyst Kenneth L. Carper contends that "Much of the knowledge used to design, construct, manufacture, and operate engineered facilities and products has been obtained through learning from failures" (Carper: 2001). Structural engineer Sir Alfred Pugsley, who is widely credited as the father of modern structural safety, thinks that "All safety rules grow out of, and are periodically amended as a result of, accidents to structures" (Pugsley: 1966, 120). More recently, a panel of four systems engineers was convened to gain a better understanding of failure's role in systems engineering and appropriate reactions to failure. The panel concluded unanimously that "It is typically recognized that failure is a common occurrence and that future success is often a consequence of our reaction to failure" (Slegers et al.: 2012, 75).

Possibly, the most articulate and more prolific among those who have advocated the primacy of failure in engineering is Henry Petroski, a civil engineer turned historian who has published sixteen books and more than fifty papers, all of them related to failure in one way or another. Even though his take on the role of failure is fairly similar to that of others mentioned above, what makes his publications notable, besides the widely acknowledged literary qualities, is the vast range of arguments and case stories from many branches of engineering that he has been able to recruit in support of his main message.

And this is a message that Petroski is fond to summarize by means of what he calls *the paradox of engineering*:

> It is an apparent paradox of science and engineering that more is learned from failures than from successes. (Petroski: 2001, 10, emphasis added)

By describing it as an *apparent* paradox, Petroski implicitly admits that his argument does not result in a logical contradiction as many of the most renowned philosophical paradoxes do.[13] Blockley and Henderson (1980, 726) share similar ideas and are equally puzzled by the counterintuitive role of failure which presents engineers with "a strange antithesis", namely:

> It is the success of engineering which holds back the growth of engineering knowledge, and its failures which provide the seeds for its future development.

Whether or not it is truly paradoxical or just counterintuitive, the claim that more is learned from failures than from successes has been echoed by many engineers who believe it represents a crucial characteristic of how engineering knowledge and practice develops. This fact alone would be enough to justify a closer look at the claim and at the evidence brought in its support. On top of that, a further reason is that, upon a moment of reflection, the claim appears to be nearly as perplexing as it is paradoxical. Even though the concepts constituting its main building blocks appear familiar at first, one might reasonably wonder: what is meant by "learning"? And what does count as "failure"? And what are "successes"?

In this paper I analyze the engineering paradox, discuss its main assumptions, and scrutinize the arguments brought in its support. The idea is to consider Petroski's paradox as a hypothesis about learning in engineering and to clarify if and how available evidence could corroborate it. To emphasize this change in perspective, from now on, it will be called the *learning hypothesis* and

---

[13] Whether or not a contradiction is an integral part of veritable paradoxes is disputed. In his recent book on the subject, Łukowski (2011, 1) contends that a paradox is "a thought construction, which leads to an unexpected contradiction". Cantini (2012), on the other hand, characterizes paradoxes as statements "claiming something which goes beyond (or even against) 'common opinion' (what is usually believed or held)" and notices that "Most paradoxes — but not all — involve contradictions".

will be expressed as follows: *In engineering, more is learned from failures than from successes.*

I start my analysis by summarizing two cases of learning that recur in Petroski's writings where they are presented as paradigmatic examples that more is learned from failures than from successes (Section 6.2). In Section 6.3, I discuss the notions of failure and success and how they are interpreted in engineering. In Section 6.4, I complete the preparatory work by dealing with the ambiguities of learning in engineering. Then, in Section 6.5, the conceptual tools prepared in the previous sections are utilized to dissect the learning hypothesis and I show that, even though Petroski contends he is dealing with just one paradox, he is actually advancing two different hypotheses about learning in engineering. Borrowing from (Vincenti: 1994), I call them the *specific-learning* and the *generic-learning hypothesis*. The conclusions are that, while the specific hypothesis has strong conceptual support and is conducive to empirical testing (in fact, a recent study claims to have found empirical evidence supporting it), the generic hypothesis rests on questionable assumptions and conflicts with currently held models of technical change.

6.2. Paradigms of learning: Roebling and Co.

In this section I will summarize two paradigmatic historical cases which, according to Petroski, clearly show that more is learned from failure than from success. Both cases are based on suspension bridge technology and appear multiple times in Petroski's writings. The narrative given here is modeled on chapter 8 and chapter 9 of Petroski's (1994) book *Design Paradigms* and supplemented with additional contributions from (Buonopane and Billington: 1993; Scott: 2001; Kawada: 2010). Main components and characteristic dimensions of suspension bridges are summarized in Figure 6.1.

6.2.1. Charles Ellet vs. John Roebling: Same engineering goal, different approaches

According to (Kawada: 2010), the first modern suspension bridge was the Jacob's Creek Bridge built 1801 in Pennsylvania by James Finley (1762-1828). Although primitive for today standards, Finley's bridge was a success and spurred interest over suspension bridges.

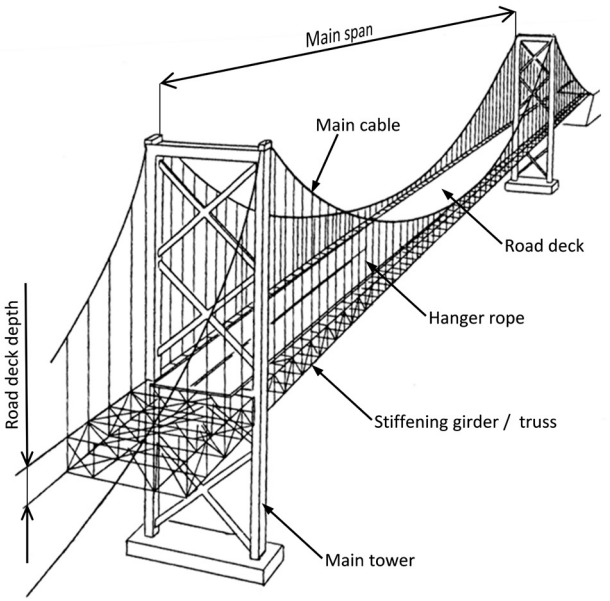

(a)

(b)

Figure 6.1: (a) Main components and characteristic dimensions of suspension bridges. Based on (Harazaki et al.: 2000) (b) A view of John Roebling's Brooklyn Bridge in New York showing his trademark diagonal cable stays radiating from the top of the tower. See main text for details.

Over the following decade, Finley designed many other suspension bridges and his work became known also in Britain and, from there, in continental Europe. After an initial tumultuous expansion, problems begun to emerge with many suspension bridges being damaged by storms or overloaded by snow. During the 1830s until the 1840s relatively few suspension bridges were built in America and in Britain as well. Nevertheless, suspension bridges remained popular in continental Europe, particularly in France. In 1834 the Grand Pont Suspendu in Fribourg was inaugurated and with a main span of 273 it became the longest in the world.

Charles Ellet (1810-1862) was born in Pennsylvania and in 1830 moved to France to study bridge engineering at the École Polytechnique. Notably, in 1831 he visited the construction site of the Grand Pont Suspendu in Fribourg. In 1832 Ellet returned to America and by that time he "had become a complete adherent of the French-style wire suspension bridge" (Kawada: 2010, 72). His first realization was the 109 m Schuylkill River Bridge which opened in 1842. The success of this bridge revamped the reputation of suspension bridges in America and boosted Ellet's fame who was awarded several major construction contracts. In 1847 Ellet began the construction of the world record breaking Wheeling Suspension Bridge, whose main span was 310 m long. Again, Ellet designed his bridge along the guidelines of the French school. The Wheeling Bridge was completed in 1849 and performed well for about five years until in May 1854 the deck was destroyed in a wind storm.

John Roebling (1806-1869) was born in Germany, where he studied civil engineering, and moved to America in 1831. Initially, he designed suspended aqueducts and later became interested in suspension bridges for road and railway transportation. In 1851 Roebling began working at a railway suspension bridge spanning the Niagara Gorge. Roebling's approach to this project was remarkably different than Ellet's. Roebling was severely critical of European suspension bridge engineers for what he considered a cavalier attitude towards structural stiffness. He studied available reports of bridges being damaged in the wind and became convinced that very flexible structures like Ellet's Wheeling Bridge were at serious risk of collapse. Therefore, he set out to make his Niagara Falls Suspension Bridge (main span 251 m) as stiff as possible without making it excessively heavy. His design solution to this difficult trade-off consisted in what he called four appliances: heavy road deck, stiffening truss, girders, and diagonal cable stays.

While Ellet's bridge was severely damaged by a storm in 1854, Roebling's four appliances conceived from the study of failures equipped his bridge with enough stiffness to withstand the loads imposed by stormy weather as well as those from railway traffic.

6.2.2. From Brooklyn to Tacoma Narrows and beyond

After the success of Niagara Falls Suspension Bridge, Roebling became the most authoritative American suspension bridge engineer. In 1867, he started design work on the Brooklyn Suspension Bridge spanning the East River in New York. With a main span of 486 m, the bridge was the longest suspension bridge in the world from its opening in 1883 until 1903.

The Brooklyn Bridge features all four appliances (heavy deck, stiffening truss, girders, and cable stays) that Roebling developed in designing the Niagara Falls Suspension Bridge. These features became standard in suspension bridge design over the following decades.

Things began to change only in the first decade of the twentieth century as a result of theoretical advancements as well as the introduction of new structural materials. On the theoretical side, a major departure from the established tradition was the development to the so-called deflection theory. One of its authors was Leon Moisseiff (1872-1943) who utilized the new theory in designing the Manhattan Bridge. Completed in 1912, the bridge is located a few hundred meters upstream of the Brooklyn Bridge. Although with a main span of 451 m the Manhattan Bridge was smaller than its venerable neighbor, it was rightly regarded as a groundbreaking accomplishment. The introduction of the deflection theory allowed Moisseiff to design a much lighter bridge with considerable savings in structural materials and construction costs. Moreover, the bridge featured the first major use of nickel steel (which is 50% stronger than carbon steel) in a suspension bridge. The Manhattan Bridge "launched Moisseiff on a long and distinguished career as the most prolific suspension bridge engineer of his generation; he would contribute to virtually every major American suspension span built in the first 4 decades of the 20th century" (Scott: 2001, 17).

One of implications of the deflection theory is that it allows utilizing the main suspension cables as contributors to the structure's overall stiffness thereby making heavy stiffening trusses and stay cables (i.e., Roebling's appli-

ances) redundant. Moisseiff and Othman Ammann (1879-1965) took full advantage of this characteristic in the design of the George Washington Bridge in New York. The bridge was completed in 1931 and with a main span of 1067 m it was almost twice as long as the current record holder. Besides its size, the bridge was remarkable for its extreme slenderness made possible by the deflection theory. Its road deck is just three meters deep which gives a depth to span ratio of 1 to 350. As a comparison, with its heavy stiffening truss the Brooklyn Bridge's deck is 5.3 m deep resulting in a depth to span ratio of 1 to 90. As noted by (Buonopane and Billington: 1993, 972), in the case of the George Washington Bridge a stiffening truss similar to the one utilized in the Brooklyn Bridge "would have been virtually unbuildable".

The trend towards slender bridges culminated in the construction of the Tacoma Narrows Bridge. When completed in July 1940, its 853 m main span made it the third longest in the world. Its road deck was even shallower than George Washington Bridge resulting in the same depth to span ratio of 1 to 350.

Although designed to withstand winds up to 160 kmh, the bridge was destroyed in November 1940 by winds blowing at just 68 kmh. The engineering community was shocked. Apparently no one could understand the mechanism that induced the bridge to oscillate so violently given the relatively modest speed of the wind. The George Washington Bridge was as slender and flexible as Tacoma, yet it never showed signs of instability in the wind. The spectacular collapse of Tacoma clearly showed that an unexpected interaction between aerodynamic loads and the bridge response to them initiated a positive feed-back loop where ever increasing stresses where applied to the structure until it came apart. The efforts to understand how the collapsed happened resulted in a period of extensive research on the aeroelastic behavior of bridges. After Tacoma, it became common to perform wind tunnel tests of new bridges.

A better understanding of aeroelastic phenomena allowed engineers to anticipate the behavior of bridges under wind loads and devise appropriate design solutions. A new generation of bridges appeared where the road deck was made of a streamlined steel box girder shaped like an airfoil. The first major suspension bridge taking full advantage of the new understanding of aeroelastic phenomena was the Severn Bridge in Britain which was inaugurated in 1966. A main span of 988 m and a road deck just 3 m thick imply that the Severn's Bridge depth to span ratio of 1 to 324 is comparable to the Tacoma Narrows' figure. Yet, the Severn Bridge has endured successfully much stronger winds

than those which brought down Tacoma Narrows Bridge. For many years, and especially after Tacoma, the depth to span ratio was regarded by bridge engineers as a vital design parameter. However, "the principles underlying the Severn's Bridge design were so different as to render this traditional measure almost irrelevant" (Scott: 2001, 151).

6.3. Defining failures and successes in engineering

Admittedly, the cases of failure mentioned so far are all examples of structural damage (HMS *Alarm*) or, even worse, structural collapse (e.g., Wheeling Suspension Bridge, Tacoma Narrows). The notion of failure in engineering has a broader reach, though. Because of his academic background and because the discipline has a long history he can tap into, Petroski tends to favor examples from civil engineering which, quite often, involve buildings or other structures being physically damaged. However, Petroski's interest is by no means limited to civil engineering and to collapses. In fact, he is acutely aware that a narrow view on failure inevitably limits potential learning opportunities. To clarify the point he takes the hypothetical example of a skyscraper which, although structurally sound (i.e., it is in no danger of collapse), should be regarded as a failure because it oscillates noticeably when winds blow from a particular direction thereby making the occupants of the upper floors seasick (Petroski: 2011, 104).

Indeed, one of the problems with the notion of failure is to establish its reach. As shown in (Del Frate et al.: 2011), definitions of failure abound in the engineering literature and vary from being tightly coupled to a specific field to being very broad in scope. Even though proposals from authoritative standardization institutions like the *International Electrotechnical Commission* have gained some prominence, special purpose definitions continue to emerge and there is hardly any indication of widespread consensus (Tam and Gordon: 2009).[14]

Of the many alternatives available in the literature, in his book *Success through Failure* (2006) Petroski endorses a definition originally proposed in

[14] The failure definition stipulated by the *International Electrotechnical Commission* can be found, together with a series of other failure-related concepts, in the *International Electrotechnical Vocabulary* and reads as follows: "Failure: Termination of the ability of an item to perform a required function" (IEC 60050(191): 1990).

(Leonards: 1982, 108) and later also adopted by the Technical Council on Forensic Engineering of the US Society of Civil Engineers:

> Failure is an unacceptable difference between expected and observed performance.

Admittedly, by utilizing vague notions such as "unacceptable difference" and "expected performance", Leonard's definition leaves ample room for interpretation. Despite its vagueness, supporters of Leonard's definition insist in high-highlighting its merits. Carper (2001) praises Leonard's definition because, differently from others that tend to equate failures with sudden events like structural collapses, it takes a different approach and extends the concept over a different class of phenomena which, though less dramatic than collapses, are nevertheless commonly regarded as failures. In the building industry, for instance, a roof leaking or a façade degrading faster than specified are routinely treated as failures. Yet, such kinds of phenomena are clearly different from typical examples of failures like ruptures and collapses. These differences can be captured by distinguishing two concepts of failure, *event-oriented* and *goal-oriented* concepts. Event-oriented concepts are informed by the idea of failures as occurrences having a rather precise location in space and time. The occurrence may consist in the interruption of a required function, as in the definition proposed by Birolini (2007, 3):

> A failure occurs when the item stops performing its required function.

Or can be characterized in physical terms as in (Collins: 1993, 6):

> Mechanical failure might be defined as any change in the size, shape, or material properties of a structure, machine, or machine part that renders it incapable of satisfactorily performing its intended function.

In its various forms, the event-oriented concept presupposes a list of technical specifications that is to say, a set of product attributes that can be measured and that are associated with predefined thresholds.[15] A failure ensues when those

[15] The engineering literature typically distinguishes between *product specifications* and *technical specifications*. Product specifications or requirements are sets of desired attributes expressed in quantitative manner of a product being developed. In short, they are a list of desiderata. Technical specifications document the actual attributes of finished products and may differ to

thresholds are trespassed. Rausand and Øien (1996) illustrate the concept by considering a shut-down valve, a kind of safety valve which in normal conditions is open and that closes the passage of fluid in emergency situations. To be effective as safety devices, shut-down valves need to close rapidly. On the other hand, they should not deploy abruptly because hazardous shock waves may ensue as a consequence. As a rule of thumb, from the moment the emergency signal is received, a shut-down valve should take 6 to 14 seconds to close the passage of fluid. If observed performance falls outside these boundaries (by being either too fast or too slow), the valve is said to have failed. The failure event itself can be the final outcome of a gradual process of corrosion or fatigue; still *failure* refers only to the event in which the critical parameter (e.g., valve closing time) trespasses the acceptable limits.[16] The resulting concept of failure, then, is strictly binary: either the item performs within the limits or not.

Goal-oriented concepts, of which Leonard's definition is an example, take a different approach which results in a significant shift in the ontology of failures: from being classified as event-like entities to the state-like category. As shown by Leonard's definition, failures are considered to be states or conditions whose most notable characteristic is a gap or inability with respect to an ideal state. Furthermore, differently from event-oriented concepts which presuppose precisely defined performance parameters, goal-oriented concepts tend to be more qualitative and based on high-level properties of products such as their intended goals.

Because of its vague terminology, Leonard's definition may not be the most suitable to express this approach. The engineering literature, however, offers more perspicuous definitions of failure inspired by a goal-oriented approach. One example is the definition proposed in Del Frate et al (2011, 271)

> Failure is the inability of an engineering process, product, service or system to meet the design team's goals for which it has been developed.

---

some extent from product specifications. With respect to failure, the thresholds that matter are those found in technical specifications lists.

[16] A more detailed analysis of the event-oriented concepts of failure can be found in (Del Frate: 2012) where it is decomposed in three concepts: *function-based*, *specification-based*, and *material-based failure*.

One reason this definition improves over Leonard's is that it replaces the vague term *expectations* with the more precise term *goals*. The latter has been the subject of in-depth analysis by engineering design theorists and, even though not entirely consensual, it plays a prominent role in many models of engineering design and product development activities. In this paper I will use as a reference the work by Dym and Little (2008). As a starter, Dym and Little distinguish between, on the one hand, the goals of a design project and, on the other hand, the goals a product is expected to achieve. The latter are those that matter with respect to failure and success. These goals provide a translation of customer-needs into qualitative statements of expected product properties and capabilities. The goal-setting process begins at a very early stage of the development process. Ideally, it precedes and provides guidance to the process of establishing product requirements in which qualitative and approximate statements are translated into quantitative and precise statements. In practice, similarly to other processes involved in product development, they can partially overlap and the final fixation of goals may occur when the development process is already in full swing.

Dym and Little emphasize that an important aspect of the goal-setting process consist in organizing multiple goals into a hierarchical tree structure in which a top-level goal is decomposed into sub-goals of differing levels of importance and scope. While goals located in the lower branches may be translated rather easily into product requirements, higher-level goals are typically more problematic and might retain their qualitative nature. Nevertheless, metrics can be defined that provide an indication of the product's ability to meet those goals. Dym and Little consider the example of a device for which *being durable* is a high-level goal. Differently from a goal like say, *being low-weight*, which can translate directly into grams or kilograms (a low-weight phone is clearly lighter than a low-weight city bus), durability does not have a natural unit of measure. In the case of electronic appliances like mobile phones, for example, drop tests are often used as one of the means to assess durability: test specimens are dropped multiple times until cracks begin to emerge. Drop test results can be seen as proxies of the desired property. Usually, however, multiple proxies are needed resulting in qualitative scores over different metrics. This implies that, while in event-oriented concepts failure is seen as binary, goal-oriented concepts allow for gradations of failure depending on how close the product comes to meet its goals in full.

Higher-level goals are also important because they somehow define the identity of products, not just what they do but what they actually are. Products that share largely identical higher-level goals can be seen as members of the same kind and the same metrics can be utilized to assess them. This aspect of goals will play a crucial role in spelling out the details of the learning hypothesis in Section 6.5.

The definition of failure given by Del Frate et al (2011) jointly with Dym and Little's conceptualization of goals has a further favorable quality when compared to Leonard's: it leads naturally to a definition of success. So far, this section has dealt almost exclusively with the concept of failure. A similar bias can be found in Petroski's publications and in most of the engineering literature on learning. In fact, definitions of success are hard to come by as though common sense intuitions would suffice. Thus, this section concludes by turning Del Frate's et al definition of failure into the definition of success that will be used in the rest of this paper:

> Success is the ability of an engineering process, product, service or system to meet the design team's goals for which it has been developed.

6.4. Ambiguities of learning in engineering

Next to failure, the concept of learning should be clarified for it is as much elusive. A traditional way of conceptualizing learning is by means of an *acquisition metaphor*. Sfard (1998, 5) has noted that:

> Since the dawn of civilization, human learning is conceived of as an acquisition of something. Indeed, the *Collins English Dictionary* defines learning as "the act of gaining knowledge". [...]. Concepts are to be understood as basic units of knowledge that can be accumulated [analogously with] the activity of accumulating material goods.

This view of learning implies a triadic relation between an epistemic agent (e.g., a chemistry student), a source of knowledge (e.g., a teacher of chemistry) and a unit of knowledge (e.g., the chemical formula of water). It is assumed that knowledge can be stored either in human minds or external devices like books from where it can be transferred into a recipient. In this conceptualization of learning, therefore, the acquisition process leaves the units of knowledge unchanged and just modifies their distribution across a population of learners.

Thus, seen at the population level, learning consists in the process by which precompiled knowledge diffuses within a society. This is the conceptualization Mokyr subscribes to in his book *The Gifts of Athena* (2002, 4–5) where he stipulates that "learning or diffusion would be defined as the transmission of existing knowledge from one individual or device to another".

It is immediately evident that this way of conceptualizing learning does not fit with the case stories discussed above. Consider again the case of HMS *Alarm*: the knowledge acquired through the failure investigation was not retrieved from a stockpile of precompiled knowledge. Instead, it was *generated* by the investigators during the process of explaining what happened to the iron nails utilized to fasten the copper sheathing to the ship's hull. Similarly, the collapse of Tacoma Narrows Bridge originated a spate of studies on the dynamic responses of structures to wind loads which reshaped drastically existing knowledge about suspension bridges.

The idea of generation, however, does not necessarily mean generation of *new* knowledge in the sense of unprecedented knowledge previously unavailable within the community. As noticed by (Jovanovic and Rob: 1989, 570) "knowledge is obviously unevenly distributed in any economy" and some engineers may have to learn from experience (including the experience of failure) what others learned from textbooks. The uneven distribution results also from engineers intentionally preventing some of their knowledge falling in the hands of the competition. In those cases, engineers may resort to the practice of reverse engineering which can be seen as an attempt to re-generate the knowledge originally used by the makers of a specific product.

A survey of the engineering literature shows that the notion of learning as used by engineers encompasses both the process of generation of knowledge as well as its diffusion. Notably, the ambiguity can be detected in Walter Vincenti's book *What Engineers Know* (1990), a classical source philosophers of technology frequently tap into on matters of engineering knowledge. Vincenti classifies engineering knowledge into six *categories of knowledge*: fundamental design concepts, criteria and specifications, theoretical tools, quantitative data, practical considerations, and design instrumentalities. Furthermore, he identifies seven *knowledge-generating activities* which contribute to one or more categories: transfer from science, invention, theoretical engineering research, experimental engineering research, design practice, production, and direct trial (which includes operation of the engineering artifact). In describing how these activities

work and their effects on engineering knowledge, Vincenti employs two terms: *growth of knowledge* and *learning*.[17] The former refers to changes in knowledge that occur at the level of the engineering community and subsumes generation of new knowledge (e.g., addition to the existing stock of precompiled knowledge) and its diffusion through the community.

The term *learning*, on the other hand, performs double duty and can refer both to collective level phenomena – e.g., "The learning process over the intervening years provides an example of how an *engineering community* translates an ill-defined problem..." (51, emphasis added) – and to learning by individuals – e.g., "Such [practical] considerations are mostly learned on the job rather than in school or from books; they tend to be carried around, sometimes more or less unconsciously, in designers' minds" (217).

6.4.1. Multiple levels of learning

The fact that learning can occur at multiple levels (together with the challenges this fact poses with respect to promoting effective learning) is a vital topic of research in areas like Safety science and Organizational theory. Recently, a special issue of the journal *Safety Science* aptly titled *The gift of failure:* Carroll and Fahlbruch (2011) gathered thirteen papers from an international workshop on *New approaches to analyzing and learning from events and near-misses*. One of the key areas of investigation was, indeed, that of multilevel learning. Contributions in the special issues as well as elsewhere in the literature have proposed different models. The model discussed in (Jacobsson et al.: 2011), for instance, identifies four levels: local level (e.g., the individual operator within a process plant), process unit level, site level (e.g., the process plant), and higher level (e.g., the whole corporation). In (Hovden et al.: 2011) the levels of learning are spelled out thus: individual, company management, sector/trade, authorities, and technology.

Although the details vary, the overall picture is clear in indicating that an adequate analysis of learning presupposes a hierarchy of *learners* whose base is constituted by individual learners. Given the scope of this paper, there is no need to delve into the details of the models discussed in the literature. Therefore, I

[17] Other expressions that recur only sparingly are: *cognitive growth*, and *cognitive change*.

will distinguish only two levels of learning, i.e., *individual* and *collective*, which suffice for the aim of clarifying the learning hypothesis.

Individual learning can be conceived as the process by which an individual acquires new knowledge as a result of *direct* or *indirect* experience of either failure or success. According to Petroski's historical reconstruction, the latter is how Roebling came to develop the four appliances that made his bridges successful. He formulated a criterion as to what counts as failure, namely a suspension bridge being damaged by the wind. Then, he gathered trustworthy information about the design of those bridges and how failures unfolded. From these studies he concluded that lack of stiffness was the culprit and he made estimations on the appropriate amount of stiffness needed to prevent reoccurrence. Charles Ellet, on the other hand, epitomizes learning from success: he went to France to study the French approach to suspension bridge design and he acquired new design concepts and construction procedures which he reenacted in his own designs.

As noted by Mokyr "collective knowledge as a concept raises serious aggregation issues: how do we go from individual knowledge to collective knowledge [...]?" (2002, 7). The same aggregation issues, of course, surface in the case of learning. One option that has been widely exploited both in Safety science and in Organization theory consists in bypassing the aggregate aspect altogether by postulating that collective entities[18] possess cognitive capabilities of their own. In analogy to human beings, collective entities can be said to have the ability to acquire new knowledge (and possibly to forget it too). This is the approach adopted, for instance, by (Hopkins: 2008) in his investigation of the 2005 explosion at BP's Texas City refinery which killed 15 workers and injured more than 170 others. Hopkins' interpretation can be easily seen in the following passage: "In the previous chapter, we saw that Texas City did not learn from the process incidents occurring at the site [in previous times] because of a general lack of focus on process safety. In this chapter, I want to argue that the failure to focus on process safety involved such a serious failure to learn the lessons

[18] The notion of collective entity covers a lot of ground in terms of size and cohesiveness. They can be as small as a design team made up of two engineers working in close cooperation or as large and loosely connected as the whole engineering community. In between are organizations of any size and with largely different hierarchical structures.

already available that Texas City can be said to have suffered from some kind of learning disability" (65-66).

In a review paper on organizational learning, Friedman et al (2005) call this approach to collective entities as if they were individuals *anthropomorphic*, and contend it does pose some problems. When used parsimoniously it is a helpful heuristic that "has offered fertile ground" (20) for research. Abuse, on the other hand, can turn it into a sterile source of mystification that ignores the internal complexity of collective entities and how knowledge is created and circulated within them.

Even if controversial, the anthropomorphic approach is intuitively attractive and appears to capture effectively large scale changes in engineering knowledge like those that followed the collapse of Tacoma Narrows Bridge. In fact, the moral that Petroski draws from that episode aligns neatly with that approach for he believes the suspension bridge engineering community as a whole learned from the failure of Tacoma Narrows the hidden dangers of wind loads and how to control them. Similarly, Vincenti's treatment of the growth of engineering knowledge is underpinned by the idea of a stored up body of knowledge that is the shared possession of the engineering community. Consider, as an example, his view of engineering standards like the boiler code promulgated by the American Society of Mechanical Engineers: "Such universal specifications, like the criteria on which they are based, become part of the stored up body of knowledge about how things are done in engineering" (212).

6.4.2. Analysis is not learning

Be it at the individual or at the collective level, the engineering literature is unanimous in describing learning from experience (of either failures or success) as a process. Unsurprisingly, different ways of spelling out the stages of this process have been proposed. Still, disagreement occurs mostly about the details while there is substantial agreement about the main aspects. To start with, failures and successes need to be identified. Spectacular collapses like Tacoma can be misleading in giving the impression that failures come bundled with a neon sign, as it were, that says "Failure here". As a matter of fact, failures can be deceptive and in the absence of clear identification criteria they can escape detection.

According to the reconstruction given in (Cowan et al.: 2006) this is what happened with the radiation therapy overdose accidents caused by functional failures of Therac-25 machines. When the first radiation overdose accident occurred the 3rd of June 1986, the local technician dismissed as "impossible" the patient's complaint that she had been burned. Even afterwards, when the patient manifested clear symptoms of burning injuries, thereby suggesting radiation overdose as a likely explanation, machine failure was never seriously taken into account. The whole event was "treated as a one-off fluke" and as a consequence, Cowan et al. comment, no learning occurred. Only after a string of similar accidents had happened, it became clear that something went amiss and the machine was called into question. Thus, there can be a lag between the moment in which the engineering artifact deviates from the established goals and the realization that this deviance constitutes a failure in need of an explanation.

Sure enough, virtually all accounts of learning include an analysis stage whose aim is to investigate the focal event and develop an explanation. Most of them link explanation to the identification of *causal factors*. Unsurprisingly, engineering discussion on what causal factors are has produced a vast literature. On the physical-chemical side, engineers have successfully isolated a number of *failure mechanisms* whose ranks are constantly increasing thanks to technical and scientific advances. Dasgupta and Pecht (1991, 521) define failure mechanisms as "the physical processes by which *stresses* cause damage to the elements comprising the system" (emphasis in the original) and classify them into five categories based on the nature of the stresses which trigger the mechanism: mechanical, thermal, electrical, radiation, and chemical (e.g., the bi-metallic corrosion process that caused the HMS *Alarm* failure belongs to the chemical category).

When it comes to the human side, the picture is far less clear and unanimous. If failure analysis determines that a specific failure mechanism, say corrosion, is responsible for the demise of an item, it is only natural to wonder why that mechanism was let to happen. Was there a mistake in material selection? Has the item been used outside its intended operational environment? This kind of questions brings to the fore an altogether different set of causal factors pivoting on the role of humans either as individuals or in groups. How to characterize these factors, parse them into different levels, and work out the causal relations involved are hotly debated questions that have resulted in a highly fragmented area of research. Regardless of the debates on the role and

nature of human factors, good explanations are expected to take advantage of the best engineering knowledge available and to *generate* knowledge as to how, what, and why the focal event did occur. By circulating that knowledge (e.g., by means of official failure reports) engineers and the wider public alike acquire a new piece of knowledge that was previously unavailable.

An interesting aspect of this new knowledge concerns its relation to knowledge available *before* the focal event. It is often assumed without much argumentation that the knowledge about the causes of the event *matches* the lack in knowledge that was somehow responsible for setting the stage of the event. Petroski contends that:

> If the cause of a failure is understood, then any other similar structures should come under close scrutiny and the incontrovertible lesson of a single failed structure is what *not* to do in future designs. (Petroski: 1985, 97, emphasis in the original)[19]

A claim which is echoed by Hummerdal et al (2013, 404) thus:

> Learning from failure can be defined as the act of creating a difference between what was known before and after the failure. These differences are often formalized in abstract knowledge (like an accident report) and made available for future remembrance.

Claims like these fit well with episodes such as HMS *Alarm* and Tacoma Narrows where failure analysis generates a piece of knowledge previously missing from the engineering toolbox and one that would have steered the design process towards a different solution. According to Vincenti's categorization of engineering knowledge, those episodes end up contributing to the available set of *fundamental engineering concepts*.

Now consider the second example of by-metallic corrosion mentioned in the Introduction: the evaporator end-plate case. Nothing of substance was added to already available fundamental engineering concepts by investigating it. Trethewey and Chamberlain do not provide further details on the mishap; still it would not be surprising that the evaporator was designed by engineers who were

[19] Another quote of the same tenor: "When failures do occur, engineers necessarily want to learn the causes. Understanding of the reason for repeated failures – structural or otherwise – that jeopardize the satisfactory use and therefore the reputation of a product typically leads to a redesigned product" (Petroski: 2001, 13).

already cognizant of galvanic corrosion phenomena. Perhaps they just copy-pasted the detail from a similar design and then it slipped through the design review process.[20] It does not mean that analysis of similar episodes is futile, though. Even if *collective* engineering knowledge can remain unaffected, realizing that such mistakes can happen may result in improvements to *local* engineering knowledge and practices, for instance by strengthening design review procedures. According to Vincenti's taxonomy, those changes amount to revisions affecting categories of knowledge such as *practical considerations* or *design instrumentalities*. The point could be summarized thus: in situations like HMS *Alarm*, post-failure knowledge matches the pre-failure knowledge gap in fundamental engineering concepts and affects *directly* how engineers will solve similar design problems in the future; in the evaporator case, knowledge generated from failure analysis affects *indirectly* future design solutions by reducing the likelihood of mistakes or oversights.

So far, the discussion has dealt with generation of knowledge from failure episodes. What about episodes of success? Petroski repeatedly dismisses the possibility that successes play any contributing role to engineering knowledge:

> While engineers can learn from structural mistakes what not to do, they do not necessarily learn from successes how to do anything but repeat the success without change. (Petroski: 1985, 98)

This claim presupposes a failure-biased notion of engineering knowledge and seems to underestimate the analytic abilities of engineers. While Petroski is willing to concede that engineers have the analytic abilities to unearth causal factors responsible of failure events and extract valuable knowledge from them, he denies that successes can be conducive of analogous beneficial results. Many examples can be found, however, of engineers acquiring new knowledge from the study of success cases. Ellet realized that French suspension bridges had something to teach and that is why he went there to study. More recently, American automakers were taken by surprise by the ability of their Japanese

---

[20] Copying from previous work is standard practice in engineering. Taylor (2007, 81) notes that: "A very large part of the design of a process plant involves copying. This may be of complete units, of parts of plants, or just individual design details. Copying is carried out not out of laziness or even out of a special drive to minimize design cost. Some companies even establish standard designs which they insist be copied".

competitors, chiefly Toyota, to manufacture highly reliable cars at much lower cost. Thus, they set out to learn from the Japanese how to improve their products. To be sure, their aim was not merely that of cloning Toyota's cars: by studying Toyota's methods they were looking for new knowledge to be implemented in the realization of their own products.[21]

Let us return to the learning process. Prompted either by a failure or by a success an epistemic gap has been exposed. The next step consists in carrying out an analysis whose goal is to bridge that gap by providing answers as to what, how, and why of the focal episode. Most importantly, the analysis should be able to generate *lessons*. The term *lesson learned*, especially with respect to failure events, is widespread in the engineering literature. Gordon (2008, 31) defined it as "information that has a real impact on operation; valid in that it is factually correct; and applicable in that it identifies a process or decision that reduces or eliminates the potential for the recurrence of an incident or reinforces a positive result".

Although clearly designed with operational safety in mind, Gordon's definition highlights several interesting aspects that apply to learning in engineering generally. First of all, however, it should be noted that this definition misses and important point about lessons learned. To qualify as a lesson, besides being valid and applicable the information needs to *originate from the analysis of a focal episode* (either failure or success). Lessons are *tagged*, so to speak, when they are generated. In principle, the same piece of knowledge could be arrived at through transfer from science or generation through engineering experimentation as well as from analysis of a failure or success episode. Only the latter though, makes it a lessons learned. With the passage of time that piece of knowledge will become an integral part of the body of knowledge shared by the community of practitioners and its nature of lesson learned will fade away. In fact, nowadays no engineer looks at design criteria against bi-metallic corrosion as lessons learned from failure, and probably only a few are aware what the frigate HMS *Alarm* has to do with it.

Another aspect worth noting of Gordon's definition is the emphasis on the fact that lessons are meant to be *applicable* in a very specific sense: lessons from

---

[21] According to (Ward et al.: 1995, 43) one crucial factor of Toyota's success in making "better cars quickly and cheaply" consists in the rather counter intuitive strategy of delaying "decisions and provide their suppliers with hard, specifications very late in the process".

failure are expected to have the effect of preventing recurrence, whereas lessons from success should promote further positive outcomes. The application aspect is indeed so important that many authors are explicit in saying that until the lessons are turned into practice learning cannot be considered complete. That is what Carroll and Fahlbruch (2011) mean by claiming that "analysis is not learning": analysis and generation of lessons constitute a vital stage of learning. Yet, "once the useful information from an incident has been defined and extracted, the knowledge *must* be implemented" (Jacobsson et al.: 2011, 335, emphasis added). Models of learning include, in parallel with implementation, a *diffusion* stage where lessons are circulated throughout the relevant community and organization.

Wrapping up, *learning from X*, where *X* stands either for failure or for success, is shorthand for *learning from the analysis of X* and consists in a process that begins with the identification of an epistemic gap associated with the occurrence of *X* and concludes with implementation and diffusion of the engineering knowledge generated in the course of analyzing *X*.

The exploration of learning undertaken in this section, then, arrives at a conclusion which corresponds nicely with an important result from Vincenti's study of engineering knowledge which he summarizes as follows:

> The inseparability of knowledge and its practical application is in fact a distinguishing characteristic of engineering. (Vincenti: 1990, 207)

6.5. The learning hypothesis disambiguated

Equipped with definitions and clarifications from the preceding sections, we can now turn again to Petroski and see why his learning hypothesis is ambiguous. An appropriate starting point is the following extract from the book *To Engineer is Human*:

> Thus the lessons of failures generally pinpoint weak links. [...]. The weak link can be avoided or strengthened in future designs, and the science of that genre of weak-linked structures can generally be said to have benefited in a way that years or even decades of skywalks hanging, bridges standing, or DC-10s flying did not. For this reason it is important that engineers study failures at least as much, if not more than successes, and it is important that the causes of structural failures be as openly discussed as can be. Should a young engineer look for models in weak-linked structures while they are still functioning, he could indeed design weak links into his own structures. However, if the cause of a failure is understood, then any other

> similar structures should come under close scrutiny and the incontrovertible lesson of a single failed structure is what *not* to do in future designs. That is a very positive lesson, and thus the failure of an engineering structure, tragic as it may be, need never be for naught. (Petroski: 1985, 97, emphasis in the original)

This passage shows that Petroski is unaware of the ambiguities surrounding the concept of learning explored in Section 6.4 and of the problems they pose to his hypothesis. As signaled by the reference to "the science of [...] weak-linked structures", he starts out by dealing with learning at the most general level, the level of an entire engineering community, and claims that failures play a preponderant role. Immediately afterwards, he contends that the same situation obtains with respect learning at the individual level, in particular with respect to young engineers. Now, the point I am trying to make here is not that Petroski's hypothesis is erroneous. The issue I am dealing with precedes any discussion about the claim's rightness or wrongness. I argue that, given the ambiguities concerning the concept of learning and given the range of phenomena that Petroski aims to cover (e.g., learning both at the individual and at the collective level), it makes good sense to question whether the hypothesis that "more is learned from failures than from successes" can be taken at face value or whether it needs to be disambiguated first. More precisely, I argue that, despite the appearances that the hypothesis is always one and the same, Petroski's narratives conceal at least two different hypotheses about learning each one based on a different concept of learning.

It might be tempting to resolve the ambiguity by splitting the hypothesis in two versions: one version dealing with learning at the individual level and the other dealing with learning at the collective level. However, as mentioned in Section 6.4, collective learning is problematic for it covers a wide range of epistemic agents, from small and well defined organizations (e.g., closely knit design teams), to large groups of loosely connected people, to multicultural communities of practitioners. More importantly, the individual-collective dichotomy neglects the fact that engineering activities are carried out simultaneously at different organizational levels. Correspondingly, learning can happen at any level and lumping together into the collective learning category anything that involves more than two individuals amounts to a gross simplification.

A more promising approach can be found in Vincenti's (1994) historical study on landing gear technology where two different modes of learning are identified which are called *specific learning* and *generic learning* respectively.

Vincenti's distinction does not depend exclusively on the level of learning. In fact, while generic learning deals with learning at the community level, specific learning could be seen as somewhat hybrid with respect to the level of learning for it covers both individual engineers and engineering organizations. It is worth stressing, however, that in the scale going from the individual to the community, the organizations envisaged in specific learning lie somewhat in the middle. Thus, they are more restricted and less fluctuating than a whole community of practitioners.

Besides the level-of-learning issue, the two modes diverge on a number of aspects, including time scale of the learning process and outcome of learning (i.e., what is learned). To clarify these additional aspects and to explain how Vincenti's modes of learning can help disambiguating the learning hypothesis I will proceed as follows. In Section 6.5.1, I will summarize Vincenti's paper and the two modes of learning he presents there. I will also explain why the notion of goal discussed earlier in Section 6.3 is needed. Then, in Section 6.5.2 and 6.5.3, I will show that Vincenti's two modes of learning allow to disambiguate Petroski's hypothesis in a meaningful way although with rather opposite outcomes: with respect to specific learning it leads to an empirically testable statement, while the generic version of the hypothesis appears to rest on questionable assumptions.

6.5.1. Learning about landing gears: specific and generic learning

Nowadays, all high-performance aircraft are equipped with retractable landing gears. The first examples of retractable landing gears appeared in the late 1920's. Yet, in the early 1930's, many of the fastest and more technically advanced aircraft had still fixed landing gears. Among these were the airplanes designed by one of the most innovative engineers in aviation history, Jack Northrop (1895-1981), which were equipped with so-called trouser gears that is to say, fixed landing gears encased in streamlined fairings to reduce drag.

In his paper, Vincenti sets out to investigate the period during which the engineering community experimented with different kinds of landing gear technology, e.g., retractable, fixed, trouser, etc.; see Figure 6.2. For a while, all these arrangements were considered valid solutions to the design problem of providing airplanes with the ability to land safely.

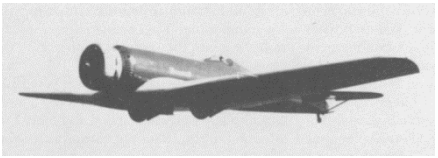

Partially retractable landing gear on a Boeing Monomail

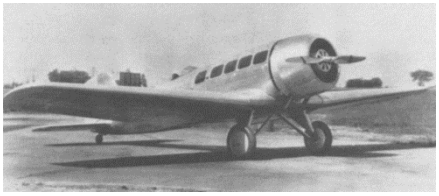

Fixed tripod landing gear on a Northrop Alpha

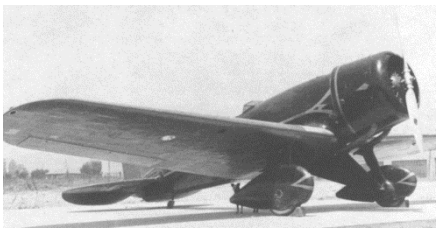

Wheel pants on a Lockheed Sirius

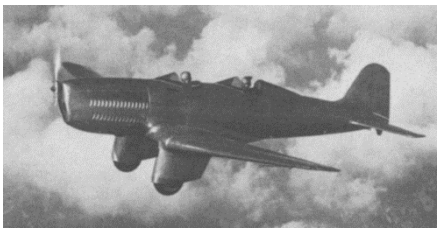

Trouser gear on a Northrop Beta

Figure 6.2: High-performance airplanes from the early 1930s showing different kinds of landing gears. From (Vincenti: 1994)

As usual with engineering, solving the problem involved a tension between conflicting requirements, chiefly: robustness, light weight, reliable operation, ease of maintenance, and low aerodynamic drag.

Vincenti's historical account shows that initially it was far from clear that retraction represented the best compromise for high-performance aircraft:

> Though engineers nowadays take it for granted, it had to be *learned* at some time. (Vincenti: 1994, 24, emphasis added)

Within this transition period, Vincenti distinguishes two learning processes which he calls *specific* and *generic*. The specific-learning process refers to the design and testing activity which allowed individual airplane designers to acquire knowledge needed to solve *specific* landing gear problems for particular airplanes. Although there is a higher level design goal which is shared among these designers of high performance airplanes, details of the conflicts between requirements work differently for each airplane. For instance, the innovative wing structure layout he had invented confronted Northrop with additional constraints which made a retractable gear less attractive. In the early stages of the transition period, by studying and experimenting with the available options (either directly or through the engineering literature), Northrop learned that a trouser gear was the optimal solution. In fact, his airplanes were among the fastest of the lot and were well received by the market.

Other designers solved the trade-offs differently: some became early adopters of retraction and some decided for so-called wheel pants (streamlined fairings enclosing just the wheels). As a result, for a period the market of high performance airplanes had several alternatives on offer. In the second half of the 1930's, things begun to change and retractable gears became increasingly popular until, by the end of the decade, they established themselves as "long-term generic solution" to landing gear design for high performance airplanes. What factors did prompt this change? Two factors played a crucial role according to Vincenti's reconstruction. First, because of more powerful engines airplane speeds went up, thereby increasing the drag penalty imposed by fixed landing gears. Second, engineers learned how to improve the reliability of retraction mechanisms.

By generic learning, Vincenti refers to a community wide learning process which spans the whole transition period.

> The design community followed the [generic learning] process to find out whether high-speed airplanes *as a whole* ought or ought not to have retractable gear; by doing so it solved a generic problem for a class of aircraft. (Vincenti: 1994, 25, emphasis added)[22]

A specific-learning episode can be considered successful if, given the current status of engineering knowledge, including lessons from failures and successes, and given the specific design goal at hand, the design solution implements those lessons and meets the stated goals. In practice, because of the many uncertainties at stake (e.g., status and distribution of available engineering knowledge, circularity between available knowledge and establishment of design goals), assessing individual episodes of learning in isolation is problematic. Comparative studies that take into account multiple episodes provide more accurate assessments. At the beginning of his study Vincenti wonders whether Northrop

---

[22] In his paper, Vincenti proceeds to show that both learning processes can be accommodated within a *blind-variation and selective-retention* model of learning. According to that model, the main difference between the two processes lies in the *locus of selection*. In the case of specific learning, the selection process among alternatives is located within the mental processes of an individual engineer or the deliberation process within a team of designers. Instead, the locus of selection "for the generic solution must lie in the design community" (Vincenti: 1994, 26). This part of the paper, though, is not immediately relevant for the analysis of the learning hypothesis.

did fail to learn how to design landing gears for high performance aircrafts. To answer this question, he decides to compare Northrop's designs with work from other contemporary American engineers. The comparison shows that, although Northrop's solution would not prove the most effective in the long term, with respect to its specific goals it was a viable solution and Northrop did implement available lessons.

The bottom line is that, in order to be meaningful, such kind of assessments must be based on episodes of learning that are reasonably homogenous with respect to design goals and available engineering knowledge. It would be unreasonable to compare Northrop's design goals with the objective of designing the landing gear for, say, a large transport airplane. Similarly, design choices made on the basis of knowledge available in the early 1930s should not be directly compared to design choices that rely on knowledge existing a decade later.

6.5.2. The learning hypothesis with respect to specific learning

In the second half of the passage quoted above, Petroski speculates about a hypothetical young engineer who is faced with design choices and can look either at lessons from failures and from successes. Successes, Petroski claims, can be deceptive: Ellet's Wheeling Suspension Bridge (see Section 6.2.1) was completed in 1849 and for almost five years it fulfilled the design goals that were set for it. The storm that took it down on May 1854 *revealed* that, hidden in its design, there were serious flaws.

But, did the collapse *reveal* anything that was not known before? In the same years Ellet was building his bridge, Roebling was questioning the French approach to suspension bridge design. He was convinced that it did not provide adequate stiffness to withstand high winds. Thus, the collapse of Ellet's bridge was not a revelation to Roebling. Both Ellet and Roebling were accomplished bridge engineers. According to Petroski's narrative, the reason they arrived at such different solutions to the same higher level design goal (i.e., to build a long span suspension bridge) derived from their approach to lessons from failures and successes. While Ellet's design solution was inspired by the successes he witnessed during his visit to France, Roebling analyzed carefully past failures and derived design guidance from them. Roebling learning process was successful because he was able to prevent reoccurrence. Ellet, on the other hand, derived the wrong the lessons from success and was unable to replicate them.

It can be seen, then, that this comparison between Ellet and Roebling constitutes an example of Vincenti's *specific learning*: the two designers share a similar goal and rely on the same stock of background engineering knowledge. Petroski's study of the history of engineering provided him with multiple examples of the same pattern. Hence his conclusion:

> It follows, therefore, that having as wide and deep an acquaintance as possible with past failures should be at least desirable, if not required, of all engineers engaged in design. Understanding from case histories how and why errors were made in the past cannot but help eliminate errors in future designs. And *the more case histories a designer is familiar with* or the more general the lessons he or she can draw from the cases, *the more likely* are patterns of erroneous thinking to be recognized and generalization reached about what to avoid. (Petroski: 1994, 6, emphasis added)

This passage, particularly the words in italics, reveals what may be called "the learning hypothesis with respect to specific learning" or, for simplicity's sake, the *specific-learning hypothesis*. This passage makes clear that, when it comes to specific learning, the claim that "more is learned from failures than from successes" should be interpreted as follows:

> The more case histories of failure epistemic agents are familiar with and the more general the lessons they draw from the cases, the more likely they are to prevent failure

The specific-learning hypothesis could be translated in graphical terms as shown in Figure 6.3.

At the top of the diagram sits a higher-level design goal which is shared among a series of engineering projects. In terms of the conceptualization of goals presented in Section 6.3, it corresponds to the upper part of a goal tree a new product is expected to achieve. By achieving it, the new product is also expected to satisfy *customer needs*. For instance, the customer may be looking for a way to connect two sites located across a river and the answer may be to build a suspension bridge. A design goal, besides providing guidance and motivation for the upcoming development process, dictates a series of *metrics* (see Section 6.3) that allow to assesses the outcome and determine to which degree the finished product achieves the goal set for it. Moreover, it offers indications as to previous engineering realizations that may hold valuable lessons either from successes or from failures. The hypothesis contends that outputs of various engineering projects pursued by different agents (who could be either individual engineers or

engineering organizations), although having different lower-level goals, can be meaningfully compared as long as their higher-level goals coincide. Furthermore, these agents can be parsed into two groups A and B. While members of the two groups are assumed to share the same level of overall technical competence (i.e., are cognizant of the fundamental design concepts that apply to the goal at hand), what tells them apart is how they tap into lessons from failures and from successes. Group A is constituted by agents who make an effort to gain "as wide as possible an acquaintance with past failures" in order to prevent reoccurrence and invest comparatively less into analyses of successes. Agents in group B, on the other hand, are more inclined towards imitating past successes than studying failure episodes. The ratios of learning are represented in Figure 6.3 by arrows of different thickness.

The hypothesis predicts that products developed by members of group A are "more likely" to achieve the design goal, thereby also preventing reoccurrence of failure. In the terminology examined in Section 6.4, they have successfully completed the learning process that is to say, they have learned the lessons. On the contrary, output from group B is, on average, more likely to result in products unable to meet the design goal: success has not been replicated as hoped.

On what sort of conceptual or empirical grounds do these predictions rest? Petroski tends to emphasize a single advantage granted by learning from failures over learning from success: failures teach what *not* to do. If, as we have seen in Section 6.4 learning from failure means avoiding recurrence, just knowing that a certain design solution has failed may leave still open a plethora of potential solutions some of which could be even worse. In short, knowing what not to do does not necessarily bring you closer where you want to be. Sure enough, in many cases the analysis of failure can provide essential clues towards a viable alternative. The case of HMS *Alarm* readily comes to mind. Notice, however, that fortuitous circumstances played a crucial role there. It was by mere chances that a number of iron nails were insulated from the copper plates because of shreds of paper trapped beneath their heads. Years before the copper sheathing experiment, the Royal Navy experimented with a similar technology where ships had their hulls sheathed with lead plates (Harris: 1966). Again corrosion failure ensued, but this time investigations were inconclusive and lead sheathing was dropped.

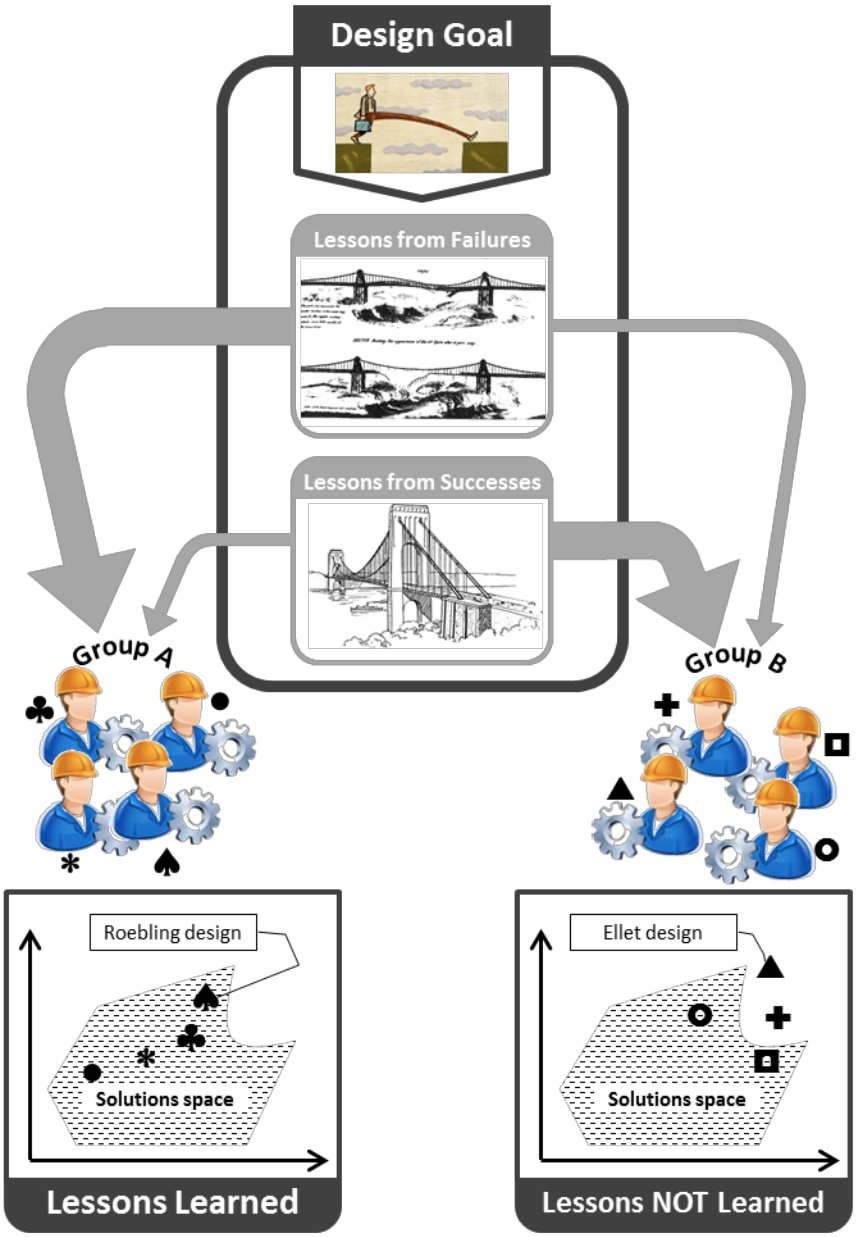

Figure 6.3: Model of the learning hypothesis with respect to specific learning

Moreover, failure analyses may arrive at conclusions that make the problem even worse. (Whyte: 1978) recounts an emblematic episode from the development of Armstrong Siddeley's *Sapphire* jet engine in the late 1940s. The *Sapphire* was designed to increase more than threefold the thrust provided by the previous

model. This substantial performance boost implied drastic changes in almost all components including the compressor rotor drum: it used to be made of aluminum and was replaced by one made of steel. Being much heavier, the new drum imposed higher torsional loads to the main shaft especially during engine ignition. Although engineers designed a new sturdier shaft thought able of coping with the increased load, the new shaft broke during testing. Engineers concluded they had underestimated the torsional loads and, as a countermeasure suggested by the failure analysis, they increased the shaft cross-section thereby making it even stronger. Long story short, the shaft failed again until there was no more space to design one with even a larger diameter. It was at this point that the chief development engineer realized that they made a wrong assumption right at the beginning and then, by holding to it, they did draw the wrong lessons from testing failures. By designing a sturdy and strong shaft at the beginning, they inevitably made it also heavy and stiff which resulted in substantial shock loads being generated during engine ignition. To prevent breaking, the engineer concluded, "the shaft should be weakened (by reducing it's [sic] diameter) to give it some torsional spring" (38). Initially, the proposal was met with a good deal of skepticism by the rest of the team. Given the lack of alternatives, however, it was tried and proved completely successful.

Petroski develops his argument eloquently and backs his hypothesis with abundance of historical examples. However, impressive as it might be, his collection of anecdotes has been put together following subjective criteria and does not qualify as a rigorous empirical study. Furthermore, his reasoning based on common sense intuitions about knowledge and learning, though credible, is not conclusive either. Thus, pending more solid empirical evidence or stronger conceptual motivations the case for the alleged primacy of learning from failures would remain undecided.

Suitably, Madsen and Desai (2010) have published recently a study set out to determine whether empirical evidence supports the hypothesis that organizations do learn more from failures than from successes. Madsen and Desai note that, even though the topic recurs often in organizational theory literature, "existing evidence that failure is more important than success for organizational learning is entirely anecdotal" (452). Scholars have put forward contrasting arguments. Sitkin (1992) offers a short compendium that summarizes *benefits* and *liabilities* of failures and successes. If successes can induce overconfidence, failures may result in conservatism. Failures are likely to promote a healthy

process of reassessment of current knowledge; however, they may engender a sense of urgency that drives agents into tunnel vision and neglect of broader issues. Similarly, building upon lessons learned from successes can set a path of incremental improvement, but in the long run it may turn into stagnation and mere repetition. Sitkin contends that, all things considered, the most beneficial approach to optimize learning resides in harvesting the fruits of small scale failures. Like most other contributions, however, Sitkin's paper is mostly speculative and contemplates systematic empirical testing only as future work.

Madsen and Desai decided to take a different approach by "disaggregating organizational experience into failure experience and success experience and comparing the contribution of each to organizational performance" (452). In order of doing so, their study focuses on organizations whose products share comparable higher level design goals (with correspondingly similar metrics to assess failure and success) and rely on a largely uniform body of knowledge. Thus, their study fits with the assumptions behind the specific-learning hypothesis and the diagram in Figure 6.3. The core business of the engineering organizations surveyed by Madsen and Desai consists in developing and operating orbital launch vehicles; the NASA is a prominent example. The top-most goal their products are expected to achieve is "to place a payload (one or more satellites) into orbit around the earth" (458). With it comes a rather clear cut failure criterion: either the payload is placed into its intended orbit or it is not.[23]

From their study, Madsen and Desai conclude that on average organizations do learn more from failures than from successes. In particular, they claim that the balance tilts in favor of learning from failure because of the contribution of what they term *visible failures* that is to say, launches that dramatically (and sometimes spectacularly) miss to achieve the design goal and result in substantial losses. Madsen and Desai maintain that knowledge earned from visible failures *persists longer* than knowledge from success and, more importantly, from small scale failures. To illustrate this point they contrast two episodes from the Space Shuttle history. One is notorious: the 2003 *Columbia* disaster when the orbiter disintegrated during reentry. The official investigation determined that

[23] Madsen and Desai note that some launches may seem to fall in between failure and success in that the satellite, while successfully injected into the right orbit, is seriously damaged in the process (e.g., by colliding with the rocket during the separation phase). They decided to consider these launches to be failures.

few minutes after lift-off a piece of insulating foam detached from the external tank and hit the orbiter's left wing at very high speed thereby damaging several thermal protection tiles. The fact is that foam shedding was a known problem NASA engineers had been struggling with since the inception of the Space Shuttle program. Over the years, engineers attempted various solutions from changing the foam's chemical composition to revising foam application procedures. Although never completely solved, it was assumed that foam shedding was unlikely to cause serious damage to the Shuttle and did not constitute a major threat.

That was the situation until the launch of *Atlantis* in 2002, just a few months (and two Space Shuttle launches) before *Columbia*, when a chunk of foam did cause structural damage, not to the orbiter, but to a metal link connecting the external tank and the left solid rocket booster. Luckily, the *Atlantis* survived the mishap and was able to complete its mission without further accidents. Still, the episode clearly falsified the assumption that flying with foam shedding was safe and should have alerted NASA managers and engineers about the need of urgent corrective action. In short, in its current shape the Space Shuttle was unsafe and unable to achieve its goal. Instead, NASA arrived at the opposite conclusion. Given the small scale of failure consequences, which did not compromise the launch, "the *Atlantis* mission was viewed as a success" (451) and investigation into it was not given high priority. Indeed, by the time of *Columbia*'s launch in 2003 it had not been completed.

Although the two foam shedding episodes where almost identical (both of them originating from the same area of the external tank and having roughly the same size), their consequences where dramatically different and so were reactions by NASA. Whereas *Atlantis* did not suffice to undermine NASA unwarranted assumptions, *Columbia* was the epicenter of radical and long lasting organizational changes. This disparity can be found in many other cases and according to Madsen and Desai it explains the role of visible failures in learning. Lessons learned from visible failures are often "codified and embedded in formalized organizational memory systems. [...]. On the other hand, because success reinforces existing bases of organizational knowledge, organizational decision makers are unlikely to alter formal organizational memory systems in response to success" (456).

In concluding their paper, Madsen and Desai acknowledge that their results, being based on data from a rather exotic form of engineering, may have limited

generalizability. So far, however, the first empirical study on learning from failures and success appears to corroborate the *specific-learning hypothesis*. Moreover, by making clear the decisive role of visible (or high-profile) failures in tilting the balance of learning towards the failure side, it clarifies a factor that though present in Petroski's writings, is somehow diluted in the abundance of fairly similar arguments he makes. When Petroski claims that "Tacoma Narrows Bridge proved *more instructive* than the success of all the bridges that had performed satisfactorily – or nearly so – over the preceding decades" (2001, 11, emphasis added), he is indeed making the same point as Madsen and Desai that high-profile failures result in drastic changes to several categories of engineering knowledge. Those changes can be shown to be effective in preventing recurrence of similar episodes: lessons have been learned.

6.5.3. The learning hypothesis with respect to generic learning

Let us return to the passage from Petroski with which this section started out. The first part of that passage reads as follows:

> Thus the lessons of failures generally pinpoint weak links. [...]. The weak link can be avoided or strengthened in future designs, and *the science of that genre* of weak-linked structures can generally be said to have benefited in a way that years or even decades of skywalks hanging, bridges standing, or DC-10s flying did not. (Petroski: 1985, 97, emphasis added)

Note that here, Petroski is not dealing with specific learning anymore: the gist of the argument has shifted from individual engineers or design teams learning how to deal with specific problems (e.g., Northrop busy studying how to design a landing gear for his airplanes), to *the science* of an engineering structure and how it develops. Although his hypothesis remains the same (i.e., more is learned from failures than from successes) a crucial change has been made in that the local differences between learners, which constitute the backbone of the learning hypothesis with respect to *specific learning*, have been obliterated and replaced by an undifferentiated body of knowledge that moves uniformly in a specific direction. The direction of change being influenced more by failures than from successes. Another revealing quote from the same book is the following:

> To understand what engineering is and what engineers do is to understand how failures can happen and how they can contribute more than successes *to advance technology*. (Petroski: 1985, xii, emphasis added)

To summarize, while the specific-learning hypothesis deals with epistemic agents being able to achieve varying results depending on their attitude towards lessons from failures and from successes, the learning hypothesis with respect to generic learning, or *generic-learning hypothesis* for short, consists of an explanation of technical change in terms of knowledge generated from failures and from successes. Petroski tacitly assumes that technological evolution can be meaningfully compartmentalized into *technical genres* (e.g., hanging skywalks, airplane cargo doors, suspension bridges) each characterized by a set of *parameters* that remains fairly stable over time. Of course, observed *performance*s along those parameters do vary over time. The set of characteristic parameters for say, suspension bridges, has always included aspects such as main span and carrying capacity. Suspension bridges from different epochs achieve different levels of performance: Finley's bridge crossing Jacob's Creek in Pennsylvania (completed in 1801) had a main span of just above 20 m; the current record holder, the colossal Akashi Kaikyō Bridge in Japan (completed in 1998), has a main span of 1991 m. Nevertheless, since they belong to the same *technical genre*, they are characterized by the same set of parameters. Thus, chronologically and geographically distant items can be compared directly and graphs can be drawn where observed performance is plotted against time.

Now, *value criteria* are needed in order to understand which way is *forward* or *advancing*. Often times the exercise is forthright for there is an intuitively compelling directional pattern: in suspension bridge engineering technical advance means longer spans, in aviation it means faster airplanes, and in computers greater computational power. Indeed, the intuition is so compelling that Petroski takes it for granted and does not investigate any further the idea of technical advancement. Instead, what he does in his narratives is to look closely at the history of a *genre* seeking for those factors that can explain how and why advancement happened. Be it aluminum cans, paper clips, or suspension bridges those factors turn out to depend on failures.

Petroski maintains that, similarly to historical development in science, technical advancement follows a dynamics of "recurrent revolutions against 'normal science' [as] Thomas Kuhn so convincingly demonstrated" (Petroski: 2006, 177). Periods of *normal design* in technology are those during which a design solution that has been found able to achieve a certain goal spreads throughout the community of practitioners and becomes routine or, borrowing again from Kuhn, a *paradigm* (e.g., slender suspended bridges designed accord-

ing to deflection theory, see Section 6.2.2). Although during this period the design solution is put under increasingly higher demands, its main assumptions or, in Vincenti's terminology, its fundamental design concepts are maintained. Gradually, the range of products made possible by the design solution is used up, as it were, until a failure occurs which proves that the limits have been trespassed. At this point, Petroski's model predicts that a new paradigm will emerge thanks to lessons learned from failure and it will put the train of technical advancement back on the same track of performance growth (i.e., longer spans, higher speed, and so on).

Petroski's model of technical change is summarized in Figure 6.4. In the upper part there are five silhouettes of suspension bridges of increasingly larger size. Though resembling real bridges the silhouettes are for mere illustrative purposes and are not meant to be historically accurate. The pair of Cartesian diagrams in the middle represents what Petroski calls *normal design*, the historical phase in which learning is driven mostly from imitation of previous successes: bridges become larger and larger, yet there is relatively little growth in engineering knowledge. The box at the bottom represents the paradigm shift that ensues in the wake of a major failure, here illustrated by the destructive oscillations which brought down Tacoma Narrows Bridge. This time the diagram shows a sudden and significant increase in engineering knowledge and afterwards a new phase of normal design taking over. It is worth stressing here that Petroski's view on paradigm shift departs remarkably from the account given by Kuhn. While for Kuhn a paradigm shift represents a chiasm between two consecutive periods of normal science whose main concepts are assumed to be incommensurable, in Petroski's account a paradigm shift actually bridges the gap between two normal design phases. As mentioned above, Petroski assumes that technical genres survive unscathed paradigm shifts: pre-shift and post-shift suspension bridges belong to the same historical tradition. Post-failure designs, however, are more advanced or, to put it differently, they rely on a more advanced body of engineering knowledge.

Crucially, according to the generic-learning hypothesis, such noticeable advances in knowledge are seldom, if ever, achieved by learning from successes and they are almost invariably the outcome of learning from failures. In Petroski's words:

> This again is the paradox of design: Things that succeed teach us little beyond the fact that they have been successful; things that fail provide incontrovertible evidence

that the limits of design have been exceeded. Emulating success risks failure; studying failure increases our chances of success. (Petroski: 2006, 114)

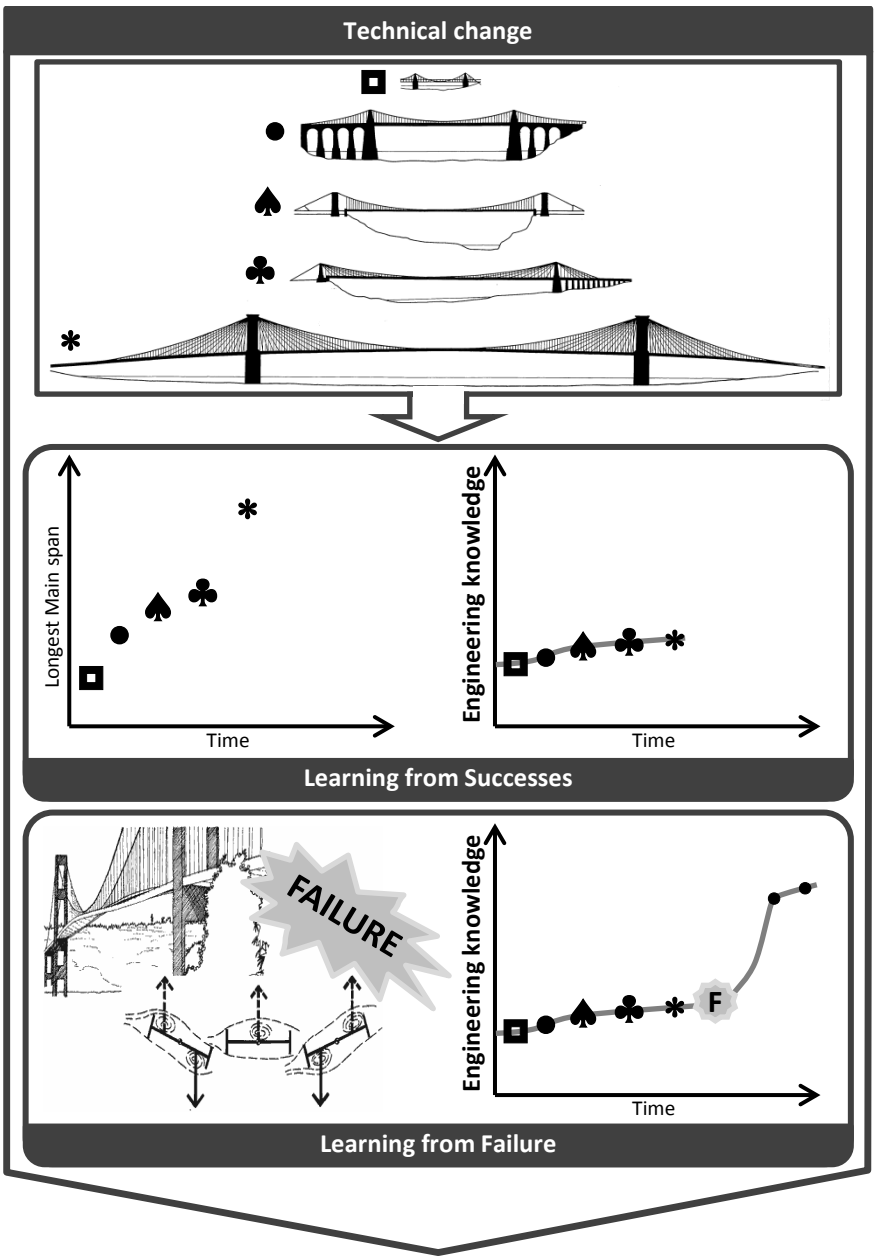

Figure 6.4: Model of the leaning hypothesis with respect to generic learning

Arguably, Vincenti's categories of engineering knowledge mentioned in Section 6.4 may help clarifying Petroski's view of paradigm shift. There are many analogies between what Petroski calls an *engineering* or *design paradigm* and Vincenti's knowledge category of *fundamental design concepts*. Thus, Petroski seems to suggest that failures benefit engineering knowledge by forcing engineers to rethink and strengthen fundamental design concepts. For instance, deflection theory was put into question because of Tacoma Narrows collapse and aeroelasticity theory became a vital aspect of bridge design. On the other hand, those performance improvements that may take place during normal design phases do not affect extant fundamental design concepts and any growth in engineering knowledge it is limited to other categories like *practical considerations* or *design instrumentalities*.

Comparing specific learning and generic learning by looking at Figure 6.3 and Figure 6.4 allows to appreciate that different concepts of learning are involved, thereby justifying the need for disambiguating Petroski's learning hypothesis into specific and generic. Besides the issue of the *level of learning* that has already been discussed, there are at least two further aspects. First, the learning hypothesis with respect to specific learning is *probabilistic*. Recall Petroski's quote to the effect that "the more case histories [of failure] a designer is familiar with [...], *the more likely* are patterns of erroneous thinking to be recognized" (Petroski: 1994, 6, emphasis added). This aspect is rendered in Figure 3 by having two groups of agents, A and B, whose outcomes display different distributions, with the outcomes of group A being more likely to fall within the solutions space. Thus, it is in this probabilistic perspective that the claim "*more* is learned from failure" should be interpreted. In fact, that is also the reading adopted in Madsen and Desai's empirical study. When it comes to generic learning, the probabilistic aspect is replaced by an accounting view, as it where, in which failures are believed to add more to the existing stock of knowledge than successes do.

That brings about the second relevant difference: What does count as an instance of successful learning? With respect to specific learning, learning from events (either failures or successes) is said to have occurred when knowledge acquired from the study of focal events has been implemented into practice, thereby resulting in prevention of further failures or iteration of success (see Section 6.4.2). A key part is played by engineering goals and the associated metrics that characterize the products involved in the focal events. Since pre-

event products and post-event products share the same higher-level goals they can be assessed by the same set of metrics, which provide engineers with objective criteria to decide whether progress has indeed been made. Think of the orbital vehicle example: after a failure has happened, a rocket is redesigned and changes are implemented to avoid recurrence. When the redesigned rocket is operated its performance is compared to that of the previous design by means of the same set of metrics. It is precisely the fact that goals and metrics are largely constant that makes specific learning *specific*. It unfolds within a fairly stable set of boundaries and goals largely acknowledged by the agents taking part to it. Generic learning, on the other hand, allows for such a broader range of changes in goals, boundaries, and knowledge that it becomes far less clear what sort of criteria should be utilized to decide whether learning has occurred. To solve this problem, Petroski has developed what I would call *the technical-champion approach*. It consists in postulating that the level of technical advancement of any epoch can be represented by the undisputed *champion* of its days, e.g., the longest bridge, the fastest airplane, the tallest skyscraper and so on. In technology as in sports, successive technical champions set the performance bar increasingly higher. Hence, a path emerges which in hindsight can be utilized to tell where technology has been heading all the time. The following quote illustrates the technical-champion approach in action:

> The design process, whether it be applied to bridges or anything else, is timeless. It proceeds through persistence from failure to success. The failure of a single piece of stone to span a great distance led to the use of multiple spans. The failure of a single cast-iron beam to span greater than about thirty or forty feet (and only rarely as far as fifty feet) led to the trussed girder. The failure of the trussed girder to span a hundred feet led to the wrought-iron tube and the open truss. Other limitations led to the use of the suspension and cantilever bridges. The failure of the cantilever to withstand the rigors of construction led to its curtailment and to the dominance of the suspension bridge. And until the extension of the cable-stayed bridge into once unheard of realms, there appeared to be no alternative for long-span structures. (Petroski: 2006, 161)

Petroski's approach appears to conflict with recent accounts of technical change. Historians and sociologists of technology have shown that within an engineering community multiple traditions of practice can coexist whose activities adhere to different sets of metrics or, where these sets are largely similar, rank them differently. Even more importantly, sets of metrics themselves are far from fixed: they change over time by gaining new members and dropping older ones. As a

result, technical change unfolds simultaneously along multiple trajectories and at different rates such that singling out a single trajectory as the hallmark of this complex process is likely to produce a severely distorted picture. In recent accounts of technical change, for instance in Bijker's study on the development of the bicycle (1995), evolutionary trees have supplanted linear models. Evolutionary trees attempt to provide a more balanced representation of technical change by including products that, although ultimately abandoned, for a while were considered competitive or even hailed as the best technology available.

Another objection Bijker raises against linear models is that they tend to ignore the broader context in which technical change occurs and give the impression, like in Petroski's quote above, that technical change unfolds in a vacuum. Developments in bridge engineering, for instance, cannot be properly understood without considering contemporary developments in means of transportation. In the early 1920s, increasing popularity of the automobile and the corresponding decline in railway transportation created ideal conditions for suspension bridges, which cope better with moderate and evenly distributed traffic than the heavy and localized loads generated by passing trains.

In concluding this section, it may be worthwhile to briefly contrast Petroski's model of technical change with the historical study of the hard disks drive industry by Christensen (1997). Christensen's study is relevant here because his starting point closely matches Petroski's in observing that hard drive manufacturers "have established a trajectory of performance improvement over time" (9). Over the about thirty year period covered by the study, from 1967 until 1995, the information recording density (measured in megabits per square inch of disk surface) has increased by 35 percent per year, on average. The hard drive trajectory, although compressed in a far shorter interval, replicates the increase in main span of suspension bridges. Moreover, in both fields the trend towards greater performance has been made possible by a series of technical innovations. Christensen, however, notes that behind this seemingly linear trajectory there is a tumultuous process constituted of technical innovations some of which "were straightforward technology improvements; others were radical departures" (11). A good example of the latter is the introduction of 5.25-inch drives in 1980. At the time, two architectures were available on the market: expensive 14-inch drives with capacities around 200 megabyte (MB) were installed on mainframe computers, while cheaper 8-inch drives capable of storing around 60 MB were installed on so-called minicomputers.

Compared to 8-inch drives, the newly arrived 5.25-inch drives not only had an inferior capacity (around 10 MB), they were also slower and had a higher cost per megabyte. Consequently, they were not interesting for minicomputer manufacturers. However, a new application had just emerged that valued hard drives according to a different set of metrics: desktop computers. For desktop computers, physical size, weight, and purchasing costs were essential characteristics and 5.25-inch drives performed better than larger and more expensive 8-inch drives. Once established in the desktop market, the growth of 5.25-inch drives was so fast that it took manufacturers of 8-inch drives by surprise. After few years, capacities of 5.25-inch drives caught up with 8-inch drives and invaded that market too. Of the four leading 8-inch drive makers, Christensen concludes, only one survived to become a significant manufacturer of 5.25-inch drives. A similar pattern unfolded towards the end of the same decade when, in 1989, the 2.5-inch drive architecture was introduced. The market was dominated by 5.25-inch and 3.5-inch drives and, if judged according to the same parameters, the new entry was not competitive. This time, notebooks came to rescue the newcomer. With notebooks, performance parameters like low weight, small physical size, ruggedness, and low power consumption became prominent.

Wrapping up, Christensen's study starts out by acknowledging the indisputable trajectory of hard drives towards greater capacity. A closer look at technical innovations and new market opportunities brings him to realize that a simple linear narrative would be inadequate. Notably, to summarize his historical narrative he draws a diagram where multiple "intersecting trajectories", one for each architecture, are represented, see Figure 6.5. Even if such a diagram drastically simplifies the complex process of technical change, it is surely more accurate than a single-trajectory narrative according to which technical change unfolds along a single dimension, capacity, and extant hard drives are supplanted by new more capacious models.

6.6. Conclusion

The learning hypothesis as formulated by Petroski and subscribed by many engineers appears deceptively simple: *in engineering more is learned from failures than from successes*. To reveal its internal workings required quite a bit of conceptual reverse engineering. First, its main components, the notions of *failure*, *success*, and *learning* are rather complex machinery themselves. Each one of them

admits of many definitions and their meanings are often discussed in the literature within and without engineering. Second, Petroski's elaborations on the learning hypothesis create a vast array of connections spread over the whole history of engineering, from the Egyptian pyramids to the Space Shuttle, and encompassing virtually all disciplines. To prepare the ground for the analysis, two episodes from the history of suspension bridge engineering have been summarized in Section 6.2. As clarified in the following sections, the first episode pivoting on the contrasting approaches of Charles Ellet and John Roebling constitutes a case of *specific* learning; the second episode, spanning about one hundred years from the conception of the Brooklyn Bridge to the construction of the Severn Bridge exemplifies *generic* learning.

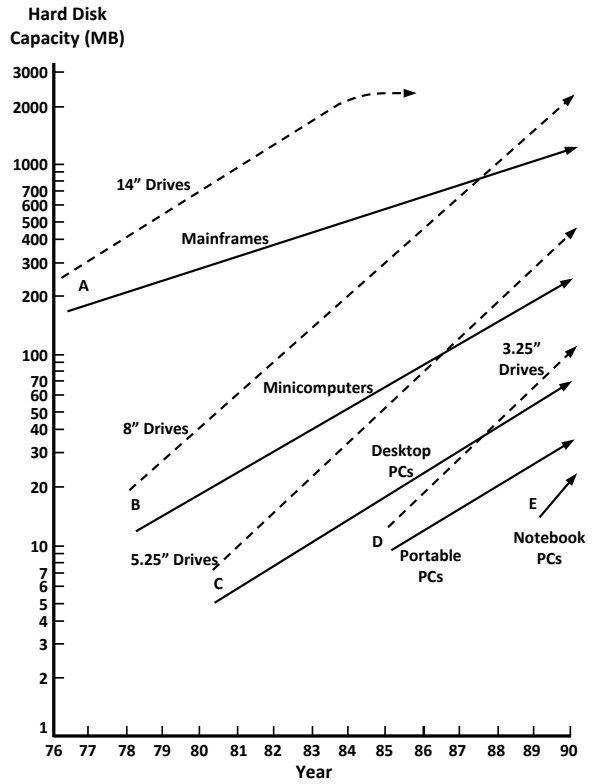

Figure 6.5: Representation of technological trajectories in the hard drive industry. Dashed arrows: trajectory of capacity increase by the 8-, 5.25-, and 3.5-Inch architectural generations; solid arrows: trajectories of capacity demanded in the Mainframe, Mini, Desktop PC, and Portable PC markets. From (Christensen: 1997)

In my analysis, I have attempted to balance the aim of conceptual clarity with the broadness in scope envisioned by Petroski. One early result of this approach has been the adoption of a *goal-oriented* notion of failure instead of an *event-oriented* one (Section 6.3). In turn, the concept of design goal has been explicated by borrowing the ideas of *goal tree* and assessment *metrics* from Dym and Little's textbook (2008). Similarly, the discussion in Section 6.4 has touched upon multiple aspects and interpretations of learning in engineering.

Starting from an intuitive understanding of learning as transmission of precompiled knowledge, the analysis progressed to consider learning as generation of knowledge as well the issue of multiple levels of learning. All these aspects converge into the concept of *lessons learned* where they merge with a fundamental characteristic of engineering knowledge that is to say, "the inseparability of knowledge and its practical application" (Vincenti: 1990, 207).

Section 6.5 begins by illustrating a conceptual tool developed by Vincenti in his paper on landing gear technology (1994). There, Vincenti distinguishes between a *generic-learning* process, which consists in the historic transition made by an engineering community from one set of technologies (i.e., retraction, trouser gears, wheel pants) to a different set (i.e., retraction is the dominant technology for high-performance airplanes), and a *specific-learning* processes consisting in individual engineers (or engineering organization) learning how to implement a technology which is, at the same time, implicated in generic learning. While Vincenti maintains that the two learning processes can be reconciled by his blind-variation selective-retention model, to Petroski they are instantiations of the same fundamental paradox of learning. Indeed, he appears to believe that no fundamental conceptual divide exists between them as proven by the extensive quote at the beginning of Section 6.5 where he seamlessly moves from one to the other.

In the rest of Section 6.5 I have shown that Petroski's belief is unwarranted. The concepts of failure, success, and learning can be combined, as illustrated in Figure 6.3, to produce a *specific-learning hypothesis* which is empirically testable. In fact, Madsen and Desai (2010) have recently published a study where they test a version of the hypothesis (i.e., the orbital launch vehicle version) and claim to have found evidence for it. However, when the same components are assembled together to generate a generic-learning hypothesis, the output is a drastically simplified representation of technical change that clashes with recent accounts offered by historians and sociologists of technology.

Although Petroski's paradoxical view of learning in engineering is only partly successful, an important lesson can be learned from the arguments he developed to defend it. It is tempting to think that failure can be avoided by focusing on achievement of success. Although it is safe to say that a successful product is one that does not meet with failures, pursuing success by closely following models of success can make engineers less responsive to valuable lessons from failures. Each design is different and there are no reliable procedures to tell engineers that a new design is well within the boundaries of a previous successful one. Apparently minor changes can introduce unforeseen behaviors in a product otherwise identical to its predecessor. Switching from a thinking process hooked on success to an approach of proactive search for vulnerabilities and weaknesses may greatly help engineers in finding those failure episodes that hold valuable lessons. There is no shortage of lessons learned in the engineering literature, yet to realize that a specific lesson does apply to the case at hand one needs to make a cognitive effort. It may be easier to see the connection if one is already alert to the possibility of failure.[24]

[24] I would like to thank Peter Kroes and Maarten Franssen for their support and valuable comments during the writing of this chapter. A shorter version of this paper was presented at the Fifth International Conference on Engineering Failure Analysis (ICEFA 2012, Den Haag, The Netherlands) where I was very pleased of the opportunity to discuss my paper and my research project with several experienced failure analysts. In particular, I wish to thank Richard Clegg (conference co-chair), Colin Gagg (conference co-chair), Emiel Amsterdam (conference co-chair), and Fabrizio D'Errico. I am grateful to Russell Wanhill for the generosity shown in sharing his knowledge and for his ability in clarifying complex technical matters. Special thanks go to Stan Lynch who provided stimulating comments and to whom I am indebted for pointing out the HMS *Alarm* historical case which introduces this paper.

Bibliography

Abdelhamid, T. S. and Everett, J. G. (2000) 'Identifying Root Causes of Construction Accidents', in: *Journal of Construction Engineering and Management* 126 (1): 52–60.

Affonso, L. O. A. (2006) *Machinery Failure Analysis Handbook*, Gulf Publishing Company, Houston, Tex.

Ale, B. (2009) *Risk: An Introduction*, Routledge, New York, NY.

Aliya, D. (2002) 'The Failure Analysis Process: An Overview', in: Becker, W. T. and Shipley, R. J. (eds.), *ASM Handbook, Vol. 11: Failure Analysis and Prevention*, ASM International, Materials Park, OH: 315–323.

American Institute of Chemical Engineers (2009) *Inherently Safer Chemical Processes: A Life Cycle Approach*. 2nd ed., Center for Chemical Process Safety, New York, NY.

Andersen, B. and Fagerhaug, T. (2006) *Root Cause Analysis: Simplified Tools and Techniques*, American Society for Quality.

Andersen, P. H. and Drejer, I. (2009) 'Together We Share? Competitive and Collaborative Supplier Interests in Product Development', in: *Technovation* 29 (10): 690–703.

Anderson, E. G. junior, Davis-Blake, A., Erzurumlu, S. S., Nitin, J. R., and Parker, G. G. (2008) 'Effects of Outsourcing, Offshoring and Distributed Product Development Organizations and Coordinating the NPD Process', in: Loch, C. H. and Kavadias, S. (eds.), *Handbook of New Product Development Management*, Butterworth-Heinemann/Elsevier, Oxford, UK: 259–289.

Andreasen, M. M. and Hein, L. (1987) *Integrated Product Development*, IFS, Bedford.

Apple Inc. (2010) Statement by Apple on White iPhone 4 [online] in: *Apple*. Available from: http://www.apple.com/pr/library/2010/07/23iphonestatement.html [Accessed: 5 Jan 2011].

——— (2011) Apple iPod classic technical specifications. [online] in: *Apple*. Available from: http://www.apple.com/ipodclassic/specs.html [Accessed: 5 Jan 2011].

Avizienis, A., Laprie, J.-C., Randell, B., and Landwehr, C. (2004) 'Basic Concepts and Taxonomy of Dependable and Secure Computing', in: *IEEE Transactions on Dependable and Secure Computing* 1 (1): 11–33.

Badenius, D. (1991) 'New Definitions of Basic R&M Terms', in: *Microelectronics Reliability* 31 (2-3): 525–535.

Barella, S., Cincera, S., Boniardi, M., Bellogini, M., Gelati, S., and Montanari, A. (2011) 'Failure Analysis of Tuna Cans', in: *Journal of Failure Analysis and Prevention* 11 (4): 446–451.

Barros, D. B. (2013) 'Negative Causation in Causal and Mechanistic Explanation', in: *Synthese* 190 (3): 449–469.

Bauer, J. E., Duffy, G. L., and Westcott, R. (2006) *The Quality Improvement Handbook*. 2nd ed., ASQ Quality Press, Milwaukee, Wis.

Bazu, M. and Bajenescu, T. (2011) *Failure Analysis: A Practical Guide for Manufacturers of Electronic Components and Systems*, John Wiley & Sons, Chichester, West Sussex, UK.

Becker, W. T. and Shipley, R. J., eds. (2002) *ASM Handbook, Volume 11: Failure Analysis and Prevention*. 10th ed., ASM International, Materials Park, OH.

Becker, W. T., Shipley, R. J., and Aliya, D. (2005) 'Use of the Term Defect', in: *Journal of Failure Analysis and Prevention* 5 (2): 16–20.

Behrendt, S., Jasch, C., Peneda, M. C., and van Weenen, H. (1997) *Life Cycle Design: A Manual for Small and Medium Sized Companies*, Springer, Berlin.

Bellgran, M. and Säfsten, K. (2010) *Production Development*, Springer, London, UK.

Benner, L. (1975) 'Accident Investigations: Multilinear Events Sequencing Methods', in: *Journal of Safety Research* 7 (2): 67–73.

Bhaumik, S. (2009) 'A View on the General Practice in Engineering Failure Analysis', in: *Journal of Failure Analysis and Prevention* 9 (3): 185–192.

Bieniawski, Z. T., Denkhaus, H. G., and Vogler, U. W. (1969) 'Failure of Fractured Rock', in: *International Journal of Rock Mechanics and Mining Sciences & Geomechanics Abstracts* 6 (3): 323–341.

Bijker, W. E. (1995) *Of Bicycles, Bakelites, and Bulbs: Toward a Theory of Sociotechnical Change*, MIT Press, Cambridge, Mass.

Birolini, A. (2007) *Reliability Engineering: Theory and Practice*. 5th ed., Springer, Berlin.

Blache, K. M. and Shrivastava, A. B. (1994) 'Defining Failure of Manufacturing Machinery and Equipment', in: *Proceedings of Annual Reliability and Maintainability Symposium (RAMS)*, Presented at the Annual Reliability and Maintainability Symposium (RAMS), Anaheim, CA, USA: 69–75.

Blischke, W. R. and Murthy, D. N. P. (2000) *Reliability: Modeling, Prediction, and Optimization*, John Wiley and Sons, New York, NY.

Blockley, D. I. and Henderson, J. R. (1980) 'Structural Failures and the Growth of Engineering Knowledge', in: *ICE Proceedings* 68 (4): 719–728.

Borgo, S. and Leitão, P. (2007) 'Foundations for a Core Ontology of Manufacturing', in: Sharman, R., Kishore, R., and Ramesh, R. (eds.), *Ontologies: a handbook of principles, concepts and applications in information systems*, Springer: 751–775.

Borgo, S. and Vieu, L. (2009) 'Artefacts in Formal Ontology', in: Meijers, A. (ed.), *Handbook of Philosophy of Technology and Engineering Sciences*, 273–308.

Borst, W. N. (1997) Construction of engineering ontologies for knowledge sharing and reuse, University of Twente, Enschede, The Netherlands.

Bottazzi, E. and Ferrario, R. (2005) 'A Path to an Ontology of Organizations', in: *Proceedings of International EDOC Workshop on Vocabularies, Ontologies and Rules for The Enterprise*, Presented at the VORTE 2005, Centre for Telematics and Information Technology, University of Twente 2005, Enschede, The Netherlands: 9–16.

——— (2011) 'Faulty Institutional Objects. A Threat for the Infallibilist (and the Fallibilist as Well)', in: *Seventh European Conference of Analytic Philosophy*, Presented at the ECAP 7, Milan.

Le Bourhis, E. (2008) *Glass: Mechanics and Technology*, Wiley-VCH, Weinheim.

Boyle, T. A., Kumar, V., and Kumar, U. (2006) 'Concurrent Engineering Teams II: Performance Consequences of Usage', in: *Team Performance Management* 12 (5/6): 125–137.

Bradsher, K. (1997) 'Light Trucks, Heavy Risk: A Special Report; A Deadly Highway Mismatch Ignored', in: *The New York Times*, 24 Sepp. A1.

——— (2000) 'Study of Ford Explorer's Design Reveals a Series of Compromises', in: *The New York Times*, 7 Dec.

——— (2002) *High and Mighty: SUVs – The World's Most Dangerous Vehicles and How They Got That Way*, Public Affairs, New York, NY.

Buonopane, S. G. and Billington, D. P. (1993) 'Theory and History of Suspension Bridge Design from 1823 to 1940', in: *Journal of Structural Engineering* 119 (3): 954–977.

Busby, J. S. (2001) 'Characterizing Failures in Design Activity', in: *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture* 215 (10): 1417–1424.

Buys, J. R. and Clark, J. L. (1995) Events and Causal Factors Analysis. No. SCIE-DOE-01-TRAC-14-95, Scientech, Inc., Idaho Falls.

Cannon, M. D. and Edmondson, A. C. (2005) 'Failing to Learn and Learning to Fail (intelligently): How Great Organizations Put Failure to Work to Innovate and Improve', in: *Long Range Planning* 38 (3 SPEC. ISS.): 299–319.

Cantini, A. (2012) 'Paradoxes and Contemporary Logic', in: Zalta, E. N. (ed.), *The Stanford Encyclopedia of Philosophy*, The Metaphysics Research Lab, Stanford, CA.

Carlson, J. and Murphy, R. R. (2005) 'How UGVs Physically Fail in the Field', in: *IEEE Transactions on Robotics* 21 (3): 423–437.

Carper, K. L., ed. (2001) *Forensic Engineering*. 2nd ed., CRC Press, Boca Raton, FL.

Carroll, J. S. (1995) 'Incident Reviews in High-Hazard Industries: Sense Making and Learning Under Ambiguity and Accountability', in: *Organization & Environment* 9 (2): 175–197.

Carroll, J. S. and Fahlbruch, B. (2011) '"The Gift of Failure: New Approaches to Analyzing and Learning from Events and near-Misses." Honoring the Contributions of Bernhard Wilpert', in: *Safety Science* 49 (1): 1–4.

CCPS (2003) *Guidelines for Investigating Chemical Process Incidents*. 2nd ed., American Institute of Chemical Engineers, New York, NY.

Chillarege, R. (1996) 'What Is Software Failure?', in: *IEEE Transactions on Reliability* 45 (3): 354–355.

Christensen, C. M. (1997) *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business Press.

Clark, K. B. (1991) *Product Development Performance: Strategy, Organization, and Management in the World Auto Industry*, Harvard Business School Press, Boston, Mass.

Collings, D. (2008) 'Lessons from Historical Bridge Failures', in: *Proceedings of the Institution of Civil Engineers: Civil Engineering* 161 (SPEC. ISS. 2): 20–27.

Collins, J. A. (1993) *Failure of Materials in Mechanical Design: Analysis, Prediction, Prevention*. 2nd ed., Wiley, New York.

——— (2003) *Mechanical Design of Machine Elements and Machines: A Failure Prevention Perspective*, Wiley, New York, NY.

Collins, J. A. and Daniewicz, S. R. (2006) 'Failure Modes: Performance and Service Requirements for Metals', in: Kutz, M. (ed.), *Mechanical engineers' handbook – Materials and mechanical design*, Wiley, Hoboken, N.J. 860–924.

Collins, J. A., Hagan, B., and Bratt, H. M. (1976) 'The Failure-Experience Matrix: A Useful Design Tool', in: *Journal of Engineering for Industry* 98 (3): 1074–1079.

Cowan, R., Fauchart, E., Foray, D., and Gunby, P. (2006) 'Learning from Disaster', in: Pyka, A. and Hanusch, H. (eds.), *Applied Evolutionary Economics And the Knowledge-based Economy*, Edward Elgar Publishing: 40–70.

Cowan, R. S. (1985) 'How the Refrigerator Got Its Hum', in: MacKenzie, D. A. and Wajcman, J. (eds.), *The Social shaping of technology: how the refrigerator got its hum*, Open University Press: 202–218.

Cox, S. (1998) *Safety, Reliability, and Risk Management: An Integrated Approach*. 2nd ed., Butterworth-Heinemann, Oxford.

Le Coze, J. (2008) 'Disasters and Organisations: From Lessons Learnt to Theorising', in: *Safety Science* 46 (1): 132–149.

Crawford, C. M. (1992) 'The Hidden Costs of Accelerated Product Development', in: *Journal of Product Innovation Management* 9 (3): 188–199.

Cummins, R. (1975) 'Functional Analysis', in: *The Journal of Philosophy* 72 (20): 741–765.

Dale, B. G., ed. (2003) *Managing Quality*, Blackwell Publishing, Malden, MA.

Daley, D. T. (2008) *The Little Black Book of Reliability Management*, Industrial Press, New York, NY.

Dasgupta, A. and Pecht, M. G. (1991) 'Material Failure Mechanisms and Damage Models', in: *IEEE Transactions on Reliability* 40 (5): 531–536.

Davis, J. R., ed. (1992) *ASM Materials Engineering Dictionary*, ASM International, Materials Park, OH.

Dekker, S. (2005) *Ten Questions About Human Error: A New View of Human Factors and System Safety*, Lawrence Erlbaum Associates, Mahwah, N.J.

Dennies, D. P. (2002) 'The Organization of a Failure Investigation', in: *Journal of Failure Analysis and Prevention* 2 (3): 11–16.

Department of Energy (1992) DOE Guideline: Root Cause Analysis Guidance Document. Guideline, Department of Energy, Washington, DC.

Department of the Air Force (2008) Air Force Instruction 91-204. Safety Investigations and Reports. Instruction No. 91-204, Department of the Air Force, Washington, DC.

Dodson, B. (1999) *Reliability Engineering Handbook*, Marcel Dekker, New York, NY.

Doerner, D. (1980) 'On the Difficulties People Have in Dealing With Complexity', in: *Simulation & Gaming* 11 (1): 87–106.

Dougherty, D. (1992) 'Interpretive Barriers to Successful Product Innovation in Large Firms', in: *Organization Science* 3 (2): 179–202.

Dym, C. L. and Little, P. (2008) *Engineering Design: A Project Based Introduction*. 3rd ed., Wiley.

Egan, G. R. (2006) 'The Significance of Defects in Welded Long-Span Bridge Structures', in: *Annals of the New York Academy of Sciences* 352 (1): 177–191.

Ehrlenspiel, K. (1995) *Integrierte Produktentwicklung: Methoden fuer Prozessorganisation, Produkterstellung und Konstruktion*, Carl Hanser Verlag, München.

Erden, M. S., Komoto, H., Van Beek, T. J., D'Amelio, V., Echavarria, E., and Tomiyama, T. (2008) 'A Review of Function Modeling: Approaches and Applications', in: *Artificial Intelligence for Engineering Design, Analysis and Manufacturing: AIEDAM* 22 (2): 147–169.

Ezrin, M. (1996) *Plastics Failure Guide: Cause and Prevention*, Hanser Verlag, Munich.

Fashandi, A. and Umberg, T. (2003) 'Equipment Failure Definition: A Prerequisite for Reliability Test and Validation', in: *Proceedings of the IEEE/CPMT International Electronics Manufacturing Technology (IEMT) Symposium*, 357–358.

Feld, J. and Carper, K. L. (1997) *Construction Failure*, John Wiley and Sons, New York, NY.

Ferjencik, M. (2010) 'Root Cause Analysis of an Old Accident in an Explosives Production Plant', in: *Safety Science* 48 (2): 1530–1544.

Ferrario, R. and Guarino, N. (2009) 'Towards an Ontological Foundation for Services Science', in: Domingue, J., Fensel, D., and Traverso, P. (eds.), *Future Internet - FIS 2008*, Springer, Berlini: 152–169.

Ferry, T. S. (1988) *Modern Accident Investigation and Analysis*, Wiley-IEEE.

Finlow-Bates, T. (1998) 'The Root Cause Myth', in: *TQM Magazine* 10 (1): 10–15.

Fischhoff, B. (1975) 'Hindsight Is Not Equal to Foresight: The Effect of Outcome Knowledge on Judgment under Uncertainty', in: *Journal of Experimental Psychology: Human Perception and Performance* 1 (3): 288–299.

Ford, D. N. and Sterman, J. D. (2003) 'The Liar's Club: Concealing Rework in Concurrent Development', in: *Concurrent Engineering* 11 (3): 211–219.

Fortune, J. and Peters, G. (1995) *Learning from Failure. The Systems Approach*. 1st ed., John Wiley & Sons, New York, NY.

Del Frate, L. (2012) 'Preliminaries to a Formal Ontology of Failure of Engineering Artifacts', in: Donnelly, M. and Guizzardi, G. (eds.), *Formal Ontology in Information Systems: Proceedings of the Seventh International Conference (FOIS 2012)*, Presented at the FOIS 2012, IOS Press, Graz, Austria: 117–130.

Del Frate, L., Franssen, M., and Vermaas, P. E. (2011) 'Towards a Trans-Disciplinary Concept of Failure for Integrated Product Development', in: *International Journal of Product Development* 14 (1-4): 72–95.

Frawley, D. J. (2002) *ISO 9001 QMS. Policies, Procedures and Forms*, Bizmanualz.com, Inc., St. Louis, Mo.

Friedman, V. J., Lipshitz, R., and Popper, M. (2005) 'The Mystification of Organizational Learning', in: *Journal of Management Inquiry* 14 (1): 19–30.

Gagg, C. R. and Lewis, P. R. (2007) 'Wear as a Product Failure Mechanism – Overview and Case Studies', in: *Engineering Failure Analysis* 14 (8): 1618–1640.

——— (2009) 'In-Service Fatigue Failure of Engineered Products and Structures – Case Study Review', in: *Engineering Failure Analysis* 16 (6): 1775–1793.

Garvin, D. A. (1984) 'What Does "Product Quality" Really Mean?', in: *Sloan Management Review* 26 (1): 25–43.

Gordon, H. (2008) 'Integrating Learning into Safety', in: *Professional Safety* 53 (9): 30–34.

Graça, M. L. A., Hoo, C. Y., Silva, O. M. M., and Lourenço, N. J. (2009) 'Failure Analysis of a 300M Steel Pressure Vessel', in: *Engineering Failure Analysis* 16 (1): 182–186.

Grant, R. (1996) 'Toward a Knowledge-Based Theory of the Firm', in: *Strategic Management Journal* 17: 109–122.

Grantham Lough, K. A., Stone, R. B., and Tumer, I. Y. (2008) 'Failure Prevention in Design Through Effective Catalogue Utilization of Historical Failure Events', in: *Journal of Failure Analysis and Prevention* 8 (5): 469–481.

Griffin, A. (1997) 'PDMA Research on New Product Development Practices: Updating Trends and Benchmarking Best Practices', in: *Journal of Product Innovation Management* 14 (6): 429–458.

Grimvall, G., Holmgren, Å. J., Jacobsson, P., and Thedéen, T., eds. (2010) *Risks in Technological Systems*, Springer, London.

Groot Boerle, D. J. (2002) 'EMC and Functional Safety, Impact of IEC 61000-1-2', in: *IEEE International Symposium on Electromagnetic Compatibility, EMC 2002*, Presented at the EMC 2002, 353–358.

Guarino, N., Oberle, D., and Staab, S. (2009) 'What Is an Ontology?', in: Staab, S. and Studer, R. (eds.), *Handbook on Ontologies*, Springer, Berlin: 1–17.

Gutting, G. (1984) 'Paradigms, Revolutions, and Technology', in: Laudan, R. (ed.), *The nature of technological knowledge. Are models of scientific change relevant*, Springer: 47–65.

Haasl, D. F. (1965) 'Advanced Concepts in Fault Tree Analysis', in: *System Safety Symposium*, Presented at the System Safety Symposium, Seattle, Wash.

Halpern, J. Y. and Pearl, J. (2005) 'Causes and Explanations: A Structural-Model Approach. Part I: Causes', in: *The British Journal for the Philosophy of Science* 56 (4): 843–887.

Haque, B. (2003) 'Problems in Concurrent New Product Development: An in-Depth Comparative Study of Three Companies', in: *Integrated Manufacturing Systems* 14 (3): 191–207.

Harazaki, I., Suzuki, and Okukawa (2000) 'Suspension Bridges', in: Chen, W. F. and Duan, L. (eds.), *Bridge Engineering Handbook*, Taylor & Francis.

Harland, D. M. and Lorenz, R. (2005) *Space Systems Failures*, Praxis, New York, NY.

Harris, J. R. (1966) 'Copper and Shipping in the Eighteenth Century', in: *The Economic History Review* 19 (3): 550–568.

Hart, G. C. (1982) *Uncertainty Analysis, Loads, and Safety in Structural Engineering*, Prentice-Hall, Englewood Cliffs, NJ.

Hart, H. L. A. and Honoré, T. (1985) *Causation in the Law*, Oxford University Press.

Hatamura, Y., ed. (2008) *Learning From Design Failures*, Springer, New York, NY.

Hattangadi, A. A. (2000) *Electrical Fires and Failures: A Prevention and Troubleshooting Guide*, McGraw-Hill, New York, NY.

Hausman, D. M. (1998) *Causal Asymmetries*, Cambridge University Press, Cambridge, UK.

Hedlund, G. and Nonaka, I. (1993) 'Models of Knowledge Management in the West and Japan', in: Lorange, P., Chakravarthy, B., Roos, J., and van de Ven, A. (eds.), *Implementing Strategic Processes: Change, Learning, and Co-Operation*, Blackwell Business, Oxford: 117–144.

Heinrich, H. W. (1980) *Industrial Accident Prevention: A Safety Management Approach*. 5th ed., McGraw-Hill, New York.

Helper, S. and Sako, M. (1995) 'Supplier Relations in Japan and the United States: Are They Converging?', in: *Sloan Management Review* 36 (3): 77–84.

Henderson, R. M. and Clark, K. B. (1990) 'Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms', in: *Administrative Science Quarterly* 35: 30–39.

Hendrick, K. and Benner, L. (1987) *Investigating Accidents with STEP*, CRC Press.

Henshaw, J. M., Wood, V., and Hall, A. C. (1999) 'Failure of Automobile Seat Belts Caused by Polymer Degradation', in: *Engineering Failure Analysis* 6 (1): 13–25.

Hjort, H., Hananel, D., and Lucas, D. (1992) 'Quality Function Deployment and Integrated Product Development', in: *Journal of Engineering Design* 3 (1): 17–29.

Hohns, M. (1985) 'Learning from Failures: Procedural Changes in the Design and Construction Process to Reduce Failures', in: *Reducing Failures of Engineered Facilities*, Presented at the Reducing Failures of Engineered Facilities, American Society of Civil Engineers, New York, NY: 75–83.

Hokstad, P. and Rausand, M. (2008) 'Common Cause Failure Modeling: Status and Trends', in: Misra, K. B. (ed.), *Handbook of Performability Engineering*, Springer-Verlag, London: 621–640.

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

Hollnagel, E., Nemeth, C. P., and Dekker, S., eds. (2008) *Resilience Engineering Perspectives: Remaining Sensitive to the Possibility of Failure*, Ashgate, Aldershot, UK.

Hoopes, D. G. and Postrel, S. (1999) 'Shared Knowledge, Glitches, and Product Development Performance', in: *Strategic Management Journal* 20 (9): 837–865.

Hopkins, A. (2008) *Failure to Learn: The BP Texas City Refinery Disaster*, CCH Australia Limited.

Houkes, W. and Vermaas, P. E. (2010) *Technical Functions: On the Use and Design of Artefacts*, Springer, Dordrecht, New York.

Hovden, J., Størseth, F., and Tinmannsvik, R. K. (2011) 'Multilevel Learning from Accidents - Case Studies in Transport', in: *Safety Science* 49 (1): 98–105.

Hubka, V. and Eder, W. E. (1996) *Design Science*, Springer, Berlin.

Hull, E., Jackson, K., and Dick, J. (2010) *Requirements Engineering*. 3rd ed., Springer, London.

Hummerdal, D., Wilhelmsson, A., and Dekker, S. (2013) 'Learning from Failure', in: Lee, J. D. and Kirlik, A. (eds.), *The Oxford Handbook of Cognitive Engineering*, Oxford University Press: 404–412.

Iansiti, M. (1995) 'Technology Integration: Managing Technological Evolution in a Complex Environment', in: *Research Policy* 24 (4): 521–542.

ICAO (2009) *Safety Management Manual*. 2nd ed., International Civil Aviation Organization, Montreal, Canada.

IEC 60050(191) (1990) *International Electrotechnical Vocabulary (IEV), Chapter 191 – Dependability and Quality of Service*, International Electrotechnical Commission, Genève, Switzerland.

IEC 60812 (2006) *Analysis Techniques for System Reliability - Procedure for Failure Mode and Effects Analysis (FMEA)*, International Electrotechnical Commission, Genève, Switzerland.

IEEE 1220 (2005) *Standard for Application and Management of the Systems Engineering Process*, Institute of Electrical and Electronics Engineers, New York, NY.

IEEE 610.12 (1990) *Standard Glossary of Software Engineering Terminology*, Institute of Electrical and Electronics Engineers, New York, NY.

Isermann, R. (2005) *Fault-Diagnosis Systems: An Introduction from Fault Detection to Fault Tolerance*. 1st ed., Springer, Berlin.

Ishii, K. (1995) 'Life-Cycle Engineering Design', in: *Journal of Mechanical Design* 117 (B): 42–47.

ISO 9001 (2008) *Quality Management Systems – Requirements*. 4th ed., International Organization for Standardization, Genève, Switzerland.

ISO/IEC 15288 (2008) *Systems and Software Engineering – System Life Cycle Processes*. 2nd ed., International Organization for Standardization, Genève, Switzerland.

ISO/IEC TR 24748-1 (2010) *Systems and Software Engineering – Life Cycle Management – Part 1: Guide for Life Cycle Management*. 2nd ed., International Organization for Standardization, Genève, Switzerland.

Jacobsson, A., Ek, Å., and Akselsson, R. (2011) 'Method for Evaluating Learning from Incidents Using the Idea of "Level of Learning"', in: *Journal of Loss Prevention in the Process Industries* 24 (4): 333–343.

Janssen, M., Zuidema, J., and Wanhill, R. J. H. (2004) *Fracture Mechanics*, Taylor & Francis.

Jespersen, B. and Carrara, M. (2011) 'Two Conceptions of Technical Malfunction', in: *Theoria* 77 (2): 117–138.

Job, M. (1996) *Air Disaster. Volume 2*, Aerospace Publications, Weston Creek, Australia.

Johnson, C. W. (2003) *Failure in Safety-Critical Systems: A Handbook of Accident and Incident Reporting*, University of Glasgow Press, Glasgow.

Johnson, W. G. (1980) *MORT Safety Assurance Systems*, Marcel Dekker, New York, NY.

Jones, S. S. (2004) *Basics of Telecommunications*. 5th ed., International Engineering Consortium, Chicago, Ill.

Jovanovic, B. and Rob, R. (1989) 'The Growth and Diffusion of Knowledge', in: *The Review of Economic Studies* 56 (4): 569–582.

Kahn, K. B., ed. (2005) *The PDMA Handbook of New Product Development*, Wiley, Hoboken, NJ.

Kaminetzky, D. (1991) *Design and Construction Failures: Lessons from Forensic Investigations*, McGraw-Hill, New York, NY.

Kaplan, S. (1990) 'Bayes Is for Eagles', in: *IEEE Transactions on Reliability* 39 (2): 130–131.

Kawada, T. (2010) *History of the Modern Suspension Bridge: Solving the Dilemma Between Economy and Stiffness*, ASCE Press.

Keller, R. T. (2001) 'Cross-Functional Project Groups in Research and New Product Development: Diversity, Communications, Job Stress, and Outcomes', in: *The Academy of Management Journal* 44 (3): 547–555.

Kieselbach, R. (1997) 'Bursting of a Silo', in: *Engineering Failure Analysis* 4 (1): 49–55.

King, L. W. (1915) *The Code of Hammurabi*.

Kinnersley, S. and Roelen, A. (2007) 'The Contribution of Design to Accidents', in: *Safety Science* 45 (1-2): 31–60.

Kitamura, Y. and Mizoguchi, R. (1999) 'An Ontological Analysis of Fault Process and Category of Faults', in: *Proceedings of tenth international workshop on principles of diagnosis*, Presented at the DX-99, Loch Awe, Scotland: 118–128.

Kletz, T. A. (1998) *Process Plants: A Handbook for Inherently Safer Design*, Taylor & Francis, Philadelphia, PA.

Koji, Y., Kitamura, Y., and Mizoguchi, R. (2005) 'Ontology-Based Transformation from an Extended Functional Model to FMEA', in: *Proceedings of the International Conference on Engineering Design*, Presented at the ICED 2005, Melbourne, Australia.

Kortge, G. D. and Okonkwo, P. A. (1989) 'Simultaneous New Product Development: Reducing the New Product Failure Rate', in: *Industrial Marketing Management* 18 (4): 301–306.

Kuntz, M., Leitner-Fischer, F., and Leue, S. (2011) 'From Probabilistic Counterexamples via Causality to Fault Trees', in: Flammini, F., Bologna, S., and Vittorini, V. (eds.), *Computer Safety, Reliability, and Security*, Springer, Berlin: 71–84.

Ladkin, P. (2000) 'Causal Reasoning about Aircraft Accidents', in: Koornneef, F. and Meulen, M. van der (eds.), *SAFECOMP 2000*, Springer, Berlin: 344–360.

Lai, R. (2010) White iPhone 4 delay: the challenges faced by Apple's glass supplier [online] in: *Engadget*. Available from: http://www.engadget.com/2010/07/18/white-iphone-4-delay-the-challenges-faced-by-apples-glass-supp/ [Accessed: 5 Jan 2011].

Lang, S. Y. T., Dickinson, J., and Buchal, R. O. (2002) 'Cognitive Factors in Distributed Design', in: *Computers in Industry* 48 (1): 89–98.

Laprie, J. C. (1985) 'Dependable Computing and Fault-Tolerance', in: *Digest of Papers FTCS-15* 2–11.

Latin, H. and Kasolas, B. (2002) 'Bad Designs, Lethal Profits: The Duty to Protect Other Motorists against SUV Collision Risks', in: *Boston University Law Review* 82: 1161–1223.

Latino, R. J. and Latino, K. C. (2006) *Root Cause Analysis: Improving Performance for Bottom-Line Results*. 3rd ed., CRC, Boca Raton, FL.

Leonards, G. (1982) 'Investigation of Failures', in: *Journal of the Geotechnical Engineering Division* 108 (2): 185–246.

Levin, M. and Kalal, T. T. (2003) *Improving Product Reliability: Strategies and Implementation*, John Wiley and Sons, New York, NY.

Lewis, D. K. (1973) *Counterfactuals*, Harvard University Press, Cambridge.

Lewis, P. R. (2000) *Polymer Product Failure*, iSmithers Rapra Publishing, Shrewsbury, UK.

Lewis, P. R., Reynolds, K., and Gagg, C. R. (2003) *Forensic Materials Engineering: Case Studies*, CRC, Boca Raton, FL.

Liker, J. K., Sobek, D. K., Ward, A. C., and Cristiano, J. J. (1996) 'Involving Suppliers in Product Development in the United States and Japan: Evi-

dence for Set-Based Concurrent Engineering', in: *IEEE Transactions on Engineering Management* 43 (2): 165–178.

Łukowski, P. (2011) *Paradoxes*, Springer, New York, NY.

Lullies, V. (2000) 'Knowledge Management Is the Key Prerequisite for the Improvement of New Product and Process Development', in: Jürgens, U. (ed.), *New Product Development and Production Networks*, Springer, Berlin: 427–439.

MacIntosh, R. (2010) 'The Accident "CAUSE" Statement – Is It beyond Its Time?', in: *ISASI Forum* 43 (2): 5–9.

Mackie, J. L. (1974) *The Cement of the Universe; a Study of Causation*, Clarendon Press, Oxford.

Madsen, P. M. and Desai, V. (2010) 'Failing to Learn? The Effects of Failure and Success on Organizational Learning in the Global Orbital Launch Vehicle Industry', in: *Academy of Management Journal* 53 (3): 451–476.

Marks, A. P. (1989) 'The Sinclair C5 – An Investigation into Its Development, Launch, and Subsequent Failure', in: *European Journal of Marketing* 23 (1): 61–71.

Márquez, A. C. (2007) *The Maintenance Management Framework: Models and Methods for Complex Systems Maintenance*, Springer, London.

Masanet, E. and Horvath, A. (2007) 'Assessing the Benefits of Design for Recycling for Plastics in Electronics: A Case Study of Computer Enclosures', in: *Materials and Design* 28 (6): 1801–1811.

Le May, I. and Deckker, E. (2009) 'Reducing the Risk of Failure by Better Training and Education', in: *Engineering Failure Analysis* 16 (4): 1153–1162.

McDonough, E. F. (2000) 'Investigation of Factors Contributing to the Success of Cross-Functional Teams', in: *Journal of Product Innovation Management* 17 (3): 221–235.

McKinnon, R. C. (2000) *Cause, Effect, and Control of Accidental Loss with Accident Investigation Kit*, CRC Press, Boca Raton, FL.

Melchers, R. E. (1999) *Structural Reliability Analysis and Prediction*. 2nd ed., Wiley, New York, NY.

Merrison, A. W. (1973) *Inquiry into the Basis of Design and Method of Erection of Steel Box-Girder Bridges : Report of the Committee*, H.M. Stationery Office, London, UK.

Miller, M. C. and Guimaraes, T. (2005) 'Addressing Some HRM Issues to Improve Performance of Cross-Functional Teams in Concurrent Engineering', in: *Proceedings of the IEEE International Engineering Management Conference 2005*, Presented at the Engineering Management Conference 2005, Piscataway, NJ: 260–264.

MIL-STD-1629A (1980) *Procedures for Performing a Failure Mode, Effects and Criticality Analysis*, US Department of Defense, Washington, DC.

MIL-STD-721C (1981) *Definition of Terms for Reliability and Maintainability*, US Department of Defense, Washington, DC.

Mobley, R. K. (1999) *Root Cause Failure Analysis*, Newnes, Boston.

Mokyr, J. (2002) *The Gifts of Athena: Historical Origins of the Knowledge Economy*, Princeton University Press, Princeton, NJ.

Mueller, N. (2006) Van Halen Fell Silent On Top of the World [online] in: *Washingtonpost.com*. Available from: http://www.washingtonpost.com/wp-dyn/content/article/2006/10/21/AR2006102100113.html [Accessed: 7 Jan 2011].

Murphy, J. F. (2008) 'What I Learned as an Investigator with the CSB – Effective Investigations', in: *Process Safety Progress* 27 (4): 266–273.

My Digital Life (2010) iPhone 4 White Version Delays Due To Manufacturing Challenge [online] in: *My Digital Life*. Available from: http://www.mydigitallife.info/2010/07/26/iphone-4-white-version-delays-due-to-manufacturing-challenge/ [Accessed: 18 Mar 2011].

NASA (2006) 'NPR 8621.1B - NASA Procedural Requirements for Mishap and Close Call Reporting, Investigating, and Recordkeeping w/Change 5 (03/15/2010)',.

NATO (2008) *NATO Reliability & Maintainability Terminology Applicable to ARMPs (ARMP - 7)*, NATO, Brussels.

Nederlands Aviation Safety Board (1994) El Al Flight 1862. Boeing 747-258F 4X-AXG. Bijlmermeer, Amsterdam. October 4, 1992. Aircraft Accident Report No. 92-11, Nederlands Aviation Safety Board, Den Haag, NL.

Nieuwhof, G. W. E. (1984) 'The Concept of Failure in Reliability Engineering', in: *Reliability Engineering* 7 (1): 53–59.

NTSB (1990) United Airlines Flight 232 McDonnell Douglas DC-10-10 Sioux Gateway Airport. Sioux City, Iowa. July 19, 1989. Aircraft Accident Report No. NTSB/AAR-SO/06, National Transportation Safety Board, Washington, DC.

O'Connor, P. D. T., Newton, D., and Bromley, R. (2002) *Practical Reliability Engineering*, John Wiley and Sons, New York, NY.

Olsen, G. R., Cutkosky, M., Tenenbaum, J. M., and Gruber, T. R. (1994) 'Collaborative Engineering Based on Knowledge Sharing Agreements', in: *Procedings of the 1994 ASME Database Symposium*, Minneapolis, MN: 1–12.

Olsson, K. G. F. (1976) *Systematisk konstruktion: En studie med syfte att systematisera innehåll och metoder i samband med produktkonstruktion*, Institutionen för Maskinkonstruktion, Lunds Tekniska Högskola, Lund, Sweden.

Pahl, G., Beitz, W., Feldhusen, J., and Grote, K.-H. (2007) *Engineering Design: A Systematic Approach*. 3rd ed., Springer, London.

Paradies, M. and Busch, D. (1988) 'Root Cause Analysis at Savannah River Plant', in: *IEEE Conference on Human Factors and Power Plants*, 479–483.

Pecht, M. G. (2006) 'Establishing a Relationship Between Warranty and Reliability', in: *Electronics Packaging Manufacturing, IEEE Transactions on* 29 (3): 184–190.

Perrin, E., Kirwan, B., and Stroup, R. (2006) 'A Systemic Model of ATM Safety: The Integrated Risk Picture', in: *Risk Analysis and Safety Performance in Aviation*, Presented at the Conference on Risk Analysis and Safety Performance in Aviation, Atlantic City, NJ.

Petroski, H. (1985) *To Engineer Is Human: The Role of Failure in Successful Design*. 1st ed., St. Martin's Press, New York, NY.

——— (1994) *Design Paradigms: Case Histories of Error and Judgment in Engineering*, Cambridge University Press, Cambridge, UK.

——— (1996) *Invention by Design: How Engineers Get from Thought to Thing*, Harvard University Press, Cambridge, Mass.

——— (2001) 'Success and Failure in Engineering', in: *Journal of Failure Analysis and Prevention* 1 (5): 8–15.

——— (2006) *Success through Failure: The Paradox of Design*, Princeton University Press, Princeton.

——— (2011) *An Engineer's Alphabet: Gleanings from the Softer Side of a Profession*, Cambridge University Press, New York, NY.

Pfaender, H. G. (1996) *Schott Guide to Glass*, Chapman & Hall, London.

Piésold, D. D. A. (1991) *Civil Engineering Practice: Engineering Success by Analysis of Failure*, McGraw-Hill, London.

Pinto, M. B., Pinto, J. K., and Prescott, J. E. (1993) 'Antecedents and Consequences of Project Team Cross-Functional Cooperation', in: *Management Science* 39 (10): 1281–1297.

Plumbridge, W. J. (2009) 'New Avenues for Failure Analysis', in: *Engineering Failure Analysis* 16 (5): 1347–1354.

Poteet, S., Patel, J., Giammanco, C., Whiteley, I., Xue, P., and Kao, A. (2008) 'Words Are Mightier Than Swords... and Yet Miscommunication Costs Lives!', in: *Proceedings of the Second Annual Conference of the International Technology Alliance*, Presented at the ACITA'08, London.

Prasad, B. (1996) 'Toward Definitions of a Concurrent Product Design, Development, and Delivery (PD 3) System', in: *Concurrent Engineering* 4 (2): 102–109.

Prasad, D., McDermid, J., and Wand, I. (1996) 'Dependability Terminology: Similarities and Differences', in: *Aerospace and Electronic Systems Magazine, IEEE* 11 (1): 14–21.

Pugsley, A. G. (1966) *The Safety of Structures*, Edward Arnold, London.

Qi, H., Ganesan, S., and Pecht, M. G. (2008) 'No-Fault-Found and Intermittent Failures in Electronic Products', in: *Microelectronics Reliability* 48 (5): 663–674.

Ramachandran, V., Raghuram, A. C., Krishnan, R. V., and Bhaumik, S. K. (2005) *Failure Analysis of Engineering Structures : Methodology and Case Histories*, ASM International, Materials Park, OH.

Ramesh, B. and Tiwana, A. (1999) 'Supporting Collaborative Process Knowledge Management in New Product Development Teams', in: *Decision Support Systems* 27 (1-2): 213–235.

Ratay, R. (2009) *Forensic Structural Engineering Handbook*. 2nd ed., McGraw Hill Professional, New York, NY.

Rauniar, R., Doll, W., Rawski, G., and Hong, P. (2008) 'Shared Knowledge and Product Design Glitches in Integrated Product Development', in: *International Journal of Production Economics* 114 (2): 723–736.

Rausand, M. and Høyland, A. (2004) *System Reliability Theory: Models, Statistical Methods, and Applications*. 2nd ed., Wiley-Interscience, Hoboken, NJ.

Rausand, M. and Øien, K. (1996) 'The Basic Concepts of Failure Analysis', in: *Reliability Engineering and System Safety* 53 (1): 73–83.

Reason, J. (1990) *Human Error*, Cambridge University Press, Cambridge, UK.

——— (2008) *The Human Contribution: Unsafe Acts, Accidents and Heroic Recoveries*, Ashgate, Farnham, UK.

Rees, R. (1997) 'What Is a Failure', in: *IEEE Transactions on Reliability* 46 (2): 163.

Riley, W. F., Sturges, L. D., and Morris, D. H. (2006) *Mechanics of Materials*. 6th ed., Wiley, New York, NY.

Roche, C. (2000) 'Corporate Ontologies and Concurrent Engineering', in: *Journal of Materials Processing Technology* 107 (1-3): 187–193.

Roe, C. L. (1996) 'Project Management and Systems Engineering in an IPD Environment', in: *INCOSE 1996 – 6th Annual International Symposium Proceedings*, Presented at the INCOSE 1996, INCOSE, Seattle, Wash.

Roelen, A. L. C., Lin, P. H., and Hale, A. R. (2011) 'Accident Models and Organisational Factors in Air Transport: The Need for Multi-Method Models', in: *Safety Science* 49 (1): 5–10.

Rösler, J., Harders, H., and Bäker, M. (2007) *Mechanical Behaviour of Engineering Materials: Metals, Ceramics, Polymers, and Composites*, Springer, Berlin.

Ross, B., McDonald, B., and Vijay Saraf, S. E. (2007) 'Big Blue Goes Down. The Miller Park Crane Accident', in: *Engineering Failure Analysis* 14 (6): 942–961.

SAE J1739 (2002) *Potential Failure Mode and Effects Analysis in Design (Design FMEA), Potential Failure Mode and Effects Analysis in Manufacturing and Assembly Processes (Process FMEA), and Potential Failure Mode and Effects Analysis for Machinery (Machinery FMEA)*, Society of Automotive Engineers, Warrendale, PA.

Sage, A. P. (1995) 'Systems Engineering and Systems Management for Reengineering', in: *Journal of Systems and Software* 30 (1-2): 3–25.

Scherp, A., Saathoff, C., Franz, T., and Staab, S. (2011) 'Designing Core Ontologies', in: *Applied Ontology* 6 (3): 177–221.

Schlager, N. (1994) *When Technology Fails : Significant Technological Disasters, Accidents, and Failures of the Twentieth Century*, Gale Research, Detroit.

Scott, R. (2001) *In the Wake of Tacoma: Suspension Bridges and the Quest for Aerodynamic Stability*, ASCE Publications, Reston, Va.

Scutti, J. J. (2002) 'Introduction to Failure Analysis and Prevention', in: Becker, W. T. and Shipley, R. J. (eds.), *ASM Handbook. Vol 11: Failure Analysis and Prevention*, ASM International, Materials Park, OH: 3–23.

Scutti, J. J. and Aliya, D., eds. (2000) *Failure Prevention through Education: Getting to the Root Cause*, ASM International, Materials Park, OH.

Sfard, A. (1998) 'On Two Metaphors for Learning and the Dangers of Choosing Just One', in: *Educational Researcher* 27 (2): 4–13.

Sheridan, T. B. (2008) 'Risk, Human Error, and System Resilience: Fundamental Ideas', in: *Human Factors* 50 (3): 418–426.

Sherr, I. (2011) Apple Says White iPhone 4 Is Coming This Spring [online] in: *WSJ.com*. Available from: http://blogs.wsj.com/digits/2011/04/14/apple-says-white-iphone-4-is-coming-this-spring/ [Accessed: 14 May 2011].

Sims, B. (1999) 'Concrete Practices: Testing in an Earthquake-Engineering Laboratory', in: *Social Studies of Science* 29 (4): 483–518.

Sitkin, S. B. (1992) 'Learning through Failure: The Strategy of Small Losses', in: *Research in Organizational Behavior* 14: 231–266.

Slegers, N. J., Kadish, R. T., Payton, G. E., Thomas, J., Griffin, M. D., and Dumbacher, D. (2012) 'Learning from Failure in Systems Engineering: A Panel Discussion', in: *Systems Engineering* 15 (1): 74–82.

Stamatis, D. H. (1995) *Failure Mode and Effect Analysis: FMEA from Theory to Execution*, ASQ Quality Press, Milwaukee, Wisc.

Stone, R. B. and Wood, K. L. (2000) 'Development of a Functional Basis for Design', in: *Journal of Mechanical Design, Transactions of the ASME* 122 (4): 359–370.

Sudhakar, K. V. and Cruz Paredes, J. (2005) 'Failure Mechanisms in Motor Bearings', in: *Engineering Failure Analysis* 12 (1): 35–42.

Suess, M. E. (1992) 'Bacteria-Induced Corrosion of a Stainless Steel Chemical Trailer Barrel', in: Esaklul, K. A. (ed.), *Handbook of case histories in failure analysis*, ASM International: 70–73.

Syan, C. S. and Menon, U. (1994) *Concurrent Engineering: Concepts, Implementation and Practice*, Chapman & Hall, London.

Tam, A. S. . and Gordon, I. (2009) 'Clarification of Failure Terminology by Examining a Generic Failure Development Process', in: *International Journal of Engineering Business Management* 1 (1): 33–36.

Tawancy, H. M., Ul-Hamid, A., and Abbas, N. M. (2004) *Practical Engineering Failure Analysis*, M. Dekker, New York, NY.

Taylor, J. R. (2007) 'Understanding and Combating Design Error in Process Plant Design', in: *Safety Science* 45 (1-2): 75–105.

The Dutch Safety Board (2009) About the Safety Board - De Onderzoeksraad voor veiligheid [online] in: *The Dutch Safety Board*. Available from: http://www.onderzoeksraad.nl/en/index.php/over/ [Accessed: 20 Nov 2009].

Thomas, D. A., Avers, K., and Pecht, M. G. (2002) 'The "Trouble Not Identified" Phenomenon in Automotive Electronics', in: *Microelectronics Reliability* 42 (4-5): 641–651.

Thomas, S. J. (2005) *Improving Maintenance and Reliability through Cultural Change*. 1st ed., Industrial Press, New York, NY.

Trethewey, K. R. and Chamberlain, J. (1995) *Corrosion for Science and Engineering*, Longman.

Tumer, I. and Stone, R. B. (2003) 'Mapping Function to Failure Mode during Component Development', in: *Research in Engineering Design* 14 (1): 25–33.

Vajna, S. and Burchardt, C. (1998) 'Dynamic Development Structures of Integrated Product Development', in: *Journal of Engineering Design* 9 (1): 3–15.

Van der Vegte, W. F., Kitamura, Y., Mizoguchi, R., and Horváth, I. (2002) 'Ontology-Based Modeling of Product Functionality and Use – Part 2: Considering Use and Unintended Behavior', in: *Proceedings of The Third International Seminar and Workshop Engineering Design in Integrated Product Development*, Presented at the EDIProD 2002, Zielona Góra, Poland: 115–124.

Vermaas, P. E. (2009) 'The Flexible Meaning of Function in Engineering', in: *Proceedings of the 17th International Conference on Engineering Design (ICED'09)*, Presented at the ICED 2009, Design Society, Stanford, CA: 113–124.

Vincenti, W. G. (1990) *What Engineers Know and How They Know It: Analytical Studies from Aeronautical History*, The Johns Hopkins University Press, Baltimore.

——— (1994) 'The Retractable Airplane Landing Gear and the Northrop "Anomaly": Variation-Selection and the Shaping of Technology', in: *Technology and Culture* 35 (1): 1–33.

Viswanadham, P. and Singh, P. (1998) *Failure Modes and Mechanisms in Electronic Packages*, Chapman & Hall, New York, NY.

Van Vuuren, W. (1999) 'Organisational Failure: Lessons from Industry Applied in the Medical Domain', in: *Safety Science* 33 (1-2): 13–29.

Wahl, I. (2006) *Building Anatomy: An Illustrated Guide to How Structures Work*, McGraw-Hill Professional, New York, NY.

Walker, M. B. (2009) 'Causation: What Is It and Does It Really Matter?', in: *ISASI Forum* 42 (2): 4–8.

Walker, M. B. and Bills, K. M. (2008) Analysis, Causality and Proof in Safety Investigations. No. AR-2007-053, Australian Transport Safety Bureau, Canberra, Australia.

Wanhill, R. J. H. (2003) 'Milestone Case Histories in Aircraft Structural Integrity', in: Milne, I., Ritchie, R. O., and Karihaloo, B. L. (eds.), *Comprehensive structural integrity*, Elsevier: 61–72.

Wanhill, R. J. H. and Oldersma, A. (1997) Fatigue and Fracture in an Aircraft Engine Pylon. No. NLR TP 96719, National Aerospace Laboratory NLR, Amsterdam, The Netherlands.

Wanyama, W., Ertas, A., Zhang, H.-C., and Ekwaro-Osire, S. (2003) 'Life-Cycle Engineering: Issues, Tools and Research', in: *International Journal of Computer Integrated Manufacturing* 16 (4-5): 307–316.

Ward, A., Liker, J. K., Cristiano, J. J., and Sobek, D. K. (1995) 'The Second Toyota Paradox: How Delaying Decisions Can Make Better Cars Faster', in: *Sloan Management Review* 36: 43–61.

Wardhana, K. and Hadipriono, F. C. (2003) 'Study of Recent Building Failures in the United States', in: *Journal of Performance of Constructed Facilities* 17 (3): 151–158.

Wetzel, L. (2011) 'Types and Tokens', in: Zalta, E. N. (ed.), *The Stanford Encyclopedia of Philosophy*, The Metaphysics Research Lab, Stanford, CA.

Whyte, R. R. (1978) *Engineering Progress through Development*, Mechanical Engineering Publications, London.

Womack, J. P., Jones, D. T., and Roos, D. (2007) *The Machine That Changed the World: The Story of Lean Production*, Free Press, New York, NY.

Wood, R. H. and Sweginnis, R. W. (2006) *Aircraft Accident Investigation*. 2nd ed., Endeavor Books, Casper, WY.

Wulpi, D. J. (1999) *Understanding How Components Fail*. 2nd ed., ASM International, Materials Park, OH.

Xing, L. and Amari, S. V. (2008) 'Fault Tree Analysis', in: Misra, K. B. (ed.), *Handbook of Performability Engineering*, Springer, London: 595–620.

Yang, G. (2007) *Life Cycle Reliability Engineering*, John Wiley & Sons, Hoboken, NJ.

Yellman, T. W. (1999) 'Failures and Related Topics', in: *IEEE Transactions on Reliability* 48 (1): 6–8.

——— (2006) 'Redundancy in Designs', in: *Risk Analysis* 26 (1): 277–286.

Summary

This dissertation is the result of a research project taking place at the cross-roads between engineering and philosophy of technology. Its main subject is the engineering concept of failure, a concept that is rooted in the very foundations of engineering activity (i.e., realizing products that do work) and that, of course, can get a fair amount of attention from the users of the engineering products as well. Its main aim is to investigate how engineers make sense of and define this important concept and how they utilize it in practice. This is done by surveying a wide variety of sources in the engineering literature, from textbooks to historical case studies, from accident reports to research papers. A close inspection of this literature reveals a series of assumptions and conceptual distinctions which have not been fully spelled out and appreciated so far. Hence, the analysis proceeds to delineate the preliminaries of a conceptual framework capable of rationally organizing the multiplicity of approaches retrieved from the literature. Furthermore, in line with the growing interests in sustainability and diffusion of integrated approaches to product development, this framework somewhat expands the reach of the notion of failure and aims to take into account life cycle aspects of engineering products.

Chapter 2 documents how the concept of failure has been given different and even contrasting definitions within the engineering community. Research has already shown that the lack of a common terminology can have detrimental effects on knowledge sharing among individuals as well as between organizations. Chapter 2 focuses on knowledge-sharing among members of cross-functional design teams which are increasingly being adopted in industry and in which communication problems are can arise by cultural and educational divides. Its aim is to analyze a set of thirty definitions retrieved from the engineering literature and to identify those that are better suited for facilitating communication in cross-functional teams. The chapter discusses the advantages afforded by cross-functional teams (which have led to the popularity of this organizational solution) as well as the potential drawbacks, particularly those related to conceptual and terminological barriers. A result of this discussion is the identification of four criteria which can be utilized to single out those definitions of failure that are more likely to facilitate knowledge-sharing. The criteria

are: *accuracy*, *completeness*, *flexibility*, and *clarity*. It is shown that none of the existing definitions meets all four criteria, even though six of them come close by meeting three criteria. The analysis, however, provides valuable insights into potential improvements. In fact, the chapter's concluding section offers a tentative definition of failure potentially capable of meeting all four criteria. This definition, which will return in subsequent chapters, reads as follows: "Failure: The inability of an engineering process, product, service or system to meet the design team's goals for which it has been developed".

In Chapter 3 it is argued that many well-established definitions of failure share four basic assumptions: *missing functionality*, *utilization context*, *item level*, and *negativity assumptions*. Jointly, these assumptions define a traditional view on failure that could be described as *event-oriented* and can be contrasted with a *goal-oriented* view that, though less widely represented, is also present in the engineering literature. The chapter shows goal-oriented concepts of failure violate the first three assumptions, and because of that they are capable of capturing a range of situations and events that many engineers are inclined to classify as failures, even though those events do not fit easily with the traditional event-oriented view. Interestingly, goal-oriented concepts are well suited for Life Cycle Engineering (LCE). Event-oriented concepts of failure reflect a view of engineering product development that focuses on end-user's needs and on products' functional characteristics that are expected to satisfy them. On the other hand, LCE takes into account needs and requirements of multiple stakeholders whose interests may lie in anyone of the life cycle stages, from supply of raw materials to manufacturing and recycling. Thus, product properties which are not directly related to functional performance and yet have an impact on stakeholders' interests now become relevant with respect to failure judgments.

The most relevant difference between event-oriented and goal-oriented views is that, according to the latter, failure is no longer conceptualized as a discrete occurrence in the utilization phase of an item. Instead, failure judgments depend on the ability of products to achieve predefined goals that may involve anyone of the stages in the life cycle. It turns out then, that the definition proposed at the end of Chapter 2, which instantiates a goal-oriented view, is also well suited as a definition of failure for LCE.

In recent years, formal ontologies are being developed in various scientific and technical domains for the purpose of facilitating knowledge representation and sharing. Chapter 4 deals with the problem of representing engineering

knowledge about failures in a way that can be easily shared, archived, and retrieved. In particular, it carries out preliminary work that would allow formal ontologists to formalize the concept of failure by identifying the main ontological commitments underling event-oriented concepts of failure. The chapter distinguishes between three sub-concepts, *Function-based*, *Specification-based*, and *Material-based* failure. By means of an exemplary case story, it is shown that the three sub-concepts are mutually independent: an event that classifies as a failure given say, a function-based concept, could be classified otherwise by the other two. Nevertheless, at the most abstract level these three sub-concepts are based on the same ontological outline and share the same fundamental ontological commitments. Given the basic ontological categories of *occurrent*, *continuant*, and the *participation* relation, all three sub-concepts conceptualize failure events as *atomic occurrents* in which physical items *participate*. Physical items belong to the ontological category of *continuants*. States or conditions, on the other hand, belong to the *occurrent* category. Two states in particular are singled out in the representation of failures: *functioning states*, that is to say those states in which items are performing as expected, and *fault states* which obtain when performance deviates from expectations.

Chapter 5 examines the concept of *root cause* of failure events and seeks to understand whether it is possible to reconcile the different views expressed in the engineering literature, particularly between the need to understand why a failure happened and how reoccurrence could be prevented. In this chapter, failure investigations are analyzed as constituted of two sub-investigations. One is a backward looking investigation whose aim is to unearth the causal structure of those events which eventually culminated into the failure event. The underlying concept of cause is deterministic and token-based, meaning that the causal factors are deterministically linked clearly identifiable entities or events. The second sub-investigation is characterized by a probabilistic and type-based concept of cause. The causal factors identified by the backward looking investigation provide the grounds for developing potential failure scenarios that may happen in the future thereby initiating the forward looking sub-investigation. Its aim is to understand which factors are likely to reoccur and where corrective measures are more likely to be effective. By pursuing it, investigators attempt to establish probabilistic causal connections between types or categories of events which are based on already known causal factors. Differently from claims about the causal connections that hold the sequence of events together, which may

have strong empirical support, claims about *future* causal connections and scenarios envisaged by the forward looking investigation are less certain and can only be expressed by means of probabilities. Still, for the investigation to achieve tangible improvements, it should motivate why a certain countermeasure (e.g., redesign of a component vs. revision of maintenance procedures) is going to be *most* beneficial in preventing reoccurrence. The factor targeted by that countermeasure *is* the root cause, which Chapter 5 proposes to conceptualize as a U-turn between the backward looking and the forward looking sub-investigations. The root cause of a failure, then, is that element of the factors and causes which, if corrected in future scenarios, is the most likely to prevent similar events from happening again.

The dissertation concludes with a chapter on learning from failures. More precisely, Chapter 6 deals with the belief, which is shared by many engineers, that *in engineering more is learned from failures than from successes*. By looking closely at the case stories and at the arguments advanced in its support (especially those found in Henry Petroski's works), it is shown that this belief can be understood as a twofold hypothesis, a *specific-learning hypothesis* and a *generic-learning hypothesis*, which depends on two different interpretations of learning. According to the specific-learning hypothesis, the epistemic agent (i.e., the subject who learns) is either an individual engineer or a well-identifiable group of engineers (e.g., a design team or an engineering organization). The adjective *specific* indicates that the design goal facing the epistemic agent comes with clearly specified metrics for success and failure. The generic hypothesis is more ambitious and far reaching in that it interprets learning as the cognitive changes occurring at the level of a whole engineering community. To put it differently, its scope is no less than technical change on a large scale.

Chapter 6 shows that, given reasonable interpretations of the key concepts of *failure*, *success*, and *learning*, it is possible to derive an interpretation of the specific-learning hypothesis which is empirically testable. In fact, recent studies by organizational theorists have found empirical evidence in its support. On the contrary, when the same concepts are employed to analyze the generic-learning hypothesis, the result is a drastically simplified representation of technical change which conflicts with recent accounts offered by historians and sociologists of technology.

Samenvatting

Dit proefschrift is de uitkomst van een onderzoeksproject op het kruispunt van ingenieurspraktijk en techniekfilosofie. Het centrale onderwerp is het begrip 'falen' zoals dat in de ingenieurswetenschappen en ingenieurspraktijk gangbaar is, een begrip dat medebepalend is voor de activiteiten van ingenieurs (namelijk, het realiseren van producten die wél werken) en dat vanzelfsprekend ook in de belangstelling van de gebruikers van technische producten staat. Het belangrijkste doel is om te onderzoeken hoe ingenieurs dit kernbegrip opvatten en definiëren en hoe ze het in de praktijk gebruiken. Dit heb ik gedaan door het onderzoeken van een grote verscheidenheid aan bronnen in de ingenieursliteratuur, van studieboeken tot gevalsbeschrijvingen en van ongevalsrapporten tot academische artikelen. Een nadere bestudering van deze literatuur heeft een reeks aannames en conceptuele onderscheiden opgeleverd die tot dusver niet volledig duidelijk gemaakt en op waarde geschat zijn. Daarom gaat de analyse over in een aanzet tot een begrippenkader dat in staat is om de veelheid van benaderingen die ik in de literatuur heb aangetroffen rationeel te organiseren. Bovendien breidt dit kader, in lijn met het groeiende belang van duurzaamheid en de verbreiding van geïntegreerde benaderingen van productontwikkeling, de reikwijdte van het faalbegrip uit en laat het toe rekening te houden met aspecten die de levenscyclus van technische producten betreffen.

Hoofdstuk 2 beschrijft hoe het begrip 'falen' verschillende en zelfs tegengestelde definities heeft gekregen binnen de ingenieursgemeenschap. Onderzoek heeft ruimschoots aangetoond dat het ontbreken van een gemeenschappelijke terminologie nadelige effecten kan hebben op het delen van kennis tussen individuen en tussen organisaties. Hoofdstuk 2 richt zich op het uitwisselen en delen van kennis tussen de leden van cross-functionele ontwerpteams, welke in toenemende mate worden ingezet in de industrie en waarin communicatieproblemen kunnen ontstaan door culturele en educatieve verschillen. Het doel is om een dertigtal definities, verkregen uit de technische literatuur, te analyseren en vervolgens te bepalen welke definities beter geschikt zijn om communicatie in cross-functionele teams soepel te laten verlopen. In het hoofdstuk bespreek ik welke voordelen cross-functionele teams bieden – voordelen die hebben geleid tot de populariteit van deze organisatorische oplossing – alsmede de potentiële

nadelen, met name die met betrekking tot conceptuele en terminologische barrières. Deze analyse resulteert in de formulering van vier criteria aan de hand waarvan definities van falen die de kans op kennisdeling vergroten geselecteerd kunnen worden. Deze criteria zijn: *nauwkeurigheid*, *volledigheid*, *plooibaarheid* en *duidelijkheid*. Ik toon aan dat geen van de bestaande definities aan alle vier de criteria voldoet, en dat zes van de definities enigszins in de buurt komen door aan drie criteria te voldoen. Niettemin biedt de analyse waardevolle inzichten om tot betere definities te komen. In de afsluitende paragraaf van het hoofdstuk stel ik een voorlopige definitie van falen voor die, in potentie, aan alle vier de criteria voldoet. Deze definitie, die zal terugkeren in de volgende hoofdstukken, luidt als volgt: "Falen: Het onvermogen van een technisch proces, product, dienst of systeem om de doelstellingen van het ontwerpteam ter realisering waarvan het is ontwikkeld daadwerkelijk te realiseren".

In Hoofdstuk 3 betoog ik dat veel gevestigde definities van falen de volgende vier uitgangspunten delen: *ontbrekende functionaliteit, gebruikscontext, itemniveau* en *negativiteitsaannames*. Gezamenlijk definiëren deze aannames een traditionele visie op falen die kan worden omschreven als *gebeurtenisgericht* en kan worden afgezet tegen een *doelgerichte* opvatting die, hoewel minder ruim vertegenwoordigd, ook aanwezig is in de ingenieursliteratuur. Ik laat zien dat doelgerichte opvattingen van falen de eerste drie aannames schenden, en daardoor in staat zijn om allerlei situaties en gebeurtenissen die veel ingenieurs als vormen van falen op zullen vatten inderdaad als zodanig te classificeren, terwijl deze gevallen niet goed te rijmen zijn met de traditionele gebeurtenisgerichte opvatting van falen. Doelgerichte opvattingen zijn daarom uitstekend geschikt voor Life Cycle Engineering (LCE). Gebeurtenisgerichte opvattingen van falen passen bij een perspectief op de ontwikkeling van technische producten als gericht op de behoeften van eindgebruikers en op de functionele eigenschappen van producten waardoor ze, naar verwachting, in die behoeften kunnen voorzien. LCE daarentegen richt zich op de behoeften en eisen van verschillende belanghebbenden, wier belang bij elke fase van de levenscyclus kan liggen, vanaf de levering van de grondstoffen en de productie tot aan recycling. Zo kunnen producteigenschappen die niet direct gerelateerd zijn aan de functionele prestaties en die toch raken aan de belangen van belanghebbenden relevant worden voor een oordeel dat er van falen sprake is.

Het meest relevante verschil tussen de gebeurtenisgerichte en de doelgerichte opvatting is dat in het laatste geval falen niet langer opgevat wordt als een

afzonderlijke gebeurtenis in de *gebruiks*fase van een item. In plaats daarvan hangt het oordeel af van het vermogen van een product om aan vooraf gedefinieerde doelen die betrekking kunnen hebben op elke fase van de levenscyclus te voldoen. Het blijkt dan dat de definitie aan het einde van Hoofdstuk 2, die een doelgerichte opvatting inhoudt, ook zeer geschikt is als een definitie van falen voor LCE.

In de afgelopen jaren zijn in verschillende wetenschappelijke en technische domeinen formele ontologieën ontwikkeld die het vergemakkelijken van kennisrepresentatie en kennisdeling beogen. Hoofdstuk 4 behandelt het probleem van het representeren van ingenieurskennis met betrekking tot falen zo dat deze gemakkelijk kan worden gedeeld, gearchiveerd en teruggehaald. Meer in het bijzonder doe ik hier voorbereidend werk om formele ontologen in staat te stellen het begrip van falen te formaliseren door het identificeren van de belangrijkste ontologische uitgangspunten en keuzes die gebeurtenisgerichte opvattingen van falen gemeen hebben. In het hoofdstuk wordt onderscheid gemaakt tussen drie deelbegrippen, *op functie gebaseerd*, *op specificatie gebaseerd* en *op materiaal gebaseerd* falen. Door middel van een voorbeeld laat ik zien dat de drie deelbegrippen onderling onafhankelijk zijn: een gebeurtenis die geclassificeerd wordt als falen aan de hand van, zeg, een op functie gebaseerd begrip van falen kan anders worden ingedeeld aan de hand van de andere twee. Niettemin gaan deze drie begrippen op het meest abstracte niveau op dezelfde ontologische grondschets terug en maken ze dezelfde fundamentele ontologische keuzes. Onder gebruikmaking van de fundamentele ontologische categorieën *occurrent* en *continuant* en van de *participatierelatie* betoog ik dat alle drie de deelbegrippen gevallen van falen conceptualiseren als *atomaire gebeurtenissen* waarin fysieke items *participeren*. Fysieke items behoren tot de ontologische categorie van de *continuanten*. Toestanden of voorwaarden behoren echter tot de categorie van de *occurrenten*. In de representatie van falen onderscheid ik twee toestanden in het bijzonder: *functioneringstoestanden* ofwel toestanden waarin de items presteren zoals van hen verwacht wordt, en *faaltoestanden* ofwel toestanden die optreden wanneer de prestatie afwijkt van de verwachtingen.

In Hoofdstuk 5 onderzoek ik het begrip '*hoofdoorzaak* van falen' en ga ik na of het mogelijk is twee verschillende gezichtspunten in de ingenieursliteratuur met elkaar te verzoenen, met name de opvatting die zich richt op de noodzaak om te begrijpen waarom een geval van falen optrad en die welke zicht richt op de vraag hoe herhaling kan worden voorkomen. In dit hoofdstuk analyseer ik

onderzoeken naar falen als bestaande uit twee deelonderzoeken. Eén daarvan is terugblikkend en heeft als doel om de causale structuur bloot te leggen van de gebeurtenissen die uiteindelijk culmineerden in het falen. Het onderliggende oorzaakbegrip is deterministisch en geïndividualiseerd, wat betekent dat oorzakelijke factoren duidelijk identificeerbare entiteiten of gebeurtenissen zijn die deterministisch met elkaar verbonden zijn. Het tweede deelonderzoek wordt gekenmerkt door een probabilistisch oorzaakbegrip dat gebaseerd is op types. De causale factoren die het terugblikkende onderzoek identificeert bieden de basis voor het ontwikkelen van mogelijke rampscenario's die in de toekomst kunnen gebeuren, aldus het vooruitblikkende deelonderzoek initiërend. Het doel daarvan is te begrijpen welke factoren waarschijnlijk opnieuw zullen optreden en waar corrigerende maatregelen het meeste kans hebben effectief te zijn. Op deze manier proberen onderzoekers om probabilistische causale verbanden tussen soorten gebeurtenissen of categorieën van gebeurtenissen vast te stellen, die zijn gebaseerd op reeds bekende oorzakelijke factoren. Voor beweringen over de causale verbanden die de reeks van opeenvolgende gebeurtenissen samenbinden kan sterke empirische steun bestaan. Beweringen over *toekomstige* causale verbanden en scenario's daarentegen, die het resultaat van het vooruitblikkende onderzoek vormen, zijn minder zeker en kunnen alleen in termen van waarschijnlijkheden worden uitgedrukt. Desalniettemin dient het onderzoek, wil het concrete verbeteringen opleveren, te motiveren waarom een bepaalde tegenmaatregel (bijvoorbeeld het herontwerp van een component en niet de herziening van onderhoudsprocedures) het *meest* effectief is om herhaling te voorkomen. De factor die het doelwit is van de tegenmaatregel *is* de hoofdoorzaak, die ik in Hoofdstuk 5 conceptualiseer als een U-bocht tussen het terugblikkende en het vooruitblikkende deelonderzoek. De hoofdoorzaak van een geval van falen is dan dat element uit het netwerk van factoren en oorzaken dat, indien gecorrigeerd in toekomstige scenario's, het meeste kans maakt te voorkomen dat soortgelijke gevallen van falen opnieuw zullen optreden.

Het proefschrift sluit af met een hoofdstuk over het leren van falen. Preciezer gezegd buig ik me in Hoofdstuk 6 over de door veel ingenieurs gedeelde opvatting dat *in de ingenieurspraktijk meer wordt geleerd van gevallen van falen dan van gevallen van succes*. Door enkele gevalsstudies nauwgezet te analyseren en de argumenten die voor deze opvatting worden aangedragen (met name in het werk van Henry Petroski) nauwkeurig te bekijken laat ik zien dat deze overtuiging als een tweeledige hypothese kan worden opgevat, een *hypothese over specifiek leren*

en een *hypothese over generiek leren*, afhankelijk van welke van twee verschillende interpretaties van leren wordt gekozen. Volgens de hypothese over specifiek leren is degene die iets leert hetzij een individuele ingenieur hetzij een duidelijk identificeerbare groep van ingenieurs (bijvoorbeeld een ontwerpteam of een organisatie binnen de ingenieurswereld). Het bijvoeglijk naamwoord *specifiek* geeft aan dat er voor het ontwerpdoel waar de ingenieur(s) mee geconfronteerd wordt (worden) duidelijk gespecificeerde metrieken voor succes en falen bestaan. De hypothese over generiek leren is ambitieuzer en verder reikend in de zin dat leren daarin wordt vereenzelvigd met cognitieve veranderingen op het niveau van de hele ingenieursgemeenschap. Anders gezegd, het leren betreft hier technische verandering op grote schaal.

In Hoofdstuk 6 laat ik zien dat redelijke interpretaties van de kernbegrippen *falen*, *succes* en *leren* het mogelijk maken om tot een interpretatie van de hypothese over specifiek leren te komen die empirisch toetsbaar is. Sterker nog, recente studies in de organisatiekunde hebben empirisch bewijs gevonden dat deze hypothese ondersteunt. Wanneer daarentegen dezelfde begrippen worden gebruikt om de hypothese over generiek leren te analyseren is het resultaat een sterk vereenvoudigde weergave van technische verandering die in strijd is met het beeld dat historici en sociologen van de techniek hier tegenwoordig van schetsen.[25]

[25] Ik ben Maarten Franssen en Christine van Burken erg dankbaar voor de Nederlandse vertaling van de Samenvatting.

About the author

Luca Del Frate was born in Italy in 1972. He received his technical education in high-school. He attended a technical college from which he obtained a diploma in Aeronautical Technologies. He matured work experience as production engineer in the metal-working and in the glass-working sectors where he was in charge of production planning and supervision. He has also a Philosophy degree from the University of Padua (Italy) where he graduated in 2004 with a dissertation on evolutionary explanations of social behavior. In 2008 he moved to the Netherlands to pursue his PhD research at Delft University of Technology.

Simon Stevin Series in the Philosophy of Technology
Delft University of Technology & Eindhoven University of Technology
Editors: Peter Kroes and Anthonie Meijers

Books and Dissertations

Volume 1: Marcel Scheele, *'The Proper Use of Artefacts: A philosophical theory of the social constitution of artefact functions'*, 2005

Volume 2: Anke van Gorp, '*Ethical issues in engineering design, Safety and sustainability*', 2005

Volume 3: Vincent Wiegel, '*SophoLab, Experimental Computational Philosophy*', 2006

Volume 4: Jeroen de Ridder, '*Technical Artifacts: Design and Explanation*', 2006

Volume 5: Melissa van Amerongen, *'The Interpretation of artifacts; A critique of Dennett's design stance*', 2008

Volume 6: Krist Vaesen, '*A Philosophical Essay on Artifacts and Norms*', 2008

Volume 7: Giacomo Romano, '*Thoughtful Things. An investigation in the descriptive epistemology of artifacts'*, 2009

Volume 8: Dingmar van Eck, '*Functional Decomposition: On Rationality and Incommensurability in Engineering*', 2011

Volume 9: Auke Pols, '*Acting with Artefacts*', 2011

Volume 10: Marieke van Holland, '*Extended Cognition and the Mark of the Cognitive: Prospects of a proper function-based approach*', 2013

Volume 11: Luca Del Frate, '*Failure: Analysis of an Engineering Concept*', 2014

Research Documents

Peter Kroes and Anthonie Meijers (eds.), *'Philosophy of Technical Artifacts*', 2005

Simon Stevin (1548-1620)

'Wonder en is gheen Wonder'

This series in the philosophy and ethics of technology is named after the Dutch / Flemish natural philosopher, scientist and engineer Simon Stevin. He was an extraordinary versatile person. He published, among other things, on arithmetic, accounting, geometry, mechanics, hydrostatics, astronomy, theory of measurement, civil engineering, the theory of music, and civil citizenship. He wrote the very first treatise on logic in Dutch, which he considered to be a superior language for scientific purposes. The relation between theory and practice is a main topic in his work. In addition to his theoretical publications, he held a large number of patents, and was actively involved as an engineer in the building of windmills, harbours, and fortifications for the Dutch prince Maurits. He is famous for having constructed large sailing carriages.

Little is known about his personal life. He was probably born in 1548 in Bruges (Flanders) and went to Leiden in 1581, where he took up his studies at the university two years later. His work was published between 1581 and 1617. He was an early defender of the Copernican worldview, which did not make him popular in religious circles. He died in 1620, but the exact date and the place of his burial are unknown. Philosophically he was a pragmatic rationalist for whom every phenomenon, however mysterious, ultimately had a scientific explanation. Hence his dictum 'Wonder is no Wonder', which he used on the cover of several of his own books.