MARY KAY O'CONNOR
PROCESS SAFETY CENTER
TEXAS A&M ENGINEERING EXPERIMENT STATION

22nd Annual International Symposium
October 22-24, 2019 | College Station, Texas

Fire and Gas Hazard Mapping Continues to Require Engineering Judgment

James McNay BSc MIFireE CFSP MIET[2] and William Pittman, PhD*[1]
[1]Fire and Gas Detection Consultant, Micropack (Detection) Americas
[2]Managing Director, Micropack (Engineering) Ltd
*Presenter E-mails: williampittman@micropackamericas.com

Abstract

It is becoming increasingly common to use computer applications and software tools to aid engineers in designing fire and gas detector layouts. These tools help designers optimize layouts as well as verify, quantify and document the level of performance the system can be expected to achieve. The process of designing and assessing layouts in this way is called hazard mapping.

In an age where “optimization” is increasingly becoming associated with computerized approaches that seek to minimize or maximize objective functions, many in the industry see these software tools and assume or wonder if the layout is generated automatically by the mapping software. Upon learning that a human engineer must still design the detector layout and place the detectors in the facility model, many ask why it hasn’t been automated and how the system can be considered optimized if it isn’t designed with the assistance of an algorithm.

Fire and gas detector layout design requires the use of a great deal of engineering judgement which cannot currently be replicated or replaced by optimization algorithms. There are many practical concerns that must be addressed in the design of such systems that are difficult to model and automate. There are many rules of thumb, soft restrictions and best practices that are sometimes bent or broken to address the specific need of an application or facility. This paper reviews many of these issues in an attempt to make the case that the expert judgment of a human engineer is and will continue to be essential to the design of truly optimized fire and gas detection systems.

Introduction

Ever since fire and gas detectors were first brought to market the industry has faced the challenge of designing detector layouts and assessing, either qualitatively or quantitatively, the coverage achieved by the system. Two important and inter-related questions must always be addressed:

1. How much coverage is need? Alternatively: What level of coverage is adequate or acceptable?
2. Does the detector layout currently in use or planned for installation deliver that level of coverage?

Fire and gas hazard mapping studies, so-called “coverage assessments,” attempt to help engineers and operators answer the second question, but the answers generally are not straight-forward and do not lend themselves to “cookie-cutter” approaches.

The goal of a design engineer is, as always, to design an “optimized” system, with “optimized” in this case usually meaning that the system provides an acceptable level of coverage - as defined by whoever or whatever answered the first question - with the lowest possible number of detectors. In principle, this should deliver the system that fulfills the requirements of the facility or operator at the lowest possible cost.

The problem with “optimized” these days is that, when most people hear it, they tend to assume that the entire design process is being handled by a computer with an objective function, an algorithm or some weighted set of criteria that it’s trying to minimize or maximize. When discussing fire and gas hazard mapping, it is not uncommon for some to ask, “are the detector locations being chosen by a human designer or by a computerized optimization approach.” The answer to this is, the detectors are placed by a human design engineer.

The question that inevitability follows is usually some form of “why not have the computer do it?” or “then how can you be sure that the system is optimized?”

For flame detection coverage assessments, the approach is generally to calculate the percentage of the volume with different levels of coverage. The designer specifies the volume(s) or space(s) where-in flame detection is expected or required. Software tools are then used to model the coverage provided by a detector layout and determine if one or more detectors have a clear line of sight to that volume, if the volume is within a detector’s field of view, and if the detector is close enough to detect a fire of the specified size. This information is used to calculate what percentage of the volume of interest has coverage and at what level - usually 1ooN or 2ooN.

At first glance it might seem to most that this problem would lend itself to computerized optimization programs. All the system must do is to either take a fixed number of detectors and arrange them in a way that maximizes coverage. Alternatively, the system can start from nothing and add detectors at the location that incrementally improves the coverage the most until the coverage target is met or exceeded. That sounds easy. Right? Not exactly.

Is a target percentage coverage, in and of itself, excluding other considerations, an appropriate test of effectiveness or acceptability? No.

The design for flame detector layouts requires an appreciation of many practical concerns that are not necessarily easy to factor into computer-based optimization systems. This paper reviews and discusses some of these issues, sometimes illustrating the point using example 3D coverage assessments and results. Because of these issues, the authors regard it as unlikely in the near term that any computerized system or algorithm will be able to provide an effective and reliable substitute for an experienced engineer using sound judgement to design a fire and gas detector layout. Some of these issues will impact the design of both flame and gas detector layouts, some are more limited to flame detection.

What Optimization Criteria are Used and How are They to be Ranked?

For flame detection coverage assessments, most guidance documents currently in use assign a rank or “grade” to equipment based on the perceived fire risk associated with it - usually into “low”, "medium," "high" and / or "special" categories. The standards put forward target coverage levels and performance targets for each risk rank or grade. The Performance Targets will usually include a target fire size that the detector layout should be able to detect in a timely manner (usually 10s or less) for each risk grade - smaller target fire sizes are selected for higher risk equipment with 50 KW RHO being typical for high risk equipment, 250 KW RHO being typical for medium risk equipment, and 500-1000 kW RHO being typical for low risk equipment.

Many operators will design a FGS to generate an alarm in response to "unconfirmed," 1ooN, detection and take an automated control action in response to "confirmed," 2ooN, detection.

The combination of voting schemes and risk ranking means, in an area where-in equipment with all three risk grades are present, a designer will not have a single coverage value consider. Rather, there could be eight values to consider - 1ooN and 2ooN coverage for high, medium, and low risk equipment, and the module overall. This forces a designer to ask, which value should we optimize based on? If we rank and prioritize the different values in a fixed hierarchy, how do we arrange them? Is 1ooN coverage more important or 2ooN? If two detector locations are being considered and one improves 1ooN coverage more but the other increases 2ooN coverage more, which do you select?

To illustrate how this can happen, consider the situation shown in Figure 1 below where there are 3 vessels in an area with a medium risk grade assigned: Figure 1 and all subsequent figures in the paper have been generated using HazMap3D, a hazard mapping and coverage assessment application developed and used by Micropack.

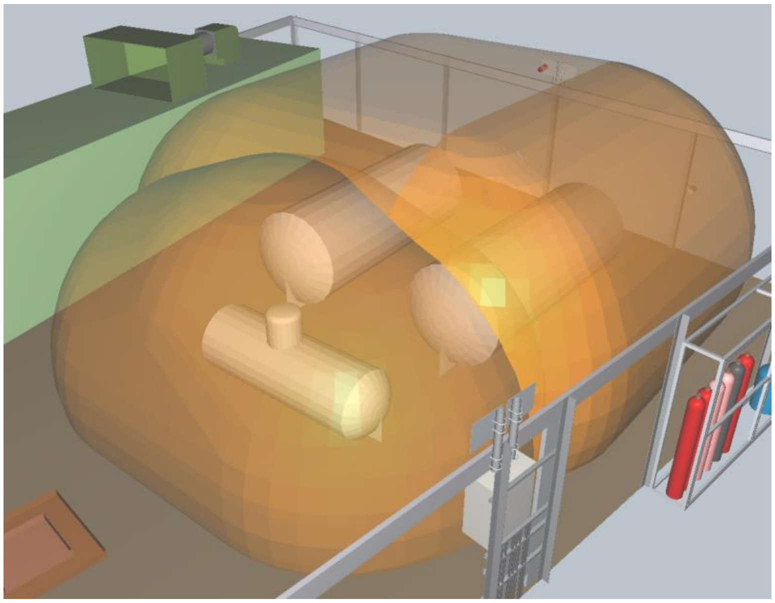

Figure 1: Three vessels assigned a "medium risk" grade.

To provide coverage in this case, the design engineer might choose to recommend a detector layout using 3 detectors. An assessment of such a 3-detector layout, assessed in a manner consistent with TR 84.00.07, indicated that the layout would give 81% 2ooN coverage, giving the 80% control action coverage recommended in ISA TR 84.00.07 for medium risk assets. Alarm-only 1ooN coverage was assessed at about 98%.

Table 1 provides a summary of the coverage levels achieved for 1ooN, 2ooN and 3ooN voting. Table 2 also shows, in the “individual” column, what percentage of the graded volume each detector can see, and, in each row, what coverage is reduced to for 1ooN, 2ooN and 3ooN voting if the detector named on that row fails.

Table 1: Summary of voted coverage and detector coverage contributions for the three vessels using a three-detector layout.

	Individual	**1ooN**	**2ooN**	**>2ooN**
All Detectors		98.2	81.1	50.1
Det01	81.1	93.0	55.3	0.0
Det02	74.2	92.1	63.0	0.0
Det03	74.1	92.4	62.9	0.0

From this, we can see a few things. Det03 covers the smallest percentage of the graded volume, although it comes in only slightly lower than Det02. Unconfirmed 1ooN coverage increases more when Det02 is added to the layout (rising from 92.1 to 98.2) but confirmed 2ooN detection rises more when Det 03 is added to the other 2 (rising from 62.9 to 81.1%). Det01 sees the largest fraction of the graded volume (81.1%) and is clearly the most important detector for providing 2ooN coverage, but it adds the least amount of the three to 1ooN coverage. So even this simple example, where-in only one risk grade is present, shows that choosing where to add another detector or the next detector to a layout may not be obvious when only looking at % coverages to see which one gives you more coverage.

What is more important? Coverage for equipment and assets that are deemed “High Risk” (HR), “Medium Risk” (MR), or “Low Risk” (LR)? If two detector locations are being considered and one improves HR coverage more but the other increases MR coverage more, which do you select? It might seem easy to say, “choose the one that increases HR coverage more, it’s HR!” However, what if one option gives you 0.1% higher HR coverage where the other gives you 2.0% higher MR coverage? Most would likely take 2.0% MR coverage over a paltry 0.1% increase in HR coverage, all else equal. How would one program an algorithm to make a choice between options like that?

This is not to say that different optimization systems cannot be ranked or that a weighting system cannot be developed and programmed into a computer program. However, the complexity of the considerations may make designing and implementing such a system extremely difficult with detector coverage and hazard mapping. A design produced in such a manner also would not represent a true, objective, global optimum. It would simply represent the optimum using that weighting system for the selected criteria. It is not credible to think that such a system would always develop a layout that experienced designers would universally agree is truly optimized.

Concerns regarding prioritizing 1ooN coverage or 2ooN coverage will impact flame detector and gas detector layouts. Designers are far more likely to encounter issues arising from multiple risk grades being present in an assessment area with flame detection than with gas detection. In flame detection the individual pieces of equipment are assigned risk grades where, with gas detection, the entire assessment is usually assigned a single risk grade based on the levels of congestion and / or confinement.

Overall Coverage Levels Say Little about the Nature of The Gaps

A percent coverage alone says nothing about the size, number, or location of the gaps. When the full volume of an area is considered at once (coverage across the entire volume shown simultaneously), percentage coverage results can appear lower than expected while the area appears to have suitable coverage overall. This can result from having a large number of small gaps that collectively add up to a significant fraction of the graded volume. These gaps may be in small blind spots that are located adjacent to areas with good coverage and which are not large enough to hide a fire that is likely to occur in that unit or likely to cause significant damage.

Even where gaps exist that are large enough to conceal a large fire, not all such gaps area equally problematic. A large gap in coverage on the back side of a storage tank, where there are no pumps or pipe connections, is not as concerning as a coverage gap where transfer pumps and pipe connections are located next to the tank. A gap in coverage behind the driver of a pump is not as concerning as a coverage gap on the liquid handling side of the pump. Oftentimes, the dominant fire scenario in such situations is a liquid pool fire and the decks around tanks and pumps is usually slightly sloped. This will cause the liquid pool to flow and expand, often out of the gap in detector coverage, into an area of the deck where detection will occur with minimal delay.

In these cases, it is of critical importance to analyze the areas of low coverage and make a judgement upon the suitability of the detection arrangement, based upon knowledge of the area and equipment, identified fire scenarios, the size of the coverage gaps relative to the size of fires expected in the area, and the location of the gaps. For gas detection coverage assessments, designers may also need to consider issues like whether an area is naturally ventilated or subject to forced ventilation. If forced ventilation is used in an enclosed space and air-flow patterns are more predictable, a designer may be able to excuse or accept some coverage gaps on that basis.

Achieving Target Coverage Levels is Often Very Difficult in Some Environments

As the foregoing example in the previous section shows, achieving high levels of 1ooN or 2ooN coverage with flame detectors is often easy in areas that are very open and where there is not a high degree of congestion. However, on offshore facilities space is at a high premium and even at some onshore facilities congestion is quite high, sometimes because the available land area was small or because equipment was clustered together to lower the amount of piping that had to be installed in the unit.

Where congestion is high and there are many large and small visual obstructions in the process area, providing high levels of flame detector coverage is often quite difficult when conducting assessments with detailed, accurate 3D models. Small and large diameter piping, manifolds, decking and steel supports combine in such spaces to severely limit a detectors ability to achieve a clear line of site over long distances. In such process areas and facilities, it quickly becomes difficult, if not practically impossible, to achieve the 70-90% coverage that is usually specified in company-specific performance targets, even with only 1ooN coverage.

For an example of this, consider the following figure, which shows images of a process area in a 3D model of an offshore platform:

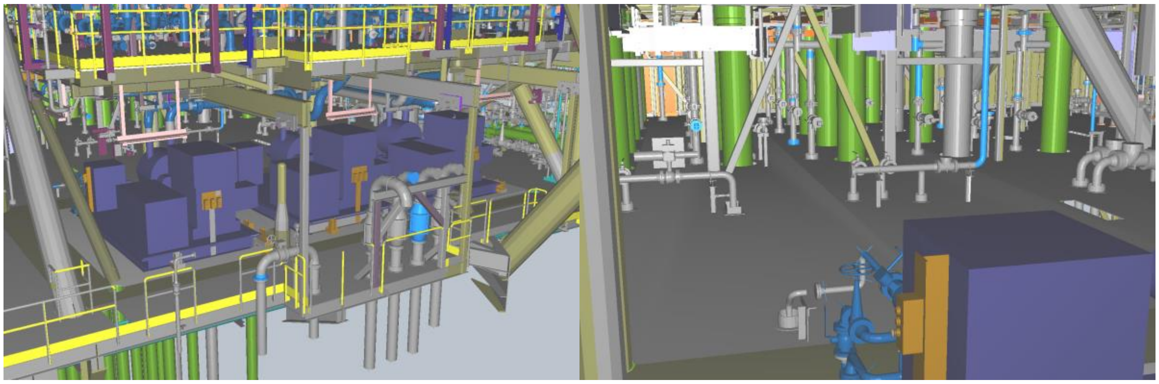

Figure 2: Screen shots of a coverage assessment area in a 3D model of an offshore platform.

Figure 3 is provided to give the reader an idea of the flame detection / fire risk grading that was applied in this area.

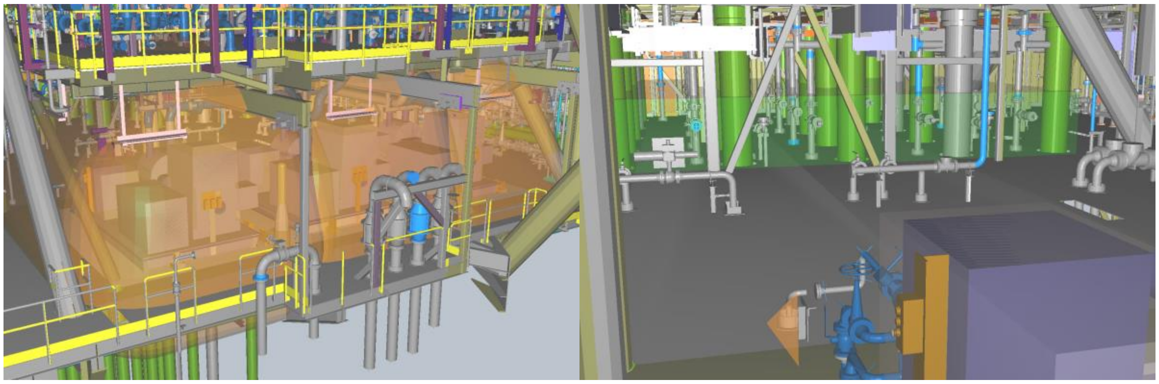

Figure 3: Screen shots showing the grading applied to a section of an offshore platform

The tight confines and high level of congestion is a kind of nightmare scenario for flame detection design. The process area shown in Figures 2 and 3 contains only medium and low risk equipment and measures only 32 meters by 17.5 meters with a deck height of about 3.5 meters. Based on the grading applied and the size of the space alone, one might think that such a space could easily be covered with 3-4 detectors, similar to the situation with the three tanks above. However, because of the very high level of congestion within this area, and because the detectors cannot be set back, away from the equipment being covered as with the prior example, a layout using nine (9) detectors still produces lower coverage than what is achieved in the previous example, and this nine-detector layout achieves coverage that is still far from what many would consider "ideal." Table 2 summarizes the coverage achieved by that nine-detector layout.

Table 2: Summary of voted coverage by risk grade for the nine-detector layout.

	Green	Orange	Brown	Red
MR	62%	22%	2%	15%
LR	48%	31%	0%	21%
Overall	58%	25%	1%	16%

Table 3 gives the contributions of each detector to the overall layout, similar to Table 1 in the previous example:

Table 3: Summary of voted coverage and detector coverage contributions for the nine-detector layout.

	Individual	**1ooN**	**2ooN**	**>2ooN**
All Detectors		82.5	57.7	30.6
Det01	5.2	80.7	56.2	29.3
Det02	1.7	81.8	57.0	30.3
Det03	37.7	78.7	45.0	18.9
Det04	27.3	79.6	51.1	21.2
Det05	23.3	79.8	52.3	23.8
Det06	33.4	79.8	47.7	20.0
Det07	33.1	77.2	48.8	20.9
Det08	11.3	80.1	53.7	27.3
Det09	12.0	79.7	53.3	27.5

It is worth noting that, while in the previous example each flame detector covered 70-80% of the graded volume, in this area, 20-35% is more typical and some detectors cover only 5-12% of the covered volume. That these nine detectors were placed in what the designer felt were the most promising locations. Each incremental detector would therefore be expected to, on average, add less to the assessed coverage than the nine used here.

The nine-detector layout only achieved 2ooN coverage for 58% and 1ooN for 83% of the module overall. How many additional detectors might be required to raise 2ooN coverage for the LR equipment from 48% to the 70-90% that many operators list as a target? How many additional detectors might be required to raise MR 2ooN coverage to 80%? A human designer knows when a point of diminishing returns has been reached and when it is no longer practical, feasible, necessary or wise to keep adding detectors in a space. A computer program would keep adding detectors until the target coverage was achieved unless logic were built in to make it stop adding detectors if adding another detector would increase coverage by less than a pre-specified critical amount - perhaps 1.0% of the graded volume. A human designer would likely look at Table 3, realize that Det01 and Det02 aren't doing much, remove them, and recommend the installation of a seven-detector layout that provides coverage that is as good as is reasonably achievable. Depending on the operator's risk tolerance, Det08 and Det09 might be eliminated as well, leaving only the five detectors which individually cover 23% or more of the graded volume.

For gas detection, a designer might design a system to detect a certain percentage of scenarios assessed as part of a CFD dispersion study within a certain period of time after the release starts. Depending on the number of scenarios assessed and the results of the CFD analysis, the allowed time to detection and the targeted coverage percentage a computer algorithm tasked with designing that system could propose a very large number of detectors.

Coverage targets in the range of 70-90% were developed when operators assessed coverage based on a single 2D slice or plot plan rather than a full, detailed 3D model. Using this approach, accounting for the impact of small diameter piping and similar obstructions was difficult and the high coverage factors were specified to add a layer of conservatism. As assessments using 3D models become more common, this layer of conservatism and reliance on % coverage as a measure of the quality of coverage provided by a system may no longer be necessary or appropriate.

The Challenge of Mounting Location Selection

Detectors should not be mounted an any random location that isn't currently occupied by other equipment. One of the most important things to consider in designing a detector layout - whether or not the mounting location is appropriate - is not necessarily easy to program into a computer program, in part because of the number of considerations that must be made in deciding if a location is "suitable."

Detectors ideally should not be positioned randomly hanging in space in a model. Yes, this is done sometimes on the assumption that a new post or support is or can be placed in that location to allow the detector to be mounted in that location. However, designers tend to prefer to mount detectors on existing structures where the layout is being designed for an existing facility. Where a layout is designed for a green-field unit, operators are still not going to want to build extra supports and run extra wiring solely for the FGS. The goal will always be to use existing or otherwise already planned / required structures as much as possible.

There are other instances where a computer program might, out of "ignorance," if we can call it that, propose mounting a detector in a disallowed or otherwise reserved space. Operators will also not be thrilled if a computer algorithm designs a layout that puts a detector on or next to a railing that it can't be used for because that railing is required for maintenance activities.

Detectors cannot be mounted to an object / surface that vibrates too much. Flame detectors ideally should not be too close to the floor - they will be blocked too easily by workers, temporary equipment and other temporary obstructions. Flame detectors ideally should not be too high above the deck. This makes them much harder to access for testing and maintenance when needed. Because of all of this, most standards indicate that detectors should be situated between two and four meters above the local deck (ALD). The tilt on a flame detector should ideally be within the range of 20-30 degrees. However, these are not hard and fast rules. Flame detectors are sometimes installed only 1.5 meters above the deck. At other times they may be 6-8 meters above the deck when used in spaces with very high ceilings. Many manufacturers and operators allow 10-40 degrees when needed and in some cases 45-50 degree down-angles are used - but only when deemed truly necessary. It's generally wise to position gas detectors so that they are unlikely to be splashed, clogged, or otherwise fouled by liquid sprays or other contaminants.

If 2ooN coverage is desired or required, designers might catch a computer algorithm placing two or more detectors right on top of each other with maybe a one- or two-meter difference in elevation (stacking detectors), or placing detectors very close to each other, with flame detectors facing in the same or similar directions. On the surface, this is an easy way to get 2ooN coverage in a process area, which is why one might see it happen. The problem with this practice is that having two flame detectors so close together with nearly identical fields of view drastically increases the odds of both detectors being blocked by a temporary obstruction and the risk of both flame detectors being triggered by the same false alarm stimulus. If there is an automated suppression system triggered by the flame detection system, this could lead to large numbers of spurious activations. It might be possible to overcome this with a requirement that two detectors cannot be placed within, say, 5 or 10 meters of each other, or that two detectors placed within 5 to 10 meters of each other must have at least a 60° or 90° difference in their "pan" or compass orientation. The problem with this, however, is that there may be special or unique circumstances where-in a designer might feel it best to ignore or loosen this restriction - especially if the assessment area is relatively small.

Many of these rules of thumb need to be flexible, with designers allowed to “break the rules” a little in some cases to deal with the nuances and unique needs of a specific application and facility. How do you make a computer understand all of this?

There have been optimization approaches that allow a user to program in acceptable / allowed mounting locations for detectors.[2] This approach still does not really create an exhaustive list of all possible mounting locations. It just gives a list of mounting locations that the designer thought were promising enough to enter in as options. It would also be difficult to allow the computer to consider the possibility of adding new posts or structures to mount to, where absolutely needed, because the designer would then be tasked with somehow telling the computer what voids or open spaces it can suspend detectors in. All of this immediately destroys the prospect of having a truly 100% computer optimized design. Whether this is accomplished by telling the computer what locations, surfaces, or areas are allowed or by marking disallowed locations, this is potentially time consuming.

If all of that could somehow be overcome or if a designer simply decided that whatever list of mounting locations they provided would be “good enough” for the computer to choose from, how would a user or programmer get the computer to consider the relative desirability of a location? How would this be factored into the design? What happens if a situation arises in which one additional detector is needed to reach the specified target coverage level and there are three detector locations where adding a detector will increase coverage by the necessary amount? What if, hypothetically, one of these locations is significantly easier or safer to access for maintenance purposes - possibly because of noise levels, toxic gas hazards, elevation relative to grade, or some other consideration, but this location gives marginally lower / worse coverage than the other two locations? Most operators would rather have a detector in a location where it is easier to clean and maintain, but how do you make the computer understand this and select that mounting location? Even if mounting detectors only in locations that would be relatively easy to access for testing and maintenance would require the addition of two detectors rather than just 1, many operators would rather spend more on CAPEX to install the system in order to reduce future OPEX.

Effective Flame Detection Distance is Not a Single, Fixed Value

It was previously noted that many operators will specify different target fire sizes which the FGS will be tasked with detecting depending on assessed or assigned risk levels. Owing to the inverse square law that governs the intensity of light sources in relation to distance to the source, this means that a flame detector can be farther away from a low risk asset than it can be to a high-risk asset while still providing effective coverage.

To use an example, an operator might set the target fire size for “high risk” assets at 50 KW RHO and the target fire size for low risk assets at 500 KW RHO. In this case, with the target fire size for a low risk asset being 10 times larger, a detector will be able to provide coverage for low risk assets from more than three times the distance at which it can cover high risk assets.

To show how this might look, we can consider the hypothetical case of a set of three vessels with two flame detectors monitoring them from a distance of about 15-21 meters. The generic detectors have a range to a 1 sq. ft test fire (roughly 40 kW RHO) of 17.5 meters.

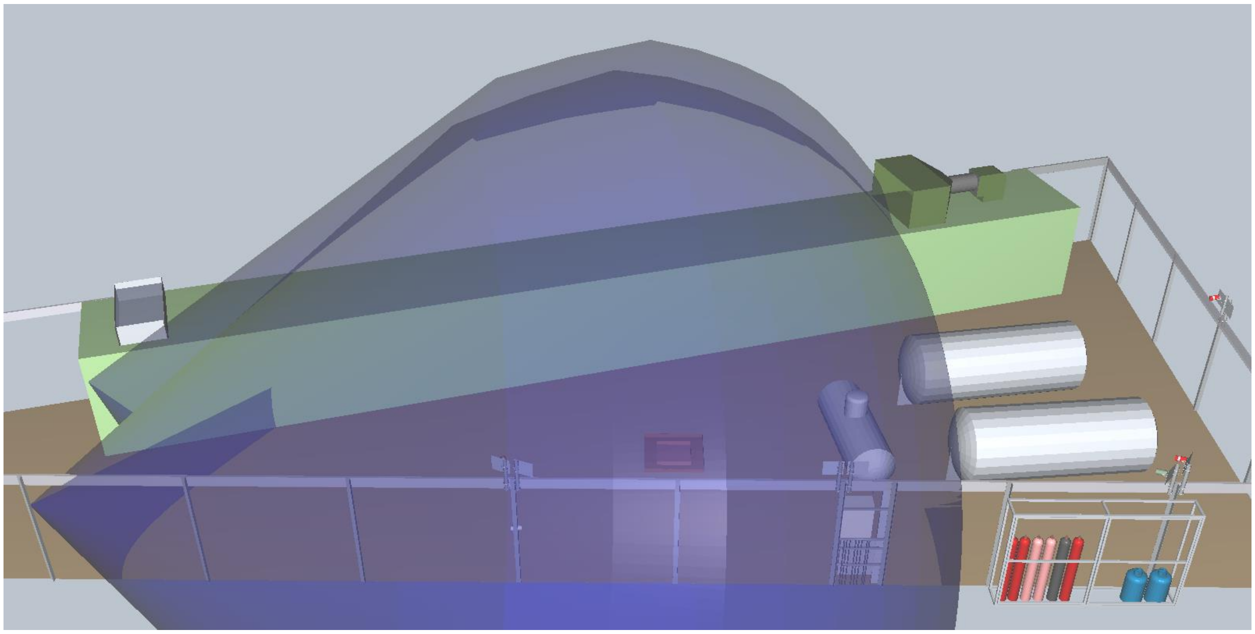

Figure 4: Image depicting the location and orientation of two detectors relative to the three vessels.

These vessels may be graded as HR, MR, or LR, as shown below.

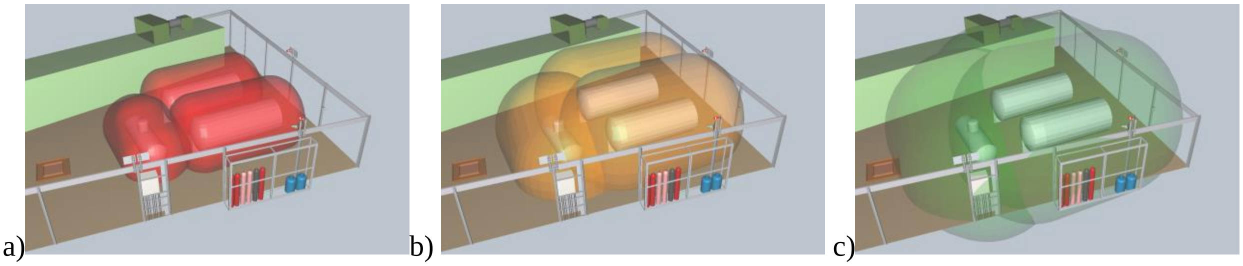

a) b) c)

Figure 5: Images depicting the three vessels graded a) "high risk," b) "medium risk," and c) "low risk."

The operator might set target fire sizes for alarm and control action as outlined in the table below with two voting detectors required for control action.

Table 4: Grade definitions for the geographic coverage assessments

Grade	Grade Color	Fire Size		Votes Required	
		Alarm	Control Action	Alarm	Control Action
HR	🟥	10	10	1	2
MR	🟧	10	50	1	2
LR	🟩	100	250	1	2

Figure 6 and Table 5 below summarize the results obtained using the two detector layout in Figure 4 with the three different grading schemes shown in Figure 5.

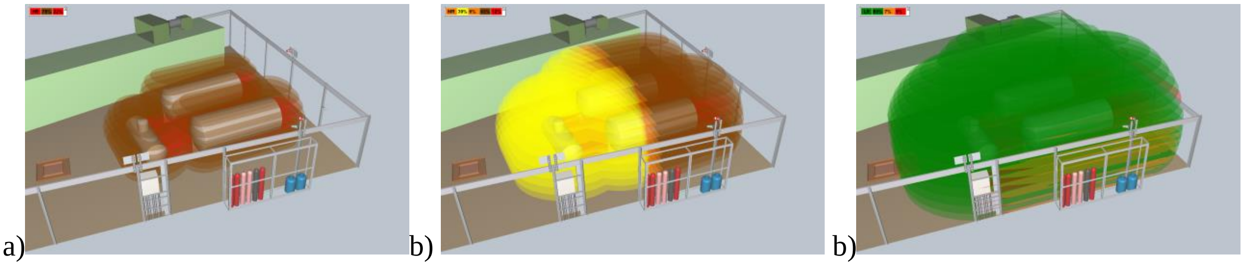
a) b) b)

Figure 6: Results of the geographic coverage assesments where-in all three vessels are graded a) “high risk,” b) “medium risk,” and c) “low risk.”

Table 5: Summary of results for the geographic coverage assesments conducted with the three vessels.

	Coverage Sufficient for Control Action	Coverage Sufficient for Delayed Control Action	Coverage Sufficient for Alarm	Coverage Sufficient for Delayed Alarm	No Coverage
All HR	-	-	-	78%	22%
All MR	-	39%	6%	43%	12%
All LR	85%	-	7%	-	9%

The results clearly show the impact that risk-based performance targets and target fire sizes can have on the suitability or adequacy of a proposed layout. The detectors are easily able to provide coverage for LR assets at this range but cannot provide any coverage for HR assets at the same distance.

This would, again, further complicate any attempt to generate an optimization algorithm. It may prove simple enough for an algorithm to consider a single, variable target fire size and a single effective detection range for a detector. However, would such a program be able to deal with several different target fire sizes and effective ranges simultaneously? This could occur in an area where there are pieces of equipment with different assigned risk grades, as shown in Figure 7, or it could result from an operator using more than one make and model of detector with different effective ranges, as sometimes happens.

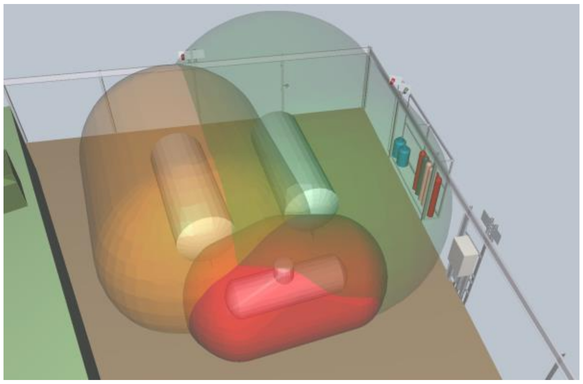

Figure 7: Image depicting an alternative grading scheme where-in all three vessels have been assigned different risk grades.

Conclusion

The authors have presented and reviewed a list of practical considerations involved in the design of fire and gas detection systems - particularly flame detection. These considerations, which include rules of thumb and recommendations that are flexible and not always tightly enforced in the design process, are difficult, if not impossible, to model using current technology and algorithms. A computerized design optimization algorithm is not going to understand or have an awareness of the reasons these rules or guidelines exist and will not be able to exercise judgement in deciding when it is okay to bend or break these guidelines in the way that a human engineer can. At the present time and for the foreseeable future, computer-based sign using optimization algorithms cannot and should not be used as a substitute for the nuanced judgment of an experienced fire protection engineer. Machine learning and artificial intelligence systems may eventually enable a computer to approach the level of judgement required to allow for fully computer-generated flame detector layouts, but such technology is not available today.

References

1. ISA TR 84.00.07. Guidance on the Evaluation of Fire and Gas System Effectiveness.
2. Benavides-Serrano, AJ; MS Mannan, CD Laird. “A quantitative assessment on the placement practices of gas detectors in the process industries.” Journal of Loss Prevention in the Process Industries. 35 (2015) 339-351.