Process Considerations when Reworking Area Array Packages

Patrick McCall
PACE Incorporated
Laurel, Maryland, USA

Introduction:

Over the last 2 to 3 years, standard area array packages have become the package of choice in both design and manufacturing. There are several obvious reasons for this movement away from traditional leaded packages such as increased I/O per unit of area on the PCB, wider capability and functionality, and higher production yields. Today, we are in the midst of another shift in packaging technology, the movement towards micro electronic packages (MEPs). MEPs offer the same features and functions as standard BGA devices and more. Additionally, they allow electronics manufacturers to build products the consumer market is demanding. Some of the demands being placed on consumer electronics manufacturing:

1. The demand for smaller products,
2. The demand for higher power and increased performance in products,
3. The combining of functions in products such as the combining of internet access, email, palm-top computing, and video in cellular/mobile communications devices.

Because of the pressure demands identified above, BGA applications are becoming CSP and Flip Chip applications that are being enhanced further by the development of micro-via technology. The telecommunications and consumer electronics industries are truly driving the packaging and interconnect technology development in microelectronics. As a result, CSPs and Flip Chips are now an integral part of most communication and personal electronics devices. MEPs production is expected to increase by 2 to 3 times in 2000 and will increase again by 1.5 to 2 times in 2001. As the volume of these packages increases year after year, package cost will come down and the specialized technology to manufacture, attach, and apply MEPs will become more accessible to the electronics industry in general. Already, most major contract manufacturers and OEMs are using MEPs and are successfully incorporating them into their production lines.

Even though a handful of industries are truly pushing the technology envelope when it comes to maximizing I/O count on small, densely populated PCBs, the wide spread acceptance of micro-electronics will ultimately benefit all industries. Over the last decade, we have seen products that once contained numerous PCBs evolve into products with one PCB that are 75% or more percent smaller with capability not yet imagined 10 years ago.

In the quest for the next generation of size reduction, CSPs and Flip Chips will replace BGAs and SMT devices in an increasing number of IC applications and will become even more wide spread. With each generation of smaller and smarter products, new challenges are created when reworking or repairing assemblies. CSPs and Flip Chips bring with them special process considerations when rework and repair is required. Higher levels of process control and controlled heat application are required.

An issue related to reworking CSPs and FCs is structural stability of the interconnection of the device and PCB. One of the ways to add strength and robustness to micro-interconnects is to use underfill. When underfill is used, rework is usually impacted significantly with respect to process limitations, additional procedures, higher training requirements for technicians, and increased time per repair; all of which means higher rework costs. As a result, we are now beginning to see the introduction of reworkable underfills. Technology evolution will always pose challenges and the challenges will be solved in due time.

As MEPs find their way into every product, repair and rework centers will have to deal with them. There are some subtle differences between reworking standard packages and MEPs. Some of the areas that will be explored are:

1. Removal of Old Package and Site Preparation
2. Placement
3. Flux/Paste Application
4. Bottom Side Heating
5. Heating Process Considerations

Characteristics of Standard Packages and MEPs:

Standard packages can vary in size from 5.0 mm square to 50 mm square and larger. They have solder ball diameters between .5 and 1.05 mm with pitches between .9 and 1.5 mm. These packages are generally thicker than 1.5 mm. Standard packages usually contain a silicon die mounted to a substrate that is encapsulated with a plastic or resin material. Ceramic encapsulation and substrates are also common.

MEPs can vary in size from 1.5 to 20 mm square. Solder ball diameters between .1 and .4 mm are common and pitches can range from .25 to .8 mm. MEPs are usually less than 1 mm thick and come in a variety of package configurations and some types have no package. An example of this type of MEP is a Flip Chip. These are essentially the silicon die that may or may not have a protective coating on the top. CSPs have a small package around them, by definition; the package can be no larger than 1.2 times the size of the silicon die. The encapsulating material is usually some type of polymer.

CSPs can actually fall into either category, depending on the package, ball diameter, and pitch. The number of interconnects can vary widely and is not a defining characteristic.

Removal of old package and Site Preparation:

When removing packages from PCBs, the approach should be guided by the reason for the removal and whether or not the component will be reused. When removing a package because it has failed, a removal profile may be developed that uses faster ramp rates and shorter times. However, when a device is intended to be reused, a removal profile similar to the installation profile with slow and even heating or the installation profile itself should be used. This is also the case when dealing with PCBs and packages that are made from materials that must adhere to ramp rate guidelines.

Regardless of the approach, proper pre-heating and thorough PCB warming is critical to the success. Removing packages by only applying heat from the top should be avoided. This technique can easily result in pads being damaged or lifted. The amount of heat required to bring the solder balls to melting temperature through topical application will often overheat the rework site and will usually exceed recommended heating ramps. It also causes significant temperature differences between the top and bottom of the device and PCB, which can cause twisting, flexing and damage to micro-vias and other delicate circuitry as well as delamination of the package.

It is extremely important that all the solder joints under the package are liquidus before lifting the package from the PCB. Rework equipment that utilize devices to automatically lift a package after the removal profile is completed are common and generally work well. However, if they are used, the removal profile must be validated through the use of thermocouples to ensure proper temperatures are reached. Should one or a few solder joints not reach solder melt temperatures and the automated head lifts up, pads can and will be pulled off the PCB.

Another consideration when vacuum cups are used inside of the nozzle is that the vacuum cups shield a portion of the component resulting in uneven heating. Typically, this is not a major issue in larger packages because they have longer heating cycles that allow the entire package to reach a homogenous temperature. Because MEPs have shorter heating cycles, are thinner and have different materials compositions than their standard cousins, they are more likely to develop cool spots under the vacuum cup which can result in non-liquidus solder joints. To avoid pad damage, this can be accounted for by increasing the duration of the reflow phase. The use of thermocouples when developing and validating profiles should become standard practice in every rework facility.

Once a package has been removed, the rework site must be properly prepared for the next installation. Excess solder on the land patterns with pads larger than .8 mm can usually be removed using one of a variety of conductive desoldering tools available. Other techniques include the use of solder wick or the use of hot gas heating. Removing excess solder from land sites with pads less that .5 mm requires a very delicate hand. If conductive desoldering tools are selected, a light touch and low temperatures must be used. Temperatures over 200 degrees C combined with pressure can easily damage pads or the solder mask. Hot gas removal systems or conductive systems that use Teflon tips at low temperatures and incorporate hot gas assisted heating are preferred. The use of solder wick is not recommended for MEP land preparation.

In general, PCBs that have vias through them are not good candidates for conductive desoldering techniques. While the desoldering tip is moved across the land pattern, the vias can become filled with solder. These are not easy to clean out and attempting to do so usually results in damage to the via, solder mask or the PCB.

Bumping of the lands is another technique used to remove excess solder or to add additional solder to the rework site. This is quite common when placing leaded devices. This technique uses a soldering iron tip to reflow the solder on the pads and level it as well. Typically, flux is applied to the land site, and then a tinned tip is drawn across the array. The surface tension of the solder is relied upon to control the volume of solder left behind on the pad.

If pads are bumped, they will be rounded, not flat. Therefore, a sticky or gel flux must be used in to hold the package in position. During Soak when the flux is activated and driven off, the package will usually move and can easily become misregistered to the land array. While bumping can work fairly well on standard packages, it is not recommended for MEPs land site preparation. If bumping is used on MEPs, a number of issues must be considered. The amount of solder involved with a MEPs interconnect is very small. Using old solder can result in poor quality joints and adding too much can result in bridging. Often, MEPs pads are so small in diameter, that the surface tension of the solder on the tip will just pull the solder over the pad, leaving none behind. Additionally, drawing a hot iron tip over the MEPs array can damage and lift pads. The most significant issue is that the solder bumps will not be perfectly level. When installing packages with ball diameters as small as .1 mm, it doesn’t take much to end up with an unlevel package.

Once the excess solder has been removed, it is important to properly clean the PCB. The pads must be clean and free from old solder and flux residue before proceeding with the replacement of the package. When installing MEPs, proper cleaning is just as important to the process as the installation profile. Shorter reflow cycles, small amounts of flux used for installation, coupled with the small volumes of solder present in each joint does not make for a forgiving reflow configuration.

Placement:

A variety of package sizes, solder ball diameters and pitches exist in standard package configurations as well as MEPs.

Placing standard packages is relatively easy and can be accomplished with commonly available techniques. Placement by hand is a viable option for these packages by using a template or the silkscreen around the land pattern as a guide. A vision overlay system (VOS) can be used to ensure proper alignment and is usually preferred by most operators. Typically, a minimum magnification of 35X is required to ensure proper alignment. A placement precision of .1 to .2 mm in the Z travel is usually adequate.

To place MEPs successfully a VOS must be used. A minimum magnification of 80x is required and quite often 100x is needed to accurately align flip chips. Placement precision in the Z travel should be 50 microns or less. When placing flip chips, 25 microns is needed.

The level of precision required for placement in the Z travel mechanism is directly related to the diameter of the solder ball itself. Precision rating requirements should be calculated by using the smallest solder ball diameter that will be reworked. To ensure proper installation, the solder ball must cover at least 50% of the pad on the PCB. In other words, when the device is placed on the PCB the accuracy must be at least 50% in order to take advantage of the self-aligning properties of array packages. See Figure 1. Keep in mind that the more accuracy that can be achieved, the better the result.

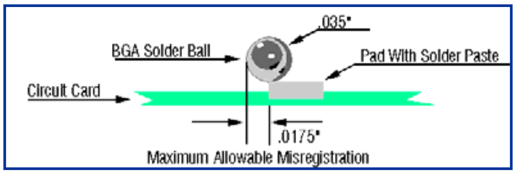

Image courtesy of IBM

Figure 1

For example, if placing a package that has solder ball diameters of .5 mm, we can assume that the maximum we can be off by is .25 mm (50% of ball diameter) when the package is placed on the board. This means that the precision tolerance of the placement system in the Z travel has to be less than .25 mm over the distance of travel, otherwise component registration could be off by more than 50%.

In theory, the accuracy can be off by as much as 50%. However, a good rule of thumb is to use 30% to ensure proper placement. If we look at our example again using the 30% rule, we would want to have a Z travel precision of .15 mm (.5 mm * 30%), over the distance of travel.

VOSs are readily available with many different options. Typically, a VOS is comprised of a prism that is used to collect two images, one from above and one from below. The images are projected onto a series of mirrors where they are then projected into the lens of a camera. The images are displayed on some type of video monitor and appear as two separate images overlaid on one another. Either the component or the board is repositioned until the ball and land array patterns match exactly. See Figure 2.

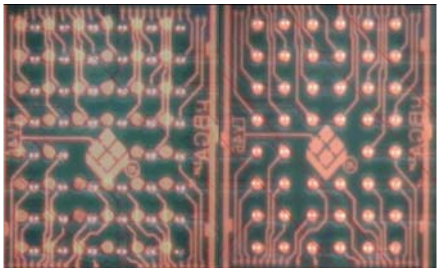

CSP not aligned CSP aligned

Figure 2.

Most VOSs used today also have a “split vision” capability, which allows the viewer to see only two opposite corners of the images at a higher magnification. This is valuable when placing devices that have hundreds of solder balls. To attempt to align more than 300 data points is difficult. If the images can be “split” and magnified, the user can simply align two opposite corners of the package while focusing on less than 100 points of data. See Figure 3.

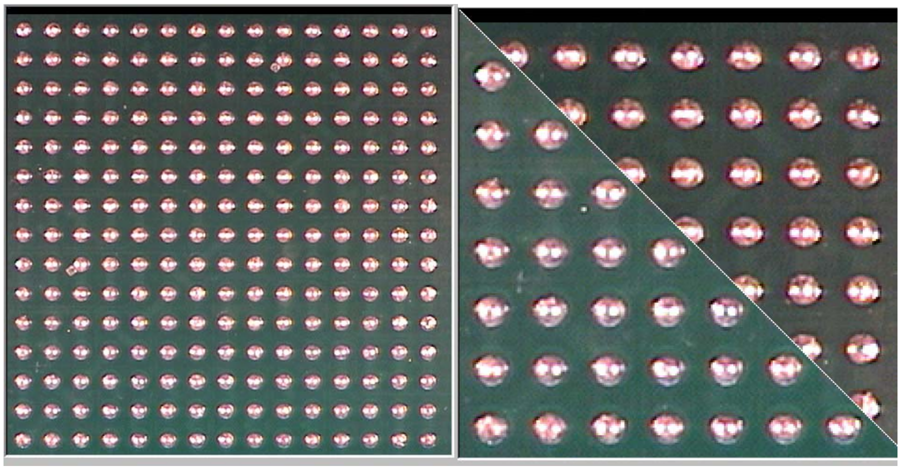

Figure 3: Full view and Split Screen View

Proper care and maintenance of the VOS is important to placement success. Calibration checks should be performed at regular intervals as indicated by the manufacturer. When adjustment is required, it should only be completed by qualified individuals.

An equally important piece in the placement puzzle is the table/board holder. The table/board holder should be stable, durable, adjustable, and able to hold a variety of board shapes, either directly or by using optional carriers. Board holders should also be spring-loaded and be capable of holding a wide array of board sizes. When reworking MEPs the ability to securely hold small boards, down to 20mm wide, is desirable. Additionally, PCBs containing MEPs are usually thinner and require proper support from the underside to ensure planarity.

Flux/Paste Application:

Flux must be used to have successful package installation. It may be used by itself or it may be combined with solder in the form of solder paste. Applying the proper amount of flux to the rework site is critical. Too little flux and the solder will not flow correctly, too much flux and out-gassing can occur which result in voids within the solder joint or flux will be left over which can cause resistivity problems as well as have corrosive effects on the solder joint after reflow. There are a variety of methods for applying flux. Some of the methods are to use a brush or pen applicator, to use a piston driven mechanism to dispense a gel or sticky flux, and to use a flux applicator tool.

When applying flux to area array packages, the ideal amount of flux to apply is enough to cover 1/3 of the solder ball. This ensures that enough flux is present to clean and remove oxides while eliminating the potential for out-gassing and excessive flux being left behind. With the introduction of gel flux, applying precise amounts of flux to balls can be accomplished and repeated consistently.

The best method for applying flux is to use an application device known generically as a “gel-flux applicator”. These devices are essentially a metal block that has a “well” milled into it with a depth that correlates to 1/3 the solder ball diameter. Gel flux is applied and a squeegee is used to distribute the flux evenly across the well. The component is then picked up using a vacuum tool or pick and set down into the even layer of gel flux. This process closely mimics automated placement machines and ensures that the proper amount of flux has been applied to the solder balls and consistent amounts have been applied to each of the balls. The component is then ready for placement.

Solder paste should always be used when:

- Solder balls on the bottom of the package are made from 90/10 solder.
- PCBs have bare copper pads on them,

- The package contains an elastomer layer for CTE differences,
- It is critical to have the same joint geometry as the production assembly and solder paste was used in initial production,
- No clean flux can not be used,
- A specification exists for the joint stand-off height for function or cleaning.

Solder paste can be applied using a variety of techniques. Some of the more common techniques include using spot stencils to apply solder paste to the PCB, using component stencils to apply solder paste to the component, and using dispensing equipment to apply dots of solder paste to each individual pad.

Solder paste with ball sizes of #4 should be used for small pad printing. When applying solder paste to standard packages, typically, a paste thickness of .17 to .23 mm is desired. When using solder paste for MEPs, a paste thickness of .1mm to .16mm is desired.

Stencil aperture shape is also an important consideration. Typically, stencils with apertures larger than .2 mm are round and lend themselves to good release. Apertures in stencils designed for paste deposits smaller than .2 mm are often diamond shaped. The diamond shape allows the solder paste to release from the stencil easier than a circular one.

The Impact of Bottom-Side Heating:

Bottom side heating is usually associated with the “pre-heat” phase of a profile. However, bottom side heating is as important in Soak and Reflow phases if the installation is to be successful while exposing the PCB and package to a minimal amount of thermal stress.

In the “Pre-heat” phase, bottom side heating ensures homogenous temperatures across the board. This keeps the PCB from warping, twisting, or flexing, during the process, which is essential for maintaining planarity of the installation site. Heat application from the bottom of the PCB during the “pre-heat” phase is also used to warm the entire PCB so that heat is not drawn away from the installation site during the rest of the process.

During “Soak”, the bottom side heater should continue to operate while a relatively small amount of heat is added from the top heater. The combination of top and bottom side heat application allows the installation site and package to reach a temperature of between 140 and 160 degrees C and to stabilize. The stabilization should be maintained for 40 to 60 seconds, allowing the flux to activate and driving off any volatiles in the flux. This is important as it eliminates the potential for out-gassing and prepares the package and PCB for reflow.

During “Reflow”, heat is usually applied from the top heater. Top heaters generally operate between temperatures of 200 and 350 degrees C. Obviously, this can place a lot of thermal stress on the top of the package. It is important to apply heat slowly and evenly so the entire package warms to reflow temperature uniformly. Temperature differences across a package of as little as 7 to 10 degrees C can cause damage.

The set temperature of the bottom heater can be maintained or increased. Increasing the temperature by as little as 30 degrees C during reflow will have a dramatic impact on the package profile. One of the many benefits of adding additional heat from the bottom or “Spiking” during the reflow phase is that lower temperatures can be applied from the top.

Installations should always be achieved with the lowest temperatures possible. This ensures the safety of the package as well as PCB. Additionally, by subjecting the package to lower temperatures, there is less chance for temperature overshooting which can result in significant temperature variations between the top of the package, bottom of the package, the solder joint and PCB. See Figures 4 and 5.

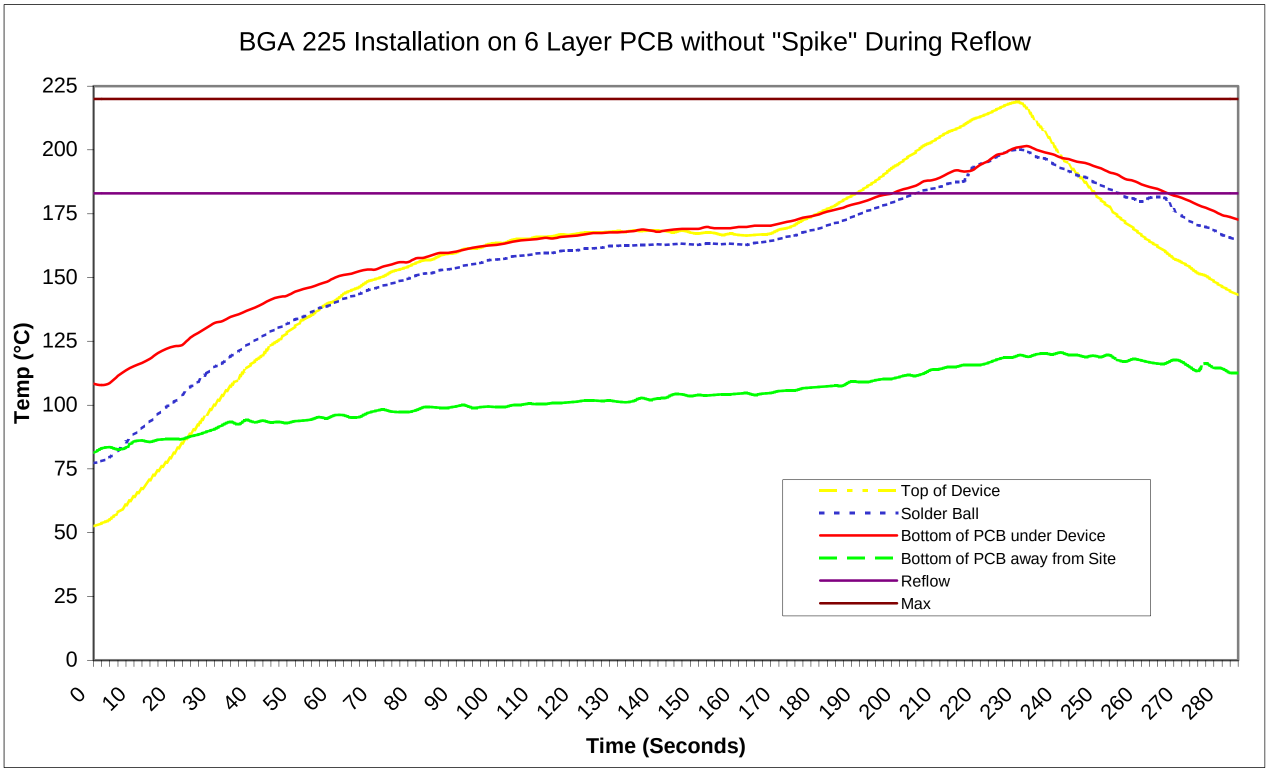

Figure 4.

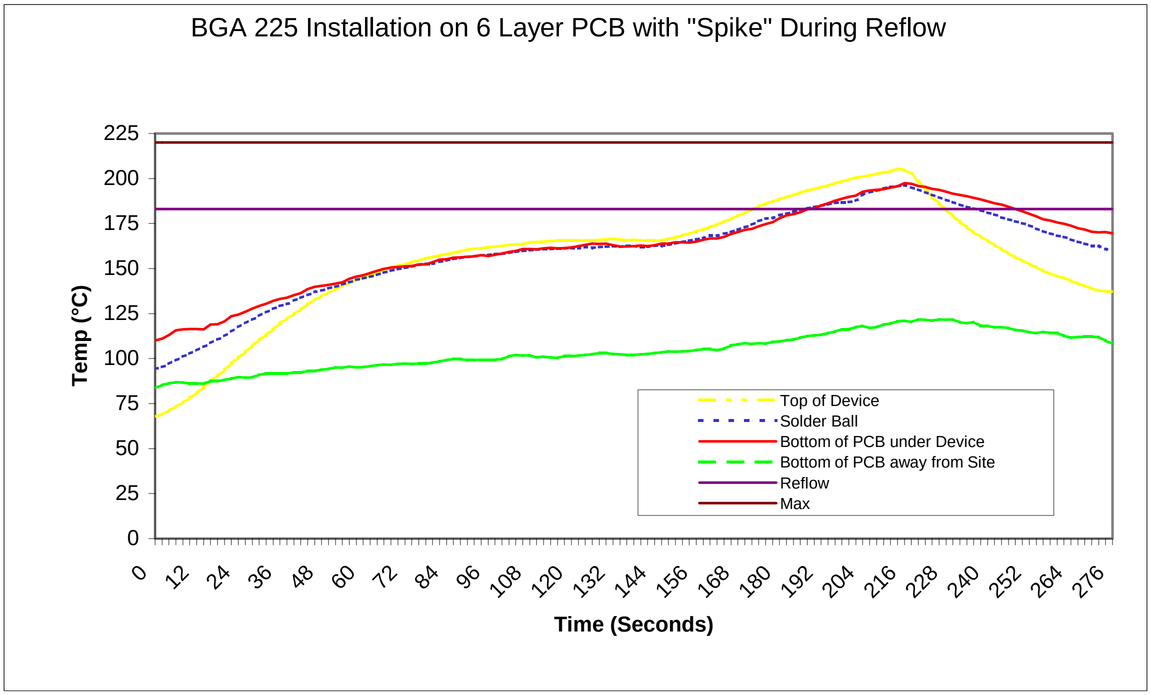

Figure 5.

As can be seen in Figures 4 and 5, reflow was achieved successfully in both scenarios. However, in Figure 4 the temperature variation between the top of the package and the solder ball is significant enough to damage the package. Refer to Table 1 for spike impact analysis.

Table 1

	No Spike (°C)	Spike (°C)
Max reflow temp (top of Device)	218.5	205.4
Max temp of solder ball	200.2	196.1
Max temp of PCB under Device	201.4	197.5
Delta between top of Device and Solder ball @ max reflow temp	18.3	9.3
Delta between Bottom of PCB and solder ball @ max temp	1.2	1.4
Delta between Top of Device and Bottom of PCB @ max temp	17.1	7.9

In another comparison, the effect of the spike can be seen in lower reflow temperatures as well as in more even heating between outer and middle solder balls in the array pattern. Holes were drilled through the center of two lands on a 6 layer PCB, one on an outer row of balls and one in the center of the component. Thermal couples were positioned and secured even with the top of the PCB so the solder balls on the component would contact the thermocouples without affecting the position of the package. Figures 6 and 7 show the effect of the spike in this scenario.

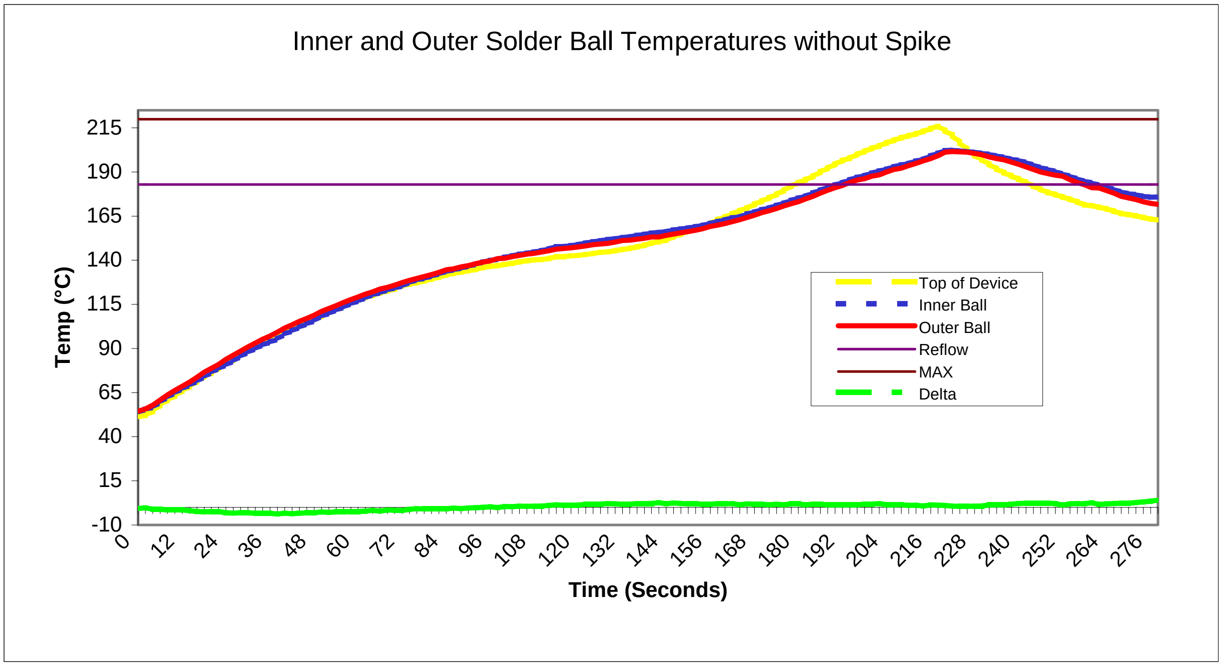

Figure 6.

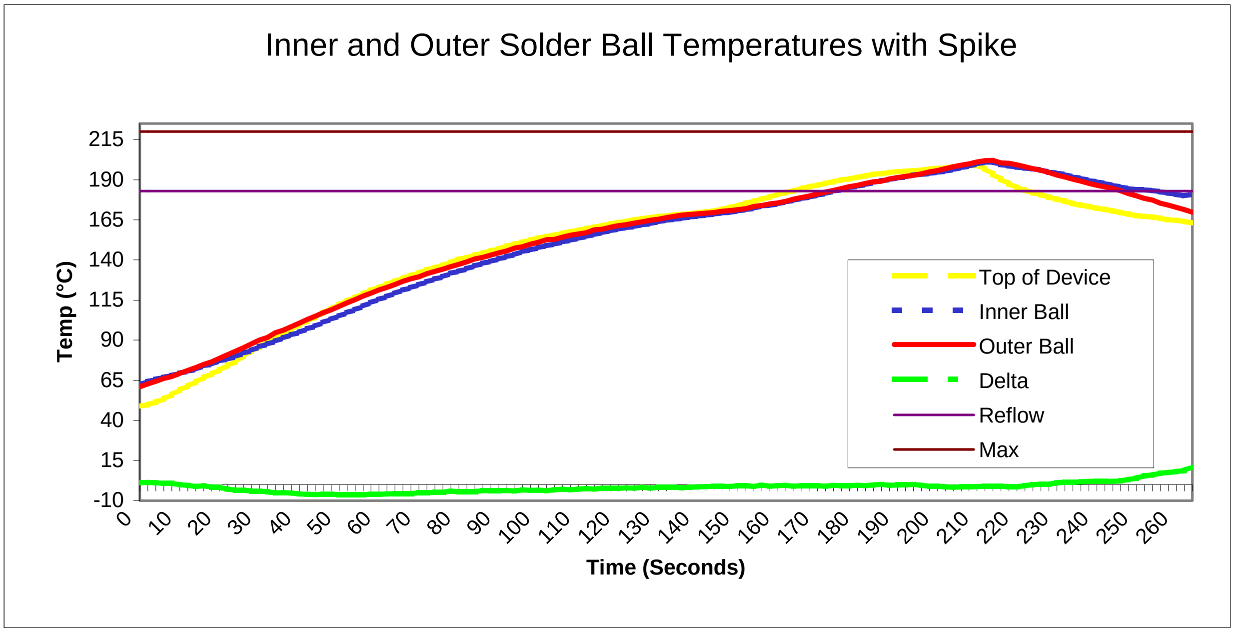

Figure 7.

By controlling the application of heat from the top and bottom of the installation site, more even heating can be achieved. This can be seen in lower temperature variances between the top and bottom of the package and between solder balls located on the outer edge and in the middle of the array pattern. See Table 2.

Table 2.

	No Spike (°C)	Spike (°C)
Max reflow temp (top of Device)	216	199.01
Max temp of Inner ball	201.9	201.8
Max temp of Outer Ball	201.66	202.15
Delta between top of Device and Inner Solder ball @ max reflow temp	14.1	2.79
Delta between top of Device and Outer Solder ball @ max reflow temp	14.34	3.14
Delta between Inner and Outer ball @ max temp	.24	.35

Heating Process Considerations:

While standard BGAs and MEPs look similar in design, they are actually quite different and demand different profile configurations when it comes to installation and removal profiles. Physically, they differ in several areas: materials, dimension, mass, and tolerances to temperature. In general, profiles used to install MEPs are shorter than profiles for standard packages. Also, lower temperatures can and should be used when installing MEPs. Because the MEPs package is the size of the silicon die, or slightly larger, there is no significant encapsulation for protection, heat is applied/transferred directly to the silicon.

The smaller mass of the MEP solder balls, coupled with a very thin package allows heat to be transferred through the component very quickly. When standard BGAs are installed, heat must be driven through and around the package, which requires a longer heat application period. Therefore, if a profile for a standard package is used on a MEP, the component will become super heated which is something that must be avoided.

As was mentioned previously, vacuum cups in the nozzle can affect profile results. To demonstrate their effect, a micro BGA 46 was installed twice on a PCB from a cellular phone using the same profile parameters. A hole was drilled through the center of a pad and a thermocouple was installed flush with the top of the PCB, contacting a solder ball. The only difference between the two trials was the addition of the vacuum cup. See Figure 8.

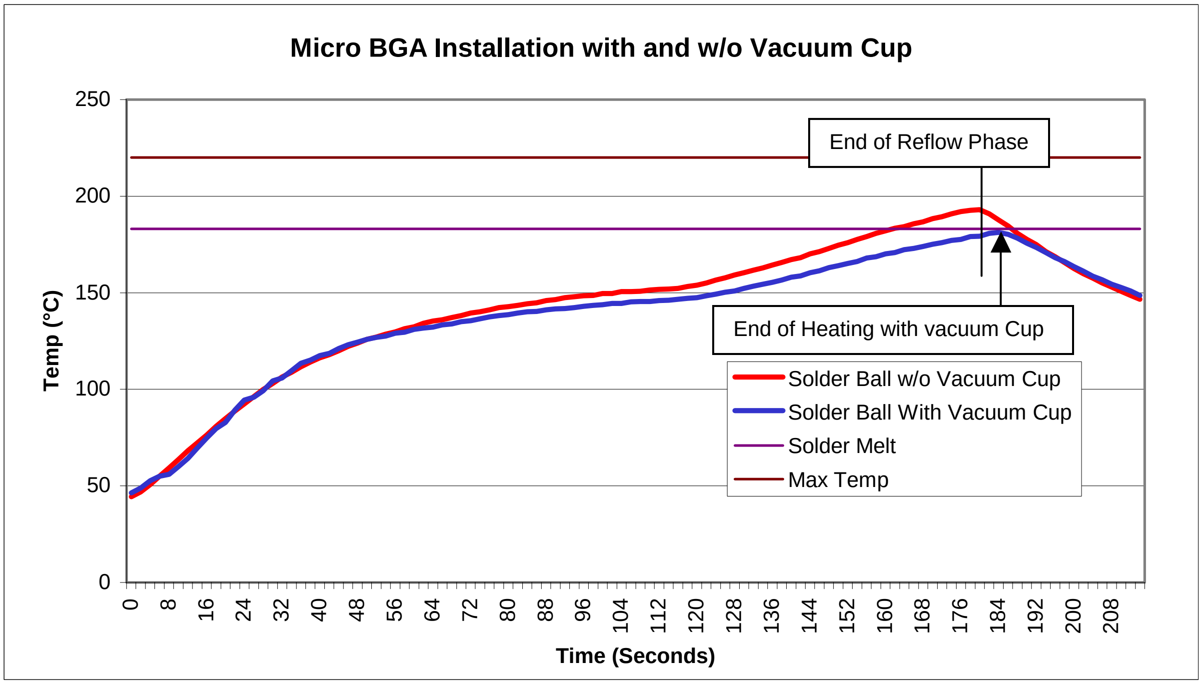

Figure 8

The effects of the vacuum cup is dramatic:

1. Soak and Reflow ramp rates are slower
2. Final Soak Temp is lower
3. Heating continued beyond the reflow phase
4. Cooling was inhibited
5. The solder balls did not reflow.

In this case, the vacuum cup used was round so the edges and corners of the micro BGA were exposed. The exposed parts of the package were subjected to more heat and were exposed to more heat than the portion of the package under the vacuum cup. This results in higher temperature variations across the entire package.

There are many methods for applying heat as well as different methods to monitor and control it. Arguments can be made for all. Heaters used in area array rework equipment are almost always controlled through closed loop structures. This means that the heater is cycled on and off based on the condition of the thermal sensor. In many cases the variation between set temperatures and actual temperatures between area array rework equipment is due to the position of the sensors.

Sensor placement can dramatically affect the result of set temperatures, which is why a valid profile configuration developed on one rework machine does not usually successfully translate to another. Ultimately the set values on a particular machine are not critical. What is critical, however, is the effect of the set temperatures on the package and PCB. Profiles should always be created and validated using feedback from thermocouples.

Many systems allow profile settings to be adjusted while a profile cycle is running. Usually a PCB will be populated with thermocouples in a number of locations. It is good practice to place thermocouples:

- on the top of the device,
- on at least one solder ball, 2 is better (inner and outer ball), and
- on the bottom-side of the PCB, both, directly under the rework site and away from the rework site.

Positioning thermocouples in these locations allows:

- monitoring of temperature variations between the top and bottom of the package,
- monitoring of temperature variations between the middle and outer edge of the package, as well as
- monitoring of the thermal environment the solder balls are exposed to.

Monitoring the temperature directly below the rework site is important as PCBs can be subjected to too much heat, as can packages. This can cause damage to circuitry, micro vias, and delamination of the PCB itself. The purpose of the thermocouple on the bottom of the PCB away from the rework site is to ensure the entire board is warmed properly and is not exposed to high temperatures.

Profiling can be accomplished using a site containing a previously installed package, or by performing an actual installation. Either method can be used to develop a reliable profile. However, there are some issues to be aware of with each.

If using a site where a package is already installed, the placement of the thermocouples is important. They must have contact with the existing solder joints. This is best accomplished by drilling through the bottom of the board into a solder joint and affixing the thermocouple with an adhesive or other means to secure it. Affixing a thermocouple to the top of the package that is already soldered is easier in most cases, especially when MEPs are involved.

When developing profiles through an actual installation, it is important to make sure the solder balls are touching the thermocouples through out the entire process. Should a thermocouple lose contact with the solder balls, bad data will be collected. Additionally, it is often difficult to affix a thermocouple to the top of a loose package and successfully maintain contact with the PCB during the process.

Feedback from the thermocouples will assist the operator in adjusting the time and temperature parameters. In general, the following guidelines should be adhered to when developing profiles.

Ramp and Maximum Temperatures

Acceptable ramp rates and maximum temperatures can be obtained from the package manufacturer. It is wise to select a maximum temperature that allows a margin of safety to the manufacturer's specification.

Pre-Heat

1. If a "step profile" is desired, the top of the PCB should reach a stable temperature of 95 to 105 ºC during pre-heat. Stability is the main goal here. When plotting a curve, the trace should level off at this temperature.

2. If a linear slope is desired, preheat is merged with the soak phase. In this method, the package and PCB are warmed at a constant rate (the ramp rate usually 1 to 3 ºC/s) until the desired soak temperature is reached.

Soak

Soak is a critical step in the process. It is used to activate the flux and to drive off volatiles and excess flux. It is important that a soak temperature between 145 and 165 ºC be reached. Ideally, a relatively stable temperature should be maintained for 40 to 60 seconds. This should be extended if liquid flux is brushed on to the PCB or when installing a large package. Soak also allows the entire package and PCB to come to a uniform temperature. If the entire package and PCB are the same temperature the rate of temperature increase in response to the additional application of heat during reflow will be the same throughout. This allows for uniform ramping across the entire package.

Reflow

Reflow is the second to last phase in the cycle. During this phase the solder balls reach solder melt and form the joint between the package and the pads. It is critical that all areas of the array reach solder melt together and that all balls are liquidus for at least 10 to 15 seconds. Generally, Reflow for MEPs will be a shorter phase than for standard packages. Additionally, lower temperatures can be used for MEPs as they are thinner and have less mass. When developing profiles for MEPs it is very important to not apply too much heat.

Cool Down

Cool down is the last phase of the cycle. It is used to bring the temperature of the package, solder joints, and PCB under the package below solder melt temperatures. Cooling should be controlled and a good rule of thumb is to use the same rate for cool down as for ramp up. If the ramp rate were 2 ºC/s increase, the cool down rate would be a temperature decrease of 2 ºC/s.

Once a profile has been initially defined, it is important to re-run the profile with the determined parameters in a static (non-changing) environment to ensure the results are valid.

Conclusions:

The use of MEPs will continue to expand and the ability to rework them will become a critical competency. They offer similar benefits as standard BGA packages but at a much smaller size. During rework, it is extremely important to prepare the land array properly by removing excess solder and proper cleaning. Placing MEPs on the PCB is a similar process to placing standard BGAs. However, it can require more delicate and precise equipment. The proper application of flux and paste (if used) is also critical to the success of the installation. Having the right amount is the key, too much or too little can lead to more rework.

Bottom side heating is one of the keys to success. When used properly, reflow can be achieved at lower temperatures while exposing the package and PCB to less thermal stress. Additionally, increasing the temperature of the bottom heater during reflow allows for a stable and uniform reflow. While it can be compensated for, it is important to be aware that the results of the profile can be affected if any additional mass is present inside of the nozzle or if anything is contacting the package during reflow. The use of thermocouples is paramount when developing and validating profiles. Using multiple thermocouples at multiple locations is preferred.

Guidelines for profile development exist and should be used. Manufacturers specifications should also be adhered to and it is good practice to add a margin of safety to maximum temperature and ramp rate specifications. Profiles used for standard BGA packages will not transfer directly to MEPS. However, they can be used as starting points in developing the perfect profile.

References:

1. Kevin Towle, Line Configuration for Today’s Highly Complex Manufacturing, SMT Magazine, 11/98
2. Japan Printed Circuit Association, Report on the Technology Road Map for Advanced System Integration and Packaging, 1998
3. Mike Bixenman and Treliant Fang, Wafer Solder Bumping, Advanced Packaging, 6/99
4. Mike Sauer and Kristina Bergman, Rework & Repair, SMT Magazine, 7/99
5. Paul Wood, New Considerations in CSP Rework, Circuits Assembly, 7/99
6. Mei Wang, Kazu Nakajima, et al, Investigation of the Printing Process for CSP Assembly, SMTA 1999 Conference Proceedings.