TASSC

Circa technical briefing

Circa

Risk and Contingency Calculations

Gillian Adens, Gareth Tuckwell, Robert Armstrong

Tassc Limited
www.tassc-solutions.com

First Published: April 2002
Last Updated: February 2009

Circa provides a mechanism for specifying the risks associated with a software project. The user can identify and document a number of risks, each corresponding to a particular concern or issue in the project. Each risk is potentially a time-consuming problem that may delay the project and incur additional cost. Risk assessment is therefore an essential part of accurate estimating and planning.

The risk model

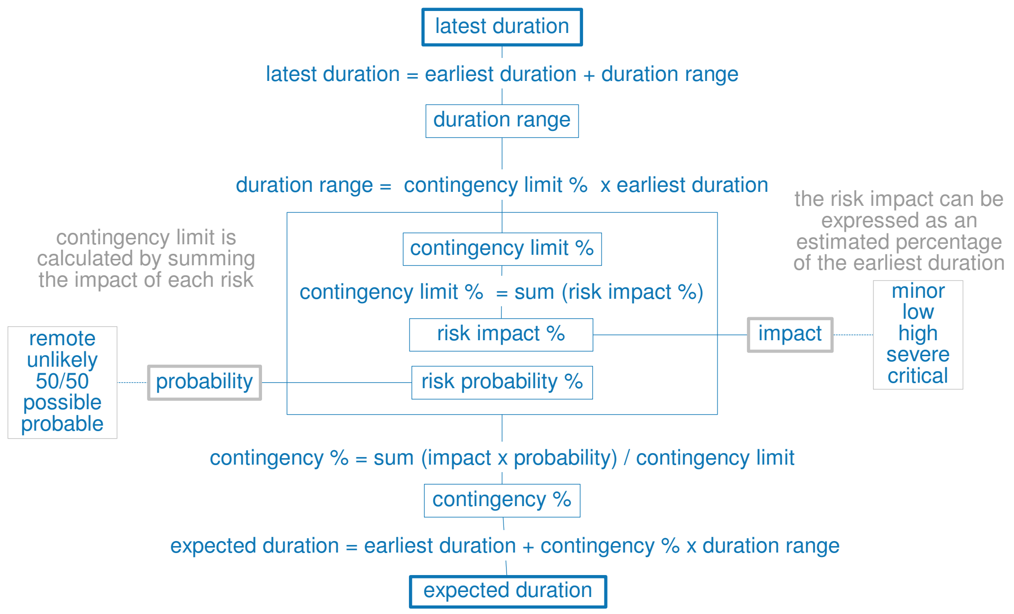

How do we quantify risks?

risks have an impact and probability

Risks are described in terms of both their impact and their probability. Each risk can be assigned a 'level of impact' and a 'probability of firing'.

impact indicates the amount of damage a risk will cause

The level of impact is the amount of damage the risk will have on the duration and cost of the project, should it fire. The impact of a risk can be recorded as minor, low, high, severe or critical. These logical labels are mapped to percentages within Circa ranging from 1% to 5%. Consequently we derive that a risk with severe impact has a potential of delaying the project by up to 4% of the overall project duration.

probability indicates the likelihood of a risk firing

The probability of firing indicates a judgement on the likelihood of the risk firing given the project circumstances. The user can select between remote, unlikely, 50/50, possible or probable. Again, these logical labels are mapped to a percentage within Circa, ranging from 15% to 85%. Consequently, we derive that a remote risk has a 15% probability of firing.

How do we calculate the impact of risks?

probability	15%	35%	50%	65%	85%	
impact	remote	unlikely	50/50	possible	probable	definite
minor	0.15%	0.35%	0.5%	0.65%	0.85%	1%
low	0.3%	0.7%	1%	1.3%	1.7%	2%
high	0.45%	1.05%	1.5%	1.95%	2.55%	3%
severe	0.6%	1.4%	2%	2.6%	3.4%	4%
critical	0.75%	1.75%	2.5%	3.25%	4.25%	5%

Contingency

the combination of probability and impact defines a risk's level of threat

The combination of probability and impact defines a risk's level of threat to the project. This is the crucial factor in calculating an appropriate contingency for the project. Therefore a serve risk is considered to have a potential impact of 4% of the project duration. If this risk will definitely fire we would want to add a 4% contingency to our project estimates. However, risks are never definite at the outset, they only become definite if they actually fire. During risk assessment we might predict that the probability of this risk is remote and then we can calculate a more appropriate contingency by taking 15% of the 4% resulting in a 0.6% contingency. This contingency is not just based on the impact; it also takes account of probability.

calculate a range of durations - earliest, latest and realistic

These calculations allow us to estimate a range of durations. The earliest possible duration is where no risks fire i.e. the original estimate based on scope and metrics. The latest possible duration, where all risks fire, is calculated by summing the potential impact of each risk and adding this percentage to the duration of the project. A third more realistic duration can be derived from summing the risk threats – this provides a sensible level of contingency based on the combination of risk impacts and probabilities.

What is the effect of risk status?

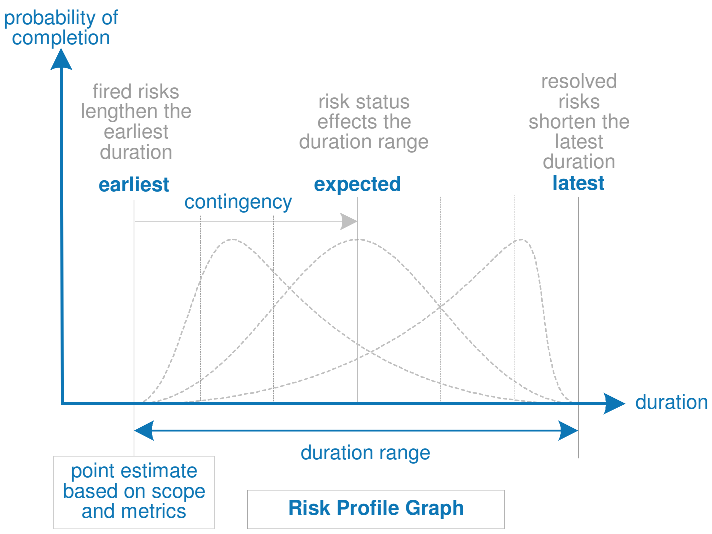

a risk can be considered to be in one of three states – active, inactive or fired

At the start of a project all potential risks are identified and quantified in terms of impact and probability. At this stage the risks are considered to be 'active'. If a risk is classified as active it may fire at any time.

As the project progresses some risks may actually fire. This effectively means that the project is delayed. The risk status is changed to 'fired' and the amount corresponding to its full potential impact is added to the project's earliest delivery date. The duration range reduces as the earliest date moves closer to the latest date.

If a risk is resolved without firing, the duration range is again reduced, but this time in a more positive way – the latest delivery date moves closer to the earliest delivery date. When risks are resolved they can be marked as 'inactive', and will no longer be counted. This effectively means that the latest delivery date is brought forward by an amount corresponding to the risk impact.

At the start of a project the range of possible estimates can be wide. As the project progresses and risks either fire or are resolved, this range tends to become narrower.

What is the meaning of risk priority?

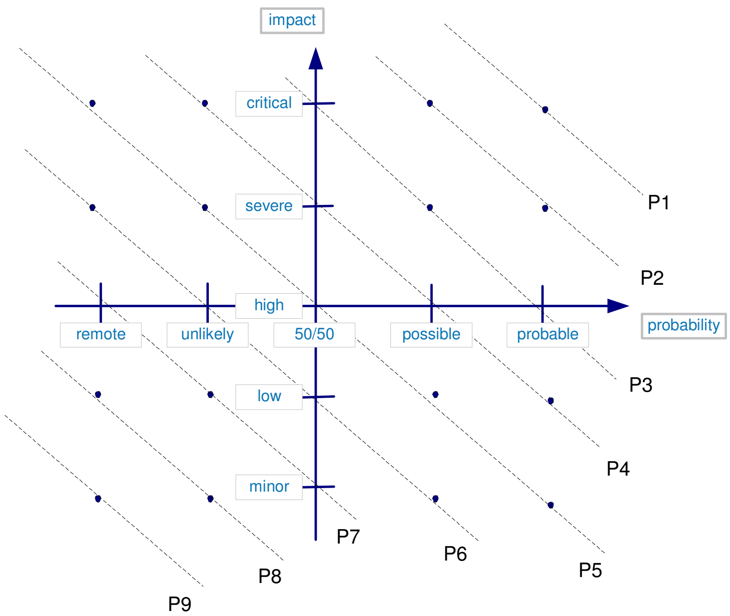

priority ranks the risks by their level of threat

Risk priority ranks the risks by their level of threat to the project. It is driven by the combination of risk impact and probability. It is essentially an indicator of the level of effort that should be put into resolving each risk. The high priority risks are those that we should actively try to resolve and that we need to develop contingency plans for, so that we have an alternative action or solution, should they fire.

plot the risk impacts against probabilities

A simple way of defining the risk priorities is to plot the risk impacts against probabilities, and draw priority lines diagonally through the graph. Those risks of highest probability and impact are of greater priority and those risks of least probability and impact are a lower priority.

What does the risk graph display?

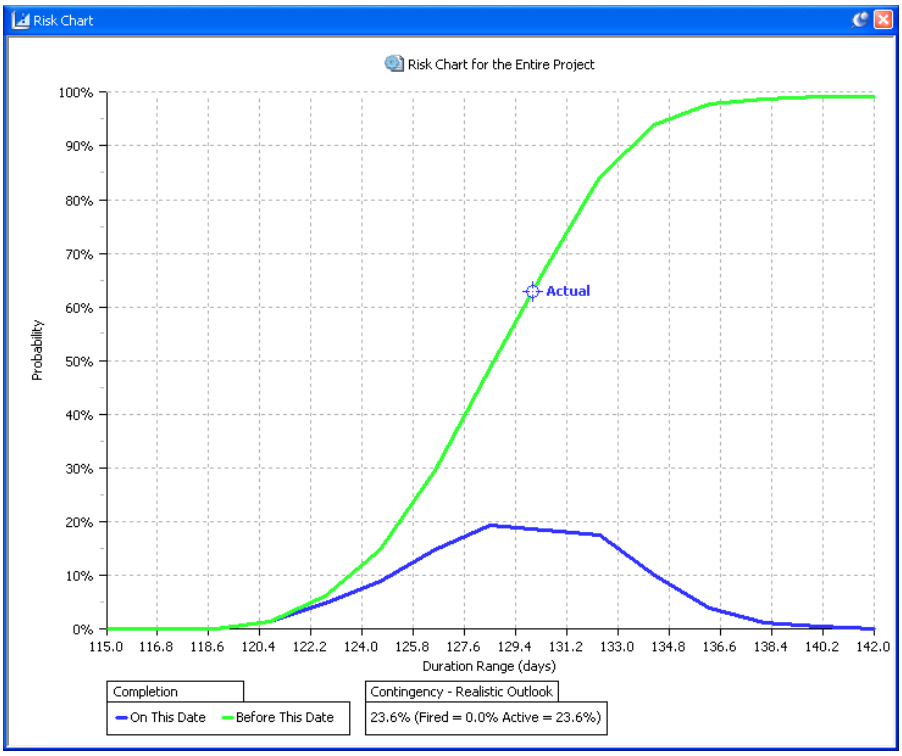

The risk data allows Circa to give a range of estimates for a project and an accuracy rating across that range, with the ultimate goal of calculating the amount of contingency (as a percentage) that should be applied to the project.

the risk graph applies a range of contingencies and plots the probability of correctness

In essence, the risk graph applies a range of project contingencies to the shortest possible duration and plots the probability of correctness. The 'probability of correctness' is a simple concept – how likely is it that any given contingency is actually going to be the correct value for this project. We plot this confidence as a probability percentage on the y-axis.

The x-axis on the graph depicts time (in days). This is the potential project duration ranging from the earliest possible completion duration to the latest possible duration. The earliest possible duration is the length of time the project will take should nothing happen to slow down the work – in other words, if no risks fire. The latest possible duration is the time the project will take if everything seen as a potential problem actually transpires – in other words, if all the risks fire.

Clearly it is most likely that the project will complete somewhere between these two bounds - some risks will fire, some will not. We also show an 'actual' expected duration based on a calculated contingency, which indicates the point that Circa considers to be the 'safest bet'.

How do we calculate the points on the risk graph?

contingency values map to the first, last and expected points on the graph

As explained earlier, a risk has a 'level of impact' and a 'probability of firing'. From this information we can derive three contingency values:

- Minimum contingency – no risks fire, so no impacts are counted (always 0%)
- Expected contingency – calculated mathematically from the probability and impact of all risks
- Maximum contingency – all risks fire, so all impacts are summed

These three contingency values map to the first, last and expected points on the graph. Effectively, giving us the range for the graph, plus one point along the curve. This point is the expected contingency and is calculated using a simple mathematical formula to find the most likely contingency.

we do not have enough information to manufacture a formula

We need to plot this information on the graph with the contingency applied to the duration on the x-axis and probability on the y-axis. The normal way of drawing a curve on a graph is by deriving a formula, which when applied to a value on the x-axis, calculates the corresponding y-axis value. We have the x-axis, or contingency range, so we need to apply a function f(x) to points along that range to derive their y-axis component.

For example: A project has a minimum contingency of 0%, expected contingency of 17% and a maximum contingency of 30%. What is the probability (y-coordinate) of 10% being the correct contingency (x-coordinate)?

Unfortunately, we do not have enough information here to manufacture a formula that can be applied to a contingency and calculate a probability of that contingency being correct!

What we do have is a collection of risks each with a probability of firing. Once the project has completed, 'some' of these risks will have fired and others will not. To calculate the contingency for the project, we need to determine which risks will fire.

randomly select which risks will fire to generate one sample

One way of doing this is by randomly selecting which risks will fire. The impacts from any risks that fire are summed to find a total impact for the project. This is easy to implement – we simply generate a random probability from 0 to 1 and add it to the probability of the risk (also 0 to 1). If the combined probabilities >= 1, then the risk can be considered to have fired. This means that risks defined with a larger probability will have a greater chance of being randomly selected and vice-versa.

repeat the procedure many times

This generates one possible contingency for the project, but not necessarily the correct contingency. As we repeat this procedure many times, we build up a sample set of potential contingencies for the project. Every result in the set will fall between the minimum and maximum contingency values.

divide the range of results into groups

We can then divide the range of results into groups of equal size. These groups will form the divisions along the x-axis.

For example: A project has a range of contingencies from minimum = 0% to maximum = 30%. These contingencies could be divided into 6 groups each of 5% size. All results falling between 0 < x < 5% will be in group 0. All results in the range 5% < x < 10% will be in group 1 and so on. We have now defined our x-axis… 0, 5, 10, 15, 20, 25, 30.

So we have generated a large set of sample project contingencies and divided the x-axis into groups. We must now count how many results fall into each group. Example: we count how many results fall into each group and find the following information:

- Group 0 - 0% < x < 5% = 0 results
- Group 1 - 5% < x < 10% = 200 results
- Group 2 - 10% < x < 15% = 400 results
- Group 3 - 15% < x < 20% = 500 results
- Group 4 - 20% < x < 25% = 600 results
- Group 5 - 25% < x < 30% = 250 results

To convert this data into actual probabilities, we simply divide the number of results in each group by the total number of results:

- Group 0 => 0 / 2000 = 0.0 probability
- Group 1 => 250 / 2000 = 0.125 probability
- Group 2 => 400 / 2000 = 0.2 probability
- Group 3 => 500 / 2000 = 0.25 probability
- Group 4 => 600 / 2000 = 0.3 probability
- Group 5 => 250 / 2000 = 0.125 probability

These probabilities will always total to 1 i.e. the project has a probability of 1 of completing in one of these groups. We can say that the project <u>will</u> complete in one of these groups. This is the information we plot on the graph. The groups form the divisions along the x-axis and the probability is the height – f(x) – of the graph at any given – x – contingency (duration).

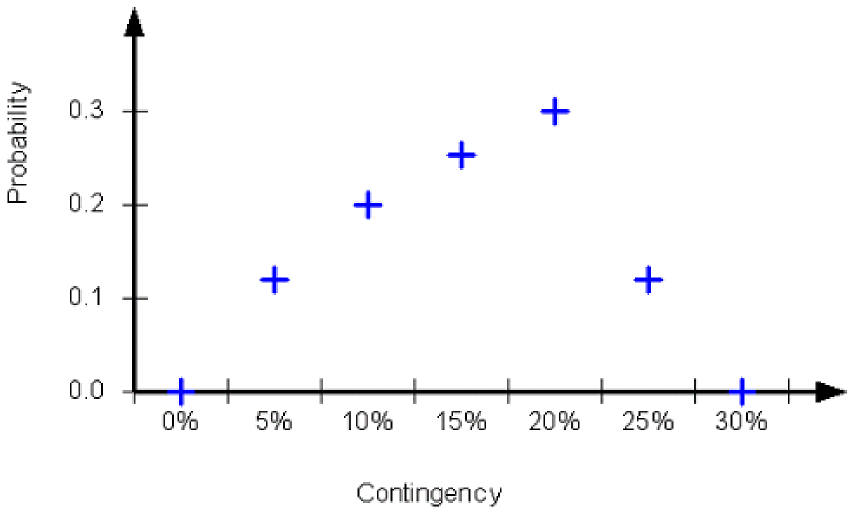

each point on the graph indicates the probability of any one duration being correct

Obviously, we are interested in project duration, so the x-axis would be scaled and labelled with duration to improve readability, the y-axis would be written as a percentage, but the shape of the graph would simply be a line drawn between these points. This solution produces an accurate and correct graph of the potential outcomes for a project.

Each point on the graph indicates the probability of any one duration being correct. This is a useful piece of information, but it would also be useful to know, what is the probability of the project finishing ‘before’ a certain duration. We can easily add another line on the graph to show this value. This second line simply displays the cumulative probabilities of all contingencies up to and including the point in question.