



AACE International Recommended Practice No. 41R-08

**RISK ANALYSIS AND CONTINGENCY DETERMINATION
USING RANGE ESTIMATING**
TCM Framework: 7.6 – Risk Management

Acknowledgments:

Dr. Kenneth K. Humphreys, PE CCE (Author)
Kevin M. Curran
Michael W. Curran
Christopher O. Gruber, CCC
John K. Hollmann, PE CCE CEP

William E. Maddex, CEP
Stephen E. Mueller, CCE EVP
Dr. Shekhar S. Patil
Robert F. Wells, CEP
John G. Zhao

RISK ANALYSIS AND CONTINGENCY DETERMINATION USING RANGE ESTIMATING

TCM Framework: 7.6 – Risk Management



June 25, 2008

INTRODUCTION

Scope

This Recommended Practice (RP) of AACE International describes the process known as range estimating, a methodology to determine the probability of a cost overrun (or profit underrun) for any level of estimate and determine the required contingency needed in the estimate to achieve any desired level of confidence. The process uses range estimating and Monte Carlo analysis techniques (as defined in RP 10S-90). The RP provides the necessary guidelines for properly applying range estimating and Monte Carlo analysis to determine probabilities and contingency in a reliable manner using any of a number of commercially available risk analysis software packages.

The RP does not recommend any particular software. Rather it describes the factors that the analyst must consider when using risk analysis software for probability and contingency determination.

Purpose

This RP is intended to provide guidelines (i.e., not a standard) for risk analysis using range estimating that most practitioners would consider to be a good practice that can be relied on and that they would recommend be considered for use where applicable.

This RP is also intended to improve communication as to what the practice called “range estimating” is. Many of the methods found in industry that are being called this are not in accordance with this RP. Practitioners should always make sure that when someone uses the term “range estimating”, that they are talking about the same practice recommended here.

Background

This RP is new. It is based upon the successful efforts of many companies to evaluate project risk and contingency using the range estimating techniques originally developed by Michael W. Curran^[1,2,3]. Users should be aware that the principles outlined in this RP must be rigorously followed in order to achieve the desired results. Failure to follow the RP’s recommendations will likely lead to significant misstatements of risk and opportunities and of the amount of required contingency. In the great majority of cases, contingency and bottom line uncertainty are understated when the RP’s recommendations are not followed.

It is AACE’s recommended practice that whenever the term “*risk*” is used, that the term’s meaning be clearly defined for the purposes at hand. In range estimating practice as described in this RP, *risk* means “*an undesirable potential outcome and/or its probability of occurrence*”, i.e. “*downside uncertainty (a.k.a. threats)*.” *Opportunity*, on the other hand is “*a desirable potential outcome and/or its probability of occurrence*”, i.e. “*upside uncertainty*.” The range estimating process for risk analysis quantifies the impact of *uncertainty*, i.e. “*risks + opportunities*”.

RECOMMENDED PRACTICE

Range Estimating

Range estimating is a risk analysis technology that combines Monte Carlo sampling, a focus on the few critical items, and heuristics (rules of thumb) to rank critical risks and opportunities. This approach is used to establish the range of the total project estimate and to define how contingency should be allocated



June 25, 2008

among the critical items (RP 10S-90). It must be understood that total project estimate does not necessarily mean a cost estimate. The range estimating technique is equally applicable to profitability analyses (e.g., return on investment, projected earnings, earnings per share). It is also applicable to schedule-risk applications provided that the ranges determined for the critical schedule tasks do not result in a change in the critical path.

It must also be noted that the process applies to estimates that are based on a defined scope. Should scope changes be needed, or scope creep develops and results in significant changes in scope, the estimate upon which the range estimate is applied must be revised to reflect such scope changes.

An exception to this rule occurs when scope changes or modifications are anticipated and when the estimate includes a line item to cover such scope changes. In no event should contingency ever be treated as a source of funds to cover scope changes.

Identifying the Critical Items

The key to performing a project risk analysis using range estimating is to properly identify those items that can have a critical effect on the project outcome and in applying ranges to those items and only to those items. It is human nature to assume, for example, that a very large item in a cost estimate is critical simply because of its magnitude. That is not the case. An item is critical only if it can change enough to have a significant effect on the bottom line. The effect need not be negative (unfavorable). What matters is its degree, either in the negative or the positive direction.

Curran^[2] has demonstrated that in virtually all project estimates the uncertainty is concentrated in a select number of critical items -- typically 20 or less. Very few things are really important. This is called variously the *Law of the Significant Few and the Insignificant Many* or the *80/20 Rule*. Others refer to it as *Pareto's Law* after the noted Italian sociologist and economist, Vilfredo Pareto. On rare occasions there may be more than 20 critical items or less than 10. If this occurs, the risk analyst should carefully reexamine the items to be certain that the critical ones have been properly identified.

A critical item is one whose actual value can vary from its target, either favorably or unfavorably, by such a magnitude that the bottom line cost (or profit) of the project would change by an amount greater than its critical variance. The bottom line's critical variance is determined from the following table:

Bottom Line Critical Variances		
Bottom Line (Cost or Profit)	Conceptual Estimates (AACE Classes 3, 4, 5)	Detailed Estimates (AACE Classes 1, 2)
Cost Δ	± 0.5%	± 0.2%
Profit Δ	± 5.0%	± 2.0%

Table 1 – Bottom Line Critical Variances

Critical items are those which can cause changes greater than the above Δs (critical variances), either in the negative or positive direction.

It is important to link or combine items that are strongly related (i.e., when one item increases or decreases, the linked item also changes either directly or inversely). Such dependencies are generally obvious or become quite apparent during the process of evaluating the critical items and establishing their ranges. As an example, if the cost of concrete is the major cost driver in more than one estimate item and concrete costs may vary over a critical range, those items for which concrete cost is the major cost driver are dependent and must be combined in the analysis.

While the above Δs may seem to be rather small to some observers, they have been proven valid on many thousands of projects and it generally is not recommended that larger values be assumed. The only exception would be if significantly more than 20 items are shown to be critical. In such a case, in keeping

with Pareto's Law, it is useful to apply larger Δ s to reduce the number of items to those which are most critical. In any event, limit any such changes to Δ s with values no larger than twice the values shown in the table.

It is very important to understand that the magnitude of an item is not important. What is important is the effect of a change in the item on the bottom line. Relatively small items are often critical while very large ones may not be critical at all. Typically, there will be only 10 to 20 critical items, even in the largest projects with hundreds or thousands of components to consider. In identifying the critical items, it is necessary to link strongly related items together, not to treat them separately.

It is also necessary in the range estimate to apply ranges only to the items which are identified as being critical. The project team must know when an item is important and when it is not. If non-critical items are ranged, the inevitable result will be a far narrower predicted range of possible project costs than actually exists, misstatements of risk and opportunity, and understatement of required contingency.

Identifying the Risks

Prior to estimating contingency or otherwise quantifying risk impacts, the risks must first be identified. This RP does not cover the methods of identifying and screening risks as described in the Total Cost Management Framework risk management process^[7]. However, in today's project management practice a "risk register" is normally established to highlight potential project risks of any type which might significantly affect the project. When determining ranges (as covered in the following section), the project team must assure that these risks are all considered in the ranges.

Determining the Ranges

To establish the ranges, the project team, including owner(s) and the contractor(s) if appropriate, determines the ranges for the critical items, and only the critical items, based upon their experience and knowledge of the project and its risks using any available databases and/or benchmarking information. Generally speaking, everyone with significant knowledge and experience about the project's items and its risks should be involved in the process. Each estimate item is a single point number that is highly likely to vary in actual practice. The project team must examine each critical item and predict its possible extreme values considering all risks, including compounding effects. It is important to understand that the range, as considered in this method, is not the expected accuracy of each item. This is a key issue. Risk analysis is not an analysis of estimate accuracy. Accuracy is dependent upon estimate deliverables and estimate maturity. Contingency, as determined via the use of risk analysis, is not a measure of estimate accuracy. Rather it is a reflection of risk at any specified or desired probability of not completing the project within the estimate. You might reasonably expect a given estimate number to be accurate say, within -10% to +20%. That is not the range needed. The range is what you don't normally expect - the extremes which could happen, not what you expect to happen. If it can happen, it must be considered. Effectively, the extremes must be predicted at a probability range of 98+% (i.e., P1/P99) without making them so high or so low as to be absurd. The extremes should not include events that would be considered as far out of scope, such as "Acts of God" or funding cuts. Using ranges which are ridiculously wide (i.e., considering very rare adverse occurrences) will lead to overstatement of risk, opportunity, or both, as will linking items which are not actually strongly correlated with each other.

Similarly, choosing narrower probability ranges (e.g., P10/P90 or P5/P95) can lead to misstatement of the level of risk. The ranges, as stated, must reflect all reasonable possibilities of occurrence short of major disasters such as Acts of God.

It is not important how the estimate is structured, for example according to a work breakdown structure (WBS), areas and units, a cost breakdown structure (CBS), disciplines or assemblies. Any consistent estimate basis may be used in determining the critical items and their ranges. No matter how the estimate is structured, the project team works down through it in order to determine which items are critical and to establish their ranges and probabilities. In order to be successful, the project team must have a full

knowledge of the project scope, the project objectives and the project plans as well as all project documentation as they go through this process. They must know the full estimate basis and, if any information is missing, they should ask for it before beginning the process.

Obviously, when the project team is examining the critical items and determining the ranges, if it is discovered that the estimating, forecasting, control or other practices being used for the project are not industry best practice, if possible corrections to the estimate and project plan should be made. If this is not possible, the added risk of using poor practices must be considered in the ranges which are established for the analysis.

Note also that extremely rare risk driven events (i.e., "Acts of God") are not part of the analysis. They and their effects are unpredictable. They are outside of the P1/P99 range. However, highly possible but unlikely events such as severe weather must be considered.

As an example, in a hurricane prone area, a hurricane might normally be expected to occur once every ten years. While not expected during a project of a few years duration, a hurricane is a reasonable possibility and the ranges should reflect the possible effects of a hurricane. However, a 100-year storm of extreme magnitude which causes massive destruction such as Hurricane Katrina in New Orleans is extremely unlikely and the possibility is not a reasonable consideration. To include such an extreme event in the analysis would lead to a probable significant overstatement of risk and of the amount of required contingency to mitigate that risk.

When the ranges are being determined by the project team, they must also make an estimate of the probability that each critical item can be completed within the estimate (e.g., do they feel that there is a 50% probability of the critical item not overrunning?... an 80% probability?). These probability estimates need not be precise figures to the nearest percent. It is sufficient to judge the probabilities within increments of 5%. Estimators generally present estimates which strive for an equal probability of overrun or underrun for each item, i.e. the most likely or modal (and mean) values. In the process of establishing the ranges, the project team will develop strong feelings as to how reasonable each estimate item really is. This in turn leads rather readily to a quantitative judgment of the probability that the particular item can be accomplished within the estimated value.

In establishing the ranges and determining the critical items, estimates should be examined "top down", not "bottom up". The first level of the breakdown, whatever its structure may be, is examined for critical extremes. Those shown to be critical are then examined at the next level down in order to determine which components of the item are causing the criticality, and so forth as necessary down through the estimate until the specific factors which lead to criticality are identified. If at one level an item is identified as being critical and examining its components does not identify anything as critical, then the first item is the critical factor for the analysis.

At this point, the project team will have the required information for conducting the risk analysis. Specifically, for each critical item, they will have:

- its estimated value
- the probability that its actual value will **not** exceed its estimated value
- its maximum possible value
- its minimum possible value

The probability should be elicited from the project team **before** the maximum and minimum possible values. This helps to preclude the team's usually incorrect assumption that the probability represents the relative proportion of where the estimated value is located within the range.

The estimated value may be less than the minimum possible value or greater than the maximum possible value. This can occur if the current state of knowledge is radically different than what it was at the time the estimated value was chosen and, for whatever reason, the estimated value cannot be revised to

reflect current knowledge. This scenario frequently occurs when a risk analysis is performed for the remainder of the life cycle of a project in execution. The current assessment of a critical item can be so pessimistic or optimistic that the entire range for that item is above or below its estimated value, in which case the probability of not exceeding its estimated value is 0 or 1, respectively.

This information becomes the input to the range estimate. All non-critical items are “backed out” and entered into the analysis as a fixed sum.

The effort to gather correct (i.e., reliable) information on critical items should not be underestimated. Based on experience gained from thousands of range estimating sessions, the necessary level of interrogation and discussion typically requires six to eight hours, even for small projects or budgets,

Probability Density Functions

Monte Carlo software for risk analysis requires identification of a probability density function (PDF) for each critical item. Not all values in a range are likely to have an equal probability of occurrence and this is reflected by an appropriate PDF. In rare instances the behavior of a critical item is **known** to conform to a specific type of PDF such as a lognormal or beta distribution, which reflects items that may skew heavily to one side of a distribution. If a distribution such as this is appropriate for the item in question, the item should be represented by that PDF. However, it generally is unlikely that the actual type of PDF that truly represents the item is known. Thus, a reasonable approximation is to use one of two distributions:

- the triangular distribution
- the double triangular distribution

In most cases, the double triangular distribution is a better approximation since it can be made to conform to the implicit skew of the project team’s probability assessment. The double triangle allows the risk analyst to use the probabilities which the project team believes are reasonable rather than letting the triangular distribution dictate a probability which, more often than not, is invalid.

While these distributions are approximations, within the anticipated level of accuracy of any estimate, Class 1 through Class 5, they are sufficiently accurate for the purposes of the risk analysis.

The Triangular Distribution

A triangular distribution looks like this:

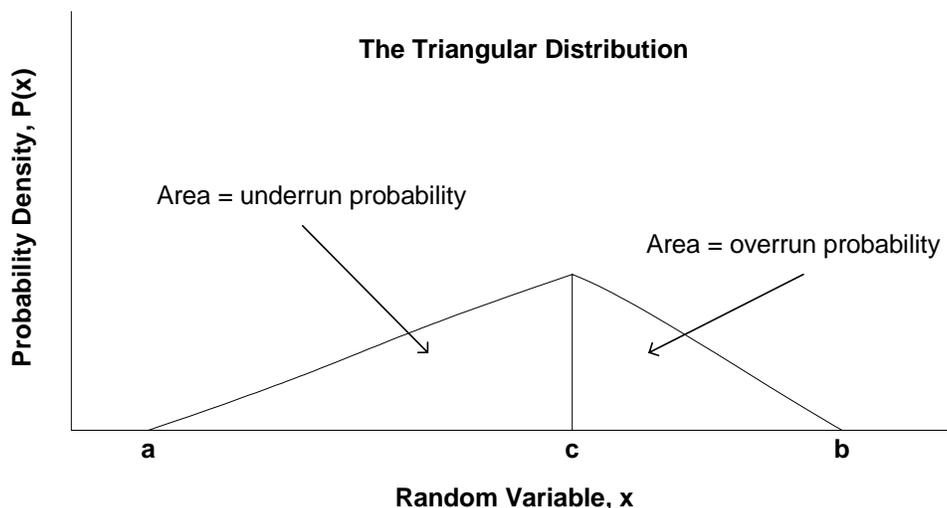


Figure 1 – Triangular Distribution

A common error is to assign the triangular distribution without verifying that it actually applies. As with any PDF, the range implies a probability in a triangular distribution. In the triangle, the areas of the two sides of the triangle to the left and to the right of the estimated value are proportional to the probability of a value being greater than or less than the estimate. If, for example, the range is 2000 (a) to 5000 (b) with an estimate of 4000 (c), the probability of being under the estimate is $(4000-2000)/(5000-2000)$ or 66.7%. If the project team believes that this implied probability is unrealistic (it is, more often than not), the triangular distribution will not be satisfactory. Far more commonly, the double triangular distribution is the more realistic choice.

The Double Triangular Distribution

The double triangular distribution looks like this:

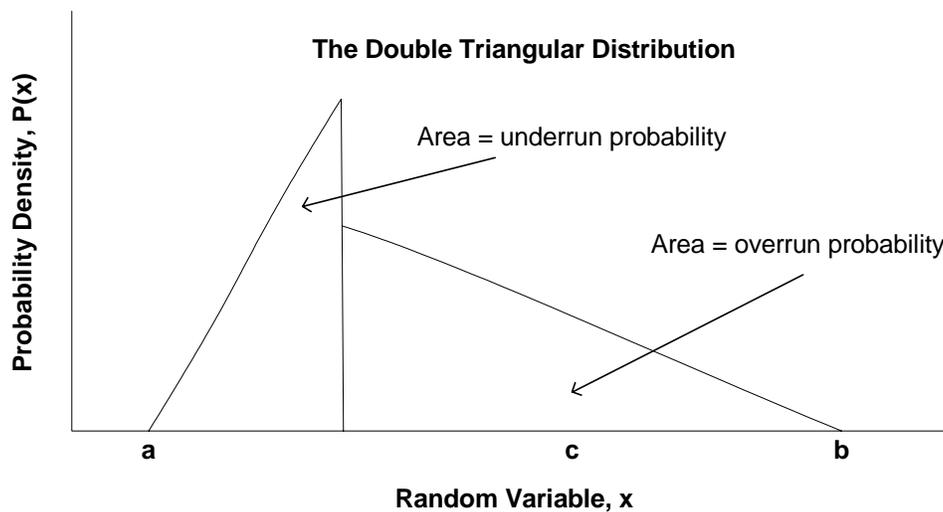


Figure 2 – Double Triangular Distribution

The double triangular distribution should be applied for any critical item for which the single triangular implied probability of not overrunning differs from the probability assessed by the project team. Assuming that the ranges properly reflect the possible extremes of cost or profit, this will generally be the case.

In the commercially available software, the double triangular distribution is not defined. It must therefore be entered into the software as a custom distribution (or as two right triangles using an "if, then" criterion to select between the two triangles depending upon the value of the randomly selected values). These options are generally available in commercial software so the specification of a custom distribution is not an obstacle to the use of risk analysis software.

The nomenclature for the two distributions above is:

- a = the minimum value of the range for the item in question
- b = the maximum value of the range for the item in question
- c = the estimated value of the item in question, i.e. what is believed to be the most likely value (the mode)
- x = a randomly selected value for the item in question
- P(x) = the probability associated with the random value of x

It must be noted that the double-triangle is really two triangular distributions, one representing values which underrun the estimate and the other, values which overrun the estimate. At the estimate value (c), the distribution is discontinuous and some software may not be able to handle a discontinuous function

such as this. In that case, an "if-then" scenario is used with the software to select the appropriate triangle for each randomly selected probability value $P(x)$, depending upon whether it is above or below the probability for c .

Contingency Determination and Probability of Overrun

The various software packages, using the inputs as described, will generate a cumulative probability distribution curve for the complete estimate. Typically, between 300 and 800 iterations are necessary in order to obtain statistically significant results using these software packages. It is recommended that 1000 iterations be used as this will, with rare exceptions, be a large enough sample for reliable results. (The rare exception is most often precipitated by abstruse discontinuous relationships amongst the critical items or in cases where one or two critical items have large and highly skewed ranges.) The curve will show the probability of the estimated cost or profit being achieved. If the probability is less than desired, the required contingency to bring the estimate to the desired probability of not overrunning is the difference in the curve's cost or profit value at the desired probability minus the value at the actual probability. Most of the software packages provide this information in tabular form as well.

It is generally better to specify the desired probability when inputting data to the software. The results will then state the contingency required to achieve the specified probability. The selection of desired probability depends upon the risk attitude of management. A good estimate should have equal probability of overrun and underrun (i.e., a 50% probability). This is a risk neutral approach, the assumption being that some projects will overrun while others will underrun and, in the long run, they will balance out.

The more conservative, risk-averse attitude used by many profit-making companies, is to specify a probability of 80% or higher that the project will not overrun. This is a safer route but by specifying a high probability, the required contingency will increase and with it the project cost. This results in a maldistribution of funds. Large contingencies on projects in the organization's project portfolio will sequester monies that could otherwise be put to productive use (e.g., funding additional projects, beefing up R&D, investing in product improvement, new equipment). Excess contingency should be released from the project as remaining work becomes smaller and the risk decreases as indicated by periodic risk analyses. This is especially important in organizations where project teams are permitted to increase scope when excess contingency is available (i.e., when project funds are expended simply because they are not required to meet the original scope).

It must be noted that the contingency which is determined is total required contingency. It does not reflect what is sometimes called "management reserve," a discretionary amount which is added to the estimate for possible scope changes or unknown future events which cannot be anticipated by the project team unless an allowance for this purpose has specifically been included in the estimate as a line item.

Before accepting the final contingency amount, the team should review the analysis and its documentation to ensure quality and effective communication of the work and its outcomes. This should include comparing the result to empirical metrics, past results and expert expectations. If the result is out of line with past experience, this does not mean it is invalid, but that it should be checked to ensure there were no errors or omissions.

All of the commercially available software will generate what is commonly known as a "tornado diagram," a graphic and ranked representation of the risks and opportunities generated by the critical items. The diagram ranks the items from that with the greatest potential effect to the one with the smallest potential effect. Some software allocates the contingency to these items in proportion to their potential effects, thus enabling contingency to be managed based upon a draw-down plan. In those cases where the software does not allocate the contingency, the contingency per item can be inferred from the tornado diagram.

For risk management purposes, contingency can be assigned to the critical items based upon their relative potential to contribute to cost variances. However, contingency should never be included in the control budget for a critical item - contingency is its own budgeted control account except in the case of

estimates which are prepared for line item bidding. Risks are dynamic and periodic range estimates should be conducted to reassess project risks and to reassign and/or release contingency as appropriate throughout the project's life.

Limitations of the Technique

Note that the method described in this RP has been used successfully for estimates of all types. However, empirically-based modeling methods often require less effort to apply in the case of Class 5 estimates for which there are only a few critical items (e.g., less than five) and for those for which the content is poorly structured, defined and understood by the project team. In such cases empirical methods (see Hackney^[5]) should be considered for such early estimates, optimally in conjunction with range estimating.

In a perfect world, the items in a project estimate truly are the most likely values and thus represent modal values. If such is the case, the statistics inherent in the range estimating process would be precise. Unfortunately the world is not perfect, nor are estimates.

Ideally, the range estimating process should be used to examine the estimate before commitments are made and, if the project team concludes that any critical project item has an unacceptable level of probability, i.e. that the estimate is not really the most likely value, the estimate should be refined to correct the problem. This is always the recommended approach to preparing estimates and conducting risk analyses and it will assist greatly in facilitating the development of a quality estimate.

However, this is often not possible. The project may have been approved at a budget level other than the estimated value, or the project may be well underway. If budget changes are not permitted, or if additional funds are not available, as is often the case, the option to revise the budget does not exist. The estimate is not as good as it should be but it is fixed and the risk analyst must work with it, no matter how bad some of the numbers may be. The range estimate nevertheless will generally yield reasonable results which are well within levels of error which most managers are willing to accept. Some items may deviate widely from the estimate. The values then may not actually be the most likely values and hence would no longer approximate the mode which statistics would decree as the proper input to the PDF which is used. In this case the distributions may not yield a reasonably precise number for the item in question. However, that is unlikely to be the case for all of the critical items and the resultant predicted PDF and contingency for the total project will generally still be reliable within acceptable levels. Effectively, the idea that the plusses will equal the minuses generally applies.

Another situation is when the budget is fixed and subsequent range estimates indicate that a critical item is now totally out of the predicted range. That can be handled by increasing the value of the critical item to the minimum value of the range, assigning a zero percent probability of being under the estimate, and to correct for the addition to the item, to enter a non-critical contingency correction of negative the amount added to the item. That retains the total equal to the budgeted amount which permitting the analysis to proceed.

When to Apply Risk Analysis

Range estimates are dynamic, not static, and should be applied regularly through all phases of project design and construction. As the estimates go from an order of magnitude Class 5 estimate to a detailed Class 1 or 2 estimate, range estimates should be conducted to refine the contingency number. Then, as the project progresses, range estimates should be done at least quarterly to track the use of contingency and to reflect project progress. This enables contingency to be released when it is no longer needed. Contingency should never be held until project completion. Periodic range estimates will indicate when contingency can and should be released. Periodic range estimates will also highlight trouble areas which have developed, or are developing, which may require corrective action and/or a revision of the project budget.

REFERENCES

1. Curran, Michael W., *Range Estimating: Coping With Uncertainty*, AACE Transactions, AACE International, Morgantown, WV, USA, 1976
2. Curran, Michael W., *Range Estimating: Measuring Uncertainty and Reasoning with Risk*, Cost Engineering, Vol. 31, No. 03, AACE International, Morgantown, WV, USA, 1989
3. Curran, Michael W., *Range Estimating: Contingencies with Confidence*, AACE Transactions, AACE International, Morgantown, WV, USA, 1989
4. Brienzo, Kenneth D., Editor, *AACE International's Professional Practice Guide to Risk, 2nd Edition*, AACE International, Morgantown, WV, USA, 2007
5. Hackney, John W., *Control and Management of Capital Projects, 2nd Edition*, Kenneth K. Humphreys, Editor, AACE International, Morgantown, WV, USA, 1997
6. AACE International Recommended Practice 10S-90, *Cost Engineering Terminology* (latest revision), AACE International, Morgantown, WV, USA
7. Hollmann, John K., Editor, *Total Cost Management Framework: An Integrated Approach to Portfolio, Program and Project Management*, AACE International, Morgantown, WV, 2006.

CONTRIBUTORS

Dr. Kenneth K. Humphreys, PE CCE (Author)
Kevin M. Curran
Michael W. Curran
Christopher O. Gruber, CCC
John K. Hollmann, PE CCE CEP
William E. Maddex, CEP
Stephen E. Mueller, CCE EVP
Dr. Shekhar S. Patil
Robert F. Wells, CEP
John G. Zhao