

A Taxonomy of Fallacies in System Safety Arguments

William S. Greenwell; University of Virginia; Charlottesville, Virginia, USA

John C. Knight; University of Virginia, Charlottesville, Virginia, USA

C. Michael Holloway; NASA Langley Research Center; Hampton, Virginia, USA

Jacob J. Pease; University of Virginia; Charlottesville, Virginia, USA

Keywords: argumentation, assurance, failure analysis, safety cases, safety management, safety tools

Abstract

Safety cases are gaining acceptance as assurance vehicles for safety-related systems. A safety case documents the evidence and argument that a system is safe to operate; however, logical fallacies in the underlying argument may undermine a system's safety claims. Removing these fallacies is essential to reduce the risk of safety-related system failure. We present a taxonomy of common fallacies in safety arguments that is intended to assist safety professionals in avoiding and detecting fallacious reasoning in the arguments they develop and review. The taxonomy derives from a survey of general argument fallacies and a separate survey of fallacies in real-world safety arguments. Our taxonomy is specific to safety argumentation, and it is targeted at professionals who work with safety arguments but may lack formal training in logic or argumentation. We discuss the rationale for the selection and categorization of fallacies in the taxonomy. In addition to its applications to the development and review of safety cases, our taxonomy could also support the analysis of system failures and promote the development of more robust safety case patterns.

Introduction

Safety cases have evolved as a valuable approach to structuring the argument that a safety-critical system is acceptably safe to operate, and they have been developed in several application domains (ref. 1). The failure of a safety-critical system indicates that the risk of operating the system might be higher than previously thought. In the most general sense, a failure is either the result of an event that was anticipated but which was predicted to have a probability of occurrence below some critical threshold, or it was the result of an unanticipated event. If an unanticipated event occurred, then there must have been a defect in the safety case since the total risk exposure for operation of the system would have been based entirely on the random occurrences of anticipated events.

Our prior analyses of accidents involving safety-critical systems, including a minimum safe altitude warning (MSAW) system whose failure contributed to a major commercial aviation accident, suggested that system safety arguments sometimes invoke incomplete or inherently faulty reasoning (ref. 2). These fallacies, if undetected, could lead to overconfidence in a system and the false belief that the system's design has obviated or will tolerate certain faults. That a safety case might contain a flaw is to be expected, but it is important to bring attention to the problem and to remove defects to the extent possible.

In this paper, we discuss a very specific source of possible defects in safety cases, namely flawed arguments. Our informal review of the safety cases built for a set of important safety-critical applications showed that flawed arguments were present in each case. Based upon our observations and a survey of known logical fallacies, we developed a taxonomy of fallacious inferences in system safety arguments.

Fallacies in System Safety Arguments

From our prior analyses of failed systems (ref. 2) we hypothesized that logical fallacies are prevalent in system safety arguments. To test this hypothesis, we sampled publicly available industrial safety arguments and then analyzed each argument in our sample for fallacies. To obtain our sample, we conducted a survey of publicly available safety arguments, which yielded eight safety arguments in the disciplines of air traffic management, automotive engineering, commuter rail transit, electrical engineering, nuclear engineering, and radioactive waste storage (ref. 3). Of these, we selected three safety arguments for inclusion in our sample: the EUROCONTROL (EUR) Reduced Vertical Separation Minimums (RVSM) Pre-Implementation Safety Case, the Opalinus Clay geological waste repository safety case, and the EUR Whole Airspace Air Traffic Management (ATM) System

Safety Case. The organization of these three arguments made them amenable to analysis by individuals who did not possess expert knowledge of the relevant application domains. The EUR RVSM and whole airspace safety cases are preliminary and do not necessarily reflect final engineering standards; however it is still appropriate to examine the arguments for fallacies so that those fallacies can be addressed prior to implementation.

Two of the authors read and independently evaluated each of the three safety cases selected for the study. The purpose of this evaluation was two-fold: (1) to determine whether fallacies appear in these safety arguments with significant frequency; and (2) to identify the types of fallacies committed. Both of the reviewers had at least a basic understanding of fallacious reasoning from classical philosophical literature and drew from that knowledge in performing their evaluations. When a reviewer came across what he believed to be faulty reasoning in an argument, he highlighted the relevant section and recorded a brief note explaining why the reasoning was problematic. Upon completing their evaluations, the reviewers compiled their results into a comprehensive set of fallacies for each of the safety cases and then achieved a consensus as to which of the comments they made were indicative of fallacious reasoning in the arguments. The following sections summarize those results for each safety case using fallacy descriptions taken from Damer (ref. 4). Note that the goal of this study was to examine the form of the safety argument, not to evaluate the safety of the associated systems.

EUR RVSM: The EUR Reduced Vertical Separation Minimums (RVSM) Pre-Implementation Safety Case concerns a proposed reduction in the minimum amount of vertical distance that must be present between any two aircraft operating in EUR airspace. RVSM would accommodate a greater density of air traffic and would thus enable EUR to meet an expected increase in demand for air travel over the next several years. The RVSM safety case is organized as a natural language document but provides a graphical version of the safety argument in Goal Structuring Notation (GSN) in an appendix (ref. 5). To make the study tractable, we limited our evaluation to the GSN portion of the safety case, which spanned nine pages.

Table 1 summarizes the fallacies the reviewers identified in the EUR RVSM safety case with one column for each reviewer. Relevance fallacies were the most prevalent in the argument and accounted for two-thirds of the fallacies identified.

Table 1 — Tally of fallacies identified in the EUR RVSM safety case

<i>Fallacy</i>	<i>Reviewer A</i>	<i>Reviewer B</i>	<i>Total¹</i>
Using the Wrong Reasons	5	15	16
Drawing the Wrong Conclusion	3		3
Red Herring	1		1
Fallacious Use of Language	2	2	4
Hasty Inductive Generalization	4		4
Omission of Key Evidence		1	1
Total	15	18	29

Although the purpose of employing two reviewers in the case study was to assemble a more complete set of fallacies and not to examine the consistency between reviewers, the disparity between the results of each reviewer is significant. Differences are present both in the quantities and types of fallacies that each reviewer identified, the most notable of which concerns the fallacy *using the wrong reasons*. All but one of the instances of this fallacy concerned an inference that the argument invoked repeatedly. Reviewer A flagged only the first few of these instances before choosing to ignore them while reviewer B flagged them all. Moreover, both reviewers identified two instances of fallacious use of language due to ambiguity in the argument; however, the specific instances they identified did not overlap, suggesting that they had trouble agreeing upon which language was ambiguous or misleading. Finally, reviewer A identified a greater variety of fallacies than did reviewer B, which may be due to A's more extensive background and experience with logic and argumentation.

¹ Instances of a fallacy that were detected by both reviewers are only reflected once in the totals.




Figure 1 — Excerpt from the EUR RVSM safety case

A recurring problem in the RVSM safety case was its use of evidence that did not support the immediate claim the argument was attempting to make. As an example, Figure 1 presents a portion of the argument concerning the role of flight training in the RVSM safety requirements. Of the four premises supporting the claim that “there is sufficient direct evidence of [flight crew] training design validity,” (St.2.3.1) only one—G.2.3.1.4—pertains to the claim. The other premises state that various aspects of the training program have been specified, but this information does not aid the reader in determining whether the training design is valid.

Opalinus Clay: The Opalinus Clay safety case concerns the feasibility of constructing a long-term radioactive waste storage facility within the Opalinus Clay—a geological formation in the Zürcher Weinland of Switzerland. The argument claims that the Clay is sufficiently stable to enable the facility to meet its safety requirements for at least the next ten million years (ref. 6). The safety case spans 22 pages and is written in bulleted natural language with major safety claims enumerated as subsections accompanied by their corresponding arguments. It includes a variety of arguments for the feasibility and safety of the facility, including “...multiple arguments for safety that:

- Demonstrate safety and compliance with regulatory protection objectives;
- Use indicators of safety that are complementary to those of dose and risk and that show that radionuclide releases and concentrations due to the repository are well below those due to natural radionuclides in the environment;
- Indicate that the actual performance of the disposal system will, in reality, be more favorable than that evaluated in quantitative analyses; and
- No issues have been identified that have the potential to compromise safety” (ref. 6).

Both reviewers agreed that the Opalinus Clay safety case was the most compelling of the three arguments they reviewed. Indeed, reviewer B did not identify any fallacies in the argument while reviewer A identified only three, one of which was later determined to be valid reasoning. Table 2 shows the fallacies identified by reviewer A.

G2.3.1.1

Table 2 – Tally of fallacies identified in the Opalinus Clay safety case

<i>Fallacy</i>	<i>Reviewer A</i>	<i>Reviewer B</i>	<i>Total¹</i>
Arguing from Ignorance	1 ²		0
Omission of Key Evidence	2		2
Total	3	0	2

One of the arguments in the safety case discusses uncertainty in the characteristics of the chosen disposal system. It states that the choice of uncertainty scenarios to consider “remains a matter of expert judgment” and then describes the process in which scenarios were developed and considered using a panel of experts. A criticism of this approach would be that scenarios suggested by some experts were not selected for consideration but should have been. To avoid this criticism, the argument should mention some of the scenarios that were suggested but excluded from consideration by the panel along with its rationale for doing so. Elsewhere, the argument claims that “uncertainties [in the risk assessment cases] are treated using a pessimistic or conservative approach,” but no evidence is provided to support the claim of conservatism. Finally, in considering possible human intrusions, the argument assumes that “...possible future human actions that may affect the repository are constrained to those that are possible with present-day technology or moderate developments thereof” (ref. 6). Although it is difficult to imagine how one would avoid making this assumption, it is possible that unforeseen future innovations will render the analysis moot.

EUR Whole Airspace: The EUR Whole Airspace Air Traffic Management (ATM) System Safety Case preliminary study was conducted to evaluate “the possibility of developing a whole airspace ATM System Safety Case for airspace belonging to EUROCONTROL member states” (ref. 7). The study proposes arguments for preserving the current safety level of EUR airspace under a unified air traffic organization instead of the patchwork of organizations that comprise EUR today. We are aware that the arguments presented in the EUR safety case are preliminary; nevertheless, we think it is appropriate to examine them because operational arguments will likely be derived from them. Like the RVSM safety case, the report presents mostly natural language arguments but does make use of GSN in some places. Again, we chose to limit our review to the major GSN elements of the report, which spanned two pages. These included two separate arguments for the safety of the whole airspace: one based on arguing over individual geographic areas and one based on reasoning about the whole airspace ATM rules. Both arguments shared the same top-level goal that “the airspace is safe.”

Table 3 contains the results of the reviewers’ evaluations of the EUR Whole Airspace argument, which reflect the combined fallacies in both the argument over geographic regions and the argument for the safe implementation of whole airspace rules.

Table 3 – Tally of fallacies identified in the EUR Whole Airspace safety case

<i>Fallacy</i>	<i>Reviewer A</i>	<i>Reviewer B</i>	<i>Total¹</i>
Red Herring	4		4
Fallacious Use of Language	2	6	6
Fallacy of Composition	2	2	2
Omission of Key Evidence	2		2
Total	10	8	14

Neither argument considered possible interactions between geographic areas, such as when an aircraft is handed off by one air traffic controller to another in an adjacent region. Even if the safety rules are respected within each region, without special considerations for interactions between regions and with external airspace, the rules might be violated in the context of the broader whole airspace system. Both reviewers flagged these arguments as instances of fallacious composition.

² The reviewers later agreed that the reasoning labeled as arguing from ignorance was not fallacious.

A Taxonomy of Safety-Argument Fallacies

All three of the case studies we reviewed exhibited common types of faulty reasoning, supporting our hypothesis that logical fallacies are prevalent in safety arguments. To facilitate detection of these fallacies, we developed a taxonomy of safety-argument fallacies based upon existing taxonomies described in the philosophical literature that we adapted according to our observations from the survey. Our specific objectives in developing the taxonomy were:

1. To cover a broad range of fallacies but only those that are relevant to safety argumentation;
2. To categorize the taxonomy so that a user may determine the set of fallacies that might pertain to a given argument without learning the entire set of fallacies in the taxonomy;
3. To define the fallacies so that they are accessible to safety professionals who have not received formal training in logic or argumentation; and
4. To design the taxonomy for extensibility.

Coverage: Several taxonomies of fallacies in general arguments exist; we surveyed those of Damer, Curtis, Dowden, Pirie, and Govier to develop our taxonomy of safety-argument fallacies (refs. 4, 8-11). These taxonomies place no restrictions on the types of arguments they consider, and so they include emotional appeals, malicious fallacies that convey acts of willful deception, and formal, syllogistic, and causal fallacies. We assumed that safety arguments do not contain emotional appeals for their acceptance or willful attempts at deception. Formal and syllogistic fallacies, which occur in deductive arguments, are unlikely to appear in safety arguments because purely deductive arguments may be expressed formally and verified mechanically. Likewise, safety arguments rarely attempt to establish causal relationships between events (with the exception of inferring a causal relationship between correlated events), and so causal fallacies are improbable. Based on these assumptions, we excluded these fallacies from our taxonomy. Table 4 summarizes the fallacies we excluded, and Table 5 provides examples of fallacies we excluded.

Table 4 – Excluded fallacies grouped by source

<i>Category</i>	<i>Damer</i>	<i>Curtis</i>	<i>Dowden</i>	<i>Pirie</i>	<i>Govier</i>
Fallacies Defined by Source	60	65	107	75	32
Emotional Appeals	-7	-3	-9	-4	-2
Malicious Fallacies	-9	-8	-15	-11	-4
Formal & Syllogistic Fallacies	-7	-19	-8	-11	-3
Causal Fallacies	-4	-1	-3	-1	-3
Other Excluded Fallacies	-7	-11	-32	-16	-3
Collapsed Fallacies	-8	-3	-12	-9	-2
Fallacies Represented	18	20	28	23	15

Table 5 – Examples of excluded fallacies

<i>Category</i>	<i>Examples</i>
Emotional Appeals	<ul style="list-style-type: none"> • Argument from Outrage • Scare Tactic • Misleading Accent • Style Over Substance
Malicious Fallacies	<ul style="list-style-type: none"> • Appeal to Force • Poisoning the Well • <i>Ad Hominem</i> • Straw Man
Formal & Syllogistic Fallacies	<ul style="list-style-type: none"> • Affirming the Consequent • Four-Term Fallacy

	• Denying the Antecedent	• Undistributed Middle Term
Causal Fallacies	• <i>Post Hoc ergo Propter Hoc</i>	• Reversing Causation
Other Excluded Fallacies	• Drawing the Wrong Conclusion	• Regression
	• Complex Question	• Scope Fallacy
	• Fallacy of the Continuum	• Special Pleading
	• Hasty Inductive Generalization	• <i>Tu quoque</i>
	• Refuting the Example	• Wishful Thinking

After excluding emotional, malicious, formal, syllogistic, and causal fallacies, two of the authors assessed the remaining fallacies individually and excluded those that were unlikely to appear in safety arguments for various reasons. These fallacies appear in the “Other Excluded Fallacies” category in Table 4, and examples appear in Table 5. For example, *wishful thinking* was excluded because it concerns arguments in which a claim is asserted to be true on the basis of a personal desire or vested interest in it being true. Such an argument is very unlikely to appear explicitly in a safety argument. *Hasty inductive generalization*, which occurs when an argument offers too little evidence in support of its claim, was excluded because its broad definition encompasses most of the other fallacies. Other fallacies such as *refuting the example* were omitted because they pertain to refutations, which seldom appear in safety cases. Finally, fallacies whose definitions differed only subtly from each other were collapsed into a single fallacy; these appear in the row marked “Collapsed Fallacies.”

There was a strong degree of overlap in the fallacies that remained from each of the five taxonomies we surveyed. We consolidated these fallacies into a final set of 33 fallacies, which is presented in Table 6.

Table 6 – The safety-argument fallacy taxonomy

<p><i>Circular Reasoning</i></p> <ul style="list-style-type: none"> Circular Argument Circular Definition <p><i>Diversivory Arguments</i></p> <ul style="list-style-type: none"> Irrelevant Premise Verbose Argument <p><i>Fallacious Appeals</i></p> <ul style="list-style-type: none"> Appeal to Common Practice Appeal to Improper/Anonymous Authority Appeal to Money Appeal to Novelty Association Fallacy Genetic Fallacy <p><i>Mathematical Fallacies</i></p> <ul style="list-style-type: none"> Faith in Probability Gambler’s Fallacy Insufficient Sample Size Pseudo-Precision Unrepresentative Sample <p><i>Unsupported Assertions</i></p> <ul style="list-style-type: none"> Arguing from Ignorance Unjustified Comparison Unjustified Distinction 	<p><i>Anecdotal Arguments</i></p> <ul style="list-style-type: none"> Correlation Implies Causation Damning the Alternatives Destroying the Exception Destroying the Rule False Dichotomy <p><i>Omission of Key Evidence</i></p> <ul style="list-style-type: none"> Omission of Key Evidence Fallacious Composition Fallacious Division Ignoring Available Counter-Evidence Oversimplification <p><i>Linguistic Fallacies</i></p> <ul style="list-style-type: none"> Ambiguity Equivocation Suppressed Quantification Vacuous Explanation Vagueness
--	---

Omission of key evidence appears in the taxonomy both category and as an entry because there are many special forms of this fallacy. *Fallacious composition* occurs when an argument attempts to infer the properties of a system from those of its components without considering interactions between the components that might violate those

properties, and *fallacious division* occurs when an argument attempts the converse. An argument *ignores available counter-evidence* when it makes a claim for which there exists refuting evidence but fails to address that evidence. *Oversimplification* describes arguments that cite evidence obtained from models of system behavior but fail to show that the models correspond to the system in question.

Categorization: We initially categorized the fallacies in our taxonomy according to Damer's categories of *relevance*, *acceptability*, and *sufficiency*. Relevance fallacies concern the use of premises that have no bearing on the truth of the claims they ostensibly support, acceptability fallacies concern the use of inherently faulty inferences, and sufficiency fallacies describe ways in which arguments can fail to provide enough evidence in support of their claims. Damer's topology exhibited three problems that contradicted our goals of making the taxonomy accessible, however. First, the categories did not correspond to the types of arguments one might encounter in a safety case, and so they provided little help in determining the set of fallacies that might pertain to a given argument. Second, the categories required a user of the taxonomy to know *a priori* the type of fallacy committed, which does not aid users who apply the taxonomy in order to determine whether an argument is fallacious. Third, the domain of safety argumentation presented special challenges in assigning the fallacies to these categories unequivocally. Many of the fallacies Damer classified as relevance or acceptability fallacies, such as fallacious composition and division, could be remedied by supplying additional evidence, suggesting that in some cases they would be better-classified as sufficiency fallacies. We also considered the topologies employed by the other taxonomies we surveyed, but they suffered similar limitations.

Since we were unable to find a suitable topology to adapt to our taxonomy, we developed our own by inferring relationships among the fallacies with respect to the types of arguments they address. Our topology—shown in Table 6—groups fallacies into eight categories, which are summarized below:

- *Circular reasoning* occurs when an argument is structured so that it reasserts its claim as a premise or defines a key term in a way that makes its claim trivially true.
- *Diversionary arguments* contain excessive amounts of irrelevant material that could distract a reader from a weakly supported claim.
- *Fallacious appeals* invoke irrelevant authorities, concepts, or comparisons as evidence.
- *Mathematical fallacies* describe common pitfalls in probabilistic and statistical inferences.
- *Unsupported assertions* are claims stated without evidence.
- *Anecdotal arguments* show that their claims hold in some circumstances but fail to generalize their validity.
- *Omission of key evidence* occurs when an otherwise complete argument omits evidence that is necessary to establish its validity.
- *Linguistic fallacies* concern the use of misleading language that might lead the reader to an unwarranted conclusion. These fallacies may appear in any informal argument.

A user of the taxonomy may compare the argument he is reviewing to each of the categories in the taxonomy to assess the argument's validity. For example, the user might first examine the structure of the argument for circularity and then evaluate the relevance of its premises. A verbose argument might contain diversionary premises, and appeals to regulatory standards, practices, conventions, or authorities such as regulatory agencies should be checked to ensure that they are relevant to the context of the argument. If the argument relies upon statistical evidence, then the user should examine the conclusions that it draws from that evidence for mathematical fallacies. Unsupported assertions and anecdotal evidence may suggest circumstances in which the argument's claims do not hold. If the argument follows an accepted pattern of reasoning, the user should verify that it has properly instantiated the pattern and not omitted evidence. Finally, the user must be wary of vague or ambiguous terms in the argument because different parts of the argument might interpret these terms differently and lead the user to an unwarranted conclusion.

Organizing the fallacies by the types of arguments in which they appear instead of the manner in which they undermine arguments addresses the shortcomings we identified with Damer's topology. For a given argument, a user may assess the argument with respect to the categories that describe it and then determine the set of fallacies that might pertain to the argument. Thus, users must only be familiar with the categories in the taxonomy and not the

Arguing from Ignorance

An argument supports a claim by citing a lack of evidence that the claim is false. The argument does not exhibit the fallacy if it cites as evidence a sufficiently-exhaustive search for counter-evidence that has turned up none.

Example: “All of the credible hazards have been identified. Aside from the hazards noted earlier, no evidence exists that any other hazards pose a threat to the safe operation of the system.”

This argument attempts to prove a negative (that there are no additional credible hazards to system operation) by noting a lack of evidence contradicting the negative. It does not cite any efforts that have been made to discover such evidence. A mere lack of evidence that a claim is false does not make it true.

Figure 2 - Sample taxonomy entry

entire set of fallacies in order to apply the taxonomy, and they are not required to know *a priori* that an argument is fallacious. This organization also improves the orthogonality of the topology because the type of argument in which a fallacy is likely to appear is relatively static, whereas the manner in which it undermines the argument depends upon the context in which the fallacy occurs.

Fallacy Definitions: For brevity we omit the definitions of the fallacies from this paper, but Figure 2 provides a sample definition, and interested readers may consult our documentation of the taxonomy (ref. 12). In defining the fallacies we followed the format used in each of the taxonomies we surveyed. Each entry in the taxonomy consists of a short name of the fallacy, a definition, safety-related examples of the fallacy, and an exposition of the examples. The examples are intended to demonstrate real-world instances of the fallacies, and in many cases they have been adapted from actual safety arguments.

Completeness & Extensibility: Despite our survey of real-world safety arguments, some of the fallacies we excluded from the taxonomy might appear with sufficient frequency to warrant inclusion, and there might exist fallacies that neither we nor the taxonomies we surveyed considered. To reduce this risk, we surveyed five different fallacy taxonomies in order to include a broad range of fallacies in our analysis, and the strong degree of overlap between the taxonomies we surveyed indicates that there is general agreement in the philosophical community as to which fallacies typically appear in arguments. Nevertheless, we designed the taxonomy so that it would be extensible. New fallacies may be added to each category, and new categories may also be added provided they respect the existing orthogonal topology. In addition, specific forms of the fallacies defined in the taxonomy, such as those that are relevant to a particular safety domain, may be added as child elements of those fallacies.

Overlap: Overlap between fallacies refers to scenarios in which an inference exhibits multiple fallacies. It can arise either when multiple aspects of an inference are fallacious or when the fallacies' definitions are not mutually exclusive. In the latter sense, overlapping fallacies are problematic if the strategies for removing them are incompatible. In Damer's classification, for example, an overlap between a relevance fallacy and a sufficiency fallacy would lead to the dilemma of either removing an inference because it was irrelevant or adding additional support to the argument in order to make it sufficient. Since the categories of our taxonomy are largely orthogonal, overlap in this sense is unlikely to occur between fallacies in different categories, and fallacies that belong to the same category share similar repair strategies. Moreover, we consolidated fallacies whose definitions contained only subtle differences in order to reduce the likelihood of overlap in our taxonomy.

Applications

The major applications of our taxonomy are to safety-case development, pre-acceptance review, and failure analysis. System developers rely upon knowledge gained from previous development efforts and observed system failures in choosing the development practices and evidence that are necessary to ensure the safety of the systems they build. Likewise, regulatory agencies rely upon industrial experience and recommendations from accident investigation boards in specifying safety standards. Thus, as systems are deployed with unknown faults and failures are observed, lessons are learned from those failures that are then incorporated into future development projects. The safety case can facilitate this process because it contains the rationale for concluding that a system was safe to operate prior to

an observed failure. If the failure was systemic, then the safety case is flawed and should be repaired. We refer to this process as the *enhanced safety-case lifecycle*, and we have developed a process for applying the fallacy taxonomy to a safety case in light of an observed failure in order to discover the fallacies in the safety argument that might have contributed to the failure (ref. 13). New fallacies may be added to the taxonomy as they are discovered, and so in addition to its applications as a preventative tool, the taxonomy is also a means by which lessons may be disseminated from failure analyses to those who develop and certify safety-related systems.

Conclusions & Future Work

As informal arguments, safety cases are prone to several forms of fallacious reasoning. Fallacies in a system's safety argument could undermine the system's safety claims, which in turn could contribute to a safety-related failure of the system. Avoiding these fallacies during system development and detecting them during review is essential if failures are to be prevented. Our assessment of three commercial safety cases revealed that they exhibited a wide variety of informal fallacies. Based on our observations, we surveyed five existing taxonomies of fallacies in general arguments to produce a new taxonomy of fallacies specific to system safety arguments. Our taxonomy is targeted at safety professionals who may lack formal training in logic and argumentation, and it is organized so that these individuals may apply it to an argument without complete knowledge of the taxonomy and without knowing *a priori* whether the argument is fallacious. Although the taxonomy is not exhaustive, it is extensible so that new fallacies may be added to it as they are discovered in safety arguments.

We plan to conduct a series of controlled trials involving students and professional engineers to evaluate the extent to which the taxonomy improves a reviewer's accuracy in detecting fallacious safety arguments. The trials will consist of asking subjects to review a set of safety arguments that have been randomly introduced and to mark which arguments they determine to be fallacious. The subjects' responses will be scored, and the results from the control group will be compared to those of the experimental group to measure the benefit afforded by the taxonomy.

Acknowledgements

This work was funded in part by NASA Langley Research Center under grant number NAG-1-02103.

References

1. Bishop, P. and R. Bloomfield. 1998. A methodology for safety case development. In Industrial Perspectives of Safety Critical Systems: Proc. Sixth Safety-Critical Systems Symposium. Birmingham. Springer-Verlag.
2. Greenwell, W.S., E.A. Strunk, and J.C. Knight. 2004. Failure analysis and the safety case lifecycle. In Proc. 7th Working Conference on Human Error, Safety, and Systems Development.
3. Dependability Research Group. 2004. Safety case repository. Department of Computer Science, University of Virginia. 2004. <http://dependability.cs.virginia.edu/research/safetycases/safetycasesexamples.php> (accessed March 31, 2006).
4. Damer, T.E. 2005. Attacking Faulty Reasoning: A Practical Guide to Fallacy-Free Arguments. Australia: Wadsworth.
5. Kelly, T. and R. Weaver. 2004. The Goal Structuring Notation - a safety argument notation. In Proc. DSN Workshop on Assurance Cases: Best Practices, Possible Outcomes, and Future Opportunities.
6. Nagra. 2002. Project Opalinus Clay: safety report. http://dependability.cs.virginia.edu/research/safetycases/Opalinus_Clay.pdf (accessed March 31, 2006).
7. Kinnorsly, S. 2001. Whole Airspace ATM safety case - preliminary study. http://dependability.cs.virginia.edu/research/safetycases/EUR_WholeAirspace.pdf (accessed March 31, 2006).
8. Curtis, G. 2006. Fallacy files. <http://www.fallacyfiles.org/index.html> (accessed March 31, 2006).
9. Dowden, B. 2005. Fallacies. In Internet Encyclopedia of Philosophy. <http://www.iep.utm.edu/> (accessed March 31, 2006).

10. Pirie, M. 1985. The book of the fallacy. London: Routledge & Kegan Paul.
11. Govier, T. 2005. A Practical Study of Argument. Australia: Wadsworth.
12. Greenwell, W.S., J.J. Pease and C.M. Holloway. 2005. Safety-argument fallacy taxonomy. <http://www.cs.virginia.edu/~wsg6p/docs/taxonomy.pdf> (accessed March 31, 2006).
13. Greenwell, W.S. and J.C. Knight. Failure analysis and the safety case lifecycle. <http://www.cs.virginia.edu/~jck/publications/greenwell.ress06.pdf> (accessed March 31, 2006).

Biographies

W. S. Greenwell, Department of Computer Science, University of Virginia, P.O. Box 400470, Charlottesville, VA 22904-4740, telephone – (434) 982-2292, facsimile – (434) 982-2214, e-mail – greenwell@cs.virginia.edu.

Mr. Greenwell is a PhD candidate in computer science and a member of the Dependability Research Group at the University of Virginia. His research interests include software safety assurance and the analysis of software failure. He is advised by Dr. John Knight.

J. C. Knight, Department of Computer Science, University of Virginia, P.O. Box 400470, Charlottesville, VA 22904-4740, telephone – (434) 982-2216, facsimile – (434) 982-2214, e-mail – knight@cs.virginia.edu.

Dr. Knight is a professor of computer science and faculty member of the Dependability Research Group at the University of Virginia. His research interests are in software dependability and specifically the areas of formal methods, the use of natural language in requirements analysis and formal specification, and the survivability of critical networked infrastructures.

C. M. Holloway, NASA Langley Research Center, 100 NASA Road, Hampton, VA 23681-2199, telephone – (757) 864-1701, facsimile – (757) 864-4234, e-mail – c.m.holloway@nasa.gov.

Mr. Holloway is a senior research engineer at NASA Langley Research Center. His primary professional interests are system safety and accident analysis for software intensive systems.

J. J. Pease, Department of Philosophy, University of Virginia, P.O. Box 400780, Charlottesville, VA 22904-4780, telephone – (434) 924-7701, facsimile – (434) 924-6927, e-mail – jpease@virginia.edu.

Mr. Pease is a graduate student of philosophy at the University of Virginia.