Optimizing Flip Chip Substrate Layout for Assembly

Pericles Kondos, Peter Borgesen, Dan Blass, and Antonio Prats
Universal Instruments Corporation
Binghamton, NY 13902-0825

Abstract

Programs have been developed to predict the expected yield of flip chip assemblies, based on substrate design and the statistics of actual manufactured boards (e.g. in pad sizes and locations, mask registration, etc.), as well as placement machine accuracy, variations in bump sizes, and possible substrate warpage. These predictions and the trends they reveal can be used to direct changes in design so that defect levels will fall below the acceptable limits. Shapes of joints are calculated analytically, or when this is not possible, numerically by means of a public domain program called Surface Evolver. The method is illustrated with an example involving the substrate for a flip chip BGA. It was found that the original design would lead to unacceptably high defect levels, but alternative designs significantly improved the yield without creating other significant problems like bridging.

Introduction

High-density flip chip applications are commonly limited by the available substrate technologies. Accordingly, considerable design efforts are focussed on the optimized routing of signal, power and ground connections. More often overlooked are some of the effects of the substrate surface layout on assembly yields.

In general, flip chip assembly tends to be rather sensitive to variations in the small dimensions of the solder bumps and contact pads, as well as to substrate warpage in reflow. Because of the large numbers involved, particularly in high density applications, manufacturing defect levels of concern are often related to the ‘tails’ of statistical distributions of these parameters. To make matters worse, the importance of even a relatively low assembly defect level is enhanced by the difficulty (at best) of a subsequent repair.

One type of assembly defect is proving sensitive, not only to the substrate tolerances but also to the accuracy of the placement machine. However, the risk of placing one of the solder bumps on top of the solder mask, or otherwise not in contact with the corresponding pad, can often be substantially reduced by the proper substrate technology selection and design.

Another potential source of defects is bridging or opens due to the combined effects of solder bump height variations, substrate pad size variations and substrate warpage in reflow. These effects depend, to some extent in contrasting fashions, on substrate technology and design, i.e. whether contact pads are mask- or pad-defined, as well as on pad size, thickness and shape.

The present work illustrates the use of available tools as applied to substrate design optimization and trade-off decisions for a particular flip chip BGA application, taking into account the constraints imposed by the available technology and cost. The overall effort involved quantification of the achievable bump height statistics, detailed modeling of the effects of pads shapes on relevant solder joint shapes, and the quantitative assessment of the sensitivity to statistical variations in the various parameters of concern.

The tools in discussion are three programs, two developed at Universal (the "yield" programs) and one in the public domain.

The Placement Yield Program

The Placement Yield Program (PYP) uses a Monte Carlo simulation technique to predict the defect levels due to in-plane variations in the substrate, and to placement machine inaccuracies. The user supplies statistics (either specs or results of actual measurements) of the relevant substrate parameters, including variations in mask registration, mask opening dimensions, fiducial and pad locations, pad and mask opening shapes and dimensions, etc. Also needed are the placement machine performance parameters, i.e. variations in x, y, and theta. Up to 4 different combinations of shapes can be used for the same die site. Essentially, the PYP assumes that a defect will occur when the center of a bump is placed on top of the solder mask or the laminate, rather than on the pad. The case when the center of the bump is not placed on top of the pad but a part of the bump still touches the side of the pad can be counted as a defect or not according to the user's wish. Experimentation is ongoing to determine when to use which defect criterion.

The most recent version also considers the possibility of a bump being prevented from touching its pad by the interference of the solder mask wall. A variety of configurations allow the solder mask to interfere, depending on whether the contact is with the sides or the edge of the mask and the pad. The placement yield program does not yet consider all of these, but as the case study below shows, for each configuration it is possible to calculate geometrically the maximum distance from the center of the pad that the center of the bump can be placed without the mask interfering. This distance defines an effective pad size, which can be subsequently used in the PYP in place of the actual pad dimensions. The program does not take into account variations in bump height, warpage, etc. These are handled by the Assembly Yield Program. Its output is an estimate of defect levels, with a breakdown of the results according to reason for failing to form a joint and location of failure.

Surface Evolver

The Surface Evolver is a public domain program that calculates equilibrium shapes for liquids. Extensive documentation explaining its use is given on the website of the program's author (Brakke 1999) so no details will be given here. Suffice it to say that the program can be used to calculate the shape of a solder joint of a given volume for given boundary conditions (contact pads).

In all the cases described in the present work, the assumption is that the liquid solder wets perfectly to the metal pads and not to any other surfaces, and that the solder joint shape is determined solely by minimization of the surface energy. Independent estimates have shown that gravity has no measurable influence for the dimensions of relevance to flip chip. The solder joint shape is therefore independent of the value of the surface tension and knowledge of the wetting angles is not necessary.

Surface Evolver divides every surface into triangles called "facets", formed by "edges" that connect "vertices". The accuracy of the results depends on how fine the mesh of triangles is; usually, a crude triangulation is used in the beginning and this is then successively refined (usually, but not always, this splits each facet into four smaller triangles). As the surface evolves, some facets shrink in size, while others grow. It is possible to instruct the program to split the longest edges, weed the smallest triangles, etc. All this can be performed automatically by using the appropriate script, but it can result in a very different number of facets from case to case. Although they tend to make the program more efficient, such optimizations were therefore avoided.

To estimate the accuracy of the program, comparison was made with a case where an analytical solution exists, namely when both die and substrate pads are circular. Without loss of generality, both can be assumed to be flat; if the substrate pad has a finite thickness, its volume is simply added to the solder volume. Under the assumptions mentioned above, the equilibrium shape is a truncated sphere (Figure 1). The height, h, is related to the volume of the joint and the diameters of the substrate pad, s, and the die pad, d, through the equation:

$$V = \frac{\pi}{6}\, h\left(\frac{3}{4}d^2 + \frac{3}{4}s^2 + h^2\right)$$

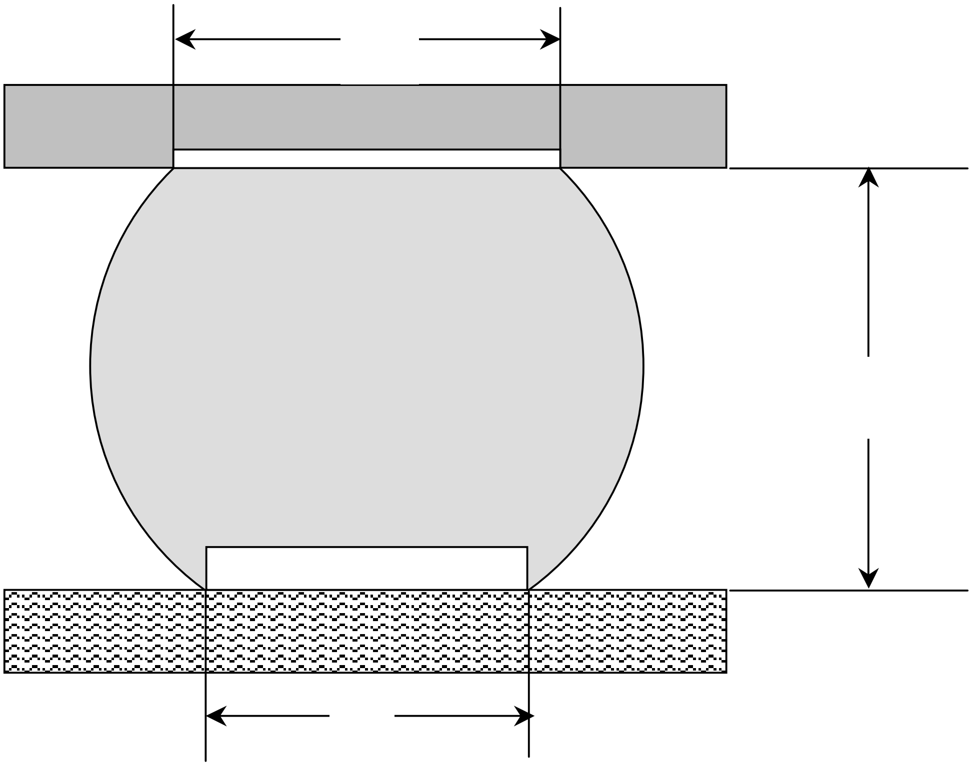

The volume of the solder was calculated from the die pad diameter and the height of the bump using the same equation with the substrate pad diameter, s, set to zero.

The following values were used for the calculation: Die pad diameter, 75 µm. Initial bump height, 100 µm. Substrate pad diameter, 50 µm. Substrate pad thickness, 15 µm. It was assumed that there was no hindrance by the solder mask, i.e. that the joint could take the equilibrium spherical shape. The following table summarizes the results.

Number of facets	**Distance from Substrate, h, µm**	**Free Surface Area, µm ^{2}**
80	99.74	36166.2
288	97.09	35279.9
1088	96.46	35063.5
4224	96.31	35009.7
Analytical Solution	96.25	34991.8

It is clear that the resulting height is close enough to the exact value for most applications, differing by approximately 0.8 µm after only two mesh refinements. However, in most cases reported here at least three refinements were used and occasionally a fourth. Running until no further change would take place after the fourth refinement would take too long, but after each refinement the free surface area was seen to fall quickly during the first few iterations. 100 iterations were therefore assumed to give the bulk of the improvement and these took only a minute or two to run. Figure 2 shows the appearance of the joint after the final (fourth) refinement, displayed at an angle to show the triangulation of the pad surfaces.

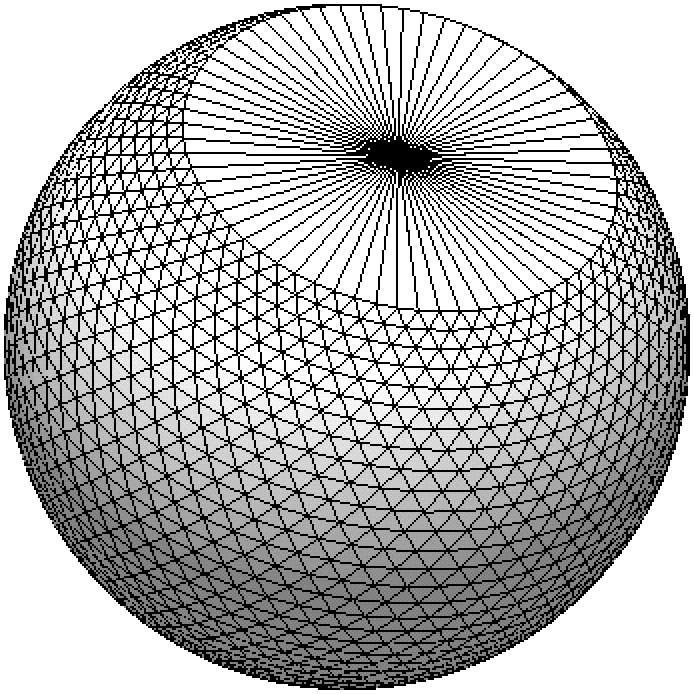

Assembly Yield Program

The Assembly Yield Program (AYP) is complimentary to the Placement Yield Program. The AYP considers out-of-plane variations, in bump height, pad thickness, solder paste volume (on the substrate) and warpage of both substrate and component. For flip chip the component warpage is, of course, negligible. In the current version of the program a "defect" occurs when a die has at least one bump that does not solder, because it never touches its pad even after the other bumps collapse and the die comes closer to the substrate. The assumption here is that as long as a bump touches its pad, it will wet it and form a joint. The risk of bridging proved negligible for the case considered but will be included in future versions of the program.

The current version of the AYP assumes circular pads and calculates the shape analytically. However, in combination with Surface Evolver, it can be used for almost arbitrarily shaped pads. For this purpose, once the average standoff has been calculated using Surface Evolver, an equivalent circular pad diameter giving the same standoff is used in the AYP. As the example shown below illustrates, the standoff is more sensitive to variations in the pad dimensions for circular pads than for the other shapes considered, so the AYP offered a conservative estimate of the defect level. A future version of the program will allow for the direct use of arbitrary pad shapes.

A Case Study

The following is an example illustrating the combined use of the three programs for the design of a flip chip BGA. The die was assumed to have 1000 bumps, arranged in a three-row perimeter array plus several rows in the inner region (Figure 3). The pitch is 150 μm, as is the distance between the perimeter rows. The die pad diameter is 75 μm, and the bump height 100 μm, with a standard deviation of 3.33 μm.

Originally, two different substrate designs were considered. The first had solder mask openings

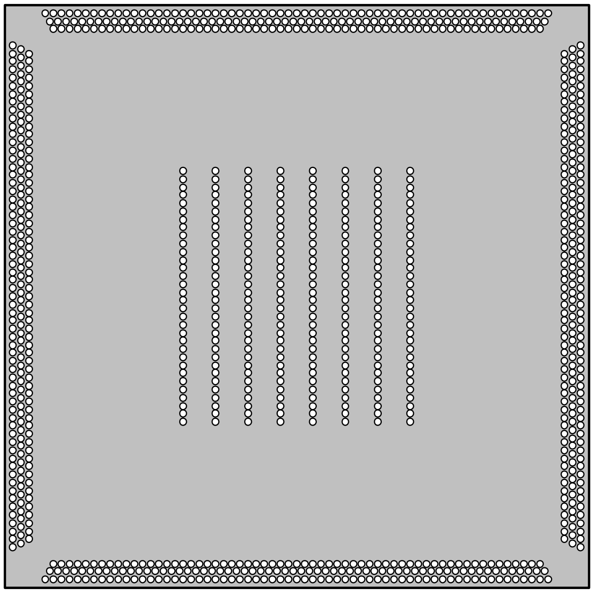

with nearly vertical walls and a nominal (average) diameter of 100 µm. The thickness of the solder mask was 25 µm. The pads were circular with a diameter of 50 µm at the bottom and a thickness of 15 µm and had 30 µm wide traces going through them. The pad side walls had a 70° taper.

A preliminary calculation, before even using the placement yield program, indicated that this design had potential problems because of the possibility of the bump being prevented from touching the pad by the solder mask, as illustrated in Figure 4. Simple geometric calculations show that the maximum allowable horizontal distance between the centers of the pad and the bump is about 16 µm in this case, giving a target (or effective pad) of just 32 µm in diameter. If the bump center is placed outside a circle with this diameter it will hit the solder mask first.

Such a small target is expected to lead to many defects. Indeed, using the placement yield program with pad diameter equal to 32 µm gives almost 2.7% defects just because of placement machine accuracy, without taking into account any effects of misregistration, variations in mask opening diameter, pad sizes etc., all of which increase the defect level. Typical machine parameters were used for this estimate, namely a standard deviation of 4 µm in x and y, and 0.033° in theta. Because of this, design No. 1 was rejected at the preliminary stage and was not considered any more.

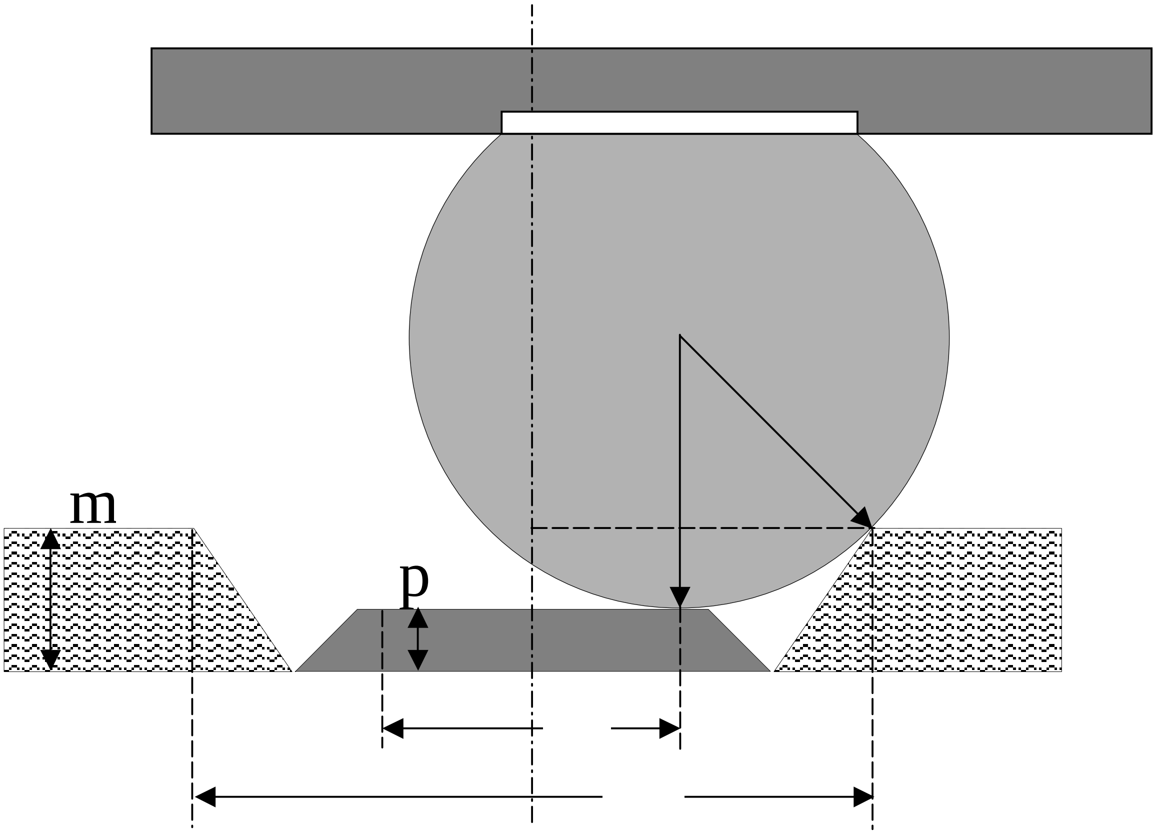

Figure 4. A Solder Bump That barely Touches the Pad Because of Mask Interference

The second design used a different method to generate the solder mask openings, and as a result the mask walls were sloping gently away from the pads, forming a "crater" of 100 μm diameter at the level of the top pad surface. The mask thus did not create any hindrance to placement, unlike the first design. In one of the cases considered the pads themselves were of the same shape as in the first design. Alternatively, the case where the circular pads were replaced with traces of width equal to 50 μm going through the circular mask openings was considered as well. Characterization of some substrates indicated that the pad diameter and trace width varied with $3\sigma = 10$ μm.

For the Surface Evolver, this alternative design of the pads was approximated with 50×100 μm^2 rectangles with vertical walls. Because of the gentle slope of the solder mask wall, it was assumed that the mask was not affecting the shape of the joint, except that it did not allow the solder to wet down the narrow sides of the traces. Of course, in reality the intersection of the circular mask opening with the pad sides is a complicated curve, but the error introduced by this approximation is small. Moreover it is practically the same over the range of the parameters used, so when two cases are compared, even this small error (common to both) tends to cancel out. It should be noted that Surface Evolver can handle complicated pad shapes; it is simply a matter of

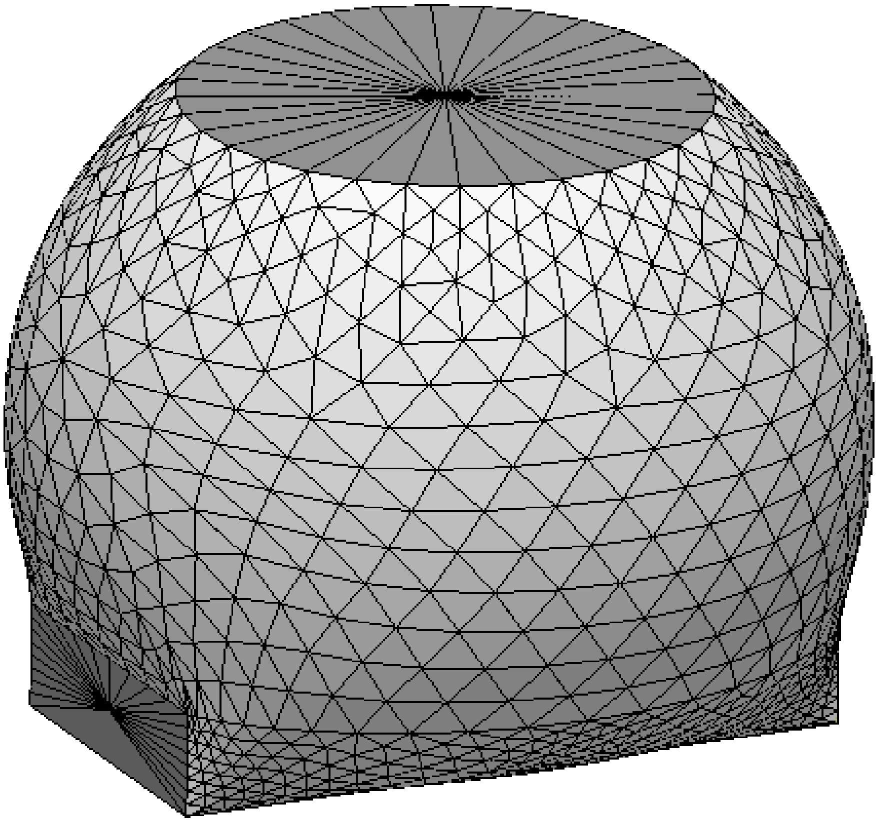

describing the shape properly. However, trying to match the shape of the pad would introduce a lot of programming complexity and increase the computing time, without offering a great improvement in accuracy. For this reason the simpler shapes were used in the present calculations.

The shape of the average joint is shown in Figure 5. In this case, the distance between the underside of the die and the top of the substrate pad came out to be 75.9 µm. Because of the assumption that as long as a bump touches its pad it will form a joint, the pad-to-die distance (PTDD) is the result of interest. The denser grid at the bottom part of the joint is due to the fact that two distinct vertical regions of unequal height were defined in the original approximation of the shape. Points lower than the top of the substrate pad had to be distinguished from points above it, and given different constraints, so that they would be prevented from moving through the pad. Three levels of refinement were used for this figure, leading to a total of 2198 facets. One extra level of refinement (8486 facets) improved the result by only 0.1 µm, to75.8 µm, indicating once more that three levels are sufficient for these shapes.

Figure 6 shows the (analytically calculated) PTDD vs. diameter for the same bump on circular pads. Curves for three different pad heights are shown. This figure shows that the above collapse (to 75.9 µm) could also be achieved on a circular pad of 79.3 µm diameter. Interestingly, such a

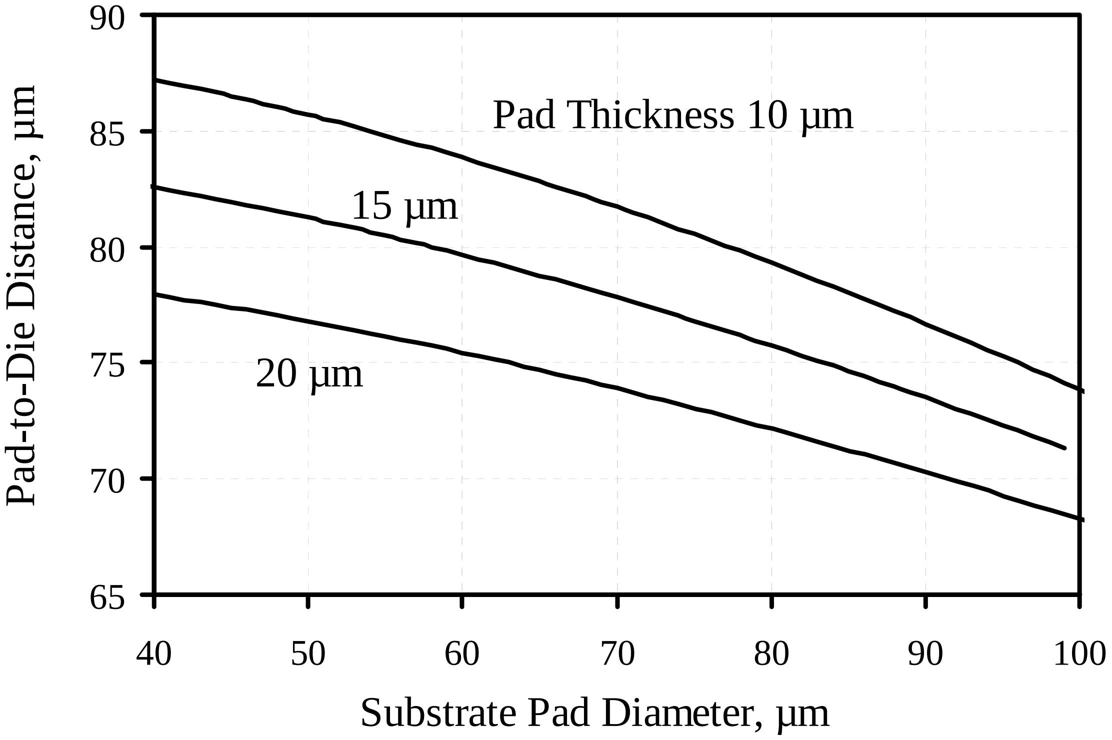

pad would have virtually the same surface area with the 50×100 μm^2 rectangular pad (which has the same area with a circle of 79.8 μm diameter). This allows a quick estimate of the collapse for parameters not very different from the ones used here, but it is not true in general, as can be seen in subsequent cases.

Using circular pads with the equivalent diameter of 79.3 μm, assuming no variations in pad diameter and a substrate without warpage, the assembly yield program predicted no defects for bump height variations up to ±20 μm (±6σ). The situation did, however, change dramatically when a modest warpage was introduced. A fixed 5 μm warpage in either direction across the length of the die gave only about 0.08 ppm defective components, but the level went up to 3.6 ppm when the warpage was increased to 7.5 μm and jumped to 77 ppm for 10 μm and 1181 ppm for 12.5 μm. These numbers indicate how sensitive the yield is to substrate warpage at reflow. Statistical variations in warpage from one substrate to another enhance this further. Even a minor variation with a standard deviation of 0.5 μm would increase the predicted defect level from 1181 to 1220 ppm for a mean warpage of 12.5 μm.

These calculations (except the first one) included a pad diameter variation with a standard deviation of 3.33 μm, which however did not significantly affect the results. It should be remembered that replacing the rectangular pad with its equivalent circle slightly increases the effect of pad size variations on the defect level. Qualitatively, this can be at least partly explained by the fact that the wetted area varies as the square of the radius for the circular pads, while it is proportional to the product of the width and the exposed length of the rectangular pad. The exposed length of the pad depends on the mask opening, whose variations are independent of the variations in pad width. For comparable scatter in pad and mask opening size the overall distribution is therefore narrower.

Incidentally, the defect levels predicted for the 50×100 μm^2 pads above show why the design with 30 μm wide traces connecting to 50 μm circular pads is not discussed further here. Because of the smaller wetted area, the collapse would be even smaller for that design, and the defect levels correspondingly higher. Rather, potential design improvements, leading to a smaller PTDD, were pursued. One option would be to increase the area of the pads, ensuring more collapse of the joints. However, design constraints prevented significant increases in the dimensions of the pads in most directions.

In a first calculation ALL the pads were assumed to be extended by 50 μm, to 50×150 μm^2, leading to a reduction in PTDD from 75.9 μm to 67.1 μm. Such an extension was, however, not actually feasible. It ***would*** be possible to extend the pads in the outer perimeter row outwards and the ones in the inner inwards, leaving the middle perimeter row unchanged. As for the pads under the central part of the die, these could of course be extended quite freely, as the distance between rows was relatively large here. Shifting one of the above pads in one direction would tend to lead to an asymmetric joint (Figure 7). If ALL the pads were shifted (in opposite directions) the asymmetry would cause a minor additional reduction in PTDD, to 66.3 μm. The fact that not all the pads can be enlarged is dealt with below. If they could, however, the assembly yields would be greatly improved. The circular pad giving the same collapse as the 50×150 μm^2 pad would

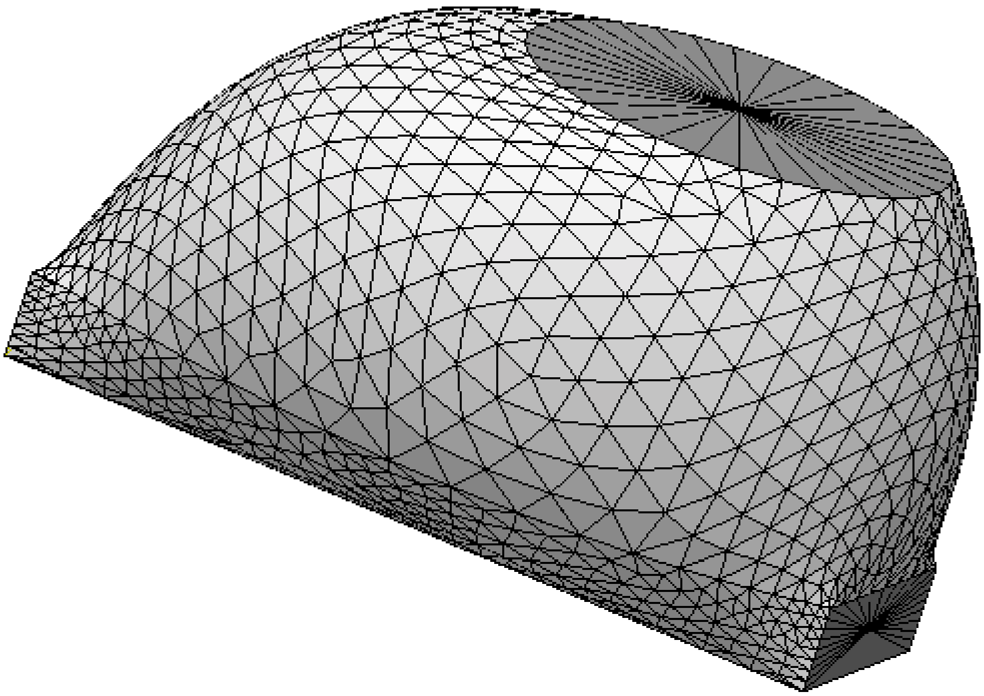

have a diameter of 118.6 μm; i.e. unlike in the 50×100 μm^2 case the equivalent circular pad has an area quite larger than the rectangular pad. According to the AYP even a warpage of 12.5 μm now leads to well less than 1 defects per *billion* components.

Another way to decrease the pad-to-die distance is to increase the height (thickness) of the pad. Of course, increasing the height means that the total volume of the joint increases as well, so that the change in PTDD will be smaller than the change in pad height. Figure 8 shows the PTDD as a function of pad height, for different (circular) pad diameters. The values are intentionally extended to unrealistic limits (for example, zero pad diameter but non-zero height!) to ensure that all the areas of interest will be included. What is immediately noticeable is that the curve for a given pad diameter is virtually a straight line (actually only the curve for zero pad diameter is exactly a straight line) and that they all pass through a common point. This common point happens to be when the pad height is half the initial height of the bump, here 50 μm, and the PTDD reaches this value. This common crossing at exactly the half-height was verified with other parameter values as well, and can undoubtedly be derived mathematically, but the matter was not further pursued. It does, however, have practical implications because we only need a single calculation in one point to draw a very good approximation of a given PTDD curve. In the case of rectangular pads, similar PTDD calculated by Surface Evolver again cross each other, but not in a single point this time. However, they too closely approximate straight lines throughout

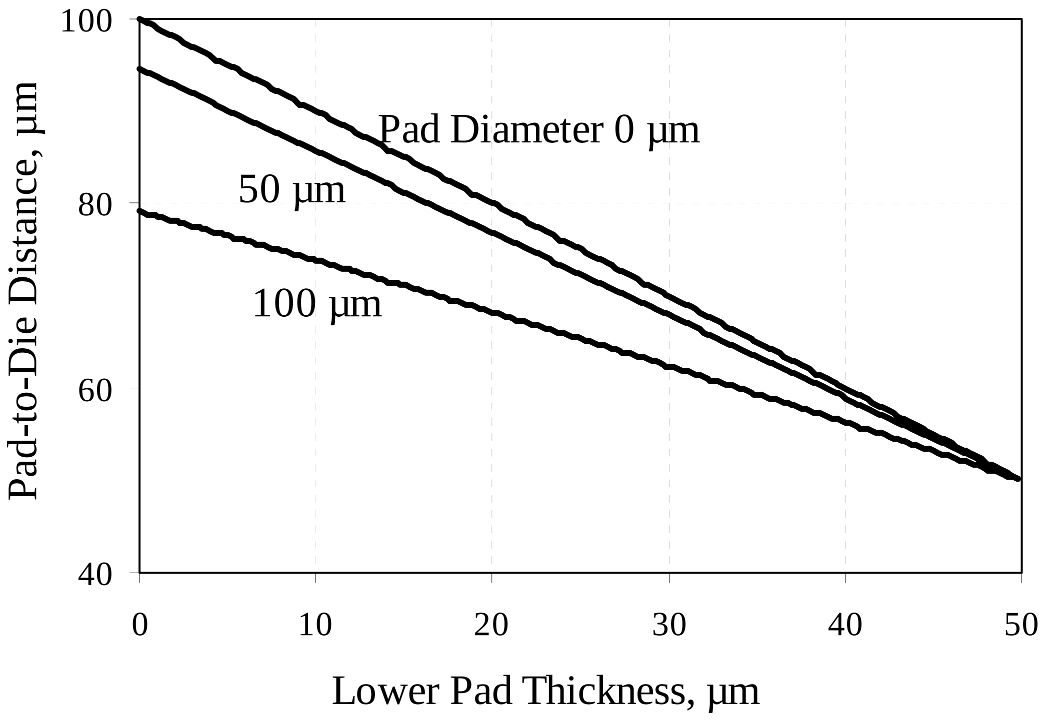

the area of interest. Thus, only two Surface Evolver calculations are required to construct a curve for a given pad shape with a reasonable accuracy.

As for the actual effect of the pad height on the defect level, increasing the height of the 50×100 µm^2 pads above from 15 to 20 µm would only reduce the PTDD by 2.7 µm to 73.2 µm. Still, the yield is significantly improved. For a constant warpage of 12.5 µm, for example, the predicted defect level is reduced from 1181 ppm to 24 ppm.

So far, as mentioned above, the fact that a certain number of pads cannot be extended in any direction was ignored for simplicity and speed of calculations. However, it is also possible to use Surface Evolver to calculate the PTDD for a collection of pads of different shape. For example, Figure 9 shows the resulting shapes when the 248 joints in the middle perimeter row still have pads with a diameter of 50 µm like in the original design. The shapes in Figure 9 appear similar to those in Figures 2 and 7, but differ in not being in equilibrium individually. It is the collection of all joints (subject to the constraint of equal heights) that is in equilibrium. The new PTDD is about 69.6 µm, i.e. about 3 µm larger than when all pads were assumed rectangular This is still sufficient to accommodate a reasonable amount of warpage. Actually, it is a worst-case estimate. In reality, the 30 µm wide traces contribute measurably to the wettable area. Anyway, an estimate like this calls for considerably more effort and computing time. In order to achieve the same accuracy as in the single pad shape case the total number of facets had to be twice as large. In addition, programming the shapes became more complicated and time consuming, since the programmer needs to keep track of every vertex, edge, and facet of the coarse initial

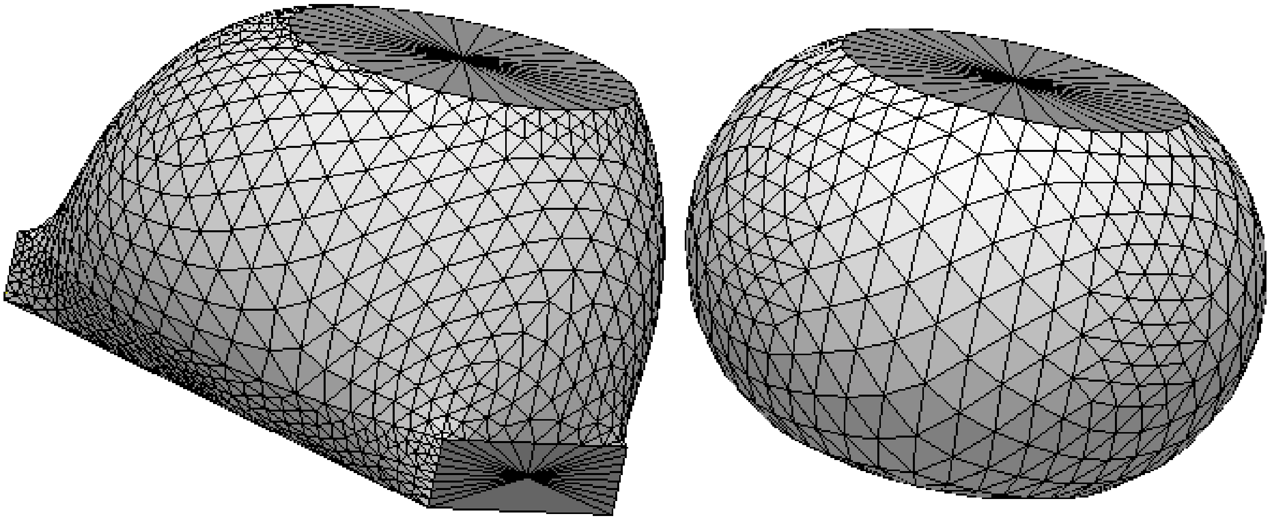

triangulation, as well as the relationship between them (for example which vertices define a particular edge) and the constraints imposed on each element. This task is tedious and errors are more likely to creep in and more difficult to detect. In principle, though, several different bodies can be considered at the same time, and the standoff that minimizes the total surface energy can be determined.

The current version of the AYP does not consider the possibility of bridging. However, as the standoff is reduced, some of the joints become more "squashed", so the possibility of bridging should not be ignored. In the case under study, the joints more at risk of forming bridges are in the middle perimeter rows, where the equilibrium heights tend to be greater than the actual gap size. The rest of the pads form joints elongated towards "safe" directions. A rough estimate suggests that about 0.1% of the die will have 2 bumps next to each other in the middle perimeter row that are both at least 3σ taller than the average, i.e. taller than 110 μm. In general, these are likely to be somewhat 'squashed' during reflow. Things will, of course, be even worse in those cases where substrate warpage leads to a smaller gap at the die perimeter (and a larger one near the center). For the sake of estimate we considered a 110 μm tall bump forming a joint to a 50 μm diameter circular pad in a gap of 64.6 μm, i.e. 5 μm below the average of 69.6 μm predicted for our improved design above. According to Surface Evolver this joint will have a diameter across the equator of about 139 μm. At a pitch of 150 μm two such joints will only be separated by 11 μm. Taking pad size variations into account as well the gap size may quite often be even smaller, but the combined probability of all the unfavorable conditions occurring at the same time remains very small. Actual bridging should therefore not be a significant issue.

Even if bridging is avoided and the assembly yield thus in principle remains acceptable, however, an 11 μm (or sometimes smaller) gap may still enhance the risk of subsequent void formation in the underfill. This is particularly unfortunate if the void remains in contact with both the joints, as that will likely lead to solder extrusion and bridging during even rather mild thermal excursions. It is therefore important to not go to extremes when trying to optimize the substrate design.

Of course, every change in design also affects the placement yield. The two yield prediction programs should therefore both be used to assess the final design. However, in general any increase in pad length, width, or diameter, can be expected to improve the placement yield, as it increases the target area and therefore the probability that it will be hit by the bumps. It is important to notice that the asymmetric extension of the pads would not create any problems similar to those created by misregistration, because the center of the site still remains at the same place.

Summary

The particular case study presented there led to the conclusion that the original design would result to unacceptable defect levels. An increase in pad length, or height, or both, is required. Given a particular acceptable defect level and an estimate of the expected substrate warpage, the required changes can be estimated and influence the design of the substrate.

In general, the combined use of the two yield programs and Surface Evolver is a powerful tool for estimating in advance expected defect levels due to statistical or systematic in-plane and out-of-plane variations in both substrate and die, thus preventing costly errors and delays. Moreover, trends in the behavior of a quantity of interest (e.g. pad-to-die distance) as various parameters are changed can be determined, providing guidance for improvement of the design. Shapes of joints and distances between them can be calculated. These programs can be used at various levels of complexity and approximation (at the cost of computing time), but reasonably accurate results can be obtained even from simplified and fast calculations.

Reference

Ken Brakke: "The Surface Evolver", Version 2.14, Aug. 18, 1999.
http://www.susqu.edu/facstaff/b/brakke/evolver/evolver.html