Flip Chip with No Flow Underfills

Pericles A. Kondos,
Ph.D., Process
Research Engineer
kondos@uic.com

Peter Borgesen, Ph.D.,
Project Manager
Flip Chip and
Optoelectronics
Packaging Research

Universal Instruments
Corporation
Binghamton, NY 13902

ABSTRACT

Reflow encapsulants or flux underfills have appeared in recent years as an alternative to capillary flow underfills, offering several cost and throughput advantages. The present work addresses issues related to the use of these in low cost flip chip assembly. Numerous materials from several manufacturers were considered. Voids in the encapsulant are a potential problem. These can originate during dispense, die placement, or reflow but can be eliminated through proper materials selection, board preparation, dispense parameters and reflow conditions. Electrical continuity was achievable by optimizing assembly parameters and reflow profile. Optimized assemblies performed quite well in thermal cycling tests.

Introduction

Conventional assembly of flip chips with eutectic Sn/Pb solder bumps involves either dispensing a liquid flux on the substrate or dipping the bumps into a thin film of a paste flux before chip placement. The latter clearly offers better 'tack' and thus less risk of the chip shifting before reflow. On the other hand, dispensing of the flux before the placement machine may improve the throughput. In either case, reflow in a nitrogen ambient is usually preferred.

After reflow, the die must be underfilled with an appropriate encapsulant, which flows underneath the chip due to capillary forces (Figure 1). No-clean flux residues tend to affect both underfill wetting/flow and subsequent reliability, leading to ongoing concerns as to materials compatibility. On the other hand, cleaning under the die is clearly not attractive. Finally, a truly optimized underfill process usually requires the maintenance of an elevated temperature on the order of 80°C to within a few degrees.

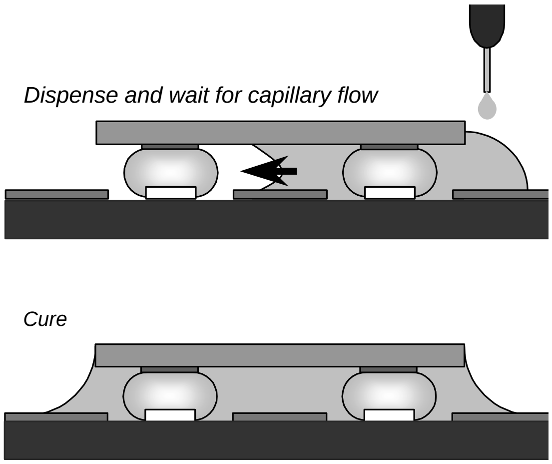

Figure 1. Underfilling with conventional capillary flow.

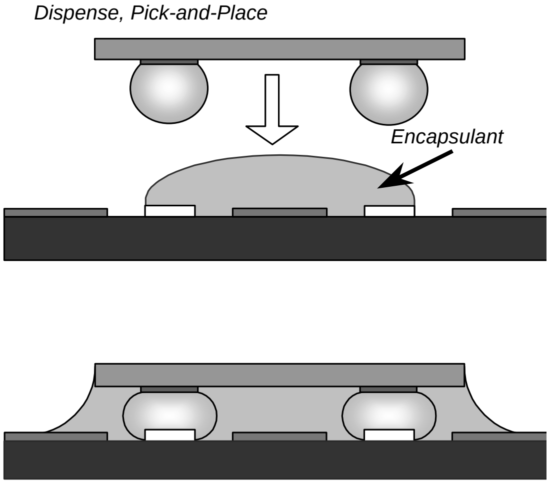

Figure 2. Assembling with a reflow encapsulant.

The use of a pre-applied reflow encapsulant (or no-flow flux underfill) offers a potentially attractive alternative, eliminating the need for nitrogen reflow as well as for a precisely controlled substrate temperature and, in the case of large die and small standoffs, long flow times (Figure 2). Compatibility of flux and encapsulant is clearly not an issue, 'tack is excellent, and the material can be dispensed before the substrate enters the placement machine.

The use of reflow encapsulants is, however, not without its own issues and limitations. Bake out requirements are not necessarily compatible with a high volume manufacturing process flow. Dispense parameters must be optimized in order to simultaneously minimize dispensing time and void formation while ensuring fluxing of all the bumps. In fact, the evolution of voids in the underfill is sensitive to details including board bake out, preparation, and handling, as well as placement and reflow parameters. The reflow encapsulants inhibit self-alignment in reflow and thus require a better placement accuracy. Depending on the specific materials properties the placement process may also need to be optimized in terms of force and holding time to ensure both throughput and assembly yields. In fact, almost any hold or slow down is likely to reduce throughput to below that achievable with a regular paste flux. In many cases, notably for materials designed to minimize or eliminate post curing, assembly yields are quite sensitive to details of the reflow profile. Finally, the very high CTEs of the reflow encapsulants affect the relative importance of the individual damage and failure mechanisms, limiting the resistance of assemblies to thermal excursions.

The present paper describes work on a range of materials, some requiring post cure and others not, and some designed specifically for one type of substrate pad metallurgy. Proper materials selection and process optimization was seen to allow for void free assemblies with quite high reliability. The potential robustness and limitations of the process in manufacturing is discussed.

Experimental

The observations presented here are based on numerous experiments performed with many reflow encapsulants in different stages of development and maturity; some of these materials are not available any longer. Suppliers of these materials included Dexter, Emerson & Cuming, Kester, Multicore, Questech, and Sumitomo. The materials covered a wide range of properties (e.g. rheology) and reflow conditions.

The liquids were dispensed by means of an Asymtek Millennium series dispenser, using an auger rotary pump (DV-6000). Three different Universal placement machines with different features and accuracies down to s < 3 mm were used. The reflow ovens included a Heller 1800 8-zone, and several Vitronics (7 and 10 zones). The results did not seem to depend in any way on the oven used.

Die with sizes ranging from 138 mil (3.5 mm) to 550 mil (14 mm) were assembled onto 16 - 62 mil (0.4 - 1.6 mm) thick FR-4 and BT substrates from a variety of suppliers. Both perimeter array and area array die, with pitches ranging from 4 to 18 mil (100 to 457 mm), were considered. All bumps were eutectic Sn/Pb solder. Both Ni/Au and OSP coated copper substrate pads were investigated. A typical pad configuration was a 2 - 3 mil (51 - 76 mm) wide trace through an 8 mil (203 mm) wide solder mask trench opening, but pads in individual solder mask openings were considered as well.

Assemblies were analyzed for soldering and extrusion by cross sectioning and by using a FeinFocus X-ray microscope, and for voiding and delamination from chip and substrate using a Sonoscan Scanning Acoustic microscope (C-SAM).

Assemblies were tested for popcorning by moisture exposure followed by 3 reflows with peak temperatures of 230oC. Most of them were cycled in Liquid-to-Liquid-Thermal-Shock (LLTS) between -55oC and 125oC with 10-second transfers and 5-minute holds. Others were, however, exposed to a variety of Air-to-Air cycling or shock tests.

Assembly Issues

Board Preparation

Even with traditional underfills the substrate needs to be dry before underfilling and curing, otherwise moisture from the board may come out during the curing cycle and form voids. Voids are, of course, undesirable, especially if they are near joints, because they can lead to later extrusions and bridgings; popcorning is also a potential problem. In addition, the properties of the cured encapsulant may be affected by the presence of water. For this reason, unless the substrate can be trusted to be dry, bake out of the boards is required before underfilling.

With reflow encapsulants, board dryness is even more critical, because of the much higher temperature the assembly sees during reflow (as high as 235°C), compared to traditional underfill curing (typically of the order of 150°C). Not only does the higher temperature increase the rate of moisture transfer from the board, but it can also cause water to come out that would have still remained bound inside the board at lower temperatures. Connected to this is the shorter time boards can be kept in ambient conditions after drying without more voids appearing in the final product. Figure 3 shows the effect of exposing the substrate to moderately humid conditions before dispensing the encapsulant. Even one hour of exposure is enough to see significant increase in the number of voids.

Dispensing

Reflow encapsulants, being unfilled, are as a rule less viscous than traditional underfills. Even so, their viscosity varies widely; the present work is based on studies of materials with a range from 1.2 to 120 Pa.s. Although the viscosity of the material is not as critical a parameter as with traditional underfills, which have to flow by capillary action in small gaps in a reasonable amount of time, it should still be taken into account when the dispense step is designed.

The requirement from the dispense step is the formation of a single glob of material, of the appropriate volume and shape, which will allow the fluxing of all the bumps during reflow and form proper fillets afterwards, with the minimum number and size of voids.

All the patterns dispensed used the "weight control" mode of the dispenser. The weight to be dispensed is given as input, and the machine uses this and a flow rate it measures independently to calculate the dispense time. If the needle is too close to the board, the material is prevented from coming out of the tip; the flow rate is reduced and smaller amount than the desired is dispensed. This problem is more pronounced when a high viscosity material is used. Another

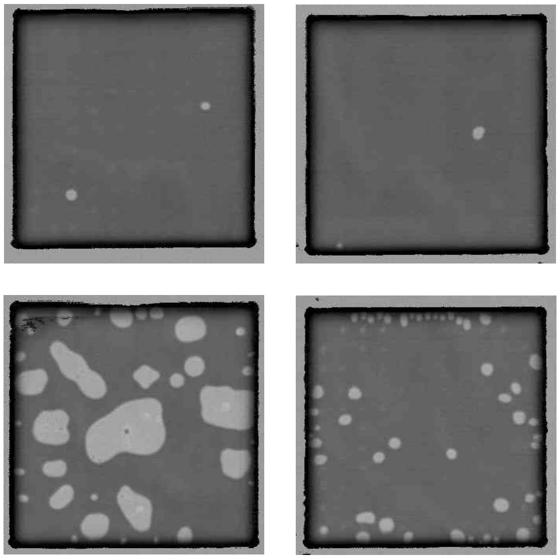

Figure 3. Voids in the encapsulant after exposure of substrate to 30°C / 60% R.H. Clockwise from upper left:

viscosity-dependent problem is the tendency of viscous liquids to form a "tail" as the needle retracts. This tail eventually breaks but if the needle has moved in the meantime, strings of encapsulant can land elsewhere on the board. To avoid this, longer retract distances and lower retract speeds must be used for viscous materials, but this slows down the step and reduces throughput.

The simplest dispense pattern is a single drop at the center of the site. Its advantage is that, since the dispense head is not moving, a higher flow rate can be used, speeding up the process. The tip of the needle can remain at a relatively large distance from the substrate, minimizing its effect on the flow rate. The needle needs to be retracted only once at the end of the step. But this pattern is not the optimum for avoiding assembly voids, and cannot be used at all for large die because the material will not reach the corner bumps.

Other patterns that can be used include cross, x, or an "area fill" (the Asymtek fills an area by describing a square-shaped "spiral", but other paths can be programmed if desired). Of those, the area fill spreads the material better, but requires the longest path for the dispense head. As a consequence, either the flow rate must be very small or the movement of the dispenser very fast. Small flow rate increases the time needed for dispensing the liquid and may affect the accuracy of the amount dispensed. On the other hand, the needle needs to be lifted only once, at the end of the pattern, just like when a single drop is dispensed. Other patterns (e.g. an asterisk formed by an x and a cross) might have a shorter total path, but any time savings will be lost when a material that tends to form tails as the needle retracts is used, because the retract speed must be set to a low value.

At the other size extreme, an area fill cannot be used for large die, because after dispensing the material will be spread too thin and tend to shrink back into a more compact shape. It may even break into smaller pieces, which virtually guarantees large voids. But other patterns can form large voids as well. Figure 4 shows voids appearing under a 550 mil (14 mm) square die. The pattern dispensed was a complicated one, in the shape of a small area fill with a cross and an x superimposed on it, the latter with arms extended all the way to the edges of the die.

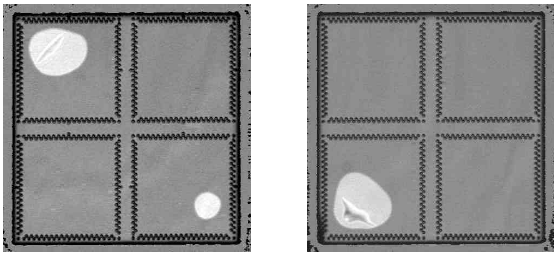

Figure 4. Voids under a large die formed because of the dispense pattern.

This pattern produced several void-free assemblies, but occasionally it broke up and large voids were formed. On the other hand, when a single glob of material was dispensed near the center of the site, the material never reached the corner bumps and they didn't solder, even when a larger amount of material was used.

Placement

The presence of a more or less viscous liquid over the die site affects several aspects of the die placement and must be taken into account when the placement parameters are decided. Even before the die is placed, the dispensed encapsulant may affect the choice of features for local placement correction. In some instances, instead of fiducials, it is preferable to use pad site recognition for local correction. If some of the pads are covered, this method cannot be used. Conversely, if it is necessary to use pad site recognition, care must be taken so that the encapsulant does not spread so much after dispense that it covers the pads.

As the die is placed on the substrate, it first comes in contact with the encapsulant at an area near its center. As it keeps moving downwards, it compresses the liquid and pushes it outwards. As it moves, the liquid has to go around or over "obstacles" (e.g. bumps, traces, etc.) and fill the solder mask openings. This motion generates and entraps bubbles, in particular behind bumps, in a more-or-less downstream direction. Figure 5 shows bubbles near bumps; the die was placed on a glass slide and it is viewed from underneath. The grainy appearance of the image is due to the fact that this particular encapsulant contained solid particles that dissolved at higher temperature. The bubbles near the center bumps (the ones arranged in a triangular shape) are comparable in size with the bumps themselves;

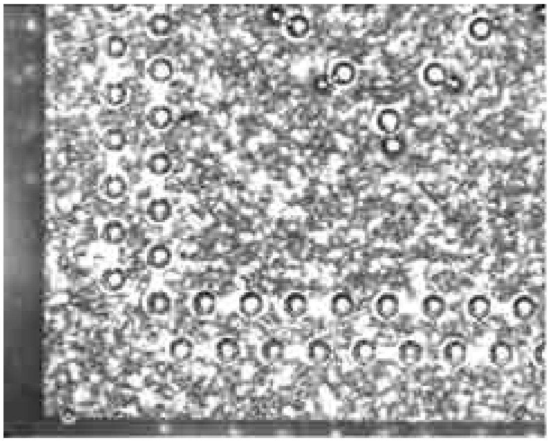

Figure 5. Placement bubbles in the encapsulant near the bumps.

the size of the bubbles can be estimated by comparison with the pitch of the perimeter bumps, which is equal to 10 mil (254 mm). The perimeter bumps had similar bubbles, but they were detached and can be seen rising in the fillet (out of focus); one bubble from a corner bump can be seen half under the die and half outside in the lower left corner. Most of the detached bubbles left a tiny bubblet behind them, near the respective bumps, but these cannot be seen in Figure 3. Such behavior is typical of every placement and cannot be avoided. Even when a clean glass slide was placed on a glob of material on another glass slide, several bubbles were seen in the liquid. Of course, the more irregularities the two surfaces (die and substrate) have the more bubbles will form during placement. Once the material has reached the edges of the die, it must wet the sides and "climb" them to form a fillet. During that time, the die must be held in place, or it might temporarily float on the drop of liquid and settle in the wrong place later. The time it takes for this to happen depends on the system, and in particular it is strongly dependent on the viscosity of the liquid. Videotaping the liquid as it wetted the sides of the die indicated that the time required for the liquid to stop visibly moving varied from about 0.1 seconds to well over 1, depending on the encapsulant. This holding time reduces the throughput of the pick-and-place machine although the delay is at least partly offset by the absence of a flux-dipping step.

Not only must the die usually be held for a certain amount of time before the placement head is withdrawn, but it must be pressed down with a certain force during that time as well. The viscous fluid underneath the die resists flow, and it must be squeezed out. In particular, when there are individual pad openings, the encapsulant filling them must flow through the narrow space between each bump and the opening walls, making space for the bump. However, even when the individual openings are replaced by traces through trenches (which allow freer motion of the encapsulant), a substantial placement force must still be used. Another result of this force is that the bumps are pressed down on the substrate and partially coined. It should be remembered that in general only a few bumps (in theory only three) touch their pads originally, so the force on each one is relatively large and coins them, allowing more and more bumps to come in contact with their pads. Bringing as many bumps as possible close to their pads before reflow is believed to be essential in order to avoid electrical opens. When a conventionally fluxed component goes through reflow and the bumps melt and wet their pads, there is nothing preventing the collapse of the die, so it can start even with a small number of bumps wetting their pads originally. As soon as the die starts collapsing, more and more bumps will touch their pads and pull the die downwards. Provided the bumps are all sufficiently fluxed and there is nothing preventing them from touching their pads (e.g. mask misregistration, a very short bump, etc.), the result will be a 100% interconnected assembly. In contrast to this, reflow encapsulants resist the downward motion and collapse of the component. If only a few bumps are in contact with their pads initially, the force exerted by them after they melt and form joints may be too feeble to squeeze more liquid from underneath the die and allow collapse. This is especially true if polymerization of the encapsulant has already started and the material has thickened appreciably.

Incidentally, the liquid resists not only downward motion of the die, but lateral motion as well, although to a smaller extent. Therefore, there should be reduced self-centering of the component during reflow, and accurate placement is more important when reflow encapsulants are used in comparison with conventionally fluxed components.

The force required for electrical continuity of the assembly varies with the viscosity of the encapsulant, more viscous liquids requiring higher force. For example, a die with 88 bumps and one of the most viscous materials have been giving consistently good results when a placement force of 800 g was used, but when the force was reduced to 500 g for two die, both gave opens. On the other hand, one should be careful not to use excessive force. Apart from the obvious danger of damaging the die, too high a force will cause the substrate to bend and then bounce when the force is released, perhaps even leading to movement of the component, but certainly leaving the bumps coined to a surface that no longer exists.

Reflow

Reflow is, perhaps, the trickiest step when reflow encapsulants are used, because several processes take place during it, some of them with conflicting requirements. The material must remain fluid enough during reflow so that it will not prevent joints from forming and the die from collapsing, and yet by the time the board comes out of the reflow oven, it must have cured to the desired extent (in many cases, completely). It is not surprising then that some of these materials are quite sensitive to variations in the reflow profile, leading to opens if the temperature becomes too high too soon, or even if the heating rate is too high.

There are several reflow encapsulants that require a more-or-less traditional SMT profile (Figure 6, curve labeled "Standard"). This offers a definite advantage if other surface mount components are to be built at the same time. In addition, these materials are expected to be the least sensitive to variations in the reflow profile, since they are "slow" to start polymerizing. However, they all require a certain amount of post-curing after reflow, typically 30 to 40 minutes, but occasionally even longer.

At the other extreme, there are materials that are not compatible with the "soak" period of the traditional reflow. Instead they require a continuous ramping to reflow conditions, followed by a few minutes at a relatively lower temperature. The post-heat period makes the profile look like a mirror image of the traditional SMT profile (Figure 6, curve labeled "Inverted"). These materials do not require post-cure; essentially, the post-curing has been moved inside the reflow oven.

In between these two extremes, there is a whole range of reflow profiles, which can be called "intermediate" (one such profile is included in Figure 6). They are characterized by a short or completely absent soak period (occasionally soaking at a lower temperature than usual). There is no post-cure period inside the reflow oven. Some of them require post-curing of various lengths, while others are supposed to cure to a sufficient extent during reflow.

The last two groups of materials start curing faster, and are therefore more sensitive to variations in reflow profile. Figure 7 shows a case where the material started gelling too early, preventing this particular bump from wetting its pad, although it is clear from the shape of the bump that the other joints have collapsed and the molten solder was pressed onto the pad, trying to form a joint. A thin film of material can be

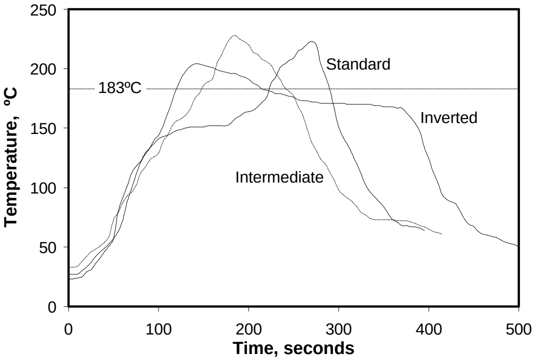

Figure 6. The three different types of reflow profiles

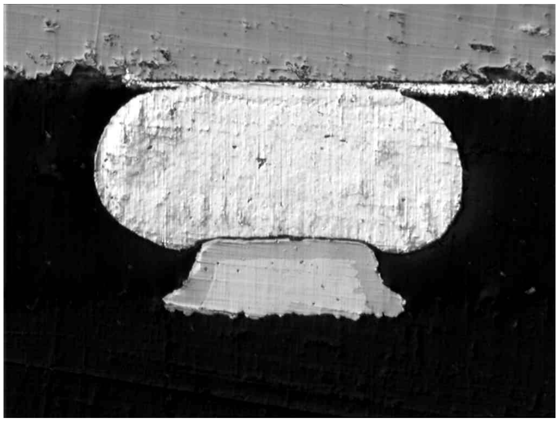

Figure 7. A joint that failed to form due to early onset of gelling of the encapsulant.

seen separating the bump from its pad. In general, joints formed with these materials often have a kind of irregular shape (Figure 8) indicating that the encapsulant started gelling while the solder was still liquid.

One question concerning these no-postcure materials is the effect of the exact profile on the degree of curing. Since the material must cure almost completely in the short interval between the end of collapse and the assembly cool down (before it even comes out of the oven), small variations in the cooling rate might affect the degree of curing to a significant extent. This effect will not be discernible when the assembly comes out of reflow, but it will later affect the reliability, since an insufficiently cured material will have reduced strength. Encapsulants that use a similar reflow profile but require post-cure should be affected by the cooling rate to a much smaller extent.

Voids can also appear and disappear during reflow. When the solder joints collapse, material is squeezed still further out and it tends to take bubbles still remaining behind the perimeter bumps with it. Once these bubbles reach the fillet, and since the material should still be quite fluid, they can rise to the surface of the liquid, or dissolve in it. As a rule, the fillet comes out of reflow void-free. On the other hand, bubbles that are under the body of the die or near center bumps cannot possibly reach the fillet during collapse. The only way they can disappear is by dissolving in the liquid encapsulant, and indeed bubbles have been observed to dissolve in a drop of encapsulant held between glass slides. This dissolution is helped by small initial size of the bubbles and by elevated temperature, but it slows down as the material thickens and stops when it gels or if moisture comes out of the substrate and diffuses into the still-existing bubbles. This behavior shows another advantage of materials requiring a traditional reflow profile. The relatively long soak period at an elevated temperature provides ideal conditions for bubbles to dissolve.

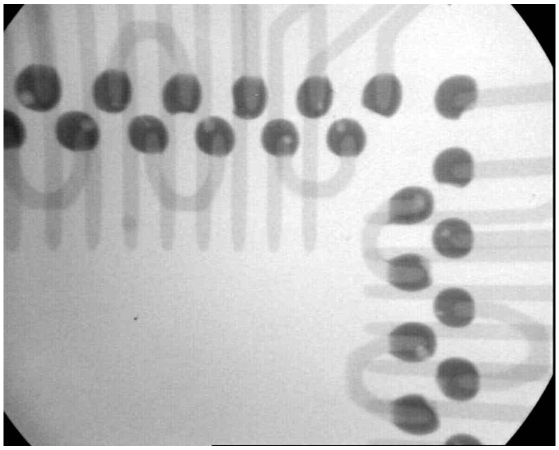

Figure 8. X-ray picture showing irregularly shaped joints with a "fast" (no-postcure) reflow encapsulant.

If the board is not sufficiently dry, new voids will form even if all the assembly bubbles have dissolved during the earlier stages of the reflow cycle.

Reliability

Capillary flow underfills are engineered to have a coefficient of thermal expansion (CTE) matching that of the solder joints; this is accomplished by the addition of appropriate filler particles to the liquid encapsulant. Reflow encapsulants do not contain fillers, because these would interfere with the formation of joints. As a result, they all have a CTE much higher than that of the solder joints and they should not be expected to perform as well in the usual thermal cycling tests. In fact, their performance depends on several factors, including design and process parameters. However, recent experiments have demonstrated better performances than was achieved with regular encapsulants a couple of years ago. With proper process control and the best materials ¼" die on 40 mil thick FR-4 substrates have survived over 3000 LLTS cycles (-55°C to 125°C). The performance in other tests, notably aging and moisture, requires further investigation.

Summary

Reflow encapsulants, within their limitations, can replace traditional capillary flow underfills in many cases, simplifying the assembly process and reducing costs. With proper board preparation and dispense and reflow conditions, voids in the final product can be minimized or eliminated; many materials consistently provide 100% electrical continuity if the proper assembly parameters are used. Narrow process windows could be an issue for some materials, but others are less sensitive. Performance in thermal cycling is not as good as with traditional underfills, but it keeps improving and it already exceeds the requirements for many applications.