CHEMICAL ENGINEERING
September 2007
www.che.com

Environmental Manager

Tolerable Risk

While determined risk is generally well understood, tolerable risk can be the missing link to complete risk assessments

Mike Schmidt
Emerson Process Management

Safety instrumented systems (SIS) are in the spotlight these days. More companies have come to recognize the importance of recently published consensus standards, and recent fatal accidents tied to the failure of instrumented safeguards have brought further attention. The result is that facilities all over the world are looking at their safety interlocks and concluding that what they really need is a full-blown SIS.

As some work to upgrade their SISs, they find plenty of guidance available for the early steps — process hazard analysis (PHA), consequence analysis, and layers of protection analysis (LOPA) — and also guidance for the later steps — component selection, failure probability and SIL (safety integrity level) verification calculations, and system-design guidelines. A disconnect comes when end users have developed their lists of safety instrumented functions (SIFs) and are asked to assign an SIL to each one (Table 1).

Detailed design of an SIS cannot begin until each SIF has been assigned an SIL. This is the point in the process where many are tempted to throw up their hands in despair and want to arbitrarily assign everything as SIL 3. This is not the answer. At the very least, it fails to comply with the letter and the spirit of the standards, but worse, it misallocates resources to overdesign when these resources could be better used to reduce risk elsewhere.

The disconnect occurs because SIL assignment depends on a comparison of two values: determined risk, about which much has been written; and tolerable risk, which is frequently invoked but rarely presented in terms that are useful for SIL assignment. Instead, the literature usually contains some variation of this statement: "Each organization is responsible for establishing its own risk-tolerance criteria." Some organizations have corporate risk-tolerance criteria that speak to the total corporate risk, but many have no risk-tolerance criteria at all. It is still a rare organization that has established risk-tolerance criteria that are actually useful at the level of a process unit, where the SIS resides.

WHAT IS AN SIS?

Various acronyms are used in discussions about safety instrumented systems (SISs). While these acronyms are defined in the text of this article, a little more explanation is warranted.

The acronym "SIL" stands for safety integrity level. Each SIL represents the difference between the process risk without an SIF, and the risk deemed tolerable. An "SIF" is a safety instrumented function and consists of the sensors, logic solvers, and final elements necessary to detect a particular hazardous condition and cause the process to go to a safe state. It's important to note that SILs apply to each SIF, not to the SIS as a whole.

An "SIS" is a collection of all the sensors, logic solvers, and final elements necessary to address all the SIFs in a process, hence the word "system". SILs establish the required reliability of an SIF. Once a safety instrumented function is designed, reliability calculations can show if the SIF is reliable enough to reduce the risk to a tolerable level.

Safety integrity levels range from 1 to 4, but it should be noted that SIL 4 is not recognized in all SIS standards. Instead, there is an expectation that the process will be re-designed to reduce the risk. An SIS can also include SIFs that are not SIL rated. Various users refer to them as "N/R" (not rated) or "SIL 0" — handy terms, albeit not defined in any standards. ❑

Recommended approach

A typical approach to SIL assignment, but certainly not the only one, is to use a risk matrix. Risk is a product of consequence

TABLE 1. SAFETY INTEGRITY LEVELS

Safety Integrity Level	Probability of Failure on Demand (PFD_{AVG})	Risk Reduction Factor (RRF)
SIL 4	10^{-4} > PFD > 10^{-5}	10,000 < RRF 100,000
SIL 3	10^{-3} > PFD > 10^{-4}	1,000 < RRF < 10,000
SIL 2	10^{-2} > PFD > 10^{-3}	100 < RRF < 1,000
SIL 1	10^{-1} > PFD > 10^{-2}	10 < RRF < 100

TABLE 2. RISK MATRIX

		Consequences (per event)				
		< 0.01 x Serious	> 0.01 x Serious	> 0.1 x Serious	> 1 x Serious	> 10 x Serious
Likelihood	*f* > High	N/R	SIL 1	SIL 2	SIL 3	SIL 4
	High > *f* > 0.1 High	N/R	N/R	SIL 1	SIL 2	SIL 3
	f < 0.1 High	N/R	N/R	N/R	SIL 1	SIL 2

f =Benchmark frequency

and likelihood. As shown in Table 2, consequence and likelihood categories, like the SILs themselves, are separated by orders of magnitude. The challenge for each company then, is to define the consequences that are considered "serious" and the benchmark frequency (*f*).

Many companies already use some type of matrix, usually to prioritize recommendations from HazOps (hazard and operability reviews) or other PHAs. An example of this is shown in Table 3, where a rating of 1 would mean something should be done immediately, and a 10 would mean that no action is required.

Typical likelihood categories include terms such as frequent, occasional, seldom, remote and unlikely — but what do they mean? A company may leave it to the individual participants in the PHA to decide for themselves.

Some companies define these terms, but may use frequency intervals that are not uniformly spaced. A typical set might be:

- Event occurs once a year
- Event occurs once every 10 years (~1 order of magnitude)
- Event occurs once every 50 years (~ ½ order of magnitude)
- Event occurs once every 150 years (~ ½ order of magnitude)
- Not likely to occur

The same holds true for the severity in the consequence categories; a typical (nonlinear) arrangement might be:

- Loss of life; damage over $1 million
- Lost-time injury; damage over $500,000
- Medical treatment; damage less than $500,000
- Minor injury; near miss; poor quality
- No injury, impact on process

Converting an existing matrix

To convert an existing matrix to an SIL-assignment matrix, the first step is to adjust the likelihood categories so they are one order of magnitude apart. After that, consequence categories should be adjusted so they are also one order of magnitude apart. The intervals could be 1.5, 15 and 150 years, or 5, 50 and 500 years, as long as they are an order of magnitude apart.

TABLE 3. AN EXAMPLE OF A MATRIX FOR PHAS

		Severity (Consequence)				
		5	4	3	2	1
Likelihood	1	5	4	3	2	1
	2	8	7	6	4	2
	3	9	8	7	6	3
	4	10	9	8	7	4
	5	10	10	9	8	5

TABLE 4. RELATIVE COST OF CONSEQUENCES

Fatalities	1.0x
Serious injuries	0.1x
Injuries (reportables)	0.01x
First aids (non-reportables)	0.001x

The next step is to define the consequence categories. The universally used scale — and absolutely required by the SIS standards — is of consequences to personnel. But other important consequence scales, like community, societal or environmental consequences, may also show up in one of these matrices. Just keep in mind that as an SIF for reducing the risk of a specific hazard is evaluated, the consequences of that hazard must be evaluated on each scale. The event will have the same frequency, regardless of the types of consequences considered, and the consequence that results in the highest SIL rules.

While some organizations want to consider site and operability consequences of a hazard in their analyses, there are serious disadvantages to keep in mind. One, it puts a dollar scale on the same matrix as environmental, health and safety (EHS) consequences, and it is not a big leap for a plaintiff's attorney to equate the EHS consequences to the dollar amounts. Also, functions that strictly protect against economic loss don't need to be burdened with the documentation requirements of the SIS. Organizations are usually better off to leave the economically driven functions in the basic process control system and avoid giving them SIL ratings.

Consequence categories

If classifying likelihoods into categories that are separated by orders of magnitude is hard, it's nothing compared to separating consequences into categories. Once a benchmark frequency is defined as a number, dividing or multiplying it by ten is straightforward. Development of consequence categories, however, requires agreement on what consequence is ten times worse than another.

In the early twentieth century, H.W. Heinrich [1] proposed the idea that for each fatality there were many incidents with less serious consequences, and that for each incident with less serious consequences, there were even more near misses. From that simple observation grew the notion that if we could be more alert to near misses and reduce them, we would necessarily reduce the incidence of more severe consequences. This led to the now familiar, safety pyramid (Figure 1).

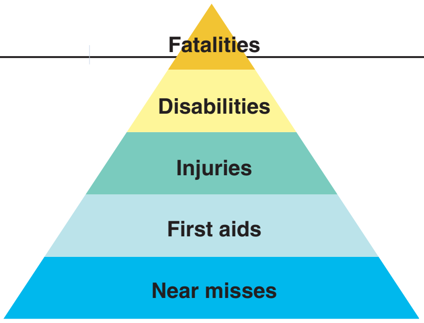

FIGURE 1. The Accident Triangle or Safety Pyramid depicts the simple observations of H.W. Heinrich [*1*]

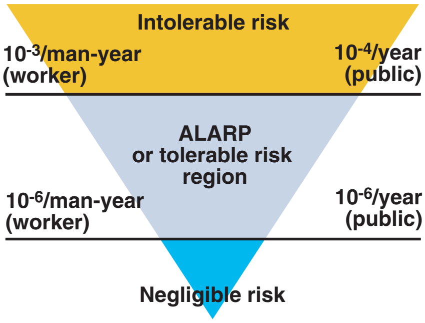
Intolerable risk
10^{-3}/man-year (worker)
10^{-4}/year (public)
ALARP or tolerable risk region
10^{-6}/man-year (worker)
10^{-6}/year (public)
Negligible risk

FIGURE 2. The ALARP principle, developed in the U.K., states that risks should be reduced to a level that is "as low as reasonably practicable"

While the premise of the safety pyramid has recently been called into question, Heinrich's original observation is still valid: near misses, first aids, and fatalities are each separated by about an order of magnitude (Table 4).

Many have published cost estimates associated with increasing levels of injury, with costs ranging from a few hundred dollars for a first-aid injury to a million dollars for a fatality. What the various scales have in common is that severity of consequences, when reduced to dollars, seems to be separated by orders of magnitude.

Once likelihood and consequence categories that are separated by orders of magnitude have been established, an SIL-assignment matrix can be developed. An example of a blank table is given in

TABLE 5: SIL ASSIGNMENT MATRIX, WITHOUT SILS

		Consequences (per event)				
		5 < 0.1 injuries	4 ≥ 0.1 injuries ≥ 1 first aid ≥ 1 near miss	3 ≥ 0.1 disability ≥ 1 Injury ≥ 10 first aids	2 ≥ 0.1 fatality ≥ 1 disability ≥ 10 injuries	1 ≥ 1 fatality ≥ 10 disabilities
Likelihood	1. Frequent (>1 event/1.5 years)					
	2. Occasional (>1 event/15 years)					
	3/4. Seldom/Remote (>1 event/150 years)					
	5. Unlikely (<1 event/150 years)					

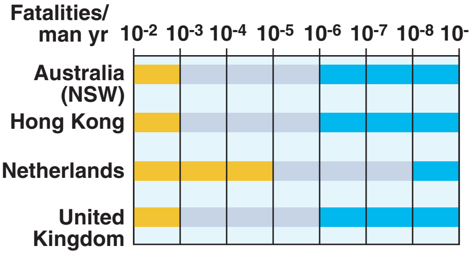

FIGURE 3. Government mandates defining tolerable risk are often referred to as bright lines. The color coding here refers back to the 3 levels in the ALARP principle as shown in Figure 2. The U.S. does not set tolerable risk levels, nor offer guidelines

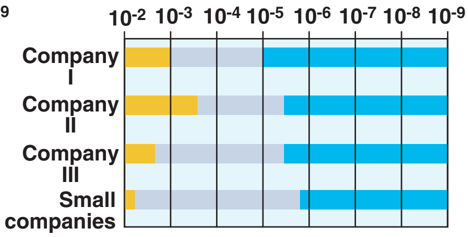

FIGURE 4. CPI benchmarks show that large, multinational companies tend to set levels consistent with government mandates. Smaller companies tend to operate in wider ranges and implicitly at higher levels of risk

TABLE 6: VOLUNTARY RISKS

Smoking	9.7 x 10^{-3}/yr or 9,700 fatalities per year per million smokers
Automobile accident	1.4 x 10^{-4}/yr or 140 fatalities per year per million people
Lightning strike	1.5 x 10^{-7}/yr or 15 fatalities per 100 years per million people

Table 5. All that remains is to associate an SIL with each cell in the matrix.

What is a tolerable risk?

SILs represent the residual risk between what is left after all other layers of protection have been applied to a hazard, and the tolerable risk. But what is the tolerable risk? Consider this: Whether or not your company has established a tolerable risk criterion for your facility, what do you believe it should be? How infrequently would a fatality occur for you to be able to describe your facility as safe? Once a year? Once every three years? Once every 5,000 years? Once every million years?

The ALARP principle (Figure 2) was developed by the Health and Safety Executive of the U.K. It states that risks should be reduced to a level that is "As low as reasonably practicable." The ALARP principle first divides risk into three areas: Intolerable Risk at the high end, Negligible Risk at the low end, and the Tolerable Risk that falls between the two. Risk in the middle region can be tolerated as long as all cost-effective measures to reduce risk have been put into place. Anyone operating a process with risks in the tolerable risk region must demonstrate that they have achieved the lowest risk possible, taking into consideration cost versus risk reduction.

Government mandates

Some countries have mandated upper and lower levels for tolerable risk. Most countries that have set these levels use the average level of risk faced by the entire workplace population, regardless of occupation, as the upper limit for tolerable risk. In many developed countries, that is the equivalent of one fatality per year per 1,000 workers. This upper boundary represents the level of fatality rate already tolerated in industries such as underground mining or deep-sea fishing. The lower boundaries are generally two or three orders of magnitude lower.

The Netherlands took a similar approach, but then declared that workplace risk should be 1% of general risk, because it is not voluntary risk. This results in levels that are 100 times lower than those chosen by most countries that set levels.

These levels, as shown in Figure 3, are sometimes called "bright lines," especially when discussed in the U.S., where there has been no success in setting tolerable risk levels. The argument was made that "a strict bright line approach to decision making is vulnerable to misapplication since it cannot explicitly reflect uncertainty about risks, variation of susceptibility within a population, community preferences and values, or economic considerations — all of which are legitimate components of any credible risk management process." In other words, in the U.S., tolerable risk remains very much subject to the political process.

CPI benchmarks

Almost all major companies in the chemical process industries (CPI) have guidelines for tolerable risk, but they are generally considered highly confidential. No company is going to announce to the public, for example, that they would tolerate up to one fatality per year per ten thousand workers.

Some companies, especially smaller ones, do not have explicit guidelines. However, the risk they tolerate can be inferred from the risk reduction measures they have in place. Smaller companies without explicit guidelines generally tolerate greater levels of risk than their larger counterparts and this isn't inappropriate.

A small company, struggling to establish itself or to make payroll, is going to look at the amount of risk that it can and must tolerate to get through the next year very differently from a company that is looking out over the next 100 years. A small company with a single plant will look at a single event that occurs once every 250 years and may very well decide that once every 250 years might as well be never. A large company, with 50 similar plants around the world would look at that same risk as once every 250 years for each plant, but once every 5 years for the corporation, and make a very different decision about whether the risk is tolerable (Figure 4).

Voluntary and 'natural' risks

The levels mandated by governments, or self-imposed by companies, can best be understood in terms of some of the risks to which people voluntarily expose

TABLE 7. COMPLETED RISK MATRIX FOR EXAMPLE A

		5 < 0.1 injuries	4 ≥ 0.1 injuries ≥ 1 first aid ≥ 1 near miss	3 ≥ 0.1 disability ≥ 1 Injury ≥ 10 first aids	2 ≥ 0.1 fatality ≥ 1 disability ≥ 10 injuries	1 ≥ fatality ≥ 10 disabilities
Likelihood	1. Frequent (1 event/1.5 years)	N/R	SIL 1	SIL 2	SIL 3	Redesign
	2. Occasional (1 event/15 years)	N/R	N/R	SIL 1	SIL 2	SIL 3
	3/4. Seldom/Remote (1 event/150 years)	N/R	N/R	N/R	SIL 1	SIL 2
	5. Unlikely	N/R	N/R	N/R	N/R	SIL 1

themselves. That is not to say that risks like these should ever be compared to process risks in a public setting, which would likely inflame the audience. But, for one's own understanding, it is helpful to know that the upper level for ALARP is typically set at 1 x 10^{-3} per year, or almost ten times lower than the risk to which smokers expose themselves. If a chemical process had the same risk as smoking, it would be considered an intolerable risk.

On the other hand, dying from a lightning strike would generally be considered a negligible risk. If you consider your own personal risk management strategies, chances are that you have not invested in lightning protection at your home. Even in your plant, the lightning protection is for electrical equipment, not personnel.

The risk of a fatality in a car accident is within the ALARP region. Of course driving has the potential for fatalities, but the risk is tolerable because the convenience or benefit is worth it. If it wasn't, society would insist on lower speed limits and more safety features in cars.

Example A: 1 fatality/100 years

Let's get back to SIL assignment. While each of us has a different level of personal tolerable risk, when asked how small the average frequency for a fatality would need to be for a plant to be considered safe, it is not at all uncommon for people to pick 100 years. This is longer than any one person's career but short enough to be meaningful in terms of the life of a facility. So consider a facility with a tolerable fatality rate of one per 100 years.

Further assume that at this facility, the workforce includes 250 workers who would be exposed to process risks:

(1 year/250 man-years) (1 fatality/100 years) = 1 fatality/25,000 man-years
= 4 x 10^{-5} fatalities/man-year
= total tolerable risk

In this case the total tolerable risk is one fatality per 25,000 man-years, or 4 x 10^{-5} fatalities per man-year, which falls in the middle of the tolerable risk region.

Process risk is only part of risk

Note that the tolerable risk established above is for all fatalities in the plant. However, not all fatalities result from process safety risks, that is from process hazards. There are still the slips, trips and falls, the housekeeping issues, the electrical safety and confined-space entry issues, and transportation-related fatalities.

For Example A, assume that half of the risk comes from process safety risks or process hazards. If total tolerable risk is 4 x 10^{-5} fatalities/man-year, then process safety risk is half of that, or 2 x 10^{-5} fatalities/man-year.

Although we've assumed that no more than half of the risk comes from the process hazards, we need to consider all of the process hazards that could lead to a fatality. Unless there is only one potentially fatal hazard, not all of the process safety risk can be allocated to a single hazard. One approach is to assume workers are each exposed to a certain number of potentially fatal hazards, say five, and then distribute the "risk budget" over those five hazards in establishing SIL selection criteria. Other hazards, although not expected to be fatal, can then use the same criteria. If total process safety risk is 2 x 10^{-5} fatalities/man-year, then:

N/R: ≤ 4 x 10^{-6} fatalities/man-year
SIL 1: ≤ 4 x 10^{-5} fatalities/man-year
SIL 2: ≤ 4 x 10^{-4} fatalities/man-year
SIL 3: ≤ 4 x 10^{-3} fatalities/man-year
SIL 4 or re-design: > 4 x 10^{-2} fatalities/man-year

Putting a stake in the ground

With the tolerable risk criteria settled, then the next step is to assign an SIL to one of the cells in the matrix. For this example, let's look at cell 2–2 in Table 5.

We want to consider the risk at the middle of the cell. Since the scale is logarithmic, the middle of the cell on the likelihood scale is the log (geometric) mean of 1.5 and 15 years, or 4.7 years. The log mean of 0.1 fatalities per event and 1 fatality per event is 0.32 fatalities per event.

Doing the math, the center of our target cell has a risk of 2.7 x 10^{-4} fatalities per man-year. Going back to our tolerable risk criteria, that is less than 4 x 10^{-4}, so it translates to SIL 2.

Once we can assign SIL 2 to that cell, completing the rest of the matrix is easy (Table 7). The cell to the left is one order of magnitude less risk, so it is SIL 1. Everything to the left of that is not rated. Similarly, the cell below is one order of magnitude less risk, so it is also SIL 1. Everything below that cell is not rated. Filling in the rest of the matrix follows the same logic.

When we're finished, we have a risk matrix that can be used for SIL assignment for all SIFs, whether or not fatal consequences are anticipated.

Interestingly enough, when the risk matrix is finished, it spells out the criteria for assigning an SIL to an SIF without ever explicitly acknowledging a tolerable risk. The tolerable risk is only implied.

Example B: 1 x 10^{-6} fatalities/man-yr

Now consider Plant B. It is very similar to Plant A, with the same number of exposed workers and the same number of potentially fatal hazards. In this case, however, senior management has declared that risk needs to be reduced to the level considered "negligible" on the ALARP diagram (Figure 2). That is, less than one fatality per million man-years.

The portion of total risk to an individual attributable to process safety risk would again be half, or 5 x 10^{-7} fatalities per man-year. Once again, with each individual worker exposed to about five potentially fatal hazards, the threshold for avoiding an SIL-rated SIF is 1 x 10^{-7} fatalities per man-year.

Likewise, SIL 1 is less than 10^{-6}, SIL 2 is less than 10^{-5}, and SIL 3 is less than 10^{-4}. In the CPI, anything larger than 10^{-4} would need to be redesigned to reduce the hazard before a safety instrumented function for risk reduction was even considered.

TABLE 8. COMLETED RISK MATRIX FOR EXAMPLE B

		5 < 0.1 injuries	4 ≥ 0.1 injuries ≥ 1 first aid ≥ 1 near miss	3 ≥ 0.1 disability ≥ 1 Injury ≥ 10 first aids	2 ≥ 0.1 fatality ≥ 1 disability ≥ 10 injuries	1 ≥ fatality ≥ 10 disabilities
Likelihood	1. Frequent (1 event/1.5 years)	SIL 2	SIL 3	Redesign	Redesign	Redesign
	2. Occasional (1 event/15 years)	SIL 1	SIL 2	SIL 3	Redesign	Redesign
	3/4. Seldom/Remote (1 event/150 years)	N/R	SIL 1	SIL 2	SIL 3	Redesign
	5. Unlikely	N/R	N/R	SIL 1	SIL 2	SIL 3

Looking at the middle of the same cell, we still end up with a risk of 2.7 x 10^{-4} fatalities per man-year. The plant is the same, so the risk is the same. What's different is the tolerable risk. In the case of Plant B, the tolerable risk is set at a much lower level, and so in this case, the risk is not quite low enough to meet the SIL 3 rating:

(1 event/4.7 years) (0.32 fatalities per event) (1 yr/250 man-years) = 2.7 x 10^{-4} *fatalities/man-year*

SIL rating = SIL 4 or redesign

As in the first example, the rest of the matrix is easily populated. Because of the difference in tolerable risk, the SIL assignments have shifted toward the lower left-hand corner of Table 8.

Variations to consider

A variation to consider, depending on risk philosophy, is what part of a cell to benchmark to. In our examples, we took the center of a cell. If the upper right corner is used instead, the result will be an SIL assignment that is higher by one level. This is clearly a more conservative approach.

Another variation to consider is the number of categories for both likelihood and for consequence. Typically, a 5 x 5 matrix is the largest workable matrix, with a larger matrix, the team doing risk assessment will probably spend more time debating categories than is warranted. By the same token, if the matrix is smaller than 3 x 3, it is unlikely to have enough detail to be worthwhile.

In general, it is a good idea to extend the matrix far enough to the left to get a column of all non-rated SIFs, or else all the SIFs, regardless of risk, will end up being SIL rated. By the same token, there is no reason to extend the matrix farther to the left than a column that is all "N/R", or farther to the right than is all "redesign".

Business results achieved

When a company establishes a risk matrix for SIL assignment, the matrix receives a fairly broad distribution. To the extent that the company is uncomfortable publicizing its tolerable risk criteria, the risk matrix approach to SIL assignment does not require an explicit statement of tolerable risk.

Because likelihood and consequence only need to be estimated to the nearest order of magnitude, most SIL assignments can be made without using detailed and expensive quantitative risk-analysis techniques. It is usually easy to get agreement, based solely on the experience of the risk-assessment team, that an event is likely to occur between 10 and 100 years, rather than spend the time determining whether it will be once every 37 years or once every 54 years.

Since it is usually easy to get agreement in the middle regions of consequence and likelihood, quantitative risk analysis can be reserved for the more extreme situations, such as hazards that are either expected to be very, very rare, or to have very, very large consequences. And even in those cases, the quantitative analysis only needs be done if there is a question about the SIL assignment of a particular SIF.

Finally, risk matrix is easy to apply, and will result in hazards and their associated risks being treated similarly from project to project. This will result in much less under-specification with intolerably high risk, and much less over-specification with resulting misallocation of limited safety resources. ■

Edited by Dorothy Lozowski

Reference

1. Heinrich, H.W., "Industrial Accident Prevention: A Scientific Approach," 4th ed., McGraw-Hill, N. Y., 1959.

Author

Mike Schmidt is a principal consultant with Emerson Process Management at the Refining and Chemical Industry Center in St. Louis (641 Lambert Pointe Drive, Hazelwood, MO 63042; Phone: 314-872-6069; Fax: 314-872-8686; Email: mike.schmidt@emersonprocess.com). His responsibilities include facilitating HazOps and other PHAs, consequence analysis, facilitating LOPAs, reviewing and preparing SRSs (SIS safety requirements specifications), and performing SIL calculations. Schmidt consults on process design and optimization and teaches in all these areas. In addition to being a registered PE in several states, Schmidt is a CFSE. He has been with Emerson for over ten years, and has been working in the CPI since 1977. The majority of his career has been in operating companies, including Union Carbide, Rohm and Haas, and Air Products. He is the author of several articles on process design and process safety, and is currently serving on the CCPS committee that is preparing a guideline book on tolerable risk and on the API committee on overfill protection.