Documenting The Design Of Safety-Critical, User Interfaces *

Chris Johnson,
Department of Computing Science,
University of Glasgow,
Glasgow, United Kingdom, G12 8QQ.
Telephone: +44 41 339 8855 Fax: +44 41 330 4913.
E-mail: johnson@dcs.glasgow.ac.uk

August 4, 1995

Abstract

A number of techniques have been developed to map out the 'design space' of interactive systems. For instance, design rationale notations represent the various options that must be considered during interface development: the style of interaction; the choice of input and output media etc. Such notations offer important benefits. They can be shown to customers and other designers. This helps to identify important development decisions. They can be shown to regulatory authorities in order to demonstrate that sufficient attention has been paid to particular design issues. Unfortunately, these notations are semi-formal. They are, typically, descriptive and do not support proof techniques. This leads to a number of important weaknesses. They cannot be used to determine whether criteria and options have been used consistently during the development of an interactive system. Nor do they provide any procedures to determine whether interface design decisions have, indeed, been guided by a rationale. These are important limitations when regulatory authorities, such as the UK Health and Safety Executive, expect designers to demonstrate the consistency and coherence of their decisions. The following pages address these problems and show how decision theory can be recruited to support design rationale during the development of human computer interfaces for safety-critical systems.

Keywords: *Safety-Critical Interfaces, Design Rationale, Decision Theory, Design*

*Submitted to Interacting With Computers

1 Introduction

The Commission of the European Community [4], the Japanese Fifth Generation Initiative [25] and United States' Presidential Task Forces [20] have all cited human intervention as a primary factor in the cause and exacerbation of accidents. In consequence, statutory authorities such as the UK Health and Safety Executive and the US Nuclear Regulatory Commission require documentary evidence to justify the decisions that guide the development of interactive systems. Typically, this is done through the generation of large, natural language 'safety-cases'. These documents are often poorly structured, the evidence that they contain may be contradictory and they can suffer from significant omissions [14]. This paper argues that design rationale notations can be enhanced to avoid these limitations.

1.1 Design Rationale

The semi-formal languages of design rationale provide a common format for the justifications that support particular development decisions during the design of human computer interfaces. They provide a structure that can be used to clarify the arguments that are frequently lost amongst the prose descriptions of natural language documents [17]. For example, Figure 1 illustrates the

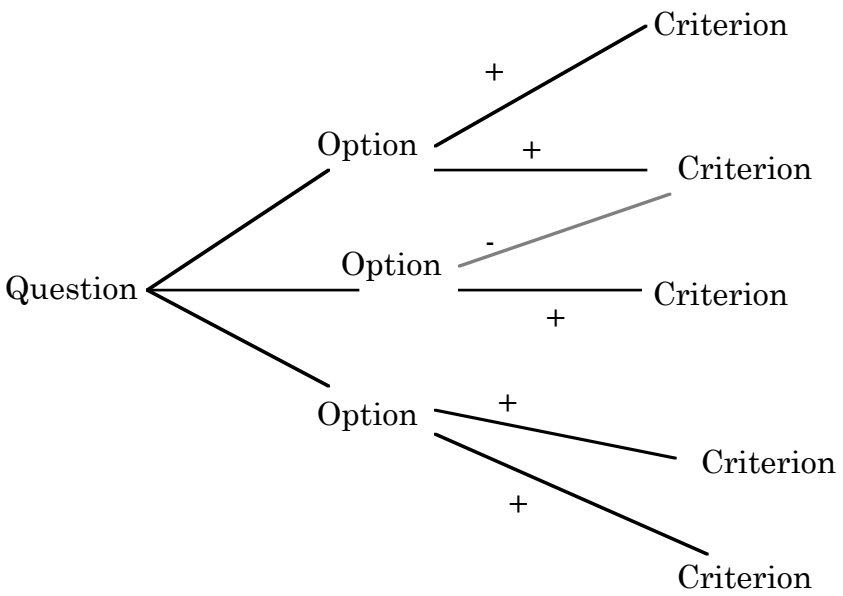

Figure 1: A Simple QOC diagram.

syntax of the Questions, Options and Criteria (QOC) notation, originally developed by Rank Xerox [23]. These diagrams are built by identifying the key questions that must be addressed during interface development. The construction of a rationale then proceeds by linking the

options to the questions that they answer. The criteria that strengthen or weaken options are then linked using solid or broken lines. QOC diagrams are intended to support decision making by explicitly representing the reasons for and against design options. They can be shown to new members of design groups or to colleagues in other development teams. This helps to ensure that the same criteria are applied throughout the design of an interactive system.

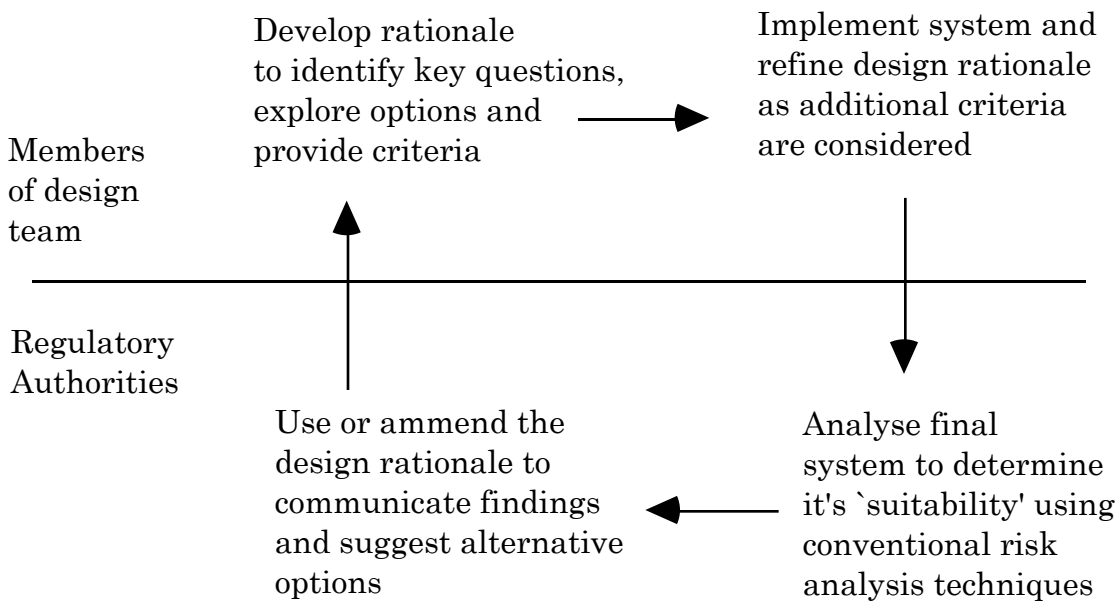

Figure 2: The Role Of Design Rationale In Safety-Critical Development.

Figure 2 shows how semi-formal notations, such as QOC and the more elaborate gIBIS [2], might be recruited within the development of safety-critical systems. The explicit identification and representation of key design questions not only helps to guide the development of a system, it can also help regulatory authorities to identify critical stages in the design of an interactive application. This approach has already been used with some success by a number companies, including Boeing [15]. The use of a design rationale does not, however, guarantee the implementation of a safe and successful user interface. Semi-formal notations, such as QOC, provide descriptive records of the factors that guide the development process. It is entirely possible for designers to select a poorly supported option even if other, better supported options are available. Similarly, it is entirely possible for criteria to be used inconsistently, by supporting an option in one context and weakening it in another. This flexibility offers important benefits

for the development of many human-computer interfaces. It does, however, create problems for the development of safety-critical interfaces. In such domains, designers must be able to demonstrate that their decisions are both consistent and coherent [7].

1.2 Decision Theory

In contrast to design rationale notations, such as QOC, decision theory is based around a series of formal models. These models have been developed in the fields of economics and management science to reason about the ways in which people select various options from a range of alternatives [8]. The intention is not to map out all of the complex factors that affect human decision making but to focus on the interaction between issues of cost and probability that affect commercial project management. Decisions are characterised in terms of the rewards that might accrue from selecting a particular option. In the general case, these rewards might be amounts of money, increases in safety, reduction in development time etc. Each reward has an associated probability. This represents the likelihood that a reward will actually be obtained from selecting a particular option. The following paragraph introduces a formal model that has been derived from this work [21]. Readers who are more concerned with the pragmatic application of decision theory than with its formal basis should move to the paragraph that starts immediately after the box.

A decision can be represented by a vector where o represents a design option, each x^i is a potential reward, in our case this would be the criteria supporting a design decision, from a set X and each p^i is the associated probability of obtaining that reward; $(o; x^1, p^1; ...; x^n, p^n)$. Decisions are represented by preference relations, $\succ$, between the various sets of criteria that are associated with a particular option:

$$\forall x \in X : F_u(x) = P([0, x]) \tag{1}$$

Risk is modeled through the set $D(X)$ where the probability, P, of obtaining a reward in the range $[0, x]$ is determined by a cumulative distribution function $F_u \in D(X)$, the number $F_u(x)$ is interpreted as the probability of receiving an amount less than or equal to x.

$$\forall F_u \in D(X) : V(F_u) = \sum_{i=1}^{n} u(x^i)p^i \tag{2}$$

$$\forall F_u, G_u \in D(X) : F_u \succ G_u \leftrightarrow V(F_u) > V(G_u) \tag{3}$$

In the first formulae, the utility of a potential reward is derived by the application of a function $u : X \rightarrow R$. The value of a lottery, F_u, which is of the form $(x^1, p^1; ...; x^n, p^n)$ is modeled by $V : D(X) \rightarrow R$. In the second formulae, decisions are represented as risk preferences between lotteries. This is represented by the $\succ$ relation that describes an ordering over the elements of $D(X)$.

The previous clauses provide a number of potential benefits for the development of safety-critical systems. They are amenable to formal analysis in a manner that is not possible using notations such as QOC. This is important because a range of government bodies, including the UK Ministry of Defense [18] and NASA [18], are increasingly advocating the use of mathematical specification techniques. A prime motivation behind the use of formal notations is that specifications can be refined to a point where it is a trivial step to derive an implementation. If human-computer interaction techniques cannot be represented within this refinement process then they run the risk of being relegated to an afterthought in the development of safety-critical systems.

A further benefit from recruiting decision theory to support the development of interactive systems is that notions of cost or utility, u, and probability, p_i, are explicitly represented [8, 5]. This is vital for the development of safety-critical systems where human-reliability analysis forms a necessary stage in interface development. Unfortunately, formulae such as (3) cannot easily be used to support multi-disciplinary design; many clients, designers and users will not have the mathematical background necessary to apply elements of decision theory. In contrast, the following paper shows how the formal models can support the application of design rationale without introducing additional constructs into the graphical notation. It is intended that the formal aspects of the approach should only be used at a meta-level to address those problems that prevent QOC from being effectively applied to safety-critical interfaces.

1.3 Structure Of The Paper

Section 2 presents the example that is used to illustrate this paper. A communications system for an underground train network has been chosen. This is appropriate because the interface development for this application requires close cooperation between human factors and systems engineers. Failures in the design and maintenance of such a system exacerbated the London Underground accident at Kings' Cross [3]. Section 3 uses this application to argue that some concept of utility or value must be attached to criteria in design rationale documents. This is necessary if regulatory authorities and other development teams are to understand the priorities that guide interface design. Decision theory is then used to show how designers can explicitly represent rankings between design options based on the value of their criteria. Section 4 builds on this argument and demonstrates that risk can be considered during the application of design rationale. This provides an important link between HCI development techniques and established practice in probabilistic risk assessment including human reliability analysis. Section 5 shows how the rankings derived from decision theory can be used to identify weaknesses in QOC diagrams. In particular, they are used to identify missing criteria. Section 6 builds on this by using the elements of decision theory to gradually refine the level of detail that is represented in design rationale. This is important if interface designers are to convince regulatory authorities that they have considered different options at a sufficient level of detail. Section 7 then demonstrates that these techniques can represent consistency criteria between different design rationale documents. This is essential if development teams are to demonstrate that the same options have been considered for similar design questions. Section 8 identifies areas for further work and presents the conclusions that can be drawn from this research.

2 An Example Application

It is essential to provide effective communications systems for the controllers, station managers and rescue services that operate modern Underground transportation systems. Following the London Underground, Kings' Cross fire, considerable resources have been allocated to improve the flow of information between these user groups [3]. High level requirements for such a system will be used to illustrate the remainder of this paper. This example is appropriate because it typifies many of the problems that complicate the development of safety-critical, human-machine interfaces. Its successful development relies upon the close cooperation of many different experts. For instance, human factors engineers must determine whether textual communications,

standard audio links of full video conferencing applications provide the best means of coordinating the various activities in the work-system. Systems engineers must determine whether they can support the communications infrastructure that is necessary to meet the information needs of system operators. These development tasks are further complicated by the need to consider the interests of 'external' agencies such as fire-fighting and medical services as well as regulatory authorities and government agencies. With such potential for disagreement, it is vital that designers have some means of ensuring the completeness and the consistency of semi-formal design rationales during the development of safety-critical interfaces.

3 Lack Of Utility

One of the problems that complicates the multi-party design of safety-critical applications, such as the Underground communications system, is that different groups may possess different sets of priorities for interface development. Unfortunately, design rationale provides no means of representing the different values, or utilities, that developers associate with particular options. For example, Figure 3 represents the arguments for and against the use of different media in the

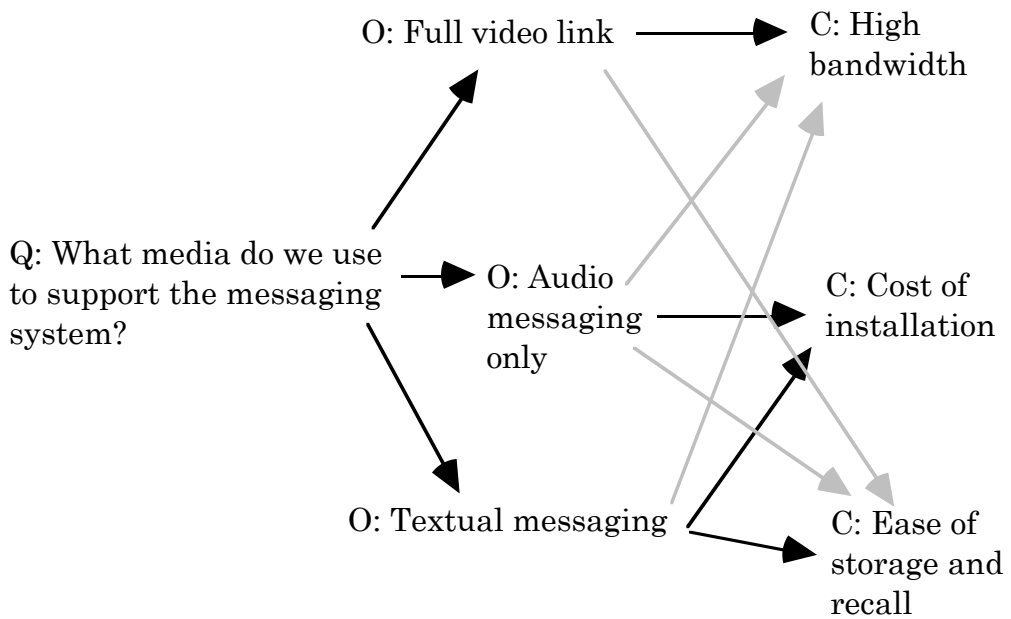

Figure 3: QOC For The Communications System

presentation of our example communications system. Video is supported by its broad bandwidth; users can exploit these links to show their colleagues documents or directly view areas of the network. This form of communication is not supported by the likely cost of the system nor is it

easy to store and retrieve video-based communication from other operators. In contrast, both audio and text based communication are easily recalled using standard techniques. They are both relatively low-cost options but neither supports broadband communication. It is important to emphasise that this diagram represents a simplification of the available design options. Video, text and audio links are not exclusive alternatives. A voice messaging system might be combined with video-based communications. Figure 3 is, however, sufficient to illustrate some of the problems that occur when design rationale is used to support the development of safety-critical interfaces. For example, regulatory authorities might consider the cost of the communications system to be much less important than the provision of broadband communications in Figure 3. Under such circumstances, the rationale would justify a decision to select the video option. However, if the cost were a significant factor for the development team then either the audio or text options might be selected. The utility or value assigned to each criteria, therefore, has a profound impact upon the interpretation of design rationale documents. It is vital that such assessments be made explicit during the development of large-scale projects. In the pathological case, one development team might push for low cost, reduced bandwidth options while another developed high-cost broadband communications.

The concept of utility in decision theory can be used to weight the criteria that are associated with particular options. Our intention is specifically **not** to alter the existing flexibility of design rationale notations but to demonstrate that it is possible to introduce the additional level of detail that is necessary during interface development for safety-critical applications. For instance, the model introduced in Section 1.2 can be used to explicitly represent the utility or value associated with the options in QOC diagrams. This provides means of supplying the contextual information about shared objectives that was missing in Figure 3. The following rankings represents a situation in which ease of recall and low cost are higher objectives than a rich bandwidth:

$$(audio_only; low_cost; easy_recall) \succ (full_video; rich_bandwidth) \quad (4)$$

$$(text_only; low_cost; easy_recall) \succ (full_video; rich_bandwidth) \quad (5)$$

The previous two clauses state that the low cost and ease of recall associated with audio and text links are assigned higher priority than the rich bandwidth of video links. On the basis of the existing QOC and the previous clauses there is no means of distinguishing any preferences between audio and text links.

Such an assessment would clearly justify a decision to adopt either the audio or the text based

information system. In a safety-critical applications, however, the additional information provided by video communications might be a justified expense. This would reverse the priorities shown in the rankings given by (4) and (5):

$$(full_video; rich_bandwidth) \succ (audio_only; low_cost; easy_recall) \quad (6)$$

$$(full_video; rich_bandwidth) \succ (text_only; low_cost; easy_recall) \quad (7)$$

The previous two clauses state that the low cost and ease of recall associated with audio and text links are assigned a lower priority than the rich bandwidth of video links.

These formulae illustrate the dangers that can arise from the use of QOC diagrams without some means of representing and reasoning about design priorities. Clearly, it makes little sense to say that the rich bandwidth of video communication is preferred to the low cost and easy recall of text links (7) and that the low cost and easy recall of text links are preferred to the rich bandwidth of video communication (5). The previous rankings help to identify such conflicting priorities in a manner that is not possible using design rationale notations, such as QOC.

4 Lack Of Risk

Preference relations, such as (6) and (7), help to explicitly represent shared priorities for multi-disciplinary design teams. They do not, however, provide a panacea for the use of design rationale in safety-critical systems. In particular, they ignore the probabilistic information that is represented in the models drawn from decision theory. We are concerned to represent such information because design is not a risk-free activity [9]. Development teams are seldom guaranteed the rewards that they anticipate from their decisions [19]. This can be illustrated by the

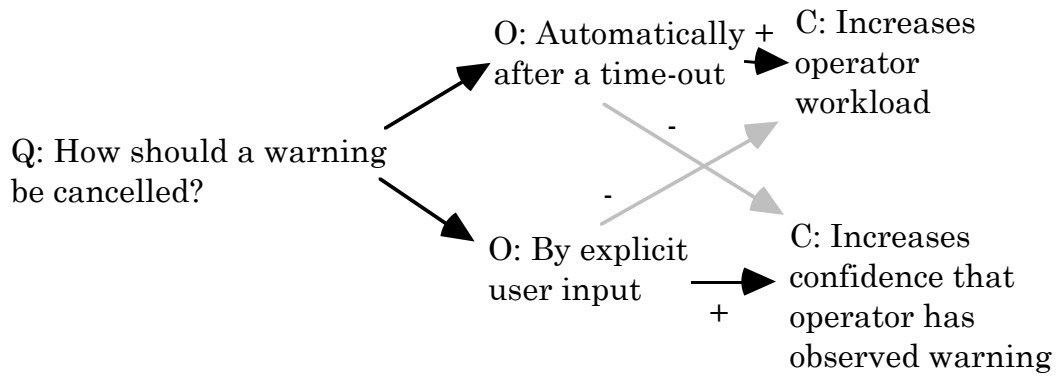

Figure 4: QOC For Warning Confirmation.

QOC diagram shown in Figure 4. The option of explicitly requiring operators to cancel warnings is intended to ensure that they have observed the display. This option is not supported by the argument that explicit confirmation reduces operator workload. Such a diagram is deceptive. It is important to emphasise that design teams have **no** guarantee that a particular display option will satisfy any of its criteria. For instance, empirical studies have shown that operators may not always understand the information that is presented to them even if they have to confirm that they have understood the display [26]. The distractions and interruptions of many control rooms can result in a 90% probability that operators will actually observe and act upon a warning after they have confirmed it. In contrast, the automatic cancellation of an error message after some time-out ensures a 60% success rate. Such data is commonly available from standard sources such as the United States' Nuclear Regulatory Commission's Handbook of Human Reliability [24]. The point here is not that human reliability figures should be absolutely accurate. They are profoundly affected by particular characteristics of the operating environment and by the demands of particular control tasks. The importance of these figures lies in their role as worse-case estimates and design objectives; 'operators must immediately respond to 90% of all warnings' [5]. Rasmussen and Pedersen [22] argue that it is vital to explicitly represent these estimates in order to focus the subsequent development of interactive systems:

$$(operator_confirmation; 0.9, successful_observation) \succ (automatic_cancellation, 0.6, successful_observation) \quad (8)$$

The 90% chance that explicit operator confirmation of error messages will elicit immediate intervention is preferred to the 60% chance that automatic cancellation will elicit the same response.

It is important to emphasise that the introduction of probabilistic information need not overcomplicate the task of interface development. If designers are certain that an option is supported by a particular criterion then they can assign it a probability of 1.0. For example, there will be an increase in workload if operators must confirm the observation of error messages. This additional load would persuade designers to reverse their previously expressed preference between the competing options:

$$(automatic_cancellation, 0.6, successful_observation) \succ$$
$$(operator_confirmation; 0.9, successful_observation; 1.0, higher_workload) \quad (9)$$

The 60% probability that automatic cancellation will elicit immediate operator intervention is preferred to the 90% chance that explicit operator confirmation of error messages will elicit the same response with a certain increase in operator workload.

The important point here is that the preference relation, $\succ$, cannot simply be constructed from subjective utilities. In most complex design enterprises, interface designers must also consider the probability that an option will satisfy a particular criteria. It is misleading to assert that particular options will always provide the criteria that are associated with them. Decision theory provides means of explicitly representing this uncertainty. It also provides an important link between HCI development techniques and the existing methods of reliability analysis.

5 Incoherence

The use of design rationale in the development of safety-critical systems is further complicated by the lack of effective decision procedures. There are few heuristics that can be used to determine whether a particular option is appropriate given a QOC or gIBIS diagram. This creates problems because other development teams cannot determine which option has been chosen from the design rationale alone. In contrast, rankings such as (9) explicitly represent this information. Figure 5 provides a further example of this benefit from the integration of decision theory and design rationale. This diagram represents the design space for alerting users to the arrival of a new message. Even at this high level of abstraction, it can be seen that purely audible or visual warnings have the weakest support. A visual notification is easily ignored. A purely audible warning is less easily ignored but can easily be mis-understood if a range of audio warnings are used. A combined approach offers the best solution. The audible warning attracts the user's attention and visual cues provide reinforcement. Unfortunately, there is no guarantee that a design team will interpret the rationale in this way. It is perfectly possible for a purely visual or audio alarm to be implemented despite the argumentation captured by Figure 5. It might seem unlike that any designer would hold the such a preference. There are, however, good reasons why it might not be possible to exploit both audible and visual warnings. For instance, screen space is limited in most safety-critical applications. The redundancy implied by both forms of warning might not be justified:

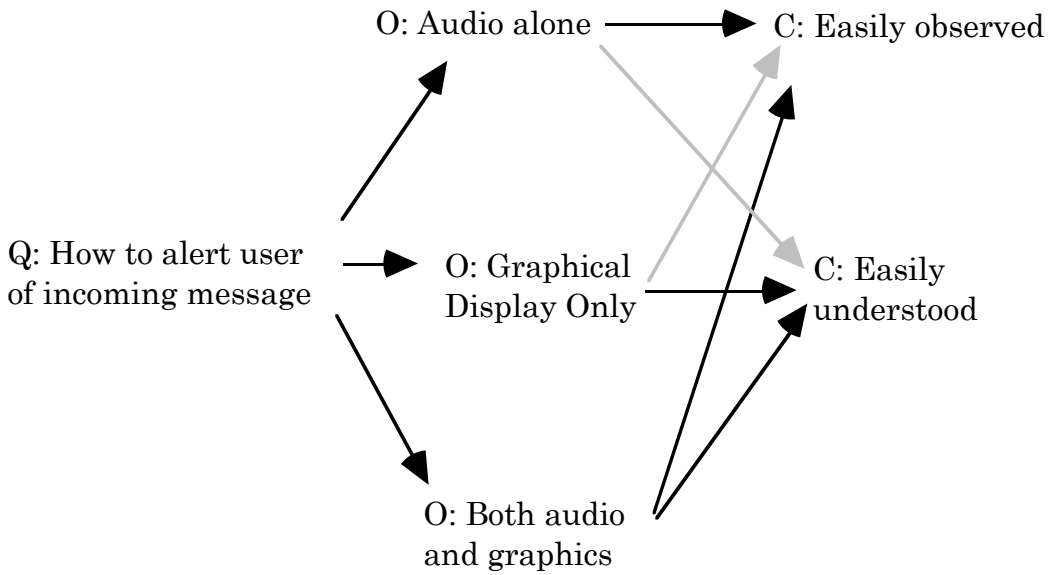

Figure 5: A QOC diagram illustrating incoherence

$$(audible_warning; 0.9, successful_observation; 0.7, operator_comprehension) \succ$$
$$(visual_warning, 0.6, successful_observation; 1.0, operator_comprehension) \succ$$
$$(audio_visual_warning; 0.9, successful_observation; 1.0, operator_comprehension) \quad (10)$$

The 60% chance that a visual warning will be observed and the 100% chance that the operator will understand the warning is preferred to the 90% chance that an audio warning will be heard and the 70% chance that the operator will understand the warning. This in turn is preferred to a combined audio and visual warning with a 90% probability that the operator will observe the warning and a 100% chance that they will understand it.

This demonstrates the manner in which decision theory and design rationale support the development of safety-critical interfaces. The ranking captured in (10) cannot easily be explained in terms of Figure 5. This suggests that the original rationale must be amended in order to explain the low ranking assigned to the audio-visual warning. Figure 6 shows how this can be done. The explicit representation of design preferences encourages designers to review the range of options and criteria that are represented in the graphical structures of design rationale.

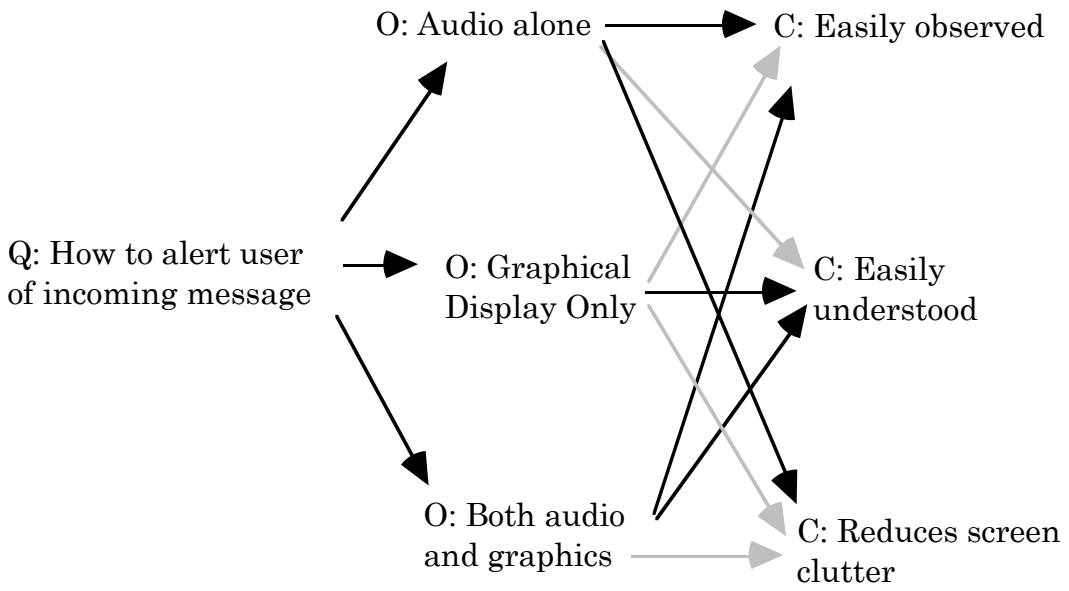

Figure 6: The amended QOC diagram

6 Inconsistency

Figure 7 illustrates the problems that inconsistency creates for the use of design rationale in the development of safety-critical interfaces. The QOC diagram presents arguments for and against the use of different media in the messaging system within station centres. The textual messaging system is **not** represented as a low cost option. This is inconsistent with the criteria shown for the text option in line centres, see Figure 3. It is important to emphasise that this inconsistency need not be 'a bad thing'. There may be good reasons why text-based systems are an inexpensive option for line centres but not for stations; line control rooms might already have terminals and low-bandwidth communications links. The problem here is that it is impossible for regulatory authorities to reconstruct this argument from the QOC diagram alone. Similarly, additional detail is required if the members of other development teams are to understand this apparently contradictory use of the same criteria to both support and weaken text based communication. Hypertext links provide means of introducing this additional detail into QOC diagrams. Figure 8 illustrates this approach. Together with Chan, we have implemented a hypertext browser that supports such annotations [1]. Unfortunately, experience suggests that these tools cannot easily be used to keep track of the many notes and links that would be required to explain inconsistencies in large-scale design rationale documents. The tractability of the networks is quickly lost as the number of annotations increases. In contrast, the logic

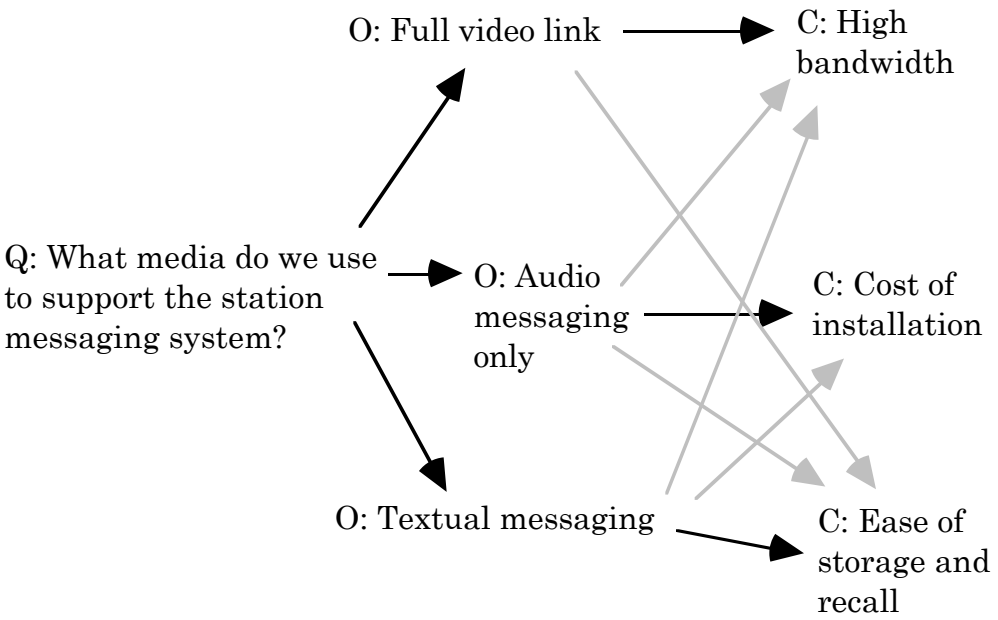

Figure 7: QOC Showing Inconsistency

foundations of decision theory provide a precise and concise means of explicitly representing the causes of inconsistency:

$$supports(text_only, low_cost, line) \Leftarrow$$
$$terminals_installed \wedge communication_lines_secure \quad (11)$$

This states that the text only option is a low cost opportunity for line centres if terminals are installed and there are secure communications lines.

A further benefit of this approach is that additional detail can be gradually introduced for terms, such as *communications_lines_secure*. This provides human factors and systems engineers with means of gradually refining their rationales so that they contain sufficient detail to convince external agencies and regulatory authorities. For example, it can be specified that secure communications lines require both redundant network connections and an auxiliary power supply. Such details can, in turn, be refined to introduce yet more detail as design progresses. At each stages, these options might be supported by the development of additional QOC diagrams:

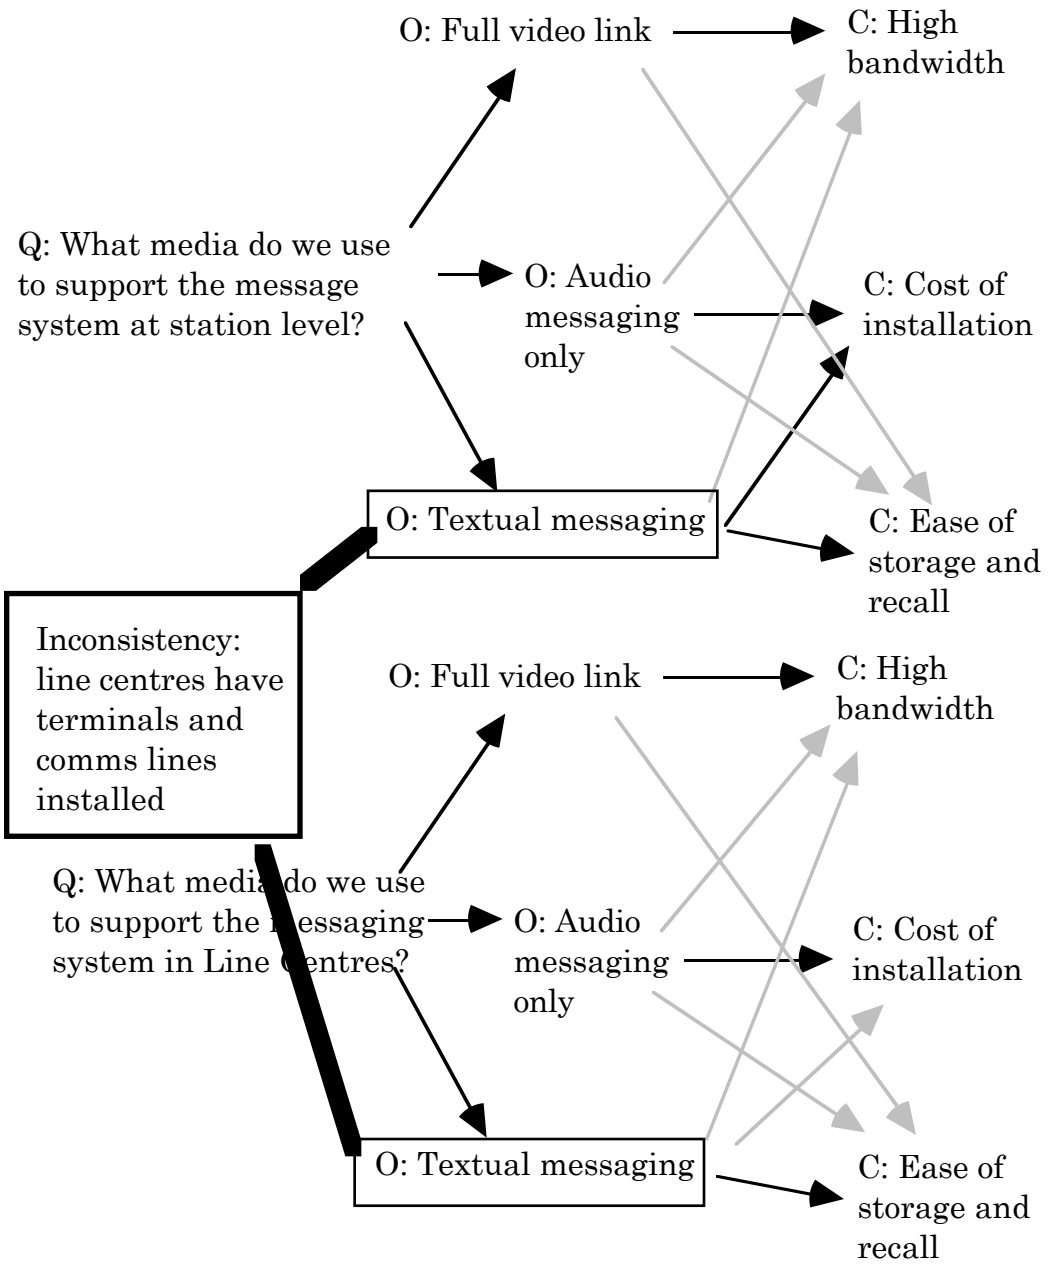

Figure 8: Providing Contextual Information To Explain Inconsistency

$$communication_lines_secure \Leftarrow redundant_channels \wedge redundant_power \quad (12)$$

This states that there are secure communications lines if redundant channels are provided and there exists secondary power supplies.

It is difficult to envisage how such additional detail could be introduced into the natural language descriptions of Figure 8 without introducing considerable additional complexity into a design rationale.

7 Incompleteness

Design rationales can be incomplete; options that are considered in one situation might be ignored for other similar questions. This creates problems if regulatory authorities use the resulting documents to determine whether sufficient attention has been paid to critical development decisions. The absence of particular options might be taken to indicate a lack of rigor during system

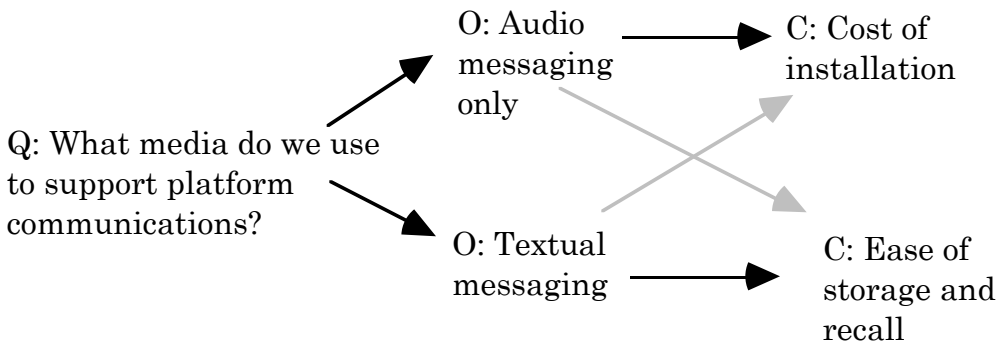

Figure 9: An incomplete QOC diagram

development. Figure 9 presents an incomplete QOC diagram. Only two options are considered for the on-platform communications systems: text and audio based messaging. This represents an incomplete rationale because the video option that was considered for both station and line communications is not considered for platform personnel. This is a significant omission; if video communications were installed at higher levels in the command structure then it should at least be considered at platform level. In terms of decision making, incompleteness raises a number of problems. By restricting the number of options that are available to a designer it increases the chances of discrepancies in the development process. For instance, video equipment might be provided everywhere in the system except on the platforms. This would effectively prevent some users from accessing important sources of information. The problems that are created by such incompleteness would be exacerbated if different sub-system were developed by different concurrent engineering groups, as proposed by Klein [15]. In our example, designers would have to support the transfer of information between video and text or speech based communications systems.

The design task here is to identify those questions which are sufficiently similar that the same options should be considered for all of them. Unfortunately, no automated procedures currently

exist for determining 'similarity' between design questions. In our example, experience and intuition suggest that the same media ought to be considered as options for the platform, line and station communications systems. We are, however, developing a set of heuristics to support this task. For example, the design of sub-systems that perform a similar function can often be guided by the same set of options. In the absence of such heuristics, decision theory provides a concise means of representing the criteria for completeness:

$$
\begin{aligned}
&complete_qoc(platform, line) \Leftarrow \\
&\quad considers(line, [(full_video; rich_bandwidth) \\
&\qquad (audio_only; low_cost; easy_recall) \\
&\qquad (text_only; low_cost; easy_recall)]) \wedge \\
&\quad considers(platform, [(full_video; rich_bandwidth) \\
&\qquad (audio_only; low_cost; easy_recall) \\
&\qquad (text_only; low_cost; easy_recall)] \qquad (13)
\end{aligned}
$$

This states that the rationale for the platform communications system is complete when compared with that for the line centres if both diagrams consider video, audio and text based options.

It is important to emphasise that designers can use clauses such as (13) to prove that the same options have been considered in both cases. This, in turn, provides means of showing regulatory authorities that the same degree of rigor has been used to examine the design space for similar questions. The use of decision theory also provides benefits for the members of other development teams. For instance, clauses such as (13) can be used to record the inter-dependencies between different areas of a design rationale. If new options are considered for the line centres and not the platforms then the resulting rationales may be incomplete with respect to each other:

$$
\begin{aligned}
&incomplete_qoc(platform, line) \Leftarrow \\
&\quad considers(line, [(full_video; rich_bandwidth) \\
&\qquad (audio_only; low_cost; easy_recall) \\
&\qquad (audio_text; low_cost; rich_bandwidth) \\
&\qquad (text_only; low_cost; easy_recall)]) \wedge \\
&\quad considers(platform, [(full_video; rich_bandwidth) \\
&\qquad (audio_only; low_cost; easy_recall) \\
&\qquad (text_only; low_cost; easy_recall)]
\end{aligned} \tag{14}
$$

This states that the rationale for the platform communications system is incomplete when compared with that for line centres if both diagrams consider video, audio and text based options but the platform rationale does not consider the use of mixed text and audio communications.

It is not clear how such completeness requirements could be captured in QOC diagrams without introducing a number of additional constructs into the graphical representations. This would considerably reduce the flexibility and accessibility of the semi-formal notation. Clauses, such as (14), provide an alternative that enables designers to retain the conceptual simplicity of design rationale and yet at the same time provide powerful reasoning tools for the analysis of safety-critical design documents.

8 Conclusion

This paper has argued that design rationale suffers from a number of problems that limit its application to support the design of safety-critical interfaces. We have illustrated this argument by performing a detailed analysis of the QOC notation. A particular problem with the graphical components of this semi-formal approach is that they do not record the priority or value that is associated with design options. This makes it difficult for regulatory authorities and other development teams to understand the objectives that guided the design of an interactive system. A further limitation is that there is no means of representing probability in design rationale. This can result in undue confidence being placed in the criteria that are promised for a particular option. We have shown that these weaknesses can be avoided by ranking the options in design rationale documents. This helps to explicitly represent design priorities and provides an

important link between HCI development techniques and established practice in probabilistic risk assessment, such as human reliability analysis. There are further benefits. The rankings derived from decision theory can also be used to determine the consistency and completeness of QOC diagrams. They can identify missing criteria and can be used to demonstrate that the same options have been considered for similar design questions.

Much work remains to be done. This paper describes part of a wider initiative to develop *literate design techniques*. Together with Phil Gray and Gilbert Cockton, Computing Science Department, Glasgow University, we are integrating formal, semi-formal and informal notations to support all stages of the design cycle. The intention is to build up a web of documents that record all stages in the development of interactive systems. This parallels Knuth's web for 'literate programming' [16]. The success of this work depends upon the provision of adequate tool support. At present, we have implemented a number of logic-based programming tools. These can be used to establish a direct correspondence between the elements of design rationale and the rankings derived from decision theory [12, 10]. They can also be used to automatically test the consistency and coherence of a rationale. Brevity prevents a complete description of this system, although an introduction is provided in [6, 11].

Our intention has been **not** to alter the existing components of design rationale. Previous attempts by the author to annotate these diagrams have reduced the readability and the flexibility of the resulting documents. In contrast, our intention has been to demonstrate that formal and semi-formal notations can be integrated to support the development of safety-critical, user interfaces. The rankings from decision theory and the graphical structures of QOC documents fill different but necessary roles in the design cycle. Such integrated approaches raise a number of questions for further research . In particular, social and managerial problems are created by the use of these techniques. The mathematical knowledge required to understand the model of decision theory will prevent many users and designers from accessing those documents. This, in turn, can make it difficult to validate the rankings against informal, vernacular descriptions [14, 13]. We have begun to address this problem by establishing clear links between the elements of formal analysis and the more accessible, graphical notations of design rationale.

Acknowledgements

This work is funded by the United Kingdom's Science and Engineering Research Council, grant number GR/J07686.

References

[1] T. Chan. A graphical tool to support the QOC notation. Technical report, Department of Computer Science, University of York, York, United Kingdom, 1994. M.Sc thesis.

[2] J. Conklin and M.L. Begeman. gIBIS: A tool for all reasons. *Journal Of The American Society For Information Science*, pages 200–213, May 1989.

[3] D. Fennell. *Investigation Into The Kings Cross Underground Fire*. Department of Transport, London, United Kingdom, 1988.

[4] International Atomic Energy Agency and The Commission of the European Community. *Critical Survey of Research On Human Factors And The Man-Machine Interaction*, IAEA-SM-26B/29, Vienna, Austria, 1984.

[5] C.W. Johnson. A probabilistic logic for the development of safety-critical interactive systems. *International Journal Of Man-Machine Studies*, 39(2):333–351, 1993.

[6] C.W. Johnson. Specifying and prototyping dynamic human-computer interfaces for stochastic applications. In J.L. Alty, D. Diaper, and S. Guest, editors, *People And Computers VIII*, pages 233–248. Cambridge University Press, Cambridge, United Kingdom, 1993.

[7] C.W. Johnson. The formal analysis of human-computer interaction during accidents investigations. In G. Cockton, S.W. Draper, and G.R.S. Weir, editors, *People And Computers IX*. Cambridge University Press, Cambridge, United Kingdom, 1994.

[8] C.W. Johnson. Decision theory and safety-critical interfaces. In K. Nordby, P.H. Helmersen, D. Gilmore, and S.A. Arnesen, editors, *Human Computer Interaction - Interact '95*, pages 127–132, London, United Kingdom, 1995. Chapman And Hall.

[9] C.W. Johnson. The economics of interface design. In K. Nordby, P.H. Helmersen, D. Gilmore, and S.A. Arnesen, editors, *Human Computer Interaction - Interact '95*, pages 19–25, London, United Kingdom, 1995. Chapman And Hall.

[10] C.W. Johnson. Integrating human factors and systems engineering to reduce the risk of operator 'error'. *To appear in the Journal Of Safety Science*, 1995.

[11] C.W. Johnson. Literate specification: Using design rationale to support formal methods in the large-scale development of human-machine interfaces. Technical report, Department of

Computing Science, University of Glasgow, Glasgow, United Kingdom, 1995. Submitted to Human Computer Interaction Journal.

[12] C.W. Johnson and M.D. Harrison. Using temporal logic to support the specification and prototyping of interactive control systems. *International Journal Of Man-Machine Studies*, 36:357–385, 1992.

[13] C.W. Johnson and M.D. Harrison. Software engineering for human computer interaction. *SIGCHI Bulletin*, 26(2):46–48, 1994.

[14] C.W. Johnson, J.C. McCarthy, and P.C. Wright. Using a formal language to support natural language in accident reports. *Ergonomics*, 38(6):1265–1283, 1995.

[15] M. Klein. Capturing design rationale in concurrent engineering teams. *IEEE Computer*, 26(1):39–47, 1993.

[16] D.E. Knuth. Literate programming. *Computer Journal*, 27(2):97–111, 1984.

[17] A. MacLean, R.M. Young, V.M.E. Bellotti, and T.P. Moran. Questions, options, and criteria: elements of design space analysis. *Human-Computer Interaction*, 6(3-4):201–50, 1991.

[18] Ministry Of Defence. *Requirements for the Procurement of Safety Critical Software*, MOD DEF-STAN 00-55, London, United Kingdom, 1991.

[19] J. Nielsen. *Usability Engineering*. Academic Press, New York, United States of America, 1993.

[20] President's Task Force On Aircraft Crew Compliment, United States' Government. *Report On Aircraft Crew Compliment*, Washington DC, United States of America, 1981.

[21] C. Puppe. *Distorted Probabilities And Choice Under Risk*. Lecture Notes In Economics And Mathematical Systems, No 363. Springer Verlag, Berlin, Germany, 1991.

[22] J. Rasmussen and O.M. Pedersen. Human factors in probabilistic risk analysis and in risk management. Technical Report Risø-N-18-83, The Risø National Laboratory, Roskilde, Denmark, 1983.

[23] S. Buckingham Shum. Cognitive dimensions of design rationale. In D. Diaper and N.Hammond, editors, *People And Computers VI: Proceedings Of HCI'91*. Cambridge University Press, Cambridge, United Kingdom, 1991.

[24] A.D. Swain and H.E. Gutman. Handbook of human reliability analysis with emphasis on nuclear power plant applications. Technical Report NUREG/CR-1278, United States' Nuclear Regulatory Commission, Washington DC, Unites States of America, October 1980.

[25] I.A. Watson. Review of human factors in reliability and risk assessment. In J. Burgoyne, editor, *The Assessment And Control Of Major Hazards*, pages 323–337. Pergamon Press, Oxford, United Kingdom, 1985.

[26] C.D. Wickens. *Engineering Psychology And Human Performance*. C.E. Merrill Publishing Company, London, United Kingdom, 1984.