CHAPTER 1

Selecting Safety Integrity Levels: Introduction

The purpose of a *safety instrumented system* (SIS) is to reduce the risk that a process may become hazardous to a tolerable level. The SIS does this by decreasing the frequency of unwanted accidents. The amount of risk reduction that an SIS can provide is represented by its *safety integrity level* (SIL), which is defined as a range of probability of failure on demand. An SIS senses hazardous conditions and then takes action to move the process to a safe state, preventing an unwanted accident from occurring.

The method organizations use to select SILs should be based on their risk of accident, an evaluation of the potential consequences and likelihoods of an accident, and an evaluation of the effectiveness of all relevant process safeguards. Implementing an SIS, and therefore selecting an SIL, should involve considering relevant laws, regulations, and national and international standards. In the United States, the "Process Safety Management" (PSM) section of the OSHA standard OSHA 29 CFR Part 1910.119 requires organizations to provide assurance of the mechanical integrity of all their emergency shutdown systems and safety critical controls. The "Seveso Directive" (96/82/EC) promulgates similar requirements in the European Union. In the United States, ISA—The Instrumentation, Systems, and Automation Society promulgated industry standard ANSI/ISA-84.01-1996 to promote compliance with the PSM regulation. The International Electrotechnical Commission (IEC) created a similar document, IEC 61508, which is an umbrella standard that covers numerous industries. IEC standard 61511 is the process-sector specific standard that falls under the IEC 61508 umbrella. This standard, when ratified, will be reviewed by ISA SP84 and accepted as a replacement for ANSI/ISA-84.01, possibly with some modification. The IEC standard

84.01-1996 and IEC 61508 and 61511, require that an SIL be selected. These standards are the basis of organizations' efforts to comply with the local and national laws and regulations that govern processes that contain significant risks. Many "authorities having jurisdiction," who are responsible for enforcing these laws and regulations, tend to view complying with such international standards as equivalent to complying with "good and generally recognized engineering practices" clauses.

1.1 Safety Integrity Level

Safety integrity levels (SILs) are categories based on the *probability of failure on demand* (PFD) for a particular *safety instrumented function* (SIF). The categories of PFD range from one to three, as defined by ANSI/ISA-84.01-1996, or one to four as defined by IEC 61508 and 61511. Table 1.1 shows the PFD ranges and associated risk reduction factor (RRF) ranges that correspond to each SIL.

Table 1.1 Safety Integrity Levels and Corresponding PFD and RRF

SIL	PFD Range	RRF Range
4	$10^{-4} \rightarrow 10^{-5}$	10,000 → 100,000
3	$10^{-3} \rightarrow 10^{-4}$	1,000 → 10,000
2	$10^{-2} \rightarrow 10^{-3}$	100 → 1,000
1	$10^{-1} \rightarrow 10^{-2}$	10 → 100

The SIL is the key design parameter specifying the amount of risk reduction that the safety equipment is required to achieve for a particular function in question. If an SIL is not selected, the equipment cannot be properly designed because only the action is specified, not the integrity. To properly design a piece of equipment, two types of specifications are required: a specification of what the equipment does and a specification of how well the equipment performs that function. The safety integrity level addresses this second specification by indicating the minimum probability that the equipment will successfully do what it is designed to do when it is called upon to do it.

choosing the same size as the piping and selecting equal percentage trim), your selection would not be optimal. You would have no guarantee that the valve would be able to pass the proper flow rate, and you would almost be guaranteed to have selected a valve that is oversized, and thus overpriced. You could improve performance and lower capital expenditures by investing the effort required to select a piece of equipment that not only performs the proper function, but also has the required performance characteristics.

Selecting safety integrity level involves giving a numerical target upon which subsequent steps in the safety life cycle are based. Thus SIL selection offers an important guide when you are selecting equipment and making maintenance decisions. The SIL is documented along with the SIS operational requirements and logic as part of the safety requirements specification. This specification provides the foundation for all of the safety life cycle activities an organization later conducts.

IMPORTANT: The process we are referring to as SIL selection in this book has been described by many other terms, including *SIL determination* and *SIL classification*. We specifically chose *SIL selection* because it describes the overall process most clearly. *Determination* is a vague term allowing too many variations in connotation. *SIL classification* implies that the process does not involve making a decision and that every situation is the same if you know its category. *Selection* is the clearest and most descriptive term because it emphasizes the act of choosing the correct value based on clear criteria.

1.2 Safety Instrumented Functions

In this book, we will adopt the terminology of IEC 61511, wherein a safety instrumented function (SIF) is an action a safety instrumented system takes to bring the process or the equipment under control to a safe state. This function is a single set of actions that protects against a single specific hazard. A safety instrumented system (SIS), on the other hand, is a collection of sensors, logic solvers, and actuators that executes one or more safety instrumented functions that are implemented for a common purpose, such as a group of functions protecting the same process or implemented on the same project. Note that the term *SIF* often refers to the equipment that carries out the single set of actions in response to the single hazard, as well as to the particular set of actions itself. Here are some examples:

- SIF 1: High reactor temperature closes the two reactor feed valves.
- SIF 2: High column pressure or high column temperature closes a valve in the steam to the reboiler.
- SIF 3: High column pressure closes the two reactor feed valves.

The logic for all safety functions is performed in a safety PLC. This PLC would then combine with all of the equipment associated with each SIF to constitute the SIS.

Figure 1.1 Safety Instrumented Functions versus Safety Instrumented Systems

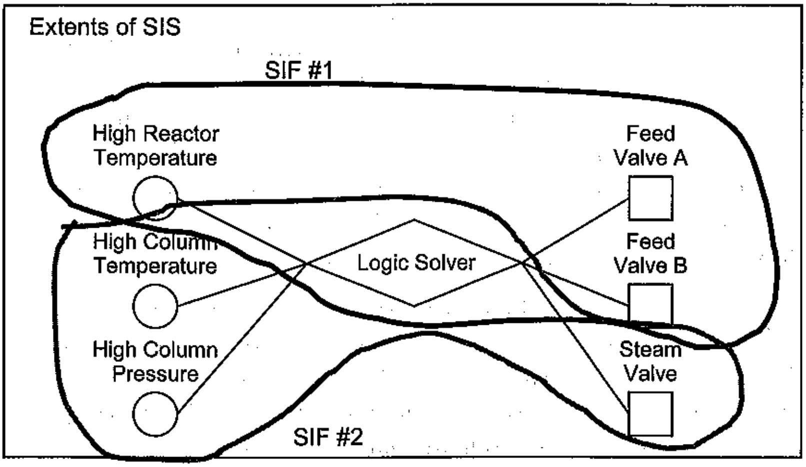

You may implement one or more SIFs in a SIS, as shown in figure 1.1. ANSI/ISA-84.01-1996 uses the terms *SIF* and *SIS* in a somewhat interchangeable and confusing way. IEC 61511 makes the distinction between SIF and SIS very clear. As figure 1.1 shows, a safety function can include multiple inputs and outputs. SIF 1 is executed with two outputs, that is, the two reactor feed valves, and SIF 2 has two inputs, that is, the high pressure and high temperature measurements. It is also important to note that a multiple SIF system can include common equipment. For instance, in figure 1.1, both SIFs use the same logic solver. In instances where common equipment is used in multiple SIFs, the common equipment item should be designed to meet the SIL of the SIF that has the highest requirements.

⚠ **IMPORTANT:** The SIL belongs to the specific safety instrumented function (SIF), not to the entire safety instrumented system (SIS). When an equipment item is common to multiple SIFs, it should be designed to meet the highest SIL requirements of

1.3 SIL Selection and Risk

The reason an organization should use a systematic methodology, which includes layer of protection analysis, to select safety integrity level is to make the choice that best reduces risk. A good decision during this phase of the safety life cycle will ensure that the safety system specified will be cost-effective while still providing appropriate loss prevention. To make the best decision about safety integrity level, an SIS designer needs to completely understand not only the potential likelihood of an unwanted event, but also the possible consequences of that event. Viewing either of these two facets of the risk equation in isolation will yield poor results. Once the risk is known, one must determine how to reduce that risk to a tolerable level. The amount of risk that an organization is willing to tolerate will determine the amount of risk reduction it needs. Many proposed risk reduction projects are financially impractical because the amount of risk reduction they provide is grossly disproportionate to their cost. SIS designers must weigh the amount of risk reduction an SIF achieves against the equipment's cost. Good designs will optimize the return on investment.

Since safety integrity level is defined by the amount of risk reduction an SIS provides, it is important to understand what is meant by risk. There are many different types of risk, and it means different things to different people, but risk has a particular meaning in the context of SIL selection. Here we define risk as a measure of the likelihood and consequences of adverse effects when a process goes out of control and its hazards are realized. Risk is the product of both *likelihood* and *consequence*. The total risk can only be known when both the likelihood and consequences are known. Knowledge of either in isolation is simply not enough to properly solve risk reduction problems.

The risk reduction process is illustrated in figure 1.2. Before selecting an SIL, you must evaluate the inherent risk of the process. The starting point for SIL selection, or the baseline risk, is the level of risk that exists after considering all non-SIS mitigation measures (e.g., relief valves, dikes). Once the baseline risk is determined, you can employ an SIS to further reduce the risk by decreasing the likelihood of an incident. Each additional SIL step, by definition, reduces the likelihood of harm by an order of magnitude. For example, if the baseline likelihood were 10^{-2} per year, an SIL 3 system would reduce the likelihood by three orders of magnitude to 10^{-5} per year. The appropriate SIL is the one that reduces the risk to a tolerable level, or *as low as reasonably practicable* (ALARP). The ALARP concept is explained in more detail in section 3.1. In figure 1.2, SIL 1 would be appropriate if the cost of SIL 2 could not be reasonably justified. SIL 2 is acceptable without further analysis so the cost effectiveness of SIL 3 need not be investigated. This is because the gen-

Figure 1.2 Effect of Safety Instrumented Systems on Process Risk
(Courtesy Process Safety Progress, American Institute of Chemical Engineers, © 1999 AIChE)

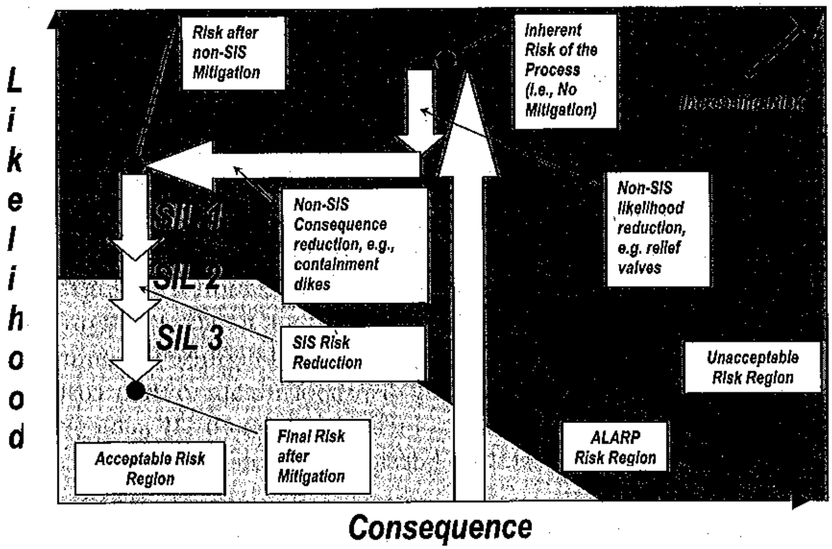
Likelihood
Risk after non-SIS Mitigation
Inherent Risk of the Process (i.e., No Mitigation)
Increasing Risk
SIL 1
SIL 2
SIL 3
Non-SIS Consequence reduction, e.g., containment dikes
Non-SIS likelihood reduction, e.g. relief valves
SIS Risk Reduction
Unacceptable Risk Region
Final Risk after Mitigation
ALARP Risk Region
Acceptable Risk Region
Consequence

eral acceptability of SIL 2 directly implies that the additional SIL level could not be cost-justified.

1.3.1 Consequence

If a process goes out of control, the potential energy contained in the process, in the form of temperature, pressure, and chemical reactivity, may be released and cause harm. This harm is the result of the impact of the event on various receptors. The consequence of a process incident could result in impacts to one or more receptors, such as:

- People
- Property
- Environment
- Business (production and profits)

Estimating the magnitude of the consequence is the first step in the process of selecting an SIL. Many consequences are small enough to be acceptable regardless of how often they occur. If this is the case, the anal-

Figure 1.2 Effect of Safety Instrumented Systems on Process Risk
(Courtesy Process Safety Progress, American Institute of Chemical Engineers, © 1999 AIChE)

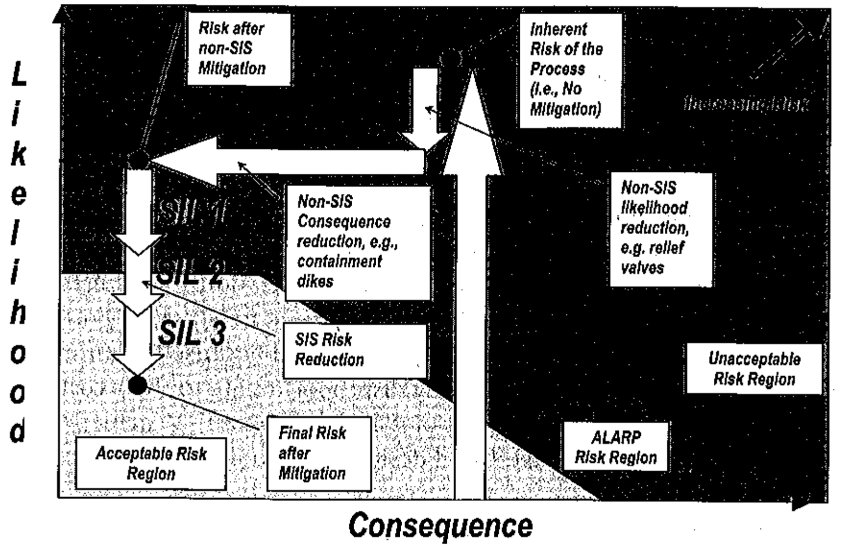
Likelihood
Risk after non-SIS Mitigation
Inherent Risk of the Process (i.e., No Mitigation)
Increasing Risk
SIL 1
SIL 2
SIL 3
Non-SIS Consequence reduction, e.g., containment dikes
Non-SIS likelihood reduction, e.g. relief valves
SIS Risk Reduction
Unacceptable Risk Region
ALARP Risk Region
Acceptable Risk Region
Final Risk after Mitigation
Consequence

eral acceptability of SIL 2 directly implies that the additional SIL level could not be cost-justified.

1.3.1 Consequence

If a process goes out of control, the potential energy contained in the process, in the form of temperature, pressure, and chemical reactivity, may be released and cause harm. This harm is the result of the impact of the event on various receptors. The consequence of a process incident could result in impacts to one or more receptors, such as:

- People
- Property
- Environment
- Business (production and profits)

Estimating the magnitude of the consequence is the first step in the process of selecting an SIL. Many consequences are small enough to be acceptable regardless of how often they occur. If this is the case, the analysis can stop at this stage. Consequence analysis is typically performed by risk analysts in the loss prevention department of an organization or by consultants, especially if the consequence is due to the release of a flammable or toxic chemical. Understanding the consequences of these releases requires a thorough understanding of analytical chemistry, physics, and meteorology. Such knowledge is often beyond the scope of control system and process engineers. We consider consequence analysis in more depth in chapter 6.

The practice in industry has been to limit the scope of risk analysis to such consequences as the potential for injury to workers and the public as well as environmental impact. This scope is required by the mechanical integrity clauses of the U.S. OSHA and EPA process safety management (PSM) and risk management program (RMP). However, the potential financial losses resulting from property damage and business interruption are often just as significant, especially in the case of fires and explosions. If an organization addresses only risks to safety and environment, it loses from its analysis the financial benefits it could gain from reducing risks to property and business. Omitting these two impacts may cause a company to select a lower (less demanding) SIL, and thus fail to achieve significant, long-term benefits.

1.3.2 Likelihood

The likelihood of harm from a process hazard is the combination of the frequency and probability of all contributing events. It is unlikely that a single point failure will result in a hazardous consequence, since most process plants are designed with multiple layers of protection. For example, the analysis of the series of events leading to harm could involve the following:

- Frequency of process failure
- Probability of safeguards failing, resulting in a hazardous release
- Probability that the release will result in a harmful impact

To determine the overall likelihood of harm, you must divide the sequence into basic events that are narrow enough in scope that historical experience will yield a reasonably good estimate of their frequencies and probabilities. The level of detail for a likelihood analysis can range from a simple qualitative estimate (based on engineering judgment) to a quantitative analysis that uses sophisticated fault propagation modeling techniques in combination with historical data. These procedures are discussed further in chapters 7 and 8.

1.3.3 Tolerable Risk and SIL Assignment

The United States does not require companies to use formal risk decision criteria. However, company-based criteria can help ensure that an SIS achieves the desired level of risk reduction. Many companies have adopted internal risk guidelines to ensure an appropriate level of protection for workers and the surrounding community. Often, government and corporate risk guidelines only set limits for safety (i.e., potential fatalities and injuries) and environmental impacts, while ignoring the financial risks to property and production. We consider these government and corporate risk guidelines in more detail in chapter 3.

The final step in the SIL selection process is to compare the risk that is estimated to exist in the process with the amount of risk an organization is willing to tolerate. The difference between the existing risk and the tolerable risk determines the amount of risk that must be reduced. As mentioned earlier, SIL is a measure of required risk reduction. Once the required risk reduction is known, it can be converted into a probability of failure on demand (PFD). This is the PFD required of the safety instrumented function, and thus its required safety integrity level.

1.4 Qualitative versus Quantitative SIL Selection

The standards governing SIS design require that SILs be selected, but they do not place any requirements on the methods that should be used for that selection. As a result, many SISs are either overdesigned or underdesigned because poor risk analysis techniques used during the SIL selection process resulted in incorrect SIS design specifications. Qualitative techniques are often used to determine the risk upon which the selection of an SIL is based. A company's overall risk analysis is much more accurate when it uses layer of protection analysis (LOPA) to improve the accuracy of its likelihood estimate and when it uses quantitative consequence modeling. When the existing layers of protection are accounted for more accurately, the integrity requirements of the SIS can be relaxed. As a result, the SIS has both lower initial installation costs and lower overall life cycle costs. This accurate accounting also provides the organization with a clearer understanding of what contributes to the risk reduction and how. This knowledge can then lead to more informed decisions regarding what SISs and what non-SIS risk reduction methods are most effective and appropriate for the situation at hand.

There are two general ways to perform risk analysis: qualitatively and quantitatively. *Qualitative methods* use engineering judgment and personal experience as the basis for making decisions. Often, qualitative methods use rules and checklists to determine how processes should be designed. These rules and checklists are developed by a group of engineers over many years, and are based on a sort of "institutional memory" of the accidents and near misses that have occurred at a plant. Although qualitative methods are based on experience, there are no precise numerical records to verify the professional judgment of the qualitative analyst.

Quantitative methods strictly use historical data and mathematical rela-

The final step in the SIL selection process is to compare the risk that is estimated to exist in the process with the amount of risk an organization is willing to tolerate. The difference between the existing risk and the tolerable risk determines the amount of risk that must be reduced. As mentioned earlier, SIL is a measure of required risk reduction. Once the required risk reduction is known, it can be converted into a probability of failure on demand (PFD). This is the PFD required of the safety instrumented function, and thus its required safety integrity level.

1.4 Qualitative versus Quantitative SIL Selection

The standards governing SIS design require that SILs be selected, but they do not place any requirements on the methods that should be used for that selection. As a result, many SISs are either overdesigned or underdesigned because poor risk analysis techniques used during the SIL selection process resulted in incorrect SIS design specifications. Qualitative techniques are often used to determine the risk upon which the selection of an SIL is based. A company's overall risk analysis is much more accurate when it uses layer of protection analysis (LOPA) to improve the accuracy of its likelihood estimate and when it uses quantitative consequence modeling. When the existing layers of protection are accounted for more accurately, the integrity requirements of the SIS can be relaxed. As a result, the SIS has both lower initial installation costs and lower overall life cycle costs. This accurate accounting also provides the organization with a clearer understanding of what contributes to the risk reduction and how. This knowledge can then lead to more informed decisions regarding what SISs and what non-SIS risk reduction methods are most effective and appropriate for the situation at hand.

There are two general ways to perform risk analysis: qualitatively and quantitatively. *Qualitative methods* use engineering judgment and personal experience as the basis for making decisions. Often, qualitative methods use rules and checklists to determine how processes should be designed. These rules and checklists are developed by a group of engineers over many years, and are based on a sort of "institutional memory" of the accidents and near misses that have occurred at a plant. Although qualitative methods are based on experience, there are no precise numerical records to verify the professional judgment of the qualitative analyst.

Quantitative methods strictly use historical data and mathematical relationships to estimate risk. These tools and techniques are based on scientific principles, which makes them easier to compare alternative designs because they provide objective, measurable criteria.

Both methods have their strengths and limitations. Qualitative methods are effective in situations where the process has a long history and its risk reduction techniques are well established. Quantitative methods are more objective and allow the proactive application of risk reduction to novel situations.

1.4.1 Problems with Qualitative Techniques

The risk analysis methods that organizations use today to select SILs rely on either qualitative or quantitative techniques or, in some cases, a combination of the two. Qualitative techniques are very subjective and can be susceptible to large errors because of the psychological traps that can cloud human thought.

Table 1.2 Common Psychological Traps of Qualitative Analysis

Relying too much on first thoughts	>>	The anchoring trap
Focusing on the current design	>>	The status quo trap
Protecting earlier choices	>>	The sunk cost trap
Focusing on dramatic events	>>	The recallability trap
Neglecting relevant information	>>	The base rate trap
Slanting probabilities and estimates	>>	The prudence trap

Researchers have repeatedly shown that humans are actually quite horrible judges of the frequency of events that occur at long intervals. Despite this, human experts are expected to evaluate the difference between two events whose frequency is less than once in one thousand years (or ten lifetimes)! It's no wonder their results are often poor.

Luckily, the recallability trap and prudence trap conspire to create risk estimates that are too conservative rather than too aggressive. Although conservative estimates may mean the part of a plant in question is judged to be slightly safer than it really is, the resulting overdesigned systems require much more capital to install and maintain than is necessary. Studies have shown that more than 50 percent of a typical refinery's safety functions are overengineered. The extra capital spent for marginal improvement in a few arbitrary areas could always be spent more wisely elsewhere in the plant to improve safety more equitably over a broader range of situations. It is thus safer and more efficient to look at the entire