The UK Nuclear Industry Guide to:

Appropriate Conservatism in Safety Cases

This Nuclear Industry Guide was produced by the Safety Case Forum and published on behalf of the Nuclear Industry Safety Directors Forum (SDF)

JULY 2015

Revision History

Issue Number	Revision Date	Changes
1	July 2015	First issue.

It is recognised that – through experience of using this Safety Case Forum Guide – there may be comments, questions and suggestions regarding its contents.

In the first instance, any such comments should be sent to the following:

alec.bounds@arcadis-uk.com
dave.graham@magnoxsites.com
keith.murphy@awe.co.uk

Executive Summary

Conservatism is ‘being on the safe side’, for example estimating a potential radiation dose in an accident scenario to be higher than the best estimate of the dose. Conservatism is required for a safety case methodology called Design Basis Analysis (DBA). However, *over*-conservatism is relatively common.

The symptoms of an overly-conservative safety case include:

- The safety case is larger and more complex than expected, with excessive production costs.
- Many shortfalls are identified, while a similar (or higher hazard) facility elsewhere that has gone through a valid safety case process has very few or no shortfalls.
- The requirement for safety measures (both engineered and operational) is very high in comparison to a similar (or higher hazard) facility elsewhere.
- The safety case owner and other stakeholders find the safety case to be not credible.

Perhaps surprisingly, over-conservatism, as well as under-conservatism, can be the enemy of safety.

Overall conservatism is what is required for Design Basis Analysis; it is not necessary for every element in the analysis to be demonstrably conservative. Combining conservative values of multiple input values (e.g. inventory) can easily lead to a ridiculously unrealistic output value (e.g. dose), *without the knowledge of the analyst*.

Techniques for removing or managing under/over-conservatism that are significant throughout this SCF Guide include:

- explicit identification of optimisms and conservatisms within safety case documents;
- author self-checking of documents for under-conservatisms or over-conservatisms;
- good communication, and avoiding ‘silo working’, in particular it is important to seek the views of facility operators during fault identification and fault development;
- awareness of the multiplicative effect of over-conservatisms;
- use of sensitivity analysis;
- appropriate use of expert judgement; and
- appropriate challenge of assumptions and conclusions (asking ‘Does it look right?’), within a team-working environment.

Safety Directors Forum

In a sector where safety, security and the protection of the environment is, and must always be the number one priority, the Safety Directors' Forum (SDF) plays a crucial role in bringing together senior level nuclear executives to:

- Promote learning;
- Agree strategy on key issues facing the industry;
- Provide a network within the industry (including with government and regulators) and external to the industry;
- Provide an industry input to new developments in the industry; and,
- To ensure that the industry stays on its path of continual improvement.

It also looks to identify key strategic challenges facing the industry in the fields of environment, health, safety, quality safeguards and security (EHSQ&S) and resolve them, often through working with the UK regulators and DECC, both of whom SDF meets twice yearly. The SDF members represent every part of the fuel cycle from fuel manufacture, through generation to reprocessing and waste treatment, including research, design, new build, decommissioning and care and maintenance. The Forum also has members who represent the Ministry of Defence nuclear operations, as well as "smaller licensees" such as universities and pharmaceutical companies. With over 25 members from every site licence company in the UK, every MoD authorised site and organisations which are planning to become site licensees the SDF represents a vast pool of knowledge and experience, which has made it a key consultee for Government and regulators on new legislation and regulation.

The Forum has a strong focus on improvement across the industry. It has in place a number of subject-specific sub-groups looking in detail at issues such as radiological protection, human performance, learning from experience and the implementation of the new regulatory framework for security (NORMS). Such sub groups have developed a number of Codes of Practice which have been adopted by the industry.

SDF Codes of Practice and Guides are available on this link:
http://www.nuclearinst.com/Publications

Safety Case Forum

This document is produced by the Safety Case Forum, which is a sub-group of the Safety Directors' Forum. The Safety Case Forum was established in June 2012 and brings together a wide range of representatives of nuclear operators, from all the Licensees and Authorisees across the United Kingdom, including:

- Civil, commercial and defence activities;
- Design, operation and decommissioning of nuclear facilities;
- Research facilities.

The purpose of the Safety Case Forum is to provide guidance that is useful to, and will benefit the widest possible range of UK nuclear operators.

Such guidance is not mandatory, nor does it seek to identify minimum standards. It aims to provide a tool kit of methods and processes that nuclear operators can use if appropriate to their sites and facilities.

These guides are intended to improve the standardisation of approach to the delivery of fit for purpose safety cases, while improving quality and reducing the cost of production. They are

designed to cater for all stages of a facility's life cycle and for all processes within that life cycle. This includes any interim, continuous and periodic safety reviews, allowing for the safe and efficient operation of nuclear facilities.

When using the information contained within these guides, the role of the Intelligent Customer shall always remain with the individual nuclear operator, which shall retain responsibility for justifying the arguments in their respective Safety Cases. The Office for Nuclear Regulation is a consultative member of the Safety Case Forum.

The following companies and organisations are participating members of the Safety Case Forum:

SCF Codes of Practice and Guides are available on this link:
http://www.nuclearinst.com/SDF-safety-cases

Disclaimer

This UK Nuclear Industry Guide has been prepared on behalf of the Safety Directors Forum by a Technical Working Group. Statements and technical information contained in this Guide are believed to be accurate at the time of writing. However, it may not be accurate, complete, up to date or applicable to the circumstances of any particular case. This Guide is not a standard, specification or regulation, nor a Code of Practice and should not be read as such. We shall not be liable for any direct, indirect, special, punitive or consequential damages or loss whether in statute, contract, negligence or otherwise, arising out of or in connection with the use of information within this UK Nuclear Industry Guide.

This Guide is produced by the Nuclear Industry. It is not prescriptive but offers guidance and in some cases a toolbox of methods and techniques that can be used to demonstrate compliance with regulatory requirements and approaches.

Contents

Introduction

Aims

Safety Case Forum Guides are produced by representatives of nuclear operators (nuclear site licensees and other companies with nuclear operations in the UK). Their purpose is *to provide guidance that is useful to a wide range of UK nuclear operators*. Such Guides do not set mandatory requirements on any nuclear operator, nor do they identify minimum standards. Guides provide a tool kit of methods and processes that nuclear operators can use if appropriate to their sites and facilities. The responsibility for justifying arguments in safety cases remains with nuclear operators.

The Safety Case Forum reports to the UK nuclear industry Safety Directors Forum. The companies represented at the Safety Case Forum include companies that cover:

- civil and defence activities
- design, operation and decommissioning of nuclear facilities
- low hazard and high hazard nuclear facilities

The ***purpose*** of this Guide is firstly to help safety case practitioners be aware of conservatism, including under-conservatism and over-conservatism, and the potential effects of getting the level of conservatism wrong.

The second purpose, which builds on the first purpose, is to help safety case practitioners use conservatism when it is required, and to achieve a balance between under-conservatism and over-conservatism. The guidance should also help other people involved with safety cases, e.g. those who check or review safety cases.

Scope

The scope focuses on Design Basis Analysis (DBA), including safety assessment and engineering substantiation. Conservatism is not required for probabilistic safety analysis (PSA) or for severe accident analysis (SAA), which should normally use best-estimate approaches. However, conservatism does get used in PSA and SAA (sometimes inappropriately), and this is included in the scope. In addition, the effect of conservatism on design and on normal doses is included. DBA, PSA and SAA as safety assessment techniques are excluded from this SCF Guide.

The scope of the guide includes some statistical analysis relating to log-normal distributions. The conclusions from this analysis are clear in the Guide, so it is not necessary to be a statistician in order to understand the key points of the Guide.

The examples throughout this SCF Guide are based on real examples in the nuclear industry. There are more examples of over-conservatism than under-conservatism in this SCF Guide, mainly because a tendency to over-conservatism is currently apparent in many safety cases. However, examples of under-conservatism are also given, and general advice includes both over-conservatism and under-conservatism.

An alternative to using a conservative approach is to use ‘Best Estimate Plus Uncertainty’ (BEPU). This is described for nuclear power plants in Reference 3. It is not specifically included in this SCF Guide, although some BEPU-related concepts are included.

Terminology/Key Concepts

Conservatism

In safety analysis, conservatism is an approach where the use of models, data and assumptions would be expected to lead to a result that bounds the best-estimate (where known) on the safe side. For example, the conservatively assessed dose to a worker in a fault scenario may be 60 mSv, while for the same scenario the best-estimate dose may be 15 mSv.

The above definition (taken from Reference 1) assumes that 'best-estimate' is well understood. In fact, it is not clear whether it means:

- the **mean**, i.e. the arithmetic mean (the sum of the values divided by the number of values), not the geometric mean;
- the **median** (the numerical value separating the higher half of a data sample, a population, or a probability distribution, from the lower half); or
- the **mode** (the number which appears most often in a set of numbers).

This is explored in Appendix A, which states that the mean is the most useful type of average for safety case practitioners.

Worst case is another phrase used. It implies that there is no credible case that is worse. In some cases, a worst-case value may be adopted as the conservative value, but it is important to note that *worst case* may be overly conservative.

In this Guide, the phrase '*a* conservatism' refers to a specific instance of applying the conservative approach.

An 'over-conservatism' is a conservatism that is meant to address an uncertainty, but its scale is in fact greater than is actually required to address the uncertainty. See examples in the section 'Achieving a Balance between Under-conservatism and Over-conservatism, including Examples'.

An 'under-conservatism' is a conservatism that is meant to address an uncertainty, but its scale is in fact not great enough to address the uncertainty. See examples in the section 'Achieving a Balance between Under-conservatism and Over-conservatism, including Examples'.

An 'optimism' is the opposite of a conservatism, and refers to an approach where the use of models, data and assumptions would be expected to lead to a result that is on the unsafe side of the best-estimate (where known).

Uncertainty

In safety analysis, there are usually many uncertainties. There are known uncertainties, e.g. in a statistical distribution (which cannot be avoided), but there are also unknown uncertainties (which need to be avoided and minimised as far as practicable). There can be uncertainties about:

- fault development;
- the level of consequence arising from an accident scenario, even a well-defined one;
- the engineering performance of structures, systems and components; and
- many other aspects.

In reality, there are very few elements of a safety argument where there is no uncertainty at all. Thus ***management of uncertainties is a fundamental part of safety analysis***.

Appendix A describes the characteristics of statistical uncertainty, and how this affects safety cases. There are also uncertainties that do not directly relate to parameters, such as assumptions about how a fault sequence progresses - these uncertainties can be more significant than statistical uncertainties.

Application/Readers Guide

SCF Guides are written for suitably qualified and experienced safety case practitioners. Thus the basics of safety cases and safety assessment as they affect a particular SCF Guide are not necessarily explained in SCF Guides. Safety case practitioners are expected to use their judgement in applying SCF Guides, taking account of company-specific, site-specific, building-specific and scenario-specific factors.

Introduction to Subject Matter

There are two main reasons why a conservative approach is adopted in safety cases:

- To address uncertainty.
- For convenience.

Uncertainty is addressed by using a conservative approach. Paragraph 607 of Reference 1 states: "*In DBA, any uncertainties in the fault progression and consequence analyses are addressed by the use of appropriate conservatism.*" This conservatism is important in order to demonstrate safety effectively, and it is most important when there are cliff-edge effects. For example, if a liquid might have a range of activities, failure to take account of the range (e.g. by simply using the average activity) might mean that a DBA consequence threshold, e.g. 20 mSv to a worker is presented as simply not exceeded, whereas in fact there is a significant probability that it is exceeded, and the DBA (and the whole safety assessment) may then be inadequate. But if at the top of the range of uncertainty, the liquid is so radioactive that it is significantly self-heating, then this is a cliff-edge effect that must be addressed, as there will be additional fault sequences that would otherwise be omitted, resulting in a seriously inadequate safety assessment. Thus under-conservatism can be a major weakness of a safety case. Principle SC.5 in Reference 1 states '***Safety cases should identify areas of optimism and uncertainty, together with their significance, in addition to strengths and any claimed conservatism***', and this important Principle is supported in this SCF Guide.

The 'convenience' angle is an important factor in relation to conservatism. For example, when assessing sealed sources, it may be convenient to treat all the sources as if they have the highest inventory, in a single bounding case. As long as this does not cause any problems, e.g. requiring a too onerous approach to the lower inventory sources, this conservatism is not only harmless, it's also an approach to be recommended, simplifying the safety case significantly, and reducing the amount of effort required to make the safety case. Another common 'conservatism for convenience' is to assume that consequences for PSA are the same as consequences for DBA. Understanding the level of conservatism in references and data sources is important for understanding the overall conservatism in any particular case.

It is sometimes unclear whether a conservatism is used to address uncertainty, or for convenience, or for a mixture of the two. And in some cases, the author does not

make it clear that a conservatism has been used. In general, conservatisms (and areas of optimism or uncertainty) should be identified as such.

In general, it is helpful to reduce the unquantified conservatism in order to demonstrate a larger quantified margin to a safety case limit. This is especially beneficial if the margins are eroded or threatened for some reason.

Relevant legislation

This document has been generated with consideration of relevant health and safety legislation. Where appropriate legislation has been referenced, but the primary legislation that has influenced this document is:

- Health and Safety at Work etc Act 1974
- The Nuclear Installations Act 1965

The Nuclear installations Act 1965 requires the licensing of nuclear sites. Standard licence conditions apply, including Licence Condition 14 (Safety documentation).

Guiding Principles/Concepts

Management of uncertainties is a fundamental part of safety analysis.

Over-conservatism, as well as under-conservatism, can be the enemy of safety.

The effect of combining individual conservatisms is not just cumulative; it can be multiplicative, resulting in major over-estimates in assessed results. Combining conservative values of multiple input values can easily lead to a ridiculously unrealistic output value, without the knowledge of the analyst.

It is important to achieve a balance between under-conservatism and over-conservatism: *appropriate conservatism*.

Over-conservatism

The symptoms of an overly-conservative safety case include:

- The safety case is larger and more complex than expected, with excessive production costs.
- Many shortfalls are identified, while a similar (or higher hazard) facility elsewhere that has gone through a valid safety case process has very few or no shortfalls.
- The requirement for safety measures (both engineered and operational) is very high in comparison to a similar (or higher hazard) facility elsewhere.
- The safety case owner and other stakeholders find the safety case to be not credible.

Many of these symptoms might have other causes apart from over-conservatism, so caution must be applied.

Reference 3 states: ‘The use of a conservative methodology may be so conservative that important safety issues may be masked.’ If there are elements of the PSA that are conservative while the other elements are best-estimate, the understanding of where the highest risks lie may be skewed. Another effect of an overly-conservative safety case can be that the safety case suffers from a lack of credibility. Even more

seriously, it can result in a solution that does not give an overall risk that is As Low As Reasonably Practicable (ALARP):

- If the accident analysis is overly-conservative, there may be additional requirements for maintenance workers to carry out examination, inspection, maintenance and/or testing, and this may result in maintenance workers receiving significant doses over time. This may apparently be an overall risk benefit, but once the conservatisms are removed, it may become clear that the holistic risk is actually higher because of the safety case requirements.
- Decommissioning may be delayed because of the complexity of implementing the safety case, whereas a more reasonably conservative safety case would have allowed quicker implementation, and earlier removal of the hazard.

Thus over-conservatism, as well as under-conservatism, can be the enemy of safety.

Unfortunately, it is relatively easy to construct a safety case that is overly-conservative. This derives from combining conservative assumptions. For example, in a consequence analysis associated with an airborne release of activity, there are usually many factors (including inventory, release fractions, decontamination factors etc.) that have some uncertainty associated with them. Taking only these three parameters, if each of these uses a conservative value that is ten times worse than the best-estimate, the assessed dose to a worker would be 1,000 times the best-estimate dose, and most safety case professionals would regard the assessed dose in this case as overly-conservative. ***The effect of combining individual conservatisms is not just cumulative; it can be multiplicative, resulting in major over-estimates in assessed results.*** The significance of this multiplicative effect is often not appreciated. The same effect applies in theory to under-conservatisms, but these are rarer, and any under-conservatisms are usually combined with over-conservatisms rather than with other under-conservatisms.

Appendix A, which refers out to Reference 2, concludes that ***combining conservative values of multiple input values can easily lead to a ridiculously unrealistic output value, without the knowledge of the analyst.***

In one case of over-conservatism, the dose rate from fuel on a ramp was over-estimated, and a new gamma interlock system was installed to prevent anyone entering the area and receiving a significant dose. After installation, the system was checked by bringing fuel up the ramp. However, the door to the area still opened allowing operator access. This was not because of system failure, it was because the dose rate from the fuel was too low to trip a gamma interlock or be of any real concern to operators. This type of event can bring the whole safety case process into disrepute - not a good result when a safety case author was simply trying to be conservative!

Achieving a Balance between Under-conservatism and Over-conservatism, including Examples

Before considering conservatisms that might be over-conservatisms, it is worth reviewing the whole fault analysis to see if there are any ‘convenient conservatisms’ (see section on ‘Introduction to subject matter’) that can easily be removed. This review can be done before any of the safety argument is written down, after it’s all been written down, or at any point in-between. Sometimes a convenient conservatism was made at the time in the belief that the safety case would not be adversely affected, but with subsequent developments, the convenient conservatism

is now decidedly inconvenient! What's more, it may be hidden because it was perceived earlier to be an inconsequential conservatism. Thus a trawl through the assessment might be necessary to identify any convenient conservatisms. It's obviously easier if these have been clearly identified at the time of assuming the convenient conservatism.

Paragraph 103 of Reference 1 states that the safety case should provide a proportionate justification that includes appropriate conservatism but without undue pessimism. This statement is supported in this SCF Guide.

Examples of achieving a balance between under-conservatism and over-conservatism are divided into the following sections:

- Achieving a balance in design
- Achieving a balance in fault progression
- Achieving a balance in consequence assessment
- Achieving a balance in engineering substantiation
- Achieving a balance in other areas
- Achieving an overall balance

Note that the advice in 'Achieving an overall balance' is worth bearing in mind right from the start, rather than waiting until the 'wrong answer' has been reached at the end of the safety case production process. It may be appropriate to include the issue of avoiding under-conservatism and over-conservatism in a project start-up meeting. Where there are large uncertainties, it may well be appropriate to identify the major contributors to overall uncertainty early on, and obtain data that reduces this uncertainty before decisions have to be made.

Achieving a Balance in Design

'Design' here can refer to a new facility or to design of a decommissioning project or to design of a modification.

Inventory can be particularly important here. For example, in a facility processing self-heating liquor, it was important to provide enough cooling for the highest heat-loading of the liquor, but it appears that little or no consideration was given to the potential for over-cooling, which was thought to be primarily an operational issue. In a tank, four cooling coils were provided. When it came to operation, over-cooling was encountered, so two of the four cooling coils were 'mothballed' and the other two coils were operated at low flow. An instruction was provided for un-mothballing a coil, but no consideration was apparently given (at the time) to the possibility of a mothballed coil developing a leak such that radioactive liquor could fill the coil, which would then flow out of the shielding at the time of un-mothballing, causing significant dose rates. This fault was later identified in the periodic review of safety, the instruction was withdrawn, and no unusual dose rates occurred, but the potential fault arose from an over-conservatism in the first place.

Uncertainty must be addressed not only by safety professionals, but also by designers. Sometimes using the worst of all variables would lead to a design that is too onerous, and it may be more appropriate to focus the design on expected parameter values, but with provision for 'outliers', that may be more operationally focussed. While operational safety measures are relatively low in the fundamental safety hierarchy, it may still be appropriate to go down this route if for example the 'too onerous design' would delay decommissioning by years. Uncertainty affecting the design basis should be highlighted and conservatisms and/or optimisms should

be made clear, so that the overall hazard management strategy can accommodate it. It may well be necessary to carefully distinguish between unusual non-fault conditions (inherently of a low probability) and accidental conditions (generally of a lower probability, but only because of the provision of safety measures).

For design safety cases, it is important to show that expected doses will be within limits and ALARP. This involves prediction of expected dose rates (based on shielding calculations) in certain plant areas, along with an estimate of occupancy in those areas. There are many uncertainties, but dose estimates can be improved by learning from experience of similar facilities now in operation. The same opportunity rarely presents itself for accident dose estimates, making it all the more difficult to achieve appropriate conservatism - see the section 'Achieving a balance in consequence assessment'

Another uncertainty arises from inventory estimates. In another example, it was considered that a robot would be needed to transfer packages out of a store, which was needed to enhance security. The need for a robot was based on the dose rate from the package with the highest dose rate. However, average doses to workers almost by definition relate directly to mean package dose rates, and often even the dose to the most exposed worker is not unduly influenced by the highest dose rate package. By reviewing the distribution of package inventories and dose rates, it was possible to show that the transfer of packages could be achieved safely by employing manual transfer, thus achieving the security enhancement significantly earlier than if a robot had been employed.

Achieving a Balance in Fault Progression

Development of the fault progression is a key stage in fault analysis, and one in which under-conservatism or over-conservatism can easily arise. This is best illustrated by means of an example:

An existing store contains packages of radioactive material that is self-heating. Cooling is provided by means of a ventilation extract from the store. The fault identification process has identified the possibility of the vent extract failing and the radioactive material over-heating. Specialists are asked to determine if releases from the radioactive material may occur from over-heating in the case of ventilation extract failure. Their answer, in 4 weeks' time, is that they cannot prove that a release cannot occur.

At this point, most safety assessors will feel that it's important to progress with safety analysis assuming that over-heating may occur, especially if there are timescale pressures. This is a conservative approach, but after several weeks of fault analysis, it is found that safety assessment criteria cannot be met unless an additional fan, fed from a different electricity supply and backed up by a separate generator, is provided, and even then the assessed risk would be uncomfortably high. Strict operator instructions are required, and restrictions are to be placed on fan maintenance that are difficult for the safety case owner to live with. And a new detailed emergency instruction is required, that will need to be practised annually.

Only then are the specialists consulted again, and it is found that:

- The specialist analysis focussed only on natural ventilation, and did not consider in detail at what temperature releases would actually occur from the packages.

- The specialist analysis effort was minimal, and less conservative assumptions could be made if more time is devoted to the topic. In particular, heat loss through conduction has not been modelled at all.
- It was assumed that temperature would keep on rising essentially forever (until equilibrium is reached, as this was easier to model), with no restoration of forced cooling.

Once the specialist analysis is done in more detail, removing some conservatisms, it is shown that no release from the packages could reasonably occur.

While this example is at least partly about poor communication (and about a transactional approach that divides professionals who need to talk to each other), it is also about a failure to understand the importance of conservatisms at this vital stage in fault analysis. It is important to ask the right specialist and understand their speciality, for example a different specialist may be required for heat loss from the specialist for a temperature threshold for release from material. Asking for an *accurate* thermal model might in this case have been appropriate, but in addition, getting the right level of resource applied to the technical question posed, is important, as is enabling technical personnel to see the bigger picture - to understand the impact of their answer on the safety case and the operation of the relevant facility. All this is only likely to work if the specialist feels part of an overall team having a common aim. Involvement of the Safety Case Owner may be required to ensure appropriate team working across departmental boundaries.

Amending the scenario somewhat, if the safety assessor finds that releases from the packages have relatively low consequences, and safety criteria are easily met without any modification to the facility (and without overly-restrictive safety designations), it would not be necessary to seek further specialist advice, and it *would* be a reasonable assumption that releases from packages can occur. It would never be shown that releases cannot occur, but this would not matter for the safety case or its users, though risk estimates would be potentially skewed. This assumption (of releases from packages occurring) should be identified in the safety case as a potential conservatism, for clarity.

Involvement of facility operators (both operational workforce and their management) in fault identification and in fault development can help to make the fault analysis more realistic, helping to avoid both under-conservatism and over-conservatism. Caution is required here, because sometimes facility operators say that something can't happen when it can but only extremely infrequently, or it can't happen because of deterministic safety features that need to be designated on the basis of the fault potentially happening. But more often, the facility operator view is crucial in determining a more realistic fault development.

It has been known for a safety case author:

- to be informed of two fires that have happened in a facility in the last ten years;
- to argue that fires won't happen again because steps have been taken to prevent them from happening; and
- to then use a frequency of 1E-4/y for fires in that facility.

This is an example of under-conservative use of information from facility operators, and of failing to learn from experience. It should be obvious that there may well be other initiators of fires that haven't yet happened.

No matter how inconvenient, there should be a healthy respect for operational experience. Use of site-specific and plant-specific reliability data can help to make safety cases more realistic, and avoid both under-conservatism and over-conservatism. When extrapolating data, great care is needed.

Another example of over-conservatism in fault development is to assume that any earthquake bigger than the design basis earthquake is bound to result in catastrophic collapse of the structure. It is likely that there will be 'graceful degradation' rather than catastrophic collapse for an earthquake that is only a little larger than the DBE; if credit is taken for this, it should be substantiated. On the other hand, in general cliff-edge effects beyond design basis need to be considered for the PSA. For example, a tsunami that is only a little higher than a design basis tsunami wall could have catastrophic effects.

If novel technology is being employed, it is particularly important to beware of potential cliff-edge effects or unusual failure modes, as there is a higher risk of under-conservatism in these cases.

When modelling potential events that have not occurred, it is easy to make incorrect and under-conservative (or even optimistic) assumptions about the availability of equipment, or that a control room is still inhabitable (this can be in the form of an implicit assumption, when it is assumed that action can be initiated from the control room).

Most of the thoughts in this section relating to fault progression apply also to the process of fault identification, where it is important not to dismiss credible faults. If they are wrongly dismissed at that stage in the process, it is perfectly possible for the under-conservatism to remain.

Achieving a Balance in Consequence Assessment

Different consequence assessments make different assumptions, but for example for an aerial release off-site there are a large number of elements in the calculation (and associated uncertainties):

- The inventory (how much of which radionuclides).
- The physical and chemical form of the material.
- The Release Fraction in the postulated accident.
- The settling of particles of different sizes.
- The model assumed for dispersion of the 'cloud' of airborne activity.
- Weather category or categories.
- The breathing rate for members of the public, and their distance from the release point.
- Dry deposition rates and wet deposition rates.
- Models for transfer of radionuclides in the environment to the food chain.
- Consumption rates for different types of food.
- The activity median aerodynamic diameter (AMAD) of the airborne material when it is breathed in.
- The model assumed for converting activity inhaled / ingested into dose.

Many of these assumptions include a degree of conservatism, most of which cannot be easily quantified. Some assumptions can be quantified, for example, some models for food consumption assume 97.5th centile consumption rates. Release Fractions and Decontamination Factors are usually set on the conservative side, but it is not generally known by how much.

Principle FA.7 in the SAPs (Reference 1) expects conservative consequences for DBA, not worst-case consequences. Within the DBA section of Reference 1, the only inclusion of the word 'worst' relates to 'the worst normally permitted configuration of equipment outages for maintenance, test or repair' (paragraph 631(c)).

One licensee, that routinely produces two sets of consequences (worst case and best estimate) reviewed the assumptions in the calculations behind these two sets.

For internal dose assessments, there was a factor between these sets of around 200. When the assumptions behind the best estimate were reviewed, it was found that the best estimate itself contained many conservatisms, and gave a value that was considered after review to be appropriately conservative. This implies that the worst case consequences were overly conservative. In general, it may be appropriate to assess true best estimate consequences so that a comparison can be carried out.

It is worth checking the dose resulting from calculations, to see if it matches expectations. If a >2 Sv dose is calculated for release of a material generally considered to be low hazard, it may be that there is an error in the calculation waiting to be discovered. In general, author self-checking of documents for under-conservatism or over-conservatism is good practice.

If an assessed dose is just above a consequence threshold, e.g. in the range 20-40 mSv, it is usually worth looking intensively for any unnecessarily excessive conservatisms that can be removed while still leaving sufficient conservatism to address any uncertainty, as reducing the dose below the threshold may result in safety requirements that are more appropriate to the actual risk, as well as making the rest of the safety assessment significantly easier. If a calculated value is just below a threshold (19 mSv, say), and if the calculations have used considerable and demonstrable conservatism to produce a sub-threshold value, then that value should be used, rather than it being artificially 'bumped-up' to above the threshold, simply because it is close to it – 19 is always less than 20, and the 'real' value is more likely to be closer to 1.9 or even 0.19.

An example of over-conservatism in consequence analysis is not to take credit for deterministic safety features, e.g. vessels, cell walls, etc. As long as they would not be affected by the scenario (e.g. in an explosion), credit should be taken for these, unless there are major difficulties with their substantiation.

Reference 4 gives guidance on conservative exposure durations for unmitigated worker doses in Design Basis Analysis.

In terms of the ***uncertainty associated with inventory***, this can occasionally be described well in statistical terms - this can be described as 'known uncertainty'. For example, the distribution of alpha activity in 10,000 consignments of waste can be described as follows:

- Average (mean) activity: 0.3 GBq
- 90^{th} centile: 2.0 GBq
- 99^{th} centile: 3.4 GBq
- Highest activity to date: 29 GBq

In this case, the distribution has a very long tail at the high activity end, with the highest activity to date being a hundred times the average activity. This is not unknown in radioactive waste inventories.

In many other cases ('unknown uncertainty'), the statistical uncertainty associated with a parameter is simply not known, for example when there is just one sample taken of a material (and it's not reasonably practicable to take more samples). Although the distribution of sample results from a single material should not be as variable as in the above example, the concepts of 'average' or '90^{th} centile' just don't apply in this type of uncertainty. To avoid under-conservatism, it would normally be appropriate to increase the apparent activity concentration somewhat from the analysis values, depending on an understanding of potential variability between samples, if more than one were to be taken.

More commonly, the uncertainty associated with inventory is somewhere between the two examples above.

Where there is good data on inventory, assuming that there are no cliff-edge effects (see the section 'Introduction to subject matter'), the hard question still remains as to whether to use the mean, or maximum of results so far, or something else:

- There is an argument for using the mean, if it is known that other elements of the relevant calculations contain enough conservatism to give overall conservative results. There is the possibility that the average may move over time, potentially to a higher value, leaving the inventory element of the calculations potentially optimistic. Having any element of optimism in the calculation can be difficult to defend, but is defensible as long as overall conservatism is maintained.
- 'The maximum of results so far' is more conservative, but if a higher result becomes available (e.g. due to a separate sampling campaign), there may be a demand to amend calculations. In this sense, this is not a robust value to choose if it is being argued that the maximum value is key.
- It is possible to argue for the 90th centile value or the 99th centile value - both have been used historically in different companies. Comparing the inventory with the average value (e.g. 'the 90th centile value of 2.0 GBq is more than six times the average value of 0.3 GBq') can be persuasive in achieving acceptance that it is reasonably conservative to use the 90th centile value, while acknowledging that higher values are expected one in ten times. One advantage is that even if the distribution shifts significantly over time to a higher range, the 90th centile value chosen is unlikely to be lower than the new mean, robustly maintaining a degree of conservatism in the inventory element of the calculations.

The context is very important here. If an average normal dose is required, the mean may well be best. In safety assessment of potential accidents, it is more difficult to justify use of the mean.

One criticism that can be levelled at using any value lower than 'the maximum of results so far' is that it is possible that the value used gives a dose (e.g. 19 mSv to a worker) below a DBA threshold while use of the maximum of results so far would give a dose (e.g. 160 mSv, using the example of waste consignments above) above a DBA threshold. This could result in a failure to apply DBA appropriately (or at all). One way to resolve this is to define a small subset (e.g. the top 1%) of the inventory that is treated as a different bounding case. Where it is known which containers have the top 1% of inventory, it may be possible to treat them differently, e.g. with an additional safety measure that wouldn't be practicable to apply to all containers. Even if this isn't possible, or if the inventory isn't known at the relevant process stage, the initiating event frequency can take account of the conditional probability (in this case 1%) that the container is a high inventory container. This fault sequence is quite likely to 'land' outside the DBA region, due to the lower initiating event frequency. This approach is also suitable for PSA, where the risk from the top 1% is likely to be smaller than the risk from the bottom 99%. Note also the possibility of cliff-edge effects (see the section 'Introduction to subject matter').

Reference 2 states that if uncertainties cannot be fully assessed, using mean values for nearly all input parameters, in combination with 'bounding' values for one or two variables ensures an output magnitude that is comparable to the actual 95th percentile. Reference 2 states that this approach offers the best alternative for incorporation of parameter uncertainties while providing reasonably conservative results. However, in some cases (e.g. some external dose scenarios) there are only a few parameters that need to be combined, so the Reference 2 view is modified in this SCF Guide, which recommends that ***appropriate conservatism in consequences can be achieved by using conservative values (e.g. equivalent to approximately 95th percentile) typically for two parameters, with mean values being used for all other parameters.***

The key point is that there is overall conservatism in the calculations, not that every element in the calculations (including inventory) has to be conservative by a certain degree or at a certain centile.

Inventory information can at times be confusing and apparently contradictory, in which case the above mathematical points cannot be applied. For example, in a store of legacy waste drums with very little provenance information, it may be thought that the drums contain uranium, and this is supported by sample analysis of a leak from one drum showing that it contains uranium, with no traces of plutonium. The drums, which may have come from one or more other sites, cannot easily be sampled as that would involve movement across site to another building, with transport and other risks. However, all the drums have been assayed in a drum assay monitor that shows significant positive plutonium values, including for the drum that subsequently leaked. Each gram of plutonium is roughly equivalent to *one million* g of uranium in internal dose terms, mainly due to the much higher specific activity of plutonium. It is known that the positive plutonium values may arise from difficulties that the assay monitor has in screening out U238, which is likely to be prevalent in the waste. For the drum that has leaked, it may well be reasonable to justify an assumption of uranium in the safety assessment, but for the rest of the drums, the conservative approach of assuming that the positive plutonium values are realistic is likely to be the only justifiable approach, even if this makes the safety case more difficult and complex. This conservatism should be identified in the text.

Achieving a Balance in Engineering Substantiation

Before engineering substantiation became a fundamental part of safety cases, it was generally assumed that a Structure, System or Component (SSC) would be capable of meeting its safety function. This was an under-conservatism (failing to address uncertainties) that has now been resolved, and now the main risk appears to be over-conservatism. However, it is still a risk that not all relevant SSCs are identified, which would be an under-conservatism, e.g. not identifying SSCs having normal safety functions.

One way of obtaining assurance that a safety function is provided with sufficient confidence is to compare the SSC with relevant modern standards. If these standards are achieved, the safety function is usually substantiated. If they are not achieved, it may be easier to declare that the safety function is not substantiated, which is an incorrect conservative approach. But knowing that a wall does not meet modern standards does not mean that it is incapable of providing the structural stability or the shielding function that is required by the safety case. It is as easy to be overly conservative in engineering substantiation as it is in safety assessment.

In one case, the safety function of structural stability of a roof was not substantiated in the event of over-depression in the old building (which was served by a ventilation system drawing air from several buildings). Yet, no restriction was placed on the number of people who could work on the roof (which was an operating roof), indicating no real concern with structural stability. When combining that with known leak paths in the building fabric (so the depression would not be large), not confirming structural stability of the roof in these circumstances appears to have been over-conservative. Appropriate challenge, along with team-working, can resolve issues like this. Time spent on the challenge can be far more productive than working with the over-conservative assumption, but not always.

Sometimes a substantiating engineer may have an opinion that the safety function is highly likely to be provided, but will not *write* that judgement in a document. This can result from a misunderstanding of the confidence level required, or, particularly if the work is being done by an external contractor, concern over liability. Moving to a

position where it is regarded as acceptable to write expert judgement into documents (along with a description of how the expert judgement is reached) is important in avoiding over-conservatism. The following advice may be helpful in approaching the formalisation of engineering judgement:

- Use a team approach to clarify the logic, from safety assessment through safety function wording to substantiation.
- Identify issues where your experience and expertise tells you that there is both a solution and an efficient route to that solution.
- Formulate your rationale (including logical argument based on analysis of the design and on ageing effects).
- Discuss your rationale with your peers both within the project and within your engineering discipline.
- Agree with your peers the approach, the level of underpinning, and the level of peer review.
- If appropriate discuss with regulators.
- Write up the solution using the agreed approach.
- Obtain formal agreement through signatures on your document.

It is of course possible that engineering judgement could show that a safety function is not substantiated.

Substantiating engineers must have enough independence from the safety case owner to be able to identify any reasonable shortfalls against safety functions. Many important safety improvements have been made because engineers have identified insufficient confidence in SSCs meeting safety functions. In addition, many unnecessary safety improvements have been identified because engineers have misunderstood safety functions or have been given inadequately defined safety functions. Thus accurate definition of safety functions (for example by using appropriate performance requirements) is key to avoiding under-conservatism and over-conservatism.

Achieving a Balance in Other Areas

The scope of a safety case sometimes conservatively includes an operation that is no longer carried out. The analysis of this operation may be long and complex, and difficult due to lack of current plant experience. It may be more appropriate to 'ban' the operation rather than to analyse it, e.g. by excluding from scope and requiring a modification proposal/case to allow the operation to take place if desired in future. On the other hand, it is possible that an infrequent operation is excluded from a safety case that really should be included. Close working with representatives of operations is required.

Bounding cases can be significant in avoiding complexity, and thereby avoiding over-designation of safety measures. However, bounding too many scenarios within one bounding case can result in safety measures applying unnecessarily to too many operations. Again, a balance is required.

It is often thought that DBA requires a conservative value for initiating fault frequency, but paragraph 629 of Reference 1 does not require this, at least not for internal faults: 'Initiating fault frequencies should be determined on a best-estimate basis with the exception of natural hazards where a conservative approach should be adopted.' The exception for natural hazards is mainly to reflect uncertainties in the underlying data used when defining the most extreme events.

A frequent example of conservatism in probabilistic safety assessment (where conservatism is not required) is to use the equation for assessing worker risk like this:

Risk of death to an individual worker per year = the sum across all accidents of:

accident dose to most exposed worker (Sv)

x 0.05 deaths/Sv

x frequency of the relevant accident in the facility per year.

The above equation conservatively assumes that:

- every worker is present in the facility at the time of every accident; and
- every worker present in the facility is exposed to the same extent as the most exposed worker.

The first bullet is clearly conservative for shift workers (typically 5 workers are required to provide 1 worker present on the facility at all times), and is still conservative for day (non-shift) workers, but less so. The second bullet is also a conservatism, especially in the case of airborne releases for facilities that are well compartmentalised and there are many workers in the facility (mainly in other compartments). Given that the worker risk criterion is a probabilistic criterion, so best-estimate is appropriate, significant worker risk reduction factors (e.g. 10) may be appropriate in some instances, but individual justification is required.

Achieving an Overall Balance

As well as attempting to achieve a balance in all the above areas, an holistic approach may also be required. A key question when considering the overall analysis is ‘Does it look right?’ - the answer to this question can usefully prompt a review. It is usually advantageous to bring together all the people involved in a safety case, especially when the overall conclusion of the safety case doesn’t feel right. It may be that a tendency to under-conservatism or over-conservatism by several different groups would be identified, allowing an overall resolution to be reached.

In order to achieve an overall balance, it is important that potential over-conservatisms are recognised within an overall safety argument, especially where the outcome appears to fail to give risk that is ALARP. In order to support this, it is helpful if the individual contributions to the overall argument identify in their own documentation (e.g. a hazard analysis) where potential under-conservatisms and over-conservatisms exist.

Carrying out a sensitivity analysis, e.g. by carrying out a best-estimate analysis as well as a conservative analysis, may be useful in demonstrating the depth of conservatism.

When considering what is ALARP, over-conservatism in one aspect, e.g. the risk from accidents, can dominate thinking such that an overall ALARP balance is not achieved. The tendency in ALARP arguments just to emphasise conservatisms (sometimes to the extent that the concepts of ‘ALARP’ and ‘conservatism’ can become confused) does not achieve an acceptable ALARP argument (which needs to consider options) - it would be better to remove over-conservatisms within the main safety analysis, thus providing an appropriate basis for the overall ALARP argument.

In some cases, the same conservatism can be quoted as a response to a series of emergent issues or unexpected observations. Making an effort to see the big picture

may reveal that all of the original conservatism has been used up, such that the overall argument is now under-conservative or even optimistic.

Summary of Key Points

Management of uncertainties is a fundamental part of safety analysis. Safety cases should identify areas of optimism and uncertainty, together with their significance, in addition to strengths and any claimed conservatisms (and potential conservatisms).

Conservatism is an important part of safety cases. ***Overall conservatism*** is what is required for Design Basis Analysis; it is not necessary for every element in the analysis to be demonstrably conservative.

Over-conservatism, as well as under-conservatism, can be the enemy of safety.

Areas where under-conservatism or over-conservatism are particularly important include:

- fault progression;
- consequence assessment; and
- engineering substantiation.

In addition, the overall balance is very important - see the section on 'Achieving an overall balance'. It is important to understand how the conservatisms might skew ALARP decision-making.

The effect of combining individual conservatisms is not just cumulative; it can be multiplicative, resulting in major over-estimates in assessed results. Combining conservative values of multiple input values (e.g. inventory) can easily lead to a ridiculously unrealistic output value (e.g. dose), *without the knowledge of the analyst*. In consequence or risk assessment (where several parameters are multiplied together), appropriate conservatism in consequences can be achieved by using conservative values (e.g. equivalent to approximately 95th percentile) typically for two parameters, with mean values being used for all other parameters.

Techniques for removing or managing under/over-conservatism that are significant throughout this SCF Guide include:

- explicit identification of optimisms and conservatisms within safety case documents;
- author self-checking of documents for under-conservatisms or over-conservatisms;
- good communication, and avoiding 'silo working', in particular it is important to seek the views of facility operators during fault identification and fault development;
- awareness of the multiplicative effect of over-conservatisms;
- use of sensitivity analysis;
- appropriate use of expert judgement; and
- appropriate challenge of assumptions and conclusions (asking 'Does it look right?'), within a team-working environment.

References

1 HSE, Safety Assessment Principles for Nuclear Facilities, 2014 Edition, Revision 0.

2 National Nuclear Security Administration (USA), Technical Bulletin 2014-1, Achieving Reasonable Conservatism in Nuclear Safety Analyses, Kamiar Jamali.

3 IAEA SSG-2, Deterministic Safety Analysis for Nuclear Power Plants, 2009.

4 UK Nuclear Industry Safety Case Forum Guide, Conservative Exposure Durations for Unmitigated Worker Doses in Design Basis Analysis, Alec Bounds, January 2014.

Appendix A – Statistical Uncertainties

An individual parameter will typically have a statistical distribution of values. This distribution is often significantly skewed, as in the diagram below, which shows two separate log-normal distributions that both have the same median value, but one distribution has greater uncertainty than the other. Parameter distributions of both types are encountered in the nuclear industry. Thus the high end of the distribution can be significant, such that the concept of a maximum value becomes almost meaningless. Of the 3 types of average, it is clear that the value most likely to be encountered (the mode) could be significantly optimistic compared to the mean. Of the three averages, the mean can be regarded by safety case practitioners as a more useful type of average for skew distributions, as its value increases towards higher percentiles of the underlying distribution with increasing levels of uncertainty (see diagram below). The mean takes account of all values, including values that may be considered 'abnormal'.

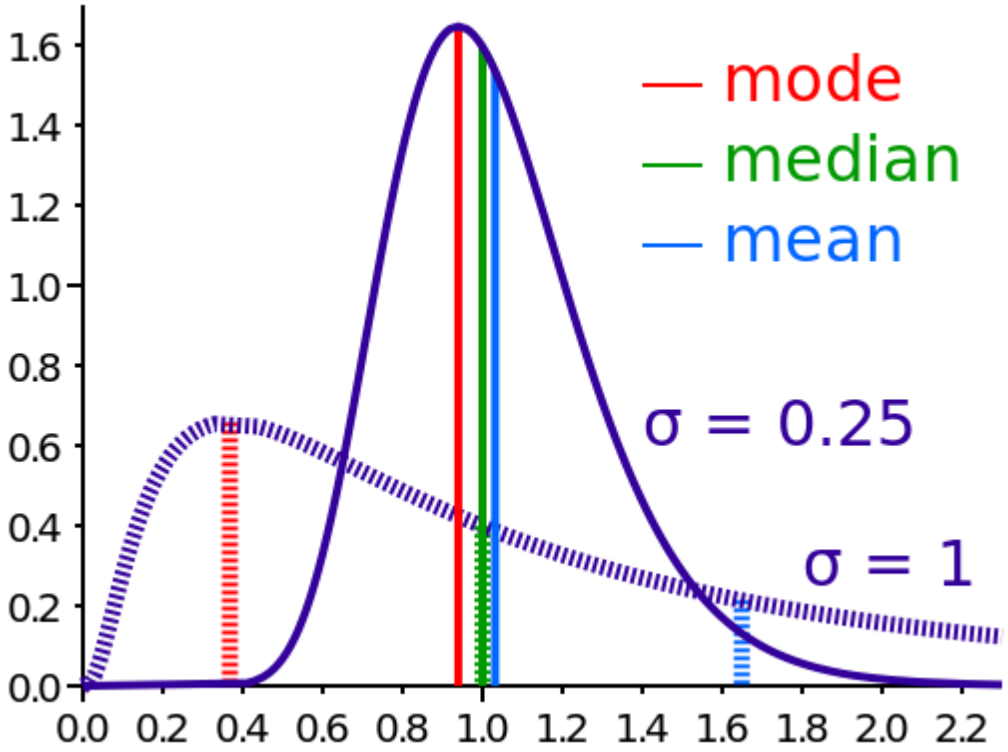

A real example of points in a skewed distribution is given in the section 'Achieving a balance in consequence assessment'.

The situation in safety cases is more complex than in the above diagram because parameters (and thus distributions of values of those parameters) are multiplied together. If the individual distributions are normally distributed, the product of the distributions will be log-normally distributed. Furthermore, if the individual parameters are log-normally distributed, the product of those distributions will also be log-normally distributed, but with a higher level of uncertainty (giving a longer tail) - see Reference 2. Thus log-normal distributions are likely to occur in dose assessments.

Reference 2 analyses over-conservatism from a mathematical perspective. It states that conservatism in nuclear safety analyses can lead to extreme conservatism

without the knowledge of the analyst. The divergence to extreme conservatism occurs rapidly. This phenomenon is often beyond the expert analysts' intuitive feeling, but it can be demonstrated mathematically. The upper-bound is generally associated with the overall 95th percentile of the output when parameter uncertainties are fully propagated in a mathematical model. It is shown that the product of the upper-bounds can be orders of magnitude above the actual upper bound of the output. In other words, ***combining conservative values of multiple input values can easily lead to a ridiculously unrealistic output value, without the knowledge of the analyst.***